McKinsey & Company

Social Sector Practice

Public health's inflection point with generative AI

Public health organizations can responsibly use generative AI to improve service delivery, bolster outbreak preparedness, accelerate R&D, and enhance health outcomes for communities.

by Steve Davis, Saumya Singh, Nikhil Srinidhi, and Matt Wilson

February 2024

Throughout history, scientific innovation has propelled major advancements in public health, from the development of the world's first vaccine in 1796[1] to the discovery of penicillin in 1928.[2] Most recently, the public health ecosystem mobilized and executed a massive effort to develop a safe and effective COVID-19 vaccine in roughly 300 days (compared with a typical timeline of ten years or more[3]) and achieve coverage of more than 70 percent worldwide in about two years.[4]

Now, a technological innovation—generative AI (gen AI)—could lead to powerful new advancements in public health and healthcare. Gen AI has the potential to improve a public health organization's ability to do the following:

— *Engage with patients and people in various communities.* Examples include more-efficient personalized inquiry responses, enhanced chatbot performance, and improved access to quick, reliable information.

— *Synthesize concise insights.* Gathering large volumes of information and distilling insights can improve decision making and save time, for example, to accelerate writing and reviewing of grants. Gen AI is expected to lead to productivity gains of 15 to 20 percent in pharmaceutical and medical-product R&D, and a substantial portion of these gains could come from automatically drafting clinical study reports that accelerate regulatory submissions.[5] This could translate to savings of approximately $100 million to $200 million in worldwide tuberculosis R&D alone.[6]

— *Create more tailored content.* For example, using gen AI to draft and edit text, images, and other media could lead to a 5 to 15 percent savings in public-engagement spending through more-efficient content generation.[7] If applied to 2021 US federal government spending on vaccine-related campaigns, this could have led to annual savings of between $85 million and $300 million, according to McKinsey analysis.[8]

— *Produce efficiencies in generating and updating code.* Generative AI–based tools enable tremendous productivity gains by helping developers expedite manual and repetitive work, jump-start first drafts of code, support code completion, accelerate updates to existing code, and more readily tackle new challenges (for example, by helping developers rapidly brush up on an unfamiliar code base or language).

Gen AI is increasingly being applied to many functions in the pharmaceutical, medical-device,[9] and healthcare fields.[10] Beyond healthcare entities (for example, care delivery organizations, labs, and payers) and pharmaceutical, biotechnology, and medical-device manufacturers, a wide variety of public health stakeholders could experience a positive effect from gen AI. This includes government public health–related institutions and policy makers (for example, federal, state, local, tribal, and territorial agencies, as well as health information exchanges); donors and multilateral institutions; implementation partners (nongovernmental, community-based, and faith-based organizations); and researchers (universities and product development partners).

However, a concurrent focus on risk management, change management, and skills building is necessary to not only implement the right gen AI pilots but also sustain value as the pilots scale. Considering gen AI's distinct potential for public health stakeholders and its relevance to other

[1] "Innovations through public health history," *Harvard Public Health Magazine,* Spring 2022.
[2] Ibid.
[3] "100 Days," CEPI, accessed January 29, 2024.
[4] Of the world population, 70.6 percent have received at least one dose of a COVID-19 vaccine. However, coverage varies substantially across high- and low-income countries. "Coronavirus (COVID-19) vaccinations," Our World in Data, accessed January 12, 2024.
[5] "The economic potential of generative AI: The next productivity frontier," McKinsey, June 14, 2023.
[6] Based on 2021 figures for annual tuberculosis R&D spending from *Tuberculosis research funding trends, 2005–2021,* Treatment Action Group, December 6, 2022.
[7] "The economic potential," June 14, 2023.
[8] Exact savings potential might vary depending on breakdown of applicable spend.
[9] "The economic potential," June 14, 2023.
[10] Shashank Bhasker, Damien Bruce, Jessica Lamb, and George Stein, "Tackling healthcare's biggest burdens with generative AI," McKinsey, July 10, 2023.

related industries, we have identified four public health domains with clear use cases (exhibit): service delivery and operations; resilience, preparedness, and outbreak response; product R&D; and foundations for public health action, including data and technology enablers, talent enablers, and policies and standards. This article outlines potential use cases for gen AI in each of these public health domains, considerations for organizations adopting gen AI, risks specific to public health to mitigate and manage, and tips for getting started. (For an overview of gen AI, including how to use it responsibly, see McKinsey's article "What every CEO should know about generative AI."[11])

Gen AI use cases in public health

Public health leaders could benefit by learning from innovation and investment in adjacent industries while accounting for risk and implementation considerations specific to public health. Life sciences and private sector healthcare delivery, two sectors adjacent to public health, have already begun to seize the gen AI opportunity and consider how to mitigate its inherent risks. Gen AI is expected to deliver a productivity lift of 3 to 5 percent of global industry revenue, or $60 billion to $110 billion globally, for the pharmaceutical and medical-product industries.[12] Gen AI will likely contribute to a bigger productivity impact in healthcare than in many other industries.[13] Private payers, hospitals, and physician groups have well-articulated use cases and budding adoption.[14]

New technology offerings and industry innovations are supporting these efforts.[15] For example, Epic and Microsoft have partnered with several health systems—including UC San Diego Health, UW Health in Wisconsin, and Stanford Health Care—to develop a tool to automatically draft responses to patient messages in online portals.[16] The rapid pace of innovation in gen AI for healthcare and life sciences suggests that the global public health ecosystem can quickly adapt and learn from early use cases. Recent public health–specific guidance on design, development, training, validation, retraining, and deployment of AI from public health institutions such as the World Health Organization (WHO)[17] and national and international authorities can help ensure safe scaling of these early use cases. Released in December 2023, the European Union AI Act and the US Department of Health and Human Services' (HHS) rules regulating AI transparency have added more guardrails to public health organizations' use cases.[18]

[11] "What every CEO should know about generative AI," McKinsey, May 12, 2023.
[12] "The economic potential," June 14, 2023.
[13] Ibid.
[14] "Tackling healthcare's biggest burdens," July 10, 2023.
[15] "Embracing generative AI in health care," *The Lancet Regional Health – Europe,* July 2023, Volume 30.
[16] Heather Landi, "HIMSS23: Epic taps Microsoft to integrate generative AI into EHRs with Stanford, UC San Diego as early adopters," *Fierce Healthcare,* April 17, 2023.
[17] *Regulatory considerations on artificial intelligence for health,* WHO, 2023.
[18] Ibid.; Ben Leonard, "HHS strikes first in regulating the new artificial intelligence tools in health care," POLITICO Pro, December 13, 2023; "EU AI Act: first regulation on artificial intelligence," European Parliament, updated December 19, 2023.

Public health leaders could benefit by learning from innovation and investment in adjacent industries.

Exhibit

Potential use cases for generative AI in public health can be categorized into four key domains.

Domain: Service delivery and operations			
Supply chain and financing	Accelerate grant writing by improving speed and quality of research, drafting, and tailoring of grant applications	Aggregate nontraditional sources of qualitative data (eg, social media and speech-to-text feedback surveys) to identify issues in near real time (eg, shortages of malaria bed nets)	Capture insights from unstructured data to inform models of risk scenarios and mitigation strategies, structurally improving investor confidence in global health innovation opportunities
Frontline healthcare worker enablement	Develop individualized training content and simulations for physicians, nurses, and health-care workers tailored to needs (eg, region, area of specialty, and demographics served)	Synthesize care coordination notes across patient's care team (eg, nurses, primary care physicians, doulas, and midwives involved in antenatal care)	Enhance incomplete or low-quality patient data (eg, build a visual of patient tissue or organs from an MRI scan) to aid in the diagnostic process
Community and patient engagement	Generate campaign communications (eg, images and text) tailored to local language and context	Provide enhanced telemedicine and point-of-care support through gen AI–enhanced chatbots to improve healthcare in rural areas (eg, antenatal care and MNCH[1] outcomes)	Generate personalized treatment plans and patient guidance based on individual patient data, delivered in a format tailored to lifestyle and literacy level
Domain: Resilience, preparedness, and outbreak response			
Predictions and preparations for threats	Summarize trends and latest developments from scientific literature search to improve existing models and shape new approaches to various conditions	Automatically generate content for response frameworks and scenarios for table-top exercises to test and improve readiness	Synthesize learnings from after-action review reports to generate best practices and inform response frameworks for future public health threats
Detection and monitoring of potential threats	Leverage image processing across multiple data sources (eg, thermal mapping, satellite images) to generate insights and flag aberrations	Identify early-warning factors or signals (eg, synthesize emergency department visit notes, triage notes, and signals from social media for syndromic surveillance) to enhance syndromic surveillance and ensure rapid awareness of potential threats	
Investigation of and response to threats	Address operational challenges in the first-responder community (eg, assist emergency responders in quickly generating tailored communications in different languages from existing collateral) to improve timeliness and effectiveness of aid	Generate and enhance tailored health guidance that health officials can review and disseminate to the public during an outbreak	Improve decision support across the public health ecosystem through answering questions about and extracting content from daily after-action review reports, medical knowledge, electronic medical records, and other knowledge sources

[1]Maternal, newborn, and child health.

McKinsey & Company

Exhibit (continued)

Potential use cases for generative AI in public health can be categorized into four key domains.

Domain: Product R&D			
Discovery and research	Synthesize insights from translational research for disease elimination (eg, monitoring uptake of a disease preventive therapy across different LMIC[2] contexts or across organizations)	Understand and summarize biomedical research papers to identify papers that are most relevant for global health	Enhance indication finding and drug repurposing through representational learning from patient histories
Development	Use deep learning to design protein molecules and advance new protein-based vaccines for diseases (eg, leveraging data from vaccine or drug libraries and medical legal review)	Integrate, clean, and synthesize data across trial sites to improve data quality and expedite analysis or summaries for dossier submission	Automatically generate case report forms based on digital protocol, and run AI-based automated data quality checks and cleaning
Regulatory approval	Generate documentation needed for regulatory approval using natural language processing and through analyzing prior submissions and synthesizing real-time drug development or clinical trial data (eg, auto-drafting CSRs,[3] MLRs,[4] IND,[5] and NDA[6] submissions)		
Domain: Foundations for public health action			
Data and technology enablers	Create synthetic data sets to test and develop solutions without risking privacy breaches or providing personally identifiable information	Automate administrative tasks to reduce manual burden (eg, HR chatbot for self-service, processing large amounts of unstructured data)	Correct and troubleshoot code to accelerate engineering productivity by more than 50 percent; accelerate pair programming with generative AI "programmer" (eg, Codex, AlphaCode, Codegen, and Copilot)
Talent enablers	Enable engineering talent to search for creative solutions, decipher unfamiliar syntax, and find the correct algorithm when unsure of how to proceed	Create personalized learning journeys and recommend learning resources based on current skills and future goals	
Policies and standards	Analyze and synthesize policy documents across countries to determine best practices and tailor global health interventions to regional contexts	Simulate various scenarios and predict the impact of potential policies and standards to aid decision making	Generate draft data policies and standards (eg, by synthesizing trends and impact of prior or other policies and standards across different geographies)

[2]Lower- and middle-income countries.
[3]Clinical study reports.
[4]Medical legal reviews.
[5]Investigational new drug.
[6]New drug application.

McKinsey & Company

Streamline service delivery and operations

Gen AI could streamline service delivery and operations for frontline public health and community health workers, reducing administrative burdens, optimizing supply chain efficiencies, and ensuring effective, tailored engagement with diverse communities. Some of these gen AI applications, when adopted in combination with other automation technologies, could lead to productivity increases of 15 to 40 percent for healthcare workers.[19] This could unlock about seven billion to 26 billion hours of work capacity for the approximately 25 million to 35 million healthcare workers in 127 lower- and middle-income countries (LMICs).[20] Gen AI–enhanced tools that effectively build on innovations in natural language processing could be used to do the following:

— *Engage the public in seamless conversations about important public health topics (for example, finding nearby locations for vaccine boosters).* Of the 48 recipients of Grand Challenges Catalyzing Equitable Artificial Intelligence (AI) Use grants in August 2023, 66 percent deployed gen AI–based projects that involved some form of chatbots.[21]

— *Develop individualized training content and simulations for physicians, nurses, doulas, and other healthcare workers tailored to needs.* For example, training could be tailored by region, languages spoken, area of specialty, and demographics served.

— *Accelerate grant writing by improving the speed and quality of researching, drafting, and tailoring grant applications.* For a demo prototype developed, it was estimated that 15 to 25 weeks could be saved by one country's grant-writing team using gen AI technology to draft select sections of one specific grant.[22] The technology could synthesize background documents such as national strategic plans to prepare a first draft of a context section of a grant.

— *Assess the effectiveness of past global health campaigns and automatically generate content and material for future ones.* For example, gen AI applications could help create tailored messages for a large-scale public health effort such as the Zero Malaria campaign, which aims to increase malaria awareness and prevention with an eye toward eliminating the disease completely. In 2021, the campaign reached about 1.4 billion people and generated more than 24.0 million digital interactions.[23] Quickly created, tailored messaging—supported by gen AI and a human in the loop to check the output—could enhance the campaign's reach and engagement and further the goal of eliminating malaria.

Bolster resilience, preparedness, and outbreak response

Gen AI, and AI more broadly, can support a broad set of use cases across public health resilience, preparedness, and outbreak response. Early-adopter use cases have included using GPT-4 to analyze and interpret epidemiological data, clinical records, and research literature to help predict outbreaks or accelerate medical response.[24] Meanwhile, farmers are using AI-optimized tools to upload photos of sick animals to receive advice

[19] "The economic potential," June 14, 2023.

[20] This includes community health workers, medical doctors, nursing and midwifery personnel, and dentist roles, according to "Global Health Workforce statistics database," WHO, accessed January 12, 2024. The lower bound (seven billion work hours) is the potential impact of gen AI alone, and the upper bound (26 billion work hours) is the potential impact of gen AI when coupled with other automation technologies.

[21] McKinsey analysis of "Awards," Global Grand Challenges, accessed January 12, 2024.

[22] This assumes the grant-writing team consists of 12 to 20 members (five high-level staff members and seven to ten junior and consultant staff members) working at different capacities (15 to 20 percent and 60 percent, respectively). Based on McKinsey analysis and expert interviews with ministries of health participating in the grant-writing process.

[23] Several recipients of Grand Challenges Catalyzing Equitable Artificial Intelligence (AI) grants (for example, the Child Health Research Foundation in Dhaka, Bangladesh) are using epidemiological data, demographic data, and medical records collected from a catchment area to adapt published code with the help of GPT-4 to predict the impact of introducing specific vaccines in Bangladesh. "Awards," accessed January 12, 2024.

[24] Joseph Mulabbi, "Using artificial intelligence to predict disease emergence in Uganda," Global Grand Challenges, July 11, 2023.

and assess risks, while hot spot maps informed by AI help health officials better monitor potential zoonotic disease outbreaks.[25]

Gen AI could also be applied in applications that closely align with broader public health resilience strategies, including the following:

— synthesizing unstructured text from daily after-action review reports across vaccination sites during an ongoing vaccination campaign

— generating synthetic data[26] to augment training of epidemiological models and improve forecasting of public health threats

— processing images from multiple data sources (for example, thermal mapping and satellite images) to flag aberrations and help generate simulations to enhance wastewater surveillance to predict potential future illnesses, for example[27]

When applied at scale and in a public health emergency, gen AI use cases may need to be further evaluated for risks. Using gen AI for use cases related to outbreak response may also require further emphasis on transparency and documentation as well as rigorous and continuous evaluation of data quality—such as in the use case above on the synthesis of daily after-action review reports.

Accelerate product R&D

Life sciences companies are investing hundreds of millions of dollars into AI (including gen AI) across the R&D value chain to speed up time to market. For example, gen AI use cases related to product R&D can include the following:

— automating the generation of draft documents for submissions that are 80 percent correct, including clinical study reports and other common technical documents from the trial protocol and tables, lists, and figures

— helping clinical trial managers address operational bottlenecks in ongoing trials more quickly and easily

— helping medical writers draft higher-quality health assessment questionnaire responses informed by insights from regulatory guidance

— supporting indication identification and prioritization for existing and novel mechanisms of action

Global public health stakeholders have an opportunity to build on such advances to more effectively tackle diseases with a major global health burden such as HIV, malaria, and tuberculosis. For example, gen AI could be used to synthesize insights from translational research[28] across different LMIC contexts, understand and summarize biomedical research papers to identify those most relevant for global health (by using BioGPT, for example), and accelerate the generation of clinical study reports. Moreover, gen AI use cases in product R&D could support collaborative global public health efforts such as the 100 Days Mission, which seeks to reduce the time it takes to develop a vaccine for a new pathogen.[29]

---

[25] Ibid.; Zameer Brey, "Announcing $5 million in Grand Challenges AI grants to spur local innovation for global good," Bill & Melinda Gates Foundation, August 9, 2023.

[26] Synthetic data is data that has been artificially created by computer algorithms, unlike data that has been collected from real-world events.

[27] "Tackling healthcare's biggest burdens," July 10, 2023.

[28] Translational research is the effort to translate results in basic research to results that directly benefit people.

[29] The 100 Days Mission is a global commitment by the Coalition for Epidemic Preparedness Innovations (CEPI) to accelerate the development of, and access to, safe and effective vaccines against new pathogens within 100 days of identification, building on the ability to do so for COVID-19 (for more, see *100 Days Mission implementation report - 2023,* International Pandemic Preparedness Secretariat, January 24, 2024). The Medicines and Healthcare products Regulatory Agency in the United Kingdom granted emergency authorization for the Pfizer-BioNTech vaccine Comirnaty, a COVID-19 mRNA vaccine, on December 2, 2020, 307 days from January 30, the day WHO declared COVID-19 a public health emergency of international concern. The timeline shown by CEPI is about 326 days (for more, see the 100 Days CEPI website).

Build foundations for public health action

Gen AI could strengthen public health action and decision making more broadly by unleashing the full potential of data, technology, and technical talent and accelerating standards adoption. For example, attracting, retaining, and training high-quality technical talent is often difficult for government public health agencies due to budget constraints.[30] The use of gen AI could lead to productivity gains of up to 50 percent for developers and support their training and development.[31] Other industries are already using gen AI to assist developers with writing and translating code, which public health organizations could build upon to accelerate the development of technology.

Early-adopter use cases can include the following:

- Accelerate pair programming with gen AI programmer assistance and code-completion tools (such as OpenAI Codex, AlphaCode, Codegen and GitHub Copilot). For example, in one study, software developers using Copilot completed tasks 56 percent faster than those not using the tool.[32]

- Integrate AI capabilities into existing health software systems to improve their usability or make them more efficient and faster.

- Automate administrative tasks to reduce manual burden, such as by processing large amounts of unstructured data to support repetitive HR tasks.

Five steps to get started with gen AI in public health

Many public health organizations recognize the future potential for gen AI in their missions and want to start learning and investing, but they struggle to figure out where to start, how to manage risks, and how to navigate this fast-moving field. Like their business counterparts, public health leaders across government, multilateral, and nongovernment organizations can benefit from a sequenced and structured approach to gen AI, tailoring as needed based on the specific context and organizational readiness.[33]

Identify high-impact use cases that advance existing strategic priorities

Broadly considering "art of the possible" gen AI use cases—those that could fundamentally enable the broader public health mission of the agency and its closest collaborative partners—is a worthwhile exercise. However, it is also important to quickly home in on a small subset of "lighthouse" use cases (perhaps one to three) that could most effectively advance existing near-term strategic priorities.

Budgets in public health must cater to multiple priorities and potential annual fluctuations in funding. For example, the US Centers for Disease Control and Prevention (CDC) stood to reduce its funding by about $1.3 billion because of a federal debt ceiling negotiation in 2023,[34] and between 2010 and 2020, US spending on state public health departments fell by 17 percent.[35] Consequently, as such agencies make difficult choices around priorities and budgets, focusing on how gen AI pilots lead to measurable efficiencies and improvements in delivery of strategic priorities could help build a business case for at-scale gen AI implementation. Public health agencies and donors also worry that costs could inhibit initial work on gen AI. Indeed, building a foundation model from scratch could be cost-prohibitive for most public health agencies and donors. Instead, public health organizations could choose to leverage advancements in industry to get started while managing costs and measuring public health impact. As of January 2024, most early adopters in public health were choosing to be "shapers"—integrating large language models with internal data and systems to generate more customized results—or "takers" for specific niche

[30] Anita Dutta, Nora Gardner, Megan McConnell, and Angela Sinisterra-Woods, "Transforming public sector hiring with data-enabled talent 'win rooms,'" McKinsey, July 5, 2023.
[31] "Unleashing developer productivity with generative AI," McKinsey, July 27, 2023.
[32] Peter Cihon et al., "The impact of AI on developer productivity: Evidence from GitHub Copilot," arXiv, February 13, 2023.
[33] "What every CEO should know," May 12, 2023.
[34] Brenda Goodman and Deidre McPhillips, "CDC facing major funding cuts, with direct impact on state and local health departments," CNN, July 14, 2023.
[35] William Eger and Margaret House, "Confronting a legacy of scarcity: A plan for reinvesting in U.S. public health," Stat, June 28, 2021.

use cases such as coding, in which they use publicly available models (with little or no customization) through a chat interface or an API.[36]

Further, other feasibility considerations—such as the effort needed to mitigate risks associated with using gen AI—can help organizations systematically winnow the list of potential gen AI use cases. For example, one large multilateral public health organization is developing a proof of concept for a gen AI–powered knowledge base for its internal users to easily query and synthesize unstructured and disaggregated information stored in disparate document repositories, and filter by geography, type of health emergency, resources, and constraints. Given the comparatively higher risks of implementing gen AI in anything that is public- or patient-facing, the organization chose to start by prioritizing an internal-facing use case.

Assess organizational readiness to implement gen AI

Public health organizations can assess whether their current technology and capabilities could initially support a proof of concept and, eventually, a fully scaled implementation. To develop gen AI use cases and fully scale them, organizations must have the right data, technology, and talent capabilities. Gen AI offers unique advantages compared with other forms of AI because it can work on unstructured data, which is often abundant in public health settings. So, while it might be relatively easy to get started with a gen AI pilot from a data perspective, public health organizations still need to ensure there is appropriate capability building and risk mitigation to support an enterprise-wide rollout (such as training users across the organization on prompt engineering).

For example, one large US public health agency is currently focusing its internal AI task force to define the staff and technical infrastructure it needs to put in place to make the most of its gen AI investments. As of September 2023, several US state governments (such as California, Kansas, and Pennsylvania) had issued executive orders governing the use and adoption of gen AI, including by state public health agencies.[37] This included, for example, guidelines on creating sandbox infrastructure environments for collaborating with vendors on gen AI pilots and provisions for training materials on gen AI usage and certifications for employees.

Consider risks unique to implementing gen AI in public health

Developing and deploying gen AI comes with several challenges and risks that need to be evaluated and managed as part of any organization's journey. These include, but are not limited to, fairness and bias of models; privacy, intellectual property infringement, and regulatory compliance concerns; the interpretability and usability of models; the need for human oversight and accountability to override incorrect recommendations; and performance inaccuracy stemming from misinformation or hallucinations, in which the gen AI model presents an incorrect response based on the highest-probability response.[38] Each of these general risks introduces specific considerations in the public health context. For example, gen AI models could hallucinate and produce content that is biased, factually wrong, or illegally scraped from a copyrighted source. To combat this, appropriate training for the human or humans in the loop will be needed, including on how to check the model output and recognize when to use or discard the output generated.

Every organization using these new tools has a responsibility to identify the risks and adopt strategies and approaches to mitigate them, but this work need not deter public health organizations from getting started with gen AI uses cases that support their missions.

Consider policy, health equity, and security implications specific to gen AI in public health

Early adopters can look to major stakeholders in the global public health ecosystem for guidance

[36] For more on costs associated with gen AI, see "Technology's generational moment with generative AI: A CIO and CTO guide," McKinsey, July 11, 2023.

[37] Sage Larrazo, "Will U.S. states figure out how to regulate AI before the feds?," Yahoo Finance, September 29, 2023; "AI chatbots are invading your local government—and making everyone nervous," *WIRED,* September 11, 2023; Skylar Rispins, "Kansas' generative AI policy sets flexible guardrails," StateScoop, September 8, 2023.

[38] For more, see "Technology's generational moment," July 11, 2023.

on the responsible use of AI—gen AI in particular. WHO lays out six core principles aimed at protecting autonomy and safety, providing transparency and accountability, and ensuring inclusiveness and equity, among other goals.[39] *Artificial intelligence in global health,* a report from the Rockefeller Foundation and USAID, similarly emphasizes the importance of privacy, security, and transparency.[40]

In addition, WHO's recently published paper *Regulatory considerations on artificial intelligence for health* details several precautions public health organizations can take and outlines actions countries can consider in regulating the use of AI in health.[41] These include externally validating data, documenting the intended use of applications, complying with previously issued guidance,[42] conducting a premarket assurance test of the model's safety and performance, and doing real-world performance monitoring. This is a quickly evolving, innovative field. Public health organizations can take steps to continually monitor evolving guidance on management of risks from WHO and regional regulations such as the EU AI Act and HHS rules about transparency.

Stakeholders could take steps to address the particular challenges of gen AI adoption—including cultural barriers and infrastructure challenges such as limited broadband and internet access—to help ensure effective and equitable adoption. For example, a recent *Nature* article proposes grounding a gen AI research agenda in the African context with locally tailored, evidence-based policies.[43]

When defining and training gen AI models, public health organizations could ensure the proper inclusion of underrepresented populations—such as older populations, tribes, or specific vulnerable populations they serve. They could also take into account the specific risks, cultural sensitivities, and local healthcare infrastructure and policies that might affect the health outcomes in any particular use case.

Invest in partnerships

Public health agencies do not have to work on their own to reap the benefits of gen AI. Partners from industry and academia could help design, test, and scale initial gen AI use cases. For example, the UK government announced partnerships with industry partners DeepMind, OpenAI, and Anthropic to support the safe adoption of AI.[44] In the United States, several federal public health–related agencies, such as the National Institutes of Health, participate in coalitions that aim to define guidelines for and health equity implications of health organizations adopting AI.[45]

Gen AI is ushering in an era of boundless possibilities for public health. From optimizing care delivery and better preparing for future threats to accelerating the development of therapeutics for high-burden diseases, gen AI could be a transformative force, enhancing our ability to tackle challenges in global, national, state, and local public health systems. Ultimately, it could lead to improved health outcomes for communities.

[39] *Ethics and governance of artificial intelligence for health,* WHO, June 28, 2021.
[40] *Artificial intelligence in global health: Defining a collective path forward,* Rockefeller Foundation and USAID, April 2019.
[41] *Regulatory considerations,* 2023.
[42] Melanie J. Calvert et al., "Reporting guidelines for clinical trial reports for interventions involving artificial intelligence: The CONSORT-AI extension," *Nature Medicine,* September 2020, Volume 26, Number 9.
[43] Rachel Adams et al., "A new research agenda for African generative AI," *Nature Human Behaviour,* October 2023, Number 7.
[44] Laurie Clarke, "OpenAI, DeepMind will open up models to UK government," Politico, June 12, 2023.
[45] "Providing guidelines for the responsible use of AI in healthcare," Coalition of Health AI, accessed January 16, 2024.

Steve Davis is a senior adviser in McKinsey's Seattle office, **Saumya Singh** is an associate partner in the New Jersey office, **Nikhil Srinidhi** is an associate partner in the Berlin office, and **Matt Wilson** is a senior partner in the New York office.

The authors wish to thank Marie-Renée B-Lajoie, Shriya Beesam, Sam Bourton, Ankit Fadia, Jessica Kahn, Nafees Khan, Jessica Lamb, Anthony Nguyen, Mark Nichols, Keerthana Sundar, Wouter Versluijs, Chaitanya Adabala Viswa, and others in QuantumBlack, AI by McKinsey and the McKinsey public health domain for their contributions to this article.