COALITION FOR INNOVATION

HealthTech Blueprint for the Future

Coalition for Innovation, supported by LG NOVA

Jami Diaz, Director Ecosystem Community & Startup Experience
William Barkis, Head of Grand Challenges & Ecosystem Development
Sokwoo Rhee., Executive Vice President, LG Electronics, Head, LG NOVA

Blueprint Chairs

Alex Fang, CleanTech Chair
Sarah Ennis, AI Chair
Alfred Poor, HealthTech Chair

Authors

John Barton, Victor L. Brown, H. Timothy Hsiao, John Hsu, Sylwana Kaźmierska, Ricardo Machado, Ann M. Marcus, Nicholas Matias, Qiana Martin, Alfred Poor, Refael Shamir, Mark Wesson

The views and opinions expressed in the chapters and case studies that follow are those of the authors and do not necessarily reflect the views or positions of any entities they represent.

Senior Editor, Alfred Poor

October 2025

Preamble

The Coalition for Innovation is an initiative hosted by LG NOVA that creates the opportunity for innovators, entrepreneurs, and business leaders across sectors to come together to collaborate on important topics in technology to drive impact. The end goal: together we can leverage our collective knowledge to advance important work that drives positive impact in our communities and the world. The simple vision is that we can be stronger together and increase our individual and collective impact on the world through collaboration.

This "Blueprint for the Future" document (henceforth: "Blueprint") defines a vision for the future through which technology innovation can improve the lives of people, their communities, and the planet. The goal is to lay out a vision and potentially provide the framework to start taking action in the areas of interest for the members of the Coalition. The chapters in this Blueprint are intended to be a "Big Tent" in which many diverse perspectives and interests and different approaches to impact can come together. Hence, the structure of the Blueprint is intended to be as inclusive as possible in which different chapters of the Blueprint focus on different topic areas, written by different authors with individual perspectives that may be less widely supported by the group.

Participation in the Coalition at large and authorship of the overall Blueprint document does not imply endorsement of the ideas of any specific chapter but rather acknowledges a contribution to the discussion and general engagement in the Coalition process that led to the publication of this Blueprint.

All contributors will be listed as "Authors" of the Blueprint in alphabetical order. The Co-Chairs for each Coalition will be listed as "Editors" also in alphabetical order. Authorship will include each individual author's name along with optional title and optional organization at the author's discretion.

Each chapter will list only the subset of participants that meaningfully contributed to that chapter. Authorship for chapters will be in rank order based on contribution: the first author(s) will have contributed the most, second author(s) second most, and so on. Equal contributions at each level will be listed as "Co-Authors"; if two or more authors contributed the most and contributed equally, they will be noted with an asterisk as "Co-First Authors". If two authors contributed second-most and equally, they will be listed as "Co-Second Authors" and so on.

The Blueprint document itself, as the work of the group, is licensed under the Creative Commons Attribution 4.0 (aka "BY") International License: https://creativecommons.org/licenses/by/4.0/. Because of our commitment to openness, you are free to share and adapt the Blueprint with attribution (as more fully described in the CC BY 4.0 license).

The Coalition is intended to be a community-driven activity and where possible governance will be by majority vote of each domain group. Specifically, each Coalition will decide which topics are included as chapters by majority vote of the group. The approach is intended to be inclusive so we will ask that topics be included unless they are considered by the majority to be significantly out of scope.

We intend for the document to reach a broad, international audience, including:

- People involved in the three technology domains: CleanTech, AI, and HealthTech
- Researchers from academic and private institutions
- Investors
- Students
- Policy creators at the corporate level and all levels of government

Chapter 4:
The Ethical Use of AI in Healthcare

Authors: Ann M. Marcus, John Barton

The Rise of AI in Healthcare and Emerging Ethical Concerns

Artificial intelligence (AI) is rapidly transforming healthcare, bringing remarkable advances in diagnostic tools, personalized treatment plans, administrative efficiency, and remote patient monitoring. Yet with this progress comes a pressing ethical reckoning. As AI systems are increasingly trusted with decisions that directly impact human health, questions of equity, accountability, and privacy rise to the fore.

Concerns have emerged over how these technologies may unintentionally dehumanize care, erode patient trust, or worsen existing disparities. As AI assumes a more prominent role in decision-making – especially in sensitive contexts such as diagnosis and risk assessment – some patients and practitioners are growing uneasy over the transparency of its processes and the consequences of algorithmic errors. Moreover, reliance on biased or incomplete data sets threatens to replicate historical injustices within the healthcare system.

Who Is Affected? Identifying Stakeholders and Vulnerable Groups

The ethical landscape of AI in healthcare implicates a diverse group of stakeholders:

- **Patients** bear the brunt of AI decisions. From misdiagnoses to insurance denials, the consequences can be life-altering. Concerns over loss of autonomy, consent, and data security loom large, particularly when AI tools operate opaquely or without meaningful human oversight.
- **Healthcare professionals** face a shifting role as AI systems influence or even override clinical judgment. This dynamic can create tension between professional responsibility and technological authority.
- **Marginalized communities**—including people of color, low-income individuals, and non-native language speakers—are especially vulnerable. Underrepresented in medical datasets, they face higher risks of algorithmic misjudgment and reduced access to high-quality care.
- **Payers and policymakers** are grappling with the implications of AI in underwriting, pricing, and eligibility decisions, often without clear guidance on fairness or legal liability.

How AI Is Being Used, and by Whom

AI is no longer confined to back-office functions; it now plays a central role across the healthcare continuum:

- **Diagnosis and Treatment**: From radiology to dermatology, AI systems interpret scans, flag anomalies, and recommend therapies. While such tools can augment physician capabilities, their accuracy varies and often depends on how representative their training data is.
- **Home Testing and Observation**: Wearables and remote monitoring tools use AI to detect changes in vital signs or behaviors. While convenient, these technologies collect vast personal data, sometimes blurring boundaries between medical oversight and surveillance.

- **Administrative and Insurance Uses**: AI automates claims processing, fraud detection, and resource allocation. However, these efficiencies may come at the cost of human discretion, compassion, and fairness, especially for patients with complex or atypical profiles.

Who Is Getting Shortchanged? Equity and Justice

AI's promise of precision medicine is not evenly distributed. Access to advanced tools often correlates with institutional wealth and geographic location. Hospitals in underserved communities may lack the funding, staff, or infrastructure to adopt cutting-edge technologies, exacerbating existing care disparities.

Bias in algorithmic design and deployment further compounds the problem. If an AI system is trained on data that underrepresents certain populations, its decisions may systematically disadvantage those groups which can lead to missed diagnoses, inappropriate treatments, or denial of care.

Additionally, the communities most affected by AI systems often have the least influence over how those systems are designed and governed. This power imbalance challenges the democratic development of ethical, patient-centered technology.

Is AI for Diagnosis and Home Monitoring Responsible?

Responsible deployment of AI hinges on several principles:

- **Informed Consent**: Patients should be fully aware of when AI is being used, what data is being collected, and how decisions are made. They must retain the right to opt out.
- **Privacy and Surveillance Concerns**: Tools that monitor behavior or health at home can inadvertently collect non-medical information. Without updated regulatory protections – including reforms to HIPAA – such data could be exploited for commercial, disciplinary, or profiling purposes.
- **Transparency and Explainability**: Patients and clinicians must be able to understand and trust AI decisions. Black-box algorithms that offer no rationale undermine confidence and can erode the therapeutic relationship.

Insurance and Regulatory Challenges

AI is changing how insurers evaluate risk and make coverage decisions. While automation promises speed and efficiency, it also risks embedding structural biases into critical decisions about access to care. People with certain demographics, geographies, or social histories may be unfairly penalized.

Determining legal responsibility for AI-driven errors remains murky. If a diagnostic tool recommends a harmful course of action, who is liable: the developer, the provider, or the system itself? Existing medical and legal frameworks are ill-equipped to answer these questions.

There is a clear need for comprehensive policy updates that center on equity, patient rights, and algorithmic accountability.

Balancing Benefits and Challenges

AI can enhance healthcare delivery in powerful ways: reducing physician burnout, enabling earlier interventions, and tailoring treatments to individual biology. It holds particular promise for remote and underserved communities where it could close gaps in provider availability and diagnosis speed.

But these benefits must be weighed against real challenges. Without deliberate ethical oversight, AI could become another mechanism of exclusion. Data privacy, algorithmic bias, and the erosion of clinician-patient trust are not theoretical risks; they are already surfacing in practice.

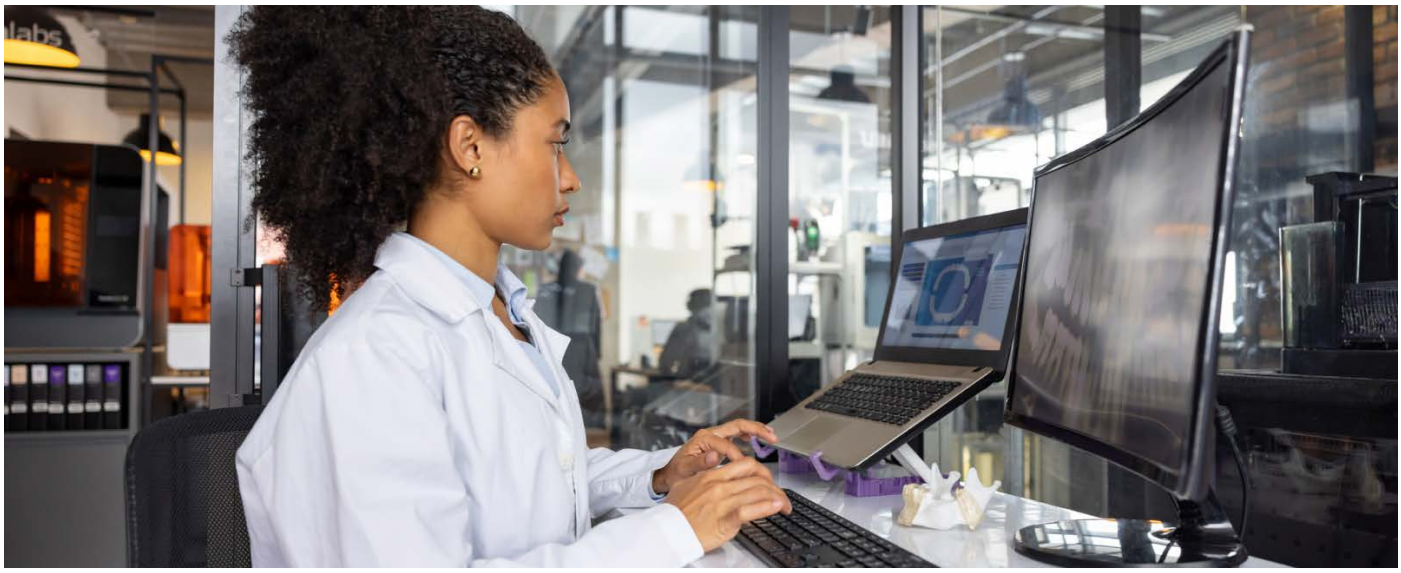

Recommendations for Responsible AI Use

To guide the ethical use of AI in healthcare, we offer the following recommendations:

1. **Develop Inclusive AI Policies**: Engage stakeholders from diverse backgrounds to co-create fair and equitable systems.
2. **Enhance Transparency and Accountability**: Ensure AI decision-making is understandable and traceable. Assign liability clearly.
3. **Strengthen Data Privacy Protections**: Update laws and frameworks including HIPAA to address the scope and scale of modern data collection.
4. **Promote Public and Professional Education**: Equip clinicians, patients, and policymakers with the knowledge needed to understand both the promise and pitfalls of AI tools.
5. **Engage Diverse Stakeholders**: Prioritize participatory design processes that include the voices of those most at risk of harm.

Charting a Responsible Path Forward in Healthtech

The integration of AI into healthcare is not just a technological shift; it is a cultural, ethical, and systemic transformation. As these tools become more deeply embedded in diagnostics, treatment, monitoring, and administration, the stakes grow higher for every stakeholder involved.

For **healthcare practitioners**, AI should be a support, not a substitute to compensate for insufficient staffing or the lack of other resources. Clinical experience and human judgment remain irreplaceable, especially when navigating ambiguity or addressing patients' unique contexts. Practitioners must have access to transparent tools they can trust, along with the training to use them effectively and ethically.

Administrators and healthcare system leaders have a responsibility to ensure AI adoption aligns with institutional values of equity, quality, and accountability. Procurement decisions should consider not only performance metrics but also the representativeness of training data, explainability, and compliance with emerging standards in algorithmic fairness.

For **health technology developers**, innovation must go hand in hand with inclusion. This means

engaging early and often with diverse populations, clinicians, and ethicists; ensuring datasets reflect the full spectrum of humanity; and building systems that are interpretable, secure, and adaptable to local needs. Responsible AI is not a regulatory burden; it is a design imperative.

Patients and communities—especially those historically marginalized in healthcare—must be centered in the AI development and deployment process. They deserve transparency, consent, and the right to opt out. Most importantly, they must have a voice in shaping the systems that will increasingly shape their care.

Ultimately, AI in healthcare can be a force for tremendous good — unlocking efficiency, insight, and access. But without intentional safeguards and inclusive design, it risks becoming another mechanism of inequity and harm. The future of ethical healthtech depends on collaboration across domains, transparency at every level, and a steadfast commitment to human dignity.

Now is the time to reimagine not just what AI *can* do in healthcare, but what it *should* and *should NOT* do, and to or for whom.

Sources

https://pmc.ncbi.nlm.nih.gov/articles/PMC11977975/

https://iris.who.int/bitstream/handle/10665/341996/9789240029200-eng.pdf

https://pmc.ncbi.nlm.nih.gov/articles/PMC8826344/

https://www.cdc.gov/pcd/issues/2024/24_0245.htm

https://www.frontiersin.org/journals/surgery/articles/10.3389/fsurg.2022.862322/full

https://www.facs.org/for-medical-professionals/news-publications/news-and-articles/bulletin/2023/february-2023-volume-108-issue-2/ethical-concerns-grow-as-ai-takes-on-greater-decision-making-role/

https://hitrustalliance.net/blog/the-ethics-of-ai-in-healthcare

Author (In order of contribution)

Ann M. Marcus, Director, Ethical Tech & Communications, WeAccel
Ann M. Marcus is a Sonoma-raised, Portland-based communications strategist and ethical technology analyst focused on smart cities, community resilience, and public-interest innovation. She leads the Marcus Consulting Group and serves as director of ethical technology and communications at WeAccel.io, a public-good venture advancing mobility, communications, and energy solutions for communities. Ann has advised public and private organizations—including Cisco, the City of San Leandro, Nikon, AT&T, and InfoWorld—on trust-based data exchange, digital public infrastructure, resilience strategy, AI and more. Her current projects include a California senior evacuation program, a Portland robotics hub, and digital energy resource initiatives with utilities in Portland and the Bay Area.

John Barton, Founder/Executive Director; AI Strategist & Architect
John Barton, Founder & Executive Director of the Spectrum Gaming Project, is an AI strategist and governance architect focused on building ethical systems for underserved markets. With a Master's in Counseling and decades in community education, he has delivered over 10,000 trainings in neurodiversity, education, and innovation. Based in Appalachia, his work has been recognized and adopted by the American Bar Association, the ACLU of West Virginia, AmeriCorps VISTA Leaders, and the WV Community Development Hub.

For more information about the Coalition for Innovation,
including how you can get involved, please visit coalitionforinnovation.com.

View the Next Chapter