AIRIA: A Human-AI Framework for Geospatial Regulatory Interpretation Alignment

Sreedhar Ailu
sree.6506@gmail.com
ORCID: 0009-0004-4289-445X

Abstract — Multi-vendor enterprise systems operating under regulatory mandates face a structural software quality failure that existing governance frameworks do not address: interpretation divergence. When multiple vendors independently translate the same regulatory requirement into system behavior, each may produce an implementation that is internally correct yet incompatible with the others at system integration boundaries. In regulated healthcare exchange programs under the Affordable Care Act (ACA), this failure is compounded by geospatial risk — regulatory terms such as service area, county boundary, and rating region carry implicit spatial data model dependencies that different vendors resolve through incompatible geographic representations, producing divergent eligibility determinations for the same enrollees.

This paper introduces AIRIA: the AI-Assisted Regulatory Interpretation Alignment framework, a Design Science Research artifact comprising a five-layer human-AI architecture that intercepts interpretation divergence before it propagates into acceptance criteria and implementation. The five layers are: (1) Regulatory Ingestion, (2) Ambiguity Detection, (3) Geospatial Interpretation Analysis, (4) Regulatory Knowledge Graph construction, and (5) Divergence Detection and Alerting. AI performs computationally intensive consistency work across these layers while human governance structures retain all binding decision authority.

This paper proposes research artifacts and evaluation pathways addressing three research questions: whether geospatial regulatory terms can be systematically classified by spatial data model divergence risk; whether a knowledge graph schema can represent regulatory interpretation lineage with explicit geospatial nodes; and whether AI-assisted governance can detect multi-vendor divergence before implementation. AIRIA is grounded in more than 13 years of practitioner observation across healthcare exchange and government-regulated enterprise implementations involving multiple independent vendor organizations, presented as a conceptual architecture with a five-phase empirical validation roadmap.

Index Terms—*regulatory interpretation alignment, multi-vendor governance, AI decision support, geospatial eligibility, knowledge graphs, acceptance testing, information systems*

I. INTRODUCTION

Geographic information science has extensively documented the challenges of spatial data model divergence across organizational boundaries — coordinate system mismatches, boundary definition conflicts, and semantic heterogeneity in spatial feature definitions [1], [2], [3]. This paper identifies a new context: regulatory compliance in multi-vendor enterprise systems. When regulatory text specifies a geographic concept — service area, county, rating region — without specifying a spatial data model, vendors independently select incompatible models, and the result is a governance failure with regulatory compliance consequences.

Despite extensive work on spatial interoperability [1], [3], little attention has been given to governance mechanisms required when competing spatial models arise from regulatory interpretation itself. This paper addresses that gap.

The primary proposed artifact of this Design Science Research study is a governance framework for the identification, classification, and resolution of geospatial regulatory ambiguity arising from competing spatial data models. This is a problem geographic information science is well positioned to investigate — one that existing regulatory knowledge graph frameworks (FIBO, Akoma Ntoso, HL7 FHIR) do not address.

A. Research Questions

This paper addresses three research questions. RQ1: Can geospatial regulatory terms be systematically classified by spatial data model divergence risk, and can that classification be automated using NLP methods? RQ2: Can a knowledge graph schema represent regulatory interpretation lineage — from ambiguous term through governance decision to acceptance criterion and test case — with explicit geospatial reference nodes? RQ3: Can an AI-assisted governance framework detect multi-vendor interpretation divergence before implementation, reducing late-cycle defect discovery?

The paper does not claim that AIRIA answers these questions empirically. Rather, it proposes artifacts — a taxonomy, a knowledge graph schema, scoring models, and a governance architecture — and evaluation pathways through which these questions may be investigated. This positioning follows the Design Science Research paradigm, in which artifact specification and evaluation roadmap are themselves recognized contributions [4].

B. Proposed Research Artifacts

- A GIS-grounded taxonomy of regulatory geospatial term types with spatial data model divergence risk profiles.
- A five-layer AI architecture for regulatory ambiguity detection, geospatial interpretation analysis, knowledge graph construction, and divergence detection and alerting.
- A Regulatory Knowledge Graph (RKG) schema with explicit geospatial reference nodes.
- A Geospatial Ambiguity Risk Score (GARS) and Alignment Confidence Score (ACS) with defined validation paths.
- A five-component human governance structure retaining full binding decision authority.
- A five-phase empirical validation roadmap with measurable success criteria.

C. Design Science Research Positioning

AIRIA is presented as a Design Science Research (DSR) artifact in the sense established by [4] and [5]. AIRIA has not been deployed in a controlled implementation. Projected outcomes are derived from reasoning about observed failure modes, not from experimental measurement. All numerical values in worked examples are illustrative only and are not empirically derived.

II. The Regulatory Interpretation Problem

A. The Structural Failure Mode

Regulatory language is designed for legal precision, not operational specificity. Terms such as 'timely,' 'service area,' and 'prior notice' are legally well-formed but operationally underspecified. When multiple vendors independently translate such terms into system behavior, they produce divergent implementations that are each internally correct and mutually incompatible. This divergence is invisible within any single vendor's development process, becoming visible only at UAT — the most expensive remediation point in the SDLC.

B. Why Geospatial Terms Exhibit High Divergence Risk

Among the regulatory term categories observed in healthcare exchange implementations, geospatial terms consistently exhibited the greatest interpretation divergence risk. Two vendors can agree on the semantic meaning of 'service area' and still produce incompatible implementations because they chose different spatial data models — CMS county polygons versus carrier-defined service polygons.

C. Prevalence and Observed Failure Patterns

The failure patterns described here were observed directly across multiple independent program engagements in U.S. healthcare exchange systems operating under ACA regulations, involving state-based and federal marketplace programs. Each program involved three or more independent vendor organizations. Interpretation divergence was repeatedly observed across reviewed program categories, consistently manifesting at or before UAT.

While precise frequency statistics are unavailable due to client confidentiality constraints, practitioner experience suggests that requirements involving geospatial terms, eligibility conditions, and temporal specifications are particularly exposed. This assessment is observational and requires validation through the corpus and pilot studies specified in Section IX.

TABLE I Failure Pattern Taxonomy

ID	Pattern and Effect
P1	Interpretation Divergence — Multiple vendors derive different operational behaviors from the same mandate. Creates divergent data contracts invisible within any single vendor's test suite.
P2	Acceptance Criteria Misalignment — QA teams write criteria against their own interpretation. UAT produces contradictory results.
P3	Integration Amplification — Minor interpretation differences produce critical-severity defects at integration points requiring architectural rework.
P4	Late Discovery During UAT — UAT is the first multi-vendor integration checkpoint. Divergences discovered at maximum remediation cost.

III. Related Work

A. NLP for Regulatory Text Analysis

Natural language processing research on legal and regulatory documents spans rule-based and neural classification approaches. Lam et al. [6] established that regulatory texts exhibit characteristic linguistic patterns amenable to rule-based extraction. LegalBERT [7], CaseHOLD [8], and BERT [9] provide strong benchmarks for legal NLP. AIRIA Layer 2 builds on this literature, adding per-term ambiguity risk scores calibrated for multi-vendor operational divergence. The rationale for transformer-based approaches is tractability at scale — rule-based extraction requires manual authoring of patterns per term class, while a fine-tuned transformer generalizes from labeled examples. Both are viable Phase 2 prototype candidates; the validation study will compare them empirically.

B. Knowledge Graphs for Regulatory Compliance

Knowledge graph applications in regulatory compliance have been explored across financial services (FIBO [10]), healthcare (HL7 FHIR [12]), and legal informatics (Akoma Ntoso [11]; LegalRuleML [13]). GeoSPARQL [14] provides an RDF/OWL vocabulary for geospatial knowledge representation. The AIRIA RKG extends this tradition by adding an interpretation lineage layer connecting regulation nodes to vendor interpretation nodes, acceptance criteria nodes, and test case nodes — not present in existing compliance knowledge graph schemas.

C. Geospatial Data Integration

Geographic information science research has extensively documented spatial data model divergence challenges [1], [2], [15]. Boundary conflation is well-studied [16]. Geospatial semantic interoperability has been addressed through ontology alignment [3], [17] and GeoSPARQL [14].

Geospatial semantic uncertainty research [21], [22] addresses the inherent vagueness of geographic concepts. The Modifiable Areal Unit Problem (MAUP) [23], [24] describes how results vary depending on spatial unit of analysis. AIRIA Layer 3's GARS can be understood as a domain-specific operationalization of MAUP sensitivity for regulatory eligibility contexts. Reasoning over spatial context for AI-assisted analysis — an active area of recent GIS research [32] — provides additional theoretical grounding for the geospatial ambiguity detection approach.

D. Multi-Vendor Integration Literature

Research on coordination failures in distributed software development [25], [26] documents how organizational separation produces integration defects consistent with the failure patterns in Section II. The AIRIA framework addresses the upstream regulatory interpretation layer that these frameworks assume has already been resolved. Standards such as IEEE 829 [18] and ISO/IEC 29119 [19] assume acceptance criteria exist and are consistent — AIRIA addresses the upstream precondition they assume but do not provide.

E. AIRIA Novelty

The key structural absence in all prior frameworks is the interpretation governance layer: the process of detecting, surfacing, and resolving divergent operational interpretations of regulatory text before those interpretations become code. Table II positions AIRIA against five existing frameworks following the DSR novelty verification method described by Gregor and Hevner [20].

TABLE II Capability Matrix

Capability	FIBO	AN	FHIR	Geo S.	LRuleML	AIRIA
Regulatory representation	✓	✓	✓	✓	✓	✓
Geospatial representation	✗	✗	✗	✓	✗	✓
Ambiguity detection	✗	✗	✗	✗	~	✓
Spatial model divergence	✗	✗	✗	✗	✗	✓
Interpretation lineage	✗	✗	✗	✗	✗	✓
Pre-impl. divergence alert	✗	✗	✗	✗	✗	✓

✓ present ~ partial ✗ not addressed. AN = Akoma Ntoso; GeoS. = GeoSPARQL; LRuleML = LegalRuleML

IV. The AIRIA AI Architecture

The AIRIA AI architecture processes regulatory text through five sequential layers before human governance structures engage. Each layer is independently evaluable. All worked example values are illustrative only — they are not empirically derived from deployed implementations.

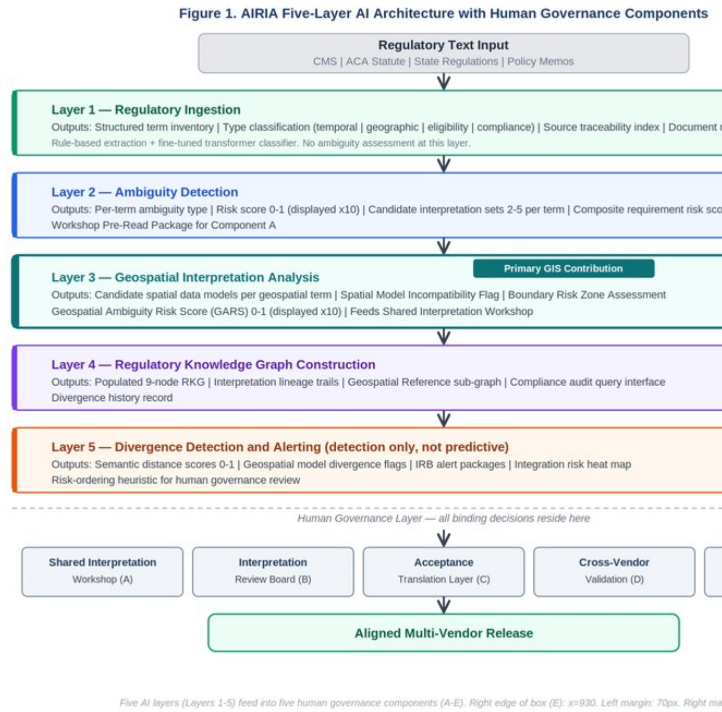

Fig. 1. AIRIA five-layer AI architecture with human governance components. Layer 3 (geospatial interpretation analysis) represents the primary IJGIS-focused contribution. All binding decisions reside in the human governance layer.

A. Layer 1: Regulatory Ingestion

Layer 1 ingests and classifies regulatory source documents into structured term records. Inputs include CMS regulations, ACA statute, and state-level exchange regulations. AI outputs include a structured term inventory, type classification (temporal | geographic | eligibility | compliance | threshold | procedural), source traceability index, and document metadata. Layer 1 uses rule-based extraction augmented by a fine-tuned transformer classifier.

B. Layer 2: Ambiguity Detection

Layer 2 classifies each extracted term by ambiguity type and generates per-term risk scores. AI outputs include Per-term Ambiguity Type (temporal | scope | threshold | geospatial | definitional gap), Per-term Ambiguity Risk Score on 0–1 scale (displayed ×10), Candidate Interpretation Sets (2–5 per flagged term), and Workshop Pre-Read Package. All binding decisions remain with human governance.

Worked Example (illustrative): Input: 'timely' | illustrative regulatory notice requirement | Risk Score: 0.87 [HIGH]. Candidates: (A) within 24 hours [HIPAA]; (B) within 3 business days [CMS QHP]; (C) within 5 calendar days [state Medicaid].

C. Layer 3: Geospatial Interpretation Analysis

Layer 3 detects geospatially typed regulatory terms and enumerates candidate spatial data models. AI outputs include Per-term Geospatial Type (administrative boundary | carrier-defined area | census unit | proximity zone | jurisdictional overlap), Candidate Spatial Data Models with source authority, Spatial Model Incompatibility Flag, Boundary Risk Zone Assessment, and GARS value on 0–1 scale (displayed ×10). Layer 3 does not select the authoritative spatial data model — that is a governance decision. This layer operationalizes MAUP sensitivity [23], [24] as a per-term governance risk signal.

Worked Example (illustrative): Input: 'service area' | illustrative ACA/QHP requirement. Models: (A) CMS FIPS county polygon; (B) carrier service polygon; (C) USPS ZIP/ZCTA. Flag: RAISED. GARS: 0.94 (9.4/10) [CRITICAL] — illustrative value.

D. Layer 4: Regulatory Knowledge Graph Construction

Layer 4 populates a machine-queryable knowledge graph linking regulation, interpretations, spatial models, acceptance criteria, and test cases. Outputs include a Populated RKG with typed nodes and directional edges, Interpretation Lineage Trails, Geospatial Reference Sub-graph, and Compliance Audit Query Interface. Table III defines the nine-node schema.

TABLE III RKG Node Schema (Nine Types)

Node Type	Example Instance	Outbound Edges
Regulation	ACA §1311(c)(1)	→ CONTAINS_TERM
Ambiguity Flag	timely [HIGH, 0.87]	→ GENERATES
Candidate Interpretation	timely := 3 days	→ ADOPTED_AS
Binding Definition	timely := 3 days [IRB-004]	→ GOVERNS
Geospatial Reference	service area := CMS county	→ INSTANTIATES
Spatial Data Source	CMS QHP shapefile v2024	→ PROVIDES_GEOMETRY
Vendor Interpretation	Vendor A: carrier polygon	→ DIVERGES_FROM
Acceptance Criterion	Notice within 3 days	→ REALIZED_BY
Test Case	TC-047: verify by day 3	→ VALIDATES

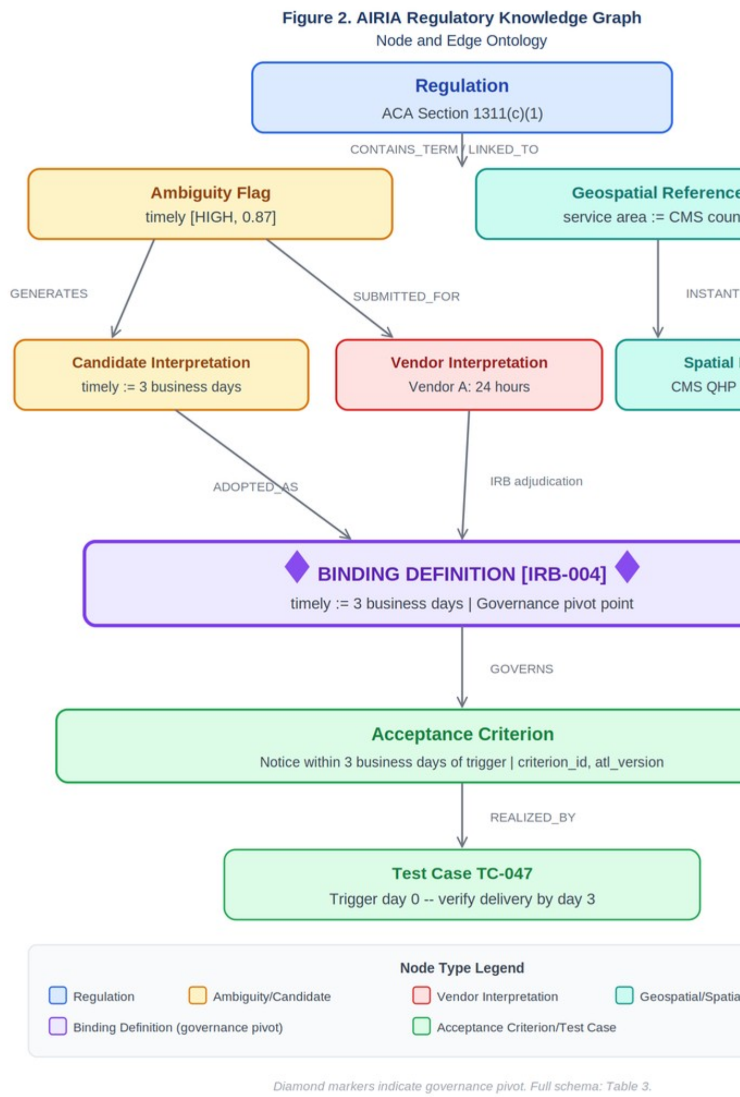

Fig. 2. AIRIA Regulatory Knowledge Graph node and edge ontology. The Binding Definition node, marked by diamond symbols, is the governance pivot point. Full schema: Table III.

E. Layer 5: Divergence Detection and Alerting

Layer 5 detects semantic distance between vendor-submitted interpretations and surfaces structured alerts to governance structures before implementation begins. This layer is a detection and alerting mechanism — not a predictive model. It does not predict program outcomes, defect counts, or release failures. AI outputs include

Semantic Distance Scores (0–1), Geospatial Model Divergence Flags, Divergence Alert Packages, Integration Risk Heat Map, and Integration Risk Prioritization (risk-ordering heuristic, not a predictive model).

V. Human Governance Components

The five AI layers produce structured outputs consumed by five human governance components. All binding decision authority resides with the human governance layer. AIRIA employs AI as a decision-support mechanism; all binding regulatory interpretations remain the responsibility of human governance authorities.

Component A — Shared Interpretation Workshop: cross-vendor facilitated session reviewing AI-flagged terms and evaluating candidate interpretations. Component B — Interpretation Review Board (IRB): standing governance body adjudicating interpretive disputes and issuing binding definitions. Component C — Acceptance Translation Layer (ATL): derives vendor-specific acceptance criteria from IRB binding definitions. Component D — Cross-Vendor Validation Protocol: pre-UAT joint execution of AI-identified highest-risk integration scenarios. Component E — Release Readiness Alignment Review: gate review combining defect metrics with AI-generated ACS.

A. Alignment Confidence Score

The ACS is a composite metric providing a release readiness signal. All weights are illustrative starting estimates requiring validation before operational use. Weight rationale: Interpretation agreement (0.30) addresses root cause P1; acceptance criteria alignment (0.25) addresses P2; geospatial consistency (0.20) is elevated due to eligibility misclassification severity; validation coverage (0.15) reflects process execution; open divergence (0.10) is an inverse health indicator.

ACS = (0.30×interp_agreement) + (0.25×criteria_alignment) + (0.20×geo_consistency) + (0.15×validation_coverage) + (0.10×(1−open_divergence_rate)). Scale: 0.0 to 1.0.

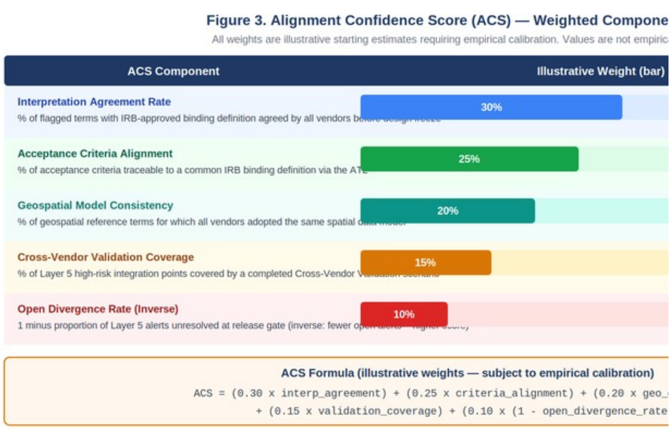

Fig. 3. Alignment Confidence Score (ACS) weighted component model. All weights are illustrative and require empirical calibration (Phase 4 of validation roadmap). Values are not empirically derived.

VI. GIS Contributions

A. Regulatory Geospatial Term Taxonomy

One of the most concrete GIS contributions of this paper is the identification of regulatory geospatial terms as a distinct linguistic category. The regulatory geospatial term problem is distinct from general geographic named entity recognition in three ways: (1) terms are operational, specifying a spatial computation; (2) terms carry compliance consequences; and (3) terms are systematically ambiguous by design. Table IV presents the five-category taxonomy.

TABLE IV Regulatory Geospatial Term Taxonomy

Category	Regulatory Examples	Divergence Mechanism
Administrative boundary	county, state, ZIP, census tract	Multiple boundary authorities produce non-identical polygons. Related to MAUP [23].
Carrier-defined area	service area, network coverage area	Carrier polygons proprietary; CMS tables discrete; vendors disagree on partially-covered counties.
Proximity zone	within X miles, local area	Euclidean, network, and drive-time distances produce different polygons.
Jurisdictional boundary	state jurisdiction, tribal land	Legally defined but geometrically contested at edges.
Temporal-spatial reference	open enrollment window, plan year	Time-bounded geographic rules require temporal versioning alignment.

B. Spatial Ambiguity Risk Scoring (GARS)

$GARS(t) = w_1 \cdot M(t) + w_2 \cdot D(t) + w_3 \cdot H(t) + w_4 \cdot (1-A(t))$. Scale: 0–1 internal; displayed ×10. Suggested weights: $w_1=0.20$, $w_2=0.35$, $w_3=0.30$, $w_4=0.15$ — principled starting estimates, not validated. D(t) receives highest weight (0.35) because spatial boundary displacement directly determines eligibility misclassification magnitude. Terms with GARS ≥ 0.70 are provisionally flagged for priority review using this illustrative threshold. Both threshold and weights require calibration against labeled outcomes before operational use.

C. Geospatial Interpretation Lineage

The RKG schema includes Geospatial Reference nodes linked to Spatial Data Source and Binding Definition nodes. Across multiple program implementations and regulatory cycles, this becomes a geospatial interpretation lineage infrastructure: a queryable record of which spatial data models have been used to implement which regulatory terms under which governance decisions.

D. Validation Scenario

Regulatory anchor: 45 CFR §155.420 — Special Enrollment Periods (illustrative scenario). Ambiguous operational concept: a permanent move resulting in access to new QHPs across a service-area boundary. The following vendor interpretations are hypothetical operationalizations, not quotations from the CFR.

Vendor A interpretation (hypothetical): Permanent move operationalized as county-boundary crossing (FIPS polygon). 60-day window starts on effective date of address change. Vendor B interpretation (hypothetical): Permanent move operationalized as change from one carrier service polygon to a non-overlapping carrier polygon. Window starts on application submission date.

Without AIRIA: conflict discovered at UAT — week 34 of a 40-week program. Remediation requires architectural rework (both vendors). Three-week schedule slip. All values illustrative.

With AIRIA: Layer 3 alert at week 4. IRB Binding Definitions: BD-047 (permanent move = CMS FIPS county boundary crossing); BD-048 (window opens on eligibility determination date); BD-049 (authoritative spatial source = CMS QHP county table). Divergence resolved before design freeze. No schedule impact. All values illustrative.

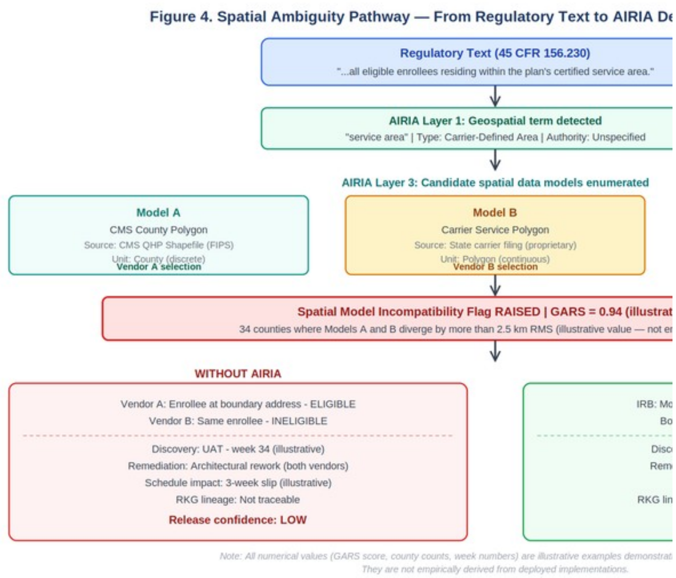

Fig. 4. Spatial ambiguity pathway for 'service area.' Three candidate spatial data models produce different eligibility determinations. AIRIA Layer 3 detects divergence before design freeze; the IRB selects the authoritative model. All values are illustrative.

TABLE V Validation Scenario Outcomes (Illustrative)

Dimension	Without AIRIA	With AIRIA
Discovery point	UAT — week 34	Layer 3 alert — week 4
Remediation	Architectural rework	Acceptance criteria update
Schedule impact	3-week slip	None
RKG lineage	Not traceable	BD-047 → AC-089 → Test cases

Release confidence	Low	High

VII. Process Comparison

A. Current State

Regulatory mandate issued → Independent vendor interpretation (no shared checkpoint) → Divergent acceptance criteria → UAT → Integration failure discovered → Late rework → Eroded release confidence.

B. AIRIA State

Regulatory mandate issued → Layer 1: Regulatory Ingestion → Layer 2: Ambiguity Detection → Layer 3: Geospatial Analysis → Human: Shared Interpretation Workshop + IRB → Layer 4: RKG Construction → Human: ATL Acceptance Criteria → Layer 5: Divergence Detection and Alerting → UAT → Aligned criteria → Release Readiness Review → Aligned multi-vendor release.

VIII. Limitations and Threats to Validity

A. Deployment Status

AIRIA has not been deployed in any controlled or real-world implementation. Every process outcome described is a projection derived from reasoning about observed failure patterns, not an empirically measured result. The validation scenario in Section VI-D is illustrative. All numerical values are illustrative examples constructed from publicly available ACA regulatory text. Readers should not cite AIRIA's projected outcomes as evidence of framework effectiveness.

B. Threats to Internal Validity

The four failure patterns in Section II are derived from direct practitioner observation. They have not been validated through systematic review of documented defect records. The GARS formula components and their suggested weights are reasonable starting estimates, not empirically calibrated values. The ACS formula components and weights (0.30, 0.25, 0.20, 0.15, 0.10) are illustrative. Layer 2 and Layer 3 are described at a functional architecture level — performance characteristics are unknown.

C. Threats to External Validity

All practitioner observations are drawn from U.S. healthcare exchange programs under the ACA. The failure patterns, taxonomy categories, and governance mechanisms may not generalize to other regulated domains.

D. What This Paper Does and Does Not Claim

This paper claims: a coherent explanation of interpretation divergence; a five-layer AI architecture specified with sufficient detail; an RKG schema with geospatial reference nodes; a regulatory geospatial term taxonomy; and a research agenda with measurable success criteria. This paper does not claim: that AIRIA has been proven to reduce defects; that GARS or ACS are validated instruments; that the AI layers can be built without further engineering research; or that validation scenario outcomes represent actual program results.

IX. Future Research and Validation Roadmap

A. Phased Validation Roadmap

TABLE VI AIRIA Phased Validation Roadmap

Ph.	Name	Success Criteria
1	**Corpus Development**	Cohen's κ ≥ 0.75 for geospatial term category assignment. Labeled corpus available for NLP training.
2	**NLP Prototype**	F1 ≥ 0.80 per ambiguity category. GARS values correlate with human expert ratings (Spearman ρ ≥ 0.70).
3	**RKG Implementation**	Lineage traversal for ≥ 10 requirements. Schema accommodates all nine node types from Table III without modification.
4	**Controlled Pilot**	Statistically significant reduction in P4-pattern defects. ACS predictive of release outcome (ROC AUC ≥ 0.75).
5	**Cross-Domain Study**	Taxonomy covers ≥ 85% of geospatial terms in FINRA/SEC domain corpus.

X. Conclusion

The AIRIA framework addresses a structural governance failure that existing testing and requirements engineering frameworks do not: the divergent interpretation of regulatory mandates across independent vendor teams in multi-vendor enterprise systems. The failure is not a testing problem — it is an interpretation alignment problem that must be resolved upstream of implementation.

The framework's five AI layers provide a computationally tractable path to surfacing interpretation risk before it propagates. The human governance components ensure that binding decisions remain with people who carry regulatory and contractual accountability. The geospatial dimension is not rhetorical: when vendors operationalize service areas, county boundaries, and rating regions through incompatible spatial data models, the result is an eligibility classification problem with regulatory compliance consequences native to geographic information science.

At the policy level, AIRIA's taxonomy and GARS model have implications for regulatory text design: regulators could reduce spatial ambiguity by explicitly identifying an authoritative spatial data source and boundary model for ACA service area determinations. This represents a direct implication for CMS drafting practice and state exchange authority guidance documents.

AIRIA should therefore be evaluated not as a completed solution, but as a research artifact intended to stimulate empirical investigation into geospatial regulatory interpretation governance. Future work will evaluate AIRIA through retrospective analysis of historical regulatory interpretation cases and prospective pilot implementations within regulated multi-vendor enterprise environments.

The central claim of this paper is precise: AI does not determine regulatory interpretation. AI makes the humans who govern interpretation faster, more consistent, and better-prepared — creating a testable pathway for preventing the failure chain. AIRIA employs AI as a decision-support mechanism; all binding regulatory interpretations remain the responsibility of human governance authorities.

Submission Compliance Statements

Figures and Tables: All figures (Figs. 1–4) are original publication-quality graphics created by the author. All tables (Tables I–VI) are original works. No third-party images or spatial data products have been reproduced.

Originality Statement: This manuscript has not been previously published in this form and is not under review at any other journal. The preprint cited as [45] represents an earlier practitioner-focused version.

Data Availability: No empirical datasets were used. Regulatory text referenced is publicly available from the U.S. Government Publishing Office.

Conflicts of Interest: The author declares no conflicts of interest.

References

[1] M. Duckham and M. Worboys, "An algebraic approach to automated geospatial information fusion," Int. J. Geogr. Inf. Sci., vol. 19, no. 5, pp. 537-557, 2005.

[2] F. Fonseca, M. Egenhofer, P. Agouris, and G. Camara, "Using ontologies for integrated geographic information systems," Trans. GIS, vol. 6, no. 3, pp. 231-257, 2002.

[3] F. Harvey and W. Kuhn, "A linguistic framework for assessing ontology mistranslations," in Geographic Information Science, Springer, 2002, pp. 136-150.

[4] A. R. Hevner, S. T. March, J. Park, and S. Ram, "Design science in information systems research," MIS Q., vol. 28, no. 1, pp. 75-105, 2004.

[5] S. T. March and G. F. Smith, "Design and natural science research on information technology," Decis. Support Syst., vol. 15, no. 4, pp. 251-266, 1995.

[6] H. P. Lam, M. Hashmi, and G. Governatori, "Enabling automated regulation compliance checking," Requirements Eng., vol. 24, no. 4, pp. 453-484, 2019.

[7] I. Chalkidis et al., "LEGAL-BERT: The muppets straight out of law school," in Findings ACL: EMNLP 2020, pp. 2898-2904.

[8] Y. Koreeda and C. D. Manning, "CaseHOLD: How holdings are treated in case law datasets," in Proc. ACL 2021, pp. 4603-4614.

[9] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers," in Proc. NAACL 2019, pp. 4171-4186.

[10] D. Graux et al., "FIBO: The Financial Industry Business Ontology," in Proc. Joint Workshop Financial Narrative Processing, 2018, pp. 56-61.

[11] G. Barabucci et al., "Multi-layer markup in Akoma Ntoso," in Approaches to Legal Ontologies, Springer, 2010, pp. 239-257.

[12] T. Benson and G. Grieve, Principles of Health Interoperability, 3rd ed., Springer, 2016.

[13] T. Athan et al., "LegalRuleML: Design principles," in Reasoning Web, Springer, 2015, pp. 151-188.

[14] R. Battle and D. Kolas, "Enabling the geospatial Semantic Web with GeoSPARQL," Semantic Web, vol. 3, no. 4, pp. 355-370, 2012.

[15] M. F. Worboys and E. Clementini, "Integration of imperfect spatial information," J. Vis. Lang. Comput., vol. 12, no. 1, pp. 61-80, 2001.

[16] A. Saalfeld, "Conflation automated map compilation," in Proc. AUTO-CARTO 9, 1993, pp. 167-177.

[17] W. Kuhn, "Geospatial semantics: Why, of what, and how?" J. Data Semantics, vol. 3, pp. 1-24, 2005.

[18] IEEE, IEEE Std 829-2008 (Reaffirmed 2013), IEEE, 2013.

[19] ISO/IEC/IEEE, ISO/IEC/IEEE 29119-1:2013, ISO, 2013.

[20] S. Gregor and A. R. Hevner, "Positioning design science research," MIS Q., vol. 37, no. 2, pp. 337-355, 2013.

[21] P. F. Fisher, "Models of uncertainty in spatial data," in Geographical Information Systems, vol. 1, 1999, pp. 191-205.

[22] M. F. Goodchild and L. Li, "Assuring quality of volunteered geographic information," Spatial Stat., vol. 1, pp. 110-120, 2012.

[23] S. Openshaw, The Modifiable Areal Unit Problem, Geo Books, 1984.

[24] A. S. Fotheringham and D. W. S. Wong, "The modifiable areal unit problem in multivariate analysis," Env. Plan. A, vol. 23, no. 7, pp. 1025-1044, 1991.

[25] J. D. Herbsleb and D. Moitra, "Global software development," IEEE Softw., vol. 18, no. 2, pp. 16-20, 2001.

[26] M. Cataldo, J. D. Herbsleb, and K. M. Carley, "Socio-technical congruence," in Proc. ESEM 2008, pp. 2-11.

[27] B. Nuseibeh and S. Easterbrook, "Requirements engineering: A roadmap," in Proc. ICSE 2000, pp. 35-46.

[28] M. J. Egenhofer and D. M. Mark, "Naive geography," in Spatial Information Theory, LNCS 988, Springer, 1995, pp. 1-15.

[29] B. Smith and D. M. Mark, "Geographical categories: An ontological investigation," Int. J. Geogr. Inf. Sci., vol. 15, no. 7, pp. 591-612, 2001.

[30] M. Ruiz-Llorente and A. Ruiz, "Geospatial data harmonization for health services," Int. J. Health Geogr., vol. 13, no. 1, pp. 1-12, 2014.

[31] M. F. Goodchild, "Citizens as sensors: Volunteered geography," GeoJournal, vol. 69, no. 4, pp. 211-221, 2007.

[32] R. Zhu, K. Janowicz, L. Cai, and G. Mai, "Reasoning over spatial context for geospatial AI," Trans. GIS, vol. 26, no. 3, pp. 1222-1242, 2022.

[33] J. Z. Pan et al., Eds., Exploiting Linked Data and Knowledge Graphs, Springer, 2017.

[34] P. Hitzler, M. Krotzsch, and S. Rudolph, Foundations of Semantic Web Technologies, CRC Press, 2010.

[35] M. Palmirani and G. Governatori, "Modelling legal knowledge for GDPR," in Proc. JURIX 2018, pp. 101-110.

[36] M. J. Bommarito and D. M. Katz, "A quantitative analysis of the practice of law," J. Empirical Legal Stud., vol. 15, no. 3, pp. 626-659, 2018.

[37] H. Zhong et al., "How does NLP benefit legal system?" in Proc. ACL 2020, pp. 5218-5230.

[38] M. Dragoni et al., "Combining NLP approaches for rule extraction from legal documents," in Proc. Workshop Mining Legal Texts, 2016.

[39] N. Reimers and I. Gurevych, "Sentence-BERT," in Proc. EMNLP 2019, pp. 3982-3992.

[40] T. Mikolov, K. Chen, G. Corrado, and J. Dean, "Efficient estimation of word representations," arXiv:1301.3781, 2013.

[41] A. Cockburn, Writing Effective Use Cases, Addison-Wesley, 2001.

[42] D. Zowghi and C. Coulin, "Requirements elicitation: A survey," in Engineering and Managing Software Requirements, Springer, 2005, pp. 19-46.

[43] J. W. Ross, P. Weill, and D. C. Robertson, Enterprise Architecture as Strategy, Harvard Business Review Press, 2006.

[44] J. A. Zachman, "A framework for information systems architecture," IBM Syst. J., vol. 26, no. 3, pp. 276-292, 1987.

[45] S. Ailu, "Closing the Regulatory Interpretation Gap in Enterprise UAT," Preprint, Zenodo, 2026. doi: 10.5281/zenodo.19870009.

[46] Office of the National Coordinator for Health IT, "Report on Health Information Blocking," HHS, 2015.

[47] C. E. Witt, "Implementing the Affordable Care Act," Congressional Research Service, 2014.