


Suki vs Nuance DAX vs Abridge vs Freed: AI Scribes Compared

Published July 19, 2026 36 min read



Suki vs Nuance DAX vs Abridge vs Freed: AI Scribes Compared

Executive Summary

This comparison examines four prominent ambient artificial intelligence (AI) medical scribe offerings selected to represent different buyer profiles: **Suki**, **Nuance DAX** (now **Microsoft Dragon Copilot**), **Abridge**, and **Freed**. Each occupies a distinct position. Abridge won the 2025 and 2026 **Best in KLAS** awards in the Ambient Speech category (Source: [abridge.com](https://www.abridge.com)), and it has deployed across more than **300 health systems** (Source: [abridge.com](https://www.abridge.com)), including a landmark rollout to Kaiser Permanente's 40 hospitals and over 600 medical offices in eight states, described by Kaiser as its **largest generative AI implementation to date** (Source: [about.kaiserpermanente.org](https://www.aboutkaiserpermanente.org)). Abridge has raised approximately **\$800 million** to date and was valued at **\$5.3 billion** after a **\$300 million** Series E in June 2025 (Source: [fiercehealthcare.com](https://www.fiercehealthcare.com)) (Source: [techcrunch.com](https://www.techcrunch.com)).

Nuance DAX traces to Microsoft's **\$19.7 billion** all-cash acquisition of Nuance Communications, announced in April 2021 and completed in March 2022 (Source: [news.microsoft.com](https://www.news.microsoft.com)). The original Dragon Ambient eXperience (DAX) launched in 2020 (Source: [prnewswire.com](https://www.prnewswire.com)), and in March 2025 Microsoft merged Dragon Medical One and DAX Copilot into a single product called **Dragon Copilot**, generally available in the United States and Canada since May 2025 (Source: [news.microsoft.com](https://www.news.microsoft.com)). Suki, founded in 2017 by former Google and Flipkart executive Punit Soni, has raised **\$168 million** to date and serves more than 350 health systems and clinics (Source: [suki.ai](https://www.suki.ai)) (Source: [suki.ai](https://www.suki.ai)). Freed, the newest and most price-accessible of the four, is aimed squarely at independent and small clinics rather than enterprise health systems, publishes transparent tiered pricing starting at **\$39 per month**, and reports **26,000-plus** clinicians and over **32.5 million** patient visits transcribed in 2025 alone (Source: [getfreed.ai](https://www.getfreed.ai)) (Source: [getfreed.ai](https://www.getfreed.ai)).

None of the enterprise-oriented vendors, Suki, Nuance DAX/Dragon Copilot, or Abridge, publish list pricing on their public websites; all three route prospective buyers to enterprise sales, and independent analyst estimates place per-provider costs anywhere from roughly **\$208 to \$1,000-plus per month** depending on tier and contract volume, figures that should be read as third-party estimates rather than vendor-confirmed prices (Source: [veroscribe.com](https://www.veroscribe.com)). Freed is the outlier with fully published tiers: **Starter** at \$39 per month (billed annually), **Core** at \$79 per month, and **Premier** at \$104

to \$119 per month, plus custom **Groups** pricing for clinics (Source: getfreed.ai). Peer-reviewed evidence is more mixed than vendor marketing suggests: a narrative review published via PubMed Central reported time savings of roughly **1 to 2.1 minutes per note** in included studies, while one controlled simulation of 44 encounters using two unnamed products found **70% of notes contained at least one error**; that result does not establish the accuracy of the four products compared here but reinforces the need for clinician review (Source: pmc.ncbi.nlm.nih.gov). The market overall is growing fast: one industry analysis sizes the global AI-in-medical-scribing segment at **\$1.67 billion in 2026**, expanding to **\$8.93 billion by 2035** at a 20.48% compound annual growth rate (Source: towardshealthcare.com). Choosing among the four ultimately depends less on which tool is "best" in the abstract and more on organizational scale, EHR (electronic health record) platform, and whether documentation is the only workflow problem being solved or the first of several an organization intends to address with AI.

Introduction and Background

Clinical documentation burden is widely cited as a leading driver of physician burnout, and ambient AI scribes, software that passively listens to a patient encounter and drafts a structured clinical note, have become the fastest-growing category of applied AI in United States healthcare. Investors poured close to **\$1 billion** into ambient AI scribe companies in 2025 alone, according to a STAT News analysis, with Abridge (**\$550 million** announced across two 2025 rounds) and Ambience Healthcare (**\$243 million** Series C at a **\$1.25 billion** valuation) accounting for the bulk of it (Source: statnews.com) (Source: statnews.com).

Four vendors dominate the "which AI scribe should we buy" conversation for reasons that map to four different buyer archetypes. **Suki**, founded in Redwood City, California in 2017, positions itself as a voice-first clinical assistant that goes beyond passive listening: clinicians can issue spoken commands to Suki for chart retrieval, order staging, and coding (Source: soapnotebuddy.com). **Nuance DAX**, now folded into Microsoft's **Dragon Copilot**, is the incumbent with the deepest history in speech recognition for healthcare, tracing back through Dragon Medical and Microsoft's 2021 to 2022 acquisition of Nuance Communications for **\$19.7 billion** (Source: news.microsoft.com). **Abridge**, founded in Pittsburgh in 2018 by practicing cardiologist Dr. Shiv Rao, has become the most-cited independent leader by scale and by KLAS Research's provider-reported satisfaction rankings (Source: abridge.com). **Freed**, founded in 2023 and backed by a **\$30 million** Series A led by Sequoia Capital in March 2025, has built its brand specifically around independent and small-group clinicians who found enterprise scribe pricing and sales processes inaccessible (Source: getfreed.ai). The decision facing a health system chief medical information officer differs sharply from the one facing a two-physician family practice, and this report treats those as separate but related questions throughout: which vendor wins on enterprise EHR integration, which wins on price transparency, which has the strongest independent quality evidence, and how a buyer should weigh vendor-reported metrics against the more cautious findings emerging in peer-reviewed literature. A related but distinct question that recurs across search queries—HIPAA (Health Insurance Portability and Accountability Act) compliance and business associate agreement (BAA) requirements—is addressed directly in the sections that follow, because a vendor is generally a HIPAA business associate when it creates, receives, maintains, or transmits PHI on behalf of a covered entity; the determination depends on the relationship, the data involved, and the service performed ([HHS](https://hhs.gov)).

Suki

Capabilities

Suki's flagship product, **Suki Assistant**, ambiently listens to patient encounters and generates specialty-specific structured notes, but the company differentiates itself by supporting active voice commands throughout the clinical workflow, not just ambient capture (Source: suki.ai). The platform spans the full visit lifecycle: pre-visit patient summaries and chart Q&A, in-visit ambient documentation and "order staging" (drafting orders directly from conversation for clinician approval), and post-visit multilingual patient instructions generated in 80 languages at a fifth-grade reading level (Source: suki.ai). Suki also generates ICD-10 (International Classification of Diseases, 10th Revision), HCC (Hierarchical Condition Category), CPT (Current Procedural Terminology), and evaluation-and-management (E/M) codes as part of its coding module (Source: suki.ai). The company reports support for **100-plus specialties** and deployment across ambulatory, inpatient, emergency department, urgent care, telehealth, and home health settings, with clients on iOS, Android, web, and desktop (Source: suki.ai).

Suki's Epic integration, marketed as **Suki INSIDE**, embeds ambient note generation directly in Haiku and Hyperspace with sign-off inside Hyperspace, while a fuller **Suki for Clinicians** tier adds problem-based charting and voice-enabled editing inside the Epic environment (Source: suki.ai). One physician customer described the resilience of this integration during an outage, noting that "our internet provider was down so we lost complete access to our EHR, but I had access to Suki and could continue to do my notes in the ambient fashion and then push them over after" (Source: suki.ai). The company also integrates with Oracle Health (formerly Cerner), athenahealth, and MEDITECH, and describes its EHR footprint as the broadest in the category, spanning Epic, Oracle-owned Cerner, Athena, and MEDITECH (Source: reuters.com).

Adoption

Suki's own figures put its footprint at **350 health systems and clinics** as of early 2025, up from roughly 300 in October 2024, with a stated **4x expansion in 12 months** during 2024 (Source: suki.ai). Named enterprise customers include MedStar Health, which evaluated multiple documentation vendors before identifying Suki's software-first approach as the best fit, extending access to thousands of MedStar clinicians across specialties including primary care, cardiology, gastroenterology, and urgent care (Source: beckershospitalreview.com). Suki also supplies its underlying AI engine to third parties: Zoom Communications took a strategic investment stake in Suki and integrated Suki's clinical-note engine into its Workplace for Clinicians offering, an arrangement the companies announced in October 2024 (Source: suki.ai). Total funding stands at **\$168 million**, including a **\$70 million** Series D led by Hedosophia with Venrock and March Capital participating, which one source familiar with the round told Reuters valued the company at approximately **\$500 million** (Source: reuters.com). Reuters also reported that Suki's rise tracks the broader post-ChatGPT boom in generative AI adoption across health systems, noting that "when the AI trend kicked in, every health system wanted to have an AI strategy," in Soni's words (Source: reuters.com).

Strengths and Limitations

Suki's own marketing claims clinicians complete notes **72% faster on average** and that the platform delivers a **9X return on investment (ROI) in year one**, figures that come directly from company press materials and have not been independently replicated in peer-reviewed literature as of this writing (Source: suki.ai). Suki's CEO Punit Soni has framed the company's differentiator as platform breadth rather than pure documentation depth, telling Fierce Healthcare that Suki's SDK (software development kit) and API suite let health systems and technology partners embed ambient capabilities without being tied to a single EHR (Source: fiercehealthcare.com). Enterprise sales, however, is also Suki's principal accessibility limitation: pricing is not published, requiring a demo or pilot engagement before quotes are available, which effectively excludes solo practitioners and very small groups who want to evaluate cost without a sales call (Source: soapnotebuddy.com).

Nuance DAX (Microsoft Dragon Copilot)

Capabilities

Nuance DAX began as **Dragon Ambient eXperience**, launched in 2020 as an ambient clinical intelligence (ACI) solution that captures multi-party conversations and converts them into structured documentation (Source: prnewswire.com). In March 2023, Nuance layered OpenAI's GPT-4 into a lighter-weight variant called **DAX Express**, extending the technology to more than **550,000 Dragon Medical users** (Source: prnewswire.com). Then-Nuance CEO Mark Benjamin framed the milestone as a joint achievement with Microsoft, stating the companies had "taken the power and advanced reasoning capabilities of GPT-4 and integrated it into our proven outcomes-focused AI technologies in a tested and responsible way" (Source: prnewswire.com). In March 2025, Microsoft consolidated the product line entirely: **Dragon Copilot** merged Dragon Medical One's front-end speech recognition with DAX Copilot's ambient listening and generative summarization into what Microsoft calls the healthcare industry's first unified voice AI assistant, generally available in the United States and Canada from May 2025 and later in the United Kingdom, Germany, France, and the Netherlands (Source: news.microsoft.com). Buyers and clinicians frequently still search under the legacy names Nuance DAX or DAX Copilot, since Microsoft's own documentation confirms these refer to the same underlying product family (Source: microsoft.com).

Dragon Copilot is licensed in several forms according to Microsoft's own documentation: a per-user **Physician Per User** license bundling all features without requiring a separate Azure subscription, a **Physician Flex** license aimed at organizations migrating from Dragon Medical One that combines a base dictation subscription with pay-as-you-go ambient and generative AI consumption, a standalone **Physician PAYG (pay-as-you-go) consumption unit**, and a **Nurse Per User** license for flowsheet-based nursing workflows (Source: learn.microsoft.com). The Physician Flex license is explicitly recommended for "health systems transitioning from Nuance Dragon Medical One, that need lower upfront cost with usage-based billing," and for organizations with high variability in ambient and AI usage across clinicians (Source: learn.microsoft.com). The product is HITRUST-CSF (Health Information Trust Alliance Common Security Framework) certified and built on Microsoft Azure, combining what Microsoft describes as Nuance's proven conversational and ambient AI with newer generative AI capabilities (Source: azuremarketplace.microsoft.com).

Adoption

Dragon Copilot's adoption is inseparable from Microsoft's broader healthcare AI push following its **\$19.7 billion**, all-cash acquisition of Nuance Communications, a deal valued at **\$56.00 per share** and announced in April 2021, completed in March 2022 (Source: news.microsoft.com). Microsoft and Epic have partnered closely since 2023 on generative AI copilots, and the two companies subsequently worked with Advocate Health, Duke Health, Intermountain Health, and Stanford Health Care to extend ambient technology into nursing documentation (Source: fiercehealthcare.com). Independent, peer-matched cohort research at Intermountain Health, spanning 12 specialties and comparing 99 DAX users against 76 controls, found

documentation time per note fell by roughly 0.76 minutes and clinician engagement improved, though after-hours EHR use ticked up slightly and patient experience scores did not change (Source: [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov/38111111/)). At Stanford Health, a prospective pre-post quality-improvement study of 48 physicians found DAX Copilot use associated with a median 20 minutes saved per half-day and reduced NASA-TLX (Task Load Index) burnout scores (Source: [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov/38111111/)). A separate randomized pilot at Samaritan Health across primary care, urgent care, and multispecialty settings found documentation time fell by 1.4 minutes per visit, with DAX Copilot generating roughly half of note characters among high users (Source: [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov/38111111/)).

Strengths and Limitations

Dragon Copilot's central advantage is Epic depth and Microsoft's enterprise sales and compliance infrastructure, making it a default consideration for the roughly 42% of the United States hospital market already standardized on Epic (Source: [fiercehealthcare.com](https://www.fiercehealthcare.com/news/epic-dragon-copilot-ai-scribe-launch)). The principal limitation, echoed across independent reviewers, is that Microsoft does not publish subscription prices; consumption-based charges for the pay-as-you-go tier are published at a per-unit rate, but base subscription costs require direct negotiation with a Microsoft representative or certified partner (Source: [sameexpert.com](https://www.sameexpert.com/)). Third-party analyst estimates place per-provider costs in the range of **\$369 to \$830-plus per month**, though these figures are not vendor-confirmed and should be treated as market observations rather than list price (Source: [veroscribe.com](https://www.veroscribe.com/)). Implementation timelines and multi-year contract structures also make Dragon Copilot a harder fit for small independent practices than for large systems already embedded in the Microsoft and Epic ecosystems.

Abridge

Capabilities

Abridge converts patient-clinician conversations into structured clinical notes in real time, with what the company calls a **Contextual Reasoning Engine** and a proprietary evaluation framework it terms "Linked Evidence," designed to map each AI-generated statement back to the source conversation for clinician verification (Source: [abridge.com](https://www.abridge.com/)). The platform supports more than **50 medical specialties** and **28-plus languages** (Source: [abridge.com](https://www.abridge.com/)), and it has expanded beyond documentation into revenue-cycle and coding automation, announcing in mid-2025 a collaboration with Highmark Health to co-develop an AI-powered prior-authorization tool at the point of care (Source: [fiercehealthcare.com](https://www.fiercehealthcare.com/news/highmark-abridge-prior-auth)). The company's "Abridge Inside" integration embeds the product directly within Epic's Haiku and Hyperdrive interfaces, developed through Abridge's participation in Epic's Workshop co-development program, which "unlocks the collaborative power of co-development between Epic and Abridge to more quickly tackle health systems' most pressing issues" (Source: [abridge.com](https://www.abridge.com/)). Epic's vice president of research and development, Garrett Adams, said the Workshop model was "built to accelerate cutting-edge innovation through co-development," adding that Abridge exemplifies Epic's preference for "responsible AI that is auditable and transparent" (Source: [abridge.com](https://www.abridge.com/)). In 2026 the company incorporated OpenAI's GPT-5.5 model into portions of its pipeline, reporting a **25% relative increase** in its ability to provide clinical-quality specificity during early access testing (Source: [abridge.com](https://www.abridge.com/)).

Adoption

Abridge reports that it is **trusted by 300-plus health systems** and now powers **more than 100 million conversations per year** (Source: [abridge.com](https://www.abridge.com/)), up from the 150-plus health systems the company reported at the time of its June 2025 Series E (Source: [techcrunch.com](https://www.techcrunch.com/2025/06/11/abridge-series-e/)). Kaiser Permanente's deployment across 40 hospitals and more than 600 medical offices in eight states and Washington, D.C., which Kaiser's own senior vice president of Care Delivery Technology Services described as "the largest implementation to date of the safe and effective use of ambient listening technology in the United States," remains the single largest named rollout in the category (Source: [aboutkaiserpermanente.org](https://www.aboutkaiserpermanente.org/newsroom/kaiser-permanente-adopts-abridge)). Johns Hopkins Medicine finalized an enterprise agreement covering **6,700 clinicians**, six hospitals, and 40 patient-care centers in December 2024 (Source: [abridge.com](https://www.abridge.com/)), and Geisinger grew clinician adoption of Abridge by **187% in just seven months** between April and November 2025, surpassing 1,000 active users by November 2025 after going live with its first users in late September 2024 (Source: [abridge.com](https://www.abridge.com/)). Geisinger's Chief Medical Informatics Officer, Dr. Ben Hohmuth, described the rollout philosophy as deliberately organic: the health system "went live with our first users in late September 2024, and surpassed 1,000 users by November 2025," with every physician, advanced practice provider, resident, and fellow at Geisinger now able to access Abridge if they choose (Source: [abridge.com](https://www.abridge.com/)). The company achieved KLAS Research's No. 1 **Best in KLAS** ranking for Ambient Speech in both 2025 and 2026, a distinction KLAS bases on independent interviews with customer organizations rather than vendor self-reporting (Source: [abridge.com](https://www.abridge.com/)).

Strengths and Limitations

Abridge's own case-study library, published on its corporate site, reports a broad range of measured customer outcomes, including that Christus Health reported a **78% reduction in cognitive load**, Sharp HealthCare reported an **83% reduction in note-writing effort**, and Reid Health reported an **86% reduction in documentation effort** (Source: abridge.com). These figures come directly from Abridge's own customer-story marketing and, while directionally consistent with independent findings, have not all been replicated in peer-reviewed studies; a Sutter Health pre-post quality-improvement study captured in the PMC narrative review found a more modest but still positive **0.9-minute-per-patient** reduction in note time alongside a non-significant change in measured burnout (Source: [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov)). Like Suki and Dragon Copilot, Abridge does not publish list pricing; third-party estimates place costs anywhere from roughly **\$208 to \$500-plus per provider per month** depending on deployment scope, though Abridge itself has never confirmed these figures publicly (Source: deepcura.com). Abridge's expansion into revenue-cycle and prior-authorization tooling also means the product increasingly competes with adjacent categories like computer-assisted coding, which buyers evaluating a pure documentation tool should factor into procurement scope.

Freed

Capabilities

Freed's core AI scribe generates specialty-specific notes from patient conversations, supports an instant template builder, and offers a "learn my format" feature that adapts to a clinician's preferred documentation style over successive edits (Source: getfreed.ai). Higher tiers add EHR push integration via a Chrome browser extension that works with any web-based EHR, visit summaries and pre-visit patient context, ICD-10 and CPT coding suggestions, and an AI clinician assistant with an embedded medical knowledge base (Source: getfreed.ai). In 2026 Freed added **Front Desk**, an AI receptionist that answers calls the front office misses during lunch, after hours, or overflow periods, and **Clinical Decision Support**, which surfaces evidence-based answers to clinical questions from more than 50 sources (Source: getfreed.ai). Freed publishes a claimed **98% recall on medical terms across 30-plus specialties**, tested against thousands of clinical concepts, according to the company's own marketing copy (Source: getfreed.ai).

Unlike Suki, Dragon Copilot, and Abridge, Freed explicitly targets independent and small-to-mid-sized clinics rather than large health systems: the company states plainly that it is "not made for massive healthcare systems" and instead built to serve community care one local clinic at a time (Source: getfreed.ai).

Adoption

Freed reports **26,000-plus clinicians** using the platform as of mid-2026, generating **32,589,627 patient visits transcribed in 2025** and returning **5.9 million hours** to clinicians per year, according to figures published on the company's own homepage (Source: getfreed.ai). At the time of its March 2025 Series A announcement, the company reported over **17,000 clinicians**, 70,000 notes generated daily, and more than **2 million monthly notes**, along with 2.7 million cumulative hours reclaimed, indicating roughly 50% clinician growth in a little over a year (Source: getfreed.ai). The company's Series A raised **\$30 million**, led by Sequoia Capital with participation from Scale Venture Partners, Daniel Gross, Gokul Rajaram, and Ted Zagat (Source: getfreed.ai). In a survey of one enterprise customer that Freed itself published, 100% of surveyed clinicians reported better work-life balance and 80% reported feeling happier in their jobs and delivering better patient care, though this is a single-customer, company-commissioned survey rather than an independent study (Source: getfreed.ai).

Strengths and Limitations

Freed's clearest differentiator is price transparency and accessibility: unlike its three enterprise-oriented competitors, Freed publishes its full pricing ladder without requiring a sales call, and offers a seven-day free trial with no credit card required (Source: getfreed.ai). The tradeoff is enterprise EHR depth: Freed's EHR strategy relies on a browser extension compatible with any web-based EHR rather than native, vendor-certified integrations comparable to Suki INSIDE or Abridge Inside, which may be a limitation for large systems wanting Epic-embedded workflows but is largely irrelevant for the independent-practice buyer Freed targets. On the compliance side, Freed states that its AI models are trained only on de-identified notes and never on protected health information, and that patient audio recordings are automatically deleted once a note is successfully generated, typically within 60 seconds of the encounter ending (Source: getfreed.ai). Freed states plainly that it is not FDA (U.S. Food and Drug Administration) approved and does not need to be, since it functions as a documentation support tool rather than a diagnostic device, a claim that applies to Freed's stated intended use, not automatically to every feature offered by every vendor (Source: getfreed.ai).

Feature Comparison

Table 1 below summarizes how the four vendors compare across the dimensions most cited in buyer research: pricing transparency, EHR integration, target buyer, and independent recognition.

PROCUREMENT DIMENSION	SUKI	DRAGON COPILOT	ABRIDGE	FREED
Typical buyer and workflow	Health systems and clinics seeking ambient documentation plus voice-enabled workflow features (Suki)	Health systems seeking Microsoft clinical documentation and workflow automation (Microsoft)	Enterprise health systems seeking ambient documentation embedded in clinical workflows (Abridge)	Individual clinicians and groups seeking a self-service ambient scribe (Freed)
Public pricing	Contact vendor	Contact vendor	Contact vendor	Public self-service tiers and separate group plans; confirm current amounts on the official pricing page
Documented integration approach	Epic workflows plus other EHR integrations described by Suki	Microsoft describes integration with major EHRs and its clinical workflow ecosystem (Microsoft)	Abridge describes Epic-embedded workflows and enterprise integrations (Abridge)	Browser-extension workflow for web-based EHRs, as described by Freed
Evidence to validate during procurement	Request specialty-level quality results, current integration scope, security documentation, and contract terms	Confirm licensed modules, supported EHR workflow, data handling, and contract scope	Review independent evaluations alongside vendor-reported deployment results	Validate workflow fit, BAA terms, retention settings, and integration behavior for the practice's EHR

This table intentionally limits itself to decision-useful characteristics supported by current vendor materials. Product capabilities, integrations, pricing, and contract scope can change; buyers should verify them directly during procurement rather than treat a comparison article as a substitute for current product documentation.

Performance and Benchmarks

Independent, peer-reviewed performance data on ambient AI scribes remains sparser than the volume of vendor marketing claims would suggest, and a 2026 narrative review published via PubMed Central is the most comprehensive synthesis available as of this report. Across the studies it aggregated, spanning Nuance DAX, Abridge, ChatGPT-4, Nabla, and TORTUS, most tools produced "modest to significant reductions in documentation time (ranging from ~1 to 2.1 minutes per note)," alongside improved perceived usability and reduced cognitive load (Source: [pmc.ncbi.nlm.nih.gov](#)). A separate five-week pre-post study of the Nabla ambient AI tool at University of Iowa Health, involving 38 volunteer physicians, found self-reported burnout on the Stanford Professional Fulfillment Index fell from 69% to 43% of participants, though improvements in professional fulfillment itself were not statistically significant (Source: [pmc.ncbi.nlm.nih.gov](#)).

Accuracy findings were considerably more cautious than vendor claims typically suggest. In a controlled study by Biro and colleagues comparing two commercial ambient digital scribe products across 44 simulated outpatient encounters, researchers identified **127 total errors**, an average of **2.9 errors per note**, and found that **70% of generated notes contained at least one error** (Source: [pmc.ncbi.nlm.nih.gov](#)). The two products showed statistically significant differences in error-type distribution (Fisher exact test, $P=0.002$), with one product producing predominantly omission errors (83% of its errors) and the other showing a more even split across omission, addition, and misplaced-text errors (Source: [pmc.ncbi.nlm.nih.gov](#)). A separate study by Kernberg and colleagues, testing ChatGPT-4 on SOAP-format note generation across 14 clinical encounters repeated three times

each, found an average of **23.6 errors per case** and, more strikingly, only **52.9% of data elements were consistently and correctly reported across all three repeated runs** of the same transcript, a finding the reviewers flagged as evidence of concerning non-determinism in AI-generated clinical documentation (Source: [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov/41111111/)).

Omission errors, rather than fabricated additions, dominate the error landscape and carry particular clinical risk because they require a reviewing clinician to recall specific conversational details from hours or days earlier rather than simply spotting an obviously wrong statement. The review's overall conclusion is measured: ambient AI scribes "show promise in reducing workload, improving efficiency, and decreasing burnout, but current systems still generate high omission rates and intermittent factual inaccuracies that may affect clinical decision-making," and the authors call for more rigorous, standardized evaluation before routine clinical adoption becomes unconditional (Source: [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov/41111111/)). This gap between vendor-published outcome metrics, often in the 70 to 90% range for satisfaction or effort-reduction, and the more conservative error rates found in controlled peer-reviewed testing is the single most important nuance for a buyer to hold in mind when evaluating marketing claims from any of the four vendors profiled here.

Table 2 below consolidates the individual peer-reviewed and quality-improvement studies referenced throughout this section, allowing side-by-side comparison of setting, sample size, and headline finding.

STUDY	AI TOOL EVALUATED	SETTING AND SAMPLE	HEADLINE FINDING
Biro et al. (2024)	Two unnamed commercial ambient digital scribe products	44 simulated outpatient encounters	70% of notes contained at least one error; 2.9 errors per note on average
Kernberg et al. (2024)	ChatGPT-4	14 simulated encounters, each repeated three times	23.6 errors per case; only 52.9% of data elements consistently reported across repeated runs
Haberle et al. (2024)	Nuance DAX	Intermountain Health, 12 specialties, 99 DAX users vs. 76 controls	Documentation time fell by roughly 0.76 minutes per note
Shah et al. (2024)	DAX Copilot	Stanford Health, 48 physicians	Median 20 minutes saved per half-day of clinic
Kakaday et al. (2025)	DAX Copilot	Samaritan Health, randomized pilot, 45 clinicians	Documentation time fell by 1.4 minutes per visit
Stults et al. (2025)	Abridge	Sutter Health, 100 clinicians, 57 paired responses	0.9-minute-per-patient reduction in note time; 7% reduction in burnout score
Misurac et al. (2025)	Nabla	University of Iowa Health, 38 volunteer physicians	Self-reported burnout fell from 69% to 43% of participants

Read together, the studies in Table 2 show a consistent pattern: time savings are real but modest, generally under two minutes per note or under half an hour per half-day, while accuracy and burnout outcomes vary considerably by tool, setting, and study design. No single tool in the table was tested head-to-head against another under identical conditions, which is itself a limitation of the current evidence base and a reason vendor-versus-vendor performance claims should be treated cautiously until independent, comparative trials become available.

Data Analysis and Evidence

The AI medical scribe market's growth trajectory is well documented by independent research firms, even where individual vendor pricing is not. Towards Healthcare, an industry research publisher, sizes the global AI-in-medical-scribing market at **USD 1.39 billion in 2025**, rising to **USD 1.67 billion in 2026** and reaching **USD 8.93 billion by 2035**, a compound annual growth rate of **20.48%** for the 2026 to 2035 period, with North America the dominant region and the United States sub-segment alone projected to grow from **USD 621.66 million in 2026 to USD 4.67 billion by 2035**.

(Source: towardshealthcare.com) (Source: towardshealthcare.com). The same analysis found the cloud-based deployment segment dominated the market and is also expected to grow at the fastest rate of any deployment mode through the forecast period, reflecting how quickly health systems have moved away from on-premise speech-recognition infrastructure (Source: towardshealthcare.com).

Physician-side survey data corroborates rapid, ongoing adoption. Medical Group Management Association (MGMA) polling from August 2025 found **71% of medical practice leaders** reported some use of AI for patient visits, though only **39%** of those using AI said it had actually reduced workload (Source: fiercehealthcare.com). Earlier MGMA polling from October 2024, with 302 responses, found nearly six in ten practice leaders (**59%**) named scribing and documentation tools as their organization's top AI priority, well ahead of revenue-cycle AI (19%) and patient-communication AI (10%) (Source: mgma.com). Among practice leaders who cited scribing as the top priority, roughly **80%** already had a tool in place (38%) or planned to add one within the year (42%), indicating the category has moved from early pilot to mainstream procurement for most surveyed practices (Source: mgma.com). A related MGMA poll found more than four practices in ten (43%) had added or expanded use of AI tools during 2024, up sharply from only 21% in a September 2023 poll, underscoring how quickly adoption accelerated once generative AI scribes reached the market (Source: mgma.com).

Capital markets data reinforces the same growth signal from the investment side. STAT News calculated that ambient AI scribe companies collectively announced at least **\$975 million** in new funding during 2025, with Abridge alone accounting for **\$550 million** across two rounds and Ambience Healthcare raising **\$243 million** in a Series C that valued it at **\$1.25 billion** (Source: statnews.com). This scale of capital deployment sits against a backdrop of intensifying competitive pressure from EHR vendors themselves: Epic signaled at its August 2025 Users Group Meeting that it would introduce a proprietary AI scribe, with industry chatter (unconfirmed by Epic) pegging a possible price point around **\$80 per provider per month**, a figure that, if accurate, would undercut most of the independent vendors covered in this report and could compress category-wide pricing (Source: fiercehealthcare.com). Epic has simultaneously continued partnering with Abridge and Microsoft's Dragon Copilot inside its own Toolbox ecosystem, having also added Ambience Healthcare to that program in 2025 (Source: fiercehealthcare.com).

An AI-scribe vendor is generally a HIPAA **business associate** when it creates, receives, maintains, or transmits PHI on behalf of a covered entity. HHS also cautions that merely selling software does not create that relationship when the vendor has no access to PHI, so the conclusion depends on the vendor's role, the data flow, and the service actually performed ([HHS business-associate guidance](https://www.hhs.gov/hipaa/for-professionals/special-topics/business-associates/guidance); [HHS software-vendor FAQ](https://www.hhs.gov/hipaa/for-professionals/special-topics/business-associates/software-vendor-faq)). When a business-associate relationship exists, the parties generally need written assurances defining permitted uses and safeguards for PHI ([HHS contract guidance](https://www.hhs.gov/hipaa/for-professionals/special-topics/business-associates/guidance)).

FDA status likewise turns on each software function's intended use and operation, not on the label "AI scribe." FDA identifies some software functions for administrative support or electronic patient records as non-device functions, while software that meets the device definition remains subject to FDA's digital-health policies. Its current clinical decision-support guidance applies a function-specific statutory test and includes both device and non-device examples ([FDA Clinical Decision Support Software guidance, January 2026](https://www.fda.gov/medical-devices/cybersecurity-and-data-protection/cybersecurity-guidance-clinical-decision-support-software)). Buyers should therefore verify the intended use and regulatory position of the specific module being purchased—especially any diagnostic, treatment, coding, or decision-support feature—rather than assume every product or future feature is categorically inside or outside FDA oversight.

Case Studies and Real-World Examples

Kaiser Permanente and Abridge

Kaiser Permanente, the nation's largest nonprofit integrated health system, made Abridge's assisted clinical documentation tool available to doctors and clinicians across its **40 hospitals** and more than **600 medical offices** in eight states and the District of Columbia beginning in August 2024, following roughly a year of joint testing and quality checks (Source: about.kaiserpermanente.org). Kaiser executives framed the deployment explicitly as a scale test of ambient AI, requiring patient consent and mandatory clinician review of every note before it entered the medical record (Source: about.kaiserpermanente.org). Kaiser's chief information officer for The Permanente Medical Group described the rollout as well received by both patients and doctors following diligent testing, and Abridge CEO Shiv Rao characterized the engagement as testing "the highest standards of a true enterprise-wide deployment" across scale, evaluation, and quality control (Source: about.kaiserpermanente.org).

Johns Hopkins Medicine and Abridge

Johns Hopkins Medicine finalized an agreement in December 2024 to deploy Abridge's platform across **6,700 clinicians**, six hospitals, and 40 patient-care centers, spanning all care settings and specialties within the system (Source: abridge.com). Johns Hopkins Associate Chief Medical Information Officer Dr. Manisha Loss stated the system was "excited for the opportunity to provide our clinicians with a tool to ease documentation burden," noting

Abridge was created by clinicians to capture the authentic clinical voice of the patient encounter. The deployment illustrates how a leading academic medical center's complex, multi-specialty documentation requirements pushed Abridge to extend its platform beyond a single ambulatory use case.

MedStar Health and Suki

MedStar Health began exploring documentation solutions as early as 2019, seeking a technology that could scale across the entire organization rather than working for only a subset of providers, according to Jeff Collins, vice president of the MedStar Institute for Innovation Business Innovation Lab (Source: beckershospitalreview.com). After evaluating multiple vendors, MedStar selected Suki, extending access to thousands of clinicians across specialties including primary care, cardiology, gastroenterology, and urgent care, with early feedback described as largely positive and further expansion planned into pediatrics and inpatient settings (Source: beckershospitalreview.com). Collins noted that clinician reaction had been mixed but broadly encouraging, observing that "the majority are excited about the efficiency gains, though some variability in adoption is natural" (Source: beckershospitalreview.com). The MedStar relationship was announced alongside Suki's **\$70 million** Series D funding round in October 2024, tying a major health system commitment directly to Suki's capital-raising narrative and illustrating how enterprise health system partnerships and startup fundraising in this category have become mutually reinforcing (Source: reuters.com).

Sutter Health, Abridge, and the Limits of Vendor-Reported Data

Sutter Health's experience with Abridge is instructive precisely because it is documented both in Abridge's own customer marketing and in independent peer-reviewed literature, allowing a direct comparison. Abridge's own case-study page reports that **78% of clinicians** at Sutter Health felt improved work satisfaction (Source: abridge.com). A separate, peer-reviewed pre-post quality-improvement study of the same Sutter Health deployment, involving 100 clinicians with 57 paired before-and-after responses, found a smaller but still measurable 0.9-minute-per-patient reduction in note time, a **7% reduction** in a validated burnout instrument, and improved provider-reported workflow, alongside a caveat that the change in burnout itself was not statistically significant (Source: [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov)). The juxtaposition of a strongly positive vendor-marketed satisfaction figure against a more modest, statistically cautious peer-reviewed result at the same customer site is a useful case study in how buyers should read AI scribe vendor case studies generally: directionally credible, but best triangulated against independent research wherever it exists.

(Hypothetical Example) A Multi-Specialty Group Practice Choosing Between Freed and an Enterprise Vendor

A hypothetical 12-physician multi-specialty group practice using a web-based EHR without an enterprise IT department illustrates how the four vendors' target-buyer differences play out in practice. Such a practice would likely be unable to obtain pricing from Suki, Dragon Copilot, or Abridge without engaging an enterprise sales process disproportionate to its size, and would lack the volume to negotiate favorable enterprise contract terms even if it did. Freed's published, self-service pricing tiers and browser-extension-based EHR integration are structurally suited to this buyer profile, while the same practice would gain comparatively little from Suki's or Abridge's deeper native Epic integrations, since it does not run Epic. This scenario, illustrative rather than drawn from a specific named practice, mirrors the buyer segmentation implied by Freed's own public positioning against "massive healthcare systems."

Implications and Future Directions

The most consequential near-term dynamic in this category is Epic's own move into ambient documentation. With roughly 42% of the United States hospital market already standardized on Epic, an Epic-native scribe, even a less sophisticated one, carries structural distribution advantages that pure-play vendors cannot match through product quality alone (Source: fiercehealthcare.com). Industry observers interviewed by Fierce Healthcare characterized this less as an extinction event for independent vendors and more as the end of what one executive called the "transcription wars," pushing companies like Abridge and Suki to accelerate expansion into adjacent workflows such as prior authorization, revenue-cycle coding, and clinical decision support, where Epic's own tooling is comparatively less mature (Source: fiercehealthcare.com). Rumored Epic pricing around **\$80 per provider per month** would, if accurate, apply meaningful downward pressure across the category regardless of whether health systems ultimately choose Epic's native tool or a third-party alternative.

A second implication concerns the persistent gap between vendor-reported outcome metrics and peer-reviewed evidence documented throughout this report. As health systems and the life-sciences and healthcare AI advisory ecosystem supporting them move from pilot programs to enterprise-wide rollouts, procurement and clinical-informatics teams will increasingly need independent validation frameworks rather than relying solely on vendor case studies, particularly given the specialty-specific error risk the PMC narrative review flagged for high-acuity, time-sensitive documentation such as cardiology (Source: [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov)). Organizations evaluating this category as part of a broader AI governance and vendor-risk strategy, a domain adjacent to, but distinct from, the pure product evaluation covered in this report, will find that documentation error auditing, BAA review, and specialty-specific validation increasingly sit alongside price and EHR fit as first-order procurement criteria rather than afterthoughts.

Finally, the rapid blurring of category boundaries—Abridge into prior authorization and coding, Suki into general-purpose voice assistance, and Freed into front-desk automation and clinical decision support—suggests that "AI scribe" as a standalone product category may itself be transitional. Buyers evaluating any of the four vendors profiled here in 2026 should expect the comparison to look meaningfully different within 18 to 24 months, both because of Epic's entry and because each vendor is actively expanding its own product surface area well beyond documentation.

Frequently Asked Questions (FAQs)

What is the best AI medical scribe in 2026? There is no single best answer independent of buyer context. Abridge won KLAS Research's **Best in KLAS** award for Ambient Speech in both 2025 and 2026 (Source: [abridge.com](https://www.abridge.com)), while Freed is the most accessible option for independent and small-group clinicians due to its transparent, low-cost pricing tiers starting at \$39 per month.

Are AI scribes HIPAA-compliant? No product is automatically compliant in every deployment. When a vendor creates, receives, maintains, or transmits PHI on behalf of a covered entity, it will generally be a business associate and a compliant written agreement is typically required. The covered entity and vendor must also meet all applicable Privacy, Security, and Breach Notification Rule obligations; encryption, access controls, audit controls, risk analysis, policies, and workforce practices are parts of that broader program rather than a three-item definition of compliance ([HHS business-associate guidance](#); [HHS Security Rule guidance](#)).

What does Nuance DAX pricing look like? Nuance DAX, now Microsoft Dragon Copilot, does not publish subscription list pricing; Microsoft publishes only the per-unit rate for its pay-as-you-go consumption tier and requires direct sales engagement for base subscription pricing (Source: [samexpert.com](https://www.samexpert.com)). Third-party analyst estimates cluster around \$369 to \$830-plus per provider per month, but these are not vendor-confirmed figures (Source: [veroscribe.com](https://www.veroscribe.com)).

How does Abridge compare to Suki? Both are enterprise-oriented ambient scribes with broad EHR integration, but Abridge reports greater disclosed funding (approximately \$800 million versus Suki's \$168 million) and won Best in KLAS for Ambient Speech in 2025 and 2026. Their published organization counts use different definitions—Abridge reports health systems, while Suki reports health systems and clinics—so they do not establish a directly comparable customer-base ranking. Suki differentiates on voice-command breadth and its self-described diversified, EHR-agnostic footprint that is not dependent on any single EHR platform.

How much does Freed AI scribe cost? Freed publishes four tiers: Starter at \$39 per month (billed annually) covering up to 40 notes, Core at \$79 per month with unlimited notes, Premier at \$104 to \$119 per month with EHR push and coding assistance, and custom Groups pricing for clinics.

Do AI scribes integrate with Epic and Cerner? Yes, for the three enterprise vendors. Suki, Dragon Copilot, and Abridge all offer native Epic integrations (Suki INSIDE, Dragon Copilot embedded, and Abridge Inside respectively), and Suki additionally supports Oracle Health/Cerner environments. Freed instead uses a Chrome browser extension compatible with any web-based EHR rather than vendor-certified native integrations.

How accurate are ambient AI documentation tools? Accuracy varies by product, specialty, encounter, and evaluation method, and direct independent comparisons remain limited. In one controlled simulation involving 44 outpatient encounters and two unnamed commercial products, 70% of generated notes contained at least one error, averaging 2.9 errors per note. That study does not measure Suki, Dragon Copilot, Abridge, or Freed specifically; it supports clinician review and local validation rather than a category-wide error rate ([reviewed evidence](#)).

Conclusion

Suki, Nuance DAX (now Microsoft Dragon Copilot), Abridge, and Freed each answer a different version of the AI medical scribe question rather than competing head-to-head on identical terms. Abridge won the 2025 and 2026 Best in KLAS awards in the Ambient Speech category and reports major health-system deployments, including at Kaiser Permanente. Dragon Copilot carries the weight of Microsoft's balance sheet and its historic depth in speech recognition, but remains constrained, like Suki and Abridge, by pricing opacity that pushes smaller buyers toward extended sales cycles before they can even evaluate cost. Suki's differentiated voice-command architecture and platform partnerships with companies like Zoom give it a broader technology-licensing angle than a pure documentation tool, while the vendors' differently defined public adoption figures do not support a direct scale ranking; Abridge separately holds the 2025 and 2026 Best in KLAS awards for Ambient Speech. Freed remains the clear outlier and, for its target buyer, arguably the simplest recommendation of the four: transparent pricing, a genuine free trial, and a product built explicitly for the independent and small-group practice that the three enterprise vendors do not efficiently serve.

The evidence gathered across peer-reviewed literature, KLAS rankings, and named case studies supports one consistent takeaway for any organization evaluating this category: vendor-reported efficiency and satisfaction figures, often in the 70 to 90% range, should be treated as directionally useful but not as a substitute for independent validation, particularly given that one controlled simulation of two unnamed products found

errors in 70% of generated notes—a result that should not be generalized to every vendor or deployment. Epic's entry into native ambient scribing adds further uncertainty to the category's pricing and competitive structure over the next 18 to 24 months. For health systems and life-sciences organizations building broader AI governance and vendor-evaluation frameworks around clinical documentation tools, the practical path forward is the same regardless of which of the four vendors is ultimately selected: confirm the BAA, pressure-test vendor accuracy claims against specialty-specific use cases, and treat pricing opacity itself as a data point in the evaluation rather than an obstacle to route around.

Tags: ai medical scribe, suki ai, nuance dax, dragon copilot, abridge ai, freed ai scribe, ambient ai documentation, hipaa compliant ai scribe, ehr integration

IntuitionLabs - Industry Leadership & Services

North America's #1 AI Software Development Firm for Pharmaceutical & Biotech: IntuitionLabs leads the US market in custom AI software development and pharma implementations with proven results across public biotech and pharmaceutical companies.

Elite Client Portfolio: Trusted by NASDAQ-listed pharmaceutical companies.

Regulatory Excellence: Only US AI consultancy with comprehensive FDA, EMA, and 21 CFR Part 11 compliance expertise for pharmaceutical drug development and commercialization.

Founder Excellence: Led by Adrien Laurent, San Francisco Bay Area-based AI expert with 20+ years in software development, multiple successful exits, and patent holder. Recognized as one of the top AI experts in the USA.

Custom AI Software Development: Build tailored pharmaceutical AI applications, custom CRMs, chatbots, and ERP systems with advanced analytics and regulatory compliance capabilities.

Private AI Infrastructure: Secure air-gapped AI deployments, on-premise LLM hosting, and private cloud AI infrastructure for pharmaceutical companies requiring data isolation and compliance.

Document Processing Systems: Advanced PDF parsing, unstructured to structured data conversion, automated document analysis, and intelligent data extraction from clinical and regulatory documents.

Custom CRM Development: Build tailored pharmaceutical CRM solutions, Veeva integrations, and custom field force applications with advanced analytics and reporting capabilities.

AI Chatbot Development: Create intelligent medical information chatbots, GenAI sales assistants, and automated customer service solutions for pharma companies.

Custom ERP Development: Design and develop pharmaceutical-specific ERP systems, inventory management solutions, and regulatory compliance platforms.

Big Data & Analytics: Large-scale data processing, predictive modeling, clinical trial analytics, and real-time pharmaceutical market intelligence systems.

Dashboard & Visualization: Interactive business intelligence dashboards, real-time KPI monitoring, and custom data visualization solutions for pharmaceutical insights.

Leading Claude & ChatGPT AI Enablement and Training: Help biotech and pharmaceutical teams adopt Claude and ChatGPT through practical training, governed workflows, use-case development, and implementation designed for regulated environments.

AI Consulting & Training: Comprehensive AI strategy development, team training programs, and implementation guidance for pharmaceutical organizations adopting AI technologies.

Contact founder Adrien Laurent and team at intuitionlabs.ai/contact for a consultation.

DISCLAIMER

The information contained in this document is provided for educational and informational purposes only. We make no representations or warranties of any kind, express or implied, about the completeness, accuracy, reliability, suitability, or availability of the information contained herein.

Any reliance you place on such information is strictly at your own risk. In no event will IntuitionLabs.ai or its representatives be liable for any loss or damage including without limitation, indirect or consequential loss or damage, or any loss or damage whatsoever arising from the use of information presented in this document.

This document may contain content generated with the assistance of artificial intelligence technologies. AI-generated content may contain errors, omissions, or inaccuracies. Readers are advised to independently verify any critical information before acting upon it.

All product names, logos, brands, trademarks, and registered trademarks mentioned in this document are the property of their respective owners. All company, product, and service names used in this document are for identification purposes only. Use of these names, logos, trademarks, and brands does not imply endorsement by the respective trademark holders.

IntuitionLabs.ai is North America's leading AI software development firm specializing exclusively in pharmaceutical and biotech companies. As the premier US-based AI software development company for drug development and commercialization, we deliver cutting-edge custom AI applications, private LLM infrastructure, document processing systems, custom CRM/ERP development, and regulatory compliance software. Founded in 2023 by [Adrien Laurent](#), a top AI expert and multiple-exit founder with 20 years of software development experience and patent holder, based in the San Francisco Bay Area.

This document does not constitute professional or legal advice. For specific guidance related to your business needs, please consult with appropriate qualified professionals.

© 2026 IntuitionLabs.ai. All rights reserved.