



# Keeping Patients First:

A Blueprint for AI in U.S. Healthcare

April 2026

# Introduction

For many Americans, healthcare feels like a system you have to fight your way through. It is expensive, hard to access, and hard to understand. Many people struggle to identify available providers, have limited time with clinicians, and face challenges navigating complex health information. The result is a system that too often becomes “sick care,” where people delay help until they have no choice, and where preventable problems become emergencies.

In this environment, people are using AI to reclaim more agency over their health. A recent OpenAI survey [found](#) that three in five U.S. adults said they have used AI in recent months to understand symptoms, prepare for a visit, and better comprehend clinical instructions. For many people turning to ChatGPT, care is far away, unavailable after hours, or difficult to navigate. As a case in point, based on anonymized ChatGPT message data, most health-related conversations happen outside normal clinic hours, and users in underserved rural areas send an average of nearly 600,000 healthcare-related messages each week.

Clinicians are also increasingly turning to AI to stay current with fast-moving medical developments, accelerate administrative tasks, and spend more time on patient care. According to the American Medical Association, more than four in five American physicians reported using AI for at least one use case in 2026, more than double the 2023 rate (38 percent).

AI will not, on its own, reopen a shuttered hospital, restore a discontinued obstetric unit, or bring back care in healthcare deserts. But it can make a near-term contribution by helping people interpret information, prepare for care, and navigate gaps in access, while also helping clinicians reduce administrative burden.

Looking to the future, patients can also use AI to take a more proactive role in their health by tracking health data, navigating provider and insurance networks, exploring symptoms and treatment options, and knowing when to reach out to clinicians before a concern becomes an emergency. Clinicians, in turn, can spend more time on patient-centered care as they responsibly leverage AI for routine and administrative tasks, creating more room for judgment, empathy, and careful decision-making. As these changes take hold, healthcare can become more affordable and accessible, with less friction in scheduling, documentation, follow-up, and care coordination, and with reliable support for when clinicians are not immediately available. Taken together, these shifts can improve patient and clinician experiences because more people are able to seek and find care providers, and because care itself becomes more responsive to patient needs.

AI may also serve as a catalyst for change in the broader healthcare system. For decades, clinicians, researchers, administrators, and policymakers have grappled with how to improve a system that is



essential, complex, and highly regulated. This blueprint does not attempt to resolve every longstanding debate in healthcare policy. Instead, it offers principles for responsible AI adoption that can operate within the institutions we already have, support better care, and drive improved experiences. These principles work within our existing system, where the FDA regulates medical devices at the federal level, states regulate the practice of medicine through licensing and professional standards, and where the needs, voices, and expertise of patients and providers should always come first. These principles include:

## Empowering patients through data portability

Patients should have real agency over their health information and confidence that their sensitive data is handled responsibly. In practice, this means being able to easily access their records, bring them together across systems, and use them to understand their care, coordinate across providers, and advocate for themselves.

Data portability is essential for AI to be useful. When patients can securely access and direct their data, AI can help them interpret complex information, identify trends, and prepare for more productive conversations with clinicians. Without that access and control, AI lacks the complete picture needed to provide meaningful, personalized support—leaving patients less informed and empowered. Ensuring that patients enjoy broad, secure access to their health data while protecting privacy and allowing information to move safely across systems requires:

*Reinforcing and enforcing patients' right to access and to use their own medical information.* Over the past decade, meaningful progress has been made toward granting patients effective control over their health data through policies such as the 21st Century Cures Act Information Blocking Rule, which prohibits providers from interfering with patients' ability to access, exchange, and use their electronic medical information. This policy helps patients bring together data that is dispersed across multiple providers; however, its application across the entire healthcare system remains incomplete. Policymakers should therefore clarify that all healthcare providers – even those that do not participate in government healthcare programs such as certain laboratories and pharmacies – should refrain from the information blocking practices that prevent patients from accessing their data. Making this clarification would ensure that patients can direct AI tools to help organize their information, including their laboratory results, comprehensive pharmacy dispensing and medical records to surface relevant insights, and support more informed decisions about their care.

*Deploying strong privacy and security protections for consumer health data.* As patients gain greater access to and control over their health data—including the ability to aggregate and direct it across systems—it is essential that this information remains protected. The Health Insurance Portability and Accountability Act (HIPAA) provides strong protections when health data is used within the healthcare system, and AI companies operating in that setting should comply with it. Outside of that setting, companies offering AI-enabled consumer health tools should also apply strong protections for health



data, including robust encryption for data at rest and in transit, meaningful user control over their data, and transparency about whether and how data may be used for model training.

*Supporting patient generated health data portability.* Giving patients access to their clinical records is only the first step toward meaningful control over health information. Policymakers should also extend portability rights to patient-generated health data—such as wearable biometric data, home monitoring results, and independent laboratory tests— so that it can move with the patient, in a machine-readable format, just as clinical records do. Patient-generated health data often captures important signals about daily health and wellness, prevention, and condition management that may not appear in clinical records alone. Enabling patients to assemble, analyze, and share their health records, including data generated by wearable devices, would allow individuals to identify risks earlier, manage conditions more effectively, and engage more productively with clinicians.

## Rare-disease families show why patient-directed data portability matters.

After 31 years of scattered records, inconclusive labels, and dead ends, Betsy Minium used Probably Genetic, whose OpenAI-powered intake system elicits long, free-form symptom histories from patients and caretakers, converts that into structured phenotype data and connects those families to experts with the right expertise to make a diagnosis. That process ended decades of confusion and led Betsy to a provider who made a confirmed Pitt-Hopkins syndrome diagnosis for Betsy's daughter Kali in less than a month. When families can bring fragmented information together in a form AI can review, they can more quickly reach the right test, specialist, expert-delivered diagnosis and treatment sooner.

Full story [here](#)

*Granting patients access to non-traditional medical data, especially information that is essential for coordinating care like medical appointments.* Managing care frequently requires navigating logistical questions such as when the next appointment is scheduled, which tests or results must be brought to a specialist visit, or what questions should be raised during consultations. Yet, patients sometimes cannot access or transfer scheduling information, referral details, and other operational data that shape how care is delivered. Policymakers should expand portability requirements so patients can access and direct this coordination data alongside their clinical records. While such information may not always qualify as traditional medical data, it is nonetheless essential for navigating complex treatment pathways and maintaining continuity of care. Making these elements portable would allow AI systems operating



under patient direction to help them prepare for visits, coordinate across providers, and manage their care with greater clarity and confidence.

*Removing data portability friction and implementation delays.* Policymakers should prevent information blocking practices that delay the implementation of measures that reduce friction in data portability. For example, delayed participation in and full implementation of nationwide frameworks for health information sharing, such as the Trusted Exchange Framework and Common Agreement (TEFCA), or slow adoption of identity-proofing standards such as Identity Assurance Level 2 (IAL2), can create unnecessary friction that prevents patients from securely accessing their records. Ensuring timely participation in and implementation of these requirements is essential to realizing the intent of interoperability policy: empowering patients to use their data to better understand, manage, and take control of their health.

## New research shows how AI can help patients be better informed and stay on track with their health goals.

As part of the U.S. Department of Health and Human Services (HHS) plan, “Ending the HIV Epidemic in the U.S. by 2030”, US health policy focuses on diagnosing all individuals with HIV as early as possible after infection and preventing new HIV transmissions. A recent study covering over 155,000 HIV-negative adults eligible for PrEP (HIV prevention medication) demonstrated that patients who were granted access to OpenAI-based chatbot and engaged with it had 3x higher rates of PrEP initiation than those who didn't. Follow-up care rates nearly doubled (57% vs. 32%), and appointment attendance jumped from 54% to 66%. Importantly, 80% of AI users came from racial and ethnic minority communities, which are disproportionately affected by HIV, and historically the hardest to reach. 56% of clinic staff reported saving time on routine tasks, freeing them up for more complex care. Of note, the OpenAI-based chatbot delivered standardized, evidence-informed education and behavioral reminders related to PrEP adherence, appointment scheduling, and general sexual health.

### More details in the [study](#).

Narayan, A., Blasingame, M., Warmuth, S. et al. AI-augmented communication improves HIV PrEP initiation and persistence in populations disproportionately impacted by HIV. *npj Digit. Med.* (2026). <https://doi.org/10.1038/s41746-026-02519-3>



## Strengthening human-centered care through transparent AI use

AI should strengthen the relationship between clinicians and patients, and expand the time providers can spend on care. Licensed healthcare professionals should be able to use AI, consistent with licensing requirements and professional standards, where it improves care delivery – by reducing administrative burden, expanding access, and helping patients receive timely support that might not otherwise be available. These gains are especially important as healthcare systems face growing fiscal pressure: AI can help clinicians and care teams do more with limited resources and improve access and affordability without compromising quality. For instance, clinicians should be able to use AI to handle routine administrative work – freeing up more of their time to focus on direct patient care. Clinicians should also be able to integrate AI into clinical workflows where it supports better care, with clear, commonsense transparency for patients. To us, that means:

*Prohibiting AI from impersonating licensed healthcare professionals.* Trust in healthcare depends on clear accountability and honest interfaces. Individuals should understand when they are talking to an automated system rather than a licensed healthcare professional. Automated systems should therefore be prohibited from representing themselves as licensed clinicians. Preserving this distinction reinforces that human practitioners remain ultimately responsible for diagnosis, treatment decisions, and medical accountability. Clear boundaries help patients interpret information appropriately while preserving the clinician’s central role in care decisions.

*Allowing licensed healthcare professionals to deploy AI for administrative use.* Clinician burnout remains a major challenge across the U.S. healthcare system, with 45 percent of physicians reporting at least one symptom of burnout in 2023, according to the American Medical Association. Administrative and routine tasks continue to divert time away from patient care, even though many of these functions—such as transcription, documentation, note drafting, and summarizing clinical records—can be meaningfully supported by AI. Licensed healthcare professionals should be allowed to use AI for these administrative purposes so they can spend more time on diagnosis, care planning, and direct engagement with patients. Policies should protect clinicians’ ability to responsibly integrate AI into their workflows while keeping these uses under clinician supervision. Disclosure requirements should not interfere with routine workflow tools, particularly where AI is assisting clinicians behind the scenes rather than interacting directly with patients.

*Enabling clinician-directed use of AI within professional standards.* AI tools will increasingly support clinicians in activities that extend beyond administrative work, including helping interpret information, identify care gaps, and assist with aspects of clinical decision-making. Policymakers should allow licensed healthcare professionals to deploy these tools within the scope of their existing professional standards and responsibilities, which already require clinicians to explain diagnoses, treatment options, and care decisions to patients. Transparency about AI’s role in care should therefore be integrated into clinicians’ existing ethical obligations to present medical facts accurately and to disclose relevant information sensitively and respectfully, tailored to the patient’s needs and preferences, rather than



imposed through rigid process rules that require clinicians to repeatedly disclose their internal use of AI tools. Blanket disclosure requirements risk substituting procedural compliance for meaningful transparency, creating notification fatigue without improving patient understanding or outcomes.

*Design pilot AI affordability measures with public reporting on what works.* States and the federal government should establish AI pilots in Medicaid, Medicare, public hospitals, and other settings with rising costs. These pilots should test whether AI can expand service availability, shorten delays, and lower avoidable costs while preserving quality and human accountability. Pilots should have a clear ongoing public reporting of common metrics like patient out-of-pocket burden, access and wait times, clinician time returned to patient care, and health outcomes. Governments should publish results on a regular cadence to scale what works, address what does not, and ensure AI-driven efficiency flows to patients and public programs.

## Modernizing regulation to enable safe and scalable AI-supported care

AI can strengthen the clinician–patient relationship by extending clinicians’ reach, reducing delays, and helping patients receive timely support that might not otherwise be available. But realizing these benefits depends on whether AI can be used in real-world care. Policymakers should ensure that clinicians can use AI in ways that are safe and scalable by creating clear pathways for deployment—enabling providers to test and adopt AI in clinical settings while establishing consistent national standards. Fragmented or unclear regulatory requirements can slow adoption and limit the impact of otherwise beneficial tools.

Enabling safe, scalable deployment requires a coordinated approach across state and federal policy. States should enable supervised experimentation—particularly in how AI can be integrated into care delivery—while federal oversight should cover AI systems that qualify as medical devices. At the same time, the federal government should modernize how it evaluates AI-enabled systems, while allowing states to support near-term learning and adoption. To us, that means:

*Enabling responsible innovation through regulatory sandboxes.* As AI moves from administrative support into clinical use, providers and clinicians should help lead decisions about how these tools are safely integrated into care, because they understand workflow, patient context, and clinical accountability. That work should be open and collaborative, with input from patients, regulators, developers, safety experts, and community health organizations.

Policymakers should support this work by creating structured, supervised pathways that allow providers to test and adopt AI systems in clinical settings, generate evidence, and refine standards over time while protecting patients. AI is a fast moving technology and it is critical to get the regulation right, through experimentations and early pilots that responsibly evaluate what works. Regulatory sandboxes offer one approach, by enabling time-limited testing with appropriate safeguards, monitoring, and limited flexibility



from specified state administrative or professional-practice requirements, while preserving patient protections, clear accountability, and escalation to qualified professionals when risk requires it. To implement this approach, policymakers should:

- Create federally aligned regulatory sandboxes that allow states, health systems, and clinicians to test AI-supported care models and generate evidence under common national guardrails, without creating state-by-state market authorization or displacing FDA oversight of medical devices.
- Provide clear regulator authority to suspend or terminate deployments that present unacceptable risk.
- Limit participation to defined, time-bound periods, with renewal contingent on continued compliance and acceptable risk.
- Encourage partnerships among AI providers and qualified clinical, research, community, rural, and safety-net organizations, including academic medical centers where appropriate, to support evidence generation and advance state health priorities.

*Modernizing and clarifying FDA oversight so AI-enabled medical systems can be evaluated and adopted safely at scale.* FDA's existing frameworks and policies leave open questions about the scope of its oversight for AI tools that assist clinicians, guide treatment decisions, or support care delivery. FDA should clarify how and when AI-based systems are not subject to medical device regulation or to its enforcement discretion. For higher-risk systems intended for diagnosis or treatment, FDA should establish workable review pathways that allow innovators to demonstrate that AI systems are safe and effective as medical assistants. To implement this approach, FDA and HHS should:

- Clarify regulatory pathways for specialized, multi-function, and generalist AI-enabled medical software across a range of supervision models, so defined clinical tasks can be evaluated based on intended use, evidence, safeguards, and post-deployment monitoring. FDA has a long history of reviewing a wide range of medical devices on a disease-by-disease or medical specialty basis, which may not map neatly onto AI tools with comprehensive capabilities. Outline a transparent and workable policy for demonstrating safety and effectiveness through appropriate data, representative clinical scenarios, and post-deployment monitoring, without requiring exhaustive condition-by-condition submissions where risk can be appropriately bounded..
- Preserve and clarify the 21st Century Cures Act framework for clinical decision support software, including practical examples for provider-facing AI tools used in low-risk clinical contexts. Clear guidance for tools that assist – rather than replace – clinical judgment and allow clinicians to independently review the basis for recommendations can support responsible adoption, reduce uncertainty for clinicians and developers, and help expand access to timely care.



# Conclusion

AI presents a significant opportunity to improve healthcare in the United States—empowering patients to better understand and manage their health, enabling clinicians to refocus on human-centered care, and helping governments address growing budget pressures. Realizing this potential will require policies that give patients control over their data through meaningful portability, allow clinicians to adopt AI responsibly, and build trust through balanced approaches to transparency and efficiency within coherent federal and state regulatory frameworks. Ultimately, this vision will only succeed if the benefits of AI are broadly shared with patients across the country—improving access, affordability, and outcomes—rather than being concentrated among large health systems, insurers, or pharmaceutical companies.

