



Generative AI business use cases

Healthcare

Table of contents

Introduction	1
Addressing risk	2
Unlocking potential	3
How to use this document	4
Business use cases	
Prior authorizations + appeals	5
Discharge instructions + follow-up care guidance	9
Treatment planning	13
Medical coding	17
Additional generative AI use cases for Healthcare	18
Diagnostic processes	21
Intra-Clinician communication	22
Patient documentation and record-keeping	22
Research + ongoing education	23
Administrative, non-patient care tasks	23
Conclusion	24
Appendix	25

Generative AI is transforming the healthcare industry. By alleviating administrative burdens, accelerating patient research and diagnoses, and aiding in solving complex problems, this technology enables providers to prioritize the most important aspect of their job — **improving patient lives.**

22 hours per doctor per month

time savings with Liminal-secured generative AI

\$16.6M

Potential annual efficiency gain from Liminal-secured generative AI

58%

of hospitals are looking to deploy and use generative AI within the next year

Address the risks

As healthcare organizations look to embrace the benefits of generative AI, they must address risks concerning data privacy, security and sovereignty. Implementing a layered strategy combining policy, process, and technology is key to securing patient information and complying with data protection regulations with HIPAA, CCPA, and others.

01

PHI/PII exposure

Liminal protects against the inadvertent sharing of compliance-defined terms with Large Language Models (LLMs).

02

Sensitive Data + IP leaks

Liminal allows you to define and safeguard unique intellectual property and sensitive corporate data types across all interactions with generative AI.

03

Inappropriate exchanges

Liminal defends against the ingestion or output of offensive, discriminatory, or derogatory content.

Liminal exists to help organizations unlock the **potential of generative AI** by addressing the data privacy and security risks inherent with this technology.

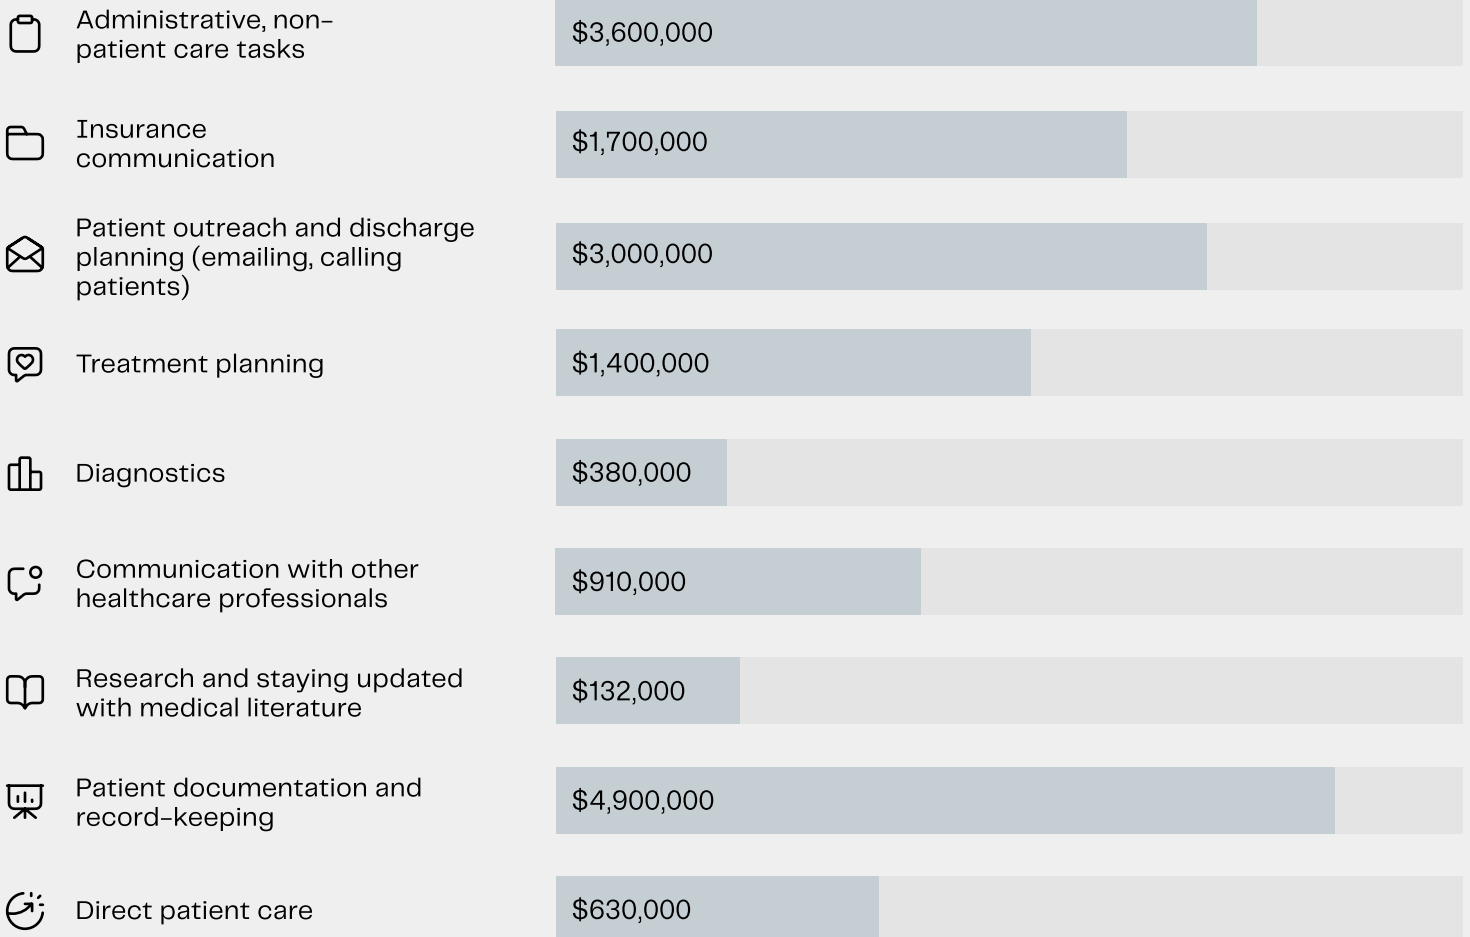
With Liminal, organizations have complete control over the data submitted to large language models (LLMs). Whether that be through direct interactions, through the consumption of off-the-shelf software with generative AI capabilities, or via the generative AI-enabled applications built in-house, Liminal's unique horizontal platform helps ensure compliance-defined data like PHI, PII, and PCI - as well as organization-specific data like intellectual property and trade secrets - aren't leaked outside your organization.

Across every generative AI model, in every application you use, and in every application you're building.

Unlock the potential

This document showcases several real-world applications of generative AI deployed in the healthcare industry today, where the aforementioned data risks have been addressed with Liminal. For each scenario, we describe the use case, the problem it addresses, its benefits, and the ROI that a Liminal-protected generative AI solution can offer.

ROI from Liminal ranges anywhere from **5x–30x on a single use-case**, and each subsequent use case compounds your return.





To help **estimate the implementation effort** and **assess the business value** of deploying specific use cases, we've developed a high-level schema for evaluation.

Effort

Simple

Generative AI solution can be implemented in a matter of hours

Moderate

Solution implementation can typically be completed within a few weeks

Detailed

Planning and solution deployment is more complex and can take several months

Value

Low

ROI on the Liminal-enabled generative AI solution is 1x-3x

Medium

ROI on the Liminal-enabled generative AI solution is 4x-6x

High

ROI on the Liminal-enabled generative AI solution is greater than 7x

**Please note: these implementation timelines refer to the work required to research and deploy a generative AI solution. The deployment of Liminal is a one-time process and can be completed in under an hour.*



Use case

Prior Authorizations + Appeals

Implementation Effort Simple

Business Value Impact High

 Insurance communication

Department/Resource: Internal Medicine Physicians, Specialists

With Liminal, physicians can conservatively **save 10 minutes per day** on the creation of appeals and prior authorization communications, which equates to nearly **40 working hours per physician per year** that can instead be devoted to delivering care.



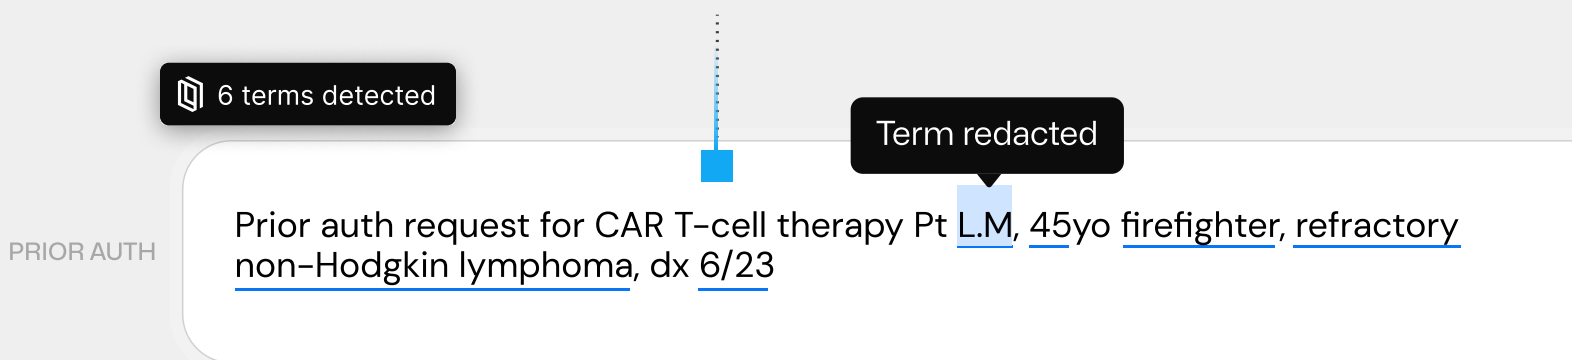
Prior Authorizations + Appeals

What

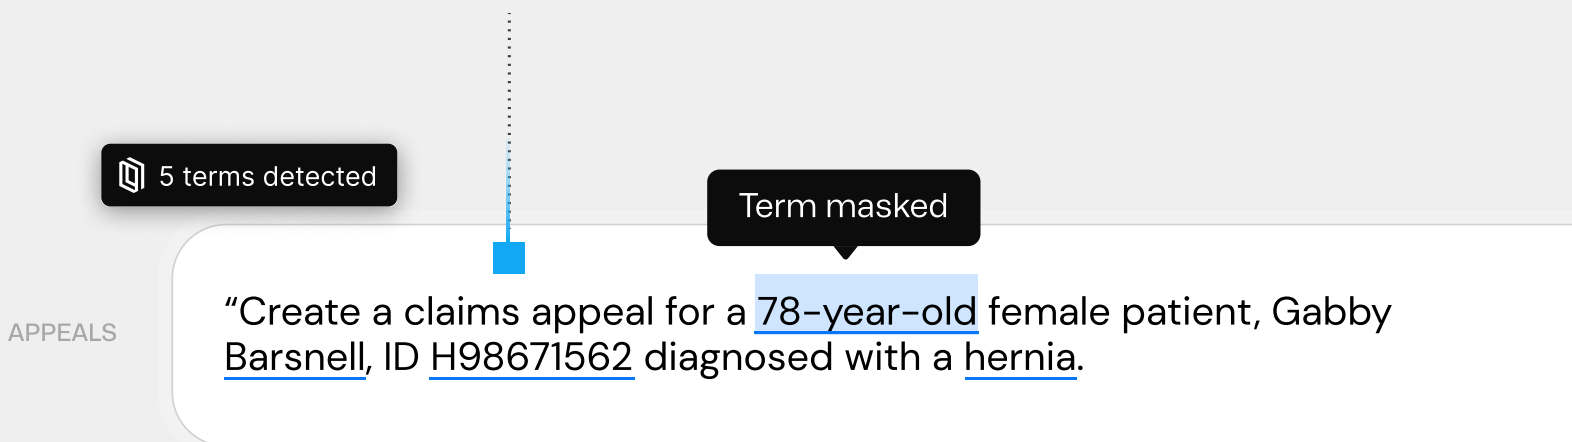
Patient-specific communications to and from insurance providers are critical to ensuring appropriate and timely treatment. Obtaining authorization and addressing coverage denials - while necessary - are pervasive, arduous tasks that detract from clinicians' ability to spend time with patients.

How

Leveraging Liminal Spaces, your practitioners can now quickly and efficiently use generative AI to safely prepare the personalized communications necessary to ensure appropriate ongoing patient care.



The terms underlined in blue are identified data types that have been detected as sensitive, and to which organizational policy is being applied. The hover state boxes over each identified term highlights how that data is being treated prior to its submission to the generative AI tool.





Prior Authorizations + Appeals

Internal Medicine Physicians, Specialists

Output

The cleansed prompt is then submitted, and the generative AI model generates an output. Through a process called rehydration, outputs are reviewed upon return, and the protected terms are re-seeded back to their submission state. Each of the identified terms that was detected and protected prior to submission has been highlighted for display purposes.



Entire time spent writing a claims appeal:

Less than 1 minute

ACME Hospital Space
Prior Authorization Request ↗
New Chat +

[Date]

[Insurance Company Name] [Insurance Company Address] [City, State, Zip Code]

Re: Prior Authorization Request for CAR T-Cell Therapy

Dear [Insurance Company Name],

I am writing to request prior authorization for firefighter, who has been diagnosed with Diagnosis. I believe that treatment option for his condition is the most appropriate.

Patient Information:

- Patient Name: Mr. L.M.
- Date of Birth: [Patient Information]
- Diagnosis: Refractory

Treatment Information:

- Proposed Treatment: [Patient Information]
- Date of Diagnosis: Jurisdiction

Clinical Justification: [Proposed Treatment], and why CAP is the most appropriate clinical evidence or guideline.

ACME Hospital Space
Appeals letter ↗

[Doctor's Name] [Doctor's Office Address] [City, State, ZIP] [Phone Number] [Email Address]

[Current Date]

Humana, Inc. P.O. Box 14601 Lexington, KY 40512-4601

Subject: Medical Necessity Appeal for Patient Gabby Barsnell, ID: H98671562

Dear Sir/Madam,

I am writing on behalf of my patient, Gabby Barsnell, who was recently denied coverage for a medically necessary surgery to repair a hernia. I am formally requesting a review and reversal of this denial. It is my professional opinion that this surgery is medically necessary and that the refusal of insurance coverage may put my patient's health at serious risk.

Gabby Barsnell is in her late 70s and the hernia, if left untreated, could lead to severe complications that can threaten her life. Her condition has progressively worsened, causing her extreme discomfort.



Prior Authorizations + Appeals

Internal Medicine Physicians, Specialists

Detailed Assumptions and Impact

5k	total employees
300	Internal Medicine Physicians
100	Specialists
\$100	Specialist organizational cost per hour
\$100	Internal Medicine Physician organizational cost per hour
220	number of working days per Physician and Specialist per year
30	minutes spent on appeals and authorizations per Physician and Specialist per day

Creating Appeals and Authorization Requests

Without Liminal	30 mn per day/per doctor	110 hrs per year/per doctor
With Liminal + Generative AI model (GPT-4, Gemini, or any other model)	33% faster	40 hrs saved per doctor per year

\$1.6M
annual efficiency gain

14,500
hours saved

 Implementation time: less than an hour*


*Really, just 1 hour to get this into the hands of your team. For more details on implementation, see Appendix A

Use case

Discharge Instructions + Follow-Up Care Guidance

Implementation Effort Simple

Business Value Impact High

 Patient outreach and discharge planning

Department/Resource: Nurses

With Liminal, **nurses can save on average 12 minutes each day** on the creation of personalized discharge instructions, translating to over **40 hours per nurse per year** that can be redirected towards patient care.

What

Personalized, treatment-specific discharge instructions and follow-up care guidance not only enhance patient safety and satisfaction but also increase post-procedure compliance. Moreover, they play a crucial role in reducing repeat emergency department visits*. However, creating these personalized instructions can be time-intensive.

Utilizing Liminal’s HIPAA-compliant platform, nurses can securely leverage generative AI to lessen the time required to produce personalized patient discharge instructions.

How

Leveraging Liminal Spaces, your nurses can now securely use generative AI to streamline the preparation of the personalized discharge instructions necessary to ensure appropriate ongoing patient care.

15 terms detected

Name: Danny Hernandez
DOB: 1/25/61
SSN: 141-00-8771
MRN: 000123
Treatment Notes:
Summarize discharge info for pt DH, 63/M, 123 Elm, Springfield. Post-angioplasty. Rx: Aspirin, Lipitor, Metoprolol. Monitor for CP, SOB. F/U in 2 Wks.



The terms underlined in blue are identified data types that have been detected as sensitive, and to which organizational policy is being applied. The hover state boxes over each identified term highlights how that data is being treated prior to its submission to the generative AI tool.



Discharge Instructions + Follow-Up Care Guidance

Nurses

Output

The cleansed prompt is then submitted, and the generative AI model generates an output. Through a process called rehydration, outputs are reviewed upon return, and the protected terms are re-seeded back to their submission state. Each of the identified terms that was detected and protected prior to submission has been highlighted for display purposes.



Entire time spent writing discharge instructions:

Less than 1 minute

ACME Hospital Space

Discharge summary



Dear Danny Hernandez,

In relation to your recent hospitalization and procedure, we have provided the following summary and discharge instructions:

Discharge Summary for Patient: DH

- Age: Early 60s
- Gender: Male
- Location: Pennsylvania
- Recent Procedure: Angioplasty

Prescribed Medications:

- Aspirin: Please take as directed by your physician, typically one per day.
- Lipitor: This is to manage your cholesterol levels. Follow the prescribed dosage from your physician.
- Metoprolol: This is for your blood pressure. Follow the prescribed dosage from your physician.

Symptom Monitoring:

- Chest Pain (CP): Contact your healthcare provider if you experience any unusual or persistent chest pain.
- Shortness of Breath (SOB): Monitor your breath. If it becomes more challenging to breathe, even during rest, contact your healthcare provider.





Nurses

Detailed Assumptions and Impact

5k	total employees
1250	Nurses
\$60	Nurse organizational cost per hour
220	number of working days per Nurse per year
60	minutes spent on Discharge Instructions per Nurse per day:

Discharge Instructions + Follow-Up Care Guidance

Without Liminal	60 mn per day/per nurse	220 hours per year/per nurse
With Liminal + Generative AI model (GPT-4, Gemini, or any other model)	20% faster	40 hours saved per nurse per year

\$3.0M
annual efficiency gain

50,000
hours saved

 Implementation time: less than an hour*

*Really, just 1 hour to get this into the hands of your team. For more details on implementation, see Appendix A



Use case

Treatment Planning

Department/Resource: Internal Medicine Physicians, Specialists

Implementation Effort Simple

Business Value Impact High

Treatment planning

With Liminal, the average time to create treatment plans is **reduced by 20%**. This efficiency saves internal medicine physicians up to **30 hours per year** and specialists about **22 hours annually**, allowing for a greater focus on delivering exceptional care.



What

Tailored patient treatment plans are the cornerstone of delivering personalized, efficient, and coordinated care, while also improving patient engagement and satisfaction. Additionally, they support cost-effectiveness and compliance with healthcare standards, benefiting both patients and healthcare providers. As such, physicians spend a significant amount of time developing the appropriate treatment steps - especially when working with new patients.

Leveraging Liminal, clinicians can save time and overcome blank state paralysis by securely using generative AI to create the appropriate treatment plan for patients.

How

Leveraging Liminal Spaces, your practitioners can now use generative AI to safely create treatment plans tailored to the individual needs of patients.

17 terms detected

Create a treatment plan for a 51-year-old female patient, Anna Nguyen, MRN 27134124 diagnosed with ICD-10 code I25.10: Atherosclerotic heart disease without angina pectoris. She has hypertension, type 2 diabetes, high LDL cholesterol, and uses metformin, lisinopril, and atorvastatin. Focus on lifestyle changes, additional medications, and specialist referrals. Include objectives, risks, and a follow-up schedule, adhering to current clinical guidelines for atherosclerotic heart disease.

The terms underlined in blue are identified data types that have been detected as sensitive, and to which organizational policy is being applied. The hover state boxes over each identified term highlights how that data is being treated prior to its submission to the generative AI tool.



Output

The cleansed prompt is then submitted, and the generative AI model generates an output. Through a process called rehydration, outputs are reviewed upon return, and the protected terms are re-seeded back to their submission state. Each of the identified terms that was detected and protected prior to submission has been highlighted for display purposes.



Entire time spent writing a treatment plan:

Less than 1 minute

ACME Hospital Space

Treatment Planning



Treatment Plan for Anna Nguyen, 27134124

- Objectives:** Aim to manage her current conditions ([atherosclerotic heart disease](#), [hypertension](#), [diabetes](#)) and lower her LDL cholesterol levels to prevent the progression of atherosclerosis, lead a healthy lifestyle, and minimize associated risks.
- Lifestyle Changes:**
 - Diet: Recommend a heart-healthy diet low in saturated fats, trans fats, cholesterol, and sodium. Encourage her to consume plenty of fruits, vegetables, whole grains, and lean proteins. Discuss referral to a dietitian for personalized dietary advice.
 - Exercise: A goal of at least 150 minutes of moderate-intensity or 75 minutes of high-intensity cardiovascular exercise weekly. Additionally, strength training exercises on two or more days a week.
 - Smoking & Alcohol: If applicable, stress the importance of smoking cessation and advise moderation in alcohol consumption.
 - Weight Management: Encourage maintaining a healthy weight. BMI should ideally be within [18.5-24.9](#). **
- Medication Adjustment**:**
 - Continue current medications ([metformin](#) for [diabetes](#), [lisinopril](#) for [hypertension](#), and [atorvastatin](#) for high LDL cholesterol levels).
 - Consider the addition of low-dose aspirin therapy for heart disease prevention after reviewing her risk for gastrointestinal bleeding.
- Specialist Referrals:**
 - Cardiologist for regular evaluation of heart function and progress of [atherosclerosis](#).
 - Endocrinologist for optimizing [diabetes](#) treatment.
 - Dietitian for further dietary advice.
- Risks:** Inaction or inappropriate management can lead to the progression of [heart disease](#), leading to possible [myocardial infarction](#) ([heart attack](#)), [stroke](#), and other serious conditions.





Detailed Assumptions and Impact

5k	total employees
300	Internal Medicine Physicians
100	Specialists
\$100	Internal Medicine Physician organizational cost per hour
\$160	Specialist organizational cost per hour
220	number of working days per Physician and Specialist per year
40	minutes spent on treatment planning per Internal Medicine Physician per day
30	minutes spent on treatment planning per Specialist per day

Treatment planning

Without Liminal	40 mn per day/per doctor	150 hours per year/per doctor
With Liminal + Generative AI model (GPT-4, Gemini, or any other model)	20% faster	20-30 hours saved per year per doctor

\$1.3M
annual efficiency gain

11,000
hours saved per year

Implementation time: less than an hour*

*Really, just 1 hour to get this into the hands of your team. For more details on implementation, see Appendix A


Use case

Medical Coding

Department/Resource: **Administrative Staff**

Implementation Effort **Simple**

Business Value Impact **High**

 Insurance communication

With Liminal, administrative staff members can save an average **12 minutes per day** on these functions, which is over **40 working hours per year** on medical coding functions.



What

Revenue cycle management is a critical function in any healthcare organization that helps ensure financial stability and ultimately, the ability to continue providing patient services. Medical coding for appropriate documentation and billing - while vital - can be labor intensive and time consuming.

By utilizing Liminal’s HIPAA compliant platform, administrative staff can employ generative AI to quickly transform physician notes and lab results into the appropriate universal medical codes necessary to facilitate record-keeping and invoicing processes.

How

Leveraging Liminal Spaces, your administrative staff can now safely use generative AI to more efficiently research and assign the appropriate universal medical codes to reflect the outcomes of a patient visit.

Powered by Liminal

Need correct ICD-10 procedure code for patient Jason Klein
 DOB: 4/9/1951
 SSN: 587-00-8545 Sensitive Information
 Notes: Pt J.K., transcatheter mitral valve repair w/ MitraClip, bifoliate
valve, 4/15.

The terms underlined in blue are identified data types that have been detected as sensitive, and to which organizational policy is being applied. The hover state boxes over each identified term highlights how that data is being treated prior to its submission to the generative AI tool.



Output

The cleansed prompt is then submitted, and the generative AI model generates an output. Through a process called rehydration, outputs are reviewed upon return, and the protected terms are re-seeded back to their submission state. Each of the identified terms that was detected and protected prior to submission has been highlighted for display purposes.



Entire time spent writing a treatment plan:

Less than 1 minute

ACME Hospital Space

Treatment Planning [↗](#)



The ICD-10 procedure code for "transcatheter mitral valve repair with MitraClip" is 02UG3JZ.



Detailed Assumptions and Impact

5k	total employees
500	Administrative Staff
\$30	Administrative Staff Organizational Cost Per Hour
220	number of working days per Administrative Staff member per year
60	minutes spent on Medical Coding per staff member per day

Medical Coding

Without Liminal	60 mn per day / per staff member	220 hours per year / per staff member
With Liminal + Generative AI model (GPT-4, Gemini, or any other model)	20% faster	40 hours saved per team member per year

\$653K

annual efficiency gain

22,000

hours saved per year

 Implementation time: less than an hour*

*Really, just 1 hour to get this into the hands of your team. For more details on implementation, see Appendix A

Additional Generative AI Use Cases for Healthcare

The use cases detailed in this document offer a solid foundation for getting started with generative AI. As your organization experiences productivity improvements from initial implementations, more applications will emerge.

To assist in further maximizing generative AI deployments, here are several additional use cases for consideration.

Diagnostic processes

Generative AI, especially in image-based specialties like radiology or pathology, can rapidly analyze images and provide diagnostic support, reducing the time physicians spend on analysis.

5k	total employees
150	Specialists
\$160	Specialist organizational cost per hour
220	number of working days per Specialist per year

Implementation effort	Business value impact	Time Spent on Task Without Gen AI (per resource)	Time Savings with Liminal (per resource)	Annual Capacity Savings	Annual Efficiency Gain
Moderate	Medium	15 mn/day	4 mn/day	2,200 hours	\$350,000



Intra-Clinician Communication

Generative AI can help draft and process communications back and forth amongst clinicians, including the automated creation of referral letters and patient summaries.

5k	total employees
300	Physicians
100	Specialists
\$100	Physician organizational cost per hour
\$160	Specialist organizational cost per hour
220	number of working days per team member per year

Implementation effort	Business value impact	Time Spent on Task Without Gen AI (per resource)	Time Savings with Liminal (per resource)	Annual Capacity Savings	Annual Efficiency Gain
Moderate	Medium	30 mn/day	4 mn/day	5,800 hours	\$667K

Patient Documentation and Record-Keeping

Generative AI can automate data entry and assist in organizing patient information, significantly reducing the time physicians spend on documentation.

5k	total employees
300	Physicians
100	Specialists
\$100	Physician organizational cost per hour
\$160	Specialist organizational cost per hour
220	number of working days per team member per year

Implementation effort	Business value impact	Time Spent on Task Without Gen AI (per resource)	Time Savings with Liminal (per resource)	Annual Capacity Savings	Annual Efficiency Gain
Complex	High	120 mn/day	20 mn/day	29,300 hours	\$3.3M



Research + Ongoing Education

Generative AI tools can quickly synthesize and summarize relevant research, helping physicians stay updated with less time investment.

5k	total employees
300	Physicians
100	Specialists
\$100	Physician organizational cost per hour
\$160	Specialist organizational cost per hour
220	number of working days per team member per year

Implementation effort	Business value impact	Time Spent on Task Without Gen AI (per resource)	Time Savings with Liminal (per resource)	Annual Capacity Savings	Annual Efficiency Gain
Simple	Low	60 mn/week	5 mn/week	1,000 hours	\$132K

Administrative, Non-Patient Care Tasks

From marketing content to organizational documentation, Generative AI can help expedite the creation of thoughtful and engaging materials, enabling marketing and administrative teams to be more productive.

5k	total employees
500	Admin + Marketing Staff
\$30	Admin + Marketing Staff organizational cost per hour
220	number of working days per team member per year

Implementation effort	Business value impact	Time Spent on Task Without Gen AI (per resource)	Time Savings with Liminal (per resource)	Annual Capacity Savings	Annual Efficiency Gain
Complex	High	150 mn/day	50 mn/day	91,600 hours	\$2.75M

Conclusion

Generative AI is a productivity game-changer for the healthcare industry. The scenarios highlighted in this document represent a fraction of the potential applications that can be deployed alongside your clinicians and staff.

Successful implementations will help drive operational efficiency, advance medical breakthroughs, and ensure your team can focus on providing the best possible patient care.

Security is the first step in any Generative AI journey

As the value of generative AI continues to proliferate, it's vital for organizations to develop a holistic data privacy and security strategy that encompasses policy, process, and technology.

Liminal serves as this essential technology layer, offering a comprehensive security solution for healthcare organizations looking to deploy generative AI. Liminal's unique horizontal platform provides cybersecurity and risk professionals complete control over data management in any generative AI interaction.

Liminal works across every generative AI model, in every application you use, and in every application you're building - and can be deployed in less than an hour.

To explore your specific use cases and the impact Liminal can help drive, visit liminal.ai/start

Appendix A

To assist in estimating the effort needed for deploying specific use cases within your organization, we've created a high-level schema to outline the time required for implementation.

Please note: these timelines refer to the work required to research and deploy a generative AI use case solution. The initial deployment of Liminal is a one-time process and can be completed in under an hour.

Simple Generative AI solution can be implemented in a matter of hours

Moderate Solution implementation can typically be completed within a few weeks

Detailed Planning and solution deployment is more complex and can take several months

Each of the detailed use cases presented in this document are classified as Simple implementations.

For each, the most efficient strategy is to leverage Liminal Spaces to securely enable direct chat interactions with Azure OpenAI, Google Gemini, GPT-3.5, GPT-4, Claude 2, or any other generative AI model.

Setup

- 1 License a generative AI model (or models) for your organization
- 2 Share the API keys for those models with Liminal via the Liminal Admin Console
- 3 Determine which users have access to which models (can be done via IdP integration or on a one-off basis)
- 4 Set your security policies around what types of data can be shared, and how that data should be treated

Appendix A (cont.)

From here, your designated users can use Liminal Spaces to securely leverage generative AI for the streamlined creation of appeals letters, requests for authorization, discharge instructions, treatment plans, and medical coding documentation - all while helping ensure your security policies are enforced and your sensitive data remains protected.

To get started, request a demo, or discuss your specific use cases, visit **liminal.ai/start**.



Liminal empowers regulated enterprises to securely deploy and use generative AI. With Liminal, organizations have complete control over data privacy, security, and sovereignty - across any generative AI model, in every application you use, and in every application you build. Learn more about horizontal security and Liminal by visiting [**liminal.ai**](https://liminal.ai).