

Ensuring Regulatory Compliance of Algorithmic Decisions Through Structured Explanation Trace Schemas

Kaleshwar Aryasomayajula¹

¹Sr Specialist Software Developer, Charles Schwab, San Tan Valley, Arizona, USA

Abstract—This article examines how structured explanation trace schemas can support regulatory compliance in algorithmic decision-making. The problem is acute in high-stakes domains where automated outputs affect access to credit, employment, services, or public benefits. At the same time, the evidentiary basis of each decision remains fragmented across logs, model reports, and policy files. The study aims to determine which documentation elements must be captured at the decision level, how they should be organized, and how they can be translated into an implementable compliance architecture. The materials consist of ten sources published between 2021 and 2025 on AI documentation, auditability, machine-readable compliance artifacts, and traceability infrastructures. Comparative analysis, conceptual synthesis, and typologization were applied. The analytical section establishes that compliance-ready explainability depends on linking rationale, provenance, model state, policy version, and human intervention records within a structured trace. The practical value lies in the proposed implementation model, the monitoring metrics, and the documentation logic for regulated algorithmic environments.

Keywords—algorithmic decision-making, regulatory compliance, explanation trace schema, AI documentation, auditability, traceability, model governance, machine-readable documentation, algorithmic accountability, explainable AI.

I. INTRODUCTION

Algorithmic decisions are increasingly used in contexts where mistaken outputs have immediate legal, financial, or social consequences. Yet compliance work in such settings is still often built from disconnected artifacts. A model card may describe the intended use, a dataset document may record provenance, a monitoring dashboard may contain performance indicators, and an application log may preserve timestamps. When a decision is challenged, these materials rarely form a coherent evidentiary chain. The operational weakness arises from the absence of a structured record that binds explanation to the exact decision event.

This study aims to develop an analytical model of a structured explanation trace schema as a compliance instrument for algorithmic decisions. Three objectives organize the inquiry. The first is to identify why current documentation approaches leave a gap at the moment of decision. The second is to determine which elements a compliance-oriented trace schema must contain. The third is to propose an implementation and monitoring logic suitable for regulated decision environments.

The novelty of the article lies in treating explanation trace schemas as an intermediate compliance layer. In that layer, upstream documentation on data, models, and system design is linked to downstream functions such as audit, dispute handling, internal control, and regulatory review.

II. MATERIALS AND METHODS

The source base comprises 10 publications issued between 2021 and 2025. The corpus was assembled through targeted screening of recent literature on AI documentation, auditability, machine-readable compliance instruments, and traceability infrastructures. The selection included foundational work on dataset documentation [3], comparative studies of

documentation practices and principles [1; 2; 6], frameworks for machine-readable risk and lifecycle documentation [4; 7; 8], and studies that move from documentation toward auditable trace architectures [5; 9; 10]. The logic of inclusion followed the article's research design. Each source had to contribute either to the diagnosis of documentation insufficiency, to the specification of structured fields for compliance, or to the design of an operational trace infrastructure.

The study combines comparative analysis, source analysis, conceptual synthesis, typologization, and analytical generalization. Comparative reading was used to isolate convergent and divergent positions across documentation frameworks. Conceptual synthesis was applied to derive a unified definition of the explanation trace schema. Typologization was used to separate upstream artifacts, decision-level records, and post-decision audit instruments. Analytical generalization supported the transition from literature-based findings to an implementable compliance model.

III. RESULTS

Recent literature no longer treats AI documentation as a single artifact. It describes a diversified documentation ecology whose units differ by scope, audience, and level of automation. A systematic review of documentation practices reports that data and model cards still dominate the field, while system cards remain less common, and documentation production remains largely manual [1]. A large-scale study of 32,111 model cards reaches a parallel conclusion from a different angle. Model cards are widely present for heavily downloaded models, yet their informativeness is uneven, with limitations and evaluation among the weakest sections [6]. Read together, these studies expose a structural compliance problem. The field has produced many descriptive documents, but their distribution across lifecycle stages remains uneven, and the

most compliance-sensitive content is frequently incomplete. In such conditions, explanation at the level of a single decision becomes fragile because the necessary evidence is dispersed before the decision is even made.

The answer proposed by documentation theory is to make documentation more comparable, legible, actionable, and robust [2]. That move matters because compliance work depends on consistency across cases. A record that differs radically from one system, team, or release to another becomes difficult to audit, compare, and test against internal policy. At the same time, the comparative literature warns against assuming that a single static template is sufficient for every situation [2]. A regulated decision environment requires standardized fields, yet it also requires sensitivity to intended use, stakeholders, and the lifecycle stage at which evidence is generated [1; 2]. The literature, therefore, pushes in two directions at once. It favors standardization but resists flattening the entire compliance problem into a single generic form.

Foundational documentation work on datasets already points toward the logic needed for a trace schema. Datasheets were designed to record motivation, composition, collection procedures, intended uses, and maintenance conditions for datasets [3]. That contribution remains indispensable because no decision record can be reliable if it cannot point back to the data conditions from which the model emerged. Still, dataset documentation stays upstream. It does not by itself tell an auditor which model instance was used, which policy threshold was active, which features were decisive, or whether a human reviewer intervened in a particular case. More recent work closes that distance. AI Cards propose a framework that combines technical specifications, context of use, and risk management in both human-readable and machine-readable form [4]. The move toward machine readability is decisive. Once documentation is represented in a structured form, it becomes possible to exchange it across the AI value chain, update it when conditions change, and use it in conformity assessment tasks [4].

Automation research reinforces the same direction. The literature review on documentation practices notes that full automation remains rare [1]. That observation explains why many organizations still treat compliance documentation as an *ex post* exercise. By contrast, recent work on automatic generation of models and data cards shows that documentation pipelines can improve completeness, objectivity, and faithfulness when generation is grounded in structured retrieval workflows [7]. The significance of the present study is direct. A compliance-oriented explanation trace schema should be generated from structured sources already present in the system, or from registries designed to be queried at the moment of decision.

A clearer picture emerges when one compares the functions assigned to documentation across several recent frameworks. The CLeAR framework supplies principles for documentation quality [2]. AI Cards introduce a machine-readable structure for intended use and risk management [4]. TechOps extends the logic across data, model, and application levels and explicitly links documentation to lifecycle oversight and auditable operational practice [8]. The review of documentation practices

adds evidence that the field is moving toward more holistic and automated forms, even though manual effort remains dominant [1]. Their convergence is substantial. Compliance-oriented documentation must cover more than technical performance. It has to register who the system is for, what risks are anticipated, which component produced the output, which constraints govern use, and how the record remains usable across time. Their divergence is instructive as well. Principle-based frameworks clarify what good documentation should look like. Machine-readable frameworks clarify how documentation can circulate and be checked. Lifecycle templates clarify where evidence should be captured. None of these alone fully resolves decision-level accountability. Together, they indicate that a separate decision-event layer is needed.

On that basis, the explanation trace schema can be defined more narrowly and more rigorously. It is a structured record attached to a concrete algorithmic output and linked to the surrounding documentation environment. At minimum, it should contain a unique decision identifier, timestamp, system identifier, model, model version reference, data or feature provenance reference, decision outcome, confidence or score (where applicable), explanation payload, policy threshold or rule applied, human intervention status, documentation links, and post-decision state. Recent traceability work clarifies why these fields must be grouped. The AI Model Passport treats traceability as a lifecycle property and frames model identity, verification, versioning, and monitoring as a unified documentation problem [5]. BAXDT goes further by placing explanation summaries, contextual metadata, and cryptographic integrity within a single decision-trace object [9]. Work on knowledge-graph-based auditing reaches the same conclusion via a different route, arguing that audit traces need semantic structure to be integrated, queried, and interpreted across complex systems [10]. The literature supports a strong claim: compliance rests on explanation joined to provenance, state, and governance metadata in a form that remains inspectable.

Figure 1 synthesizes the position of the structured explanation trace schema within the broader compliance chain. It is the runtime node that binds those upstream materials to the concrete decision event and to downstream review activities.

The architectural literature adds one more step that is often underappreciated in debates about explainability. A trace schema must remain durable, queryable, and resistant to silent alteration. The AI Model Passport responds to that need through unique identity, version control, and lifecycle monitoring [5]. BAXDT responds via a unified JSON trace, with its cryptographic hash separately preserved, thereby strengthening evidentiary integrity [9]. Knowledge-graph-based auditing makes relations between artifacts explicit so that auditors can follow connections among models, datasets, system components, and control requirements [10]. These approaches differ in technical realization, yet they converge on the same operational lesson. A compliance-ready explanation trace schema is a structured evidence object.

A second comparative lesson follows from the relationship between static documentation and dynamic compliance. Datasheets, model cards, and system-level templates remain

indispensable because they capture upstream conditions [1; 2; 3; 6]. Machine-readable documentation and lifecycle templates extend their utility by enabling interoperability, updating, and automated checking [4; 7; 8]. Traceability frameworks then translate those materials into decision-level evidence [5; 9; 10]. This layered logic resolves a recurrent tension in the literature. Static cards are too broad to answer a concrete dispute on their own. Event logs are too narrow if they are detached from the system's meaning. The explanation trace schema closes that gap by turning a single decision into a traceable node within a wider compliance network.

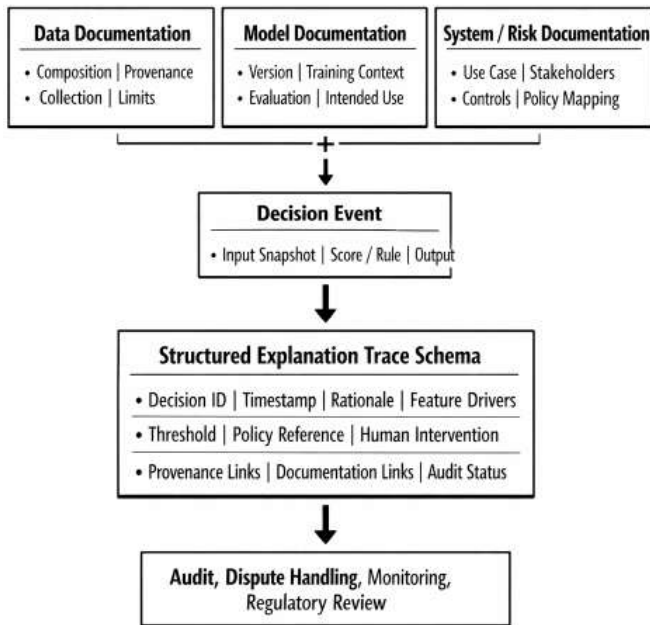


Fig. 1. Lifecycle position of a structured explanation trace schema in compliance-oriented algorithmic decisioning, adapted from [2; 4; 8]

The analytical synthesis of the reviewed literature therefore supports three findings. First, current documentation approaches improve transparency, yet they often fail to preserve a reconstructable chain at the level of individual decisions [1; 2; 6]. Second, a compliance-oriented trace schema must combine rationale with provenance, system state, governance references, and intervention records [4; 5; 8; 9; 10]. Third, the schema becomes operationally reliable only when it is machine-readable, connected to upstream documentation, and supported by retrieval or registry mechanisms that reduce manual completion and post hoc reconstruction.

IV. DISCUSSION

The reviewed literature suggests that regulatory compliance in algorithmic decision-making is best understood as a record-formation problem. Explanations fail in practice when they are produced late, detached from system state, or written in forms that auditors cannot verify. For that reason, the explanation trace schema should be implemented as a standing infrastructure for decision evidence. Organizations working in high-stakes environments should separate three questions from the start: what must be documented before deployment, what must be captured at decision time, and what must remain

measurable after deployment. Once these layers are separated, compliance design becomes tractable.

The first Table I compares the organizational layers that a structured explanation trace schema has to connect. Its purpose is to distinguish stable documentation registers from event-level traces and from repositories that enable later review.

TABLE I. Compliance architecture of a structured explanation trace schema [1–10]

Layer	Primary purpose	Core contents	Main owner	Update rhythm
Data documentation layer	Preserve provenance and data-use constraints	Dataset identity, source, collection logic, preprocessing, limitations, permitted use	Data governance team	On dataset change
Model documentation layer	Preserve model state and evaluation history	Model ID, version, training context, validation metrics, known limits	ML engineering team	On model release or retraining
System and policy layer	Bind technical behavior to organizational rules	Intended use, risk controls, human oversight rules, decision thresholds, escalation logic	Product, compliance, legal	On policy or workflow change
Decision trace layer	Capture the individual algorithmic event	Decision ID, timestamp, input abstraction, output, explanation payload, threshold applied, reviewer action	Runtime service with governance controls	At each decision
Evidence storage layer	Keep records durable and queryable	Immutable store, hash or integrity marker, retention label, access rights	Platform and security teams	Continuous
Review and recourse layer	Support internal and external scrutiny	Audit queries, case reconstruction, dispute packets, override review, exception history	Compliance, audit, operations	Periodic and case-triggered

The comparison shows that the explanation trace schema occupies a narrow but decisive place. It translates them into case-bound evidence. If the decision trace layer is weak, the surrounding layers remain informative yet legally fragile. If the decision-trace layer is strong but isolated, the record becomes technically precise but institutionally thin. Compliance depends on linkage.

The second table addresses a different problem. A trace schema can exist formally and still fail in practice if its fields are missing, stale, slow to retrieve, or inconsistent with current policy. Table II sets out a monitoring model for compliance readiness.

The metrics clarify a broader point. Compliance cannot be assessed solely through the formal existence of templates. It has to be measured through the system's ability to reproduce and justify a decision under scrutiny. In that sense, the most revealing indicators are operational properties of the evidence-formation process. Completeness without policy alignment is weak. Fast retrieval without provenance linkage is weak. Integrity without reviewer capture is weak. The monitoring model works because it treats compliance as a chain whose weakest link can be identified.

TABLE II. Monitoring metrics for explanation trace quality and compliance readiness [1–10]

Metric	What it measures	Practical interpretation	Management use
Trace completeness rate	Share of decisions with all mandatory fields populated	Low values indicate broken capture logic or missing integrations	Repair field mapping and block incomplete releases
Provenance linkage rate	Share of decisions linked to valid data and model identifiers	Weak linkage makes reconstruction unreliable	Tighten registry integration
Policy-version match rate	Share of decisions tied to the active policy release	Mismatch signals governance drift	Review rollout controls
Human intervention capture rate	Share of overrides or reviews recorded in structured form	Missing entries hide accountability at the most sensitive points	Redesign reviewer workflow
Explanation retrieval latency	Time needed to assemble a review-ready trace	High latency undermines dispute handling and audit response	Improve storage and indexing
Reproducibility rate	Share of sampled cases that can be reconstructed without manual inference	Low rates indicate evidentiary weakness	Expand captured runtime state
Integrity verification rate	Share of records that pass hash or immutability checks	Failures weaken legal defensibility	Strengthen integrity controls
Exception recurrence rate	Frequency of repeated schema failures by system segment	Repetition indicates structural defects, not isolated errors	Trigger targeted remediation

Implementation should proceed in a defined sequence. An organization first needs a decision taxonomy that distinguishes automated, assisted, and human-overridden decisions. It then needs a canonical field dictionary that specifies which elements are mandatory for every case and which are conditional on domain, model type, or legal regime. The next step is to connect those fields to the source systems so that data provenance, model version, policy version, and runtime output are automatically pulled. After that, explanation payloads should be normalized into a limited set of representational forms such as ranked factors, rule references, or structured reason objects. Human review actions have to be recorded in the same schema. Retention, access control, and integrity protection come after schema stabilization.

In practical terms, the explanation trace schema works best when it is treated as a living compliance object. It should be reviewed during model updates, policy changes, and incident analysis. The internal audit should sample it. It should be used

in training for reviewers and case handlers. Most of all, it should be designed for retrieval by people who were not present when the decision was made. That requirement changes the trace's logic. The schema has to privilege field clarity, evidentiary sufficiency, and cross-functional readability.

V. CONCLUSION

The review showed that current documentation practices improve transparency across datasets, models, and systems, yet they frequently leave decision-specific evidence fragmented and unevenly captured.

The analysis established that a defensible schema must combine rationale, provenance, model state, policy reference, and human intervention records into a single structured decision object.

The discussion developed a layered architecture and a metric-based monitoring model through which explanation traces can function as operational evidence for audit, dispute handling, and regulatory review.

REFERENCES

- [1] Arnold, S., Yesilbas, D., Gröbner, R., Riedelbauch, D., Horn, M., & Weinzierl, S. (2024). Documentation practices of artificial intelligence. *arXiv*. <https://arxiv.org/abs/2406.18620>
- [2] Chmielinski, K., Newman, S., Kranzinger, C. N., Hind, M., Vaughan, J. W., Mitchell, M., Stoyanovich, J., McMillan-Major, A., McReynolds, E., Gray, M. L., Hudson, M., Chang, A., & Esfahany, K. (2024). *The CLEAR documentation framework for AI transparency*. Shorenstein Center on Media, Politics and Public Policy. https://shorensteincenter.org/wp-content/uploads/2024/05/CleAR_KChmielinski_FINAL.pdf
- [3] Gebru, T., Morgenstern, J., Vecchione, B., Vaughan, J. W., Wallach, H., Daumé III, H., & Crawford, K. (2021). Datasheets for datasets. *Communications of the ACM*, 64(12), 86–92. <https://doi.org/10.1145/3458723>
- [4] Golpayegani, D., Hupont, I., Panigutti, C., Pandit, H. J., Schade, S., O'Sullivan, D., & Lewis, D. (2024). AI cards: Towards an applied framework for machine-readable AI and risk documentation inspired by the EU AI Act. In *Annual Privacy Forum* (pp. 48–72). Springer Nature Switzerland.
- [5] Kalokyri, V., Tachos, N. S., Kalantzopoulos, C. N., Sfakianakis, S., Kondylakis, H., Zaridis, D. I., Colantonio, S., Regge, D., Papanikolaou, N., ProCancer-I Consortium, Marias, K., Fotiadis, D. I., & Tsiknakis, M. (2025). AI model passport: Data and system traceability framework for transparent AI in health. *Computational and Structural Biotechnology Journal*, 28, 386–404. <https://doi.org/10.1016/j.csbj.2025.09.041>
- [6] Liang, W., Rajani, N., Yang, X., Ozoani, E., Wu, E., Chen, Y., et al. (2024). What's documented in AI? Systematic analysis of 32K AI model cards. *arXiv*. <https://arxiv.org/abs/2402.05160>
- [7] Liu, J., Li, W., Jin, Z., & Diab, M. (2024). Automatic generation of model and data cards: A step towards responsible AI. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)* (pp. 1975–1997).
- [8] Lucaj, L., Loosley, A., Jonsson, H., Gasser, U., & van der Smagt, P. (2025). TechOps: Technical documentation templates for the AI Act. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, 8(2), 1647–1660.
- [9] Parlak, İ. E. (2025). Blockchain-assisted explainable decision traces (BAXDT): An approach for transparency and accountability in artificial intelligence systems. *Knowledge-Based Systems*, 329(Part B), 114402. <https://doi.org/10.1016/j.knosys.2025.114402>
- [10] Waltersdorfer, L., & Sabou, M. (2025). Leveraging knowledge graphs for AI system auditing and transparency. *Journal of Web Semantics*, 84, 100849. <https://doi.org/10.1016/j.websem.2024.100849>