

Legal and Ethical Reasoning in Generative AI Models via Prompt-Chaining and Principle-Constrained Alignment with Trusted Federated Explainability

Mohan Siva Krishna Konakanchi

mohansivakrishna16@gmail.com

Abstract—Generative AI systems increasingly support legal and compliance workflows, including policy analysis, document drafting, risk triage, and ethical review. Yet deploying these systems in high-stakes environments requires more than fluent text generation: outputs must be grounded in structured reasoning, must align with legal and ethical principles, and must be auditable under organizational governance. Two practical obstacles complicate real-world adoption. First, legal and ethical reasoning often requires multi-step deliberation: identifying facts, classifying issues, applying rules, considering exceptions, and communicating uncertainty. Second, relevant data and expertise are fragmented across organizational silos (regions, business units, law firms, and regulators), limiting centralized training while increasing integrity risks such as poisoning, low-quality updates, and accountability gaps. Additionally, organizations must manage a recurring tension: increasing explainability (e.g., richer rationales, evidence summaries, and traceability) can reduce throughput and sometimes degrade performance, whereas high-performing black-box generation can be hard to justify in audits.

This paper proposes *LEAP-Chain* (Legal and Ethical Alignment via Prompt-Chaining), a framework that combines structured prompt-chaining with principle-constrained alignment and a trust metric-based federated learning (FL) governance layer. *LEAP-Chain* decomposes a user request into a chain of constrained reasoning stages (fact extraction, issue spotting, principle retrieval, risk assessment, and response drafting), where each stage emits compact intermediate artifacts suitable for audit. To support cross-silo learning without centralizing sensitive documents, *LEAP-Chain* introduces: (i) a trust metric that scores participants using provenance attestations, update consistency, evaluation reliability, and policy compliance; (ii) trust-aware robust aggregation that reduces poisoning and limits low-quality contributions; and (iii) *federated explainability*, where explanation quality and stability are measured locally and shared as summaries. Finally, *LEAP-Chain* provides a budgeted controller that quantifies and optimizes the explainability–performance trade-off by allocating explanation effort to high-impact requests and high-risk domains.

We evaluate *LEAP-Chain* using a controlled prototype simulation of legal/compliance tasks across heterogeneous silos with non-IID distributions, variable policy regimes, and adversarial/faulty contributors. Results indicate that structured chaining reduces hallucinated assertions and improves consistency under ambiguity, trust-aware federation improves robustness and lowers policy-violation rates compared to standard FL baselines, and moderate explanation budgets deliver stable, actionable

rationales with limited loss in task effectiveness. We conclude with deployment guidance for responsible legal and ethical reasoning in generative systems under enterprise governance.

Index Terms—legal reasoning, ethical reasoning, generative AI, alignment, prompt chaining, governance, federated learning, trust metrics, explainable AI, accountability

I. INTRODUCTION

Generative language models are increasingly embedded in business processes that touch legal compliance and ethical risk: privacy impact assessments, contract review, regulatory summaries, policy drafting, and employee guidance. These workflows require not only accuracy but also disciplined reasoning and defensible explanations. Errors can be costly: incorrect compliance advice can trigger regulatory exposure, and unethical recommendations can harm users and brand trust. High-stakes domains therefore demand systems that are safe-by-design, auditable, and robust to adversarial and operational failures.

Two gaps remain in many deployments. First, a single prompt response is often inadequate for legal and ethical reasoning. Legal analysis typically proceeds through multiple steps: gather facts, identify applicable legal issues, apply rules and precedents, consider exceptions, and communicate uncertainty. Ethical analysis similarly requires identifying stakeholders, mapping harms and benefits, and applying principles (e.g., non-maleficence, fairness, transparency). Treating these as a single generation problem encourages confident but brittle outputs.

Second, legal and compliance data is fragmented across silos. Enterprises store sensitive contracts, case summaries, internal policies, and regulatory interpretations across regions and departments. Centralizing this data for training or evaluation is frequently disallowed. Federated learning (FL) provides a pathway to improve shared models without pooling raw data [13], [17]. However, standard FL emphasizes privacy, not integrity. Legal and ethical reasoning systems require stronger governance: participant influence must be accountable, low-quality updates must be limited, and the overall learning loop must be auditable.

A. Prompt-Chaining as Structured Deliberation

We use *prompt-chaining* as an operational pattern: decompose a complex request into a sequence of constrained sub-prompts, each producing a small artifact. Unlike unconstrained chain-of-thought disclosure, we focus on compact, policy-safe intermediate products (e.g., issue lists, principle references, and uncertainty flags) that can be logged and reviewed. This approach draws from broader trends in structured reasoning and modular decision systems, where intermediate representations improve controllability and debugging.

B. Principle-Constrained Alignment

Many organizations adopt principle-based ethical guidelines and compliance policies. We generalize “constitutional” alignment as *principle-constrained alignment*: generation is guided by an explicit set of principles (privacy, non-discrimination, transparency, safety) and legal compliance constraints. Rather than relying solely on training-time alignment, we enforce principles at inference time through checks and refinements, and at training time through federated policy signals.

C. Explainability–Performance Trade-off

Explanations are essential: legal stakeholders require justification and traceability. Yet explanations can be costly and can constrain outputs. Explainability methods (e.g., local attributions and rule-like rationales) provide primitives [8]–[10], but in high-stakes decisions interpretability concerns remain [12]. We therefore treat explainability as a resource that can be budgeted: the system should explain more when risk is higher.

D. Problem Statement

We address: **P1**: How to improve legal and ethical reasoning reliability using structured prompt-chaining and principle constraints. **P2**: How to enable cross-silo improvement via FL while ensuring integrity and accountability with trust metrics. **P3**: How to quantify and optimize explainability versus performance under operational budgets.

E. Contributions

This paper proposes *LEAP-Chain* and contributes:

- A prompt-chaining pipeline for legal and ethical reasoning with compact intermediate artifacts suitable for audit.
- A principle-constrained alignment mechanism that enforces legal/ethical guardrails during generation and revision.
- A trust metric-based federated governance framework that scores silos on provenance, update consistency, evaluation reliability, and policy compliance, using trustaware robust aggregation [15], [16].
- A budgeted explainability controller that explicitly manages the explainability–performance trade-off with stability checks.
- A prototype simulation evaluation showing robustness gains and improved safety outcomes under heterogeneity and integrity failures.

II. RELATED WORK

A. AI Safety, Alignment, and Human Feedback

Alignment research in the 2010s emphasized reducing harmful behavior through problem formulations and scalable oversight. Work on concrete AI safety problems framed robustness, reward hacking, and distribution shift as central risks [1]. Preference learning and human feedback approaches demonstrated how to align behavior with human judgments [3]. Cooperative inverse reinforcement learning and related ideas examined learning from humans under uncertainty [2]. Scalable oversight and aligned agent behavior were also explored through training signals and safety framing [4]. While these works do not directly prescribe prompt-chaining, they motivate explicit constraints, oversight, and auditability for high-stakes reasoning.

B. Ethics, Accountability, and Transparency

The 2010s also produced extensive analysis of algorithmic accountability and ethical constraints. Researchers emphasized that transparency and governance must be treated as system-level requirements, not afterthoughts [6]. Ethical analysis frameworks and guidance for AI governance highlight that stakeholder harm, fairness, and accountability are central [7]. These perspectives motivate our principle-constrained alignment and trust-based governance layer.

C. Explainability and Interpretability

Model-agnostic and gradient-based explanation techniques provide tools for interpreting predictions [8]–[10]. Anchors produce rule-like explanations emphasizing precision and human usability [11]. Concerns about using post-hoc explanations for high-stakes decisions motivate interpretable designs and constrained systems [12]. LEAP-Chain adopts explanation budgets and stability checks to support auditable reasoning without overwhelming operational constraints.

D. Federated Learning, Privacy, and Robustness

Federated averaging formalized communication-efficient training over decentralized data [13]. Later surveys highlighted open problems including non-IID data, heterogeneity, and security [17]. Secure aggregation protects privacy of participant updates [14]. Byzantine-robust methods address adversarial updates [15], [16]. However, legal/ethical reasoning requires more than statistical robustness: participant influence must be accountable. We therefore introduce explicit trust metrics and audit artifacts.

E. Machine Learning for Law and Legal Text

The 2010s saw growth in legal NLP: datasets and models for legal judgment prediction and classification tasks. For example, work predicting outcomes at the European Court of Human Rights demonstrated potential and limitations [18]. Studies in legal analytics and computational law emphasized the challenges of capturing legal reasoning and context [19]. These works motivate our emphasis on uncertainty and trace-ability.

F. Auditability Infrastructure

Permissioned blockchain systems and secure logging infrastructures support tamper-resistant records and provenance [20], [21]. LEAP-Chain uses an audit plane to record commitments and rationales for governance without exposing sensitive content.

III. LEAP-CHAIN FRAMEWORK OVERVIEW

LEAP-Chain integrates three modules: (i) structured prompt-chaining, (ii) principle-constrained alignment, and (iii) trusted federated explainability.

A. System Components

C1: Chaining Orchestrator. Decomposes requests into a sequence of stages with constrained outputs.

C2: Principle Library. A curated set of legal and ethical principles represented as short, auditable rules and checklists (e.g., privacy, non-discrimination, informed consent, uncertainty disclosure). In enterprise deployments, this library is versioned and approved by governance stakeholders.

C3: Memory and Context Store. Stores project-specific policies, jurisdiction metadata, and approved templates. Sensitive raw documents remain silo-local.

C4: Federated Governance Layer. Trains shared components (e.g., risk classifiers, issue spotters, retrieval scorers) without centralizing raw text.

C5: Explainability Layer. Produces compact rationales and stability metrics; shares only summaries across silos.

B. Operational Philosophy

LEAP-Chain aims for *disciplined reasoning*:

- reduce overconfident claims by forcing uncertainty signaling,
- separate *fact extraction* from *normative judgment*,
- treat legal advice boundaries explicitly (“not legal advice” in user-facing deployments),
- produce artifacts that are reviewable and auditable.

IV. PROMPT-CHAINING FOR LEGAL AND ETHICAL REASONING

We define a chain as a fixed sequence of stages; each stage produces a small artifact.

A. Stage S1: Fact Pattern Extraction

Given user input, S1 extracts:

- relevant facts (entities, jurisdiction, timeline),
- missing information prompts,
- assumptions explicitly labeled as assumptions.

Output is a compact bullet list with uncertainty flags.

B. Stage S2: Issue Spotting and Task Classification

S2 identifies the category:

- privacy compliance, employment policy, consumer protection, contract clause, etc.,
- ethical risk themes: fairness, safety, transparency, consent.

This stage selects which principle subsets apply.

C. Stage S3: Principle Retrieval and Constraint Assembly

S3 retrieves relevant principles from the library and con-

structs a constraint checklist:

- **hard constraints:** must not recommend illegal acts, must avoid discriminatory guidance,
- **soft constraints:** prefer safer alternatives, include disclaimers, encourage professional review when needed.

Principles are cited by identifier and version for audit.

D. Stage S4: Risk Assessment and Uncertainty Characterization

S4 produces a risk summary:

- high/medium/low risk label,
- key uncertainty drivers (missing facts, jurisdiction ambiguity),
- recommended next questions for clarification,
- boundaries of what the system can safely conclude.

E. Stage S5: Draft Response Generation

S5 generates the user-facing response with:

- structured sections (facts, applicable principles, analysis, actions),
- explicit uncertainty statements,
- safe alternatives and escalation guidance.

F. Stage S6: Self-Check and Revision

S6 applies a checklist-based review:

- verify that constraints from S3 are satisfied,
- check for ungrounded legal claims,
- check for harmful or discriminatory recommendations,
- rewrite to reduce overconfidence.

The system records a compact “check report” as an audit artifact.

G. Why Chaining Helps

Chaining improves reliability by:

- separating extraction and judgment,
- forcing explicit principle retrieval,
- making uncertainty and escalation systematic,
- enabling targeted evaluation at each stage.

V. PRINCIPLE-CONSTRAINED ALIGNMENT MECHANISM

We operationalize principle-constrained alignment via two enforcement modes.

A. Inference-Time Constraint Enforcement

At inference time:

- generation is conditioned on principle checklists,
- responses are revised if violations are detected,
- high-risk outputs trigger stronger templates and disclaimers.

B. Training-Time Policy Signals

During learning (including federated learning):

- outputs violating principles are penalized via local feedback signals,
- explanation stability metrics influence whether updates are trusted,
- policy compliance becomes part of the trust score.

C. Policy Versioning and Governance

Principles evolve. LEAP-Chain treats the principle library as a versioned artifact. All outputs and evaluations record:

- principle version,
- jurisdiction context (if provided),
- risk tier that determined enforcement strength.

This supports audit and incident response.

VI. TRUSTED FEDERATED LEARNING FOR CROSS-SILO LEGAL/ETHICAL REASONING

A. Federated Targets

We federate components that are beneficial across silos but do not require sharing raw text:

- issue spotting classifier (S2),
- risk classifier (S4),
- retrieval scorer for principle selection (S3),
- violation detector for constraints (S6),
- summarization templates for compact audit artifacts.

B. Threat Model

Silos may be:

- **honest:** reliable evaluations and policy compliance,
- **faulty:** unstable pipelines or noisy labels,
- **malicious:** poisoning updates to relax constraints or bias outputs.

Accountability evasion includes missing provenance, unverifiable evaluation, and incomplete reporting.

C. Trust Metric (Operational Definition)

Each participant i receives a trust score $T_i \in [0, 1]$ composed of:

- **Provenance and reproducibility (P_i):** signed attestations for config, policy version, and evaluation protocol.
- **Update consistency (U_i):** anomaly detection on update magnitude and drift across rounds.
- **Evaluation reliability (E_i):** stability across reruns and variance bounds.
- **Policy compliance (C_i):** rate of principle violations detected locally.
- **Audit completeness (A_i):** completeness of submitted evidence summaries and check reports.

Guardrails and penalties. Trust decreases sharply for:

- missing provenance,
- repeated policy violations,
- inconsistent evaluation reporting,
- persistent outlier updates across rounds.

D. Trust-Aware Robust Aggregation

Aggregation influence is:

$$\text{Influence} = \text{data/usage weight} \times \text{trust weight}.$$

After gating low-trust participants, robust aggregation reduces remaining outliers using trimmed or selection-based methods [15], [16]. Secure aggregation can protect privacy of updates [14].

E. Audit Plane for Accountability

The audit plane records:

- global model lineage per round,
- trust score rationale summaries,
- policy/principle version per round,
- hashes of evaluation sets and stage-wise metrics,
- number of gated participants and reason categories.

Permissioned ledger approaches can provide tamper resistance [20], [21].

VII. FEDERATED EXPLAINABILITY AND TRADE-OFF OPTIMIZATION

A. Explainability Targets

LEAP-Chain explains:

- **Output rationale:** what principles and facts influenced the response.
- **Stage-level artifacts:** issue list, risk tier, uncertainty drivers.
- **Governance actions:** why a silo update was gated or down-weighted.

B. Explanation Methods

We use lightweight artifacts plus optional model attributions:

- **Trace explanations:** show which principles and which extracted facts were used.
- **Attribution explanations:** local attributions with LIME/SHAP/IG primitives [8]–[10].
- **Rule-like anchors:** high-precision rationale snippets when feasible [11].

Only summaries and stability statistics are shared across silos to preserve privacy.

C. Explanation Quality Metrics

Operational metrics:

- **Stability:** top-k principle/fact agreement under small perturbations.
- **Actionability:** whether rationale maps to editable principles or missing facts.
- **Fidelity:** local alignment between explanation and model behavior.

D. Budgeted Controller

An *explanation budget* determines:

- which requests receive deep explanations (high-risk, high-impact),
- whether stability checks are enforced,
- which explanation method is used (trace vs attribution).

The controller optimizes a simple utility notion:

Utility increases with task effectiveness and explanation quality, and decreases with latency and operational cost.

This makes the explainability–performance trade-off explicit and governable.

VIII. METHODOLOGY

A. Prototype Evaluation Setup

We evaluate LEAP-Chain using a controlled simulation representing an enterprise legal/compliance deployment:

- $N = 20$ silos (regions/business units) with distinct policy regimes and document styles,
- non-IID request distributions: different common issues per silo,
- heterogeneous data quality: some silos have noisy labels or incomplete policies,
- faulty and adversarial contributors.

B. Tasks

We simulate three task families:

- **T1 Policy Q&A:** answer employee questions using internal policy summaries.
- **T2 Drafting assistance:** draft compliant language with disclaimers and constraints.
- **T3 Risk triage:** classify and escalate high-risk requests (privacy, discrimination, safety).

C. Baselines

We compare:

- **B1 Single-shot generation:** one-pass response without chaining or principle checks.
- **B2 Chaining only:** staged prompts without trust-based federation.
- **B3 FedAvg governance:** federated updates without trust scoring [13].
- **B4 Robust-only FL:** robust aggregation without trust [16].
- **LEAP-Chain:** chaining + principle constraints + trust-aware robust FL + budgeted explainability.

D. Metrics

We measure:

- **Task effectiveness:** normalized success score (correct structure, appropriate escalation).
- **Policy violation rate:** fraction of outputs violating principle constraints.
- **Hallucination proxy:** ungrounded legal claims (detected by rule-based checks).

- **Robustness drop:** degradation under adversarial/faulty silos.
- **Explanation stability:** top-k stability for rationale factors.

E. Integrity Failure Injection

We include:

- **Faulty silos (4):** unstable evaluation and noisy labels.
- **Adversarial silos (2):** poisoned updates to reduce violation detection and weaken constraints.

IX. EXPERIMENTS

A. Federated Training Protocol

We simulate 30 federated rounds. Each round updates the shared classifiers and checkers (issue spotting, risk classifier, violation detector). Raw requests remain silo-local. Silo evaluations produce signed summaries and stability statistics.

B. Explainability Budget Regimes

- **E1 Low:** deep explanations for top 5% high-risk requests.
- **E2 Medium:** deep explanations for top 20% with stability checks.
- **E3 High:** deep explanations for all responses.

X. RESULTS

To avoid IEEE formatting issues, tables have minimal columns.

A. Effectiveness and Safety Gains from Chaining

Table I compares single-shot vs chaining.

TABLE I
CHAINING VS SINGLE-SHOT: EFFECTIVENESS AND VIOLATIONS

| Method | Effectiveness | Viol. Rate |
|------------------|---------------|-------------|
| B1 Single-shot | 0.74 | 0.10 |
| B2 Chaining only | 0.81 | 0.06 |

Chaining improves effectiveness and reduces violations by enforcing a structured flow and systematic constraint checks.

B. Robustness Under Integrity Failures

Table II reports outcomes under faulty/adversarial silos.

TABLE II
ROBUSTNESS UNDER INTEGRITY FAILURES

| Method | Robust Drop | Viol. Rate |
|----------------------|-------------|-------------|
| B3 FedAvg governance | 0.06 | 0.08 |
| B4 Robust-only FL | 0.04 | 0.07 |
| LEAP-Chain | 0.02 | 0.03 |

Trust-aware gating reduces the impact of poisoned updates and improves safety. Robust-only aggregation helps but does not enforce accountability signals such as provenance and evaluation reliability.

TABLE III
EXPLAINABILITY BUDGET TRADE-OFF (LEAP-CHAIN)

| Budget | Effectiveness | Expl. Stability |
|-----------|---------------|-----------------|
| E1 Low | 0.83 | 0.60 |
| E2 Medium | 0.82 | 0.76 |
| E3 High | 0.81 | 0.79 |

C. Explainability–Performance Trade-off

Table III shows how budgets affect effectiveness and explanation stability for LEAP-Chain.

Moderate budgets yield large stability gains with limited performance loss, supporting a governance approach that prioritizes deep explanations for high-risk cases.

XI. DISCUSSION

A. Practicality for Legal and Compliance Workflows

LEAP-Chain is designed to complement, not replace, human legal review. It supports:

- consistent drafting with disclaimers and escalation triggers,
- structured risk triage with uncertainty disclosure,
- auditable artifacts for governance review.

B. Why Trust and Accountability Are Essential

In legal/ethical reasoning, a small number of poisoned or low-quality updates can cause widespread harm by weakening constraint enforcement or increasing ungrounded claims. Trust metrics ensure:

- influence is conditional on provenance and evaluation reliability,
- policy violators are down-weighted,
- audit logs support post-incident investigation.

C. Interpretable-First vs Hybrid Alignment

Post-hoc explainability may be insufficient in high-stakes contexts [12]. LEAP-Chain supports:

- **interpretable-first mode:** rely on trace explanations and strict templates, reducing model freedom,
- **hybrid mode:** allow higher flexibility but require budgeted explanations and stability checks.

D. Limitations

Jurisdiction and legal specificity. Legal reasoning depends heavily on jurisdiction and facts; any generic system must avoid overclaiming.

Evaluation realism. Our evaluation is a controlled prototype simulation; real deployments require curated test suites and human review.

Trust gaming. Participants may optimize trust metrics; guardrails and periodic audits mitigate but do not eliminate this risk.

Explanation faithfulness. Explanations can be imperfect; stability checks improve reliability but do not guarantee causality.

XII. CONCLUSION

This paper introduced LEAP-Chain, a framework for legal and ethical reasoning in generative AI systems using structured prompt-chaining and principle-constrained alignment under trusted federated governance. LEAP-Chain decomposes requests into auditable stages, enforces explicit principles during generation and revision, and enables cross-silo improvement without centralizing sensitive data via a trust metric-based federated learning layer with robust aggregation. It further provides a budgeted controller to quantify and optimize the explainability–performance trade-off, allocating explanation effort to high-impact and high-risk cases. Prototype simulation results suggest improvements in effectiveness, significant reductions in policy violations under integrity failures, and stable explanations under moderate budgets. Future work includes stronger domain-specific evaluation suites, richer provenance attestations, and longitudinal studies of stakeholder trust and compliance outcomes in real enterprise deployments.

ACKNOWLEDGMENT

The author thanks the research community for foundational work on AI safety, explainability, federated learning, and computational law that informed this framework perspective.

REFERENCES

- [1] D. Amodi *et al.*, “Concrete problems in AI safety,” *arXiv preprint arXiv:1606.06565*, 2016.
- [2] D. Hadfield-Menell, S. Russell, P. Abbeel, and A. Dragan, “Cooperative inverse reinforcement learning,” in *Proc. NeurIPS*, 2016.
- [3] P. F. Christiano *et al.*, “Deep reinforcement learning from human preferences,” in *Proc. NeurIPS*, 2017.
- [4] J. Leike *et al.*, “Scalable agent alignment via reward modeling: A research direction,” *arXiv preprint arXiv:1811.07871*, 2018.
- [5] S. Russell, *Human Compatible: Artificial Intelligence and the Problem of Control*. Viking, 2019.
- [6] B. D. Mittelstadt, P. Allo, M. Taddeo, S. Wachter, and L. Floridi, “The ethics of algorithms: Mapping the debate,” *Big Data & Society*, vol. 3, no. 2, 2016.
- [7] L. Floridi *et al.*, “AI4People—An ethical framework for a good AI society: Opportunities, risks, principles, and recommendations,” *Minds and Machines*, vol. 28, no. 4, pp. 689–707, 2018.
- [8] M. T. Ribeiro, S. Singh, and C. Guestrin, “Why should I trust you?: Explaining the predictions of any classifier,” in *Proc. ACM KDD*, 2016.
- [9] S. M. Lundberg and S.-I. Lee, “A unified approach to interpreting model predictions,” in *Proc. NeurIPS*, 2017.
- [10] M. Sundararajan, A. Taly, and Q. Yan, “Axiomatic attribution for deep networks,” in *Proc. ICML*, 2017.
- [11] M. T. Ribeiro, S. Singh, and C. Guestrin, “Anchors: High-precision model-agnostic explanations,” in *Proc. AAAI*, 2018.
- [12] C. Rudin, “Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead,” *Nature Machine Intelligence*, vol. 1, no. 5, pp. 206–215, 2019.
- [13] H. B. McMahan *et al.*, “Communication-efficient learning of deep networks from decentralized data,” in *Proc. AISTATS*, 2017.
- [14] K. Bonawitz *et al.*, “Practical secure aggregation for privacy-preserving machine learning,” in *Proc. ACM CCS*, 2017.
- [15] P. Blanchard, E. Mhamdi, R. Guerraoui, and J. Stainer, “Machine learning with adversaries: Byzantine tolerant gradient descent,” in *Proc. NeurIPS*, 2017.
- [16] D. Yin, Y. Chen, K. Ramchandran, and P. Bartlett, “Byzantine-robust distributed learning: Towards optimal statistical rates,” in *Proc. ICML*, 2018.
- [17] P. Kairouz *et al.*, “Advances and open problems in federated learning,” *arXiv preprint arXiv:1912.04977*, 2019.



- [18] N. Aletras, D. Tsarapatsanis, D. Preo, tiuc-Pietro, and V. Lampos, "Predicting judicial decisions of the European Court of Human Rights: A natural language processing perspective," *PeerJ Computer Science*, vol. 2, e93, 2016.
- [19] D. M. Katz, M. J. Bommarito, and J. Blackman, "A general approach for predicting the behavior of the Supreme Court of the United States," *PLoS ONE*, vol. 12, no. 4, e0174698, 2017.
- [20] E. Androulaki *et al.*, "Hyperledger Fabric: A distributed operating system for permissioned blockchains," in *Proc. EuroSys*, 2018.
- [21] B. Putz, F. Pernul, and G. Kablitz, "A secure and auditable logging infrastructure based on a permissioned blockchain," *Computers & Security*, vol. 87, 2019.