

2024-2026

ORGANIZATIONAL STRATEGY



HEALTH AI

The Global Agency for Responsible AI in Health

TABLE OF CONTENTS

Executive Summary	3
Background	4
Redefining the Organization's Role	7
Vision, Mission, Principles	8
Strategic Approach	10
Intended Impact	18
Other Strategic Considerations	21



EXECUTIVE SUMMARY

This document sets forth a three-year strategic vision for HealthAI to guide the next phase of the organization. The document captures the organizational vision, mission, strategic approach, and goals to clarify the role of HealthAI in the larger global public health landscape.

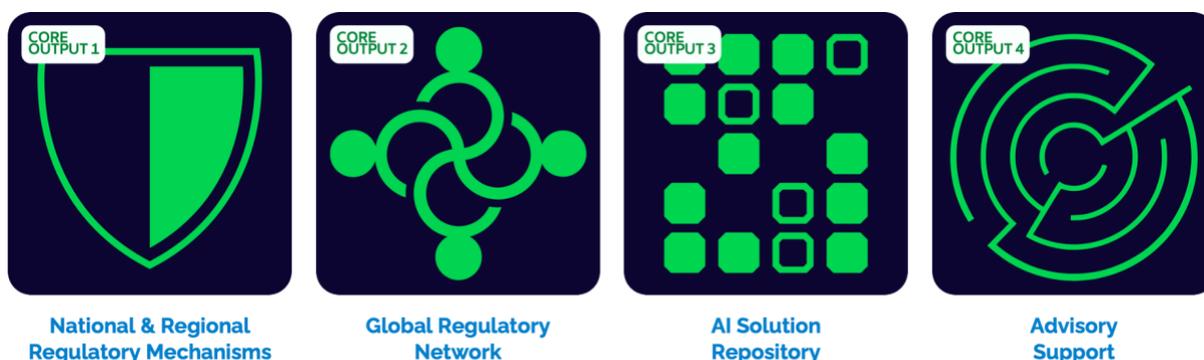
Founded in 2019, the International Digital Health & AI Research Collaborative (I-DAIR) was initially focused on addressing the urgent need to support collaborative, equitable research and development of artificial intelligence (AI) and digital health solutions. During its first three years, I-DAIR took an exploratory approach to better define the gaps in the landscape and better understand how the organization could support the positive benefits of AI and other digital technologies in health.

During these years, external factors shifted the landscape dramatically. The COVID-19 pandemic brought the world's attention to global health. Many health systems leaned on available digital technologies to maintain routine services, adding or augmenting existing telemedicine and virtual care services, while also applying data-driven decision making in their pandemic response. In addition, the incredible pace of innovation in AI resulted in both excitement over the potential of these technologies and trepidation about the ethical implications they have.

As I-DAIR began the development of its long-term vision, it became clear that the organization itself must evolve to address the gaps in AI governance for health. Part of this evolution included a refresh of the organizational identity—transitioning I-DAIR to HealthAI.

Captured in this strategy, HealthAI refined its mission to **advance the development and adoption of Responsible AI solutions in health through the collaborative implementation of regulatory mechanisms and global standards.**

This mission will be achieved through four Core Outputs:



The remainder of this document provides additional details on the organization's vision, mission, and strategic approach. It also provides concrete activities that HealthAI will pursue to achieve its strategic goals. The intended audience of this document is HealthAI staff members, Board of Directors, funders, and critical partners, and aims to guide the activities and approaches the organization pursues over the next three years.



BACKGROUND

Over the last two decades, artificial intelligence (AI) has continued to mature from exploratory and early-stage development into a vital sector of innovation. Machine learning and other AI approaches have led to the advent of facial recognition, driverless cars, language translation, and many other applications. Supported by technological advancement, AI systems are now a normal part of everyday life. The rise of generative AI, and the implications this technology has on the future of every sector, has led to a global awakening to the power and potential of AI.

This surge of innovation has also led to the rapid development of AI-driven health solutions. As with the digitalization of health systems, AI poses incredible opportunities for improving the health and well-being of individuals and communities around the world. However, this emergent technology comes with unique challenges and risks that must be addressed.

The Governance of AI in Health

There are multiple dimensions to the governance of AI in health systems and services. Private sector companies and other technology developers oversee the internal governance of a particular AI solution. Governments are responsible for the governance of technologies used within national health systems, while global normative bodies like the World Health Organization (WHO) and the International Telecommunications Union (ITU) set standards and ethical principles that are adopted by others. Finally, civil society also has the important role of advocating for and monitoring the ethical use of AI.

However, many established governance mechanisms have not kept pace with the rapid evolution of AI technologies. In some countries and regions, the regulation of AI has been integrated into existing government agencies and included in data privacy regulations, but this has been primarily limited to the countries with the highest Gross Domestic Product and most advanced technology sectors. According to a report from the WHO, some national regulatory agencies “may have neither the capacity nor the expertise to approve the use” of AI technologies. There is a serious risk of deepening inequity in both access and outcome between “early adopter” countries and countries that do not have the resources or flexibility to match the pace of technological innovation.

This lack of national governance mechanisms contributes to the slow adoption of AI solutions within health systems. Governments are hesitant to approve technologies without stronger evidence of technology’s safety and efficacy; technology developers do not have a clear path for certification and approval from regulatory agencies; and private sector companies are left to develop ethical frameworks without the broad, inclusive network or experience needed, resulting in frameworks that may be too narrow, incomplete, or misaligned with the public good.

Strengthening the governance of AI in health is necessary to safeguard the future of health. To realize the potential of these technologies and keep citizens safe, we must ensure these technologies contribute effectively to global progress on health and well-being.



Toward the Global Governance of AI in Health

Global efforts to strengthen the governance of AI in health have increased in the past five years. These efforts seek to establish common ethical and technical standards for AI technologies, provide support for legislative, regulatory, and policy development, and begin identifying and addressing critical ethical and safety concerns related to AI. With a global lens, these efforts aim to create cohesion across countries and enshrine a human rights lens into the use of AI in health.

Most notably, the [Focus Group on Artificial Intelligence for Health](#) (FG-AI4H) was established by the WHO and ITU in 2018 to begin exploring the governance of AI. Over the course of multiple years, the FG-AI4H developed resources that advance a collaborative, inclusive perspective on ethics, regulations, technology, and clinical use cases related to AI-enabled health systems and services.

A 2021 report authored by this group, titled ***Ethics and Governance of Artificial Intelligence for Health*** provides a set of ethical principles to guide the development, regulation, and implementation of AI solutions in health. The six principles serve as a foundation for the responsible use of AI, and include:

- **Protecting Human Autonomy** to ensure that individuals remain central in the collection, analysis, and use of data;
- **Promoting Human Well-Being and Safety**, echoing the broader medical principle of “do no harm;”
- **Ensuring Transparency, Explainability, and Intelligibility** that will ensure all stakeholders have the relevant information about the AI solutions to make informed decisions;
- **Fostering Responsibility and Accountability** to ensure AI solutions deliver on their potential while meeting ethical requirements;
- **Ensuring Inclusiveness and Equity**, proactively contributing to the ultimate goal of health equity without further exacerbating the digital divide;
- **Promoting Responsive and Sustainable AI** that are designed to meet the health needs of communities and individuals over time.

This same report outlines the need for global governance of AI solutions in health:

“With the increase of AI standards and laws around the world and diffusion of how and where AI ethics is managed, additional international oversight and enforcement may be necessary to ensure convergence on a core set of principles and requirements that meet ethical principles and human rights obligations.”

The need for global governance was also reinforced in [a study from the Global Digital Health Partnership](#), which concluded that “*international collaboration and coordination on AI governance for health care is needed to ensure coherent solutions and allow countries to support and benefit from each other’s work.*” Such global governance would provide support to country governments, multilateral organizations, private sector companies, civil society organizations, international organizations, and others working in the space of AI-enabled healthcare.



At the 2023 AI for Global Good Summit, the [Global Initiative on AI for Health](#) was announced, convened by the WHO, ITU, and World Intellectual Property Organization (WIPO) as the successor to the FG-AI4H. This initiative will build on the work of the Focus Group and address the need for global governance, in coordination with the WHO Global Initiative for Digital Health.

Other global efforts to support the ethical use of AI are beginning to emerge that address the need for global coordination. Many global health organizations, private sector companies, and academic institutes are collaborating to provide technical standards, ethical standards, and other guidance in this space.

Defining ‘Responsible AI’

The term ‘Responsible AI’ refers to artificial intelligence technologies that align with requirements set by normative agencies and other leaders in the sector, with specific focus on ethical, human-centric attributes.

HealthAI generally defines ‘Responsible AI’ as:

AI solutions that are ethical, inclusive, rights-respecting, and sustainable.

Attributes of Responsible AI include:

- Protection of and respect for human autonomy, agency, and oversight;
- Promotion of human well-being and safety;
- Commitment to “do no harm”;
- Technical robustness and safety;
- Adherence to laws and ethics;
- Transparency, Explainability, and intelligibility;
- Responsibility and accountability;
- Inclusivity and equity;
- Sustainability;
- Societal and environmental well-being.

These attributes are applied across all aspects of AI technologies, from the technical development of the solution to the use and management of data, the implementation and stewardship of the technology, uses of the technology, and the ultimate result of its application.

This definition is derived from the WHO publication titled [Ethics and Governance of Artificial Intelligence for Health](#), the International Development Research Center’s [AI for Global Health Initiative](#), a framework developed by the European Commission’s High-Level Expert Group on AI, described in the [Ethics Guidelines for Trustworthy Artificial Intelligence](#), and a journal publication from [Information Systems Frontier](#).



REDEFINING THE ORGANIZATION'S ROLE

The International Digital Health & AI Research Collaborative (I-DAIR) was founded in 2019 to address an urgent need to support collaborative, equitable research and development of AI and digital health solutions, in particular through the strengthening of research capacities in low- and middle-income countries.

However, the significant landscape changes in the intervening years required an examination of the organization's role in the global health sector. The COVID-19 pandemic brought focus on health systems globally, demonstrated areas of weaknesses in existing health systems, and showcased the ability of digital technologies to impact the provision of health services. At the same time, advancements in artificial intelligence have accelerated, creating new opportunities and concerns for health systems globally.

Responding to this new reality and the expressed needs of governments around the world, I-DAIR undertook a strategy renewal and rebranding process. Through a collaborative and consultative process, a consensus was reached to refocus the organization on the implementation of regulatory and policy approaches for AI in health.

This new direction builds on the organization's first three years of work, utilizing the partnerships and expertise gained through the organization's pathfinder projects. The outputs, findings, and lessons from I-DAIR's existing portfolio will guide the implementation of the new strategy from the [Trusted Research Infrastructure](#), [Global Research Maps](#), and to the ongoing work on [Responsible AI for Health](#). These projects—along with other organizational efforts—will serve as a foundation for the new mission and approach.

With a new name—**HealthAI**—and new strategic direction, the organization is positioned as a premier implementation partner for country governments, normative agencies, the private sector, and other stakeholders on the topic of Responsible AI in health. HealthAI will provide essential services to ensure the adoption of global standards and best practices across all stakeholders, a critical but nascent area of work that will continue to grow in importance.

HealthAI benefits from a staff of world-class thought leaders and innovators, who will continue advancing the field of AI in health. The organization will continue deepening relationships with partner organizations across the global, bringing perspectives from diverse contexts to bear in the development and implementation of global governance standards. Partnerships with country governments will allow HealthAI to engage with key stakeholders at the national and regional levels, while championing equity at the global level.

HealthAI will also expand its partnerships with the WHO, through the Global Initiative on AI for Health and Global Initiative for Digital Health, serving as an implementing partner for the normative function of these initiatives.

The remainder of this strategy provides details on the organizational approach and envisioned activities of HealthAI over the next three years.



VISION, MISSION, PRINCIPLES

Vision

HealthAI envisions a world where artificial intelligence produces equitable and inclusive improvements in health and well-being for all individuals and communities.

This vision requires the collective efforts of governments, international bodies, technologists, social scientists, health workers, and communities themselves to promote the ethical and responsible use of AI solutions in health while fostering innovation. This will require new approaches to regulation and policy, structures that allow health systems to respond nimbly to technological developments, and stronger partnerships between public and private sector entities. Progress must also be made to reduce the global digital divide so that all individuals and communities benefit from the advancements AI solutions will bring.

AI-enabled healthcare will require continuous cycles of learning and adapting as both technologies and health needs change over time. HealthAI's vision requires a participatory approach that centers on equity—where all stakeholders have a voice in the creation, regulation, and implementation of AI solutions in health.

Mission

HealthAI advances the development and adoption of Responsible AI solutions in health through the collaborative implementation of regulatory mechanisms and global standards.

As a partner to the WHO, governments, international organizations and the private sector, HealthAI supports the contextual application of global technical and ethical standards for AI in health. The standards and recommendations put forth by the Global Initiative on AI for Health and other normative, multilateral agencies serve as the foundation for HealthAI's approach, which is then further tailored to the expressed needs of country health systems.

Serving as a resource to and a bridge between WHO, countries, and global efforts, HealthAI will help countries enshrine regulatory best practices into country-level policies and legislation, as well as support the assessment and validation of specific AI health solutions against established global standards. Working across geographies, HealthAI can serve as a curator of information and insights—identifying common needs across countries, building connections across sectors, and serving as an early warning mechanism for adverse events related to AI in the health sector.

As a neutral partner, HealthAI will also work directly with private sector technologists to provide guidance at all stages of technological innovation and define mechanisms for assessing and validating AI solutions against ethical standards.



Core Principles

HealthAI bases all its work in three core principles. These principles inspire the organizational approach, ground the organization's decision-making, and define the way the organization and its staff conducts business. These values are reflected in the organization's Code of Conduct but goes further to define how the organization operates.



Cultivate Trust

Trust is both an essential component for and critical barrier to the acceptance and adoption of AI solutions within health systems. Concerns over safety, privacy, accuracy, and effectiveness must be addressed for these solutions to be widely used. The successful cultivation of trust will result in regulatory mechanisms that build confidence in the effectiveness of AI solutions and provide protections from potential harms.



Catalyze Innovation

Regulation, when done thoughtfully, can spur new advancements in technology and support stakeholders who are involved in the design and implementation of these technologies. Often seen as an impediment to innovation, regulation must be flexible and responsive while also providing clear guidance on topics of equity, ethics, and safety. The successful catalyzation of innovation within the regulatory space will result in continued technological development of AI solutions with the greatest potential for health impact that are supported by regulatory mechanisms that protect individuals and communities.



Center Equity

Equity must be firmly entrenched in all aspects of the AI for Health sector to prevent the continued widening of the digital divide. Rapid development of this sector in some high-income geographies and other "early adopter" countries risks further marginalizing communities in other regions, so addressing equity at every stage of the regulatory process will help capture and apply best practices and promising technologies within and across all countries. The successful prioritization of equity will result in regulatory mechanisms that support healthy equity and ensure AI solutions do not exacerbate the digital divide.



STRATEGIC APPROACH

HealthAI has developed a strategic approach in alignment with its Vision, Mission, and Principles. Outlined below, this approach includes a theory of change that clarifies the intended impact of the organization and its work, as well as additional details on each core output.

Organizational Theory of Change

As outlined in Figure 1, the theory of change captures the intended impact of the organization’s mission: that responsible and trusted AI solutions are adopted by a wide range of countries and health systems, supported by regulatory mechanisms that safeguard against potential harms and reinforce the health impact of these AI solutions.

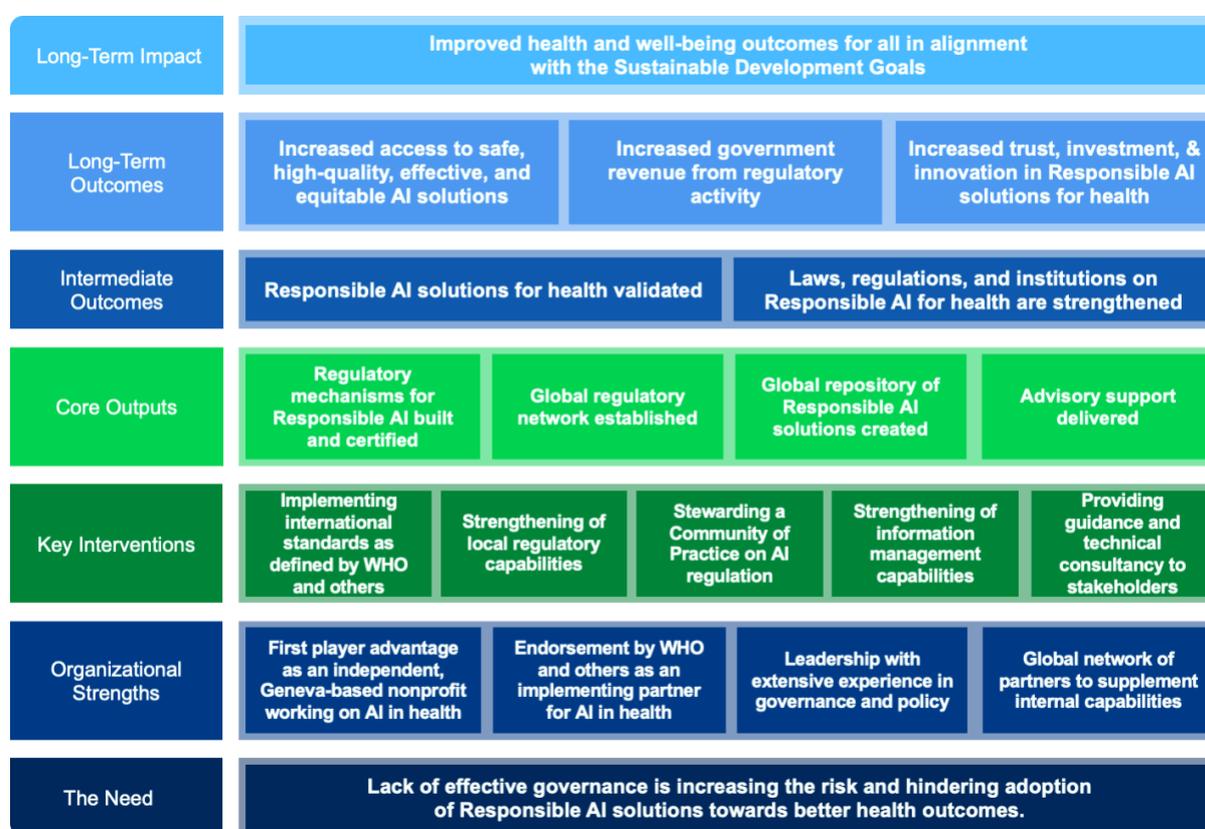


Figure 1. HealthAI Theory of Change

With improvements to the regulatory mechanisms at both a global and country level, governments, health workers, and patients alike will have more confidence in AI solutions, making them more acceptable and desirable to implement. With greater insight into AI solutions that have been validated against global standards, countries will be able to identify appropriate solutions more quickly and with greater certainty in their efficacy. Private sector partners will have clarity about the regulatory requirements and a better understanding of



how AI is used within health systems and services. Finally, taking a global approach to the regulation of AI solutions in health will support greater equity by sharing best practices and lessons learned as well as serving a critical role as an early warning system for adverse events related to the use of AI in health.

Core Outputs

HealthAI aims to become the premier partner and standard bearer for ethical and equitable AI solutions in the health sector. In pursuit of this overarching goal, HealthAI has outlined four Core Outputs that define the work the organization will undertake over the next three years. Each output is supported by key interventions which will be used to develop a detailed implementation plan (separate from this strategy).

The Core Outputs are designed to work together to support the entire ecosystem of AI regulation in the health sector (Figure 2):

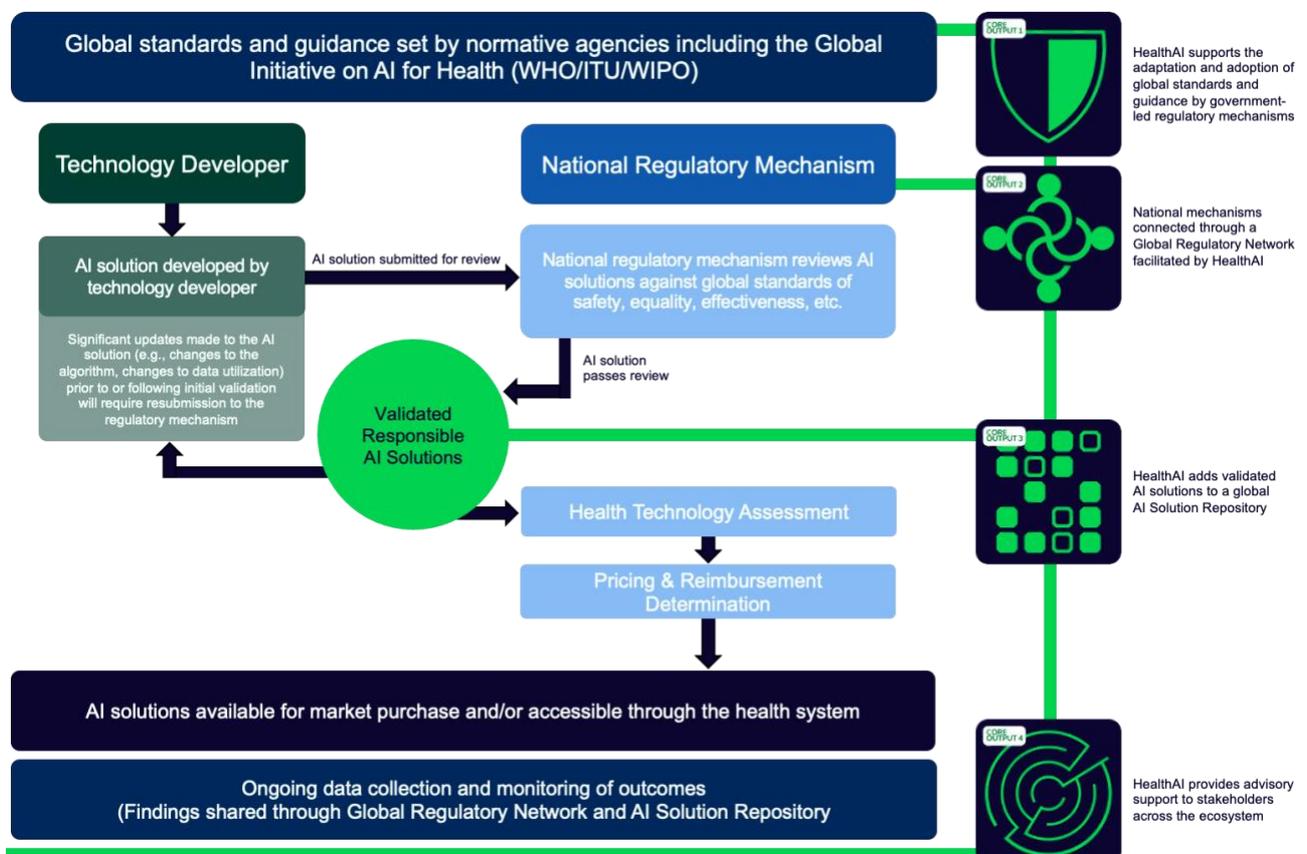


Figure 2 HealthAI's Core Outputs mapped against the regulatory process.





Core Output 1

National and regional validation mechanisms related to Responsible AI for health are built and certified.

HealthAI will support the establishment and strengthening of in-country, government-led regulatory mechanisms to accelerate the standards-based validation of AI technologies for use in health systems and services. HealthAI will provide its expertise to these regulatory mechanisms, support the implementation of existing auditing tools, and provide guidance on the use of data for AI solutions validation. Further, HealthAI will provide technical assistance to stakeholders with interest in regulatory guidelines and tools.

Regulatory Mechanisms

Guidance and standards provided by the WHO and other global normative bodies require in-country mechanisms for implementation. In many countries, there is no singular body that has the expertise required to evaluate AI solutions with both a health and technology perspective.

HealthAI will work directly with country governments to establish an in-country regulatory body, inspired by current models of medical validation and assessment conducted by regulatory agencies for medicines and medical devices. Depending on the contextual need, countries may opt for a stand-alone institution, a unit within a preexisting structure, or other models designed to fit the country's existing structures and processes. Each regulatory body will centralize the process of AI solution assessment and certification for the health sector, in partnership with relevant government ministries (e.g., Ministries of Health, Ministries of ICT). The regulatory mechanism will be managed through a government agency and will benefit from the tools, expertise, and capacity strengthening provided by HealthAI.

The role of these regulatory mechanisms will be to:

- Evaluate AI tools against safety and effectiveness standards to guarantee certified AI solutions meet the globally accepted standards set by WHO and other normative bodies;
- Provide formal pathways to market for AI solutions, including guidance and information on solution reimbursement and market authorization processes specific to the country;
- Create new revenue sources through validation fees and other cost recovery models;
- Establish and/or monitor data sharing regulation in support of data sovereignty;
- Coordinate with regulators in various other geographies;
- Foster an ecosystem of development, investment, research, and adoption of AI technologies in health systems and health care with a diverse group of experts and stakeholders.



Role of HealthAI

HealthAI will provide technical assistance to countries and other stakeholders interested in the development or support of regulatory mechanisms for Responsible AI in health. This technical assistance may take the form of training, workshops, bootcamps, consulting and advisory services, secondments, mentoring and coaching. HealthAI may provide technical assistance directly, engage third party experts, or support the acquisition of services through calls for consulting services. HealthAI will develop a fee-for-service model to allow for a range of service types including discrete scopes of work, project or initiative level support, and multi-stakeholder engagements.

HealthAI activities may include:

- Development of an investment case, a management, staffing and operational model, and advocacy resources to identify resources for the establishment of a regulatory mechanism;
- Facilitation of regulatory mechanism development through capacity strengthening, technical support, and provision of auditing and validation tools;
- Development of certification standards for country-level regulatory mechanisms to allow for ongoing evaluation;
- Initial certification and annual recertification of regulatory mechanisms to ensure compliance with predefined standards;
- Leveraging its Advisory Council to provide additional expertise from a diverse range of fields and perspectives;
- Ongoing technical support as needed.





Core Output 2

A global regulatory network for knowledge sharing and adverse events is established.

HealthAI will facilitate knowledge sharing between regulatory agencies in different countries and with members of the global regulatory community. This will serve as a mechanism for peer-to-peer learning, identification of best practices, and ultimately enable rapid notification of adverse events arising from an AI-driven health solution. Countries will share findings from AI solution assessments and certification, allowing other countries to streamline certification of the same technology and to identify AI solutions that require refinement or re-evaluation based on global standards.

Regulatory Network

The global regulatory network will serve as a connector between regulatory mechanisms at the country level, facilitating cross-country collaboration, providing a platform for real-time problem solving, and ensuring close working relationships between experts across the world. The information sharing and capacity strengthening enabled by the network will allow for greater coordination, rapid notification of adverse events, and wider adoption of regulatory best practices. The global regulatory network will also serve as a key partner to stakeholders at the global level, private sector companies, and others seeking to engage with a broad and inclusive group of regulators on the topic of Responsible AI for health.

The role of the global regulatory network will be to:

- Connect individuals working within country-level regulatory mechanisms with one another and with other stakeholders;
- Facilitate information exchange across geographies;
- Identify opportunities for coordination and collaboration both within the network and between the network and other stakeholders;
- Serve as an early warning mechanism for adverse events;
- Provide information about regulatory approvals for particular technologies;
- Provide capacity strengthening opportunities for network members.

Role of HealthAI

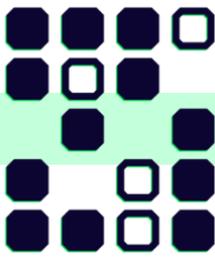
HealthAI will serve as the secretariat for the global regulatory network, managing the operational components of the network and facilitating information exchange between network members. Working at the global level, HealthAI will also help develop connections between the regulatory network and other stakeholders including the WHO, private sector companies, and other interested parties. HealthAI may also host training, workshops, and other capacity strengthening activities.



HealthAI activities may include:

- Secretariat support for the global regulatory network;
- Maintenance of a listserv and group discussion space for information sharing;
- Biannual meetings for members to engage with one another;
- Facilitation of cross-country learning exchanges.





Core Output 3

Global public repository of validated AI solutions for health is created.

HealthAI will create and maintain a global repository of validated AI solutions for health, allowing countries to evaluate solution options against local health needs. HealthAI will develop general use cases as an independent, neutral resource for countries to explore potential technologies and applications relevant to their specific context. This will also surface unmet needs as insight and inspiration for technology developers.

AI Solution Repository

An AI Solution repository will be developed to collate the results of in-country assessment, validation, and certification processes to provide a publicly available resource on AI solutions. This resource will help categorize AI solutions into different health domains and use cases, allowing countries to quickly understand the available tools for specific health system and service needs. In addition, the repository will capture findings from in-country regulation processes, noting AI solutions that have gained certification in specific countries as well as those that raised possible concerns or areas for improvement. This will help countries streamline their regulatory processes for technologies that have been previously assessed by other countries, as well as providing in-country regulators with supporting information and evidence. The repository will also be a resource for technologist and private sector companies interested in identifying gaps in existing solutions as they develop new technologies and applications.

The role of the AI Solution Repository will be to:

- List AI solutions that have been reviewed by in-country regulatory mechanisms, including supporting information and evidence;
- Categorize AI solutions by health system domain or function.

Role of HealthAI

HealthAI will build and maintain the global repository and serve as administrator to the platform. This will include setup and maintenance of the requisite software, gathering and collating information from in-country regulatory mechanisms, managing access permissions and security, and developing the data structures required to catalog AI solutions for health. HealthAI may opt to develop a tiered approach to access, providing a potential revenue source while also making detailed information available to a larger group of stakeholders.

HealthAI activities may include:

- Development and maintenance of a repository of validated AI solutions;
- Development of general use cases for AI in health, including validated AI solutions that apply to each potential use.





Core Output 4

Advisory support on policy and regulations is delivered.

Alongside partners like the Harvard-Charite Global Health Policy Lab, HealthAI will contribute to available global resources related to policy and legislation for Responsible AI for Health. By performing analyses of existing government regulation, HealthAI will provide technical guidance and insights into global trends and best practices. Both public and private stakeholders will be able to develop effective, contextually relevant strategies, policies, and regulations for Responsible AI in health. Through diverse stakeholder and citizen engagement, HealthAI will further the democratization of AI for health policy making to cultivate trust and improve inclusiveness.

Advisory Support

HealthAI may offer advisory services on a number of topics related to Responsible AI in health. The organization will ensure flexibility both in terms of topic areas and advisory models to support the widest possible range of stakeholders. Some examples of advisory services may include:

- **Guidance on Policy Development:** Many governments have begun the process of developing or expanding policies and regulations related to the use of AI in health. Some “early adopter” countries have made solid progress, but many other countries would benefit from HealthAI support and expertise. HealthAI will support country governments on the drafting, review, and implementation of national strategies, policies, and legislation related to the use of AI in health systems and services.
- **Guidance on Technology Development:** HealthAI will partner with private sector companies, technology developers, and other interested stakeholders to integrate global standards of ethics and equity at every stage of the development cycle. By working directly with technologists, HealthAI can help the private sector align with global standards early, creating a foundation of ethics and equity in their tools. This includes all development work up to and including validation. Strict protocols will be in place to ensure HealthAI remains neutral in the final validation of any AI solutions where consultation was provided, but this will provide an additional fee-for-service mechanism for the organization.
- **Guidance on Data Use:** HealthAI will also support the WHO and other normative bodies in the development and dissemination of standard guidance and incentive structures to promote privacy-protecting data sharing agreements and inclusive data platforms. With this guidance in place, HealthAI will be able to provide expertise to country governments, private sector developers, and other stakeholders on the responsible use of training data with a particular focus on Generative AI applications.



INTENDED IMPACT

Outlined in the Theory of Change, HealthAI has designed its strategic approach to maximize the impact of its direct work, as well as the impact that the results of this work will have on the global health ecosystem. Below is an initial summary of the organization's intended impact, including both broad contributions to the achievement of the Sustainable Development Goals and specific benefits for governments, partners, and other stakeholders.

Long-term Impact

HealthAI's work will contribute to the **improved health and well-being outcomes for all in alignment with the Sustainable Development Goals.**

AI has immense potential to augment existing approaches to health systems and services, and to create entirely new approaches. For example, AI solutions can support drug discovery through rapid analysis of chemical compounds for specific markers of treatment viability. AI solutions have begun to augment rapid review of imaging for the diagnosis of diseases like tuberculosis or certain retinal conditions. In service provision, chatbots powered by AI can answer patient questions, saving time for overburdened health workers. There are many other examples of how AI is already contributing to improve health and well-being—and as innovation continues this list will only continue to grow.

These applications of AI have the potential to address the three dimensions of Universal Health Coverage (UHC). They can reduce the costs of existing services by introducing efficiencies and reducing the burden on health workers. They can extend the reach of services to more segments of the population, including communities that are traditionally marginalized. They can also introduce new types of services, diagnostics, and treatments. Importantly, AI solutions can accelerate progress across all three areas in ways that other health or system strengthening interventions cannot.

A critical dimension to the work of HealthAI is centering equity in the regulation and adoption of Responsible AI solutions. The digital divide within and between countries threatens the collective progress of health and well-being, so it is imperative that any intervention related to the use of AI technologies in health address this divide. This requires a global perspective that can capture learnings from 'early adopter' countries and adapt them for other geographies. In addition, the contextualization of AI regulation and investments in in-country innovation will be required to protect against 'digital colonialism'.

Long-term Outcomes

HealthAI has identified three primary areas of impact over the long-term:

Increased access to safe, high-quality, effective, and equitable AI solutions

With strong in-country regulatory mechanisms, HealthAI will support the availability of Responsible AI solutions within health systems. By defining regulatory pathways for AI technologies, country governments will have greater confidence in the adoption and



implementation of AI solutions as well as support defining processes for market access. With a focus on safety, efficacy, and equity, these regulatory mechanisms will also provide a critical layer of protection for patients—ensuring that Responsible AI solutions accelerate improvements to health and well-being.

Specific areas of impact of HealthAI's work include:

- Ensuring all available AI solutions are safe for use, comply with quality standards, and effectively deliver their intended health outcome or system improvement;
- Providing stakeholders the necessary information on market access authorization and reimbursement processes;
- Protecting patient safety through an early-warning mechanism that alerts countries of AI-related adverse events across a broad geographic network;
- Streamlining information sharing between countries to propagate the availability of proven Responsible AI solutions.

Increased government revenue from regulatory activities

HealthAI is committed to developing sustainable regulatory mechanisms and processes to ensure long-term success. Working with country stakeholders and other funding partners, HealthAI will develop a business model for each regulatory body that addresses the specific economic context while also creating future revenue streams to support regulation.

Specific areas of impact of HealthAI's work include:

- Generating new sources of revenue for regulatory agencies and government budgets to allow for sustained funding for the regulatory mechanisms as well as additional investment capacity for the health system;
- Accelerating approval processes across countries through information sharing, leading to cost savings and bureaucratic streamlining.

Increased trust, investment, and innovation in Responsible AI solutions for health

Through strong, agile regulation of AI solutions, HealthAI will increase the trust and acceptability of these technologies. Regulation will help address concerns over data use, human centricity, and efficacy, ensuring that the Responsible AI solutions that make it to market are safe and effective. With these safeguards in place, the use of AI technologies in health systems and services will be more widely accepted by health workers and patients alike. There will also be greater incentive to invest in these technologies in both the public and private sectors. Finally, technologist and other innovators will have the guidance needed to develop new technologies that meet ethical and safety standards.

Specific areas of impact of HealthAI's work include:

- Protecting national data sovereignty and ensuring that all real-world data is collected, shared, and used under regulatory rules and through approved data centers;



- Ensuring compliance with WHO and other internationally-defined Responsible AI standards;
- Supporting validation processes that account for the usage of real-world data and enable feedback from civil society;
- Fostering an ecosystem that promotes investment in the research, development, and adoption of Responsible AI solutions for health.



OTHER STRATEGIC CONSIDERATIONS

While not directly represented in the theory of change, there are additional strategic considerations that will have bearing on the work of HealthAI. These considerations touch all aspects of the organization but merit specific attention here:

Inclusivity

Particular attention will be paid to ensure marginalized voices are included in the work of HealthAI. The organization is committed to identifying partners and contributors that represent the communities in which work is conducted, including intentional partnerships with civil society organizations of all types. Because of the focus of the organization, the intentional inclusion of two dimensions will be essential:

Gender: There is inherent tension between the technology sector and the health sector with regards to gender. The technology sector remains largely male-dominated, while the majority of health workers are female (though there is a gender imbalance between frontline health workers and high-level decision makers). In addition, the burden of healthcare often falls on women. Women also have unique health needs that are currently facing social and political machinations. Furthermore, gender bias is often present in datasets and can be exacerbated through training models used in the development and application of AI. It is therefore vital that girls, women and individuals with other gender identities are part of all regulatory, governance, and implementation processes and that special care is taken to identify and address biased data.

Youth: The impacts of policy and regulation can persist for decades, and the implications of such regulation will have greater long-term impact on youth. As digital natives, youth often set technology trends and have a unique perspective on how technologies like AI should be regulated and applied. They also have important insight into the ways technology shapes their daily lives and the novel ways they engage with it. To ensure the sustainability and long-term positive impact of the organization's work, youth must be an essential stakeholder.

Sovereign Data

HealthAI believes that individuals and communities should control their own data and have agency over how it is used. Therefore, the organization prioritizes data sovereignty in its technical and policy approach. To realize this, the organization will utilize data governance frameworks that give consent, control, and protection to the individuals and communities to whom the data belongs. Storage and processing of data will be done locally whenever possible, while exploring tools like differential privacy and federated learning to garner insights without compromising raw data. Adherence to regulations like the Global Data Protection Regulation (GDPR), enacting regional data repositories, and ensuring ethical data stewardship of the organization and its parts will be central to its mission.



Interoperability, Cloud Computing, and Other Technical Considerations

Throughout its work, HealthAI will apply global standards and best practices related to the underlying technologies used. This includes the use of open standards, APIs and interfaces that allow for technical interoperability. The use of proprietary software will be intentionally monitored and limited where possible to prevent the concentration of power into singular software systems. The organization will use diverse platforms and encourage portability and adherence to standards to prevent lock-in. Finally, the organization will work with in-country stakeholders to understand the acceptability and efficacy of cloud computing systems for data storage and analysis, in accordance with local laws and the organizational view on data sovereignty and protection.

Generative AI

The recent release of several commercial AI platforms has led to particular concern around the ethics and regulation of Generative AI technologies. Technologies like [ChatGPT](#), [DALL-E2](#), and [Google Bard](#) use natural language processing to create text and images from prompts. These tools quickly gained broad notoriety which has led to important conversations about the impact of Generative AI on society.

Concerns have been raised about the training data used for such tools—both from a data bias perspective as well as from an intellectual property perspective. Biased data has influenced the outputs of Generative AI tools, underlining the issue of entrenched bias within datasets even when applied in seemingly neutral ways. Questions of intellectual property and unethical use of datasets have grown prominent as individuals and communities see their work used in the training of Generative AI algorithms, often without consent or notice.

While the impact of Generative AI technologies on health systems and services is still unknown, it is paramount that all ethical concerns are addressed. As such, HealthAI will work with its partners to assess Generative AI tools before they are brought to market, define terms and authorizations for clinical trials involving Generative AI, and support technology developers to embed Responsible AI principles and standards throughout the lifecycle of development.

HealthAI will also stay abreast of developments in Generative AI in order to work directly with global normative bodies and country governments on the appropriate governance mechanisms for Generative AI. This rapidly evolving space will require ongoing monitoring and evaluation to reduce potential harm and prioritize ethical approaches.

