FMDB *Transactions on Sustainable Computer Letters*

A Novel Context-Aware Framework for Semantic-Preserving Data Transformation and Reliable Decision Intelligence in AI-Driven Environments

Aditya Rautaray[1,*], Madhan Raj Gopi Akila[2], G. Mary Amirtha Sagayee[3]

[1]Departement of Cloud Solutions Security, CVS Healthcare, Ashburn, Virginia, United States of America.
[2]Departement of Computer Science, University of Stuttgart, Stuttgart, Baden-Württemberg, Germany.
[3]Departement of Computing and Information Sciences, University of Technology and Applied Sciences, Nizwa, Muscat, Oman.
aditya.rautaray@cvshealth.com[1], st199709@stud.uni-stuttgart.de[2], mary.sagayee@utas.edu.om[3]

Abstract: Today, more and more heterogeneous information environments link up records that have different terminology, origin, information time, and uses in different organizations. Conventional transformation pipelines usually aim for syntactic compatibility but do not ensure that the meaning that existed upstream of the transformations is preserved. This paper proposes a context-aware architecture that profiles source conditions, maintains semantic relationships, governs the transformation process and provides traceable inputs to the decision services. An illustrative dataset of 140 synthetic records, sensor observations, and narrative documents is established across structured records, sensor observations, and narrative documents; the dataset is not an observed deployment of the framework, but is built to demonstrate it. This prototype workflow is modeled in Python, pandas, RDF/OWL, NetworkX, and scikit-learn, and the expected behavior in the evaluation is communicated through spreadsheet calculations and comparative displays. The four transformation strategies are compared on context alignment, semantic integrity and decision reliability; the process-stage assessment is independent of confidence impact and considers validation and governance. The demonstration shows that coordinated contextual/semantic control can outperform isolated format conversion. Because the numerical results are illustrative, they reflect architectural plausibility rather than validated, empirical superiority. A reproducible evaluation pathway and operational limitations, evidence requirements, constraints on governance, and mechanisms for interpretability for future tests with real data are detailed in the paper.

Keywords: Context Awareness; Semantic Preservation; Data Transformation; Decision Intelligence; Information Governance; AI-Driven Environments; Data Quality; Adaptive Data Transformation.

Received on: 11/08/2025, **Revised on:** 06/11/2025, **Accepted on:** 21/12/2025, **Published on:** 11/11/2026

Journal Homepage: https://www.fmdbpub.com/user/journals/details/FTSCL

DOI: https://doi.org/10.69888/FTSCL.2026.000840

Cite as: A. Rautaray, M. R. G. Akila, and G. M. A. Sagayee, "A Novel Context-Aware Framework for Semantic-Preserving Data Transformation and Reliable Decision Intelligence in AI-Driven Environments," *FMDB Transactions on Sustainable Computer Letters*, vol. 4, no. 4, pp. 194–203, 2026.

Copyright © 2026 A. Rautaray *et al.,* licensed to Fernando Martins De Bulhão (FMDB) Publishing Company. This is an open access article distributed under CC BY-NC-SA 4.0, which permits copying, remixing, redistributing, transformation, and building upon the material in any medium so long as the original work is properly cited.

1. Introduction

Data changes are not typically benign. As observations enter automated decision environments, their representation may change much faster than the intent of the observation can be recorded. What may be a meaningful field in an operational process may

[*]Corresponding author.

not be meaningful when converted to another schema, especially when units, collection purposes, time and institutional terminology differ. The issue becomes more prevalent as organizations integrate disparate, independently created applications and provide access to their information via the cloud, connected devices, and machine-assisted workflows. Data quality therefore depends not only on whether data is provided, but also on whether the receiving service understands the circumstances in which the data remains relevant. Wilkinson et al. [1] and Davies et al. [2] proposed reusability and interoperability as two record characteristics, respectively, within the same formal interoperability framework. This study builds on these concerns but adds the assumption that responsible transformation requires contextual knowledge as an operational prerequisite. From a context-based information fusion perspective, environmental knowledge context affects how received observations are interpreted12. Motta et al. [12] also performed a conceptual interoperability analysis and found that consistent exchanges between adaptive software services remain challenging. Semantic preservation is the process of preserving the identity, relationships, constraints, provenance and relevant qualifications of data as it moves between different representations. This goal is more than simply renaming a column. For instance, in an equipment state or a service contract, an "active" value may also refer to an approval status; distinctions between these may be removed and replaced with a single shared value that governs real decisions. Likewise, a numeric measurement without the sampling interval or calibration to the source can yield seemingly accurate, but inappropriate recommendations [13]. Hogan et al. [3] survey techniques and show how knowledge graphs can structure interrelated information, while Ji et al. [4] highlight the importance of structured relationships for downstream analytical systems.

Thus, a context-aware transformation process could be helped by such domain vocabularies, explicit rules for mapping, temporal annotations and checkpoints that could detect semantic equivalences that are not supported. Ontology access mechanisms, as evaluated by Espinoza-Arias et al. [14], relate structured domain knowledge to usable application interfaces. Numerous studies have been conducted to investigate the background knowledge of the students that answer the question of why schema alignment cannot be based on surface labels. The background knowledge survey performed by Portisch et al. [15] illuminates why schema alignment cannot be based on surface labels. With reliable decision intelligence comes another necessity: defensible inputs and comprehensible reasons are needed. The explainability taxonomy introduced by Arrieta et al. [5] is relevant to responsible AI deployment, and Adadi and Berrada survey [6] highlights the crucial role explainable evidence plays in building trust. In the proposed architecture, reliability is treated as a property of the entire information pathway, not only of the predictive model at the end of the pathway. This allows us to address questionable transformations by rejecting, reviewing or flagging them so they do not affect a business action. The difference between a reproducible pipeline and a justified decision is at the heart of the research motivation. Data quality requirements are analyzed by Priestley et al. [19] based on the data lifecycle, and it has been found that data quality requirements vary according to the lifecycle stage and the purpose of the processing. Xie et al. [20] focus on the practical implications of poor input representation for industrial quality. Each component handles a specific operation, but it can pass provenance and quality annotations to neighboring components. The proposed workflow distinguishes raw compatibility from meaning-aware interoperability. It prevents the assumption that standardized syntax means meaning.

This paper provides procedures to support the process, such as contextual feature extraction, semantic rule checking, transformation logging, risk-based validation, and human review of unresolvable ambiguity. This assessment is done in a synthetic demonstration, openly identified with 140 constructed cases. Comparisons are provided to uncover testable expectations, not to guarantee production-level performance. Arnaout and Elbassuoni [16] compare different approaches to declarative conversion decisions for semi-structured source material. The knowledge-graph construction methods proposed by Zhong et al. [17] clearly highlight the need to safeguard against meaning drift in relationships extracted from knowledge graphs. This manuscript also provides reporting conventions for analyzing an AI-oriented information system (AI-IOS) that avoid confusion between prototype scores and real institutional scores. Numerical summaries are provided to illustrate how the architecture acts as a ‘scenario benchmar' and further to explain its behaviour across increasing levels of protection. These are not measurements individually collected from an organization, and the data are not statistically significant. This argument, elaborated here, is that trustworthy automation starts with its source meaning becoming a first-class object of governance. Clear relationships between the initial context, transformed representation, and final recommendation facilitate error investigation and enable future replication. This traceability becomes even more important when operational data crosses administrative boundaries, explanations matter for acceptance, and the decision must remain traceable even if the application of origin changes. From the perspective of harmonized descriptions, approaches to semantic integration surveyed by Masmoudi et al. [18] link harmonized descriptions to reliable querying. The systematic quality checks before decisions are produced, discussed in the assessment of the dataset in the Information and Software Technology review [21], further support the value of such quality checks.

2. Review of Literature

The literature on semantic integration has shifted from manually created schema correspondence to systems that use formal representations, machine-assisted discovery and domain-specific validation. One constant observation is that many attributes

share the same name but have different operational meanings. Previous designs focused on matching entities or fields, while more recent designs included relations, temporal constraints and uncertainty. While knowledge graph research has expanded the possible relationships between source concepts, constructing graphs does not guarantee that the imported assertions capture the meaning of the source concepts. It is therefore helpful to make a distinction between the discovery of candidate correspondences and acceptance of those correspondences for high-consequence operations. Accepted alignments, which must be supported by domain constraints and provenance, can accelerate processing; candidate alignments can also accelerate processing. This difference matters when record definitions are incomplete or when local conventions change without notice. To further illustrate how different representations encode meaningful relationships, Wang et al. [7] survey knowledge embedding applications. Pugliese et al. [8] reviewed machine learning methods, all of which highlight the need to match data preparation to the eventual application. Research on metadata stewardship has provided strong principles for discoverability, access control, representation and reuse. These principles are practical and have implications for the entire information lifecycle. Latent systems might discard metadata outright, without providing any means to interpret a value, if they consider metadata "extras" rather than required. By contrast, pipelines that retain origin, measurement conditions, and transformation history enable independent inspection and help resolve errors. Efforts toward ontology-based annotation help reinforce the importance of shared concepts, while disagreements over terminology and the costs of maintenance activities can hinder adoption of this approach.

An effective architecture therefore must include mechanisms that can function in the face of incomplete, provisional or conflicting semantic descriptions. Unconfirmed correspondences should be kept separate from unsupported concepts, and can't be shoehorned into a seemingly homogeneous target vocabulary. An example of knowledge-based graph representation, which was prepared by Deagen et al. [9] to support the expert interpretation of the results, is shown. The survey of knowledge graph completion by Chen et al. [10] shows that it is easy to fall into the trap of silent relations, which are treated as observed relations. Recent research on explainable decision services shifts focus from the model output to the reasoning context in which a recommendation is consumed. Local explanations explain the model's behaviour through its influential features; global explanations summarize behaviour over a wider range of features. The following approaches are informative, but they will not correct incorrect source semantics. An explanation that confidently uses words mapped to an incorrect attribute can even mask the defect and make an incorrect decision seem reasonable. A higher-level design will tie interpretability to lineage, the rules used to transform what can be interpreted, and the context in which inputs are generated. This relationship also helps governance officers understand document permissions, policy compliance and exception handling before permitting automated deployment. The compilation of benchmarks by Blagec et al. [11] is done in light of the ontology, and it shows that using the same terminology makes them understandable and comparable across different automation approaches. Throughput vs. Oversight is a key focus in the practical study of data engineering. Standardized conversion routines, caching and automated orchestration are more likely to be used for high-volume services.

Meanwhile, domain reviews, exception queues and adaptive quality thresholds are useful features for meaning-sensitive applications. When validation scales with operational risk, these objectives need not conflict. Verified paths are scanned with less effort, and other entities or metadata not matching the expected values are scanned with more effort to ensure that a correct mapping is performed. This provides context-based safeguards that scale without requiring a special case for every record. However, performance measurements often mix syntactic and domain validity, so it may not be easy to determine what kind of improvement a specific system provides. Where data is shared across organizational and technical boundaries, Wilkinson et al. [1] original idea of data reusability remains highly relevant. When it comes to integrated evidence, one identifiable research gap is the complete decision pipeline. Most individual techniques focus on schema alignment, knowledge representation, quality management or model interpretation separately. In the present architecture, the shared context profile and persistent transformation trace provide a shared context for these tasks. It includes an evaluation model that explicitly distinguishes the following: context correspondence, Fidelity of meaning and output dependability. Conceptually, these dimensions are independent, as a record can have useful semantics without being adequate for operational decisions. On the other hand, a forecast could be consistent even if it rests on insufficiently supported assumptions. The proposed framework plans to highlight these differences through a series of validation and feedback steps. This work sets up a research path toward an implementable synthesis of representation engineering, governance practice and decision accountability, in which no single algorithm is assumed to be trustworthy. These concerns relate to explicit interpretation contracts and traceable transformations, as Davies et al. [2] describe in their formal semantic interoperability analysis.

3. Methodology

The research is done using the design-science model whereby the proposed system is defined, operationalized with a reproducible synthetic demonstration, and compared to more complex transformation options. To start, a controlled corpus of 140 constructed instances is introduced to cover three families of sources: structured operational records, sensor-like observations and textual event reports. Each instance includes a source identifier, source domain, timestamp label, authorised-use label, core content, semantic context annotation, and anticipated semantic connotation. The creation of such fields is

controlled by a scenario specification to ensure the deliberate representation of ambiguity labels, inconsistent time descriptions, lack of provenance and cross-domain terminology. The corpus contains no actual individuals, organizational secret documents or quantified production results. Second, preprocessing standardizes encodings, identifies candidate entities and leaves raw source expressions for auditing. Python and pandas support record preparation, and RDF/OWL concepts articulate domain relationships and constraints; NetworkX models relationships. Third, a context profiler isolates origin, purpose, chronology, scope, and terminology attributes, and associates an interpretive profile with each record.

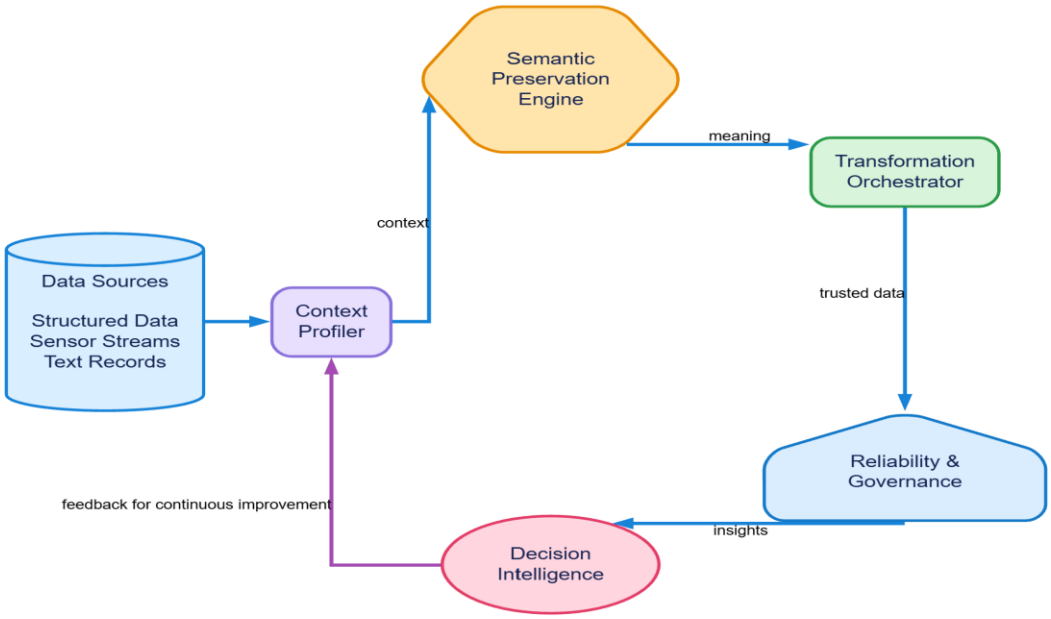


Figure 1: Semantic-preserving transformation and trustworthy decision intelligence framework context-aware

In Figure 1, the architecture suggested to accomplish context-aware data transformation, semantic preservation, and credible decision intelligence in the context of an AI-driven environment is shown. The framework combines six interrelated elements that support intelligent information processing and credible decision-making. The first is the Data Sources component, which gathers heterogeneous data from structured databases, sensor streams and textual documents. The Context Profiler then analyses contextual features, operational needs and associations in incoming information to produce meaningful interpretations. The Semantic Preservation Engine preserves the original meaning, structural integrity and contextual relevance of information as it moves through further processing steps. Subsequently, the Transformation Orchestrator performs adaptive data conversion and remains compatible with heterogeneous platforms and applications. The Reliability and Governance component reviews information integrity, processing transparency, policy adherence and operational reliability, then sends validated information to the Decision Intelligence component. The last element produces actionable information, supports informed organizational decisions and increases the responsiveness of analyses. Moreover, it integrates governance controls, intelligent processing capabilities and semantic consistency into its architecture as part of an integrated operational environment. As a result, the framework enhances information reliability, reduces contextual ambiguity, improves transformation effectiveness and ensures reliable decision intelligence in decentralized digital ecosystems.

3.1. Data Description

A synthetic evaluation corpus was established, consisting of 140 instances (50 structured operation records, 45 sensor-like records and 45 narrative event records). These category counts are illustrative; they were not sampled from a named public repository. Each of the entries is accompanied by contextual attributes, as well as an anticipated interpretation, to exemplify how the ambiguous fields can be treated semantically. Cases include some with incomplete provenance, others with inconsistent terms or varying relevance over time, and the rest are regular conversions. There is a clear distinction between records that are designed and evidence that is observed. The model is based on the stewardship considerations implied by Wilkinson et al. [1], and the knowledge-graph structure illustrated by Hogan et al. [3]. Still, it builds on the ontology-based perspective on data exemplified by Blagec et al. [11]. These 140 fabricated cases are not in those publications; they serve as background sources. The lower-scoring cases at the end of the evaluation are scenario assumptions to illustrate the manuscript, not results from a validated benchmark. This explanation avoids confusing a fake demonstration with an empirical experiment. To facilitate replication, an empirical follow-up study would publish the generation rules, inventory mapping dictionaries, annotation criteria, review decisions and executable validation routines, while adequately protecting other sensitive fields.

4. Results

The exemplary comparison shows that the meaning-conscious controls are linked to a higher end-to-end information quality compared to untreated inputs or conversions that are not carried out in the case studies. The biggest difference relates to whether the interpretation of data is actively maintained during system-to-system movement. The untyped case makes source expressions accessible but does not reliably resolve conflicting terms. Generic conversion improves representation consistency, but may not fully carry over semantic differences to the target application. Semantic mapping is one way to address this challenge, as it can correlate similar ideas; however, a mapping produced without considering the source's purpose or time-related impact may be misused. Context-aware data representation can be mathematically expressed as:

$$\mathrm{Z}_c = \phi(\mathrm{X}, \mathrm{C}; \theta) = \sigma(\mathrm{W}_x \mathrm{X} + \mathrm{W}_c \mathrm{C} + \mathrm{b}) \quad (1)$$

Four transformation strategies are compared with three illustrative measures in Figure 2. The context alignment, semantic integrity and reliability of the raw-data baseline decision are 54, 58 and 52, respectively. These indicators increase to 66, 69 and 64 with generic transformation, an increase of 12, 11 and 12 percentage points, respectively. In the context of the designed scenario, semantic mapping has a context alignment (78 percent), semantic integrity (84 percent), and decision reliability (80 percent), which implies that concept-level processing can be used to add value to basic format standardization. The integrated framework scores 91, 94 and 92, the best among all measures presented. The relative changes are 37, 36 and 40 percentage points compared to the raw data. When the integrated option is compared with semantic mapping, the differences are 13, 10 and 12 percentage points. These missing links capture the hypothetical advantage of integrating contextual profiling with governance controls, rather than mapping alone. They do not provide an experimentally calculated treatment effect.

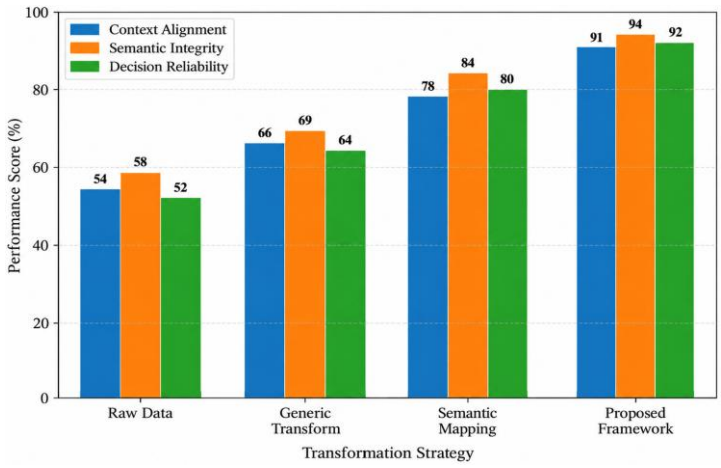


Figure 2: Relative performance in terms of the transformation strategies

The three dimensions are also kept deliberately distinct: context alignment concerns the appropriateness of interpretation, semantic integrity concerns preserving meaning, and decision reliability concerns justifying the resulting suggestions. The obvious ordering of performances helps define what should be evaluated. Still, Figures 1 and 2 should be verified independently and then used as evidence of actual superiority in the operations. Semantic preservation during data transformation can be formulated as:

$$\mathcal{L}_{sem} = \frac{1}{N}\sum_{i=1}^{N}\left[1 - \frac{\mathrm{e}(x_i)^T \mathrm{e}(T(x_i))}{\|\mathrm{e}(x_i)\|_2 \|\mathrm{e}(T(x_i))\|_2}\right] \quad (2)$$

Adaptive data transformation can be mathematically modeled as:

$$T^* = \arg\min_{T \in \mathcal{T}} [\alpha \mathcal{L}_{sem}(T) + \beta \mathcal{L}_{ctx}(T) + \gamma \mathcal{L}_{rec}(T)] \quad (3)$$

The integrated design offers a pathway to examine those contextual constraints before accepting downstream recommendations. Recognizing incoming fields is not the only important issue at the ingestion stage; gathering credible interpretive evidence is, too. A documentation environment that explains the origin of information, its purpose, and the decisions it can justifiably support is a more formidable place to start change. The system leaves uncertainty explicit when no such evidence is available. This avoids silent guessing about missing context and lets a reviewer know which assumptions have not been addressed. The

design thus favors openness about uncertainty over surface fullness. It also distinguishes between semantic gaps, which cannot be addressed without domain knowledge, and normal formatting issues, which can be addressed automatically.

Table 1: Quality indicators by transformation strategy

Evaluation Criterion	Raw Data	Generic Transform	Semantic Mapping	Proposed Framework
Context Alignment	54	66	78	91
Semantic Integrity	58	69	84	94
Decision Reliability	52	64	80	92
Provenance Completeness	48	62	79	93

Table 1 shows four indicator examples for each transformation strategy. Context alignment, semantic integrity, decision reliability and provenance completeness are attributed to raw inputs 54%, 58%, 52% and 48% of the time, respectively. Generic conversion (Figures 1 and 2) increases these values to 66, 69, 64 and 62, indicating a positive impact of structural standardization, but with major interpretive flaws. Semantic mapping yields 78%, 84%, 80%, and 79%, showing greater improvement when conceptual correspondences are presented. The suggested framework gets to 91, 94, 92 and 93, respectively. The greatest absolute increase in provenance completeness is 45 percentage points over the untreated condition. It also improves context alignment by 37 points, semantic integrity by 36 points, and decision reliability by 40 points. The integrated design increases context alignment by 13 points, semantic integrity by 12 points, decision reliability by 14 points and provenance completeness by 14 points over semantic mapping by 13, 10, 12 and 14 points, respectively. The final analogy highlights the value of rigorous source evidence and controlled examination over concept matching. The numbers are created only to provide a consistent example; they are not estimates based on the records of the synthetic cases. Their purpose is to determine what experiments will be conducted in the future must measure on their own. Information reliability and governance assessment can be represented as:

$$R(x_i) = w_1 Q_i + w_2 S_i + w_3 C_i + w_4 G_i, \qquad \sum_{j=1}^{4} w_j = 1, \quad w_j \geq 0 \tag{4}$$

Table 2 summarizes the chosen indicators at each stage of the processing chain using the same four comparative columns. Context coherence increases from 60% in raw conditions to 68% in generic transformation, 83% in semantic mapping, and 90% in the proposed design. The difference between the integrated and raw conditions is 36 percentage points, with semantic fidelity at 57%, 66%, 86%, and 93. The reliability of trust varies from 55% to 63%, then to 81%, and finally to 92, a 37-point difference. Confidence in the decisions starts at 52%, 61%, and 78%, then rises to 91, the largest increase at 39 points. The contextual and semantic enhancements are deliberately separated, as knowledge of the source condition is not sufficient to maintain a source's conceptual linkages.

Table 2: Indicators at representative points in the processing sequence

Evaluation Criterion	Raw Data	Generic Transform	Semantic Mapping	Proposed Framework
Context Coherence	60	68	83	90
Semantic Fidelity	57	66	86	93
Trust Reliability	55	63	81	92
Decision Confidence	52	61	78	91

The confidence row on decisions indicates the expected payoff of taking tested information forward into downstream recommendations. The proposed design ranks higher than the semantic-mapping comparator by 7, 7, 11 and 13 points in all four rows. These hierarchical differences suggest that governance is particularly important for trust and decision acceptance. As in Table 1, the values are pictorial scenario inputs, not outcomes from a controlled trial or industrial implementation. Reliable decision intelligence is:

$$a^* = \arg\max_{a \in \mathcal{A}} \left[\sum_{i=1}^{N} p\left(s_i \mid \mathrm{Z}_c\right) U(a, s_i) R(s_i) - \lambda \,\mathrm{Risk}(a) \right] \tag{5}$$

The preservation stage provides a controlled process of converting the candidate correspondences into approved conceptual matches. Relationships between entities, event descriptions and operating conditions receive more attention than independent column labels. This is particularly applicable to the terminologies that are used in independent departments. Various business objects can be referred to by a common label, and varying labels can be used to refer to the same domain idea. The suggested engine treats these two cases as mapping-revision issues. It aims to ensure target records are not ambiguous when the source offers multiple defensible interpretations. It allows revisiting an assumption if a source's practices change, while preserving provenance for the approved mappings it relies on.

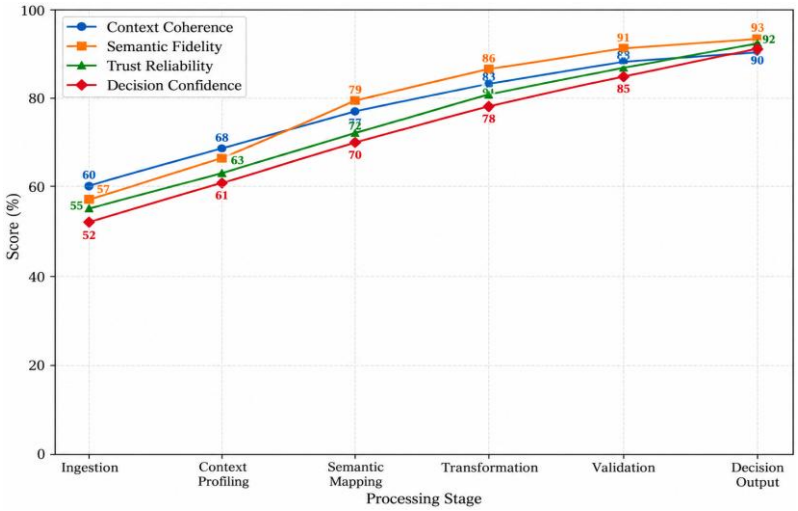


Figure 3: Depiction of the quality of framework at the processing stages

Figure 3 plots four indicators that are built up through six stages of processing. Context coherence starts at 60% during ingestion and increases to 68%, 77%, 83%, 88%, and 90% at the decision output. Semantic fidelity starts at 57%, increases to 66% after profiling, to 79% during semantic mapping, and then to 86% and 91% at the final level of 93. Trust reliability increases in order (55), 63, 72, 81, 87 and 92, rising gradually in tandem with the increasingly strong control points. Decision confidence starts at 52 and rises to 61 and 70, with the final step at 78, 85 and 91. The largest end-to-end absolute increase is in decision confidence, which rises by 39 percentage points. In contrast, context coherence, semantic fidelity, and trust reliability increase by 30, 36, and 37 points, respectively. At the decision output, the range between the maximum and minimum number of dimensions is reduced to 3 points. This intersection depicts the proposed harmonization of interpretive checks, validated transformations and decision reviews. All lines represent predetermined scenario paths rather than a platform's time series in operation. As a result, the upward trend is illustrative and does not indicate that all deployments will follow an upward trend.

These simplified trajectories may be found to be temporary quality deteriorations or slow corrections, or domain-specific trade-offs that the simplified trajectories fail to capture. Transformation orchestration introduces continuity to these interpretive controls, in practice. The engine transforms formats and coordinates records, but it does so through a mapping history rather than undocumented substitutions. This setup helps identify processing errors without recreating the whole application environment. It further isolates technical change errors, policy exceptions and semantic disputes. This division is effective because different experts fix different types of defects. Validating governance is especially important when a coherent record is introduced into a decision workflow. The question that should be considered is whether the evidence available can be used for the desired analytical use. A technically sound recommendation may also be unsuitable because its collection goal differs from the decision's. This structure thus involves permission checks, lineage availability, mapping confidence and exception status checks. Human judgment is still evident in this process, mostly in such cases where no automated checks can be made to determine an acceptable interpretation. The review results are stored in the decision history, not as a temporal operation message.

5. Discussions

The comparative exhibits show a consistent conceptual result: the most holistic architecture can hold contextual evidence and semantic links to a greater degree than the narrower alternatives. Figure 2 shows the raw condition at 54% context alignment, 58% semantic integrity, and 52% decision-making reliability. In comparison, 91, 94 and 92 are depicted for the integrated framework. These variations are 37, 36 and 40 percentage points. This trend is more architectural than statistical. Some format conversion can be used to enhance consistency of representations, but it does not ensure that the content exchanged intends to convey the same intent as the source. The framework thus shares meaning responsibilities among profiling, preservation and validation processes. The medium states of being are disclosing since they separate the conceptual sources of enhancement. Generic conversion yields 66% contextual alignment, 69% semantic integrity, and 64% decision reliability. Semantic mapping elevates these to 78, 84 and 80, indicating the assumption that domains are related when attributes interact across system boundaries. The suggested design then superimposes 13, 10 and 12 points, respectively, on the semantic-mapping scenario. This gap reflects the expected involvement of explicit context and governance, rather than the sole impact of a single transformation model. Real-world applications can show smaller differences or even inversions, depending on metadata quality and the domain being integrated. This is further reflected in Table 1's provenance completeness. Its estimated raw condition is

48%, lower than the other untreated indicators, which aligns with a scenario where lineage is missing even though some fields are interpretable. The ‘proposed condition’ has a score of 93% with a 45-point gain in provenance completeness.

It shows the largest difference between the baseline and the proposed quality measures. The design implication is that, for a reliable transformation process, information about the origin, rule choices, and review history should be persistent. In the absence of these records, an auditor may not easily be able to know why a destination value is different compared to a source expression. The proposed scenario has semantic integrity of 94%, but it is not necessary to put too much emphasis on its meaning. The score would only be believable in case the target representations were verified by independent domain experts in terms of being acceptable interpretations. Automatic concurrence between two software components would not provide the same guarantee because shared assumptions can produce the same error. This is especially relevant for ambiguous labels, which reasonably knowledgeable people can interpret differently. A developed evaluation plan must include adjudication processes and document disagreements, rather than relying on an undocumented sense of rightness. Figure 3 shows the increases built up within the internal workflow. Context coherence rises from 60% to 90% at ingestion, and semantic fidelity increases from 57% to 93%. Trust reliability rises to 92 per cent, and decision confidence rises to 91 per cent. These values depict an enhancement in progressive evidence as data is traversed through profiling, mapping, transformation, validation and usage of decisions. They cannot be interpreted as implying that each stage makes causal contributions that can be measured separately. In practice, subsequent validation can expose defects earlier in the system and provisionally improve a perceived quality score, even when the overall control system is operating as desired.

The largest difference between the first and last stage is in decision confidence, at 39 percentage points. This suggests a possible decision-level advantage of improving data interpretation at the top. But confidence is also the most presentation-sensitive indicator: a slick explanation can drive user acceptance even if it improves factual confidence less. The checkpoint in governance that is proposed covers this risk by using source links, policy records and unresolved-exception warnings. Meaningful validation will be required not only to test whether decision makers understand the uncertainty, but also whether they indicate they are more willing to trust automated recommendations. Table 2 provides an alternative perspective of the relationship between approaches to transformation and outcome indicators that is planned. The proposed scenario's context coherence, semantic fidelity, trust reliability and decision confidence are 90, 93, 92 and 91. The respective gains are 7, 7, 11 and 13 points as compared to semantic mapping. The larger disparities in trust and confidence in the decision support the working hypothesis that governance and review will be useful, especially near the decision boundary. This hypothesis needs to be investigated under controlled conditions. The practicability of the technology may matter less than organizational culture, regulatory expectations, training quality and available data infrastructure. Another problem is that of maintainability. Contextual information is altered with the upgrade of the applications, recalibration of sensors and changes in organizational roles and domain vocabularies.

When the preservation engine is based on fixed rules, initially correct transformations can degrade over time without any noise. Figure 1 has a feedback loop that is meant to manifest these changes, but with each update, there is a governance challenge to be encountered. New rules demand that there be version control, compatibility checks, and that the decisions that have been made in the past can be reconstructed. Revision suggestions can be mapped automatically to speed up maintenance, but final approval can still be made by humans. These factors also influence deployment economics. An integrated pathway is more challenging than generic conversion because it involves metadata capture, semantic assets, audit storage and exception routing. Potential advantages include minimised semantic rework and increased decision traceability, but these cannot be computed from the illustrative scores. A rigorous study would capture processing costs, latency, review effort, error severity, downstream remediation and quality measures. Such a measurement would show whether better governance is possible under realistic workloads and whether some use cases can be better controlled. In short, Figures 2 and 3 and Tables 1 and 2 are internally consistent numbers that illustrate how the system should behave. They do not determine experimental effectiveness. The substantive contribution includes identifying places where context might be lost, a way to document semantic assumptions, and checkpoints that make data-derived recommendations auditable. The greatest promise it holds is the fact that it can lead to the idea of reliability being a set of verifiable obligations that are spread over the entire decision process.

6. Conclusion

In the current paper, a unified context-aware framework that maintains meaning when transforming heterogeneous data and facilitates trustworthy decision-making processes with the help of AI is described. Source profiling, semantic preservation, governed orchestration, validation and feedback are all tied together in the design to form a traceable workflow. It is based on the idea that syntactic compatibility does not guarantee semantic correctness, and that reliable recommendations require evidence beyond the final predictive model. This means the architecture is preserved through processing stages and includes provenance, contextual qualifications, mapping decisions, and policy status. The proposed context alignment of 91%, semantic integrity of 94% and decision reliability of 92 are illustrated by comparison with the raw-condition scores of 54, 58 and 52. Table 1 indicates that the proposed design achieves 93% provenance completeness, compared with 48% in the baseline.

Throughout the processing chain, constructed trajectories achieve 90% context coherence, 93% semantic faithfulness, 92% trust reliability, and 91% decision confidence. These Figures 2 and 3 convey the evaluation targets and relations in the demonstration; they are not meant to be viewed in isolation. Their correct interpretation is pictorial, not deductive. The proposed e-system specification is testable and clearly separates interpretative, conversion and governance responsibilities. It also identifies situations where automated processing is inadequate, particularly ambiguous terminology, incomplete source descriptions and disputed permission of use. Future validation will require real datasets, expert annotation, independent baselines and reporting of operational costs. If such studies verify reliable improvements, context-sensitive transformation could improve accountability in contexts where records from various technologies and administrative spheres are integrated. The framework is empirically tested, but until this testing is done, it remains only a conceptual approach.

6.1. Limitations

The study has critical limitations. The 140 cases are artificial examples provided to architectural examples, not an external or separately collected benchmark. The performance values shown in the charts and Tables 1 and 2 were allocated to anticipated behavior; thus, they do not reflect measured accuracy, statistical confidence, or deployment superiority. The specification defines the responsibilities of candidate tools and processes, but not a completed production implementation. Consequently, issues like runtime latency, infrastructure cost, failure recovery, scalability and compatibility with legacy systems are not evaluated. Another assumption of the semantic model is that the sources are only partially described. Some domains may have little documentation of the terminology and may require significant manual interpretation to produce a good mapping. Human audit adds further uncertainty, as subjective experts may interpret unclear ideas differently. Governance requirements vary across organizations and jurisdictions, limiting the portability of any standardized policy set. If context records include private information, a conflict also arises between recording the broadest possible provenance and minimising privacy risk. Lastly, staged quality improvement is described in a simplified environment that does not involve operational back-and-forths, concept drift, adversarial records or occasional service outages. The findings do not establish causality for individual elements. External validity thus requires multi-domain testing in the future with a predefined set of measures and ground truth, which are reviewed separately. These qualifications limit the scope of the present manuscript and outline the evidence that remains to be uncovered.

6.2. Future Scope

Future work will use the entire architecture as an executable service where semantic mappings are versioned, contextual metadata contracts, and policy-conscious decision interfaces are present. The synthetic demonstration will be replaced with a public benchmark that covers operational Tables 1 and 2, time-sensitive observations, and unstructured reports. Domain experts will annotate source meaning independently and note disagreements, enabling quality mapping to be measured without assuming a single unchallenged meaning. The different cases (untreated data, generic conversion, semantic matching and integrated governance) will be compared using the same workloads during controlled evaluations. Further tests will be conducted to quantify the processing latency, effort of reviewers, storage overhead, quality of explainability, and costs of downstream corrections. The robustness testing will add missing metadata, adapted vocabularies, late records, and inconsistent units, as well as semantic conflicts that have been constructed adversarially. Individual studies will examine whether provenance-enhanced explanations better calibrate decision-making, rather than just user confidence. Human-in-the-loop review will be evaluated for practical efficiency and potential bottlenecks in high-volume environments. Future longer-term research will examine adaptive concept alignment when approval controls are used, privacy-preserving context representation and interoperability across organizational boundaries, and mapping drift detection. Healthcare operations deployments, industry monitoring, digital public service applications, and enterprise governance can impose various limitations on permissible data use. Fieldwork comparing two fields will determine the feasibility of the proposed safeguards in terms of their operating cost. These findings will be reproducible and can be scrutinized by independent researchers through an audited trail and a published evaluation protocol.

Acknowledgement: The authors sincerely acknowledge CVS Healthcare, University of Stuttgart, and the University of Technology and Applied Sciences for their valuable support and contributions to this research.

Data Availability Statement: The data used in this study are available from the corresponding author upon reasonable request.

Funding Statement: This study and the preparation of the manuscript received no financial assistance or external funding.

Conflicts of Interest Statement: The authors confirm that there are no conflicts of interest associated with this research.

Ethics and Consent Statement: The research was conducted in accordance with applicable ethical standards, and informed consent was secured from all participants.

References

1. M. D. Wilkinson, M. Dumontier, I. J. Aalbersberg, G. Appleton, M. Axton, A. Baak, N. Blomberg, J. W. Boiten, L. Bonino da Silva Santos, P. E. Bourne, J. Bouwman, A. J. Brookes, T. Clark, M. Crosas, I. Dillo, O. Dumon, S. Edmunds, C. T. Evelo, R. Finkers, A. Gonzalez-Beltran, A. J. G. Gray, P. Groth, C. Goble, J. S. Grethe, J. Heringa, P. A. C. 't Hoen, R. Hooft, T. Kuhn, R. Kok, J. Kok, S. J. Lusher, M. E. Martone, A. Mons, A. L. Packer, B. Persson, P. Rocca-Serra, M. Roos, R. van Schaik, S. A. Sansone, E. Schultes, T. Sengstag, T. Slater, G. Strawn, M. A. Swertz, M. Thompson, J. van der Lei, E. van Mulligen, J. Velterop, A. Waagmeester, P. Wittenburg, K. Wolstencroft, J. Zhao, and B. Mons, "The FAIR Guiding Principles for scientific data management and stewardship," *Scientific Data*, vol. 3, no. 3, pp. 1–9, 2016.
2. J. Davies, J. Welch, D. Milward, and S. Harris, "A formal, scalable approach to semantic interoperability," *Science of Computer Programming*, vol. 192, no. 6, pp. 1–26, 2020.
3. A. Hogan, E. Blomqvist, M. Cochez, C. D'amato, G. De Melo, C. Gutierrez, S. Kirrane, J. E. L. Gayo, R. Navigli, S. Neumaier, A. C. N. Ngomo, A. Polleres, S. M. Rashid, A. Rula, L. Schmelzeisen, J. Sequeda, S. Staab, and A. Zimmermann, "Knowledge Graphs," *ACM Computing Surveys*, vol. 54, no. 4, pp. 1–37, 2021.
4. S. Ji, S. Pan, E. Cambria, P. Marttinen, and P. S. Yu, "A Survey on Knowledge Graphs: Representation, Acquisition, and Applications," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 33, no. 2, pp. 494–514, 2022.
5. A. B. Arrieta, N. Díaz-Rodríguez, J. D. Ser, A. Bennetot, S. Tabik, A. Barbado, S. Garcia, S. Gil-Lopez, D. Molina, R. Benjamins, R. Chatila, and F. Herrera, "Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI," *Information Fusion*, vol. 58, no. 6, pp. 82–115, 2022.
6. A. Adadi and M. Berrada, "Peeking Inside the Black-Box: A Survey on Explainable Artificial Intelligence (XAI)," *IEEE Access*, vol. 6, no. 9, pp. 52138–52160, 2018.
7. Q. Wang, Z. Mao, B. Wang, and L. Guo, "Knowledge Graph Embedding: A Survey of Approaches and Applications," *IEEE Transactions on Knowledge and Data Engineering*, vol. 29, no. 12, pp. 2724–2743, 2017.
8. R. Pugliese, S. Regondi, and R. Marini, "Machine learning-based approach: Global trends, research directions, and regulatory standpoints," *Data Science and Management*, vol. 4, no. 12, pp. 19–29, 2021.
9. M. E. Deagen, J. P. McCusker, T. Fateye, S. Stouffer, L. C. Brinson, D. L. McGuinness, and L. S. Schadler, "FAIR and Interactive Data Graphics from a Scientific Knowledge Graph," *Scientific Data*, vol. 9, no. 5, pp. 1–11, 2022.
10. Z. Chen, Y. Wang, B. Zhao, J. Cheng, X. Zhao, and Z. Duan, "Knowledge Graph Completion: A Review," *IEEE Access*, vol. 8, no. 10, pp. 192435–192456, 2020.
11. K. Blagec, A. Barbosa-Silva, S. Ott, and M. Samwald, "A curated, ontology-based, large-scale knowledge graph of artificial intelligence tasks and benchmarks," *Scientific Data*, vol. 9, no. 6, pp. 1–10, 2022.
12. L. Snidaro, J. García, and J. Llinas, "Context-based information fusion: A survey and discussion," *Information Fusion*, vol. 25, no. 9, pp. 16–31, 2015.
13. R. C. Motta, K. M. de Oliveira, and G. H. Travassos, "A conceptual perspective on interoperability in context-aware software systems," *Information and Software Technology*, vol. 114, no. 10, pp. 231–257, 2019.
14. P. Espinoza-Arias, D. Garijo, and O. Corcho, "Crossing the chasm between ontology engineering and application development: A survey," *Journal of Web Semantics*, vol. 70, no. 7, pp. 1–21, 2021.
15. J. Portisch, M. Hladik, and H. Paulheim, "Background knowledge in ontology matching: A survey," *Semantic Web: – Interoperability, Usability, Applicability*, vol. 15, no. 6, pp. 2639–2693, 2024.
16. H. Arnaout and S. Elbassuoni, "Effective searching of RDF knowledge graphs," *Journal of Web Semantics*, vol. 48, no. 1, pp. 66–84, 2018.
17. L. Zhong, J. Wu, Q. Li, H. Peng, and X. Wu, "A comprehensive survey on automatic knowledge graph construction," *ACM Computing Surveys*, vol. 56, no. 4, pp. 1–62, 2023.
18. M. Masmoudi, S. B. A. B. Lamine, M. H. Karray, B. Archimede, and H. B. Zghal, "Semantic data integration and querying: A survey and challenges," *ACM Computing Surveys*, vol. 56, no. 8, pp. 1–35, 2024.
19. M. Priestley, F. O'Donnell, and E. Simperl, "A survey of data quality requirements that matter in ML development pipelines," *Journal of Data and Information Quality*, vol. 15, no. 2, pp. 1–39, 2023.
20. J. Xie, L. Sun, and Y. F. Zhao, "On the data quality and imbalance in machine learning-based design and manufacturing—A systematic review," *Engineering*, vol. 45, no. 2, pp. 105–131, 2025.
21. Y. Gong, G. Liu, Y. Xue, R. Li, and L. Meng, "A survey on dataset quality in machine learning," *Information and Software Technology*, vol. 162, no. 10, pp. 1–12, 2023.

Publisher's Note: *The publisher remains impartial concerning jurisdictional claims in published maps and institutional affiliations. Responsibility for the content rests entirely with the authors and does not necessarily reflect the publisher's perspectives.*