



INTERNATIONAL JOURNAL OF ADVANCE RESEARCH, IDEAS AND INNOVATIONS IN TECHNOLOGY

ISSN: 2454-132X

Impact Factor: 6.078

(Volume 12, Issue 2 - V12I2-1266)

Available online at: <https://www.ijariit.com>

A Proof-Carrying Medical Legal Right Advisor where Natural Language Generation is Intentionally Non-Intelligent

Nambarm Athoiba Khuman
singharajdeep59@gmail.com
SRM Institute of Science and
Technology, Tamil Nadu

Naman Kumar Sharma
namankrsharma1275@gmail.com
SRM Institute of Science and
Technology, Tamil Nadu

Soshya Joshi
soshyaj@srmist.edu.in
SRM Institute of Science and
Technology, Tamil Nadu

ABSTRACT

This paper presents PC-MLRA (Proof-Carrying Medical-Legal Rights Advisor), a deterministic system designed to provide structured awareness of patient rights and professional ethics in medical contexts. The system maps user queries to legally grounded responses through a fixed, rule-based pipeline without machine learning or generative language models. Responses are generated using predefined templates linked to a static legal knowledge base derived from the NHRC Charter of Patients' Rights (2019) and the IMC Ethics Regulations (2002). Each output can include a structured proof trace that records the matched intents, activated clauses, and template selection, enabling auditability and reproducibility. The system is evaluated with respect to determinism, legal clause coverage, trace completeness, and safety under ambiguous queries.

Keywords: Medical-Legal Rights Awareness, Deterministic Systems, Proof-Carrying Systems, Rule-Based Classification, Template-Based Generation, Patient Rights, Medical Ethics, Auditability, Safety-Critical AI.

INTRODUCTION

Healthcare interactions often involve statutory patient rights and professional obligations that are not immediately accessible to patients or practitioners. Legal documents such as patient rights charters and medical ethics regulations are formal, distributed across multiple sources, and rarely presented in an interactive or auditable manner. While automated assistants have been used to summarize legal information, many rely on probabilistic or generative techniques that may produce non-reproducible outputs or statements that are difficult to verify against authoritative sources. In medical-legal contexts, such behavior raises concerns regarding traceability, reliability, and safety.

PC-MLRA (Proof-Carrying Medical-Legal Rights Advisor) is designed as a deterministic alternative for patient rights awareness. Instead of employing machine learning or generative language models, the system maps user queries to legally grounded responses through a fixed, rule-based pipeline. Queries are processed using explicit intent indicators, statutory priority rules, and a static legal knowledge base derived from the NHRC Charter of Patients' Rights (2019) and the Indian Medical Council (IMC) Ethics Regulations (2002). Responses are assembled using predefined templates to ensure consistency and bounded output behavior.

A distinguishing feature of the system is its proof-carrying response structure. Each output can include a structured trace that records the normalized query, matched intents, activated legal clauses, and selected response template. This design enables reproducibility, inspection, and post-hoc audit of system decisions without reliance on hidden state or stochastic inference.

The contributions of this work are threefold: (1) a deterministic, rule-based pipeline for surfacing statutory patient rights; (2) a proof-carrying response model that exposes the legal basis of each output; and (3) an evaluation framework centered on determinism, legal completeness, and safety under ambiguity. The remainder of the paper describes the system design, knowledge representation, and evaluation methodology.

RELATED WORK

Patient Rights and Medical Ethics Frameworks

In India, the NHRC Charter (2019) and IMC Regulations (2002) lay out the enforceable rights and ethical standards that govern healthcare. International bodies like the WHO have pushed for rights-based approaches to digital health, but concrete implementation mechanisms remain absent. These frameworks draw ethical boundaries without offering computational methods that could translate abstract legal language into something patients can actually use.

Rule-Based and Deterministic Systems in Healthcare

Machine learning gets most of the attention these days, but rule-based systems still hold their ground in safety-critical clinical applications. The reason is straightforward: transparency and verifiability matter when decisions affect lives. Studies indicate that rule-based NLP methods can actually outperform probabilistic approaches when it comes to pulling safety-critical information from medical records. Deterministic architectures offer verification capabilities that opaque neural models simply cannot match, which explains why they continue to appear in regulated environments.

Legal Question Answering and Informatics

Legal question answering has moved from logic-based systems to neural architectures over the years, but this shift has brought its own set of problems around jurisdictional specificity and accuracy validation. Recently, the field has seen a return to ontology-driven, deterministic designs that produce clearer legal conclusions. These newer frameworks care more about provenance than plausibility, pushing back against the “black box” limitations that plague data-driven models.

Explainability and Transparency

When AI systems operate in legal and medical domains, explainability stops being optional. Post-hoc explanations for opaque models simply do not cut it for high-stakes decisions; interpretability needs to be built in from the ground up. Regulatory standards now treat transparency and traceability as essential ingredients for trustworthy AI. The concept of proof-carrying artifacts, which originated in programming languages research, gives us a theoretical foundation for connecting system outputs to verifiable justifications.

Ethics, Governance, and Safety

Large Language Models produce impressively fluent text, but they also introduce serious risks in clinical settings. Hallucinations and output instability are not minor concerns. Regulatory bodies have started raising alarms about how such variability undermines accountability. Constrained generation approaches offer a way forward by limiting expressive range in exchange for predictability. In this view, constraints are not limitations to work around but safety features worth embracing.

SYSTEM ARCHITECTURE

PC-MLRA employs a strictly compartmentalized topology rather than a monolithic design. The workflow consists of independent logic gates that generate audit trails at every computational junction. Each layer produces immutable artifacts that become mandatory inputs for subsequent layers, isolating decision points and ensuring auditable response lineages.

Architecture Overview

The system comprises four primary layers:

- i. **User Input Layer:** Collects queries and renders responses via web interface
- ii. **Intent Classification Layer:** Performs deterministic intent scoring and statutory priority resolution
- iii. **Knowledge Retrieval Layer:** Maps intents to NHRC/IMC clauses in static knowledge base
- iv. **Template-Based Response Layer:** Generates constrained responses using fixed templates

Component Responsibilities

- i. **Web UI:** Interactive chat interface for query submission and response rendering
- ii. **Flask API:** HTTP endpoint management, request validation, response serialization
- iii. **PCMLRA Core Orchestrator:** Coordinates between web layer and processing pipeline
- iv. **Response Assembler:** Implements end-to-end pipeline including normalization and proof construction
- v. **Intent Classifier:** Rule-based intent detection with statutory priority enforcement
- vi. **Knowledge Base Loader:** Provides access to static NHRC/IMC clause repository
- vii. **Template Engine:** Generates responses using fixed templates with constrained variable substitution
- viii. **Proof Trace:** Captures evidence chain for each response

Deterministic Dataflow

- i. User submits query via Web UI to /api/query end-point
- ii. Flask API validates and dispatches to PCMLRA Core
- iii. Minimal normalization: trim whitespace, collapse multiple spaces
- iv. Intent classification using explicit indicators: keywords (+2), verbs (+1), negative patterns (+3)
- v. Confidence computation = $\min\left(\frac{\text{score}}{6.0}, 1.0\right)$
- vi. Statutory priority resolution for conflicting intents
- vii. Clause lookup from static knowledge base via intent- clause mappings
- viii. Template selection based on resolved intent and context
- ix. Template rendering with constrained variable substitution
- x. ProofTrace construction and serialization

Runtime and Deployment

PC-MLRA is deployed as a Flask-based web application with primary endpoint /api/query. Auxiliary routes include /chat for UI rendering, /api/health for service status, /api/knowledge/search for keyword lookup, and /api/examples for demonstration queries. Configuration flag ACADEMIC_EXPLAINER_ENABLED controls proof trace display. Optional Google Sheets integration provides external logging without affecting core functionality.

Proof Trace Structure

The Proof Trace serves as a machine-readable record of processing decisions. It captures the complete decision path, enabling independent audit and reproduction of responses. The structure includes:

query: Original user input.

- i. **normalized_query:** Query after minimal normalization.
- ii. **matched_intents:** List of intents matched, each with deterministic confidence score.
- iii. **matched_clauses:** Legal clauses activated, including identifier, title, and citation.
- iv. **template_used:** Identifier of the template selected for response generation.
- v. **variables_filled:** Count of variables substituted into the template.

```

{
  "query": "Can I get a copy of my medical records
    ?",
  "normalized_query": "Can I get a copy of my
    medical records?",
  "matched_intents": [
    { "intent": "right_to_records", "confidence":
      "citation": "NHRC Charter (2019)"
    }
  ]
}

```

Listing 1. Example ProofTrace JSON

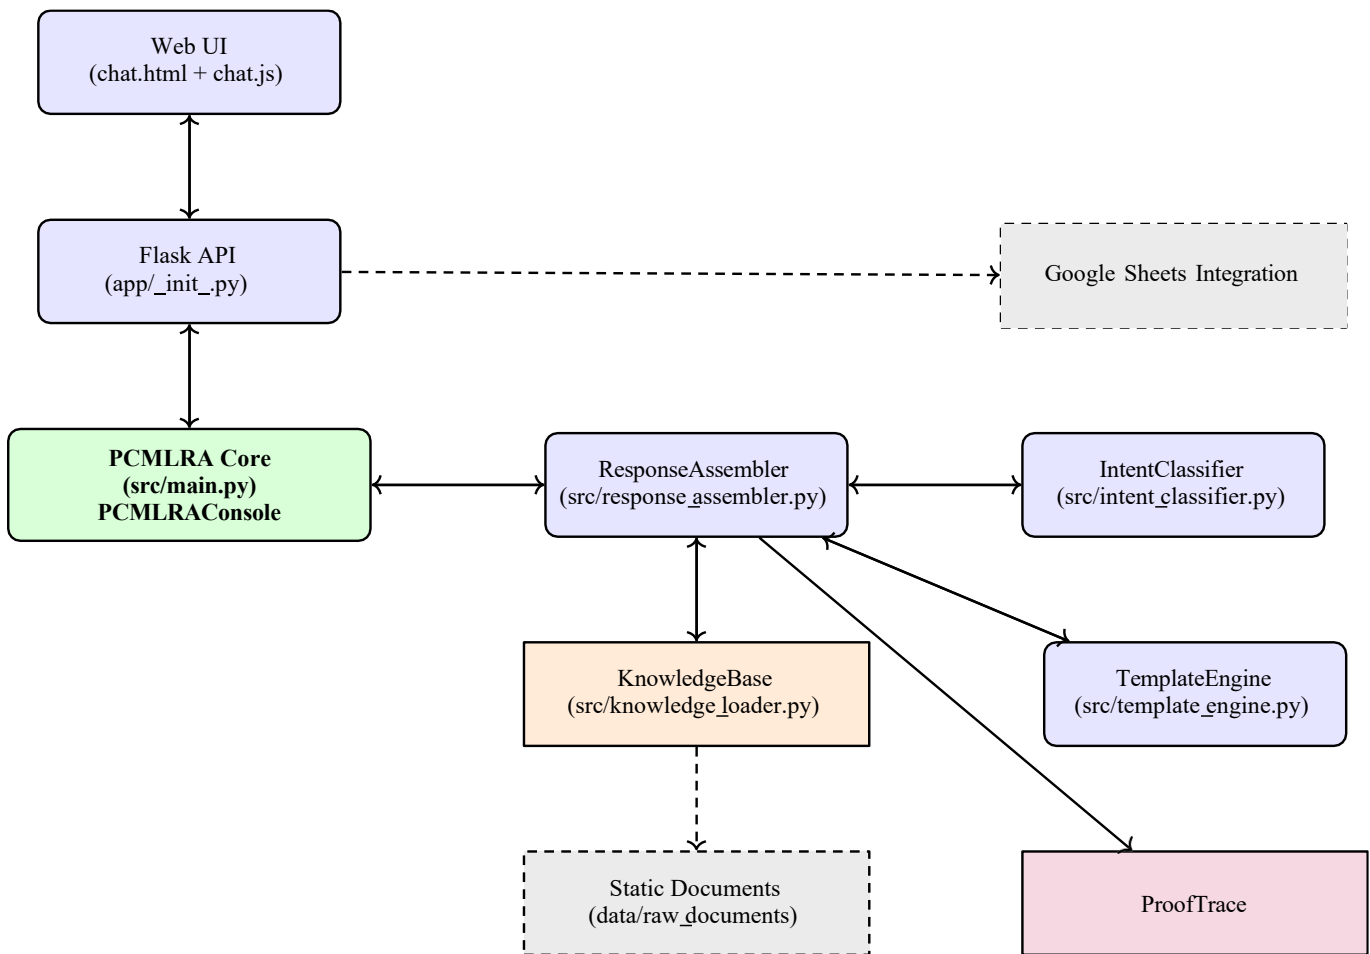


Fig. 1. PC-MLRA deterministic, proof-carrying dataflow. Solid arrows: core pipeline; dashed arrows: optional integrations. The confidence score is computed as $confidence = \min(score/6.0, 1.0)$, where score is the sum of indicator matches (keywords: +2, verbs: +1, negative patterns: +3). This value represents normalized rule activation strength, not a statistical probability.

Security and Safety

Safety is enforced architecturally through deterministic rules, template-based generation, and proof-carrying design. No machine learning models or generative services are used. Session context is handled in-memory, with optional external logging separated from response logic to limit data persistence.

METHODOLOGY

Design Rationale: Determinism and Auditability

PC-MLRA was built with determinism as a core principle. The system ensures that the same input always produces the same output, regardless of when or where it runs. This matters for several reasons. It makes systematic auditing possible. It allows meaningful comparison across different cases. And it means every response can be traced back to specific decision rules and legal sources. Safety here is not left to statistical filtering; it is built into the architecture through constrained behavior.

Input Handling and Scope

The system takes unstructured free-text queries as they come. It does not ask for clarification, reformulate what the user typed, or engage in multi-turn conversations. Each query stands alone and moves through a single-pass pipeline. There is no attempt at semantic understanding, paraphrasing, or linguistic enrichment. Queries are treated as strings to be matched against legal indicators, nothing more.

Minimal Normalization

Normalization is kept deliberately minimal. The system strips leading and trailing whitespace and collapses multiple consecutive spaces into one. That is it. No lowercasing, no grammar correction, no spelling normalization, no deduplication. The idea is to preserve the original query structure while ensuring consistent token spacing for matching. Anything more would risk introducing assumptions that the system was not designed to make.

Deterministic Intent Classification

Intent classification uses rule-based scoring with fixed indicators:

- i. Keywords: +2 per match
- ii. Verbs: +1 per match
- iii. Negative patterns: +3 per match

Confidence is computed as confidence = $\min \left(\frac{\text{score}}{6.0}, 1.0 \right)$ representing normalized rule activation strength, not probability.

Statutory Priority Resolution

When multiple intents match, conflicts are resolved via explicit statutory priority list reflecting legal hierarchy. Emergency care intents precede autonomy-related ones; research protections override general patient rights. This ensures legally dominant rights surface consistently regardless of raw indicator scores.

Clause Mapping to Legal Sources

Resolved intents map to legal clauses via static knowledge base containing NHRC Charter (2019) and IMC Regulations (2002). Clause activation is binary (matched/not matched) without semantic expansion or interpretation. Each activated clause includes identifier, title, and citation for proof trace exposure.

Template Selection and Variable Binding

Responses use fixed, predefined templates selected deterministically based on intent and context. Variable substitution is strictly constrained to values derived from query or assembled context. No inference of missing values, quantity approximation, or context fabrication occurs.

IMC Ethics Awareness Handling

IMC provisions are informational only. When professional conduct signals are detected, relevant IMC clauses append to responses as advisory context without influencing intent classification, priority resolution, or NHRC rights enforcement.

Proof Trace Construction

For each response, ProofTrace captures: raw query, normalized query, matched intents with confidence, matched clauses, selected template, and variables filled. This enables independent inspection, reproduction, and audit without accessing hidden state or external models.

KNOWLEDGE BASE DESIGN

Design Objectives

The knowledge base is static, immutable, and fully deterministic, supporting auditability, reproducibility, and hallucination resistance. All legal knowledge resides in fixed JSON artifact knowledge_base_complete.json, eliminating runtime variability and undocumented content.

Source Documents

Two authoritative instruments are represented:

- i. NHRC Charter of Patients' Rights (2019): Enforceable patient rights
- ii. IMC Ethics Regulations (2002): Professional conduct and ethics standards

Explicit separation prevents conflation during processing.

Clause Schema

Each clause includes: unique identifier, document source, section identifiers, title, statutory text, paraphrase, rights/obligations fields, actors, exceptions, category labels, keywords, intent matches, template identifiers, and citation formats. Supplementary metadata captures timeframes or response obligations when defined in source text.

Intent-Clause Binding

Intent detection and clause storage are separate layers. Intents are detected upstream via rule-based mechanisms; knowledge base provides static clause definitions with explicit intent identifiers. Clause activation is transparent, declarative operation driven solely by JSON-encoded mappings.

NHRC-IMC Separation Policy

Strict separation ensures NHRC clauses determine rights while IMC clauses provide contextual awareness. IMC provisions do not override NHRC rights, influence statutory priority, or affect intent selection. Ethical guidance supplements rather than alters legal rights determination.

Clause Activation Semantics

Activation is binary based on explicit intent-clause mappings. No graded relevance scoring, semantic expansion, similarity matching, or interpretive reasoning occurs. This ensures predictability and prevents unintended legal interpretations.

Safety and Immutability Guarantees

Static JSON knowledge base ensures reproducibility across executions. Traceability is maintained via proof trace references to clause identifiers and citations. Output restriction to pre-approved statutory text eliminates hallucination risk.

EVALUATION

Evaluation Objectives

Evaluation verifies design compliance: (1) rule-based intent detection, (2) complete deterministic clause surfacing, (3) correct ethics awareness without statutory influence, (4) complete proof traces, (5) identical inputs producing identical outputs, (6) hallucination resistance with vague queries.

Intent Detection Metrics

Using information retrieval metrics without probabilistic interpretation:

$$\text{Precision} = \frac{TP}{TP + FP} \quad (1)$$

$$\text{Recall} = \frac{TP}{TP + FN} \quad (2)$$

$$\text{F1-score} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \quad (3)$$

where TP=True Positives, FP=False Positives, FN=False Negatives. Top-k criterion evaluates detection correctness.

Clause Coverage Accuracy

Measures legal completeness: Coverage = $\frac{|C_e \cap C_d|}{|C_d|}$

where Ce=expected clauses, Cd=detected clauses. Binary matching; partial matches not considered.

Ethics Awareness Evaluation

Independent from statutory rights: Ethics Accuracy = $\frac{|E_e \cap E_d|}{|E_d|}$ where E_e =expected IMC clauses, E_d =appended IMC clauses. Verifies correct attachment without decision-making influence.

Proof Trace Completeness

Verifies auditability: Proof Completeness = $\frac{|F_p|}{|F|}$ where F=query, normalized_query, matched_intents, matched_clauses, template_used, F_p =presented fields.

Determinism Verification

Core invariant: $O1(q) = O2(q) = \dots = On(q)$ for query q. Any deviation in intents, clauses, templates, or proof traces constitutes violation.

Safety Under Ambiguity

Hallucination resistance: Safety Score = $1 - \frac{\sum H(q)}{|Q_a|}$ where Q_a =ambiguous queries, $H(q) = 1$ indicates hallucination (invented facts, inferred conditions, speculative advice, unstated provisions). Conservative responses not considered failures.

Confidence Score Interpretation

Deterministic, non-probabilistic: raw score S from key- word (+2), verb (+1), negative pattern (+3) matches. Confidence = $\min_{-6.0}^S 1.0$ representing normalized rule activation strength.

Evaluation Scope and Limitations

Testing used small manually curated dataset due to tooling constraints. Focus on correctness, legal completeness, determinism, and safety rather than statistical generalization. No population-level performance claims.

EXPERIMENTAL RESULTS

PC-MLRA is not evaluated via conventional experimental paradigms (large datasets, statistical benchmarks) due to its deterministic nature. Instead, results refer to observed behavioral invariants across executions.

Observed Behavioral Properties

- i. **Deterministic Output:** Identical queries produce identical intents, clause mappings, templates, and proof traces
- ii. **Complete Legal Traceability:** Every response exposes legal basis via explicit clause identifiers and citations
- iii. **Safe Handling of Ambiguity:** No invention of facts, medical conditions, or legal claims for vague queries
- iv. **Separation of Rights and Ethics:** IMC references append as informational context without altering statutory rights

Illustrative Case Analysis

Qualitative inspection confirms conservative, transparent behavior. Example: queries with billing dispute and emergency care indicators prioritize emergency-related rights via statutory priority rules, regardless of raw indicator counts. Professional misconduct indicators without patient-rights violations trigger ethics awareness without enforceable legal claims.

DISCUSSION

PC-MLRA occupies a design space distinct from predictive or generative assistants commonly used in medical and legal domains. Rather than optimizing for linguistic flexibility or adaptive behavior, the system prioritizes determinism, trace-ability, and legal grounding. This trade-off reflects the constraints of medical-legal contexts, where reproducibility and verifiability are often more critical than expressive range.

The proof-carrying design enables inspection of every response through explicit mappings between query, intent, clause, and template. This structure allows outputs to be audited without relying on post-hoc explanation techniques. The strict separation between statutory rights (NHRC) and professional ethics (IMC) further ensures that advisory context does not override enforceable legal provisions.

At the same time, the deterministic architecture introduces limitations. The system does not generalize beyond predefined indicators, interpret novel legal scenarios, or engage in multiturn clarification. Queries outside the encoded rule space result in conservative responses. This constraint is intentional: in regulated environments, bounded and transparent behavior may be preferable to adaptive but opaque reasoning.

Overall, the results suggest that rule-based, proof-carrying systems remain viable for narrowly scoped legal awareness applications, particularly where accountability and reproducibility are primary requirements.

LIMITATIONS

- i. **Scope:** Legal rights awareness system, not legal advice, clinical recommendations, or decision support
- ii. **Semantic Interpretation:** Avoids paraphrase detection, contextual inference, linguistic generalization
- iii. **Static Knowledge Base:** Manual updates required for statutory amendments or new interpretations
- iv. **Jurisdictional Restriction:** Limited to Indian NHRC/IMC frameworks; not generalizable to other jurisdictions
- v. **Ethics Handling:** IMC provisions informational only; cannot override or modify rights-based outputs
- vi. **Conversational Limitation:** Single-turn processing without dialogue, context retention, or clarification
- vii. **Coverage Gap:** Excludes NCCE Charter (2021) institutional responsibilities and operational rights
- viii. **Evaluation Scale:** Small curated dataset; design verification rather than empirical performance

CONCLUSION

PC-MLRA is a deterministic, rule-based system built to provide structured awareness of patient rights and professional ethics in medical contexts. It maps unstructured queries to responses rooted in statutory sources by combining explicit intent indicators, priority rules, a static legal knowledge base, and template-driven generation. By sidestepping machine learning, probabilistic inference, and generative models, the system ensures that the same query always produces the same result. This approach builds in auditability, reproducibility, and the ability to inspect exactly how any response was reached. Anyone can look back and see precisely how a particular answer came together. This design makes auditability, reproducibility, and behavioural inspection possible. Rather than treating ethical guidance as something that could override statutory rights, the system maintains a clear separation—NHRC provisions determine rights while IMC material provides supplementary context.

What this work shows is that transparency, legal safety, and deterministic behavior do not have to take a back seat to expressive or adaptive capabilities. For regulated domains where accountability is non-negotiable, that trade-off is not a limitation but a deliberate and necessary choice.

FUTURE WORK

- i. Scale knowledge base with strict versioning protocols while maintaining auditability
- ii. Develop tooling for systematic knowledge base maintenance and traceable updates
- iii. Enhance proof trace visualization for improved explain- ability
- iv. Implement structured multilingual support via parallel static templates and translations
- v. Explore offline/institutional deployments for controlled environments
- vi. Formal verification of rule coverage and expanded design- verification datasets

ACKNOWLEDGMENT

Thanks to academic mentors for methodological rigor guidance, and to NHRC/IMC for public legal documents enabling this work. Acknowledgment of open-source software ecosystem supporting implementation.

REFERENCES

- [1] National Human Rights Commission (India), *Charter of Patients' Rights*. New Delhi, India: Ministry of Health and Family Welfare, Government of India, 2019.
- [2] National Council for Clinical Establishments, *Charter of Patients' Rights and Responsibilities*. Ministry of Health and Family Welfare, Government of India, 2021.
- [3] Medical Council of India, *Indian Medical Council (Professional Conduct, Etiquette and Ethics) Regulations, 2002*. Gazette of India, Extraordinary, Part III, Sec. 4, New Delhi, India, 2002.
- [4] World Health Organization, *Ethics and Governance of Artificial Intelligence for Health*. WHO Guidelines, 2021.
- [5] World Health Organization, *Patient Safety Rights Charter*. WHO Publications, 2024.
- [6] High-Level Expert Group on Artificial Intelligence, *Ethics Guidelines for Trustworthy AI*. European Commission, 2019.
- [7] U.S. Food and Drug Administration, *AI/ML-Based Software as a Medical Device Action Plan*. FDA Guidance Document, 2021.
- [8] E. Tabassi, *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. National Institute of Standards and Technology (NIST), 2023.
- [9] 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.
- [10] A. F. T. Winfield, "IEEE P7001: A proposed standard on transparency," *Frontiers in Robotics and AI*, vol. 8, 2021.
- [11] A. Kiseleva, D. Kotzinos, and P. De Hert, "Transparency of AI in healthcare as a multilayered system of accountabilities," *Frontiers in Artificial Intelligence*, vol. 5, 2022.
- [12] Y. Ning, M. Liu, and N. Liu, "Advancing ethical AI in healthcare through interpretability," *Patterns*, 2025.
- [13] E. Hildt, "What is the role of explainability in medical AI? A case-based approach," *Bioengineering*, 2025.
- [14] A. Abdallah, R. Piryani, and A. Jatowt, "Exploring the state of the art in legal question answering systems," *Journal of Big Data*, 2023.
- [15] J. Collenette, K. Atkinson, and T. Bench-Capon, "Explainable AI tools for legal reasoning," *Artificial Intelligence*, 2023.
- [16] H. de Martim, "Deterministic legal agents: Canonical primitives for auditable legal reasoning," arXiv preprint, 2025.
- [17] H. de Martim, "Ontology-driven deterministic reasoning over legal norms," *Artificial Intelligence and Law*, vol. 34, no. 1, pp. 45–68, 2026.
- [18] G. C. Necula, "Proof-carrying code," in *Proc. ACM SIGPLAN–SIGACT Symp. Principles of Programming Languages*, 1997, pp. 106–119.
- [19] X. Qian, "Rule-based natural language processing in clinical safety analysis," *JMIR Aging*, 2025.
- [20] D. Zakim, "Computerized history-taking improves data quality for clinical decision-making," *PLOS ONE*, 2021.
- [21] E. C. Rijcken, "Machine learning versus rule-based methods in health- care text classification," *Natural Language Processing*, 2025.
- [22] A. Alnattah, "Rule-based systems in clinical decision-making: A scoping review," *Healthcare*, 2024.
- [23] M. Lyu, "Natural language generation in healthcare: A review," arXiv preprint, 2025.
- [24] Y. Hua, "From rule-based systems to large language models in mental health chatbots," *World Psychiatry*, 2025.
- [25] Q. Du, "Rule-based versus large language model-based healthcare chatbots," *Journal of Medical Internet Research*, 2025.
- [26] B. Mesko' and E. J. Topol, "The imperative for regulatory oversight of generative AI in healthcare," *npj Digital Medicine*, 2023.