Secure Deployment of Large Language Models in HEOR: Governance, Infrastructure, and Risk Mitigation Best Practices

MSR130

Barinder Singh, Haseeb Raza, Inderpreet Singh Marwaha, Shubhram Pandey, Ritesh Dubey, Rajdeep Kaur
Pharmacoevidence, Mohali, India

INTRODUCTION

- Large Language Models (LLMs) are rapidly gaining adoption across Health Economics and Outcomes Research (HEOR), with growing recognition that responsible deployment requires dedicated guidance on security, privacy, and reproducibility
- However, the use of on consumer-grade artificial intelligence (AI) interfaces may expose organizations to substantial risks spanning data privacy, regulatory compliance, intellectual property protection, and AI-specific vulnerabilities such as prompt injection and model retraining on proprietary inputs
- This study aimed to propose a regulatory-aligned framework that lays out secure and compliant LLM deployment principles for responsible use within HEOR workflows

METHODS

- A regulatory-aligned framework was developed to outline principles for the secure, compliant, and scalable deployment of LLMs in HEOR
- Healthcare data protection regulations and regional data-sovereignty requirements were reviewed to ensure compliance with applicable privacy and data governance standards
- Enterprise security and infrastructure standards were evaluated, including internationally recognized information security frameworks and zero-trust architecture principles
- Emerging AI governance practices were reviewed: prompt-injection defense, model retraining policies, output supervision, and permanent audit logging norms
- Potential security risks were systematically assessed across the HEOR AI lifecycle, covering data access, model interaction, and output management
- Identified risks were consolidated into core mitigation principles, organized as five deployment pillars and visualized as a sequential filtering funnel (**Figures 2**)

RESULTS

- The proposed framework established a five-pillar architecture for secure and regulatory compliant LLM deployment within HEOR workflows (**Figure 2**):
  - **Enterprise-ready infrastructure** using private, high-security environments that excluded public LLM interfaces lacking data privacy protections or retraining safeguards.
  - **Authentication Control** that **e**nforces strict role-based access (RBAC) through robust multi-factor authentication (MFA), enabling project-specific access to confidential data
  - **Regulatory compliance mechanisms** aligned with data protection rules and organizational AI governance requirements, including end-to-end encryption and permanent audit logging
  - **Automated governance agents** that integrates real-time guardrails and supervisory agents that enforce usage policies and pre-emptively block unsafe or non-compliant outputs
  - **AI-specific risk mitigation controls** designed to address emerging threats, including prompt injection and unintended model behaviors
- All LLM requests were processed sequentially through these 5 control layers to generate secure, traceable, and audit-ready outputs, with verification performed at each checkpoint (**Figure 2**)
- Together, these measures mitigate critical risks tied to data sovereignty, regulatory non-compliance, and AI-specific attack surfaces across the HEOR LLM lifecycle (**Figure 3**)

Figure 1. From Successful GenAI Proof-of-Concepts (POCs) to Trusted Enterprise Deployment

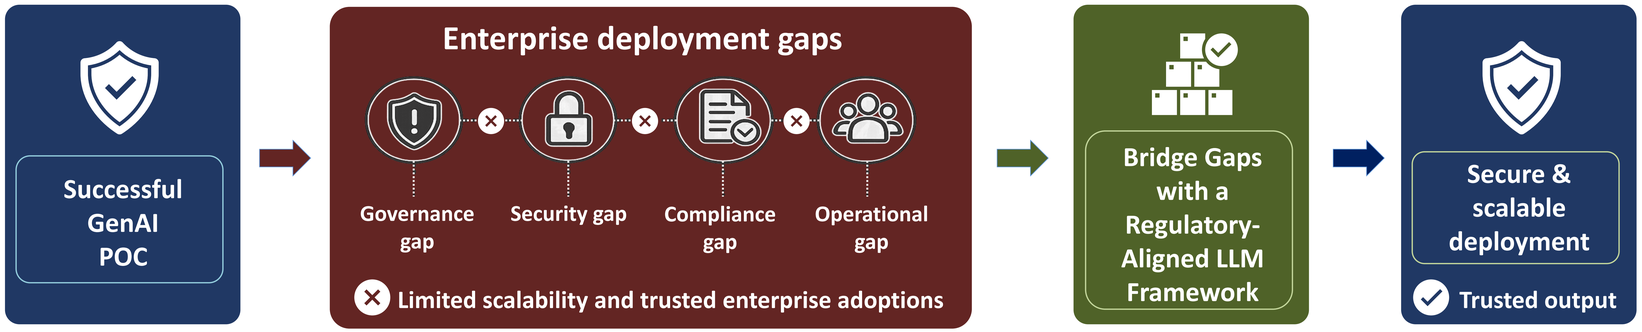

Figure 2. Five Pillars of Secure Validation Pipeline for Enterprise HEOR LLM Deployment

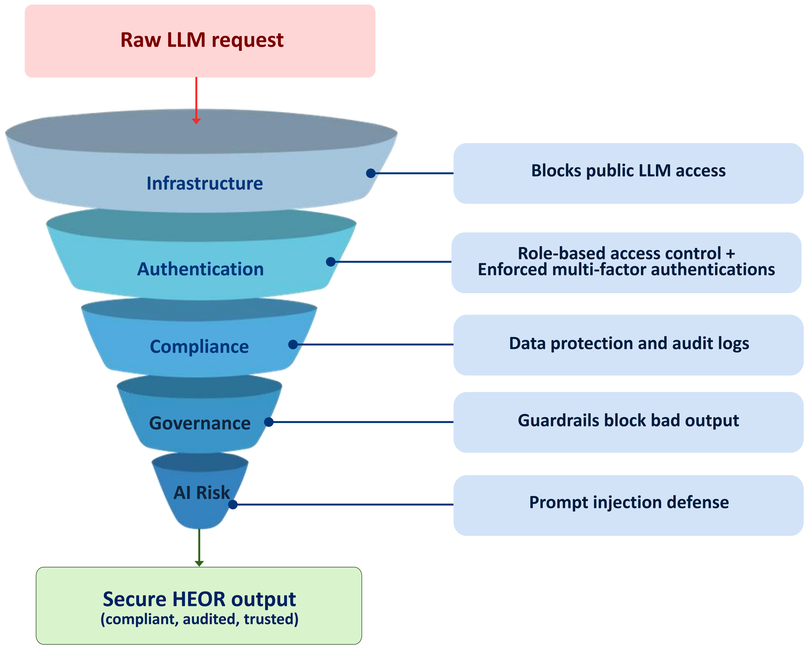

Figure 3. Risk Profile of HEOR LLM Deployment Without vs With the Proposed Five-Pillar Framework

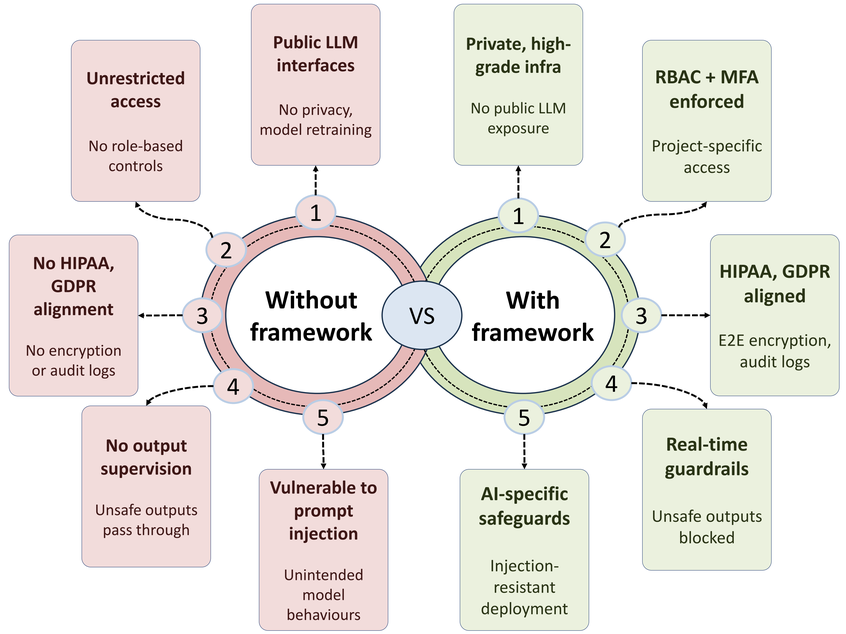

- A successful GenAI proof-of-concept does not automatically translate into scalable enterprise deployment.
- Production-grade HEOR applications require secure IT grade infrastructure, governance frameworks, regulatory compliance, access controls, and AI-specific risk mitigation to ensure safe, trusted, and scalable end-user adoption

CONCLUSIONS

- Compliant infrastructure, strong access controls, and AI-specific safeguards together enable secure and responsible LLM adoption in HEOR, going beyond model performance considerations
- The five-pillar framework provides actionable, regulatory-aligned guidance for safe and scalable operationalization of LLMs within HEOR environments and is aligned with ISPOR Working Group recommendations on AI security and privacy

ISPOR, May 17-20, 2026, Philadelphia, PA, United States