Vital Signs: AI Healthcare

Clinical AI and its regulation

HealthSeed

HEALTHSEED INTELLIGENCE

Edition 002 • 11 August to 13 September 2026 • Published for the HealthSeed ecosystem

Two FDA authorisations and a Medicare refusal

THIS EDITION'S SIGNAL

The FDA and the MHRA's commission pushed the demand for evidence past the day a product launches, and Medicare refused to pay for a cleared device.

Two regulators spent this cycle building approvals that do not finish. A company holding a twelve-year-old FDA clearance asked Medicare for a billing code, put its objection in writing when the answer came back no, and was refused again, while a different class of software now collects a four-figure national rate. Four large studies showed that medical software reads images accurately, and one of them followed what then happened to the patient.

Underneath those events is a change in who defines sufficient evidence, and when. Approval used to close an evidence programme. It is becoming a state that keeps asking, because an authorisation can stay open and go on requiring data from the field, and a billing code has to be argued for before anyone pays.

Payment follows the same logic more exactly than most plans assume. Where a specialty organised itself and carried a billing code through, the money arrived. Where a device fell outside the category Medicare may pay for at all, no coding effort would have reached it.

$1,000

The 2026 Medicare rate for AI coronary plaque analysis, where a specialty built the code

43

AI-related codes in CPT after the AMA added ten on 9 September 2026

IN THIS EDITION

01 The approval terms

On 3 September the FDA approved a piece of heart software and agreed in advance that the company could keep changing it. A week later a commission reporting through the MHRA asked for approval to be handed out in steps. Both decisions push the demand for evidence past the day a product launches, which lands on any company whose evidence work stops when the submission is filed. The response is to name someone now who owns how the product performs in use, while it is still an internal choice and not a condition someone else has written.

The FDA approved two AI devices their makers may keep changing

The product is the STEMI AI ECG Model from Powerful Medical, granted through the De Novo route, used when a device is new enough that no similar product exists to compare it against (FDA, 3 September 2026). The record states that a predetermined change control plan was authorised, meaning permission agreed in advance to update the software later without going back for a new approval. Eight days earlier the agency had cleared Heartvue.Proton from Heartvue.ai under the 510(k) route, where approval rests on a product being close enough to something already on the market, and that clearance carried a plan of the same kind (FDA, 26 August 2026).

Neither record states what the plan permits. That scope sits in the decision summary, which could not be retrieved. This edition therefore describes nothing about what either company may now change. Powerful Medical's own website still calls the model investigational and pending approval, and the agency database says otherwise. The discrepancy is unresolved.

The MHRA's commission asked for staged approval, and no government response exists yet

The National Commission into the Regulation of AI in Healthcare reported through the MHRA on 10 September, and asked for three things (MHRA, 10 September 2026):

- Evidence requirements covering a product's whole life, and not only the file that wins the approval
- A public register showing how a tool performs once it is deployed
- Approval granted in steps, with each step released by evidence the regulator has specified in advance

These are recommendations only. The publication states that "a cross-government response will follow separately", so no commitment and no timetable exist yet.

Regulation (EU) 2026/1744 set 2 August 2028 for high-risk systems sitting inside regulated products, which is the Annex I category an AI component in a medical device falls into (Official Journal, 24 July 2026). That date gives a company time to build the function, and it is not permission to arrive without one.

CORRECTION TO THE RECORD

The FDA marked its AI-Enabled Medical Devices list as updated on 4 September, and parts of the trade press read the refresh as a wave of new clearances. The most recent final decision on that list is dated 29 June 2026. If your competitive tracking depends on that page, move it to the De Novo and 510(k) databases this week, because the two authorisations that landed in this window appear only there.

Applies to you if: you hold a cleared AI product and nobody owns the post-market performance evidence that approval in stages or a public register would require.

02 Payment

Medicare now pays more than a thousand dollars when software measures the build-up inside the arteries of the heart from a scan. In the same cycle it looked at an approved home urine-flow device and paid nothing at all. The two outcomes turned on questions that have nothing to do with how well either product works, which is why a company can hold a clearance and still have no route to revenue. What that asks for, before the clearance strategy is fixed, is an answer to which class of thing the payer may pay for at all, and to who would build the code.

Twelve payer and coding decisions in thirty-four days

Every in-window action by a payer or a coding body in the AI Healthcare inventory. Vendor and product announcements filed under the same theme are excluded. Source: HealthSeed intelligence inventory, 13 September 2026.

DATE	BODY	WHAT IT DID
08-12	Centers for Medicare & Medicaid…	CMS updated its ACCESS Model accepted-applicants page, confirming more than 150…
08-12	Centers for Medicare & Medicaid…	CMS updated the ACCESS model, an outcome-aligned Original Medicare payment model that…
08-13	Aidoc	Aidoc announced that CMS had approved New Technology Add-on Payment eligibility for its…
08-14	Centers for Medicare & Medicaid…	In an updated B1 2026 HCPCS determination, CMS assigned the AI-powered MenHealth…
08-17	American Medical Association (CPT…	The AMA updated its announcement that the CPT Editorial Panel had revised Appendix S,…
08-19	American Medical Association (AMA)	The AMA published revised CPT Appendix S guidance classifying AI-enabled medical…
08-21	G-BA Innovationsfonds…	The Innovations Committee recorded a decision on CASSANDRA, a funded inpatient research…
08-31	HM Treasury, Cabinet Office, and…	The UK Government launched the £100 million Sovereign AI R&D Procurement Scheme,…
09-06	Centers for Medicare & Medicaid…	CMS updated its billing-and-coding article for AI-enabled CT quantitative coronary…
09-07	Centers for Medicare & Medicaid…	CMS updated a second AI-QCT/AI-CPA billing-and-coding article used by a Medicare…
09-09	American Medical Association (AMA)	The AMA released the CPT 2027 code set, adding 10 AI-related CPT codes and bringing the…
09-10	Centers for Medicare & Medicaid…	CMS updated the WISeR model page, confirming that the Medicare payment-integrity model…

Paid

- CPT 75577, permanent code since 1 January 2026
- Above $1,000 on the 2026 national price list for work in a doctor's office, $950.50 in hospital outpatient
- NTAP for Aidoc CARE Body CT, up to $137.53 a case from 1 October 2026
- Final coverage decisions from five of seven regional contractors

Not paid

- HCPCS A9279, a code that already existed
- No Medicare benefit category, because the device does not treat
- Not covered under Part B, and no Medicare payment
- Pricing indicator 00

Medicare pays over $1,000 for AI plaque analysis because cardiology built the code

The case against this movement comes first, and the panel above carries it. On 1 January 2026 the temporary trial codes for AI coronary plaque analysis were retired, and the service moved to CPT code 75577, a permanent billing code as against the temporary kind used while a service is still being evaluated (Society of Cardiovascular Computed Tomography; American College of Radiology). Payment for clinical AI plainly exists. It exists because a specialty society assembled the evidence and carried a code through a process that takes years.

CMS found no benefit category for a cleared device, so it pays nothing

BE Technologies applied for a new code for MenHealth, a home urinary flowmeter that the company says uses sound analysis to produce clinical-grade measurements, and which the FDA cleared in 2014. The Centers for Medicare and Medicaid Services assigned it to the existing code A9279, covering a monitoring device "not otherwise classified". BE Technologies disagreed on the record, arguing that the product performs a standardised diagnostic test. CMS finalised its determination unchanged, finding no Medicare benefit category, meaning the class of thing Medicare is legally allowed to pay for at all, because the device "does not provide treatment for a disease or injury". Items under A9279 "are not covered under Medicare Part B", and the payment determination reads "No Medicare payment" (CMS, First Biannual 2026 HCPCS cycle, 14 August 2026).

MenHealth failed an earlier test than the coding one. CMS decides the benefit category on what a device does. A device that measures without treating falls outside the equipment category before any question of a code arises.

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

Two codes, one cycle. 75577 pays above $1,000 nationally. A9279 carries no benefit category and pays nothing. Before your next clearance strategy is signed off, establish which benefit category your product falls into and name the society that would sponsor a code for it. If neither answer exists today, your revenue model is a clinical plan with no payer attached to it.

Applies to you if: your financial model assumes a reimbursement code follows a clearance, in any market where the payer and the regulator are separate institutions.

03 Evidence

Four large studies of medical software were published in six weeks, three measuring how well the software spots what it is looking for and one measuring what happened to the people in front of it. That split matters commercially, because the first satisfies a regulator and the second is what a payer or hospital asks for when deciding whether to keep paying. A company whose main study is still built around the first kind should add a measure of what changes in the service, before the protocol locks.

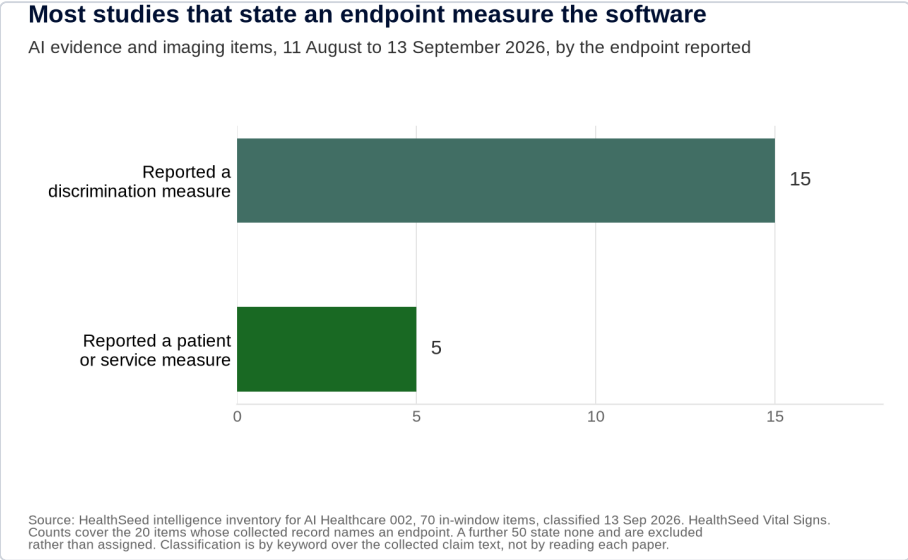

Figure 2. Source: HealthSeed intelligence inventory, 70 in-window items, classified 13 September 2026.

Alibaba, MetaOptima and Loyola published scores the software earned

The largest came from Shengjing Hospital of China Medical University and Zhejiang University with Alibaba's DAMO Academy, in *Nature Medicine* on 19 August. Its Liver DiagnOsis Network read contrast-enhanced CT scans for more than 10,000 patients in a single-arm trial, meaning there was no comparison group, so it shows what happened with the software and not what would have happened without it. The system flagged 51 liver lesions overlooked in the original reports, 15 of them cancer. Thirty-seven reports were amended (Nature Medicine, 19 August 2026). The paper is paywalled. Neither the patient count of 10,333 nor the accuracy score of 0.952 could be confirmed, so both come from the collected pack.

Royal Cornwall Hospitals Trust and MetaOptima reported that DermDx found 296 of 298 confirmed cancers, a sensitivity of 99.3%, across 1,024 lesions on the NHS suspected skin-cancer pathway. Sensitivity across malignant and premalignant lesions together was 98.1%, and specificity was 57.6%, which is how often it correctly leaves alone people who do not have the condition (Skin Health and Disease, 27 August 2026). The images were collected as patients came through the clinic, then analysed afterwards and blinded. The software never influenced anyone's care.

Radiologists at Loyola University Chicago tested Siemens Healthineers' XProstate research prototype on 202 patients against their own reading. The software scored 0.76 for separating cancer from benign tissue and the radiologists 0.73, on a scale where 0.5 is a coin toss and 1.0 is perfect. The gap of 0.03 is not statistically significant, and the authors report a p-value of 0.38 (Medical Physics, 21 August 2026).

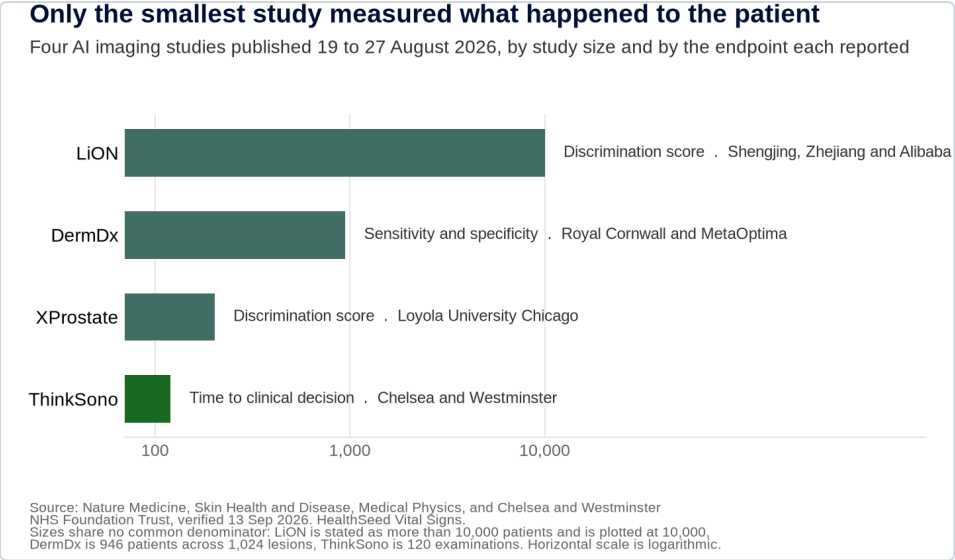

Figure 1. Source: Nature Medicine, Skin Health and Disease, Medical Physics, and Chelsea and Westminster NHS Foundation Trust, verified 13 September 2026.

Only Chelsea and Westminster measured what happened to the patient

Over a three-month pilot the trust's AI-supported handheld ultrasound pathway for suspected blood clots in the leg produced 120 examinations and ruled out clots in 97 of them. It discharged 80% of patients after a single scan. The typical time to a clinical decision was under 30 minutes. For patients needing further investigation, the average wait from first presentation to diagnosis fell from 19 hours to 9 (Chelsea and Westminster, 26 August 2026). None of those numbers describes the model, and all of them describe the service around it.

WHERE WE COME OUT

The case for discrimination endpoints is strong and we grant it. They are what a regulator asks for, they cost less, and an imaging trial powered for a clinical outcome can take years longer. We still come out with Chelsea and Westminster. If your primary endpoint is currently a score, add a service measure to the protocol before it locks, because a payer will ask for one and retrofitting it means paying for the study twice.

Applies to you if: you are building the evidence plan for a diagnostic or triage product and the primary endpoint is currently a discrimination statistic.

04 Documentation

Three health systems moved software that listens to appointments and writes the notes from trial into everyday use inside four weeks, and one of the efficiency figures being used to justify that shift turns out to have come from the vendor rather than from the health service that ran the trial. Deployment is running ahead of the evidence underneath it, and the exposure sits with whoever signed the contract as much as with the vendor. Anyone buying or selling this software should get an accuracy monitoring clause into the contract before the next site goes live.

Thirty-four documentation items in the window, and eleven name the vendor

Every in-window ambient documentation and administrative AI item in the inventory, dated. Source: HealthSeed intelligence inventory, 13 September 2026.

08-11 LCMC Health and Qualified Health
08-11 Vivos Therapeutics, Inc.
08-11 Cleveland Clinic
08-11 Cleveland Clinic
08-12 PointClickCare
08-12 Suvi Health
08-12 Suvi Health
08-13 athenahealth
08-13 athenahealth
08-17 Cerbo
08-17 Hertfordshire Community NHS Trust
08-17 Evaxion A/S and Duke University School of…
08-18 PhyxUp Health
08-19 Summa Health, Ohio
08-19 Matic
08-21 Penn Medicine
08-23 Cleveland Clinic Foundation
08-24 DeepCura
08-24 FDB (First Databank) and Tebra
08-24 LUVN and Tandem Health
08-25 NHS clinicians get more consultation time…
08-25 Suki
08-25 American Academy of Orthopaedic Surgeons…
08-25 Aiforia Technologies Plc and Stratipath
08-26 University of Wisconsin–Madison and UW Health
08-28 Penn Highlands Healthcare; Doximity
08-28 Penn Highlands Healthcare and Doximity
08-28 Doximity; Penn Highlands Healthcare
08-31 Hertfordshire Community NHS Trust (UK)
09-03 Commure Ambient AI becomes commercially…
09-03 Metro South Health
09-08 Dove Press / *Journal of Healthcare…
09-09 Avo and MEDITECH
09-09 Nivaran, formerly ScribeEMR

Cleveland Clinic and two NHS trusts moved AI note-taking into everyday use

Cleveland Clinic published its account of deploying Ambience to more than 4,000 clinicians, reporting 70% use at the level of the individual appointment among established users, and 60% agreeing that the tool had increased their likelihood of remaining in practice (npj Health Systems, 11 August 2026). Royal Wolverhampton and Walsall selected Heidi Health after an eight-month pilot, rolling it out across the emergency department, medical same-day emergency care, dermatology, ENT and audiology, and obstetrics and gynaecology (Walsall Healthcare, 25 August 2026).

Metro South's headline saving came from the vendor, not from its trial

Metro South Health in Queensland announced a service-wide rollout of the same vendor's system, reporting that 92% of participating clinicians spent less time on administration. The figure of 59 minutes saved per clinician per day has been widely repeated as a Metro South finding. The outlet that first reported it has since corrected the record, saying the figure came from Heidi Health's own research and was not a finding of the health service's trial (MobiHealthNews). The 92% did come from the trial, which covered 70 clinicians in 2025.

Wisconsin researchers found only moderate patient trust in the tools being deployed

Patients are reaching the question before the industry has settled it. A study of patient acceptability and trust from the University of Wisconsin-Madison reports that 60% of those surveyed held only moderate trust in these tools. The article is behind a paywall. That figure comes from the collected pack and not from the paper itself (Applied Clinical Informatics, 26 August 2026).

HEALTHSEED PERSPECTIVE

The retention finding comes from a health system reporting on its own deployment, alongside the vendor it partnered with, and one institution's experience does not generalise. That is granted before the position that follows. It is still the number to carry into a procurement meeting, because a workforce argument survives a budget review that a time-saving argument does not, and this cycle has just shown how a time-saving number can trace back to the vendor. Get an accuracy monitoring clause into the contract now, on whichever side of it you sit.

Applies to you if: you are selling documentation AI into a health system, or buying it, and the contract has no accuracy monitoring clause.

05 Capability and deployment

Two federal awards, two published evaluations and one European surgical rollout landed inside three weeks, and together they circle the same question: what a clinical AI system does when it should not answer. The money is going into architectures that keep clinical authority outside the model, while the evaluations found that models do not reliably decline in the places they are most likely to be wrong. A company building an agentic product should be able to say in one sentence where its clinical authority sits.

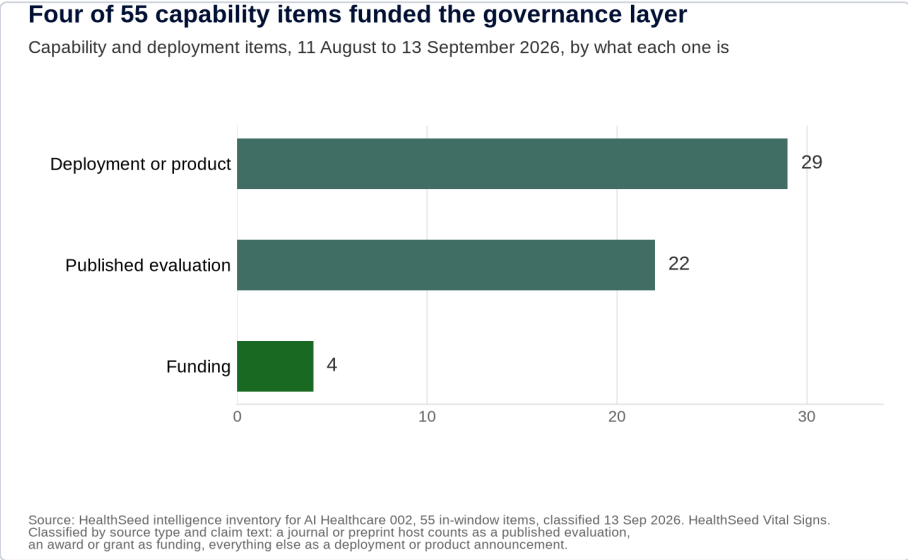

Figure 3. Source: HealthSeed intelligence inventory, 55 in-window items, classified 13 September 2026.

ARPA-H put $16.7 million behind keeping clinical authority outside the model

On 9 September ARPA-H made two awards under its ADVOCATE programme, up to $9 million to UpDoc and up to $7.7 million to Atman Health, both for agentic cardiovascular care. UpDoc's record describes "a clinical AI agentic system whose conversational intelligence is separated from clinical authority by a clinician-built rules system that validates every proposed action against approved protocols before execution" (ARPA-H, 9 September 2026). Atman Health's is built on an evidence-based clinical decision engine behind a voice-first interface (ARPA-H, 9 September 2026).

Two evaluations found models declining in the wrong places

The *Journal of Managed Care and Specialty Pharmacy* tested five models against 250 clinician-curated medication lists, each holding one clinically relevant interacting pair. Accuracy ran from 54.1% to 83.7%. The finding that matters sits elsewhere: alignment between a model's refusal behaviour and its likelihood of error was weak to moderate, and prompting for caution did not reliably recalibrate it (JMCP, 2026;32(9):1076). Beth Israel Deaconess released mental-health subsets of HealthBench on 25 August, 610 conversations screened from 5,000 and validated through two rounds of blinded clinician review, reporting a statistically tied frontier cluster and measurable refusal behaviour in two of twenty models (arXiv, 25 August 2026).

Medtronic ran AI-planned spine surgery at seven European sites

On 10 September Medtronic reported more than 23 procedures with its Stealth AXiS system across seven sites in Germany, the United Kingdom, Italy and Spain, combining "AI-driven surgical planning, real-time navigation and robotic-assisted execution" (Medtronic, 10 September 2026). The Saudi Food and Drug Authority separately authorised Dental IQ and SAARIA, the first Saudi-developed AI medical software products of their kind (SFDA, 10 September 2026).

HEALTHSEED PERSPECTIVE

Two evaluations and two funding decisions is a thin base, and neither evaluation tested a deployed product under clinical governance, which is the configuration ARPA-H is paying to build, and both points are granted before the position that follows. The pattern is still worth acting on, because a buyer who has read either paper will ask what your system does with a question it should decline, and "the model handles it" is the answer that loses the room. Write the refusal path into the product description before the first procurement conversation.

Applies to you if: you are building or buying an agentic clinical product and cannot yet say, in one sentence, where its clinical authority sits.

The argument

Five institutions acted this cycle, and each asked for the same thing in a different language. The FDA approved software that its maker may keep changing, so the approval now depends on what the company does after launch. The MHRA's commission put it more plainly still, asking for approval in stages and a public register showing how a tool performs once it is in use. Medicare ruled that a device sat outside the category it may pay for at all, which is a question no clearance answers. Health systems buying documentation software want to know it still works in their own building. Patients want to be told it is being used at all.

A well-informed reader could object that this is ordinary institutional behaviour, and that it was ordinary before this cycle opened. The objection is fair, and each event is unremarkable on its own. What changed is that the evidence these institutions want has become specific enough to design a study around, and expensive enough that designing for it late means paying for the work twice.

What follows for an operator is narrow, and it is not a strategy. Two questions belong on next week's agenda. Which benefit category does the product fall into, and who would sponsor a billing code for it? Who owns the performance data after launch, by name, with time in their week for it? A clearance answers neither, and both answers cost more the longer they wait.

About HealthSeed

HealthSeed is a Swiss healthcare venture studio and commercialisation partner. We work alongside biotech, medtech, diagnostics, digital health, and AI companies as they enter and scale across Europe, the United States, and the Middle East and North Africa, backed by an Expert Community of more than 35 specialists and a network of partners who have run the functions they advise on. We also build ventures and AI products of our own, so we read this market as participants. Vital Signs is where we publish what we are reading and what we think it means, with a primary source behind every claim.

HealthSeed AG, Weidmannstrasse 5, 8046 Zürich, Switzerland. CH-020.3.056.106-6.

Sources

Every claim here carries a source that can be opened. Where a figure could not be confirmed against the source itself, the entry says so and the body says so in the line where the figure occurs.

1. FDA, De Novo DEN250044, STEMI AI ECG Model, Powerful Medical, decision 3 September 2026. Primary regulator, retrieved. https://www.accessdata.fda.gov/scripts/cdrh/cfdocs/cfpmn/denovo.cfm?id=DEN250044
2. FDA, 510(k) K260811, Heartvue.Proton, Heartvue.ai Inc., decision 26 August 2026, change control plan authorised. Primary regulator, retrieved. https://www.accessdata.fda.gov/scripts/cdrh/cfdocs/cfpmn/pmn.cfm?ID=K260811
3. MHRA, National Commission into the Regulation of AI in Healthcare, recommendations, 10 September 2026. Primary regulator, retrieved. https://www.gov.uk/government/publications/national-commission-into-the-regulation-of-ai-in-healthcare-recommendations-for-a-future-regulatory-framework
4. Official Journal of the European Union, Regulation (EU) 2026/1744, OJ L, 2026/1744, 24 July 2026. Primary, retrieved. https://eur-lex.europa.eu/eli/reg/2026/1744/oj/eng
5. FDA, AI-Enabled Medical Devices list, marked updated 4 September 2026, most recent final decision 29 June 2026. Primary regulator, retrieved. https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-enabled-medical-devices
6. Society of Cardiovascular Computed Tomography, CPT 75577 and 2026 Medicare rates. Secondary coverage, retrieved. https://scct.org/news/684550/CMS-updates-Medicare-coverage-for-AI-Coronary-Plaque-Analysis.htm
7. American College of Radiology, final local coverage determinations on AI-enabled CT. Secondary coverage, retrieved. https://www.acr.org/News-and-Publications/Medicare-Contractors-Release-Final-LCDs-on-Artificial-Intelligence-Enabled-CT-Technology
8. CMS, HCPCS Level II Final Coding, Benefit Category and Payment Determinations, First Biannual (B1) 2026 cycle, applicant BE Technologies, device MenHealth, code A9279, 14 August 2026. Primary regulator, read at first hand, pages 18 to 20. https://www.cms.gov/files/document/2026-hcpcs-application-summary-biannual-1-2026-non-drug-non-biological-items-services.pdf
9. *Nature Medicine*, "Large-scale AI-guided liver malignancy diagnosis: multicenter study and a single-arm trial", Shengjing Hospital, Zhejiang University and Alibaba DAMO Academy, 19 August 2026. Published evidence, paywalled; patient count and accuracy score not confirmed. https://www.nature.com/articles/s41591-026-04589-y
10. *Skin Health and Disease*, "Performance of the 'DermDx' algorithm in triaging suspected skin cancers: results of the AI-SCSS study", Anderson et al., Royal Cornwall Hospitals NHS Trust and MetaOptima, 27 August 2026. Published evidence, retrieved. https://academic.oup.com/skinhd/advance-article/doi/10.1093/skinhd/vzag122/8771688
11. *Medical Physics*, "Real-world assessment of a deep learning neural network algorithm for prostate cancer detection in MRI using true de novo data", Shea et al., Loyola University Chicago and Siemens Healthineers, 2026;53(8):e70614. Published evidence, retrieved. https://pmc.ncbi.nlm.nih.gov/articles/PMC13431142/
12. Chelsea and Westminster Hospital NHS Foundation Trust, ThinkSono pathway, 26 August 2026. Health system announcement, retrieved and confirmed verbatim. https://www.chelwest.nhs.uk/about-us/news/ai-powered-ultrasound-helps-most-suspected-blood-clot-patients-receive-a-clinical-decision-within-30-minutes-at-chelsea-and-westminster-hospital
13. Cleveland Clinic Foundation, "Accelerating ambient AI scribe enterprise-scale deployment", *npj Health Systems*, 11 August 2026. Published evidence, retrieved and confirmed verbatim. https://pmc.ncbi.nlm.nih.gov/articles/PMC13462757/
14. Walsall Healthcare NHS Trust, ambient voice technology rollout with Heidi Health, 25 August 2026. Health system announcement, retrieved and confirmed. https://www.walsallhealthcare.nhs.uk/news/2026/08/25/new-digital-ambient-voice-technology-avt/
15. MobiHealthNews, correction to the 59-minute figure attributed to Metro South Health. Secondary coverage. https://www.mobihealthnews.com/news/anz/queensland-health-services-ai-scribe-saves-doctors-59-minutes-daily
16. *Applied Clinical Informatics*, "Ambient Artificial Intelligence Scribes in Ambulatory Care: Patient Acceptability and Trust", Nguyen et al., University of Wisconsin-Madison, 2026;17(4):701-710, 26 August 2026. Published evidence, publisher returned 403; trust figure not confirmed. https://doi.org/10.1055/a-2936-9935
17. Saudi Food and Drug Authority, marketing authorisation for Dental IQ and SAARIA, 10 September 2026. Primary regulator, retrieved and confirmed. https://www.sfda.gov.sa/en/news/19462
18. ARPA-H, ADVOCATE award to UpDoc, TREAT-CVD, up to $9 million, 9 September 2026. Primary, retrieved and confirmed verbatim. https://arpa-h.gov/explore-funding/awards/4702
19. ARPA-H, ADVOCATE award to Atman Health, up to $7.7 million, 9 September 2026. Primary, retrieved and confirmed verbatim. https://arpa-h.gov/explore-funding/awards/4700
20. *Journal of Managed Care & Specialty Pharmacy*, "Navigating uncertainty matters: Evaluating large language models for drug-drug interaction identification", 2026;32(9):1076. Published evidence, retrieved and confirmed. https://doi.org/10.18553/jmcp.2026.32.9.1076
21. Flathers et al., "HealthBench-Psych: A Mental Health Subset of OpenAI's HealthBench", Beth Israel Deaconess Medical Center, 25 August 2026. Published evidence, retrieved and confirmed verbatim. https://arxiv.org/abs/2608.25071
22. Medtronic, first procedures in Europe with the Stealth AXiS Surgical System, 10 September 2026. Company announcement, retrieved and confirmed verbatim. https://news.medtronic.com/Medtronic-announces-first-procedures-in-Europe-with-Stealth-AXiS-TM-Surgical-System