AI Adoption in US Hospitals Statistics: Survey vs. Deployment

IntuitionLabs.ai · Adrien Laurent · 2026-09-05 · ai adoption hospitals · hospital ai statistics · ama physician ai survey · aha ai adoption survey · generative ai healthcare adoption · clinical ai deployment · predictive ai ehr · ai adoption survey methodology · ambient ai documentation · healthcare ai statistics 2026

Executive Summary

There is no single correct answer to "how many US hospitals use AI," because the surveys and studies that produce these figures measure different populations, technologies, years, and stages of deployment, and this report reconciles them rather than adding another blended number. For predictive AI (machine-learning risk scoring) integrated into the electronic health record (EHR), a federal data brief found **71%** of hospitals in 2024, up from **66%** in 2023 (healthit.gov). For generative AI specifically, an independent 2024-data study published in *JAMA Network Open* in December 2025 found only **31.5%** of hospitals were current users, with another **24.7%** planning adoption within a year (pmc.ncbi.nlm.nih.gov). For any AI across five operational categories in 2022, a *Health Affairs Scholar* study found just **18.7%** (academic.oup.com). At the physician level, the American Medical Association (AMA)'s three-wave national survey found the share of physicians reporting any awareness or use of AI rising from **38%** (2023) to **66%** (2024) to **81%** by early 2026 (ama-assn.org).

Adoption is also sharply uneven. Federal data show large hospitals (400+ beds) at **96%** predictive-AI adoption in 2024 versus **59%** for small hospitals, and system-affiliated hospitals at **86%** versus **37%** for independent hospitals (healthit.gov). This report's central methodological contribution is a source-by-source comparison table that records each survey's sample, its exact definition of "adoption," and the deployment stage it actually captures, distinguishing tools that are merely licensed or piloted from those in routine active use. The clearest illustration is ambient AI clinical documentation: a 2025 survey of 43 health systems found **100%** had at least piloted it, yet only **14%** had it fully deployed (essentialhospitals.org), and the Peterson Health Technology Institute found that even where hospitals offered ambient scribes to all clinicians, ongoing active use typically settled at **20% to 50%**, with one system reporting only **15%** (phti.org).

Industry-panel surveys generally report higher figures than the federal, probability-sampled data: KLAS Research found the share of organizations reporting AI use climbed from under half to more than two-thirds across 2025 (engage.klasresearch.com), and Deloitte's late-2025 survey of 120 health care C-suite executives found **49%** still experimenting with generative or agentic AI against roughly a third operating it at scale (deloitte.com), a gap attributable in large part to sampling toward larger, more digitally mature organizations. Separately, the **1,615** devices on the Food and Drug Administration's AI/ML-Enabled Medical Device List, current as of September 4, 2026, measure marketing authorization, not hospital deployment, a distinction the FDA itself states explicitly (fda.gov). Readers and researchers citing an adoption statistic from this space should carry its sample, year, and deployment-stage definition alongside the figure itself, an approach consistent with the adoption-measurement principle that provisioned or licensed use is not the same as active, routine use (intuitionlabs.ai).

Introduction and Background

Artificial intelligence (AI) has moved from an innovation-lab curiosity to a line item that hospital boards, physicians, and federal regulators now track annually. Yet a reader trying to answer a simple question, how many U.S. hospitals actually use AI, quickly finds that the honest answer depends on which of a dozen credible surveys is consulted, what year it covers, and what activity it counts as "use." One federal data brief reports that **71%** of hospitals use predictive AI in 2024 (healthit.gov). A peer-reviewed analysis of a different year of the same government survey program found that only **18.7%** of hospitals had adopted any AI at all across five operational workflows (academic.oup.com). A national physician survey puts the share reporting any awareness or use of AI in their practice at **81%** as of early 2026 (ama-assn.org). None of these figures is wrong. They measure different populations (hospitals versus individual physicians), different technologies (narrow predictive risk scores versus generative AI versus "any AI"), different years, and, most consequentially, different points on the path from a signed vendor contract to a clinician who uses a tool routinely.

IntuitionLabs published an earlier analysis of AI adoption in U.S. hospitals in 2025 that catalogued sepsis-detection, documentation, and imaging deployments behind the headline federal figures (intuitionlabs.ai). This report does not restate that ground. It instead builds a reproducible, source-by-source comparison of every major U.S. hospital and physician AI adoption survey published between 2023 and early 2026, records each source's sample, exact definition, and the deployment stage it actually captured, and separates pilot programs and purchased-but-idle licenses from routine clinical use and demonstrated outcomes.

The scope is deliberately narrow and quantitative: U.S. hospitals and the physicians who work in them, as measured by government data briefs, peer-reviewed studies, and named industry surveys, **as of September 2026**. Anecdote and single-institution case studies are used only where they illustrate a measurement point already

established by a broader dataset. The goal is that a reader, or another writer citing this piece, can trace any adoption percentage back to its originating sample and know exactly what it does and does not demonstrate.

Methodology: Defining Pilot, Purchase, Active Use, and Demonstrated Outcome

Survey percentages in this space routinely conflate four distinct facts, and this report treats them as separate categories throughout.

- **Purchased or licensed**: a hospital has a contract or license for an AI product, whether or not it is switched on for any clinician.
- **Piloted**: a limited group of users or a single department has tried the tool, typically without an enterprise rollout decision.
- **Active or routine use**: the tool is deployed and used as part of normal clinical or administrative workflow, at least by a defined population of eligible users.
- **Demonstrated outcome**: a controlled or quasi-controlled comparison has measured a change in a clinical, operational, or financial metric attributable to the tool.

Few surveys ask separately about all four. The **Office of the National Coordinator for Health Information Technology (ONC)**, reorganized in 2025 under the Assistant Secretary for Technology Policy (ASTP), asks hospitals whether predictive AI is "integrated into their EHR," a self-report closer to active use than to a pilot, but it does not verify clinician-level usage (healthit.gov). A 2025 survey of 43 health systems published in the *Journal of the American Medical Informatics Association* is unusual in asking the question directly, using a four-stage ladder of no activity, developing or piloting, deployed in limited areas, and deployed fully (essentialhospitals.org). That study found that **100%** of respondent health systems had at least begun piloting ambient AI documentation, yet only **14%** had it fully deployed, the largest gap of any use case it measured (essentialhospitals.org).

IntuitionLabs, an adjacent consultancy that advises life-sciences and healthcare organizations on AI governance and measurement rather than selling clinical AI software itself, uses a similar four-tier distinction in its own adoption-measurement work: **"Adoption is a progression. Provisioned users are not active users. Active users are not necessarily repeat users"** (intuitionlabs.ai). The framing is generic to enterprise AI programs, not specific to hospitals, but it names the exact conflation this report is built to avoid: an executive's press release about "AI use" and a chief nursing officer's report of clinicians opening the tool during a shift are not measuring the same thing.

A second, subtler issue is that two independent research teams analyzing the **same** underlying dataset in the **same** year can still publish different percentages, because of differences in analytic sample construction and statistical weighting. Later sections document one such case directly.

Hospital-Level Survey Data: ONC, AHA, and Health Affairs

The most frequently cited hospital-level figures trace back to a single instrument, the **American Hospital Association (AHA) Information Technology Supplement**, fielded annually alongside the AHA Annual Survey and analyzed both by ONC/ASTP directly and by independent academic teams using the same underlying microdata.

ONC/ASTP's Data Brief No. 80 (published September 2025) reports that predictive AI integrated into the EHR rose from **66%** of hospitals in 2023 to **71%** in 2024 (healthit.gov). The 2024 wave drew responses from **2,253** non-federal acute-care hospitals, a **51%** response rate, while the 2023 wave drew **2,547** hospitals at a **58%** response rate (healthit.gov). A separate *Health Affairs* study by Nong and colleagues, analyzing the **same 2023 wave** with an analytic sample of **2,425** hospitals, reported **65%** adoption of AI or predictive models integrated with the EHR, a one-point difference from ONC's own 2023 figure attributable to differences in weighting and sample construction rather than a change in the underlying phenomenon (healthaffairs.org). That study also found that among hospitals using predictive models, only **61%** locally evaluated them for accuracy and **44%** for bias, and that system-affiliated hospitals were substantially more likely to use predictive models than independent hospitals (healthaffairs.org).

Table 1 below lays out every major hospital- and physician-level source used in this report, its sample, its exact definition of adoption, and the deployment stage it actually measured, so the discrepancies above are traceable rather than mysterious.

Table 1. *Hospital and Physician AI Adoption Sources Compared, 2022 to 2026*

Source	Data Year(s)	Sample	What "Adoption" Means	Deployment Stage Captured	Reported Figure
AMA Physician AI Survey, three waves (ama-assn.org) (ama-assn.org)	2023, 2024, 2026	n=1,081 (2023); n=1,183 (2024); n=1,692 (2026)	Physician self-report of any awareness or use of AI; explicit use is reported separately as incorporating one or more named use cases	Awareness or use; not a formal pilot/production split	38% (2023) to 66% (2024) to 81% (2026)
ONC/ASTP Data Brief 80 (AHA IT Supplement) (healthit.gov)	2023, 2024	N=2,547 hospitals (2023); N=2,253 (2024)	Predictive AI (ML/statistical models) integrated into the EHR	Hospital reports the tool is in use; does not verify clinician-level usage	66% (2023) to 71% (2024)
Health Affairs (Nong et al.), same survey program, independent analysis (healthaffairs.org)	2023	N=2,425 analytic sample	AI/predictive models integrated with the EHR	Same "in use" self-report; adds model source and local evaluation detail	65%
Health Affairs Scholar (Baten et al.) (academic.oup.com)	2022	1,107 of hospitals in the AHA Annual Survey	Binary yes/no across 5 non-clinical operational workflows (demand prediction, workflow optimization, task automation, staffing, scheduling)	Any reported use in 1 of 5 categories; no depth or frequency measure	18.7% "any AI adoption"
JAMA Network Open (Everson, Nong, Richwine) (pmc.ncbi.nlm.nih.gov)	2024	N=2,174 hospitals	Generative AI/LLM integrated into the EHR, 3-tier classification	Distinguishes current users ("early adopters") from planned adopters ("fast followers") from	31.5% current, 24.7% planned within 1 year

Source	Data Year(s)	Sample	What "Adoption" Means	Deployment Stage Captured	Reported Figure
				no near-term plan ("delayed adopters")	
AJMC (Yang and Graetz) (ajmc.com)	June 2025	6,561 US hospitals; 2,784 Epic-EHR hospitals	A specific ambient AI documentation tool licensed among Epic-EHR hospitals	Tool licensed/enabled at the hospital level, not a clinician usage rate	62.6% of Epic hospitals
JAMIA (Poon et al.) (essentialhospitals.org)	2025	43 US health systems	4-stage ladder: no activity, developing/piloting, deployed in limited areas, deployed fully, per use case	Explicitly separates piloting from full deployment	100% at least piloting ambient documentation; 14% fully deployed
KLAS Research (klasresearch.com) (engage.klasresearch.com)	Interviews Jan 2025, tracked through Dec 2025	256 healthcare organizations	Self-reported organizational "use" of AI in any function	Organizational claim, not a use-case deployment stage	Rose from under half to over two-thirds across 2025
Deloitte 2026 US Health Care Outlook (deloitte.com)	Fielded 2025, published Dec 2025	120 C-suite health care executives	Self-assessed maturity: experimenting, not yet adopted, or operating at scale	Explicitly separates experimentation from scaled use	49% still experimenting; about one-third operating at scale
HIMSS/Medscape AI Adoption by Health Systems Report (himss.org)	2024	800+ mixed respondents (physicians, nurses, IT staff, executives)	Broad individual self-report of "leveraging AI" in the organization	No stage distinction; individual perception, not a hospital census	86%

Table 1 shows why "how many hospitals use AI" cannot have one honest number. The **18.7%** Baten figure and the **71%** ONC figure are two years and two definitions apart, not a contradiction; the **65%** and **66%** Health Affairs/ONC figures are the same year and nearly the same definition, yet still differ by a point due to methodology; and the KLAS, Deloitte, and HIMSS figures describe organizational self-perception rather than a verified hospital census at all. A reader should always ask which row of this table a headline number belongs to before comparing it to another.

Adoption also varies sharply by hospital type. *Table 2* summarizes the 2024 ONC/ASTP breakdown by hospital characteristic, the most granular federally sourced cut available.

Table 2. *Predictive AI Adoption by Hospital Characteristic, 2024 (ONC/ASTP Data Brief 80)*

Hospital Characteristic	2024 Adoption
Small (<100 beds)	59%
Medium (100 to 399 beds)	80%

Hospital Characteristic	2024 Adoption
Large (400+ beds)	96%
Rural	56%
Urban	81%
Independent hospital	37%
System-affiliated hospital	86%
Critical access hospital (CAH)	50%
Non-CAH	80%

(figures from healthit.gov and aha.org)

The gap between large and small hospitals (96% versus 59%) and between system-affiliated and independent hospitals (86% versus 37%) is larger than the gap between any two years of national data. The **AHA Center for Health Innovation**'s own November 2025 commentary on this same dataset also notes that among hospitals using predictive AI in 2024, only **82%** locally evaluated the tools for accuracy, **74%** for bias, and **79%** conducted post-implementation monitoring, meaning governance lags adoption even among the hospitals furthest along (aha.org). Critical access hospitals, defined by the AHA as small facilities generally located at least 35 miles from another hospital, reported just **50%** adoption in 2024 against **80%** for non-critical-access hospitals (aha.org). The fastest-growing use cases between 2023 and 2024 were administrative rather than clinical: billing automation rose **25 percentage points** and appointment scheduling rose **16 points**, both faster than clinical risk-prediction use cases, which were already common and grew only modestly (aha.org).

Generative AI specifically, as opposed to older predictive models, is measured separately. A *JAMA Network Open* study by Everson, Nong, and Richwine, published December 12, 2025 and also using the 2024 AHA IT Supplement (N=2,174 hospitals), found **31.5%** of hospitals were "early adopters," meaning they currently used a large language model integrated into the EHR, **24.7%** were "fast followers" planning adoption within a year, and **43.7%** were "delayed adopters" with no near-term plan (pmc.ncbi.nlm.nih.gov). Combining current and near-term adopters, the study projected that roughly half of hospitals would use generative AI by the end of 2025. Early-adopter status varied enormously by hospital type: **53.9%** of major teaching hospitals were early adopters versus **25.9%** of nonteaching hospitals, and hospitals on the Epic EHR platform were early adopters at **48.8%** versus **15.6%** for hospitals on Oracle Health (pmc.ncbi.nlm.nih.gov). A related study using a different EHR-vendor lens, published in *AJMC* in January 2026 using June 2025 data, found that **62.6%** of the **2,784** Epic-using hospitals (of 6,561 US hospitals studied) had adopted an ambient AI documentation tool, and that adoption skewed heavily toward nonprofit and metropolitan hospitals over government-owned hospitals (ajmc.com) (ajmc.com).

Physician-Level Adoption: The AMA Survey Series, 2023 to 2026

The **American Medical Association (AMA)** runs the most closely watched physician-level AI survey, now in its third wave. The 2023 wave (n=1,081) found **38%** of physicians using AI; the 2024 wave (n=1,183), fielded through November 2024, found **66%**, a **78%** relative increase (ama-assn.org); its press materials describe the comparison

window as running "from August 2023 to November 2024" (ama-assn.org). The most recent wave, fielded January 15 to February 2, 2026 among nearly **1,700** physicians, found **81%** reporting any awareness or use of AI in their practice, more than double the 2023 rate (ama-assn.org) (ama-assn.org).

AMA's methodology defines explicit AI use as **selecting one or more named AI use cases** from a structured list, a lower bar than routine daily use but a higher one than mere awareness (ama-assn.org). A comparability caveat in the 2026 report matters for anyone stacking these three numbers on a chart: the 2023 and 2024 waves counted only fully completed surveys, while the 2026 wave counted qualified partial responses as well, a methodology change the report itself discloses (ama-assn.org). The survey also shows depth of use increasing alongside breadth: the average number of distinct AI use cases per physician rose from **1.1 in 2023** to **2.3 in 2026**, and physician uncertainty about which AI tools their own practice offered fell from **18%** in 2024 to **9%** in 2026, both proxies for a shift from experimentation to routine familiarity (ama-assn.org) (ama-assn.org).

The specific use cases physicians report also shifted. In 2024, the leading applications were documentation of billing codes, medical charts, or visit notes; creation of discharge instructions and care plans; translation services; and assistive diagnosis (ama-assn.org). By 2026, the single most-used application was summarizing medical research or standards of care, used by nearly **40%** of physicians, a **26-point** jump from 2024, while creation of discharge instructions rose to **30%** of physicians from roughly 20% (ama-assn.org). Notably, none of the three AMA waves names "ambient documentation" or "AI scribe" as a standalone category; that use case is folded into broader documentation categories, which means AMA's figures cannot be compared directly to the ambient-scribe-specific hospital studies discussed in the next section.

Physician sentiment moved alongside use. In 2026, more than three-quarters of physicians said AI provides an advantage in their ability to care for patients, up from **65%** in 2023 (ama-assn.org), though a substantial **40%** minority in 2026 still reported being equally excited and concerned about the technology, indicating that rising use has not eliminated ambivalence (ama-assn.org). The 2024 wave found physicians' top requirements for expanding AI use were a designated feedback channel (88%), data privacy assurances (87%), and EHR integration (84%), a ranking that helps explain why hospital-level infrastructure, not just tool availability, gates physician adoption (ama-assn.org).

From Pilot to Production: Deployment-Stage Evidence

Ambient AI documentation, the technology that transcribes and drafts clinical notes from a patient encounter, is the clearest illustration of the gap between adoption headlines and routine clinical use, because several independent sources measured the same category with materially different yardsticks. *Table 3* compares them directly.

Table 3. *Ambient AI Documentation: Pilot, License, and Active-Use Evidence, 2025*

Source	Sample	What Was Measured	Result
JAMIA (Poon et al.) (essentialhospitals.org)	43 US health systems	Deployment stage per use case (pilot vs. limited vs. full)	100% at least piloting; 14% fully deployed
Peterson Health Technology Institute (PHTI) (phti.org)	8 large health systems	Share of eligible clinicians actively using the tool, once offered	Typically 20 to 50% active use; one system reached 75 to 80%

Source	Sample	What Was Measured	Result
PHTI, single-system detail (phti.org)	1 health system (Providence)	Active use versus systemwide offer	Offered to all physicians; about 15% used it on an ongoing basis
Bain & Company/Bessemer/AWS Adoption Index (bain.com)	400+ senior healthcare leaders	Systemwide deployment vs. piloting/implementing	30% deploying systemwide; 60% piloting or implementing
Bain & Company/KLAS (bain.com)	228 healthcare executives	Full rollout vs. partial rollout	Roughly 1 in 5 providers at full rollout, another 2 in 5 partial
AJMC (Epic hospitals) (ajmc.com)	2,784 Epic-EHR hospitals	Hospital-level licensing/enablement	62.6% licensed the tool (not a usage rate)

The pattern across every row is consistent: licensing or piloting an ambient AI tool is close to universal among large systems, while sustained, hospital-wide active use by the majority of eligible clinicians is not. PHTI's own review of eight large health systems found that even where a system offered ambient AI scribes to all clinicians, ongoing active use typically settled at **20% to 50%** of eligible users, with one organization reaching **75% to 80%** (phti.org). At one named system, Providence, only about **15%** of clinicians used the tool on an ongoing basis despite it being offered to all physicians (phti.org). PHTI's broader conclusion is that financial return on investment for ambient scribes "remains unclear" even at systems with meaningful active use, a distinct and later-stage question from whether the tool is deployed at all (phti.org).

The Poon *JAMIA* survey's top-cited barrier to moving from pilot to full deployment was tool immaturity, cited by **77%** of the 43 health systems, ahead of financial concerns at **47%** (essentialhospitals.org), suggesting that the pilot-to-production gap is driven less by hospital budgets than by clinicians' and IT leaders' confidence in the tools themselves. A November 2025 *JAMA Health Forum* analysis using the Census Bureau's Business Trends and Outlook Survey took a still narrower view, finding that only **8.3%** of health-care businesses reported currently using AI "in producing goods or services" in 2025, a figure the authors explicitly contrasted with the Poon study's finding that full implementation ranged from **0% to 48%** depending on use case, underlining how differently "current use" can be operationalized across instruments (pmc.ncbi.nlm.nih.gov).

A preprint (not yet peer-reviewed) posted to medRxiv in mid-2025, analyzing AI implementation signals across **3,560** US hospitals, found that hospital-level AI implementation is significantly geographically clustered into hotspots and coldspots rather than evenly distributed, a pattern consistent with the system-affiliation and teaching-status effects documented above, though the finding should be treated as preliminary pending peer review (medrxiv.org).

Analysis of Key Segments: Health Systems, Payers, and Market Research

Beyond the government-linked hospital surveys, several research firms track AI adoption across health systems and payers using their own executive and organizational panels, generally reporting higher and faster-rising figures than the AHA-based data, in part because their samples skew toward larger, more digitally mature organizations.

KLAS Research interviewed **256** healthcare organizations in January 2025 about current AI adoption and investment plans (klasresearch.com); by its December 2025 update, KLAS reported the share of organizations describing themselves as AI users had climbed from under half to more than two-thirds across the year (engage.klasresearch.com). Among concrete use cases, KLAS identified **ambient speech and documentation** as the single most commonly cited clinical AI application, while explicit mentions of **agentic AI**, AI systems that autonomously plan and execute multi-step tasks, remained a small fraction of reported use cases relative to the term's prominence in industry marketing (engage.klasresearch.com) (engage.klasresearch.com).

A joint **Bain & Company/KLAS** survey of **228** provider and payer executives in October 2025 found **70%** of providers now had an AI strategy in place or in development, up from **60%** a year earlier (bain.com), and identified ambient documentation as the furthest-along use case, with roughly **1 in 5** providers at full rollout and another **2 in 5** at partial rollout (bain.com). A separate April 2025 Bain survey with Bessemer Venture Partners and Amazon Web Services, covering more than 400 senior healthcare leaders, found providers were the furthest-along segment (ahead of payers and life-sciences companies) in moving AI proofs-of-concept into production, at **35%** (bain.com).

Deloitte's 2026 US Health Care Outlook, based on **120** C-suite health care executives surveyed in late 2025, found **49%** of organizations still experimenting with generative or agentic AI and roughly a third operating it at scale, while **83%** of health system executives (versus **97%** of health plan executives) expected generative or agentic AI to add value to clinical functions in 2026 (deloitte.com) (deloitte.com). A separate Deloitte report published February 2026, surveying **100** health-system and health-plan technology executives fielded in September 2025, found **61%** already building, implementing, or budgeting for agentic AI initiatives, and **98%** expected at least **10%** cost savings from agentic AI within two to three years (deloitte.com) (deloitte.com). These are expectations, not measured outcomes, and should be read as such.

HIMSS and Medscape's 2024 AI Adoption by Health Systems Report, based on more than 800 mixed respondents including physicians, nurses, IT professionals, and executives, found **86%** already using AI in their organizations, though **72%** cited data privacy as a significant risk (himss.org). This figure sits well above the ONC and Health Affairs hospital-level numbers, consistent with a mixed, self-selected respondent panel rather than a probability sample of hospitals. A McKinsey survey, restated by *Becker's Hospital Review* because McKinsey's own report page could not be independently verified this session, reportedly found generative AI implementation among US healthcare organizations rising from **25%** in late 2023 to **47%** in 2024 and **50%** by the end of 2025 (beckershospitalreview.com); this figure should be treated as a secondary restatement rather than a verified primary source until McKinsey's own page can be confirmed.

Data Analysis and Evidence

Reconciling these sources requires separating four questions that are routinely blended into one: how many organizations have acquired an AI tool, how many have piloted it, how many use it routinely, and how many have proven it works. Cutting the evidence this way produces a coherent picture rather than a contradictory one.

On acquisition and licensing, the supply side of the market is large and growing quickly. The **Food and Drug Administration (FDA)**'s AI/ML-Enabled Medical Device List, content current as of **September 4, 2026**, contains **1,615** authorized devices (fda.gov). FDA itself cautions that the list is **"not a comprehensive resource of AI-enabled medical devices"** and exists to identify devices **"authorized for marketing,"** not devices in active clinical use (fda.gov) (fda.gov). FDA also discloses that devices are added to the list primarily through keyword-matching

against authorization documents rather than verified functional review, a further reason the device count is a supply-side indicator, not an adoption metric (fda.gov). A device being FDA-authorized says nothing about how many, if any, US hospitals have purchased or activated it, exactly the gap this report's deployment-stage tables are built to surface.

On measured active use, the range narrows considerably once the definition is held constant. For predictive AI narrowly defined as ML-based risk scoring integrated into the EHR, ONC's 2024 figure of **71%** is the most defensible hospital-level estimate available, because it is drawn from a probability sample with a stated response rate (healthit.gov). For generative AI specifically, the comparable, methodologically similar figure is the *JAMA Network Open* study's **31.5%** current-use rate for 2024, roughly **40 percentage points** lower than the predictive-AI figure for the same underlying survey population, illustrating that "AI adoption" collapsed into one number obscures a wide gap between older predictive tools and newer generative ones (pmc.ncbi.nlm.nih.gov). Independent-panel surveys from KLAS, Bain, and HIMSS report notably higher figures (two-thirds or more), a divergence best explained by sampling: these panels weight toward large, digitally mature health systems that are not representative of the roughly 6,000-hospital national population the AHA and ONC surveys attempt to cover.

On demonstrated outcomes, the evidence is the thinnest of the four categories by design, since outcome studies require a control or comparison group that most hospitals cannot construct for a single AI tool. The clearest outcome-adjacent finding in this report's sources is PHTI's observation that financial ROI for ambient AI scribes "remains unclear" even at systems with substantial active clinician use, meaning the active-use figures in Table 3 should not be read as evidence of proven value (phti.org). Readers seeking outcome-specific evidence for a particular clinical use case should treat this report's adoption figures as a necessary but not sufficient precursor to that separate body of literature.

Implications and Future Directions

Three implications follow directly from the reconciliation above. First, any single adoption percentage cited without its source, year, sample, and definition should be treated as incomplete, not because the underlying research is unreliable but because the category "AI adoption" spans at least four operationally distinct phenomena that this report has separated into acquisition, piloting, active use, and demonstrated outcome. Second, the gap between piloting and full deployment, most starkly illustrated by the Poon *JAMIA* survey's **100%**-piloting-versus-**14%**-fully-deployed finding for ambient documentation, suggests that future surveys and purchasing decisions alike would benefit from asking about deployment stage explicitly rather than a single yes/no adoption question (essentialhospitals.org). Third, the resource gradient by hospital size, teaching status, and system affiliation documented in Table 2 shows no sign of narrowing in the newer generative AI data, where the teaching-hospital and Epic-versus-Oracle gaps are, if anything, wider than the older predictive-AI gaps, a pattern with equity implications for rural and independent hospitals that federal and state health IT policy will likely continue to track.

For hospital and health system leaders, the practical takeaway is that a vendor's or a peer institution's adoption claim should be interrogated for which of Table 1's rows it corresponds to before it informs a purchasing or resourcing decision. IntuitionLabs' own adoption-measurement work, while addressed to life-sciences and healthcare AI programs generally rather than to hospital deployments specifically, frames this as writing down a decision, a workflow hypothesis, and success or stop thresholds before results are known, precisely to avoid the trap of redefining "adoption" after a pilot has already produced a convenient-sounding number (intuitionlabs.ai). As more health systems reach the JAMIA survey's "deployed fully" stage for use cases like ambient documentation,

later editions of the surveys tracked in this report should show the pilot-to-production gap narrowing; whether it does is an empirical question this report's method makes straightforward to check in future years, using the same source-by-source approach rather than a single blended percentage.

Frequently Asked Questions (FAQs)

How many hospitals use AI? It depends entirely on the definition. For predictive AI (ML-based risk scoring) integrated into the EHR, ONC's 2024 data brief found **71%** of hospitals (healthit.gov). For generative AI specifically, a *JAMA Network Open* study of the same year found **31.5%** currently using it (pmc.ncbi.nlm.nih.gov). For any AI across five operational categories in 2022, a Health Affairs Scholar study found **18.7%** (academic.oup.com). See Table 1 above for the full comparison.

What is the AMA physician AI adoption survey, and what does it measure? It is a recurring national survey of US physicians, now in three waves (2023, 2024, 2026), that asks about named AI use cases in practice; it measures individual physicians, not hospitals, and reported **38%**, **66%**, and **81%** reporting any awareness or use across the three waves respectively (ama-assn.org).

What is the AHA AI adoption survey? There is no single AHA-branded AI survey; the relevant instrument is the AHA Information Technology Supplement, fielded alongside the AHA Annual Survey, whose underlying hospital-level data is analyzed both by ONC/ASTP directly and by independent academic teams, sometimes producing slightly different published percentages for the same year (healthit.gov) (healthaffairs.org).

What is the difference between clinical AI deployment and a pilot? A pilot is a limited trial, often with one department or a small group of volunteer users, without an enterprise rollout decision. Deployment or active use means the tool is part of normal workflow for a defined population of eligible users. The Poon *JAMIA* survey found **100%** of 43 health systems had at least piloted ambient documentation, but only **14%** had it fully deployed (essentialhospitals.org).

How many hospitals use generative AI specifically? Using the 2024 AHA IT Supplement, a *JAMA Network Open* study found **31.5%** of hospitals were current ("early adopter") users of generative AI integrated into the EHR, with another **24.7%** planning adoption within a year (pmc.ncbi.nlm.nih.gov).

Do more FDA-authorized AI devices mean more hospitals are using AI? Not directly. The FDA's AI/ML-Enabled Medical Device List, at **1,615** devices as of September 4, 2026, tracks marketing authorization, and FDA explicitly states it is not a measure of comprehensiveness or clinical deployment (fda.gov) (fda.gov).

Why do AI adoption survey methodologies vary so much? Surveys differ in sample (hospitals versus physicians versus executives), in definition (any AI versus predictive AI versus generative AI versus a single named use case), in deployment stage measured (licensed versus piloted versus routinely used versus outcome-proven), and in response rate and weighting, even when two studies use the same underlying dataset, as the one-point ONC/Health Affairs discrepancy for 2023 predictive AI adoption illustrates (healthaffairs.org).

Conclusion

No single number answers "how many US hospitals use AI," and this report has tried to show why that is a feature of the evidence rather than a flaw in it. Predictive AI integrated into the EHR reached **71%** of hospitals by 2024 according to the federal data brief most analysts cite; generative AI specifically reached roughly a third on a comparable measure the same year; and any AI across a broader set of operational categories stood at under a fifth of hospitals just two years earlier, using a stricter binary standard. The share of physicians reporting any awareness or use of AI, a related but distinct measure, reached **81%** by early 2026 in the AMA's latest survey. Independent-panel surveys from KLAS, Bain, Deloitte, and HIMSS report higher figures still, generally because their samples weight toward the largest and most digitally mature organizations rather than the full national hospital population that the AHA and ONC instruments attempt to represent.

The more durable finding, and the one least visible in headline percentages, is the gap between acquiring or piloting an AI tool and using it routinely at scale. Ambient AI documentation is licensed or piloted at nearly every large health system studied, yet full deployment and sustained active use by a majority of eligible clinicians remain the exception rather than the rule, and demonstrated financial or clinical outcomes remain thinner still. Readers, researchers, and hospital leaders who cite an adoption statistic from this space should carry forward its sample, its year, and its deployment-stage definition alongside the number itself, using a source comparison of the kind built in Table 1 rather than a single figure lifted out of context.
