

Neuro-Symbolic Compliance: Integrating LLMs and SMT Solvers for Automated Financial Legal Analysis

Yung-Shen Hsia*, Fang Yu*, and Jie-Hong Roland Jiang†

*Department of Management Information Systems, National ChengChi University, Taipei, Taiwan
Emails: 113356046@g.nccu.edu.tw, yuf@nccu.edu.tw

†Department of Electrical Engineering, National Taiwan University, Taipei, Taiwan
Email: jhjiang@ntu.edu.tw

Abstract—Financial regulations are increasingly complex, hindering automated compliance—especially the maintenance of logical consistency with minimal human oversight. We introduce a Neuro-Symbolic Compliance Framework that integrates Large Language Models (LLMs) with Satisfiability Modulo Theories (SMT) solvers to enable *formal verifiability* and *optimization-based compliance correction*. The LLM interprets statutes and enforcement cases to generate SMT constraints, while the solver enforces consistency and computes the *minimal factual modification* required to restore legality when penalties arise. Unlike transparency-oriented methods, our approach emphasizes logic-driven optimization, delivering verifiable, legally consistent reasoning rather than post-hoc explanation. Evaluated on 87 enforcement cases from Taiwan’s Financial Supervisory Commission (FSC), the system attains 86.2% correctness in SMT code generation, improves reasoning efficiency by over 100×, and consistently corrects violations—establishing a preliminary foundation for optimization-based compliance applications.

Index Terms—Financial precedent analysis, Retrieval-Augmented Generation (RAG), LLM Agent, SMT Solver, Neuro-Symbolic, Legal tech, Legal understanding analysis, Constraint solving, Financial Supervisory Commission (FSC), Optimization, Legal consistency, MaxSMT

I. INTRODUCTION

As legal regulations grow in complexity and change more frequently, automating compliance has become a key challenge in Legal Tech [1], [2]. A central difficulty lies in translating legal provisions into machine-readable formats while preserving their semantic richness [3]. Meanwhile, AI has accelerated transformation in finance, supporting tasks from anomaly detection to regulatory monitoring [4]–[6].

Formal methods like SAT/SMT solving and model checking [7], [8] have recently been applied to legal reasoning by encoding laws and facts as constraints [9]. Large Language Models (LLMs) perform well in interpreting legal text [10], [11] but often lack precision for high-stakes tasks. This motivates hybrid neuro-symbolic approaches that combine the interpretability of symbolic reasoning with the language capabilities of LLMs [12], [13].

Neuro-symbolic AI enables compositional generalization and transparency by integrating logic-based reasoning with

deep learning. It has shown success in areas such as visual reasoning, robotic manipulation, and concept grounding [13]. These capabilities are especially relevant in compliance domains, where decisions must be both explainable and legally defensible.

Recent efforts in combining LLMs with formal reasoning systems have underscored the growing demand for verifiable AI reasoning. For instance, the LLM-Based Formalized Programming (LLMFP) framework [14] demonstrates that translating natural language problems into Satisfiability Modulo Theories (SMT) formulations enables completeness and soundness guarantees that purely neural approaches cannot ensure. Similarly, Chen et al. [15] show that even advanced prompting strategies, such as Dual Chains of Logic (DCoL), still fall short of achieving consistent logical inference, as LLMs inherently lack self-verification mechanisms. These findings suggest that while LLMs excel in flexible natural language understanding, symbolic solvers provide the rigor required for provable correctness.

Building on this insight, Our work integrates LLMs with SMT-based formal verification to bridge interpretability and verifiability—aiming to enhance both transparency and logical consistency in regulatory compliance tasks.

Complementing this trend, the study on *Formal Verification of LLM-based Agents* [16] further emphasizes that LLM reasoning processes should be formally checkable, and proposes strategies such as *Multiple LLMs Debating* and *Self-Correction* to enhance consistency and reliability. This collective direction highlights a broader consensus that coupling LLMs with formal methods is essential for developing trustworthy and verifiable AI systems.

Yet, few frameworks target the unique characteristics of financial legal enforcement in specific jurisdictions. Taiwan’s regulatory ecosystem, for example, involves multilingual documentation, localized statutes, and detailed enforcement case records. These pose challenges that are not well addressed by existing compliance tools.

In this work, we present a compliance framework that combines LLMs with SMT-based verification, tailored to

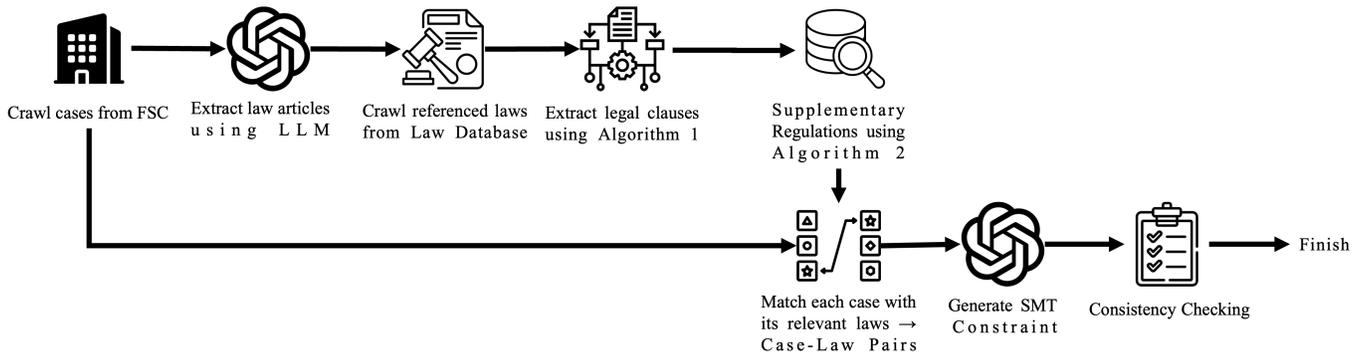


Fig. 1: SMT Constraint Generation Pipeline

Taiwan’s financial legal context. A RAG pipeline integrates statutes, enforcement cases, and violation records from the FSC. Legal norms are encoded as constraints, and generative models aid interpretation to support violation prediction and legal guidance.

We evaluate the system on 87 real-world FSC enforcement cases, demonstrating its ability to detect unlawful contract terms and generate rule-consistent recommendations. The SMT layer enhances interpretability and accuracy, offering a potential foundation for more transparent and verifiable Legal Tech applications.

II. RELATED WORK

Numerous studies have explored the automation of regulatory compliance through various computational techniques. LLMs have gained attention in this space. In financial auditing, Berger et al. [17] evaluated LLMs for automated compliance checks, comparing models such as LLaMA-2 and GPT-4. GPT-4 excelled in multilingual tasks, while LLaMA-2 was more cost-effective for detecting non-compliance. However, challenges like variability and false positives persist. Similarly, Liao et al. [18] propose a Neuro-Symbolic Reasoning framework, where LLMs generate logical forms that are verified by symbolic solvers. Ning et al. [19] expand on this by integrating multi-agent planning and external symbolic tools for enhanced reasoning and explainability. Shi et al. [20] also present a dual-system neuro-symbolic architecture, combining heuristic LLM inference with logic-based verification to improve accuracy and interpretability. These efforts show growing convergence between language models and formal methods for regulatory tasks.

SMT solvers and symbolic approaches offer another line of research. Judson et al. [21] introduced *soid*, which integrates Z3 with symbolic execution via KLEE, providing legal practitioners with GUI-supported formal accountability tools. Feng et al. [22] translated legal constraints into SMT using Metric First-Order Temporal Logic (MFOTL), enabling efficient compliance checks in healthcare and finance. Similarly, Tsigkanos et al. [23] applied SMT to model economic networks, ensuring consistency through logical constraints.

Hybrid LLM–SMT systems aim to bridge natural and formal language. The Semantic Self-Verification (SSV) framework [24] combines LLM-generated constraints with SMT solver feedback, improving accuracy on formal reasoning tasks. This integration minimizes manual effort while enhancing correctness. Other complementary methods support this landscape. Beach et al. [25] proposed a semantic framework that decouples legal expertise from implementation, empowering non-technical users in compliance modeling. Separately, counterfactual explanations [26], [27] have been studied for making algorithmic decisions transparent and contestable under data protection laws such as GDPR.

Our work builds upon these trends by applying LLM–SMT integration to regulatory compliance. Rather than focusing solely on translation or dialogue modeling, we construct an end-to-end pipeline that supports constraint synthesis, verification, and interactive correction, offering interpretable and verifiable legal analysis.

III. METHODOLOGY

A. SMT Constraint Generation

We first discuss our SMT constraint generation pipeline. Figure 1 illustrates the overall process of transforming unstructured legal and case documents into formalized SMT constraints. The pipeline begins with data collection and preprocessing, where legal articles and case texts are parsed and structured. Subsequently, key legal elements and logical relationships are extracted to construct high-quality constraints suitable for automated reasoning.

1) *Case Crawling and Preliminary Article Extraction:* We collected financial regulatory case summaries from public repositories, each containing structured metadata and a narrative describing the violation. We use an LLM to extract the legal provisions explicitly mentioned in the case description.

2) *Legal Article Retrieval and Parsing:* For each cited provision, the corresponding legal text was retrieved from an official law database and parsed using a rule-based algorithm (Algorithm 1). The extracted articles were structured into dictionaries mapping article numbers to their titles, clauses, and full content.

Each case was then mapped to its cited articles, forming a structured case-law association dataset suitable for semantic retrieval and reasoning.

Algorithm 1 Article and Clause Extraction from Raw Legal Documents

Input: Raw document D

Output: Dictionary A mapping article numbers to titles, clauses, and content

```

1: Initialize empty text  $T \leftarrow ""$ 
2: Open  $D$  using pdfplumber; extract and concatenate text into  $T$ 
3: Split  $T$  into lines  $L$ 
4: Initialize empty dictionary  $A$ ; set  $a \leftarrow \text{None}$ ,  $c \leftarrow \text{None}$ 
5: for each line  $l$  in  $L$  do
6:   Normalize  $l$  (strip and collapse whitespace)
7:   if  $l$  matches section heading pattern then
8:     continue
9:   end if
10:  if  $l$  matches "Article  $n$ " then
11:     $a \leftarrow n$ ; initialize  $A[a]$  with title, empty clauses, and content
12:     $c \leftarrow \text{None}$ 
13:    else if  $a \neq \text{None}$  then
14:      if  $l$  matches clause pattern (e.g., " $m$ . text") then
15:        Append clause text to  $A[a][\text{clauses}]$ ; set  $c \leftarrow \text{text}$ 
16:      else if  $c \neq \text{None}$  then
17:        Append  $l$  to last clause in  $A[a][\text{clauses}]$ 
18:      else
19:        Append  $l$  to  $A[a][\text{content}]$ 
20:      end if
21:    end if
22: end for return  $A$ 

```

3) *Supplementary Article Retrieval (RAG-Based Enrichment)*: To ensure comprehensive coverage of legal context, a Retrieval-Augmented Generation (RAG) module expands each base article L_0 by retrieving semantically related provisions and enforcement rules from a legal database \mathcal{D} . This procedure improves SMT constraint completeness by including contextually supportive clauses.

The process, detailed in Algorithm 2, consists of four stages:

- 1) **Query Generation**: An LLM generates a diverse query set $\mathcal{Q} = \{q_1, \dots, q_m\}$ to capture different semantic aspects of L_0 .
- 2) **Hybrid Retrieval**: Each query q performs both BM25 and vector searches:

$$R_{\text{bm25}} = \text{BM25Search}(\mathcal{D}, q, K),$$

$$R_{\text{vec}} = \text{VectorSearch}(\mathcal{D}, q, K).$$

Results are merged as $R = R_{\text{bm25}} \cup R_{\text{vec}}$, and scored by a hybrid relevance function:

$$s_{\text{hybrid}}(d, q) = \alpha \text{sim}_{\text{vec}}(d, q) + (1 - \alpha) \text{score}_{\text{bm25}}(d, q),$$

where $\alpha = 0.8$ balances semantic and lexical relevance.

- 3) **Reranking and Filtering**: The top- K candidates are reranked by a Cross-Encoder (FlagReranker) and filtered via LLM_FilterUseful to retain non-redundant, contextually relevant clauses.
- 4) **Aggregation**: Filtered clauses C from all queries are merged and deduplicated into a final supplementary set \mathcal{S} .

This retrieval module, implemented as the *LegalSearchEngine*, integrates semantic and lexical retrieval, hybrid scoring, reranking, and LLM filtering. It employs *text-embedding-ada-002* for embeddings, *ChromaDB* for vector indexing, and *BAAI/bge-reranker-v2-m3* for reranking.

Algorithm 2 Supplementary Article Expansion via Retrieval-Augmented Generation

Input: Base article L_0 ; database \mathcal{D} ; weight $\alpha \in [0, 1]$; retrieval size K

Output: Supplementary set \mathcal{S}

```

1:  $\mathcal{Q} \leftarrow \text{LLM\_GenQueries}(L_0)$ 
2:  $\mathcal{S} \leftarrow \emptyset$ 
3: for each  $q \in \mathcal{Q}$  do
4:    $R_{\text{bm25}} \leftarrow \text{BM25Search}(\mathcal{D}, q, K)$ 
5:    $R_{\text{vec}} \leftarrow \text{VectorSearch}(\mathcal{D}, q, K)$ 
6:    $R \leftarrow R_{\text{bm25}} \cup R_{\text{vec}}$ 
7:   for each  $d \in R$  do
8:     Compute hybrid score:
9:      $s_{\text{hybrid}}(d) = \alpha \cdot \text{sim}_{\text{vec}}(d, q) + (1 - \alpha) \cdot \text{score}_{\text{bm25}}(d, q)$ 
10:    end for
11:    $R_{\text{top}} \leftarrow \text{TopK}(R, s_{\text{hybrid}}, K)$ 
12:    $R' \leftarrow \text{Rerank}(R_{\text{top}}, \text{CrossEncoder}, q)$ 
13:    $C \leftarrow \text{LLM\_FilterUseful}(R', L_0)$ 
14:    $\mathcal{S} \leftarrow \mathcal{S} \cup C$ 
15: end for return  $\text{Unique}(\mathcal{S})$ 

```

4) *SMT Constraint Generation*: After obtaining the case-article pairs, the system proceeds to generate SMT (Satisfiability Modulo Theories) constraints that formally represent the relationship between the factual scenario and the corresponding legal provisions. Once the constraints are generated, a series of validation and consistency checks are performed to ensure their correctness and logical soundness.

a) *Build Law Constraints*: Legal reasoning requires that all actions conform to statutory norms. To formalize these rules, we use LLM interpretation to convert legal provisions into Boolean expressions representing their logical structure. Each statutory clause is expressed as a set of propositional formulas—typically involving conjunctions, disjunctions, and implications—to capture the conditional relationships defined by the law.

To illustrate the process, we present a real case from the Financial Supervisory Commission (FSC).¹ Using our extraction and RAG modules (Section III-A3), relevant provisions

¹https://www.fsc.gov.tw/en/home.jsp?id=131&parentpath=0,2&mcustomize=multimessages_view.jsp&dataserno=202407120004&dtale=Penalty

were retrieved from the *Insurance Act* (Articles 143-4 and 143-6)² and its *Enforcement Rules* (Articles 2–5).³ These articles define the capital adequacy classification and corresponding supervisory measures. The encoded SMT hard constraints are summarized below.

Let $r = \frac{\text{own_capital}}{\text{risk_capital}} \times 100$. Then the capital level classification is defined as:

$$\text{capital_level} = \begin{cases} 4, & r < 50 \text{ or } \text{net_worth} < 0, \\ 3, & 50 \leq r < 150, \\ 2, & 150 \leq r < 200, \\ 1, & r \geq 200. \end{cases}$$

Supervisory measures are expressed as:

$$\begin{aligned} \text{cL2_exec} &\Leftrightarrow (\text{plan_submitted} \wedge \text{plan_executed}), \\ \text{L3_exec} &\Leftrightarrow (\text{L2_exec} \wedge \text{person_removed} \wedge \text{asset_approved}). \end{aligned}$$

The penalty rule is defined compactly as:

$$\begin{aligned} \text{penalty} &= 1 \text{ if } (\text{level} = 4 \wedge \neg \text{L4_exec}) \\ &\quad \vee (\text{level} = 3 \wedge \neg \text{L3_exec}) \\ &\quad \vee (\text{level} = 2 \wedge \neg \text{L2_exec}), \\ \text{penalty} &= 0 \text{ otherwise.} \end{aligned}$$

The resulting set of formulas constitutes the law constraint base, which serves as the foundation for SMT encoding and subsequent logical verification.

b) Build Compliance Fact Constraints: In addition to statutory constraints, each regulatory case provides factual conditions that must be formalized into logical predicates for reasoning. We construct a set of *compliance fact constraints* $\mathcal{H}_{\text{case}}$ to represent the factual state of the case—such as capital adequacy ratios, remediation actions, or supervisory measures actually executed. These factual predicates are encoded in Boolean or numerical form and combined with the statutory constraint base \mathcal{H}_{law} to form a complete logical representation of the case.

For the FSC case described above, the extracted factual inputs are:

```
own_capital = 111.09,
risk_capital = 100.0,
net_worth = 2.97,
plan_submitted = true,
plan_executed = false
```

For variables not explicitly mentioned in the case description, the system treats them as *free variables* rather than assigning default Boolean values. This design allows the solver to explore feasible valuations dynamically during constraint satisfaction, ensuring that the reasoning process remains faithful to the available factual scope while avoiding unwarranted assumptions.

These factual assignments constitute the case-specific constraint set $\mathcal{H}_{\text{case}}$, which, when merged with the law constraint

base \mathcal{H}_{law} , produces the complete SMT problem instance:

$$\mathcal{H}_{\text{law}} \cup \mathcal{H}_{\text{case}}.$$

This unified representation allows the solver to verify whether the factual state of the case satisfies the statutory framework or violates one or more encoded provisions. It therefore provides a logically grounded test instance for subsequent consistency checking and compliance optimization.

B. SMT Constraint Consistency Check and Optimization

1) Law Constraint Consistency Check.: After encoding all statutory provisions as hard constraints, we perform a consistency validation to ensure the legal logic is self-coherent. The set of encoded legal formulas

$$\mathcal{H}_{\text{law}}$$

is submitted to the SMT solver. A satisfiable (*SAT*) result indicates that the statutory provisions are logically consistent and contain no internal contradictions. This helps ensure that the encoded law forms a logically consistent foundation. Conversely, an *UNSAT* result would reveal inconsistency within the legal encoding, prompting iterative refinement until the system achieves a consistent *SAT* result.

2) Case Illegality Check.: Once the statutory constraints are verified to be internally consistent, we integrate the case-specific facts and perform a second verification phase. This Case Illegality Check determines whether the factual conditions correctly trigger a statutory violation by evaluating the combined constraint set:

$$\mathcal{H}_{\text{law}} \cup \mathcal{H}_{\text{case}} \cup \{\text{Penalty} = \text{False}\}. \quad (1)$$

As shown in Eq. (1), the combined constraint set incorporates both statutory and case-specific hypotheses, with the penalty condition explicitly fixed to *False* to test for potential violations.

At this stage, an *UNSAT* result is expected, indicating that the factual configuration is incompatible with lawful conditions and thus represents a valid non-compliant scenario. The solver demonstrates that no assignment can satisfy both the statutory and factual constraints simultaneously, confirming the existence of a legal violation.

For the FSC case, the solver indeed returns *UNSAT*, revealing contradictions among:

```
insurance:capital_level,
meta:penalty_conditions,
insurance:level_3_measures_executed
```

These entries form the *unsat core*, pinpointing the minimal subset of constraints responsible for inconsistency. The conflict indicates violation of Articles 143-4 and 143-6 due to incomplete corrective measures, consistent with the actual administrative disposition.

If the solver instead returns *SAT*, the logical formulation fails to capture the intended violation, and the system initiates an iterative refinement process to regenerate or adjust the encoded constraints until the expected *UNSAT* outcome is

²<https://law.moj.gov.tw/ENG/LawClass/LawAll.aspx?pcode=G0390002>

³<https://law.moj.gov.tw/LawAll.aspx?pcode=G0390003>

achieved. The *Case Illegality Check* therefore validates that the encoded case indeed represents a legitimate breach under the statutory framework, but it does not attempt to restore satisfiability or modify any factual predicates at this stage.

After a valid *UNSAT* result is obtained, the system transitions to the optimization phase. In that subsequent step, Z3’s *Optimize()* module is applied to the same constraint set, where case-specific predicates are relaxed into soft constraints while statutory provisions restored.

3) *Optimization for Minimal Compliance Solution*: Following the verification and adjustment phases described above, the optimization stage formalizes the search for the minimal compliance solution. Here, we treat the previously identified *UNSAT* configurations as optimization targets and employ Z3’s *Optimize()* module to derive the smallest set of factual changes that restore full logical satisfiability.

After constructing both hard and soft constraints, we employ Z3’s *Optimize()* module to compute the *minimal compliance solution*. This procedure searches for a model that satisfies all statutory (hard) constraints while modifying as few factual (soft) constraints as possible, thereby achieving legal compliance with minimal intervention. All factual statements extracted from a regulatory case are encoded as soft constraints and appended to the solver through the `add_soft()` interface. Each soft constraint s_i is assigned a weight $w_i = 1$, unless regulatory priorities specify otherwise (e.g., penalties may receive higher weights). The hard constraint set \mathcal{H} represents immutable statutory obligations that must hold under all conditions.

The *Optimize()* module operates internally through a weighted MaxSAT approach, aiming to maximize the total satisfaction of soft constraints while maintaining the logical consistency of all hard constraints. Formally, the optimization problem can be written as:

$$\max_{\mathbf{x}} \sum_{s_i \in \mathcal{S}} w_i \cdot s_i(\mathbf{x}) \quad \text{s.t.} \quad \mathcal{H}(\mathbf{x}) = \text{true}, \quad (2)$$

where $s_i(\mathbf{x}) \in \{0, 1\}$ indicates whether s_i is satisfied under model \mathbf{x} .

If no fully satisfiable assignment exists, Z3 incrementally relaxes violated soft constraints by introducing auxiliary Boolean selectors $\lambda_i \in \{0, 1\}$ and minimizing their activation:

$$\min_{\mathbf{x}} \sum_{i=1}^{|\mathcal{S}|} w_i \lambda_i, \quad \text{s.t.} \quad (\mathcal{H}(\mathbf{x}) \wedge \bigwedge_i (s_i(\mathbf{x}) \vee \lambda_i)) = \text{true}. \quad (3)$$

As shown in Equation (2), the optimization seeks to maximize the satisfaction of all soft constraints under the statutory framework. When this is not achievable, Equation (3) defines the relaxation procedure that minimizes the total violation cost by activating the smallest number of λ_i variables.

This mechanism effectively searches for the *maximally satisfiable subset (MSS)* of all factual conditions, corresponding to the smallest logical modification required to restore legal compliance. Conceptually, Z3 iteratively relaxes the lowest-weighted constraints until all hard constraints are consistent,

Algorithm 3 Minimal Compliance via Weighted MaxSMT

Input: Hard constraints \mathcal{H} ; soft constraints $\mathcal{S} = \{s_i\}$ with weights $\{w_i\}$

Output: Minimal compliance model \mathbf{x}^* ; violation set Δ ; cost $C(\Delta)$

- 1: $O \leftarrow \text{Optimize}()$ ▷ Initialize MaxSMT optimizer
 - 2: Add all $h \in \mathcal{H}$ via $O.\text{add}(h)$
 - 3: Add all $s_i \in \mathcal{S}$ via $O.\text{add_soft}(s_i, w_i)$
 - 4: Run status $\leftarrow O.\text{check}()$
 - 5: **if** status = *unsat* **then**
 - 6: **return** No feasible compliance: hard constraints inconsistent
 - 7: **end if**
 - 8: $\mathbf{x}^* \leftarrow O.\text{model}()$ ▷ Optimal minimal-violation model
 - 9: Evaluate each s_i under \mathbf{x}^* ; collect $\Delta \leftarrow \{s_i \mid s_i(\mathbf{x}^*) = \text{false}\}$
 - 10: Compute $C(\Delta) \leftarrow \sum_{s_i \in \Delta} w_i$
 - 11: **return** $\mathbf{x}^*, \Delta, C(\Delta)$
-

preserving maximal factual assumptions under the fixed statutory framework.

In practice, the solver identifies minimal factual revisions—such as flipping a Boolean or adjusting a threshold—yielding \mathbf{x}^* that satisfies \mathcal{H} and minimizes $\sum_i w_i \lambda_i$. If \mathcal{H} is inconsistent, the solver returns *unsat*.

For the FSC case, the optimization result suggests a single factual flip:

```
improvement_plan_executed: false → true
```

This indicates that the firm would regain legal compliance by executing the improvement plan initially submitted, satisfying the requirements under Articles 143-4 and 143-6 of the *Insurance Act*.

To enhance interpretability, the optimization results include both the optimal model \mathbf{x}^* and the difference set Δ , which lists the factual variables altered during optimization. Each element of Δ identifies which factual condition violated compliance and how it was adjusted, forming an explicit *correction trace* interpretable by auditors or legal analysts. This trace directly supports post-hoc explanations and model accountability.

Each violated constraint indicates a specific non-compliant element, allowing transparent traceability and human-auditable reasoning. When integrated with the language model’s explanation module, each violation in Δ is translated into natural language statements such as: “The company failed to submit a remediation plan within the statutory period.” This hybrid reasoning procedure aligns with the legal principle of *minimal intervention*, ensuring that compliance is restored through the smallest set of factual corrections while maintaining logical and statutory integrity.

The optimization aims to maintain logical soundness and interpretability, forming the foundation for transparent, verifiable compliance automation.

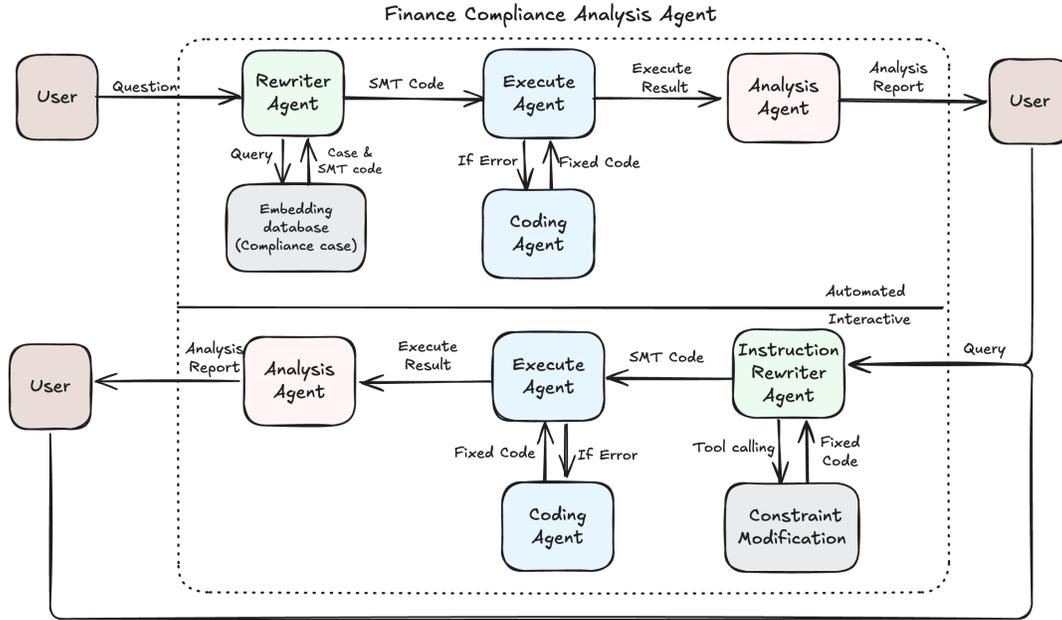


Fig. 2: System architecture of the Finance Compliance Analysis Agent.

Optimization Result When $check()$ returns sat , the model \mathbf{x}^* produced by $Optimize()$ minimizes $\sum_i w_i \lambda_i$ among all models satisfying \mathcal{H} . Therefore, the induced delta set Δ is a weighted-minimal set of factual violations. Under uniform weights, this coincides with a minimum-cardinality revision; with nonuniform weights, it yields a minimum-cost revision. In both cases, \mathbf{x}^* constitutes the desired *minimal compliance solution*.

C. Multi-Agent Compliance Platform

Building upon the constraint-generation and verification framework, we designed a modular multi-agent platform that operationalizes the reasoning process.

1) *Design Principles*: The design of the Finance Compliance Analysis Agent adheres to four core principles: (1) Modularity, allowing agents to be replaced or upgraded independently; (2) Explainability, ensured through structured analysis and symbolic reasoning; (3) Human-in-the-loop adaptability, enabling iterative refinement via natural language; and (4) adaptability, supporting future expansion into additional regulatory domains or jurisdictions.

2) *System Architecture with AI Agents*: Figure ?? illustrates the architecture of the *Finance Compliance Analysis Agent*, a modular multi-agent system designed to assess financial legal compliance. The workflow begins with a user query that sequentially passes through several specialized agents, each responsible for a distinct stage of reasoning and verification.

The Rewriter Agent reformulates the user query into a structured legal task and retrieves relevant precedents from the embedding-based *Compliance Case Database*, which stores pre-generated SMT code for constraint evaluation. The Execute Agent then runs the retrieved SMT code using the SMT solver to verify compliance with the encoded legal constraints,

supporting both automated evaluation and interactive refinement.

In cases where execution errors occur—such as syntax inconsistencies, ill-formed logic, or unbound variables—the Coding Agent collaborates with the Execute Agent to repair the SMT code. It inspects solver traces and conversation history to identify stable prior versions or generate corrective patches, thereby enabling a self-healing feedback loop between the solver and the agent.

The Analysis Agent interprets solver outputs and transforms them into structured compliance reports comprehensible to human auditors. When users provide feedback or suggest rule modifications, the Instruction Rewriter Agent translates these inputs into updated SMT logic for re-evaluation. Finally, the Modify Constraint Module performs low-level edits on SMT formulations, such as toggling logical conditions, injecting parameters, or fixing variable assignments.

Together, these coordinated agents form a robust and adaptable workflow that integrates symbolic reasoning, automated correction, and interactive human guidance for end-to-end legal compliance analysis.

TABLE I: Summary of LLM-generated SMT code

Method	Accuracy	Avg. Time (s)
LLM Generate SMT Code	0.8621	36.85

3) *Agent Collaboration and Workflow*: The system operates through coordinated collaboration across reasoning stages. The Rewriter Agent reformulates user queries and retrieves relevant SMT cases, which the Execute Agent verifies via the SMT solver. When errors occur, the Coding Agent activates the self-healing mechanism—tracing the source of failure, performing

TABLE II: Descriptive statistics of variables and constraints across 87 cases.

	#Var	Real	Int	Bool	#Con	Hard	Soft	H_Ratio	S_Ratio
Average	12.85	3.58	1.97	7.31	30.81	19.30	11.51	0.63	0.37
Min	4	0	0	1	11	6	4	0.38	0.24
Max	40	21	7	21	63	39	25	0.76	0.63
Std	5.60	3.40	1.74	4.11	9.62	6.35	4.61	0.08	0.08

rollback, or generating logical patches—until the constraints yield a valid satisfiable or unsatisfiable result.

Once successful execution is achieved, the Analysis Agent summarizes the results into interpretable compliance reports. For interactive exploration, users may modify assumptions or logic; the Instruction Rewriter Agent translates these inputs into structured tool calls dispatched to the Modify Constraint module, which updates and revalidates the SMT code.

This modular, agent-oriented workflow ensures interpretability, fault tolerance, and adaptability, supporting both automated and human-in-the-loop compliance reasoning in evolving financial regulatory environments.

IV. EXPERIMENTAL EVALUATION

To validate our framework, we conducted an empirical analysis focusing on financial penalty cases issued by the FSC of Taiwan. In total, 87 real-world cases were collected and analyzed. We also developed an interactive user interface using Gradio to facilitate real-time interaction with the system. As part of our infrastructure, we curated a comprehensive regulatory database comprising 75 distinct laws and 3,753 statutory provisions extracted from the Taiwan Law Database.

This experiment is designed to address the following research questions:

- **RQ1:** To what extent can LLMs effectively generate SMT constraints in the legal domain?
- **RQ2:** Does the integration of LLMs with SMTs enhance the model’s capability to reason about illegal terms in real-world regulatory cases, compared to using LLMs alone?
- **RQ3:** Does the integration of LLMs with SMTs improve the accuracy and consistency of legal compliance restoration in real-world cases, relative to standalone LLMs?

We evaluate these questions through both qualitative and quantitative assessments, focusing on reasoning performance, legal alignment, and correction accuracy.

All supplementary materials—including detailed SMT formulations, additional case studies, and implementation scripts—are publicly available on GitHub [28].

A. RQ1: SMT Constraint Synthesis

We conducted generation experiments on 87 regulatory cases. Using a few-shot learning approach, we prompted the LLM to generate SMT constraints tailored to each case. The generated code was then manually reviewed to verify its logical soundness and contextual correctness. If the code was valid, it was categorized as *correct*; otherwise, we prompted the model to iteratively revise its output until the issues were

resolved, and such cases were classified as *incorrect*. Human inspection was conducted solely to verify factual consistency between the generated SMT constraints and the original case context. No manual editing or correction was performed on model outputs, and to ensure fair comparison, human intervention was limited to factual verification only—without influencing any intermediate reasoning or solver correction steps. This procedure guarantees that all evaluations objectively reflect the model’s intrinsic capability.

Common errors included omissions of impacted clients in compliance outcomes, failure to zero out fines where applicable, and logically inconsistent constraint sets resulting responses from the solver. Through manual correction and validation, we ensured that all final SMT codes were accurate and could serve as ground truth for downstream evaluation.

The generation accuracy and response times are summarized in Table I.

Scale Analysis of Case Studies. To clarify the scale of our case studies, we analyzed the number of variables and constraints across all 87 cases. The detailed statistics are summarized in Table II. On average, each case contains about 13 variables (#Var) and 31 constraints (#Con), with the largest instance involving 40 variables and 63 constraints. Roughly 63% of the constraints are hard constraints (H_Ratio), while 37% are soft (S_Ratio), reflecting a balance between strict regulatory rules and flexible planning considerations.

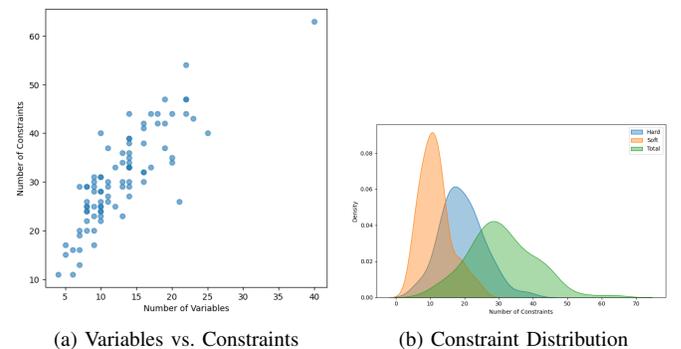


Fig. 3: Relationship between variables and constraints (a), and the distribution of hard, soft, and total constraints (b).

As illustrated in Fig. 3, the scatter plot between the number of variables and constraints reveals a clear positive correlation. This observation indicates that as the number of variables increases, the constraint set grows proportionally, suggesting that the proposed framework is suitable for small- to moderate-scale SMT problems.

Fig. 3 further shows the distribution of hard, soft, and total constraints. The hard constraints are primarily concentrated between 15 and 25, while the soft constraints fall mostly within the range of 8 to 15. Consequently, the total number of constraints typically ranges from 25 to 40, with a few outliers exceeding 60.

TABLE III: Illegal term analysis: Comparison of LLM and LLM+SMT.

Method	Avg. Illegal Terms per Case	Avg. Time (s)
LLM + SMT	5.08	0.021
LLM Only	6.40	7.01

B. RQ2: Illegal Term Reasoning

To support compliance in real-world legal cases, our system is designed to identify specific illegal terms—statutory violations—within each scenario. Highlighting these terms helps users recognize which actions must be taken to align with relevant regulations. We evaluated this capability by comparing the number of identified illegal terms and the average response time between the baseline LLM and our proposed LLM+SMT approach, as shown in Table III.

Our SMT-based method leverages *assert* and *track* commands during code generation. When contradictions arise—i.e., Z3 returns an *unsat* result—this indicates a violation between statutory requirements and the case description. For instance, if the law mandates remedial actions for capital inadequacy but the case omits such steps, the system flags this inconsistency.

Since Z3 returns only a minimal *unsat* core by default, we redefine all constraints—including originally soft ones—as hard constraints, forcing the solver to expose multiple *unsat* cores. We then extract all components traced to hard constraints across iterations to define the full set of “illegal terms.” Remaining elements from soft constraints are treated as free variables and iteratively reintroduced until the system reaches a satisfiable state. The total number of unique illegal terms is computed via set union over these rounds.

To validate this mechanism, we also queried the LLM directly using the same case and legal text. However, the LLM’s output tends to lack logical consistency and interpretability. By contrast, our SMT-enhanced method not only produces comparable results but does so with over two orders of magnitude improvement in reasoning efficiency and improved explainability.

In a real insurance case,⁴ our SMT module identified the following illegal terms in 0.191 seconds: `req_net_worth`, `self_capital_sum`, `risk_capital_sum`, `def_improvement_plan`, `req_final_CAR`, `def_pre_CAR`, and `def_total_improvement`. These correspond to formal statutory violations that must be resolved to restore compliance.

⁴See the FSC penalty announcement for Sang Shang Meibang Life Insurance Co., Ltd.: https://www.fsc.gov.tw/ch/home.jsp?id=131&parentpath=0,2&mcustomize=multimessages_view.jsp&dataserno=202407120004&dtable=Penalty

In contrast, the LLM took approximately 5 seconds and returned a longer, less precise list, including: Capital Adequacy Ratio below statutory level, Net Worth Ratio below regulatory requirement, Significant capital inadequacy, Failure to submit a feasible capital improvement plan, Failure to ensure capital adequacy in 2024, High uncertainty in planned private placement of hybrid instruments, Inadequate backup plan, Violation of Insurance Act Article 143-6 Paragraph 1 Subparagraph 1, Restricted transactions with related parties, and Noncompliance with capital adequacy regulations.

Closer inspection shows these additional terms fall into three categories: (1) **Contextual Misclassifications**—observations that indicate financial risk but are not legal violations (e.g., uncertainty in private placement); (2) **Misinterpretation of Advisory Clauses**—soft recommendations wrongly flagged as mandatory (e.g., lacking a contingency plan); and (3) **Vague or Composite Labels**—overgeneralized statements that lack traceability to specific provisions.

These findings highlight the key strengths of our SMT-based approach: precise legal modeling, symbolic consistency checking, and the ability to isolate actionable violations. In domains such as financial regulation, where auditability and precision are critical, this advantage is particularly valuable.

C. RQ3: Compliance Restoration

Following the analysis of illegal term identification, we proceed to evaluate the models’ capabilities in correcting real-world cases. This task investigates what changes are necessary to revise a case such that the resulting legal outcome would be penalty-free. We performed this correction analysis on all 87 cases.

For the SMT-based approach, we increased the weight of the penalty-related soft constraints and set the target penalty value to zero. This allows the SMT solver to return a satisfiable assignment that eliminates the penalty, which serves as the *ground truth* for comparison. Importantly, the ground truth in this task does not refer to a single unique assignment but rather to the satisfiability status (*SAT* or *UNSAT*) of the case. For satisfiable instances, multiple valid assignments may exist, and any assignment that satisfies all constraints is accepted as correct. Thus, our evaluation does not require the LLM to replicate a specific SMT-generated assignment; it only needs to produce a valid satisfiable solution when one exists.

In total, 63 cases were satisfiable, among which only 13 were accompanied by a sample assignment for reference, while the others were not restricted to a single solution. This setup aligns with real-world legal reasoning, where multiple compliant revisions of a case may coexist.

In this experiment, the LLM is provided with the soft constraint variables and their original values. It is allowed to freely decide whether to flip the value of a given variable or assign a new one. The LLM also receives the corresponding case text and statutory references to support its reasoning process.

To further evaluate model robustness, we randomly selected 24 cases and introduced contradictory constraints into the

TABLE IV: Performance of case Compliance Restoration.

Method	Accuracy	Precision	0	F1 Score	Avg. Time (s)
LLM + SMT	1.0000	1.0000	1.0000	1.0000	0.0040
LLM Only	0.3333	0.6190	0.2063	0.3080	1.4700

TABLE V: Confusion matrix for LLM without SMT Solver.

Prediction	Actually Satisfiable	Actually Unsatisfiable
Satisfiable	13 (TP)	8 (FP)
Unsatisfiable	50 (FN)	16 (TN)

solver. For example, we may define a capital adequacy ratio requirement of over 200%, where the ratio is calculated as own capital divided by risk-weighted capital. If we then add a constraint setting own capital to zero, the system becomes unsatisfiable. These newly introduced constraints are also provided to the LLM. If the LLM correctly reasons that the problem has no feasible solution, it should return an unsatisfiable judgment.

The performance comparison is shown in Table IV, and the confusion matrix for the LLM-only baseline is shown in Table V.

From the results, we observe that the LLM tends to default to predicting an unsatisfiable outcome, even for cases where a valid solution exists. This bias leads to a significant number of false negatives. Furthermore, the LLM’s compliance restoration lacks consistency and completeness compared to the SMT solver. In contrast, the SMT-based method achieves high accuracy in compliance correction tasks and substantially faster inference, completing each correction in approximately 0.004 seconds—over 300 times faster than the average LLM response time.

D. Summary of Experimental Findings

Across all three research questions, the results consistently demonstrate the effectiveness of integrating SMT solvers into the legal reasoning pipeline. For **RQ1**, the LLM generated syntactically and semantically correct SMT code in over 86% of cases, enabling precise modeling of legal constraints. In **RQ2**, the SMT-enhanced method achieved comparable illegal term identification to the baseline LLM while offering faster responses and greater logical clarity. For **RQ3**, the Compliance Restoration powered by LLM+SMT reached perfect

accuracy, outperforming the baseline across all major metrics. Overall, the hybrid LLM+SMT framework proves both robust and interpretable, showing strong potential for practical legal reasoning and compliance correction.

Cross-Jurisdictional Adaptability. Although our evaluation focuses on Taiwan’s financial regulatory environment, the proposed framework is *jurisdiction-agnostic*. By adapting the retrieval corpus and prompt exemplars to reflect the case precedents and statutory structures of target jurisdictions, the system can be extended to common law settings such as the United States or the United Kingdom. Future work will explore multilingual and precedent-driven adaptations to further evaluate cross-jurisdictional applicability.

V. CONCLUSION

We propose **Neuro-Symbolic Compliance**, a hybrid framework integrating LLM-based understanding with SMT-based reasoning through a multi-agent design. Each agent models domain-specific expertise, enabling interpretable and verifiable analysis of financial regulation. Applied to Taiwan’s FCS cases, our system outperforms LLM-only baselines in reasoning consistency and inference speed. This work demonstrates how agent-based neuro-symbolic AI can bridge natural legal language and machine-verifiable logic toward scalable regulatory automation.

VI. ACKNOWLEDGEMENT

This work is partially supported by the National Science and Technology Council (NSTC), Taiwan, under Grant No. 113-2221-E-004-010-MY2 and 114-2221-E-002-183-MY3, and by the NTU-Delta Electronics Innovation Research Funding Project under Grant No. NTU-DE-114FR032, and by Chunghwa Telecom Co., Ltd. under the commissioned project “National Chengchi University Financial AI Research Project: Application of Traditional Chinese Large Language Models and Artificial Intelligence Agent Platforms for Investment Analysis Report Generation.”

TABLE VI: Internationalization Potential: Minimal Engineering Modifications Required for Legal Domain Transfer

Module	Purpose	Should it be modified?
RAG Corpus	Retrieve relevant legal text and precedents	Yes (It should be replaced with a local legal corpus.)
Prompt Template	Control query format and language context	Yes (It should be adjusted for language and legal phrasing.)
Clause Parser	Map legal clauses to logical templates	Partial (Parsing patterns should be redefined.)
Legal Mapping Dictionary	Link clause labels to statutes	Yes (It should be updated for the target jurisdiction.)
Z3 Rule Engine	Enforce symbolic legal constraints	No (It can be reused without modification.)
SMT Relaxation Loop	Perform minimal correction search	No (It remains domain-independent.)
Gradio UI	Provide user interface	No (It can be directly reused.)

REFERENCES

- [1] T. Libal and T. Novotná, “Towards transparent legal formalization,” in *Explainable and Transparent AI and Multi-Agent Systems: Third International Workshop, EXTRAAMAS 2021, Virtual Event, May 3–7, 2021, Revised Selected Papers 3*. Springer, 2021, pp. 296–313.
- [2] I. Ponkin and A. Redkina, “Digital formalization of law,” *International Journal of Open Information Technologies*, vol. 7, no. 1, pp. 39–48, 2019.
- [3] H. Surden, “Computable law and artificial intelligence,” *U of Colorado Law Legal Studies Research Paper Forthcoming, Cambridge Handbook of Private Law and Artificial Intelligence (forthcoming 2024)*, 2024.
- [4] S. E. Mohsen, A. Hamdan, and H. M. Shoaib, “Digital transformation and integration of artificial intelligence in financial institutions,” *Journal of Financial Reporting and Accounting*, 2024.
- [5] R. J. Bolton and D. J. Hand, “Statistical fraud detection: A review,” *Statistical science*, vol. 17, no. 3, pp. 235–255, 2002.
- [6] D. W. Arner, J. Barberis, and R. P. Buckley, “The evolution of fintech: A new post-crisis paradigm,” *Geo. J. Int’l L.*, vol. 47, p. 1271, 2015.
- [7] C. Barrett, C. L. Conway, M. Deters, L. Hadarean, D. Jovanović, T. King, A. Reynolds, and C. Tinelli, “cvc4,” in *Computer Aided Verification: 23rd International Conference, CAV 2011, Snowbird, UT, USA, July 14–20, 2011. Proceedings 23*. Springer, 2011, pp. 171–177.
- [8] L. De Moura and N. Bjørner, “Z3: An efficient smt solver,” in *International conference on Tools and Algorithms for the Construction and Analysis of Systems*. Springer, 2008, pp. 337–340.
- [9] C. A. Bradley and T. W. Morrison, “Presidential power, historical practice, and legal constraint,” *Columbia Law Review*, pp. 1097–1161, 2013.
- [10] S. Athul, A. Saxena, J. Mahajan, L. Panikulangara, S. Kulkarni, and S. Bang, “Legalmind system and the llm-based legal judgment query system,” in *2024 International Conference on Trends in Quantum Computing and Emerging Business Technologies*. IEEE, 2024, pp. 1–5.
- [11] D. Shu, H. Zhao, X. Liu, D. Demeter, M. Du, and Y. Zhang, “Lawllm: Law large language model for the us legal system,” in *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*, 2024, pp. 4882–4889.
- [12] P. Hitzler, A. Eberhart, M. Ebrahimi, M. K. Sarker, and L. Zhou, “Neuro-symbolic approaches in artificial intelligence,” *National Science Review*, vol. 9, no. 6, p. nwac035, 2022.
- [13] J. Mao, J. B. Tenenbaum, and J. Wu, “Neuro-symbolic concepts,” *arXiv preprint arXiv:2505.06191*, 2025.
- [14] Y. Hao, Y. Zhang, and C. Fan, “Planning anything with rigor: General-purpose zero-shot planning with llm-based formalized programming,” *arXiv preprint arXiv:2410.12112*, 2024.
- [15] M. Chen, G. Li, L.-I. Wu, R. Liu, Y. Su, X. Chang, and J. Xue, “Can language models pretend solvers? logic code simulation with llms,” in *International Symposium on Dependable Software Engineering: Theories, Tools, and Applications*. Springer, 2024, pp. 102–121.
- [16] Y. Zhang, Y. Cai, X. Zuo, X. Luan, K. Wang, Z. Hou, Y. Zhang, Z. Wei, M. Sun, J. Sun *et al.*, “The fusion of large language models and formal methods for trustworthy ai agents: A roadmap,” *arXiv preprint arXiv:2412.06512*, 2024.
- [17] A. Berger, L. Hillebrand, D. Leonhard, T. Deußer, T. B. F. De Oliveira, T. Dilmaghani, M. Khaled, B. Kliem, R. Loitz, C. Bauckhage, and R. Sifa, “Towards automated regulatory compliance verification in financial auditing with large language models,” in *2023 IEEE International Conference on Big Data (BigData)*, 2023, pp. 4626–4635.
- [18] H. Liao, S. He, Y. Xu, Y. Zhang, K. Liu, and J. Zhao, “Neural-symbolic collaborative distillation: Advancing small language models for complex reasoning tasks,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 39, no. 23, 2025, pp. 24 567–24 575.
- [19] M. Ning, Z. Zhou, Q. Wang, X. Huang, and K. Huang, “Gns: Solving plane geometry problems by neural-symbolic reasoning with multimodal llms,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 39, no. 23, 2025, pp. 24 957–24 965.
- [20] J. Shi, X. Ding, H. Zhao, T. Liu, and B. Qin, “Bridging neural and symbolic reasoning: A dual-system framework for interpretable question answering,” in *ICASSP 2025 - 2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2025, pp. 1–5.
- [21] S. Judson, M. Elacqua, F. Cano, T. Antonopoulos, B. Könighofer, S. J. Shapiro, and R. Piskac, “soid: A tool for legal accountability for automated decision making,” in *International Conference on Computer Aided Verification*. Springer, 2024, pp. 233–246.
- [22] N. Feng, L. Marsso, M. Sabetzadeh, and M. Chechik, “Early verification of legal compliance via bounded satisfiability checking,” in *International Conference on Computer Aided Verification*. Springer, 2023, pp. 374–396.
- [23] C. Tsigkanos, A. Arleo, J. Sorger, and S. Dustdar, “How do firms transact? guesstimation and validation of financial transaction networks with satisfiability,” in *2019 IEEE 20th International Conference on Information Reuse and Integration for Data Science (IRI)*. IEEE, 2019, pp. 15–22.
- [24] M. Raza and N. Milic-Frayling, “Almost sure reasoning: Generating verified formalizations with language models and logical solvers.”
- [25] T. H. Beach, Y. Rezgui, H. Li, and T. Kasim, “A rule-based semantic approach for automated regulatory compliance in the construction sector,” *Expert Systems with Applications*, vol. 42, no. 12, pp. 5219–5231, 2015.
- [26] S. Wachter, B. Mittelstadt, and C. Russell, “Counterfactual explanations without opening the black box: Automated decisions and the gdpr,” *Harv. JL & Tech.*, vol. 31, p. 841, 2017.
- [27] H. Andreas, M. Armgardt, and M. Gunther, “Counterfactuals for causal responsibility in legal contexts,” *Artificial intelligence and law*, vol. 31, no. 1, pp. 115–132, 2023.
- [28] Y.-S. Hsia and F. Yu, “Neuro-symbolic-compliance,” 2025. [Online]. Available: <https://github.com/shiayongshen/Neuro-symbolic-compliance>