



AI-driven Innovation in Medical Imaging

Focus on Lung Cancer and Cardiovascular Diseases

Comte, V., Bertolini, L., Consoli, S., Leoni, G., Zanca, F.,
Querci, M., Ceresa, M., Wiesenthal, T.

2025



This document is a publication by the Joint Research Centre (JRC), the European Commission's science and knowledge service. It aims to provide evidence-based scientific support to the European policymaking process. The contents of this publication do not necessarily reflect the position or opinion of the European Commission. Neither the European Commission nor any person acting on behalf of the Commission is responsible for the use that might be made of this publication. For information on the methodology and quality underlying the data used in this publication for which the source is neither Eurostat nor other Commission services, users should contact the referenced source. The designations employed and the presentation of material on the maps do not imply the expression of any opinion whatsoever on the part of the European Union concerning the legal status of any country, territory, city or area or of its authorities, or concerning the delimitation of its frontiers or boundaries.

Contact information

Name: Valentin Comte

Email: valentin.comte@ec.europa.eu

EU Science Hub

<https://joint-research-centre.ec.europa.eu>

JRC142224

EUR 40368

PDF ISBN 978-92-68-28966-2 ISSN 1831-9424 doi:10.2760/6281523 KJ-01-25-352-EN-N

Luxembourg: Publications Office of the European Union, 2025

© European Union, 2025



The reuse policy of the European Commission documents is implemented by the Commission Decision 2011/833/EU of 12 December 2011 on the reuse of Commission documents (OJ L 330, 14.12.2011, p. 39). Unless otherwise noted, the reuse of this document is authorised under the Creative Commons Attribution 4.0 International (CC BY 4.0) licence (<https://creativecommons.org/licenses/by/4.0/>). This means that reuse is allowed provided appropriate credit is given and any changes are indicated.

For any use or reproduction of photos or other material that is not owned by the European Union permission must be sought directly from the copyright holders.

How to cite this report: Comte, V., Bertolini, L., Consoli, S., Leoni, G., Zanca, F. et al., *AI-driven Innovation in Medical Imaging - Focus on Lung Cancer and Cardiovascular Diseases*, Publications Office of the European Union, Luxembourg, 2025, <https://data.europa.eu/doi/10.2760/6281523>, JRC142224.

Contents

- Abstract 4
- Executive Summary 5
- 1. Introduction 11
- 2. Policy context 13
 - 2.1. Regulatory Context 13
 - 2.1.1. General Data Protection Regulation (GDPR) 13
 - 2.1.2. Medical Devices Regulation (MDR) 13
 - 2.1.3. Artificial Intelligence Act (AI Act) 14
 - 2.1.4. European Health Data Space (EHDS) 14
 - 2.2. European Initiatives Supporting AI in Medical Imaging 15
 - 2.2.1. European Cancer Imaging Initiative 15
 - 2.2.2. AI4HI Network 15
 - 2.2.3. Testing and Experimentation Facility for Health (TEF-Health) 17
 - 2.2.4. AI Factories 17
 - 2.2.5. European Research Infrastructure Consortia (ERICs) 18
 - 2.2.6. European Digital Infrastructure Consortia (EDICs) 18
 - 2.2.7. EIC Portfolio: Supporting AI-Driven Innovation in Medical Imaging 18
- 3. AI for Medical Imaging: Where We Are and What Comes Next 20
 - 3.1. Types of Medical Image Data 20
 - 3.2. Fundamental AI techniques in Medical Imaging 21
 - 3.2.1. Convolutional Neural Networks 22
 - 3.2.2. Transformer Networks 23
 - 3.2.3. Multimodal Models 25
 - 3.3. Main Application Areas of AI in Medical Imaging 26
 - 3.3.1. Segmentation 27
 - 3.3.1.1. CNNs in Medical Image Segmentation 27
 - 3.3.1.2. Vision Transformers (ViTs) in Medical Image Segmentation 28
 - 3.3.2. Detection 29
 - 3.3.3. Registration 30
 - 3.3.4. Classification 32

3.3.5.	Generative AI in Medical Imaging.....	33
3.3.5.1.	Image Enhancement and Reconstruction.....	33
3.3.5.2.	Disease Progression Modelling.....	34
3.3.5.3.	Synthetic Medical Image Generation.....	34
4.	Barriers and Drivers.....	35
4.1.	Infrastructure and Interoperability.....	35
4.2.	Privacy Preserving Learning on Medical Imaging.....	36
4.2.1.	Federated Learning.....	36
4.2.2.	De-identification and Anonymisation.....	36
4.2.3.	Differential Privacy.....	37
4.2.4.	Homomorphic Encryption.....	37
4.2.5.	Synthetic Data Generation.....	37
4.3.	Explainable Artificial Intelligence (xAI) in Medical Imaging.....	37
5.	Selected Use Cases.....	39
5.1.	AI-Driven Diagnosis and Analysis of Lung Lesions.....	39
5.1.1.	Chest Anatomy Segmentation.....	39
5.1.2.	Lung Nodules Detection.....	40
5.1.3.	Follow-up of Nodules Over Time.....	42
5.1.3.1.	Siamese Neural Network.....	42
5.1.3.2.	Cascaded Registration.....	43
5.1.4.	Conformal Risk Control for Pulmonary Nodule Detection.....	44
5.1.5.	Web Platform and Tools for Clinical Imaging Workflows.....	45
5.1.5.1.	Anonymisation.....	46
5.1.6.	Discussion.....	47
5.2.	Biomechanics-Driven Neural Network for Cardiovascular Disease Identification.....	48
5.2.1.	Deep Cascaded Registration.....	48
5.2.2.	Biomechanical Regularisation.....	49
5.2.3.	Segmentation with Multi-Frame Propagation.....	50
5.2.4.	Classification of CVDs.....	51
5.2.5.	Discussion.....	53
6.	Conclusions.....	55

References 57

List of figures..... 65

List of tables..... 67

Abstract

Artificial Intelligence (AI) is nowadays increasingly used in the healthcare sector to support the **prevention**, early **diagnosis**, **monitoring**, and **outcome prediction** of **non-communicable diseases** (NCDs), which account for most deaths and disease burden across Europe. By automating tasks such as organ segmentation, lesion detection, disease classification, and uncertainty quantification, AI technologies offer the potential to reduce clinician workload, improve diagnostic accuracy, and enable earlier interventions.

This report reviews the most widely adopted **AI methods in medical imaging**, focusing on segmentation, detection, and classification tasks, and assesses their technological maturity and readiness for clinical uptake. It illustrates key technical and systemic requirements through **two practical use cases**. The first use case is an end-to-end AI pipeline for lung cancer imaging, which includes full chest multi-organ segmentation from CT scans, pulmonary nodule detection and segmentation, longitudinal tracking of nodules, and uncertainty quantification. To support large-scale, privacy-preserving research and clinical deployment, this use case also develops a complete workflow for data anonymisation, secure upload, and web-based visualisation and validation by clinicians. The second use case focuses on cardiovascular disease classification, leveraging a biomechanics-informed model that extracts physiologically meaningful features from cine cardiac MRI sequences to enhance explainability and trustworthiness.

Our findings demonstrate that while AI holds strong promise to transform healthcare delivery for NCDs, achieving clinical deployment requires careful attention to **data access**, model **transparency**, and robust **validation**. Based on technical outcomes and operational experience, we offer concrete recommendations to support EU initiatives aiming at the safe, effective, and trustworthy development of AI in healthcare.

Authors

Valentin Comte, Lorenzo Bertolini, Sergio Consoli, Gabriele Leoni, Federica Zanca, Maddalena Querci, Mario Ceresa, Tobias Wiesenthal.

Acknowledgments

We would like to thank Claudius Griesinger (JRC.F.2), Aleksandra Wesolowska, Matthieu Ruthven, Szymon Bielecki, Kyriacos Hatzaras (DG CNECT), Yiannos Toliás, and Andreia Oliveira (DG SANTE) for their helpful comments and support during the preparation of this report.

Executive Summary

Why This Report?

Europe is facing an increasing burden from **non-communicable diseases** (NCDs), such as **cancer** and **cardiovascular conditions**, which are responsible for **80% of the disease burden** in EU countries and represent the **leading causes of avoidable premature deaths**¹. Healthcare systems are under increasing strain from demographic changes, rising chronic disease prevalence, and workforce shortages. Artificial Intelligence (AI) and Deep Learning (DL) offer critical opportunities to improve **early detection**, accelerate **diagnostic** workflows, and **personalised medicine**.

What Does The Report Focus On?

The report concentrates on the following major AI tasks in medical imaging:

- **Segmentation** – automatic annotation of anatomical structures and pathological regions,
- **Detection** – identification of lesions and abnormalities, such as pulmonary nodules,
- **Classification** – diagnosis and stratification of disease conditions,
- **Registration** – spatial alignment of medical images across time points, modalities, or subjects,
- **Synthetic image generation** – creation of realistic medical images using generative models to augment datasets or simulate rare cases.

We consider the following two practical case studies in order to illustrate these AI applications:

- **Lung Cancer Imaging:** We develop an end-to-end AI system for thoracic Computed Tomography (CT) imaging that performs fully automated multi-organ segmentation of the chest, as well as the detection, segmentation, and longitudinal tracking of pulmonary nodules. To enhance reliability in clinical decision-making, we integrate conformal prediction techniques—specifically Conformal Risk Control—to provide statistically valid, patient-level confidence estimates on detection sensitivity. This system operates within a fully privacy-preserving framework, incorporating a lightweight anonymisation tool, secure cloud-based data transfer, and an interactive web platform designed for clinical expert review and validation.
- **Cardiovascular Disease Imaging:** We develop a biomechanics-informed neural network for cardiac disease classification, leveraging DL-based image registration to estimate myocardial deformation from cine-cardiac Magnetic Resonance Imaging (MRI). Physiologically meaningful features such as local strains and biomechanical moduli are extracted to enhance the explainability and trustworthiness of classification results. This approach integrates physical priors into AI models, providing greater transparency and clinical interpretability.

¹ https://health.ec.europa.eu/non-communicable-diseases/overview_en

Policy Context

The topic of this report aligns with major EU strategies on health innovation and digital transformation, including **Europe's Beating Cancer Plan**², the **European Health Data Space (EHDS) Regulation**³, and **Horizon Europe research priorities**⁴. It directly supports the objectives of the **European Cancer Imaging Initiative**⁵, the flagship initiative under Europe's Beating Cancer Plan, which aims to improve access to cancer imaging data across Europe. Implemented through the **European Federation for CAncer IMages (EUCAIM) project**⁶, which seek to build a **federated, secure infrastructure** for imaging data across Europe to foster the development of **reliable and generalisable AI tools**. The growing computational demands associated with developing and deploying advanced AI models, particularly for high-resolution and longitudinal imaging data, highlight the urgent need for strategic investment in AI-optimised infrastructure. European **AI Factories**⁷—sovereign high-performance computing facilities dedicated to AI research and application—are critical to support researchers and clinicians, ensure technological autonomy, and maintain Europe's competitiveness in trustworthy medical AI.

Key Conclusions

AI technologies in medical imaging have reached a high level of maturity, particularly in segmentation and detection tasks. These methods are already enhancing clinical workflows and diagnostic accuracy, with demonstrated benefits in lung cancer screening and cardiovascular disease classification. Their integration into healthcare systems reflects AI's growing relevance for addressing Europe's increasing burden of non-communicable diseases.

Nonetheless, classification tasks pose greater challenges, largely because they directly inform clinical decision-making and therefore require a higher degree of explainability and transparency. While AI models often achieve strong technical performance, their lack of interpretability and limited clinical validation remain major obstacles to trust and adoption. The use cases presented in this report illustrate how integrating domain-specific knowledge and explainable AI methods can improve both predictive accuracy and clinical relevance.

Persistent systemic barriers must be addressed to fully realise the potential of AI in this domain. These include the scarcity of large, annotated datasets, fragmented data infrastructures, and a lack of harmonised standards for evaluating AI performance. European initiatives such as EUCAIM and EHDS play a vital role in enabling secure, federated, and interoperable access to imaging data for AI development, while ensuring compliance with data protection regulations. In addition, the growing computational demands of modern AI models call for sustained investment in sovereign AI Factories that can provide the infrastructure necessary for trustworthy medical AI.

² https://commission.europa.eu/strategy-and-policy/priorities-2019-2024/promoting-our-european-way-life/european-health-union/cancer-plan-europe_en

³ https://health.ec.europa.eu/ehealth-digital-health-and-care/european-health-data-space-regulation-ehds_en

⁴ https://research-and-innovation.ec.europa.eu/funding/funding-opportunities/funding-programmes-and-open-calls/horizon-europe_en

⁵ <https://digital-strategy.ec.europa.eu/en/policies/cancer-imaging>

⁶ <https://dashboard.eucaim.cancerimage.eu/>

⁷ <https://digital-strategy.ec.europa.eu/en/policies/ai-factories>

Finally, trust in AI must be built through clinician-centred innovation. Active involvement of healthcare professionals in the design and validation of AI tools, along with a focus on usability, transparency, and clinical relevance, will be essential to support adoption and ensure that AI contributes meaningfully to improving patient care across Europe.

Main Findings

How Are AI Methods Used in Healthcare?

AI and DL methods are extensively employed for image segmentation, abnormality detection, and disease classification. Practical use cases demonstrate successful applications in lung cancer nodule detection and longitudinal tracking, as well as in CVD classification using biomechanical modelling.

What Are the Challenges?

Deploying AI in medical imaging faces several persistent challenges. Access to large, high-quality, and well-annotated datasets remains limited, often due to fragmentation, privacy constraints, and the high cost of expert annotation. Interoperability issues across healthcare systems hinder the integration of AI tools into clinical workflows, while variation in data formats and annotation standards complicate model training and validation. In addition, AI models still lack transparency, making it difficult for clinicians to trust their outputs. Finally, the growing complexity of AI systems is driving up computational demands, highlighting the need for scalable and accessible computing resources to support development and deployment.

What Are The Policy Needs?

To enable the deployment of AI in clinical settings, several policy enablers must be addressed in parallel. Secure access to high-quality, annotated datasets is critical, supported by trusted processing environments that ensure compliance with data protection and ethical standards. Robust computational infrastructure is also needed to support the development, validation, and clinical integration of advanced AI models. Initiatives such as the AI Factories are working to meet this need. Interoperability remains a key requirement to ensure that AI systems can be integrated across different healthcare providers and national systems. This calls for harmonised data formats, annotation protocols, and model evaluation frameworks. Enhanced coordination between AI developers, clinicians, and regulators is crucial to ensure that AI tools are clinically meaningful, technically robust, and aligned with regulatory expectations. EU-level infrastructures such as EUCAIM help facilitate this coordination by supporting cross-border collaboration and data sharing in areas like cancer imaging, creating the necessary foundations for real-world deployment.

Glossary

This glossary provides definitions of key technical and regulatory terms used throughout the report. It is designed to help a broad audience, including policymakers, clinicians, and researchers, understand the terminology related to artificial intelligence and medical imaging.

AI (Artificial Intelligence): A field of computer science that enables machines to perform tasks that typically require human intelligence, such as learning, reasoning, and pattern recognition.

Annotation: The process of labelling data (e.g. medical images) with information that can be used to train AI algorithms, such as identifying tumours or organs.

Anonymisation: The process of removing or modifying personal identifiers from data to ensure individuals cannot be identified.

Attention Mechanism: A component in neural networks (particularly Transformers) that allows the model to dynamically focus on relevant parts of the input data when making predictions.

Bias (in AI): Systematic error in AI models that can lead to unfair or inaccurate outcomes for certain groups or data types.

Benchmark Dataset: A standardised dataset used to evaluate and compare the performance of different AI models.

Classification (in medical imaging): The use of AI to automatically categorise medical data or images into clinically relevant groups, such as diseased vs. non-diseased, or by specific disease type.

CNN (Convolutional Neural Network): a deep learning model particularly effective for image analysis.

CT (Computed Tomography): A medical imaging method using X-rays to produce detailed cross-sectional images of the body.

Dice Score: A statistical metric used to measure the similarity between two samples. In medical imaging, it is often used to evaluate the overlap between predicted and ground-truth segmentations.

DL (Deep Learning): A type of machine learning based on artificial neural networks, especially effective for tasks such as image analysis.

DICOM (Digital Imaging and Communications in Medicine): A standard for storing, transmitting, and handling medical imaging data.

EHDS (European Health Data Space): An EU initiative to promote safe access to and sharing of health data across Member States for healthcare delivery, research, and innovation.

EUCAIM (European Federation for Cancer Images (EUCAIM): The project implementing the European Cancer imaging Initiative, Europe's Beating Cancer Plan flagship, to create a federated infrastructure for cancer imaging data.

xAI (Explainable AI): AI systems designed to make their operations and predictions understandable to humans, critical in high-stakes domains like healthcare.

Explainability: The extent to which an AI system can provide understandable justifications for its outputs. This may involve ante hoc approaches, where the model is designed to be explainable from the outset, or post hoc techniques that generate explanations after the fact, such as highlighting influential input features or generating natural language rationales. Explainability is especially important for complex or opaque models (e.g. deep neural networks), where interpretability alone is insufficient.

Feature Selection: A process used to identify the most relevant variables (features) in a dataset to use for building predictive models.

FL (Federated Learning): A machine learning approach that trains algorithms across decentralised devices or servers holding local data, without exchanging that data.

GAN (Generative Adversarial Network): a type of neural network used to generate synthetic data.

GDPR (General Data Protection Regulation): The EU regulation that governs data protection and privacy for individuals within the EU.

GPU (Graphics Processing Unit): A hardware component that accelerates computation, especially used in deep learning to process large volumes of data efficiently.

Ground Truth: The actual, verified data used as a benchmark to evaluate the accuracy of model predictions, such as expert-annotated medical images.

Imaging Modality: A specific technique or method used to create images of the body for clinical analysis, e.g., MRI, CT, X-ray.

Interoperability: The ability of different information systems and devices to exchange and use data in a coordinated manner.

Interpretability: The degree to which a human can understand the internal logic of an AI system based on its structure and parameters. An interpretable model, such as a decision tree or linear regression, allows users to follow how specific inputs lead to specific outputs, without requiring additional tools or explanations. It is typically associated with models that are inherently transparent and simple enough for direct human comprehension.

MDR (Medical Device Regulation): The EU regulation that governs the safety and performance of medical devices, including AI-based diagnostic software.

MRI (Magnetic Resonance Imaging): A non-invasive imaging technique that uses magnetic fields and radio waves to create detailed images of organs and tissues.

NCDs (Non-Communicable Diseases): diseases not passed from person to person, such as cancer and cardiovascular diseases.

PACS (Picture Archiving and Communication System): A medical imaging technology used for storing, retrieving, presenting, and sharing images produced by various modalities.

Radiomics: The extraction of large amounts of quantitative features from medical images, which can be used to support diagnosis or predict treatment response.

Registration (in medical imaging): The process of aligning two or more images into a common coordinate system, which is crucial in longitudinal or multi-modality imaging studies.

Segmentation (in medical imaging): The process of partitioning a medical image into different regions, such as organs or lesions, for further analysis.

Synthetic Data: Data artificially generated to resemble real-world data, often used when real data is limited or sensitive.

Training Data: A dataset used to teach an AI model by adjusting its internal parameters.

Validation Data: A subset of data used to fine-tune model parameters during training to prevent overfitting.

Testing Data: A separate dataset used to evaluate the performance of an AI model after training is completed.

Transformer: A neural network architecture, initially developed for natural language processing. It uses self-attention to process sequences and has recently been applied to imaging tasks, such as classification and segmentation.

Trustworthiness: The overall reliability, safety, and ethical integrity of an AI system. A trustworthy system performs as intended, protects data privacy, complies with legal and regulatory standards, and earns the confidence of users and stakeholders.

X-ray: A basic imaging modality using electromagnetic radiation to view the inside of the body, especially bones.

US (Ultrasound): A medical imaging technique using high-frequency sound waves to produce images of structures within the body.

ViT (Vision Transformer): a deep learning architecture that applies transformer models to image analysis.

1. Introduction

In recent years, artificial intelligence (AI) and deep learning (DL) have significantly impacted various industries and public sectors, particularly healthcare. Europe's healthcare systems are facing increasing pressure due to the growing burden of non-communicable diseases (NCDs), including cancer and cardiovascular conditions, which are responsible for 80% of the disease burden in the EU countries and represent the leading causes of avoidable premature deaths across EU today. At the same time, a shortage of healthcare workers and under-resourced systems are putting strain on clinical workflows and patient care. In this context, AI and DL technologies are seen as promising solutions to improve prevention, diagnosis, and monitoring of NCDs, particularly through medical imaging.

The integration of AI and DL in medical imaging not only holds promise for enhancing diagnostic accuracy but also for optimising clinical workflows, thereby addressing the critical shortage of healthcare workers. By automating image analysis and supporting clinical decision-making, these technologies can potentially reduce the workload on healthcare professionals and improve patient outcomes.

As healthcare systems look to innovate, understanding and integrating these advanced technologies becomes crucial in transforming medical practices and addressing systemic challenges. Recent expert discussions, including a dedicated workshop on AI and medical imaging [1], identified several emerging technologies with high transformative potential. Among them, generative AI and digital twins⁸ stand out. Generative AI can help address data scarcity by producing synthetic data, while digital twins offer dynamic, patient-specific simulations to support personalised treatment planning and more precise modelling of disease progression. Experts also pointed to multimodal data integration and explainable AI as crucial directions for development, especially in clinical settings where trust and transparency are essential. As AI continues to evolve, its diverse applications in healthcare are attracting interest from both the public and private sectors. Investment in AI-driven healthcare solutions is accelerating, driven by the urgency to improve patient care and the efficiency of healthcare systems.

This report was developed by the Joint Research Centre (JRC) to support the European Commission and its strategic policy activities in the field of AI and healthcare. It aims to provide evidence-based input that may inform future funding priorities under programmes such as Horizon Europe⁹ and the European Innovation Council (EIC)¹⁰, including targeted EIC Challenges and Digital Europe initiatives like EUCAIM. The goal of the report consists in exploring the capabilities and limitations of AI in healthcare and identify key policy needs.

What Policy Problem Does It Address?

While AI models for medical image analysis are being rapidly developed, their clinical adoption remains limited. This gap between innovation and implementation is due to a combination of factors:

⁸ <https://www.edith-csa.eu/>

⁹ https://research-and-innovation.ec.europa.eu/funding/funding-opportunities/funding-programmes-and-open-calls/horizon-europe_en

¹⁰ https://eic.ec.europa.eu/index_en

insufficient clinical validation¹¹, data fragmentation¹², lack of interoperability¹³, concerns over explainability¹⁴, and uneven regulatory frameworks¹⁵ [2] [3] [4]. These challenges risk slowing down the uptake of tools that could otherwise strengthen health systems, reduce diagnostic delays, and support more personalised care.

Decisionmakers require a clearer understanding of which AI technologies are ready for use, what their limitations are, and how to prioritise investment and coordination efforts to accelerate safe and impactful deployment.

The objectives of this report are the following:

- **Review** the most relevant AI and DL methods currently used in medical image analysis, particularly for annotation, detection, and classification.
- **Assess** their maturity and level of readiness for clinical uptake.
- **Analyse** two practical use cases: one focused on the detection of lung lesions in lung CT, and the other on CVDs classification.
- **Identify** key barriers to clinical implementation, including issues related to data access and anonymisation, interoperability, model evaluation, and trustworthiness.
- **Provide** practical, policy-relevant recommendations that support responsible innovation and guide EU-funded initiatives such as the European Cancer Imaging Initiative, the European Health Data Space, and future EIC Challenges.

The report is structured as follows: Section 2 outlines the policy context, including key European regulations and strategic initiatives supporting AI in medical imaging. Section 3 provides a technical overview, covering data types, core AI techniques, application areas, and current barriers and enablers to deployment. Section 4 presents two practical use cases, lung cancer detection and cardiovascular disease classification, demonstrating how AI can be applied in real-world clinical scenarios. Finally, Section 5 offers conclusions and policy reflections on how to advance trustworthy, effective, and human-centric AI in healthcare across Europe.

¹¹ <https://cordis.europa.eu/project/id/101057699>

¹² <https://www.european-health-data-space.com/>

¹³ https://cordis.europa.eu/programme/id/H2020_DT-TDS-05-2020

¹⁴ https://ec.europa.eu/futurium/en/system/files/ged/ai-and-interpretability-policy-briefing_creative_commons.pdf

¹⁵ <https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai>

2. Policy context

The deployment of AI in medical imaging within the European Union (EU) is governed by a complex, increasingly integrated legal framework. Several interrelated legislative instruments aim to ensure that the use of AI respects fundamental rights, patient safety, and data protection, while also fostering innovation and cross-border collaboration.

2.1. Regulatory Context

2.1.1. General Data Protection Regulation (GDPR)

The **General Data Protection Regulation (GDPR)**¹⁶ (Regulation (EU) 2016/679) provides the foundation for personal data protection across all EU sectors, including healthcare. Under GDPR, health data is classified as a special category of personal data, subject to stricter processing conditions. Consent or another lawful basis (e.g. public interest in public health or scientific research) is required for the use of health data in AI systems. Moreover, GDPR provides patients with rights such as access, rectification, erasure, and data portability.

For AI in medical imaging, GDPR has direct implications:

- It mandates **transparency** and **explainability** in data processing.
- It limits the use of **automated decision-making** (Article 22), unless specific conditions are met.
- It introduces **data minimisation** and **purpose limitation** requirements, which can be challenging for training AI models that require large, diverse datasets.

While GDPR grants data portability, it primarily applies to raw personal data (e.g., imaging files). The **European Health Data Space (EHDS)** proposal would complement this by extending portability to inferred and processed data, such as diagnoses derived through AI algorithms.

2.1.2. Medical Devices Regulation (MDR)

The **Medical Devices Regulation (MDR)**¹⁷ (Regulation (EU) 2017/745) defines the requirements for software, including AI systems, that are intended for medical purposes such as diagnosis or monitoring. When AI algorithms are used in medical imaging and influence clinical decisions, they are considered **medical devices** and must be CE-marked.

This includes compliance with:

- **Safety and performance requirements**, including risk management and clinical evaluation.
- **Post-market surveillance** and continuous performance monitoring.

¹⁶ <https://eur-lex.europa.eu/eli/reg/2016/679/oj/eng>

¹⁷ <https://eur-lex.europa.eu/eli/reg/2017/745/oj/eng>

- Demonstration of **robustness and generalisability** in clinical settings.

AI models used for image classification, lesion detection, or segmentation fall under MDR's scope and must undergo conformity assessment procedures appropriate to their risk class.

2.1.3. Artificial Intelligence Act (AI Act)

The **AI Act**¹⁸ is the EU's first horizontal regulation targeting AI systems. It categorises AI applications by risk level. AI systems used for **medical diagnosis** are considered **high-risk**, requiring:

- High-quality datasets to minimise bias.
- Transparent and explainable decision-making processes.
- Robust monitoring, traceability, and human oversight mechanisms.
- Compliance with existing sectoral laws, such as GDPR and MDR.

The AI Act establishes a harmonised framework to ensure trust in AI systems used in healthcare, while supporting innovation. Importantly, it introduces **conformity assessments** and **registration in an EU database**, aiming to make high-risk AI more trustworthy and auditable.

2.1.4. European Health Data Space (EHDS)

The **European Health Data Space (EHDS)**¹⁹ is a sector-specific initiative designed to enable secure sharing and secondary use of health data. Once adopted, it will:

- Guarantee **patients' access** to their electronic health records across borders.
- Establish infrastructures such as **MyHealth@EU**²⁰ (for primary use) and **HealthData@EU**²¹ (for secondary use, e.g. AI development and training).
- Mandate **interoperability and common standards**, critical for cross-border AI deployment.
- Support **secondary use** of health data for research and innovation, under clear governance frameworks and ethical oversight.

EHDS is expected to interact with GDPR, MDR, and the AI Act, serving as a *lex specialis* in the health domain. It addresses gaps in data portability under GDPR by extending rights to inferred and processed data. Furthermore, it includes obligations when AI systems interoperate with EHR systems.

¹⁸ <https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng>

¹⁹ <https://www.european-health-data-space.com/>

²⁰ https://www.interregeurope.eu/sites/default/files/2022-04/EW0220140ENN.en_.pdf

²¹ <https://ec.europa.eu/newsroom/sante/newsletter-archives/61867>

2.2. European Initiatives Supporting AI in Medical Imaging

2.2.1. European Cancer Imaging Initiative

The European Federation for Cancer Images (EUCAIM)²² is the implementation project of the European Cancer Imaging Initiative²³, a flagship action under Europe's Beating Cancer Plan²⁴. Funded by the Digital Europe Programme²⁵, EUCAIM aims to develop a federated and interoperable infrastructure for cancer imaging data across Europe. By connecting imaging repositories and clinical centres, it supports the creation of a user-friendly and secure platform that enables access to high-quality, annotated datasets. This infrastructure is designed to facilitate the development, testing, and validation of AI tools for cancer diagnosis, prognosis, and treatment planning, advancing personalised medicine and clinical research throughout the EU.

EUCAIM is grounded in the principles of privacy, interoperability, and FAIR principles²⁶. By connecting imaging repositories and clinical sites across Member States, it will provide researchers and developers access to high-quality, annotated datasets while ensuring compliance with the GDPR and future regulations under the EHDS. This infrastructure facilitates federated learning, allowing algorithms to be trained across distributed datasets without centralising sensitive patient data.

EUCAIM builds upon the technical and methodological foundations laid by earlier EU-funded projects under the AI4HI network. In particular, the CHAIMELEON project provided a valuable blueprint by establishing a secure repository of harmonised and annotated imaging data and defining protocols for data sharing and model validation.

2.2.2. AI4HI Network

The AI for Health Imaging (AI4HI) network²⁷ brought together five large-scale EU-funded projects, each contributing to the development of AI solutions for cancer imaging. Although most of these projects have now concluded, they played a critical role in:

- Harmonising multimodal imaging data across different institutions and cancer types.
- Designing ethical and legal frameworks for cross-border data sharing.
- Developing and validating AI algorithms for lesion detection, segmentation, and diagnosis.
- Promoting federated learning and privacy-preserving AI techniques.

The AI4HI network includes:

²² <https://dashboard.eucaim.cancerimage.eu/>

²³ <https://digital-strategy.ec.europa.eu/en/policies/cancer-imaging>

²⁴ https://commission.europa.eu/strategy-and-policy/priorities-2019-2024/promoting-our-european-way-life/european-health-union/cancer-plan-europe_en

²⁵ <https://digital-strategy.ec.europa.eu/en/activities/digital-programme>

²⁶ <https://www.go-fair.org/fair-principles/>

²⁷ <https://ai4hi.net/>

- **CHAMELEON**²⁸: Focused on lung, breast, prostate, and colorectal cancer. It created a secure, harmonised imaging repository and workflows for training and validating AI tools.
- **EuCanImage**²⁹: built a FAIR-compliant cancer imaging platform linking imaging with biological data to enable secure, multi-centre AI development, evaluation, and responsible data sharing for personalised oncology.
- **ProCancer-I**³⁰: built a high-performance AI platform and the largest interoperable prostate cancer MRI repository to support precision care across diagnosis, metastasis detection, and treatment response prediction, aiming to reduce overdiagnosis and overtreatment through clinically validated models.
- **INCISIVE**³¹: developed an AI-based toolbox and a federated imaging repository to improve cancer diagnosis, prediction, and follow-up using multimodal data, while ensuring secure, ethical, and interoperable data sharing across Europe.
- **PRIMAGE**³²: built a cloud-based platform using in-silico imaging biomarkers and tumour modelling to support personalised diagnosis, prognosis, and treatment follow-up for paediatric cancers.
- **RadioVal**³³: aims to validate AI tools for breast cancer treatment by assessing their accuracy, fairness, usability, and robustness, with the goal of delivering trustworthy, explainable, and clinically usable solutions that support clinicians in real-world decision-making.

These projects addressed key challenges such as:

- **Data fragmentation** and heterogeneity across centres.
- Limited **generalisability** of AI models.
- The need for rigorous **clinical validation** and **interpretability**.
- Ensuring **legal and ethical compliance** for secondary use of health data.

Together, the AI4HI projects have shaped the landscape for AI in medical imaging in Europe. Their insights and infrastructures now serve as a foundation for EUCAIM, which aims to consolidate and scale these efforts into a long-term, sustainable platform that accelerates AI innovation in cancer care.

²⁸ <https://chameleon.eu/>

²⁹ <https://eucanimage.eu/>

³⁰ <https://www.procancer-i.eu/>

³¹ <https://incisive-project.eu/>

³² <https://www.primageproject.eu/>

³³ <https://radioval.eu/>

2.2.3. Testing and Experimentation Facility for Health (TEF-Health)

TEF-Health³⁴ is part of the broader network of European Testing and Experimentation Facilities (TEFs)³⁵, created to accelerate the development and adoption of **trustworthy AI** in the healthcare sector. TEFs are a core implementation instrument foreseen by the **AI Act**.

TEF-Health helps **AI technology providers** bring their solutions to market faster by offering **testing, validation, and compliance services** in clinically relevant settings. This includes access to real-world data, healthcare environments, and expert guidance to assess safety, performance, usability, and ethical compliance.

Under **Articles 74 and 75 of the AI Act**, TEFs are explicitly recognised as tools to:

- **Support conformity assessment** for high-risk AI systems in healthcare
- **Provide technical and scientific expertise** to developers and notified bodies
- **Enable safe and trustworthy deployment** of AI in line with EU harmonisation legislation

By combining **clinical infrastructure** with **regulatory alignment**, TEF-Health plays a critical role in bridging the gap between research, regulation, and real-world adoption of AI in medical imaging and broader healthcare applications.

2.2.4. AI Factories

AI Factories³⁶ are an emerging European initiative designed to scale the development of AI across sectors using high-performance computing (HPC). In the health domain, AI Factories provide the infrastructure needed to train and validate large-scale medical imaging models efficiently and ethically.

Key contributions of AI Factories include:

- **Access to supercomputing resources** for model training and experimentation
- **Support for federated and privacy-preserving AI workflows**
- **Scalable pipelines** that help develop robust and generalisable AI systems

These factories complement initiatives like EUCAIM by providing the computational backbone necessary to handle vast, diverse imaging datasets across the EU.

³⁴ <https://tefhealth.eu/home>

³⁵ <https://digital-strategy.ec.europa.eu/en/policies/testing-and-experimentation-facilities>

³⁶ <https://digital-strategy.ec.europa.eu/en/policies/ai-factories>

2.2.5. European Research Infrastructure Consortia (ERICs)

Several **ERICs**³⁷ play a foundational role in supporting data governance, access, and harmonisation across the life sciences and medical imaging ecosystem. These long-standing infrastructures ensure that data-driven AI research in health remains sustainable, ethical, and interoperable.

Relevant ERICs include:

- **BBMRI-ERIC**³⁸ (Biobanking and Biomolecular Resources): Offers access to biological samples and associated metadata essential for linking imaging to clinical outcomes.
- **ECRIN**³⁹ (Clinical Research): Supports the design and conduct of multinational clinical trials, providing pathways for evaluating AI models in real-world clinical contexts.
- **Euro-BioImaging**⁴⁰: Facilitates access to imaging technologies and expertise across Europe, with emphasis on FAIR data practices and integration with AI development pipelines.

Together, these ERICs provide the **structural and legal frameworks** necessary to enable secure, cross-border data use, making them key enablers of AI innovation in medical imaging.

2.2.6. European Digital Infrastructure Consortia (EDICs)

European Digital Infrastructure Consortium (EDIC)⁴¹ is an instrument made available to Member States under the Digital Decade Policy Programme 2030 to speed up and simplify the setup and implementation of multi-country projects. Eleven EU Member States are discussing a possible establishment of a European Digital Infrastructure Consortium (EDIC) to ensure sustain operation of the Cancer Image Europe platform developed by the EUCAIM project.

EUCAIM EDIC is envisaged as a digital infrastructure facilitating access to cancer imaging and related clinical data for secondary use in research and innovation, from hospital networks and research repositories across the EU. Beyond providing access to data, it would offer a secure, reliable and trustworthy environment for AI experimentation, creating a framework in Europe that makes it more efficient to conduct multi-country collaborative projects and AI validation studies, in respect of applicable data protection rules.

2.2.7. EIC Portfolio: Supporting AI-Driven Innovation in Medical Imaging

The European Innovation Council (EIC) supports a broad and fast-growing portfolio of over 30 projects advancing AI-driven innovation in medical imaging. These span the entire clinical pathway—from early screening and diagnostics to image-guided interventions and monitoring—and cover key clinical domains such as cardiology, oncology, neurology, and respiratory and emergency care. Technologies range from AI-enhanced MRI (including portable and low-field systems), cardiac ultrasound and chest

³⁷ https://research-and-innovation.ec.europa.eu/strategy/strategy-research-and-innovation/our-digital-future/european-research-infrastructures/eric_en

³⁸ <https://www.bbmri-eric.eu/>

³⁹ <https://ecrin.org/>

⁴⁰ <https://www.eurobioimaging.eu/>

⁴¹ <https://digital-strategy.ec.europa.eu/en/policies/edic>

CT analysis, to imaging biomarkers, 3D visualization for surgery, and cloud-native platforms for real-time clinical decision support. These projects reflect a strategic effort to combine cutting-edge imaging technologies with advanced artificial intelligence, contributing to faster, more precise, and more accessible diagnosis and treatment planning. Within this broader effort, the EIC has built a strong group of deep tech SMEs working at the intersection of medical imaging and AI. European infrastructures such as EUCAIM, TEF-Health, Euro-BioImaging and BBMRI-ERIC may offer valuable support to these SMEs by providing access to high-quality data, clinical expertise, and environments for testing, validation, and integration. These platforms can play a critical role in accelerating the deployment of trustworthy AI solutions in clinical settings, fostering interoperability, regulatory alignment, and cross-border collaboration for better healthcare outcomes across Europe.

3. AI for Medical Imaging: Where We Are and What Comes Next

AI has rapidly evolved into a powerful tool for medical image analysis, offering the potential to improve diagnostic accuracy, reduce clinical workload, and support more personalised treatment strategies. DL methods have demonstrated strong performance in tasks such as image segmentation and registration, abnormality detection, and disease classification. These capabilities are especially relevant for addressing the growing burden of NCDs, including cancer and cardiovascular conditions. The use of AI in medical imaging builds on decades of research in pattern recognition, computer vision, and radiology. However, the past decade has seen a shift from traditional image processing techniques to data-driven models that learn directly from large volumes of annotated medical images. This section provides an overview of the main categories of AI methods used in medical imaging today. It highlights their applications across different stages of the imaging pipeline, including image segmentation, feature detection, and disease follow-up and classification. In addition, it examines recent advances in generative AI, multimodal data integration, and the use of explainable AI techniques to address concerns around transparency and trust.

3.1. Types of Medical Image Data

Medical imaging encompasses a range of modalities, each providing different types of anatomical or functional information to support clinical decision-making (**Figure 1**). These imaging techniques are essential across nearly every medical specialty, from initial diagnosis to treatment planning and monitoring.

Magnetic Resonance Imaging (MRI): MRI leverages magnetic fields and radio waves to generate detailed images, primarily of soft tissues. It is particularly effective in **neuroimaging, cardiovascular, musculoskeletal,** and **liver** imaging. Functional and dynamic MRI techniques, such as **fMRI** and **cine-MRI**, enable the visualisation of tissue activity and motion, respectively.

Computed Tomography (CT): CT combines X-ray data from multiple angles to produce cross-sectional views of the body. It is commonly used to examine the **chest, brain, lungs,** and **kidneys,** offering high-resolution images that are particularly useful in trauma care, oncology, and cardiovascular diagnostics.

Positron Emission Tomography (PET): PET imaging captures metabolic activity using radioactive tracers. It is widely employed in **oncology, neuroimaging, cardiology,** and for assessing **infected tissues.** When combined with CT or MRI, PET enables precise anatomical and functional correlation.

Single-Photon Emission Computed Tomography (SPECT): Like PET, SPECT captures functional information using gamma-emitting radioisotopes. It is widely applied in **cardiology, neurology,** and **bone imaging,** offering a more accessible alternative to PET in many clinical settings.

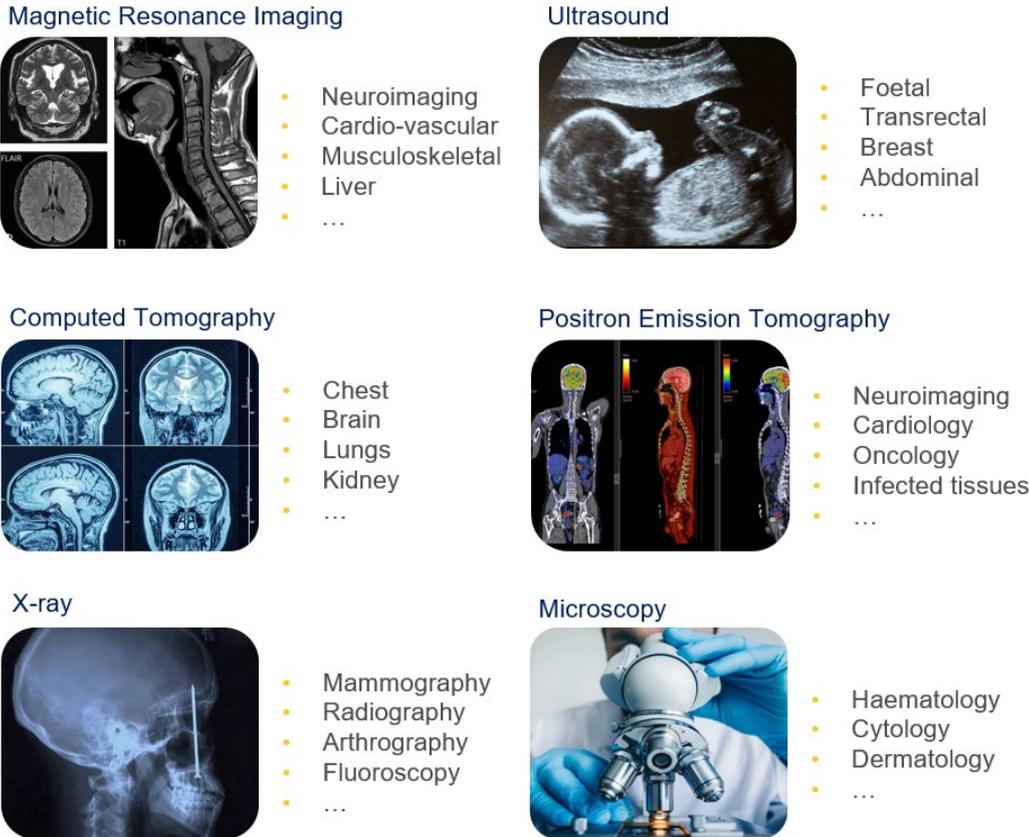
X-ray: X-ray imaging remains one of the most accessible and widely used modalities. It is fundamental in **mammography, radiography, arthrography,** and **fluoroscopy.** While less detailed than other techniques, its speed and availability make it ideal for routine assessments and emergency diagnostics.

Ultrasound (US): US uses high-frequency sound waves to generate real-time images. It is a preferred choice for **foetal, transrectal, breast,** and **abdominal** imaging due to its non-invasive nature and absence of ionising radiation. Its portability also makes it indispensable in point-of-care settings.

Microscopy: Microscopy enables visualisation at the cellular and sub-cellular level and is central to **haematology, cytology, and dermatology**. Advances in digital and AI-assisted microscopy are expanding its applications in pathology, allowing automated cell classification and disease detection from high-resolution histological slides.

Each imaging modality plays a complementary role in modern medicine. The integration of artificial intelligence in these domains promises to further enhance diagnostic precision, accelerate workflow efficiency, and support the transition towards personalised medicine.

Figure 1: Overview of major medical imaging modalities and their primary clinical applications.



Source: own elaboration.

3.2. Fundamental AI techniques in Medical Imaging

The application of AI to medical imaging has been largely driven by the development and success of machine learning (ML) and, more recently, DL approaches. These techniques enable computers to automatically detect patterns in complex data such as X-rays, Computed Tomography (CT) scans, or Magnetic Resonance Imaging (MRI), reducing the need for hand-crafted rules and enabling faster, often more accurate interpretations.

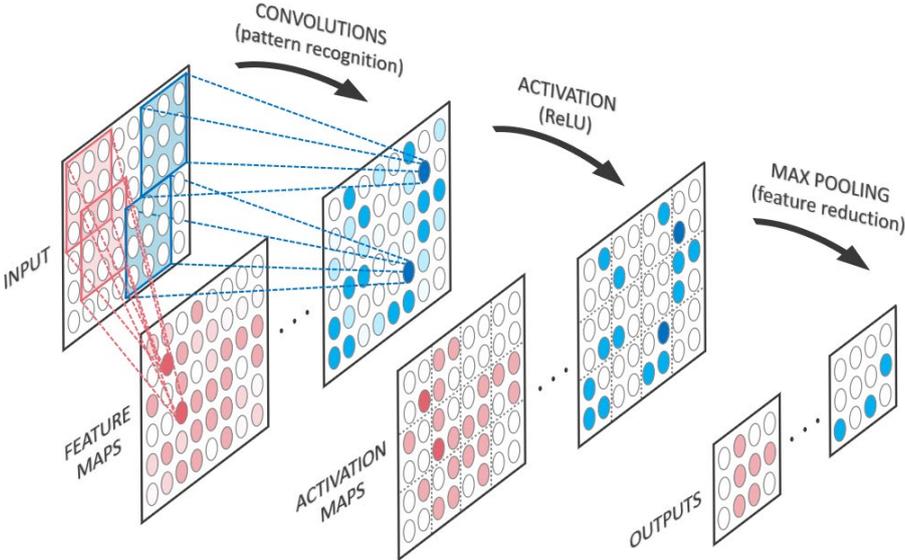
At the core of many recent breakthroughs is the use of Convolutional Neural Networks (CNNs). [5] CNNs are designed to process visual data by learning to recognise spatial hierarchies in images, such as edges, shapes, and textures, making them particularly suited to tasks like tumour detection, organ segmentation, or identifying signs of disease. CNNs have become the foundation for most modern AI

models in radiology and diagnostic imaging. More recently, transformer-based models, originally developed for natural language processing, have begun to show promise in medical imaging by capturing long-range dependencies and global context within images. These models, either standalone or integrated with CNNs in hybrid architectures, are being explored for complex classification tasks and have demonstrated competitive performance in areas such as multi-organ analysis and pathology detection. As AI methods evolve, so do efforts to make their predictions interpretable and trustworthy. A growing field of research in explainable AI (xAI) is focused on helping clinicians understand how and why AI systems reach a given conclusion, which is particularly important in high-stakes domains like diagnosis or treatment planning.

3.2.1. Convolutional Neural Networks

In recent years, CNNs have become the backbone of AI applications in medical imaging, due to their ability to automatically learn hierarchical patterns and features from complex, high-dimensional visual data. Unlike traditional image processing methods that rely on manually designed filters, CNNs learn these representations directly from data, through an end-to-end training process. This capability is especially valuable in medical contexts, where subtle variations in tissue texture, shape, or contrast can indicate the presence of pathology. CNNs operate in a way that mirrors the human visual system. Early layers focus on low-level features such as edges or textures, while deeper layers progressively encode these into more abstract representations like shapes or anatomical structures. This hierarchical processing allows CNNs to capture spatial patterns at multiple scales, a key advantage when analysing detailed medical images such as MRI, CT, or X-Ray. A Convolutional Neural Network (CNN) is built from repeated convolutional blocks (**Figure 2**), each composed of a convolutional layer, an activation function, and a pooling layer. The convolutional layer applies small, trainable filters that slide across the input image to produce feature maps, capturing salient local patterns such as edges or textures. Each unit in the convolutional layer processes information from a restricted region of the image—the receptive field—enabling the network to learn spatially localised features effectively.

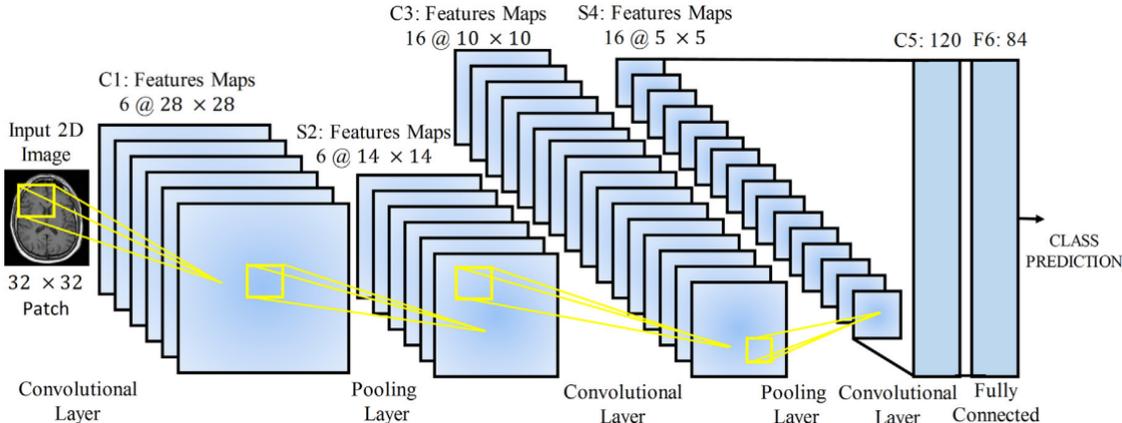
Figure 2: A convolutional block.



Source: own elaboration.

Figure 3 presents a complete CNN architecture for medical image classification, illustrating how input images are processed through multiple layers of convolutions and pooling, before reaching fully connected layers that produce the final prediction. This architecture is now widely used and adapted for various medical imaging tasks such as tumour classification, organ segmentation, anomaly detection, and image registration. Together, these capabilities make CNNs highly effective for automating and enhancing clinical workflows. By detecting subtle patterns that may not be visible to the human eye, CNNs contribute to more precise diagnostics and support clinical decision-making across a wide range of medical applications.

Figure 3: A simple CNN for disease classification from MRI images.

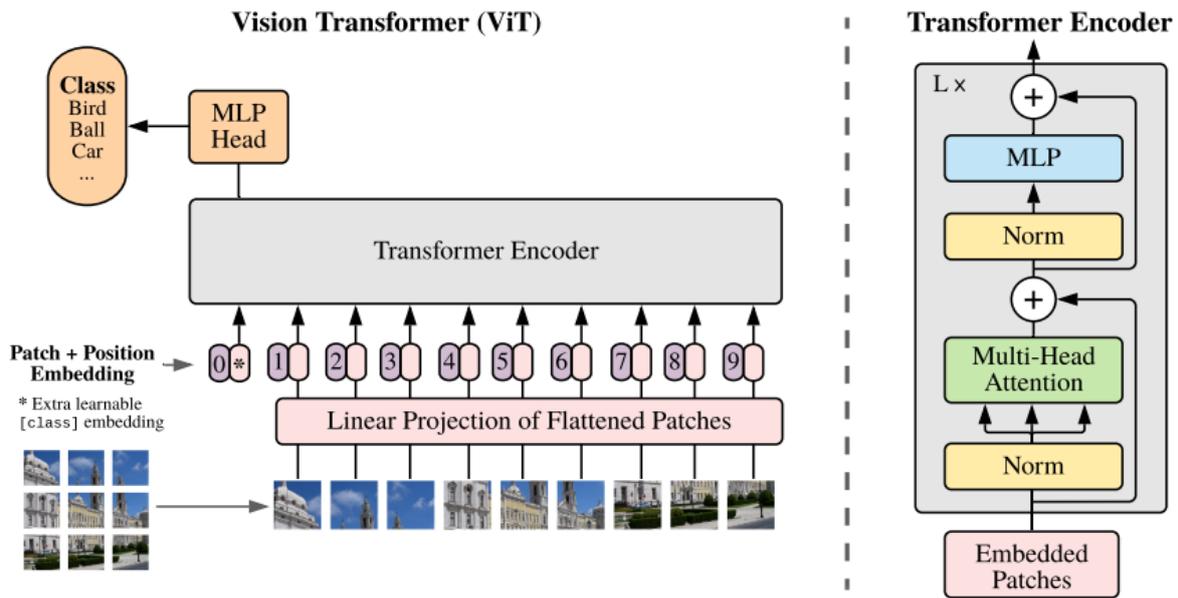


Source: Anwar et al. (2018) [6].

3.2.2. Transformer Networks

The Vision Transformer (ViT) [7] is a DL architecture that adapts the Transformer model, originally developed for natural language processing [8], to image analysis. Unlike CNNs, which process local regions of an image using hierarchical filters, ViT divides an image into fixed-size patches and processes them as a sequence, enabling it to capture global image context through self-attention mechanisms (**Figure 4**). This approach has shown competitive performance to CNNs on vision tasks. In medical imaging, ViT models have been successfully applied to tasks such as tumour classification, organ segmentation, retinal disease detection, and histopathology image analysis.

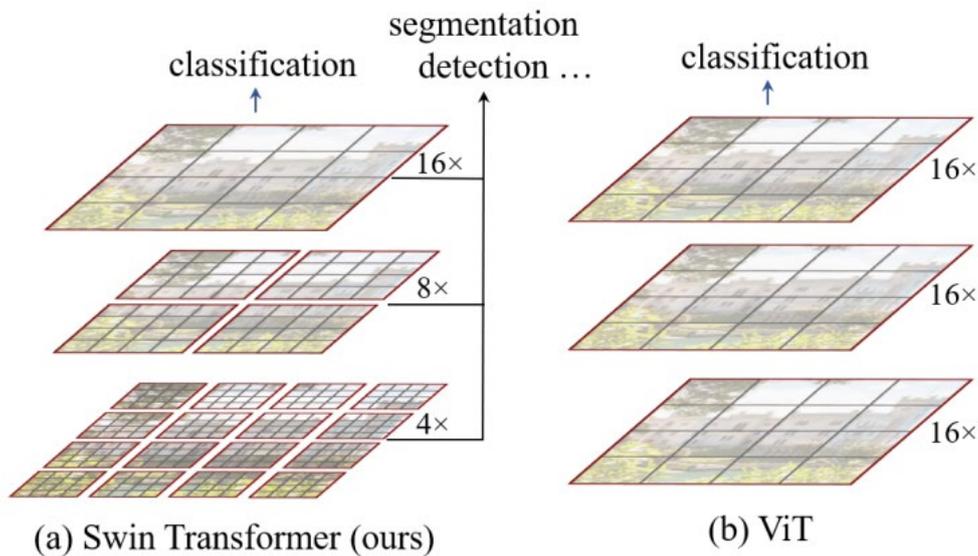
Figure 4: ViT model overview.



Source: Dosovitskiy et al. (2020) [7].

While ViTs have shown strong performance in medical imaging, their computational demands and limited scalability to high-resolution images have encouraged the development of more efficient models. One of the most notable advances is the Shifted Window (Swin) Transformer [9].

Figure 5: Swin Transformer hierarchical strategy versus ViT.



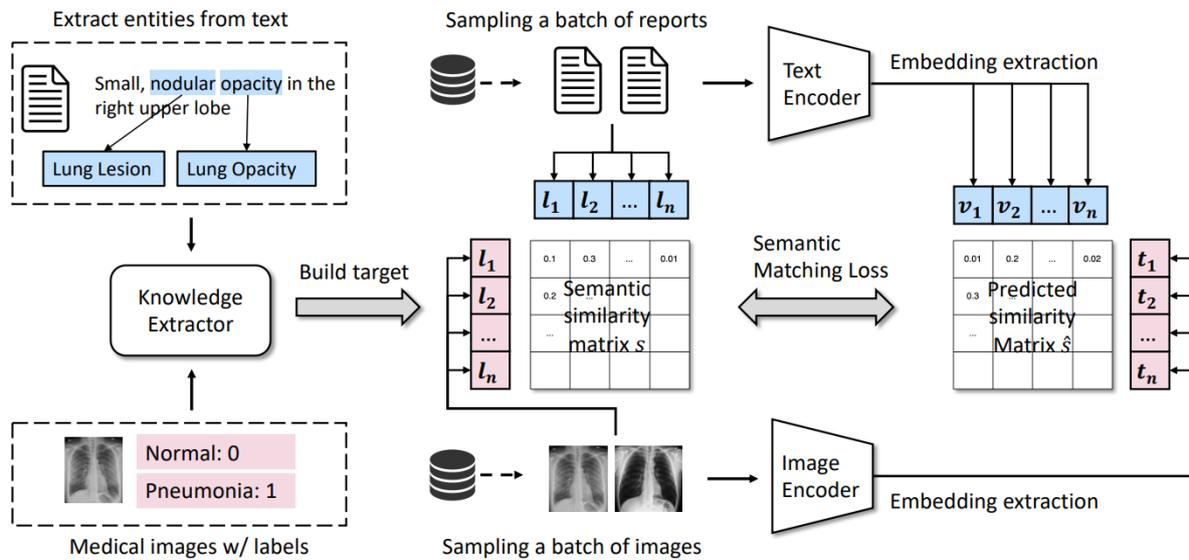
Source: Liu et al. (2022) [9].

Swin Transformer introduces a hierarchical structure with local self-attention in shifted windows, allowing it to capture both fine details and global context more efficiently than ViTs (**Figure 5**). This design makes it especially suitable for medical imaging tasks such as organ segmentation, tumour detection, and 3D image analysis.

3.2.3. Multimodal Models

Multimodal models in healthcare, particularly those integrating medical images and text, represent a promising frontier in clinical AI. These systems aim to replicate the integrative reasoning of human clinicians, who routinely synthesise diverse inputs such as radiology scans, clinical notes, and laboratory results to inform diagnosis and treatment. Their rapid development has been propelled by advances in large vision-language models (VLMs) and multimodal large language models (M-LLMs), which are capable of processing and reasoning across varied data modalities. By combining visual and textual information, multimodal models can perform complex tasks such as diagnosis, radiology report generation, and clinical question answering. This is especially relevant to medical practice, which is inherently multimodal. Unlike unimodal models that handle only one type of input, these systems provide responses that are contextually grounded in both image content and textual descriptors, enabling more nuanced and accurate decision support [10]. One of the most advanced applications of multimodal AI in healthcare is Medical Visual Question Answering (MedVQA). In this task, a system is given a medical image and a natural language question, for example, “What abnormality is visible on this chest X-ray?”, and it must produce a free-text answer. Recent innovations include the PMC-VQA dataset and the MedViNT model [11], which use over 227,000 question-answer pairs based on nearly 149,000 images from various modalities. These tools go beyond earlier multiple-choice formats by enabling open-ended responses through the alignment of pre-trained vision encoders with large language models via instruction tuning. A new generation of specialised models has emerged to address the challenges of multimodal learning in clinical contexts. MedCLIP [12], for example, builds on the CLIP framework [13] by applying contrastive learning to unpaired radiology images and their corresponding reports (**Figure 6**). To mitigate the problem of false negative pairs, common in medical data due to overlapping visual and textual content, it introduces a semantic similarity loss guided by domain knowledge. This enhancement leads to improved zero-shot performance in both classification and retrieval tasks. Med-Flamingo [14] takes a different approach by extending the OpenFlamingo architecture [15] to enable few-shot learning in medicine. Trained on interleaved sequences of medical images and text from sources like PubMed and medical textbooks, it supports complex reasoning tasks such as rationale generation, clinical decision-making, and open-ended visual question answering.

Figure 6: Overview of the MedCLIP framework.



Source: Wang et al. (2022) [12].

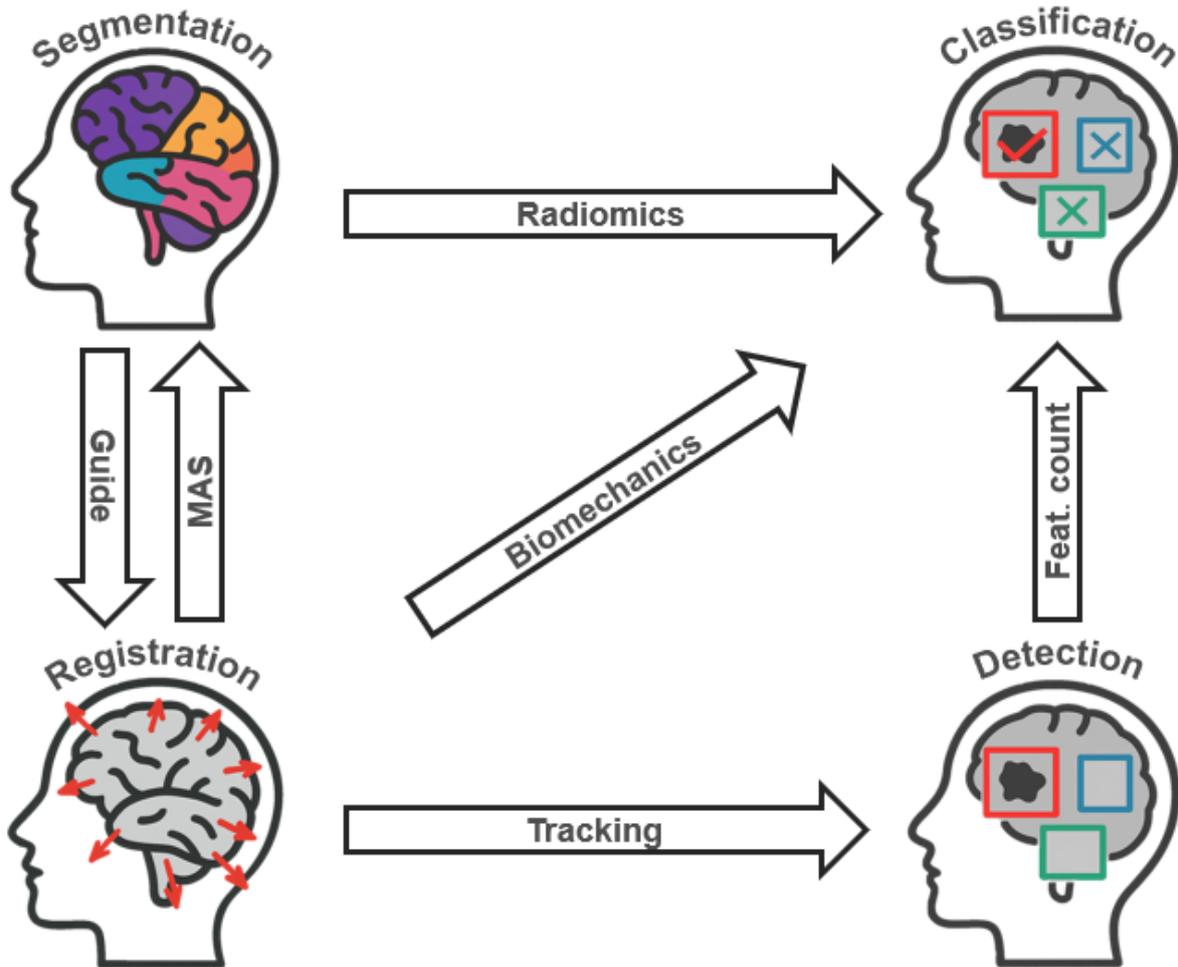
Despite their promise, several challenges hinder widespread deployment. High-quality annotated multimodal datasets are still scarce, which limits generalisability and increases the risk of bias. The interpretability of these models remains a concern, especially in high-stakes environments like clinical care, and ongoing research is exploring ways to make their reasoning processes more transparent. Ethical and regulatory questions also remain a challenge, as these systems must ensure privacy, fairness, and clinical reliability. Lastly, integrating multimodal AI into clinical workflows will require technical interoperability and compliance with health information standards.

3.3. Main Application Areas of AI in Medical Imaging

AI plays a central role across multiple domains in medical imaging, supporting the entire diagnostic and treatment planning pipeline. Segmentation algorithms are widely used to delineate anatomical structures or pathological regions, often serving as a preprocessing step for other tasks. Detection models, typically based on object detection architectures, automatically localise abnormalities such as tumours, lesions, or organ anomalies. Registration methods align images across time, modalities, or patients, enabling longitudinal studies, image fusion, and atlas construction. Classification is used to assign diagnostic labels or risk scores based on features extracted from images and often relies on deep convolutional neural networks. Beyond these core areas, AI is increasingly being used for synthetic data generation, such as creating realistic anatomical variations or augmenting rare pathological cases using generative models.

Figure 7 illustrates the interplay between key AI applications in medical imaging. Segmentation provides radiomics information for classification and can guide registration. Registration, in turn, supports multi-atlas segmentation (MAS) and enables biomechanical modelling that can be used for classification. Detection can also feed into classification through feature count. Together, these interconnected tasks demonstrate the integrated nature of AI workflows in clinical imaging.

Figure 7: Interconnected AI workflows in medical imaging. Segmentation, detection, classification, and registration interact closely in medical imaging pipelines. MAS: multi-atlas segmentation Feat. Count: features count.



Source: Own elaboration.

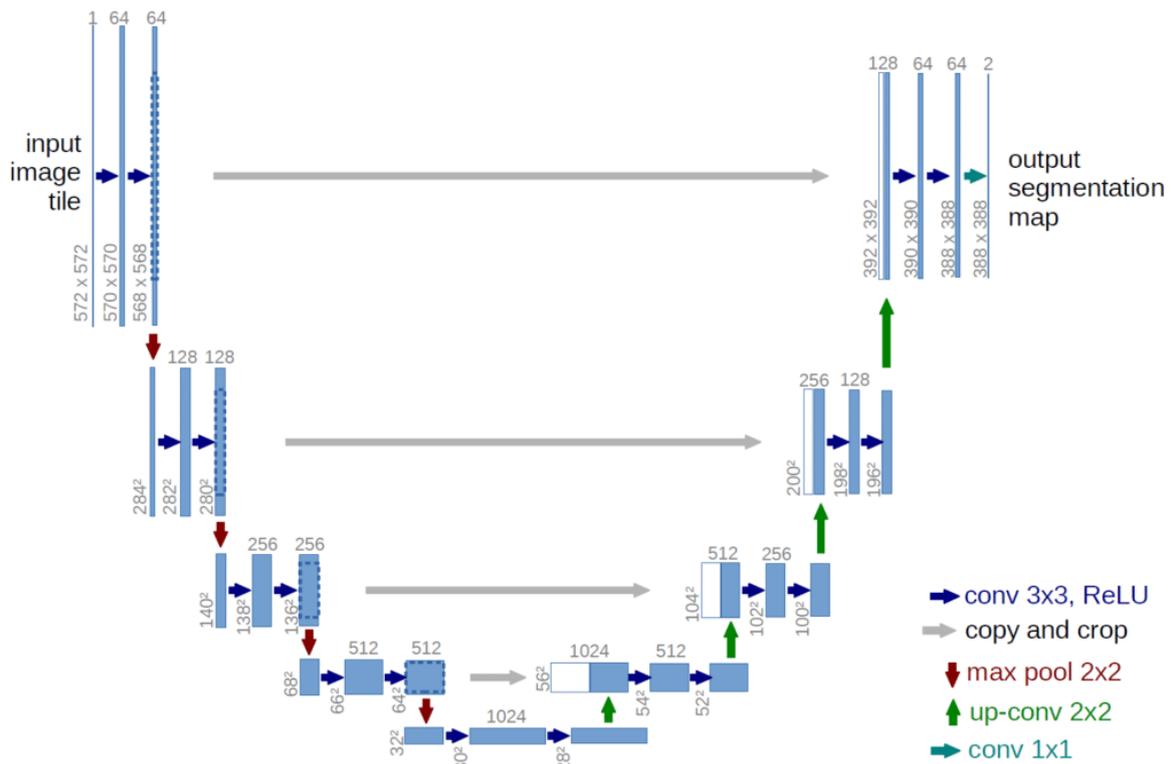
3.3.1. Segmentation

DL has revolutionised medical image segmentation by enabling models to learn complex patterns directly from data, significantly enhancing the accuracy and efficiency of delineating anatomical structures. This advancement is crucial for various clinical applications, including diagnosis, treatment planning, and disease monitoring. Two prominent architectures have emerged in DL-based medical image segmentation: CNNs and ViTs.

3.3.1.1. CNNs in Medical Image Segmentation

One of the most influential CNN architectures in medical image segmentation is the U-Net [16] (**Figure 8**). It features a symmetric encoder-decoder structure with skip connections that bridge corresponding layers of the encoder and decoder paths. The encoder captures contextual information through successive convolutional and pooling layers, while the decoder reconstructs the segmentation map using upsampling operations. Skip connections facilitate the transfer of fine-grained spatial information, enhancing the precision of segmenting intricate structures.

Figure 8: U-Net architecture.



Source: Ronneberger et al. (2015) [16].

The success of U-Net has led to the development of numerous variants and extensions tailored to different segmentation tasks, including adaptations for 3D data [17], improved training stability, and integration with residual connections [18]. Among these, nnU-Net (No New U-Net) [19] stands out for its robustness and adaptability. Rather than proposing a new architecture, nnU-Net provides a fully automated pipeline that configures and optimises a standard U-Net model based on the characteristics of each dataset, making it one of the most widely used tools for medical image segmentation today.

3.3.1.2. Vision Transformers (ViTs) in Medical Image Segmentation

While CNNs remain the dominant architecture in medical image analysis due to their strong inductive biases and ability to capture local features, their reliance on localised operations can limit their effectiveness in modelling long-range dependencies within images. This limitation is particularly evident when analysing complex structures that span large areas, such as organs with irregular shapes or diffuse pathological features [20].

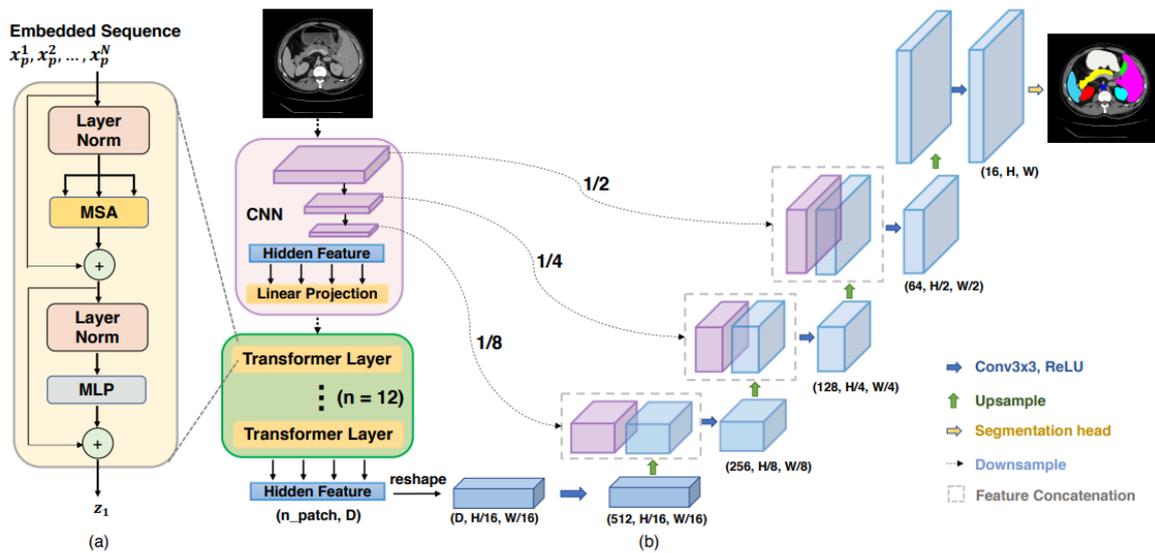
To overcome this, ViTs [7] have been introduced into the medical imaging field. ViTs partition an image into non-overlapping patches, embed these patches linearly, and add positional encodings before processing them through layers of self-attention. This approach enables ViTs to model global context effectively, which is advantageous for tasks requiring holistic understanding of image content. However, ViTs often require large datasets for pretraining, as they lack the inductive biases that make CNNs efficient on small medical datasets.

To combine the strengths of both architectures, hybrid models have been proposed. One prominent example is the Swin Transformer [9], which introduces a hierarchical architecture using shifted

windows. This design enables both local attention within patches and global interactions across them, offering a balance between computational efficiency and contextual modelling.

Another successful hybrid approach is TransUNet [21] (**Figure 9**), which integrates transformer modules into the encoder of a U-Net architecture. CNN layers extract fine-grained local features, while transformers capture long-range dependencies from tokenised feature maps. The decoder then combines these multiscale representations using skip connections, improving both spatial precision and contextual understanding. TransUNet has achieved state-of-the-art performance on several medical image analysis benchmarks by leveraging the complementary advantages of CNNs and transformers.

Figure 9: TransUNet architecture.



Source: Chen et al. (2021) [21].

In summary, while CNN-based models such as U-Net continue to be effective for many medical image segmentation tasks due to their efficiency and ability to capture local features, transformer-based and hybrid models provide powerful tools for incorporating broader contextual information. The choice of architecture often depends on the specific requirements of the task, the size and diversity of the dataset, and available computational resources.

3.3.2. Detection

Object detection in medical imaging focuses on locating and classifying abnormalities, such as tumours, lesions, and nodules; by predicting bounding boxes around regions of interest. This task is crucial for applications like early cancer screening, lesion monitoring, and surgical planning, where precision in localisation directly impacts clinical outcomes.

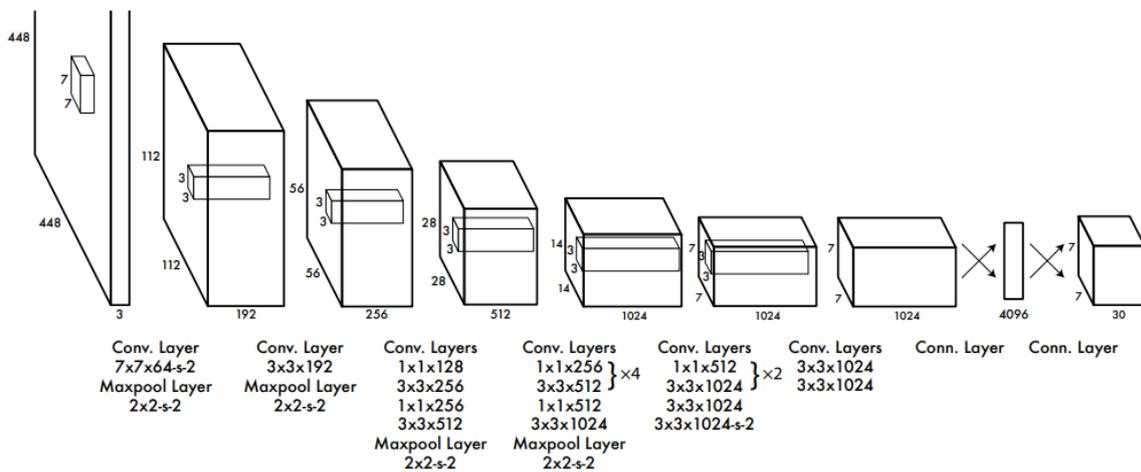
Modern object detection approaches using deep learning are typically divided into two-stage and single-stage detectors. Two-stage detectors first generate region proposals and then classify these regions, while single-stage detectors perform both tasks in a single, unified step.

Among two-stage methods, Region Based CNN [22] (R-CNN) introduced the idea of applying CNNs to object proposals generated by selective search. Fast R-CNN [23] improved efficiency by processing the entire image with a CNN and then pooling features for each proposal. Faster R-CNN [24] further

advanced this framework by integrating a Region Proposal Network (RPN) directly into the model, enabling end-to-end training. These detectors are widely used in medical contexts where accuracy is prioritised, such as tumour or polyp detection. Enhancements like Feature Pyramid Networks [25] (FPN) were developed to improve detection across multiple object scales. FPNs use a top-down architecture with lateral connections to produce semantically strong features at all scales, which is particularly useful in medical imaging where lesions can vary greatly in size.

Single-stage detectors offer faster inference by eliminating the region proposal step. YOLO [26] (You Look Only Once, **Figure 10**) reframes detection as a regression problem, predicting bounding boxes and class labels directly from images. Later iterations like YOLOv3 and YOLOv4 increased accuracy through deeper backbones and multi-scale prediction heads. Similarly, Single Shot Multibox Detector [27] (SSD) predicts objects from multiple feature maps, allowing efficient detection of objects at different scales.

Figure 10: YOLO architecture.



Source: Redmon et al. (2016) [26].

RetinaNet [28] is a single-stage detector designed to help address the class imbalance issue often found in dense object detection. In medical imaging, where the background often dominates the image and positive samples (e.g., tumours) are sparse, this imbalance leads to suboptimal training. RetinaNet addresses this using a novel focal loss function, which down-weights well-classified examples and focuses learning on hard, misclassified ones. Its effectiveness has made it a popular choice for tasks like lesion detection in CT or MRI where both speed and accuracy are critical.

These detection methods have been successfully applied to a wide range of imaging modalities and clinical problems. For example, Faster R-CNN has been used for lymph node detection in MRI [29], SSD for polyp detection in gastrointestinal endoscopic images [30], and RetinaNet for lung nodule detection in lung CT [31]. However, detection models still face challenges in clinical deployment due to limited annotated data, domain shifts across scanners or institutions, and the need for explainable outputs.

3.3.3. Registration

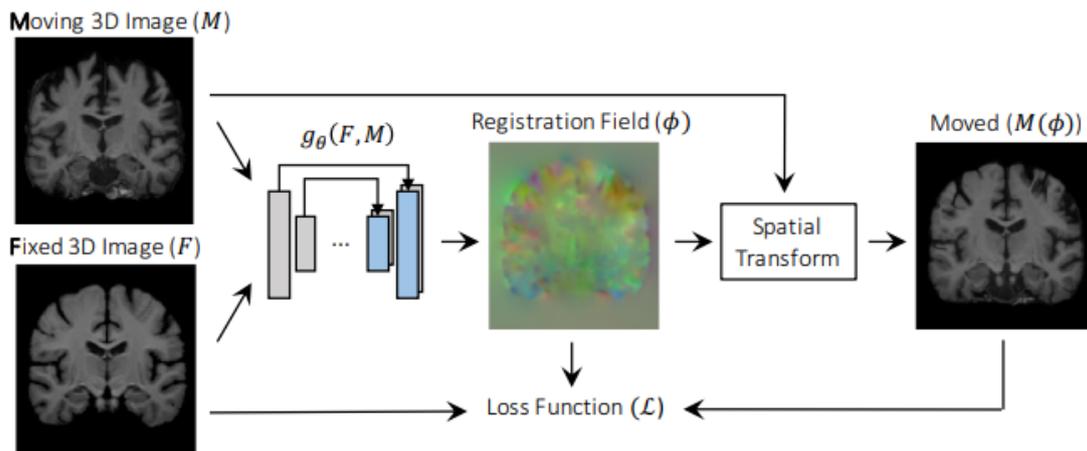
Medical image registration is a fundamental step in many clinical workflows, enabling the alignment of images acquired at different times, from different patients, or across imaging modalities. This process supports tasks such as longitudinal disease monitoring, multi-modal image fusion, image-

guided surgery, and radiotherapy planning [32]. Traditional registration methods rely on iterative optimisation of similarity metrics and transformation models. While effective, they are computationally demanding and sensitive to local minima, particularly in the case of large deformations or multi-modal images [33].

Deep learning has recently emerged as a powerful alternative for medical image registration, offering fast, data-driven models that can predict spatial transformations in a single forward pass. This paradigm shift allows for near real-time performance and has been shown to improve robustness and accuracy in various clinical settings. Unlike conventional methods that require careful hand-crafted features and optimisation loops, deep learning approaches can learn complex deformation patterns directly from image data [32].

One of the earliest and most widely adopted frameworks in this field is VoxelMorph [34] (**Figure 11**), which uses a U-Net-like architecture [16] with an additional spatial transformation layer [35] to estimate deformation fields between image pairs. Trained in an unsupervised manner, VoxelMorph minimises image similarity losses and deformation smoothness constraints. Variants of VoxelMorph and other encoder-decoder architectures have been adapted to different applications, including intra- and inter-patient registration, as well as multi-modal and multi-resolution tasks.

Figure 11: VoxelMorph architecture.

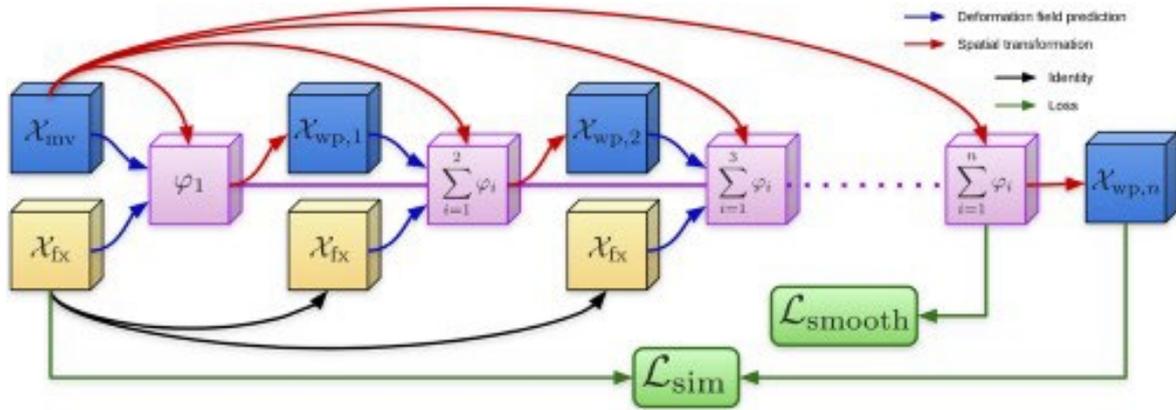


Source: Balakrishnan et al. (2019) [36].

Recent developments have introduced attention-based models and transformers to better capture long-range spatial dependencies. For example, TransMorph [37] replaces convolutional components with Swin Transformer blocks [9], improving performance on complex 3D registration problems such as brain MRI. While these models offer advantages in capturing global context, they remain computationally intensive.

A significant advancement in deep learning-based registration is the cascaded architecture proposed by Comte et al. (2025) [36]. Their model decomposes the deformation into a sequence of incremental transformations; each predicted at a different spatial resolution (**Figure 12**). Unlike earlier cascaded methods, it accumulates deformation fields rather than applying repeated warping, preserving image quality. A multi-resolution similarity loss ensures accurate alignment at both global and local scales. This approach achieves state-of-the-art results on brain MRI, outperforming methods like VoxelMorph and TransMorph in both accuracy and deformation smoothness.

Figure 12: Deep Cascaded Registration Framework.



Source: Comte et al. (2025) [36]

Deep learning-based registration has been successfully applied across a variety of anatomical regions and imaging modalities. In brain imaging, it supports applications such as atlas construction [38], population studies [39], or longitudinal monitoring of pulmonary disease progression [40]. For multi-modal registration tasks like aligning MRI with CT or PET, deep learning approaches improve robustness by learning modality-invariant features, outperforming traditional intensity-based methods [41].

Despite significant progress, medical image registration continues to face important challenges. Domain shifts across institutions, including differences in scanners, imaging protocols, and patient populations, often lead to reduced model performance when applied outside the original training context. Moreover, the absence of ground-truth deformation fields makes validation inherently difficult. Standard metrics such as landmark error or anatomical overlap provide limited insight into whether estimated deformations are anatomically plausible. One common issue is voxel folding, which can be partially assessed by measuring the proportion of voxels with a non-positive Jacobian determinant. While this metric is widely reported, it does not fully capture the anatomical or clinical validity of deformation fields. To address these limitations, recent work is incorporating anatomical priors and biomechanical constraints to guide models toward more plausible outputs.

3.3.4. Classification

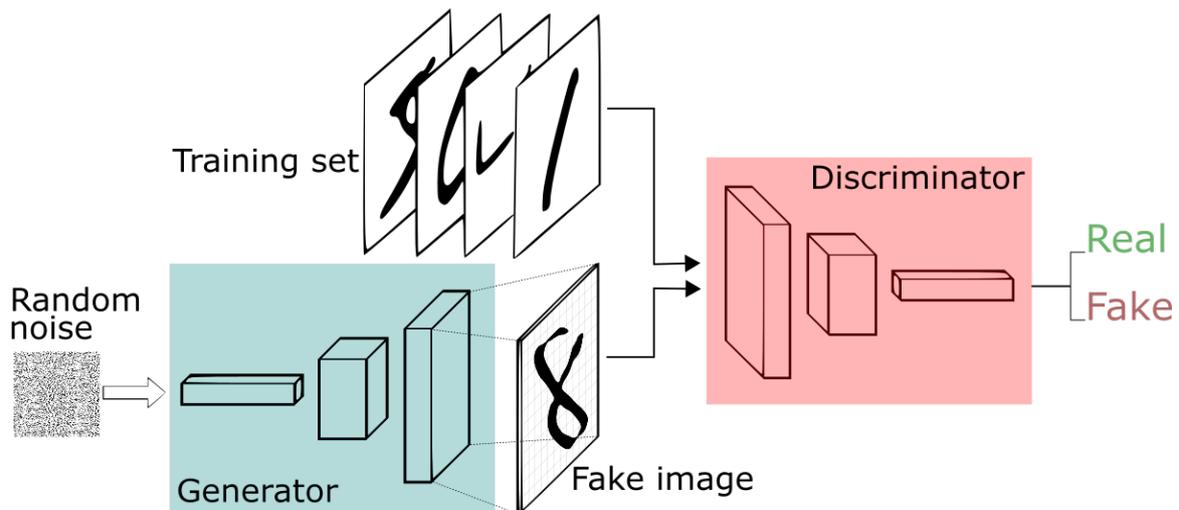
AI, particularly DL, has revolutionised medical image classification by improving diagnostic accuracy, efficiency, and scalability. CNNs have been central to this shift, offering powerful tools for analysing complex medical images. Innovations in CNN architecture, such as transformer hybrids [42], advanced activation functions, and ensemble methods, have enabled increasingly robust classification models across various imaging modalities [43]. However, these models still face challenges related to interpretability and limited applicability in low-resource clinical settings.

A major obstacle to the widespread adoption of AI in clinical workflows is the lack of transparency in decision-making. The black-box nature of deep learning models raises concerns among medical practitioners. To address this, xAI techniques have gained traction. Methods like saliency maps, concept-based explanations, and inherently interpretable models aim to make classification decisions more understandable and trustworthy [44]. This is particularly important in high-stakes domains like oncology, where AI models are increasingly used for diagnosis, subtype classification, and prognostic assessment. Studies have shown that when supported by rigorous feature selection and ensemble strategies, these models can effectively complement human expertise [45].

In parallel, substantial progress is being made to reduce the need for large annotated datasets. Transfer learning allows models trained on large image corpora to be fine-tuned on smaller medical datasets with strong performance. Self-supervised learning (SSL) is also gaining ground as a way to extract informative representations from unlabelled data, enabling better generalisation across clinical settings and imaging protocols [46].

Generative Adversarial Networks [46] (GANs, **Figure 13**), though primarily known for image synthesis, have also played a valuable role in classification. By generating synthetic examples and augmenting limited datasets, GANs help mitigate class imbalance and enhance the robustness of classification models, especially in rare disease scenarios [47].

Figure 13: Schematic representation of a Generative Adversarial Network (GAN).



Source: <https://sthalles.github.io/intro-to-gans/>

Despite these advances, concerns about reproducibility, methodological consistency, and clinical validation remain. A meta-analysis of diagnostic accuracy studies highlights significant variability in study design and reporting, calling for standardised evaluation protocols to assess real-world performance [48].

3.3.5. Generative AI in Medical Imaging

Generative AI has become a powerful paradigm in medical imaging, enabling new applications that range from data augmentation to cross-modality synthesis, disease modelling, and reconstruction. By learning data distributions and anatomical priors, generative models can simulate clinically plausible variations of medical images, enhance diagnostic accuracy, and reduce the dependency on large labelled datasets. The recent shift from classical GANs to more stable and expressive architectures such as diffusion models and transformers is further expanding the boundaries of what is possible in this domain.

3.3.5.1. Image Enhancement and Reconstruction

Generative models are increasingly used for image enhancement tasks, such as denoising [49], super-resolution [50], and image reconstruction [51] from undersampled or corrupted scans. These applications are particularly impactful in MRI and CT imaging, where scan time and radiation dose are critical concerns. Models trained to predict high-fidelity images from low-quality inputs can support faster acquisitions and lower radiation exposure. For example, latent diffusion models like those

implemented in the MONAI Generative Models platform have demonstrated improved reconstruction accuracy and stability across modalities [52]. By embedding anatomical priors into the generative process, these models ensure that reconstructed outputs are both realistic and clinically meaningful. This capability is especially valuable in low-resource settings where imaging infrastructure may be limited or where image quality is compromised.

3.3.5.2. Disease Progression Modelling

Generative models can simulate the natural evolution of disease over time by generating hypothetical follow-up images under different scenarios. This approach has potential in personalised medicine, enabling visualisation of disease progression or response to treatment. Models can be conditioned on clinical metadata, such as patient demographics or pathology status, to simulate patient-specific outcomes. These techniques have been applied to predict tumour growth, simulate postoperative scans, or evaluate the effect of treatment protocols. By generating counterfactual image sequences, clinicians gain a visual aid to support decision-making and assess the efficacy of interventions.

3.3.5.3. Synthetic Medical Image Generation

Synthetic image generation has become an essential tool to address the lack of annotated medical data, enabling data augmentation, anonymisation, and support for rare disease modelling. GANs are the most widely used models in this field, capable of producing realistic, high-resolution images across various modalities. For instance, Guibas et al. (2017) [53] proposed a dual-GAN architecture that first generates anatomical masks and then maps them to realistic images, promoting diversity and photorealism. Comparative studies have shown that advanced architectures like StyleGAN [54] and SPADE-GAN [55] can achieve high visual fidelity and improve downstream tasks such as segmentation [56]. Newer approaches based on diffusion models offer improved training stability and diversity. Pan et al. (2023) [57] introduced a Swin Transformer-based diffusion model that outperformed traditional GANs across several modalities. Similarly, hybrid models like ResViT combine CNNs and transformers to leverage both local and global features, showing strong performance in Cross-modal synthesis [58], which involves generating images in one modality (e.g., CT) based on another (e.g., MRI), thereby enabling multimodal integration and improving accessibility. For instance, GAN-based frameworks have been developed to synthesise PET images from MRI, facilitating hybrid imaging in settings where dual-modality scanners are not available. Such models are also used to simulate contrast-enhanced scans from non-contrast ones, reducing patient exposure to contrast agents. Hybrid approaches combining GANs with transformers, like ResViT [58], allow spatially informed synthesis that preserves fine-grained structures and cross-modal correlations. These synthetic modalities can be used to augment training datasets, enable multimodal segmentation, or serve as input to downstream diagnostic algorithms. Although synthetic data cannot yet replace real datasets, it has proven valuable for training and validating AI models, particularly when combined with real data. Ongoing research is improving the quality, diversity, and clinical utility of synthetic medical images.

4. Barriers and Drivers

The uptake of AI in medical imaging is shaped by both accelerating innovation and persistent practical challenges. Technological advances in deep learning, generative modelling, and multimodal integration are enabling new tools that support earlier diagnosis, improve imaging workflows, and reduce radiologist workload. These developments are reinforced by rising imaging volumes, clinician shortages, and systemic pressures on healthcare systems, which increase the demand for automation and decision support. Collaboration between research institutions, healthcare providers, and industry is driving momentum, and an increasing number of regulatory clearances are building confidence in clinical deployment. These trends are consistent with findings from the JRC's recent *Science for Policy* report *IMAGING THE FUTURE*, which highlights emerging opportunities and structural bottlenecks in medical imaging and AI.

Yet, several barriers continue to slow adoption. A core limitation is access to annotated, high-quality medical imaging data. Most clinically relevant AI models rely on supervised learning, which requires expert-labelled datasets. However, annotation is particularly time-consuming for 3D or longitudinal data and typically must be performed by medical imaging specialists. Since annotation is rarely part of routine clinical practice, this task is often relegated to research settings with limited scalability. In parallel, data silos, privacy constraints, and limited interoperability across institutions hinder data sharing and reuse. Interpretability remains a challenge for many AI models, reducing trust and limiting acceptance in clinical practice. While regulatory frameworks are evolving, uncertainties about validation requirements, liability, and cross-border compliance contribute to hesitancy. Integration into clinical workflows also depends on adequate technical infrastructure, clinician training, and cultural adaptation, all of which require resources and time. Finally, issues such as algorithmic bias and unequal access to AI solutions reinforce the need for inclusive design, diverse datasets, and transparent governance.

Addressing these challenges will require continued coordination among AI developers, clinicians, regulators, and patients to ensure that AI systems are not only technically robust but also clinically meaningful, ethically sound, and widely accessible

4.1. Infrastructure and Interoperability

The effective deployment of AI in medical imaging does not rely solely on model performance; it critically depends on the availability of robust and interoperable infrastructures that support large-scale, privacy-preserving, and ethically governed data access. As medical imaging generates high volumes of sensitive data, often fragmented across institutions and countries, the need for scalable, secure, and standardised infrastructures is central to both research and clinical applications. A major barrier is data fragmentation. Imaging data are typically confined to local Picture Archiving and Communication Systems (PACS) and electronic health records (EHRs), which differ widely in format, structure, and access policies. This fragmentation obstructs the creation of large, diverse, and harmonised datasets required to train and validate generalisable AI models. As highlighted by Kondylakis et al. (2023) [59], several European projects have made significant strides to address this through the development of data-sharing infrastructures within the AI4HI network, such as CHAIMELEON, EuCanImage, INCISIVE, ProCAncer-I, PRIMAGE, and more recently EUCAIM. These projects employ various data federation strategies, ranging from centralised cloud repositories to decentralised and hybrid approaches, all aimed at facilitating cross-institutional AI development. Common to all is the emphasis on interoperability, achieved through the adoption of international

standards such as DICOM⁴² for imaging data and HL7 FHIR⁴³, SNOMED CT⁴⁴, LOINC⁴⁵, and OMOP-CDM⁴⁶ for clinical and metadata. For example, CHAIMELEON and ProCancer-I use OMOP-CDM with oncology extensions to ensure semantic consistency in structured clinical variables, while EuCanImage and PRIMAGE rely on MIABIS⁴⁷ extensions to standardise metadata across biobanks. These efforts in semantic harmonisation, quality control, anonymisation, and secure access have laid critical groundwork for EUCAIM, which builds directly on the experience and tools developed by these earlier projects to create a federated, pan-European cancer imaging platform that is FAIR (Findable, Accessible, Interoperable, and Reusable) by design. The goal of EUCAIM is to provide a common data space for cancer images, enabling researchers and clinicians to securely access, annotate, and analyse data across borders while ensuring compliance with European regulations.

4.2. Privacy Preserving Learning on Medical Imaging

Protecting patient privacy while harnessing the full potential of AI in medical imaging is a critical challenge. Medical imaging datasets are often sensitive and subject to strict regulatory frameworks such as the GDPR. Several privacy-preserving techniques have emerged to enable model development without compromising personal data integrity. Below, we explore the most prominent approaches, drawing from the current literature.

4.2.1. Federated Learning

Federated learning (FL) is one of the most promising approaches to decentralised model training. Rather than pooling data centrally, FL allows institutions to collaboratively train AI models by sharing only model updates or gradients, not the data itself. This paradigm reduces the risk of data breaches and keeps raw data within institutional firewalls. In medical imaging, FL has demonstrated robust performance across heterogeneous datasets. For example, Sheller et al. (2020) [60] applied FL to brain tumour segmentation and showed comparable performance to centrally trained models, while maintaining data privacy. Similarly, Kumar et al. (2021) [61] used FL for COVID-19 diagnosis from chest CT scans across multiple hospitals. However, FL is not without challenges. Differences in imaging protocols across institutions (non-IID data) and the risk of information leakage from model updates still pose open research problems.

4.2.2. De-identification and Anonymisation

De-identification, often the first step in privacy preservation, involves removing direct identifiers (e.g., name, ID number) from imaging metadata. Advanced anonymisation also addresses facial features in head and neck scans, which can be reconstructed into identifiable 3D models. However, standard de-identification methods are increasingly seen as insufficient in the face of linkage attacks or re-identification via auxiliary data sources [62].

⁴² [DICOM](#)

⁴³ [Index - FHIR v5.0.0](#)

⁴⁴ [What is SNOMED CT | SNOMED International](#)

⁴⁵ [Home – LOINC](#)

⁴⁶ [OMOP-CDM | Documentation du SNDS & SNDS OMOP](#)

⁴⁷ [MIABIS - BBMRI-ERIC](#)

4.2.3. Differential Privacy

Differential privacy (DP) is a mathematical framework that introduces noise to data or model outputs, ensuring that individual data points cannot be distinguished. DP is particularly useful in settings where some level of data sharing is necessary, such as in model evaluation or publication. Kaissis et al. (2021) [63] explored the application of differential privacy in training medical imaging models and observed that, while it enhances privacy guarantees, it can degrade model performance if not carefully calibrated. This trade-off between privacy and utility remains a central concern.

4.2.4. Homomorphic Encryption

Homomorphic encryption (HE) allows computations to be performed directly on encrypted data. This enables a third party (e.g., a cloud service provider) to process sensitive medical images without ever accessing their unencrypted form. Despite its theoretical appeal, HE remains computationally intensive and is rarely used in practical clinical applications [64]. To overcome the computational burden of full homomorphic encryption (HE), Kaissis et al. (2021) [65] proposed a hybrid approach that applies HE only during communication in federated learning. In this framework, only the locally computed model updates are encrypted before being sent to the central server. This allows secure aggregation of encrypted updates without ever exposing individual patient data or raw model parameters. By limiting the use of HE to the communication layer rather than applying it to the entire training process, this method achieves a balance between security and computational feasibility.

4.2.5. Synthetic Data Generation

Synthetic data generation is increasingly explored as a strategy to mitigate privacy concerns in medical imaging. Using generative models, particularly GANs, synthetic images can be created to supplement or even replace real patient data for tasks such as training, validation, and algorithm development [66]. While synthetic data can help reduce reliance on sensitive datasets and avoid direct exposure of personal health information, its role in preserving privacy remains an open question. If synthetic images closely resemble real data distributions, there is a risk of inadvertently encoding identifiable patient characteristics, particularly when training data is limited or poorly anonymised [68]. Differential privacy mechanisms are sometimes integrated during generation to further reduce this risk, but they also tend to reduce image fidelity. Another critical consideration is regulatory acceptance. Although synthetic data is useful for pre-training or augmenting datasets, regulators currently require validation on real clinical data for safety and efficacy assessments. At present, synthetic data alone is unlikely to be accepted as a substitute for real-world evidence in approval processes, but it can play a valuable supporting role.

4.3. Explainable Artificial Intelligence (xAI) in Medical Imaging

Explainable Artificial Intelligence (xAI) plays a crucial role in the clinical adoption of AI tools for medical imaging, addressing the need for transparency, trust, and accountability in high-stakes healthcare environments. In medical imaging, where AI models often serve as decision-support systems for diagnoses, treatment planning, and monitoring, the ability to understand and communicate how these systems arrive at their predictions is fundamental to clinical validation and user acceptance.

From a technical standpoint, xAI techniques can be broadly categorised into *post-hoc* and *intrinsic* methods. Post-hoc approaches are applied after model training and include techniques such as saliency maps, Grad-CAM, SHAP (SHapley Additive exPlanations), and LIME (Local Interpretable Model-agnostic Explanations). These methods help visualise or approximate how input features—like image

pixels or clinical attributes—contribute to a model’s output. For example, Lundberg et al. (2020) [67] introduced TreeExplainer to provide exact SHapley values for tree-based models, demonstrating how many local explanations can be aggregated to form a globally interpretable structure that is faithful to the original model.

However, post-hoc methods often face criticism for their potential lack of faithfulness, as they do not directly influence model training. This raises concerns about whether the explanations genuinely reflect the model’s internal logic or merely offer plausible narratives. As an alternative, intrinsic methods aim to make models interpretable by design, such as through attention mechanisms, prototype learning, or by enforcing sparse and human-readable representations. Yet, these often come at the cost of reduced predictive performance, particularly for complex image-based tasks.

Furthermore, recent European policy initiatives emphasise the need for trustworthy and transparent AI systems, particularly in high-stakes domains like healthcare. While the European AI Act does not explicitly mandate explainable AI, it includes requirements for transparency, human oversight, and risk management in high-risk applications, which include many medical AI tools. These provisions contribute to a broader policy environment that encourages the development of interpretable and auditable AI systems. In this context, xAI techniques are increasingly explored as practical means to meet these expectations and support clinical accountability.

As AI becomes increasingly embedded in radiology and digital pathology, explainability is not just a technical challenge but a sociotechnical imperative. Clinicians must be able to trust and understand AI recommendations, especially in edge cases or when results conflict with their clinical judgment. To this end, xAI will remain a key enabler of responsible, trustworthy, and clinically viable AI in medical imaging.

5. Selected Use Cases

Artificial intelligence is transforming medical imaging by enabling more accurate, efficient, and consistent diagnostic support across diverse clinical contexts. This section presents two practical use cases developed by JRC.F7 that demonstrate how trustworthy AI systems can be designed to support diagnosis, monitoring, and decision-making in real-world healthcare settings. The first use case addresses lung cancer screening and follow-up, a task of high clinical relevance given the disease's prevalence and mortality. The second one focuses on CVD classification using a biomechanics-informed approach. In the lung imaging use case, the AI pipeline covers the entire clinical workflow—from anatomical segmentation and nodule detection to longitudinal analysis and calibrated uncertainty estimation. The system is supported by a secure, web-based platform for data anonymisation and upload, AI processing, visualisation and validation. The objective is to reduce radiologist workload while ensuring outputs remain interpretable, robust, and aligned with clinical needs. In the cardiovascular domain, the second use case introduces a biomechanics-driven neural network applied to cine-cardiac MRI sequences. The approach leverages DL-based registration enhanced with Neo-Hookean strain energy regularisation to estimate biomechanically plausible deformations. Derived biomechanical features are then used for accurate and explainable classification of cardiac pathologies. By embedding physical priors into the model architecture, the system provides outputs that are both clinically meaningful and trustworthy.

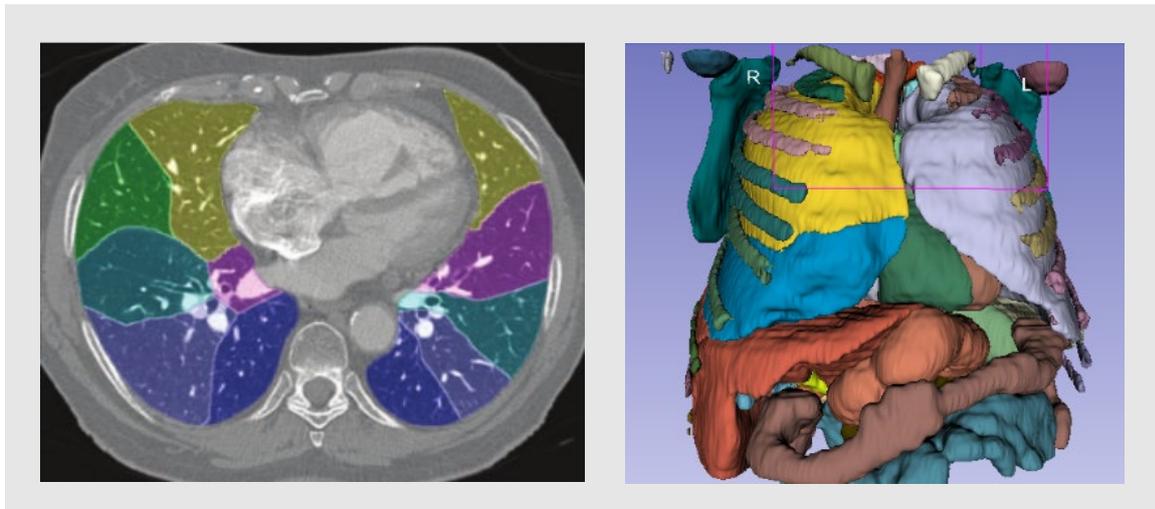
5.1. AI-Driven Diagnosis and Analysis of Lung Lesions

The analysis of lung lesions in CT imaging involves several complex tasks that benefit from AI-driven automation, including the segmentation of anatomical structures, the detection and classification of pulmonary nodules, and the monitoring of their progression over time. Given the heterogeneity of lung diseases and the subtle visual differences between benign and malignant findings, traditional rule-based systems often fall short. AI models offer the potential to address these challenges by learning from large, annotated datasets and generalising across patient populations. This section outlines the technical foundations and implementation of several complementary components: volumetric segmentation using U-Net, lesion identification via CNNs and RetinaNet, longitudinal tracking through Siamese networks and deep registration, uncertainty quantification with conformal risk control, and clinical deployment through a secure, web-based platform. Together, these elements form a robust and scalable framework that supports the clinical management of pulmonary nodules with improved accuracy, consistency, and interpretability.

5.1.1. Chest Anatomy Segmentation

For the segmentation of medical images, which involves identifying individual voxels within the volumetric image that belong to one or more objects of interest, the most commonly used CNN architecture is U-Net [68], which effectively handles the problem of imbalanced classes typical of such tasks (e.g., very few malignant nodules and many benign lesions) and can scan an entire image in a single pass without losing the overall context of the image, unlike previous patch-based approaches. The segmentation of the entire body into 104 anatomical labels was performed using nnDetection [69] (**Figure 14**) for volumetric (3D) segmentation, which was trained with the Totalsegmentator [70] dataset consisting of 1,204 whole-body CT scans.

Figure 14: Left: Axial view of the various segmented lung lobes. Right: 3D view of the segmentation of all organs presents in the CT volume.

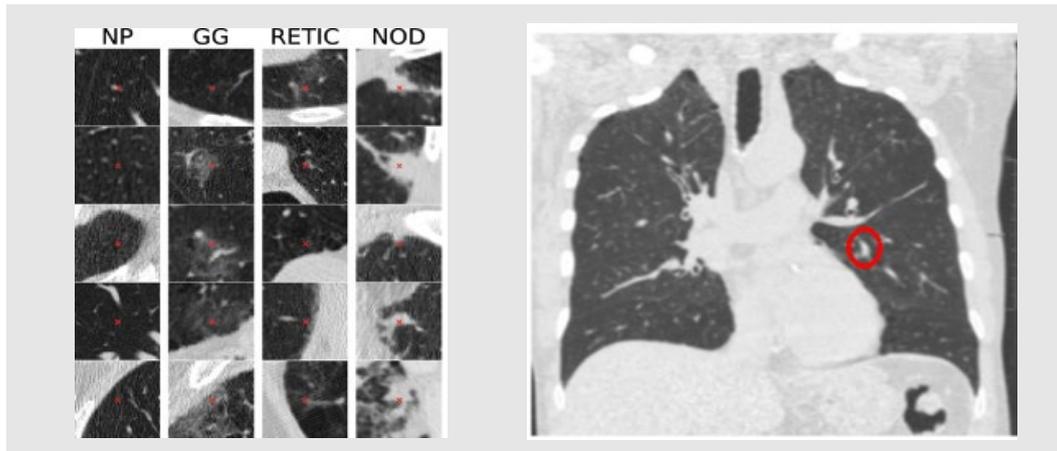


Source: own elaboration.

5.1.2. Lung Nodules Detection

Interstitial lung diseases (ILD) constitute a heterogeneous group of over 200 lung disorders (**Figure 15**) primarily affecting the lung parenchyma but may also involve airways or vascular manifestations. There is increasing recognition that lung tumours and some forms of ILD, particularly idiopathic pulmonary fibrosis (IPF), may be preceded by early radiographic findings observed in chest CT scans. The visual presence of these lung abnormalities is associated with a range of manifestations, such as reduced lung volume and the development of adenocarcinomas. Therefore, their rapid identification and monitoring of evolution over time are of paramount importance for patients. They may present with varied texture and consistency, ranging from easily identifiable solid nodules in a CT scan to "ground-glass" nodules, the recognition of which is much more complex.

Figure 15: Left: Different Lung Anomalies Visible on CT Images - from Left to Right: Normal Parenchyma (NP), Ground Glass (GG), Reticular (RETIC), Solid Nodule (NOD). Right: example of a solid lung nodule.

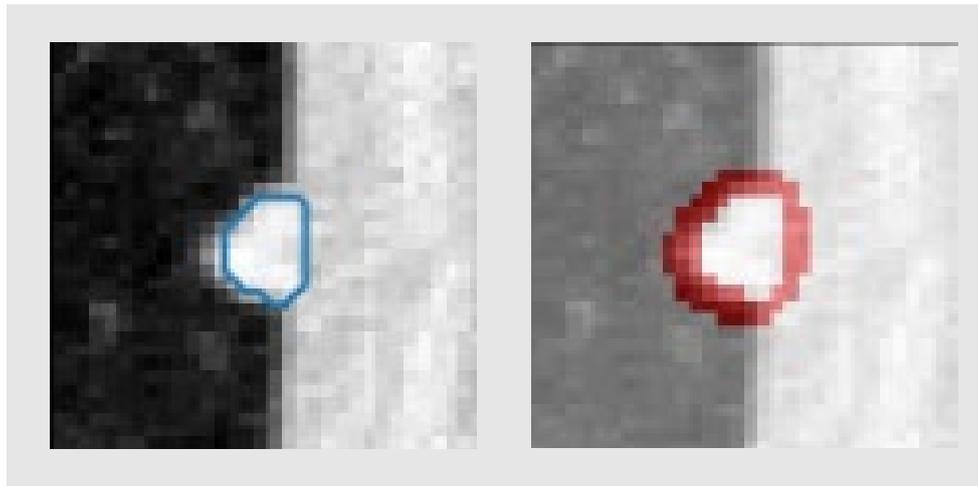


Source: own elaboration.

The field of computer vision has evolved rapidly, offering a wide range of techniques applicable to image classification, object detection, image retrieval, semantic segmentation, and human pose estimation [71]. The application of these techniques to medical scans such as CT, MRI, PET, and US is now a mature field. For image classification, which involves identifying the presence of an object or disease, CNNs and ViT [72] are widely used in image processing. In 2017, Demir et al. (2019) [73] developed the Inception V3 model to classify clinical images related to skin cancer exams into benign and malignant variants. Kang et al. (2017) [74] proposed a method to enhance the performance of 2D CNNs using a multi-view 3D CNN for lung nodule classification incorporating spatial contextual information with the assistance of the 3D Inception-Resnet architecture. Lesion detection aims to identify specific coordinates within a 3D volume, thus involving a combination of information identification and localised classification. In 2016, Hwang and Kim proposed a Self-Transfer Learning (STL) method for detecting nodules in chest radiographs and lesions in mammography [75]. There is broad consensus that successful training of deep networks requires many thousands of annotated training samples. For this project, using the nnU-Net [19] architecture, we were able to rely on synthetic data augmentation techniques to utilise the available annotated samples more efficiently. The architecture consists of a contracting pathway to capture context and a symmetric expanding pathway that allows for precise localisation.

Using images from the public LUNA16 dataset, we trained the network to recognise lung tissue and solid lesions as indicated in **Figure 16**.

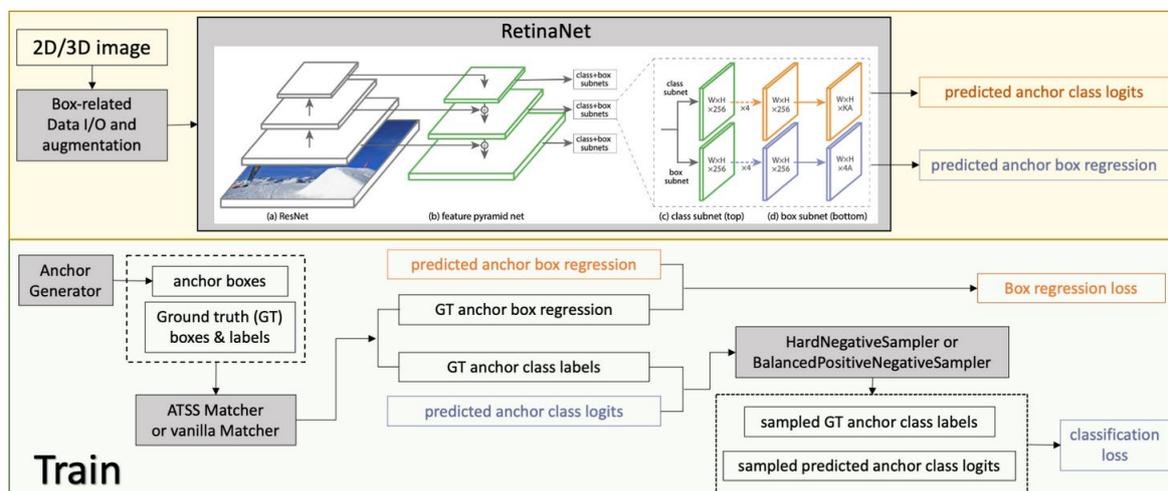
Figure 16: Initial identification of malignant lung lesions (left) and its subsequent re-identification in follow-up analyses (right).



Source: own elaboration.

Additionally, we employed Monai's CT lung nodule detection algorithm, which relies on RetinaNet [31] and is trained on the public LUNA16 dataset, for comparison purposes. The network's overall structure is depicted in **Figure 17**. Through visual examination, we noted a considerable number of false positives located outside the lung area. Consequently, we implemented lung segmentation to automatically filter out these spurious detections.

Figure 17: Main features of the second nodule detection algorithm based on RetinaNet.



Source: <https://monai.io/model-zoo.html>

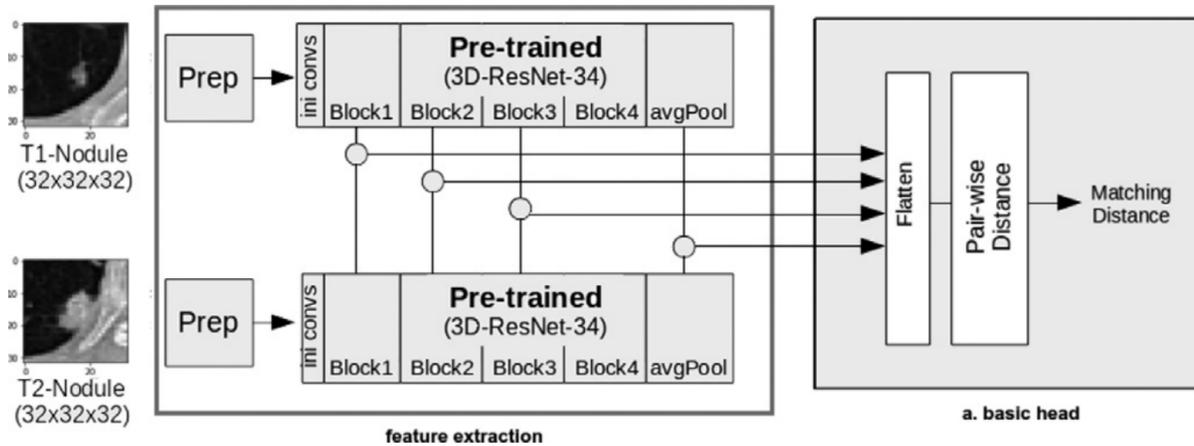
5.1.3. Follow-up of Nodules Over Time

5.1.3.1. Siamese Neural Network

Lung cancer follow-up is a complex task, prone to errors, and requires considerable time for clinical radiologists. Several lung CT scanning images taken at different times for a given patient must be

individually inspected, searching for possible cancerous nodules. Radiologists primarily focus their attention on the size, density, and growth of nodules to assess malignancy. In this study, we employ a novel method based on a 3D Siamese neural network for the re-identification of nodules in a pair of CT scans of the same patient (**Figure 18, Figure 19**).

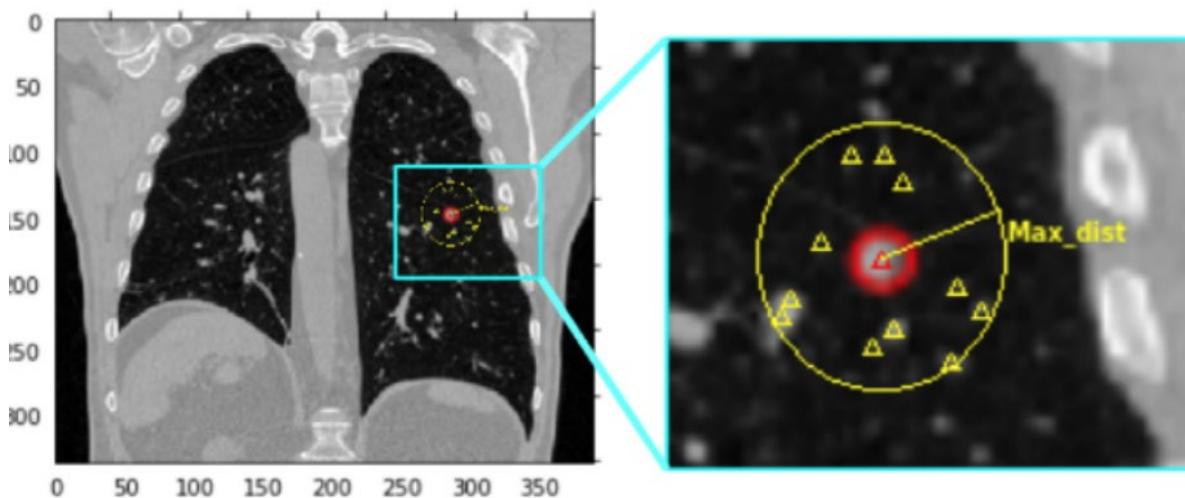
Figure 18: Schema of the initial module for identifying lung lesions in the follow-up analysis.



Source: Rafael et al. (2021) [76].

The network has been integrated into a two-phase automated pipeline to detect, match, and predict nodule growth given pairs of CT scans. Results on an independent test set reported a nodule detection sensitivity of 94.7%, a temporal matching accuracy of nodules of 88.8%, and a growth detection sensitivity of 92.0% with a precision of 88.4%.

Figure 19: Example of how the nodule re-identification algorithm in the follow-up selects candidates.



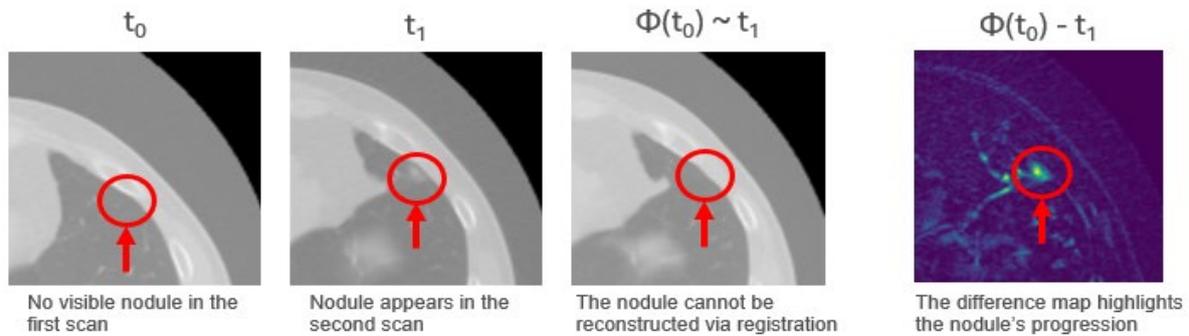
Source: Rafael et al. (2021) [76].

5.1.3.2. Cascaded Registration

Medical image registration can also be effectively used to track the progression of pulmonary nodules over time. By aligning follow-up scans from different time points, registration models, such as the deep learning cascaded framework shown in **Figure 12** enable the comparison of anatomical

structures with high spatial precision. As illustrated in **Figure 20**, when a nodule is not visible in an initial scan but appears in a later scan, image registration can reveal discrepancies between the predicted transformation and the actual follow-up image. The resulting difference map highlights the nodule's appearance and growth, providing a valuable tool for early detection and monitoring progression.

Figure 20: Follow-up of nodule's progression using image registration.



Source: own elaboration.

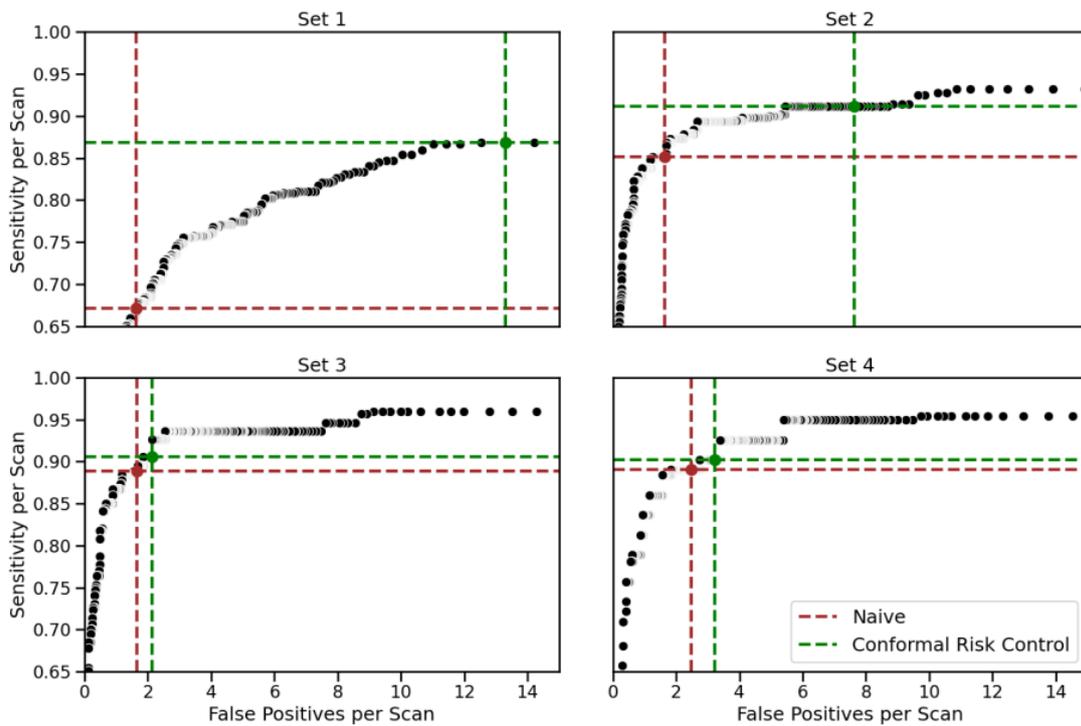
5.1.4. Conformal Risk Control for Pulmonary Nodule Detection

Reliable deployment of AI in clinical settings requires more than strong average performance, it requires guarantees about how the model behaves in individual cases. In safety-critical tasks such as pulmonary nodule detection, where diagnostic errors can be serious and inter-rater variability is common, clinicians need models that are both accurate and transparent about uncertainty.

To meet this need, we employed Conformal Risk Control [77] (CRC), a post hoc calibration method that wraps around any trained prediction model. CRC generates prediction sets that are guaranteed to contain the correct answer with a user-specified probability, ensuring a defined level of sensitivity per scan. This approach shifts the focus from aggregate performance metrics to patient-level reliability, better aligning with clinical decision-making needs.

We applied CRC to a state-of-the-art pulmonary nodule detection model and evaluated it across four datasets (Set 1–4) derived from the LIDC-IDRI collection [78]. As shown in **Figure 20**, CRC consistently achieved or closely approached the target 90% sensitivity per scan, even in the more challenging sets. For example, it reached 91.35% sensitivity at 2.25 false positives per scan among nodules annotated by at least three radiologists. In contrast, a Naive thresholding strategy—which applies a fixed confidence cutoff—underperformed in Sets 1 and 2, failing to meet the sensitivity target due to lack of calibration.

Figure 21: Sensitivity per scan vs. false positives per scan for four LIDC-IDRI subsets. CRC (green) meets the 90% sensitivity target more reliably than the Naive strategy (red), especially in challenging cases.



Source: Hulsman et al. (2025) [81].

These results demonstrate that CRC not only improves sensitivity but also adapts to the difficulty of the data. When nodules are ambiguous or consensus among radiologists is low, CRC automatically produces larger prediction sets to reflect higher uncertainty. This behaviour increases transparency and trust, as clinicians are alerted when the model is uncertain, rather than being misled by overconfident predictions. While CRC introduces more false positives in some cases, especially in noisier datasets, it provides formal guarantees that are critical for safe deployment in healthcare. Its ability to balance reliability and interpretability makes it a promising solution for integrating AI into real-world clinical workflows, where trust and accountability are as important as accuracy. These findings were reported in Hulsman et al. (2025) [81].

5.1.5. Web Platform and Tools for Clinical Imaging Workflows

To enable **large-scale, privacy-preserving clinical research** and facilitate **AI-assisted diagnostics** in lung cancer, existing tools and platforms can support secure data handling, annotation, and visualisation workflows. These tools are crucial for integrating AI models into research and clinical environments in a compliant and user-friendly manner.

5.1.5.1. Anonymisation

Anonymisation is a critical step to ensure compliance with data protection standards such as the GDPR and the forthcoming EHDS regulation. Widely used tools and libraries, such as DCMTK⁴⁸, GDCM⁴⁹, and DICOM-anonymizer⁵⁰, support automated DICOM metadata anonymisation and can be integrated into local clinical environments.

These tools enable healthcare providers to:

- Remove or replace patient identifiers from DICOM headers
- Maintain data utility for AI training and clinical use
- Securely prepare datasets for federated learning or centralised analysis

Common anonymisation workflows typically involve uploading DICOM files, performing metadata stripping or substitution based on configurable rules, and securing data transmission using encrypted protocols.

Visualisation

For image review and annotation, tools like Orthanc, an open-source DICOM server and viewer, can be used to store, manage, and visualise medical imaging data, including lung CT scans, organ segmentations, and nodule detection results. As shown in **Figure 22**, Orthanc supports interactive visualisation of AI outputs, making it a practical choice for integrating automated segmentations and detection overlays in a clinical or research workflow.

Orthanc allows users to:

- View images in multiple planes (axial, coronal, sagittal)
- Overlay AI-generated annotations (e.g., segmentations, nodules)
- Enable expert review and interaction with AI outputs to ensure clinical relevance

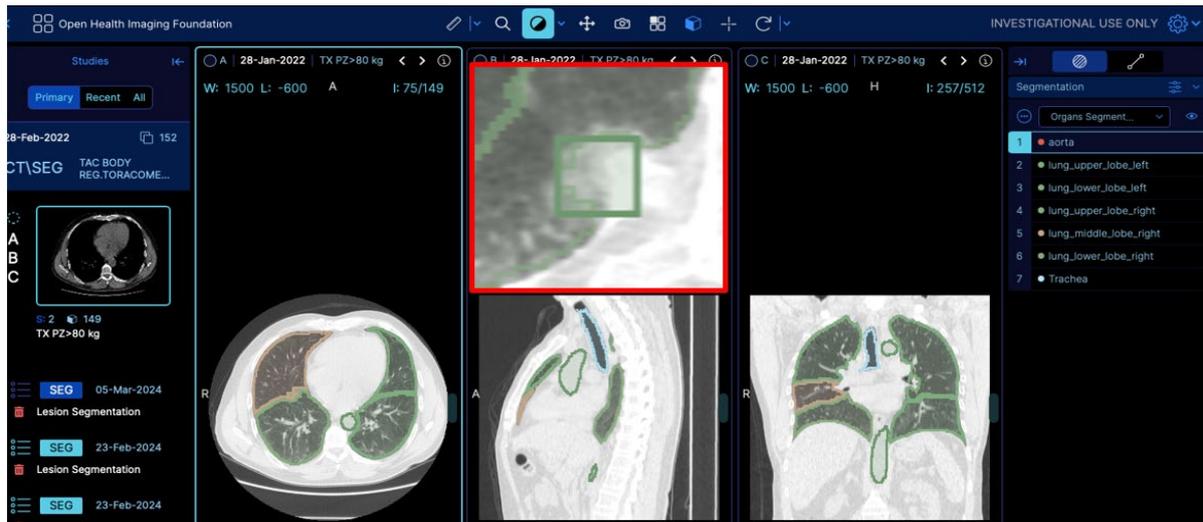
Such tools help bridge the gap between technical AI development and practical clinical use by enabling both transparency and expert oversight.

⁴⁸ <https://github.com/DCMTK/dcmtdk>

⁴⁹ <https://github.com/malaterre/GDCM>

⁵⁰ <https://dicomanonymizer.com/>

Figure 22: Online Viewer.



Source: own elaboration.

5.1.6. Discussion

The development of an AI pipeline for lung lesion detection, longitudinal follow-up, and uncertainty quantification illustrates the transformative potential of artificial intelligence in enhancing early diagnosis and continuous monitoring of lung cancer. By integrating segmentation models, detection networks, Siamese-based matching algorithms, and conformal risk estimation within a unified framework, this approach achieves strong technical performance and demonstrates the feasibility of AI-assisted clinical workflows. However, it also sheds light on the broader systemic challenges that must be addressed to ensure that such tools are trustworthy, generalizable in diverse local environments, and clinically deployable at scale. A key barrier remains the limited availability of large, diverse, and high-quality imaging datasets for training and validation. Clinical data are often siloed within national health systems and hospital PACS, with limited interoperability and variable standards. This fragmentation hampers the development of robust AI models and increases the risk of bias and poor generalisation across populations, scanner types, and clinical protocols. The establishment of European infrastructures such as EUCAIM and EHDS is crucial in this regard. These initiatives aim to harmonise imaging data sharing across Member States while upholding FAIR principles, enabling developers to access distributed datasets through federated learning and other privacy-preserving technologies, without compromising patient confidentiality. In parallel, initiatives like TEF-Health and AI Factories provide essential support for scalable testing, validation, and training of high-risk AI systems in regulated environments. TEF-Health plays a critical role in implementing the AI Act, by offering clinically relevant settings for performance evaluation and compliance testing. AI Factories complement this by leveraging Europe's high-performance computing infrastructure to accelerate the training of large-scale medical imaging models. Beyond infrastructure and access, the legal and ethical frameworks that govern health data use must strike a careful balance: protecting individual privacy (as required under the GDPR) while enabling responsible scientific use. Promoting secure data sharing, anonymisation or pseudonymisation, and secured processing environments, along with streamlined ethical approval pathways, is necessary to reduce friction for AI research and clinical innovation. Finally, while the models developed here show high average performance, trustworthy AI also requires transparency, interpretability, and calibrated risk communication. Techniques such as conformal risk control offer individual-level confidence guarantees, which are critical in clinical contexts where uncertainty must be explicitly communicated. Broader adoption of such methods,

alongside mechanisms for expert oversight and iterative validation, will be essential to integrate AI safely into clinical decision-making. In conclusion, this work highlights both the technical readiness of AI for lung cancer applications and the systemic conditions needed to support its transition into real-world use. Delivering trustworthy, clinically accepted AI solutions will depend not only on algorithmic excellence but also on a coordinated European effort to build the legal, infrastructural, and ethical foundations for equitable and sustainable deployment.

5.2. Biomechanics-Driven Neural Network for Cardiovascular Disease Identification

CVDs are the leading cause of mortality worldwide, accounting for nearly 18 million deaths annually [79]. Traditional diagnostic methods based on manual interpretation of medical images are time-consuming, error-prone, and variable. DL approaches, particularly CNNs, have significantly advanced automatic classification of CVDs, offering rapid and consistent analysis. Cine-cardiac MRI, which captures dynamic heart motion, is especially valuable for AI applications. The ACDC dataset [80] serves as a benchmark for automated classification of five cardiac pathology categories. Previous studies achieved strong results using radiomics or anatomical segmentation features [81] [82] [83] but largely ignored the dynamic and biomechanical properties of cardiac tissues. Deformable image registration offers a way to bridge this gap by estimating non-linear tissue motion, and the incorporation of biomechanical priors such as Neo-Hookean elasticity leads to more physiologically plausible deformation fields. Recent approaches [84] [85] have demonstrated the benefits of biomechanics-informed registration for analysing cardiac dynamics. Building on these ideas, this study proposes a novel method that extracts biomechanical parameters from cine-MRI sequences using a biomechanics-constrained DL registration model. These features, combined with conventional radiomics, enable more accurate and explainable classification of cardiac diseases, enhancing both performance and clinical trustworthiness.

5.2.1. Deep Cascaded Registration

The registration strategy we use is based on Comte et al. (2025), who proposed a cascaded deep learning model for the registration of 3D medical images [36]. Unlike traditional multi-stage approaches, where sequential CNNs progressively refine deformations, the cascaded model processes inputs simultaneously through stacked networks, computing the loss only at the final output to encourage collaborative optimisation.

The registration task seeks a spatial transformation $\varphi : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ that warps a moving image X_{mv} to align with a fixed image X_{fx} ,

so that:

$$X_{wp} = X_{mv} \circ \varphi \approx X_{fx}$$

At each cascade level, a network predicts a deformation field φ_i and the cumulative transformation is obtained by summing the intermediate fields:

$$X_{wp,n} = X_{mv} \circ \sum_{i=1}^n \varphi_i$$

The model optimises a loss combining an image similarity term and a regularisation term, based on the negative local normalised cross-correlation (NCC) between the warped and fixed images:

$$\mathcal{L}_{sim} = -NCC(X_{fx}, X_{wp})$$

This loss is computed for patches of varying sizes, promoting image similarity at multiple scales.

5.2.2. Biomechanical Regularisation

In DL-based medical image registration, deformation fields are often regularised by penalising the sum of their local gradients. This is typically expressed as the integral or sum over the domain of the squared gradient norm of the deformation field $\sum ||\nabla\varphi(x)||^2$. This approach ensures that the predicted transformations are smooth and invertible. However, it does not consider the mechanical properties of biological tissues, particularly cardiac tissue, which exhibits nonlinear and anisotropic behaviour under stress. As such, the deformations produced may lack physiological realism and be difficult to interpret in a clinical setting. To address this limitation, a biomechanical regularisation strategy based on Neo-Hookean strain energy is introduced. This approach models the myocardium as a hyperelastic material, allowing for the description of large, nonlinear deformations. The transformation of each tissue voxel is defined by a function $\varphi(x)$, which maps its position in the reference configuration to a new location in the deformed configuration. The deformation gradient tensor F is defined as the spatial derivative of φ with respect to x :

$$\mathbf{F} = \partial\varphi / \partial x$$

From this, the right Cauchy-Green deformation tensor is computed as:

$$\mathbf{C} = \mathbf{F}^T \mathbf{F}$$

This tensor characterises local stretch without sensitivity to rigid-body motion. The volume change induced by the deformation is captured by the determinant of F :

$$J = \det(\mathbf{F})$$

The strain energy density of a Neo-Hookean material depends on both the shape and volume changes of the tissue. It is given by the following expression:

$$\Phi(\varphi) = \frac{\mu}{2} \left(I_{1C} J^{-\frac{2}{3}} - 3 \right) + \frac{\kappa}{2} (J - 1)^2$$

Here, μ and κ represent the shear and bulk moduli of the tissue, I_{1C} is the first invariant of the tensor C (i.e., the trace of C), and J accounts for the local volumetric expansion or compression. The term involving I_{1C} penalises distortional changes, while the term involving J penalises volume changes, thus allowing for the independent control of tissue stretching and compression.

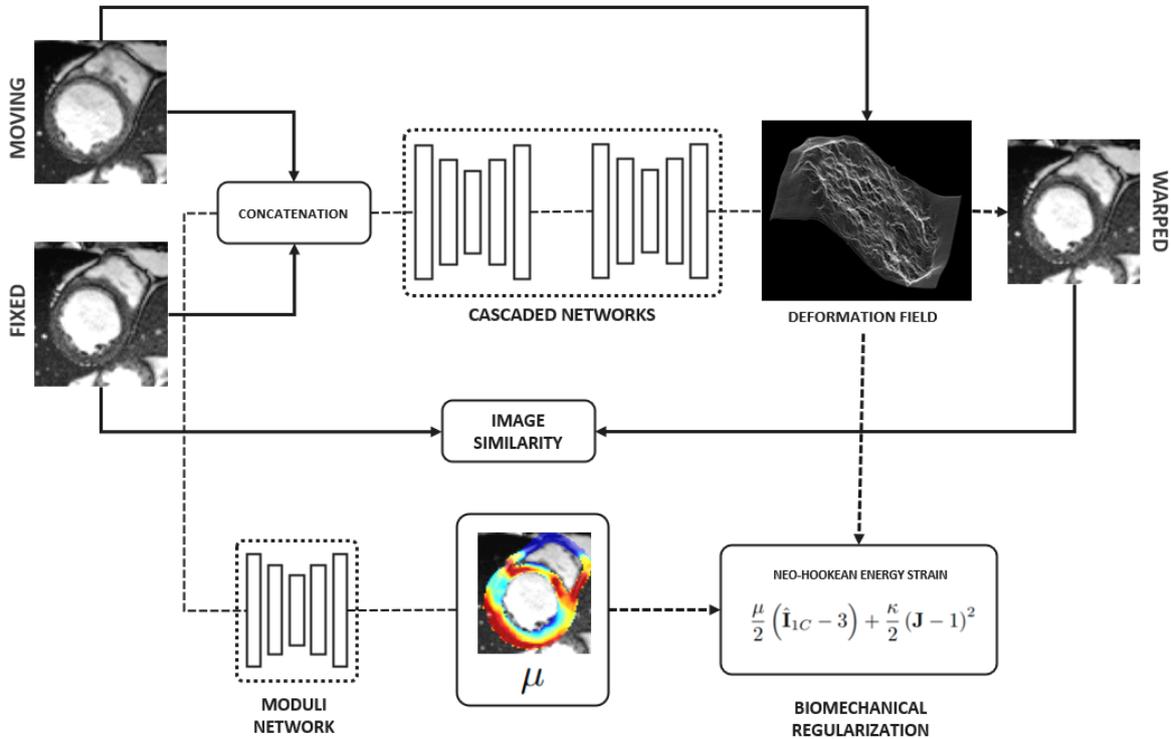
This biomechanical regularisation term is added to the overall training objective of the model. The full loss function is:

$$\mathcal{L} = \mathcal{L}_{sim} + \lambda \cdot \Phi(\varphi)$$

where \mathcal{L}_{sim} is a similarity loss, often based on normalised cross-correlation between the fixed and warped images and λ is a regularisation parameter that controls the influence of the biomechanical prior. By incorporating this physically grounded energy model into the learning objective, the registration framework produces deformation fields that are both mathematically stable and physiologically plausible (**Figure 23**). This enhances not only the numerical robustness of the model but also its explainability and clinical trustworthiness, as the transformations can be interpreted in

terms of realistic tissue mechanics. This makes the method particularly suitable for cardiac applications, where the deformation patterns of the myocardium carry diagnostic significance.

Figure 23: Overview of the biomechanics-driven framework.



Source: own elaboration.

5.2.3. Segmentation with Multi-Frame Propagation

We compare the segmentation results obtained using two label propagation techniques in **Table 1**. The first is the direct label propagation approach, in which segmentation labels are transferred directly between the End-Systole (ES) and End-Diastole (ED) frames. The second is the multi-frame propagation technique, where labels are sequentially propagated from ED to adjacent frames, and all resulting labels are repropagated back to the ED frame. A local weighted voting strategy is then applied to merge the propagated labels, improving robustness and consistency across the sequence (**Figure 24**).

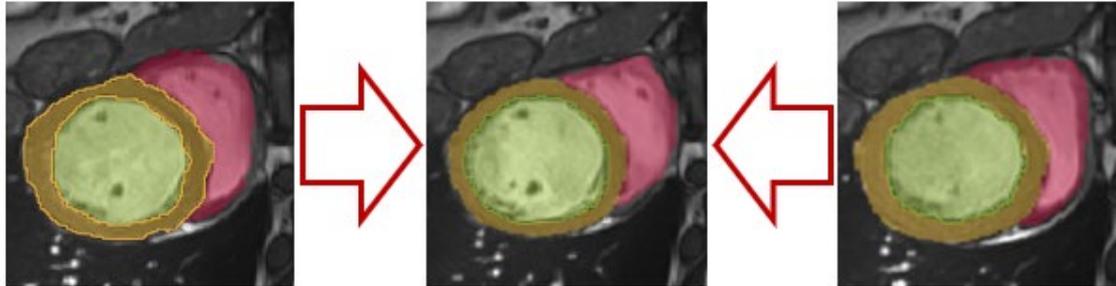
Table 1: Average Dice scores on the three anatomical labels at ES and ED using registration ED to ES frames, and vice versa, multi-frame registration, and nnU-Net.

Method	ES			ED		
	LV	RV	MYO	LV	RV	MYO
nnU-Net	0.968	0.946	0.902	0.931	0.899	0.919
Registration	0.970	0.878	0.920	0.913	0.804	0.870
Multi-Frame Registration	0.974	0.912	0.960	0.956	0.857	0.922

Source: own elaboration

Figure 24: Segmentation using multi-frame propagation.

Label propagation + fusion (local weighted voting)

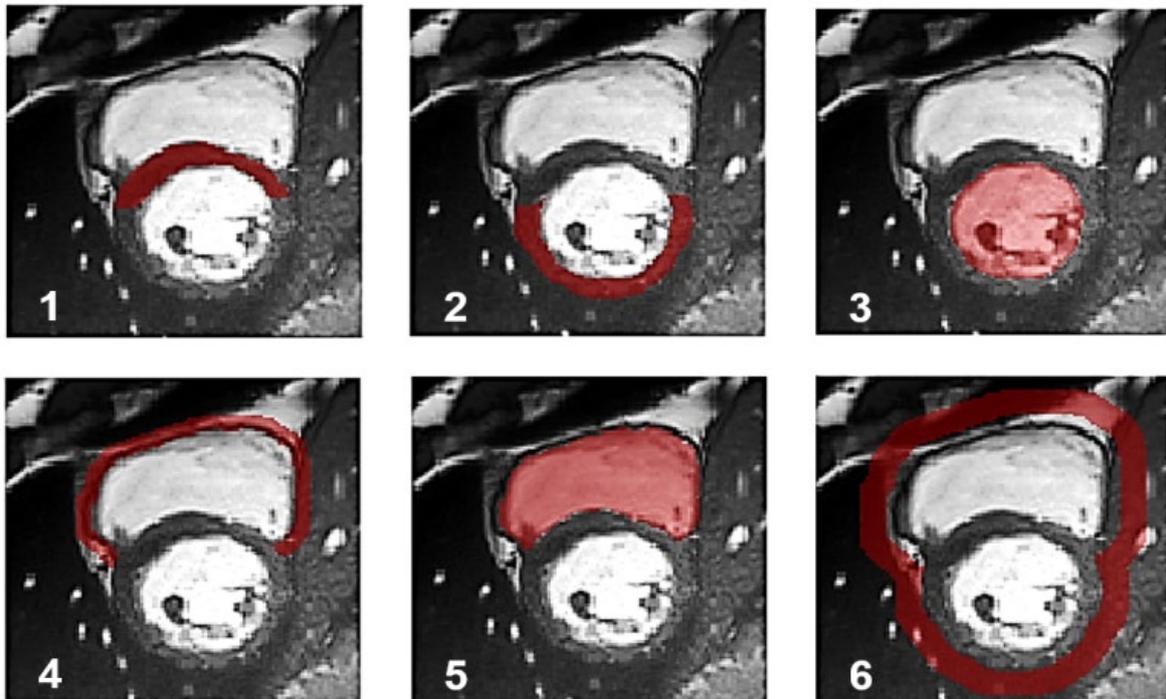


Source: own elaboration.

5.2.4. Classification of CVDs

A comprehensive set of features is extracted from the locally estimated biomechanical parameters, shear modulus μ and bulk modulus κ , computed at both End-Diastole (ED) and End-Systole (ES). For each of six anatomical regions (**Figure 25**), we derive statistical descriptors including the mean, standard deviation, and the 10th and 90th percentiles. This biomechanical feature set is further complemented with volumetric measurements obtained from the ground-truth segmentations at ED and ES. To identify the most informative features for cardiovascular disease classification, a greedy forward-backward feature selection strategy is employed.

Figure 25: Labels used to compute the local biomechanical features.



Source: own elaboration.

To evaluate the contribution of biomechanical features to disease classification, we constructed a feature pool comprising values derived from local estimates of the shear modulus μ , bulk modulus κ , the local modulus of the deformation field φ , and volumetric measurements. The performance of five different classification algorithms was assessed: logistic regression, multilayer perceptron, support vector classifier, random forest, and nearest neighbour, each combined with a feature selection procedure. The classification accuracies achieved by each method are reported in Table 2, with logistic regression reaching the highest performance, attaining an accuracy of 0.98 on the training set and a perfect 1.0 on the test set.

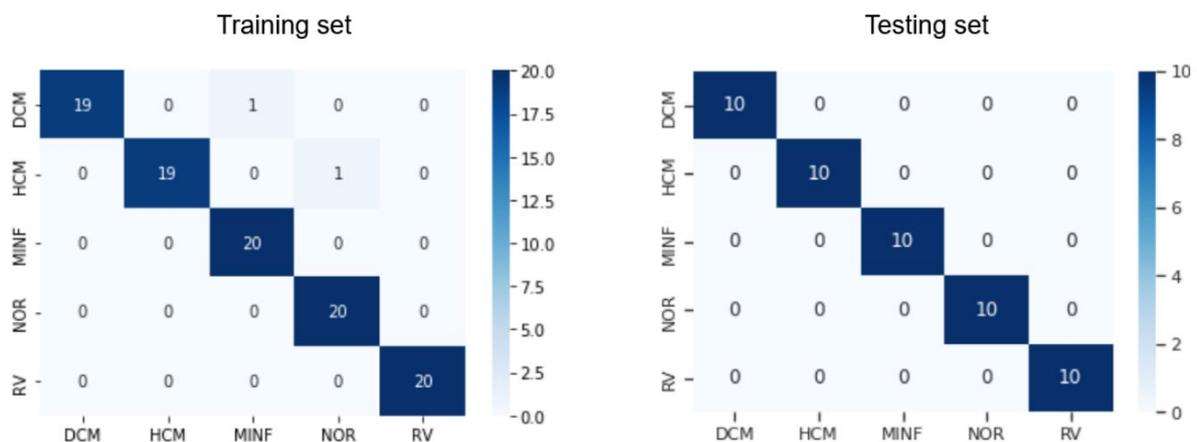
Table 2: Classification accuracy over the train and test set for different classification methods.

Classifier	Accuracy train	Accuracy test
Logistic Regression	0.98	1
Multi-Layer Perceptron	0.91	0.96
Support Vector Classifier	0.90	0.92
Random Forest	0.89	0.92
Nearest Neighbour	0.87	0.88

Source: own elaboration

Figure 26 displays the confusion matrices for the classification task across the five cardiac disease categories. These matrices provide a detailed assessment of the model's performance by illustrating how accurately each case is assigned to its correct class. On the training set, the model demonstrates strong performance, with only two misclassifications observed. Specifically, one case of dilated cardiomyopathy (DCM) is incorrectly classified as myocardial infarction (MINF), and one instance of hypertrophic cardiomyopathy (HCM) is misclassified as normal (NOR). Nevertheless, the model achieves perfect classification on the independent test set, correctly identifying all cases.

Figure 26: Confusion matrices of the classification on the training set and testing set.

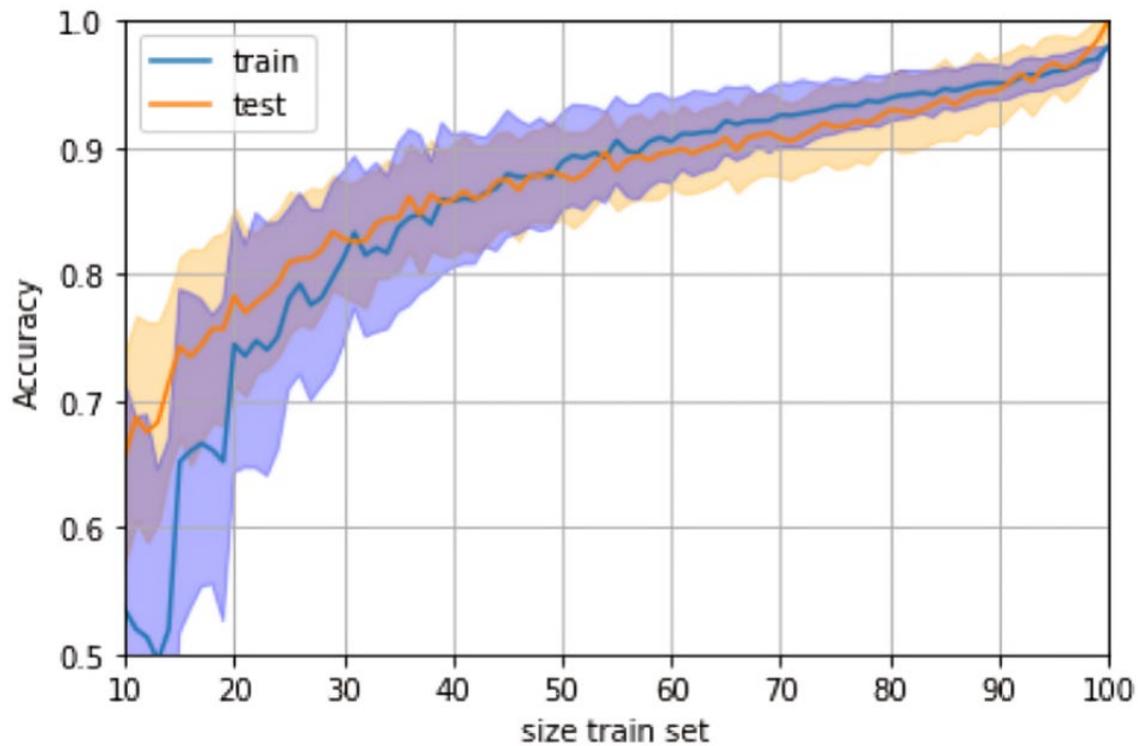


Source: own elaboration.

Figure 27 examines the generalisation ability of the classification model by showing how accuracy varies with different training set sizes. Using the selected feature set and a Logistic Regression classifier, fifty random subsets were generated for each training size, and the average classification accuracy was computed. The solid line represents the mean accuracy, while the shaded area shows one standard deviation. The results demonstrate that the model maintains strong performance even with limited data, achieving accuracy above 0.7, 0.8, and 0.9 for training sets of 20, 30, and 60 subjects, respectively. The narrow variability, particularly at larger training sizes, indicates that the

model is stable and reliable. This robustness highlights its potential for clinical deployment, where data availability may be variable and resilience across patient populations is essential.

Figure 27: Accuracy as a function of the training set size.



Source: own elaboration.

5.2.5. Discussion

The second use case investigated the integration of biomechanical modelling into AI-based classification of cardiovascular diseases (CVDs), offering a novel and promising direction for building interpretable and trustworthy AI systems in medical imaging. While conventional image-based deep learning approaches can deliver strong predictive performance, they often function as opaque “black boxes”, posing challenges for clinical adoption, particularly in high-stakes settings where transparency and accountability are paramount. This study addresses the challenge by incorporating a biomechanics-informed regularisation scheme based on the Neo-Hookean strain energy model into the AI framework. By constraining the learned deformation fields to align with the physical behaviour of cardiac tissues, the model produces features that are not only discriminative but also physiologically meaningful. Compared to models using standard volumetric features, the inclusion of these biomechanical quantities leads to significantly improved classification accuracy. Crucially, this approach enhances the explainability of AI predictions. Instead of relying on abstract latent features, clinicians can interpret the model’s outputs using familiar mechanical concepts, such as tissue stiffness, strain patterns, and regional deformation. This physical interpretability builds clinical trust, enables expert validation, and facilitates dialogue between automated analyses and medical judgment. The use case reinforces that explainability and trustworthiness are essential, for the successful integration of AI into cardiovascular diagnostics. High accuracy alone is not sufficient, clinicians must be able to understand, interrogate, and act upon AI outputs with confidence. By blending data-driven learning with domain-specific physical priors, the method provides a tangible path forward for developing clinically viable AI tools. Looking ahead, further work should focus on

validating these models across larger, more heterogeneous datasets and diverse clinical settings to ensure generalisability. Equally important is the development of user-centric interfaces that present biomechanical features in ways that are accessible, actionable, and seamlessly integrated into existing clinical workflows—maximising both their interpretive value and practical utility. This approach reflects a growing recognition—across research, policy, and clinical communities—that trustworthiness must be a foundational principle in the development of medical AI. Beyond raw performance, AI systems must demonstrate explainability, robustness, ethical soundness, and respect for human oversight.

These values are strongly echoed in the European approach to AI in healthcare, which promotes the development of technologies that are:

- Transparent and interpretable, enabling expert review and patient trust
- Technically and clinically validated, supporting safe and effective deployment

Ultimately, this use case underscores that the path to real-world impact is not only through algorithmic advances, but through a broader commitment to trust, safety, and usability—core principles that are increasingly central to the European vision for AI in health. Future efforts should continue to build on this foundation, advancing AI that is not only powerful, but also clinically meaningful, ethically grounded, and designed with human users in mind.

6. Conclusions

This report reviewed the **current landscape of AI in medical imaging**, analysing **recent progress** and **remaining challenges** across key computational tasks such as segmentation, detection, classification, registration, and synthetic data generation. It also situated these developments within the broader **European policy landscape**, focusing on **regulatory instruments** (GDPR, MDR, AI Act, EHDS) and **EU initiatives** including the European Cancer Imaging Initiative, TEF-Health, AI Factories, and ERICs and the EIC portfolio on AI-driven innovation in medical imaging. To ground the analysis in practice, the report examined **two selected use cases**: one on **lung cancer screening and longitudinal monitoring**, the other on **cardiovascular disease (CVD) classification using biomechanics-informed neural networks**. These cases illustrate how **AI methods can be tailored to real clinical workflows**, offering not only technical improvements but also insights into the systemic conditions necessary for **trustworthy deployment**.

In the first use case, a comprehensive pipeline was developed for **lung cancer screening** and follow-up. It combined multiple steps: chest segmentation, lung nodule detection, longitudinal follow-up, and conformal prediction to quantify uncertainty. This approach demonstrates that **early diagnosis and follow-up of pulmonary nodules** can benefit significantly from AI when multi-step image analysis tasks are integrated and clinically validated. The case also underscored the need for **secure and scalable infrastructure** to manage longitudinal CT data. Overall, this use case demonstrated:

- The potential of AI to support **early diagnosis and longitudinal monitoring** when image analysis tasks are integrated into a single, clinically validated workflow.
- The value of **uncertainty quantification** for enabling risk-based clinical decisions, in line with the AI Act's requirements for transparency and traceability in high-risk medical applications.
- **The importance of secure and scalable infrastructure** to manage medical imaging data, particularly to support the secondary use of data for research and innovation. This will be facilitated by the European Health Data Space (EHDS), which aims to enable the GDPR-compliant and trustworthy reuse of health data across borders.

In the second use case, a **hybrid deep learning model** was used to **classify cardiovascular diseases** based on **biomechanical features** derived from image registration. Registration techniques were not only used for spatial alignment, but also to estimate **realistic** deformation fields. These were regularised using a neo-Hookean energy model to produce interpretable biomechanical biomarkers. This case illustrated:

- The potential of **hybrid approaches** that integrate deep learning with **domain knowledge** to enhance both model performance and clinical relevance.
- How grounding AI outputs in physiological parameters can facilitate **transparency** and support **human oversight**, helping clinicians interpret and trust model predictions in practice.

Together, the use cases and broader analysis point to five enabling conditions for scaling AI in medical imaging across Europe:

- **Hybrid and domain-informed approaches** offer a promising path towards clinically meaningful and interpretable AI tools. Integrating medical knowledge, physics-based priors, and uncertainty estimation can enhance trust and relevance in real-world settings.

- **Transparency and clinician oversight** should be embedded from the start. User-centred design and interactive interfaces can support the safe use of AI by aligning tools with clinical workflows and decision-making practices.
- **Federated infrastructure and annotated data access** remain critical bottlenecks. Scalable, secure, and standardised systems for sharing annotated imaging data are essential to train robust and generalisable models while preserving privacy. European projects and initiatives such as the European Cancer Imaging Initiative, and the EHDS regulation aim to address these needs by enabling cross-border access to imaging data while upholding GDPR compliance.
- **Interoperability and standardisation** must be reinforced across data formats, annotation protocols, and model evaluation. Aligning with EU-wide and international frameworks will support reproducibility and regulatory consistency.
- **Ecosystem coordination and long-term investment** are needed to translate innovation into deployment. Efforts like TEF-Health, AI Factories, ERICs and the proposed EUCAIM EDIC offer the infrastructure and governance needed to coordinate AI development for medical imaging, validation, and deployment across Europe. They provide frameworks for supporting collaborative research, development and validation, as well as controlled environments for regulatory testing and offer a foundation for long-term public investment. Sustained coordination between research, policy, and clinical domains will be essential to build a cohesive and trustworthy AI ecosystem.

In summary, AI in medical imaging holds transformative potential for early diagnosis, personalised treatment, and efficient healthcare delivery. Yet, the path forward is not purely technical. It requires a broad and integrated strategy: grounded in trust, supported by infrastructure, aligned with regulation, and guided by clinical needs. Europe's continued investment in federated infrastructures, regulatory foresight, and human-centred AI design puts it in a strong position to lead in the development of safe, effective, and inclusive medical AI solutions.

References

- [1] European Commission, Joint Research Centre, Mochan, A., Farinha, J., Bailey, G., Rodriguez, L., Zanca, F. and Polvora, A., Imaging the Future - Horizon scanning for emerging technologies and breakthrough innovations in the field of medical imaging and AI, Publications Office of the European Union, Luxembourg, 2024, <https://data.europa.eu/doi/10.2760/1426669>, JRC140099.
- [2] Wilkinson, M. D., Dumontier, M., Aalbersberg, I. J., Appleton, G., Axton, M., Baak, A., ... & Mons, B. (2016). The FAIR Guiding Principles for scientific data management and stewardship. *Scientific data*, 3(1), 1-9.
- [3] Agrawal, R., & Prabakaran, S. (2020). Big data in digital healthcare: lessons learnt and recommendations for general practice. *Heredity*, 124(4), 525-534.
- [4] Villarreal, E. R. D., García-Alonso, J., Moguel, E., & Alegría, J. A. H. (2023). Blockchain for healthcare management systems: A survey on interoperability and security. *IEEE Access*, 11, 5629-5652.
- [5] Sarvamangala, D. R., & Kulkarni, R. V. (2022). Convolutional neural networks in medical image understanding: a survey. *Evolutionary intelligence*, 15(1), 1-22.
- [6] Anwar, S. M., Majid, M., Qayyum, A., Awais, M., Alnowami, M., & Khan, M. K. (2018). Medical image analysis using convolutional neural networks: a review. *Journal of medical systems*, 42, 1-13.
- [7] Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., ... & Houlsby, N. (2020). An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- [8] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I. (2017). Attention is all you need. *Advances in neural information processing systems*, 30.
- [9] Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., ... & Guo, B. (2021). Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision* (pp. 10012-10022).
- [10] Meskó, B. (2023). The impact of multimodal large language models on health care's future. *Journal of medical Internet research*, 25, e52865.
- [11] Zhang, X., Wu, C., Zhao, Z., Lin, W., Zhang, Y., Wang, Y., & Xie, W. (2023). Pmc-vqa: Visual instruction tuning for medical visual question answering. *arXiv preprint arXiv:2305.10415*.

- [12] Wang, Z., Wu, Z., Agarwal, D., & Sun, J. (2022, December). Medclip: Contrastive learning from unpaired medical images and text. In Proceedings of the Conference on Empirical Methods in Natural Language Processing. Conference on Empirical Methods in Natural Language Processing.
- [13] Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., ... & Sutskever, I. (2021, July). Learning transferable visual models from natural language supervision. In International conference on machine learning (pp. 8748-8763). Pmlr.
- [14] Moor, M., Huang, Q., Wu, S., Yasunaga, M., Dalmia, Y., Leskovec, J., ... & Rajpurkar, P. (2023, December). Med-flamingo: a multimodal medical few-shot learner. In Machine Learning for Health (ML4H) (pp. 353-367). PMLR.
- [15] Awadalla, A., Gao, I., Gardner, J., Hessel, J., Hanafy, Y., Zhu, W., ... & Schmidt, L. (2023). Openflamingo: An open-source framework for training large autoregressive vision-language models. arXiv preprint arXiv:2308.01390.
- [16] Ronneberger, O., Fischer, P., & Brox, T. (2015). U-net: Convolutional networks for biomedical image segmentation. In *Medical image computing and computer-assisted intervention—MICCAI 2015: 18th international conference, Munich, Germany, October 5-9, 2015*.
- [17] Çiçek, Ö., Abdulkadir, A., Lienkamp, S. S., Brox, T., & Ronneberger, O. (2016). 3D U-Net: learning dense volumetric segmentation from sparse annotation. In *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2016: 19th International Conference, Athens, Greece, October 17-21, 2016, Proceedings, Part II 19* (pp. 424-432). Springer International Publishing.
- [18] Li, X., Chen, H., Qi, X., Dou, Q., Fu, C. W., & Heng, P. A. (2018). H-DenseUNet: hybrid densely connected UNet for liver and tumor segmentation from CT volumes. *IEEE transactions on medical imaging*, 37(12), 2663-2674.
- [19] Isensee, F., Jaeger, P. F., Kohl, S. A., Petersen, J., & Maier-Hein, K. H. (2021). nnU-Net: a self-configuring method for deep learning-based biomedical image segmentation. *Nature methods*, 18(2), 203-211.
- [20] Maurício, J., Domingues, I., & Bernardino, J. (2023). Comparing vision transformers and convolutional neural networks for image classification: A literature review. *Applied Sciences*, 13(9), 5521.
- [21] Chen, J., Lu, Y., Yu, Q., Luo, X., Adeli, E., Wang, Y., ... & Zhou, Y. (2021). Transunet: Transformers make strong encoders for medical image segmentation. arXiv preprint arXiv:2102.04306.
- [22] Girshick, R., Donahue, J., Darrell, T., & Malik, J. (2014). Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 580-587).

- [23] Girshick, R. (2015). *Fast r-cnn*. In *Proceedings of the IEEE international conference on computer vision* (pp. 1440-1448).
- [24] Ren, S., He, K., Girshick, R., & Sun, J. (2015). *Faster r-cnn: Towards real-time object detection with region proposal networks*. *Advances in neural information processing systems*, 28.
- [25] Lin, T. Y., Dollár, P., Girshick, R., He, K., Hariharan, B., & Belongie, S. (2017). *Feature pyramid networks for object detection*. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 2117-2125).
- [26] Redmon, J., Divvala, S., Girshick, R., & Farhadi, A. (2016). *You only look once: Unified, real-time object detection*. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 779-788).
- [27] Liu, W., Anguelov, D., Erhan, D., Szegedy, C., Reed, S., Fu, C. Y., & Berg, A. C. (2016). *Ssd: Single shot multibox detector*. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part I 14* (pp. 21-37). Springer International Publishing.
- [28] Lin, T. Y., Goyal, P., Girshick, R., He, K., & Dollár, P. (2017). *Focal loss for dense object detection*. In *Proceedings of the IEEE international conference on computer vision* (pp. 2980-2988).
- [29] Lu, Y., Yu, Q., Gao, Y., Zhou, Y., Liu, G., Dong, Q., ... & Yang, S. (2018). *Identification of metastatic lymph nodes in MR imaging with faster region-based convolutional neural networks*. *Cancer research*, 78(17), 5135-5143.
- [30] Chen, X., Zhang, K., Lin, S., Dai, K. F., & Yun, Y. (2021). *Single shot multibox detector automatic polyp detection network based on gastrointestinal endoscopic images*. *Computational and Mathematical Methods in Medicine*, 2021(1), 2144472.
- [31] Harsono, I. W., Liawatimena, S., & Cenggoro, T. W. (2022). *Lung nodule detection and classification from Thorax CT-scan using RetinaNet with transfer learning*. *Journal of King Saud University-Computer and Information Sciences*, 34(3), 567-577.
- [32] Fu, Y., Lei, Y., Wang, T., Curran, W. J., Liu, T., & Yang, X. (2020). *Deep learning in medical image registration: a review*. *Physics in Medicine & Biology*, 65(20), 20TR01.
- [33] Boveiri, H. R., Khayami, R., Javidan, R., & Mehdizadeh, A. (2020). *Medical image registration using deep neural networks: a comprehensive review*. *Computers & Electrical Engineering*, 87, 106767.
- [34] Balakrishnan, G., Zhao, A., Sabuncu, M. R., Guttag, J., & Dalca, A. V. (2019). *Voxelmorph: a learning framework for deformable medical image registration*. *IEEE transactions on medical imaging*, 38(8), 1788-1800.

- [35] Jaderberg, M., Simonyan, K., & Zisserman, A. (2015). *Spatial transformer networks*. *Advances in neural information processing systems*, 28.
- [36] Comte, V., Alenya, M., Urru, A., Recober, J., Nakaki, A., Crovetto, F., ... & Ballester, M. A. G. (2025). *Deep cascaded registration and weakly-supervised segmentation of fetal brain MRI*. *Heliyon*, 11(1).
- [37] Chen, J., Frey, E. C., He, Y., Segars, W. P., Li, Y., & Du, Y. (2022). *Transmorph: Transformer for unsupervised medical image registration*. *Medical image analysis*, 82, 102615.
- [38] Iqbal, A., Khan, R., & Karayannis, T. (2019). *Developing a brain atlas through deep learning*. *Nature Machine Intelligence*, 1(6), 277-287.
- [39] Zhang, Y., Shi, J., Wei, H., Han, V., Zhu, W. Z., & Liu, C. (2019). *Neonate and infant brain development from birth to 2 years assessed using MRI-based quantitative susceptibility mapping*. *Neuroimage*, 185, 349-360.
- [40] Ding, Y., Feng, H., Yang, Y., Holmes, J., Liu, Z., Liu, D., ... & Liu, W. (2023). *Deep-learning based fast and accurate 3D CT deformable image registration in lung cancer*. *Medical physics*, 50(11), 6864-6880.
- [41] Li, Y., Daho, M. E. H., Conze, P. H., Zeglache, R., Le Boité, H., Tadayoni, R., ... & Quéllec, G. (2024). *A review of deep learning-based information fusion techniques for multimodal medical image classification*. *Computers in Biology and Medicine*, 108635.
- [42] Kuang, H., Wang, Y., Liu, J., Wang, J., Cao, Q., Hu, B., ... & Wang, J. (2024). *Hybrid cnn-transformer network with circular feature interaction for acute ischemic stroke lesion segmentation on non-contrast ct scans*. *IEEE Transactions on Medical Imaging*, 43(6), 2303-2316.
- [43] Mohammed, F. A., Tune, K. K., Assefa, B. G., Jett, M., & Muhie, S. (2024). *Medical image classifications using convolutional neural networks: A survey of current methods and statistical modeling of the literature*. *Machine learning and knowledge extraction*, 6(1), 699-735.
- [44] Patrício, C., Neves, J. C., & Teixeira, L. F. (2023). *Explainable deep learning methods in medical image classification: A survey*. *ACM Computing Surveys*, 56(4), 1-41.
- [45] Pudjihartono, N., Fadason, T., Kempa-Liehr, A. W., & O'Sullivan, J. M. (2022). *A review of feature selection methods for machine learning-based disease risk prediction*. *Frontiers in Bioinformatics*, 2, 927312.
- [46] Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., ... & Bengio, Y. (2020). *Generative adversarial networks*. *Communications of the ACM*, 63(11), 139-144.

- [47] Jeong, J. J., Tariq, A., Adejumo, T., Trivedi, H., Gichoya, J. W., & Banerjee, I. (2022). Systematic review of generative adversarial networks (GANs) for medical image classification and segmentation. *Journal of Digital Imaging*, 35(2), 137-152.
- [48] Aggarwal, R., Sounderajah, V., Martin, G., Ting, D. S., Karthikesalingam, A., King, D., ... & Darzi, A. (2021). Diagnostic accuracy of deep learning in medical imaging: a systematic review and meta-analysis. *NPJ digital medicine*, 4(1), 65.
- [49] Khader, F., Müller-Franzes, G., Tayebi Arasteh, S., Han, T., Haarbuerger, C., Schulze-Hagen, M., ... & Truhn, D. (2023). Denoising diffusion probabilistic models for 3D medical image generation. *Scientific Reports*, 13(1), 7303.
- [50] Umirzakova, S., Ahmad, S., Khan, L. U., & Whangbo, T. (2024). Medical image super-resolution for smart healthcare applications: A comprehensive survey. *Information Fusion*, 103, 102075.
- [51] Ahishakiye, E., Bastiaan Van Gijzen, M., Tumwiine, J., Wario, R., & Obungoloch, J. (2021). A survey on deep learning in medical image reconstruction. *Intelligent Medicine*, 1(03), 118-127.
- [52] Pinaya, W. H., Graham, M. S., Kerfoot, E., Tudosi, P. D., Dafflon, J., Fernandez, V., ... & Cardoso, M. J. (2023). Generative ai for medical imaging: extending the monai framework. arXiv preprint arXiv:2307.15208.
- [53] Guibas, J. T., Viridi, T. S., & Li, P. S. (2017). Synthetic medical images from dual generative adversarial networks. arXiv preprint arXiv:1709.01872.
- [54] Sauer, A., Schwarz, K., & Geiger, A. (2022, July). Stylegan-xl: Scaling stylegan to large diverse datasets. In *ACM SIGGRAPH 2022 conference proceedings* (pp. 1-10).
- [55] Park, T., Liu, M. Y., Wang, T. C., & Zhu, J. Y. (2019). Semantic image synthesis with spatially-adaptive normalization. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 2337-2346).
- [56] Skandarani, Y., Jodoin, P. M., & Lalande, A. (2023). Gans for medical image synthesis: An empirical study. *Journal of Imaging*, 9(3), 69.
- [57] Pan, S., Wang, T., Qiu, R. L., Axente, M., Chang, C. W., Peng, J., ... & Yang, X. (2023). 2D medical image synthesis using transformer-based denoising diffusion probabilistic model. *Physics in Medicine & Biology*, 68(10), 105004.
- [58] Dalmaz, O., Yurt, M., & Çukur, T. (2022). ResViT: residual vision transformers for multimodal medical image synthesis. *IEEE Transactions on Medical Imaging*, 41(10), 2598-2614.

- [59] Kondylakis, H., Kalokyri, V., Sfakianakis, S., Marias, K., Tsiknakis, M., Jimenez-Pastor, A., ... & Lekadir, K. (2023). Data infrastructures for AI in medical imaging: a report on the experiences of five EU projects. *European radiology experimental*, 7(1), 20.
- [60] Sheller, M. J., Edwards, B., Reina, G. A., Martin, J., Pati, S., Kotrotsou, A., ... & Bakas, S. (2020). Federated learning in medicine: facilitating multi-institutional collaborations without sharing patient data. *Scientific reports*, 10(1), 12598.
- [61] Kumar, R., Khan, A. A., Kumar, J., Golilarz, N. A., Zhang, S., Ting, Y., ... & Wang, W. (2021). Blockchain-federated-learning and deep learning models for covid-19 detection using ct imaging. *IEEE Sensors Journal*, 21(14), 16301-16314.
- [62] Wiens, J., Saria, S., Sendak, M., Ghassemi, M., Liu, V. X., Doshi-Velez, F., ... & Goldenberg, A. (2019). Do no harm: a roadmap for responsible machine learning for health care. *Nature medicine*, 25(9), 1337-1340.
- [63] Kaissis, G. A., Makowski, M. R., Rückert, D., & Braren, R. F. (2020). Secure, privacy-preserving and federated machine learning in medical imaging. *Nature Machine Intelligence*, 2(6), 305-311.
- [64] Gong, Y., Chang, X., Mišić, J., Mišić, V. B., Wang, J., & Zhu, H. (2024). Practical solutions in fully homomorphic encryption: a survey analyzing existing acceleration methods. *Cybersecurity*, 7(1), 5.
- [65] Kaissis, G., Ziller, A., Passerat-Palmbach, J., Ryffel, T., Usynin, D., Trask, A., ... & Braren, R. (2021). End-to-end privacy preserving deep learning on multi-institutional medical imaging. *Nature Machine Intelligence*, 3(6), 473-484.
- [66] Kaabachi, B., Despraz, J., Meurers, T., Otte, K., Halilovic, M., Kulynych, B., ... & Raisaro, J. L. (2025). A scoping review of privacy and utility metrics in medical synthetic data. *NPJ digital medicine*, 8(1), 60.
- [67] Lundberg, S. M., Erion, G., Chen, H., DeGrave, A., Prutkin, J. M., Nair, B., ... & Lee, S. I. (2020). From local explanations to global understanding with explainable AI for trees. *Nature machine intelligence*, 2(1), 56-67.
- [68] Ronneberger, O., Fischer, P., & Brox, T. (2015). *U-net: Convolutional networks for biomedical image segmentation*. In *Medical image computing and computer-assisted intervention–MICCAI 2015: 18th international conference, Munich, Germany, October 5–9, 2015, proceedings, part III 18* (pp. 234-241). Springer international publishing.
- [69] Baumgartner, M., Jäger, P. F., Isensee, F., & Maier-Hein, K. H. (2021). *nnDetection: a self-configuring method for medical object detection*. In *Medical Image Computing and Computer Assisted Intervention–MICCAI 2021: 24th International Conference, Strasbourg, France*,

September 27–October 1, 2021, Proceedings, Part V 24 (pp. 530–539). Springer International Publishing.

- [70] Wasserthal, Jakob, et al. "Totalsegmentator: Robust segmentation of 104 anatomic structures in ct images." *Radiology: Artificial Intelligence* 5.5 (2023).
- [71] Chai, J., Zeng, H., Li, A., & Ngai, E. W. (2021). Deep learning in computer vision: A critical review of emerging techniques and application scenarios. *Machine Learning with Applications*, 6, 100134.
- [72] Ali, H., Mohsen, F., & Shah, Z. (2023). Improving diagnosis and prognosis of lung cancer using vision transformers: a scoping review. *BMC Medical Imaging*, 23(1), 129.
- [73] Demir, A., Yilmaz, F., & Kose, O. (2019, October). Early detection of skin cancer using deep learning architectures: resnet-101 and inception-v3. In *2019 medical technologies congress (TIPTEKNO)* (pp. 1-4). IEEE.
- [74] Kang, G., Liu, K., Hou, B., & Zhang, N. (2017). 3D multi-view convolutional neural networks for lung nodule classification. *PloS one*, 12(11), e0188290.
- [75] Hwang, S., & Kim, H. E. (2016). Self-transfer learning for weakly supervised lesion localization. In *Medical Image Computing and Computer-Assisted Intervention– MICCAI 2016: 19th International Conference, Athens, Greece, October 17–21, 2016, Proceedings, Part II* 19 (pp. 239–246). Springer International Publishing.
- [76] Rafael-Palou, X., Aubanell, A., Bonavita, I., Ceresa, M., Piella, G., Ribas, V., & Ballester, M. A. G. (2021). Re-identification and growth detection of pulmonary nodules without image registration using 3D Siamese neural networks. *Medical image analysis*, 67, 101823.
- [77] Angelopoulos, A. N., Bates, S., Fisch, A., Lei, L., & Schuster, T. (2022). Conformal risk control. *arXiv preprint arXiv:2208.02814*.
- [78] Armato III, S. G., McLennan, G., Bidaut, L., McNitt-Gray, M. F., Meyer, C. R., Reeves, A. P., ... & Clarke, L. P. (2011). The lung image database consortium (LIDC) and image database resource initiative (IDRI): a completed reference database of lung nodules on CT scans. *Medical physics*, 38(2), 915–931.
- [79] Nawaz, M. S., Shoaib, B., & Ashraf, M. A. (2021). Intelligent cardiovascular disease prediction empowered with gradient descent optimization. *Heliyon*, 7(5).
- [80] Bernard, O., Lalande, A., Zotti, C., Cervenansky, F., Yang, X., Heng, P. A., ... & Jodoin, P. M. (2018). Deep learning techniques for automatic MRI cardiac multi-structures segmentation and diagnosis: is the problem solved?. *IEEE transactions on medical imaging*, 37(11), 2514–2525.

- [81] *Isensee, F., Jaeger, P. F., Full, P. M., Wolf, I., Engelhardt, S., & Maier-Hein, K. H. (2018). Automatic cardiac disease assessment on cine-MRI via time-series segmentation and domain specific features. In Statistical Atlases and Computational Models of the Heart. ACDC and MMWHS Challenges: 8th International Workshop, STACOM 2017, Held in Conjunction with MICCAI 2017, Quebec City, Canada, September 10-14, 2017, Revised Selected Papers 8 (pp. 120-129). Springer International Publishing.*
- [82] *Cetin, I., Sanroma, G., Petersen, S. E., Napel, S., Camara, O., Ballester, M. A. G., & Lekadir, K. (2018). A radiomics approach to computer-aided diagnosis with cardiac cine-MRI. In Statistical Atlases and Computational Models of the Heart. ACDC and MMWHS Challenges: 8th International Workshop, STACOM 2017, Held in Conjunction with MICCAI 2017, Quebec City, Canada, September 10-14, 2017, Revised Selected Papers 8 (pp. 82-90). Springer International Publishing.*
- [83] *Khened, M., Alex, V., & Krishnamurthi, G. (2018). Densely connected fully convolutional network for short-axis cardiac cine MR image segmentation and heart diagnosis using random forest. In Statistical Atlases and Computational Models of the Heart. ACDC and MMWHS Challenges: 8th International Workshop, STACOM 2017, Held in Conjunction with MICCAI 2017, Quebec City, Canada, September 10-14, 2017, Revised Selected Papers 8 (pp. 140-151). Springer International Publishing.*
- [84] *Qin, C., Wang, S., Chen, C., Qiu, H., Bai, W., & Rueckert, D. (2020). Biomechanics-informed neural networks for myocardial motion tracking in MRI. In Medical Image Computing and Computer Assisted Intervention–MICCAI 2020: 23rd International Conference, Lima, Peru, October 4–8, 2020, Proceedings, Part III 23 (pp. 296-306). Springer International Publishing.*
- [85] *Zhang, X., You, C., Ahn, S., Zhuang, J., Staib, L., & Duncan, J. (2022, September). Learning correspondences of cardiac motion from images using biomechanics-informed modeling. In International Workshop on Statistical Atlases and Computational Models of the Heart (pp. 13-25). Cham: Springer Nature Switzerland.*
- [86] *Huang, S. C., Pareek, A., Jensen, M., Lungren, M. P., Yeung, S., & Chaudhari, A. S. (2023). Self-supervised learning for medical image classification: a systematic review and implementation guidelines. NPJ Digital Medicine, 6(1), 74.*
- [87] *Umirzakova, S., Ahmad, S., Khan, L. U., & Whangbo, T. (2024). Medical image super-resolution for smart healthcare applications: A comprehensive survey. Information Fusion, 103, 102075.*
- [88] *Ahishakiye, E., Bastiaan Van Gijzen, M., Tumwiine, J., Wario, R., & Obungoloch, J. (2021). A survey on deep learning in medical image reconstruction. Intelligent Medicine, 1(03), 118-127.*

List of figures

Figure 1: Overview of major medical imaging modalities and their primary clinical applications.....	21
Figure 2: A convolutional block.....	22
Figure 3: A simple CNN for disease classification from MRI images.....	23
Figure 4: ViT model overview.....	24
Figure 5: Swin Transformer hierarchical strategy versus ViT.....	24
Figure 6: Overview of the MedCLIP framework.....	26
Figure 7: Interconnected AI workflows in medical imaging. Segmentation, detection, classification, and registration interact closely in medical imaging pipelines. MAS: multi-atlas segmentation Feat. Count: features count.....	27
Figure 8: U-Net architecture.....	28
Figure 9: TransUNet architecture.....	29
Figure 10: YOLO architecture.....	30
Figure 11: VoxelMorph architecture.....	31
Figure 12: Deep Cascaded Registration Framework.....	32
Figure 13: Schematic representation of a Generative Adversarial Network (GAN).....	33
Figure 14: Left: Axial view of the various segmented lung lobes. Right: 3D view of the segmentation of all organs presents in the CT volume.....	40
Figure 15: Left: Different Lung Anomalies Visible on CT Images - from Left to Right: Normal Parenchyma (NP), Ground Glass (GG), Reticular (RETIC), Solid Nodule (NOD). Right: example of a solid lung nodule.....	41
Figure 16: Initial identification of malignant lung lesions (left) and its subsequent re-identification in follow-up analyses (right).....	42
Figure 17: Main features of the second nodule detection algorithm based on RetinaNet.....	42
Figure 18: Schema of the initial module for identifying lung lesions in the follow-up analysis.....	43
Figure 19: Example of how the nodule re-identification algorithm in the follow-up selects candidates.....	43
Figure 20: Follow-up of nodule's progression using image registration.....	44
Figure 21: Sensitivity per scan vs. false positives per scan for four LIDC-IDRI subsets. CRC (green) meets the 90% sensitivity target more reliably than the Naive strategy (red), especially in challenging cases.....	45
Figure 22: Online Viewer.....	47
Figure 23: Overview of the biomechanics-driven framework.....	50

Figure 24: Segmentation using multi-frame propagation.....	51
Figure 25: Labels used to compute the local biomechanical features.....	51
Figure 26: Confusion matrices of the classification on the training set and testing set.....	52
Figure 27: Accuracy as a function of the training set size.....	53

List of tables

Table 1: Average Dice scores on the three anatomical labels at ES and ED using registration ED to ES frames, and vice versa, multi-frame registration, and nnU-Net..... 50

Table 2: Classification accuracy over the train and test set for different classification methods..... 52

Getting in touch with the EU

In person

All over the European Union there are hundreds of Europe Direct centres. You can find the address of the centre nearest you online (european-union.europa.eu/contact-eu/meet-us_en).

On the phone or in writing

Europe Direct is a service that answers your questions about the European Union. You can contact this service:

- by freephone: 00 800 6 7 8 9 10 11 (certain operators may charge for these calls),
- at the following standard number: +32 22999696,
- via the following form: european-union.europa.eu/contact-eu/write-us_en.

Finding information about the EU

Online

Information about the European Union in all the official languages of the EU is available on the Europa website (european-union.europa.eu).

EU publications

You can view or order EU publications at op.europa.eu/en/publications. Multiple copies of free publications can be obtained by contacting Europe Direct or your local documentation centre (european-union.europa.eu/contact-eu/meet-us_en).

EU law and related documents

For access to legal information from the EU, including all EU law since 1951 in all the official language versions, go to EUR-Lex (eur-lex.europa.eu).

EU open data

The portal data.europa.eu provides access to open datasets from the EU institutions, bodies and agencies. These can be downloaded and reused for free, for both commercial and non-commercial purposes. The portal also provides access to a wealth of datasets from European countries.

Science for policy

The Joint Research Centre (JRC) provides independent, evidence-based knowledge and science, supporting EU policies to positively impact society



EU Science Hub

[Joint-research-centre.ec.europa.eu](https://joint-research-centre.ec.europa.eu)



Publications Office
of the European Union