Engineering Inferential Composition Control for Federated RAG in Data Spaces

Carlos Mario Braga[1], **Manuel A. Serrano**[1], **Eduardo Fernández-Medina**[2]

[1]Alarcos Research Group – University of Castilla-La Mancha – Ciudad Real – Spain

[2]GSyA Research Group – University of Castilla-La Mancha – Ciudad Real – Spain

carlosmario.braga1@alu.uclm.es,{manuel.serrano,eduardo.fdezmedina}@uclm.es

Abstract. *Federated Data Spaces increasingly rely on Retrieval-Augmented Generation (RAG) to provide intelligent access to distributed and sovereign data sources. However, current RAG systems typically assume that retrieved fragments from different providers can be merged into a single generation context without semantic conflict. In regulated and heterogeneous settings, this assumption can trigger inferential errors, such as treating conditional or alternative requirements as jointly mandatory, posing risks to trustworthiness and governance. This paper reframes semantic interoperability in federated RAG systems as a problem of inferential composition rather than conceptual alignment. We introduce a lightweight engineering framework that adds an explicit reasoning control layer to the RAG pipeline. The framework uses semantic roles and pragmatic signals to guide selective context separation and conservative fragment composition, preventing unsafe inferences without requiring ontologies or shared vocabularies. We evaluate the approach through a reproducible case study in a regulated data-sharing scenario, showing that inferential composition control prevents reasoning failures under adversarial queries, while preserving benign integration. These findings highlight the need for explicit reasoning control in RAG-based systems operating in federated and regulated environments.*

1. Introduction

Data Spaces (DSs) are emerging as a paradigm for information sharing among independent organizations, particularly in regulated domains such as healthcare, industry, public administration, and finance [Möller et al. 2024, Penedo 2024]. Their core objective is to enable federated and sovereign data ecosystems in which data providers retain control over data usage and preserve organizational and semantic autonomy, avoiding centralized data management [Möller et al. 2024, Braga et al. 2026]. In the absence of mandatory global semantic alignment, DSs are therefore characterized by the coexistence of heterogeneous interpretative frameworks that software systems operating over them must explicitly handle.

In parallel, the widespread adoption of large language models (LLMs) has driven the rise of Retrieval-Augmented Generation (RAG) techniques for intelligent access to distributed information [Lewis et al. 2020]. By combining fragment retrieval from heterogeneous sources with natural language generation, RAG systems improve traceability and reduce hallucinations compared to purely generative models [Schneider 2024, Yang et al. 2025]. However, most RAG pipelines assume that retrieved content can be

merged into a single constructed generation context without semantic conflict. This assumption of semantic neutrality is problematic in federated environments, where fragments that are individually valid within their original context may lead to unintended or misleading inferences when combined, such as presenting alternative or conditional requirements as jointly applicable [Lewis et al. 2020, Schneider 2024, Edge et al. 2024].

Semantic interoperability is typically addressed through strategies such as conceptual unification, semantic alignment, or knowledge structuring, often relying on shared ontologies, knowledge graphs, or, more recently, on delegating conflict resolution to large language models [Möller et al. 2024, Schneider 2024]. While these approaches enhance traceability and knowledge organization, they usually require semantic coordination that is hard to sustain among autonomous providers in federated DSs. Moreover, they remain limited by the assumption that retrieved fragments can be safely integrated into a single constructed generation context.

This paper reframes semantic interoperability in federated RAG systems as a problem of inferential composition, rather than of conceptual alignment or meaning unification. The focus shifts to explicitly controlling how retrieved fragments are combined during generation, in order to prevent invalid inferences from contextually incompatible content while preserving safe cross-context integration.

Building on this reformulation, we propose a lightweight framework that enforces inferential composition control through selective context separation and conservative fragment integration. It uses the semantic role of each fragment as its core unit of reasoning, drawing on pragmatic and inferentialist perspectives. Crucially, the approach avoids reliance on ontologies or shared knowledge graphs, allowing integration into federated environments without global semantic commitments.

The contributions of this work are threefold:

- A conceptual reframing of semantic interoperability in federated RAG systems, focusing on inferential composition rather than meaning unification or alignment.
- A lightweight and model-agnostic framework for controlling how retrieved fragments are combined, using semantic roles and pragmatic signals without requiring ontologies or shared vocabularies.
- A qualitative evaluation through a federated case study, showing that explicit composition control prevents reasoning errors under adversarial queries without degrading benign integration.

The remainder of the paper is structured as follows. Section 2 reviews related work on semantic interoperability, DSs, and RAG systems. Section 3 introduces the problem reformulation and the proposed framework. Section 4 describes the evaluation through a canonical case study. Section 5 discusses the results, and Section 6 concludes the paper.

2. Context and Related Work

This section reviews prior work on semantic interoperability in DSs and RAG, with emphasis on how existing approaches handle semantic heterogeneity and reasoning in federated contexts. In these settings, challenges extend beyond structural or terminological differences, as the key issue lies in how context-dependent fragments are implicitly combined during automated reasoning. While prior work has addressed conceptual alignment,

knowledge structuring, governance, and prompting techniques, the problem of inferential composition remains largely unexplored.

A dominant line of work addresses semantic interoperability through shared ontologies, common conceptual schemas, or alignment mechanisms. Widely adopted in software engineering and data science, these approaches are presented in recent surveys as the prevailing solution to conceptual heterogeneity across sources [J. Carlos et al. 2025]. Representative proposals combine semantic and structural information to define explicit mappings between vocabularies [Hao et al. 2023]. However, they rely on the assumption that aligned concepts can be safely combined during automated reasoning, an assumption that becomes fragile in federated DSs, where semantic coordination across autonomous providers is costly and constrained [Möller et al. 2024, Penedo 2024].

Another line of work uses knowledge graphs and structured semantics to enrich RAG systems, aiming to improve grounding and explainability [Schneider 2024]. Proposals like Graph-RAG organize retrieved fragments into graph structures before generation [Edge et al. 2024]. While this improves contextual organization, it still presents a unified context to the generative model, leaving the resolution of semantic conflicts and inferential combinations to the LLM.

A third line of work integrates RAG into federated DSs, focusing on architectural aspects such as sovereignty, traceability, and responsibility [Braga et al. 2026]. These proposals define explicit roles and interaction flows among providers, consumers, and federators. However, they typically treat semantic reasoning implicitly, assuming that fragments can be combined during generation if provenance is preserved, without mechanisms to control the inferential composition of context-dependent content.

Finally, some approaches rely on LLMs to handle semantic conflicts via prompting strategies such as chain-of-thought, instruction tuning, or guided reasoning [Schneider 2024, Diao et al. 2024]. While these techniques improve local coherence, they ultimately leave the management of inferential composition to the model, without ensuring that semantic boundaries between context-dependent fragments are preserved.

Beyond these four lines of work, related literature on contextual integrity [Nissenbaum 2004], normative reasoning in multi-agent systems [Governatori 2005], and policy-aware NLP [Ashley 2017] has addressed reasoning under legal or ethical constraints using formal logic or structured models. These approaches offer foundational insights but typically rely on symbolic representations or declarative policy frameworks. In contrast, the present work focuses on the operational composition of retrieved language fragments in RAG pipelines, without requiring access to internal model structures or external ontologies. As such, the proposed framework complements but does not replicate those paradigms.

In summary, existing approaches address key aspects of the problem, including conceptual alignment, context structuring, governance, and linguistic coherence, as shown in Table 1. However, they either rely on semantic coordination, difficult to sustain among autonomous providers, or shift conflict resolution to the generation process. In both cases, they assume that retrieved fragments can be coherently composed in a single generation context. This leaves open the central question of how to explicitly manage the *inferential composition* of context-dependent content, which we address next.

Table 1. Summary of related approaches and their limitations regarding inferential composition.

Approach	Main Contribution	Limitation
Ontologies and semantic alignment	Establish explicit correspondences between heterogeneous vocabularies, supporting conceptual unification and improving traceability.	Depend on semantic coordination among providers and assume that post-alignment integration is semantically neutral, which limits robustness in federated and heterogeneous settings.
Knowledge graphs	Structure contextual information to enhance grounding and explainability through graph-based or formal semantic representations.	Provide a unified context to the generative model and delegate inferential resolution to it, requiring centralized or shared representations.
Federated RAG architectures	Address governance aspects such as data sovereignty, traceability, and distributed retrieval across autonomous providers.	Handle semantic reasoning implicitly, assuming fragments can be directly combined during generation as long as provenance is preserved, without controlling inferential interactions.
Prompting and reasoning-guided approaches	Improve local coherence and guide model reasoning through techniques like chain-of-thought or instruction tuning.	Do not impose structural constraints on how fragments are composed, limiting robustness against implicit conflicts or conditional alternatives.

3. Conceptual Framework for Inferential Composition Control

This section presents one of the main contributions of the paper, namely a reformulation of semantic interoperability in RAG systems for federated DSs. Instead of aligning meanings across heterogeneous sources, the proposed approach focuses on explicitly controlling the *inferential composition* of retrieved fragments during answer generation.

Figure 1 illustrates this reformulation by contrasting the conventional view of semantic interoperability with the proposed focus on inferential composition. It highlights the shift in problem framing, analytical unit, source of semantic failure, and key assumptions in federated RAG systems.

This reformulation builds on a pragmatic view of meaning, where meaning is not a fixed property of isolated statements but depends on their use, context, and authorized inferences [Austin 1975]. From this perspective, semantic errors arise not only from conceptual mismatches between sources, but from using statements outside their original context. For example, a fragment may express a binding obligation in a legal domain but function as a non-binding recommendation in an ethical one, leading to conflicting inferences if this distinction is not explicitly maintained.

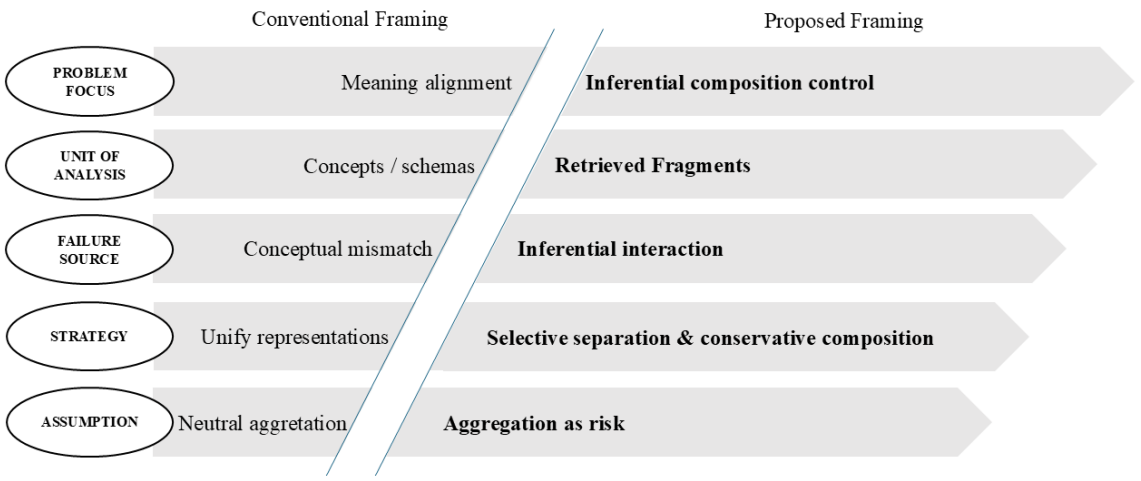

Figure 1. Reformulation of semantic interoperability in federated RAG systems. Comparison between the traditional meaning-alignment perspective and the proposed inferential composition approach.

A related perspective from dynamic semantics views interpretation as a context that evolves through successive discourse updates [Kamp and Reyle 2013]. In this view, not all information can be merged into a single context without risking inconsistency. Adding new fragments may not refine meaning, but instead override it or require keeping separate alternatives, especially when those fragments come from different normative frameworks.

Finally, from an inferentialist perspective, the meaning of a statement depends not on its correspondence with the world, but on the normative commitments and inferences it carries [Brandom 1994]. Applied to RAG systems, this view suggests that the problem lies not in semantic heterogeneity itself, but in composing fragments that authorize incompatible inferences from distinct and mutually exclusive contexts.

Building on these theoretical foundations, we adopt an operational notion of *minimal contextual dependence*. A fragment is contextually dependent when its combination with others may authorize incorrect inferences during automated reasoning. This dependence is not defined through ontologies or conceptual models but identified through observable signals within the fragment itself, such as normative language, mutually exclusive alternatives, or implicit references to institutional frameworks.

For example, a fragment that states a condition of applicability and another that expresses a conflicting exception may both be valid in isolation. However, if merged into a single generation context, they can trigger invalid inferences unless their distinction is preserved. In contrast, transversal fragments whose meaning remains stable across contexts can be safely integrated without introducing ambiguity or inconsistency during generation.

To manage contextual dependence during automated reasoning, the approach uses the *semantic role* of each fragment as its functional unit. A semantic role is not defined by the fragment's conceptual content in isolation, but by the inferential function it plays when combined with other fragments in response to a given query and context.

A single fragment may take on different semantic roles depending on the query and the interpretative setting. For example, a requirement might act as a binding condition in a legal domain or as a non-binding guideline in a best-practices scenario, leading to different inferences unless its role is made explicit. This abstraction allows precise control over fragment composition during generation without relying on ontologies or global semantic models.

To make this idea operational, we define a minimal and functional typology of semantic roles. The goal is not to provide an exhaustive or normative classification, but to capture the main inferential patterns that lead to composition failures in federated and regulated contexts. These roles reflect common distinctions found in normative, ethical, and technical sources, and provide a sufficient basis for selective separation and conservative composition in the framework.

Table 2 presents the semantic roles defined in this work, along with illustrative examples and the observable signals used for their identification.

Table 2. Functional typology of semantic roles with examples and identification criteria

Semantic Role	Example Fragment	Identification Signal
Descriptive	"Data must be encrypted during transmission."	Factual tone and technical formulation
Normative	"It is mandatory to obtain explicit consent."	Deontic language such as *must* or *shall*
Orientative	"Researchers should consider anonymization to minimize risk."	Non-binding modal verbs such as *should* or *may*
Alternative	"Either consent is obtained, or a public interest basis is documented."	Explicit conditional or mutually exclusive phrasing

The notions of minimal contextual dependence and semantic roles are especially well suited to federated environments like DSs, where semantic heterogeneity reflects organizational sovereignty rather than a defect to be corrected. In these settings, each provider operates within its own legal, technical, or institutional domain. Contextual dependence helps identify fragments that should not be combined beyond their original scope, while semantic roles guide their use during generation without requiring shared semantics or prior alignment.

When retrieved fragments from different providers are treated as minimal units of analysis, they can be assigned semantic roles based on contextual criteria. These roles follow a functional typology and are determined using observable signals within the fragment, indicating its intended inferential use and guiding how it should be handled during automated reasoning.

To address the main sources of conflict observed in regulated domains, we adopt a functional and extensible typology of semantic roles. Rather than listing all possible fragment functions, the typology focuses on composition patterns that most directly affect inferential behavior in federated settings. It differentiates between four types of fragments, each with a distinct inferential role. Descriptive fragments contribute contextual

or factual information. Normative fragments introduce binding obligations linked to legal or institutional frameworks. Alternative fragments express mutually exclusive conditions or options. Finally, orientative fragments convey non-binding principles or recommendations.

Semantic roles are assigned using operational signals observable within the fragments themselves, rather than through ontological taxonomies or prior semantic alignment. These signals include deontic or prescriptive language, explicit conditionals or exclusions, mutually exclusive formulations, and implicit references to institutional sources. Although this method does not provide an ontological definition of fragment meaning, it offers a practical alternative for controlling inferential composition. It enables query- and context-sensitive role assignment, and supports explicit decisions on how fragments are combined during automated reasoning.

From a technical standpoint, semantic roles can be assigned through various means without tying the approach to any specific method. These may include lightweight linguistic cues, classification aided by large language models, or metadata provided by the data providers. Regardless of the method, what matters is not how roles are assigned, but how they are used to control inferential composition during automated reasoning.

Based on semantic roles, the framework applies two complementary mechanisms to control inferential composition. The first is selective context separation, used when combining fragments with incompatible roles could lead to invalid inferences. In such cases, the fragments are not discarded but placed in separate generation contexts, preventing the system from treating them as cumulative or simultaneously applicable.

In contrast, conservative composition allows the integration of fragments with compatible semantic roles, as long as their inclusion does not introduce ambiguity or global inconsistency. From an inferentialist standpoint, this mechanism supports only shared and uncontroversial inferences. It is applied selectively, mainly to descriptive or transversal fragments whose meaning remains stable across contexts.

For instance, two fragments that state alternative legal bases for data processing are kept in separate contexts through selective separation. A descriptive fragment describing a technical safeguard can be conservatively integrated into both contexts without disrupting their conditional structure.

Decisions about separation and integration are guided by the notion of *semantic tension*, which serves as an operational indicator of potential inferential conflict between fragments. Tension is assessed locally, by comparing fragment pairs to detect problematic combinations that require separation, and globally, across the constructed context, to ensure that integration does not increase conflict or compromise inferential coherence during generation.

Although semantic tension is applied qualitatively in this work to support separation and integration decisions, it can be preliminarily formalized to make the concept more explicit. A tentative definition of the tension $T(f_i, f_j)$ between two fragments f_i and f_j can be written as a weighted sum of discrete conflict indicators:

$$T(f_i, f_j) = \alpha_1 \cdot \text{RoleMismatch}(f_i, f_j) + \alpha_2 \cdot \text{DeonticConflict}(f_i, f_j) + \alpha_3 \cdot \text{SourceDivergence}(f_i, f_j)$$

Each term represents a binary or ordinal-valued signal capturing semantic incompatibility, such as mismatched roles, conflicting deontic expressions, or divergence in provenance. The weights α are tunable parameters that can be adjusted according to domain-specific priorities.

This formulation is not meant to be exhaustive, but rather to capture the core intuition that some fragment combinations carry disproportionate inferential risk. It also opens a clear path to extend the framework with more systematic and explainable metrics for context construction.

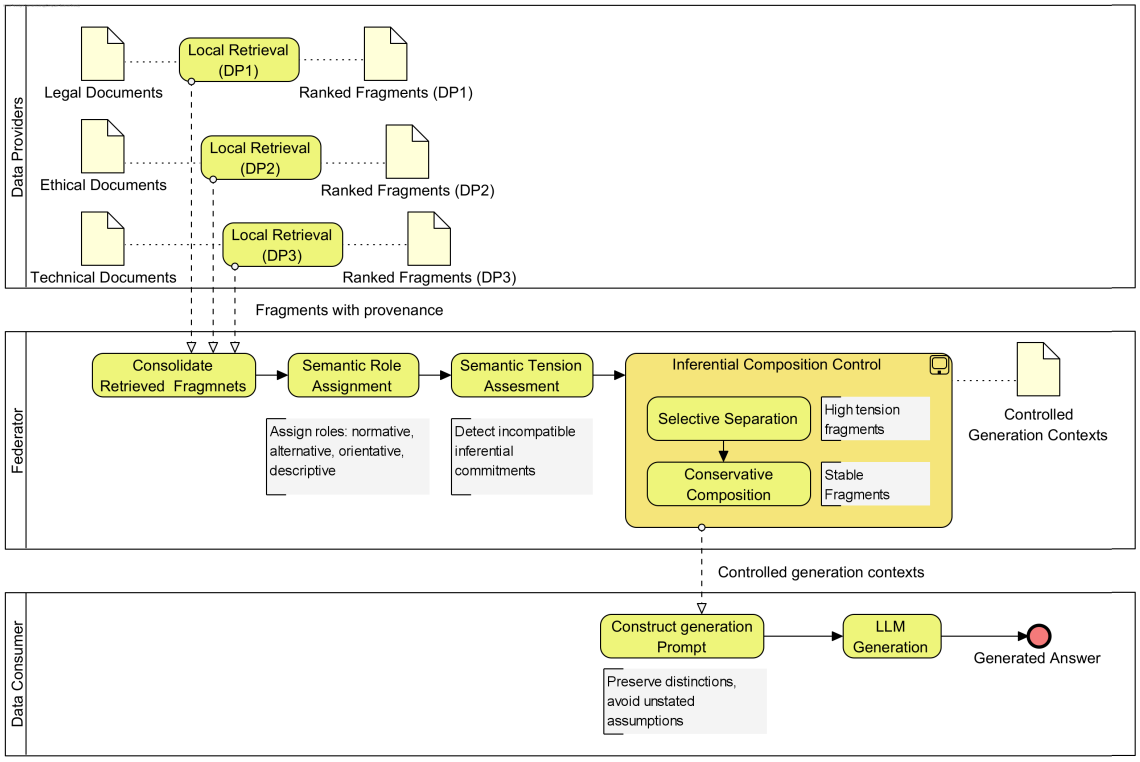

Figure 2. Inferential composition control in federated RAG. Overview of the complete RAG pipeline, from retrieval to generation, incorporating semantic role assignment and controlled inferential composition.

The combined use of selective separation and conservative composition leads to a state of equilibrium, referred to as *semantic homeostasis*, where the reasoning context maintains inferential stability during generation while absorbing local conflicts without collapsing semantic diversity. This controlled reconfiguration prevents the accumulation of incompatible commitments and supports coherent reasoning. Figure 2 illustrates the core workflow of the proposed framework. While the figure shows three data providers to align with the evaluation scenario in the next section, the model generalizes to any number of providers.

4. Case Study and Evaluation

The evaluation aims to determine how different strategies for inferential composition, when applied to the same set of retrieved fragments, influence automated reasoning in federated RAG systems. Rather than assessing the normative or linguistic quality of the

responses, the focus is on the operational impact of context construction within the RAG pipeline.

The evaluation examines how the same set of retrieved fragments is organized into generation contexts under three configurations. It traces how each context construction path enables or restricts different types of inferences during generation. By making intermediate steps explicit, such as fragment groupings, semantic role assignments, and controlled context partitions, the evaluation enables systematic observation of inferential collapse and its prevention in a controlled and reproducible setting.

The evaluation is carried out in a regulated domain focused on the use and sharing of personal data for research. This setting is particularly suited to reveal inferential composition failures caused by the coexistence of legal, ethical, and technical frameworks. The domain is instantiated through a dedicated corpus built from three data providers, each reflecting a conceptually homogeneous legal, ethical, or technical perspective.

Each provider contributes three complete documents, preserving semantic coherence within sources while introducing heterogeneity only across providers. This design ensures that inferential conflicts result from cross-provider composition, not from inconsistencies within individual sources.

For each query, fragments are retrieved automatically using standard vector-based search over the full corpus, producing a fixed set of results without manual curation. This same set is reused across all evaluated configurations, ensuring that observed differences are due solely to context construction rather than retrieval variation.

The evaluation includes three queries, each targeting a different semantic behavior. The first serves as a negative control and focuses on integrating transversal information without contextual dependence. It asks: *Which technical and governance measures should be implemented to protect personal data in research projects?*

The remaining two are adversarial queries designed to expose complementary failure modes of inferential composition. One introduces an explicit normative conflict between legal obligations and ethical considerations: *Is it mandatory to simultaneously comply with legal obligations and ethical considerations for any use of personal data in research?* The other induces implicit inferential collapse via universal quantification: *Which requirements must be met in all cases to share personal data in research projects, regardless of the applicable legal basis?*

The validation builds on the federated RAG model introduced in prior work on DSs [Braga et al. 2026], reusing its distributed retrieval architecture and access governance mechanisms. For each query, all configurations use the same set of retrieved fragments, ensuring that any observed differences result solely from inferential composition decisions.

All experiments use the same large language model (GPT-4 via OpenAI API), with deterministic decoding and identical prompts and generation parameters across configurations. Queries and source documents are in Spanish to reflect the regulatory domain under study. Since the evaluation focuses on inferential composition, this language choice does not affect the validity of the reasoning behavior observed.

The validation process is fully reproducible. The implementation, along with

all experimental artifacts, is available as open source at `https://github.com/GSYAtools/DSRAG-inferential-control`.

The evaluation compares three context construction strategies. In the baseline configuration, all retrieved fragments are aggregated into a single generation context without semantic control. The hard separation variant groups fragments by data provider and preserves provenance, but does not manage inferential interactions. The proposed configuration applies semantic role assignment and tension-based control to enable selective separation and conservative composition, as introduced in Section 3.

The evaluation relies on qualitative observables that capture the inferential effects of different context construction strategies. These include the improper accumulation of alternative requirements, the preservation of legitimate alternatives, the controlled integration of transversal information, and the resulting structural coherence of the generated reasoning. A quantitative benchmark would not capture these structural effects, which depend on context-sensitive inference rather than accuracy scores. The following scenarios illustrate how these behaviors emerge under each configuration, using the same set of retrieved fragments.

For the control query, which targets the integration of transversal information without relevant contextual dependencies, all three configurations produce substantively equivalent responses. Across configurations, the generated answers enumerate standard technical and governance measures such as access control, encryption, monitoring, and security by design principles, without introducing improper inferential relations between fragments. For example, all responses include statements equivalent to: *"Technical measures to protect personal data in research projects include access control (restricting use to authorized individuals), encryption (both at rest and in transit), activity logging, and the application of security by design principles."* This confirms that inferential composition control does not introduce collateral effects or unnecessary intervention when no semantic tensions are present in the retrieved information.

For the adversarial query involving an explicit normative conflict, which asks whether legal obligations and ethical considerations must be fulfilled simultaneously when using personal data for research, the retrieved set includes fragments expressing legally binding requirements, ethical guidance, and transversal technical safeguards. Among these, normative and orientative fragments are inferentially critical, as combining them without role distinction may cause normative collapse, for example, treating ethical recommendations as mandatory.

In the baseline configuration, all fragments are merged into a single generation context without semantic distinction. As a result, legal and ethical fragments coexist as if jointly applicable, allowing the model to present both as mandatory. The generated response states, for example: *"Yes, the use of personal data in research must comply with both legal obligations and ethical considerations."* This illustrates an inferential collapse, as the response fails to differentiate between binding and non-binding norms.

In the hard separation setting, fragments are grouped by provider and provenance is preserved. However, since inferential roles are not controlled, the model still combines normative and ethical fragments in a single prompt. This again leads to responses that treat ethical principles as mandatory alongside legal obligations, despite their distinct

normative weight.

In contrast, the proposed configuration assigns semantic roles to the retrieved fragments and evaluates semantic tension before generation. Legal obligations are classified as normative, ethical principles as orientative, and technical safeguards as descriptive. Semantic tension between normative and orientative fragments triggers selective separation, while descriptive fragments are conservatively composed into each context. This results in a controlled prompt that maintains normative boundaries. The generated response reflects this distinction: *"Legal obligations are binding requirements, while ethical considerations guide responsible decision making but are not legally enforceable."*

Figure 3 schematically compares how each configuration organizes retrieved fragments prior to LLM invocation. The baseline merges all fragments into a single generation context, while hard separation preserves provenance without enforcing semantic constraints. In contrast, the proposed approach assigns semantic roles to retrieved fragments, separates incompatible contexts, and assembles a controlled prompt that preserves inferential boundaries.

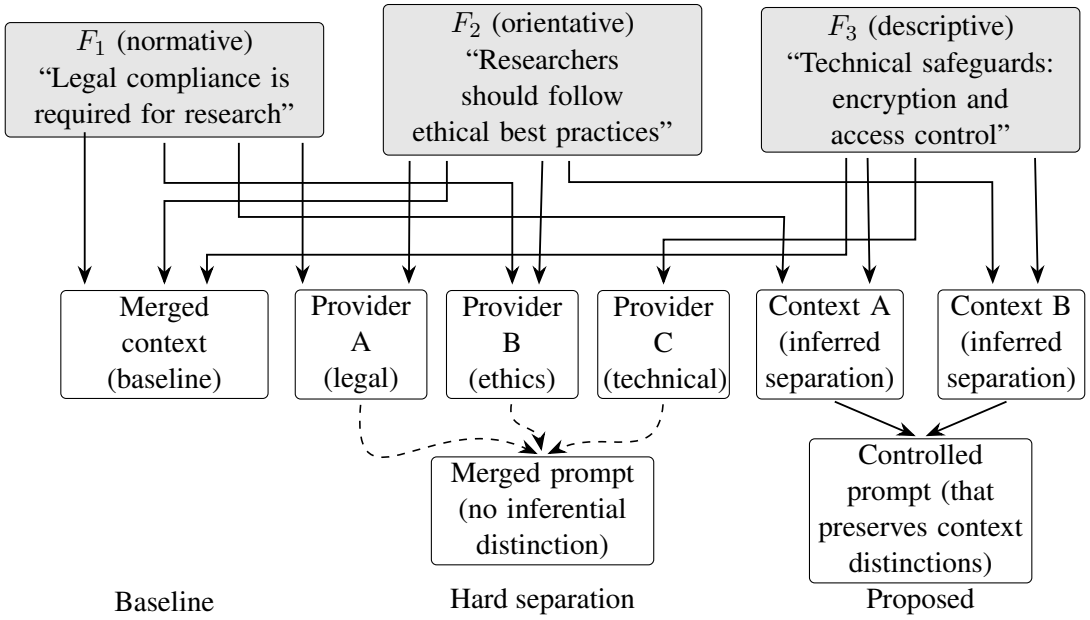

Figure 3. Schematic comparison of prompt construction strategies in federated RAG. Baseline merging, hard separation by data provider, and the proposed role-based inferential composition are contrasted.

For the adversarial query with implicit conflict, which asks which requirements must be fulfilled in all cases to share personal data for research, the inferential challenge does not stem from explicitly competing normative frameworks, but from the interaction between universal linguistic quantification and conditionally applicable legal requirements. The retrieved fragments include descriptions of alternative legal bases for data processing, binding obligations tied to specific legal contexts, and transversal technical safeguards. When combined indiscriminately, these fragments can mask the conditional scope of legal requirements and induce a collapse into a universal reading.

In the baseline configuration, all fragments are merged into a single context. The

universal formulation of the query encourages the model to treat all retrieved requirements as simultaneously applicable. As a result, conditions that are valid only under certain legal bases, such as consent or public interest, are presented as mandatory in all cases. For example, the model states: *"All data sharing must comply with consent requirements, risk assessment, access restrictions, and pseudonymization"*, even though some of these apply only under specific legal grounds.

Under hard separation, fragments are grouped by data provider and their provenance is preserved, but no inferential control is enforced. This again allows fragments tied to distinct legal bases to be interpreted cumulatively. The generated response still collapses alternative conditions, failing to preserve their contextual constraints.

In contrast, the proposed configuration assigns semantic roles that differentiate between conditional legal bases, normative obligations, and transversal safeguards. Semantic tension between conditional requirements and the query's universal framing triggers selective context separation. Transversal fragments, such as technical safeguards, are conservatively integrated into each context without altering their conditional boundaries. As a result, the generated response preserves the distinct applicability of alternative legal bases. For instance, the model states: *"The reuse of personal data for research is only lawful when a valid legal basis, a legitimate purpose, and compliance with data protection principles are all present. Specific requirements, such as consent, depend on the legal basis being used."*

Figure 4 visualizes how the same fragments are handled under each configuration. Baseline and hard separation collapse alternative legal bases into a universal interpretation, while inferential composition control preserves conditional structure by separating roles and controlling integration.

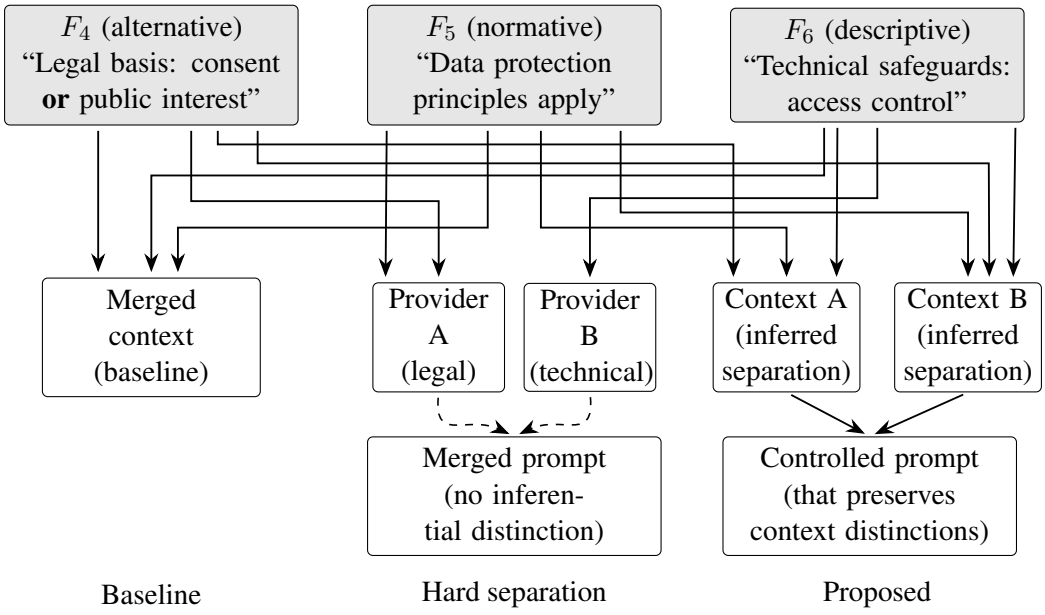

Figure 4. Schematic comparison of fragment composition under an adversarial query with implicit conflict. Baseline, hard separation, and inferential composition control are contrasted.

Taken together, these results demonstrate that explicit inferential composition con-

trol is not intended to improve the linguistic form of responses, but to constrain the space of authorized reasoning during generation. By structuring how fragments are combined, the method prevents inferential errors that arise when contextually dependent content is merged without control. Across both explicit and implicit adversarial scenarios, the approach consistently preserves normative distinctions and conditional applicability, supporting more reliable and auditable reasoning in federated RAG systems.

5. Discussion

This work proposes a shift in how semantic interoperability is approached in RAG systems for federated DSs, reframing the problem from conceptual alignment or prompting to the control of inferential composition during automated reasoning. Rather than unifying meanings or delegating conflict resolution to the model, the goal is to control how contextually dependent fragments are combined before generation, while respecting the semantic diversity of federated environments.

The case study illustrates the practical implications of this reformulation. In adversarial scenarios, baseline RAG collapses implicit normative alternatives, while hard separation avoids such collapses at the cost of excessive fragmentation. The proposed method does not seek to improve linguistic form or determine normative correctness, but to constrain the space of authorized inferences. Its primary effect is the structural prevention of inferential collapse, especially under implicit conflict.

From an engineering perspective, the approach is lightweight, modular, and explainable. It relies on observable signals and explicit mechanisms for separation and composition, allowing integration into existing RAG architectures without ontologies, global knowledge structures, or changes to the language model. Semantic roles and tension provide transparent criteria for context construction, enabling auditing and reasoning traceability by clarifying which fragments are integrated, which are separated, and under what conditions.

From a regulatory perspective, inferential composition control acts as a structural safeguard for generative systems operating in regulated and sovereign DSs. By externalizing context construction and constraining inferences at the orchestration level, the approach supports explainability and traceability without requiring access to proprietary model internals. This is especially relevant in domains where reasoning coherence and accountability have normative implications.

This study is limited by design to a qualitative case analysis focused on exposing structural failures of inferential composition. As such, it does not aim to produce large-scale quantitative conclusions. The semantic role typology is deliberately minimal, and the notion of semantic tension is intentionally flexible to preserve a lightweight and extensible framework. Issues such as normative correctness, fairness, and model bias remain beyond the scope of this work.

Future work includes extending the evaluation to additional domains and query types, as well as introducing quantitative measures of inferential stability. Other directions involve dynamically expanding the semantic role typology and integrating composition control with complementary mechanisms such as external verifiers or structured knowledge sources. Overall, this work lays the foundation for a research line focused

on the operational management of semantic interoperability in federated and regulated reasoning systems.

6. Conclusions

This work has reframed the challenge of semantic interoperability in federated DSs as a problem of inferential composition during automated reasoning with RAG systems. Unlike approaches based on ontologies, knowledge graphs, or prompting, we identify the core issue in how contextually dependent fragments are combined before generation.

The validation in a regulated domain shows that inferential composition control is not intended to improve linguistic output or resolve normative correctness, but to prevent specific forms of inferential collapse in federated RAG systems lacking semantic control. The evaluation demonstrates observable effects on automated reasoning behavior. By enabling explicit control and auditing at the orchestration level, the approach supports regulatory requirements such as traceability and documentation (Article 52) and systemic risk mitigation for generative AI (Article 55) under the European AI Act.

This work lays the foundation for a research line centered on the operational management of semantic interoperability in federated and regulated generative systems. Future efforts will focus on formalizing the framework, extending validation across domains and query types, developing quantitative metrics of inferential stability, and integrating with external verifiers or structured knowledge sources.

Acknowledgments

This publication would not have been possible without the support of the following projects: Di4SPDS (PCI2023145980-2), funded by MCIN/AEI/10.13039/501100011033 and by the European Union (Chist-Era Prgram); KOSMOS-UCLM (PID2024-155363OB-C44) funded by MCIN/AEI/10.13039/501100011033/ FEDER, UE; AURORA (SBPLY/24/180225/000074), funded by the Consejería de Educación, Cultura y Deportes de la Junta de Comunidades de Castilla La Mancha, and Fondo Europeo de Desarrollo Regional FEDER (Programa Operativo de Castilla-La Mancha 2021–2027); RADAR (2025-GRIN-38447) funded by FEDER; and Grant RED2024-154240-T funded by MICIU/AEI/10.13039/501100011033.

References

Ashley, K. D. (2017). *Artificial intelligence and legal analytics: new tools for law practice in the digital age*. Cambridge University Press.

Austin, J. L. (1975). *How to do things with words*. Harvard university press.

Braga, C. M., Serrano, M. A., and Fernández-Medina, E. (2026). Guided and federated rag: Architectural models for trustworthy ai in data spaces. In *Intelligent Data Engineering and Automated Learning – IDEAL 2025*, pages 363–374, Cham. Springer Nature Switzerland.

Brandom, R. (1994). *Making it explicit: Reasoning, representing, and discursive commitment*. Harvard university press.

Diao, S., Wang, P., Lin, Y., Pan, R., Liu, X., and Zhang, T. (2024). Active prompting with chain-of-thought for large language models. In *Proceedings of the 62nd Annual*

Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pages 1330–1350.

Edge, D., Trinh, H., Cheng, N., Bradley, J., Chao, A., Mody, A., Truitt, S., Metropolitansky, D., Ness, R. O., and Larson, J. (2024). From local to global: A graph rag approach to query-focused summarization. *arXiv preprint arXiv:2404.16130*.

Governatori, G. (2005). Representing business contracts in RuleML. *International Journal of Cooperative Information Systems*, 14(02n03):181–216.

Hao, Z., Mayer, W., Xia, J., Li, G., Qin, L., and Feng, Z. (2023). Ontology alignment with semantic and structural embeddings. *Journal of Web Semantics*, 78:100798.

J. Carlos, M., Ana Alice, B., Rui M., S., and Paulo J., M. (2025). Semantic mediation: a literature review on semantic interoperability through ontologies. *Procedia Computer Science*, 263:734–743. International Conference on Industry Sciences and Computer Science Innovation (iSCSi'24).

Kamp, H. and Reyle, U. (2013). *From discourse to logic: Introduction to modeltheoretic semantics of natural language, formal logic and discourse representation theory*, volume 42. Springer Science & Business Media.

Lewis, P. et al. (2020). Retrieval-augmented generation for knowledge-intensive NLP tasks. *Advances in neural information processing systems*, 33:9459–9474.

Möller, F. et al. (2024). Industrial data ecosystems and data spaces. *Electronic Markets*, 34(1):41.

Nissenbaum, H. (2004). Privacy as contextual integrity. *Wash. L. Rev.*, 79:119.

Penedo, A. C. (2024). The Regulation of Data Spaces under the EU Data Strategy: Towards the 'Act-ification'of the Fifth European Freedom for Data? *European Journal of Law and Technology*, 15(1).

Schneider, J. (2024). Explainable generative AI (GenXAI): A survey, conceptualization, and research agenda. *Artificial Intelligence Review*, 57(11):289.

Yang, J., Shu, L., Duan, H., and Li, H. (2025). RDguru: A conversational intelligent agent for rare diseases. *IEEE Journal of Biomedical and Health Informatics*, 29(9):6366–6378.