SACRA

EQUITY RESEARCH

UPDATED
07/17/2025

OpenEvidence

TEAM

Jan-Erik Asplund
Co-Founder
jan@sacra.com

Marcelo Ballve
Head of Research
marcelo@sacra.com

DISCLAIMERS

This report is for information purposes only and is not to be used or considered as an offer or the solicitation of an offer to sell or to buy or subscribe for securities or other financial instruments. Nothing in this report constitutes investment, legal, accounting or tax advice or a representation that any investment or strategy is suitable or appropriate to your individual circumstances or otherwise constitutes a personal trade recommendation to you.

This research report has been prepared solely by Sacra and should not be considered a product of any person or entity that makes such report available, if any.

Information and opinions presented in the sections of the report were obtained or derived from sources Sacra believes are reliable, but Sacra makes no representation as to their accuracy or completeness. Past performance should not be taken as an indication or guarantee of future performance, and no representation or warranty, express or implied, is made regarding future performance. Information, opinions and estimates contained in this report reflect a determination at its original date of publication by Sacra and are subject to change without notice.

Sacra accepts no liability for loss arising from the use of the material presented in this report, except that this exclusion of liability does not apply to the extent that liability arises under specific statutes or regulations applicable to Sacra. Sacra may have issued, and may in the future issue, other reports that are inconsistent with, and reach different conclusions from, the information presented in this report. Those reports reflect different assumptions, views and analytical methods of the analysts who prepared them and Sacra is under no obligation to ensure that such other reports are brought to the attention of any recipient of this report.

OpenEvidence

Visit Website

AI copilot for doctors to assist in making critical decisions at the point of care

#ai

REVENUE
$50,000,000
2025

VALUATION
$3,500,000,000
2025

FUNDING
$100,000,000
2025

Details

HEADQUARTERS
Cambridge, MA

CEO
Daniel Nadler

SACRA

Revenue

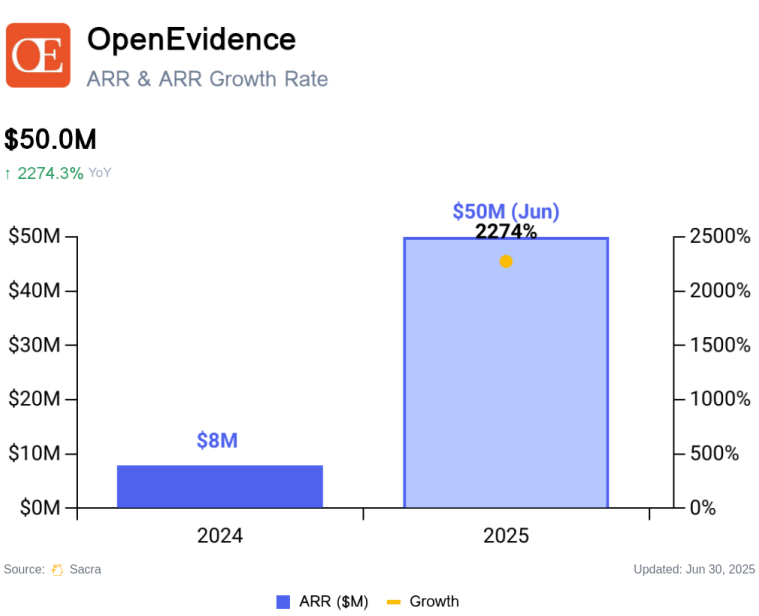

Sacra estimates OpenEvidence reached $50 million in annualized revenue as of June 2025, growing 30% month-over-month from $7.9 million in December 2024. The company monetizes through pharmaceutical and medical device advertisements, achieving CPMs of $70-150 compared to $5-15 for typical social media platforms, translating to approximately $124 in average revenue per user (ARPU).

The company's free access model has driven rapid physician adoption by bypassing lengthy hospital procurement cycles, contributing to its capture of 40% of U.S. physicians. Strategic content partnerships with The New England Journal of Medicine and JAMA in early-to-mid 2025 have strengthened the platform's credibility and data moat.

For comparison, Doximity (NYSE: DOCS) generates $570 million in trailing twelve-month revenue with 80% from advertising at $228 ARPU, while UpToDate parent Wolters Kluwer (EURONEXT: WKL) achieves $595 million in revenue through a traditional $500 per seat licensing model. OpenEvidence's advertising-based approach positions it to maintain rapid growth through both direct physician adoption and potential enterprise expansion, operating in a healthcare advertising market that supports premium pricing for targeted physician audiences.

Valuation

OpenEvidence raised $210 million in July 2025 led by GV (Google's venture arm) and Kleiner Perkins at a $3.5 billion valuation for a ~70x multiple on $50M in annualized revenue.

This follows the company's $1 billion valuation achieved in its Series A round led by Sequoia Capital in February 2025, where it raised $75 million and brought total funding to over $100 million at that time.

Key investors across both rounds include GV, Kleiner Perkins, Sequoia Capital, Coatue, Conviction Partners, Thrive Capital, Breyer Capital, and Mayo Clinic through their Platform Accelerate program.

Product

With 1.5M new medical papers published every year—doubling the total body of medical literature every five years—doctors struggle to ingest the data needed to create up-to-date treatment plans. That challenge, combined with the emergence of large language models in 2017, inspired the founding of OpenEvidence in 2021 as a search engine and chatbot trained on open-sourced medical papers and textbooks, purpose-built to answer complex clinical questions like *"what antibiotic is best for a trach-dependent pneumonia patient with cerebral palsy"*—with citations to the latest research.

The core workflow begins with physician verification through NPI scanning or hospital email confirmation, ensuring only licensed healthcare providers gain access. Clinicians then input natural language queries about patient scenarios, such as treatment options for specific conditions or drug interactions. The platform's retrieval engine searches exclusively through licensed medical content from sources like NEJM, JAMA, specialty guidelines, and drug labels.

A medical-tuned large language model generates responses with inline citations, allowing users to tap on references to view original study abstracts or guideline paragraphs. The system includes built-in clinical calculators that auto-populate when relevant, enabling direct integration into clinical documentation. Color-coding indicates evidence strength levels to help clinicians assess recommendation quality.

OpenEvidence 2.0 expanded beyond clinical search to include administrative functions like generating prior authorization letters, patient instructions, and ICD-10 coding suggestions. The platform also offers workflow modules for order-set recommendations and discharge summary drafting, with mobile-first design optimized for bedside use during hospital rounds.

The technology runs on a vertical LLM trained exclusively on licensed medical texts, avoiding public internet data contamination. A retrieval-augmented generation pipeline requires deterministic citation linking, rejecting answers that cannot be properly sourced. The system incorporates reinforcement learning from clinician feedback and maintains HIPAA compliance through edge encryption without retaining patient health information.

Business Model

OpenEvidence operates a B2B2C freemium model that provides free access to verified U.S. physicians while monetizing through advertising and enterprise subscriptions. By making their product (1) free, and (2) touching no patient records, OpenEvidence bypassed the ~18-month sales cycles in healthcare en route to capturing 40% of U.S. physicians—creating the bottom-up leverage necessary to now go upmarket, plug into the EHR, and sell into hospitals on enterprise contracts as they build agents with the capacity to extend into the rest of the diagnosis and treatment workflow, from auto-drafting clinical notes (Abridge, Ambience) and prior authorizations (Cohere Health) to checking drug interactions in real time (Lexicomp, EURONEXT: WKL).

The company's value delivery mechanism combines a specialized medical AI platform with exclusive content licensing agreements, creating a differentiated clinical decision support tool. The go-to-market strategy targets individual clinicians through direct app downloads and hospital system procurement. Revenue streams include targeted advertising to pharmaceutical companies and medical device manufacturers, premium enterprise features for health systems, and API licensing for clinical decision support integration.

Where UpToDate (EURONEXT: WKL) charges hospitals $500/seat for their expert-curated, encyclopedia-style reviews of the recent literature, OpenEvidence gives its chatbot away for free—monetizing through pharma and med-device ads at $70–$150+ CPM ($124 ARPU) compared to $5–15 for the average ad on Facebook or Instagram.

The core monetization logic shifts from advertising-supported free access toward enterprise per-seat pricing as the platform integrates with electronic health record systems. Early FHIR-based pilots with Epic installations demonstrate how embedding within clinical workflows can increase average revenue per user significantly compared to standalone app usage.

Competition

EHR gatekeepers

Epic and Oracle-Cerner represent the most significant competitive threat through vertical integration of AI capabilities directly into electronic health record systems. Epic's Clinical Insights pilot combines GPT-4 with UpToDate content, providing physicians with AI-generated summaries within their existing Hyperspace workflow. This approach leverages Epic's dominant market position and reduces switching costs for healthcare providers already embedded in their ecosystem.

Oracle-Cerner's Clinical Digital Assistant bundles AI reasoning capabilities with their Millennium platform, emphasizing voice interaction and clinical decision support. These EHR-native solutions threaten OpenEvidence's mobile and web entry points by offering integrated experiences that don't require separate app downloads or workflow changes.

Content incumbents

UpToDate leads the established clinical reference market with 2 million global users across 44,000 organizations, now adding AI-powered question-answering capabilities that cite their extensive topic database. Their advantages include unmatched clinical trust, established procurement relationships, and comprehensive specialty coverage built over decades.

Elsevier ClinicalKey and EBSCO's DynaMed are similarly upgrading their reference platforms with interactive chat features linked to structured medical monographs. These incumbents compete on evidence fidelity and institutional relationships but face user experience challenges compared to OpenEvidence's consumer-style interface design.

Pure-play AI copilots

Glass Health focuses specifically on differential diagnosis generation and clinical order sets, targeting the diagnostic reasoning workflow that represents a core use case for clinical AI. DeepEvidentia emphasizes rapid model development and claims significantly lower costs for replicating clinical reasoning capabilities.

Hippocratic AI and other specialized medical AI companies are building competing platforms with different approaches to clinical decision support, creating a crowded field of startups racing to establish market position before larger technology companies or healthcare incumbents dominate the space.

Along with other vertical AI companies like Harvey (law) and Hebbia (finance), OpenEvidence in medicine is part of the next wave of LLM applications that win on the basis of their proprietary data from (1) partnerships, (2) users, and (3) integration into workflows—rather than competing with OpenAI, Anthropic, and Google on general-purpose model capabilities.

TAM Expansion

Workflow integration

OpenEvidence’s next phase of growth comes from expanding beyond clinical search into the broader set of medical workflows, building agents with the capacity to extend into the rest of the diagnosis & treatment workflow, from auto-drafting clinical notes and prior authorizations to surfacing context-specific drug interaction risks and generating discharge summaries.

Where the current ecosystem is fragmented—doctors jump between UpToDate for research, Lexicomp for drug data, Abridge for notes, and Cohere Health for prior auth—OpenEvidence aims to consolidate all of that into a single interface that responds in real time to complex queries, embedded directly into the clinical workflow via integrations with Epic, Cerner, and other EHR systems using HL7 FHIR APIs.

These integrations unlock a new enterprise TAM where OpenEvidence can command 5–10x higher ARPU compared to its ad-supported model by moving from reference tool to real-time, in-workflow system. Already, physicians use OpenEvidence to generate ICD-10 codes, write prior auth letters, and suggest order sets at the bedside. Going forward, discharge summaries, differential diagnoses, and AI-assisted charting are within reach—turning the platform into a unified intelligence layer across hospital systems.

This workflow expansion reinforces the product’s defensibility: embedding OpenEvidence into daily routines raises switching costs, while each new use case generates structured usage data that improves the underlying model and makes the product more indispensable—tightening the feedback loops that power OpenEvidence’s flywheel.

Customer base expansion

The platform currently serves approximately 25% of U.S. physicians but can expand to 5.2 million nurses and advanced practice providers who share similar clinical information needs. This expansion would effectively double the domestic user base while tapping into a underserved market segment lacking dedicated AI tools.

Pharmaceutical medical science liaisons, payer medical policy teams, and life science contract research organizations represent adjacent customer segments facing similar literature synthesis challenges. These enterprise users could support higher-value subscription models compared to individual clinician access.

Geographic and enterprise integration

International expansion targets 15 million physicians globally, with English-first markets like the UK, Canada, and Australia representing immediate opportunities with lower regulatory barriers. Multilingual model development could unlock European Union and Latin American markets where clinical AI adoption is accelerating.

FHIR-based integration with Epic and other EHR systems transforms OpenEvidence from a reference tool into a workflow system embedded within clinical documentation. This integration enables per-seat enterprise pricing models that could increase average revenue per user by 5-10x compared to advertising-supported individual access, while creating switching costs that improve customer retention.

Risks

Regulatory liability: Clinical AI tools face increasing scrutiny from medical malpractice insurers and regulatory bodies as healthcare providers rely more heavily on algorithmic recommendations for patient care decisions. Any high-profile case where OpenEvidence's recommendations contribute to adverse patient outcomes could trigger liability claims and regulatory restrictions that fundamentally alter the company's business model and market access.

EHR integration dependence: OpenEvidence's long-term revenue growth depends heavily on successful integration with dominant EHR platforms like Epic and Oracle-Cerner, which control clinical workflow access for most U.S. healthcare providers. These EHR vendors have strong incentives to develop competing AI capabilities internally or partner with established players like Microsoft, potentially blocking third-party integrations or demanding unfavorable revenue-sharing arrangements that undermine OpenEvidence's unit economics.

Content licensing costs: The company's competitive advantage relies on exclusive licensing agreements with major medical publishers like NEJM and JAMA, but these content costs could escalate rapidly as publishers recognize the value of their data for AI training. Rising licensing fees combined with potential competition from publishers developing their own AI platforms could compress margins and force OpenEvidence to raise subscription prices beyond what individual clinicians or health systems are willing to pay.

DISCLAIMERS

This report is for information purposes only and is not to be used or considered as an offer or the solicitation of an offer to sell or to buy or subscribe for securities or other financial instruments. Nothing in this report constitutes investment, legal, accounting or tax advice or a representation that any investment or strategy is suitable or appropriate to your individual circumstances or otherwise constitutes a personal trade recommendation to you.

This research report has been prepared solely by Sacra and should not be considered a product of any person or entity that makes such report available, if any.

Information and opinions presented in the sections of the report were obtained or derived from sources Sacra believes are reliable, but Sacra makes no representation as to their accuracy or completeness. Past performance should not be taken as an indication or guarantee of future performance, and no representation or warranty, express or implied, is made regarding future performance. Information, opinions and estimates contained in this report reflect a determination at its original date of publication by Sacra and are subject to change without notice.

Sacra accepts no liability for loss arising from the use of the material presented in this report, except that this exclusion of liability does not apply to the extent that liability arises under specific statutes or regulations applicable to Sacra. Sacra may have issued, and may in the future issue, other reports that are inconsistent with, and reach different conclusions from, the information presented in this report. Those reports reflect different assumptions, views and analytical methods of the analysts who prepared them and Sacra is under no obligation to ensure that such other reports are brought to the attention of any recipient of this report.

Published on Jul 17th, 2025