ComplianceGpt: An LLM-Assisted Regulatory Compliance Verifier

Priscilla Kyei Danso, Stony Brook University, New York, USA

I. Introduction

The widespread concern over mass data collection by social networks, corporations, and data brokers has prompted the establishment of regulatory frameworks designed to safeguard personal data and enhance privacy protections. Frameworks like the Health Insurance Portability and Accountability Act (HIPAA) of 1996, the European Union's General Data Protection Regulation (GDPR), the California Consumer Privacy Act (CCPA), etc. However, the complexity, jurisdictional variations, and frequent updates make them difficult to adhere to. Large Language Models (LLMs) offer potential solutions by helping interpret and navigate complex legal texts, but they have significant limitations. They are, however, prone to inaccuracies and "hallucinations," producing unreliable or misleading responses that risk non-compliance. In contrast, formal verification approaches can effectively assess compliance by translating regulations into logical formulas using fragments of First-Order Logic (FOL) or First-Order Temporal Logic (FOTL). These also have the limitation of accessibility; they are not user-friendly for general audiences who must engage with formal logic languages to interact with these systems.

To bridge this gap, we introduce ComplianceGPT, a hybrid system that combines specialized LLMs with a dedicated logic-based compliance checker, called précis. ComplianceGPT aims to provide an intuitive, user-friendly interface for navigating regulatory compliance while leveraging the rigour and reliability of formal verification. The system will process natural language queries (e.g., "Can a doctor send my medical record to a third party under HIPAA?") and translate them into FOL formulas (e.g., $disclose(p_1, p_2, q, \text{medical\_records}) \wedge \mathsf{inrole}(p_1, \text{doctor}) \wedge \mathsf{inrole}(p_2, \text{patient}) \wedge p_2 = q$, in which p_1 denotes the information sender, p_2 denotes the information receiver, and q denotes the subject whose PII is being released by p_1 to p_2) using a specialized vocabulary derived from the regulation's representation. Figure 1 depicts our system architecture.

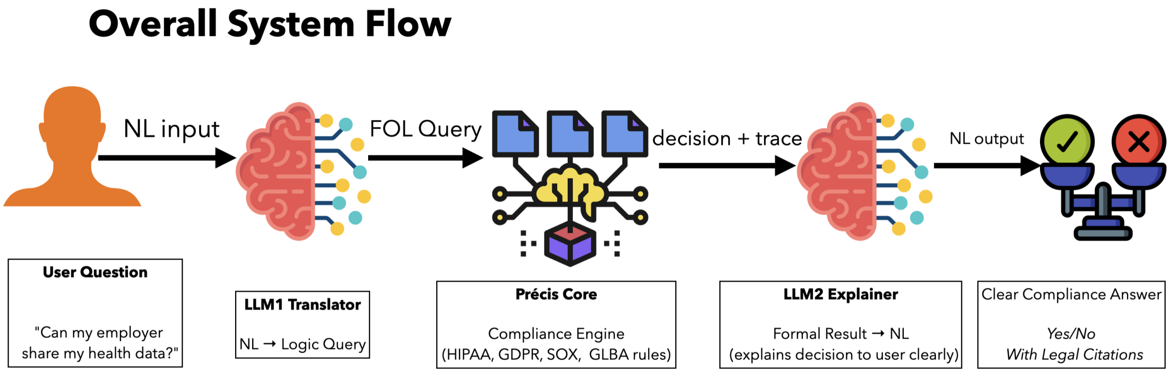
Overall System Flow
NL input
FOL Query
decision + trace
NL output
User Question
"Can my employer share my health data?"
LLM1 Translator
NL → Logic Query
Précis Core
Compliance Engine (HIPAA, GDPR, SOX, GLBA rules)
LLM2 Explainer
Formal Result → NL (explains decision to user clearly)
Clear Compliance Answer
Yes/No With Legal Citations

Fig. 1. System Architecture

The system architecture of ComplianceGPT is a pipeline-based approach separating translation, reasoning, and explanation. This is a complex and logic-driven system. It uses two separate LLMs to translate a natural language (NL) query into formal logic and then translate the formal result back to natural language. The core of this system is the précis, a compliance engine that executes the logic query against a database of rules. The user's natural language query is processed by LLM1, which translates the query into a FOTL representation understood by précis. The compliance engine receives the FOTL query and runs it against a set of predefined rules to generate a formal decision and a traceability of the reasoning. This traceability information is subsequently processed by LLM2 and translated into a clear, natural language explanation for the user, including legal citations. This output would be a simple "Yes/No" answer with legal citations, based on the compliance decision.

I am currently: 1) designing a structured intermediary representations that bridge natural language and formal logic, enabling better compositional reasoning and error detection; 2) developing systems where specialized agents (fact extraction, formula synthesis, verification, explanation) collaborate to solve complex compliance queries; 3) building systems that engage in dialogue with users to clarify ambiguous requirements before generating formal specifications.

My long-term goal is to develop an automated multi-regulatory agentic system that performs the above-mentioned. I envision a future where every AI system comes with formal guarantees. In conclusion, my infinitesimal contribution to the overarching goal is to make AI systems that are not just powerful (since they presently are), but provably safe, compliant, and aligned with human values.