



**Journal of Innovation**

March 2024 | 24<sup>th</sup> Edition



# Driving Healthcare Transformation Through Generative AI

**Authors:**

Shyam V. Nath  
Deloitte  
[ShyNath@Deloitte.com](mailto:ShyNath@Deloitte.com)

Sunil Dadlani  
Atlantic Health  
[Sunil.Dadlani@AtlanticHealth.org](mailto:Sunil.Dadlani@AtlanticHealth.org)

**CONTENTS**

---

- 1 Overview ..... 3**
  - 1.1 Introduction ..... 3**
  - 1.2 Purpose..... 6**
  - 1.3 Scope ..... 7**
- 2 Industry Landscape ..... 7**
  - 2.1 Literature Review ..... 7**
- 3 Generative AI for Patient Care ..... 8**
  - 3.1 Generated Patient Summary..... 10**
    - 3.1.1 Process Optimization ..... 10
    - 3.1.2 Benefits ..... 10
  - 3.2 Personalized Patient Instructions ..... 12**
  - 3.3 Referrals to Other Physicians ..... 14**
  - 3.4 GenAI-Assisted Insurance Claim Explanation ..... 15**
- 4 High Level Architecture ..... 16**
  - 4.1 Technical Architecture ..... 17**
- 5 Discussion ..... 18**
- Acknowledgements ..... 20**

**FIGURES**

---

- Figure 1-1: Hype cycle for artificial intelligence, 2023. .... 4
- Figure 1-2: What CEOs talked about in Q3 2023 (vs. Q2 2023)..... 5
- Figure 3-1: Comparing medical domain to medical LLMs. .... 9
- Figure 4-1: Flow of generative AI-assisted patient care. .... 17
- Figure 4-2: Technical architecture of a GenAI-based patient care system. .... 18

---

# 1 OVERVIEW

---

Generative AI has the potential to revolutionize patient care by improving efficiency, quality, and patient experience. This paper explores how generative AI can be used to enhance patient outcomes, quality of care and improve the experience of the healthcare professionals. We will look at scenarios such as generating patient summaries to reduce the preparation time it takes for the physician and healthcare professionals ahead of a patient visit or procedure.

Likewise, generating personalized patient instructions, based on the specific needs of the patients as well as follow-up activities after the visit or discharge from the hospital. We will look at driving clinical efficiencies by generating advanced notes based on the discrete data captured during the patient visits, such as vitals, form entries and dictation from the physician.

Such advanced notes can help drive referrals to other physicians as well as assist with the medical billing. Finally, from the patient perspective, in generating explanations of the insurance claims. This paper will then go into high level architecture and technology considerations for Generative AI solutions in the healthcare industry.

## 1.1 INTRODUCTION

The origin of Artificial Intelligence (AI) traces back to the 1950s. While Generative AI (referred as GenAI, henceforth in this paper) does not have a specific origin, Ian Goodfellow and his colleagues who wrote the paper in 2014 on GAN's titled "Generative Adversarial Nets" are recognized as the most influential force behind GenAI. GenAI is still evolving and there are many contributors for the same. The paper titled "Attention is All You Need" by Vaswani et al from 2017, is considered the seminal work for Transformer Learning. The Transformer Architecture, introduced in this paper, is the foundation for Large Language Model (LLM), including GPT-3, LaMDA and BERT<sup>1</sup>.

As we look at the current landscape of AI, as shown by Figure 1-1, GenAI is at the peak of the Hype Curve for AI in 2023, according to Gartner<sup>2</sup>. As a result, it is quite logical to explore the use of GenAI in the healthcare industry. While there is no single or formal definition of GenAI, it generally refers to the use of AI to produce text, images, insights, software code and similar content, based on training information. The GenAI models are often trained on large datasets of existing information to allow the models to learn about the patterns and relationships in the data. Let's look at LLMs. Once trained, the LLMs can be used for application such as:

- Summarization of textual information,
- Text or idea creation,

---

<sup>1</sup> <https://arxiv.org/abs/1706.03762>

<sup>2</sup> <https://www.gartner.com/en/articles/what-s-new-in-artificial-intelligence-from-the-2023-gartner-hype-cycle>

## Driving Healthcare Transformation Through Generative AI

---

- Question answering about the textual information, and
- Chatbots to carry out initial conversations with humans.

The two well-known types of LLMs are:

- Language Model for Dialogue Applications (LaMDA) by Google AI: LaMDA is well known for dialogue or conversational applications and the ability to generate poems, scripts, musical pieces, emails, letters and software code. Example Google Bard.
- Generative Pre-Trained Transformer (GPT) by OpenAI: GPT-3 and GPT-4 are designed for a wider range of tasks and can follow instructions, answer questions and generate creative textual content. Example ChatGPT.

LaMDA is trained using large datasets of text and code while GPT-3/GPT-4 is trained using textual datasets only. GPT-4V with voice and image capabilities was announced in September 2023. LaMDA has about 1.5 billion parameters while GPT-3 has 175 billion. The estimated parameters for GPT-4 are in the range of 200-300 billion.

A parameter in LLM is a variable that the model learns from its training data. Parameters can represent a variety of things, such as the weights between neurons in the model or the biases of the neurons. The parameters are weights to control the model's output for a given input. The parameters of LLM should not be confused with the hyper-parameters in traditional AI models. LLM can be considered as Foundation model for natural language generation and comprehension. The "foundation" here is natural language.




Figure 1-1: Hype cycle for artificial intelligence, 2023<sup>3</sup>.

---

<sup>3</sup> <https://www.gartner.com/en/articles/what-s-new-in-artificial-intelligence-from-the-2023-gartner-hype-cycle>

## Driving Healthcare Transformation Through Generative AI

GenAI being at the peak implies that it is a technology phenomenon that is generating a lot of excitement and interest, but there is a lot of uncertainty about its potential business outcomes and its long-term viability. GenAI is so hyped up today can it be used to create new products and services, automate tasks, and generate creative content. The opportunity of GenAI applies to healthcare<sup>4, 5, 6, 7</sup> as well.




Figure 1-2: What CEOs talked about in Q3 2023 (vs. Q2 2023)<sup>8</sup>.

The healthcare industry is heavily regulated in the United States. There are several reasons for this high level of regulation in the opinion of the authors, namely:

- Ensuring patient safety and effective care, given human lives are at risk, in the practice of medicine and surgery.
- Ensuring access to quality care beyond any kind of socio-economic factors like race, gender, affordability, religious or political beliefs.

<sup>4</sup><https://www.mckinsey.com/industries/healthcare/our-insights/tackling-healthcares-biggest-burdens-with-generative-ai>

<sup>5</sup> <https://blog.google/technology/health/cloud-next-generative-ai-health/>

<sup>6</sup><https://www.healthcareitnews.com/news/generative-ai-can-be-applied-nearly-every-healthcare-use-case-you-can-think>

<sup>7</sup><https://www.fiercehealthcare.com/ai-and-machine-learning/us-patients-believe-generative-ai-can-improve-access-affordability>

<sup>8</sup> <https://iot-analytics.com/what-ceos-talked-about-q3-2023/>

- Promoting fair competition in the commercial sector to encourage research and innovation while balancing it with the public sector stepping up to keep overall costs in control.

Figure 1-2 clearly shows that CEOs across the different industries were enamored by GenAI in early parts of 2023<sup>9, 10</sup>. The rapid growth of Artificial Intelligence including GenAI, presents an interesting situation for the healthcare industry which usually is conservative in adoption of emerging technologies, due to its regulated nature. Healthcare is evidence-based, which implies a process of making decisions about patient care that is based on the best available scientific evidence. It helps to ensure that patients receive the highest quality care possible and reduce the risk of patients being harmed by ineffective or harmful treatments.

The Food and Drug Administration (FDA) is responsible for regulating the AI-powered medical devices or similar solutions. Likewise, the Health Insurance Portability and Accountability Act (HIPAA) governs the use of the information to train AI models. In the absence of formal regulations guiding the use of GenAI in healthcare, the onus lies on the early adopters of this technology to ensure that Protected Health Information (PHI) is not at risk. Healthcare providers cannot share sensitive patient information with tools like ChatGPT. See this article for details<sup>11</sup>.

This paper will look at the application and use cases of Generative AI to the Healthcare industry from both medical professional and patient's perspective. The intended audience for this paper includes both clinicians and technologists, who have some level of techno-functional understanding of the healthcare industry in clinical settings.

### 1.2 PURPOSE

In recent years, the pandemic as well as the overall aging population in the US, has stressed the healthcare system. In addition, the supply chain crisis and The Big Resignation, has made matters worse<sup>12</sup>.

Automation of insurance authorization for imaging and radiology have been some of the successful use cases of AI, to boost productivity. Prior Authorization Intelligent Automation, which is an AI-based application, has seen encouraging outcomes at Atlantic Health Systems. They are able to submit authorizations in time 50% below the industry average<sup>13</sup>. Now 70% of all

---

<sup>9</sup> <https://iot-analytics.com/product/trend-report-what-ceos-talked-about-q2-2023/>

<sup>10</sup> <https://iot-analytics.com/wp/wp-content/uploads/2023/09/What-CEOs-talked-about-in-Q3-2023-vs-Q2-2023-vf.png>

<sup>11</sup> <https://www.bricker.com/insights-resources/publications/chatgpt-in-healthcare-navigating-the-hipaa-cups>

<sup>12</sup> <https://www.cio.com/article/482327/atlantic-health-streamlines-insurance-authorization-with-intelligent-automation.html>

<sup>13</sup> <https://www.physicianspractice.com/view/dealing-with-prior-authorization>

## Driving Healthcare Transformation Through Generative AI

---

authorizations are handled by intelligent automation and less than 1% of appointments are cancelled due to delayed or denied authorizations at Atlantic Health. Here we will look at further adoption of GenAI, in similar healthcare settings.

### 1.3 SCOPE

We will deep dive into four use cases of GenAI related to patient care:

- Generated Patient Summary
- Personalized Patient Instructions
- Referrals to Other Physicians
- GenAI-assisted Insurance Claim Explanation

These uses cases will be discussed in detail, in this paper.

## 2 INDUSTRY LANDSCAPE

---

We will look at some examples of the AI related attempts and initiatives, in the healthcare domain, in the last two decades. These will help to show how the adoption of GenAI can improve the efficiency in healthcare settings such as in hospital and physician group practices, with the welfare of patients at the core of it.

### 2.1 LITERATURE REVIEW

Let us do a limited literature review here to see how companies, researchers and practitioners have used or attempted to use AI - both discriminative AI and GenAI. We will look at IBM Watson Health that started several years ago, and more recent work such as the Nature publication on GenAI for healthcare, UCSF-BERT<sup>14</sup> and Mayo Clinic.

- IBM Watson Health: While IBM Watson did not directly use present day LLMs, it used similar techniques for natural language processing. It had limited success and eventually spun it off as Merative in January 2022. It includes Merative Vida<sup>15</sup>, Watson – Oncology, Genomics and Drug Discovery.
- Nature Magazine, June 2023 Paper titled “Health system-scale language models are all-purpose prediction engines”: The researchers state that LLMs could be used to read the notes written by physicians, to access a detailed description of a patient’s medical condition, thereby supporting the decision support at the point of care<sup>16</sup>.

---

<sup>14</sup> <https://browse.arxiv.org/pdf/2210.06566.pdf>

<sup>15</sup> <https://www.merative.com/content/dam/merative/documents/ebook/achieving-the-benefits-of-ai.pdf>

<sup>16</sup> <https://www.nature.com/articles/s41586-023-06160-y>

## Driving Healthcare Transformation Through Generative AI

---

- UCSF-BERT: The University of California at San Francisco (UCSF) has used the open-source Bidirectional Encoder Representations from Transformers (BERT) and trained it with the corpus of 75 million deidentified clinical notes, to create a medical LLM<sup>17</sup>. They also evaluated the UCSF-BERT model on two tasks that use the patient's discharge summaries from UCSF. The first task does the ICD-9 diagnostic code assignment, and the second task does the therapeutic class inference.
- Mayo Clinic: Mayo Clinic has done an early pilot for combining the power of LLMs with Microsoft 365 CoPilot, involving hundreds of doctors, clinical staff and healthcare workers<sup>18</sup>. GenAI will help Mayo Clinic physicians automate many form-filling tasks and allow them to focus on patient care. Dr Bradley Erickson is working on generating Brain MRI<sup>19</sup> at the Radiology Informatics Lab at Mayo Clinic.

These examples show that there is increased investment and experimentation with GenAI, in the healthcare industry. In addition, HCA Healthcare is working with Google on GenAI for hospitals<sup>20</sup> and UNC Health<sup>21</sup> is working with Epic and Microsoft on integration of LLMs with Epic software for EMR. As a result, many healthcare stakeholders are feeling the pressure to not be left behind.

### 3 GENERATIVE AI FOR PATIENT CARE

---

In this section we will go into the details of each of the four use cases of GenAI for medical patient care and creating efficiencies in the system with the patient at the center of it. Before we get into the details of these four cases, that are being piloted at Atlantic Healthcare Systems, we look into our point of view of the relationship between the medical domain and medical LLMs in Figure 3-1.

---

<sup>17</sup> <https://arxiv.org/ftp/arxiv/papers/2210/2210.06566.pdf>

<sup>18</sup> <https://newsnetwork.mayoclinic.org/discussion/mayo-clinic-to-deploy-and-test-microsoft-generative-ai-tools/>

<sup>19</sup> <https://www.mayo.edu/research/labs/radiology-informatics/research/projects>

<sup>20</sup> <https://investor.hcahealthcare.com/news/news-details/2023/HCA-Healthcare-Collaborates-With-Google-Cloud-to-Bring-Generative-AI-to-Hospitals/default.aspx>

<sup>21</sup> <https://news.unhealthcare.org/2023/05/unc-health-works-with-epic-on-integration-of-generative-artificial-intelligence-ai-tools/>



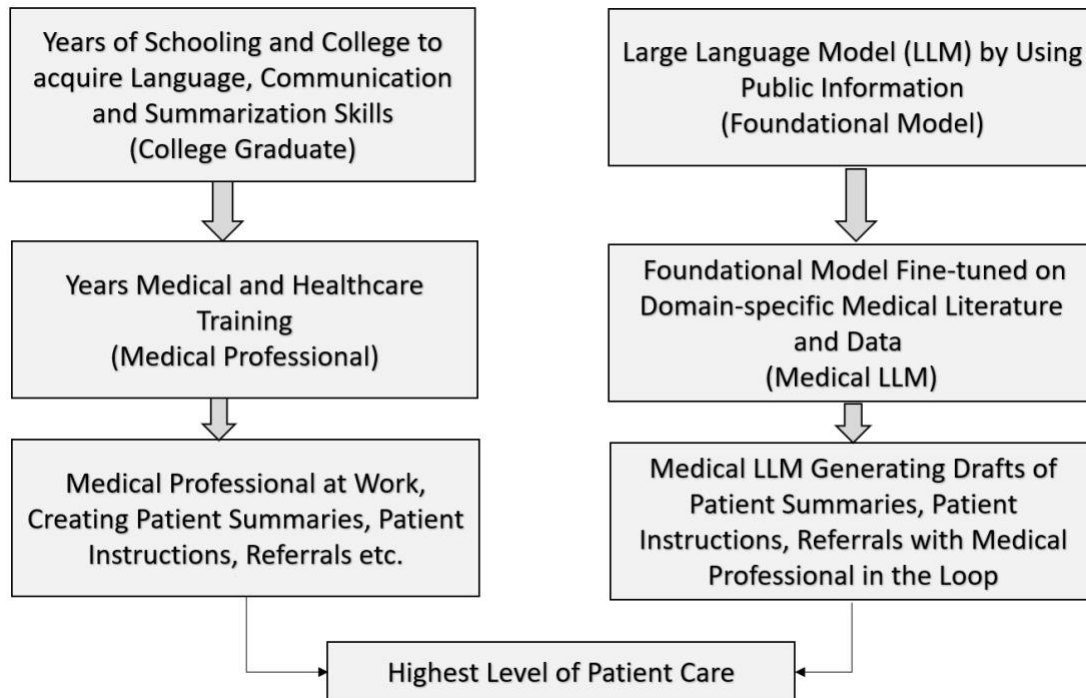


Figure 3-1: Comparing medical domain to medical LLMs.

The medical professionals in the US typically go through regular schooling, then 4-years of college before the medical degree programs. In the education leading up to an undergraduate degree, they pick up broad subjects, language, and communication skills. The Foundational LLMs try to simulate similar knowledge by scraping the corpus of knowledge from the web. The huge computation power and large volumes of input training data, tried to create models that can act like humans in limited sense, using natural language for interaction. In the next stage, medical professionals go through specialized education such as physicians would do MD, followed by Residency and optionally Fellowship or similar part for nursing, therapist, or physician assistant, to learn about the healthcare principles.

Similarly, Foundational LLMs can be trained and fine-tuned using medical data, to create medical LLMs. Such training data is very scarce and subject to data privacy and security, making this a challenging task. Finally, when trained healthcare professionals start working in a medical setting, this is comparable to a deployed Healthcare GenAI solution which has access to specific patient data and can create patient summaries and similar drafts for review by trained healthcare professionals, to augment them. Use of Retrieval-Augmented Generation (RAG) also applies at this stage of the solution. RAG allows use of enterprise information to enhance the results of the LLM, without exposing the internal data to the public LLMs or outside the organization. RAG is further described in 3.1.2.

Next, let's look at each of the four use cases that we will deep-dive into.

### 3.1 GENERATED PATIENT SUMMARY

GenAI is making inroads into healthcare by automating tasks and optimizing processes such as the generation of patient summaries. These summaries can be crucial for physicians and other healthcare professionals to quickly understand a patient's medical history, current conditions, and treatment plans before visits or procedures. Prior attempts with limited success, for similar use case without the use of GenAI is described here<sup>22</sup>. Let us look at how the use of GenAI enhances this scenario.

#### 3.1.1 PROCESS OPTIMIZATION

The process optimization involves three steps:

- **Data Aggregation:** Medical LLMs can aggregate data from multiple sources such as EHRs (Electronic Health Records), lab results, and imaging studies.
- **Natural Language Processing:** Advanced natural language processing techniques are employed to interpret clinical notes, and other free-text fields, converting them into structured data or summarizing the essential points. The fine tuning of the Foundational LLM, with clinical data, helps to achieve this.
- **Summarization:** The model generates concise yet comprehensive patient summaries that encapsulate all the necessary clinical information. The process creates a draft for the medical professional such as physician assistance to review it and approve it for use by the physician before a patient visit. The human in the loop ensures patient safety and compliance.

#### 3.1.2 BENEFITS

The use of GenAI in creating patient summaries offers several distinct advantages. First, automation significantly enhances efficiency, freeing up valuable time for healthcare professionals to devote more attention to direct patient care. Second, the generated summaries bring a level of consistency in both format and detail, which makes it easier for healthcare providers to quickly understand the patient's medical situation. Third, these models can be customized to cater to the specific informational needs of various medical specialties, thereby increasing the summaries' relevance and utility. Lastly, these GenAI systems can be configured to provide real-time updates, ensuring that healthcare professionals have access to the most current patient information. Collectively, these benefits not only improve the quality of healthcare delivery but also contribute to more informed and timely decision-making.

While the advantages of using GenAI for patient summaries are clear, the technology faces several challenges that need to be addressed. First and foremost is the issue of data quality: the

---

<sup>22</sup> <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7225507/>

reliability of generated summaries hinges on the accuracy and completeness of the data used to train the models. Subpar or incomplete data can yield summaries that are misleading or incorrect. Another challenge is explain-ability, as it's crucial for healthcare professionals to understand the rationale behind the generated summaries to fully trust and validate the information being presented. Hallucinations in the GenAI is another challenge. Additionally, there are regulatory and ethical obstacles to consider. Compliance with existing healthcare regulations and ethical guidelines is essential to ensure that the technology does not inadvertently cause harm or violate standards. However, the absence of clear regulations makes it very difficult to determine if such applications will be deemed fully compliant and safe for the patient.

GenAI has the potential to significantly improve the efficiency and quality of healthcare delivery by automating the generation of patient summaries. However, its successful implementation requires overcoming technical and regulatory challenges. In the near term, humans in the loop, in consuming the generated patient documents, will reduce the risks. To reduce the impact of hallucinations, which are sections of the generated outputs that are not real, do not match any data the LLM has been trained on, or do not make logical sense, use of RAG is being tested. RAG combines the retrieval of curated information such as clinical data in our case and generative models to improve the performance of medical LLMs on knowledge-intensive and domain-specific tasks. The seminal paper on RAG by Patrick Lewis et al, can be found here<sup>23</sup>. Our end-goal is to optimally use the combination of fine-tuning and RAG, to improve the reliability of the generated documents and minimize the dependence on the human in the loop.

Next, we look at the samples of the generated summaries and outputs from our GenAI systems. Here is an example of a generated medical patient summary using pre-trained Medical LLM using transformer-based neural network architecture.

**Patient Name:** John Smith

Age: 52

Gender: Male

Chief Complaint: Chest pain and shortness of breath

History of Present Illness:

Mr. Smith is a 52-year old male with a 20 pack year smoking history who presents with acute onset of chest pain and shortness of breath for the past 3 hours. The pain is described as "heavy" and "crushing" in nature, 8/10 in intensity, located in the substernal area with radiation to his left arm. Onset was sudden when he was doing yardwork. He also complains of diaphoresis and feeling lightheaded. He has nausea but no vomiting. He took 324mg of aspirin by EMS during transport with mild relief.

Past Medical History:

Hypertension, Hyperlipidemia, Type 2 Diabetes Mellitus, Obesity

Past Surgical History:

Appendectomy

---

<sup>23</sup> <https://browse.arxiv.org/pdf/2005.11401.pdf>

## Driving Healthcare Transformation Through Generative AI

---

### Medications:

Lisinopril 20mg daily, Atorvastatin 40mg daily, Metformin 1000mg BID, Aspirin 81mg daily

### Social History:

20 pack year smoking history. Quit 5 years ago. No alcohol or recreational drug use.

### Family History:

Father with history of CAD, had MI at age 62.

### Review of Systems:

Positive for chest pain, shortness of breath, diaphoresis, lightheadedness, and nausea. Negative for all other systems.

### Vital Signs:

BP 148/92, HR 112, RR 28, Temp 98.5F, SpO2 92% on room air

### Physical Exam:

General: Appears uncomfortable, diaphoretic

Cardiovascular: Tachycardia, regular rhythm. No murmurs.

Pulmonary: Tachypneic. Lung sounds clear bilaterally.

Extremities: No peripheral edema.

Neurologic: Alert and oriented x3.

This patient summary highlights the key details of Mr. Smith's history and examination pertinent to his chief complaint of chest pain and concern for acute coronary syndrome. The AI extracted relevant medical details and formatted them into an organized summary. The workflow routes this to a professional like physician's assistant as a draft, who verifies it and approves it for the physician to look at before seeing a patient. This can help physicians quickly review the case before deciding on next steps in management and treatment. The AI can synthesize patient data into summaries like this, saving a physician's time.

## 3.2 PERSONALIZED PATIENT INSTRUCTIONS

GenAI and medical LLMs hold significant promises for enhancing the quality and efficiency of healthcare through personalized patient instructions and advanced clinical notes. By integrating data from multiple sources like EHR's, diagnostic tests, and physician notes, AI algorithms can create a nuanced patient profile. This allows the system to generate customized post-visit instructions that take into account individual conditions, history, and other contextual factors, like age and lifestyle. As new data becomes available, these instructions can be dynamically updated, ensuring they remain relevant and actionable.

Additionally, during a patient's visit, real-time data capture of vitals, form entries, and physician dictations enables the AI to generate advanced, structured clinical notes. These notes not only save time for healthcare professionals but also minimize human error, providing more accurate and insightful documentation for follow-up care. These auto-generated notes can include inferred insights, potential areas of concern, and recommended follow-up activities, thereby acting as a valuable tool for healthcare providers in decision-making.

## Driving Healthcare Transformation Through Generative AI

---

Automating these processes allows clinicians to focus more on direct patient care while also achieving a balance between standardization and personalization in healthcare delivery. Such advances can be integrated seamlessly into existing EHR systems, promoting interoperability, and facilitating a more collaborative approach to patient care. However, to unlock the full potential of this technology, it's crucial to address challenges related to data quality, the explain-ability of AI decisions, and regulatory compliance.

Here is a personalized post-discharge instruction sheet for patient John Smith:

**Patient Name:** John Smith

Discharge Diagnosis: Acute myocardial infarction

Discharge Medications:

Aspirin 81mg daily

Clopidogrel 75mg daily

Atorvastatin 40mg at bedtime

Lisinopril 10mg daily

Metoprolol succinate 25mg twice daily

Nitroglycerin 0.4mg sublingual as needed for chest pain

Follow-up Appointment:

See Dr. Jones in the cardiology clinic in 1 week. The appointment is on Tuesday, September 15th at 10:30am.

Post-Discharge Instructions:

Activity:

No heavy lifting or strenuous exercise for 4 weeks.

Increase activity gradually. Walk for 5-10 minutes 2-3 times per day and increase weekly.

Avoid baths for 2 weeks. Showers are okay.

Use caution on stairs or steep inclines initially. Take them slowly.

Do not drive until cleared by your doctor at follow-up visit.

Diet:

Follow a heart healthy low-fat, low-sodium diet. Limit red meat intake.

Eat plenty of vegetables, fruits, whole grains, and lean protein.

Restrict sodium to less than 2000mg daily.

Limit caffeine intake to 1-2 cups of coffee per day.

Medications:

Take medications exactly as prescribed. Do not stop them abruptly.

Avoid OTC ibuprofen, naproxen or herbal supplements, as they can interact with your prescriptions.

Use nitroglycerin tablet under tongue for chest pain. Call 911 if pain persists after 3 tablets.

Notify Dr. Jones if you experience medication side effects like rash, swelling, or shortness of breath.

Symptoms to monitor:

Chest pain, pressure, tightness. Use nitroglycerin as instructed.

Shortness of breath at rest or with activity.

Dizziness, palpitations, or irregular heartbeat.  
Excessive bruising or bleeding.

Call 911 or go to ER if you have:  
Ongoing chest pain or pressure  
Fainting or loss of consciousness  
Shortness of breath not relieved by rest

Please call the cardiology clinic with any questions or concerns. Seek care immediately if you experience any severe symptoms. We wish you a speedy recovery.

### 3.3 REFERRALS TO OTHER PHYSICIANS

Generative AI is becoming a pivotal tool in healthcare for streamlining the process of referring patients to other physicians or specialists. By analyzing comprehensive patient data, which may include clinical notes, diagnostic test results, and treatment histories, the AI system can identify conditions or gaps in care that require specialized attention. For example, if the AI detects consistently high blood sugar levels and associated symptoms from the patient's records, it may suggest a referral to an endocrinologist for more specialized diabetes management.

These automated referral suggestions can be embedded directly into the advanced clinical notes generated by the AI system, making it easier for the primary care physician to both recognize the need for and communicate these additional specialist consultations. This automation not only speeds up the referral process but also enhances its accuracy, helping ensure that patients receive the specialized care they need in a timely manner. In summary, generative AI contributes to making the physician referral process more efficient and precise, which in turn facilitates better and more targeted patient care.

Here is a sample cardiologist referral letter for patient John Smith:

Date: 09/12/2023

**Referring provider:** Dr. Jane White, Family Medicine

**Referral to:** Dr. James Roberts, Cardiology

**Re:** JOHN SMITH, 52 year old male with acute myocardial infarction

Reason for Referral:

Mr. John Smith is a 52 year old male with a recent hospitalization for an ST-elevation myocardial infarction involving the left anterior descending artery. He underwent successful percutaneous coronary intervention with a drug-eluting stent to the proximal LAD.

He was stabilized and discharged on appropriate medical therapy including dual antiplatelet therapy, statin, ACE inhibitor, and beta blocker. He needs close cardiology follow up to ensure optimal medical therapy and monitor for complications.

**Pertinent Medical History:**

20 pack-year smoking history, quit 5 years ago  
Hypertension, hyperlipidemia, type 2 diabetes mellitus  
Family history of early coronary artery disease

### Diagnoses:

ST-elevation myocardial infarction of left anterior descending artery (ICD-10 code I21.4)

Coronary artery disease (ICD-10 code I25.10)

Hypertension (ICD-10 code I10)

Hyperlipidemia (ICD-10 code E78.0)

Please evaluate Mr. Smith for post-MI management including:

Risk factor modification

Medical therapy optimization

Activity recommendations

Further diagnostic testing if appropriate

Cardiac rehab referral

Let me know if any additional information is needed. I appreciate your assistance in caring for my patient. Please send consultation notes so I may best coordinate Mr. Smith's care.

Sincerely,

Dr. Jane White

86 Main St.

Santa Ana, USA 12345

(201) 321-6234

### 3.4 GENAI-ASSISTED INSURANCE CLAIM EXPLANATION

Navigating the intricacies of insurance claims can be a daunting task for patients, often marred by complex terminology and confusing billing codes. Generative AI has the capability to significantly alleviate this issue by generating comprehensible, plain-language summaries of insurance claims. When a claim is processed, the AI system can scrutinize the billing codes, listed services, and other associated data to craft a straightforward explanation. This summary would outline the nature of each medical service, the amount covered by insurance, and any out-of-pocket costs the patient should expect.

In more advanced setups, the AI can even provide contextual insights. For instance, it might highlight unusual scenarios where a typically covered service was not paid for in a particular case or offer suggestions on how patients might minimize future costs. These simplified explanations could be shared with patients through an online portal, mobile application, or directly via email.

By leveraging generative AI in this way, patients are better equipped to understand their medical bills, reducing uncertainty and stress. This not only clarifies the financial aspects of healthcare for patients but also smooths out interactions with healthcare providers and insurance companies by preemptively resolving potential questions or confusions. Therefore, from the patient's perspective, generative AI can play a key role in demystifying and simplifying the often-complex process of understanding insurance claims.

Here is a sample generative AI-assisted explanation of insurance benefits for John Smith's recent hospitalization:

## Driving Healthcare Transformation Through Generative AI

---

Dear Mr. Smith,

This is an **explanation of your insurance benefits** for your recent hospital stay from 9/8/2023 to 9/10/2023 at Mercy Hospital for treatment of a heart attack.

### Hospital Charges:

The total billed charges from the hospital are \$35,000. This includes costs for the emergency department, lab tests, cardiac catheterization procedure, operating room, hospital room stay, medications, and more.

### Insurance Payment:

Your insurance, UAC Health Plan, has negotiated discounted rates with Mercy Hospital. The discounted allowed amount they will pay is \$22,000.

Your plan paid \$17,600 directly to the hospital under your in-network benefits which covers 80% of allowed costs after you meet the annual deductible. Your deductible of \$2000 was reached earlier this year.

### Your Responsibility:

That leaves a remainder of \$4,400 which is your 20% coinsurance responsibility under the plan. The hospital will bill you separately for this amount owed.

### Additional Notes:

The cardiac rehab program you will need is covered at 100% under your plan benefits.

Your ambulance transport will be billed separately. It is covered at 80% after deductible.

Your emergency room copay was \$250 and will appear separately.

Let me know if you have any other questions! I am happy to explain your plan's benefits in detail so you understand the costs associated with your care. Please reach out if you need any help with appeals or negotiating payment arrangements.

Sincerely,

[GenAI Assistant]

UAC Health Plan

## 4 HIGH LEVEL ARCHITECTURE

---

Let's look at the high-level flow of the GenAI assisted patient care in Figure 4-1. The figure shows the four different use cases we have covered in this paper. The first three use cases shown in the Figure are for the productivity and efficiency of the medical professionals as they provide service to the patients. The fourth use case is for the patients to develop a better understanding of the insurance benefits applicable to them.






Figure 4-1: Flow of generative AI-assisted patient care.

### 4.1 TECHNICAL ARCHITECTURE

In this section, we look at a similar solution, covered earlier in the literature review, to take a look into the technical architecture of such a patient care solution. Figure 4-2 shows the architecture for such a patient care solution based on GenAI:

- Data Collection – Shows how to collect curated clinical data such as physician’s notes and EHR’s.
- Pre-training – Shows how the LLM is fed the massive amount of medical text data to learn the statistical relationships between words, medical terms and phrases.
- Fine tuning – Optimization of the LLM for specific tasks and use cases related to patient care.
- Deployment of the Solution – Taking the solution to production so that it can be easily used by the stakeholders either directly or via email alerts and content.

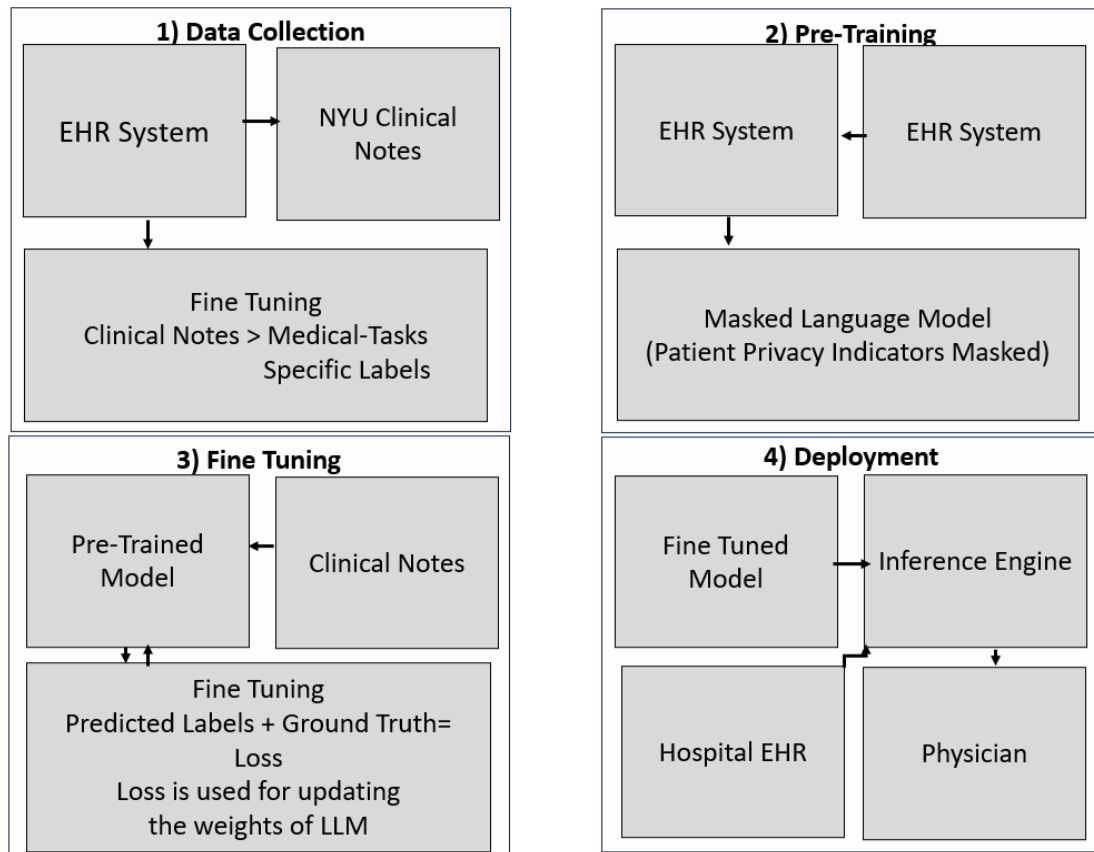


Figure 4-2: Technical architecture of a GenAI-based patient care system<sup>24</sup>.

Figure 4-2 shows how an LLM that is pre-trained, is fine-tuned with tasks specific data, to cater to the use cases that it is expected to serve. Next, such a fine-tuned model is deployed alongside a Hospital's EHR system. Such a deployment then has access to the relevant patient information, to generate summaries, personalized patient instructions at the time of discharge and referrals to specialists. Likewise, it can generate an explanation of benefits for the patients when they receive their medical bills.

## 5 DISCUSSION

In this paper we looked at a GenAI-based solution for use in the healthcare industry. Such a medical LLM solution will drive efficiency in the healthcare-resource constrained industry which has seen a lot of burnouts during and after the pandemic. The first three use cases show how it can reduce the human effort for the healthcare professionals as they take care of the patients.

Such a system can help to improve and augment patient care in hospitals and in physician practices. Since this solution involves healthcare physicians and paraprofessionals to review the GenAI provided patient-facing content, the human-in-the-loop principle ensures patient safety

<sup>24</sup> <https://www.nature.com/articles/s41586-023-06160-y>

## Driving Healthcare Transformation Through Generative AI

---

and compliance. As more such GenAI-based healthcare and clinical decision support solutions will be piloted and adopted, the industry will regulate itself to improve the patient privacy and safety considerations.

Healthcare data often contains Personal Identifiable Information (PII) which is subject to leakage via the medical LLMs<sup>25</sup>. The use of synthetic data <sup>26,27</sup> is being actively explored, for training the medical LLMs, to overcome the shortage of relevant healthcare data, due to privacy and security concerns. Such synthetic data can be additionally used to create the digital twin of the hospital and patients, to augment the GenAI-based solutions.

Mayo Clinic is generating synthetic data for histopathology<sup>28</sup>. The privacy preserving AI is emerging as a focus area for researchers, in response to the threats posed to PII by medical LLMs<sup>29</sup>. Google is working on Sensitive Data Protection to help secure GenAI workloads including those in healthcare<sup>30</sup>.

FDA introduced regulations for software as a medical device (SaMD) in 2018<sup>31</sup>. In late 2021, FDA introduced the AI/ML in SaMD Action Plan<sup>32</sup>. We will continue to see similar evolution on the digital health regulatory space. In September 2021, the US Department of Health and Human Services introduced Trustworthy AI<sup>33</sup>. National Institute of Standards and Technology (NIST) created a Generative AI Public Work Group in June 2023<sup>34</sup>.

The most recent one is the Executive Order by the US President on Oct 30, 2023<sup>35</sup>. The Executive Order also references the NIST's AI Risk Management Framework<sup>36</sup>. As the Government agencies and bodies like Industry IoT Consortium (IIC) continue to shape this landscape, we will see increased collaboration between the health care industry and software industry to roll out GenAI based solutions benefitting the patients and different stakeholders of the industry.

---

<sup>25</sup> <https://arxiv.org/abs/2302.00539>

<sup>26</sup> <https://www.nature.com/articles/s41746-023-00927-3>

<sup>27</sup> <https://www.nature.com/articles/s41746-023-00888-7>

<sup>28</sup> <https://www.mayo.edu/research/faculty/tizhoosh-hamid-r-ph-d/bio-20530617>

<sup>29</sup> <https://www.sciencedirect.com/science/article/pii/S001048252300313X>

<sup>30</sup> <https://cloud.google.com/blog/products/identity-security/how-sensitive-data-protection-can-help-secure-generative-ai-workloads>

<sup>31</sup> <https://www.fda.gov/medical-devices/digital-health-center-excellence/software-medical-device-samd>

<sup>32</sup> <https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-and-machine-learning-software-medical-device>

<sup>33</sup> <https://www.hhs.gov/sites/default/files/hhs-trustworthy-ai-playbook.pdf>

<sup>34</sup> [https://airc.nist.gov/generative\\_ai\\_wg](https://airc.nist.gov/generative_ai_wg)

<sup>35</sup> <https://www.whitehouse.gov/briefing-room/presidential-actions/2023/10/30/executive-order-on-the-safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence/>

<sup>36</sup> <https://nvlpubs.nist.gov/nistpubs/ai/nist.ai.100-1.pdf>

## **ACKNOWLEDGEMENTS**

---

The views expressed in the *OMG Journal of Innovation* are the author's views and do not necessarily represent the views of their respective employers nor those of the Object Management Group® (OMG®).

© 2024 The OMG logo is a registered trademark of Object Management Group®. Other logos, products and company names referenced in this publication are property of their respective companies.

- Return to *OMG Journal of Innovation landing page* for more articles and past editions.