ChatGPT Health and Why Foundation Models Still Cannot Crack Healthcare Alone

JAN 09, 2026 • PAID

Abstract

OpenAI launched ChatGPT Health in January 2025, partnering with Color Heal bwell Connected Health while collaborating with major health systems to build domain-specific clinical AI. The move signals that even the most sophisticated foundation model companies recognize healthcare requires specialized training, regulatory expertise, and clinical validation beyond what general-purpose LLMs provide. Key points include:

- ChatGPT Health targets clinical documentation, patient communication, and c coordination through partnerships rather than direct deployment

- Color Health brings clinical validation capabilities while bwell Connected Hea provides consumer engagement infrastructure and health plan relationships

- Foundation model competitors including Anthropic, xAI, and Google Gemini f strategic choices about whether to pursue similar healthcare-specific products o enable partners to build on their platforms

- The partnership model validates that healthcare AI requires domain expertise, regulatory positioning, and distribution channels that foundation model compar lack internally

- Market opportunities exist across multiple layers of the stack for specialized companies that can navigate healthcare's unique requirements

Table of Contents

Why General Purpose Models Keep Failing in Clinical Settings

The past two years saw hundreds of health systems and digital health companies attempt to deploy ChatGPT and similar foundation models for clinical use cases. Most of these initiatives failed or got stuck in pilot purgatory because general-purpose language models lack the reliability, specificity, and regulatory position required for actual clinical deployment. The gap between passing medical licens exams in controlled settings and functioning safely within real clinical workflow proved far wider than the initial hype suggested. Health system CIOs who jumpe

foundation models in 2023 spent 2024 explaining to their boards why pilot proje did not scale despite impressive demos.

The failure pattern repeated across organizations. Initial pilots showed promisir results with cherry-picked use cases and heavy human oversight. Attempts to sca beyond controlled environments exposed problems with hallucinations, inconsis outputs, inability to handle edge cases, and lack of integration with existing workflows. Legal and compliance teams raised concerns about liability exposure regulatory ambiguity, and patient safety risks that pilot coordinators had not adequately addressed. The technology worked well enough to be interesting but well enough to be deployable at scale without significant additional developmen

OpenAI apparently learned these lessons through their own partnership discuss and pilot failures. ChatGPT Health exists as a distinct product line precisely bec the company recognized that healthcare requires fundamentally different architecture, training approaches, and go-to-market strategies compared to gene purpose AI. The decision to partner with domain experts like Color Health and I Connected Health rather than attempting direct deployment acknowledges that foundation model companies lack the clinical expertise, regulatory knowledge, a healthcare distribution channels needed to succeed independently.

This matters enormously for health tech investors and entrepreneurs because it validates several hypotheses about how AI will actually penetrate healthcare. Fir domain specificity beats generalization in clinical settings. Models need training clinical datasets, fine-tuning for specific workflows, and validation against real-v use cases rather than academic benchmarks. Second, partnerships between AI companies and healthcare domain experts will dominate the market rather than plays by foundation model providers. Third, the technical challenge of building models represents only part of the problem, with regulatory navigation, liability frameworks, and workflow integration often proving more difficult than model development itself.

What OpenAI Actually Built and the Partnership Strategy Behind It

ChatGPT Health is not just GPT-4 with some medical textbooks in the training The product involves specialized fine-tuning on clinical datasets, integration wit healthcare-specific standards like SMART on FHIR, and partnerships designed t provide clinical validation and regulatory expertise that OpenAI lacks internally company positioned the offering around several core use cases including clinical documentation, patient communication, care coordination, and administrative automation. Each use case requires different levels of clinical risk tolerance, regulatory oversight, and workflow integration, creating a portfolio approach ra than a single product.

The partnership with Color Health provides OpenAI with clinical credibility an validation capabilities that would take years to build internally. Color operates clinical genomics programs and health screening services that give them deep expertise in regulatory compliance, clinical validation, and quality assurance. Th experience navigating FDA oversight for genetic testing and working with healt systems on population health programs positions them to guide ChatGPT Healt through similar regulatory and clinical validation processes. Color also brings relationships with health systems and payers that provide deployment pathways access to real-world clinical data for training and testing.

The collaboration with bwell Connected Health adds a different dimension focu consumer engagement and health plan relationships. bwell built a platform that aggregates health data from multiple sources including EHRs, claims systems, a wearables to provide consumers with unified access to their health information. technology infrastructure handles the complex data integration, identity resolut and consent management required to connect disparate healthcare data sources. ChatGPT Health, bwell provides the pipes and infrastructure to enable patient- AI applications that can access comprehensive health histories while navigating privacy requirements and data access restrictions.

The bwell partnership matters particularly for consumer-facing use cases like pa communication, care navigation, and health information access. Health plans an employers increasingly want tools that engage members outside traditional clini settings to improve adherence, reduce avoidable utilization, and enhance satisfa However, building effective consumer health applications requires solving for da access, identity management, and consent workflows that are completely differe from clinical applications. bwell spent years building this infrastructure and establishing relationships with health plans and employers who can distribute consumer health applications at scale.

OpenAI also announced collaborations with health systems including Kaiser Permanente and Cleveland Clinic that provide access to clinical data for validati and deployment pathways for testing products in real care settings. These partnerships follow a familiar pattern in healthcare AI where foundation model companies partner with prestigious academic medical centers to gain credibility access to data. The health systems benefit from early access to cutting-edge technology and potential competitive advantages if the tools prove effective. Hov these collaborations often produce more press releases than actual deployed pro because navigating institutional review boards, privacy requirements, and risk management processes takes far longer than the partnership announcements sug

The multi-partner strategy acknowledges that no single organization possesses a capabilities required to deploy clinical AI successfully. OpenAI brings the founc models and brand recognition. Color provides clinical validation and regulatory expertise. bwell contributes consumer engagement infrastructure and health pla distribution. Health systems offer deployment environments and clinical data. E partner fills gaps that the others cannot address independently, creating a compl web of relationships that will either prove synergistic or create coordination challenges that slow progress.

Color Health's Role in Clinical Validatio and Regulatory Navigation

Color Health's involvement in ChatGPT Health extends beyond a typical techno partnership because Color brings regulatory and clinical validation expertise tha OpenAI completely lacks. Color spent years building relationships with FDA an establishing processes for clinical validation of genetic tests and population hea screening programs. Their experience navigating 510(k) submissions, clinical val studies, and quality management systems provides a playbook for bringing AI-en clinical tools to market in a compliant manner.

The clinical validation piece matters more than most people realize because healthcare providers will not adopt AI tools without rigorous evidence of safety efficacy. Academic publications, benchmark performance metrics, and pilot resu help but do not substitute for formal validation studies that measure real-world outcomes in representative patient populations. Color knows how to design thes studies, work with institutional review boards, recruit representative patient col and generate evidence that satisfies skeptical clinicians and risk-averse hospital committees. OpenAI could attempt to build these capabilities internally but doi would require hiring dozens of clinical research professionals and spending yea learning healthcare's unique evidence generation requirements.

Color's existing relationships with health systems and payers also provide crucia distribution channels for ChatGPT Health. Color's genomics and screening prog already operate within dozens of health systems, creating established procureme relationships, contracting frameworks, and technical integrations that new vend struggle to replicate. Rather than starting from zero with health system sales cyc that stretch twelve to twenty-four months, ChatGPT Health can potentially leve Color's existing footprint to accelerate deployment. Whether this actually happe depends on execution and whether the AI products deliver sufficient value to jus expansion beyond Color's current genomics-focused relationships.

The regulatory strategy Color brings becomes particularly important as FDA an other agencies develop frameworks for overseeing clinical AI. FDA categorizes s AI applications as medical devices requiring premarket review while excluding c based on factors like whether the tool provides diagnostic conclusions versus information for clinicians to interpret. Determining which side of this line a

particular AI application falls requires deep regulatory expertise and often invol preliminary discussions with FDA to gauge their likely interpretation. Color has expertise and these relationships, reducing the risk that ChatGPT Health pursu cases that later get categorized as requiring extensive FDA review and approval.

Color also understands the operational realities of deploying clinical tools withi health systems in ways that pure technology companies do not. Clinical workflo vary enormously across organizations based on EHR platforms, staffing models, patient populations, and institutional culture. Products that work perfectly at academic medical centers with tech-savvy clinicians and robust IT support often at community hospitals with older systems and limited resources. Color's experi deploying genomics programs across diverse care settings provides insights into operational requirements for successful deployment that will inform ChatGPT Health's product development and implementation approach.

bwell Connected Health and the Consumer Engagement Infrastructure Play

bwell Connected Health's role in ChatGPT Health focuses on enabling consume facing applications through their health data aggregation platform and health pl relationships. bwell built infrastructure that connects to hundreds of data sourc including EHRs, claims systems, pharmacies, labs, and wearables to create unifie health records accessible to consumers. The technical complexity of this data aggregation involves solving for incompatible data formats, inconsistent identifi varying consent frameworks, and API limitations across hundreds of source syst bwell spent years building and maintaining these integrations, creating a defens moat that would be extremely difficult and expensive for OpenAI to replicate independently.

The consumer health data aggregation market has seen significant investment ar competition over the past five years with companies like Human API, Particle H and Ciox Health pursuing similar strategies. bwell differentiated by focusing on

health plan relationships and building a consumer-facing experience layer on to
the data aggregation infrastructure. Health plans need tools to engage members
various purposes including care gap closure, utilization management, chronic di
management, and member retention. bwell's platform enables health plans to of
members unified access to their health data while layering in plan-specific conte
outreach, and engagement features.

For ChatGPT Health, bwell provides the infrastructure to enable AI-powered pa
interactions that can reference comprehensive health histories while navigating
complex privacy and consent requirements. Effective patient-facing health AI re
access to clinical data, claims history, medication lists, and lab results to provide
relevant and personalized responses. However, accessing this data requires navig
HIPAA compliance, individual consent workflows, and technical integration wit
multiple source systems. bwell already solved these problems, allowing ChatGP
Health to focus on the AI interaction layer rather than rebuilding data infrastru
from scratch.

The health plan distribution channel bwell brings also matters significantly for
consumer health applications. Health plans reach over 250 million Americans th
commercial insurance, Medicare Advantage, and Medicaid managed care. Plans
actively seek tools that improve member engagement, HEDIS scores, and Star Ra
while reducing avoidable utilization. AI-powered care navigation, symptom
assessment, and health information tools align well with health plan priorities if
can demonstrate measurable impact on utilization and quality metrics. bwell's
existing health plan relationships provide pathways to deploy ChatGPT Health
consumer applications at scale far faster than building health plan sales channel
from scratch.

The consumer engagement use cases also present different regulatory and liabili
profiles compared to clinical decision support tools. Patient-facing applications
provide health information, answer questions about coverage and benefits, or he
navigate the healthcare system generally fall outside FDA regulation as long as t
avoid providing diagnostic or treatment recommendations. This regulatory
positioning allows faster deployment and iteration compared to clinical tools tha

require FDA review. However, the liability risks around providing incorrect or misleading health information to patients remain significant, requiring careful product design, appropriate disclaimers, and probably some form of liability protection for deploying health plans.

bwell's involvement signals that OpenAI views consumer engagement and patien facing applications as equally important to clinical workflows in their healthcare strategy. This makes sense given the massive inefficiencies in how consumers int with the healthcare system and the alignment with health plan incentives to imp engagement and reduce avoidable utilization. However, consumer health AI face different challenges around driving adoption, demonstrating value, and changin behavior compared to clinician-facing tools. Many consumer health apps achieve engagement despite sophisticated features because they fail to integrate into dai routines or provide compelling enough value to drive sustained usage.

The Economics of Healthcare AI and Where Value Actually Accrues

Understanding where money gets made in healthcare AI requires examining pay flows and incentive structures across different stakeholders. Health systems pay tools that reduce costs, increase revenue, or improve quality metrics that affect reimbursement. Payers invest in technologies that reduce medical costs, improve quality scores, or enhance member retention. Patients rarely pay directly for healthcare AI tools, limiting direct-to-consumer business models. These paymen dynamics determine which use cases attract investment and which struggle to fi sustainable business models despite strong clinical value propositions.

Clinical documentation represents the most economically straightforward use ca because it addresses a clear pain point with measurable ROI. Physicians spend approximately two hours on documentation for every hour of direct patient care according to multiple time motion studies. This documentation burden contribu physician burnout while creating opportunity costs where physicians could see additional patients or improve care quality. AI-powered documentation tools tha

reduce documentation time by even thirty to forty percent generate clear econor value through some combination of increased physician capacity, reduced burno and turnover, and improved work-life balance.

The willingness to pay for documentation tools is well-established with health systems already spending on medical scribes, transcription services, and physici time dedicated to after-hours charting. Companies like Abridge, Ambience Healthcare, Suki, and Nuance DAX compete actively in this market with produc generating real revenue and demonstrating clear ROI. Market sizing suggests se billion dollars in addressable spend just for clinical documentation, with room f multiple winners as different products serve different specialties, care settings, a workflow preferences. OpenAI enters a competitive but validated market where product differentiation and performance determine market share.

Patient engagement and communication tools present larger total addressable markets but less proven willingness to pay and ROI measurement challenges. He plans theoretically should pay significant amounts for tools that improve membe engagement, close care gaps, and reduce avoidable utilization. However, measuri and attributing these outcomes to specific interventions remains difficult given t many factors influencing patient behavior and health outcomes. Randomized controlled trials can establish causal impact but require large sample sizes, long follow-up periods, and substantial investment. Health plans often prefer pilot programs with short timelines and quick wins, creating tension between rigorou evidence generation and commercial traction.

The revenue models for patient-facing applications also face challenges because health plans make procurement decisions but members experience the products members find the tools annoying, unhelpful, or intrusive, they will not engage regardless of how much the health plan paid. This creates principal-agent probl where the buyer and user have misaligned incentives. Successful consumer healt applications need to deliver value to both health plans and members, requiring product design that balances plan objectives like care gap closure with member like convenience and relevance. Many health plan-funded consumer health tools

achieve low engagement because they prioritize plan objectives over member experience.

Administrative automation including prior authorization, claims processing, scheduling, and eligibility verification represents enormous addressable markets clear economic waste. The American Medical Association estimates that prior authorization alone costs the healthcare system over 35 billion dollars annually i administrative overhead. However, these markets feature entrenched incumbent solutions from revenue cycle management vendors, clearinghouses, and business process outsourcing companies. Displacing existing solutions requires demonstr not just effectiveness but sufficient performance improvement to justify switchi costs and integration effort. OpenAI did not emphasize administrative automati initial ChatGPT Health announcements, possibly recognizing that competing in markets requires different capabilities than clinical AI applications.

The unit economics of healthcare AI also vary significantly across deployment m Cloud-based APIs accessed on a per-query basis create recurring revenue but fa pricing pressure as compute costs decline and competitors emerge. Licensed sof deployed on health system infrastructure generates higher revenue per customer requires more complex implementation and support. Outcome-based pricing tie measured improvements in efficiency or quality aligns incentives but requires sophisticated analytics and longer sales cycles. The optimal pricing model deper the specific use case, competitive dynamics, and customer preferences.

How Anthropic Should Respond to Stay Competitive

Anthropic faces strategic choices about whether to pursue healthcare-specific products like ChatGPT Health or focus on enabling partners to build healthcare applications on Claude. The partnership approach aligns better with Anthropic' stated emphasis on AI safety and responsible deployment because healthcare pa bring domain expertise and regulatory knowledge that Anthropic lacks. Howeve

ceding the healthcare market entirely to OpenAI and potential Google Gemini h
initiatives risks missing one of the largest and most defensible enterprise AI ma

The most likely successful strategy for Anthropic involves partnering with healt
domain experts and technology companies to build specialized applications on C
while maintaining platform flexibility. This approach leverages Anthropic's stre
in model quality and safety while avoiding the need to build healthcare-specific
capabilities internally. Potential partners could include EHR vendors looking to
enhance their platforms with advanced AI, specialized healthcare AI companies
wanting alternatives to OpenAI APIs, or healthcare delivery organizations build
proprietary tools. The key is identifying partners with complementary capabiliti
aligned incentives.

Anthropic should particularly target use cases where Claude's extended context
windows and stronger performance on complex reasoning provide competitive
advantages. Medical decision support involving synthesis of extensive patient
histories, research literature, and clinical guidelines could benefit from Claude's
ability to process longer inputs without losing coherence. Similarly, clinical rese
applications involving analysis of multiple studies or synthesis of complex datas
might favor Claude's architecture. Identifying these differentiated use cases and
partnering with domain experts to build specialized solutions would create defe
positions without requiring Anthropic to become a healthcare company.

The regulatory and safety positioning Anthropic emphasizes publicly also reson
with healthcare's risk-averse culture. Health systems and payers worry extensive
about AI safety, reliability, and unintended consequences. Anthropic's focus on
constitutional AI, interpretability, and responsible deployment aligns well with
healthcare stakeholder concerns and could provide competitive differentiation v
OpenAI's more aggressive deployment approach. Marketing Claude as the safer,
responsible choice for healthcare applications could attract customers prioritizi
risk management over cutting-edge capabilities.

Anthropic should also consider vertical integration through acquisition or
partnership with healthcare AI companies that have already achieved product-m

fit and revenue traction. Acquiring a company like Abridge or Ambience would provide immediate healthcare market presence, clinical validation, and custome relationships while bringing domain expertise in-house. This accelerated entry strategy avoids the years of relationship building and regulatory learning requir build healthcare capabilities organically. However, acquisition multiples for successful healthcare AI companies probably exceed what Anthropic would wan pay, and integration challenges could slow progress.

The developer ecosystem approach represents another viable strategy where Anthropic focuses on making Claude the best platform for healthcare developer build on rather than building healthcare products directly. This requires investi healthcare-specific documentation, reference architectures, compliance framew and integration libraries that reduce friction for healthcare developers. Creating vibrant ecosystem of healthcare applications built on Claude would generate AP revenue and platform stickiness without requiring Anthropic to become a healt expert. However, this approach cedes direct customer relationships and risks pa building competitive moats that reduce Anthropic's leverage over time.

What xAI Needs to Do If Musk Wants Healthcare Revenue

xAI faces the most challenging position among foundation model companies attempting to enter healthcare because the company lacks partnerships, healthc expertise, and the institutional credibility that healthcare stakeholders require. Musk's personal brand cuts both ways in healthcare, attracting some innovators early adopters while creating skepticism among risk-averse hospital administrat and regulators who worry about unpredictability and controversial statements. Building credible healthcare offerings requires xAI to overcome these perceptio challenges while developing domain expertise and regulatory positioning from scratch.

The most realistic path for xAI involves focusing on healthcare use cases where traditional stakeholder concerns matter less and where speed and innovation tru

institutional credibility. Consumer-facing health information applications, wellı and prevention tools, and health coaching represent categories where direct-to-consumer distribution and rapid iteration matter more than health system relationships and regulatory positioning. Musk's experience building consumer brands and his existing Twitter/X platform provide distribution advantages that accelerate consumer health product adoption despite lack of healthcare industry relationships.

xAI could also target pharmaceutical and biotechnology companies rather than healthcare providers and payers. Drug discovery, clinical trial optimization, regu submission preparation, and pharmacovigilance represent substantial markets w AI applications have demonstrated value. Pharma and biotech companies genera move faster than health systems, have higher risk tolerance for emerging techno and care more about capability and performance than institutional credibility. T characteristics align better with xAI's strengths and Musk's brand compared to t conservative, relationship-driven dynamics of health system sales.

The Twitter/X integration also creates unique opportunities for population heal surveillance, public health communication, and health-related social listening th other foundation model company can replicate. Analysis of health-related social media conversations could identify emerging disease outbreaks, track sentiment healthcare topics, and inform public health interventions. However, privacy con and the controversial nature of using social media data for health surveillance w require careful positioning and probably partnerships with public health agenci establish legitimacy.

xAI's disadvantages in healthcare are significant enough that pursuing traditiona clinical AI use cases like documentation and clinical decision support probably little sense. OpenAI, Anthropic, and Google all have better positioning, more resources, and stronger relationships in these categories. xAI attempting to com directly would waste resources chasing markets where others have insurmountal advantages. Instead, focusing on differentiated use cases where xAI's unique capabilities and distribution channels create competitive advantages represents only viable path to meaningful healthcare revenue.

The regulatory challenges also loom larger for xAI than other foundation model companies because Musk's public statements and management approach create unpredictability that risk-averse regulators and healthcare institutions dislike. Building relationships with FDA, working through clinical validation studies, a establishing processes that satisfy healthcare compliance requirements all requi patience, attention to detail, and consistent messaging that do not align with Mu typical operating style. xAI would need to create organizational separation and governance structures that insulate healthcare operations from the volatility of M other ventures.

Google Gemini's Advantages and Strategic Missteps in Healthcare

Google possesses structural advantages in healthcare AI that exceed those of any foundation model company, yet the company has repeatedly failed to convert the advantages into market leadership. Google's healthcare assets include relationsh with health systems through Google Cloud, consumer health data through Fitbi Google Health, AI capabilities through DeepMind and Google Research, and distribution through Android and Google Search. Despite these advantages, Goo has yet to establish the market presence in healthcare AI that its capabilities wo suggest possible.

The Med-PaLM initiative represented Google's most visible healthcare AI effort producing strong benchmark performance and publications in academic journal However, Med-PaLM remained primarily a research project rather than a commercialized product, illustrating Google's pattern of building impressive technology without converting it to market impact. The gap between research excellence and commercial deployment reflects organizational challenges includ risk aversion, difficulty making product decisions, and internal competition for resources between different groups working on healthcare.

Google Cloud's healthcare relationships provide distribution advantages if the company can figure out how to leverage them effectively. Google Cloud compete

health system cloud infrastructure and analytics workloads, creating existing procurement relationships and technical integrations. However, health systems c separate cloud infrastructure decisions from clinical application purchases, limi Google Cloud's ability to cross-sell AI capabilities. Additionally, healthcare organizations worry about Google's consumer data practices and search busines creating conflicts of interest around patient data, requiring careful separation a governance to build trust.

The Fitbit acquisition gave Google access to millions of consumers' health and activity data, creating training datasets for consumer health AI that competitors However, Google has done little visible with this asset, missing opportunities to consumer health applications that leverage both Fitbit data and Gemini capabili A well-designed AI health coach that combines activity tracking, personalized recommendations, and conversational interaction could compete effectively in t digital health market. Google's failure to build this despite having all necessary components illustrates the organizational dysfunction that has plagued their healthcare efforts.

Google's strategic advantages also include relationships with Android device manufacturers and the Google Play ecosystem that provide distribution for cons health applications. Over 70 percent of smartphones globally run Android, creat potential reach that exceeds any other foundation model company. However, Go has struggled to monetize Android beyond advertising, and healthcare applicati would require different business models and capabilities. Building effective con health products also requires understanding behavior change, clinical workflow healthcare economics beyond what Google's product teams have demonstrated historically.

The optimal strategy for Google in healthcare AI involves focusing on areas whe their unique advantages create defensible positions. Consumer health applicatio leveraging Fitbit data and Android distribution represent one such area. Health provider applications built into Google Cloud and Workspace tools represent a Clinical research and drug discovery applications leveraging Google's computati biology expertise through DeepMind represent a third. However, executing on t

opportunities requires organizational commitment, willingness to invest throug development cycles, and product discipline that Google has struggled to demons in healthcare.

Google also needs to address the trust deficit many healthcare stakeholders feel toward the company. Multiple failed healthcare initiatives including Google Hea shutdown, the controversial Ascension partnership, and various privacy controv created skepticism about Google's commitment and trustworthiness in healthca Rebuilding trust requires consistent engagement, transparent data practices, an following through on commitments over multiple years. OpenAI does not carry t baggage, giving them advantages despite Google's superior resources and techni capabilities.

The Infrastructure Problem That Determines Everything

Healthcare AI deployments fail more often due to infrastructure and integration challenges than model quality problems. Most health systems run on complex El platforms with proprietary data models, limited API capabilities, and institutior resistance to sharing data with external systems. Getting AI tools to function wi these environments requires solving technical integration challenges that have nothing to do with model capabilities. ChatGPT Health's SMART on FHIR integration helps somewhat by providing standardized API access, but this only addresses part of the problem.

The reality is that clinical workflows vary enormously across organizations base EHR platforms, staffing models, patient populations, and institutional practices Products that work well at academic medical centers with tech-savvy clinicians a modern EHR instances often fail at community hospitals running older systems minimal customization. Building AI tools that work reliably across this diversity requires extensive configuration options, flexible integration capabilities, and probably some amount of custom development for each deployment. This reality

favors companies with strong implementation and support capabilities over thos
with the best models.

Data quality and standardization present even larger challenges than technical integration. Clinical data exists in semi-structured formats with tremendous variability in how different organizations document similar information. Medica lists use different naming conventions, lab values have varying normal ranges, ar clinical notes follow different templates across organizations. Models trained primarily on data from a few large health systems may perform poorly when exp to data from different care settings, requiring ongoing fine-tuning and validatio increases operational complexity and costs.

The compute and hosting requirements for running large language models at sc also create practical deployment barriers. Some health systems refuse to send pa data to external cloud environments regardless of security guarantees, requiring deployment or on-premise hosting. These deployment models increase complexi costs substantially compared to cloud-based APIs. OpenAI's approach appears focused on cloud-based deployment, potentially limiting adoption among the m security-conscious organizations. Companies that can offer flexible deployment options including on-premise, private cloud, and public cloud will have advanta with security-focused customers.

Identity management and patient matching present additional infrastructure challenges for any AI application that needs to access patient data across multip systems. Healthcare lacks universal patient identifiers, forcing organizations to patients across systems using combinations of name, date of birth, address, and demographics. Match rates typically run 85 to 95 percent even with sophisticate algorithms, meaning 5 to 15 percent of patient records cannot be reliably linked across systems. This creates gaps in the data available to AI applications and pot safety issues if medications or allergies documented in unmatched records get m

The infrastructure challenges ultimately mean that technology capability repres only one component of success in healthcare AI markets. Companies need implementation capabilities, integration expertise, and ongoing support that re

different organizational capabilities than building great models. This favors companies combining technical and operational excellence or partnerships betw foundation model companies and healthcare domain experts with implementatic experience. Pure-play model companies attempting to deploy healthcare AI dire will struggle with these operational realities unless they build substantial servic capabilities.

Competition at Every Layer and Why Vertical Wins

The healthcare AI market features competition at multiple layers including foundation models, specialized medical AI models, application-layer products, a workflow integration platforms. Each layer has different competitive dynamics, different key success factors, and different opportunities for defensibility. Understanding these dynamics helps determine where to invest and what strateg create sustainable competitive advantages.

The foundation model layer will likely remain competitive with multiple viable players rather than winner-take-all dynamics. Healthcare organizations want mu model options to avoid vendor lock-in, support use-case-specific optimization, a maintain competitive tension on pricing. OpenAI, Anthropic, Google, and poten other foundation model companies will all find healthcare customers willing to u their models. However, the foundation model layer faces commoditization pressu performance converges and switching costs remain low for many use cases. Pure infrastructure plays without application-layer differentiation will struggle to cap much of the value created by healthcare AI.

Specialized medical AI models represent a middle layer where companies fine-tu foundation models or build custom architectures for specific clinical domains. Companies working on radiology interpretation, pathology analysis, clinical tria matching, or drug discovery often build specialized models trained on domain-sp datasets. These companies can create defensible positions through proprietary training data, domain expertise, and validation evidence that general-purpose m

cannot easily replicate. However, the specialized model layer also faces challenge foundation models improve and the gap between general-purpose and specialize models narrows.

The application layer where companies build complete products for specific workflows and use cases has the most opportunity for value capture. Health syst pay for outcomes and workflow improvements, not for access to models. Compa that demonstrate measurable ROI through reduced documentation time, improv coding accuracy, or better patient engagement can command premium pricing w strong retention. Application-layer companies also build switching costs throug workflow integration, training, and customization that create defensibility beyo model capabilities.

Vertical specialization consistently beats horizontal generalization in healthcare markets. Companies building AI for specific clinical workflows, conditions, or c settings outperform those attempting to address broad use cases. The depth of domain expertise, quality of training data, and workflow optimization possible w vertical focus creates superior products that health systems value more highly. T dynamic suggests market structures with dozens of specialized vendors rather th consolidation around a few horizontal platforms.

The EHR vendors including Epic, Oracle Cerner, and Meditech represent a uniq competitive layer because they control access to clinical data and workflow integration points. Epic's AI strategy involves partnerships with multiple founda model providers while maintaining control over user experience and integration positions Epic as a platform layer that captures value from AI innovation while maintaining customer lock-in. Third-party AI vendors need to offer sufficient differentiation to justify the friction of adding another vendor relationship, or th need to work through Epic's partnership programs accepting lower margins in exchange for distribution.

Smaller startups focused on specific clinical use cases often create the most valu users despite lacking the resources and brand recognition of larger competitors. Companies building AI for specific specialties, workflows, or patient population

optimize their products in ways that general-purpose tools cannot match. Dermatology AI that understands the specific documentation requirements and workflow patterns of dermatology practices will outperform general clinical documentation tools adapted for dermatology. This specialization creates opportunities for dozens of focused companies even as larger players dominate headlines.

What This Means for Health Tech Build and Investors

The ChatGPT Health announcement and surrounding competitive dynamics pro useful signals for entrepreneurs and investors building in healthcare AI. Several strategic implications emerge from analyzing OpenAI's approach and the broade market structure.

First, foundation models alone are insufficient for clinical deployment. Compan need domain-specific fine-tuning, regulatory expertise, clinical partnerships, and workflow integration to build credible products. This creates opportunities for application-layer companies that combine foundation model capabilities with healthcare domain expertise. The technical barriers to building on foundation m continue declining through better APIs, documentation, and developer tools. Competitive advantage shifts toward distribution, domain expertise, and ability navigate healthcare's unique requirements rather than raw technical capabilities

Second, partnerships between AI companies and healthcare domain experts will dominate successful go-to-market strategies. Every credible healthcare AI initia announced recently involves partnerships between foundation model companies clinical organizations, domain experts, or healthcare technology companies. Entrepreneurs should prioritize building these relationships early rather than attempting to develop products in isolation. The clinical expertise, regulatory knowledge, and distribution channels that healthcare partnerships provide cann easily replicated through internal development.

Third, vertical specialization creates more defensible businesses than horizontal platforms in most scenarios. Companies should focus on specific clinical workfl specialties, or patient populations where they can develop deep expertise and optimized products. The temptation to pursue large horizontal markets with ge purpose tools leads to mediocre products that lack differentiation and struggle t demonstrate ROI. Better to own a large share of a specific vertical than compete scraps in massive horizontal markets dominated by well-resourced incumbents.

Fourth, go-to-market strategy and distribution determine success as much as pr quality. Healthcare sales cycles are long, buying processes involve multiple stakeholders, and incumbents enjoy strong relationships and switching cost advantages. Companies need strategies to access customers through partnership channel relationships, or novel distribution approaches that bypass traditional enterprise sales. bwell's health plan relationships provide ChatGPT Health with distribution advantages that would take years to build independently. Similar thinking about distribution should inform every healthcare AI company's strate

Fifth, regulatory positioning and liability frameworks require attention from da rather than afterthoughts once products are built. The regulatory pathway and liability exposure associated with different use cases vary enormously. Clinical decision support tools providing diagnostic or treatment recommendations face different regulatory requirements than documentation or administrative automa tools. Entrepreneurs should engage early with regulatory experts and build compliance capabilities alongside products rather than hoping for regulatory cl later.

Sixth, the economic model needs to deliver clear ROI with measurable outcomes justify purchase decisions. Healthcare stakeholders pay for demonstrated value, potential or impressive technology. Companies need to build measurement frameworks and analytics capabilities that prove their products deliver the effici gains, quality improvements, or cost reductions they promise. Pilot programs sh focus on generating this evidence rather than just demonstrating technical feasi

Seventh, infrastructure and integration capabilities matter as much as model qu for successful deployment. Healthcare IT environments are complex with multip systems, inconsistent data, and limited interoperability. Products need flexible integration options, strong implementation support, and ongoing maintenance capabilities. Companies should invest in these operational capabilities even if th seem less exciting than model development.

For investors, the healthcare AI market presents opportunities across multiple s and business models. Early-stage investments in application-layer companies wi strong founding teams, clear use cases, and initial customer traction offer the be risk-adjusted returns. These companies benefit from improving foundation mod while building defensible positions through domain expertise, customer relation and workflow integration. Later-stage investments in companies with proven rev traction and demonstrated ROI can provide lower-risk exposure to market grow Infrastructure plays and horizontal platforms present higher risk given commoditization pressure and competition from well-resourced incumbents.

The market timing also appears favorable with foundation model capabilities no good enough to enable real clinical applications while many use cases remain underserved by existing solutions. The window for startups to build defensible positions before incumbents and large tech companies dominate every category remains open but probably closes over the next several years. Companies moving quickly to establish customer relationships, generate validation evidence, and bu operational capabilities can create advantages that will be difficult for later entr to overcome.

Regulatory Frameworks That Will Make or Break Clinical AI

Regulatory oversight of clinical AI remains poorly defined despite increasing deployment and growing evidence of impact. FDA categorizes some AI applicati medical devices requiring premarket review while excluding others including cl decision support tools meeting specific criteria. The distinction depends on fact

including whether the tool provides diagnostic conclusions versus information f clinicians to interpret, the level of risk associated with the use case, and the deg clinical judgment required. Determining regulatory classification requires deep expertise and often involves preliminary discussions with FDA to understand th likely interpretation.

OpenAI's focus on use cases like clinical documentation and patient communica that likely fall outside FDA device regulation represents strategic regulatory positioning. Documentation tools that assist physicians with chart preparation generally do not provide diagnostic or treatment recommendations, placing ther outside current FDA oversight. Similarly, patient communication tools that prov information and answer questions without giving medical advice avoid device classification. This positioning allows faster deployment and iteration compared clinical decision support tools that may require extensive FDA review and appro processes.

However, the regulatory landscape continues evolving as FDA, CMS, and other agencies develop frameworks for overseeing AI in healthcare. FDA issued guida clinical decision support software in 2022 but many gray areas remain about whi tools require oversight. CMS has authority over reimbursement policies that cou require specific validation or approval processes for AI-enabled services to quali payment. State medical boards could potentially regulate AI tools that affect clir practice. The regulatory uncertainty creates risk for companies investing heavily product development before rules clarify.

The liability frameworks for AI-related medical errors remain largely untested v few legal precedents establishing responsibility allocation. Questions about stan of care, informed consent, and duty to monitor AI outputs lack clear answers. If documentation tool generates an error that leads to inappropriate treatment, wh bears liability? The physician who reviewed and signed the note? The health sys that deployed the tool? The vendor that created it? The foundation model comp that provided the underlying technology? These questions will get resolved thro litigation and regulation over coming years, creating uncertainty that makes risk averse institutions hesitant.

Malpractice insurance and risk management practices developed for human clin decision-making do not translate cleanly to AI-augmented care. Insurance carrie not yet have actuarial data about AI-related risks, making it difficult to price co appropriately. Some carriers might exclude AI-related claims or charge substant premiums, increasing costs for adopting organizations. Health systems need clar about liability exposure before committing to broad AI deployment, creating pre for either regulatory guidance or industry standards that establish best practice

Privacy and security requirements under HIPAA and state laws impose significa constraints on healthcare AI development and deployment. Training models on patient data requires business associate agreements, appropriate safeguards, and limitations on secondary uses that complicate data access. Sharing patient data v cloud-based AI services raises compliance questions that health system privacy officers scrutinize carefully. Companies need security and compliance programs meet healthcare's stringent requirements, adding costs and complexity compare consumer AI applications.

CMS payment policies will ultimately determine which AI applications achieve l adoption regardless of technical capabilities or clinical evidence. Health systems invest in technologies that improve financial performance through higher reimbursement, reduced costs, or better quality scores. AI tools that improve documentation specificity and coding accuracy generate immediate ROI through better capture of hierarchical condition categories and risk adjustment. Applica that reduce readmissions or improve HEDIS scores support value-based care performance, creating alignment with health system incentives. However, CMS h not established specific payment pathways for AI-enabled services, limiting investment until reimbursement clarifies.

The regulatory environment will evolve significantly over the next several years a evidence accumulates, incidents occur, and policy frameworks mature. Compani building healthcare AI should anticipate regulatory changes and design product flexibility to adapt rather than optimizing solely for current requirements. Enga proactively with regulators, contributing to evidence development, and participa in industry standards processes helps shape regulatory frameworks in favorable

directions. Companies that wait passively for regulatory clarity will likely find themselves at disadvantage versus those actively involved in policy development.

The regulatory pathway taken by ChatGPT Health through its partnerships with Color Health and bwell Connected Health provides a template others will likely follow. Partner with domain experts who understand regulatory requirements, fo initially on lower-risk use cases outside FDA device oversight, build evidence of and effectiveness through rigorous validation studies, and maintain flexibility to as regulatory frameworks evolve. This cautious, partnership-driven approach ma proceed slower than aggressive direct deployment, but it reduces regulatory risk builds credibility with conservative healthcare stakeholders who prioritize patie safety and regulatory compliance above innovation speed.

3 Likes • 1 Restack

Previous

Nex

Discussion about this post

Comments Restacks

Write a comment...