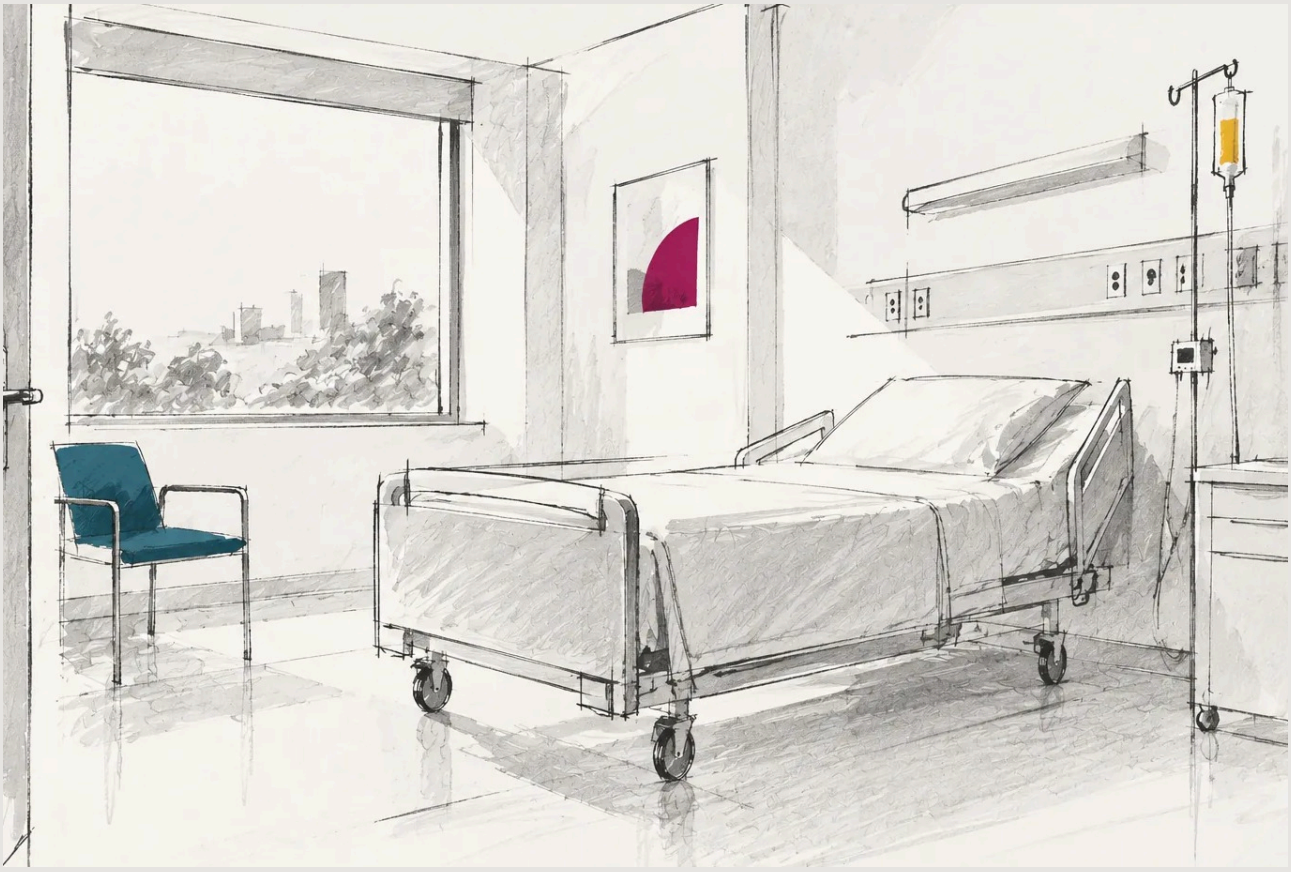


HEALTH AI IN THE EU · REPORT 1 OF 2

From Funding Rounds to Hospital Wards

The State of Health AI in the EU

RESEARCH

Júlia Keserű, DataBody Futures

IN PARTNERSHIP WITH

Civil Liberties Union for Europe (Liberties)

Contributions 3
Executive summary 4
Context 12
Methodology 15
Chapter 1: The EU's health AI ecosystem 19
1.1 Identifying leading EU health AI vendors and explicit AI procurements 21
1.2 Who supplies health AI in the EU? 23
1.3. Key revenue streams for EU health AI companies 30
1.4. Who funds the EU's health AI ecosystem? 33
Chapter 2: Compliance: expectations and reality 40
2.1. Regulatory background 40
2.2. How these laws interact, where they create tension 51
2.3. Compliance in practice: shadow AI, literacy and exemptions 60
Limitations of our analysis 63
Conclusions 65



Contributions

This research was led by Júlia Keserű, founder of DataBody Futures, who designed the study, conducted most of the analysis, and wrote the report. It was carried out in partnership with the Civil Liberties Union for Europe (Liberties), whose legal expertise shaped the analysis of the EU's regulatory framework and the information requests to hospitals in nine EU countries.

Éva Simon (Liberties) was involved in every stage of the project, from research design and legal analysis to review and editing. Tamás Szémann cleaned and structured the datasets and produced the data visualisations. Lucy Purdon designed and coordinated the data subject access requests and analysed the responses. Lindsay Ferris contributed research on the US health AI market and policy context.

We are deeply grateful to Liberties' member organisations and other local advocacy and journalist groups, who supported the hospital freedom-of-information requests in Belgium, the Czech Republic, Germany, Hungary, Italy, Malta, the Netherlands, Romania, and Spain, namely Apache.be, League of Human Rights, GFF, K-Monitor, CILD, Amphora Media, PILP, Apador and XNet. Their knowledge of national transparency laws and institutions made it possible to request information from hospitals across the EU. We also thank the volunteers who filed data access requests in their own names, and the more than 50 hospital professionals, developers, researchers and advocates who shared their time and expertise in interviews.

This research was supported by The Collaborative for AI in the Public Interest. The views expressed in this report are those of the authors.

Executive summary

Artificial intelligence (AI) tools are entering EU hospitals at every layer of the healthcare system, and developing them requires vast amounts of data from health records, wearables, mobile health apps and other sources. Health is also one of the EU's priority sectors for AI: the European Commission aims to invest in distinctly European health AI as part of a wider agenda that treats technological sovereignty as a prerequisite for Europe's independence, security and global competitiveness.

This report asks what a distinctly European health AI agenda looks like in practice. How does the push to expand computing power and compete in the global AI 'race' fit with the legal and practical realities hospitals face when choosing AI tools, and with the fundamental rights EU law grants patients?

To answer this, we test the claims behind the EU's sovereignty and simplification agendas within its health AI ecosystem. This report, the first of two, follows the money: using investment, procurement and grant data, it maps who builds, funds and supplies the AI systems that reach EU hospitals. It then examines how the law works in practice, by looking at what the EU's overlapping rules on data, medical devices and AI expect of hospitals, where they conflict, how hospitals, developers and legal experts experience them, and whether they support or undermine patients' rights.

A second report, to be published in January 2027, will go inside hospitals. Drawing on freedom-of-information requests in nine EU countries, it will examine how hospitals select, govern and deploy AI. Do they have AI strategies, and how do these relate to what vendors offer? Which risk assessments, audits and safeguards do they use? It will then follow patients' data, asking what the shift towards AI-enabled care means for consent, transparency and patients' control over their data, and what patients themselves expect.

Key findings

FINDING 01

The EU's sovereignty problem lies less in innovative potential than in scale, ownership and infrastructure.

The EU does not lack health AI start-ups, however, it does lack the capacity to scale them. Compared with the EU, the United States has roughly twice as many health AI companies, but five times as many that have raised at least USD 50 million, twenty times as many unicorns, and nearly six times the capital raised. This gives US companies far greater capacity to fund growth, including expansion into the European market.

European investment, meanwhile, remains fragmented. European investors hold most positions in the EU's best-funded health AI companies, but the vast majority back only a single company, and the few who invest repeatedly are mostly French or public. There is, in effect, hardly any cross-border group of specialist health AI investors capable of carrying EU companies to scale.

Where US capital does enter, it appears strategic. Most US investors in EU health AI companies are corporates, accelerators and industry specialists rather than independent venture funds. This suggests an interest in products, data and potential acquisitions rather than in financial returns alone. EU private capital, by contrast, tends to stay at home, with billionaire and family investors backing around a third of the leading companies, almost always in their own country. Whether this reflects sovereignty goals or simple home bias cannot be determined from the data.

For hospital AI, the single market barely exists. Almost all EU vendors with explicit AI contracts sell in only one Member State, and the few that operate across borders are mostly American. In fact, many core health AI systems come from US companies: according to interviewees, US firms supply much of the infrastructure hospitals depend on most, from electronic health records to cloud services.

KEY FACTS

- Compared with the EU, the US has nearly 2X as many health AI companies, 5X as many raising USD 50m+, 20X as many unicorns (40 vs 2) and ~6X the capital (USD 25.6bn vs 4.5bn).
- EU-based investors hold 79% of investor positions in the 29 best-funded EU vendors, but 85% of investors back only one of them, and just 17 back three or more – 12 of them French.
- Billionaire and family investors back a third of the 29 vendors, almost always in their home country.
- Only five vendors win explicit AI contracts in more than one Member State, and three of them are American; 96% of EU vendors sell in just one.

READ MORE

1.2 Who supplies health AI in the EU? · p. 23 **1.4** Who funds the EU's health AI ecosystem? · p. 33

FINDING 02

Public money seeds European health AI, but the growth, ownership and markets of these companies may lie outside the EU

Public funding plays a real but limited role in Europe's health AI ecosystem. EU research and innovation grants barely reach the leading companies, as only a handful of the best-funded EU health AI vendors received Horizon Europe or Digital Europe grants in the past years, and the total was a tiny fraction of the private capital these companies have raised. Public investors are more visible as shareholders, appearing in most of the leading companies. Several of the best-funded companies, including four of the top ten, have no recorded public investors at all.

Public and private money also back different kinds of companies. Public investors tend to support deep tech: hardware and research spin-outs that require new science and long certification processes, such as AI-driven critical care, cardiac imaging, brain implants or blood-based cancer detection. Private capital, by contrast, favours software that can scale quickly, such as ambient scribes and remote-monitoring tools. Public funding thus takes on the slower, riskier and more capital-intensive end of innovation, while private investors back products with faster routes to market. These are two separate funding tracks rather than one pipeline from public seed funding to private growth.

Public funding also rarely translates into public adoption: of the EU's best-funded health AI companies, only one appears in the above-threshold EU hospital AI procurement notices that explicitly name AI, with a single small contract. Europe's public investment in health AI therefore seeds companies whose growth, ownership and markets may lie elsewhere.

KEY FACTS

- Only 6 of 29 vendors received Horizon Europe or Digital Europe grants (2021–2025), worth about €9.6m in total, against USD 4.5bn raised.
- 4 of the 10 best-funded vendors have no recorded public investors.
- Only 1 of the 29 top-funded vendors appears in explicit health AI notices TED, with a single €42,400 contract.

READ MORE

1.3 Key revenue streams for EU health AI companies · p. 30 | **1.4** Who funds the EU's health AI ecosystem? · p. 33

FINDING 03

Hospitals are responsible for AI they often don't control.

AI features increasingly arrive in EU hospitals as bundled updates to existing systems, such as major radiology equipment or electronic health record platforms. Switching them off can mean losing vendor support or core functionality, so adoption is often driven by vendors' product decisions rather than by hospitals' own AI strategies.

In the above-threshold procurements we analysed, very little AI is bought explicitly. Documentation and administrative AI in particular tends to arrive inside broader software contracts, so it barely appears in procurement data, and hospitals do not always recognise such features as AI at all.

Where AI is bought openly in EU public procurements, contracts tend to go to incumbents, mainly large imaging manufacturers, rather than to AI specialists, while the EU's best-funded health AI start-ups rarely reach public hospitals. According to our interviews, US companies dominate the core systems on which hospitals depend most, from electronic health records to cloud infrastructure, but these rarely show up as explicit AI purchases.

Yet while vendors remain responsible for certifying their products, EU law places much of the responsibility for how these features are used on hospitals, even for tools they did not deliberately choose to adopt. Hospitals are expected to oversee tools that vendors may change without notice, rely on documentation and incident information they may struggle to obtain, and accept vendors' own descriptions of what a tool is for and how risky it is. When something goes wrong, responsibility is spread across vendors, hospitals and clinicians in ways that leave patients with little clear redress.

Some well-resourced academic hospitals have responded by building their own tools, which gives them greater control over their data and models. For most EU hospitals, however, this is not a realistic option, and dependence on vendors is a structural condition rather than a choice.

KEY FACTS

- Only 94 above-threshold Tenders Electronic Daily (TED) notices (November 2023 – October 2026) named AI explicitly.
- The most widely adopted vendors are large imaging manufacturers: GE HealthCare, Siemens Healthineers and Philips won explicit AI contracts from the most hospital buyers.
- According to interviewees, US companies supply many of the core systems hospitals depend on, from electronic health records to cloud services, but these rarely appear as explicit AI purchases.
- Interviewees also said documentation and administrative AI usually arrives inside broader software contracts, such as electronic health record platforms, and is rarely tendered or labelled as AI.

READ MORE

1.2 Who supplies health AI in the EU? · p. 23

2.2 How these laws interact, where they create tension · p. 51

FINDING 04

Shadow AI: the patient data route no one sees

One way patient data leaves EU hospitals is growing fast and almost entirely invisible: “shadow AI”, meaning clinicians using commercial generative AI tools such as ChatGPT or Gemini without institutional approval. Our interviewees named it as their most prominent concern. Its roots lie less in carelessness than in documentation burden and burnout, since doctors routinely finish their notes after hours, in so-called ‘pyjama time’, sometimes taking patient information home, and commercial chatbots offer an easy workaround, particularly where approved tools are lacking or slow to arrive.

When a patient record is entered into such a tool, it may be processed outside the EU by a provider with no data processing agreement with the hospital. Hospitals typically prohibit such use on paper but acknowledge that they cannot control what staff do on personal devices.

Since prohibition alone is unlikely to work, interviewees saw AI literacy as a prerequisite for safer use and stronger patient protection. Yet the legal incentive to provide it has weakened: most recent changes to the EU Artificial Intelligence Act (AI Act) require hospitals only to take measures that support AI literacy, not to ensure their staff reach a sufficient level. Evidence on what hospitals actually do is scarce: we found hardly any studies evaluating hospital AI training, and none showing whether it changes behaviour or reduces shadow AI use.

KEY FACTS

- In a US survey, nearly one in five healthcare professionals reported using unapproved AI tools, and one in ten had used them for direct patient care. Similar data does not exist for the EU but interviews confirmed that it's widespread.
- EU hospitals told us they prohibit entering patient data into unapproved AI tools but cannot control what staff do on personal devices.
- Since the Digital Omnibus, the AI Act no longer requires hospitals to guarantee any level of AI literacy among staff, only to take measures that support it.
- Some hospitals are filling the gap themselves, for example by making AI literacy training a condition for staff access to its AI tools.

READ MORE
2.3 Compliance in practice: shadow AI, literacy and exemptions · p. 60

FINDING 05

AI is meant to make healthcare more equal, but in the EU it risks deepening existing divides.

One key promise of health AI is that it can extend high-quality care to more patients in the EU, however, the resources needed to develop and govern it are unevenly distributed. Investment in the EU's best-funded AI health companies is concentrated in the western and northern Member States, with France alone being home to nearly half of them and attracting more than 40% of their funding, while Sweden accounts for almost a third through just two firms. None of these companies are based in Central or Eastern Europe, the AI built in the EU therefore risks being shaped mainly by the priorities, patient populations and health systems of a few wealthier Member States.

Regulation can seriously reinforce this gap. The European Health Data Space, for instance, will require standardised, interoperable data from hospitals as data holders, which those with established data platforms can readily provide. Smaller institutions, some still partly paper-based, may struggle to meet even basic requirements, let alone benefit from the system as data users.

Inequality also accumulates inside the models themselves. Training data often over-represents some patient groups, such as male or white patients, and under-represents others, so tools can perform worse for the patients who are already least well served. Detecting such bias requires sensitive attributes such as ethnicity, sex or socioeconomic status, yet existing EU regulation requires data minimisation which tends to strip out precisely this information.

To further complicate matters, if some groups opt out of secondary data use more than others, AI models training that builds on this data may also be least representative of the patients most at risk from biased models. Without deliberate efforts to close these gaps, AI could entrench, rather than reduce, the inequalities it is meant to address.

KEY FACTS

- France has 14 of 29 best-funded vendors and 43% of their funding. Sweden has 31% of funding through two companies (Neko Health and Tandem Health). None of the 29 top-funded health AI firms are based in Central or Eastern Europe.
- Interviewees said many hospitals are far from ready for data standardisation, and some still keep parts of their records on paper.
- EU laws pull in opposite directions on bias. The AI Act requires training data to be representative so that bias can be caught, while data protection rules push organisations to collect as little sensitive data as possible.

READ MORE

1.2 Who supplies health AI in the EU? · p. 23

2.2 How these laws interact, where they create tension · p. 51

FINDING 06

EU rules on health AI are so layered that hospitals struggle to use it well, and protections meant for patients can end up working against them.

Hospitals deploying AI must navigate a dense web of overlapping laws, from data protection and medical device rules to the AI Act and new rules on health data sharing. Each was written at a different time, for a different purpose and with its own idea of risk. Rather than reinforcing one another, these rules often lead to duplicated work, with some hospitals running separate compliance projects for each. Ongoing reforms, such as the EU's Digital Omnibus, add to the uncertainty, leaving hospitals to prepare for obligations whose shape may change before they apply.

The protection patients receive often depends less on the real risk a tool poses than on how it is classified, a judgement that rests largely with the companies that make it. As a result, an AI system that decides whether patients are eligible for publicly funded care must meet the AI Act's full set of rights safeguards, while a certified stroke-detection algorithm may not have to meet any of them.

Oversight is just as fragmented as several regulators supervise one part of the same system, and a single cyber incident involving an AI tool can trigger several separate reports to different authorities. Hospitals' own roles are also unstable: adapting a vendor's tool can turn a hospital into its legally responsible maker, while tools labelled as 'research' can remain outside the AI Act's obligations even when they are clearly headed for clinical use.

Overall, hospitals take pride in Europe's strong protections, yet some also see benefits forgone, such as access to large international datasets for rare diseases. Meanwhile, the safeguard many regard as most essential, training staff to use AI safely, has been weakened just as it matters most.

KEY FACTS

- Some hospitals handle each EU law as a separate compliance project, with separate teams doing much the same work.
- Interviewees said some developers deliberately stay just outside the rules for medical devices, and warned that AI note-taking tools risk acting as medical devices without being certified as such.
- Hospitals judge the risk of similar AI tools differently, often misread the rules, and tend to choose lower risk categories where they can.
- Some hospitals describe AI tools as ‘purely research’ even when they plan to use them in patient care.

READ MORE

2.1 Regulatory background · p. 40 **2.2** How these laws interact, where they create tension · p. 51

Context

Artificial Intelligence (AI) tools are entering European Union (EU) hospitals at every layer of the healthcare system.[1] They arrive through diagnostic models that analyse medical images, risk-predicting systems that flag patients whose health is likely to deteriorate, treatment planning software that supports clinical decision-making, ambient scribes that turn doctor-patient conversations into structured clinical notes, administrative tools that manage billing and capacity, and remote monitoring platforms that track patients beyond hospital walls.

One of the key drivers of AI's integration into healthcare is the **increasing demand for interoperability across diverse data sources**. This includes linking electronic health records (EHRs) with wearable devices and genetic testing services, and creating cross-border health data platforms such as the European Health Data Space.[2] To develop and refine these AI models, vast amounts of data are collected from various inputs, including digital health records, wearables, mobile health apps, and other avenues.

In her 2026 State of the Union address, President of the European Commission, Ursula von der Leyen named five 'high value' sectors where industrial AI has the greatest potential for the EU: health, transport, agri-food, advanced manufacturing, and defence and space.[3] On health, she argued that Europe needs AI that supports, rather than replaces, doctors, **calling for health AI development that is distinctly European**. In the same speech, von der Leyen called for boosting the Union's computing capacity and for new ways to use public and private funds to invest in the most promising European start-ups and companies.

Her remarks came against the backdrop of an **ambitious EU agenda that treats technological sovereignty as a prerequisite for European independence, security, and global competitiveness**. Since the 2024 Draghi report[4] warned that Europe is much better at inventing than scaling, the Commission has shifted its approach from setting rules to building more capacity for a robust AI ecosystem. In 2025, this shift materialised in a series of high-level political commitments, including the InvestAI initiative which aims to mobilise 200 billion EUR for AI investments,[5] the AI Continent Action Plan, which set out to expand Europe's computational infrastructure,[6] and the ApplyAI

1. WHO Regional Office for Europe, "New WHO/Europe Report Provides First-Ever Snapshot of AI in Health Care across European Union Member States," news release, April 20, 2026, https://www.who.int/europe/news/item/20-04-2026-new-who-europe-report-provides-first-ever-snapshot-of-ai-in-health-care-across-european-union-member-states.
2. Regulation (EU) 2025/327 of the European Parliament and of the Council of 11 February 2025 on the European Health Data Space and Amending Directive 2011/24/EU and Regulation (EU) 2024/2847, *Official Journal of the European Union* L, 2025/327 (March 5, 2025), http://data.europa.eu/eli/reg/2025/327/oj (hereafter EHDS Regulation).
3. European Commission, *State of the Union 2026: Quotes and Initiatives* (Brussels: European Commission, 2026), https://www.eeas.europa.eu/sites/default/files/2026/documents/State_of_the_Union_2026_-_Quotes_and_initiatives.pdf.pdf.
4. Mario Draghi, *The Future of European Competitiveness* (Brussels: European Commission, 2024), https://commission.europa.eu/topics/competitiveness/draghi-report_en.
5. European Commission, "EU Launches InvestAI Initiative to Mobilise €200 Billion of Investment in Artificial Intelligence," press release, February 11, 2025, https://digital-strategy.ec.europa.eu/en/news/eu-launches-investai-initiative-mobilise-eu200-billion-investment-artificial-intelligence.
6. European Commission, *AI Continent Action Plan*, COM(2025) 165 final (Brussels, April 9, 2025), https://digital-strategy.ec.europa.eu/en/library/ai-continent-action-plan.

Strategy, which plans to speed up 'trustworthy AI' adoption across priority sectors, including health.[7]

But what does a distinctly European, responsible and truly sovereign health AI agenda look like in practice? How does the push to expand computing power and compete in the global AI 'race' fit with the legal and practical realities hospitals face when deciding on AI tools, or with the fundamental rights EU law grants the Union's patients, including the right to privacy and data protection, non-discrimination, equal access to quality care, and more?

For years, the EU prided itself on having the world's most comprehensive, rights-based approach to regulating technology. Sovereignty was invoked in defence of strong rules: Europe protecting its citizens' data and rights from foreign platforms and dependence. Since the Draghi report and the EU's shift towards competitiveness, however, the same laws are increasingly treated as a geopolitical liability that holds Europe back in the global AI race. This has led to a far-reaching simplification drive: through the Digital Omnibus, the Commission has begun to revise the GDPR, the AI Act and other digital laws, simplifying and in places loosening rules that were designed to make technology more responsible and accountable.[8]

These policy and regulatory shifts rest on claims that have rarely been tested in specific parts of the AI ecosystem, namely that EU money can build EU AI champions, that strong rules make AI trustworthy, and that loosening them will make Europe more competitive without costing citizens their rights. **Healthcare is a great place to test these assumptions because both the promise of getting things right, as well as the costs of getting things wrong, are particularly high.**

The reasons for that are plentiful. Health data has become a prime target in the past decade, particularly for hackers involved in fraudulent activities, such as medical identity theft or extorting money from providers, causing substantial financial loss for healthcare providers, and often disrupting medical services on a massive scale.[9] Hallucinations have also become a pervasive problem for AI models, particularly generative AI, which can produce outputs that sound plausible but are factually wrong, logically inconsistent or not grounded in reliable sources.[10] In the healthcare setting, such errors have already contributed to incorrect diagnoses or treatment decisions and dosing mistakes,[11] and could easily become a serious liability for the entire sector.

Against this backdrop, **our research tests the claims behind the EU's sovereignty and simplification agendas within the EU's health AI ecosystem**. This report, the first of two, begins by mapping who builds, funds and supplies the AI systems that reach EU hospitals through investment, procurement and grant data; it then turns to how the law works in practice: what the EU's overlapping

7. European Commission, *Communication from the Commission: Apply AI Strategy*, COM(2025) 512 final (Brussels, October 8, 2025), https://digital-strategy.ec.europa.eu/en/library/commission-communication-apply-ai-strategy.
8. European Commission, *Proposal for a Regulation of the European Parliament and of the Council Amending Regulations (EU) 2016/679, (EU) 2018/1724, (EU) 2018/1725, (EU) 2023/2854 and Directives 2002/58/EC, (EU) 2022/2555 and (EU) 2022/2557 as Regards the Simplification of the Digital Legislative Framework, and Repealing Regulations (EU) 2018/1807, (EU) 2019/1150, (EU) 2022/868 and Directive (EU) 2019/1024 (Digital Omnibus)*, COM(2025) 837 final, 2025/0360(COD) (Brussels, November 19, 2025), https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52025PC0837 (hereafter Digital Omnibus proposal).
9. European Union Agency for Cybersecurity (ENISA), *ENISA Threat Landscape: Health Sector* (Athens: ENISA, July 2023), https://www.enisa.europa.eu/publications/health-threat-landscape.
10. Ziwei Ji et al., "Survey of Hallucination in Natural Language Generation," *ACM Computing Surveys* 55, no. 12 (2023): 248, https://doi.org/10.1145/3571730.
11. Paige Cockburn, "Error by AI Scribe during Medical Appointment Leaves Patient Devastated," *ABC News*, August 14, 2026, https://www.abc.net.au/news/2026-08-14/ai-medical-scribe-error-leaves-patient-devastated/107031672.

rules on data, medical devices and AI expect of hospitals, where they conflict, how hospitals, developers and legal experts experience them on the ground, and whether they support or undermine patients' rights in practice.

A second report, to be published in January 2027, will turn to practice. Drawing mainly on freedom-of-information requests to hospitals in nine EU countries, it will first examine how hospitals select, govern and deploy AI in practice, whether they have AI strategies, how these relate to what vendors offer, and which risk assessments, audits and safeguards they use, particularly for privacy and data protection. Finally, it will follow patients' data by asking what the shift towards AI-enabled care means for consent, transparency and patients' control over their data, how patient data is used to develop AI, how open companies and hospitals are about it, and what patients actually expect.

Methodology

For this analysis, we use the EU AI Act's definition of artificial intelligence: a machine-based system that operates with some degree of autonomy and infers from the input it receives how to generate outputs such as predictions, content, recommendations or decisions. As we explain throughout the report, this definition is very broad and covers a wide range of technologies, many of which were previously described as machine learning, computer vision, decision support or statistical modelling, and are now grouped under the same umbrella term. This makes any analysis harder, especially for comparisons over time, and it significantly shapes our findings. Wherever possible, we therefore distinguish between the different types of AI being deployed in healthcare.

For health data, we follow the broad definition used in the European Health Data Space (EHDS) Regulation,[12] which covers data directly connected to health, such as electronic health records and genomic data, as well as data indirectly related to health, including information on insurance status, professional status, education, lifestyle, wellness and behaviour. However, our analysis zooms in on clinical AI tools – systems that reach clinicians and patients as part of healthcare delivery, such as diagnostics, risk prediction, treatment planning, clinical documentation, hospital logistics, patient-facing tools, remote monitoring and robotics. We deliberately exclude AI used for drug discovery and clinical trials, areas that often fall outside the AI Act's scope, and which would distort our market findings significantly.

This research uses a mixed-methods design, combining quantitative and qualitative sources to examine the same questions from different angles. On the quantitative side, we analysed structured data that show where money and contracts flow through investment data from Dealroom, procurement records from Tenders Electronic Daily, and EU grant data from the Community Research and Development Information Service (CORDIS) on Horizon Europe and Digital Europe grants. On the qualitative side, we drew on several sources, including our own legal analysis of the EU frameworks governing health AI, 36 interviews with more than 50 hospital officials, developers, legal experts, researchers and patient advocates, freedom-of-information requests to hospitals in nine EU countries, data access requests to six to eight of the largest health AI companies in six countries, hospital press releases, correspondence with health AI developers, and a literature review of relevant academic and policy sources.

Combining these approaches let us fill some of the gaps each source left and made it possible to check certain findings against one another. For example, interviews helped explain why AI rarely appears in procurement data, while the market data put interviewees' accounts of vendor dependence into a wider context. Similarly, the legal analysis set out what hospitals and companies are required to do, while freedom-of-information and data access requests tested whether those

12. EHDS Regulation.

obligations can be verified from the outside. Where independent sources pointed in the same direction, we treated the findings as more robust. Where they diverged, we reported the difference rather than resolve it.

The most highly valued and best-funded companies were identified through Dealroom, a commercial database of startups and scaleups. [13] Results combined Dealroom's *Health* industry and *Artificial Intelligence* technology tags and were restricted to companies headquartered in the EU27.[14] Companies were included if they met at least one of two criteria: (1) a valuation of USD 500 million or more, or (2) total equity funding of USD 50 million or more.[15] Each of the 46 companies returned was reviewed individually against its product offering, and 17 were excluded. These were mainly biotechnology and pharmaceutical firms using AI for drug development, or providers of tools for drug discovery and clinical trials rather than clinical care. The final Dealroom sample comprises 29 vendors. Although the database covers companies globally, the sample was restricted to EU-headquartered firms, since Dealroom records a company's location and funding history, but not the markets in which it operates.

To see which vendors win publicly advertised hospital contracts for health AI in the EU, public procurement data from Tenders Electronic Daily (TED), the EU's official procurement journal, were also analysed. [16] The dataset covers contract award notices published between 1 November 2023 and 12 October 2026 by hospitals and health-sector purchasing bodies in the EU27, with vendors from any country.[17] Since TED publishes only contracts above EU procurement thresholds, smaller purchases are not included.[18] Notices were retrieved mainly by combining CPV codes[19] with AI-related keywords in all EU languages.[20] We counted a contract as explicit AI procurement only if its notice named AI in full, as part of what was bought, for patient care or hospital operations, and identified the winning supplier.[21] This gave 94 notices, comprising 106 vendor awards.

13. Dealroom.co, "Health AI Companies, EU27 and US Headquartered," data set (filters: industry = Health; technology = Artificial Intelligence; HQ region = European Union / United States), retrieved September 17, 2026, https://app.dealroom.co.

14. See Dealroom.co, "Industries & Sub-industries," Dealroom Knowledge Base, https://dealroom.co/knowledge/Glossary/glossary-industries; Dealroom.co, "Technologies," Dealroom Knowledge Base, https://dealroom.co/knowledge/Glossary/glossary-technologies; and Dealroom.co, "Tech Taxonomy: Classifying Innovative Companies," Dealroom blog, March 17, 2023, https://dealroom.co/blog/dealroom-taxonomy.

15. Valuation reflects investors' expectations of a company's future market position; equity funding reflects the capital available to it for development and expansion, and captures companies that have raised substantially.

16. Publications Office of the European Union, *Tenders Electronic Daily (TED)*, contract award notices, November 1, 2023–October 12, 2026, accessed via TED Search API, https://ted.europa.eu; API documentation, https://docs.ted.europa.eu/api/latest/index.html.

17. The period began shortly after the eForms standard became mandatory for TED notices (25 October 2023), so that all notices share the same structure.

18. Contracts below the EU procurement thresholds need not be published in TED. For hospitals' supply and service contracts these are roughly EUR 140,000–221,000, depending on the type of contracting authority, and are revised every two years. Many AI purchases, including pilots and software licences, are likely to fall below them. See Directive 2014/24/EU of the European Parliament and of the Council of 26 February 2014 on Public Procurement and Repealing Directive 2004/18/EC, *Official Journal of the European Union* L 94 (March 28, 2014): 65–242, art. 4, http://data.europa.eu/eli/dir/2014/24/oj.

19. Commission Regulation (EC) No 213/2008 of 28 November 2007 Amending Regulation (EC) No 2195/2002 of the European Parliament and of the Council on the Common Procurement Vocabulary (CPV), *Official Journal of the European Union* L 74 (March 15, 2008), https://eur-lex.europa.eu/eli/reg/2008/213/oj/eng.

20. Keywords covered the terms "artificial intelligence", "machine learning" and "deep learning" or "neural network" in all official EU languages.

21. Each award was checked against four conditions: AI is named in the notice; it is part of what was bought; it is for patient care or hospital operations; and the winning supplier is identified. Classification was based on the notice text, with a machine-assisted first pass reviewed by the authors.

We chose this narrow definition for two reasons. First, it can be applied consistently and checked by others. Second, its error runs in one known direction: it undercounts AI rather than mistaking other products for AI. Consequently, this approach captures only the most visible part of hospital AI procurement, and does not capture notices that only hint at AI, for example through abbreviations or specific functions (many of which are likely to be AI systems); AI that arrives inside larger contracts for electronic health records, imaging equipment or other software; AI added through updates to existing systems; purchases below EU thresholds or published only on national portals; or purchases by private providers.

In our procurement sample, vendors were attributed to their ultimate parent company where this could be established, otherwise they were attributed to the registered country of the winning entity. Where legal and operational headquarters diverge, or an EU firm is under non-EU ownership, both were recorded. Since TED attributes contracts to the winning entity, which may be a distributor of AI developed elsewhere, the EU share reflects contract holders rather than the origin of the technology supplied. The TED findings should therefore be read as describing a small and specific segment of hospital AI procurement, not the market as a whole, and treated with caution.

For each of the 29 vendors, we recorded every investor listed in Dealroom. This gave 411 investor positions held by 317 distinct investors, after spelling variants and related entities were merged. Dealroom's investor labels are inconsistent and do not record investors' headquarters, so each investor was coded manually into one of eleven types, and by the country of its headquarters: for funds, the management company; for individuals, their nationality. Dealroom does not report how much each investor contributed, so the analysis counted investor positions, not the amounts invested.

To see how far EU research funding reaches the leading vendors, we also **matched the 29 companies against Horizon Europe and Digital Europe grant data from CORDIS**, covering grants signed in 2021-2025. EU funding databases do not systematically flag AI, and definitions of what counts as an AI project vary between programmes. The matching was therefore done at company level, including subsidiaries and name variants, rather than by searching for AI projects.

For the qualitative data collection, we used **three distinct transparency tools: freedom-of-information (FOI) requests to approximately 300 hospitals in nine EU countries**, about their AI use, strategies and governance documents; **GDPR-based data subject access requests (DSARs)** to the largest health AI companies operating in Europe, to test whether EU access rights can reveal if a person's health data was used to train AI models; and **emails to almost 400 EU-based health AI developers** about their data sourcing, training and privacy practices. These methods mainly inform the second report, to be published in January 2027. However, some of what they revealed, for example about how hospitals classify AI systems and handle vendor contracts, also informed the market and legal analysis in this report.

In order to confirm and contextualise findings from the quantitative data, to suggest new research directions, and to add practical detail to our narrative, we also **conducted 36 semi-structured interviews with more than 50 people**, including hospital AI programme and compliance officers;

health AI software developers; health data and privacy researchers; and patient advocates from the EU. Interviewees were identified through existing networks, referrals, conference contacts and cold contact.

Lastly, we **relied on existing academic and policy literature** to ground the legal and market analysis, to provide comparative evidence where our own data were thin, and to capture patient and clinician perspectives that could be set against hospital and national strategies. This included legal scholarship on EU frameworks, covering how the AI Act, the Medical Device Regulation, the General Data Protection Regulation (GDPR) and the European Health Data Space apply to health AI systems, and how they interact with each other.

We also reviewed existing research on the EU's health AI ecosystem, including economic studies of venture capital, as well as policy reports from the European Commission, the European Investment Bank, the Organisation for Economic Co-operation and Development (OECD) and the World Health Organization (WHO). Additionally, the literature review covered empirical studies on how AI is deployed in clinical practice, as well as surveys of patient and public attitudes towards AI in healthcare and the sharing of health data, including recent European patient surveys. Sources were identified through PubMed, Google Scholar and SSRN, as well as EUR-Lex for legal texts, and we prioritised peer-reviewed studies and official reports published since 2020, with a focus on the EU and its member states.

In line with our commitment to the responsible and ethical use of AI, which should support rather than replace human research, analysis and writing, we used generative AI tools only for supporting tasks within this research. These tasks comprised machine translation of documents and FOI responses, cleaning and structuring complex datasets, spotting errors and inconsistencies in text and data, as well as coding the accompanying data visualisations and project website. All analysis, interpretation and conclusions are the authors' own, and all AI-assisted outputs were reviewed and verified by the research and legal teams, as well as human editors.

CHAPTER 1

The EU's health AI ecosystem

Measuring AI spending and investment is a significant research challenge. As a 2025 OECD report observes, the field lacks standardised measurement frameworks and relies on incomplete datasets.[22]

Two factors in particular complicate efforts to quantify AI spending and investment.[23] The first one is structural: AI investment flows through a wide range of instruments, from public grants to private equity, many of which are difficult to track systematically. Furthermore, AI spending, in turn, is rarely recorded as a distinct category: it is typically absorbed into broader digitalisation budgets or grouped with adjacent technologies such as robotics. Since existing economic classifications were not designed with AI in mind, they offer no straightforward way of isolating AI-focused spending from adjacent funding categories.

The second challenge is terminological, but related, as the categories used to describe AI have changed considerably over time; until recently, AI-related systems were classified according to their underlying methods, such as machine learning (ML), natural language processing (NLP) or computer vision (CV). Since the early 2020s, as more countries have introduced explicitly AI-focused strategies, these method-specific categories have increasingly been replaced by a single, generic 'AI' label. While this simplifies reporting, it limits the comparability of data over time and makes it difficult to trace how AI investment has evolved.

Consequently, **there is no single, authoritative EU-wide figure for private sector spending on 'health AI'. Available evidence is fragmented** across EU grants, national health-system budgets, hospital procurement, and private investment, with most datasets not distinguishing AI or AI-enabled systems from broader digital health or hospital infrastructure initiatives.

In the absence of such analytical grounding, the most widely cited estimates come from commercial market reports covering the broader digital health sector. These consistently document steady growth in global funding for AI-enabled digital health solutions;[24] at the same time, they also indicate a shrinkage in market breadth: according to industry estimates, both deal volume and the number of

22. OECD, *Advancing the Measurement of Investments in Artificial Intelligence*, OECD Artificial Intelligence Papers, no. 47 (Paris: OECD Publishing, 2025), https://doi.org/10.1787/13e0da2f-en.
23. Zuzanna Warso and Julia Keseru, *From Frameworks to Factories: How the EU Funds AI* (Open Future, June 2026), https://openfuture.eu/wp-content/uploads/2026/06/260623FromFrameWorksToFactories.pdf.
24. See Galen Growth, "Digital Health Funding: Global Data & Trends," HealthTech Alpha data, last updated July 2026, accessed 9 September 2026, https://www.galengrowth.com/digital-health-funding/; CB Insights, *State of Digital Health 2025 Report* (New York: CB Insights, 22 January 2026), https://www.cbinsights.com/research/report/digital-health-trends-2025/; Digital.Health, "Digital Health + AI Reports & Whitepapers," last updated 11 August 2026, accessed 13 September 2026, https://digital.health/digital-health-reports-whitepapers.

funded companies have declined by 15-20% in the past years. Taken together, these trends suggest an increasing concentration of capital, characterised by fewer but larger transactions. However, these sources often need to be treated with scientific caution.

AI INTEGRATION INTO CLINICAL CARE

Where AI sits in the hospital

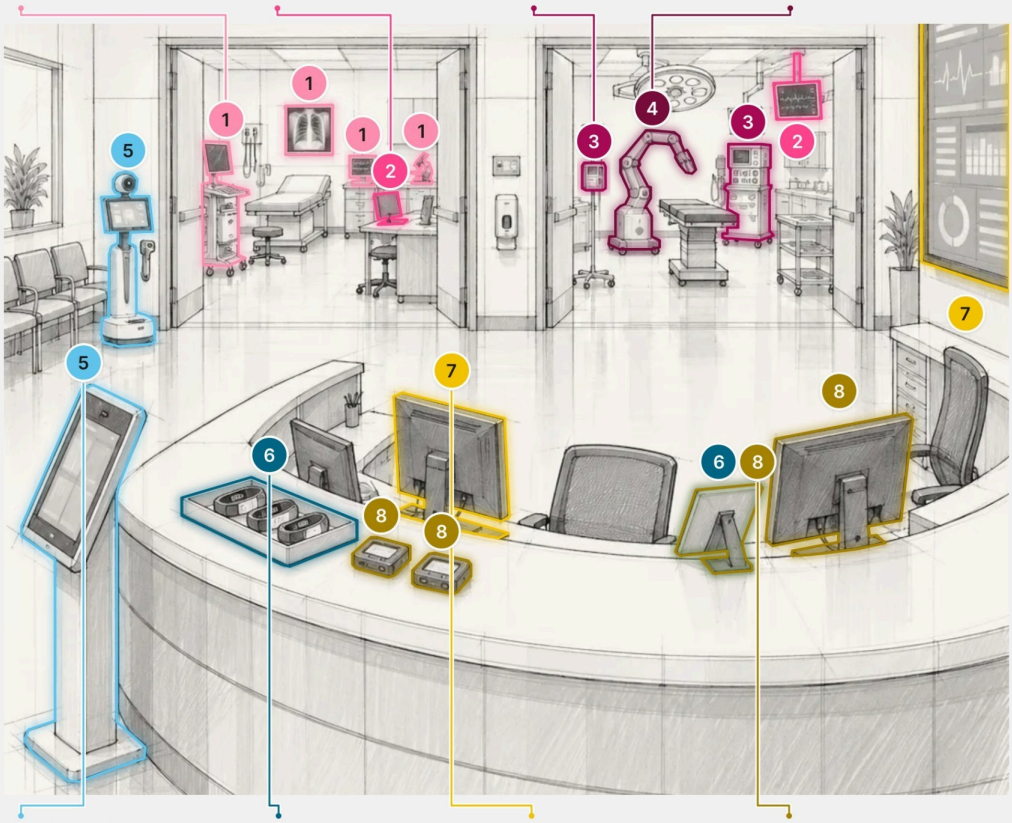
Clinical care

1 Diagnostics
AI-assisted lesion and tumour detection, screening and grading in radiology, mammography, dermatology, cardiology, pathology.

2 Risk prediction
EHR-based early warning for sepsis, deterioration, readmission, plus guideline and protocol support.

3 Treatment support
Personalised therapy design, including radiotherapy planning, pre-operative planning and precision dosing.

4 Surgical robots
Robotic surgery, rehabilitation and therapy with AI for navigation, image guidance and instrument tracking.



Facing the patient

5 Patient-facing tools
Symptom checkers, triage bots, navigation and virtual assistants, chatbots, medication reminders and mental health support.

6 Remote monitoring
Wearable- and sensor-based monitoring outside hospital (e.g. heart failure) and prescribed app-based therapies.

Running the hospital

7 Logistics & admin
Scheduling, no-show prediction, bed and capacity management, staff rostering, and patient flow.

8 Documentation & ambient
Ambient scribing and LLM drafting of visit notes, orders, discharge summaries and patient messages.

source: euhealthai.org | visualization: Tamás Szémann

In this chapter, we draw on established and credible data sources, including Dealroom investment data, procurement records from Tenders Electronic Daily (TED), the OECD.ai observatory's live data, and research funding data from CORDIS, to untangle the key trends shaping the health AI market in the EU context. Specifically, we examine who the leading vendors and health AI unicorns are, where they originate, who funds them, to what extent the EU can be said to depend on non-EU providers, what publicly advertised, explicitly AI-focused contracts reveal about how such tools are procured, who the principal investors in EU health AI companies are, and how the EU's own AI funding strategy is applied in practice.

SECTION 1.1

Identifying leading EU health AI vendors and explicit AI procurements

Our analysis further distinguishes clinical AI providers by the role AI plays in their offering. Some vendors are AI-native: their core product is an AI system, such as a radiology triage algorithm. Others are AI-embedded: they sell core clinical or administrative systems, such as electronic health records, imaging equipment, patient monitoring platforms or hospital administration software, into which AI features, such as predictive models, image reconstruction and ambient documentation tools, are increasingly integrated. A third group are AI-enabling providers, which supply the infrastructure on which AI systems run, including cloud services, compute, foundation models and data platforms, without necessarily offering clinical AI products themselves.

This report focuses mainly on AI-native and AI-embedded vendors. Both deliver AI to clinicians and patients as part of a product, and their market position is specific to the health sector. AI-enabling providers, by contrast, operate across all sectors, and their role in health AI cannot be separated from their wider market offer.

Within this scope, a further distinction excludes companies that use AI without supplying it. Manufacturers of medical imaging and monitoring equipment, for example, supply hospitals with scanners, software and systems that contain AI functions used directly in clinical practice, such as AI-assisted image reconstruction and automated detection. In these cases, AI reaches the hospital as part of the product. A biotechnology company developing antibody therapeutics, by contrast, may

use AI extensively in its research, but the product reaching hospitals and patients is a medicine, not an AI system. Given how widely AI is now used internally across all sectors, including such companies would render the category analytically meaningless.[25]

AI tools can now reach almost every layer of a hospital's work, although the extent of their use varies widely between hospitals and functions, as our later chapters on hospital AI use describe. Treating health AI as a single market would therefore obscure important differences: the vendors, investors and dependencies behind an imaging algorithm can differ sharply from those behind an ambient documentation tool or a surgical robot.

All our data sources were therefore mapped onto eight functional categories, built on the seven categories of AI application in health care used by WHO/Europe.[26] This layer of segmentation makes it possible to assess dependence segment by segment rather than treating health AI as a single market.

CATEGORY	DEFINITION AND AREAS OF APPLICATION
1. AI-assisted diagnostics	Tools used mainly for lesion and tumour detection, cancer screening support, triage of urgent findings, and grading and classification. Applications include radiology, mammography, dermatology, cardiology (ECG and echo analysis), ophthalmology, digital pathology and laboratory medicine. These systems mostly rely on computer vision or signal-processing models.
2. Risk prediction and clinical decision support	Tools that analyse structured and unstructured Electronic Health Record (EHR) data, laboratory results and vital signs to support clinical decisions, such as early warning of sepsis or patient deterioration, readmission risk, injury prediction, and guideline or protocol support. These are typically machine-learning or statistical models integrated into EHR workflows.
3. Treatment planning and therapy support	Tools that help design or personalise therapy rather than diagnose, such as radiotherapy planning and dose optimisation in oncology, pre-operative planning, and precision dosing. These often combine imaging, genomic and longitudinal EHR data, and are frequently embedded in devices or treatment planning systems.
4. Clinical documentation and ambient AI	Tools that listen to or read the clinical encounter and draft notes, orders and summaries. They transcribe and structure clinician-patient conversations into visit notes, problem lists and orders, condense long records into discharge summaries, and draft follow-up messages and patient materials. Increasingly integrated directly into EHRs and telemedicine platforms.
5. Hospital logistics and	

25. This distinction also mirrors one drawn in the EU AI Act, to be described in greater detail in Chapter 2. The Act separates providers, which develop an AI system and place it on the market under their own name, from deployers, which use an AI system in the course of their own activities. Manufacturers of AI-enabled imaging and monitoring equipment act as providers, and where their devices require third-party conformity assessment under the Medical Device Regulation, they fall under the Act's high-risk requirements. A company using AI internally for research and development would qualify at most as a deployer, and AI used solely for scientific research may fall outside the Act altogether. See Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 Laying Down Harmonised Rules on Artificial Intelligence and Amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act), *Official Journal of the European Union* L, 2024/1689 (July 12, 2024), arts. 3(3)–(4) (definitions of provider and deployer), 6(1) and annex I, sec. A (high-risk systems covered by Union harmonisation legislation, including the MDR and IVDR), and 2(6) and (8) (exclusions for scientific research and for pre-market research, testing and development), http://data.europa.eu/eli/reg/2024/1689/oj (hereafter AI Act).

26. WHO Regional Office for Europe, *Artificial Intelligence Is Reshaping Health Systems: State of Readiness across the WHO European Region*, WHO/EURO:2025-12707-52481-81028 (Copenhagen: WHO Regional Office for Europe, November 19, 2025), https://www.who.int/europe/publications/i/item/WHO-EURO-2025-12707-52481-81028. Four WHO categories carry over directly: diagnostics, remote patient monitoring (extended here to digital therapeutics), surgery and medical robotics, and the automation of logistics and administrative tasks. WHO's prognosis prediction (risk stratification) becomes risk prediction and clinical decision support. WHO groups symptom checkers with support in treatment decisions; here, treatment decisions move to a separate category for treatment planning and therapy support, while symptom checkers join WHO's conversational platforms for patient assistance in a single category for patient-facing tools. Finally, a new category for clinical documentation and ambient AI is added, reflecting its rapid growth in recent years.

admin automation	Tools that optimise how hospitals run rather than how patients are treated, including scheduling and no-show prediction, bed and capacity management, length-of-stay forecasting, staff rostering, clinical coding, supply chain and patient flow. Rarely regulated as medical devices and often bundled into EHR or enterprise platforms.
6. Patient-facing tools	Tools that interact directly with patients or citizens at the front door of care, including symptom checkers and triage bots, appointment and navigation assistants, virtual assistants that answer questions or send medication reminders, and mental health tools offering CBT-style support or psychoeducation.
7. Remote monitoring and digital therapeutics	Tools that combine wearables, sensors and connected devices with AI to monitor patients outside hospital, for example in heart failure management, vital sign tracking and early warning, as well as app-based therapies prescribed for specific conditions. These often feed into decision support or care-management platforms.
8. Surgical and medical robotics	Robotic systems for surgery, rehabilitation and therapy, increasingly incorporating AI for navigation, image guidance, instrument tracking and video analysis. Most systems in use today are surgeon-controlled rather than autonomous, so the AI component is often limited.

The boundaries between categories are not always clear-cut: ambient documentation tools are expanding into coding and administration, and real-time procedural guidance sits between treatment planning and robotics.

Furthermore, two segments fall completely outside the scope of the analysis, since neither delivers AI to clinicians and patients as a health-specific product: I) AI-enabling infrastructure, such as cloud, compute and foundation models; and II) AI for research, drug discovery and clinical trials. Both are excluded because their AI does not reach clinicians and patients as part of a health-specific product, as described in the Methodology section.

These exclusions carry significant limitations. Infrastructure is where European dependence on US technology is likely to be most pronounced, including among European clinical AI vendors that build on US cloud services and foundation models. Research and drug discovery, meanwhile, attracts a large share of health AI investment and draws extensively on hospital data. Excluding both segments means that **this report likely understates both the extent of European dependence on US technology and the overall scale of investment in health AI, as well as the volume of hospital data flowing into AI development.**

SECTION 1.2

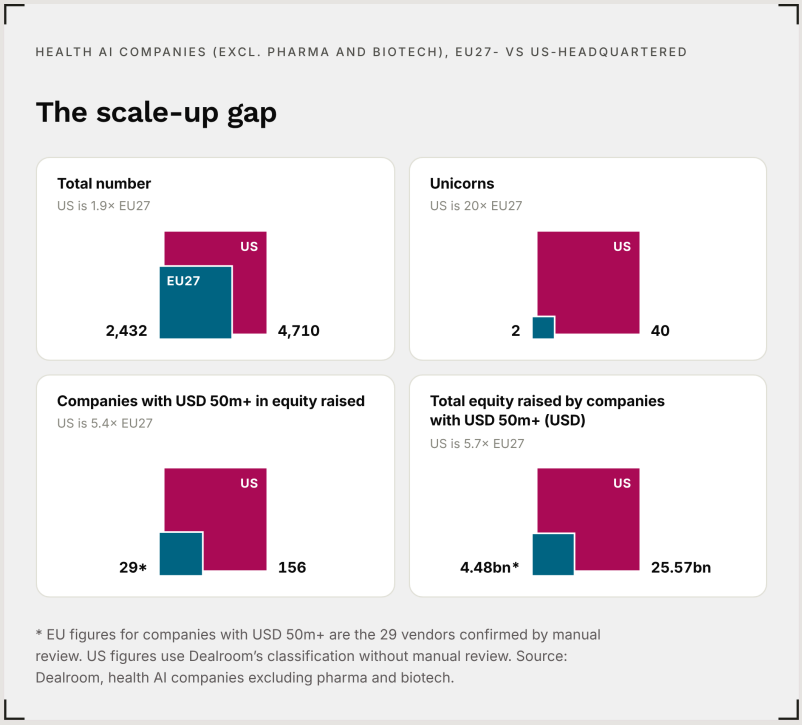
Who supplies health AI in the EU?

EU-US sectoral comparison

Dealroom data cannot show which non-EU health AI companies are active in the EU, but they do allow the overall scale of the US and EU markets to be compared. Applying the same search to US-headquartered companies, again excluding biotechnology and pharmaceutical firms, shows that the gap widens sharply with company size.[27]

27. The US comparison uses only aggregate company counts from Dealroom. The same search filters were applied (Health industry and Artificial Intelligence technology tags, excluding biotechnology and pharmaceutical firms) with headquarters set to the United States.

This snapshot suggests that the US has roughly twice as many health AI companies as the EU, but five times as many that have raised at least USD 50 million, and twenty times as many unicorns. The gap in amounts raised is similar: US health AI companies that have raised USD 50 million or more have together secured about USD 25.6 billion, nearly six times the USD 4.5 billion raised by their EU counterparts. This comparison can not measure the presence of US firms in European hospitals – it does, however, show the **far greater capacity of US health AI companies to fund growth, including expansion into the European market.**



Best-funded EU health AI companies

The five best-funded health AI vendors headquartered in the EU are Neko Health (SE, USD 1.23 billion), Owkin (FR, USD 368 million), Dental Monitoring (FR, USD 297 million), Wandercraft (FR, USD 252 million) and Ada Health (DE, USD 187 million).[28] They span four categories: Neko Health's AI-assisted full-body scans and Owkin's pathology AI fall under diagnostics, Dental Monitoring's smartphone-based orthodontic tracking under remote monitoring, Wandercraft's self-balancing rehabilitation exoskeletons under robotics, and Ada Health's symptom-assessment app under patient-facing tools. Several vendors span more than one category and were classified by their main clinical product; Owkin, for example, also develops AI for drug development but is classified by its diagnostic pathology products. Together, these five companies account for 52% of the USD 4.5 billion raised by the 29 confirmed vendors.[29]

28. Total equity funding as recorded by Dealroom, cumulative from each company's first recorded round.
29. USD 2,335 million of USD 4,482 million in total equity funding across the 29 vendors.

Only two confirmed EU-based health AI vendors from the Dealroom data qualify as unicorns, [30] privately held companies valued at USD 1 billion or more:[31] Neko Health (USD 7 billion) and Owkin (USD 1 billion). Dental Monitoring reached unicorn status in 2021, according to Dealroom, but its most recent estimated valuation is lower, at USD 370-550 million. Tandem Health (USD 400-600 million) is the only other company valued at around USD 500 million or more.[32]

In principle, the leading vendors identified in Dealroom could be matched against the TED procurement data to see whether the best-funded start-ups are also those winning hospital contracts. In practice, the two datasets are not directly comparable: they draw on different sampling frames and measure different things. Read side by side, however, the two datasets show some notable patterns.

Across the explicit AI contracts in our small TED sample, **EU-headquartered companies win the large majority, at least by number of contracts.** They account for 79% of vendor awards and 78% of buyer relationships. US companies account for 11% and 12% respectively.[33] However, these figures count contracts and buyers rather than measuring market reach. To understand which vendors have a presence in the European market that goes beyond their home country, we also looked at vendors with explicit AI contracts awarded by hospital buyers in at least two EU member states.

Within our sample, cross-border presence is extremely rare: only five vendors won explicit AI contracts in more than one member state, and only three of them, Siemens Healthineers, GE HealthCare and Philips, in three or more.[34] Almost all other vendors operated in a single national market. This suggests that, at least for publicly procured AI, **the EU's single market remains fragmented along national lines, and the EU's lead in explicit AI procurement numbers therefore rests largely on vendors selling within a single member state**, while reach across the single market is held by a handful of multinational medical-technology firms, many of them American.

The reasons for that might be simple: each additional member state brings its own procurement rules, languages and integration requirements with national health IT systems, which might explain why vendors remain confined to their home market, while cross-border reach is concentrated among a few multinational imaging manufacturers that can add AI to equipment hospitals already buy from them. Interviews pointed to one further reason why scale matters. An interviewee from a major US technology company noted that **meeting EU compliance requirements across several member states demands legal, regulatory and technical capacity that large US and European firms can afford**, but small start-ups often cannot, giving incumbents an advantage in winning contracts across borders.

30. Aileen Lee, "Welcome to the Unicorn Club: Learning from Billion-Dollar Startups," *TechCrunch*, November 2, 2013, https://techcrunch.com/2013/11/02/welcome-to-the-unicorn-club/.
31. Valuations are Dealroom's most recent recorded or estimated valuations: Neko Health as of July 2026 and Owkin as of March 2026. Where a company has not disclosed its valuation, Dealroom provides an estimated range.
32. Estimate as of September 2026. Several other vendors are estimated at USD 400-450 million: Wandercraft (up to USD 450 million), Universal Diagnostics (up to USD 422 million) and Diabeloop (USD 400 million).
33. Values are missing or not attributable for some awards, and framework agreements often state maximum ceilings rather than actual spending.
34. Siemens Healthineers won explicit AI contracts in five member states, GE HealthCare in four and Philips in three. Hologic and Medtronic won contracts in two member states each. 96% of EU-headquartered vendors with explicit AI contracts won them in a single member state only.

Equity funding to public adoption

Turning to how equity funding may potentially convert into public adoption, the two data sources point to largely separate groups of companies. **Of the 29 best-funded and most highly valued EU health AI companies identified in Dealroom, only one, Tandem Health (SE), appears in the explicit health AI notices in TED at all,** [35] having won a single contract with a Lithuanian university hospital in 2026.[36] The most frequent winners within single member states include Telefónica Soluciones (ES), Alteris (PL) and MEDIST Imaging (RO), and Medicalc (CZ).

This gap should be read with some caution. Some Dealroom companies may reach hospitals through distributors or as part of larger imaging contracts, where notices rarely name the product supplied.[37] Several of the companies, such as Neko Health, Kaia Health and Dental Monitoring, sell mainly to consumers or private providers, so their absence from public procurement is to be expected, with the lack of overlap potentially saying more about the limits of each data source than about the market itself.[38] Even so, the near-absence of well-funded firms from above-threshold hospital contracts suggests that **investor interest in EU health AI companies has not necessarily translated into large-scale public-sector adoption.**

Top categories

Comparing the most prominent functional categories in start-up funding with what hospitals procure as explicit AI in above-threshold tenders shows where investment and procurement may align and where they potentially diverge. **In the Dealroom data, AI-assisted diagnostics is by far the largest category**, accounting for 7 of the 29 confirmed vendors, and 45% of total funding. Remote monitoring and digital therapeutics (5 companies, 15% of funding), surgical and medical robotics (4 companies, 13%), clinical documentation (5 companies, 11%) and treatment planning (4 companies, 8%) follow. Risk prediction, hospital logistics and patient-facing tools are thinly represented, with one or two companies each and less than 5% of funding apiece.

Diagnostics' lead in equity funding, however, rests heavily on a single company: Neko Health, which alone accounts for 27% of all funding and almost half of total valuation.[39] Excluding the Swedish company, co-founded by Spotify's Daniel Ek, funding would be distributed fairly evenly across categories: diagnostics would then attract 24%, remote monitoring 21%, robotics 18% and clinical documentation 15%. Once Neko Health is set aside, this could also suggest that investor interest is spread across several segments of health AI, with no single category clearly dominant.

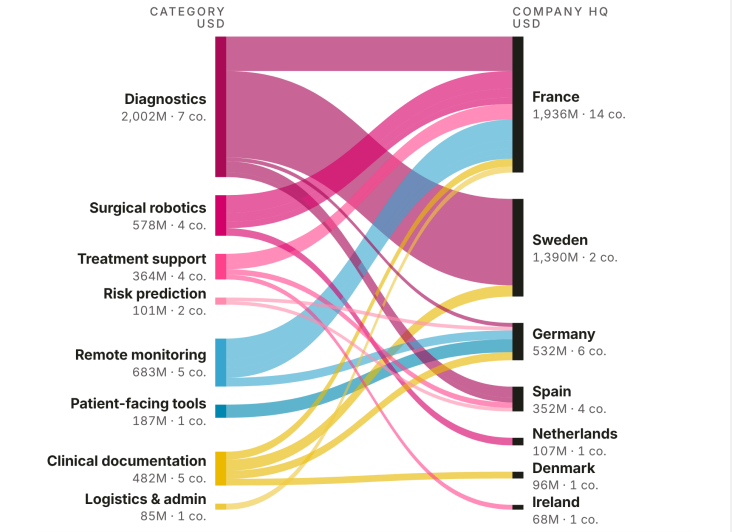
35. All 29 Dealroom vendors and their subsidiaries were matched by company name against every explicit AI notice.
36. TED notice 643196-2026, published 18 September 2026, https://ted.europa.eu/en/notice/-/detail/643196-2026. The Hospital of the Lithuanian University of Health Sciences (Kauno klinikos) used an open procedure to buy an AI-based clinical workflow management system, awarded to Tandem Health for EUR 42,400.
37. Where AI software is supplied as part of an imaging system or through an integrator, the notice typically names only the contract winner, not the developer of the AI component. Such awards are attributed to the named winner in this analysis.
38. Neko Health runs its own body-scanning clinics for self-paying customers. Kaia Health sells digital therapy programmes mainly to employers and health insurers. Dental Monitoring's customers are mainly orthodontic and dental practices, most of them private.
39. Neko Health has raised USD 1.23 billion of the USD 4.48 billion raised by the 29 vendors, and is valued at USD 7 billion of a combined USD 14.7 billion.

TOP-FUNDED HEALTH AI FIRMS IN THE EU BY CATEGORY AND COUNTRY

Health AI's front-runners

The EU's 29 best-funded health AI companies, traced from their main use in care to the country where they are headquartered. (At least USD 50M in equity funding)

Clinical care 3,045M Facing the patient 869M Running the hospital 567M



French health AI companies hold the most funding and span the most categories: 14 companies raising USD 1,936M across six categories. The single largest company, however, is Swedish: **Neko Health** alone accounts for 27% of all funding. More than half of the vendors (16) have raised less than USD 100M each.

▸ DATA TABLE · 29 CONFIRMED VENDORS

Source: Dealroom. 46 companies screened manually, 29 confirmed as clinical health AI vendors.

Among the procurements that explicitly name AI, diagnostics is even more dominant. Of the 94 explicit AI notices in TED, 72% could not be assigned to any category, because the notice text was too general.[40] Of the 26 that could, 17 (65%) involved diagnostics, mostly AI for medical imaging, pathology and endoscopy. Only a handful involved risk prediction and decision support (4), robotics (4), hospital logistics and administration (2), treatment planning (1) or documentation (1).[41] Not a single notice involved patient-facing tools or remote monitoring and digital therapeutics - however, our interviews confirmed that these tools are often bought by insurers, employers, outpatient providers or patients themselves rather than by hospitals, so they may be underrepresented in hospital procurement for that reason. (With only 26 categorised notices, these shares rest on very small numbers and should be read as indicative rather than representative of hospital AI procurement as a whole.)

40. 68 of the 94 explicit AI notices published between 1 November 2023 and 12 October 2026.
41. Since notices can cover more than one category, category counts can add up to more than 26.

Diagnostics seems to be the centre of gravity of health AI in the EU. It attracts the largest share of private funding and valuation, and it accounts for most of the AI that public hospitals buy explicitly in above-threshold tenders. This may reflect the field's maturity: diagnostic tools form the largest class of approved AI medical devices, and they can be specified and procured as discrete products, often together with imaging equipment.[42] Meanwhile, categories such as clinical documentation and remote monitoring, which together attract a quarter of start-up funding, barely appear in explicit AI procurement. Part of this pattern reflects how AI is bought: diagnostic AI is often procured as a separate module or as part of imaging equipment, where it is named in the notice, while AI in documentation and administration more often arrives inside broader software contracts that do not mention it.

To further complicate matters, much of what the procurement data cannot show is how AI reaches hospitals in practice, without ever appearing in tenders. To illustrate, interviews consistently highlighted **vendor lock-in as a key factor shaping procurement dynamics**. Once a core system is in place, such as an electronic health record platform or a major radiology or cardiology system, AI functionalities are increasingly delivered as bundled updates. Many of these core systems come from American companies: according to interviewees, US firms supply much of the infrastructure hospitals depend on most, from electronic health records to cloud services. Hospitals often have limited ability to opt out: according to interviewees, disabling these features can mean losing vendor support or core functionality. As a result, **AI adoption may often be driven less by deliberate procurement decisions than by the product decisions of large incumbent vendors.**

Geographic disparities

Health AI investment is highly concentrated geographically within the EU. **France dominates the Dealroom sample**, with 14 of the 29 confirmed vendors (48%) and 43% of total funding; Paris alone is home to 10 companies and a third of all funding. Sweden has only two companies but 31% of funding and about half (51%) of combined valuation, entirely because of Neko Health and Tandem Health. Germany ranks second by number of companies (6) but accounts for only 12% of funding, followed by Spain with four companies and 8%. The Netherlands, Denmark and Ireland have one company each.

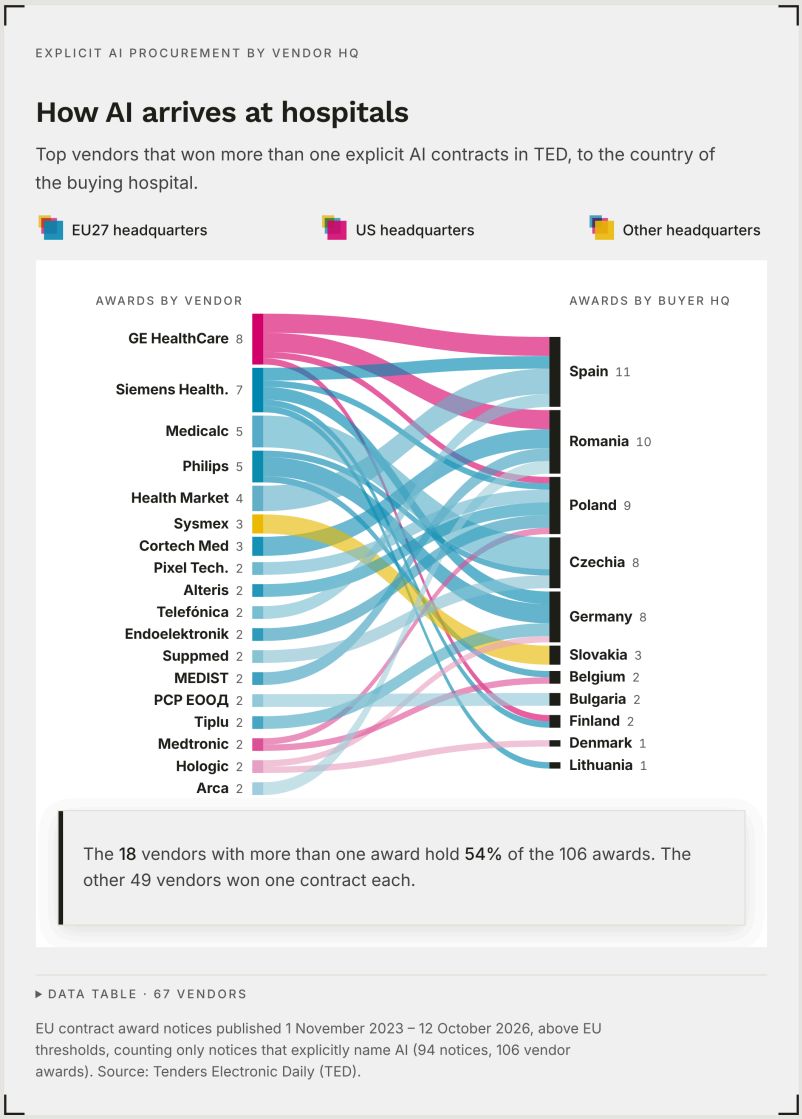
The confirmed vendors come from just seven of the 27 member states, and none from Central or Eastern Europe. The only Central and Eastern European companies returned by the original Dealroom search, Turbine AI (HU) and Sensible Biotechnologies (SK), work in drug discovery and were therefore excluded from our sample. This pattern is consistent with the WHO's observation in its report on how AI is reshaping health systems in the European region,[43] that private sector investment

42. "Radiology AI Tools Dominate Latest FDA Device Approvals," *Applied Radiation Oncology*, January 7, 2026, https://www.appliedradiationoncology.com/news/radiology-ai-tools-dominate-latest-fda-device-approvals. As of December 2025, radiology accounted for 1,039 of the more than 1,300 AI-enabled devices authorised by the US Food and Drug Administration, nearly 80% of the total. See also R. Singh et al., "How AI Is Used in FDA-Authorized Medical Devices: A Taxonomy across 1,016 Authorizations," *npj Digital Medicine* 8, no. 1 (2025): 388, https://doi.org/10.1038/s41746-025-01800-1.

43. WHO Regional Office for Europe, *Artificial Intelligence Is Reshaping Health Systems*.

in health AI startups is largely concentrated in Western Europe and Northern Europe. It also suggests that the **market for health AI is reproducing, rather than narrowing, existing regional divides within the EU**.

Procurement data, in contrast, present an almost inverted picture. **Explicit health AI notices in TED are concentrated in Central and Eastern Europe and in a handful of other markets.** Spain (17 notices), Poland (15), Romania (12), Germany (12) and the Czech Republic (11) together account for 71% of the 94 explicit AI notices. France and Italy, by contrast, published one such notice each and the Netherlands none, and only 19 of the 27 member states published any.



These differences most likely reflect how countries procure and publish contracts, including the use of central purchasing bodies, national portals and national programmes, rather than differences in actual AI adoption.[44] They also cover only contracts that explicitly name AI, and therefore undercount AI procurement overall, particularly in countries where AI is more often bought as part of broader software or equipment contracts.

44. Tünde Tátrai and Gyöngyi Vörösmarty, "Comparative Analysis of the European Centralised Public Procurement," *Central European Journal of Public Policy* 17, no. 2 (2023): 82–96, https://unipub.lib.uni-corvinus.hu/9607/1/ComparativeAnalysisoftheEuropean.pdf.

SECTION 1.3

Key revenue streams for EU health AI companies

Recent debate on AI funding in the EU has been dominated by large-scale public commitments like the InvestAI initiative announced in February 2025, which aims to mobilise €200 billion for AI investment,[45] including a dedicated fund for AI gigafactories, as well as the subsequent AI Continent Action Plan[46] and Apply AI Strategy.[47] The AI Continent Action Plan establishes a five-pillar framework covering computing infrastructure, data, AI development and adoption, skills, and regulatory simplification, with a network of AI gigafactories at the centre of its computing pillar. Meanwhile, the Apply AI Strategy serves as the practical execution roadmap, directly accelerating real-world AI integration across 11 priority industrial and public sectors, including healthcare.

Long before the term 'AI' came to cover a wide range of distinct technical approaches, the EU was already funding projects that relied on methods now described as AI, without calling them that. Through Horizon 2020 and later Horizon Europe, EU funding supported a wide range of health-related research projects in areas such as medical imaging, clinical decision support and digital pathology, typically without labelling them as AI. In parallel, companies developing AI-adjacent solutions raised private equity well before 'AI' became the dominant framing in investors' communications.

In health care, two funding streams have developed along distinct lines. First, **private sector investment has been central to accelerating recent advances in artificial intelligence (AI)**[48]. Venture capital (closed-end funds specialising in high-risk ventures - VC) plays an especially important role, as it funds early-stage, high-risk companies that would struggle to obtain financing through traditional channels such as bank lending. Private funding also reaches EU AI companies through business angels, typically wealthy entrepreneurs. Both VCs and business angels tend to invest through equity, gaining decision-making influence and a share in the company's profits or losses.[49] Furthermore, EU-based startups rely heavily on public-sector innovation and SME-support agencies, which combine grants, loans and quasi-equity to bridge financing gaps.[50]

An **OECD** analysis from early 2026, which covers a broader segment than our sample, including AI in drug discovery and biotechnology, **finds that VC backing for EU start-ups combining AI with healthcare, drugs and biotechnology rose sharply between 2020 and 2024**, from USD 331 million to nearly USD 670 million a year.[51] According to the analysis, healthcare represented 7% of total VC

45. European Commission, "EU Launches InvestAI Initiative."
46. European Commission, *AI Continent Action Plan.*
47. European Commission, *Apply AI Strategy*, 4–7.
48. OECD, "AI in Healthcare," in *Progress in Implementing the European Union Coordinated Plan on Artificial Intelligence*, vol. 2 (Paris: OECD Publishing, 2026), https://doi.org/10.1787/3ac96d41-en.
49. European Parliamentary Research Service, "Private Financing of Innovation in the EU," briefing, PE 769.516 (Brussels: European Parliament, March 2025), https://www.europarl.europa.eu/RegData/etudes/BRIE/2025/769516/EPRS_BRI(2025)769516_EN.pdf.
50. European Commission, "Funding Opportunities for Small Businesses," accessed October 15, 2026, https://commission.europa.eu/funding-tenders/how-apply/eligibility-who-can-get-funding/funding-opportunities-small-businesses_en.
51. OECD, *Progress in Implementing the European Union Coordinated Plan on Artificial Intelligence*, vol. 2.

investment in AI start-ups in the EU in 2024, and VC investment in healthcare remained "considerably lower" than in industries such as finance and banking, energy or raw materials.[52]

The OECD also notes, in a different report, that globally and as of 2025, VC investments into health-related AI firms reached USD 20 billion, approximately 8% of the USD 258.7bn total investments, which puts the EU health AI ecosystem's VC financing at roughly the same levels as global trends.[53] In short: venture funding for EU health AI roughly doubled between 2020 and 2024, but it remains a small niche: health accounts for about the same share of AI venture capital in the EU (7% in 2024) as globally (8% in 2025), so the EU lags not in how much of its AI investment goes to health but in the overall volume of AI investment.

Besides equity investment, **funding has also traditionally reached European health AI start-ups through several major EU public programmes and initiatives, ranging from early-stage research to clinical validation and market scale-up.**[54] The European Commission's AI Research and Deployment Strategy sets out how public funds are distributed to support and scale innovative ideas and solutions in the field of AI development, a strategy that has been regularly updated since its original launch in 2020.[55] Besides its more traditional research and development (R&D) grants, the EU is also increasingly trying to mobilise additional investments from the private sector and the Member States to reach an annual investment volume of €20 billion in AI innovation over the course of this decade, that is, by around 2030.[56]

A key funding stream for health AI research is Horizon Europe, the EU's main grant programme for research and innovation, with a budget of €93.5 billion for 2021-2027,[57] which the Commission has proposed to double in the EU's next long-term budget for 2028-2034.[58] Horizon Europe's health cluster funds collaborative research and innovation projects carried out by European consortia;[59] in 2025, the health cluster dedicated between EUR 471.6 million and EUR 789 million to its recent calls.[60] In the run-up to the next Horizon Europe programme, several high-level reports called for research and innovation to be embedded at the core of the EU's economy to drive sustainable growth, for free movement of knowledge to strengthen research, innovation and education across the single market, and for Horizon Europe to become more effective at aligning research and innovation with Europe's

52. As for OECD's own data collection methods, the segmentation is a single Preqin AI industry sub-category called "healthcare, drugs and biotechnology".
53. OECD, *Venture Capital Investments in Artificial Intelligence through 2025*, OECD Policy Briefs (Paris: OECD Publishing, 2026), https://doi.org/10.1787/a13752f5-en. Full 2025 and 2026 data is not available yet.
54. For an overview of current calls, see European Institute for Biomedical Imaging Research, "Open Funding Calls," https://www.eibir.org/open-funding-calls/; and Nordic Innovators, "E-health," https://nordicinnovators.com/industries/e-health/.
55. "European Research Development and Deployment of AI," European Commission, *Shaping Europe's Digital Future*, last modified July 15, 2026, https://digital-strategy.ec.europa.eu/en/policies/european-ai-research.
56. European Commission, "A European Approach to Artificial Intelligence," *Shaping Europe's Digital Future*, accessed 16 November 2026, https://digital-strategy.ec.europa.eu/en/policies/european-approach-artificial-intelligence. The target was first set in European Commission, *Artificial Intelligence for Europe*, COM(2018) 237 final, 25 April 2018.
57. European Commission, "Horizon Europe," https://research-and-innovation.ec.europa.eu/funding/funding-opportunities/funding-programmes-and-open-calls/horizon-europe_en.
58. European Commission, "Horizon Europe 2028–2034: Twice Bigger, Simpler, Faster and More Impactful," July 16, 2025, https://research-and-innovation.ec.europa.eu/news/all-research-and-innovation-news/horizon-europe-2028-2034-twice-bigger-simpler-faster-and-more-impactful-2025-07-16_en. The Commission proposes a budget of €175 billion for the new Horizon Europe programme, subject to agreement by the European Parliament and the Council.
59. "Cluster 1: Health," European Commission, Directorate-General for Research and Innovation, accessed October 27, 2026, https://research-and-innovation.ec.europa.eu/funding/funding-opportunities/funding-programmes-and-open-calls/horizon-europe/cluster-1-health_en.
60. "Horizon Europe Health Proposals: 1,444 Submissions and What They Mean for 2026," *Accelopment* (blog), August 28, 2026, https://accelopment.com/blog/funding/horizon-health-results-statistics-2025/.

strategic goals, an orientation reflected in the programme's planned close link with the new European Competitiveness Fund.[61]

Certain **Horizon grants also feed into public-private partnerships such as the Innovative Health Initiative (IHI),**[62] a partnership between the EU and industry stakeholders representing Europe's life science industries, such as EFPIA (European Federation of Pharmaceutical Industries and Associations) or EuropaBio (European Association for Bioindustries). Half of IHI's total budget for 2021-2027 (€1.2 billion of €2.4 billion) comes through Horizon Europe, while around €1 billion is committed by industry partners and up to €200 million by other contributing life science partners or associations.[63] The IHI's funding goal is to translate health research and innovation into tangible benefits for patients and society, through addressing the prevention, diagnosis, treatment or management of certain diseases affecting the EU population.

Through the Digital Europe Programme (DEP)'s AI & Health Deployment Calls, the **EU also funds the large-scale adoption and testing of digital health solutions**. With an €8.1 billion budget for 2021-2027, the DEP focuses on boosting European technological capabilities by investing in supercomputing, artificial intelligence, cybersecurity, semiconductors, advanced digital skills, and large-scale digital deployment.[64] This includes targeted grants like a €9 million call for proposals specifically earmarked for clinical AI imaging in oncology and cardiology, and €24.4 million for integrating existing national health systems with the European Health Data Space regulation.[65]

In addition to the specific research-focused grants, the EU also funds health AI systems through its Coordination and Support Actions (CSAs), designed to fund networking, training, and stakeholder coordination activities such as strategic partnerships and knowledge exchange.[66] Most relevant for health AI developers is the **GenAI4Health initiative that sits within the broader GenAI4EU program** with over €700 million committed to deploying generative models across critical sectors. GenAI4EU operates primarily through major funding vehicles like Horizon Europe, the Digital Europe Programme, and the European Innovation Council (EIC). The GenAI4Health Cluster connects funded Horizon Europe projects that develop generative AI for healthcare, biomedical research, clinical reasoning, digital twins, patient-centred innovation and deployment.[67]

The EU's support is not limited to research-focused grants or networking among those carrying out the research. Through the **European Innovation Council (EIC)**, which sits within Horizon Europe, it also actively **funds projects designed to take an idea from an early-stage lab concept all the way**

61. Draghi, *Future of European Competitiveness*; Enrico Letta, *Much More Than a Market* (Council of the European Union, 2024), https://www.consilium.europa.eu/media/ny3j24sm/much-more-than-a-market-report-by-enrico-letta.pdf; European Commission, Directorate-General for Research and Innovation, *Align, Act, Accelerate: Research, Technology and Innovation to Boost European Competitiveness* (Luxembourg: Publications Office of the European Union, 2024), https://doi.org/10.2777/9106236.
62. "Innovative Health Initiative," Innovative Health Initiative, accessed October 4, 2026, https://www.ihi.europa.eu/.
63. Innovative Health Initiative, "IHI Funding Model," https://www.ihi.europa.eu/about-ihi/ihi-funding-model.
64. "Digital Europe Programme," European Commission, accessed October 27, 2026, https://commission.europa.eu/funding-and-tenders/find-funding/eu-funding-programmes/digital-europe-programme_en.
65. "European Health Data Space Regulation (EHDS)," European Commission, accessed October 27, 2026, https://health.ec.europa.eu/ehealth-digital-health-and-care/european-health-data-space-regulation-ehds_en.
66. Thomas Berger, "Coordination and Support Actions: Strategic Networking in Horizon Europe," *EMDESK Academy*, March 28, 2026, https://www.emdesk.com/horizon-2020-horizon-europe-basics-guide/coordination-support-actions-horizon-europe-guide.
67. GenAI4Health Cluster, "Participating Horizon Europe Projects," last modified June 28, 2026, https://genai4health.eu/projects/.

to a commercially viable company.[68] This happens through three key instruments: the EIC Pathfinder supports early-stage research on ideas for radically new technologies; the EIC Transition programme provides grants to mature technologies developed in earlier research, such as Pathfinder projects; and the EIC Accelerator offers a combination of grant funding and equity investment ('blended finance') to scale high-impact start-ups and SMEs.

EU public funding for health AI also flows more indirectly, particularly through public-private partnerships (PPP) that combine EU funding with private capital. The most prominent of these are the **AI Factories and AI Gigafactories of the EuroHPC Joint Undertaking, which provide the supercomputing capacity needed to train and run AI models**. Nineteen AI Factories have been selected so far, alongside 13 smaller AI Factory Antennas, several of them with a focus on health.[69] In July 2026, EuroHPC opened a call for up to seven AI Gigafactories, with the aim to mobilise more than €20 billion in private investment, reportedly alongside around €10 billion in public funding, to strengthen Europe's own computing capacity and reduce its dependence on US infrastructure[70].

SECTION 1.4

Who funds the EU's health AI ecosystem?

To test our assumptions about how R&D funding translates into equity funding, to examine how broader sectoral trends play out in health AI specifically, and to identify patterns in investor geography and in which types of investors favour which kinds of health AI, we analysed our original Dealroom sample of leading health AI companies and their investors, alongside relevant EU public grants[71].

The Dealroom sample links the EU's 29 best-funded health AI vendors to 317 distinct investors through 411 company-investor links, ranging from venture capital funds and corporates to business angels, family offices and public bodies. Each investor was classified manually by type and by the country of its headquarters, and the resulting data were analysed at two levels: across all investor links, to show which kinds of investors are most present, and company by company, to show how the investor mix varies with a vendor's size, funding history and product category.

Independent venture funds form by far the largest investor group, with 119 of the 317 investors (38%), followed by business angels (18%), corporates (13%) and universities, research institutes and hospitals (10%). Accelerators, family offices and private equity firms each account for around 4-6%

68. European Commission, *EIC Work Programme 2026* (Brussels: European Commission, 2026), https://eic.ec.europa.eu/.
69. EuroHPC Joint Undertaking, "EuroHPC JU Selects Six Additional AI Factories to Expand Europe's AI Capabilities," October 10, 2025, https://www.eurohpc-ju.europa.eu/eurohpc-ju-selects-six-additional-ai-factories-expand-europes-ai-capabilities-2025-10-10_en. Health is among the priority sectors of, for example, the 1HealthAI factory in Spain and the Lithuanian and Dutch AI Factories.
70. EuroHPC Joint Undertaking, "EuroHPC Joint Undertaking Launches AI Gigafactories Call," July 30, 2026, https://eurohpc-ju.europa.eu/eurohpc-joint-undertaking-launches-ai-gigafactories-call-2026-07-30_en. The figure of €10 billion in public support is reported in press coverage, e.g. "EU Invests in Artificial Intelligence: Up to Seven AI Gigafactories with €10 Billion Public Support," *Fast Forward*, https://fastforward.com.cy/economy/eu-invests-artificial-intelligence-seven-ai-gigafactories-eu10-billion-public-support; the official announcement states only that funding comes from the EU and member states.
71. Dealroom's funding records are not limited to venture capital: they cover equity rounds from seed to late-stage and growth equity, as well as some non-dilutive funding such as grants, so EU and national grants can appear among a company's funding rounds and public bodies among its investors.

of investors, and most investors (85%) back only one company. Public bodies are few in number, with 19 national or regional and five EU-level, but they are amongst the most connected: the European Innovation Council backs 10 of the 29 companies, EIT Health 9 and Bpifrance 8. **Public investors are present in 20 of the 29 companies, but their weight declines as funding grows.**[72] Among the smaller vendors (like German Clinomic or Irish LUMA Vision), public bodies can make up close to half of a company's investors. Among the largest, including French Owkin and Wandercraft, or German Ada Health, they are a small minority. This reflects both larger investor groups, which dilute the public share, as well as the absence of public investors from several of the largest companies. Four of the ten best-funded vendors, Neko Health (SE), Dental Monitoring (FR), Tandem Health (SE) and Kaia Health (DE), have had no recorded public investors at all.

What sets apart these nine **companies without public backing is that many offer fast-scaling software-based solutions,**[73] including systems that listen to clinical conversations and generate structured medical documentation, such as those developed by Nabla, Tandem Health, and voize; as well as remote monitoring tools that use smartphone cameras, computer vision, and patient-reported data to manage medical conditions such as orthodontics, chronic pain, and cancer treatment, offered by firms like Dental Monitoring, Kaia Health, and Resilience. Two of the three unicorns in our sample, Neko Health and Dental Monitoring, belong to this group.[74]

Public investment, by contrast, is most prevalent in hardware companies and research spin-outs with long, multi-round funding histories. In companies like German Clinomic, Irish LUMA Vision, and Spanish firms INBRAIN Neuroelectronics and Universal Diagnostics, public bodies make up around half of all investors, and these companies have raised capital over seven to fifteen rounds. These are all highly innovative startups specialising in AI-driven critical care, advanced cardiac imaging, brain implants, and blood-based early cancer detection.

The data thus point to two distinct, though unsurprising, funding tracks rather than a single pipeline: **public instruments tend to support slower-scaling, capital- and research-intensive health AI firms**, typically offering solutions that require novel scientific or engineering development, often spun out of research institutions, and facing long development before reaching EU (and other) markets. Meanwhile, **private capital tends to coalesce around a faster-scaling segment of the health AI market**, largely software built on AI models, such as Nabla, Tandem Health, and voize, which can scale commercially at great speed.

72. Companies were ranked by total equity funding and divided into thirds: the top 10, middle 10 and bottom 9. For each company, public investors (EU institutions and national or regional public bodies) were counted as a share of all investors recorded in Dealroom, and these shares were averaged within each third. Figures count investors, not amounts invested. No recorded public investor" means that none of the investors in the Dealroom lists for the company at any stage, as recorded in Dealroom on September 17, 2026, was coded as an EU institution or a national or regional public body. Dealroom's coverage of public funding, however, is uneven, so some of these companies may have received smaller forms of public support that are not recorded.

73. Michael Ewens, Ramana Nanda, and Matthew Rhodes-Kropf, "Cost of Experimentation and the Evolution of Venture Capital," *Journal of Financial Economics* 128, no. 3 (2018): 422–42, https://doi.org/10.1016/j.jfineco.2018.03.001; Benjamin E. Gaddy et al., "Venture Capital and Cleantech: The Wrong Model for Energy Innovation," *Energy Policy* 102 (2017): 385–95, https://doi.org/10.1016/j.enpol.2016.12.035.

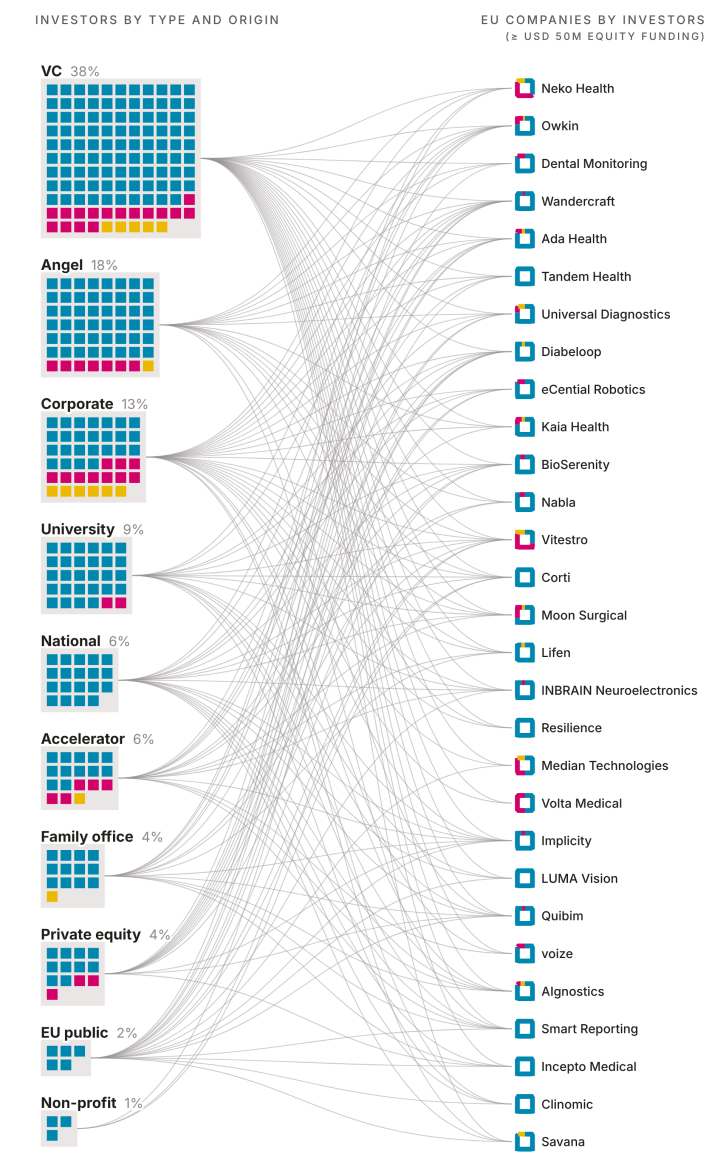
74. Unicorn status (valuation of at least USD 1 billion) as recorded by Dealroom: Dental Monitoring (2021), Owkin (2021) and Neko Health (2025).

INVESTORS IN TOP FUNDED 29 EU HEALTH AI COMPANIES

Scattered capital

Investors of the 29 EU health AI companies with more than USD 50m in equity funding, by HQ and type.

European HQ US HQ Other HQ



Each square is one of the 317 investors, grouped by type, and each donut is one of the 29 companies. Hover over a square, a company or an investor type to see its connections.

▸ DATA TABLE · 317 INVESTORS

Source: Dealroom investor data for the 29 best-funded EU health AI companies (USD 50m+ in equity), coded manually: 317 investors, 411 company-investor links. Investor origin is the country of the investor's headquarters, or for individuals their nationality. "Europe" combines EU27 with the UK, Switzerland, Norway and Iceland.

This pattern is consistent with the wider European picture. The European Investment Bank identifies **late-stage financing as the main gap in European venture markets**, and describes public instruments as supporting companies at earlier stages and catalysing private investment.[75] In our sample, public investors likewise make up a larger share of investors in less-funded companies and in companies with long, multi-round funding histories, while the largest late-stage rounds come almost entirely from private, and often non-EU, capital.

To examine where the investors behind the EU's leading health AI vendors come from, we used the investor records of the same Dealroom sample of the EU's top 29 vendors. Each investor was assigned to the country of its headquarters, distinguishing EU-based investors from those in the United States, other European countries and the rest of the world. Investor geography mattered for our analysis because it shows who shapes, from the background, which AI solutions reach the clinical market and how. In this sense, the **geography of ownership is one indicator of how far the EU's health AI ecosystem depends on capital from outside the Union, a question that has gained prominence in EU debates on competitiveness and technological sovereignty.**[76]

EU-headquartered investors hold 323 of the 411 investor positions (79%), while US-headquartered investors hold 46 (11%), and investors from other European countries, mostly the UK and Switzerland, hold a further 28 (7%). European investors also lead most rounds where Dealroom records a lead investor, although it does so for only a minority of rounds.[77] The two groups differ not only in size but also in kind. European capital is broad-based: independent venture funds form its largest group (110 positions), but it also includes all 56 public positions, from EU instruments such as the European Innovation Council and EIT Health to national promotional banks such as Bpifrance, along with almost all university, research and hospital investors.

Of the 46 US positions, only 16 belong to independent venture funds such as General Catalyst, Lightspeed and NEA. Most of the rest come from corporate investors (10), including pharmaceutical, diagnostics and healthcare-services companies such as Bristol-Myers Squibb, Johnson & Johnson, LabCorp, Quest Diagnostics and Optum and the technology groups GV and Nvidia. They also come from US industry specialists investing as angels (8) and from accelerators (6), such as Y Combinator and Google for Startups. As a result, the US share is highest among accelerators (23%) and corporates (20%), against 11% among venture funds.

Taken together, the data suggest that **EU capital dominates the investor base of Europe's leading health AI vendors by number, while US capital is concentrated in the largest companies and arrives largely through corporate and accelerator channels.** Such investors often pursue commercial as well as financial aims: access to new products, data and distribution partnerships, or a

75. European Investment Bank, *The Scale-Up Gap: Financial Market Constraints Holding Back Innovative Firms in the European Union* (Luxembourg: European Investment Bank, 2024), https://www.eib.org/en/publications/20240130-the-scale-up-gap.
76. See, for example, Draghi, *Future of European Competitiveness*, which highlights Europe's reliance on non-EU capital for scaling innovative companies; and Leevi Saari, "Capital Ecosystem of European AI: Patriotic Billionaires, Development Banks, and the Evolution of State-Finance Nexus," *Competition & Change* (2026), https://doi.org/10.1177/10245294261429545, which finds that US investors account for a third of co-investment ties in European AI start-ups.
77. Dealroom identifies lead investors for only 17 of the 29 companies, and does not specify which funding round they led. Of the 30 lead investor positions made public, 24 are held by EU investors, two by US investors (Lightspeed Venture Partners in Neko Health and Enovis in eCential Robotics), one by a UK investor, and three by investors of unknown origin.

foothold in an adjacent market, potentially with a view to future acquisition.[78] The data record who invests but not how much or on what terms, so this interpretation is not proven by the data, but it is consistent with the pattern of who invests.

The EU health AI investor base is also highly fragmented: 268 of 317 distinct investors (85%) back only one of the 29 vendors, and just 17 back three or more. The few **repeat investors are European and largely French or public**, mainly EU instruments, Bpifrance and a handful of French funds and insurers, while US investors almost always appear as one-off backers of individual companies. This mirrors the broader European AI market, where Saari finds that only 6% of investors back three or more start-ups, and suggests that the **EU still lacks a cross-border group of specialist health AI investors.**[79] Nineteen of the 29 vendors have one or more US investors, and together they account for over four-fifths (81%) of the sample's total funding. Four companies alone, Owkin, Neko Health, Vitestro and Moon Surgical, hold more than half of all US positions, while 11 others have a single US investor.

Billionaire and family capital, the actors Saari calls "patriotic billionaires"[80], appears in a third of the EU's best-funded health AI vendors, almost always in companies from the investor's home country and most visibly in France around Xavier Niel. Niel backs five French vendors, two personally (Nabla, Resilience) and three through his investment vehicle Kima Ventures (Lifen, Implicity, Wandercraft), making him the most connected individual investor in the sample: no other individual backs more than two companies. Like other billionaire investors in the sample, he acts as a minority co-investor alongside, rather than instead of, US and venture capital, and the data cannot show whether such investments reflect sovereignty goals or ordinary home bias.

It is safe to say from the data that **France dominates the EU's leading health AI vendors**: almost half of all companies with more than USD 50 million in equity funding (14 of 29) are French, and so their investor base is correspondingly French. This isn't surprising, since investors in the sample mostly back companies in their own country – what stands out is the density of the French ecosystem. Twelve of the 17 investors backing three or more vendors are French, combining a public investment bank (Bpifrance, in eight companies), large private funds (Eurazeo, Cathay Capital, LBO France), incubators (Agoranov, Wilco), billionaire capital and, distinctively, health insurers. The mutual insurer MACSF backs five vendors, and Wandercraft alone counts four French insurers among its investors (AG2R La Mondiale, MACSF, Malakoff Humanis through MH Innov' and Mutuelles Impact), more than any other company in the sample. French tech founders also reinvest in the ecosystem as angels: Julien Codorniou backs Nabla and Lifen, Simon Dawlat Nabla and Resilience. The French ecosystem also reaches into universities and regional banking: Sorbonne University, which backs BioSerenity and Moon Surgical in the sample, runs a dedicated venture fund for start-ups in areas such as AI-based medical decision support, augmented diagnostics and precision medicine, while UI

78. On the strategic aims of corporate venture capital, see Gary Dushnitsky and Michael J. Lenox, "When Does Corporate Venture Capital Investment Create Firm Value?," *Journal of Business Venturing* 21, no. 6 (2006): 753–72, https://www.sciencedirect.com/science/article/pii/S0883902605000583; on acquisitions of portfolio companies, David Benson and Rosemarie H. Ziedonis, "Corporate Venture Capital and the Returns to Acquiring Portfolio Companies," *Journal of Financial Economics* 98, no. 3 (2010): 478–99, https://www.sciencedirect.com/science/article/pii/S0304405X10001595; and on corporate accelerators, Thomas Kohler, "Corporate Accelerators: Building Bridges between Corporations and Startups," *Business Horizons* 59, no. 3 (2016): 347–57, https://doi.org/10.1016/j.bushor.2016.01.008.
79. Saari, "Capital Ecosystem of European AI."
80. Ibid.

Investissement, which backs Diabeloop and Quibim, was co-founded by French regional banking groups (Crédit Agricole, Banque Populaire and Caisse d'Epargne) and draws on insurers, pension funds and Bpifrance as investors. **France's lead in producing well-funded vendors thus coincides with the EU's most developed health AI investor network, though the data cannot show which drives the other.**

Another important cluster is Swedish investors: two vendors, Neko Health and Tandem Health, hold 31% of all funding, and neither has a public investor. They draw heavily on US venture funds (General Catalyst, Lightspeed, Seven Seven Six) and UK/Swiss funds (Atomico, Lakestar), plus Swedish wealth: Daniel Ek through Prima Materia, and the Stenbeck family's Kinnevik. **Sweden is the clearest example of the private, fast-scaling track, and it rests on a single very large company (Neko Health alone is 28% of funding).**

As for other regions, **all four Spanish vendors have public investors, and public bodies make up just under a third** (29%) of their investors: 16% national or regional and 14% EU instruments. That is the highest share of any country with more than one vendor. Spain also stands out for how many layers of public support it has: central government (CDTI, ICO, Government of Spain), regional bodies (the Catalan Institute of Finance, Avançsa, Valencia's Distrito Digital) and EU instruments. Counting EU instruments, it also has the highest share of investors from other EU countries among multi-vendor countries (31%), many of them French (Cathay Capital, MACSF, Partech), and the lowest US share among them (5%).

As for German startups, four of six German vendors have public investors, spread across federal-state banks and funds, and **Germany, together with Sweden, also has the highest share of investors from outside the EU but not from the US (17%), mostly UK and Swiss capital**: Balderton, Vitruvian, Redalpine and British angels.

Chinese capital in EU health AI investment is marginal and indirect. The clearest case is Furui Medical Science, a Shenzhen-listed healthcare group that took a 13% stake in Median Technologies in 2016 to develop its business in China. Cathay Capital, which backs five vendors, is a Paris-based firm with Sino-French origins: some of its funds have been co-anchored by Bpifrance and China Development Bank. Beyond the US and China, the **remaining non-EU investors are mostly from the UK and Switzerland:** venture funds such as Atomico, Balderton and Lakestar and a handful of angels, joined by a few strategic corporate investors from Japan (Olympus, Terumo) and Korea (Samsung, InterVest) and by Mubadala, the Abu Dhabi state investor.

Finally, we examined how equity funding aligns with where EU public funds are distributed[81]. Identifying AI-related R&D grants is notoriously difficult: EU funding databases such as CORDIS do not systematically flag AI, and definitions of what counts as an AI project vary widely between programmes.[82] At the same time, the EU's AI strategy explicitly sets out to support research, smooth

81. In order to further examine how public funding flows to these start-ups through R&D grants, we compared the top 29 health AI sample to a dataset compiled from CORDIS of Horizon and Digital Europe grants between 2021 and 2025, to see how many of those companies show up as recipients of EU grants. See European Commission, *CORDIS: EU Research Results*, https://cordis.europa.eu/.
82. Warso and Keseru, *From Frameworks to Factories*.

the transition of AI from the lab to the market, encourage adoption and fund innovation.[83] To examine how this public funding reaches the companies in our sample through R&D grants, we compared the 29 vendors against a dataset compiled from CORDIS of Horizon Europe and Digital Europe Programme grants, to see how many of them appear as recipients.

Of the 29 best-funded health AI companies in our Dealroom data, only six appear in EU grant data from the Horizon Europe and Digital Europe programmes (grants signed 2021–2025): Vitestro, INBRAIN Neuroelectronics, Implicity, Owkin, Quibim and Clinomic. Implicity is the only one that appears in both. Funding we can attribute to individual companies comes to roughly €9.6M. This is narrower than the Dealroom picture, where the European Innovation Council appears among the investors of 10 companies, since EIC equity investments and grants signed before 2021 fall outside the CORDIS dataset.

The 23 companies that don't appear at all suggest **that Europe's leading health AI companies have scaled mainly on venture capital, not EU grants**. Where EU money does reach them, it is either late-stage scale-up funding (the EIC Accelerator) or a small role in shared infrastructure projects, not early research funding that later turns into commercial growth.[84] In short, EU grants appear to have played at most a supporting role in building Europe's leading health AI companies: the public money that reaches them does so mainly through equity-linked scale-up instruments and shared infrastructure, while their growth has been driven by private capital.

83. "European Research Development and Deployment of AI"; European Commission, "AI from the Lab to the Market," *Shaping Europe's Digital Future*, accessed 19 October 2026, https://digital-strategy.ec.europa.eu/en/policies/ai-lab-market.

84. Three caveats apply here: 1) Neither dataset covers Horizon 2020 (2014-2020), when some of these companies may have received grants. 2) The Horizon file doesn't list the organisations in each project, so a company that was only a partner wouldn't show up there. 3) Six companies is too few to say anything firm about cause and effect.

CHAPTER 2

Compliance: expectations and reality

EU hospitals that develop and deploy AI systems operate under a complex and still-evolving body of law, spanning product safety, data protection, AI-specific rules, health-data sharing, cybersecurity as well as broader patient rights and transparency rules. This chapter **examines what the EU legislative framework expects of hospitals, and how far those expectations match the reality of compliance in clinical settings**. It begins by setting out the regulatory frameworks most relevant to hospital AI, then considers how these frameworks fit together, where they diverge, and how ongoing reforms are reshaping the compliance hurdles hospitals are preparing for.

SECTION 2.1

Regulatory background

The EU has no single, comprehensive framework for health AI. Instead, hospitals based in the European Union are governed by a layered set of both EU and national legislation that covers patient rights, medical device deployment, the use of artificial intelligence tools, data protection and cross-border data sharing, cybersecurity, public procurement, and access to public information, just to name a few. Of these, the following are most directly relevant to the development and deployment of AI in hospital settings, and particularly, to our research.

Patient rights

Clinical AI is increasingly assessed against the rights patients already hold: international guidance, such as the World Health Organization's 2021 principles for AI in health, treats the protection of patient autonomy, transparency and equity as preconditions for its use.[85] **The EU has no single patient-rights law, instead these rights are enshrined in the Charter of Fundamental Rights, which protects human dignity, physical and mental integrity, privacy, personal data, non-discrimination and access to healthcare,**[86] although the Charter binds Member States only when they implement EU law, so in practice most patient rights are enforced through national law.[87]

85. World Health Organization, *Ethics and Governance of Artificial Intelligence for Health: WHO Guidance* (Geneva: World Health Organization, 2021), https://www.who.int/publications/i/item/9789240029200, setting out six principles, including protecting human autonomy; ensuring transparency, explainability and intelligibility; and ensuring inclusiveness and equity.

86. Charter of Fundamental Rights of the European Union, *Official Journal of the European Union* C 326 (October 26, 2012): 391–407, arts. 1 (human dignity), 3 (right to the integrity of the person), 7 (respect for private and family life), 8 (protection of personal data), 21 (non-discrimination), and 35 (health care), http://data.europa.eu/eli/treaty/char_2012/oj.

87. Ibid., art. 51(1). Article 35 guarantees access to preventive healthcare and medical treatment only "under the conditions established by national laws and practices."

Two EU instruments add a cross-border dimension: Directive 2011/24/EU on the application of patients' rights in cross-border healthcare gives patients the right to seek care in another Member State and be reimbursed under defined conditions.[88] It requires national contact points to inform patients about the quality and safety standards of providers.[89] Regulations (EC) 883/2004 and 987/2009 on social security coordination guarantee medically necessary care during temporary stays abroad and set out the route for planned treatment in another country.[90]

These frameworks predate clinical AI, but they set expectations that AI can either support or undermine, including that patients should understand how decisions about their care are made, that the quality of care should not depend on who they are or where they are treated;[91] and that health data should be handled in ways patients can trust, including when it moves across borders.[92] For hospitals, this means an **AI purchase is not only a question of regulatory compliance but also a question of whether the tool upholds the rights patients already hold.**

The AI Act

The Artificial Intelligence Act (Regulation (EU) 2024/1689) is the EU's main instrument for governing the risks embedded in AI systems and models.[93] Unlike sector-specific rules, it applies horizontally across all industries (including healthcare), classifying AI systems by risk and attaching obligations accordingly.[94]

In healthcare, the AI Act does not replace the Medical Device Regulation (MDR - to be described in greater detail below) but overlays it. An AI system that is, or is part of, a medical device requiring assessment by a notified body under the MDR, such as software that detects pulmonary nodules on CT scans or predicts sepsis risk, is automatically considered 'high-risk' under the AI Act.[95] Because most diagnostic and decision-support software falls into these classes, much clinical AI will be high-

88. Directive 2011/24/EU of the European Parliament and of the Council of 9 March 2011 on the Application of Patients' Rights in Cross-Border Healthcare, *Official Journal of the European Union* L 88 (April 4, 2011): 45–65, arts. 1, 7 (general principles for reimbursement), and 8 (prior authorisation), http://data.europa.eu/eli/dir/2011/24/oj, as amended (consolidated version at https://eur-lex.europa.eu/eli/dir/2011/24/oj/eng).
89. Ibid., arts. 4(2)(a) and 6(3).
90. Regulation (EC) No 883/2004 of the European Parliament and of the Council of 29 April 2004 on the Coordination of Social Security Systems, *Official Journal of the European Union* L 166 (April 30, 2004): 1–123, arts. 19 (benefits in kind that become necessary on medical grounds during a stay in another Member State) and 20 (scheduled treatment), http://data.europa.eu/eli/reg/2004/883/oj; Regulation (EC) No 987/2009 of the European Parliament and of the Council of 16 September 2009 Laying Down the Procedure for Implementing Regulation (EC) No 883/2004, *Official Journal of the European Union* L 284 (October 30, 2009): 1–42, arts. 25–26, http://data.europa.eu/eli/reg/2009/987/oj.
91. Charter, arts. 21 and 35; Directive 2011/24/EU, art. 4(3) (non-discrimination on grounds of nationality).
92. Charter, arts. 7–8; Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the Protection of Natural Persons with Regard to the Processing of Personal Data and on the Free Movement of Such Data, and Repealing Directive 95/46/EC (General Data Protection Regulation), *Official Journal of the European Union* L 119 (May 4, 2016): 1–88, ch. V (transfers to third countries), https://eur-lex.europa.eu/eli/reg/2016/679/oj (hereafter GDPR); EHDS Regulation, arts. 3–10 and 90.
93. AI Act, art. 1, as amended by Regulation (EU) 2026/1744 of the European Parliament and of the Council of 8 July 2026 Amending Regulations (EU) 2024/1689, (EU) 2018/1139 and (EU) 2023/1230 as Regards the Simplification of the Implementation of Harmonised Rules on Artificial Intelligence (Digital Omnibus on AI), *Official Journal of the European Union* L, 2026/1744 (July 24, 2026), http://data.europa.eu/eli/reg/2026/1744/oj.
94. AI Act, recital 26 and arts. 2, 5 (prohibited practices), 6 (high-risk classification), 50 (transparency), and ch. V (general-purpose AI models).
95. AI Act, art. 6(1) and annex I, sec. A, points 11–12 (MDR and IVDR); recital 50; Medical Device Coordination Group and Artificial Intelligence Board, *MDCG 2025-6 / AIB 2025-1: Interplay between the Medical Devices Regulation (MDR) & In Vitro Diagnostic Medical Devices Regulation (IVDR) and the Artificial Intelligence Act (AIA)* (Brussels: European Commission, June 2025), https://health.ec.europa.eu/system/files/2025-06/mdcg_2025-6_en.pdf. See also Mateo Aboy, Timo Minssen, and Effy Vayena, "Navigating the EU AI Act: Implications for Regulated Digital Medical Products," *npj Digital Medicine* 7 (2024): 237, https://doi.org/10.1038/s41746-024-01232-3; and Felix Busch et al., "Navigating the European Union Artificial Intelligence Act for Healthcare," *npj Digital Medicine* 7 (2024): 210.

risk[96] and require third-party conformity assessments.[97] The label does not, however, extend to all medical AI: lower-risk (class I) devices and tools developed and used in-house by hospitals fall outside this obligation, although an in-house tool can still be high-risk because of how it is used, for example in emergency triage. An algorithm is also not high-risk merely because it sits inside a medical device.[98]

For high-risk systems, the Act places obligations on two main actors. Providers (developers) carry the bulk of the burden: risk management, data governance, technical documentation, logging, and human-oversight design.[99] Deployers, including hospitals, must use systems as instructed, assign trained staff to oversee them, monitor their operation, keep logs, and report serious incidents to the provider.[100] Some rights-based safeguards apply more narrowly than is often assumed: Fundamental Rights Impact Assessments (FRIAs), for instance, are required only for systems that are high-risk because of their use, such as AI used to decide eligibility for public healthcare or to triage emergency patients, and not for AI medical devices.[101]

An important exemption is that the AI Act does not apply to AI systems or models developed and put into service solely for scientific research and development,[102] nor to research, testing or development activity carried out before a system is placed on the market or put into service, although testing in real-world conditions falls outside that second exclusion and is subject to its own safeguards.[103]

The AI Act applies in stages: bans on prohibited practices and the duty to promote AI literacy have applied since February 2025, and rules for general-purpose AI models since August 2025.[104] Since 2 August 2026, the AI Office and national authorities have been enforcing the Act, which means that breaches of the prohibitions, the general-purpose AI obligations and the transparency rules can now lead to fines.[105] Transparency rules have also taken effect: chatbots must disclose that users are

96. Regulation (EU) 2017/745 of the European Parliament and of the Council of 5 April 2017 on Medical Devices, Amending Directive 2001/83/EC, Regulation (EC) No 178/2002 and Regulation (EC) No 1223/2009 and Repealing Council Directives 90/385/EEC and 93/42/EEC, *Official Journal of the European Union* L 117 (May 5, 2017): 1–175, annex VIII, ch. III, sec. 6.3 (Rule 11), http://data.europa.eu/eli/reg/2017/745/oj (hereafter MDR); Medical Device Coordination Group, *MDCG 2019-11 rev.1: Guidance on Qualification and Classification of Software in Regulation (EU) 2017/745 – MDR and Regulation (EU) 2017/746 – IVDR* (October 2019, rev. June 2025), https://health.ec.europa.eu/latest-updates/update-mdcg-2019-11-rev1-qualification-and-classification-software-regulation-eu-2017745-and-2025-06-17_en.
97. MDCG 2025-6 / AIB 2025-1.
98. AI Act, art. 6(1)(a)–(b) (the system must be a safety component of a product, or itself a product, that is "required to undergo a third-party conformity assessment"); MDR, arts. 5(5) (health institution exemption) and 52(7) (class I self-declaration); MDCG 2025-6 / AIB 2025-1. On high-risk classification by use, see AI Act, art. 6(2) and annex III, point 5(d) (emergency healthcare patient triage).
99. AI Act, arts. 8–17, in particular arts. 9 (risk management), 10 (data and data governance), 11 (technical documentation), 12 (record-keeping), 14 (human oversight), 16 (obligations of providers), and 17 (quality management system).
100. Ibid., art. 26(1)–(2) (use in accordance with instructions; human oversight by persons with the necessary competence, training and authority), (5) (monitoring and informing the provider), and (6) (keeping logs for at least six months); art. 73 (reporting of serious incidents).
101. Ibid., art. 27(1) and annex III, point 5(a) (eligibility for essential public assistance benefits and services, including healthcare services) and point 5(d) (emergency healthcare patient triage).
102. AI Act, art. 2(6); see also recital 25.
103. AI Act, art. 2(8), which excludes pre-market research, testing and development but expressly carves out "testing in real world conditions"; real-world testing is governed by art. 60 (and art. 57 on regulatory sandboxes).
104. AI Act, art. 113, third paragraph, points (a)–(b); arts. 4 (AI literacy) and 5 (prohibited practices); ch. V (general-purpose AI models).
105. European Commission, "Commission Starts Enforcing AI Act Rules and New Transparency Requirements on 2 August," press release IP/26/1714, July 31, 2026, https://ec.europa.eu/commission/presscorner/detail/en/ip_26_1714; AI Act, arts. 99 (penalties) and 101 (fines for providers of general-purpose AI models).

interacting with AI, and AI-generated content must be labelled (with a grace period until December 2026 for systems already on the market), which is directly relevant to hospitals using patient-facing assistants.[106]

New enforcement channels include a complaint tool[107] and a whistleblower tool,[108] though these cover AI systems and general-purpose AI models under the AI Office's direct competence rather than hospital deployments. Complaints about AI used in hospitals go instead to national market surveillance authorities[109]. Responsibility for enforcing the transparency rules and prohibited practices is shared between the AI Office, national competent authorities and the European Data Protection Supervisor.[110]

Even so, many of the obligations for high-risk systems prescribed by the AI Act, including clinical AI, have not yet taken effect. Their timing and scope were recently changed by the Digital Omnibus: obligations for systems that are high-risk because of their use now apply from 2 December 2027, and those for AI in medical devices, which covers most clinical AI, from 2 August 2028.[111] Further changes are under discussion, including the pending Digital Omnibus amendments to the GDPR and the Commission's proposed revision of the MDR and IVDR; we return to these developments later in this chapter.[112]

Medical Device Regulation

The Medical Device Regulation (Regulation (EU) 2017/745, MDR) governs the safety and performance of medical devices placed on the EU market.[113] In vitro diagnostics, such as laboratory tests, are covered separately by the In Vitro Diagnostic Regulation (IVDR).[114] The MDR pursues two goals at once: high standards of quality and safety for devices, and the "smooth functioning of the internal market" so that compliant products can circulate freely across the EU.[115] Its central principle is that manufacturers are responsible for ensuring, and demonstrating, that their devices are safe and perform as intended throughout their entire lifecycle, not only at the point of market entry.[116] Others in the supply chain, such as importers and distributors, have their own

106. AI Act, art. 50(1), (2), and (4), as amended by Regulation (EU) 2026/1744; European Commission, *Code of Practice on Transparency of AI-Generated Content* (2026), https://digital-strategy.ec.europa.eu/en/policies/code-practice-ai-generated-content. Under Regulation (EU) 2026/1744, providers of generative AI systems placed on the market before 2 August 2026 must comply with art. 50(2) by 2 December 2026.
107. European Commission, "AI Act Complaints Tool," *Shaping Europe's Digital Future*, last updated July 31, 2026, https://digital-strategy.ec.europa.eu/en/policies/ai-act-complaints-tool.
108. European Commission, "AI Act Whistleblower Tool," *Shaping Europe's Digital Future*, last updated July 31, 2026, https://digital-strategy.ec.europa.eu/en/policies/ai-act-whistleblower-tool.
109. The complaint tool excludes complaints about the obligations of general-purpose AI model providers under arts. 53–55, which are handled through a separate channel for downstream providers.
110. AI Act, arts. 70 (national competent authorities), 74(9) (European Data Protection Supervisor), and 75(1) (AI Office competence for AI systems based on general-purpose AI models of the same provider); European Commission, press release IP/26/1714.
111. Regulation (EU) 2026/1744, amending AI Act, art. 113. The obligations for systems that are high-risk because of their use are those for annex III systems; the obligations for AI in medical devices are those for products covered by annex I, sec. A.
112. Digital Omnibus proposal; European Commission, *Proposal for a Regulation of the European Parliament and of the Council Amending Regulations (EU) 2017/745 and (EU) 2017/746*, COM(2025) 1023 final, 2025/0404(COD) (December 16, 2025) (hereafter MDR revision proposal).
113. MDR.
114. Regulation (EU) 2017/746 of the European Parliament and of the Council of 5 April 2017 on In Vitro Diagnostic Medical Devices and Repealing Directive 98/79/EC and Commission Decision 2010/227/EU, *Official Journal of the European Union* L 117 (May 5, 2017): 176–332, https://eur-lex.europa.eu/eli/reg/2017/746/oj/eng (hereafter IVDR).
115. MDR, recitals 1–2.
116. Ibid., art. 10, including art. 10(9) (quality management system) and art. 10(10) (post-market surveillance system); annex I (general safety and performance requirements); arts. 83–92 (post-market surveillance and vigilance).

verification, traceability and reporting duties.[117] The regulation also adds transparency requirements and patient protections, including public safety summaries for high-risk devices and implant cards, which tell patients exactly which device they have received and how to live with it safely.[118]

Whether an AI tool falls under the MDR depends on its intended use or purpose, not on its underlying technology.[119] Software intended by its manufacturer for a medical purpose, such as diagnosing, predicting, monitoring or treating disease, is a medical device. General-purpose software is not, even when used in a hospital, and neither is software intended for lifestyle or wellbeing.[120] A bed-management or scheduling tool, for example, sits outside the MDR. A model that predicts which patients are likely to develop sepsis, or flags suspected strokes on CT scans, is covered by it.[121] A system that only transcribes speech may fall outside the MDR, but regulators increasingly treat tools that summarise or structure clinical content as devices, and those that suggest diagnoses or treatment options almost certainly are.[122] Many of the AI-enabled tools described in Chapter 1, including image analysis, clinical decision support and other software as a medical device, therefore fall within the MDR's scope.

Once inside, AI tools face demanding requirements. The MDR sorts devices into four risk classes (I, IIa, IIb and III).[123] Under its software-specific classification rule, tools that inform diagnostic or therapeutic decisions are at least class IIa, and they rise to IIb or III where an error could cause serious or irreversible harm.[124] A stroke-detection algorithm, for instance, is likely to sit in a higher class than a tool that tracks non-vital physiological trends, such as weight or sleep, for a clinical purpose. From class IIa upwards, and for a few categories of class I device, an independent notified body must assess the device before it can carry a Conformité Européenne, or European Conformity (CE) mark, a process that typically takes more than a year.[125] All devices require clinical evidence that they work as claimed and continuous post-market surveillance once in use.[126]

117. Ibid., arts. 13 (importers), 14 (distributors), 25 (identification within the supply chain), and 27 (Unique Device Identification system).
118. Ibid., arts. 18 (implant card) and 32 (summary of safety and clinical performance for implantable and class III devices).
119. Ibid., arts. 2(1) (definition of "medical device") and 2(12) (definition of "intended purpose"); Case C-329/16, *Syndicat national de l'industrie des technologies médicales (Snitem) and Philips France v Premier ministre and Ministre des Affaires sociales et de la Santé*, ECLI:EU:C:2017:947 (CJEU, December 7, 2017).
120. MDR, recital 19; MDCG 2019-11 rev.1.
121. MDCG 2019-11 rev.1, annex I (illustrative examples of qualification of software used in the healthcare environment). On the prevalence of CE-marked stroke and neuroradiology AI, see Kicky G. van Leeuwen et al., "Artificial Intelligence in Radiology: 100 Commercially Available Products and Their Scientific Evidence," *European Radiology* 31, no. 6 (2021): 3797–3804, https://doi.org/10.1007/s00330-021-07892-z.
122. MDR, art. 2(1); MDCG 2019-11 rev.1. See Medicines and Healthcare products Regulatory Agency (MHRA), *Ambient Voice Technology-Enabled Products: Device Qualification and Classification* (London: MHRA, 2026), https://www.gov.uk/government/publications/ambient-voice-technology-enabled-products/ambient-voice-technology-enabled-products.
123. MDR, art. 51 and annex VIII.
124. Ibid., annex VIII, ch. III, sec. 6.3 (Rule 11); MDCG 2019-11 rev.1, sec. on Rule 11. Under Rule 11, software intended to monitor physiological processes is class IIa, rising to IIb where it monitors vital physiological parameters and variations could result in immediate danger to the patient.
125. MDR, art. 52 and annexes IX–XI. Under art. 52(7), class I devices are self-declared, except sterile devices, devices with a measuring function, and reusable surgical instruments. On certification times, see European Commission, *Notified Bodies Survey on Certifications and Applications (MDR/IVDR)* (Brussels: European Commission, 2025–26); and Team-NB, *Survey on the Implementation of the MDR and IVDR* (2026).
126. MDR, art. 61 and annex XIV (clinical evaluation and post-market clinical follow-up); arts. 83–86 (post-market surveillance, PMS report, periodic safety update report). For software, see Medical Device Coordination Group, *MDCG 2020-1: Guidance on Clinical Evaluation (MDR) / Performance Evaluation (IVDR) of Medical Device Software* (March 2020).

Data protection

The General Data Protection Regulation (GDPR) governs how health data is collected, used, and shared in the EU.[127] It defines 'data concerning health' broadly, as personal data about a person's physical or mental health, including the provision of healthcare services, that reveal information about their health status.[128] The definition covers past, current, and future conditions, data from medical devices and tests, and information collected while registering for or receiving care.[129] The Court of Justice of the EU has interpreted the category expansively: data from which health information can be indirectly inferred also falls within it,[130] and so does information entered when ordering medicines online, even without a prescription.[131] This matters for AI tools, because systems that infer health status from proxies may process health data even when no medical record is involved, notably in remote monitoring and digital therapeutics (category 7), which draw on wearable signals, and in clinical documentation and ambient AI (category 4) or patient-facing tools (category 6), which analyse speech.

The GDPR classifies health data as a special category of personal data, and processing it is prohibited unless one of the exceptions in Article 9(2) applies.[132] Every processing operation must also have a legal basis under Article 6 and respect the core principles of the Regulation, including legal basis, purpose limitation and data minimisation.[133] Explicit consent is one possible exception, but it is often not the appropriate one in clinical settings, where the power imbalance between patient and institution makes it hard for consent to be freely given.[134] In practice, hospitals and health systems tend to rely instead on the provision of healthcare, public health, or scientific research, usually as further specified in national law.[135] Research processing must be accompanied by appropriate safeguards such as pseudonymisation,[136] and data protection must also be built into systems from the outset of their design.[137] High-risk processing, which covers most large-scale use of health data, requires a prior data protection impact assessment.[138] When an EU hospital contracts a vendor outside the Union, the hospital must sign a data processing agreement with the vendor, carry out a DPIA where the processing is likely to be high-risk, and rely on a lawful transfer mechanism.[139]

127. GDPR.
128. Ibid., art. 4(15).
129. Ibid., recital 35.
130. Case C-184/20, *OT v Vyriausioji tarnybinės etikos komisija*, ECLI:EU:C:2022:601 (CJEU, August 1, 2022), paras. 117–28.
131. Case C-21/23, *ND v DR (Lindenapotheke)*, ECLI:EU:C:2024:846 (CJEU, October 4, 2024).
132. GDPR, art. 9(1)–(2).
133. Ibid., arts. 5(1)(b)–(c) and 6(1).
134. European Data Protection Board, *Guidelines 05/2020 on Consent under Regulation 2016/679*, version 1.1, adopted May 4, 2020, paras. 16–20, https://www.edpb.europa.eu/our-work-tools/our-documents/guidelines/guidelines-052020-consent-under-regulation-2016679_en. These paragraphs explain that consent is unlikely to be freely given where there is a clear imbalance of power between the data subject and the controller, particularly where the controller is a public authority, and that other legal bases are generally more appropriate in such cases.
135. GDPR, art. 9(2)(a), (h), (i), and (j). On the problems with relying on consent in clinical settings, see GDPR, recital 43, and EDPB, *Guidelines 05/2020*, paras. 16–24.
136. GDPR, art. 89(1).
137. Ibid., art. 25.
138. Ibid., art. 35(3)(b).
139. Ibid., arts 28, 35, 45 and 46.

These rules apply across the lifecycle of health AI tools, from model training to deployment.[140] Training a model on patient records is a separate processing operation that needs its own legal basis and Article 9 exception. It cannot simply rely on the basis under which the data were first collected for care. The European Data Protection Board has said that **an AI model trained on personal data cannot automatically be treated as anonymous simply because that data cannot be directly extracted from it.** Whether the model contains personal data must instead be assessed case by case.[141]

The European Health Data Space

The European Health Data Space (EHDS) Regulation (Regulation (EU) 2025/327), in force since March 2025, is likely to change how hospitals handle health data more than any other legal instrument discussed so far.[142] Under the GDPR alone, reusing patient data to train AI requires hospitals and developers to identify a valid legal basis and meet national conditions under Article 9(2), which vary between Member States.[143] The European Health Data Space Regulation now provides a dedicated framework for the secondary use of electronic health data, including for training and testing AI systems.

The EHDS creates EU-wide rules for two distinct uses of electronic health data: I) **Primary use covers data used in delivering care.** Patients gain rights to access and share their records, and health systems must exchange key information, such as patient summaries, prescriptions, medical images and lab results, in interoperable formats across borders.[144] II) **Secondary use covers the reuse of the same data for research, innovation, policymaking and regulatory purposes**. The training, testing and evaluation of algorithms is explicitly listed among the permitted purposes.[145] Access is granted through permits issued by national health data access bodies, and the data are processed in secure environments rather than handed over directly.[146]

For hospitals, the EHDS creates a new role as 'health data holders'. They will have to describe the datasets they hold for national and EU catalogues, keep these descriptions up to date, and make the relevant data available to health data access bodies, generally within three months, once a permit is

140. See GDPR, art. 4(2), which defines processing as any operation performed on personal data, and art. 25(1), which requires data protection by design both when the means of processing are determined and during the processing itself. On the application of the GDPR to both the development and deployment phases of AI models, see European Data Protection Board, *Opinion 28/2024 on Certain Data Protection Aspects Related to the Processing of Personal Data in the Context of AI Models*, adopted December 17, 2024, https://www.edpb.europa.eu/our-work-tools/our-documents/opinion-board-art-64/opinion-282024-certain-data-protection-aspects_en.

141. EDPB, *Opinion 28/2024*.

142. EHDS Regulation, art. 105. The regulation entered into force twenty days after publication, on March 26, 2025.

143. GDPR, arts. 6(1), 9(2)(h)–(j), and 9(4). On how fragmented the GDPR's application to health data is across Member States, see Ciara Staunton et al., "Ethical and Social Reflections on the Proposed European Health Data Space," *European Journal of Human Genetics* 32, no. 5 (2024): 498–505, https://doi.org/10.1038/s41431-024-01543-9; and Mahsa Shabani, "Will the European Health Data Space Change Data Sharing Rules?," *Science* 375, no. 6587 (2022): 1357–59, https://doi.org/10.1126/science.abn4874.

144. EHDS Regulation, arts. 3–9 (rights of natural persons), 14(1) (priority categories), 15 (European EHR exchange format), and 23.

145. Ibid., art. 53(1)(e)(ii). This sub-point covers "training, testing and evaluation of algorithms, including in medical devices, in vitro diagnostic medical devices, AI systems and digital health applications," as a form of scientific research. For prohibited secondary uses, see art. 54.

146. Ibid., arts. 68 (data permit) and 73 (secure processing environment).

granted.[147] This requires data that is documented, standardised and extractable on request, and so implies considerable investment in infrastructure and staff. The EHDS also turns EHR systems into a regulated product category for the first time.[148]

From 2029, any EHR system that processes priority categories of personal electronic health data must include two harmonised components, which include patient summaries, electronic prescriptions, imaging reports, test results and more. The first is an interoperability component that allows data to be exchanged in a common European format, while the second is a logging component that records who accessed or changed which data. This logging underpins a new patient right: individuals can find out who accessed their electronic health data, when, and which data were accessed. These components give hospitals a way to verify interoperability and logging when procuring systems, but since most large hospitals depend on a small number of dominant, often US-based EHR vendors, the practical value of these guarantees is not yet straightforward.

The EHDS is not meant to replace the GDPR in any way, and both have to be read together by hospitals' compliance teams.[149] A data user still needs a valid legal basis for processing, and the national conditions for processing health data continue to differ between Member States. Furthermore, the regulation applies in stages, and this shapes what it means for AI in practice.[150] Most provisions apply from March 2027, by which point the governance framework should be in place. The core secondary-use regime and the first wave of cross-border data exchange (patient summaries and ePrescriptions) start in March 2029. Further categories will follow in March 2031, including medical images, lab results and discharge reports for primary use, and genomic data for secondary use.

Cybersecurity

Hospitals are among the organisations most exposed to cyberattacks in Europe.[151] In 2023, Member States reported 309 significant cybersecurity incidents in the health sector, more than in any other critical sector,[152] and ransomware accounts for more than half of analysed incidents in the sector.[153] While we don't know much about the motivations behind these attacks, the consequences extend beyond data loss: attacks can disrupt clinical services, expose highly sensitive patient information and compromise the systems on which AI tools depend,[154] or even, as in a November 2026 attack on the refrigeration systems of France's national blood service, threaten the safe storage

147. Ibid., art. 60.
148. Ibid., Chapter III (EHR systems and wellness applications).
149. Ibid., recital 52 and art. 65. See also Paulina Trigo Kramcsák and Iñigo de Miguel Beriain, "The Opt-Out Right for Secondary Uses of Personal Data in the European Health Data Space: Legal Basis, Scope, and Interplay with the GDPR," *Computer Law & Security Review* 62 (2026): 106367, https://doi.org/10.1016/j.clsr.2026.106367.
150. EHDS Regulation, art. 105. For an article-by-article overview, see Antoinette Vlieger, "EHDS Timeline: 2025–2035," EHDS Jurist, last updated July 24, 2026, https://ehds-jurist.nl/en/tijdlijn.
151. European Commission, *European Action Plan on the Cybersecurity of Hospitals and Healthcare Providers*, COM(2025) 10 final (January 15, 2025), https://health.ec.europa.eu/ehealth-digital-health-and-care/digital-health-and-care/european-action-plan-cybersecurity-hospitals-and-healthcare-providers_en.
152. Ibid.; European Commission, "Commission Unveils Action Plan to Protect the Health Sector from Cyberattacks," January 15, 2025, https://ireland.representation.ec.europa.eu/news-and-events/news/commission-unveils-action-plan-protect-health-sector-cyberattacks-2025-01-15_en.
153. ENISA, *Threat Landscape: Health Sector*. Ransomware accounted for 54 percent of the incidents analysed between January 2021 and March 2023. ENISA hasn't produced detailed sectoral data on cyberattacks since 2023.
154. Ibid. Patient data, including electronic health records, were the most targeted assets (30 percent), and 46 percent of incidents involved data breaches or leaks; COM(2025) 10, sec. 1.

of blood products.[155] As hospitals integrate a growing number of AI and data-driven applications, many of them cloud-hosted and operated by external vendors, their exposure is likely to increase further[156].

Two EU instruments are of particular relevance: The NIS2 Directive sets obligations for hospitals as operators,[157] while the Cyber Resilience Act sets requirements for the software and connected products they rely on.[158] NIS2 classifies large healthcare providers, including most hospitals, as 'essential entities', while medium-sized providers are generally treated as 'important entities'.[159] Both are required to adopt risk-management measures, among them access control, multi-factor authentication, encryption, vulnerability handling and supply-chain security.[160] For AI systems, the supply-chain security obligation is especially significant, as it requires hospitals to assess the security of vendors whose tools are connected to their systems.[161]

Public procurement

Public hospitals, and other hospitals that qualify as contracting authorities, for example because they are mainly publicly funded, acquire AI systems under the EU public procurement framework, as transposed into national law.[162] Contracts above defined value thresholds must be advertised EU-wide through Tenders Electronic Daily (TED), one of the data sources underpinning this report.[163] The framework allows contracting authorities to award contracts on the basis of quality, risk and lifecycle cost rather than price alone.[164] In principle, this makes procurement one of the main points at which hospitals can require evidence of clinical validation, cybersecurity, data-protection safeguards and vendor support for forthcoming AI Act obligations. In practice, however, as shown in Chapter 1, a substantial share of AI reaches hospitals embedded in larger EHR, imaging or software contracts, or through updates to systems already in place. This limits both the leverage procurement affords hospitals and the public visibility of AI spending.

155. "French Blood Banks Targeted by Refrigerator Cyberattack," *DW*, November 2026, https://www.dw.com/en/french-blood-banks-targeted-by-refrigerator-cyberattack/a-79857981.
156. COM(2025) 10, sec. 3.1 (supply chains and the "cloudification" of patient data).
157. Directive (EU) 2022/2555 of the European Parliament and of the Council of 14 December 2022 on Measures for a High Common Level of Cybersecurity across the Union (NIS 2 Directive), *Official Journal of the European Union* L 333 (December 27, 2022): 80–152, https://eur-lex.europa.eu/eli/dir/2022/2555/oj (hereafter NIS2).
158. Regulation (EU) 2024/2847 of the European Parliament and of the Council of 23 October 2024 on Horizontal Cybersecurity Requirements for Products with Digital Elements (Cyber Resilience Act), *Official Journal of the European Union* L, 2024/2847 (November 20, 2024), https://eur-lex.europa.eu/eli/reg/2024/2847/oj (hereafter CRA).
159. NIS2, arts. 2(1) and 3(1)–(2), and annex I, point 5 (health).
160. NIS2, art. 21(2)(d) (supply-chain security), (e) (vulnerability handling and disclosure), (h) (cryptography and encryption), (i) (access control), and (j) (multi-factor authentication).
161. Ibid., art. 21(2)(d) and (3). See also COM(2025) 10, sec. 3.1, on planned ENISA procurement guidelines for healthcare supply chains.
162. Directive 2014/24/EU, arts. 1 and 2(1)(1)–(4).
163. Directive 2014/24/EU, arts. 4 (thresholds), 49, and 51 (publication in the *Official Journal of the European Union*); Commission Implementing Regulation (EU) 2019/1780 of 23 September 2019 Establishing Standard Forms for the Publication of Notices in the Field of Public Procurement (eForms). Thresholds are revised every two years by Commission delegated regulation.
164. Directive 2014/24/EU, arts. 67 (most economically advantageous tender) and 68 (life-cycle costing).

EU HEALTHCARE COMPLIANCE TIMELINE

The compliance maze

DATE

25 May 2018 · ongoing
Controllers (usually hospitals); processors (usually vendors)
Applies (ongoing). Health data is special-category data; processing needs a lawful basis plus an Article 9(2) condition.
GDPR

2019
2020
2021
2022
2023
2024

2025

2 Feb 2025
Providers and deployers
Prohibited AI practices banned; AI literacy obligation applies.
AI Act

26 Mar 2025
Entry into force. Obligations apply from later dates.
EHDS

2 Aug 2025
GPAI model providers
Obligations for providers of general-purpose AI (GPAI) models; governance and penalties provisions apply.
AI Act

2026

26 May 2026
Manufacturers
End of transition for custom-made class III implantable devices placed on the market without a notified-body certificate.
MDR

28 May 2026
Manufacturers, authorised representatives, importers, system/procedure pack producers, sponsors; notified bodies
First four EUDAMED modules mandatory: actor registration, UDI/device registration, notified bodies and certificates, market surveillance.
MDR/IVDR

27 Jul 2026
Providers, deployers, Member States
Omnibus amendments in force: fixed high-risk dates; AI literacy reframed; new legal basis for sensitive data in bias detection (Art 4a); real-world testing extended to Annex I Section A products; national sandbox deadline moved.
AI Act

2 Aug 2026
Providers (Art 50(1)–(2)); deployers (Art 50(3)–(4))
Transparency obligations apply; AI Office and national authorities begin enforcement.
AI Act

28 Nov 2026
Manufacturers
Deadline to register legacy devices and devices placed on the market before 28 May 2026 in EUDAMED.
MDR/IVDR

2 Dec 2026
Providers

2027

End of grace period for output marking by generative AI systems placed on the market before 2 August 2026.
AI Act

26 Mar 2027
Member States; Commission
General date of application. Member States designate digital health authorities and health data access bodies; Commission adopts key implementing acts.
EHDS

28 May 2027
Notified bodies (with manufacturers)
Deadline to upload to EUDAMED certificates issued before 28 May 2026.
MDR/IVDR

2 Aug 2027
Member States; GPAI model providers
Each Member State must have at least one national AI regulatory sandbox. GPAI models placed on the market before 2 August 2025 must comply.
AI Act

2 Dec 2027
Providers and deployers
High-risk obligations apply for use-based (Annex III) systems, e.g. AI used to assess eligibility for public healthcare or to triage emergency patients.
AI Act

31 Dec 2027
Manufacturers
End of transition for legacy class III and class IIb implantable devices (except sutures, staples, dental fillings, braces, tooth crowns, screws, wedges, plates, wires, pins, clips and connectors).
MDR

2028

31 Dec 2027
Manufacturers
End of transition for legacy class D IVDs.
IVDR

2 Aug 2028
Providers (manufacturers) and deployers
High-risk obligations apply for AI that is, or is part of, a medical device or IVD requiring notified-body assessment (broadly MDR class IIa and above).
AI Act + MDR/IVDR

31 Dec 2028
Manufacturers
End of transition for remaining legacy devices (class IIb not covered above, class IIa, and class I devices requiring a notified body).
MDR

2029

31 Dec 2028
Manufacturers
End of transition for legacy class C IVDs.
IVDR

26 Mar 2029
Healthcare providers; Member States; EHR system manufacturers
Primary use, first priority categories: patient summaries, ePrescriptions and eDispensations. EHR systems processing these categories must include the European interoperability and logging components and carry a CE marking.
EHDS

26 Mar 2029
Health data holders; health data access bodies; Member States; data users
Secondary use regime applies for most data categories: data permits via health data access bodies; right to opt out of secondary use.
EHDS

2030

31 Dec 2029
Manufacturers
End of transition for legacy class B and sterile class A IVDs.
IVDR

2 Aug 2030
Providers; public-authority deployers
High-risk AI systems intended for use by public authorities, placed on the market before the high-risk application date, must comply.

2031

AI Act

26 Mar 2031
Healthcare providers; EHR system manufacturers; health data holders

2032

Primary use, second priority categories: medical images and imaging reports, laboratory results, discharge reports (including EHR system requirements). Secondary use extends to remaining categories, e.g. genetic and other omics data and clinical-trial data.

2033

EHDS

2034

26 Mar 2034
Commission; health data access bodies
Third countries and international organisations may participate in HealthData@EU.
EHDS

Abbreviations: DPIA = Data Protection Impact Assessment; EHR = Electronic Health Record; GPAI = General-Purpose AI; IVD(R) = In Vitro Diagnostic (Regulation); MDR = Medical Device Regulation; UDI = Unique Device Identification.

* Source: consolidated EU healthcare compliance timetable. For orientation only, not legal advice. Confirm exact deadlines and scope against the official regulation texts before relying on them!

Transparency

Access-to-information rules determine the extent to which hospitals' use of AI is open to public scrutiny. Requests to EU institutions are governed by Regulation (EC) 1049/2001,[165] while public hospitals fall under national freedom-of-information laws.[166] These vary considerably in scope, deadlines and exemptions, and private hospitals are frequently not covered at all.[167] The AI Act adds only limited transparency of its own: public bodies deploying certain high-risk AI systems must register that use in an EU database, but the requirement applies to Annex III use cases rather than AI medical devices, so most clinical AI will not be recorded there.[168] Therefore transparency rules will mostly depend on national requirements making the single market disrupted, and fundamental rights protection, including access to information, weak.

National laws

While EU legislation sets the overall framework, much of what hospitals are required to do in practice is determined at the national level. Member States define the conditions for processing health data under GDPR Article 9(2), including for research purposes.[169] They maintain their own patient-rights and medical-records legislation, transpose NIS2, and designate the authorities responsible for enforcing the AI Act and the MDR.[170] They will also shape key aspects of the EHDS, including exemptions from the opt-out.[171] In federal or regionalised systems such as Germany, Spain, Italy and Belgium, responsibility for healthcare is further divided between national and regional authorities.[172] As a result, hospitals operating under the same EU framework may face materially different obligations, a theme that runs through our country comparisons in Chapter 3.

165. Regulation (EC) No 1049/2001 of the European Parliament and of the Council of 30 May 2001 regarding Public Access to European Parliament, Council and Commission Documents, *Official Journal of the European Communities* L 145 (May 31, 2001): 43–48, http://data.europa.eu/eli/reg/2001/1049/oj.

166. Council of Europe Convention on Access to Official Documents (Tromsø Convention), CETS No. 205 (June 18, 2009; in force December 1, 2020).

167. Access Info Europe and Centre for Law and Democracy, *Global Right to Information Rating*, accessed October 19, 2026, https://www.rti-rating.org.

168. AI Act, arts. 26(8), 49(3), and 71. Art. 49(1) registration by providers likewise covers annex III systems only.

169. GDPR, arts. 9(2)(h)–(j), 9(4), and 89. Within the EHDS framework, recital 52 of the EHDS Regulation limits Member States' room to impose further conditions on secondary use.

170. AI Act, art. 70; MDR, art. 101.

171. EHDS Regulation, arts. 51(4) and 71(4)–(7).

172. For country-level descriptions of health system governance, see the European Observatory on Health Systems and Policies, *Health Systems in Transition* (HiT) series, country reviews for Germany, Spain, Italy and Belgium (Copenhagen: WHO Regional Office for Europe, various years).

SECTION 2.2

How these laws interact, where they create tension

With the intention of creating a robust and nuanced regulatory environment that supports hospital decision-making while upholding patients' rights to safety and security, the EU has pursued a set of legislative instruments designed to work together.[173] In practice, however, **these instruments often create more confusion and administrative burden than meaningful remedies for the real harms posed by AI systems.**

The reasons are plentiful: these laws were all written at different times, for different purposes and with different ideas of what counts as risk. The GDPR protects patient data, the MDR protects patient safety and the free movement of devices, the AI Act regulates a specific umbrella of technology across all sectors, and the EHDS builds infrastructure for sharing health data across EU borders. When a hospital buys, adapts or builds a clinical AI tool, it has to satisfy all of these legislative avenues at once, but as Clark puts it, organisations actually deploying AI systems that process personal data – such as hospitals – face the very real challenge of having to satisfy multiple frameworks "without clear regulatory guidance on how their conflicts should be resolved."[174]

There has been active scholarly debate about the many ways in which the EU's regulatory frameworks contradict one another or create unnecessary additional burdens. This section gives **a brief overview of that debate, while also offering a new lens of analysis based on the market mapping in Chapter 1.** By reading these tensions through the eight functional categories of health AI and the vendor structures through which AI reaches hospitals, we show that regulatory friction does not fall evenly – its effects differ depending on what an AI tool does, how it is procured and who supplies it. This section first examines the contradictions between these instruments, to then turn to the ways in which even adopted legislation remains a moving target, thanks in large part to the Digital Omnibus.

Differing interpretations of high risk

The first tension concerns the differing interpretations of "high risk' in the context of healthcare AI. The tension here is manyfold: the AI Act and the Medical Device Regulation **define 'high risk' in different ways, and beyond this, the term is also interpreted inconsistently within each regulation – hospitals reportedly apply the same rules differently in practice.**[175]

173. MDCG 2025-6 / AIB 2025-1. The guidance confirms that AI systems that are, or are part of, medical devices must meet the requirements of both the MDR or IVDR and the AI Act, which apply simultaneously and complement each other. Manufacturers may combine the two in a single quality management system, technical documentation and conformity assessment, but compliance with one framework does not in itself ensure compliance with the other.

174. Michael Clark, "GDPR and the EU AI Act: Five Structural Conflicts and Resolution Frameworks" (working paper, SSRN, June 2, 2026), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=6872699.

175. Based on hospitals' responses to our freedom-of-information requests. In these documents, ambient scribes were consistently classified as not high-risk, although they are not automatically so: depending on their intended purpose and functions, they can qualify as medical devices and thus as high-risk AI systems.

The MDR divides medical devices into four risk classes.[176] Class I covers lower-risk devices, many of them non-invasive, such as bandages or wheelchairs, while class III – the highest risk category – covers devices whose failure can cause serious or irreversible harm to the human body, such as pacemakers, heart valves and other implants in direct contact with the heart or central nervous system.[177] The MDR's classification logic is explicitly physical: its recitals describe the classification rules as based on "the vulnerability of the human body,"[178] taking into account factors such as how invasive a device is, how long it is used and which part of the body it affects. Societal and more abstract risks, such as discrimination, opacity or unequal access to care, play little role in how a device is classified.

The AI Act, by contrast, frames risk in broader terms, covering harm to "health, safety and fundamental rights".[179] First, it distinguishes four levels of risk: unacceptable (prohibited practices), high, limited (subject mainly to transparency obligations) and minimal. Rather than offering a single definition, though, the Act identifies high-risk systems through two lists: AI systems that are, or form part of, products such as medical devices that require third-party conformity assessment, and AI systems used in specified sensitive areas, such as access to public healthcare or emergency triage.[180]

Industry groups such as DigitalEurope go as far as describing the MDR/IVDR and the AI Act as reflecting "fundamentally different regulatory philosophies".[181]. The MDR accepts that a device may carry inherent risks as long as these are outweighed by its clinical benefit, whereas the AI Act also defines risk in terms of fundamental rights, which, according to DigitalEurope, "remain largely unspecific for medical technologies".[182] For manufacturers, DigitalEurope argues, this adds reporting obligations and parallel assessment layers on top of existing medical device requirements.

For health AI tools, this means that an AI system that is, or is part of, a medical device becomes high-risk whenever the MDR requires a notified body to assess it.[183] The AI Act's high-risk label for clinical AI is therefore inherited from the MDR's physical risk classes, even though the AI Act aims to protect a wider set of interests. What this means in practical terms is that the MDR's software rule assigns most tools that inform diagnostic or therapeutic decisions automatically as high-risk, whereas **tools such as ambient documentation assistants may easily fall outside the high-risk category unless they are being deployed to inform clinical decisions,** such as to generate risk assessments or recommend treatment. However, even when an ambient scribe is not considered high-risk under the AI Act, it still processes health data that is a special category of personal data under the GDPR, so 'not high-risk' should never be read as 'low-risk' or as exempt from data protection obligations

176. MDR, art. 51(1) and annex VIII.
177. MDR, annex VIII, ch. III, Rule 8.
178. MDR, recital 58.
179. AI Act, art. 1(1); see also art. 7(1)(b) on the criteria for assessing risks of harm to health and safety and adverse impacts on fundamental rights.
180. AI Act, art. 6(1) and annex I (AI systems that are products, or safety components of products, requiring third-party conformity assessment under EU harmonisation legislation, including the MDR and IVDR) and art. 6(2) and annex III, point 5 (access to essential public services, including healthcare, and emergency triage).
181. DIGITALEUROPE, *Q&A: Responding to Criticisms of the Annex I Merger* (Brussels: DIGITALEUROPE, April 9, 2026), 10, https://cdn.digitaleurope.org/uploads/2026/04/Annex-I-QA-31032026.pdf.
182. DIGITALEUROPE, *AI in Healthcare Policy Paper* (Brussels: DIGITALEUROPE, 2026), 2, https://www.digitaleurope.org/resources/ai-in-healthcare-policy-paper/.
183. AI Act, art. 6(1) and annex I, sec. A; MDCG 2025-6 / AIB 2025-1.

The AI Act's approach to high-risk classification has been widely criticised for raising more questions as it answers. A central difficulty is that different obligations attach to the different routes by which a system becomes high-risk: several of the Act's key safeguards apply only to systems classified as high-risk because of their use, not to those classified as high-risk because they are, or form part of, medical devices. These include the fundamental rights impact assessment, the right to an explanation of individual decisions, and the duty to inform people that a high-risk AI system is being used in decisions about them. As a result, most clinical AI, which reaches the high-risk category through the MDR route, falls outside these protections. The outcome is clearly paradoxical: an AI system that helps a public hospital decide patients' eligibility for publicly funded care triggers the full set of rights-based safeguards, while a CE-marked stroke-detection algorithm may trigger none of them. **The level of protection the AI Act affords therefore depends more on the regulatory route by which a system is classified than on the actual risks it poses to patients.**

Another fundamental issue is that classification depends heavily on the provider's own description of its AI tools: even systems that fall within a high-risk area under Annex III can be excluded from that category if the provider argues that they only perform a narrow procedural task, improve the result of work already completed by a human, or merely prepare information for a human decision.[184] Yet a tool that 'prepares information' can still rank, prioritise or filter what a decision-maker sees, and so shape real outcomes even when a human formally makes the final call. The same logic applies under the MDR, where the intended purpose declared by the manufacturer determines whether software qualifies as a medical device at all, meaning that **in both regimes, the entities subject to the regulation play a key role in determining whether their own products fall within its scope or not.** Some scholars warn that manufacturers may bypass regulation, describing systems as 'merely advisory' even where they actively influence clinical decisions, proposing a classification approach that looks at an AI system's actual capacity to influence individual clinical decisions.[185]

The documents that hospitals shared in response to our freedom-of-information requests suggest that this ambiguity is already shaping healthcare practice. Healthcare institutions interpret 'high risk' differently for comparable AI systems, and several documents reflect misunderstandings about when a tool, such as an ambient scribe, might qualify as high-risk. There also **appears to be a tendency by hospitals, understandable given the compliance burden involved, to place systems in lower-risk categories where the rules leave room for interpretation**. This is consistent with the broader concern that classification relies heavily on self-assessment, whether by vendors describing a system's intended purpose or by hospitals classifying their own tools. Against this background, scholars and civil-society organisations have argued that the current wording is too vague to be applied consistently in clinical settings, and that sector-specific guidance on high-risk classification in healthcare is urgently needed.[186]

184. AI Act, art. 6(3). The exceptions apply only to systems listed in annex III, not to those classified as high-risk under art. 6(1) as medical devices, and they do not apply to systems that perform profiling of individuals.
185. Qaiser Khan and Abdul Raffay Saeed, "Governing Medical AI under the EU AI Act: An Integrated Compliance-by-Design Framework" (2026) 8 *Frontiers in Digital Health* 1932256, 6 (s 6.1), https://doi.org/10.3389/fdgth.2026.1932256.
186. Hannah van Kolfschooten and Janneke van Oirschot, "The EU Artificial Intelligence Act (2024): Implications for Healthcare," *Health Policy* 149 (2024): 105152, https://doi.org/10.1016/j.healthpol.2024.105152; Timo Minssen, Barry Solaiman, Lea Köttering, Jakob Wested, and Abeer Malik, "Governing AI in the European Union: Emerging Infrastructures and Regulatory Ecosystems in Health," in *Research Handbook on Health, AI and the Law*, ed. Barry Solaiman and I. Glenn Cohen (Cheltenham: Edward Elgar, 2024), 311–31, https://doi.org/10.4337/9781802205657.00027.

Data: protecting patients, detecting bias

A further fundamental obstacle to harmonising these instruments lies in the **competing imperatives behind their data management requirements, and in the often conflicting philosophies that underpin them**. First, the GDPR and the AI Act clash on how much data it is appropriate to collect and retain, a tension that is particularly significant in sensitive contexts such as healthcare.

The GDPR, the oldest of these legal instruments, sets the baseline by imposing strict limits on the processing of health data through purpose limitation and data minimisation.[187] The AI Act, by contrast, requires providers of high-risk systems to take appropriate, documented measures to detect, prevent and mitigate bias throughout the system's lifecycle,[188] which pushes developers to collect and retain more data, not less. This is because detecting bias requires datasets that are large and representative enough to compare how a model performs across patient groups, including information on sensitive characteristics such as sex or ethnicity, and that are kept long enough to re-test the model as it is updated and used. Detecting and correcting bias, however, is essential if healthcare AI is to serve all patients equally: algorithmic bias can enter at every stage of an AI system's development, from data collection and preparation to model design and deployment, and in health it risks reproducing and amplifying existing inequalities in access to and quality of care.[189]

The result of the clashing philosophies of the GDPR and the AI Act is what some call the *identification paradox*: to detect bias, developers must be able to compare how a model performs across subgroups defined by sensitive characteristics. Yet many of these are themselves sensitive data that the GDPR restricts, so **preventing discrimination in health AI systems may require processing precisely the sensitive data that data protection law seeks to limit.**[190] The AI Act acknowledges this tension by allowing providers to process special categories of data where strictly necessary for bias detection and correction,[191] but this, again, leaves considerable room for divergent interpretation by national data protection authorities.

187. GDPR, art. 5(1)(b), (c), and (e).

188. The AI Act stipulates that the training, validation and testing of datasets that underpin AI systems must be relevant, sufficiently representative and, to the best extent possible, free of errors and complete. See: Regulation (EU) 2024/1689 (AI Act), Article 10(2)(f)–(g) and 10(3).

189. One of the most common forms of algorithmic bias is sampling bias, which arises when the data used to train a model does not represent the population on which it will be used. Adamson and Smith illustrate this with dermatology, where models trained predominantly on images of lighter skin perform markedly worse in detecting skin cancer in patients with darker skin, however technically advanced they are; see Adewole S. Adamson and Avery Smith, "Machine Learning and Health Care Disparities in Dermatology," *JAMA Dermatology* 154, no. 11 (2018): 1247–48, https://doi.org/10.1001/jamadermatol.2018.2348. Norori and colleagues similarly report a skin cancer algorithm that reached only around 50% accuracy for Black patients, whose images made up just 5–10% of its training data; see Natalia Norori et al., "Addressing Bias in Big Data and AI for Health Care: A Call for Open Science," *Patterns* 2, no. 10 (2021): 100347, https://www.sciencedirect.com/science/article/pii/S2666389921002026. Without systematic bias detection, such tools risk delivering a lower standard of care to precisely the groups that health systems already serve least well. See also Trishan Panch, Heather Mattie, and Rifat Atun, "Artificial Intelligence and Algorithmic Bias: Implications for Health Systems," *Journal of Global Health* 9, no. 2 (2019): 020318, https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6875681/; and Ninareh Mehrabi et al., "A Survey on Bias and Fairness in Machine Learning," *ACM Computing Surveys* 54, no. 6 (2022): 1–35, https://doi.org/10.1145/3457607.

190. See Marvin van Bekkum and Frederik Zuiderveen Borgesius, "Using Sensitive Data to Prevent Discrimination by Artificial Intelligence: Does the GDPR Need a New Exception?," *Computer Law & Security Review* 48 (2023): 105770, https://doi.org/10.1016/j.clsr.2022.105770. See also Fatma Sümeyra Doğan, "Balancing Bias Mitigation and Data Protection in AI-Driven Healthcare: Insights from the European Health Data Space, AI Act and GDPR," *University of Vienna Law Review* 9, no. 3 (2025): 99–123, https://doi.org/10.25365/vlr-2025-9-3-99.

191. AI Act, Article 10(5) (original text); Article 4a, as inserted by Regulation (EU) 2026/1744 (Digital Omnibus on AI).

A further tension arises between the GDPR's right to erasure and the long-term data obligations of AI and medical device regulation. Removing a patient's data from a model that has already been trained is technically difficult without retraining it, a challenge known as 'machine unlearning'.[192] Meanwhile, the MDR requires manufacturers to retain the technical documentation and clinical evidence behind a device for at least ten years, or fifteen for implants.[193]

The European Health Data Space (EHDS) adds a further layer of complexity to the questions of how much patient data may be collected, how, for what purposes and under what conditions. The **GDPR treats health data primarily as a patient right to be guarded – the EHDS, by contrast, treats the same data as a public-interest asset**. It obliges healthcare providers to make patient data available to national health data access bodies, which can then grant access to researchers, pharmaceutical companies or AI developers for permitted secondary uses.[194] In practice, the two systems may not differ that significantly: hospitals already often rely on GDPR's Article 9 exceptions since consent is often not the appropriate basis in clinical settings, and access will not be automatic under the EHDS either.[195]

The EHDS and the AI Act clash in their own ways. The bias-mitigation philosophy of the AI Act requires detailed, regularly updated demographic and clinical attributes, so that training data reflects the populations on which a system will be used.[196] The EHDS, by contrast, defaults to anonymised data for the secondary use of medical records, allowing pseudonymised data[197] only where the purpose of the processing cannot be achieved otherwise.[198] Anonymisation, by design, tends to remove the exact attributes needed to measure bias, and so **clinical AI developers might need to persuade EHDS-implementing health data access bodies that bias detection justifies the exception for pseudonymised data.**[199] Matters are further complicated by the lack of a common understanding of pseudonymisation, as national practice and guidance differ considerably between Member States, and EU-level guidance is still being developed.[200]

The shift from consent to opt-out raises further questions about how well these diverse data management philosophies will protect patient rights in the long run. Many warn that the EHDS's opt-out model could entrench existing inequities,[201] as opt-out rates are not evenly distributed across the

192. Lucas Bourtoule et al., "Machine Unlearning," in *2021 IEEE Symposium on Security and Privacy (SP)* (IEEE, 2021), 141–59, https://doi.org/10.1109/SP40001.2021.00019.
193. MDR, art. 10(8).
194. EHDS Regulation, ch. IV, in particular arts. 53 and 60.
195. Access to patient data under the EHDS will require a data permit, is limited to defined purposes, and is subject to safeguards such as pseudonymisation and processing in secure environments, while patients may opt out. See EHDS Regulation, ch. IV, including the provisions on data permits, secure processing environments and the right to opt out of secondary use (art. 71). See also Trigo Kramcsák and de Miguel Beriain, "Opt-Out Right for Secondary Uses of Personal Data."
196. AI Act, art. 10(2)(f)–(g) and (3).
197. Pseudonymised data has been processed so that it can no longer be linked to a specific person without additional information, such as a key linking codes to names, which is kept separately and protected. A patient's name might be replaced with a code, for example, while the hospital keeps the key. Because re-identification remains possible, pseudonymised data is still personal data under the GDPR, and all its rules continue to apply. Regulation (EU) 2016/679 (GDPR), Article 4(5).
198. EHDS Regulation, art. 66(2)–(3).
199. Doğan, "Balancing Bias Mitigation and Data Protection."
200. European Data Protection Board, *Guidelines 01/2025 on Pseudonymisation*, adopted January 16, 2025 (version for public consultation), https://www.edpb.europa.eu/our-work-tools/documents/public-consultations/2025/guidelines-012025-pseudonymisation_en. The Digital Omnibus's proposed relative definition of personal data could shift the boundary again, by determining when pseudonymised data falls outside the GDPR altogether.
201. Doğan, "Balancing Bias Mitigation and Data Protection."

population. Evidence from the UK's National Health Service (NHS) revealed significantly higher opt-out rates among patients recorded as Black (4.61%) than among White (2.87%) or Asian (1.91%) patients, as well as higher rates among women and older patients.[202] For AI developers, this means **datasets accessed through the EHDS may be least representative of precisely the groups for whom biased models pose the greatest risk**.

Humans in the loop

Both the GDPR and the AI Act aim to keep humans in the loop even when decisions are being made through significant AI involvement. The GDPR distinguishes profiling, automated decision-making with meaningful human involvement, and *solely* automated decision-making, which is generally prohibited. The AI Act requires actual, active human control where a human with proper competence, training, and authority can act as a safeguard, including in clinical decisions involving AI systems.[203] On paper, this means a clinician always has to make the final call - however, in practice, heavy workloads and opaque outputs can lead clinicians either to defer to an AI system without scrutiny, a well-documented phenomenon known as automation bias, or to disregard it altogether.[204]

Beyond these legal interpretive challenges, **hospitals also find themselves in a situation where their role in relation to AI systems can change easily, and so can their regulatory compliance obligations**. Whether a hospital is a mere user of AI or its legally responsible maker depends on how it acquires and adapts a tool, and the answer can produce odd results.

A hospital that gives a general-purpose model a medical purpose, or substantially modifies a vendor's system, becomes the provider under Article 25 of the AI Act, and possibly the manufacturer under Article 16 of the MDR, which means a lot higher documentation burdens. Khan and Saeed point out that the hospital may have no contractual lever to obtain the technical documentation it needs from the original developer, which matters a lot for the transparency and accountability of these systems, and that procurement may provide a good remedy through explicit contractual terms that require Annex XII documentation as a condition of supply.[205]

Too many cooks in the kitchen

Clinical AI in EU hospitals is overseen by many regulators, none of whom sees the whole picture, and each is organised differently from one Member State to the next. Data protection authorities, medical device authorities and notified bodies, AI market surveillance authorities, health data access bodies and cybersecurity agencies each supervise one part of the same system, and may interpret the same situation differently.[206] Reporting is equally fragmented: a single ransomware attack on an

202. Frédéric B. Piel, Brandon L. Parkes, Hima Daby, Anna L. Hansell, and Paul Elliott, "The Challenge of Opt-Outs from NHS Data: A Small-Area Perspective," *Journal of Public Health* 40, no. 4 (2018): e594–600, https://doi.org/10.1093/pubmed/fdy059.
203. The Court of Justice has held that a decision may fall within Article 22 where an automated output plays a "determining role" in it, even if a human formally takes the final step. See Case C-634/21, *OQ v Land Hessen (SCHUFA Holding)*, ECLI:EU:C:2023:957 (CJEU, December 7, 2023); and Minssen et al., "Governing AI in the European Union."
204. Kate Goddard, Abdul Roudsari, and Jeremy C. Wyatt, "Automation Bias: A Systematic Review of Frequency, Effect Mediators, and Mitigators," *Journal of the American Medical Informatics Association* 19, no. 1 (2012): 121–27, https://doi.org/10.1136/amiajnl-2011-000089.
205. Khan and Saeed, "Governing Medical AI under the EU AI Act."
206. Janos Meszaros, Jusaku Minari, and Isabelle Huys, "The Future Regulation of Artificial Intelligence Systems in Healthcare Services and Medical Research in the European Union," *Frontiers in Genetics* 13 (2022): 927721, https://doi.org/10.3389/fgene.2022.927721; van Kolfschooten and van Oirschot, "EU Artificial Intelligence Act (2024)."

AI-enabled triage tool could trigger up to five separate reports, each with its own trigger, recipient and deadline, in a sector where hospitals were the target of 42% of reported health-sector cyber incidents in 2023.[207] In 2025, health was the second most affected sector in incidents reported by Member States under NIS2, accounting for 18% of reported events, and healthcare providers were the target of nearly two-thirds of the sector's data breaches.[208]

When such incidents, or AI errors, harm patients, redress is equally fragmented. The AI Act does not itself address liability, and since the proposed AI Liability Directive was withdrawn in 2025, compensation depends mainly on the revised Product Liability Directive and national rules.[209] Because so many actors are involved, from developers and model trainers to device manufacturers and clinicians, **patients may find it exceptionally difficult to prove fault and causation.**[210]

Moving targets

As if these challenges and legal uncertainties were not enough, the regulatory landscape facing healthcare institutions that develop and deploy AI is **further unsettled by several reform proposals with significant implications for the uptake of AI in healthcare**. The most prominent is the Digital Omnibus, which the Commission presented in November 2025 to simplify the EU's digital rulebook.[211]. It consists of two proposals: one amending the GDPR and other data and cybersecurity rules, and the Digital Omnibus on AI, which amends the AI Act. Only the latter has been adopted so far.[212]

The Commission justified the AI amendments by pointing to delays in harmonised standards and in setting up national governance and conformity-assessment bodies, which had made compliance heavier than expected.[213] For hospitals, its most useful change is a fixed timeline. Under the original proposal, high-risk obligations would have applied only once the Commission judged standards and guidance to be ready, leaving hospitals with no more than a latest possible date to plan around.[214] The adopted text sets definite dates: 2 December 2027 for use-based systems such as emergency triage, and 2 August 2028 for AI medical devices.[215]

207. ENISA, *Threat Landscape: Health Sector.*
208. European Union Agency for Cybersecurity (ENISA), *ENISA Threat Landscape 2026* (Athens: ENISA, September 2026), 19 (fig. 7), 37, https://www.enisa.europa.eu/sites/default/files/2026-09/ENISA%20Threat%20Landscape%202026_Final.pdf.
209. Directive (EU) 2024/2853 of the European Parliament and of the Council of 23 October 2024 on Liability for Defective Products and Repealing Council Directive 85/374/EEC, *Official Journal of the European Union* L, 2024/2853 (November 18, 2024), http://data.europa.eu/eli/dir/2024/2853/oj.
210. Omar Tujjar, Gabriele Ientile, and Francesca Toppetti, "After the AI Liability Directive: Civil Liability for Healthcare AI and the Limits of EU Risk Regulation," *European Journal of Risk Regulation* (2026): 1–18, published online August 3, 2026, https://doi.org/10.1017/err.2026.10130.
211. Regulation (EU) 2026/1744 (Digital Omnibus on AI); European Commission, *Proposal for a Regulation of the European Parliament and of the Council Amending Regulations (EU) 2024/1689 and (EU) 2018/1139 as Regards the Simplification of the Implementation of Harmonised Rules on Artificial Intelligence (Digital Omnibus on AI)*, COM(2025) 836 final, 2025/0359(COD) (November 19, 2025).
212. Regulation (EU) 2026/1744, adopted July 8, 2026, published July 24, 2026, in force July 27, 2026.
213. Regulation (EU) 2026/1744, recitals 2–3.
214. COM(2025) 836, art. 1, point 31 (proposed amendment to AI Act, art. 113). Under the proposal, high-risk obligations would have applied six months (annex III) or twelve months (annex I) after a Commission decision confirming that standards and support tools were available, with long-stop dates of December 2, 2027, and August 2, 2028.
215. AI Act, art. 113, as amended by Regulation (EU) 2026/1744: December 2, 2027, for high-risk systems listed in annex III (including annex III, point 5(d), emergency healthcare patient triage), and August 2, 2028, for high-risk systems covered by annex I, including medical devices.

Much still depends on measures to come: the delegated acts, Commission guidelines on high-risk classification and the harmonised standards that will turn broad obligations into testable requirements.[216] However, the Omnibus also changes several obligations in ways that matter to hospitals. First, AI literacy is now an obligation of conduct rather than result, meaning hospitals must show they have taken reasonable steps, such as training, but are no longer accountable for the level of understanding staff actually reach. Second, a new legal basis for sensitive data lets providers test whether models perform worse for protected groups. Third, real-world testing is extended, including to AI medical devices, with the aim of offering a clearer route for evaluating AI devices in hospitals.

The companion proposal would change the GDPR more deeply.[217] Its central change is a relative definition of personal data, as data that a recipient[218] cannot reasonably link to a person would fall outside the GDPR for that recipient.[219] Pseudonymised patient data shared with a vendor that does not hold the re-identification key might then no longer count as personal data in the vendor's hands.[220] The proposal would also recognise legitimate interest as a legal basis for developing AI, making it easier for private developers to train models on health data.[221] However, the European Data Protection Board (EDPB) and the European Data Protection Supervisor (EDPS) have rejected this narrower definition of personal data,[222] and as of writing this report, the text had not reached trilogue.[223] Until it is adopted, hospitals are best advised to build their data governance on the GDPR as it stands.

The Commission's December 2025 proposal to revise the MDR and IVDR would move both regulations into a part of the AI Act's Annex I that is largely exempt from its requirements.[224] As a result, the AI Act's high-risk obligations, including hospitals' duties as deployers, would no longer apply directly to AI medical devices.[225] Equivalent rules would instead have to be introduced through secondary legislation under the MDR, none of which has yet been published. The stated aim is to avoid duplication by folding AI-specific requirements into medical device law. Until that happens, however, safeguards such as human oversight, logging and post-deployment monitoring, which exist precisely to catch errors and bias before they reach patients, would cease to apply.[226]

216. AI Act, arts. 6(5) (Commission guidelines on high-risk classification), 40 (harmonised standards), and 41 (common specifications); Commission Implementing Decision C(2023) 3215 of 22 May 2023 on a Standardisation Request to CEN and CENELEC in Support of Union Policy on Artificial Intelligence.
217. Digital Omnibus proposal.
218. Any organisation that receives data from another, such as a vendor, cloud provider or AI developer that receives patient data from a hospital.
219. Digital Omnibus proposal, art. 3 (proposed amendment to GDPR, art. 4(1)). The proposal draws on Case C-413/23 P, *European Data Protection Supervisor v Single Resolution Board* (CJEU, September 4, 2025).
220. Digital Omnibus proposal, proposed GDPR, art. 4(1) and new art. 41a (implementing acts on pseudonymisation).
221. Digital Omnibus proposal, proposed GDPR, art. 88c. See also EDPB, *Opinion 28/2024*.
222. European Data Protection Board and European Data Protection Supervisor, *Joint Opinion 2/2026 on the Proposal for a Regulation … (Digital Omnibus)* (adopted February 10, 2026), sec. 3, https://www.edpb.europa.eu/system/files/2026-02/edpb_edps_jointopinion_202602_digitalomnibus_en.pdf.
223. Council of the European Union, Provisional Agenda, Permanent Representatives Committee (Part 2), 6, 7 and 11 October 2026, ST 13654/1/26 REV 1 COR 1 (October 6, 2026), item 7.
224. MDR revision proposal, art. 4.
225. AI Act, art. 2(2), under which high-risk AI systems covered by annex I, sec. B legislation are subject only to arts. 6(1), 102–109, and 112; see also art. 26 on the obligations of deployers.
226. AI Act, arts. 12 (record-keeping), 14 (human oversight), 26 (obligations of deployers), and 72 (post-market monitoring). On the concerns of medical societies, see BioMed Alliance, *Position on the Revision of the MDR and IVDR* (March 2026), https://ehaweb.org/app/uploads/2026/03/BioMed-Alliance-position-MDR-IVDR-Revision.pdf.

For hospitals, the result is that it is unclear whether their duties as deployers of clinical AI will come from the AI Act or from MDR secondary legislation that does not yet exist. A 2026 regulatory-science paper reports widespread confusion among manufacturers, some of whom have abandoned AI Act implementation altogether.[227] The human oversight, logging and monitoring duties a hospital budgets for today may or may not apply to device AI in 2028. Taken together, these laws create fundamental regulatory paradoxes with competing and often contradictory demands that cannot all be satisfied at once.

Sandboxes

AI regulatory sandboxes are controlled environments, set up by competent authorities, in which developers can test AI systems under supervision before placing them on the market. They are meant to improve regulatory compliance and legal certainty, ease market entry, and regulatory learning. Following the Digital Omnibus, each Member State must operate at least one national sandbox by 2 August 2027, alongside an EU-level sandbox for EU institutions. Personal data collected for other purposes may be reused in a sandbox to develop AI in the substantial public interest, including disease detection, diagnosis, treatment and improving health systems.[228] Participation also comes with safeguards: a project-specific testing plan, regulatory oversight, an exit report, and the suspension of testing where risks to health, safety or fundamental rights cannot be mitigated.[229]

Sandboxes do not, however, suspend other rules. The GDPR continues to apply in full, including its requirements on lawful processing, transparency, data minimisation and data protection impact assessments, and participants remain legally liable for any harm caused. In practice, sandboxes may reproduce many of the tensions described above, particularly around patient consent, anonymisation, and re-identification risks, especially in trained AI models. Furthermore, a single health AI project may be supervised at once by AI authorities, data protection authorities, medical device regulators and health data access bodies.

Key safeguards, such as consent, the handling of re-identification risks in trained models, and remedies for affected individuals, remain unclear. European data protection authorities have therefore warned against treating admission to a sandbox as a compliance safe harbour, calling instead for strong involvement of data protection authorities, auditable privacy engineering and coordination across regulators.[230]

227. Niamh St John Lynch, Roisin Loughran, Martin Mc Hugh, and Fergal McCaffery, "A Single Regulatory Framework for AI-Enabled Medical Devices: Implications of the EU Simplification Digital Package and Omnibus Proposal," *Journal of Medical Device Regulation* (May 2026), https://eprints.dkit.ie/id/eprint/1042/.
228. AI Act, art. 59(1)(a)(i).
229. AI Act, arts. 57 and 59(1)(a)(i).
230. European Data Protection Board and European Data Protection Supervisor, *Joint Opinion 1/2026 on the Proposal for a Regulation as Regards the Simplification of the Implementation of Harmonised Rules on Artificial Intelligence (Digital Omnibus on AI)* (2026), https://www.edpb.europa.eu/documents/legislative-opinion/edpb-edps-joint-opinion-12026-on-the-proposal-for-a-regulation-as_en.

SECTION 2.3

Compliance in practice: shadow AI, literacy and exemptions

Given these overlapping requirements, EU hospitals often have limited visibility into, let alone control over, how AI tools enter clinical practice, while **AI compliance remains challenging even for healthcare institutions acting in good faith.** Several interviewees described running GDPR, NIS2, the AI Act and MDR as separate roadmaps with separate project teams, doing largely duplicate work.[231]

And yet, there's still very little direct research on how hospitals are actually responding to the EU's overlapping rules on health AI systems. Most of the relevant literature relies on readiness reports, such as the 2026 WHO study that looks at how EU Member States are handling regulatory alignment and capacity-building in healthcare AI adoption,[232] but these don't explain the reality of how hospitals across Europe deal with compliance in practice.

In this final section, we turn to how EU hospitals perceive and describe the tension between regulatory requirements and everyday practice, and to what they identified as the main problems with compliance.

Shadow AI

The most prominent concern raised by our interviewees was the use of shadow AI: **tools used by healthcare professionals without institutional approval or oversight.** Interviewees presented this as a response to the complexity of getting AI systems approved and deployed internally, a problem particularly visible with generative AI tools like ChatGPT or Google Gemini, which clinicians can reach within seconds.[233]

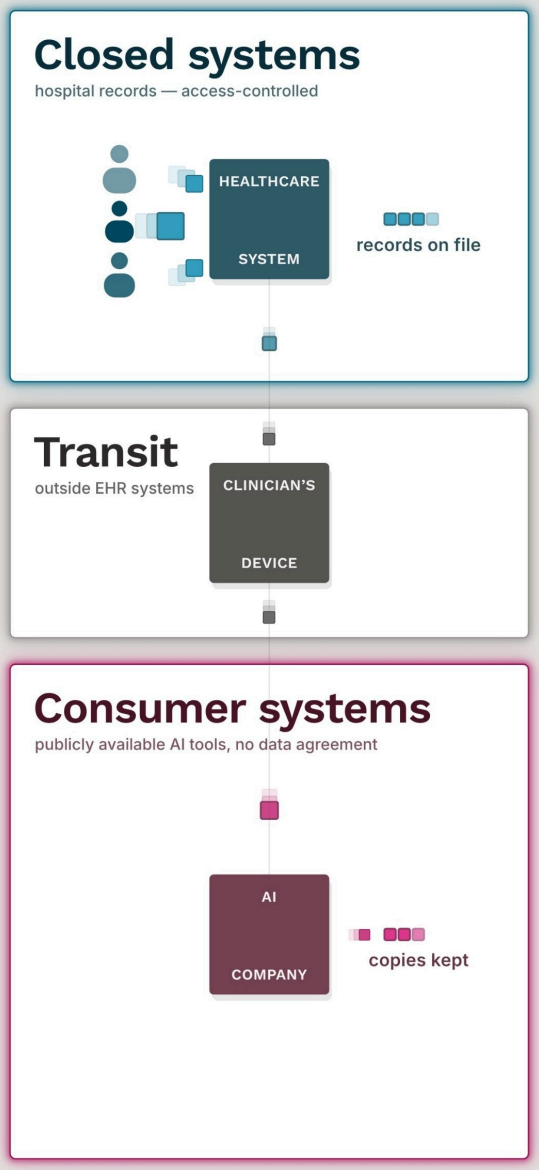
Interviewees linked this behaviour closely to documentation burden and burnout. Clinicians routinely finish their work outside working hours, so-called 'pyjama time', and several interviewees described doctors taking patient information home to complete notes, letters or case reviews.[234] Where approved tools are absent or insufficient, the pressure remains, and commercial chatbots become an easy workaround.

231. For one proposed model of such integration, see Khan and Saeed, "Governing Medical AI under the EU AI Act."
232. World Health Organization Regional Office for Europe, *Artificial Intelligence Is Reshaping Health Systems: State of Readiness across the European Union* (Copenhagen: World Health Organization, 2026), https://www.who.int/malta/publications/i/item/WHO-EURO-2026-12707-52481-81471.
233. L. Petersson et al., "Shadow AI in Swedish Health Care: Qualitative Analysis of Physicians' Free-Text Answers," *Journal of Medical Internet Research* 28 (2026): e93484, https://doi.org/10.2196/93484 (reporting that physicians mainly used ChatGPT through personal accounts and private devices); C. R. Blease et al., "Generative Artificial Intelligence in Primary Care: An Online Survey of UK General Practitioners," *BMJ Health & Care Informatics* 31 (2024): e101102, https://doi.org/10.1136/bmjhci-2024-101102.
234. H. S. Saag et al., "Pajama Time: Working After Work in the Electronic Health Record," *Journal of General Internal Medicine* 34, no. 9 (2019): 1695–96, https://doi.org/10.1007/s11606-019-05055-x; see also B. G. Arndt et al., "Tethered to the EHR: Primary Care Physician Workload Assessment Using EHR Event Log Data and Time-Motion Observations," *Annals of Family Medicine* 15, no. 5 (2017): 419–26, https://doi.org/10.1370/afm.2121.

However, when a patient record is entered into a commercial (usually unpaid version) of an generative AI tool, it creates a second, uncontrolled channel of data transfer, separate from the hospital's audited data flows. **That data may be processed and stored outside the EU by a provider with no data processing agreement with the hospital,** which raises urgent questions under GDPR's rules on special category data and international transfers, and under the Medical Device Regulation, which general-purpose tools have not been assessed against.

WHERE HEALTH DATA ACTUALLY GOES

Shadow AI in healthcare



1 / 4 · INSIDE

Patient data enters a closed system

Every patient hands over their data to the hospital. It passes through the hospital's systems, is checked, and filed in the record – access-controlled, logged, accounted for.

2 / 4 · OUT OF THE BUILDING

Some of it leaves the hospital

Sometimes a clinician copies part of a record onto a personal device or account. Outside the hospital's systems, nothing is checked and nothing is logged.

3 / 4 · CONSUMER SYSTEMS

...and goes into a generative AI tool

The clinician asks a publicly available AI tool for help with a patient record summary, a letter or a second opinion. The answer comes back into the hospital record, but the AI company may keep the input.

4 / 4 · WHAT STAYS BEHIND

The hospital's record looks complete

Everything is back where it belongs. What the hospital never sees is the second copy, held on another company's servers, outside any agreement with the hospital.

The result is a significant gap between formal governance and informal practice. Many of the hospitals we spoke to have formal rules prohibiting employees from entering patient data into unauthorised AI tools, however, they also acknowledged that the hospital cannot fully control what

employees do on personal devices or outside institutional systems. Survey evidence points the same way: in a large US study, nearly one in five healthcare professionals reported using unapproved AI tools, and one in ten had used them for direct patient care.[235]

Some **hospitals' planned move towards approved institutional tools should therefore be understood partly as a containment strategy: an attempt to channel existing informal demand into a safer, more controllable environment** rather than to eliminate it. This is consistent with the literature, which suggests that prohibition alone is unlikely to end shadow AI use and recommends providing secure, sanctioned alternatives instead.[236]

AI literacy gaps

Another key issue raised in multiple interviews was AI literacy in healthcare institutions, with many interviewees noting they see AI training as a fundamental prerequisite for safer AI use, and therefore for stronger patient protection. However, since the Digital Omnibus on AI, the AI Act no longer requires deployers to ensure a sufficient level of AI literacy among their staff, only to take measures that support it. Consequently, **EU hospitals now have weaker legal incentives to actually train their staff to use AI systems safely and responsibly**.

With regard to how hospitals currently train their staff, we found hardly any study that evaluates hospital AI training programmes, and no data on whether such training actually changes behaviour or reduces shadow AI use. Notable exceptions include a 2026 French study where, of the 1,212 medical professionals and 413 students surveyed, only 6.6% reported having had any AI-specific training, whereas 78.3% expressed a desire to participate in one.[237] In addition, a Romanian single-hospital study of 87 nurses, which measured AI legal literacy rather than general AI knowledge, concluded that higher legal literacy is strongly associated with using the AI system as the hospital's procedures required[238].

From the analytical perspective, some scholars argue that 'AI health literacy' should be treated as a distinct concept, particularly now since hospitals are deploying AI solutions at scale.[239] In the meantime, some hospitals have made AI literacy training a prerequisite for staff access to institutional AI tools, according to our interviews.

Skirting oversight

Interviewees described the minimisation **of regulatory scrutiny across hospitals less as deliberate evasion than as a response to a slow and costly compliance pathway**. For instance, hospital respondents to our freedom of information requests frequently characterised their AI tools as "purely

235. Wolters Kluwer Health, *Shadow AI: A Hidden Risk to Healthcare* (January 2026), https://www.wolterskluwer.com/en/solutions/uptodate/ai-clinical-decision-support/shadow-ai-report. The survey covered 518 US providers and administrators. Note that Wolters Kluwer sells its own AI clinical decision support product.
236. Petersson et al., "Shadow AI in Swedish Health Care."
237. A. Alati et al., "Use, Concerns, and Perspectives on AI in Health Care among French Health Professionals and Students: Web-Based Cross-Sectional Survey," *JMIR Medical Education* 12 (2026): e89152, https://doi.org/10.2196/89152.
238. Alina Doina Tănase et al., "Artificial Intelligence Legal Literacy and Procedure-Concordant Use of an AI-Enabled Early-Warning System among Hospital Nurses in Romania," *Nursing Reports* 16, no. 10 (2026): 368, https://doi.org/10.3390/nursrep16100368.
239. Olga Tzortzatou Nanopoulou, "AI Health Literacy as a Foundation for Responsible AI in Healthcare: A Framework for Trustworthy Implementation," *Frontiers in Digital Health* 8 (July 2026): 1832356, https://pmc.ncbi.nlm.nih.gov/articles/PMC13433753/.

research" and therefore exempt from the AI Act, even where the same institution's public communications indicated a mandate to integrate the tool into care in the longer term; which raises the question of how long a tool piloted on patients within a hospital can be said to remain "in the lab", and how hospitals might take advantage of the AI Act's research exemptions in order to avoid regulatory burden.

Interviewees also noted that developers were "tiptoeing around the borders of clinical decision support" to avoid the lengthy CE-marking process, which they regarded as a loss given the potential of such systems. Administrative and ambient documentation tools were similarly described as sitting outside the Medical Device Regulation, with such tools avoiding the "expensive and difficult pathway of medical devices". Some cautioned **that ambient note-taking systems risk "secretly acting as a medical device" without certification,** as large language models move from summarising conversations toward making medical inferences.

Widening gap in healthcare provision

Finally, for hospitals, many of the main compliance hurdles are likely to be practical rather than legal. The EHDS, for example, will require healthcare institutions, as data holders, to produce standardised, well-documented and interoperable health datasets.[240] Several interviewees stressed that many institutions are far from ready for this, with some still keeping parts of their records on paper.

The investment required also falls unevenly. University hospitals with long-established data platforms are well placed both to meet their data-holder obligations and to use the system for their own research and AI development, **whereas smaller general hospitals, and many in Central and Eastern Europe, may struggle to meet those obligations.**[241] The EHDS therefore risks widening the very gaps between hospitals that our findings already point to.



Limitations of our analysis

This research combined several data sources, none of them designed to track AI in healthcare, and its findings should be read with that in mind.

Our inclusion criteria restricted the analysis to vendors that supply AI to clinicians and patients as part of a health-specific product. AI-enabling infrastructure, including cloud services, compute and foundation models, falls outside the sampling frame, yet this is where dependence on US technology

240. "Health data holder" is an EHDS concept that designates the bodies obliged to make electronic health data available for secondary use, while "data controller" is a GDPR concept that designates the body determining the purposes and means of processing personal data, and hence responsible for that processing. Hospitals usually hold both roles at once.

241. WHO Regional Office for Europe, *Digital Health in the WHO European Region: The Ongoing Journey to Commitment and Transformation* (Copenhagen: WHO Regional Office for Europe, 2023).

is likely to be greatest. A diagnostic AI tool developed by a European company may still run on US cloud infrastructure, rely on a US-built foundation model and draw on clinical data held in US-built systems. Analyses that focus on AI vendors and explicit AI procurement alone are therefore likely to underestimate the depth of this dependence.

In addition, Dealroom covers the venture-backed segment of the market and underrepresents large incumbents such as imaging, device and EHR vendors. Classification depends on Dealroom's industry and technology tags, which largely reflect how companies describe themselves, reflecting a single snapshot in a fast-changing market. Its funding thresholds favour capital-intensive business models, and its valuations are estimates: for private companies, they are derived from funding rounds rather than market prices.

Our TED sample, on the other hand, captures AI only where a contract notice names it explicitly, and only for notices published since November 2023. AI bundled into larger contracts, delivered through software updates or bought below the thresholds, is largely invisible. Contracts are attributed to the winning entity, which may be a distributor of technology developed elsewhere. Explicit AI procurement should therefore be read as a lower bound for the AI hospitals actually buy, and the EU share of contracts as an upper bound for European technology. Coverage here is limited in further ways. TED includes only contracts above EU publication thresholds; below-threshold contracts, those published only on national portals and procurement by private providers are not observed.[242] A fuller picture would therefore require reviewing notices that only hint at AI, and extending the search to below-threshold and national level procurement documentation.

Investor data records which investors are present, not how much each contributed. Lead investors are recorded for only 17 of the 29 companies and are not linked to specific rounds. Dealroom's coverage of grants and soft loans is also uneven, so public support is likely understated for some companies.

The qualitative strands have their own limits. The legal analysis reflects the law as of writing this report, in a framework still being reshaped, so some findings may change fast as reforms are adopted. The 36 interviews were selected purposely rather than to be representative, drawing heavily on well-resourced institutions, particularly Dutch and Belgian university hospitals. Freedom-of-information responses depend on national transparency laws and cultures, and on whether institutions recognise what counts as AI, so they reveal as much about disclosure practices as about AI use. The nine-country sample, too, captures a range of health systems but cannot represent the whole EU.

Taken together, these limitations mean the findings are descriptive and cannot establish cause and effect. Still, independent sources often point in the same direction: European actors dominate by number while US actors cluster at the top end, activity is organised largely along national lines and concentrated in a few western and northern member states, and much hospital AI remains invisible to public scrutiny. This consistency suggests structural features of the EU health AI market that are worth testing with fuller data.

242. Directive 2014/24/EU, art. 4.



Conclusions

AI holds real promise for European healthcare: it can speed up diagnosis, ease the documentation burden on clinicians, and extend high-quality care to more patients. Our research suggests, however, that the conditions needed to realise this promise safely and fairly are not yet in place. European hospitals are adopting AI faster than they are building the governance to manage it, and much of that adoption is shaped by forces they do not control.

The first of these forces is the EU's healthcare AI market. Artificial Intelligence tools increasingly reach hospitals not through deliberate procurement but through updates to systems they already depend on, many of them supplied by large, often non-European vendors. Hospitals carry much of the legal responsibility for how these tools are used, yet they frequently lack the information, bargaining power or technical capacity to oversee them meaningfully. Europe produces promising health AI companies, but it struggles to scale them, and public investment rarely translates into visible public adoption. The result is a dependence that sovereignty goals alone cannot undo.

The second is the regulatory landscape itself, given that the EU has built one of the world's most comprehensive frameworks for protecting patients' data, safety and rights. In fact, the hospitals we spoke to take pride in it and largely reject the claim that it blocks innovation. Yet these rules were written at different times, for different purposes, and with different ideas of risk. In practice, they overlap, pull against one another and leave significant gaps. The protection patients receive often depends less on the risks a tool poses than on how it is classified, a judgement that rests largely with those who make it. Ongoing reforms add to the uncertainty, so hospitals are preparing for obligations that may change before they apply.

The third factor is uneven capacity. Well-resourced university hospitals can build their own tools, appoint AI officers and prepare for new data-sharing duties. Many general hospitals, particularly in Central and Eastern Europe, however, have little written governance and struggle to meet even basic requirements. Inequality also accumulates within the models themselves, which are trained on data that under-represents some patient groups and tested under rules that make bias hard to detect. Bias remains strikingly under-discussed in hospitals, and we found little evidence of systematic efforts to assess it in practice. Without deliberate intervention, AI risks entrenching the very divides in access to and quality of care that it is often expected to close.

Throughout, patient data travels further and less visibly than patients are likely to realise: through vendor contracts and cloud services, and increasingly through the informal use of commercial AI tools by overstretched clinicians. Neither hospitals nor patients can reliably trace where that data

ends up, and existing legal tools, such as data access requests, offer little insight into whether it has been used to train AI. As consent gives way to opt-out models of data sharing, transparency about who uses patient data, and for what, is becoming the most important safeguard patients have.

Our research also uncovered good practice: hospitals that make AI training a condition of access to approved tools, portals that show patients where their data has gone, and pilots designed to minimise the data they retain. What is missing is less will than the conditions to scale such practice: clearer and more coherent rules, stronger obligations on vendors to share information, sustained investment in hospitals' governance capacity, and a sharper focus on whether AI tools actually benefit the patients they are meant to serve.

This report is the first part of a larger study. Two further chapters, to be published in January, will build on these findings, first by examining hospitals' AI strategies and governance documents across nine EU countries: which AI systems hospitals use, how they govern them, and what drives their uptake. The second chapter will turn to patients: how the shift from consent to opt-out under the European Health Data Space will affect them, what European patients expect from AI in their care, and what our data access requests revealed about how far patients can trace the use of their data. Together with these chapters, we will publish recommendations for policymakers, hospitals, developers and patient organisations, as well as a practical toolkit to help make health AI tools more transparent and accountable.

END OF REPORT

Report 2, on how hospitals select, govern and deploy AI, follows in January 2027.