In collaboration with ZS

WORLD ECONOMIC FORUM

Patient-First Health with Generative AI: Reshaping the Care Experience

WHITE PAPER

JANUARY 2024

Images: Getty Images, Unsplash

Contents

Foreword

Shyam Bishen
Head, Centre for Health and Healthcare, World Economic Forum

Pratap Khedkar
Chief Executive Officer, ZS

In 2022-23, the World Economic Forum and ZS collaborated to create a report outlining the most promising use cases for artificial intelligence (AI) in healthcare, the stiffest barriers to adoption, a framework for building trust and adoption for AI solutions, and a call to action for stakeholders across the healthcare ecosystem and beyond. Shortly before the paper was published, OpenAI released ChatGPT, unleashing generative AI into the public discourse and opening fresh new possibilities for using AI to tackle healthcare's most intractable problems.

For example, large language models (LLMs) trained on vast troves of unstructured data can create original, human-like text based on minimal prompts. Given that somewhere around 80% of all medical data is unstructured, these models hold strong promise to address global healthcare challenges for which "traditional" predictive AI techniques are ill-suited.

This white paper looks specifically at generative AI's promise to help empower patients. It draws on interviews with digital and innovation leaders across healthcare, technology and other sectors, who have shared their views on the most promising use cases for patient engagement, their biggest barriers to adoption and key enablers for using this new modality to improve global health outcomes.

Executive summary

Generative AI offers a powerful new modality to address global healthcare challenges, but it's not a panacea.

The COVID-19 pandemic exacerbated numerous existing global healthcare problem areas – among them a severe shortage of healthcare workers,[1] widening health disparities and existential strain on health system finances. The World Health Organization estimates the global shortage of healthcare workers today at 15 million (including 10 million doctors),[2] a figure projected to decline to 10 million by 2030. Around 95% of the shortfall is in low- and middle-income countries (LMICs).[3]

This poses a significant threat to global health, as there simply are not enough healthcare providers to care for the volume of patients in need – let alone to provide the type of preventive care required to extend the health span for global populations. Powered by rapid advances in machine- and deep-learning techniques, artificial intelligence has been hailed as the answer to a broad array of endemic healthcare challenges.[4] However, significant barriers remain, as examined in the 2023 joint Forum and ZS report *Scaling Smart Solutions with AI in Health: Unlocking Impact on High-Potential Use Cases*.

While generative artificial intelligence (AI) is not a healthcare panacea, it does offer stakeholders a powerful new modality to address healthcare challenges that are difficult for existing predictive AI techniques to address. Yet, predictive AI's limitations are not the only factor spurring the development and use of generative AI in healthcare. The consumerization of health – long under way but hastened by the pandemic – is a powerful and growing force catalysing healthcare transformation. From the public's growing reliance on the internet to gather health information to US healthcare pricing transparency laws and policies across Europe to make healthcare more accessible, patients everywhere demand more convenient and accessible healthcare experiences.

Many healthcare companies and other stakeholders have struggled to keep pace with changing patient expectations for how they engage with the health system, as described in the *2024 ZS Future of Health Report*. This paper focuses on the promise of using generative AI to advance direct patient engagement, a use case with high potential to alleviate health system burden, reduce healthcare provider burnout and improve patient experiences and outcomes.

Key takeaways include:

Generative AI holds significant promise to aid healthcare consumers across their full health journeys – from dispensing reliable health information across regions and cultural contexts to ensuring patients get the appropriate level of care and helping them to manage their conditions.

The biggest barriers to the adoption of patient-facing generative AI solutions for healthcare are mistrust among doctors and the public, holes in the data foundation and scalability in low-resource environments.

Encouraging generative AI adoption in healthcare depends on instilling models with empathy and domain-specific knowledge, mitigating bias by connecting the data ecosystem and continuously fine tuning models, keeping humans in the loop and developing more cost-effective ways to train and run multi-modal foundation models.

1 Three broad categories of generative AI use cases for patient engagement

Generative AI can be applied in diverse ways in healthcare to empower patients and alleviate health system burden.

Generative AI can augment providers by offering more intuitive, human-like conversations with patients and caregivers that can help inform care decisions and spur action.

After speaking with leaders responsible for vetting and implementing generative artificial intelligence (AI) solutions across healthcare and technology, three themes emerged. One, using AI and other emerging technologies to create value sits at the top of the corporate agenda. Two, while the healthcare field is rife with experiments, there is limited agreement on which use cases are ready for deployment and hold transformational potential. Three, there is no consensus on how to group generative AI use case types.

Based on a review of the literature and qualitative interviews with digital and innovation leaders across healthcare, technology and other sectors, use cases fall into three general categories:

- **Productivity boosters:** With minimal prompting, generative AI models can create images and text sufficiently human-like as to automate many of today's manual tasks – everything from transcribing doctor-patient visits and drafting emails to summarizing clinical studies and dispensing general health information. Generative AI can also be used to help create a more robust data foundation by automating data management and extracting insights from unstructured data, a task that was extremely time-consuming, cumbersome and inefficient using earlier AI techniques.
- **Insights generators:** Generative AI enables real-time analysis of both structured datasets (such as health data) and unstructured datasets (such as physician notes), offering significant promise as a "co-pilot" to help healthcare providers quickly analyse information and interpret results. These tools also can help accelerate data integration across datasets to help doctors, nurses and health systems gain faster, more accurate insights.
- **Action drivers:** Traditional chatbots have frequently frustrated both patients and providers with incomplete, confusing or irrelevant information that fails to convert insights into action. Generative AI can augment providers by offering more intuitive, human-like conversations with patients and caregivers that can help inform care decisions and spur action.

2

How and why generative AI will revolutionize the way patients engage with their health

Generative AI is poised to upend how patients access health information, receive care and manage their conditions.

In the June 2023 report, *Scaling Smart Solutions with AI in Health: Unlocking Impact on High-Potential Use Cases*, more than 50 leaders from across healthcare, government, tech, academia and non-governmental organizations (NGOs) were interviewed. Many of them identified patient engagement as an area ripe for transformation using AI and other frontier technologies; however, they also pointed to the challenges, given the disconnected healthcare ecosystem, liability concerns inherent to engaging patients without a "human in the loop" and the technical limitations of traditional chatbots.

With the advent of generative AI, the question is: Given what is known about the connection between better patient engagement and improved health outcomes,[5] can this technology be used to safely plug the dangerous shortage of healthcare workers by communicating directly with patients across a range of healthcare interactions?

Traditionally, healthcare systems engage with individual patients mainly at a few discrete points in their health journey: when they see a doctor, when they start on therapy and when they discontinue treatment. What healthcare systems need are ways to engage with, influence and gather insights from patients across their entire health journey. Generative AI can assist with all three.

FIGURE 1 **Patient engagement and generative AI: addressing patient care gaps for better outcomes**

Health education assistance

Lifestyle support, early disease education and even prenatal care in LMICs*, with a focus on consumers and caregivers

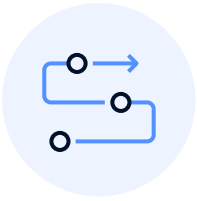

Co-pilots for patient triage

Patient care navigation in the first mile to help understand where to direct patients for diagnosis or treatment

Disease management interventions

Assist in monitoring, side effect management, adherence and keeping patients on treatment

*Low- and middle-income countries

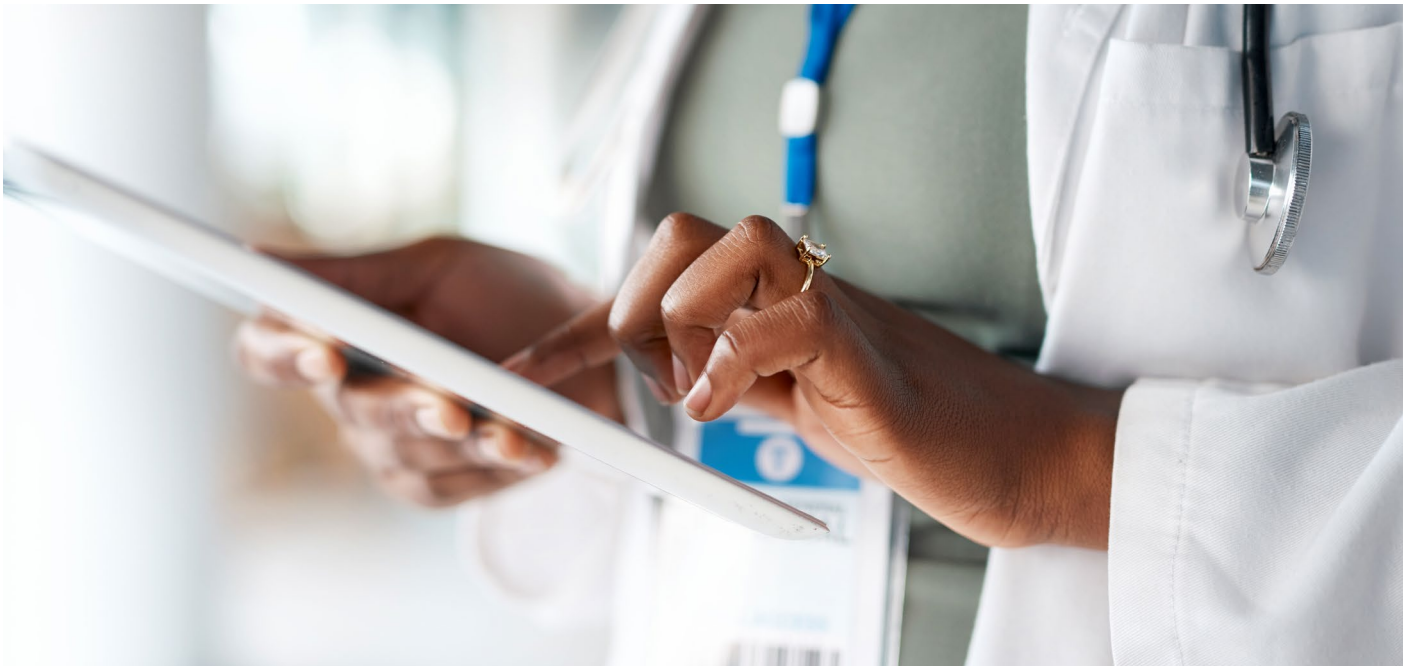

> LLMs can be trained to be polylingual, delivering answers that are not only correct – the bare minimum – but also sufficiently understandable to make them actionable.

Patient engagement use case 1: health education assistance

According to a survey of US adults conducted before ChatGPT became widely available, nearly all US healthcare consumers (94%) say they use the internet to educate themselves about their own health, self-diagnose, understand treatment protocols, manage side effects and address their mental health.[6] Respected entities such as Mayo Clinic, a US-based health system and research institution, already offer expansive, consumer-facing repositories of health information accessible through traditional search engines like Google and Bing. In the vast majority of cases, these resources should not be used in lieu of seeing a doctor. If used responsibly, however, platforms like Mayo's could significantly improve incidence of chronic illness like diabetes and cardiovascular disease through upstream lifestyle improvements and better decision-making about preventive care. However, studies have consistently found that very few health websites meet established thresholds for quality[7] and that healthcare consumers in the US,[8] Europe[9] and LMICs[10] show limited health literacy.

Generative AI may be able to provide healthcare consumers and their caregivers – more than 50 million of whom provide unpaid care for those 65 and older in the US[11] – reliable, accessible health information. (Globally, people – mostly women – log 16.4 billion hours doing unpaid care work for people of all ages.[12])

However, health education is not just about making information available, it is also about making it easy for healthcare consumers to understand (and therefore actionable). A recent survey estimated that gaps in health literacy cost the US economy as much as $238 billion per year, particularly in underserved communities and for the 8% of the population with limited English proficiency.[13] In the EU, that figure is between €120 billion and €320 billion.[14] Health literacy is even more limited in LMICs, where education levels are lower, health systems are dramatically under-resourced and access to health education is limited.[15]

Large-language models (LLMs) trained exclusively on health data that exceeds specified quality thresholds would eliminate the shortcomings that plague traditional search. Moreover, they can be trained to be polylingual, delivering answers that are not only correct – the bare minimum – but also sufficiently understandable to make them actionable no matter the user's language fluency, education level or cultural context.

BOX 1

Real-world vignette – Ada Health

Ada Health, a Germany-based medical AI company, has built a digital symptom checker based on natural language processing (NLP) enhanced by generative AI. Unlike when checking systems through a LLM like ChatGPT or Google's Bard, which are trained on the entire corpus of the World Wide Web, Ada's knowledge base is curated, clinically vetted and continuously updated. Its health assessments are deterministic, meaning identical inputs will yield identical outputs. In tests, Ada's app has proven to be safe and accurate, correctly recognizing inflammatory rheumatic diseases more often than physicians and besting competing apps in diagnosing a range of conditions.

Patient engagement use case 2: co-pilots for patient triage

The COVID-19 pandemic exacerbated an existing global healthcare provider shortage, with record levels of US doctors experiencing burnout,[18] nurses leaving the field in droves[19] and patient experiences plummeting.[20] The public is suffering the inevitable results in the form of increasing health disparities and rising healthcare costs, decreasing life expectancy in many parts of the world and increased burden of care from ageing and chronic disease.[21]

Early patient-facing triage solutions, such as MyndYou and others, have tried to solve care navigation gaps using chatbots as "virtual care assistants". These solutions offer a way for strapped healthcare providers to shift resources away from non-urgent patients to those who require timely, high-touch care. Platforms like these show strong promise but have so far failed to scale due to a lack of open-source models for the 7,000-plus languages spoken around the world and limits to predictive AI's ability to understand and reflect nuances across varied cultural contexts and healthcare ecosystems. With domain-specific and culturally diverse training, generative AI can begin addressing many of these gaps, offering a way for healthcare providers to safely stratify patients entering the healthcare system and to route their care more effectively.

In addition to stratifying patients and routing them to the appropriate level of care, generative AI can help with the more prosaic – but just as critical – task of getting patients into the health system in the first place. By automating many aspects of patient intake, AI can alleviate administrative burden and help patients navigate the "first mile" of their care journeys.

BOX 2

Real-world vignette – K Health and Cedars-Sinai

US-based healthcare innovator K Health partnered with Cedars-Sinai, a leading academic health system in the US, to offer patients 24/7 virtual access to Cedars-Sinai doctors through a new mobile app called Cedars-Sinai Connect. Patients can access Cedars-Sinai healthcare professionals 24/7 for personalized acute, chronic and preventive care, including same-day visits and evaluation for referrals to specialists within the system. The virtual care platform uses AI technology developed by K Health, which automates most aspects of patient intake and data entry, reducing administrative burden and allowing providers to focus more on patient care. Patients can share symptoms, health details and concerns through the app, which summarizes and presents them directly to clinicians, who then use their clinical expertise to treat each patient.

BOX 3

Real-world vignette – University of Rochester Medical Center

Doctors report spending as long as two hours per day reading and responding to messages sent through patient portals, a number that continues to climb as patients increasingly use these platforms. Leaders at the University of Rochester Medical Center (URMC) in Rochester, New York, have been building and testing models for years to support clinical operations – with limited success, thanks to the messiness of available data and the limitations of traditional natural language processing (NLP) techniques. Powered by ChatGPT-4 and working in a secure Microsoft Azure environment, URMC has built a model to appropriately triage messages to either a doctor, nurse or staff member. The tool was still being tested at publication time, but testing showed reliability (repeating the same action when presented with an identical input) in the high-90s and accuracy (selecting the most appropriate recipient) in the mid-80s.[22]

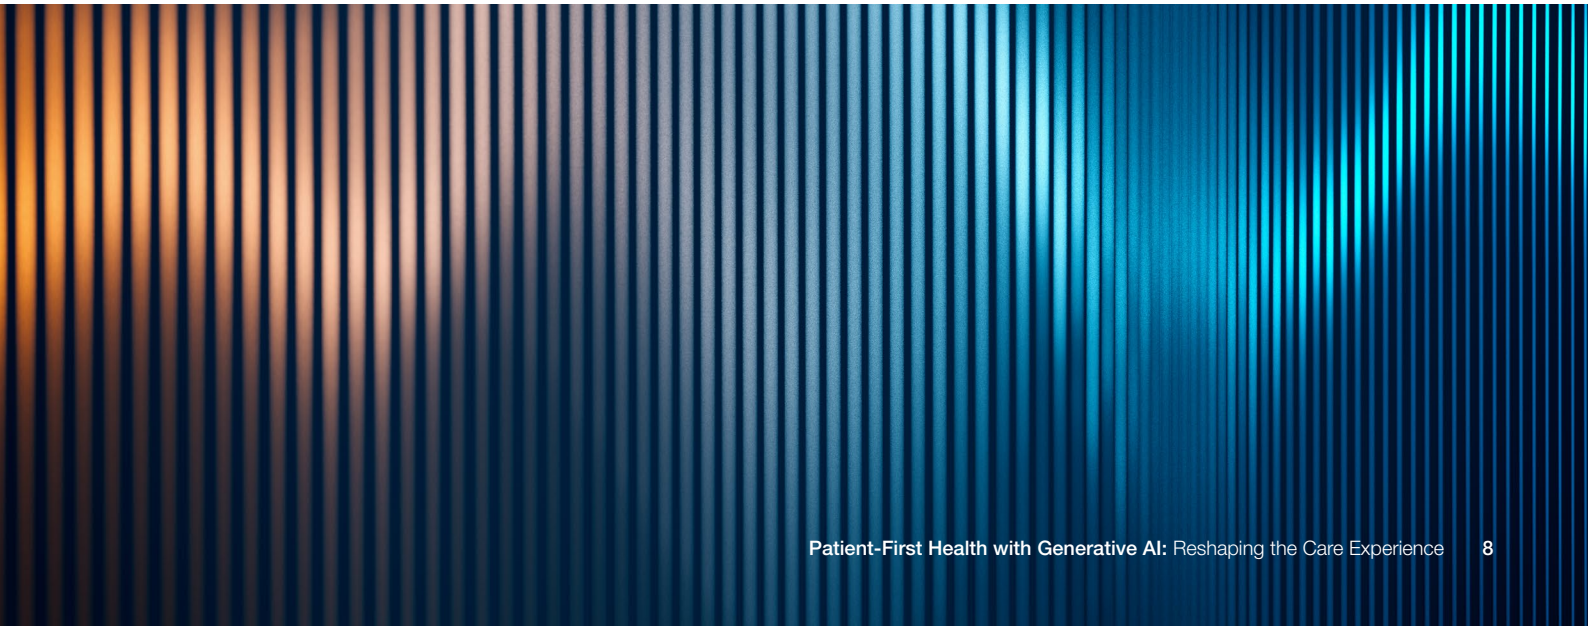

Patient engagement use case 3: disease management interventions

Once a patient has been diagnosed with a condition and has begun treatment, his or her health journey is only just beginning. Treatment adherence is a key facet of managing illness, especially chronic illness, and yet studies have found that, in the first year of treatment, as many as 60% of patients suffering from chronic illness either miss doses, take the wrong dosage or abandon treatment altogether.[23] Inadequate medication adherence (defined as taking less than 80% of the prescribed dose) is estimated to cost the US healthcare system between $100 billion and $290 billion and health systems in Europe €125 billion per year.[24]

Traditional algorithms can help predict when a patient is likely to miss a dose or drop treatment, but they struggle to suggest (let alone execute) effective interventions to keep at-risk patients on course. Generative AI can make predictive algorithms significantly more powerful by taking the output from a classical algorithm and using it to create an intervention tailored to the individual patient. This use case is relevant not only to provider organizations but also to biopharma, which spends hundreds of millions of dollars on patient-support call centres, apps and digital "nudges" to enable better disease management.

Ongoing management of disease isn't only about keeping patients on their prescribed course of treatment. Generative AI also can shoulder some of the administrative burden inherent to treating patients living with one or more chronic conditions. It is also showing promise bringing together multiple data streams, such as genetics and biometrics, to tailor treatment plans and chart personalized care paths for individual patients.

BOX 4

Real-world vignette – Mayo Clinic

Rheumatoid arthritis (RA) is a difficult disease to manage, with only about 15% of patients experiencing quick relief from its painful symptoms. RA's complexity lies in a patient's genetics, which play a big role in both the disease's progression and how well treatments work. Mayo Clinic is pioneering the use of a patient's genetic profile and developing multimodal generative AI models to quickly and accurately predict treatment effectiveness. This innovative approach is akin to translating a patient's unique genetic code into a language these models can understand. The result is a more tailored and efficient way of finding the right treatment, potentially speeding up the process of symptom relief for RA patients.

BOX 5

Real-world vignette – Amazon and Hurone AI

High-income countries earn the lion's share of investment in cancer prevention, screening and diagnosis, yet cancer remains a major health problem in LMICs. Take Nigeria, where four out of five cancer diagnoses is fatal.[25] The problem is one of resources: Rwanda, for instance, has a population of 14 million people but only 15 oncologists.[26] Amazon and Hurone AI, a US-based healthtech start-up, have partnered with the governments of Rwanda, Nigeria and Kenya to develop a generative AI-powered patient-facing application to communicate with local oncologists. Patients can track symptoms and side effects throughout treatment and remission, with generative AI recommending personalized responses and interventions for doctors to send to patients. The generative AI tools also provide doctors with personalized recommendations based on side effects, tolerability and efficacy. A beta test saved doctors 75% of time spent on specific care tasks – a potentially lifesaving efficiency boost in regions where resources are so limited.

BOX 6

Real-world vignette – Microsoft and Epic Systems

The annual volume of messages flowing through electronic patient portals doubled between 2019 and 2022 to more than 1 billion. On average, physicians spend more than 15 hours per week on administrative tasks, with almost two-thirds of that time taken up by electronic health record documentation – work that can be largely done by generative AI.

Microsoft and Epic Systems, the US's largest provider of electronic health records, are partnering to pilot the use of Azure OpenAI to respond to patient inquiries. Generative AI, which at least one study has found to provide answers displaying higher levels of empathy than those drafted by human clinicians, drafts responses based on the patient's inquiry and medical history. A doctor then has the choice to accept the responses unedited, to revise as they see fit or to draft their own response from scratch. This use case holds significant potential both to reduce provider burden and improve patient experience, which studies have consistently linked with improved patient outcomes.

3

Three barriers to safe, effective patient-facing generative AI

Adoption barriers include data issues, mistrust of outputs and barriers to scale outside wealthy countries.

In the previous joint briefing paper on AI in healthcare, technology and healthcare leaders identified numerous barriers slowing the widespread adoption of AI in healthcare, including holes in the data foundation, insufficient trust in algorithms, inadequate technological infrastructure and scalability issues.

Using generative AI to engage patients faces many of the same challenges, plus some new ones. These have been ranked from highest to lowest priority.

> **Data and trust remain the two foundational barriers for adopting generative AI in healthcare.**
>
> Mustaqhusain Kazi, Global Head, Informatics Strategy and Digital Innovation, Roche

Mistrust

Interviewees singled out trust as the most formidable barrier to widespread adoption of generative AI for patients, with one blanket issue rising above all others: inaccuracy.

Inaccuracy: For reasons that are often unclear to users, LLMs occasionally generate incorrect or misleading answers or responses in the absence of sufficient data. These hallucinations can be hard to discern, leading to confusion and misinformation. Since the risk of faulty health information is so high, it is paramount that providers integrate generative AI into workflows such that AI augments, rather than automates, clinical decision-making. In addition, careful prompt engineering by developers in initial releases, and continued refinement on new and domain-specific datasets for future releases, is needed to ensure LLMs avoid misinformation and can explain why a model responded in a specific way.

Companies like Perplexity AI are seeking to increase trust in outputs through greater transparency by building in functionality that explains how a model arrived at a given output. This can help users evaluate the degree of trust they should place in it.

Holes in the data foundation

While many organizations have made progress in building their own data management infrastructure and augmenting it with data consortia and even federated learning, additional work is needed to ensure models are trained on sufficiently robust data to produce outputs that are fit enough for safe and reliable patient use. Two data issues predominate:

Biased outputs: As with predictive AI techniques, generative AI outputs reflect underlying bias in the data on which they've been trained. This creates intense urgency to implement governance that enhances transparency across these models' entire life cycle – especially for those models, such as those from OpenAI, Google and others, that have been trained on the broad corpus of the internet.

Patient privacy: Patient health data is incredibly sensitive, so organizations must build security, privacy and ongoing data management protection into governance models. Organizations like Johnson & Johnson and Syntegra use generative AI to create "synthetic data", allowing them to use the technology while bypassing security and privacy issues altogether.

Scalability across regions and contexts

As with many novel technologies, there is a risk that generative AI exacerbates existing disparities. Multilateral, cross-border partnerships will be key to accelerating generative AI's impact outside high-income countries. Two barriers stand out:

Language: A dearth of high-quality training data in less common languages has limited the impact of text-based predictive AI techniques based on NLP. While generative AI may help address these limitations to a degree, barriers remain – most prominently, the need for context-specific prompt engineering to ensure outputs are safe and reliable across the spectrum of global populations.

Resources: Generative AI models must be trained on large volumes of data, which can be expensive to collect and store, and they require heavy computational resources to run. In addition, cloud computing services are needed to deploy and especially scale generative AI applications. All of these will prove onerous for resource-strapped LMICs, even if the technology's deployment will be a long-term resource (and life) saver.

> **Data without action is useless, and generative AI has vast potential to empower decisions, especially in middle- and low-income countries, where healthcare resources are extremely scarce.**
>
> John Sargent, Founding Partner, BroadReach Group

4

To make generative AI for patients a reality, healthcare stakeholders must take these four steps

Overcoming adoption barriers for healthcare generative AI demands action and collaboration from players across the healthcare ecosystem.

The release of ChatGPT set off a frenzy of public interest, private investment and experimentation. Moving past the razzle-dazzle phase will require sustained public-private engagement to create urgency, enable change and protect the common interest. The barriers slowing the widespread adoption of generative AI and the corresponding enablers for overcoming them mirror those for predictive AI. These are covered in depth in the June 2023 Forum-ZS insight report.

FIGURE 2 **Ecosystem and principles for generative AI acceleration**

Usable and representative data	Trust in AI and stakeholder education	Partnerships and scaling mindset	Actions by individual organizations

Note: 1. NGOs, advocacy and patient groups, distributors, group purchasing organizations, investors and other organizations in healthcare.

> All stakeholders must help educate patients and providers on generative AI's limitations, especially its potential for hallucination.

The recommendations below are intended as a generative AI-specific action plan for health and healthcare players to begin addressing barriers to the widespread adoption of patient-facing generative AI.

1. **Build trust through empathy and interactivity.** To instil trust in generative AI models, providers, payers, biopharma and tech companies must prioritize fine-tuning models to make them more empathetic – including by having doctors "test" responses to improve outputs. When these models have been trained on healthcare domain-specific data, and users can engage them in a dialogue to evaluate their outputs, then patient trust and engagement should follow.

2. **Mitigate against bias.** Providers, payers, government and biopharma all possess rich consumer and patient datasets, but they sit in disconnected data layers. Given that it's impossible to fully correct for bias inherent in models' training data, mitigating output bias requires creating a richer data layer by connecting the data ecosystem and then measuring response bias on an ongoing basis – especially in underserved populations.

3. **Keep humans in the loop.** Generative AI's uncanny ability to mimic human-like text creates a false sense of security. Yet the technology is fallible. All stakeholders must help educate patients and providers on generative AI's limitations, especially its potential for hallucination. At the same time, providers must work to integrate patient-facing tools into existing clinical workflows and implement processes to review recommendations and intervene accordingly, especially in high-risk discussions and with patients with severe disease.

4. **Plan to scale across contexts.** Running large models is extremely energy-consumptive, so technology companies must work to develop and deploy more cost-efficient means to run them. At the same time, governments, NGOs and connectors (defined here as advocacy and patient groups, distributors, group purchasing organizations, investors and other organizations in healthcare) must also develop flexible deployment models that recognize the varying needs and exigencies across regions, cultures and contexts. With generative AI and healthcare, one size does not (and cannot) fit all.

Conclusion

Patient-facing generative AI poses risks – but the risk of maintaining the status quo is graver.

A person's health is inherently precious. Any stakeholder developing, deploying or vouching for a generative AI-powered tool that could put even one patient's health at risk through faulty, confusing or ill-timed information should be cautious. However, the risk of medical or ethical malpractice stemming from the responsible use of these tools pales compared with the more general malpractice of choosing to ignore their boundless capacity for patient impact.

This is true at the patient level and even more so at the population level. Global patient demand for medical treatment outstrips supply by many orders of magnitude, leading to great loss of life, reduced lifespan for millions and growing fractures to global health and healthcare systems. Generative AI, especially when paired with predictive AI techniques, offers the greatest current hope of translating today's golden age of scientific innovation into equitable improvements to global health and healthcare.

Contributors

Lead authors

Dan Reiss
Principal, ZS

Antonio Spina
Lead, Digital and AI Health,
World Economic Forum

Acknowledgements

Thanks are extended to the following individuals for their invaluable contributions to this paper.

ZS

Alex Anokhin
Lead, Health and Healthcare Innovations

Megan Wolfe
Senior Director

Nicolas Zimmerman
Content Strategy Lead

World Economic Forum

Jitka Kolarova
Lead, Health and Healthcare Innovations

Andy Moose
Head, Health and Wellness

The authors would like to thank the following interviewees for their time and input, which were crucial to the production of this paper.

Allon Bloch
Co-Founder and Chief Executive Officer, K Health

Bernard Katz
Director, Product Management, BroadReach Group

Erica Chain
Digital Innovation Lead (Architecture, Strategy & Innovation Ecosystem Catalyst), Roche

Mustaqhusain Kazi
Global Head, Informatics Strategy and Digital Innovation, Roche

Pankaj Paul
Adviser to the Chairman and Founder, Burjeel Holdings

John Sargent
Managing Director, Founding Partner, BroadReach Group

Vivian Tan
Vice-President, Strategic Information Management & Global Relationships, Kaiser Permanente

Ruan Viljoen
Chief Technology Officer, BroadReach Group

Production

Laurence Denmark
Creative Director, Studio Miko

Martha Howlett
Editor, Studio Miko

Oliver Turner
Designer, Studio Miko

Endnotes

1. "Measuring and addressing physician burnout", *American Medical Association (AMA)*, 3 May 2023, https://www.ama-assn.org/practice-management/physician-health/measuring-and-addressing-physician-burnout.
2. Boniol, Mathieu, et al., "The global health workforce stock and distribution in 2020 and 2030: a threat to equity and 'universal' health coverage?", *BMJ Global Health*, vol. 7, no. 6, 27 June 2022, https://www.ncbi.nlm.nih.gov/pmc/articles/PMC9237893/.
3. World Health Organization (WHO), *Global strategy on human resources for health: Workforce 2030*, 2020, https://iris.who.int/bitstream/handle/10665/250368/9789241511131-eng.pdf?sequence=1.
4. World Economic Forum, *Global Health and Healthcare Strategic Outlook: Shaping the Future of Health and Healthcare*, 2023, https://www3.weforum.org/docs/WEF_Global_Health_and_Healthcare_Strategic_Outlook_2023.pdf.
5. "Patient Engagement", *Centers for Disease Control and Prevention (CDC)*, n.d., https://www.cdc.gov/healthliteracy/researchevaluate/patient-engage.html.
6. AHIMA Foundation, *Understanding Access and Use of Health Information in America*, 2021, https://ahimafoundation.org/media/ngfbggsk/oct2021_understanding_access_use_health_information_america_ahima_foundation.pdf.
7. Devine, Theresa, et al., "Making Quality Health Websites a National Public Health Priority: Toward Quality Standards", *Journal of Medical Internet Research*, vol. 18, no. 8, 2 August 2016, https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4987491/.
8. "Talking Points About Health Literacy", *CDC*, n.d., https://www.cdc.gov/healthliteracy/shareinteract/TellOthers.html.
9. "Why improving health literacy is important", *MSD*, 3 March 2022, https://www.msd.com/stories/why-improving-health-literacy-is-important/.
10. Meherali, Salima, et al., "Health Literacy Interventions to Improve Health Outcomes in Low- and Middle-Income Countries", *Health Literacy Research and Practice*, vol. 4, no. 4, 2020, pp. 251-266, https://journals.healio.com/doi/10.3928/24748307-20201118-01.
11. National Alliance for Caregiving (NAC) and AARP, *Caregiving in the U.S. 2020*, 2020, https://www.caregiving.org/research/caregiving-in-the-us/caregiving-in-the-us-2020/.
12. International Labour Organization, *Care Work and Car Jobs for the Future of Decent Work*, 2018, https://www.ilo.org/wcmsp5/groups/public/---dgreports/---dcomm/---publ/documents/publication/wcms_633135.pdf.
13. Lambert, Veronica and Deborah Keogh, "Health literacy and its importance for effective communication. Part 1", *Nursing Children and Young People*, vol. 26, no. 3, 2014, pp. 31-37, https://pubmed.ncbi.nlm.nih.gov/24708337/.
14. Baccolini, V., et al., "What is the Prevalence of Low Health Literacy in European Union Member States? A Systematic Review and Meta-analysis", *Journal of General Internal Medicine*, vol. 36, 2021, pp. 753-761, https://link.springer.com/article/10.1007/s11606-020-06407-8.
15. Meherali, Salima, et al., "Health Literacy Interventions to Improve Health Outcomes in Low- and Middle-Income Countries", *Health Literacy Research and Practice*, vol. 4, no. 4, 2020, pp. 251-266, https://journals.healio.com/doi/10.3928/24748307-20201118-01.
16. Balasubramanian, Sai, "Google Cloud and Mayo Clinic Partner To Improve Heathcare With Generative AI", *Forbes*, 7 June 2023, https://www.forbes.com/sites/saibala/2023/06/07/google-cloud-and-mayo-clinic-partner-to-improve-healthcare-with-generative-ai/.
17. Gupta, Aashima and Amy Waldron, "A responsible path to generative AI in healthcare", *Google Cloud*, 13 April 2023, https://cloud.google.com/blog/topics/healthcare-life-sciences/sharing-google-med-palm-2-medical-large-language-model.
18. "Measuring and addressing physician burnout", *AMA*, 3 May 2023, https://www.ama-assn.org/practice-management/physician-health/measuring-and-addressing-physician-burnout.
19. Mandowara, Khushi and Leroy Leo, "One-third of US nurses plan to quit profession, survey shows", *Reuters*, 1 May 2023, https://www.reuters.com/world/us/one-third-us-nurses-plan-quit-profession-report-2023-05-01/.
20. Muoio, Dave, "Pandemic ushered 'unprecedented' declines in hospital patient experience scores, study finds", *Fierce Healthcare*, 25 August 2023, https://www.fiercehealthcare.com/providers/pandemic-ushered-unprecedented-declines-hospital-patient-experience-scores-study-finds#:~:text=Compared%20to%20expected%20trends%20based,of%202021%2C%20per%20the%20study.
21. Heuveline, Patrick, "Global and National Declines in Life Expectancy: An End-of-2021 Assessment", *Population and Development Review*, vol. 48, issue 1, 2022, pp. 31-50, https://www.medrxiv.org/content/10.1101/2022.01.14.22269109v1.
22. Hagland, Mark, "At the University of Rochester Medical Center, a Breakthrough With Generative AI", *Healthcare Innovation*, 14 November 2023. https://www.hcinnovationgroup.com/analytics-ai/article/53078074/at-the-university-of-rochester-medical-center-a-breakthrough-with-generative-ai.
23. Wood, Brenda, "Medication Adherence: The Real Problem When Treating Chronic Conditions", *U.S. Pharmacist*, vol. 37, issue 4, 2021, https://www.uspharmacist.com/article/medication-adherence-the-real-problem-when-treating-chronic-conditions.
24. Kripalani, Christopher P. et al., "Medication Non-Adherence - A $290 Billion Unnecessary Expenditure", *Annals of Internal Medicine*, 2015; "€125 billion lost each year across Europe due to non-adherence to medication", *International Longevity Centre UK*, 7 April 2022, https://ilcuk.org.uk/125-billion-lost-each-year-due-to-non-adherence/.
25. WHO, *Nigeria Fact Sheet*, 2020, https://gco.iarc.fr/today/data/factsheets/populations/566-nigeria-fact-sheets.pdf.
26. "This AWS-powered startup gives Rwandan cancer patients better access to life-saving treatments", 31 January 2023, https://www.aboutamazon.com/news/aws/this-aws-powered-startup-gives-rwandan-cancer-patients-better-access-to-life-saving-treatments.

WORLD ECONOMIC FORUM

COMMITTED TO IMPROVING THE STATE OF THE WORLD

The World Economic Forum, committed to improving the state of the world, is the International Organization for Public-Private Cooperation.

The Forum engages the foremost political, business and other leaders of society to shape global, regional and industry agendas.

World Economic Forum
91–93 route de la Capite
CH-1223 Cologny/Geneva
Switzerland

Tel.: +41 (0) 22 869 1212
Fax: +41 (0) 22 786 2744
contact@weforum.org
www.weforum.org