Vital Signs: AI Healthcare

Clinical AI and its regulation

HealthSeed

HEALTHSEED INTELLIGENCE

27 July to 10 August 2026 • Published for the HealthScale community and HealthSeed ecosystem partners

THIS EDITION'S SIGNAL

The deadline everyone was planning against moved by a year. The obligations already in force did not, and the relief expires the moment you change the model.

AI Act obligations after Regulation (EU) 2026/1744

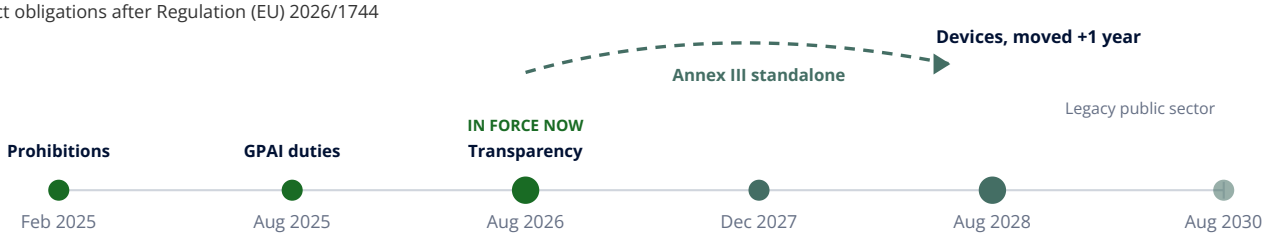

Moved

REGULATION (EU) 2026/1744, IN FORCE 27 JULY

- AI in regulated products, where a medical device sits, from 2 Aug 2027 to **2 Aug 2028**
- Stand-alone high-risk systems, from 2 Aug 2026 to 2 Dec 2027
- National regulatory sandboxes, to 2 Aug 2027

Did not move

APPLIES TODAY

- Article 50 transparency and content-marking duties, from **2 Aug 2026**
- General-purpose AI provider obligations, since Aug 2025
- Article 5 prohibitions, since Feb 2025
- Every Medical Device Regulation obligation, untouched

Three developments in a month point the same way. A regulator deferred what is hardest to meet and kept what is easiest to breach. A hospital beat a frontier model on its own archive. And the field's most credible voices published three incompatible answers to whether any of it is working.

IN THIS EDITION

01 The regulatory read

Regulation (EU) 2026/1744, the Digital Omnibus on AI, was published in the Official Journal on 24 July 2026 and entered into force on 27 July. It took effect on the third day after publication rather than the customary twentieth, because 2 August was too close to leave the text hanging.

What moved

Obligation	Was	Now
High-risk AI embedded in regulated products, Annex I. This is where an AI-enabled medical device sits	2 August 2027	**2 August 2028**
Stand-alone high-risk systems, Annex III	2 August 2026	2 December 2027
Pre-existing high-risk systems used by public authorities	not previously set	2 August 2030
National regulatory sandboxes	2 August 2026	2 August 2027

Two details worth more attention than they have had

The dates are now unconditional. The Commission's November 2025 draft tied the high-risk deadlines to an assessment that harmonised standards and support tools were ready. The final text removed that trigger. These are fixed calendar dates that cannot slip again without a fresh legislative procedure, so no further delay is coming and nobody has to wait for a readiness finding to start.

The relief lapses on significant design change. A high-risk system already on the market before the new application date sits outside the obligations only until it undergoes a significant modification. For a static device that is a genuine reprieve. For an adaptive model, or any product on a normal release cadence, it is a clock your own next release can stop.

Article 4 on AI literacy was also rewritten, from a duty to ensure a sufficient level into a duty to take measures supporting its development. It is an obligation of effort rather than of result, still binding on every deployer, with national supervision from 3 August 2026.

Official Journal, 24 July 2026. European Commission AI Act policy page. Gibson Dunn, Hunton, Cloud Security Alliance and Modulos analyses of the adopted text, all consistent on the dates.

HEALTHSEED PERSPECTIVE

The exposure is not the 2028 date. It is the gap between what a company thinks it has been given and what it actually has. A team that heard "delayed" and stood its AI Act work down is carrying three live obligations, a transparency duty that applied a week ago, and a deferral its own next release can void.

Establishing which of the three reaches your product is a matter of days. Discovering it during a conformity assessment is a review cycle.

Daren Wilson, Co-founder, Senior Partner

02 What just became possible

On 10 July a University of Michigan team published NeuroVFM in Nature Medicine: a three-dimensional visual foundation model trained on 5.24 million clinical MRI and CT volumes, drawn from 566,915 studies acquired over twenty years at Michigan Medicine.

+21.4 percentage points over **GPT-5** on critical-findings triage, at 92.6% against 71.2%, in a prospective study run across the whole health system for one week in January 2026. Roughly half the hallucination rate.

The premise is the interesting part. Frontier models trained on internet-scale public data are structurally blind to neuroimaging, because the identifiable facial features embedded in head MRI and CT keep those scans out of public datasets. The paper shows frontier models underperforming as a direct consequence, then shows that training on uncurated data generated during routine clinical care closes the gap.

NeuroVFM was trained self-supervised and vision-only, without hand labelling, without curated datasets, and without radiology report supervision. The authors call the paradigm **health system learning**.

Nature Medicine, 10 July 2026, article s41591-026-04497-1. The triage percentages come from secondary coverage of the paper rather than the paywalled text and should be confirmed against the article before use in a client document.

HEALTHSEED PERSPECTIVE

It inverts the assumed hierarchy. The prevailing assumption has been that frontier models improve faster than anything a specialist can build, and that healthcare's job is to fine-tune what arrives. NeuroVFM says that in a domain where the data cannot legally leave the institution, the institution holds an asset the frontier cannot buy.

If a hospital's own archive is a viable pretraining corpus, then twenty years of scans is not exhaust. It is capital, and hospitals have started to notice.

Safia Agueni, Co-founder, Senior Partner

The benchmark gap nobody has closed

92%
LARGE LANGUAGE MODELS ON STANDARDISED MEDICAL LICENSING EXAMINATIONS

44.8%
THE SAME MODELS ON REAL-WORLD CLINICAL TASKS, BRIDGE BENCHMARK

A separate head-to-head study published in Nature Medicine on 23 June compared general-purpose models against two cleared clinical AI tools on questions submitted by practising physicians. The general-purpose models outperformed both cleared tools across every benchmark tested. Read together: the exam score is not the capability, and clearance is not currently a proxy for performance on the query a clinician actually asks.

BRIDGE benchmark, Mass General Brigham, Nature Biomedical Engineering. Head-to-head study, Nature Medicine, 23 June 2026. Both reported via secondary coverage and worth reading directly.

03 Voices

Three published positions inside four days, arguing with each other. This is the most useful disagreement in the field right now and it is worth reading in order.

Aviv Goldenberg and Jenna Wiens

Nature Medicine

Is AI actually improving healthcare?

A qualified yes with a substantial caveat: in many cases we do not know, because a large share of deployed tools are too new for anyone to say whether they improve patient outcomes.

Nat Med 32, 1182 to 1183 (2026). DOI 10.1038/s41591-026-04329-2

Priya Abani

Chief executive, AliveCor · STAT First Opinion, 6 August 2026

Stop asking if AI is good for medicine

The question is malformed, because healthcare treats AI as a monolith when it is not. Asking whether AI improves healthcare is like asking whether lasers improve surgery: in the hands of a skilled surgeon using a validated tool they allow lifesaving precision, and in an unproven setting the question is still open. Her concern is that overgeneralisation slows the adoption of tools that demonstrably work.

Declared interest: AliveCor sells AI-powered cardiology products. She is arguing a position her company benefits from, which does not make it wrong and should be stated.

Frances Mei Hardin

Physician · STAT First Opinion, 7 August 2026

AI won't enhance physician autonomy. It will further diminish it

A direct rebuttal of the most common argument for clinical AI, that automating documentation returns time and agency to clinicians.

WHERE WE COME OUT

Abani is right about the category error and it does not resolve the problem she is pointing at. "AI is not a monolith" is true, and it is also exactly what every vendor whose product does not work would say. The distinction becomes useful only once someone specifies which tools, evaluated how, against which outcome, and that specification is what the Goldenberg and Wiens caveat says is missing for most of the field.

Hardin's objection is the one the industry keeps answering with an assertion rather than with evidence. The time-returned argument has been made for a decade with remarkably little published outcome data behind it. Someone will eventually run that study properly, and the result will move procurement more than any capability benchmark.

Safia Agueni, Co-founder, Senior Partner

04 Money

Two rounds dominated, both consumer-facing, both European in origin, and both worth reading for what the paperwork says rather than for the number.

Neko Health, $700m Series C, 15 July

The Swedish full-body scanning company raised $700 million, co-led by Lightspeed Venture Partners and O.G. Venture Partners, with Atomico, General Catalyst and Lakestar returning and Liberty City Ventures, BDT & MSD and Positive Sum joining. Reported near a $7 billion valuation against $1.8 billion at the $260 million Series B in January 2025. Disclosed funding since 2023 now exceeds $1 billion. The money opens the first US clinic, in New York.

The detail that matters is in the FDA record. Two Neko devices cleared through the 510(k) pathway in May 2026: Derma-2 as an adjunctive telethermographic system, and Spectrum-2 as a tissue-saturation oximeter. Those clearances cover two named devices for their stated intended uses. They do not clear the Neko Health Scan as a screening service, and Neko describes its US clinics as preventive health and wellness providers rather than full-service medical practices.

CORRECTION TO THE RECORD

Bunkerhill Health, 16 July. Khosla Ventures led, with Sequoia, Felicis, Optum Ventures and Y Combinator continuing. Carebricks runs at Cleveland Clinic, Intermountain Health and UTMB, where more than twenty agents are live.

The company's announcement states the Series B brings **total funding to date, including Seed and Series A, to $55 million**. At least four outlets reported a "$55M Series B".

The round size has not been disclosed. If you are benchmarking a raise against published comparables, one of yours is wrong.

TechCrunch and Axios, 15 July 2026. FDA 510(k) database, May 2026. Bunkerhill Health announcement, 16 July 2026. Misreporting visible across Becker's Hospital Review, HIT Consultant and The SaaS News.

HEALTHSEED PERSPECTIVE

Both raised on consumer demand and both carry a regulatory perimeter narrower than the proposition sold. Neko's clearances cover two sensors, not the scan. Bunkerhill's platform is described as enabling cleared algorithms, which is a different claim from the platform being cleared.

Neither is doing anything improper. Both are doing what every well-advised company in this category does, holding the product claim inside what the clearance supports while the marketing operates one level above it. If you are raising here, that gap is where diligence will concentrate, and the answer needs to exist in writing before the question arrives.

Daren Wilson, Co-founder, Senior Partner

05 Adoption reality

Wolters Kluwer's 2026 half-year report, published 5 August, states that as of 31 July more than 90% of its US Enterprise Edition customers, representing approximately 2,500 hospitals, **have signed up to adopt** the UpToDate Expert AI version, against a mid-year goal of 70%. Internationally, more than 230 sites are activated across 36 countries.

Three days later the company's own newsroom carried the same fact under a different verb: **adoption has expanded to** approximately 2,500 hospitals and health systems.

Signed up to adopt	Adopted	Used, and changed something
A contract exists. Someone has agreed to switch the feature on at some point.	The feature is enabled and available to clinicians in the workflow.	Clinicians open it, and a measurable outcome moves as a result.
PUBLISHED	NOT PUBLISHED	NOT PUBLISHED

Wolters Kluwer 2026 Half-Year Report, 5 August 2026. Wolters Kluwer newsroom, 8 August 2026. Both are the company's own disclosures and the wording differs between them.

THIS IS THE NUMBER TO HOLD UP IN YOUR OWN BOARD PACK

We are not picking on Wolters Kluwer. The half-year report is precise, and precision is what a regulated financial disclosure requires. The point is what happens to that precision three days later, and three hundred times after that, as the figure travels.

A health system that has been through two rounds of AI procurement will ask how many of the 2,500 have the feature enabled, how many clinicians opened it last month, and what changed as a result. Having answers to the second and third is now a differentiator, because almost nobody publishes them. If your own deck carries a deployment number, check which of the three it is before someone else does.

Safia Agueni, Co-founder, Senior Partner

06 The argument

Three things happened in a month and they point the same way.

A regulator deferred the obligations that are hardest to meet and kept the ones that are easiest to breach. A hospital demonstrated that its own unglamorous archive outperforms the frontier at the task that matters clinically. And the field's most credible voices published, within four days of each other, three incompatible answers to whether any of this is working.

For a company selling clinical AI in Europe over the next eighteen months, the practical consequence is that the evidence which will sell the product is not the evidence that gets it approved. Both have to be built, they are built from different work, and only one of them has a deadline attached.

The common thread is that the proxies have stopped tracking the thing.

A CE mark	An exam score	A deferred deadline	A deployment count
does not tell a buyer the tool performs.	does not tell a developer the model is clinically useful.	does not tell a compliance lead they are safe.	does not tell anyone whether a patient was better off.

HealthSeed

Do you know which of the three still applies to you?

The AI Act obligations that kept their original schedule apply now, not in 2028. Establishing which of them reaches your product is a days-long piece of work against a documented answer, and it is the one thing on this page that can be closed before the next board meeting.

https://healthseed.vc/launchpad
https://www.healthseed.vc/insights/vital-signs-ai-healthcare-august-2026

Vital Signs: AI Healthcare is produced by HealthSeed AG, Weidmannstrasse 5, 8046 Zurich, CH-020.3.056.106-6. HealthSeed takes medtech, diagnostics and digital health companies from an approved product to a paying European customer, running regulatory, reimbursement and commercial as one sequence. Every claim here carries a source you can open; where a figure comes from secondary coverage rather than the paper itself, the source line says so.