

Compliance-Aware Conversational AI for Regulated Enterprise Decision Support

Subba Nelakudhiti

AI/ML Consultant, U.S.A.
subbanelakudhiti@outlook.com

Abstract:

Conversational AI systems are now used in various settings to improve decision-making in different fields like business, healthcare, finance, human resources, customer services, etc. Though conversational AI systems hold many benefits for improving productivity, their integration into an organizational environment poses many challenges regarding how to maintain organizational Compliance. Ideal conversational AI systems are designed to maximize user experience or response accuracy; however, in many systems, regulatory requirements are often neglected during conventional design.

This paper describes Compliance-Aware Conversational AI, a new Conversational AI system particularly designed for decision support in an enterprise environment. This system is equipped with a new decision support feature whereby decision output is checked for Compliance automatically, unlike other decision support systems in modern AI. This system features a hierarchical architecture in which high-level decision support is isolated from low-level language understanding and response generation. A new feature in Compliance-Aware Conversational AI is reward planning in decision support in enterprise decision support systems.

It is evaluating the proposed approach with enterprise scenarios basing on the simulation: financial advisory support, human resources policy guidance, and data access decision assistance. Experimental results show that compliance-aware conversational agents strongly reduce policy violations while maintaining a high response relevance and user satisfaction. For the visualization of the accuracy of responses, adherence to the Compliance, and interaction latency, it shows a favorable trade-off compared to an unconstrained model for conversational interactions. These results also point towards the path of integrating compliance awareness into conversational AI as a necessary next step for its safe and scalable deployment in regulated enterprise environments.

Keywords: Compliance Aware Conversational AI, Artificial Intelligence (AI)

Introduction

Converse AI interfaces have been established as an undoubtedly important interface between humans and enterprise information systems. In addition, it enables accessibility, reduces operational workload, and promotes fast decision-making in enterprises. Many enterprises are using conversational AI interfaces as they can aid in various functions, including the interpretation of policies, acting as customer support, providing financial suggestions, and retrieving knowledge.

However, there are many enterprise domains with very tight regulatory and governance requirements. In the financial domain, there are auditing rules and disclosure regulations that must be obeyed. The healthcare domain is protecting the patient information, whereas the human resource domain prioritizes fairness and privacy policies. Conventional conversational AI systems, where training objectives are often aligned with response relevance or user satisfaction, do not inherently understand policy constraint interactions [1]. They can, therefore, easily produce violating responses, disclose confidential information, or even provide undue advice.

Received: 17-10-2025

Revised: 21-11-2025

Accepted: 29-11-2025

Published: 26-12-2025

Citation: “ Compliance-Aware Conversational AI for Regulated Enterprise Decision Support”, *ijaicn*, vol. 1, no. 1, pp. 32–35, Feb. 2025,

Copyright: © 2025 by the authors. Submitted for possible open access publication under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

For example, the existing risk reduction mechanisms may entail rule-based filters or human-in-the-loop mechanisms [2]. However, even though these mechanisms perform their intended role to an extent, they also lead to a lack of response scalability. There is an increasing requirement for integrating reasoning around Compliance into the decision-making process of the conversational AI system itself.

In this paper, a framework will be introduced on how a compliant discussion-aware artificial intelligence approach can aid regulated enterprise decision support systems. In this way, not only can information be communicated by a conversational agent, but also a balance can be achieved with regard to informativeness, safety, and regulatory oversight, thanks to a hierarchical approach where a discussion policy can be separated from a response policy to increase interpretability, adaptability, and trustworthiness of a regulated agent responding to requests posed by decision systems within an enterprise.

Simulation Report

Simulation Environment

The environment that is simulated is a regulated enterprise, where a conversational AI agent is engaged by users who require some response related to decisions. The response of the AI is assessed by how it is related, how it complies, and how it violates certain policies [3] [4]. The Compliance is done by means of restrictions that are imposed by enterprise policies, such as data access, approval, disclosure, etc.

Learning: Configuration

Three spoken systems are presented:

1. Standard Conversational AI focused on response relevance,
2. Filtered Conversational AI with post-response compliance rules, and
3. Compliance-Aware Conversational AI (Proposed), integrating Compliance into policy learning.

The compliance-aware utilizes hierarchical control. Here, one controls all the response approaches that can comply. Then one uses a low-level language. The Evaluation Metrics also include the relevancy score, policy violation rate, average response latency, and the escalating frequency to a human operator.

Real-Time Scenarios

Approaches	Response Relevance (%)	Compliance Violation (%)	Escalation to Human (%)
Standard Conversational AI	93	19	7
Filtered Conversational AI	90	8	12
Compliance Aware Conversational AI	92	4	5

Table 1: The financial advisory and the decision support for performance.




Figure 1: The financial advisory and the decision support for performance.

Financial Advisory Decision Support:

The figures indicate a significant reduction in regulatory violations by the compliance-aware design of the conversational AI without compromising the

response's relevance. Although a higher response's relevance is achieved by a regular design of the conversational AI, there is a significantly higher violation rate compared to the proposed approach. Filtered strategies result in a higher escalation of human interventions, whereas a better regulatory environment is achieved by the suggested design with minimum human intervention [5].

Human Resource Policy with Guidance Performance

Approaches	Response Relevance (%)	Compliance Violation (%)	Escalation to Human (%)
Standard Conversational AI	95	16	430
Filtered Conversational AI	94	7	520
Compliance Aware Conversational AI	96	3	460

Table 2: Illustration of response to performance through human resource policies.

From Table 2, it can be observed that the highest policy accuracy, along with the least inconsistent responses, is obtained by employing a compliance-aware conversational AI. Although filter-based models achieve this, these models face the limitation of a delayed response. Hence, the proposed approach can be considered an effective trade-off.

Approaches	Correct Access Decision (%)	Privacy Violation (%)	Audit Intervention (%)
Standard Conversational AI	91	18	High
Filtered Conversational AI	94	9	Medium
Compliance Aware Conversational AI	97	4	Low

Table 3: Data Access and the Privacy Support Decision




Figure 3: Data Access and the privacy support decision.

Data Access and Privacy Decision Support:

These results show that compliance-aware conversational AI greatly improves correct access decisions while minimally violating privacy [6]. Standard conversational AI tends to provide the maximum violation rates, thus forcing frequent audit interventions. In contrast, the proposed model is able to achieve strong access accuracy with very low audit overhead, hence proving its effectiveness in privacy-sensitive enterprise decision support scenarios.

Challenges and Solutions

Challenge 1: Balancing Quality of Decisions and Regulatory Compliance; However, such systems in an enterprise setting might have high responses in terms of user satisfaction in relation to response appropriateness, although this can inadvertently facilitate responses that are technically in violation of certain rules or set guidelines [7]. Alternatively, excessive regulation in such systems can have adverse effects in relation to response usefulness, thereby lowering user trust in such systems. This is achieved by incorporating compliant reward shaping in relation to developing an architecture of hierarchy in coordinating responses in a conversation system. Such high-level guidelines regulate system responses in line with set regulations while offering high-response information in such systems.

Challenge 2: Ambiguous Queries & Delayed Compliance Feedback

In many cases, there are a wide variety of ambiguous requests with underlying multiple intents with regard to how enterprise users phrase them, making it very hard to appropriately deduce a compliant request by a conversation-based artificial intelligence system [8]. Moreover, there are challenges since many non-compliance items can only be detected through audit data and sparse feedback, particularly after interaction completion. In contrast, there is a need to address real-time interactions with regard to examining contextual information with responses generated in real time by the system through the use of early warning items and contextual classification of intent.

Conclusion

In this regard, the present work attempted to present a new concept of a compliance-aware conversational AI framework for the specific context of regulated enterprise decisions support systems, which would, in the resulting design, successfully deal with the aforementioned challenges, particularly with regard to satisfying regulatory requirements, response quality, and response speed/efficiency. To this rather critical objective, the presented design, as also shown through the subsequent simulation outcomes for diverse real-time enterprise scenarios, achieves a considerable reduction in regulatory violations, satisfying the response level of relevancy, and minimizing the human response requirements.

References

- [1] Yan, R. (2018, July). "Chitty-Chitty-Chat Bot": Deep Learning for Conversational AI. In *IJCAI* (Vol. 18, pp. 5520-5526). <https://www.ijcai.org/proceedings/2018/0778.pdf>
- [2] Xin, D., Ma, L., Liu, J., Macke, S., Song, S., & Parameswaran, A. (2018, June). Accelerating human-in-the-loop machine learning: Challenges and opportunities. In *Proceedings of the second workshop on data management for end-to-end machine learning* (pp. 1-4). <https://dl.acm.org/doi/abs/10.1145/3209889.3209897>
- [3] Graesser, A. C., Cai, Z., Morgan, B., & Wang, L. (2017). Assessment with computer agents that engage in conversational dialogues and dialogues with learners. *Computers in Human Behavior*, 76, 607-616. <https://www.sciencedirect.com/science/article/pii/S074756321730198X>
- [4] Greengard, S. (2016). Cybersecurity gets smart. *Communications of the ACM*, 59(5), 29-31. a. <https://dl.acm.org/doi/abs/10.1145/2898969>
- [5] Ibitoye, J. S. (2018). Securing smart grid and critical infrastructure through AI-enhanced cloud networking. *International Journal of Computer Applications Technology and Research*, 7(12), 517-529. a. https://www.researchgate.net/profile/Joshua-Ibitoye-2/publication/396161458_Securing_Smart_Grid_and_Critical_Infrastructure_through_AI-Enhanced_Cloud_Networking/links/68dfe137d221a404b2a53139/Securing-Smart-Grid-and-Critical-Infrastructure-through-AI-Enhanced-Cloud-Networking.pdf
- [6] van den Broek, T., & van Veenstra, A. F. (2018). Governance of big data collaborations: How to balance regulatory Compliance and disruptive innovation. *Technological Forecasting and Social Change*, 129, 330-338. <https://www.sciencedirect.com/science/article/pii/S0040162517314695>
- [7] Kingston, J. (2017). Using artificial intelligence to support Compliance with the General Data Protection Regulation. *Artificial Intelligence and Law*, 25(4), 429-443. <https://link.springer.com/article/10.1007/s10506-017-9206-9>
- [8] Sannidhanam, A. H. (2026). LLM-Driven Workflow Optimization: Intelligent Routing in Asynchronous Distributed Systems. *International Journal of AI and Machine Learning*, 1(2), 23-33.
- [9] Distributed Data Processing Frameworks for Large-Scale Healthcare Analytics. (2019). *International Journal of Advance Industrial Engineering*, 7(04), 1-10. doi:10.14741/ijaie/v.7.4.01
- [10] Anjani Haritha Sannidhanam. (2023). Zero-Downtime AI Model Updates in Real-Time Inference Systems. *International Journal of Engineering Science & Humanities*, 13(2), 70-78.
- [11] Sannidhanam, A. H. (2021). Real-Time Claims Processing Using Event-Driven Architectures. *International Journal of Technology, Management and Humanities*, 7(01), 36-50. doi:10.21590/07.01.02