CitiusTech

HEALTHCARE LIMITLESS

Building Quality and Trust Led GenAI Applications in Healthcare

Introduction

The interest in Generative AI (GenAI) applications in business has grown exponentially. Yet, the 2023 McKinsey Global Survey reveals that just 22% of respondents feel their organizations are prepared to manage GenAI.[1] And healthcare is no different.

Generative AI has demonstrated immense potential in healthcare, promising significant benefits such as driving productivity gains, enhancing patient and provider experiences, and ultimately improving clinical outcomes.[2] Healthcare organizations have seen the GenAI revolution open doors to new dimensions in patient interaction, clinical decision support, telehealth services, health education, and health advice. GenAI deployment is delivering operational insights, improving efficiency, predicting trends, and improving population health management and spread.[3]

Despite these benefits, as of 2023, only 25% of healthcare organizations have deployed GenAI.[4] One reason for taking it slow is the lack of confidence in Generative AI output for a critical industry like healthcare.

Why is it difficult to trust GenAI applications?

Generative AI is based on Large Language Models (LLMs) and their ability to summarize, analyze, and transform vast volumes of data into resourceful assets. However, LLM-based applications tend to "hallucinate,"[5] i.e., generate incorrect responses and do not have proof of correctness. Healthcare data serves as the training material for LLMs, and the accuracy of their output hinges upon the quality of the datasets employed in their training.

There have been instances where GenAI outcome has been biased, factually incorrect, or illegally scraped from a copyrighted source.[5] It also lacks transparency and explainability on how it arrived at a certain interpretation and is riddled with privacy concerns.

In healthcare, where an incorrect decision could mean life or death, these issues are big red flags. Unfortunately, despite these issues, the perceived ease of use of these models creates a false sense of feasibility.

Inaccuracy, cybersecurity, and intellectual property infringement are the most-cited risks of generative AI adoption.

Generative AI–related risks that organizations consider relevant and are working to mitigate, % of respondents[1]

Image credits: McKinsey and Company[6]

What happens if your GenAI applications lack quality and cannot be trusted or relied upon?

Using GenAI applications that are not trustworthy can lead to severe negative consequences for healthcare players. For instance:

- Unauthorized and even inadvertent breaches of sensitive patient data compromise patient confidentiality.
- Algorithmic and training data bias can affect outcomes, leading to unfair care decisions and unequal treatment.
- The biggest concern around Generative AI is its ability to create realistic and convincing content that may not be accurate. Realistic medical images or reports that are not authentic and reliable could lead to severe consequences, and poor care decisions, endangering the patient.

These issues are a significant concern when it comes to patient safety, regulatory compliance, and reputational and potential financial damage for a healthcare organization.

Concerns regarding risk, bias, reliability, security, compliance, and privacy create friction and a lack of support or investments for building and deploying enterprise-grade GenAI solutions. As a result, the life cycle of most of the GenAI solutions ends at the pilot or proof-of-concept stage, and they are not further developed or operationalized to solve actual business problems.

Accelerating GenAI adoption in healthcare with proven Quality and Trust Solution

Leading healthcare players are now voicing their concerns, and conversations about the inaccuracy and the possible harm of unverified GenAI applications have started to emerge. Opinion leaders like Gartner also emphasize the need to create trust in GenAI applications by building solutions that incorporate safety principles and preempt regulatory and compliance issues.[7]

There have been ongoing conversations on the need to regulate GenAI due to its potential impact on the overall healthcare ecosystem.[8] Healthcare organizations must adhere to the Federal guidelines and other industry-recognized frameworks for responsible AI (RAI).

Trust, quality, reliability, predictability, and validation are essential to the success of generative AI. Without applications of generative AI being evaluated on these factors, the model's output will always be suspect and such models will lack the credibility for productionization and scaling.

However, healthcare organizations are struggling to establish governance and assurance mechanisms for the various GenAI projects and prototypes being launched across the enterprise. Presently, there are no established technology/platform agnostic frameworks to measure the quality and trust of GenAI solutions in healthcare. While there are a few RAI frameworks, they are suited for more traditional machine learning applications, or are siloed and are applicable only within a narrow context, making them inadequate for the vast expanse of healthcare applications.

There is also a lack of understanding on how to best validate GenAI, automate the validation process, and which GenAI technology stack to choose (e.g., which LLM, vector database, or search/retrieval strategy) for healthcare use cases. There is confusion about which metrics to use to evaluate these new tools from a quality and trust point of view and their relevance to the business problem at hand.

CitiusTech's GenAI Quality & Trust Solution is a first-of-its-kind that solves these challenges. It provides a decision-making framework accompanied by software libraries to choose the right quality and trust (Q&T) metrics for a GenAI solution that solves a healthcare problem. The solution builds confidence in the success of the proposed GenAI solution by anchoring quality and trust on key business outcomes. It takes a layered approach and starts with the business outcome and quality metrics associated with that. Then, it provides the right quality and trust metrics for the underlying layers, such as the LLM response interaction quality, the search and retrieval quality, and the data quality. Q&T metrics are computed for each of these layers and produce a composite picture of the overall quality and trustworthiness of the GenAI application.

The GenAI Quality & Trust Solution also synthesizes and actionizes recommendations from existing guidelines proposed by corporate and government regulatory bodies such as NIST, CHAI (Coalition for Health AI), EU GDPR, US Bill of AI Rights, and the Presidential Executive Order 14110. This solution can be leveraged for building and objectively evaluating enterprise readiness of Generative AI solutions for healthcare from a quality and trust point of view. In addition, the solution is agnostic to the technology stacks and cloud, and it can be seamlessly deployed into customer product development workflows.

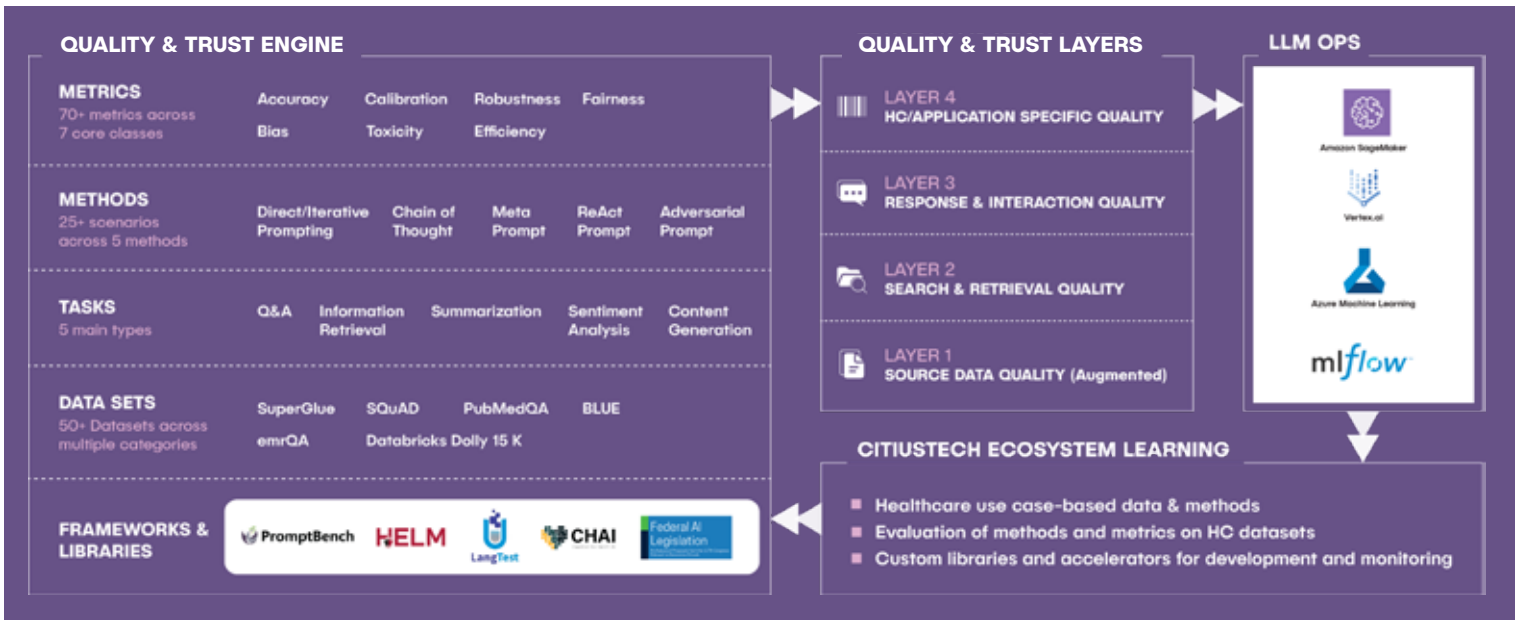

Image: Quality & Trust Solution

The CitiusTech GenAI Quality & Trust Solution will help organizations design, develop, integrate, and monitor quality and facilitate trust in Generative AI applications, providing the confidence needed to adopt and scale GenAI applications enterprise-wide. This solution can help deploy GenAI safely in almost all healthcare use cases. Some of these include:

Prior Authorization:

A GenAI Quality and Trust-powered solution helps evaluate the extraction and summarization of clinical information from the documents attached to a case for prior authorization. These might include patient history, medications, clinical guidelines, and other relevant documents. This GenAI application facilitates quick access to information like patient's symptoms, surgical history, prior treatments, allergies, and medications. It helps align with clinical guidelines, and health insurance coverage, leading to faster decisions on authorization for treatment.

Benefit Verification:

This GenAI solution efficiently summarizes patient data and enables agents with real-time access to comprehensive clinical information during calls for member benefits queries. The GenAI application powered by Quality and Trust Solution helps build confidence in the benefits verification process, which involves checking a patient's health insurance coverage before providing medical services. This helps mitigate failures in conducting benefits verification that may result in claim rejection or delayed payments by providing factual information about member benefits.

Patient Safety Narrative:

The Patient Safety Narrative is a regulatory document capturing summary of drug safety-related events during clinical trials conducted for an investigational drug. This GenAI solution helps create patient safety narratives for the risk of adverse events (AEs) and serious adverse events (SAEs) and establish a causal relationship to the investigational drug. It allows medical teams to save hours of manual review of documents, thereby reducing lengthy review cycles and variability in the narratives.

Setting the foundations of a responsible AI-led future in healthcare

With 94% of the healthcare industry yet in the nascent stage of establishing a concrete generative AI strategy, the time to incentivize quality and trust-assured GenAI initiatives is now.[9] Organizations can only realize the full potential of GenAI when they trust their deployed Generative AI solutions to deliver as expected. Objectively evaluating enterprise readiness of Generative AI for healthcare from a quality and trust point of view is the key.

A robust GenAI strategy that integrates quality and trust solutions, which synthesizes and provides actionable recommendations while aligning with existing guidelines proposed by corporate and government regulatory bodies, can change the way healthcare organizations work with GenAI. It will ensure that GenAI solutions don't end at the pilot or proof-of-concept stage and can be further developed and scaled to solve complex healthcare problems and improve patient outcomes.

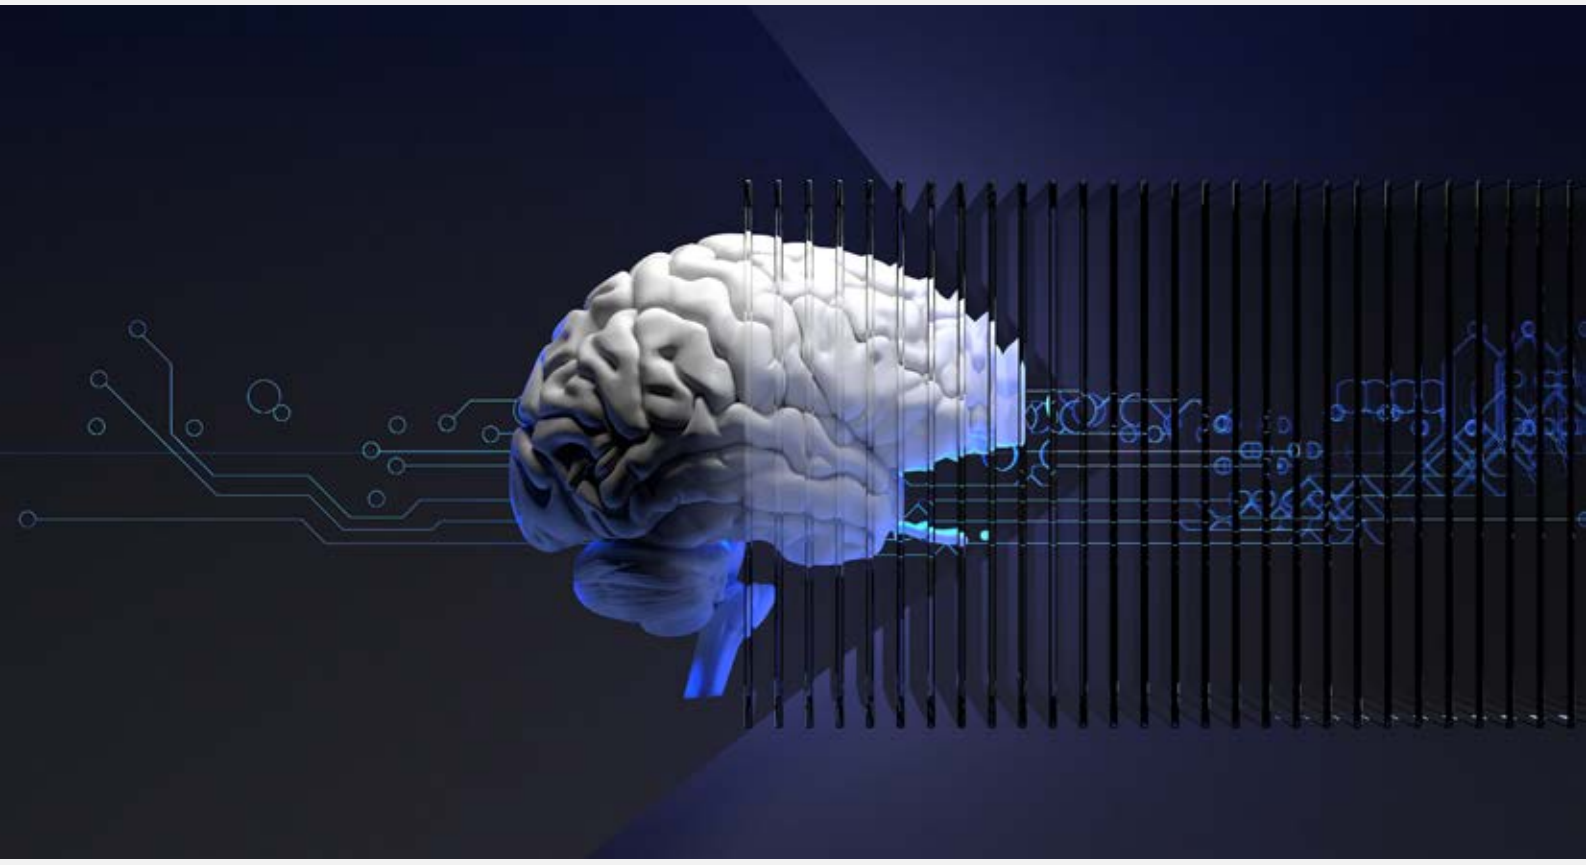

References:

1. Economic conditions outlook during turbulent times, December 2023 (mckinsey.com)
2. Global Healthcare Private Equity Report – Generative AI Will Transform Healthcare (bain.com)
3. How generative AI and large language models can close the gap between data and outcomes in healthcare (weforum.org)
4. Generative AI expected to take off in 2024, but concerns linger over cost, reliability, and security (fiercehealthcare.com)
5. The economic potential of generative AI: The next productivity frontier (mckinsey.com)
6. What's the future of generative AI? An early view in 15 charts (mckinsey.com)
7. Gartner Identifies Top Trends Impacting Technology Providers in 2024 (gartner.com)
8. Why tech companies want the government to regulate AI (yahoo.com)
9. Beyond Hype: Getting the Most Out of Generative AI in Healthcare Today (bain.com)

Authors

Kaushik Raha, PhD

Vice President, Head of AI Solutions, CitiusTech

Yogesh Parte, PhD

Assistant Vice President, AI Practice, CitiusTech

About CitiusTech

Our vision is to inspire new possibilities for the health ecosystem with technology and human ingenuity. At CitiusTech, we constantly strive to solve the industry's greatest challenges with technology, creativity, and agility. Together with the world's leading Healthcare and Lifesciences organizations and our partners, we aim to accelerate the transition to a human-first, sustainable, and digital healthcare ecosystem.

To learn more about this contact us at

GenAI@citiustech.com

Visit our website

www.citiustech.com/re-imagen-ai

CitiusTech