

Research Article

Evidence Admissibility as a Retrieval Correctness Construct for Knowledge-Based Systems

Ankur Garg¹, Jeffrey Esposito¹

1. Independent researcher

Retrieval-augmented artificial intelligence (AI) systems increasingly support decision-making in domains such as finance, medicine, law, software engineering, and scientific research, yet contemporary retrieval architectures optimize semantic relevance without evaluating evidentiary eligibility. This paper investigates whether evidence admissibility is a distinct retrieval correctness construct, independent of semantic relevance, that determines whether retrieved information is eligible to participate in downstream reasoning. We argue that relevant information is not necessarily admissible evidence: an expired regulation, a retracted study, or an overturned legal precedent may remain topically relevant while unsuitable for evidentiary use. Through cross-domain conceptual synthesis, we formalize evidence admissibility as a knowledge representation construct and develop a domain-general ontology of its recurring dimensions. We examine the recurrence of these dimensions across finance, medicine, law, software engineering, and scientific research, and introduce an Admissibility Failure Matrix identifying evidentiary failure as a retrieval outcome distinct from conventional relevance failure. A 40-document feasibility pilot across four domains provides preliminary evidence that the construct is empirically observable and that admissibility failure occurs at substantial rates. We propose that admissibility evaluation must precede relevance ranking, formalized here as the architectural principle Eligibility Before Ranking. This work contributes a previously unformalized retrieval correctness construct, its formal ontological representation, and an architectural principle with implications for trustworthy AI, knowledge engineering, and enterprise decision support.

Correspondence: papers@team.qeios.com — Qeios will forward to the authors

1. Introduction

Retrieval-augmented artificial intelligence (AI) systems increasingly support decision-making in domains such as healthcare, finance, law, software engineering, and scientific research. In these environments, system performance depends not only on reasoning quality but also on the correctness of the information retrieved to support that reasoning. As AI systems are deployed in higher-consequence settings, retrieval correctness has become a critical challenge.

Viewed from a knowledge-engineering perspective, this is fundamentally a knowledge representation problem. Before a retrieval or reasoning system can act on retrieved information, it must be able to represent, and reason over, formal criteria that determine whether that information is eligible for evidentiary use. These criteria are not reducible to learned relevance signals: they are structured, symbolic knowledge about temporal validity, lineage, authority, and other properties that govern whether information may legitimately support a decision. Admissibility, in this sense, is not merely a retrieval filter: evidence admissibility functions as a knowledge representation layer that determines which information objects are eligible to participate in downstream reasoning, independent of their semantic relevance to a query. Formalizing this layer, rather than leaving it implicit, is the central knowledge engineering task this paper undertakes.

Contemporary retrieval architectures primarily optimize semantic relevance. Documents are retrieved and ranked according to their estimated correspondence with a query, while evaluation focuses on measures such as precision, recall, and ranking effectiveness. Implicit in these approaches is the assumption that information relevant to a question is also suitable to serve as evidence in answering that question.

However, relevant information is not necessarily admissible evidence. An expired regulation may remain relevant to a compliance inquiry, a retracted study may remain relevant to a clinical question, and an overturned legal precedent may remain relevant to a legal dispute. Although each document pertains to the query, its use as evidence may produce incorrect, unsafe, or non-compliant outcomes. These examples suggest that relevance and evidence eligibility represent distinct concepts.

Despite extensive research in information retrieval, retrieval-augmented generation, provenance, trust, and knowledge governance, existing retrieval frameworks generally treat evidence validity as a filtering or ranking consideration rather than as a separate retrieval correctness construct. Consequently, current

retrieval theory lacks an explicit construct for determining whether information is eligible to be used as evidence before relevance evaluation occurs.

This paper investigates the following research question:

Does evidence admissibility represent a distinct construct from semantic relevance, and is it a necessary prerequisite for retrieval correctness in knowledge-intensive AI systems?

To address this question, we introduce evidence admissibility as a theoretical construct governing the eligibility of information for evidentiary use, and we examine whether this construct recurs across banking, medicine, law, software engineering, and scientific research. To formalize the construct, Section 4 develops a domain-general ontology of its recurring dimensions. Based on this analysis, we argue that admissibility evaluation should precede relevance evaluation and propose admissibility as a retrieval correctness construct for trustworthy AI systems.

The contributions of this paper are fourfold:

1. We introduce evidence admissibility as a previously unformalized retrieval correctness construct, distinct from semantic relevance.
2. We develop a formal ontology representing this construct and demonstrate its recurrence across multiple domains.
3. We provide preliminary empirical evidence, from a cross-domain feasibility pilot, that the construct is observable in real documents.
4. We propose Eligibility Before Ranking, an architectural principle under which admissibility evaluation precedes relevance ranking in retrieval-augmented AI systems.

The remainder of the paper is organized as follows. Section 2 reviews related work and identifies the research gap. Section 3 formalizes admissibility as a retrieval correctness problem. Section 4 presents the proposed ontology. Section 5 examines admissibility across domains. Section 6 discusses implications for retrieval correctness. Section 7 discusses architectural implications for AI system design. Section 8 discusses limitations and future research directions, and Section 9 concludes.

2. Related Work

2.1. Information Retrieval

Information retrieval (IR) is concerned with identifying and ranking information resources that satisfy a user's information need. Since the emergence of modern retrieval theory, relevance has served as the primary construct guiding retrieval effectiveness and system evaluation [1][2]. Traditional retrieval systems estimate relevance using lexical, statistical, and probabilistic representations of document-query similarity, while evaluation frameworks assess performance through relevance-oriented measures such as precision, recall, and ranking effectiveness [3][4].

Although retrieval methodologies have evolved substantially, their underlying objective has remained largely consistent. Classical approaches such as TF-IDF and BM25 estimate relevance through term occurrence and distribution patterns [5], whereas contemporary neural retrieval systems employ learned vector representations to capture semantic relationships that extend beyond exact keyword matching [6]. [7]. Benchmarking efforts similarly evaluate retrieval quality using relevance-centered metrics including mean average precision (MAP), normalized discounted cumulative gain (nDCG), and related ranking measures [8].

The emergence of large language models has further increased the importance of retrieval systems. Retrieval-Augmented Generation (RAG) architectures integrate retrieval components with generative models to provide external knowledge during response generation [9]. Subsequent retrieval-augmented architectures have demonstrated that retrieval quality directly influences downstream reasoning and generation quality, making document selection a critical determinant of overall system performance [10][11][12].

Despite these advances, the dominant notion of retrieval correctness remains fundamentally relevance-centered. Documents are generally considered desirable when they are relevant to a user's query, and retrieval effectiveness is typically evaluated according to the system's ability to identify and rank relevant information [3][8]. Consequently, relevance continues to serve as the principal theoretical and operational criterion for retrieval success across both classical and neural retrieval paradigms.

The success of relevance-centered retrieval has enabled substantial progress in search, question answering, and retrieval-augmented AI systems. However, as retrieval increasingly serves as a

foundation for evidence-based decision support, an important question emerges: Is relevance alone sufficient to determine whether retrieved information should be relied upon?

The following sections examine related work in provenance, trust, information quality, and governance that suggests additional considerations may be necessary when retrieved information is used as evidence rather than merely as information.

2.2. Retrieval-Augmented Generation and Knowledge-Intensive AI Systems

The emergence of large language models (LLMs) has substantially expanded the role of information retrieval within AI systems. While traditional retrieval systems primarily support information access, contemporary retrieval-augmented architectures employ retrieved information as an input to reasoning, generation, and decision-support processes. As a result, retrieval quality increasingly influences not only what information is presented to users, but also how AI systems formulate responses and recommendations.

Retrieval-Augmented Generation (RAG) architectures represent one of the most widely adopted approaches for integrating external knowledge into LLM-based systems ^[9]. By retrieving relevant documents and incorporating them into the model's context, RAG systems mitigate knowledge staleness and improve performance on knowledge-intensive tasks. Subsequent architectures, including REALM ^[10], Atlas ^[11], and Self-RAG ^[12], further reinforced the central role of retrieval in modern AI systems.

A common assumption underlying these architectures is that improvements in retrieval quality lead to improvements in downstream task performance. Consequently, substantial research has focused on enhancing retrieval effectiveness through improved ranking models, dense retrieval techniques, query reformulation strategies, retrieval correction mechanisms, and hybrid search architectures. Across these approaches, retrieval effectiveness remains predominantly evaluated using relevance-oriented objectives, with the expectation that more relevant retrieved information will support better reasoning outcomes.

As retrieval systems become increasingly integrated into decision-support environments, however, retrieved documents assume a role that extends beyond information access. Retrieved information often functions as evidence supporting recommendations, explanations, compliance decisions, diagnostic assessments, and other forms of knowledge-intensive reasoning. In these settings, retrieval errors may

propagate directly into downstream outputs, causing reasoning systems to rely upon information that is incomplete, outdated, superseded, unauthorized, or otherwise inappropriate for the decision context.

Recent work addressing hallucination, trustworthiness, provenance, and retrieval robustness implicitly recognizes the importance of retrieval correctness beyond ranking performance alone. Nevertheless, most retrieval-augmented architectures continue to frame retrieval quality primarily in terms of relevance and information access. Relatively little attention has been devoted to the question of whether retrieved information is eligible to serve as evidence before it is incorporated into downstream reasoning processes.

This observation does not diminish the importance of relevance. Rather, it highlights a consequence of retrieval's evolving role within AI systems. As retrieval transitions from a mechanism for locating information to a mechanism for supplying evidence used in reasoning and decision-making, additional criteria may be required to determine whether retrieved information should be relied upon. The following section examines research streams concerned with trust, provenance, information quality, and governance that begin to address aspects of this broader challenge.

2.3. Knowledge Governance and Information Governance

Beyond information retrieval, organizations have long recognized the need to govern how information is created, maintained, accessed, and used. Research in knowledge governance, information governance, records management, and enterprise knowledge management emphasizes that information assets possess characteristics that extend beyond their informational content alone ^{[13][14]}. Questions of ownership, authority, access rights, provenance, retention, and compliance frequently influence whether information may be used within a particular organizational context.

Information governance frameworks generally focus on ensuring that information remains trustworthy, compliant, secure, and fit for organizational use throughout its lifecycle. Governance mechanisms commonly address issues such as document version control, retention schedules, access restrictions, provenance tracking, policy management, and regulatory compliance. These controls operate independently of topical relevance and are instead concerned with determining whether information satisfies organizational requirements for appropriate use ^{[15][16]}.

Recent governance initiatives in artificial intelligence have extended these concerns to AI-enabled decision-support systems. Frameworks such as the NIST AI Risk Management Framework emphasize governance, transparency, accountability, data quality, and risk management as prerequisites for

trustworthy AI deployment ^[17]. Similarly, emerging AI governance literature highlights the importance of provenance, traceability, documentation, and oversight mechanisms in ensuring reliable system behavior ^{[18][19]}.

A common theme across these governance traditions is that information is rarely treated as universally eligible for use. Instead, organizations routinely evaluate conditions under which information may be accessed, relied upon, disclosed, retained, or acted upon. Access permissions determine who may view information. Retention policies determine whether information remains valid for operational purposes. Version control mechanisms identify authoritative documents. Compliance processes establish whether information satisfies regulatory requirements. These governance decisions are fundamentally concerned with information eligibility rather than information relevance.

Despite this operational emphasis on information eligibility, governance concepts remain largely separate from retrieval theory. Retrieval systems focus on identifying information relevant to a query, while governance systems focus on controlling the conditions under which information may be used. As retrieval-augmented AI systems increasingly support organizational decision-making, these traditionally separate concerns begin to converge. Retrieved information may be relevant to a user's question, yet still violate governance constraints related to authority, provenance, currency, completeness, or access rights.

This convergence suggests an unresolved challenge at the intersection of retrieval and governance. Existing governance frameworks provide mechanisms for controlling information use, while retrieval frameworks provide mechanisms for identifying relevant information. However, neither tradition provides a unified construct for determining whether retrieved information is eligible to serve as evidence within a decision-support process. The following section examines related work in information quality, provenance, and trust that further highlights this gap.

2.4. Provenance, Trust, and Trustworthy Retrieval

As retrieval systems become embedded in knowledge-intensive AI workflows, researchers have increasingly recognized that retrieval quality cannot be reduced to topical relevance alone. Recent work on trustworthy retrieval-augmented generation (RAG) identifies concerns such as reliability, privacy, safety, fairness, explainability, accountability, transparency, robustness, factuality, and attribution as central to trustworthy system behavior ^{[20][21]}. These concerns indicate that retrieved information affects

not only whether a system can answer a query, but also whether its answer can be trusted, verified, explained, and safely acted upon.

Provenance research provides one foundation for this broader view. Provenance concerns the origin, derivation, and transformation history of information artifacts, and has long been treated as important for confidence, auditability, and evidentiary traceability [22][23]. In retrieval and RAG settings, related concerns appear in work on attribution, grounding, source transparency, and citation correctness. Recent studies emphasize that generated answers should be traceable to supporting sources and that correctness and faithfulness must be evaluated separately when deploying RAG systems in information retrieval applications [20][24]. Recent empirical work has further demonstrated that faithfulness and grounding require independent evaluation even when retrieval performance remains strong, reinforcing the distinction between relevance and evidentiary reliability [25].

A related literature on credibility and trust shows that users evaluate information using factors such as source authority, expertise, transparency, consistency, and reputation rather than topical relevance alone [26][27][28][29]. This concern has reappeared in current RAG research as a system-level trustworthiness problem. Recent trustworthy RAG surveys argue that RAG systems introduce risks involving robustness, privacy, adversarial attacks, accountability, and explainability, and that these risks require broader evaluation frameworks than relevance-based retrieval metrics alone [20][21].

Taken together, this literature provides strong evidence that relevance alone does not fully determine whether information should be relied upon in reasoning or decision-making. Information may be relevant while lacking provenance, coming from an unreliable source, failing attribution checks, violating privacy expectations, or producing unfaithful downstream outputs. Current research therefore identifies many symptoms of the problem addressed in this paper: retrieved information must satisfy conditions beyond semantic relevance before it can safely support a decision.

However, these concerns remain distributed across separate conceptual frameworks. Provenance addresses origin and lineage; trust and credibility address reliability; attribution and grounding address source support; explainability addresses transparency; and trustworthy RAG frameworks address system-level risks such as safety, privacy, fairness, robustness, and accountability. These are essential contributions, but they do not yet provide a unified construct for determining whether retrieved information is eligible to serve as evidence in a particular decision context. We therefore treat this

literature not as evidence that admissibility has already been solved, but as evidence that multiple research communities have identified related fragments of a broader evidence-eligibility problem.

2.5. Evidence-Based Decision Support

The distinction between information relevance and evidence suitability is not unique to artificial intelligence. Across medicine, public policy, scientific research, governance, and regulatory compliance, decision-support disciplines have long recognized that information must satisfy standards beyond topical relevance before it can be relied upon. As AI systems become increasingly integrated into consequential decision-making, these evidentiary considerations have become correspondingly more important.

Evidence-based medicine remains one of the most mature examples of evidence-centered decision support. While the foundational principles of evidence-based medicine emphasize the use of the best available evidence in clinical decision-making ^[30], recent research has increasingly explored how AI systems can support evidence acquisition, synthesis, dissemination, and maintenance while preserving evidentiary rigor and traceability requirements ^[31].

The continued emphasis on evidence quality, provenance, and applicability reflects a broader principle that extends beyond medicine: information must not only be relevant to a decision but also suitable for use as evidence within that decision process.

Across these decision-support disciplines, a common pattern emerges. Information is rarely relied upon solely because it is topically relevant. Instead, decision-support processes routinely evaluate additional considerations, including provenance, authority, traceability, contextual applicability, evidentiary quality, integrity, and governance requirements, before information is considered suitable for use. Although these requirements are expressed through different terminology and operational practices across domains, they consistently distinguish the suitability of information for evidentiary use from its relevance to the decision at hand.

This observation is particularly significant because it suggests that the challenge emerging in retrieval-augmented AI systems may not be unique to AI. Rather, it appears to reflect a broader characteristic of evidence-based decision-making itself. While individual disciplines employ different terminology and operational mechanisms, they repeatedly confront the same underlying question:

Under what conditions is information sufficiently suitable to function as evidence?

The following section examines whether these recurring conditions can be expressed through a single conceptual framework.

2.6. Research Gap

The preceding review reveals a notable pattern across multiple research communities. Information retrieval research has developed increasingly sophisticated mechanisms for estimating and optimizing semantic relevance. Retrieval-augmented generation research has demonstrated the importance of retrieval quality for downstream reasoning and decision support. Knowledge governance research has established the need to manage information throughout its lifecycle, while provenance, trust, credibility, and information quality research have identified numerous factors that influence whether information should be relied upon. Similarly, evidence-based decision-support disciplines routinely evaluate information according to criteria extending beyond topical relevance alone.

Taken individually, these literatures provide important insights into the conditions under which information may be used effectively and responsibly. Collectively, they suggest that information suitability depends upon considerations such as provenance, authority, timeliness, applicability, integrity, transparency, traceability, and contextual fit. Recent work in trustworthy retrieval and trustworthy RAG has further expanded attention to issues including attribution, grounding, accountability, robustness, and explainability, reinforcing the view that relevance alone may be insufficient for high-consequence decision-support applications.

Despite this convergence, the existing literature remains conceptually fragmented. Information retrieval focuses primarily on relevance. Provenance research focuses on origin and lineage. Trust and credibility research focuses on reliability and confidence. Information quality research focuses on multidimensional measures of information fitness. Governance research focuses on organizational controls and lifecycle management. Trustworthy AI research focuses on system-level properties such as accountability, transparency, robustness, and safety. While these perspectives address related concerns, they generally do so within separate conceptual frameworks and at different levels of analysis.

Recent trustworthy retrieval and trustworthy RAG frameworks have expanded evaluation beyond relevance to include attribution, grounding, faithfulness, transparency, accountability, robustness, privacy, and governance considerations [20][21][24]. These developments represent an important shift toward evaluating the suitability of retrieved information for downstream reasoning. However, these frameworks generally treat such concerns as collections of desirable system properties rather than

manifestations of a unified evidence-eligibility construct. Consequently, despite substantial progress, the literature still lacks a unifying construct that explicitly separates evidence eligibility from semantic relevance.

At its core, this gap reflects the absence of a clear distinction between two fundamentally different questions:

- 1. Is the information relevant to the question being asked?**
- 2. Is the information eligible to serve as evidence in answering that question?**

The first question has been the central concern of information retrieval research for decades. The second appears repeatedly in practice through mechanisms such as provenance controls, evidence hierarchies, governance processes, compliance requirements, attribution frameworks, and decision-support protocols. Yet despite its practical importance, it has not been formalized as a distinct retrieval construct nor systematically separated from semantic relevance within retrieval theory.

This distinction becomes increasingly important as retrieval systems transition from information-access tools to evidence providers for AI-enabled reasoning and decision support. A document may be highly relevant to a query while simultaneously being expired, superseded, inapplicable to the decision context, incomplete, unauthorized, or otherwise unsuitable for evidentiary use. Existing frameworks provide mechanisms for evaluating many of these individual concerns, but they do not provide a unified conceptual framework for evaluating evidence eligibility prior to relevance optimization.

Importantly, this paper does not argue that provenance, trust, governance, information quality, or related concepts are absent from the existing literature. Rather, these research traditions identify individual determinants of evidentiary suitability, each addressing a particular aspect of when information may be relied upon. Rather than replacing these constructs, this paper introduces an orthogonal retrieval correctness construct that organizes one recurring question they frequently address: whether retrieved information is eligible to function as evidence within a decision context. Evidence admissibility is proposed as that construct, while the Universal Admissibility Ontology provides its formal representation.

The research question introduced in Section 1 motivates the next stage of this investigation. To address that question, the following section introduces evidence admissibility as a retrieval correctness construct governing the eligibility of information for evidentiary use. We then examine whether recurring

admissibility patterns can be identified across multiple domains and whether those patterns can be formalized through the Universal Admissibility Ontology (UAO).

3. Admissibility as a Retrieval Correctness Problem

The preceding review suggests that multiple research communities—including information retrieval, retrieval-augmented generation, knowledge governance, trustworthy AI, and evidence-based decision support—have identified concerns that extend beyond semantic relevance. Although these concerns are expressed through different constructs such as provenance, attribution, trust, transparency, accountability, and information quality, they converge on a common question: under what conditions may information be relied upon in support of a decision?

We argue that this question represents a distinct retrieval correctness problem. Specifically, retrieval systems increasingly operate as providers of evidence for downstream reasoning rather than as mechanisms for information access alone. Recent trustworthy retrieval and trustworthy RAG research reflects this shift by expanding evaluation beyond relevance to include grounding, faithfulness, attribution, robustness, and accountability [20][21][24]. However, these frameworks generally evaluate such concerns as system properties rather than as manifestations of a common evidence-eligibility construct.

To address this gap, we introduce **evidence admissibility** as the degree to which information is eligible to serve as evidence within a specific decision context. We propose that retrieval correctness depends not only on whether information is relevant, but also on whether it is admissible.

3.1. Relevance Versus Admissibility

Information retrieval has historically defined success in terms of relevance—the degree to which information corresponds to a user’s information need [3][8]. Relevance therefore addresses a straightforward question:

Relevance: *Does the information pertain to the question being asked?*

Admissibility addresses a different question:

Admissibility: *May the information be used as evidence in answering that question?*

The distinction is subtle but consequential. Relevance evaluates topical correspondence between a query and an information resource. Admissibility evaluates evidentiary suitability within a particular decision context.

A document may therefore remain relevant while becoming inadmissible. An expired regulation remains relevant to a compliance inquiry but may no longer be relied upon as authoritative guidance. A retracted scientific study remains relevant to a clinical question but may no longer constitute acceptable evidence. A superseded policy remains relevant to an operational inquiry but may no longer support decision-making. In each case, relevance is preserved while evidentiary eligibility is lost.

Conversely, information may be admissible yet irrelevant. A current, authoritative, and complete document may satisfy all requirements for evidentiary use while contributing little to the question under consideration. Admissibility therefore does not imply relevance, just as relevance does not imply admissibility.

We therefore treat relevance and admissibility as orthogonal constructs. Information may be relevant and admissible, relevant but inadmissible, admissible but irrelevant, or neither relevant nor admissible. This distinction is increasingly important in retrieval-augmented AI systems because retrieved information functions as evidence for downstream reasoning processes. When retrieval outputs are used to support recommendations, decisions, or generated responses, correctness depends not only on identifying relevant information but also on ensuring that such information is eligible for evidentiary use.

Accordingly, we propose that retrieval correctness in knowledge-intensive AI systems requires two sequential evaluations: relevance determines whether information pertains to the question, while admissibility determines whether that information may legitimately be relied upon as evidence.

3.2. Relevant-but-Inadmissible Evidence

The distinction between relevance and admissibility becomes apparent when information remains topically related to a query but is unsuitable for evidentiary use. Such situations occur routinely across domains and are increasingly consequential in retrieval-augmented AI systems, where retrieved information functions as evidence for downstream reasoning rather than merely as information for user consumption ^{[20][21]}.

Examples include expired regulations that are no longer legally effective, retracted scientific studies that no longer satisfy standards of scientific reliability, deprecated software documentation that no longer reflects operational reality, and overturned legal precedents that no longer constitute valid authority. In each case, the information remains relevant to the topic under consideration but fails one or more conditions necessary for evidentiary use.

These examples suggest that retrieval failure is not limited to the retrieval of irrelevant information. In knowledge-intensive environments, systems may retrieve information that is highly relevant yet unsuitable for supporting decisions, recommendations, explanations, or generated responses. Recent trustworthy retrieval and RAG research increasingly emphasizes grounding, attribution, faithfulness, and evidentiary support as determinants of downstream reasoning quality, while empirical evaluations demonstrate that retrieved information may remain relevant despite failures in these dimensions [20][24][32].

To clarify the distinction, Table 1 presents a conceptual classification of retrieved information according to relevance and admissibility.

	Admissible	Inadmissible
Relevant	<u>Quadrant I</u> Relevant and Admissible Current regulation Valid clinical guideline Active policy document	<u>Quadrant II</u> Relevant but Inadmissible Expired regulation Retracted study Deprecated API documentation Overturned legal precedent
Irrelevant	<u>Quadrant III</u> Admissible but Irrelevant Current but out-of-scope regulation Valid but unrelated policy	<u>Quadrant IV</u> Irrelevant and Inadmissible Obsolete unrelated content Unauthorized historical artifacts

Table 1. Relationship Between Relevance and Admissibility. Relevance determines whether information pertains to a question, whereas admissibility determines whether that information may be relied upon as evidence. The existence of relevant-but-inadmissible information reveals a previously unformalized class of retrieval failure in which a system successfully retrieves topically relevant content that is nevertheless unsuitable for evidentiary use. This distinction motivates the treatment of admissibility as a separate dimension of retrieval correctness.

Traditional retrieval systems primarily distinguish between relevant and irrelevant information. The existence of Quadrant II demonstrates that this distinction is insufficient when retrieved information is intended to function as evidence. A system may successfully retrieve information that is relevant while simultaneously failing to retrieve information that is admissible.

Accordingly, retrieval correctness in knowledge-intensive AI systems cannot be defined solely in terms of relevance. Correct retrieval requires both the identification of information that pertains to the question being asked and the verification that such information is eligible for evidentiary use. This observation motivates the need to identify the dimensions that govern admissibility and to determine whether such dimensions recur across domains.

3.3. Admissibility Failure Modes

The distinction between relevance and admissibility implies that retrieval outcomes cannot be fully characterized by relevance alone. While traditional retrieval evaluation distinguishes between relevant and irrelevant information, the existence of relevant-but-inadmissible information introduces an additional class of retrieval failure. Recent trustworthy retrieval and RAG research has similarly expanded evaluation beyond relevance to include grounding, attribution, faithfulness, and evidentiary support, reflecting growing recognition that downstream reasoning depends upon more than semantic correspondence ^{[20][21][24]}.

To formalize this distinction, Table 2 presents the Admissibility Failure Matrix.

Retrieval Outcome	Relevant	Admissible	Interpretation
Correct Evidence	✓	✓	Relevant and eligible for evidentiary use.
Evidentiary Failure	✓	✗	Relevant but inadmissible.
Retrieval Miss	✗	✓	Admissible but irrelevant.
Complete Failure	✗	✗	Neither relevant nor admissible.

Table 2. Admissibility Failure Matrix. Retrieval outcomes classified according to semantic relevance and evidentiary admissibility. Traditional retrieval evaluation primarily distinguishes between relevant and irrelevant information. The introduction of admissibility reveals an additional failure mode—*evidentiary failure*—in which a system successfully retrieves relevant information that is nevertheless ineligible for evidentiary use.

The matrix highlights a limitation of relevance-centered retrieval. A system may successfully retrieve information that is topically relevant while simultaneously failing to retrieve information that is suitable for evidentiary use. In knowledge-intensive environments, such failures may expose users or AI systems to expired regulations, superseded policies, retracted studies, or other forms of inadmissible evidence. Empirical evaluations of agentic RAG systems similarly report failures involving attribution, grounding, and evidentiary support despite successful retrieval, illustrating the practical significance of evidentiary failure as a distinct class of retrieval error [32].

This observation motivates a reformulation of retrieval correctness.

Definition 1 (Retrieval Correctness). *A retrieval outcome is correct if and only if the returned information is both semantically relevant to the query and admissible as evidence within the decision context in which it is used.*

Proposition 1. *Retrieval correctness in knowledge-intensive AI systems requires both semantic relevance and evidentiary admissibility.*

Under this formulation, relevance is necessary but not sufficient for retrieval success. Correct retrieval occurs only when information both addresses the question being asked and is eligible to serve as

evidence in answering that question. The next section examines the dimensions that govern admissibility across domains.

4. A Universal Admissibility Ontology

The preceding sections established that evidence admissibility is distinct from semantic relevance and that relevant-but-inadmissible information represents a recurring retrieval failure mode across knowledge-intensive domains. These observations motivate the Universal Admissibility Ontology (UAO), a formal representation of the dimensions that govern evidentiary eligibility across contexts.

If evidence admissibility is to function as an independent retrieval correctness construct, it must be represented in a form that permits systematic evaluation across heterogeneous domains. A formal ontology provides this representation by identifying the recurring dimensions that govern evidentiary eligibility independently of any single application domain. The objective of the following section is therefore not merely to enumerate admissibility criteria, but to formalize the knowledge representation required to operationalize admissibility as a retrieval correctness construct.

4.1. *Ontology Development Method*

The ontology was developed through cross-domain conceptual synthesis. Candidate dimensions were derived from the literature reviewed in Section 2 and the admissibility failure modes identified in Section 3.

Three criteria guided dimension selection. First, each dimension had to represent a factor capable of independently affecting evidentiary eligibility. Second, each dimension had to recur across multiple domains rather than reflect a domain-specific requirement. Third, each dimension had to explain at least one class of relevant-but-inadmissible retrieval failure.

Candidate dimensions were evaluated against representative examples from healthcare, law, scientific research, software engineering, enterprise governance, and regulatory compliance. Dimensions that lacked cross-domain applicability or represented special cases of broader concepts were excluded.

The resulting ontology comprises six recurring dimensions of evidentiary eligibility. Together, these dimensions provide a framework for explaining why information may be relevant yet inadmissible and for evaluating admissibility as a construct independent of relevance.

4.2. Dimensions of the Admissibility Ontology

Application of the development method described above resulted in six recurring dimensions of evidentiary eligibility. Collectively, these dimensions determine whether information may be relied upon as evidence within a particular decision context.

1. **Temporal Validity** – Whether information is valid at the time of use.
2. **Lineage Validity** – Whether information represents the appropriate version within its evolutionary history.
3. **Contextual Scope** – Whether information applies to the specific context in which it is being used.
4. **Authority and Provenance** – Whether information originates from a legitimate and authoritative source.
5. **Access and Entitlement** – Whether information is eligible to be used by the requesting entity.
6. **Completeness and Integrity** – Whether information is sufficiently complete, intact, and reliable for evidentiary use.

These dimensions are proposed as independent but complementary determinants of admissibility. Failure of any dimension may render information inadmissible despite its semantic relevance to a query. Together, they provide a framework for evaluating evidentiary eligibility separately from relevance and for explaining recurring classes of relevant-but-inadmissible retrieval failures observed across domains.

4.3. The Proposed Admissibility Ontology

Table 3 presents the proposed admissibility ontology. The ontology consists of six dimensions that collectively determine evidentiary eligibility: temporal validity, lineage validity, contextual scope, authority and provenance, access and entitlement, and completeness and integrity.

The dimensions are complementary and independently evaluable. Failure of any dimension may render information inadmissible regardless of its semantic relevance.

Together, the six dimensions provide a structured framework for evaluating admissibility and explaining relevant-but-inadmissible retrieval outcomes. The ontology therefore serves as a domain-independent model of evidentiary eligibility for knowledge-intensive AI systems.

Dimension	Admissibility Question	Example Failure
Temporal Validity	Is the information valid at the time of use?	Expired regulation
Lineage Validity	Is this the correct version within its evolution history?	Superseded policy
Contextual Scope	Does the information apply to the current context?	Guidance for another jurisdiction
Authority & Provenance	Does the information originate from an authoritative source?	Unverified third-party content
Access & Entitlement	Is the information eligible for use by this requester?	Restricted document
Completeness & Integrity	Is the information sufficiently complete and reliable?	Retracted study

Table 3. Proposed Admissibility Ontology. The ontology organizes evidentiary eligibility into six recurring dimensions. Each dimension evaluates a distinct condition under which information may become inadmissible despite remaining semantically relevant.

The ontology does not claim to enumerate every possible domain-specific admissibility rule. Rather, it identifies six recurring admissibility dimensions that were consistently observed across the domains examined in this study. Although additional domain-specific admissibility criteria may exist, these six dimensions were sufficient to characterize every relevant-but-inadmissible retrieval failure observed during the feasibility assessment.

4.4. Contextual Scope

Contextual scope concerns whether information is applicable to the specific circumstances of a decision. Information may be current, authoritative, complete, and otherwise admissible while remaining unsuitable because it falls outside the context in which the decision is being made. Admissibility therefore depends not only on the properties of information itself but also on the conditions under which that information is applied.

Although each admissibility dimension addresses a different aspect of evidentiary eligibility, they share one important characteristic: none is reducible to semantic relevance. Relevance determines whether retrieved information pertains to a query, whereas admissibility determines whether that information is suitable to function as evidence within a particular decision context. Consequently, information may remain highly relevant while simultaneously becoming inadmissible for different reasons depending on which admissibility dimension is violated. The dimensions that follow identify the principal mechanisms through which evidence eligibility is established across domains.

Scope-based admissibility constraints recur across multiple domains. Regulatory guidance applicable to one jurisdiction may be unsuitable in another. Clinical recommendations developed for specific patient populations may not generalize across demographic or treatment contexts. Likewise, retrieval and decision-support systems increasingly operate across heterogeneous users, organizations, and environments, where context-sensitive applicability directly affects the validity of retrieved evidence [32]. Recent retrieval and retrieval-augmented generation (RAG) research similarly emphasizes contextual grounding, task alignment, and environment-specific evidence selection for reliable downstream reasoning [20][21].

Within the proposed ontology, contextual scope determines whether otherwise admissible information is applicable within the operational context of the decision being supported.

4.5. Trustworthiness Dimensions of Admissibility

Beyond temporal validity, lineage validity, and contextual scope, admissibility also depends upon the trustworthiness of evidence itself. Whereas the preceding dimensions determine whether evidence is temporally and contextually applicable, the following dimensions determine whether that evidence may legitimately be relied upon. Collectively, they address three complementary questions:

- Does the evidence originate from an authoritative and traceable source?
- Is the evidence eligible for use by the requester?
- Is the evidence sufficiently complete and reliable to support a decision?

The proposed ontology formalizes these trustworthiness concerns through the dimensions of authority and provenance, access and entitlement, and completeness and integrity.

4.5.1. Authority and Provenance

Authority and provenance concern the origin and traceability of evidence. Information may be current, applicable, and highly relevant while remaining inadmissible because its source cannot be verified or because it does not originate from an appropriate authority. Admissibility therefore depends not only on informational content but also on source legitimacy and evidentiary traceability ^{[22][23]}.

Concerns regarding authority and provenance recur across multiple domains. Regulatory decisions depend upon officially issued guidance rather than unofficial interpretations. Scientific conclusions rely upon traceable and verifiable sources of evidence. Enterprise governance frameworks emphasize lineage, accountability, and auditability, while modern retrieval and AI systems increasingly incorporate attribution and source-grounding mechanisms to support trustworthy decision-making ^{[20][24]}.

Within the proposed ontology, authority and provenance determine whether evidence originates from a source that can be verified and legitimately relied upon in support of a decision.

4.5.2. Access and Entitlement

Access and entitlement concern whether information is eligible for use by a particular requester. Recent permission-aware retrieval research emphasizes that access decisions depend not only on information relevance but also on the identity, authorization, and privileges of the requesting entity ^[33]. Information may therefore be current, authoritative, complete, and highly relevant while remaining inadmissible because the requester lacks the rights necessary to access or rely upon it.

Access restrictions recur across multiple domains. Healthcare systems restrict access to patient information, legal environments protect privileged materials, and enterprise knowledge systems routinely enforce role-based access controls. Modern retrieval and RAG architectures have similarly introduced permission-aware retrieval mechanisms to ensure that retrieved information satisfies identity, authorization, and entitlement requirements ^{[33][34]}.

Within the proposed ontology, access and entitlement determine whether evidence is eligible for use by a specific requester within a particular decision context.

4.5.3. Completeness and Integrity

Completeness and integrity concern whether evidence is sufficiently complete, accurate, and reliable to support a decision. Information quality research has long recognized these properties as essential

determinants of information fitness, while recent evidence-maintenance and retrieval research emphasizes preserving evidentiary quality throughout the information lifecycle [\[31\]\[35\]](#). Information may therefore be current, authoritative, applicable, and highly relevant while remaining inadmissible because it is incomplete, corrupted, altered, or otherwise unreliable.

Failures of completeness and integrity recur across multiple domains. Retracted scientific studies, incomplete medical records, corrupted datasets, and partially documented policies may remain relevant to a question while being unsuitable for evidentiary use. Recent evaluations of retrieval-augmented systems similarly emphasize evidentiary quality, faithfulness, and grounding in determining whether retrieved information can support downstream reasoning [\[25\]](#).

Within the proposed ontology, completeness and integrity determine whether evidence is sufficiently complete and reliable to support a decision.




Figure 1. Universal Admissibility Ontology (UAO). The proposed ontology conceptualizes evidence admissibility as a multidimensional retrieval correctness construct comprising six recurring dimensions. Temporal validity, lineage validity, and contextual scope establish the temporal and contextual eligibility of evidence, while authority and provenance, access and entitlement, and completeness and integrity establish its trustworthiness. Collectively, these dimensions determine whether retrieved information is eligible to function as evidence within a specific decision context independent of semantic relevance.

5. Cross-Domain Manifestation Analysis

This section examines whether the admissibility dimensions identified in Section 4 recur across otherwise unrelated domains. If admissibility represents a general retrieval construct rather than a domain-specific artifact, similar evidentiary eligibility patterns should emerge despite substantial differences in subject matter, governance structures, and decision processes.

Dimension	Banking	Medicine	Law	Software Engineering	Scientific Research
Temporal Validity	Effective dates	Guideline currency	Current precedent	Supported version	Current evidence
Lineage Validity	Supersession chain	Guideline revision history	Appeals / precedent lineage	Version lineage	Citation lineage
Contextual Scope	Jurisdiction/product	Patient population	Jurisdiction/case type	Platform/environment	Experimental scope
Authority & Provenance	Regulator source	Issuing body	Court authority	Official vendor	Peer-reviewed source
Access & Entitlement	Internal policy access	Patient privacy	Privileged information	License tier	Data restrictions
Completeness & Integrity	Complete policy	Complete record	Complete filing	Complete documentation	Complete dataset

Table 4. Admissibility Convergence Matrix. Recurring manifestations of the proposed admissibility dimensions across banking, medicine, law, software engineering, and scientific research. Despite substantial differences in subject matter, governance structures, and decision processes, each domain exhibits mechanisms corresponding to temporal validity, lineage validity, contextual scope, authority and provenance, access and entitlement, and completeness and integrity. This convergence supports the proposition that admissibility represents a general evidentiary construct distinct from semantic relevance.

As shown in Table 4, all six admissibility dimensions recur across banking, medicine, law, software engineering, and scientific research despite substantial differences in subject matter, governance structures, and operational objectives. Although these domains employ different terminology and mechanisms, they exhibit remarkably similar evidentiary eligibility patterns.

Temporal validity appears wherever information possesses a bounded period of applicability. Banking regulations become effective and expire, clinical guidance evolves, legal precedents change through

subsequent rulings, software versions reach end-of-support, and scientific conclusions are revised as new evidence emerges [\[17\]\[19\]\[30\]](#). In each case, admissibility depends on whether information remains current within the relevant decision context.

Lineage validity similarly recurs through mechanisms governing supersession, revision, and version control. Regulatory guidance, clinical recommendations, legal precedents, software artifacts, and scientific publications all maintain explicit or implicit lineage relationships that determine which version of information is authoritative [\[22\]\[23\]](#).

Contextual scope governs whether otherwise valid information applies within a specific environment. Jurisdiction, patient population, case type, platform configuration, and experimental conditions each function as scope constraints that determine evidentiary applicability. Information may therefore remain accurate and authoritative while being inadmissible outside its intended context [\[32\]](#).

The trustworthiness dimensions show similar convergence. Authority and provenance establish source legitimacy and traceability, access and entitlement govern permissible use, and completeness and integrity ensure evidentiary reliability [\[22\]\[23\]\[25\]\[33\]\[34\]\[35\]](#). Although implemented through different governance mechanisms, these dimensions serve analogous functions across all five domains.

Collectively, these observations suggest that admissibility is not a domain-specific artifact. Rather, the same evidentiary eligibility mechanisms recur across otherwise unrelated decision-support environments. The consistency of these patterns supports the proposition that admissibility represents a general construct governing whether information may function as evidence across domains.

5.1. Empirical Feasibility Assessment

Having established the construct's cross-domain generality through conceptual analysis (Section 5), we next assessed whether evidence admissibility is also empirically observable by conducting a small feasibility pilot rather than a full validation study. Forty documents were examined across four public knowledge domains (finance, medicine, law, and software engineering; ten per domain), deliberately selected to include both known admissibility failures (e.g., retracted studies, overruled precedents, end-of-life software) and documents with no known issues. One domain (finance) was sampled via live retrieval from SEC EDGAR; the remaining three domains were assessed using documented, publicly verifiable facts about specific, well-known documents. Each document was coded across the six

admissibility dimensions using a three-point ordinal scale (0 = not triggered, 1 = triggered, 2 = explicit failure). Because this was a single-coder feasibility pilot, inter-rater reliability was not assessed.

Dimension	Finance	Medicine	Law	Software
Temporal Validity	10%	30%	40%	40%
Lineage Validity	50%	30%	40%	40%
Contextual Scope	10%	10%	30%	20%
Authority and Provenance	0%	20%	30%	10%
Access and Entitlement	0%	0%	10%	20%
Completeness and Integrity	20%	20%	0%	20%

Table 5. Cross-Domain Observability of Admissibility Failures (% of documents with score = 2). Percentage of documents exhibiting failures across the six admissibility dimensions within each evaluation domain.

Admissibility failures were observed in all four domains despite the retrieved documents remaining semantically relevant, providing preliminary feasibility evidence that the proposed admissibility dimensions are empirically observable across diverse knowledge domains.

In every domain, documents exhibiting admissibility failures remained semantically relevant to the motivating query. Failures arose from temporal expiration, version supersession, jurisdictional or population mismatch, insufficient authority, access restrictions, or incompleteness rather than from a lack of semantic relevance. These observations are consistent with the relevant-but-inadmissible failure mode introduced in Section 3.2 and provide preliminary evidence that evidence admissibility is empirically observable across diverse knowledge domains.

Given the small, non-random sample and the reliance on documented facts for three of the four domains, these findings should be interpreted as feasibility evidence of construct observability rather than prevalence estimates or validation of a measurement instrument. A larger, randomly sampled, multi-coder evaluation is proposed as future work in Section 8.

6. Implications for Retrieval Correctness

Because admissibility recurs independently across domains, retrieval correctness must extend beyond semantic relevance alone. A retrieval system that optimizes only semantic relevance may still retrieve information that is expired, superseded, unauthorized, incomplete, or otherwise ineligible for evidentiary use. If admissibility is accepted as an independent retrieval correctness construct, retrieval architectures optimized solely for semantic relevance remain incomplete for evidence-based decision support.

To illustrate the implications of admissibility failure, Table 6 presents representative examples drawn from the dimensions identified in the proposed ontology.

Failed Dimension	Representative Example	Retrieval Result	Consequence
Temporal Validity	Expired regulation	Relevant information retrieved	Incorrect compliance recommendation
Lineage Validity	Superseded policy	Relevant information retrieved	Obsolete guidance presented
Contextual Scope	Wrong jurisdiction	Relevant information retrieved	Inapplicable recommendation
Authority & Provenance	Unofficial source	Relevant information retrieved	Unsupported conclusion
Access & Entitlement	Restricted record	Relevant information retrieved	Governance or privacy violation
Completeness & Integrity	Partial dataset	Relevant information retrieved	Distorted reasoning

Table 6. Admissibility Failure Propagation. Retrieval failure is not limited to the retrieval of irrelevant information. In each example, the retrieved information remains semantically related to the query and may therefore satisfy conventional relevance-oriented evaluation criteria. Nevertheless, the resulting evidence is unsuitable for decision support because one or more admissibility dimensions have failed.

This observation exposes a limitation of relevance-centered retrieval evaluation. Metrics such as precision, recall, mean average precision, and ranking quality assess whether systems identify information related to a query [3][8]. They do not directly evaluate whether that information is eligible to function as evidence within a decision process. Consequently, retrieval systems may achieve strong relevance performance while simultaneously exposing users or downstream AI systems to inadmissible evidence.

The distinction becomes increasingly important in retrieval-augmented AI systems, where retrieved information frequently serves as evidence supporting recommendations, explanations, decisions, and generated responses [9][12]. Recent trustworthy retrieval and RAG research similarly emphasizes grounding, attribution, faithfulness, and evidentiary support as critical determinants of system reliability [20][24][32]. Under these conditions, admissibility failures propagate beyond document retrieval into downstream reasoning processes.

Collectively, the cross-domain recurrence of admissibility dimensions and the consequences of admissibility failure support the central proposition of this paper:

Proposition 1. *Retrieval correctness in knowledge-intensive AI systems requires both semantic relevance and evidentiary admissibility.*

Under this formulation, relevance remains necessary but is not sufficient for retrieval success. Correct retrieval occurs only when information both addresses the question being asked and is eligible to function as evidence within the relevant decision context. Admissibility should therefore be viewed not as an additional ranking feature, but as a prerequisite condition governing whether retrieved information may legitimately participate in downstream reasoning and decision support. Operationally, this reformulation is sequential rather than merely conjunctive: knowledge sources are first gated by admissibility evaluation, and only the resulting eligible evidence is subject to semantic relevance ranking before reaching reasoning and decision support, the architecture developed in Section 7.

7. Implications for AI System Design

The distinction between relevance and admissibility has implications extending beyond retrieval evaluation. If admissibility represents a prerequisite condition for retrieval correctness, future AI systems

must treat evidentiary eligibility as an explicit architectural concern rather than an indirect consequence of relevance optimization.

For retrieval architectures, this perspective suggests that admissibility evaluation should precede ranking and generation, an architectural principle we term Eligibility Before Ranking. Rather than optimizing over all retrievable information, systems should first determine whether candidate evidence satisfies admissibility requirements before allowing it to participate in downstream reasoning processes.

Whether statistical or neural models can successfully predict admissibility is a separate question from whether admissibility has been formally represented as a retrieval correctness construct. This paper addresses the latter. Learning that a document is frequently useful, trustworthy, or cited does not itself formalize the conditions under which that document is eligible to function as evidence. Eligibility Before Ranking therefore requires explicit representation of admissibility so that evidence eligibility can be evaluated before ranking and downstream reasoning occur.

Similar principles already appear in governance-oriented architectures that enforce policy, access, and compliance constraints prior to system action ^{[17][19]}. Figure 2 illustrates this architectural sequence.




Figure 2. Eligibility Before Ranking: The Admissibility-Gated Retrieval Pipeline. Candidate evidence drawn from knowledge sources first passes through admissibility evaluation, which filters out temporally invalid, superseded, out-of-scope, unauthoritative, inaccessible, or incomplete documents. Only the resulting eligible evidence is then subjected to semantic relevance ranking, the stage conventional retrieval architectures treat as the primary (and often only) filtering step, before passing to reasoning and, ultimately, decision support. Positioning admissibility evaluation as a prerequisite stage, rather than a ranking feature, ensures that inadmissible documents cannot reach reasoning or decision support regardless of their semantic relevance to the query.

The proposed ontology also has implications for trustworthy AI and enterprise knowledge management. Recent work emphasizes grounding, attribution, transparency, accountability, and evidentiary support as

critical determinants of trustworthy AI behavior [20][24]. From an admissibility perspective, many of these concerns may be interpreted as mechanisms for ensuring evidentiary eligibility rather than merely improving retrieval quality. Similarly, knowledge governance frameworks routinely enforce controls related to validity, provenance, entitlement, and information quality, all of which correspond to dimensions of the proposed ontology.

Finally, the ontology suggests several directions for future research, including admissibility measurement frameworks, admissibility enforcement mechanisms, governance-oriented retrieval architectures, and human-AI collaboration models that explicitly distinguish evidentiary eligibility from semantic relevance. Collectively, such efforts may contribute to retrieval systems that are not only more relevant, but also more reliable.

8. Limitations and Research Roadmap

This paper introduces admissibility as a theoretical construct distinct from semantic relevance and proposes the UAO, derived through cross-domain conceptual synthesis, as its formal representation, supported by an initial feasibility pilot. Consistent with theory-building research, the primary objective remains conceptual development rather than large-scale empirical validation [36][37], and the pilot should be understood as preliminary support for feasibility rather than as validation of the ontology's measurement properties. Consequently, several limitations remain.

First, empirical support for the ontology is currently limited to a small feasibility pilot rather than a validated measurement instrument. The pilot reported in Section 5.5 indicates that the six dimensions can be identified in real documents and that admissibility failure is observable rather than purely hypothetical, but it relied on a single coder, a small and deliberately non-random sample, and, for three of the four domains, assessment of documented facts about known cases rather than live, blinded corpus retrieval. Consequently, the reported triggering and failure rates should be treated as preliminary feasibility evidence, not as prevalence estimates or a validated measurement protocol. This work does not yet provide a psychometrically validated instrument for measuring admissibility or evaluating its impact on retrieval performance at scale. Future research should focus on a larger, randomly sampled, multi-coder study with formal inter-rater reliability assessment, followed by the development of admissibility metrics and evaluation frameworks.

Second, the ontology has been examined across a limited set of representative domains. Although banking, medicine, law, software engineering, and scientific research exhibit strong convergence, additional investigation is needed to determine whether similar admissibility patterns emerge in other knowledge-intensive environments.

Third, the ontology should be viewed as an initial theoretical model rather than a finalized taxonomy. Additional dimensions, refinements, or domain-specific specializations may emerge as empirical evidence accumulates. Continued validation is therefore necessary to assess the completeness and stability of the proposed framework.

These limitations suggest a natural research progression. The present work establishes the existence of admissibility as a distinct retrieval construct. Subsequent research must determine how admissibility can be measured, enforced, created, and scaled within operational AI systems.

Figure 3 summarizes the proposed research roadmap.




Figure 3. Research Roadmap for Admissibility. The proposed roadmap outlines a progression from conceptual identification of admissibility as a distinct retrieval construct to its operationalization within AI systems. Future research should focus on measuring admissibility, enforcing admissibility constraints during retrieval and reasoning, creating and maintaining admissible evidence, and scaling admissibility-aware capabilities across enterprise and regulatory environments.

The existence stage establishes admissibility as a construct distinct from relevance. The measurement stage develops quantitative methods for evaluating admissibility and its dimensions. The enforcement stage investigates mechanisms for preventing inadmissible evidence from participating in retrieval and reasoning processes. The creation stage explores how admissibility-aware systems can generate, maintain, and curate admissible evidence over time. Finally, the scaling stage examines how admissibility can be operationalized across large-scale enterprise, regulatory, and AI ecosystems.

Collectively, these stages define a broader research agenda for advancing admissibility from a theoretical construct to an operational capability supporting trustworthy AI, knowledge governance, and evidence-based decision support.

9. Conclusion

This paper introduced evidence admissibility as a construct distinct from semantic relevance and examined its implications for retrieval correctness in knowledge-intensive AI systems. Through conceptual analysis and cross-domain synthesis, we argued that systems may successfully retrieve information that is semantically relevant yet unsuitable for evidentiary use.

Accordingly, this paper advances Proposition 1: retrieval correctness requires both semantic relevance and evidentiary admissibility, evaluated sequentially rather than merely conjunctively. Admissibility evaluation determines which information is eligible to participate in reasoning at all; relevance then determines which eligible evidence is prioritized. We formalize this ordering as an architectural principle, Eligibility Before Ranking, for retrieval-augmented AI systems.

As AI systems increasingly participate in reasoning, decision support, and knowledge-intensive workflows, this distinction becomes increasingly important. Evidence admissibility functions as a knowledge representation layer that determines which information is eligible to participate in reasoning at all: before an AI system can determine whether information is relevant, it must first determine whether that information is admissible as evidence.

Statements and Declarations

Use of Generative AI

During the preparation of this manuscript, the authors used OpenAI's ChatGPT to assist with language refinement, structural organization, and manuscript preparation. All scientific concepts, theoretical contributions, analyses, interpretations, and final editorial decisions were developed, verified, and approved by the authors, who take full responsibility for the content of this publication.

References

1. [^]Salton G, McGill MJ (1983). *Introduction to Modern Information Retrieval*. McGraw-Hill.

2. ^aRobertson SE, Sparck Jones K (1976). "Relevance Weighting of Search Terms." *J Am Soc Inf Sci.* 27(3):129–146. doi:[doi:10.1002/asi.4630270302](https://doi.org/10.1002/asi.4630270302).
3. ^{a, b, c, d}Manning CD, Raghavan P, Schütze H (2008). *Introduction to Information Retrieval*. Cambridge University Press.
4. ^aBaeza-Yates R, Ribeiro-Neto B (2011). *Modern Information Retrieval: The Concepts and Technology Behind Search*. 2nd ed. Addison-Wesley.
5. ^aRobertson SE, Walker S (1994). "Some Simple Effective Approximations to the 2-Poisson Model for Probabilistic Weighted Retrieval." In W. B. Croft & C. J. van Rijsbergen (Eds.), *Proceedings of the 17th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval* (pp. 232–241). Springer-Verlag. doi:[doi:10.1007/978-1-4471-2099-5_24](https://doi.org/10.1007/978-1-4471-2099-5_24).
6. ^aKarpukhin V, Oguz B, Min S, Lewis P, Wu L, Edunov S, Chen D, Yih W (2020). "Dense Passage Retrieval for Open-Domain Question Answering." In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)* (pp. 6769–6781). Association for Computational Linguistics. doi:[doi:10.18653/v1/2020.emnlp-main.550](https://doi.org/10.18653/v1/2020.emnlp-main.550).
7. ^aKhattab O, Zaharia M (2020). "ColBERT: Efficient and Effective Passage Search via Contextualized Late Interaction over BERT." In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '20)* (pp. 39–48). Association for Computing Machinery. doi:[doi:10.1145/3397271.3401075](https://doi.org/10.1145/3397271.3401075).
8. ^{a, b, c, d}Thakur N, Reimers N, Rücklé A, Srivastava A, Gurevych I (2021). "BEIR: A Heterogeneous Benchmark for Zero-Shot Evaluation of Information Retrieval Models." In M. Ranzato, A. Beygelzimer, Y. Dauphin, P. S. Liang, & J. W. Vaughan (Eds.), *Advances in Neural Information Processing Systems* (Vol. 34, pp. 15285–15299). Curran Associates, Inc.
9. ^{a, b, c}Lewis P, Perez E, Piktus A, Petroni F, Karpukhin V, Goyal N, Küttler H, Lewis M, Yih W-t, Rocktäschel T, Riedel S, Kiela D (2020). "Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks." *Advances in Neural Information Processing Systems*, 33, 9459–9474. https://papers.nips.cc/paper_files/paper/2020/hash/6b493230205f780e1bc26945df7481e5-Abstract.html.
10. ^{a, b}Guu K, Lee K, Tung Z, Pasupat P, Chang M-W (2020). "REALM: Retrieval-Augmented Language Model Pre-Training." In H. Daumé III & A. Singh (Eds.), *Proceedings of the 37th International Conference on Machine Learning (ICML 2020)* (Vol. 119, pp. 3929–3938). PMLR. <https://proceedings.mlr.press/v119/guu20a.html>.
11. ^{a, b}Izacard G, Grave E, Caron M, Bevilacqua M, Joshi M, Lewis P, Dupoux E, Zhou A, Simig D, Bordes A, et al. (2023). "Atlas: Few-Shot Learning with Retrieval Augmented Language Models." *J Mach Learn Res.* 24(251):

- 1–43. <https://jmlr.org/papers/v24/23-0037.html>.
12. ^{a, b, c}Asai A, Min S, Zhong Z, Chen D (2024). "Self-RAG: Learning to Retrieve, Generate, and Critique Through Self-Reflection." In *Proceedings of the International Conference on Learning Representations (ICLR 2024)*. https://proceedings.iclr.cc/paper_files/paper/2024/file/25f7be9694d7b32d5cc670927b8091e1-Paper-Conference.pdf.
 13. ^ΔAlavi M, Leidner DE (2001). "Review: Knowledge Management and Knowledge Management Systems." *MIS Q.* 25(1):107–136.
 14. ^ΔWeill P, Ross JW (2004). *IT Governance: How Top Performers Manage IT Decision Rights for Superior Results*. Harvard Business School Press.
 15. ^ΔSmallwood RF (2014). *Information Governance: Concepts, Strategies, and Best Practices*. Wiley.
 16. ^ΔFranks PC (2018). *Records and Information Management*. ALA Neal-Schuman.
 17. ^{a, b, c}National Institute of Standards and Technology (2023). "Artificial Intelligence Risk Management Framework (AI RMF 1.0) (NIST AI 100-1)." U.S. Department of Commerce. doi:doi:10.6028/NIST.AI.100-1.
 18. ^ΔShneiderman B (2022). *Human-Centered AI*. Oxford University Press.
 19. ^{a, b, c}International Organization for Standardization (2023). "Information Technology — Artificial Intelligence — Management System (ISO/IEC 42001:2023)." ISO/IEC 42001:2023 - AI management systems.
 20. ^{a, b, c, d, e, f, g, h, i, j, k, l}Ni B, Liu Z, Wang L, Lei Y, Zhao Y, Cheng X, Zeng Q, Dong L, Xia Y, Kenthapadi K, Rossi R, Derroncourt F, Tanjim MM, Ahmed N, Liu X, Fan W, Blasch E, Wang Y, Jiang M, Derr T (2025). "Towards Trustworthy Retrieval Augmented Generation for Large Language Models: A Survey." *arXiv*. doi:doi:10.48550/arXiv.2502.06872.
 21. ^{a, b, c, d, e, f, g}Zhou Y, Liu Y, Li X, Jin J, Qian H, Liu Z, Li C, Dou Z, Ho T-Y, Yu PS (2026). "Trustworthiness in Retrieval-Augmented Generation Systems: A Survey (Version 2)." *arXiv*. doi:doi:10.48550/arXiv.2409.10102.
 22. ^{a, b, c, d}Buneman P, Khanna S, Tan WC (2001). "Why and Where: A Characterization of Data Provenance." In J. Van den Bussche & V. Vianu (Eds.), *Proceedings of the 8th International Conference on Database Theory (ICDT 2001)* (pp. 316–330). Springer. doi:doi:10.1007/3-540-44503-X_20.
 23. ^{a, b, c, d}Simmhan YL, Plale B, Gannon D (2005). "A Survey of Data Provenance in E-Science." *SIGMOD Rec.* 34(3):31–36. doi:doi:10.1145/1084805.1084812.
 24. ^{a, b, c, d, e, f, g, h}Nematov I, Kalai T, Kuzmenko E, Fugagnoli G, Sacharidis D, Hose K, Sagi T (2025). "Source Attribution in Retrieval-Augmented Generation [Preprint]." *arXiv*. doi:doi:10.48550/arXiv.2507.04480.
 25. ^{a, b, c}Tamber MS, Bao FS, Xu C, Luo G, Kazi S, Bae M, Li M, Mendelevitch O, Qu R, Lin J (2025). "Benchmarking LLM Faithfulness in RAG with Evolving Leaderboards." In *Proceedings of EMNLP 2025: Industry Track* (p

p. 799–811). <https://aclanthology.org/2025.emnlp-industry.54/>.

26. [△]Fogg BJ, Tseng H (1999). "The Elements of Computer Credibility." In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '99)* (pp. 80–87). Association for Computing Machinery. doi:[doi:10.1145/302979.303001](https://doi.org/10.1145/302979.303001).
27. [△]McKnight DH, Chervany NL (2001). "What Trust Means in E-Commerce Customer Relationships: An Interdisciplinary Conceptual Typology." *Int J Electron Commerce*. 6(2):35–59. doi:[doi:10.1080/10864415.2001.11044235](https://doi.org/10.1080/10864415.2001.11044235).
28. [△]Corritore CL, Kracher B, Wiedenbeck S (2003). "On-Line Trust: Concepts, Evolving Themes, a Model." *Int J Hum Comput Stud*. 58(6):737–758. doi:[doi:10.1016/S1071-5819\(03\)00041-7](https://doi.org/10.1016/S1071-5819(03)00041-7).
29. [△]Metzger MJ (2007). "Making Sense of Credibility on the Web: Models for Evaluating Online Information and Recommendations for Future Research." *J Am Soc Inf Sci Technol*. 58(13):2078–2091. doi:[doi:10.1002/asi.20672](https://doi.org/10.1002/asi.20672).
30. [△][♢]Sackett DL, Rosenberg WMC, Gray JAM, Haynes RB, Richardson WS (1996). "Evidence Based Medicine: What It Is and What It Isn't." *BMJ*. 312(7023):71–72. doi:[doi:10.1136/bmj.312.7023.71](https://doi.org/10.1136/bmj.312.7023.71).
31. [△][♢]Rogers EA, McGill E, Kelly MP, Sworn K, Yoganathan M, Moffat KR, Bambra C, Whitehead M, Bambra C (2025). "Leveraging AI to Optimize Maintenance of Health Evidence and Offer a One-Stop Shop for Quality-Appraised Evidence Syntheses on the Effectiveness of Public Health Interventions." *J Med Internet Res*. 27:e69700. doi:[doi:10.2196/69700](https://doi.org/10.2196/69700).
32. [△][♢][♣][♠][♣]Papageorgiou G, Sarlis V, Maragoudakis M, Tjortjis C (2025). "Evaluating Faithfulness in Agentic RAG Systems for E-Governance Applications Using LLM-Based Judging Frameworks." *Big Data Cogn Comput*. 9(12):309. doi:[doi:10.3390/bdcc9120309](https://doi.org/10.3390/bdcc9120309).
33. [△][♢][♣]Jeong J, Lee SG (2025). "Permission-Aware RAG: Identity and Access Management (IAM)-Based Access Filtering in Multi-Resource Environments." *IEEE Access*. 13:192819–192835. doi:[doi:10.1109/ACCESS.2025.3628960](https://doi.org/10.1109/ACCESS.2025.3628960).
34. [△][♢]Chen B, Tackman J, Setälä M, Poranen T, Zhang Z (2025). "Integrating Access Control with Retrieval-Augmented Generation: A Proof of Concept for Managing Sensitive Patient Profiles." In *Proceedings of the 40th ACM/SIGAPP Symposium on Applied Computing (SAC '25)* (pp. 915–919). Association for Computing Machinery. doi:[doi:10.1145/3672608.3707848](https://doi.org/10.1145/3672608.3707848).
35. [△][♢]Wang RY, Strong DM (1996). "Beyond Accuracy: What Data Quality Means to Data Consumers." *J Manage Inf Syst*. 12(4):5–33. doi:[doi:10.1080/07421222.1996.11518099](https://doi.org/10.1080/07421222.1996.11518099).

36. ^ΔGregor S (2006). "The Nature of Theory in Information Systems." *MIS Q.* **30**(3):611–642. doi:[doi:10.2307/25148742](https://doi.org/10.2307/25148742).
37. ^ΔWhetten DA (1989). "What Constitutes a Theoretical Contribution?" *Acad Manage Rev.* **14**(4):490–495. doi:[doi:10.5465/amr.1989.4308371](https://doi.org/10.5465/amr.1989.4308371).

Declarations

Funding: No specific funding was received for this work.

Potential competing interests: No potential competing interests to declare.