AI in Patient Advocacy:

Adoption, Trust, and the Next Wave of Patient-Focused Innovation

A synthesis of industry, advocacy, and patient perspective

April 2026

Rx 4 good

Methodology

Over a six-week period (October to November 2025), Rx4good conducted a structured landscape scan that combined a review of publicly available information across a defined sample of 11 life sciences companies, 12 patient advocacy organizations, and three policy/standards groups. We paired this research with interviews of patient advocacy leaders at multiple pharmaceutical companies to pressure-test the findings and understand real-world barriers and opportunities from an industry perspective. In addition, Rx4good surveyed a dozen representative patient advocacy organizations in March 2026 to better understand their experiences, AI adoption patterns, and challenges.

Executive Summary

Artificial intelligence is now mainstream technology. In 2024, 78% of organizations reported using AI, up from 55% the year before. At the same time, private investment in AI remains at record levels, and generative AI continues to attract significant funding (Stanford University Human-Centered Artificial Intelligence [HAI], 2025). Public awareness is also widespread: Pew Research Center found that 95% of U.S. adults had heard at least a little about AI as of September 2025. However, broad awareness has not translated into broad comfort, as the same Pew study found that 50% of adults said they felt more concerned than excited about AI (Kennedy et al., 2025).

Healthcare is not outside this shift. Life sciences companies are investing in AI to work faster, keep pace with competitors, and reach people in new ways. Many have also published responsible AI principles and governance positions, signaling that AI is becoming a standard business capability. Patient advocacy organizations are also closely engaged in this evolving landscape. They recognize the same promise of clearer information, more tailored support, and easier navigation, while often working within constrained resource environments that make rapid adoption of new technologies more complex.

Meanwhile, the public's AI usage is growing, but skepticism is real and, in many cases, justified. Kaiser Family Foundation (KFF) polling shows that about one in six adults (17%) uses AI chatbots at least monthly to seek health advice or information, including 25% of adults ages 18-29 (Presiado et al., 2024). Yet in October 2025, fewer than half of adults said they would trust a health care app that uses AI to manage their care (41%), and only about one-third said they would trust an AI tool that accesses their medical records to provide personalized advice (32%) (Sparks et al., 2025). In a national survey published in 2025, most respondents reported low levels of trust that their health care system would use AI responsibly (65.8%) and low trust that it would make sure an AI tool would not harm them (57.7%) (Nong & Platt, 2025).

AI adoption accelerating

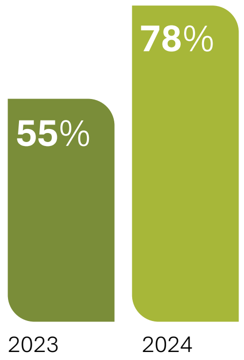

Source: Stanford University 2025.

Awareness vs Sentiment gap

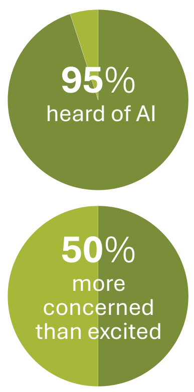

Source: Pew Research Center 2025

Executive Summary (continued)

Together, these dynamics define a central tension. AI adoption is accelerating rapidly, yet organizational capacity and public trust remain uneven. Patient advocacy organizations are responding to growing expectations for digital support while continuing to serve communities that are both hopeful and wary about how AI may affect care, privacy, and safety. This moment represents an important transition for patient advocacy, as rapid technology adoption intersects with evolving public expectations and concerns.

The opportunity for industry is not simply to deploy more tools. Instead, industry can help build the conditions necessary for responsible and equitable adoption so that AI improves patient experience without eroding trust. Based on a review of adoption patterns, public sentiment data, governance frameworks, and emerging examples in patient communities, this paper identifies five practical areas where industry can make a meaningful difference alongside their advocacy partners:

- Investing in AI readiness and literacy
- Co-creating usable governance frameworks with advocates
- Closing resource gaps for smaller organizations
- Building patient-facing AI in partnership with advocates
- Sponsoring independent evaluation and transparency to strengthen trust over time

This white paper summarizes the AI landscape in patient advocacy and the most significant opportunities for industry to support their patient advocacy partners. Finally, we offer practical steps to help organizations navigate this period of change.

AI is already reaching patients through search, chatbots, and digital tools (Presiado et al., 2024). Our research finds that trust is the gating factor. The strategic question for industry and advocacy leaders is not whether AI will be present, but whether it will be implemented in ways that are trustworthy, equitable, and aligned with patient needs. Intentional investment in governance, literacy, co-design, data stewardship, and measurement will determine whether AI strengthens patient advocacy or deepens skepticism and inequity.

> “Nothing can replace real life experience and human interaction.”
>
> *– Patient advocacy leader*

The AI in Patient Advocacy Landscape

Patient advocacy organizations exist to help individuals and families navigate serious illness, improve standards of care, and influence advancements in research and policy. AI enters this space in two primary ways. First, it can enhance internal operations, including writing, analysis, triage, and administrative efficiency. Second, and more visibly, it can shape and support patient-facing experiences such as education, navigation, and peer support.

Current use cases cluster around key advocacy efforts:

- **Summarizing and simplifying information:** Summarizing complex medical information, translating clinical terminology into accessible language, and tailoring content to a person's stage of disease.
- **Care and resource navigation:** Matching individuals to local or virtual resources such as financial assistance, transportation, and specialist directories, and helping them prepare focused questions for clinical visits.
- **Clinical trial awareness and fit:** Helping patients understand eligibility criteria and the practical implications of participation, including travel, time, and scheduling commitments.
- **Patient support at scale:** Triaging inbound questions, routing individuals to mentors or nurse educators, and supporting moderated communities while preserving human connection and safety.
- **Advocacy intelligence:** Synthesizing survey data, research insights, and policy activity to inform messaging, state-level strategy, and stakeholder engagement.
- **Operations and fundraising:** Streamlining grant writing, donor communications, and administrative work so limited staff time remains focused on patient support.

Examples of patient-facing tools are already emerging in the nonprofit sector. LUNGevity Foundation partnered with Savor Health to offer Ina®, a sophisticated chatbot-style nutrition support tool that provides on-demand guidance for people affected by lung cancer (LUNGevity Foundation, 2019). The Testicular Cancer Foundation recently launched the TC Navigator, powered by Chatbase, deploying an intelligent agent trained on TCF resources to provide in-language global virtual navigation support (Chatbase, 2026).

Other efforts—primarily research grants and collaborative data initiatives—highlight the important role patient advocacy organizations play in advancing research infrastructure and data-driven discovery. Several disease communities are investing in data platforms that may support AI-enabled insights on the research front. For example, Global Genes announced a Chan Zuckerberg Initiative grant to advance its RARE-X patient data platform and integrate AI to enhance analytics and access (Global Genes, 2025).

The Crohn's & Colitis Foundation's IBD Plexus® platform was cited as a key data asset in a collaboration with CytoReason aimed at powering advanced analytics to support research discovery in inflammatory bowel disease (Crohn's & Colitis, 2024). Similarly, The Michael J. Fox Foundation described a $14.7 million grant to support expanded analysis and publication of a Parkinson's dataset using Verily's open science tools (The Michael J. Fox Foundation, 2025).

AI shows great promise in enhancing the work of patient advocacy organizations. However, the same characteristics that make AI attractive—speed, scale, and personalization—also increase the risk profile in patient advocacy settings. Patient communities range in size, but typically are deeply personal and highly sensitive to privacy harms. The central question is not whether AI can help. It is whether it can help safely, consistently, and in ways that strengthen rather than weaken trust. Direct input from patient advocacy leaders illuminates AI adoption realities and what support is most likely to accelerate responsible progress.

Trust Is the Gating Factor

Patient skepticism toward AI in health care extends beyond hesitation about new technology. It is often rooted in lived experience, including errors in care, inequities, privacy breaches, and the high stakes inherent in health decisions. In a nationally representative survey, most respondents reported low trust that their health care system would use AI responsibly (65.8%) and low trust that it would ensure AI tools would not harm them (57.7%) (Nong & Platt, 2025). Kaiser Family Foundation (KFF) polling reinforces this caution: fewer than half of adults said they would trust an AI-enabled health app to manage their care (41%), and only about one-third would trust an AI tool that accesses medical records to provide personalized advice (32%) (Sparks et al., 2025).

Within cancer communities specifically, Cancer Support Community and NCCN reported that 44% of respondents said AI in healthcare scares them, while only 2% said they trust AI recommendations over their doctor (Fortune et al., 2025). A 2025 scoping review of studies involving adult patients with cancer found that many patients support AI when it is positioned as physician-guided, and that trust increases when clinicians endorse or supervise its use (Hilbers et al., 2025).

A skeptical public

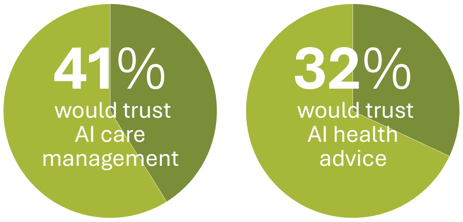

Kaiser Family Foundation 2025

Cancer communities: ↑fear, ↓trust

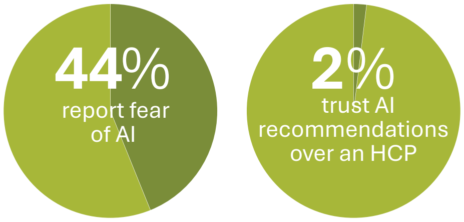

Cancer Support Community and NCCN 2025

Even in this low-trust environment, AI use is increasing: 17% of adults report using AI chatbots monthly for health information (Presiado et al., 2024). This dynamic of use without full trust creates a high-stakes environment for patient advocacy, which often sits at the intersection of patients and clinical care. Advocacy organizations provide navigation, guidance, and tools to help patients participate in their own healthcare decision-making. As AI plays an ever-growing role in the decision-making process, advocacy leaders are working to address questions and concerns, but they also echo patients' sentiments, underscoring that skepticism is rooted in issues that directly affect patient safety and equity.

In Rx4good's March 2026 survey of patient advocacy organizations, several respondents focused on medical accuracy and the risk of patients receiving incorrect answers about prescriptions or treatment questions. Others raised concerns about bias, emphasizing the need for verification and calling for peer-reviewed evidence, clarity about how outputs are generated, and assurance that AI tools work for minority groups while protecting privacy and data security.

In fact, generative models do carry known risks, including confident but incorrect answers, and biased or incomplete data can translate into biased outputs. Researchers have noted that suboptimal quality of race and ethnicity data in electronic health records can threaten fairness in health AI applications (Tsalidis et al., 2025).

Advocacy leaders say...

AI is based on cultural input. Social determinants of health are not represented correctly if at all. So patient advocacy needs to be on guard against any healthcare related AI.

We are concerned that people we represent will not get proper answers to medical questions regarding their prescriptions from AI.

Will AI work for minority groups? Is data safe? Can AI reduce costs and financial toxicity?

Verbatim responses to the Rx4good March 2026 AI Equity in Advocacy Survey validate trust as a gating factor.

These concerns are not hypothetical. They reflect the issues that patients and advocates consistently raise when asked about their fears regarding AI in health contexts (Hilbers et al., 2025; Fortune et al., 2025).

Trust is a gating factor not because advocates resist innovation, but rather because they recognize how quickly tools deemed helpful can actually be harmful when not carefully designed from the outset. For advocacy organizations, any AI-enabled experience that compromises reliability or relevance risks eroding the trust that patient communities rely on.

Human Connection Is Non-Negotiable

Rx4good's March 2026 survey revealed a consistent red line for advocacy groups: human connection, peer support, and lived experience cannot be replaced or replicated by technology. Advocacy leaders described "real life experience and human interaction" as central to their missions, particularly in rare disease and high-need communities where trust is built through shared experience. Some respondents cautioned against AI that simulates emotional support in ways that feel sufficient in the moment but may increase isolation over time.

> "The core of what we offer is human connection and peer support, and we don't see AI as something that belongs in that space, at least not yet"

The evidence is nuanced. In a cross-sectional study published in *JAMA Internal Medicine,* a panel of clinicians rated AI chatbot responses to patient questions as higher quality and more empathetic than physician responses in many cases (Ayers et al., 2023). However, scholars have cautioned against conflating fluent language with genuine therapeutic empathy. Jodi Halpern has argued that "empathic AI" is either impossible or unethical when it replaces human relationships in moments of distress (UC Berkeley Public Health, 2021). Taken together, the evidence suggests that AI may support human care, but it should not be positioned as a substitute, particularly in emotionally charged or high-stakes patient situations.

The strategic question is how AI can protect human time and capacity so advocacy teams can focus on the work that cannot be automated: peer support, community building, and patient navigation informed by lived experience.

Must-haves for earning trust:

- Disclose when and how AI is involved
- Manage expectations
- Keep humans easy to reach
- Define AI limits, i.e. what is always human
- Learn, improve and share transparently

Where Advocacy Organizations are Using AI Today

Rx4good's survey of patient advocacy organizations suggests a cautious but steadily advancing AI adoption landscape. AI use tends to show up in practical, lower-risk functions such as drafting and adapting educational materials, organizing information, summarizing content, and supporting outreach planning—more behind-the-scenes work vs direct patient support. Meanwhile, patient-facing applications remain exploratory. Most advocacy organizations appear to be in exploratory or early-use phases: testing tools, piloting small applications, or using AI intermittently. A smaller share of respondents report consistent, repeatable use in day-to-day work. None describe AI as embedded across their organizations.

Potential advocacy AI adoption barriers

(in order of significance)

Rank	Potential barrier
1	Limited budget or competing priorities
2	Concerns about data privacy and security
3	Limited internal capacity
4	Difficulty identifying the right partners or opportunities
5	Limited internal AI expertise
6	Lack of clear evidence of patient benefit

Rx4good

Taken together, the pattern reinforces a broader story: momentum is building, but responsible adoption support must meet advocacy organizations where they are today. Industry partnerships and AI capacity-building efforts must prioritize practical enablement while protecting trust, rather than flashy technology deployments. Our advocacy survey highlighted two intertwined constraints. First are hard resource limits: budget, staffing capacity, and competing priorities that make experimentation difficult even when interest is high. Second are trust guardrails: privacy and security safeguards, as well as practical guidance on which tools are trustworthy and appropriate.

Respondents also pointed to a need for partners and vendors that understand nonprofit realities and can provide support that is mission-aligned and right-sized for available resources, not designed only for large corporate environments.

Health AI Guidelines Exist, but Translation to Patient Advocacy Is Still Developing

Multiple organizations have published principles for trustworthy or responsible health AI. NIST's AI Risk Management Framework provides a risk-based approach to AI governance (National Institute of Standards and Technology [NIST], 2023). WHO has published guidance on ethics and governance of AI for health, including additional 2025 guidance focused on large multimodal models in healthcare (World Health Organization, 2025). CHAI has published a Responsible AI Guide and related tools, and CHAI and The Joint Commission have released high-level guidance for the responsible use of AI in healthcare settings (Coalition for Health AI, 2024; The Joint Commission, 2025). The National Academy of Medicine has published a discussion draft of an AI code of conduct for health and biomedical science (Bain & Company, 2024). International standards, including the OECD AI Principles, emphasize trustworthy AI and human rights (OECD, 2024).

Patient advocacy organizations could benefit from versions of these principles that also translate high-level guidance into practical operational decisions: when AI may answer questions, when human review is required, how AI use should be disclosed, what data may be collected, and how safety concerns should be managed.

Significant Opportunities for Industry to Support AI in Patient Advocacy

As patient advocacy organizations navigate an accelerating AI transition, industry is uniquely positioned to provide support through funding, convening, shared infrastructure, and responsible governance. The objective is not to push AI into advocacy settings. Rather, it is to reduce the friction and risk that can make responsible adoption more complex, while ensuring that patient trust remains a central design requirement.

5 collaboration opportunities for advocates and industry:

- AI readiness
- Realistic governance
- Filling resource gaps
- Co-creating tools
- Independent evaluation

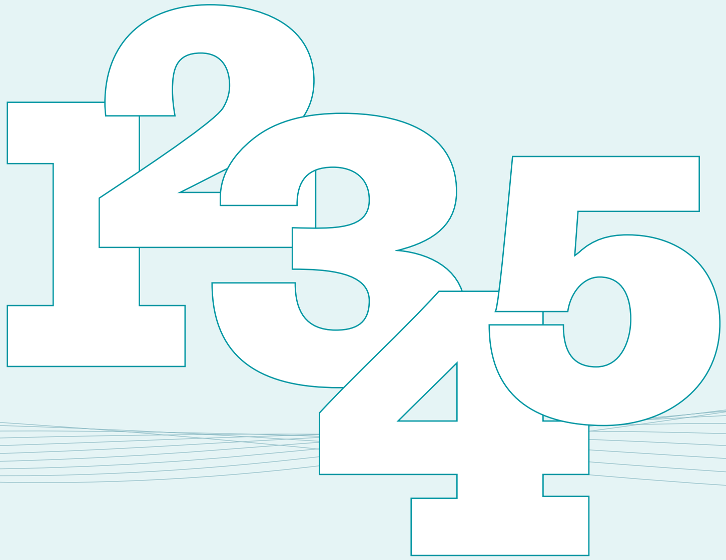

Industry Opportunity 1:

Invest in AI readiness and literacy as shared infrastructure

Many patient advocacy organizations operate with small program teams, dedicated support lines, and missions that cannot pause. For that reason, the first opportunity is to fund practical AI readiness, including concise training modules, plain-language policies, and clear guidance on tool selection and safe use. When implemented thoughtfully, these efforts can reduce fear while supporting confident and responsible use.

Strong support includes:

- Role-based training for program teams, support staff, research functions, and leadership
- Simple policy templates defining approved use cases and review requirements
- Vendor-neutral evaluation guides focused on privacy, safety testing, and accessibility
- Support for multilingual and low-literacy adaptations so benefits are not limited to a subset of patients

Impact can be measured through increased staff confidence, more consistent usage patterns, streamlined workflows, and a higher proportion of content that meets quality and disclosure standards.

Industry Opportunity 2:
Co-create guidelines and governance that advocates can use

Patient advocacy leaders can fill a critical role in shaping the evolution of AI rules and norms to ensure benefits to patients. Industry can support this by funding and advocating for patient representation in standards-setting bodies, supporting policy education for advocates, and aligning corporate commitments with patient-centered principles.

Patient-oriented rights frameworks, such as the Light Collective's Patient AI Rights Initiative, offer a practical lens for defining what responsible use should mean at the point of care and support (The Light Collective, 2024).

Practical guardrails include:

- Disclosure rules specifying when and how patients are informed that AI was used
- Content boundaries clarifying which topics require human review, including urgent safety concerns, crisis support, or medication change recommendations
- Escalation pathways ensuring timely access to a human, including after-hours handling
- Ongoing quality and safety testing for incorrect answers, bias, and unsafe advice
- Clear data rules outlining what is collected, how it is stored, whether it is used for training, and how consent is handled

Guidance continues to evolve, including new FDA principles introduced in 2026 (U.S. Food and Drug Administration, 2026). Patient communities should not be treated as downstream recipients of these decisions. They should serve as co-authors of what responsible implementation means in practice. Industry can advance this goal by supporting structured patient participation in standards efforts.

Industry Opportunity 3:

Support equitable access across advocacy organizations

A key consideration in the transition to widespread AI adoption is ensuring that organizations of different sizes and resource levels can participate fully, including many groups serving rare disease communities. Adoption differences observed in the general public demonstrate that access and confidence are not automatically equal (NIST, 2023). Industry can advance equity in advocacy by funding shared services, pooled vendor access, and security support that expand access to responsible AI capabilities.

Support includes:

- Shared service models, such as security review, privacy support, and policy development, that multiple organizations can access
- Small grants focused on foundational adoption needs, including secure tool deployment and staff training
- Vendor-neutral, patient-first technical assistance

If adoption and digital skills remain uneven, the benefits of AI will also remain uneven. Targeted grants, shared service models, and inclusive training programs can help address this challenge.

Industry Opportunity 4:

Build patient-facing AI with advocates, not for them

Patient-facing tools are most effective when grounded in lived experience rather than solely technical capability. Industry can sponsor co-design efforts that bring product teams and patient leaders together to define appropriate use cases, select effective channels, and stress-test content for tone, clarity, and cultural relevance.

Strong co-design practices include:

- Advocate and patient involvement before development begins
- Testing with diverse users, including caregivers, older adults, and non-English speakers
- Clear communication regarding limitations, information sources, and pathways to human support
- A commitment to iterative improvement based on community feedback

This consideration is particularly important for multilingual access and for topics where misinformation and fear can spread rapidly. Transparency should be the default. Patient-facing AI tools should clearly state their limits, identify information sources when possible, and make escalation to a human straightforward.

Industry Opportunity 5:

Sponsor independent evaluation and publish learnings

In many cases, claims about AI outpace available evidence. Independent evaluation enables patient communities to distinguish meaningful progress from unsupported claims and supports responsible scale.

Evaluation measures include:

- Accuracy and safety audit results tracked over time
- Equity outcomes across demographic groups and levels of digital access
- Operational impact, such as time saved or faster routing to human support
- Patient-centered outcomes, including reduced confusion and/or improved navigation success

Publishing findings, even when results are negative or mixed, can strengthen patient trust by demonstrating a commitment to learning and improvement. Transparency also equips advocacy organizations with clearer evidence when determining what to adopt and what to avoid.

Rx4good Offers Real-World Solutions

Rx4good's focus in this space is intentional. We view AI as a tool that can protect patient time, reduce friction in care navigation, and personalize support, but only when trust and equity are designed in from the outset. Our analysis identified a meaningful gap between industry's internal AI momentum and the limited public evidence of sustained, co-developed AI initiatives with advocacy partners. Closing that gap is not solely a technical task. It requires strategic foresight, disciplined planning, and sustained relationship-building.

There are several practical ways we support clients and advocacy partners engaged in this endeavor:

- **Landscape and readiness assessments:** Rapid, evidence-based evaluations of the AI environment within a disease area, including patient expectations, advocacy capacity, and risk considerations
- **Program design and co-creation:** Development of patient-facing AI or AI-enabled support programs with advocates at the table, including use-case definition, content standards, human escalation pathways, and accessibility requirements
- **Governance and trust-building:** Creation of lightweight, usable governance models aligned with recognized frameworks and clear on privacy, safety, transparency, and accountability
- **Partnership strategy:** Identification and structuring of partnerships across nonprofits, health systems, and technology providers that address capability gaps with defined roles, guardrails, and shared outcomes
- **Measurement and learning:** Practical evaluation plans and reporting frameworks that help clients understand what works, build community confidence, and scale responsibly

This moment requires deliberate and disciplined action. AI is already reaching patients through search, chatbots, and digital tools (Presiado et al., 2024). The strategic question for industry and advocacy leaders is not whether AI will be present. The question is whether it will be implemented in ways that are trustworthy, equitable, culturally meaningful and aligned with patient needs. Intentional investment in governance, literacy, co-design, data stewardship, and rigorous evaluation will determine whether AI strengthens patient advocacy or deepens skepticism and inequity.

References

Ayers, J. W., Poliak, A., Dredze, M., Leas, E. C., Zhu, Z., Kelley, J. B., Faix, D. J., Goodman, A. M., Longhurst, C. A., Hogarth, M., & Smith, D. M. (2023). Comparing physician and artificial intelligence chatbot responses to patient questions posted to a public social media forum. *JAMA Internal Medicine*, 183(6), 589–596. https://doi.org/10.1001/jamainternmed.2023.1838

Bain & Company. (2024, February 12). *40% of pharma executives are baking expected savings from generative AI into 2024 budgets*. https://www.bain.com/about/media-center/press-releases/2024/40-of-pharma-executives-are-baking-expected-savings-from-generative-ai-into-2024-budgets/

Chatbase. (2026, January) AI Agents for Non-profits: Testicular Cancer Foundation Case Study *https://www.chatbase.co/blog/ai-agents-for-nonprofits*

Coalition for Health AI. (2024). *Responsible AI guide (RAIG) and RAIG executive summary.* https://www.chai.org/workgroup/responsible-ai/responsible-ai-guide-raig-and-raig-executive-summary

Crohn's & Colitis Foundation. (2024, May 16). *CytoReason and the Crohn's & Colitis Foundation forge groundbreaking data collaboration to advance IBD research*. https://www.crohnscolitisfoundation.org/cytoreason-and-the-crohns-colitis-foundation-forge-groundbreaking-data-collaboration-to-advance-ibd

Fortune, E., Newell, A., Pink, S., & Werner, S. (2025). BIO25-023: Patient perceptions of artificial intelligence in healthcare: Findings from the cancer experience registry. *Journal of the National Comprehensive Cancer Network*: JNCCN, 23(3.5), BIO25-023. https://doi.org/10.6004/jnccn.2024.7133

Global Genes. (2025, February 18). *Global Genes awarded prestigious $1.5 million technology grant by the Chan Zuckerberg Initiative*. https://globalgenes.org/blog/global-genes-awarded-prestigious-1-5-million-technology-grant-by-the-chan-zuckerberg-initiative/

Hilbers, D., Nekain, N., Bates, A., & Nunez, J. J. (2025). Patient attitudes toward artificial intelligence in cancer care: Scoping review. *JMIR Cancer, 11,* e74010. https://doi.org/10.2196/74010

The Joint Commission. (2025, September 17). *Joint Commission and Coalition for Health AI (CHAI) release initial guidance to support responsible AI adoption across U.S. health systems*. https://www.jointcommission.org/en-us/knowledge-library/news/2025-09-jc-and-chai-release-initial-guidance-to-support-responsible-ai-adoption

Kennedy, B., Yam, E., Kikuchi, E., Pula, I., & Fuentes, J. (2025, September 17). *AI in Americans' lives: Awareness, experiences and attitudes*. Pew Research Center. https://www.pewresearch.org/science/2025/09/17/ai-in-americans-lives-awareness-experiences-and-attitudes/

The Light Collective. (2024). *Patient AI rights initiative.* https://lightcollective.org/patient-ai-rights/

LUNGevity Foundation. (2019, October 7). *LUNGevity and Savor Health® partner to bring Ina®, your personal "intelligent nutrition assistant," to lung cancer patients.* https://www.lungevity.org/news/media-releases/lungevity-and-savor-health%C2%AE-partner-to-bring-ina%C2%AE-your-personal-%E2%80%9Cintelligent

The Michael J. Fox Foundation for Parkinson's Research. (2025). *MJFF issues grant to expand, leverage important Parkinson's dataset.* https://www.michaeljfox.org/news/mjff-issues-grant-expand-leverage-important-parkinsons-dataset

National Institute of Standards and Technology. (2023). *AI risk management framework.* https://www.nist.gov/itl/ai-risk-management-framework

Nong, P., & Platt, J. (2025). Patients' trust in health systems to use artificial intelligence. *JAMA Network Open, 8*(2), e2460628. https://doi.org/10.1001/jamanetworkopen.2024.60628

References (continued)

OECD. (2024). *AI principles.* https://www.oecd.org/en/topics/sub-issues/ai-principles.html

Pfizer. (2023). *Three principles of responsibility for artificial intelligence (AI) in healthcare.* https://www.pfizer.com/news/articles/three_principles_of_responsibility_for_artificial_intelligence_ai_in_healthcare

Poushter, J., Fagan, M., & Corichi, M. (2025, October 15). *How people around the world view AI.* Pew Research Center. https://www.pewresearch.org/global/2025/10/15/how-people-around-the-world-view-ai/

Presiado, M., Montero, A., Lopes, L., & Hamel, L. (2024, August 15). *KFF health misinformation tracking poll: Artificial intelligence and health information.* KFF. https://www.kff.org/public-opinion/kff-health-misinformation-tracking-poll-artificial-intelligence-and-health-information/

Sparks, G., Montalvo, J., III, Montero, A., Mulugeta, M., Kirzinger, A., Lopes, L., & Hamel, L. (2025, October 24). *KFF health tracking poll: Public use and trust in health care apps and websites.* KFF. https://www.kff.org/public-opinion/kff-health-tracking-poll-public-use-and-trust-in-health-care-apps-and-websites/

Stanford University Human-Centered Artificial Intelligence. (2025). *The 2025 AI index report.* https://hai.stanford.edu/ai-index/2025-ai-index-report

Tsalidis, A., Bharadwaj, L., & Shen, F. X. (2025). Standardization and accuracy of race and ethnicity data: Equity implications for medical AI. *PLOS Digital Health, 4*(5), e0000807. https://doi.org/10.1371/journal.pdig.0000807

UC Berkeley Public Health. (2021, June 11). *Can we replace human empathy in healthcare?* https://publichealth.berkeley.edu/articles/spotlight/research/can-we-replace-human-empathy-in-healthcare

U.S. Food and Drug Administration. (2026, January 14). *Guiding principles of good AI practice in drug development.* https://www.fda.gov/about-fda/artificial-intelligence-drug-development/guiding-principles-good-ai-practice-drug-development

World Health Organization. (2025, March 25). *Ethics and governance of artificial intelligence for health: Guidance on large multi-modal models.* https://www.who.int/publications/i/item/9789240084759