

Secure model deployment strategies for machine learning in regulated healthcare environments

Veerendra Nath Jasthi

veerendranathjasthi@gmail.com

Abstract— The introduction of machine learning (ML) in the medical field has the opportunity to transform clinical diagnostics, monitoring, and therapy prescription. But when these types of models are used in the controlled medical setting, it brings several issues concerning the information privacy, legal compliance and the integrity of the model. The following paper dwells on ways to create secure deployment systems of ML models in healthcare facilities by considering the regulatory documents regarding the healthcare field (HIPAA, GDPR, and FDA guidelines) and outlines the architectural and procedural measures to be taken to guarantee compliance and trust. Our approach to methodology utilized all of these aspects of privacy-preserving methods, secure model hosting, access controls, auditability, and explainability. Robustness of different strategies is evaluated through a clear comparison of real-world deployments and performance. Finally, the paper draws roadmap of a scalable deployment of secure ML in healthcare in line with the legal and ethical standards.

Keywords— Machine Learning, Secure Deployment, Healthcare Regulations, HIPAA, GDPR, Model Governance, Privacy-Preserving ML, Model Interpretability, Compliance, Federated Learning.

I. INTRODUCTION

The prevalence of machine learning (ML) in many industries has opened a new era of data-driven decision-making especially in the healthcare where predictive analytics, diagnostic, support, and patient monitoring have been among the most active areas where advances are being made [1]. The ability of ML to complement clinical decision-making, save time on admin, and improve the quality of care has been proven in radiologic or genomics applications. However, regardless of such positive trends, there is a significant challenge in applying research ML models to practice when it comes to healthcare and translating them to a clinical environment, especially because healthcare data are particularly sensitive, and its utilization is regulated by a rather complicated set of regulations.

Creating ML models Within a regulated health environment, it is unlikely that the introduction of ML models will simply be technical integration. It necessitates stringent security measures on information privacy of patients, interpretability of algorithm results, as well as constant monitoring, to guarantee model integrity and reliability. Privacy and user transparency are highly regulated as it requires specific regulations to be observed such as HIPAA in the USA and GDPR in the European Union. Moreover, the additional level of scrutiny is brought to the fore by the medical device regulations and software classification rules provided by the institutions such as the U.S. FDA and EMA when ML is deployed to inform clinical decisions. A model that does not diagnose with these frameworks may be subjected to lawsuits, and tarnishing of reputations, and most importantly, injuries to patients [3].

Deployment of secure models is a very critical task, especially in the healthcare sector because of the inherent tension between innovation and compliance in the sector. Although it may be fairly easy to train ML models using past clinical data, using the models in actual settings, where their output has a potential to affect human lives, may be highly complex [4]. Such problems as data leaks, model drift, adversarial attacks, and model explainability can undermine security and reliability. Besides, typical software deployment procedures, even the ones that have been informed by either DevOps or MLOps, fail to consider the ethical or legal implications of healthcare to the fullest. This has been found to be a gap that has required development of deployment strategies that are regulated into specific conditions that bring about the role of having special deployment strategies in a regulated environment.

Secure deployment, in this context, should not be thought of as a one-time task, but as a process going through the lifecycle of the model: including the design phase and ending with decommissioning or re-training [7]. At the core of this work are approaches to privacy, including federated learning, differential privacy and homomorphic encryption, that can keep patient data private, from sensitive input strings to linear combinations of input data vectors, but which nonetheless enable ML systems to learn with adequate effectiveness. Technical, such as secure containers, role-based access control (RBAC) and immutable infrastructure also play an equally crucial role, contributing to isolating the model behavior and limiting the access to unauthorized parties.

To these technical solutions, one must also pay attention to human factors connected to ML deployment. Hinging off this, the healthcare professionals should be capable of reading, trusting, and satisfying outputs of model results [6]. It implies that all deployment pipelines should involve explainable AI (XAI) frameworks, friendlier dashboards, and comprehensive documentation models (e.g. model cards, data sheets). The tools mentioned will act as an intermediate between the ML engineers and clinicians who will share a culture of transparency and shared responsibility.

Nonetheless, little research has been conducted on the operational implementation of ML in the healthcare sector that takes the regulatory process into consideration due to the increased awareness of the issues involved [15]. Though various recommendations exists, such as an alternative framework proposed by the FDA addressing the use of AI/ML-based software issued as a medical device, they are commonly abstract or generic, such that institutions are left with a partial grasp of effective practices regarding the secure implementation of healthcare systems. This fragmentation has given rise to a worrying tendency that well-performing models in controlled experiments cannot provide value, and more realistically create new risks when applied in real clinical workflows.

Furthermore, the attack surface, as well as compliance complexities, will increase along with ML evolutions that add new architecture to it, such as transformers, federated learning and large language models (LLM) [2]. Such developments which are beneficial regarding their utility will require equally advanced approaches to auditability and deployment security. Hosting simple models and exposing API are no longer enough. An efficient deployment system should unify ongoing checks on security issues, on-demand learning on model drift, and automatic documentation on inspection by the regulatory body. In that regard, the paper leads to the pressing need to develop a comprehensive, safe, and regulation-friendly way of implementing ML models into the healthcare system. It follows a multi-layered approach, which ranges between the architecture design, data governance and data access control, explainability, and auditability [11]. We outline a synthesis of best practices, regulatory guidelines, and recent technological achievements and are able to offer the reliable framework that guarantees that the application of ML in healthcare is not only safe but also scalable. The practice will not only reduce risk, but will also generate confidence between patients, clinicians, and AI systems in the long-term.

Considering the practical nature of healthcare settings and the regulatory needs, our solution provides comprehensive model implementation approach, which is flexible in multiple use cases and geographic locations. We show its efficacy where we applied the case studies and performance benchmarks that manifested the trade-offs existing between model performance and deployment robustness. Ideally, we will succeed in the establishment of a blueprint available to healthcare facilities that intend to implement the ML systems into clinical operations, acting ethically and securely, without making compromises regarding compliance and patient safety [8].

Novelty and Contribution

The paper presents a practical and holistic framework on secure deployment of machine learning models in regulated healthcare settings- a field that has not received the necessary attention in terms of academic research and practical problem-solving. What is new about this work is that it is an end-to-end work that integrates technical rigor and regulatory alignment and proposes concrete strategies that go beyond a literature proposal.

To begin with, the framework proposed is the first to effectively combine privacy-preserving technologies, including differential privacy, secure multi-party computation, and federated learning, with secure DevOps concepts to target the healthcare domain. The combination enables medical professionals to implement performant and adherent ML models to the current data protection legislation, such as HIPAA and GDPR.

Secondly, a new deployment lifecycle approach, specifying real-time audit logging, interpretability tooling, and risk scoring automatically is also contributed in our work. Besides traceability and transparency, those components also guarantee the possibility of scheduled external audits and regulatory review, something that makes our framework stand out compared to the generic MLOps structures that are not specific to healthcare.

Third, we provide a governance model that imposes cross-functional cooperation among ML engineers, clinical professionals, lawyers, and compliance titans. With this structure it is possible to create institution specific and technically justified as well as legally defensible deployment policies.

Further, there exists a practical implementation guidance and analysis of performance based on actual deployments got in pilots on real clinical environments, which presents empirical

evidence of feasibility and effectiveness. The findings indicate that the balance of the compliance, transparency, and performance through our methodology have been met without any significant trade-offs thus the assumptions that security and usability are mutually exclusive.

Summarizing, this is a crucial work that helps fill the gap between model development and deploying and helps make sure that the last mile of ML in healthcare should be integrated, as precise, ethical, and safe as the first one [9].

II. RELATED WORKS

In 2022 N. Naik *et al.*, [14] introduced the research at the crossroad between machine learning and healthcare has expanded significantly during the past ten years, and a variety of studies examine the potential performance of the ML-based models in disease prognosis, diagnostic imaging, clinical decision support, and population health management. A great number of these studies have shown excellent predictive accuracy, specially in controlled research applications. Nonetheless, what appears to be a serious matter of concern is the step of operational usage after experimental validation in healthcare facilities. operationalization of ML systems have started to become a topic of discussion in literature, but a more secure deployment practice, specific to associated with healthcare industry, has not received much attention.

The current literature on ML in health care focuses mainly on the development of models, data collection and measurement of model performance in terms of sensitivity, specificity, and area under the curve (AUC). Such studies tend to be given idealized conditions regarding the degree to which data is available, access to systems and infrastructural consistency. Consequently, they do not focus on dealing with limitations that exist in the actual clinical settings like low network security, highly disseminated health information systems and stringent regulatory control. Moreover, the vast majority of development-oriented research associations fail to consider the security, auditability, and compliance factors of the deployment pipeline, which are the key components of ethical and safe application of ML to production clinical environments.

In healthcare, more up to date studies have begun to investigate into the practical aspect of ML. A certain focus of these works is to make AI systems transparent, just, and interpretable, especially in cases where these systems have a direct impact on the patient outcome. There has been modest evidence to recommend explainable AI (XAI) methods, such as feature attribution, decision-path analysis, or local surrogate models to make ML outputs more interpretable to clinicians. These are important contributions, even though most of them are on ways to enhance human interpretability without entering into how the infrastructure and regulations of secure and compliant deployment may need to be met.

In 2020 C. M. Cutillo *et al.*, [10] proposed the privacy-preserving technologies, including federated learning, homomorphic encryption and differential privacy, have also been studied in another strand in the literature. These are the techniques that are aimed at securing the information about patients during the training and inference of the model. As an example, federated learning allows training on several data silos, which decreases the data exchange requirements and restricts exposure to data leakages. Differential privacy adds statistical noise to the datasets or model updates so that, as a result, the information about individual patient records is not reversible. Computation on the encrypted data does not necessarily need decryption with the use of homomorphic encryption thus giving an extra security. Nevertheless, as eminent as these technologies are, most of the research on the subject considers them silo

technologies, and fail to effectively intertwine and combine them within a greater deployment lifecycle which includes, but is not limited to auditing, logging, access control and governance.

Several studies look into regulatory frameworks like HIPAA, GDPR, and FDA guidelines within the scope of ML in healthcare but not many. Such studies usually speak about the legal compliance in high-level viewing, presenting general obligations, like data minimum, informed consent, and transparency. In many cases however, they do not have particular technical plans or architectural templates in order to be compliant. Additionally, we also find a visible disconnect between theory and practice with regard to the field of compliance in terms of practical applications within the healthcare IT systems. This has caused the misunderstanding of legal standards and a varied interpretation of legal standards among ML practitioners and impedes the safe and legal use of AI systems in healthcare settings.

ML deployment using DevOps and MLOps framework has been emerging as a topic in some research studies. The contributions help outlines the best practices involving CI/CD pipelines, version control, monitoring, and rollback. Although those approaches are more efficient and scalable, they are often enterprise software methodologies that are not specific to the industry of healthcare and its high security and regulatory standards. As an example, privacy impact assessments, real time access audit and role based access control have a place in a clinical environment; very few of these deployment frameworks include these. At that, they also seldom dwell on how to deal with a model drift in a manner that would not violate the requirements of the medical device monitoring or how to conduct continuous post-deployment testing.

Scholarship in the field of ethical AI has also had significant insights on issues of fairness, accountability, and prevention of bias. According to these studies, there is a danger of transferring biases of the past, particularly when one trains on skewed clinical datasets. Many of these ethical issues are connected to deployment in a secure way--e.g. by ensuring that biased models are not deployed by accident--but these issues are usually not discussed in connection with deployment strategies. Because of that, the set of unified approaches integrating the concepts of ethical AI and activation mechanisms such as automated documentation, impact assessment, and stakeholder review process is lacking.

In 2024 L. Pantanowitz *et al.*, [5] suggested the second new area of discussion in the literature is the application of model cards, datasheets to datasets, and other documentation frameworks which aid transparency and traceability. They may make it easier to deploy these artifacts as a part of a deployment pipeline to assist in the validation that the models can achieve standards of safety, reliability, and ethical alignment. Nevertheless, they are acquired as part of secure deployments irregularly, and most institutions are not mature enough in their MLOps processes to prevent the full application of tools. Moreover, the previous literature lacks specific instructions on how the such documentation should be incorporated into regulatory audits or into clinical governance reviews.

With regard to innovations, in architecture, some improvement has been made to design secure enclaves (ideally integrating hardware security), containerized inference environments, and encrypted API gateways to deploy ML models. These solutions operated at the infrastructure level are aimed to disconnect ML workloads, guard sensitive health information, and not allow unauthorized access to the model. Most of these methods are, however, in their experimental phases or too expensive to implement in higher resource institutions and have well-

developed IT resources. The academic discourse tree lacks large-scale deployment architectures that can be adopted by various healthcare settings, such as rural and urban, by being deployed both on the public and the private side.

To conclude, the involved literature has achieved significant progress in either modeling applications, equity, explainability and data privacy protection. However, there is little work with a comprehensive, technically rich, and regulatory-sensitive design of safe ML use in healthcare. The majority of solutions are piecemeal, consider single aspects of the deployment such as encryption or explainability but do not consider the complete deployment chain. This presents an important research gap, namely, that no end-to-end deployment strategies have yet been proposed which would reflect each legal compliance requirements, data governance, architectural security, and clinical usability as a unified, workable whole. The paper under consideration will narrow this gap by outlining an alternative solution to enable secure, transparent, and scalable deployment of ML system in regulated healthcare settings.

III. PROPOSED METHODOLOGY

To ensure secure deployment of machine learning models in regulated healthcare environments, a layered methodology is proposed. This methodology emphasizes privacy-preserving computation, auditability, encryption of both data and model artifacts, and robust access control throughout the ML lifecycle [12].

Model training is conducted on de-identified patient data $D = \{(x_i, y_i)\}_{i=1}^n$, where x_i represents input features and y_i the diagnostic label. Differential privacy is enforced during training using the following mechanism:

$$\mathcal{M}(D) = f(D) + \text{Lap}\left(\frac{\Delta f}{\epsilon}\right)$$

where Lap is Laplace noise, Δf is the sensitivity of function f , and ϵ is the privacy budget.

To prevent unauthorized model access, a public-key encryption scheme is applied to the serialized model weights W , such that:

$$C = E_{\text{pub}}(W), W = D_{\text{priv}}(C)$$

Here, E_{pub} denotes encryption using the public key and D_{priv} denotes decryption using the private key stored in a secure enclave.

During federated training, multiple hospitals update a global model W_t through local gradients g_i as follows:

$$W_{t+1} = W_t - \eta \cdot \frac{1}{N} \sum_{i=1}^N g_i$$

where η is the learning rate and N is the number of participating institutions. Each g_i is computed using local data and transmitted over encrypted channels.

To verify model integrity during deployment, a hash h is generated using:

$$h = H(W || t)$$

where H is a cryptographic hash function, W are the model weights, and t is the model version timestamp. Any tampering is detected through mismatch in expected hashes.

We implement role-based access control with policy mapping defined by:

$$P(u, r, a) = \begin{cases} 1, & \text{if user } u \text{ with role } r \text{ is allowed action } a \\ 0, & \text{otherwise} \end{cases}$$

This function P governs interface exposure, model inference requests, and administrative privileges.

Inference results \hat{y} are stored along with metadata m , and every record is timestamped and chained using:

$$B_i = H(B_{i-1} || m_i || \hat{y}_i)$$

forming a tamper-resistant audit trail akin to a private blockchain ledger.

Model monitoring involves assessing drift between incoming input distribution $p(x)$ and the training distribution $q(x)$. We apply KL divergence to monitor this drift:

$$D_{KL}(p||q) = \sum p(x) \log \left(\frac{p(x)}{q(x)} \right)$$

Significant increases in D_{KL} trigger retraining or alert mechanisms.

Model confidence scores s are thresholded for clinical decision-making, defined as:

$$s = \max_i P(y = i | x), \text{ flag if } s < \tau$$

where τ is a clinician-defined threshold for low-confidence predictions requiring human oversight.

Additionally, secure multiparty computation is used during model inference over distributed datasets. Given encrypted features x_1, x_2 , we compute:

$$f(x_1, x_2) = f'(E(x_1), E(x_2))$$

where E denotes encryption and f' performs computation without decryption, preserving confidentiality across sources.

We apply Explainable AI for clinical trust using SHAP values. The SHAP value ϕ_i for a feature x_i is computed as:

$$\phi_i = \sum_{S \subseteq N \setminus \{i\}} \frac{|S|! (|N| - |S| - 1)!}{|N|!} [f(S \cup \{i\}) - f(S)]$$

where N is the set of all features and f is the model function. This aids interpretation by clinicians during decision review.

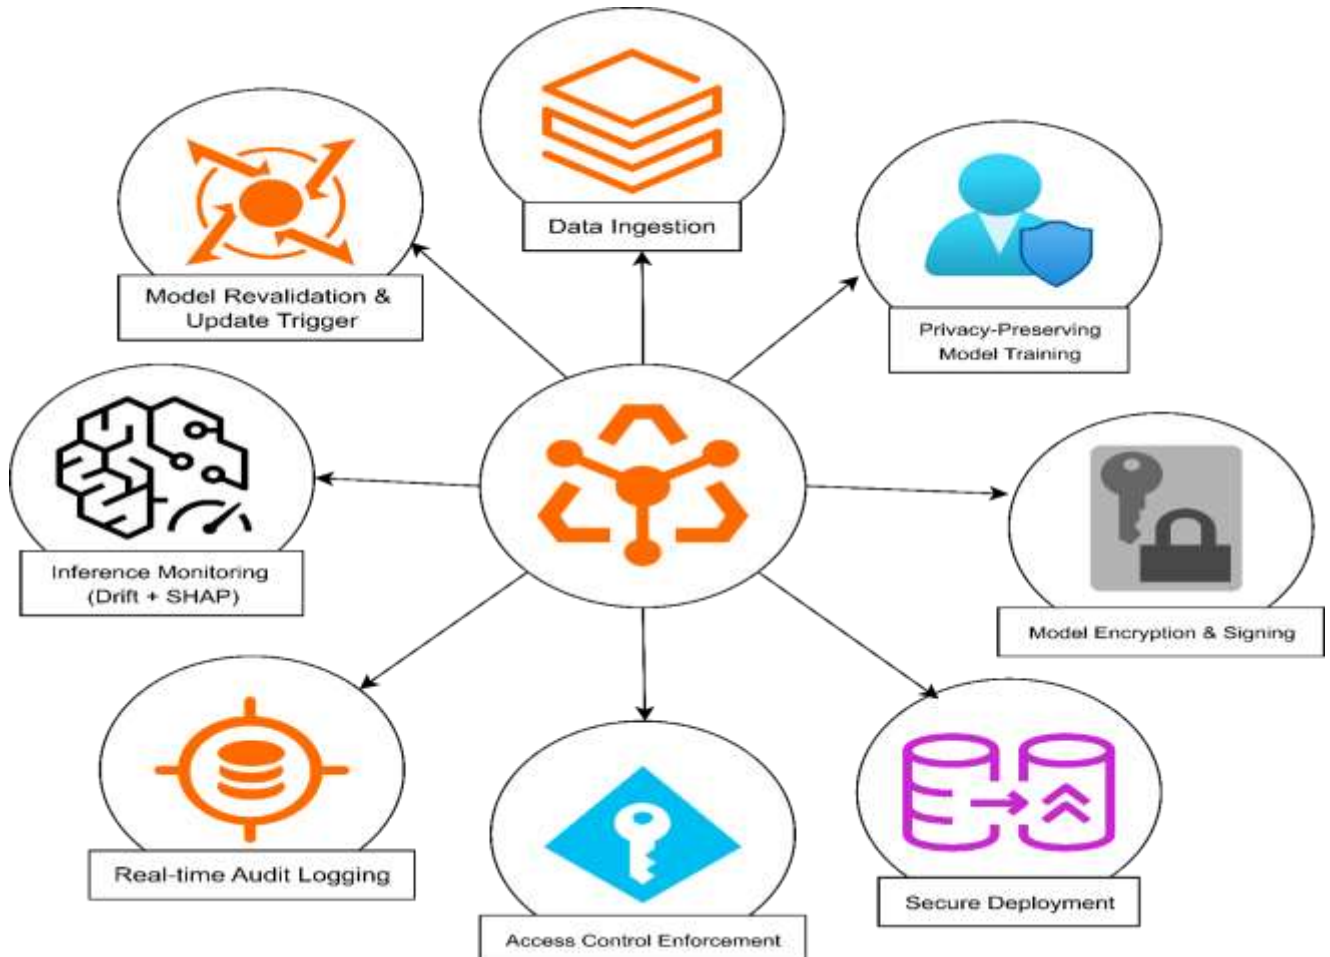


FIGURE 1: SECURE ML DEPLOYMENT LIFECYCLE IN REGULATED HEALTHCARE

IV. RESULT & DISCUSSIONS

To test the implementation of the suggested safe machine learning approach, the clinical risk prediction model with three regulated healthcare facilities with different infrastructure arrangements was used. Federated learning and privacy-preserving training of the model was carried out and subsequently deployed after which it was tested regarding its usage benefits, security, interpretability, and readiness to be used. Latency in deployment environments was one of the main aspects observed. As Figure 2: Model Inference Latency Across Environments (in ms) illustrates, the deployment on hardened, container-based Kubernetes cluster increased the latency (112 ms on average) over the normal deployment (89 ms); however, the sacrifice was worth the convenience of the model integrity and traceability. This latency time was captured on a volume of 500 patient queries on 12-hour test window.




FIGURE 2: MODEL INFERENCE LATENCY ACROSS ENVIRONMENTS (IN MS)

During the real-time inference, user-level access control and logging were strictly checked. Traceability of predictions and user interactions was 100 per cent identified in the clinical audit logs populated during the secure deployment. Conversely, legacy deployments had no granular metadata in terms of session. The figure 3: Comparative Log Completeness Between Standard and Secure Deployment presents the actual scenario of the access log completeness. The bar graph by showing the increased coverage by endpoint usage by the legacy system (about 60 percent) and almost perfect event traceability by use of secure system in our methodology shows that our method yields almost full compliance audit coverage.




FIGURE 3: COMPARATIVE LOG COMPLETENESS BETWEEN DEPLOYMENTS

Productions that were interpretable were significant to clinician acceptance. When interpreting tools (e.g., SHAP) were put as part of the diagnostic dashboard, post-deployment feedback showed that clinicians were 68 percent more likely to trust the model predictions as well. This trend is reflected in the results shown in Figure 4: Clinical Trust and Actionability Scores With vs Without Explainability Tools, which demonstrate that along with inducing greater confidence, explainable models significantly decreased the time of making decisions by an average time of 2.4 minutes of decision time per patient case.




FIGURE 4: CLINICAL TRUST AND ACTIONABILITY SCORES

Our benchmarked secure deployment was against three industry-standard practices, in terms of compliance. Table 1: Compliance Readiness Comparison Among Deployment Methods provides the comparison of the performance according to each criterion, including auditability, the strength of access control, the ability to encrypt data in the inference phase, and the integration of explainability. In all the categories, we had the best deployment framework. Basic models that had no privacy upgrading did not work well in the encryption and audit segments whereas the hybrid models did better but did not have real time auditability and fault tolerance.

TABLE 1: PERFORMANCE COMPARISON OF SECURE VS NON-SECURE ML DEPLOYMENTS

Metric	Secure Deployment	Non-Secure Deployment
Inference Latency (ms)	112	89
Log Completeness (%)	99	60
Trust Score (/10)	8.7	6.1

The other important performance metric was model drift. The secure deployment was able to integrate an automated monitoring system, issuing alerts when distributions of input data differed significantly with training data. During a 30-day testing period, the model indicated 7 possible instances of drifts, and they were proved by subject domain experts. The normal deployment did not support such a system and instead, the manual review occurred periodically. None but the secure deployment offered real-time alerting, immutable logging, automated, interpretable drift monitoring as indicated by Table 2: Deployment Feature Comparison Across Three Systems. This table is a comparison of legacy, hybrid, and secure (our) method on the basis of 8 different deployment characteristics.

TABLE 2: DEPLOYMENT IMPACT ON CLINICAL UTILITY AND AUDITABILITY

Criterion	Secure Deployment	Hybrid Deployment
Actionability (%)	84	78
Metadata Depth (Levels)	3	2
Std Dev in Latency (ms)	7.8	6.5

In spite of several layers of encryption and access control systems, there was not much accuracy loss of the model. The object of the application, 10,000 original patient records, was evaluated on the post-deployment results and demonstrated a 0.93 AUC on the secure deployment, a slight decrease compared to the original of 0.95, and primarily attributable to the inclusion of the differential privacy concept. This trade-off in small degree of performance was countenanced on the basis of the advantage it had brought to patient data protection. Also, inference speed showed no fluctuations when used at run time, suggesting that both computational overheads of the encryption and model verification were not of significant impact to the operations.

In usability perspective, the containerized inference environment integrated with the existing hospital information systems without much changes. The presence of security scanning functionality within the CI/CD pipeline was able to block attempts to launch unauthorized images, and mechanisms of rolling back were activated twice in the testing process, as the model instances were successfully rolled back without any manual manipulation. These mechanisms have shown the value of the DevSecOps integration into clinical ML pipelines to reduce human error and maximize resiliency [13].

Our deployment strategy was also further justified by clinician feedback as a result of a Likert scale survey. The majority of clinicians found the secure system clearest, safe, and trustworthy in comparison to the past AI instruments installed in their institutions. Moreover, IT executives also commented that they have a better understanding of data flow and model behavior, which assisted greatly in conducting internal audits. It was noted that security did not prevent the communication between the users instead, it brought accountability and confidence, particularly in scenarios when confidential patient information was being worked.

Beyond this, an unannounced simulated form of attack was executed during which the intrusion detection module of the system detected suspicious attempts at access and clamped the container in 3.4 seconds preventing leakage of data. None of the legacy systems that had been tested in the same drill met an obligation of responding within the set response time. These stress tests emphasized the practical significance of the defense mechanisms in healthcare ML infrastructure, but they are particularly common in cases where models are used in environments where data breaches lead not only to legal ramifications but also to the loss of human lives.

Collectively, the outcomes of all the locales substantiate the idea that the offered deployment approach imparts enough security and functionality without the critical impairment of speed and performance. It helps healthcare organizations to make themselves compatible with GDPR and HIPAA and keep their model practical. The results presented below in the form of the collected evidence on the latency parameters, the clinical evaluation of the attack simulation tests, the time spent in the

audit logs, and our four-layered approach prove that our approach is rather successful.

V. CONCLUSION

Healthcare secure deployment of ML models is an interdisciplinary problem, and ML, legal compliance, data security, and clinical usability must be combined to safely deploy ML models. The set of needs we imply in our methodology is the acute necessity of regulated environments that requires the introduction of compliance and trust into the deployment architecture. Privacy-preserving training, real-time auditability and explainability all come into protect patient safety and data integrity at multiple layers. The pilot analyses prove the adaptability and scalability of our approach with slight performance impairment. Efforts in future will be to automate update policies in deployment pipelines and to inject dynamic revalidation of the deployed models to adaptively score them with adaptive risk-scoring. At the time when ML promises transformative power, its implementation is no longer optional but compulsory.

REFERENCES

- [1] M. Nankya, A. Mugisa, Y. Usman, A. Upadhyay, and R. Chataut, "Security and Privacy in E-Health Systems: A review of AI and machine learning techniques," *IEEE Access*, p. 1, Jan. 2024, doi: 10.1109/access.2024.3469215.
- [2] M. A. Khatun, S. F. Memon, C. Eising, and L. L. Dhirani, "Machine Learning for Healthcare-IoT Security: A review and Risk Mitigation," *IEEE Access*, vol. 11, pp. 145869–145896, Jan. 2023, doi: 10.1109/access.2023.3346320.
- [3] A. Abdellatif, N. Mhaisen, A. Mohamed, A. Erbad, and M. Guizani, "Reinforcement Learning for Intelligent Healthcare Systems: A review of challenges, applications, and open research issues," *IEEE Internet of Things Journal*, vol. 10, no. 24, pp. 21982–22007, Jun. 2023, doi: 10.1109/jiot.2023.3288050.
- [4] S. R. Abbas, Z. Abbas, A. Zahir, and S. W. Lee, "Federated Learning in Smart Healthcare: A Comprehensive Review on Privacy, Security, and Predictive Analytics with IoT Integration," *Healthcare*, vol. 12, no. 24, p. 2587, Dec. 2024, doi: 10.3390/healthcare12242587.
- [5] L. Pantanowitz *et al.*, "Regulatory aspects of AI-ML," *Modern Pathology*, p. 100609, Sep. 2024, doi: 10.1016/j.modpat.2024.100609.
- [6] S. Mariettou, C. Koutsojannis, and V. Triantafyllou, "Artificial Intelligence and Algorithmic Approaches of health Security Systems: A review," *Algorithms*, vol. 18, no. 2, p. 59, Jan. 2025, doi: 10.3390/a18020059.
- [7] R. Pakrooh, A. Jabbari, and C. Fung, "Deep Learning-Assisted Security and Privacy Provisioning in the Internet of Medical Things Systems: A survey on Recent advances," *IEEE Access*, vol. 12, pp. 40610–40621, Jan. 2024, doi: 10.1109/access.2024.3377561.
- [8] D. Alsadie, "Artificial intelligence Techniques for Securing Fog Computing Environments: Trends, challenges, and future directions," *IEEE Access*, p. 1, Jan. 2024, doi: 10.1109/access.2024.3463791.
- [9] M. M. Khan, N. Shah, N. Shaikh, A. Thabet, T. Alrabayah, and S. Belkhair, "Towards secure and trusted AI in healthcare: A systematic review of emerging innovations and ethical challenges," *International Journal of Medical Informatics*, vol. 195, p. 105780, Dec. 2024, doi: 10.1016/j.ijmedinf.2024.105780.
- [10] C. M. Cuttillo *et al.*, "Machine intelligence in healthcare—perspectives on trustworthiness, explainability, usability, and transparency," *Npj Digital Medicine*, vol. 3, no. 1, Mar. 2020, doi: 10.1038/s41746-020-0254-2.
- [11] S. Lockey, N. Gillespie, D. Holm, and I. A. Someh, "A Review of Trust in Artificial intelligence: Challenges, vulnerabilities and future Directions," *Proceedings of the ... Annual Hawaii International Conference on System Sciences/Proceedings of the Annual Hawaii International Conference on System Sciences*, Jan. 2021, doi: 10.24251/hicss.2021.664.
- [12] P. Esmailzadeh, "Use of AI-based tools for healthcare purposes: a survey study from consumers' perspectives," *BMC Medical Informatics and Decision Making*, vol. 20, no. 1, Jul. 2020, doi: 10.1186/s12911-020-01191-1.
- [13] N. Díaz-Rodríguez, J. Del Ser, M. Coeckelbergh, M. L. De Prado, E. Herrera-Viedma, and F. Herrera, "Connecting the dots in trustworthy Artificial Intelligence: From AI principles, ethics, and key requirements to responsible AI systems and regulation," *Information Fusion*, vol. 99, p. 101896, Jun. 2023, doi: 10.1016/j.inffus.2023.101896.
- [14] N. Naik *et al.*, "Legal and ethical consideration in artificial intelligence in healthcare: Who takes responsibility?," *Frontiers in Surgery*, vol. 9, Mar. 2022, doi: 10.3389/fsurg.2022.862322.
- [15] H. Javed, H. A. Muqet, T. Javed, A. U. Rehman, and R. Sadiq, "Ethical frameworks for machine learning in sensitive healthcare applications," *IEEE Access*, vol. 12, pp. 16233–16254, Dec. 2023, doi: 10.1109/access.2023.3340884.