ChatGPT in healthcare: What the numbers tell us about consumer behavior and market opportunity

JAN 07, 2026 • PAID

5

Share

Abstract

OpenAI published January 2026 data revealing ChatGPT's healthcare footprint: 5% of global messages (billions weekly) concern healthcare topics, with 1 in 4 we active users and 40M+ daily users engaging on health questions. In the US speci 1.6-1.9M weekly messages focus on health insurance navigation, 600K weekly messages originate from hospital deserts, and 70% of health conversations occur outside clinic hours. Survey data shows 60% of US adults used AI for health in tl past 3 months. Provider adoption accelerated dramatically in 2024, with physici usage jumping from 38% to 66% year-over-year, and 46% of nurses using AI weel The report profiles patient self-advocacy cases (insurance appeals, medication interaction checks, urgent triage), provider efficiency tools (AI scribes, clinical decision support), and early-stage scientific applications (drug repurposing, gen analysis, clinical education simulations). Policy recommendations center on med data access, clinical trial infrastructure, FDA device pathways, and workforce transitions.

Table of Contents

The Scale Play Nobody Saw Coming

The healthcare AI narrative has been remarkably consistent for the past few yea promising pilots, limited adoption, regulatory uncertainty, reimbursement chall workflow friction. The usual suspects in digital health have been grinding throu these problems with varying degrees of success. Then OpenAI drops usage numl that reframe the entire conversation.

Five percent of all ChatGPT messages globally touching healthcare translates to billions of messages per week. Not millions. Billions. This is not a pilot program limited rollout. This is organic adoption at consumer internet scale, happening completely outside traditional healthcare purchasing cycles and without a single payor contract or provider integration. The report pegs daily healthcare-focused at over 40 million globally. For context, Epic's MyChart patient portal, considere gold standard for patient engagement in the US, has roughly 200 million patient records total, with a fraction of those representing active monthly users.

The consumer behavior pattern matters more than the absolute numbers. People not waiting for their health system to offer an AI tool or for their insurance com to approve access. They are going directly to a general-purpose AI, for free, and asking it healthcare questions at scale. This bypasses the entire traditional go-to market motion in healthcare IT, which typically involves 18-24 month sales cycle pilot programs, integration work, training, and change management.

The 70% of conversations happening outside clinic hours detail is particularly revealing. This is not people using AI as a curiosity during downtime. This is problem-solving behavior when the healthcare system is literally closed. The emergency room is always open, but people are choosing ChatGPT instead for tr information gathering, and decision support. That substitution effect has real implications for how we think about access, utilization, and where value accrues the system.

Insurance Navigation as the Wedge Product

The 1.6-1.9 million weekly messages about health insurance represent perhaps t most interesting vertical within the broader healthcare usage. Insurance navigat famously terrible, deliberately opaque, and represents a massive pain point for consumers. It also happens to be an area where AI can provide genuine value wit requiring clinical accuracy or FDA approval.

Comparing plans, understanding coverage details, decoding explanation of bene documents, figuring out prior authorization requirements, these are all informa organization and translation problems that LLMs handle reasonably well. The re includes a case study of a Seattle patient using ChatGPT to research clinical literature, build a cited literature review, and successfully appeal an insurance d for a rare autoimmune condition. That is exactly the kind of high-stakes administrative warfare that people are desperate for help with.

The commercial opportunity here extends well beyond just answering one-off questions. Insurance navigation as a category probably runs $20B+ in annual cos when you factor in call centers, broker commissions, administrative overhead, an utilization management. If AI can reduce that friction even modestly, the saving to multiple stakeholders: employers buying coverage, health plans managing administrative costs, and consumers dealing with bills and coverage gaps.

Several startups are already building in this direction. Guidewell, Rightway Healthcare, and Accolade have various takes on navigation services, mostly relyi human navigators augmented by some technology. Intercom-style AI chat interfa specifically for benefits navigation seem like an obvious next product category, whether built by incumbents or new entrants. Expect to see this show up in brok platforms, benefits administration systems, and potentially as a direct-to-consur subscription product.

The trickier question is whether this kind of AI-assisted navigation actually ben cost curve or just makes the existing dysfunction slightly more tolerable. If peop figure out their benefits faster but the underlying plan design and pricing incen remain misaligned, you have not fundamentally improved anything. The optimis case is that better information access leads to better decision-making and more consumer pressure on plans to simplify. The pessimistic case is that you have jus automated the confusion slightly more efficiently.

Rural Access and the Hospital Desert Problem

The hospital desert data deserves more attention than it is probably going to get OpenAI defined hospital deserts as locations more than 30 minutes from a gener medical or children's hospital, then measured ChatGPT usage from those areas. found nearly 600K healthcare-related messages per week from these deserts, wit Wyoming, Oregon, Montana, South Dakota, and Vermont showing the highest ra usage relative to population.

This is not surprising once you think about it, but it is validating. Rural hospital closures have been accelerating for over a decade. Ten rural hospitals per year or average are either closing completely or converting to models without inpatient Nearly half of rural hospitals operate with negative margins. Over 400 hospitals 38 states are considered vulnerable to closure by Chartis. Services like obstetrics chemotherapy have been disappearing from rural facilities for 15 years.

People in these areas have been figuring out workarounds for a while, everything driving two hours to see specialists to relying on telemedicine to just going with care. AI-based information gathering slots naturally into that workaround toolk your nearest hospital does not have neurology or cardiology and getting an appointment means a half-day trip, asking ChatGPT questions before, during, a after that trip makes complete sense.

The report profiles a family physician in Miles City, Montana, population 8,400, Oracle Clinical Assist (powered by OpenAI models) as an AI scribe to reduce documentation burden. She also uses OpenEvidence for diagnosis and guideline checks. This is the pragmatic use case: rural providers are stretched thin, seeing sicker patients who traveled farther, and they need tools that let them spend less on administrative work and more time on actual care delivery.

The challenge is that AI does not solve the structural problems causing rural ho closures in the first place. Low population density, older and sicker patient populations, lower commercial insurance penetration, and Medicare/Medicaid reimbursement that does not cover costs. AI can help the remaining clinicians b more efficient and help patients navigate around access gaps, but it is not reope closed labor and delivery units or replacing a shuttered emergency department.

For investors, the rural access angle is probably more about market sizing and T expansion than a distinct investment category. Tools that work in rural settings to be extremely low-friction, require minimal technical support, and provide cle time savings or clinical value. Anything requiring complex integration, ongoing support, or behavior change that depends on abundant staff is not going to work companies that succeed here will be building products that are simple enough t deploy at scale across resource-constrained settings, which probably means they also well-positioned for other underserved markets globally.

Provider Adoption Crosses the Chasm

The physician adoption data is where the report gets really interesting from an enterprise software perspective. American Medical Association survey data show physician AI usage for at least one specific use case jumped from 38% in 2023 to in 2024. That is a 28 percentage point increase in a single year. Half of family physicians and primary care physicians reported using AI tools for at least one w use case.

For context, EMR adoption took roughly two decades to reach comparable penet levels, driven heavily by HITECH Act meaningful use incentives worth tens of b of dollars. AI tools are hitting similar adoption rates in 1-2 years with zero regul mandate and mostly bottom-up, individual clinician adoption.

The use case breakdown matters. Twenty-one percent of physicians used AI to document billing codes, charts, or visit notes in 2024, up from 13% in 2023. Twe percent used AI to assist in diagnosis, basically flat versus 11% the prior year. Th quarters of physicians consider AI helpful for work efficiency, 72% consider it h for diagnostic ability. The biggest jumps in perceived helpfulness from 2023 to 2 were managing stress and offering personalized care.

Documentation is clearly the wedge. Ambient AI scribes that listen to patient encounters and generate draft notes are the killer app right now. Nuance DAX, S Abridge, DeepScribe, and others are all growing quickly in this category. Physic hate EMR documentation, it contributes directly to burnout, and AI scribes pro immediate, tangible time savings. This is a true painkiller, not a vitamin.

The diagnostic assistance use case being flat year-over-year is more interesting. Diagnosis is higher stakes, has more potential liability exposure, and requires hi accuracy and trust. Physicians seem willing to experiment with AI for diagnosis not yet comfortable relying on it. That could change as model accuracy improves as clinical validation studies provide more evidence, but right now the adoption looks more cautious.

The 46% of nurses using AI weekly is also notable. Nursing workflows involve extensive documentation, care coordination, patient education, and administrati tasks. AI tools that help with any of those reduce cognitive load and free up time direct patient care. Medical librarians at 53% weekly usage and pharmacists at 4 suggests that information-intensive roles are adopting faster, which makes sense

For enterprise software companies selling into health systems, this data is a mixe
signal. On one hand, widespread clinician adoption of AI tools creates demand a
makes AI-related procurement easier. On the other hand, much of this adoption
happening via free consumer tools like ChatGPT rather than through enterprise
contracts. Health systems are going to need to figure out their AI strategy quickl
either by deploying approved enterprise tools or by creating guardrails around
clinician use of external AI services.

What This Means for Digital Health Business Models

The traditional digital health business model has been some variation of: build a
clinical workflow tool, get it adopted by a few progressive health systems or phy
groups, demonstrate ROI, expand to more sites, eventually get acquired by an
incumbent or go public. This works when the product delivers clear value, integ
reasonably well with existing systems, and solves a problem that buyers are willi
pay for.

AI changes several aspects of that model. First, the consumer adoption path is n
viable in ways it was not before. If millions of people are already using general-
purpose AI for healthcare questions, building a more specialized, healthcare-foc
AI tool and going direct to consumers is suddenly much more plausible. You do
need to convince people that AI is useful for health, they are already doing it. Yo
need to offer something better than the general-purpose option.

Second, the bottoms-up clinician adoption path is much faster. If individual doc
nurses, and pharmacists are already using AI tools in their workflows, selling an
enterprise version to the health system becomes easier. The product has built-in
champions who are already using something similar. This is the Dropbox or Slac
playbook applied to healthcare IT.

Third, the regulatory and reimbursement questions become less of a barrier for
certain use cases. Administrative tools, documentation assistance, patient educa
care navigation, these generally do not require FDA clearance or CPT codes. You

build and deploy them much faster than traditional digital therapeutics or diagr tools. The market is large enough that you can build a meaningful business just administrative and workflow efficiency side without ever touching clinical decis making.

The risk is commoditization. If the underlying LLM capabilities are widely avail via API and the primary value is in the user interface or workflow integration, th differentiation becomes harder. Every EMR vendor, care management platform, patient engagement tool is going to bolt on some version of AI assistance. Standi out in that environment requires either proprietary data, specialized domain expertise, or distribution advantages.

The companies that seem best positioned are those with some combination of: e user bases they can upsell AI features to, proprietary clinical or claims data that improves model performance, deep domain expertise in a specific specialty or workflow, or strong distribution partnerships with health systems or payers. Pu play AI startups without those advantages will need to move very fast and land enterprise customers before the incumbents catch up.

The Infrastructure Build Required

The policy section of the OpenAI report is basically a wishlist for infrastructure does not currently exist but would dramatically accelerate AI applications in healthcare. Opening and securely connecting medical data, building out clinical and biomanufacturing capacity, creating workforce transition programs, clarifyi FDA regulatory pathways.

The medical data access piece is both the most important and the hardest. AI m trained on large, diverse datasets perform better. Healthcare data is famously fragmented, siloed, and difficult to access. Even data that is theoretically availab like Medicare claims or EHR data from academic medical centers, is often locke due to privacy concerns, institutional policies, or lack of technical infrastructur

The report calls for opening publicly funded medical data with strong privacy protections and creating incentives for private entities to contribute de-identifie datasets. This is directionally correct but wildly optimistic about the timeline ar feasibility. HIPAA, state privacy laws, institutional review boards, data use agreements, and general institutional risk aversion make healthcare data sharing extremely difficult. Every major health system and payer has data that would be valuable for AI training, and almost none of them are set up to share it at scale.

There are startups working on this problem. Datavant, Truveta, HealthVerity, ar others have built data infrastructure to link and de-identify healthcare data acrc sources. Researchers are experimenting with federated learning approaches whe models train on data without moving it out of secure environments. Privacy-enhancing technologies like differential privacy and homomorphic encryption a getting better but still have limitations. Progress is happening, but it is slow anc expensive.

The clinical trial infrastructure point is similarly important. AI can generate hypotheses and identify drug candidates much faster than traditional methods, l validating those candidates still requires running actual clinical trials with real patients. Trial capacity is constrained, especially for rare diseases or complex interventions. Modernizing trial infrastructure, decentralized trials, adaptive de better patient matching would all help, but these changes require coordination a regulators, sponsors, sites, and patients. Not impossible, but not fast either.

The biomanufacturing capacity question is about turning AI-designed molecule: actual drugs. You need wet labs, automation, scale-up facilities, skilled technici quality control systems. Building this out requires capital investment and time. report profiles Junevity, a biotech using AI to identify transcription factor targe Parkinson's and metabolic disease, hitting development milestones 2-3x faster a 6x cheaper than industry norms. That is impressive, but it still requires actual la work and animal studies before anything reaches patients.

For investors, the infrastructure theme suggests opportunities in picks-and-sho companies. Data infrastructure platforms, lab automation, clinical trial tech, biomanufacturing capacity. These are less sexy than drug discovery or diagnostic but arguably more durable businesses because they benefit from multiple downs applications rather than betting on a single therapeutic program.

Scientific Discovery Applications and Time Horizons

The life sciences section of the report profiles three examples: a physician explor drug repurposing for rare food allergies, a biotech founder using AI to accelerat reprogramming research, and a neuroscience PhD candidate using AI to analyze synaptic data. These are meant to illustrate the range of scientific applications, f clinical hypothesis generation to preclinical research to basic neuroscience.

The drug repurposing case is interesting because it is relatively near-term. The physician fed GPT-5 Pro a de-identified patient vignette and the model correctly ranked dupilumab as a potential treatment for food protein-induced enterocolit syndrome, a rare food allergy. The hypothesis was based on linking the drug's approved use for eczema to mucosal immunity in the gut. This is exactly the kin connection that AI might surface faster than human researchers combing throug literature.

The challenge is that surfacing a hypothesis is very different from proving it wor The physician published a case series of seven patients showing improved food a responses while on dupilumab, but that is still just case reports. To actually get dupilumab approved for this indication would require randomized controlled tr which are expensive and time-consuming. So AI might accelerate the hypothesis generation phase from years to weeks, but the validation phase still takes years.

The Junevity example is further out. They are using AI to identify transcription factors for cell reprogramming, with programs targeting type 2 diabetes, obesity Parkinson's. They claim 2-3x faster timelines and 2-6x lower costs than industry norms. Even with that acceleration, they are still in animal studies, which means are at least 5-10 years from having an approved therapy on the market. AI helps, does not eliminate the need for preclinical and clinical validation.

The neuroscience PhD example is more about research productivity than near-te
clinical applications. Using ChatGPT to generate code for data analysis and 3D
reconstruction workflows saved weeks of debugging time and allowed the resear
to focus on experimental design and interpretation. That is a real productivity g
but it is contributing to basic science research that might inform future therapi
decades from now.

The overall theme is that AI is most useful as an accelerant for existing research
development processes rather than a replacement. It helps researchers move fast
consider more hypotheses, analyze more data. But the fundamental constraints c
biology, safety testing, regulatory approval, and clinical validation remain. Expec
AI to deliver new drugs in 1-2 years instead of 10-15 is probably unrealistic. Exp
AI to make the 10-15 year timeline meaningfully faster and cheaper is reasonabl

For investors, this means adjusting expectations about time horizons and risk. A
enabled drug discovery is still drug discovery, with all the attendant risks of clin
failure. The advantage is faster iteration and potentially better target selection, l
the base rates of success in biotech have not fundamentally changed yet. Early-st
investments in AI-native biotech companies should be underwritten with the sa
assumptions as traditional biotech: long timelines, high technical risk, binary
outcomes.

Policy as Product Roadblock or Accelerant

The regulatory section of the report focuses on FDA pathways for AI medical de
both for consumer use and for clinical decision support tools. The core argumen
that the existing regulatory framework was designed for static medical devices a
does not accommodate AI systems that learn and adapt over time.

FDA has acknowledged this problem and has been working on it for several year
Digital Health Center of Excellence, Software as a Medical Device guidance
documents, and various pilot programs are all attempts to create a more flexible
regulatory approach. The challenge is that FDA is fundamentally a risk-averse

organization, and AI introduces new types of risk that are difficult to assess witł traditional methods.

For consumer-facing AI health tools, the question is what level of evidence FDA should require before allowing widespread use. Traditional medical devices go through rigorous testing, often including randomized controlled trials, before reaching the market. AI tools trained on large datasets might have different risk profiles and might benefit from post-market surveillance rather than extensive p market testing. FDA is trying to figure out where to draw those lines.

For clinical decision support tools used by physicians, the 21st Century Cures A created exemptions for certain types of software that support clinical judgment not replace it. The question is whether modern AI tools fit within those exempti represent a different category that requires oversight. The report argues that FD should clarify this to avoid stifling innovation.

The practical implication is that regulatory uncertainty creates market risk for companies building AI health products. If you are three years into developing a diagnostic AI tool and FDA suddenly decides it requires a Class III device pathw with full clinical trials, your timeline and budget just exploded. Conversely, if F creates a streamlined pathway for certain types of AI tools, that reduces time to market and capital requirements.

For investors, regulatory risk is nothing new in healthcare. Every digital health investment has some regulatory angle. The difference with AI is that the rules ar being written, so predicting outcomes is harder. Companies with experienced regulatory teams, strong clinical advisors, and willingness to engage early with F are better positioned. Companies trying to move fast and ignore regulatory ques are setting themselves up for expensive surprises.

The other policy angle is reimbursement. Even if FDA approves an AI diagnosti someone needs to pay for it. That means either getting a CPT code for the specii test, bundling it into existing care pathways, or building a direct-to-consumer m Each of those paths has its own challenges. CPT codes take years to establish an require extensive clinical evidence. Bundling into existing pathways requires convincing health systems to adopt without incremental payment. Direct-to-con requires marketing spend and consumer willingness to pay out of pocket.

Investment Implications and Market Sizing

The OpenAI data suggests a few investment themes worth paying attention to.

First, consumer health AI is real. The usage numbers are too large to ignore. Companies building consumer-facing health AI tools have a legitimate shot at sc especially if they can differentiate from free general-purpose options. The wedg probably condition-specific applications, language and cultural customization, c integration with other health data sources. Fertility, metabolic health, mental he chronic disease management all seem like plausible categories.

Second, clinical workflow automation is the near-term enterprise opportunity. A scribes, documentation tools, care coordination assistants. These have clear ROI relatively low regulatory risk, and strong product-market fit based on the adopti data. Expect continued investment and M&A activity in this category. The big question is whether ambient documentation becomes table stakes for every EMI within 2-3 years, in which case the standalone companies need to sell or expand adjacent products quickly.

Third, insurance navigation and administrative automation is underinvested rel to the pain point and market size. Most of the innovation here has been on the provider side with revenue cycle management tools. The patient-facing navigatic opportunity is large and mostly unaddressed. Startups that can embed AI naviga into broker platforms, benefits admin systems, or employer benefits packages ha clear path to revenue.

Fourth, rural and underserved markets are accessible via AI tools in ways they w not before. Companies building for these markets need to design for low technic sophistication, limited connectivity, and resource constraints. But if the product works, the addressable market is massive and relatively underserved by existing digital health companies.

Fifth, life sciences infrastructure and tools are picks-and-shovels plays with dur value. Data platforms, lab automation, clinical trial tech, biomanufacturing. The benefit from multiple downstream applications and do not depend on any single or therapeutic program succeeding.

Sixth, regulatory and reimbursement risk remains high for anything touching cl decision-making. Diagnostic AI, treatment recommendation engines, drug disco platforms all have long timelines and uncertain paths to market. Investors need underwrite these with appropriate risk adjustments and long time horizons.

The market sizing is tricky because AI is horizontal across multiple healthcare categories rather than a discrete vertical. McKinsey estimated generative AI cou create $200-360B in annual value across healthcare via use cases like clinical documentation, coding, patient engagement, drug discovery, and administrative automation. That is 5-10% of total US healthcare spending, which feels directior plausible if adoption accelerates.

For context, electronic health records are roughly a $30B annual market globally Revenue cycle management is another $20B+. Telemedicine hit $90B+ during C(peak and has settled into sustainable growth. Medical imaging AI is probably $2 currently but growing quickly. If AI-enabled healthcare follows a similar traject those categories, we are talking about tens of billions of dollars in annual revenu across multiple segments within 5-10 years.

The challenge for investors is that much of that value might accrue to large incumbents rather than startups. Epic, Oracle Cerner, Optum, CVS, Elevance, Humana all have the resources and customer bases to build or buy AI capabilitie Pure-play AI health startups need to move fast, establish strong product-market and either reach scale independently or position for acquisition before the wind closes.

The OpenAI report is fundamentally a market development play. They are publi
this data to demonstrate that AI in healthcare is real, happening at scale, and cre
measurable value. That helps them sell more API access, attract more enterprise
customers, and shape policy in favorable directions. But the underlying usage da
credible and the implications are worth taking seriously. Healthcare AI has cros
the chasm from pilot programs to mainstream adoption, at least for certain use c
Investors and operators need to adjust their mental models accordingly.

5 Likes

Discussion about this post

Comments Restacks

Write a comment...

© 2026 Thoughts on Healthcare · Privacy • Terms • Collection notice
Substack is the home for great culture