Peertechz

ENGINEERING GROUP

Trends in
Computer Science and Information Technology
OPEN ACCESS JOURNAL

ISSN: 2641-3086 Journal DOI: https://dx.doi.org/10.17352/tcsit

Opinion

Why One-Time Approval is No Longer Enough for AI Systems

Hemachandran K*

Director – AI Research Centre, Woxsen University, Hyderabad, India

Submitted : 25 April, 2026
Accepted : 28 April, 2026
Published : 29 April, 2026

***Corresponding author**: Dr. Hemachandran K, Director – AI Research Centre, Woxsen University, Hyderabad, India, Email: hemachandran.k@woxsen.edu.in

Keywords: AI governance; Continuous oversight; Model drift; Lifecycle management; AI risk management; Post-deployment monitoring; Human-in-the-loop; Responsible AI

https://www.engineegroup.us

Check for updates

Abstract

The advent of artificial intelligence (AI) systems that learn, evolve, and interact with their environment after initial deployment requires a shift from traditional "one-off" approval models of governance. This article discusses problems with one-time approvals and recommends a life-cycle approach to oversight. It explains how changes in data, operating environments, and interactions can result in model drift, new risks, and a loss of performance over time. Drawing on recent advances in AI monitoring and governance approaches, the article offers a pragmatic view that combines technical monitoring, human oversight, and organisational accountability. It presents key governance factors such as types of risk, monitoring approaches, and decision-making processes, and considers the challenges of implementation for large and small to medium-sized enterprises. The key point is that approval should be the first step, not the last, in AI governance, allowing organisations to sustain trust, compliance, and performance over the lifetime of AI systems.

Main text

For many organizations, AI governance still follows an old logic: reviewing the system, approving it, deploying it, and moving on. This model may have worked reasonably well for traditional software, but it is becoming increasingly inadequate for AI systems that continue to learn, adapt, and interact with other systems after deployment [1]. This study poses a simple but pressing question: How should organizations govern AI systems that continue to learn and interact after deployment? It proposes that one-time approval must be replaced by continuous, lifecycle-oriented oversight of AI systems.

This is where one of the most important governance challenges of 2026 will begin. A growing number of AI systems do not remain static after launch. They evolve through new data, changing environments, user interactions, and sometimes through connections with other AI-enabled tools [2]. In such settings, a one-time approval creates an illusion of control. What was safe, accurate, and compliant at the time of deployment may not remain so weeks or months later.

The problem is not that organizations are careless. In many cases, governance models have not kept up with the nature of AI. Traditional oversight mechanisms were built for systems that behaved stably and predictably. AI systems, especially those influenced by continual learning and dynamic interactions, do not always behave this way. As recent work on deployed AI monitoring notes, real-world performance can shift as operational conditions change, making post-deployment monitoring essential rather than optional [3].

Related work

Existing studies on monitoring deployed AI systems, such as the recent NIST guidance on post-deployment AI monitoring, have begun to highlight the need for ongoing performance oversight beyond initial validation [1]. Studies on continual learning models in healthcare and predictive maintenance further show how model behavior can evolve as the data and operating conditions change over time. In parallel, governance and risk reports from professional and industry bodies advocate for more dynamic approaches to AI

Citation: Hemachandran K. Why One-Time Approval is No Longer Enough for AI Systems. Trends Comput Sci Inf Technol. 2026;11(1):46-48.
Available from: https://dx.doi.org/10.17352/tcsit.000109

risk management. This study builds on these streams of work but shifts the emphasis from technical monitoring alone to the broader governance implications for boards, executives, and operational leaders.

This study makes three practical contributions. First, it clarifies why one-time approval is an inadequate governance mechanism for AI systems that learn, connect, and adapt after deployment. Second, it outlines a continuous oversight governance framework that links monitoring disciplines, human accountability, and decision-making rights across the lifecycle of an AI system. Third, it translates these ideas into concrete management issues and implementation implications that can be used by boards, risk committees, and AI governance teams.

One of the biggest concerns is the model drift. A model may perform well during testing, but its behavior can slowly change once it is exposed to real-world data. Input patterns evolve, user behavior shifts, markets change, and the context moves on. The accuracy, reliability, and fairness of the model may decline over time. This is especially serious in environments where AI systems support decisions in finance, healthcare, education, public services, and other high-impact fields. A system that was approved in January may no longer deserve the same level of trust in June [4]. For instance, a credit risk model that is calibrated under relatively stable economic conditions may begin to under estimate default risk when interest rates rise or when new customer segments start using a product. Without structured monitoring, such a drift may only be discovered after losses, complaints, or regulatory scrutiny have already occurred.

A deeper governance issue is emerging: many AI systems do not operate in isolation. They increasingly interact with APIs, decision engines, software agents, recommender systems, and other automated tools after deployment. When this occurs, new risks can emerge from the interaction itself, not just from the original model [5]. In other words, the system that was reviewed may not be the same system that actually influences outcomes in practice.

These dynamics imply that governance of AI systems must differ from traditional software governance in a few important ways.

In this governance context, risks span at least three categories: operational risks when systems fail or behave unpredictably, ethical risks when outcomes become biased or opaque, and legal or regulatory risks when behavior conflicts with emerging AI regulations or sector rules. Monitoring can be organized along similar lines, combining real-time alerts for safety-critical or customer-facing uses with periodic, deeper reviews for strategic, ethical, and compliance considerations [Table 1].

Regulatory developments have reinforced this direction. Emerging frameworks, such as the EU AI Act, sectoral guidance from financial and healthcare regulators, and voluntary AI governance standards, all emphasize documentation, transparency, risk classification, and ongoing monitoring for higher-risk- AI uses. For organizations, this means that continuous oversight is not only a good practice but also an increasingly regulatory expectation.

Table 1: comparison table: traditional vs AI governance.

Aspect	Traditional software governance	AI systems with continuous learning and interactions
System behaviour	Largely static after release	May change over time due to new data, interactions, and retraining
Approval approach	One-time approval at deployment	Ongoing approval through periodic reassessment and monitoring
Risk focus	Operational reliability and security	Operational, ethical, legal, and reputational risks
Monitoring approach	Incident-driven, periodic audits	Continuous monitoring with defined thresholds and escalation paths
Governance question	"Is the system approved?"	"Is the system still behaving acceptably under current conditions?"

Therefore, continuous oversight is crucial and can be organized as a simple four-stage- governance cycle. In practice, this cycle involves four linked activities: defining a baseline at deployment, monitoring behavior, interpreting signals, and taking action when necessary. Continuous oversight implies that governance does not conclude with deployment. It extends across the life of the system through ongoing monitoring, periodic reassessments, clear escalation thresholds, and structured reviews [6]. It asks practical questions: Is the system still performing as expected? Has the data changed? Are the outcomes still fair? Has the system started behaving differently because of the new inputs or connected tools?

Management issues

This is not a technical detail for corporate leaders. This is a management issue. Static approval provides comfort, but continuous oversight provides assurance. This difference is significant because boards, regulators, customers, and investors increasingly expect organizations to understand how their AI systems behave after deployment, not just before [7]. A governance framework that cannot see what happens in the field is unlikely to remain credible.

However, implementation is not without its challenges. Continuous oversight introduces additional costs, requires scarce skills, and depends on clear coordination between technical teams, risk functions, and business owners. Large organizations may spread these costs across multiple high-impact systems, whereas small and medium-sized enterprises (SMEs) often need lighter-weight approaches that still meet regulatory expectations. For SMEs, practical steps might include focusing on a small number of critical indicators, leveraging managed monitoring services where appropriate, and clearly documenting decisions rather than building complex governance structures.

A healthcare example illustrates this. Consider a triage support system that recommends risk levels for patients based on historical data. As new treatment protocols, demographic patterns, and documentation practices emerge, the relationships embedded in the model can shift, leading to

Citation: Hemachandran K. Why One-Time Approval is No Longer Enough for AI Systems. Trends Comput Sci Inf Technol. 2026;11(1):46-48.
Available from: https://dx.doi.org/10.17352/tcsit.000109

under-triaging of certain groups or inappropriate prioritization of cases. Continuous oversight requires that hospitals not only track headline accuracy but also examine subgroup performance, near-miss incidents, and clinician feedback over time, and then adjust or retrain the system when misalignment is detected.

The good news is that continuous oversight does not require organizations to create bureaucratic structures. It begins with a baseline definition at the point of deployment, where organizations agree on reference performance, fairness, and reliability levels against which future behavior will be judged. It continues with monitoring through mechanisms that track changes in accuracy, stability, bias metrics, anomaly rates, and operational alerts over time. The third activity is interpretation, where responsible teams review signals, distinguish noise from genuine risk, and consider how system interactions or contextual changes may contribute to emerging issues. Finally, the action step requires clear ownership; someone must be responsible for escalating issues, deciding whether a model should be retrained, constrained, paused, or withdrawn, and documenting these decisions for audit and learning purposes [8]. Typical quantitative indicators include accuracy or error rate degradation against the original baseline, drift in fairness or bias metrics for key groups, changes in anomaly rates in inputs or outputs, and patterns in user or operator feedback, such as complaint spikes, override rates, or incident reports.

Human oversight remains critical. Not all issues can be captured by automated dashboards. When AI systems are used in high-impact settings, periodic human review is necessary to assess whether the system is still aligned with business goals, ethical expectations, and regulatory obligations. Continuous oversight is not only about technology; it is about disciplined governance supported by accountable people and repeatable processes [9, 10].

Conclusion

The broader lesson is simple: AI systems that continue to learn or interact after deployment cannot be governed by static thinking alone. Approval should be seen as the beginning of oversight, not the end. Organizations that understand this early will be in a better position to build trustworthy, resilient, and future-ready AI systems. Those who do not may discover too late that the real risk was never the launch itself, but what happened quietly afterward.

AI Usage statement

The author used an AI-based language tool only for minor grammatical editing and formatting assistance. All conceptual development, theoretical framing, data collection, model specification, analysis, interpretation of results, and writing of substantive content were conducted independently by the author.

Author Contributions: The author solely contributed to the conceptualization, analysis, writing, and revision of this manuscript.

Acknowledgments

The author acknowledges the institutional support provided by Woxsen University in enabling research and discussion on AI governance and responsible innovation.

References

1. Rao AK, Keller AJ, Kalra N, Steed R, Kwegyir-Aggrey K, Klyman K, et al. Challenges to the monitoring of deployed AI systems (NIST AI 800-4). Gaithersburg (MD): National Institute of Standards and Technology; 2026. Available from: https://doi.org/10.6028/NIST.AI.800-4
2. Oxford Internet Institute. When AI systems learn during deployment, our safety evaluations break [Internet]. Available from: https://aigi.ox.ac.uk/blog-post/when-ai-systems-learn-during-deployment-our-safety-evaluations-break/
3. Smith A, Severn M. An overview of continuous learning artificial intelligence-enabled medical devices: emerging health technologies [Internet]. Ottawa (ON): Canadian Agency for Drugs and Technologies in Health; 2022 May. Available from: https://www.ncbi.nlm.nih.gov/books/NBK605105/
4. Hurtado J, Salvati D, Semola R, Bosio M, Lomonaco V. Continual learning for predictive maintenance: overview and challenges. Intell Syst Appl. 2023;19:200251. Available from: https://doi.org/10.1016/j.iswa.2023.200251
5. Deloitte. From static to dynamic AI governance [Internet]. Available from: https://www.deloitte.com/us/en/insights/industry/government-public-sector-services/static-to-dynamic-ai-governance.html
6. Morales J, Antunes L, Earl P, Edman R, Hamed J, Reynolds D, et al. Insights on implementing a metrics baseline for post-deployment AI container monitoring. In: Proceedings of the 2024 International Conference on Software and Systems Processes (ICSSP '24). New York (NY): Association for Computing Machinery; 2024. p. 46–55. Available from: https://dl.acm.org/doi/10.1145/3666015.3666018
7. Diligent Corporation. AI governance: what it is and why it matters [Internet]. Available from: https://www.diligent.com/resources/blog/ai-governance
8. Papagiannidis E, Mikalef P, Conboy K. Responsible artificial intelligence governance: a review and research framework. J Strateg Inf Syst. 2025;34(2):101885. Available from: https://dl.acm.org/doi/abs/10.1016/j.jsis.2024.101885
9. IBM. What is human-in-the-loop? [Internet]. Available from: https://www.ibm.com/think/topics/human-in-the-loop
10. European Union. Regulation (EU) 2024/1689 [Internet]. Available from: https://eur-lex.europa.eu/eli/reg/2024/1689/oj

Citation: Hemachandran K. Why One-Time Approval is No Longer Enough for AI Systems. Trends Comput Sci Inf Technol. 2026;11(1):46-48. Available from: https://dx.doi.org/10.17352/tcsit.000109