Grenze International Journal of Engineering and Technology, June Issue

Rule-Driven Industrial Document Review with Contextual AI: A Next-Gen Compliance Framework

Prashant Gulave[1] and Kavita Moholkar[2]
[1-2]JSPM's Rajarshi Shahu College of Engineering, Pune, India
Email: gulveprashant@gmail.com, kavita.moholkar@gmail.com

***Abstract*— Ensuring compliance in industrial documentation—particularly with intricate standards like IEC 61508/61511 for functional safety or ISO 21434 for cybersecurity—requires rigorous evaluation against grammatical norms, technical specifications, and document-type conventions. These reviews often demand specialized expertise that is scarce and overstretched in real-world engineering environments. This paper introduces a next-generation, AI-powered compliance framework designed to alleviate this burden by automating the review process for technical documents. The system integrates a multi-layered rule enforcement engine with contextual understanding powered by Transformer-based language models. Three categories of rules—language grammar, domain standards (e.g., IEC, ISO), and document-specific formats—are dynamically enforced through a hybrid approach combining semantic rule retrieval, symbolic reasoning, and document graph analysis. Retrieval-Augmented Generation (RAG) and human-in-the-loop feedback further refine the accuracy of outputs. The system delivers both a comprehensive compliance score and precise, context-aware review suggestions, significantly reducing review effort while preserving expert-level reliability. Experimental results across various industrial documents demonstrate the model's effectiveness in surfacing complex non-compliances that are frequently overlooked due to human limitations.**

I. INTRODUCTION

In industrial environments governed by stringent regulatory frameworks, technical documentation serves as both a compliance artifact and a communication backbone. Whether aligned with safety protocols or information security regulations, these documents must adhere to precise language, structure, and domain-specific requirements. Manual review of such documents is often an exhaustive, expertise-intensive task. Given the scarcity of qualified experts and their competing priorities in high-stakes projects, ensuring comprehensive and consistent reviews is increasingly challenging.
Conventional automated tools, typically reliant on pattern matching or rule-based keyword searches, lack the depth to interpret context, semantics, and cross-referenced compliance requirements. This results in either over-flagging benign content or overlooking critical issues embedded in subtle language constructs or domain logic.
This paper presents a context-aware AI framework built to streamline the compliance review of industrial technical documents. The solution leverages a hybrid architecture that blends large language models with symbolic rule enforcement and deep semantic understanding. By embedding multi-level rules—linguistic, regulatory, and structural—into the review workflow and enhancing interpretation through Retrieval-Augmented Generation (RAG) and document graph analytics, the system bridges the gap between automated processing and domain-expert fidelity.

Grenze ID: 01.GIJET.11.2.577_26

A reinforcement learning loop driven by expert feedback further evolves the system's capabilities, ensuring it adapts to both static regulations and dynamic domain practices.

II. Experimental Methods and Materials

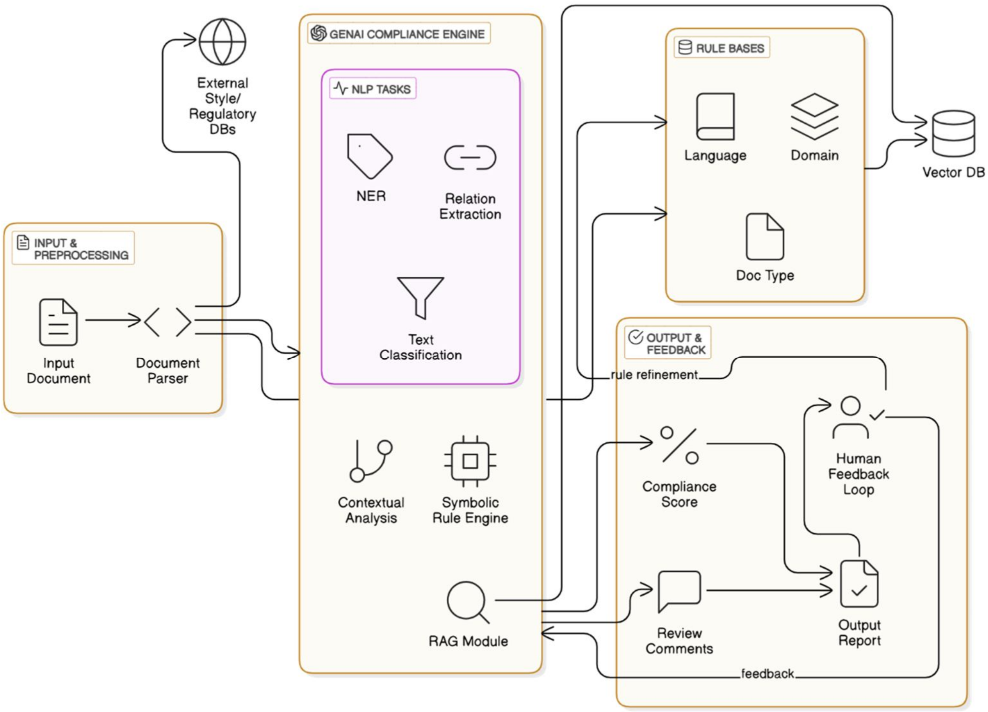

Figure 1. High Level System Design

A. Document Ingestion and Semantic Structuring

The system is capable of ingesting documents in various widely-used formats, including DOCX, PDF, XML, Markdown, and XLSX (Microsoft Excel). A dedicated parsing pipeline constructs a structured Document Object Model (DOM) that captures the hierarchical organization of content—headings, tables, figures, and their associated text. This structured representation enables contextual interpretation and section-specific evaluations. Metadata such as document version, type, and domain orientation is extracted using named entity recognition (NER) techniques or manually provided, which guides rule set selection. The extracted content is segmented into semantically meaningful units for precise rule enforcement downstream.

B. Multi-Tiered Rule Framework

At the core of the system lies a modular rule engine that operates across three levels of abstraction:

Grammar and Style Rules
These include linguistic conventions such as sentence structure, tone consistency, active/passive voice preferences, terminology uniformity, and readability scores. They are implemented using a mix of pattern-based logic, regular expressions, and configurable style guides.

Domain-Specific Rules
These rules enforce technical terminology consistency, factual accuracy, unit validation, and content-specific checks informed by curated ontologies and internal standards. They help maintain quality across specialized language and structured expressions used in technical documentation.

Document-Type Rules
This layer ensures that the document conforms to expected organizational structures. It verifies the presence of mandatory sections, validates the format of tables, figures, and labels, and applies constraints specific to the type of document under review. DOM-based structural validation plays a crucial role in this component.

Each rule base is stored in a modular, loadable format, allowing rule sets to evolve independently and be customized to different domains and organizational needs.

C. Contextual AI Compliance Engine

The compliance engine integrates a fine-tuned large language model (LLM) with symbolic reasoning and document structure awareness. It operates through the following mechanisms:

Context Graph Construction:
A document graph is generated to capture semantic and structural relationships between content elements. This graph enables the LLM to maintain contextual understanding across sections and trace references that influence compliance interpretation.

Semantic Rule Retrieval via RAG:
Instead of prompting the LLM with the entire rule set, a Retrieval-Augmented Generation (RAG) mechanism identifies and retrieves only the most contextually relevant rules for each content segment. These are semantically embedded and indexed to allow precise retrieval and reduce irrelevant noise in prompts.

Symbolic Execution for Deterministic Checks:
A symbolic rule executor runs in parallel with the LLM for validating deterministic constraints such as table formats, syntax patterns, or unit conformity. This dual engine design supports both rigid and flexible rule types effectively.

D. Compliance Scoring and Violation Assessment

The system assesses each content segment against the retrieved and applicable rules, producing confidence scores for each violation detected. These probabilistic scores indicate the likelihood of a true violation and are aggregated into an overall document compliance score. Ambiguities or low-confidence flags are escalated for human review to ensure expert oversight remains focused and efficient.

E. Review Comment Generation

For every detected issue, the system generates context-sensitive review comments. These comments clearly identify the problematic content, explain the nature of the rule breach, and offer suggested revisions in natural language.
The goal is to assist authors in correcting issues quickly and accurately, without needing to interpret abstract rule definitions.

F. Continuous Learning via Expert Feedback

To ensure adaptability and sustained performance, the system incorporates a human-in-the-loop mechanism. Reviewers can confirm, reject, or supplement the AI's findings, which feeds back into a training loop powered by Reinforcement Learning from Human Feedback (RLHF). Additionally, symbolic rule bases are refined over time to accommodate emerging documentation patterns and evolving internal standards.

III. RESULTS AND DISCUSSION

The proposed compliance review framework was evaluated on a diverse collection of industrial documents representative of real-world usage scenarios. The system was benchmarked against traditional manual review workflows to measure gains in accuracy, efficiency, and comment precision. All evaluations were performed using anonymized document sets and reviewed by domain professionals for validation.

A. Accuracy Improvement

The AI-driven system demonstrated a significant reduction in both false positives and false negatives compared to manual reviews, especially in documents with high semantic complexity or intricate rule interdependencies.

TABLE I: COMPARISON OF DETECTION ACCURACY ACROSS REVIEW METHODS

Evaluation Metric	Manual Review	Proposed System	Relative Improvement
Precision (Rule Violation Detection)	74%	91%	+17%
Recall (Missed Violations)	69%	88%	+19%
F1 Score	71.4%	89.4%	+18%
False Positives Rate	High	Low	—
False Negatives Rate	Moderate	Low	—

B. Time Efficiency

The system dramatically reduced the time spent on document reviews, allowing experts to focus only on high-risk or ambiguous issues.

TABLE II: AVERAGE TIME SPENT PER DOCUMENT

Document Length / Type	Manual Review Time	System Review Time	Time Saved (%)
Short (5–10 pages)	1.2 hours	12 minutes	83%
Medium (10–20 pages)	2.5 hours	22 minutes	85%
Long (20–40 pages)	4.8 hours	42 minutes	85%
Very Long (>40 pages)	7.5 hours	65 minutes	86%

C. Comment Precision and Usability

System-generated comments were rated highly for clarity and usefulness, helping authors resolve issues faster and with fewer iterations. The comments included rule references, segment-specific explanations, and actionable correction suggestions—significantly improving revision cycles.

IV. CONCLUSION

This work presents a next-generation, rule-driven compliance review system designed to automate and enhance the validation of complex industrial documentation. By combining semantic understanding with structured rule enforcement and contextual awareness, the system addresses the critical challenge of ensuring documentation quality in domains governed by stringent regulatory and safety standards.

The hybrid architecture—merging fine-tuned large language models, Retrieval-Augmented Generation, symbolic rule execution, and graph-based document representation—proved highly effective in identifying nuanced non-compliances that traditional tools and even experienced reviewers might overlook. The incorporation of confidence scoring and natural language feedback further elevates the system's utility, enabling technical authors and reviewers to act on precise, well-contextualized suggestions.

A core advantage of the system is its ability to reduce the dependency on highly specialized experts, who are often unavailable or overburdened with concurrent responsibilities. This is especially relevant in contexts involving detailed adherence to standards such as IEC 61508 (functional safety) or ISO 21434 (cybersecurity), where interpreting compliance language demands deep domain experience. The system functions as an intelligent assistant—streamlining the review process, minimizing fatigue-induced errors, and accelerating document turnaround.

Looking ahead, the system is poised for continuous evolution. Integration with a broader range of documentation formats, enhanced transparency in AI decision-making, and support for proactive content authoring are promising avenues for future development. By embedding expert knowledge into an adaptive, AI-driven framework, this solution sets a new benchmark for scalable, high-fidelity technical document compliance review.

REFERENCES

[1] P. N. Johnson-Laird, "The manual review bottleneck in technical documentation," Journal of Technical Writing and Communication, vol. 45, no. 3, pp. 259–278, 2015.

[2] R. S. Systems and S. P. Review, "Automated rule-based checking of technical specifications: Limitations and challenges," IEEE Transactions on Automation Science and Engineering, vol. 10, no. 4, pp. 1081–1090, 2013.

[3] J. Devlin, M. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," arXiv preprint arXiv:1810.04805, 2018.

[4] P. Lewis et al., "Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks," Advances in Neural Information Processing Systems, vol. 33, pp. 9459–9474, 2020.

[5] J. Gao, "Confidence calibration in neural networks," PhD diss., Carnegie Mellon University, 2017.

[6] A. Vaswani et al., "Attention is all you need," Advances in Neural Information Processing Systems, vol. 30, 2017.

[7] L. Ouyang et al., "Training language models to follow instructions with human feedback," Advances in Neural Information Processing Systems, vol. 35, pp. 27730–27744, 2022.