intellias

ETHICAL AI IN HEALTHCARE & LIFE SCIENCES

A pragmatic guide for CXOs in the healthcare and pharmaceutical industries

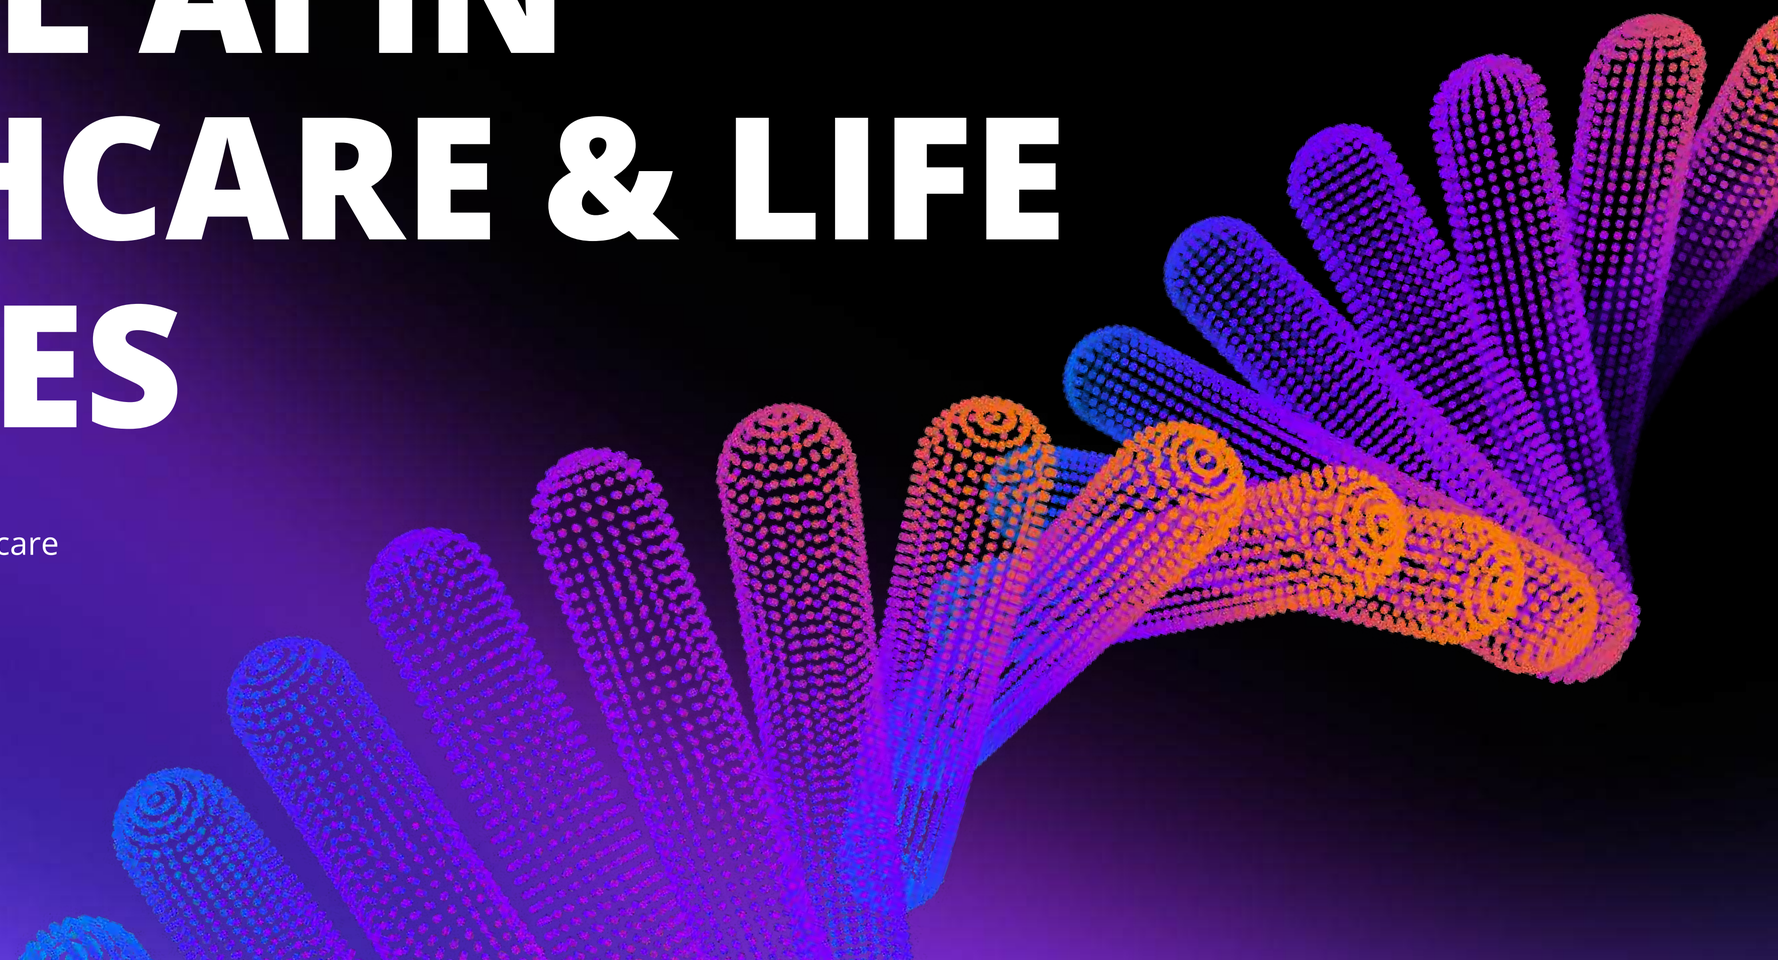

Contents

01 Executive summary

Artificial intelligence (AI) is emerging as one of the most powerful technologies in modern healthcare. It can help physicians detect diseases earlier, improve diagnostic accuracy, and expand access to specialized care. At the same time, AI raises new questions of responsibility: Whose job is it to make sure that systems remain fair, patient data stays secure, and vulnerable communities are not left behind?

The challenge for healthcare leaders is not only to deploy AI systems but to do so in ways that protect patients, empower doctors, and build public trust.

Opportunity
AI can transform healthcare by improving efficiency, detecting diseases at early stages, and expanding access to medical care.

Risks
If left unchecked, data security breaches, biased algorithms, privacy violations, and unequal access to digital tools may undermine public trust and reinforce existing inequities.

Promise
Responsible, inclusive, and transparent AI can lay the foundation for better health outcomes across all populations, ensuring that innovation benefits everyone while keeping AI-powered healthcare solutions accountable.

Embedding the principles of responsibility, inclusivity, and transparency from the start will ensure that the next-generation of digital health systems advances both technical innovation and human well-being.

02 Introduction

Inequities in diagnosis and treatment have long placed vulnerable populations at risk. For example, a Black man diagnosed with deep vein thrombosis (DVT) — a condition statistically more prevalent among Black men than any other demographic — was unaware of his heightened risk until the disease had progressed beyond prevention. Despite intervention with blood thinners, he ultimately require amputation to survive. Eventually, the disease took his life. His experience illustrates a systemic issue: a healthcare model that historically has not been grounded in equity. And this is the very issue that technology and data-driven innovation must now address.

In a primary care setting, AI can match a patient's entire medical history against data on early disease detection, then notify the physician and patient of suspected matches. With access to a system that was ethically trained on data from historically underserved patient populations, the doctor in the DVT case mentioned above could have provided the patient with information about DVT long before the condition developed. The patient could then have the opportunity to make lifestyle changes and prevent the disease that took his life.

AI adoption in healthcare

Standardization of electronic health records took decades to achieve, but AI has been adopted in healthcare within just a few years. In that time, generative AI models (such as OpenAI's popular ChatGPT) have reached more than 100 million users. That's faster than the adoption of smartphones, social media, or even the World Wide Web. As a result, AI has already become part of many medical and healthcare programs (see Figure 1). Its benefits include more personalized medical care, highly specialized diagnostic services, and access to medical and genetic research. The democratization of AI means that innovation can happen anywhere, bringing more opportunities for highly specialized diagnostic services to historically underserved populations.

Yet, democratization without governance comes at a price. Experimenting with unsanctioned AI systems may result in a breach of data protection laws and erosion of patient confidence. In their 2024 guidance on large multi-modal models in health, the World Health Organization (WHO) warns that without governance, AI systems risk undermining patient safety. Responsible governance ensures that innovation is not achieved through the use of shadow AI systems and at the expense of patient trust.

Ethical AI requires continuous vigilance against bias.

For an industry rooted in ethics, there are challenges to employing ethical AI in healthcare:

- **Privacy concerns:** Sensitive data could be processed through shadow AI systems.
- **Diagnostic accuracy:** Generative AI models have been criticized for summarizing scientific and medical information too aggressively, leaving out important details.
- **Incomplete datasets:** According to the Organisation for Economic Co-operation and Development (OECD), inequities in training data could harm historically underrepresented groups, such as the Black and LGBTQ+ communities and Native Americans.

But alongside these challenges, AI in healthcare also presents many opportunities. With a carefully designed model to evaluate AI healthcare systems for ethical righteousness across all populations, the industry can safely and effectively embrace AI as an enabler of medical progress.

Figure 1: AI solutions in healthcare

Solution	Focus area	Capabilities	Notes
Epic – Ambient Listening	Clinical documentation	Real-time transcription, EHR integration	Reduces physician burden by auto-populating records
Cerner Clinical AI	Population health & diagnostics	Predictive analytics, care coordination	Widely adopted in U.S. hospitals
Commure	Clinical workflow	Data integration, care team collaboration	Designed to unify fragmented clinical systems
Google DeepMind – Streams (NHS project)	Critical care	Early detection of acute kidney injury	High accuracy, but raised data privacy concerns
PathAI	Pathology	Computer-assisted pathology reads	Supports faster, more consistent diagnostic accuracy
Tempus	Precision medicine	Genomic data analysis, treatment recommendations	Strong oncology focus
Butterfly iQ+	Imaging	AI-enhanced handheld ultrasound	Portable diagnostics for frontline and rural care
Babylon Health	Digital triage	Symptom checker, telemedicine integration	Accessibility gains, but transparency questioned

03 The value of ethical AI in healthcare

Once a buzzword, ethical AI has transformed from theory to practice. Healthcare organizations rely on secure and ethical AI systems to analyze patient health data and suggest plans of care. Data collected from wearables and mobile devices can reveal patterns that point to serious health issues.

This information empowers doctors to make data-driven decisions about patients' care. For example, data from wearables and mobile devices could help doctors know when to adjust medications or follow-up care before a condition becomes serious.

Figure 2: Value of AI in healthcare vs. risks

Value	Risks
Cost efficiency: smarter use of limited financial resources	**Security threats:** data breaches, AI-driven hacking
Improved accessibility: extend care to remote or underserved populations	**Privacy concerns:** patient data misuse or re-identification
Innovation: new tools for clinicians and researchers	**Bias:** unequal outcomes across demographic groups
Early detection: uncovering subtle patterns and unseen problems	**Inequity:** barriers for underserved or marginalized populations
Patient-centered solutions: technology supports human interaction	**Opacity:** black box decisions without clear traceability

Ethical AI helps healthcare providers diagnose and treat illnesses more efficiently, regardless of the patient's genetics. As a result, grounding healthcare AI in ethical principles extends the benefits of modern medical technology to more people. Furthermore, ethical AI offers many underlying benefits for healthcare organizations, including:

- **Lower cost:** In situations where resources are limited, AI solutions provide a cost-effective method to deliver healthcare services.
- **Better access to care:** AI systems may make healthcare more accessible to historically underserved populations and people in remote locations.
- **More tools for experimentation:** AI gives clinicians and researchers tools to create and try new solutions.
- **Detection of unseen problems:** Through pattern recognition, AI systems can detect health anomalies sooner, which can lead to better treatment and more positive prognoses.

04 Core ethical AI considerations

Ethical AI in healthcare is a requirement for safe and equitable practice. The algorithms in AI models influence clinical decisions and patient outcomes, which means they must be designed, trained, and managed with fairness and accountability in mind. This is especially important for populations that have historically been medically underserved.

The purpose of ethical AI is to reduce inequities and ensure that new technologies improve access to care while respecting the dignity of all patients.

Responsible AI standards

ISO/IEC 42001, published in 2023, is the first international standard for AI management systems. It sets out requirements for how organizations govern AI, including policies for managing risks and ensuring accountability. The standard is complemented by the NIST AI Risk Management Framework, which provides methods for assessing and reducing risks and security vulnerabilities. Together, the ISO/IEC 42001 and the NIST framework establish a foundation for responsible AI in healthcare.

Security & privacy

The General Data Protection Regulation (GDPR) in the EU and the Health Insurance Portability and Accountability Act (HIPAA) in the US defines how health information must be collected, stored, and shared for privacy and security. Ethical AI extends these protections to development with privacy-by-design principles. Additionally, ethical AI requires patient consent to be obtained in a way that allows patients to understand and control how their information is used.

Intellectual property

The integration of AI into healthcare raises complex questions about ownership. Intellectual property rights apply to AI-generated content such as automated clinical summaries, but there is ongoing debate about attribution and control. In healthcare, intellectual property must be managed in a way that protects everyone's interests while allowing important medical knowledge to remain accessible, especially when it has a direct effect on patient health and safety.

Bias & inclusion

In healthcare, AI biases affect the reliability of care. AI models trained on incomplete or skewed datasets may not work well for all populations. Black, Native American, and LGBTQ+ communities are especially affected, as they already face disparities in healthcare access and outcomes. Ethical AI addresses bias by training models with more representative data, auditing system performance for inequities, and including affected communities in evaluation.

Breakout Box: AI-Powered Medical Translations

Clinical challenge:
Medical knowledge is often locked in single-language publications, limiting global accessibility and collaboration. Human translation can be costly and slow, especially for specialized material.

AI solution:
Intellias developed a translation model based on Helsinki architectures and fine-tuned it on medical texts to ensure terminology accuracy.

Capabilities:

- High-quality translation of medical articles into multiple languages
- Reduced dependence on human translators for technical material
- Custom training on subject-specific datasets for precision

Value:

Translation costs were cut by 50%, more people were reached, and quality was maintained.

Figure 3: Multifaceted approach for ethical implementation of AI in public health and medicine

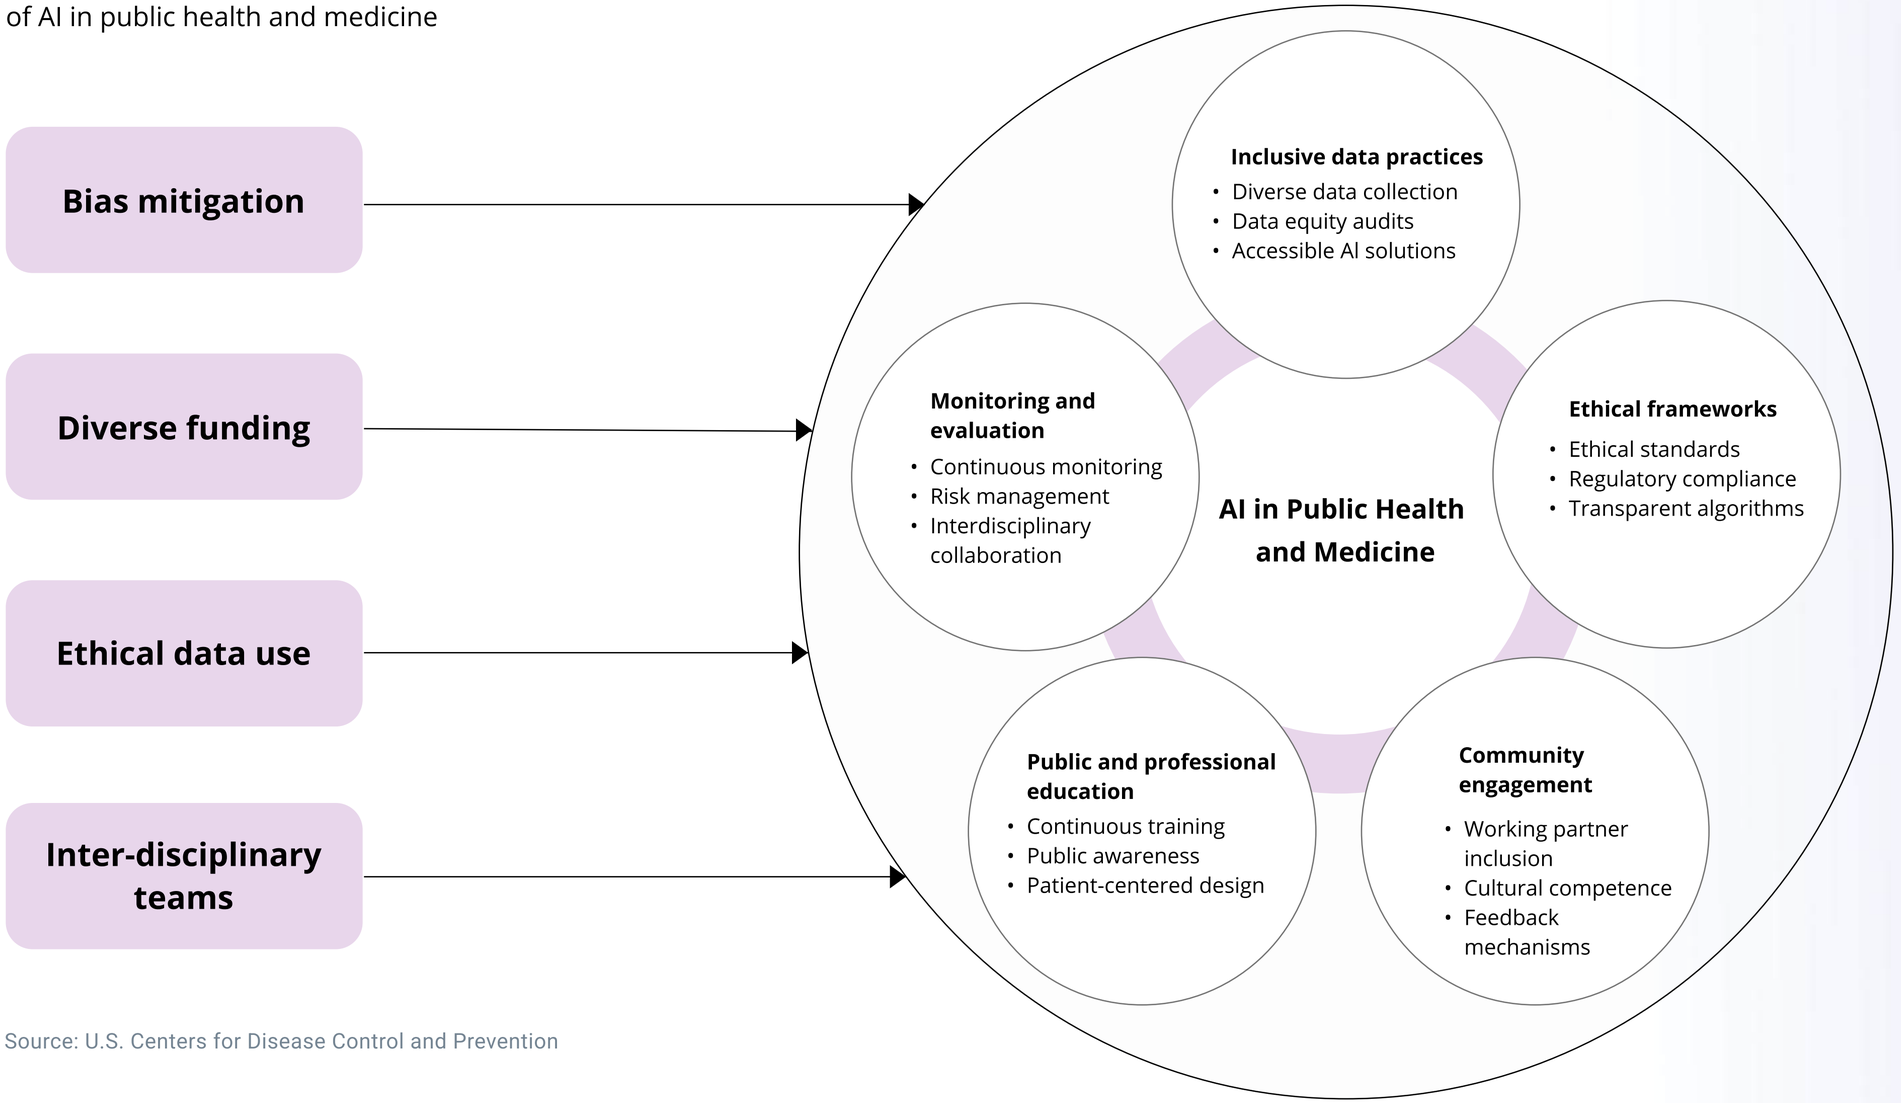

Source: U.S. Centers for Disease Control and Prevention

05 Pillars of responsible and inclusive decision-making

Ethical AI should be transparent and fair. Grounding development of healthcare AI systems in four pillars of responsible and inclusive decision-making — 1) ethics, equity, and bias mitigation; 2) regulatory compliance; 3) transparency, trust, and accessibility; and 4) interoperability — results in trusted solutions that can improve medical outcomes across all populations. They transform broad ethical goals into technical requirements. Each pillar encompasses the entire AI lifecycle, including data collection, model development, and deployment.

Ethics, equity, and bias mitigation

AI bias in healthcare distorts diagnoses and reinforces disparities. To prevent this, the development of AI models for healthcare must include training on diverse datasets, anticipating how marginalized groups interact with health systems, and testing models for disparities in prediction accuracy.

Transparency, trust, and accessibility

Trust in clinical AI grows when clinicians can interpret results without a data scientist and understand how the model made its decision. Patients also appreciate plain-language explanations from AI systems. After a model has been deployed, it must be monitored to ensure it doesn't drift and remains fair.

Regulatory compliance through an equity lens

Regulations on data protection and healthcare privacy require AI systems to demonstrate governance structures that safeguard vulnerable populations. Such measures could include rules for anonymization, access control, or informed consent, as well as model auditability, logging, and version control.

Inclusive Interoperability

Inclusive interoperability demands that AI system integration serve the needs of all patients. APIs with strict data standards share information securely, but the system must be able to handle variations in resources and availability. Systems must also support many languages and adapt to cultural norms.

Pillars of health AI systems

Health decisions must rest on the pillars of equity, inclusion, and accountability, with AI systems designed to be safe and fair for all populations.

Pillar 1

Ethics, equity & bias mitigation

- Use representative datasets spanning demographics and regions
- Validate models for disparities in prediction accuracy
- Require human oversight of clinical data

Pillar 2

Regulatory compliance through an equity lens

Align with HIPAA, FDA, GxP, and GDPR standards

- Implement anonymization and access controls by default
- Maintain auditable records of model versions and decisions

Pillar 3

Transparency, trust & accessibility

- Provide traceable outputs linked to data sources and logic
- Deliver plain-language explanations for clinicians and patients
- Monitor errors continuously with human-in-the-loop review

Pillar 4

Inclusive Interoperability

- Support secure data exchange via open APIs and standards
- Design for multilingual and multi-regional integration
- Ensure interoperability does not exclude underserved populations

06 Expanded ethical commitments

Beyond these four pillars, ethical AI requires further ethical commitments. These include commitments related to governance, engagement, auditing, sovereignty, human-centered design, and sustainability. Together, they create a holistic model of fairness and accountability.

Accountability and governance

Each stage of AI decision-making needs to be clearly defined. This is achieved through review boards whose members — including clinicians, data scientists, patients, and legal experts — represent a variety of professional and social perspectives. Recently, the US Department of Veterans Affairs (VA) formalized an AI Governance Council. Its public AI inventory demonstrates the VA's commitment to providing oversight of AI models. Meanwhile, at the Morehouse School of Medicine, the Center for Excellence in Digital Health is partnering with AIM-AHEAD to improve inclusivity in AI development.

Community engagement

When patients, clinicians, and community representatives shape AI systems, those systems begin to reflect shared values. The Canadian Institute for Advanced Research (CIFAR) showed this through its AI for Diabetes Prediction and Prevention Solution Network, in which both clinicians and community members from Peel Region — where one in six people live with type 2 diabetes — designed a diabetes-risk dashboard. Their feedback on transparency, education, and access directly transformed a technical model into one that communities could understand and trust.

Bias auditing and continuous monitoring

AI models require constant evaluation to prevent unequal outcomes. Regular audits of training data and clinical recommendations should reveal whether underrepresented groups face higher error rates, misdiagnoses, or limited access to follow-up care. Effective oversight combines quantitative checks — such as demographic error analysis — with qualitative review by clinicians and community advisors who can interpret real-world effects. Continuous re-validation ensures that, as data and populations shift, AI systems remain accurate, fair, and trustworthy in clinical use.

Human-centered design

AI in healthcare should be designed to empower clinicians, not replace them. When thoughtfully built, an AI system can strengthen the expertise of doctors and nurses. AI systems can process vast amounts of clinical data to identify patterns that might otherwise go unnoticed, thereby enhancing diagnostic accuracy and treatment planning. Furthermore, some AI systems in healthcare provide patients with personalized health information and decision support, enabling them to take a more active role in managing their health.

Data sovereignty and cultural respect

Health information is deeply personal, and cultural expectations often define how it should be protected. Safeguards must be strongest when data belongs to communities with long histories of medical exploitation. Native American and Indigenous nations have called for data sovereignty, asserting the right to control how health records are stored, shared, and interpreted. Frameworks like the CARE and OCAP principles support this, emphasizing community authority and ethical responsibility in data use. Respecting these standards helps prevent further harm and builds trust in AI-driven health systems.

Sustainability and equity of impact

Traditionally, well-funded hospitals in larger cities and private clinics in affluent neighborhoods get access to the latest healthcare technology. In many cases, rural communities lack the resources to support AI initiatives. As a result, existing disparities in population data might widen. To prevent this, members of historically underrepresented communities must advocate for fair AI practices. If a structure to ensure equity is not put in place, AI and its benefits will go to those doctors, facilities, and patients with the deepest pockets.

07 Healthcare AI in practice

To implement AI in healthcare, models must be developed and deployed within clinical environments where they will be used. This helps ensure that training data represents actual patient demographics, especially in underrepresented communities.

Data scientists must assemble diverse patient datasets from relevant sources and then use them to train and validate AI models. Members of the medical community also must review information for accuracy, trustworthiness, and bias before deployment.

For over 20 years, Intellias has worked side by side with leaders of healthcare organizations to provide engineering services and technology that meet high ethical standards. The following case studies demonstrate our dedication to ethical AI in healthcare.

Early detection assistant

Because early disease detection is difficult in primary care, a healthcare company hired Intellias to develop an AI clinical early detection assistant. The solution reviews patient records and matches them to datasets that include early signs of many different diseases. If the system suspects that a patient may have a serious condition, it will generate a referral to a specialist along with its reasoning. The system supports many languages and can be used in remote locations.

The early detection assistant includes:

- Cloud-based ML pipelines for continuous model updates
- A controlled output layer that enforces clinically justified reasoning
- BigQuery integration to manage large and varied patient datasets
- A clinical decision support system

Clinical decision support

IgA nephropathy is often underdiagnosed because early symptoms are either subtle or absent, and definitive confirmation requires an invasive kidney biopsy. To help detect the disease early, Intellias developed an AI-powered clinical decision support system. Like the early detection assistant, it generates a referral if it detects patterns for IgA nephropathy. The system has bilingual support in English and Japanese.

The clinical decision support system includes:

- A Gemini-2.5 language model paired with text-embedding-005 for structured reasoning
- A Google Cloud Services–backed vector database with ChromaDB for similarity search
- A Streamlit front end linked to Vertex AI and BigQuery for physician interaction and secure data management

08 Steps to embrace ethical AI

Adopt responsible AI standards

Healthcare organizations must ground their AI initiatives in internationally recognized standards such as ISO 42001. They provide a foundation for responsible governance, safety, and quality. Early clinical adoption avoids fragmented approaches and ensures trust is central to AI-assisted healthcare.

Embed bias checks and inclusivity reviews

Ethical AI requires continuous vigilance against bias. Bias checks include auditing the training data and continuously monitoring the model for signs of model drift that cause disparities across population groups. Inclusivity reviews also ensure that model designs reflect the diverse populations they serve.

Set up human-in-the-loop validation

Even the most advanced AI system cannot be left unchecked. AI-generated healthcare information must be reviewed regularly to ensure that the model remains unbiased for all populations. Clinicians should review every recommendation, while data analysts should periodically assess patterns of behavior and fairness.

Monitor regulatory compliance and equity of impact

Healthcare is a heavily regulated industry, and fines can cost millions. Auditing AI recommendations for compliance with HIPAA, the GDPR, and other laws and regulations is essential. This includes evaluating the downstream effects of AI adoption on patients and clinicians to ensure that AI remains beneficial.

09 Future of ethical AI in healthcare

Organizations are redefining ethical priorities to more closely align them with real practice. As new AI systems are developed, they are designed to provide explanations that can be reviewed for bias and to update themselves when problems are detected. Developers are also emphasizing sustainability and model transparency. Meanwhile, hospitals are insisting on AI consistency and protecting patient privacy through the entire AI lifecycle.

The future of AI in healthcare will be defined not only by technical progress but by the ethical choices made today. To implement ethical and pragmatic AI in healthcare, organizations should:

- **Establish ethics boards:** Create oversight bodies comprising physicians, patients, and community voices, especially those from historically underrepresented populations.
- **Urge transparent AI:** Insist on AI systems that are easy to audit and inclusive of all populations.
- **Collaborate:** Ongoing ethical oversight requires partnerships between regulators, providers, and patients.

Ultimately, the success of ethical AI in healthcare will be measured by how well patients and doctors trust AI systems to serve them fairly.

About the author

Roy Wills

Global Head of HCLS &
Alliance Partnerships at Intellias

Technology and business leader Roy Wills serves as the Global Head of HCLS & Alliance Partnerships at Intellias. In this role, Roy creates and oversees the company's strategy for serving global healthcare and life science clients with exceptional technology value. With 25 years of experience as a healthcare technology engineer, Roy manages cross-functional teams to develop people-first, scalable technology solutions for Intellias' healthcare clients.

References

1. Australian Government Department of Industry, Science and Resources. (2023). Australia's Artificial Intelligence Ethics Principles. https://www.industry.gov.au/publications/australias-artificial-intelligence-ethics-principles
2. Chakraborty, A., & Karhade, A. (2024). Cross-jurisdictional healthcare AI governance. arXiv. https://arxiv.org/html/2406.08695v1
3. Centers for Disease Control and Prevention. (2024). Health equity and ethical considerations in using artificial intelligence in public health and medicine. Preventing Chronic Disease, 21(E24). https://www.cdc.gov/pcd/issues/2024/24_0245.htm.
4. Centers for Disease Control and Prevention. (2024, August 22). Health equity and ethical considerations in using Artificial Intelligence in public health and medicine. Centers for Disease Control and Prevention. https://www.cdc.gov/pcd/issues/2024/24_0245.htm
5. Dr. Shaw, J., et al. (2024). Research ethics and AI for global health. Journal of Global Health, 14(3). https://pmc.ncbi.nlm.nih.gov/articles/PMC11025232/
6. FUTURE-AI Consortium. (2023). FUTURE-AI: Consensus guidelines for trustworthy AI in healthcare. arXiv. https://arxiv.org/abs/2309.12325
7. Gisselbaek, M., et al. (2025). Collaboration challenges in clinical AI integration. PLOS Computational Biology. https://pmc.ncbi.nlm.nih.gov/articles/PMC12105364/
8. International Organization for Standardization. (2023). ISO/IEC 42001: Artificial intelligence—Management system. Geneva: ISO.
9. National Institute of Standards and Technology. (2023). AI Risk Management Framework 1.0. https://www.nist.gov/itl/ai-risk-management-framework
10. International Organization for Standardization. (2023). ISO/IEC 42001: Artificial intelligence—Management system. Geneva: ISO.
11. Organisation for Economic Co-operation and Development. (2024). AI in healthcare: Data governance and fairness. OECD Health Policy Studies.
12. TechTarget. (2024). Shadow AI in healthcare: The hidden risk to data security. https://www.techtarget.com/healthtechsecurity/feature/Shadow-AI-in-healthcare-The-hidden-risk-to-data-security
13. The American Institute of Healthcare Compliance (AIHC). (2024). Importance of addressing shadow AI for HIPAA compliance. https://aihc-assn.org/importance-of-addressing-shadow-ai-for-hipaa-compliance/#:~:text=When%20employees%20use%20Shadow%20AI,healthcare%20organizations%20to%20significant%20risk
14. World Health Organization. (2024). Guidance on large multimodal models in health. Geneva: WHO. https://www.who.int/publications/i/item/9789240084759