Preprints.org

Article

Not peer-reviewed version

Towards Distributed Cognition: Knowledge Gravity as a Foundation for Sovereign Neuro-Symbolic AI

Antonio Salis *, Patrizia Loi , Monica Medda

Posted Date: 20 November 2026

Keywords: edge AI; neuro-symbolic AI; edge-cloud continuum; data gravity; knowledge graphs; explainable AI; sustainable AI; digital sovereignty; AI governance; knowledge gravity

Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC, OpenAlex.

Copyright: This open access article is published under a Creative Commons CC BY 4.0 license, which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Preprints.org | NOT PEER-REVIEWED | Posted: 20 November 2026 https://doi.org/10.20944/preprints202611.2231.v1

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

Towards Distributed Cognition: Knowledge Gravity as a Foundation for Sovereign Neuro-Symbolic AI

Antonio Salis [1,*], Patrizia Loi [1,2] and Monica Medda [1]

[1] Villanova.ai SpA; Research Division; Loc. Sa Illetta, SS195 km 2,3 – I-09123 Cagliari, Italy

[2] Università degli Studi di Cagliari, via Ospedale 40; 09124, Cagliari

* Correspondence: asalis@villanova.ai

Abstract

The rapid adoption of Artificial Intelligence (AI) across public administrations, industry, education, and critical infrastructures is increasing demands for computational efficiency, transparency, sustainability, and regulatory compliance. While cloud-centric AI architectures have enabled major advances in machine learning, they face growing challenges related to data movement, latency, energy consumption, explainability, cybersecurity, and digital sovereignty. At the same time, European regulatory frameworks such as the European AI Act [1], GDPR [2], NIS2 [3], the Data Act [4], and the Data Governance Act [5] impose heterogeneous requirements concerning AI risk management, data protection, cybersecurity, data governance, accountability, and access to data. Current research addresses Edge Computing, Neuro-Symbolic AI, Knowledge Graphs, Explainable AI, and Sustainable AI largely as separate domains. This paper introduces SOFIA (Sovereign, Ontology-driven Framework for Intelligent Action), a conceptual knowledge-driven neuro-symbolic framework that integrates Edge-Cloud Continuum principles, Data Gravity, Knowledge Graphs, and Governance-by-Design mechanisms. SOFIA aims to define and operationalize the concept of Knowledge Gravity, extending compute-to-data approaches by considering the placement of semantic knowledge, reasoning capabilities, policies, and regulatory constraints closer to the points where decisions are generated across the continuum. Through the integration of machine learning, symbolic reasoning, and evolving Knowledge Graphs, the framework outlines how explainable and governance-aware intelligence could be deployed on resource-constrained edge devices while reducing reliance on centralized infrastructures. In this perspective, Knowledge Gravity is proposed as a conceptual orchestration principle for distributed cognition, where knowledge assets influence placement and governance decisions. To illustrate the proposed concepts, the paper introduces a preliminary mathematical formulation and a synthetic simulation based on hypothetical deployment scenarios. These examples are intended only to demonstrate the behavior of the proposed indicators and optimization perspective and do not constitute empirical validation of the framework. By framing knowledge placement, reasoning availability, governance requirements, and sovereignty constraints within a unified architectural perspective, SOFIA contributes a research agenda for distributed, explainable, and sustainable AI across the Edge-Cloud Continuum. The framework is presented as a conceptual foundation whose assumptions, metrics, and expected benefits require future validation in real-world deployments.

Keywords: edge AI; neuro-symbolic AI; edge-cloud continuum; data gravity; knowledge graphs; explainable AI; sustainable AI; digital sovereignty; AI governance; knowledge gravity

1. Introduction

Position Statement

This paper is a position paper proposing the SOFIA framework as a new paradigm for sovereign and explainable distributed intelligence. The framework is currently conceptual and is articulated through

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org | NOT PEER-REVIEWED | Posted: 20 November 2026 https://doi.org/10.20944/preprints202611.2231.v1

architectural components developed within the Villanova initiative. The objective of this work is not to provide empirical validation but to formulate a research vision and identify future scientific directions.

1.1. Background

Artificial Intelligence is increasingly recognized as a strategic technology for digital transformation across public administration, manufacturing, healthcare, agriculture, education, and critical infrastructures. At the same time, the growth of connected devices, IoT ecosystems, cyber-physical systems, and distributed data sources is profoundly reshaping the way intelligence is generated and consumed.

The dominance of cloud-centric AI architectures has enabled substantial advances in machine learning and foundation models but has also exposed structural limitations. Data transfers toward centralized infrastructures generate communication overhead, latency, cybersecurity risks, and operational costs. Furthermore, large-scale model training and retraining contribute significantly to AI's environmental footprint and energy consumption.

Beyond technical concerns, the emergence of the European AI Act and related regulatory frameworks has highlighted the need for explainable, auditable, trustworthy, and sovereign AI systems capable of operating under governance and compliance constraints by design.

These challenges motivate the investigation of an integrated architectural perspective that combines Edge-Cloud Continuum principles, neuro-symbolic reasoning, evolving knowledge structures, and governance-by-design mechanisms to support distributed, explainable, and sovereign intelligence.

1.2. State of the Art

1.2.1. Edge-Cloud Continuum

Research on the Edge-Cloud Continuum focuses primarily on distributed execution, orchestration, workload placement, and resource optimization across heterogeneous infrastructures [6–9]. Significant progress has been achieved in enabling low-latency services and decentralized computation [10–13].

However, most existing solutions still rely heavily on cloud-hosted AI models and centralized retraining processes. Intelligence remains largely cloud-dependent despite the distributed nature of computation.

1.2.2. Data Gravity

Data Gravity has emerged as a key design principle in distributed systems [14,15]. Its central idea is that increasing data volumes naturally attract applications and computation toward data locations.

Current implementations primarily emphasize data locality, infrastructure optimization, and compute-to-data approaches. Nevertheless, limited attention has been devoted to distributed management of knowledge and reasoning capabilities across the continuum.

1.2.3. Neuro-Symbolic AI

Neuro-Symbolic AI combines neural learning with symbolic reasoning to enhance explainability, knowledge integration, and trustworthiness [16–20].

Current research demonstrates significant benefits in interpretability and decision transparency. However, most neuro-symbolic solutions are deployed in centralized environments and rarely address large-scale edge deployment scenarios [21–23].

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org | NOT PEER-REVIEWED | Posted: 20 November 2026 https://doi.org/10.20944/preprints202611.2231.v1

1.2.4. Knowledge Graphs

Knowledge Graphs provide semantic interoperability, structured knowledge representation, contextual understanding, and support for explainable reasoning [24–26].

Despite their growing adoption across multiple domains, Knowledge Graphs are typically implemented as centralized repositories [27,28]. Their integration into edge-native architectures remains limited.

1.2.5. Distributed Semantic Reasoning and Federated Knowledge Management

Recent advances in neuro-symbolic reasoning over knowledge graphs [19], large-scale knowledge graph infrastructures [25,28], ontology evolution [31,32], and distributed AI systems [13,21,22] have demonstrated the value of combining semantic representations with distributed execution environments. Similarly, federated knowledge-management approaches have addressed synchronization, consistency, interoperability, provenance, and lifecycle management of knowledge assets across organizational boundaries.

These contributions have significantly improved the ability to represent, exchange, govern, and reason over distributed knowledge. However, their primary focus remains on specific aspects of the problem, including reasoning performance, semantic interoperability, consistency maintenance, knowledge sharing, or federated coordination. A broader architectural perspective integrating these capabilities into an operational distributed-intelligence framework remains less explored.

1.2.6. Sustainable and Trustworthy AI

Recent initiatives in Green AI [29,30], knowledge and ontology evolution [31,32], and explainable, trustworthy, sovereign, and governed AI [33–39] emphasize computational efficiency, adaptive knowledge management, transparency, accountability, risk management, and regulatory compliance.

Although important progress has been achieved, sustainability, explainability, and governance are often addressed independently from edge intelligence and semantic reasoning frameworks.

1.3. Research Gap

The literature reviewed in Sections 1.2.1-1.2.6 demonstrates substantial advances in Edge-Cloud Continuum architectures, Data Gravity, Neuro-Symbolic AI, Knowledge Graphs, distributed semantic reasoning, federated knowledge management, explainable AI, and governance-aware intelligent systems. Significant progress has been achieved in distributed computation, semantic knowledge representation, reasoning transparency, knowledge synchronization, and trustworthy AI mechanisms. However, these research streams have largely evolved independently and are typically optimized for different objectives, including computational efficiency, semantic interoperability, model performance, explainability, governance, sustainability, or regulatory compliance.

As a consequence, an important gap remains between distributed computational infrastructures and distributed cognitive systems. Existing approaches provide mechanisms for distributing computation, data, services, models, or knowledge repositories, but they offer limited support for determining how semantic knowledge assets, reasoning capabilities, governance policies, and sovereignty constraints should be jointly orchestrated across heterogeneous environments in support of operational decision making.

Recent advances in distributed semantic reasoning, federated knowledge management, semantic caching, and large-scale Knowledge Graph infrastructures have addressed important aspects of knowledge distribution and lifecycle management. Nevertheless, these approaches generally focus on specific concerns such as reasoning performance, interoperability, consistency preservation, synchronization, or query efficiency. The broader question of how knowledge, reasoning, explainability, governance, and sovereignty should be dynamically positioned throughout the continuum according to operational objectives remains comparatively underexplored.

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org | NOT PEER-REVIEWED | Posted: 20 November 2026 https://doi.org/10.20944/preprints202611.2231.v1

More specifically, current literature provides limited architectural support for treating knowledge placement, reasoning availability, governance obligations, explainability requirements, synchronization costs, and digital sovereignty as components of a unified optimization and orchestration problem. While individual solutions exist for each of these dimensions, their integration into a coherent decision-centric framework remains largely understudied. Existing architectures rarely consider semantic knowledge assets, governance policies, regulatory constraints, and reasoning services as first-class entities whose placement can be optimized alongside traditional computational resources.

Building upon previous informal uses of the term *Knowledge Gravity* [40], this paper proposes SOFIA as a conceptual framework intended to address this gap. The central premise is that distributed intelligence requires more than moving computation closer to data sources. It also requires the dynamic placement of knowledge assets, semantic models, governance mechanisms, and reasoning capabilities closer to the locations where decisions are generated and operational value is created. Under this perspective, knowledge is no longer treated as a passive repository of information but as an active and deployable resource that participates directly in distributed decision-making processes.

Based on the representative literature reviewed in this paper, we did not identify frameworks that explicitly combine Edge-Cloud Continuum principles, Neuro-Symbolic World Models, continuously evolving Knowledge Graphs, governance-by-design mechanisms, runtime assurance, distributed knowledge placement, and sovereign AI execution within a single conceptual architecture. Rather than claiming uniqueness, SOFIA is proposed as an integrative synthesis of these research streams and as a conceptual framework for exploring their joint orchestration through the Knowledge Gravity principle and the broader notion of a Cognitive Continuum. The contribution of this work therefore lies primarily in the conceptual integration of existing research directions and in the formulation of a unified orchestration perspective for distributed cognition.

From this perspective, the research gap addressed by this paper is not merely the absence of another distributed AI architecture, but the lack of an integrated framework capable of jointly optimizing where knowledge resides, where reasoning occurs, how governance policies are enforced, and how decision utility can be maximized while respecting explainability, consistency, sovereignty, resource, and regulatory constraints. The SOFIA framework is proposed as an initial step toward addressing this challenge and establishing the foundations for distributed cognition across the Cognitive Continuum.

1.4. Research Question and Hypothesis

Research Question

How can a knowledge-driven neuro-symbolic edge intelligence framework based on Edge-Cloud Continuum, Data Gravity and Knowledge Gravity principles support AI sustainability, explainability, governance, and regulatory compliance processes while reducing dependence on centralized computation and continuous model retraining?

Research Hypothesis

The core hypothesis of SOFIA is that intelligence should move toward data when advantageous rather than continuously moving data toward centralized intelligence. Furthermore, continuously evolving knowledge structures may reduce the need for repeated model retraining while improving explainability, governance, and sustainability.

1.5. Contributions

This paper contributes:

1. A conceptual definition of Knowledge Gravity.
2. A novel neuro-symbolic interpretation of distributed cognition.
3. A preliminary quantitative indicator (KPQI).
4. An architectural vision for sovereign distributed AI.
5. A research agenda for future validation.

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org | NOT PEER-REVIEWED | Posted: 20 November 2026 https://doi.org/10.20944/preprints202611.2231.v1

2. SOFIA Conceptual Framework

2.1. SOFIA Architectural Framework

SOFIA framework, instantiated through a reference architecture, should be interpreted as the evolutionary extension of existing Edge-Cloud initiatives and reference architectures [41–43] toward sovereign neuro-symbolic intelligence, developed within the Villanova project in IPCEI-CIS, and aligned with the IPCEI-CIS Reference (ICRA) Version 3.0 [44].

Cognitive Decision-to-Action Loop

While many AI architectures focus primarily on decision support, SOFIA is designed to address the complete Decision-to-Action continuum. Decisions generated through neuro-symbolic reasoning may be transformed into governed actions through policy-aware orchestration mechanisms operating across edge, fog, and cloud layers.

The framework therefore distinguishes between:

- perception;
- decision;
- action planning;
- action execution;
- feedback and learning.

This closed cognitive loop enables continuous adaptation while preserving explainability, traceability, and human oversight. Future implementations will investigate how different levels of autonomy can be governed through policy-driven execution mechanisms.

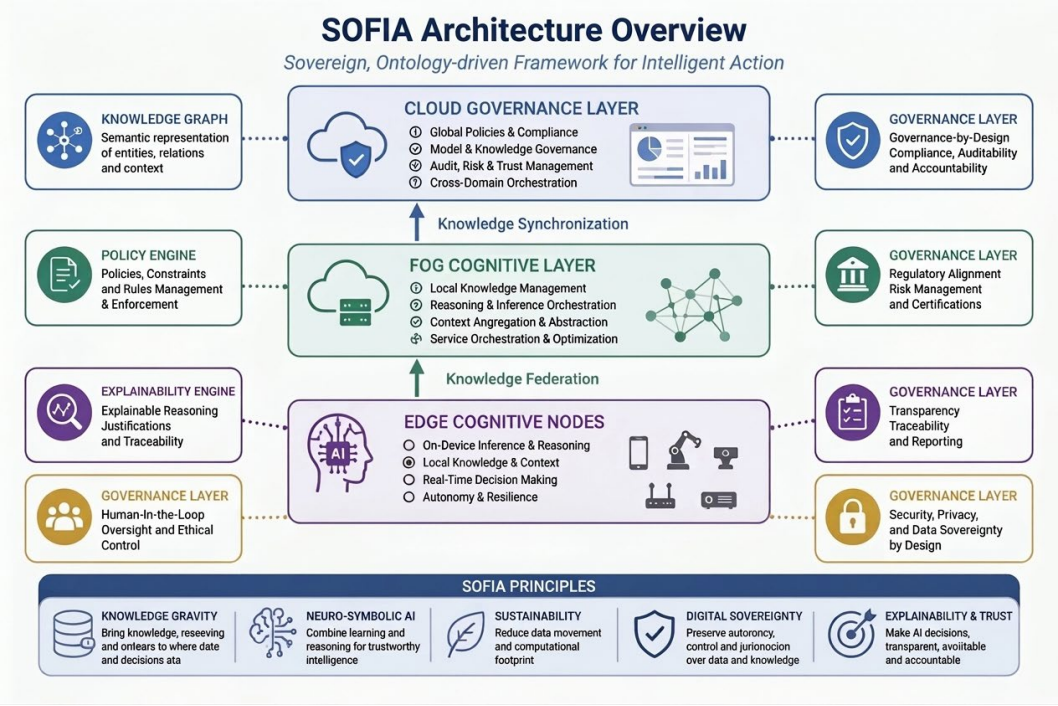


Figure 1. SOFIA Architecture Overview.

Edge Layer

The Edge Layer includes smartphones, IoT devices, embedded systems, mini-PCs, industrial gateways, and cyber-physical devices.

Functions include:

- Local inference
- Local reasoning
- Data acquisition

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org | NOT PEER-REVIEWED | Posted: 20 November 2026 https://doi.org/10.20944/preprints202611.2231.v1

- Event processing
- Context-aware decisions

Fog Layer

The Fog Layer provides:

- Semantic federation
- Knowledge aggregation
- Policy enforcement
- Resource orchestration
- Regional intelligence coordination

Cloud Governance Layer

The Cloud Layer performs:

- Large-scale model training
- Knowledge synchronization
- Lifecycle management
- Governance orchestration
- Federated coordination

Unlike conventional architectures, cloud infrastructures are not the primary intelligence center but one component of a distributed cognitive ecosystem. The Cloud Governance Layer may also host a Cognitive Control Layer responsible for global knowledge orchestration, policy management, governance supervision, and cognitive coordination across the Edge-Cloud Continuum. Rather than centralizing intelligence, this layer provides meta-cognitive functions that supervise the evolution, consistency, explainability, and compliance of distributed cognitive processes.

The Cognitive Control Layer acts as a federated control plane, which may be distributed across multiple nodes rather than implemented as a centralized component.

In future implementations, the Cognitive Control Layer may also host Knowledge Gravity optimization services responsible for evaluating placement indicators, solving placement decisions, authorizing knowledge distribution, and coordinating migration across the Cognitive Continuum.

Threats to Validity and Limitations

The framework currently presents three main limitations. First, the KPQI metric remains a preliminary formulation requiring empirical calibration. Second, the relationship between knowledge proximity and explainability has not yet been validated experimentally. Third, the illustrative use cases are synthetic and therefore cannot be interpreted as evidence of real-world performance gains.

2.2. Data Gravity and Knowledge Gravity

Data Gravity

SOFIA adopts Data Gravity by moving computation toward data sources whenever possible, minimizing data transfers and communication overhead.

Knowledge Gravity

We define Knowledge Gravity as the tendency of actionable knowledge assets to migrate toward locations where decisions are generated and operational value is created. Building on the concept of Data Gravity [14,15], Knowledge Gravity extends the notion of proximity from data placement to knowledge placement.

According to the Knowledge Gravity Principle, not only data but also knowledge assets may be attracted toward the locations where they create the greatest operational value. These assets can include ontology fragments, policies, rule sets, Knowledge Graph subgraphs, reasoning services, contextual knowledge, and cached semantic states. Their placement is influenced by factors such as decision locality, latency requirements, data sovereignty constraints, update frequency, available compute capacity, consistency costs, and privacy requirements.

Within SOFIA, Knowledge Gravity acts as a cognitive orchestration principle that guides the distribution and placement of knowledge resources across the Edge-Cloud Continuum. By relocating selected knowledge assets closer to points of decision making, it supports the evolution from an

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org | NOT PEER-REVIEWED | Posted: 20 November 2026 https://doi.org/10.20944/preprints202611.2231.v1

infrastructure-centric Edge-Cloud Continuum toward a Cognitive Continuum in which intelligence, reasoning, and governance capabilities can operate where they deliver the highest value.

Scope of this Position Paper. This work is intended as a position paper that introduces the concept of Knowledge Gravity and its potential role within distributed neuro-symbolic AI systems. The proposed definitions, metrics, and illustrative evaluations should be interpreted as a conceptual framework aimed at stimulating discussion and future research. While a preliminary mathematical formulation and example scenarios are provided, extensive empirical validation remains part of future work.

2.2.1. Formalization of Knowledge Gravity

Formal Definition

Definition 1 (Knowledge Gravity).

Let $K = \{K_1, \dots, K_m\}$ be a set of knowledge assets and $N = \{n_1, \dots, n_p\}$ a set of computational nodes distributed across the Cognitive Continuum.

Knowledge Gravity (KG) is defined as the tendency of knowledge assets to be placed, replicated, synchronized, or migrated toward the nodes where their contribution to decision utility exceeds the lifecycle costs required to maintain them. Knowledge assets may include ontology fragments, knowledge-graph subgraphs, symbolic rules, governance policies, semantic world-model components, and reasoning services. Unlike Data Gravity, which attracts computation toward data, Knowledge Gravity governs the placement of semantic knowledge and reasoning capabilities in relation to decision-making activities.

First-Order Conceptual Model

For explanatory purposes, we first consider a single knowledge asset $K_i$deployed on node n_j:

$$KG(K_i, n_j) = V_{ij} - C_{ij}$$

where V_{ij}represents the expected operational value generated by the asset on node n_j, and C_{ij}represents the lifecycle cost associated with maintaining that asset on the same node.

$$V_{ij} = \alpha L_{ij} + \beta F_{ij} + \gamma X_{ij} + \delta P_{ij}$$

where L_{ij}denotes decision-locality benefit, F_{ij}frequency of utilization, X_{ij}contribution to decision effectiveness, and P_{ij}policy and compliance relevance.

$$\alpha + \beta + \gamma + \delta = 1$$

Lifecycle cost is represented as:

$$C_{ij} = \lambda T_{ij} + \mu S_{ij} + \nu M_{ij} + \omega R_{ij}$$

where T_{ij}denotes transfer cost, S_{ij}synchronization cost, M_{ij}storage cost, and R_{ij}computational reasoning cost.

$$\lambda + \mu + \nu + \omega = 1$$

The preferred location becomes:

$$n^* = \arg\max_{n_j} (V_{ij} - C_{ij})$$

This formulation is intentionally simplified and should be interpreted as a first-order conceptual representation rather than a complete distributed optimization model. The estimation and calibration of V_{ij}remain open research questions. In practical deployments, value functions may depend on domain-specific utility models, operational objectives, risk considerations, and stakeholder preferences. The definition of general estimation procedures is left to future work.

Multi-Asset Placement and Constraints

Real distributed cognitive systems involve multiple interconnected knowledge assets deployed across multiple nodes.

Let

$$x_{ij} \in \{0, 1\}$$

indicate whether asset $K_i$is deployed on node n_j.

The global optimization problem becomes:

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org | NOT PEER-REVIEWED | Posted: 20 November 2026 https://doi.org/10.20944/preprints202611.2231.v1

$$X^* = \arg\max_X \left(\sum_{i=1}^{m} \sum_{j=1}^{p} x_{ij} V_{ij} - C_{\text{total}}(X) \right)$$

where X denotes the placement matrix and $C_{\text{total}}(X)$ includes deployment, storage, transfer, synchronization, communication, governance, and consistency costs.

Knowledge assets are often interdependent. An ontology fragment may require specific rules, entities, taxonomies, or knowledge-graph structures to remain operationally useful. Dependencies may be represented by a graph:

$$G_K = (K, E)$$

where an edge $(K_i, K_j) \in E$ indicates that asset K_i depends on asset K_j.

A feasible placement must satisfy dependency constraints such as:

$$x_{ij} \leq x_{kj}$$

when asset K_i requires asset K_k.

Synchronization costs may depend on replication patterns and topology:

$$C_{sync} = f(R, N, U)$$

where R is the number of replicas, N the network topology, and U the update frequency. Consequently, maximizing locality does not necessarily maximize system utility.

Unlike traditional placement paradigms, the objective is not simply resource optimization but the maximization of decision utility under constraints related to semantic completeness, governance compliance, synchronization effort, explainability, and sovereignty requirements.

The proposed formulation should be interpreted as a general optimization framework intended to capture the principal dimensions of Knowledge Gravity rather than as a fully specified and directly implementable optimization model. Additional constraints concerning node capacity, replication policies, geographical placement requirements, remote dependency management, and operational resource allocation remain subjects for future formalization and empirical validation.

Placement Metrics

To characterize placement quality, five placement-oriented metrics are introduced:

- Knowledge Locality Ratio (KLR)
- Decision Locality Ratio (DLR)
- Knowledge Transfer Volume (KTV)
- Knowledge Synchronization Overhead (KSO)
- Knowledge Placement Latency (KPL)

Knowledge Locality Ratio (KLR)

Measures the proportion of knowledge required by a decision process that is already available locally.

$$KLR = \frac{Local\ Knowledge\ Used}{Total\ Knowledge\ Used}$$

$$0 \leq KLR \leq 1$$

Higher values indicate stronger knowledge locality.

Decision Locality Ratio (DLR)

Measures the proportion of decisions that can be completed without remote knowledge access.

$$DLR = \frac{Locally\ Satisfied\ Decisions}{Total\ Decisions}$$

$$0 \leq DLR \leq 1$$

Higher values indicate greater decision autonomy.

Knowledge Transfer Volume (KTV)

Measures the amount of knowledge transferred across the continuum.

$$KTV = \sum_{i=1}^{n} Bytes(K_i)$$

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org | NOT PEER-REVIEWED | Posted: 20 November 2026 https://doi.org/10.20944/preprints202611.2231.v1

Lower values indicate more efficient placement strategies.

Knowledge Synchronization Overhead (KSO)

Measures the relative overhead associated with knowledge synchronization activities.

$$KSO = \frac{T_{sync}}{T_{operational}}$$

where:

- T_{sync} is synchronization time;
- $T_{operational}$ is total operational time.

Knowledge Placement Latency (KPL)

Measures the delay between a knowledge request and the availability of the requested knowledge.

$$KPL = t_{available} - t_{requested}$$

Lower values indicate more effective placement mechanisms.

KPQI Definition and Normalization

Because KTV, KSO, and KPL are measured in heterogeneous units, they are normalized using min-max normalization:

$$X_N = \frac{X - X_{\min}}{X_{\max} - X_{\min}}$$

where $X_{\min}$ and $X_{\max}$ are computed from the deployment alternatives considered within the same scenario. As a result, KPQI values are intended for relative comparison within a scenario and are not automatically comparable across different scenarios.

The Knowledge Placement Quality Index is defined as:

$$KPQI = w_1 KLR + w_2 DLR + w_3(1 - KTV_N) + w_4(1 - KSO_N) + w_5(1 - KPL_N)$$

subject to:

$$\sum_{i=1}^{5} w_i = 1$$

Higher values indicate more favorable placement configurations. KPQI should be interpreted as a placement-quality indicator rather than a direct measurement of Knowledge Gravity.

Relationship Between Global Optimization and KPQI

The Knowledge Gravity optimization framework and the Knowledge Placement Quality Index (KPQI) serve different roles within SOFIA. The multi-asset optimization model introduced in Section 2.2.1 constitutes the primary decision mechanism. Its objective is to maximize global operational utility generated by the placement of knowledge assets while minimizing lifecycle costs, subject to dependency, consistency, governance, sovereignty, security, localization, capacity, and resource constraints.

KPQI does not replace the primary Knowledge Gravity optimization problem and should not be interpreted as the objective function governing placement decisions. The constrained optimization model remains the principal decision mechanism and seeks to maximize global operational utility while satisfying all mandatory constraints.

Let f^* denote the optimal value of the primary objective function. A feasible placement configuration X is considered near-optimal when:

$$f(X) \geq f^* - \varepsilon$$

where ε is an application-dependent tolerance expressed in the same utility units as the objective function. The set of admissible near-optimal solutions is therefore:

$$\mathcal{N}_\varepsilon = \{X \in \mathcal{F} : f(X) \geq f^* - \varepsilon\}$$

where $\mathcal{F}$ denotes the set of feasible configurations satisfying all dependency, governance, sovereignty, security, localization, consistency, and resource constraints.

KPQI is applied only within this admissible near-optimal set. Its role is not to compensate for substantial losses in utility, but to express stakeholder preferences among solutions that are equivalent or sufficiently close to the optimum with respect to the primary Knowledge Gravity objective. In this sense, KPQI acts as a secondary preference and ranking criterion that helps distinguish among

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org | NOT PEER-REVIEWED | Posted: 20 November 2026 https://doi.org/10.20944/preprints202611.2231.v1

admissible alternatives exhibiting comparable decision utility while differing in placement characteristics such as locality, synchronization overhead, or latency.

Within the operational workflow, the Cognitive Control Layer and the Policy Engine first evaluate admissibility conditions and enforce all mandatory constraints before a placement alternative can be considered feasible. Candidate configurations that violate dependency relationships, sovereignty requirements, security policies, localization obligations, governance rules, consistency requirements, or resource limitations are rejected regardless of their KPQI value.

Accordingly, KPQI is primarily interpreted as a ranking and assessment metric applied to the subset of admissible solutions rather than as an independent optimization objective. A configuration with a higher KPQI score is preferred only when all mandatory constraints are already satisfied. Constraint violations cannot be compensated for by improvements in locality, latency, synchronization efficiency, or any other KPQI component. In particular, security, sovereignty, regulatory, dependency, and localization requirements are treated as feasibility constraints rather than trade-off variables.

Consequently, the optimization process operates in two conceptually distinct stages. First, feasible placement configurations are identified through constrained Knowledge Gravity optimization. Second, KPQI may be used to compare, rank, or evaluate admissible alternatives from a placement-quality perspective. In this sense, KPQI should therefore be interpreted primarily as an ex-post assessment and ranking metric. The primary placement decision remains the constrained maximization of decision utility, while KPQI serves as a secondary preference criterion among feasible deployment alternatives.

Complementary Outcome Metrics

Knowledge placement represents only one dimension of distributed cognition. Explainability, sustainability, and governance are evaluated through complementary indicators:

- Explainable Decision Coverage (EDC)
- Energy per Governed Decision (EGD)
- Governance Compliance Ratio (GCR)

EDC measures reasoning-trace coverage rather than explainability quality; GCR measures policy-validation coverage rather than governance effectiveness or overall governance quality, as high GCR values do not guarantee a direct relationship with superior governance outcomes. These metrics assess distinct constructs and should not be interpreted as direct consequences of KPQI values.

Illustrative Simulation and Sensitivity Analysis

The numerical example presented below is a synthetic sensitivity analysis intended solely to illustrate how KPQI responds to alternative weighting assumptions and placement trade-offs. The metric values assigned to Edge-1, Edge-2, and Edge-3 are hypothetical and were selected to represent plausible deployment configurations rather than measurements from real systems. Consequently, the results should be interpreted as illustrative examples of the proposed evaluation framework and do not constitute empirical validation of SOFIA, Knowledge Gravity, or their expected operational benefits.

Simulation Setup

A diagnostic knowledge graph fragment is evaluated on three alternative edge deployments:

Node	Characteristics
Edge-1	Maximizes local knowledge availability and autonomous decision-making
Edge-2	Minimizes synchronization overhead and knowledge-placement latency
Edge-3	Represents a balanced configuration

The normalized metrics used by KPQI are reported below.

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org | NOT PEER-REVIEWED | Posted: 20 November 2026 https://doi.org/10.20944/preprints202611.2231.v1

Metric	Edge-1	Edge-2	Edge-3
KLR	0.95	0.70	0.82
DLR	0.92	0.75	0.84
1-KTVN	0.88	0.72	0.80
1-KSON	0.45	0.95	0.75
1-KPLN	0.50	0.93	0.78

Edge-1 provides high knowledge locality and decision autonomy through extensive local replication but incurs increased synchronization overhead and placement latency. Conversely, Edge-2 minimizes consistency-maintenance costs and latency through a leaner local knowledge footprint, at the expense of lower decision locality. Edge-3 provides an intermediate compromise between both strategies.

Assuming equal weights:

$$w_1 = w_2 = w_3 = w_4 = w_5 = 0.2$$

the resulting KPQI values are:

Node	KPQI
Edge-1	0.7400
Edge-2	0.8100
Edge-3	0.7980

Under balanced priorities, Edge-2 achieves the highest score because its gains in synchronization efficiency and latency outweigh its lower locality.

Weight Sensitivity Analysis

To explore ranking stability, two alternative weighting schemes are considered.

Locality-Oriented Configuration

$$w = (0.35, 0.30, 0.15, 0.10, 0.10)$$

This configuration prioritizes local knowledge availability and autonomous decision capability.

Node	KPQI
Edge-1	0.8355
Edge-2	0.7660
Edge-3	0.8120

Under these conditions, Edge-1 becomes the preferred deployment option.

Efficiency-Oriented Configuration

$$w = (0.10, 0.15, 0.15, 0.30, 0.30)$$

This configuration emphasizes synchronization efficiency and rapid knowledge availability.

Node	KPQI
Edge-1	0.6500
Edge-2	0.8545
Edge-3	0.7870

Edge-2 clearly dominates when operational efficiency is prioritized.

Ranking Reversal and Weight Sensitivity

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org | NOT PEER-REVIEWED | Posted: 20 November 2026 https://doi.org/10.20944/preprints202611.2231.v1

Unlike the previous illustrative scenario, none of the three deployment alternatives dominates the others across all KPQI dimensions. Consequently, the preferred deployment depends on stakeholder priorities and on the relative importance assigned to locality-oriented versus efficiency-oriented objectives. This creates the possibility of ranking reversals, which are a natural consequence of multi-criteria decision problems involving explicit trade-offs among competing objectives.

To investigate this behavior, we consider a one-parameter family of weighting configurations in which the combined weight assigned to locality-oriented metrics (KLR and DLR) is progressively varied, while the remaining weight is distributed among the efficiency-oriented metrics (1−KTVN, 1−KSON, and 1−KPLN). Under this illustrative scenario, three distinct preference regions emerge.

Let α denote the total weight assigned to locality-oriented metrics. For the illustrative analysis, locality weights are distributed equally between KLR and DLR, while the remaining weight is distributed equally among the efficiency-oriented metrics. Therefore:

$$w_1 = w_2 = \alpha/2$$
$$w_3 = w_4 = w_5 = (1-\alpha)/3$$

The preference thresholds reported below are derived under this specific weighting family and should not be interpreted as general properties of KPQI.

Under the weighting scheme defined by $w_1 = w_2 = \alpha/2$ and $w_3 = w_4 = w_5 = (1-\alpha)/3$, two preference thresholds emerge at approximately 46.15% and 61.35%, corresponding to the points at which the preferred deployment alternative changes from efficiency-oriented to balanced, and from balanced to locality-oriented configurations.

The two transition points correspond to exact score ties between adjacent alternatives. At $\alpha \approx$ 46.15%, Edge-2 and Edge-3 achieve the same KPQI value, while at $\alpha \approx$ 61.35%, Edge-3 and Edge-1 achieve the same KPQI value. The preferred alternative changes only when α moves beyond these tie points.

Region 1: Efficiency-Dominated Priorities (0%-46.15%)

When the combined weight assigned to KLR and DLR remains below approximately **46.15%**, Edge-2 achieves the highest KPQI value. In this region, operational priorities favor efficient knowledge management, reduced synchronization overhead, and low placement latency. The superior performance of Edge-2 on the cost-related indicators more than compensates for its lower locality and decision autonomy. Consequently, Edge-2 represents the preferred deployment strategy whenever operational efficiency is the dominant objective.

Region 2: Balanced Priorities (46.15%-61.35%)

When the locality weight lies between approximately **46.15% and 61.35%**, Edge-3 becomes the highest-ranked alternative. This configuration represents a compromise between locality and efficiency objectives. The result is particularly relevant because it shows that a balanced deployment can outperform both extreme strategies when stakeholders assign comparable importance to decision autonomy and operational efficiency. In this intermediate region, the trade-offs embedded in the KPQI formulation favor the balanced characteristics of Edge-3 rather than the specialization of Edge-1 or Edge-2.

Region 3: Locality-Dominated Priorities (61.35%-100%)

When the combined locality weight exceeds approximately **61.35%**, Edge-1 becomes the preferred deployment alternative. In this region, decision autonomy, local knowledge availability, and local reasoning capabilities are given greater importance than synchronization efficiency and latency. The higher KLR and DLR values achieved by Edge-1 therefore outweigh its increased synchronization and placement costs. As a result, Edge-1 becomes the optimal strategy whenever maximizing cognitive locality is the primary design objective.

These results demonstrate that KPQI is not ranking-invariant. Rather than indicating a weakness of the metric, the observed ranking reversals reveal its ability to represent alternative operational priorities and deployment strategies. The metric therefore behaves consistently with the underlying trade-offs that characterize distributed knowledge placement problems.

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org | NOT PEER-REVIEWED | Posted: 20 November 2026 https://doi.org/10.20944/preprints202611.2231.v1

The analysis also provides an indication of sensitivity to stakeholder preferences. In this example, the width of the weight interval over which a deployment remains the preferred solution represents a deterministic sensitivity measure rather than a statistical estimate of robustness, reliability, or out-of-sample stability. Edge-2 remains preferred across a broad efficiency-oriented region, Edge-1 across a broad locality-oriented region, whereas Edge-3 dominates a narrower intermediate region. These results simply characterize how rankings respond to changes in the weighting assumptions adopted in this illustrative scenario.

It is important to emphasize that the threshold values **46.15%** and **61.35%** are not intrinsic properties of Knowledge Gravity, KPQI, or distributed cognitive systems in general. They arise from the specific synthetic metrics, normalization procedure, and weighting assumptions adopted in this illustrative example. Different deployment scenarios, metric values, constraints, or stakeholder preferences would produce different transition points and potentially different preference regions. Consequently, the numerical thresholds reported here should be interpreted as scenario-dependent illustrative results rather than universal characteristics of Knowledge Gravity.

Future work should extend this analysis through Pareto-front exploration, uncertainty analysis, Monte Carlo simulations, and multi-objective optimization techniques in order to characterize the stability and sensitivity of Knowledge Gravity placement decisions under realistic operational conditions.

Distinction from Existing Placement Paradigms

Existing placement paradigms primarily optimize the location of computation, services, data, or AI models. Typical objective functions focus on minimizing latency, communication costs, energy consumption, response time, or resource utilization. At the same time, several research streams, including distributed semantic reasoning, semantic caching, knowledge graph engineering, and federated knowledge management, address complementary aspects of knowledge representation, synchronization, consistency, and distributed reasoning.

Knowledge Gravity should not be interpreted as a replacement for these approaches. Rather, it can be viewed as an integrated extension that combines concepts originating from Data Gravity, distributed semantic reasoning, federated knowledge management, and governance-aware AI within a unified decision-centric optimization framework.

The objective is not limited to optimizing computational efficiency or knowledge accessibility, but to determining how semantically interconnected knowledge assets should be placed, replicated, synchronized, and governed to maximize decision utility while respecting operational and regulatory constraints. A generic formulation can be expressed as:

$$\max \quad DecisionUtility(K, N)$$

subject to:

knowledge locality;

- knowledge dependency constraints;
- consistency requirements;
- synchronization costs;
- governance policies;
- regulatory constraints;
- sovereignty requirements;
- storage and compute limitations;
- communication constraints.

Unlike most existing placement formulations, Knowledge Gravity explicitly incorporates ontologies, knowledge graph subgraphs, symbolic rules, governance policies, and reasoning services as optimization entities. It also considers semantic dependencies among knowledge assets and the trade-offs between decision autonomy, governance obligations, synchronization effort, explainability requirements, and operational efficiency.

In this sense, the contribution of Knowledge Gravity is not the introduction of an entirely new class of distributed systems, but the formulation of a unified optimization perspective that integrates

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org | NOT PEER-REVIEWED | Posted: 20 November 2026 https://doi.org/10.20944/preprints202611.2231.v1

knowledge placement, reasoning availability, governance-by-design, and cognitive autonomy as first-class design objectives within distributed cognitive systems, as shown in Table 1.

Table 1. Comparison of distributed-system optimization paradigms.

Paradigm	Optimized Entity	Objective Function	Typical Constraints
Data Gravity	Computation relative to data	Minimize data movement and latency	Bandwidth, storage, compute
Service Placement	Software Services	Performance and resource efficiency	CPU, memory, network
AI Model Placement	Training and inference models	Inference latency, throughput, energy	Compute, memory, energy
Knowledge Gravity	Knowledge assets, rules, ontologies, governance policies	Maximize decision utility and decision autonomy	Dependencies, consistency, synchronization, governance, sovereignty

Limitations

The proposed formulation remains conceptual and has not yet been empirically calibrated. Dependency relationships are represented only at a high level and are not yet associated with a complete graph-theoretic optimization framework. Furthermore, synchronization and consistency costs rely on simplified abstractions; realistic deployments will likely require topology-aware, stochastic, and workload-dependent models. These limitations define important directions for future research on Knowledge Gravity and the SOFIA Cognitive Continuum.

Example Research Question Enabled by Knowledge Gravity

A representative optimization question that is not explicitly formulated by existing placement paradigms is:

Given a distributed cognitive system composed of edge, fog, and cloud nodes, how should interconnected knowledge assets (ontologies, knowledge graphs, rules, policies, and digital twins) be replicated, synchronized, and migrated in order to maximize decision autonomy while respecting sovereignty, governance, consistency, and resource constraints?

2.2.2. Knowledge Gravity versus Semantic Caching and Distributed Knowledge Management

Knowledge Gravity shares certain characteristics with semantic caching, distributed knowledge management, and federated knowledge systems. However, it addresses a distinct architectural concern.

Semantic caching aims primarily at improving query performance by storing previously accessed data, reasoning results, embeddings, or semantic query responses close to anticipated consumers. The placement of cached artifacts is typically driven by access patterns, locality of reference, and performance optimization objectives.

Distributed knowledge management focuses on the storage, synchronization, governance, and lifecycle management of knowledge assets across multiple repositories or organizational boundaries. Its primary concern is knowledge consistency, sharing, accessibility, and lifecycle evolution.

Knowledge Gravity, as proposed in this paper, extends beyond both concepts. It does not merely optimize retrieval performance nor only coordinate distributed knowledge repositories. Instead, it provides an orchestration principle governing where knowledge assets should reside in relation to decision-making activities.

Under the Knowledge Gravity principle, knowledge assets may include:

- ontology fragments;
- knowledge graph subgraphs;

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org | NOT PEER-REVIEWED | Posted: 20 November 2026 https://doi.org/10.20944/preprints202611.2231.v1

- symbolic rules;
- governance policies;
- regulatory constraints;
- reasoning services;
- semantic world-model components.

Their placement is determined by an optimization objective balancing:

- decision locality;
- explainability requirements;
- governance relevance;
- sovereignty constraints;
- synchronization costs;
- operational efficiency.

Consequently, semantic caching can be interpreted as a possible mechanism supporting Knowledge Gravity, while distributed knowledge management provides part of the enabling infrastructure. Neither concept alone captures the broader objective of dynamically positioning knowledge and reasoning capabilities according to where operational decisions are generated.

In this sense, Knowledge Gravity acts as a cognitive orchestration principle for distributed intelligence, whereas semantic caching addresses performance optimization and distributed knowledge management addresses knowledge lifecycle coordination.

2.2.3. Orchestration and Consistency Costs in Distributed Knowledge Gravity

While Knowledge Gravity promotes the migration of semantic knowledge assets toward decision points, knowledge placement is not cost-free. The orchestration of distributed knowledge assets introduces additional overhead associated with synchronization, consistency maintenance, policy propagation, reasoning-state updates, and governance enforcement across heterogeneous nodes.

In particular, distributed Knowledge Graphs may experience increasing synchronization costs as the number of participating nodes grows or as connectivity becomes intermittent. In highly dynamic environments, ontology updates, rule modifications, policy changes, and semantic state transitions must be propagated across the continuum, potentially generating communication overhead that offsets part of the benefits achieved through local decision-making.

The challenge is particularly evident in balanced edge environments, such as the Edge-3 scenario introduced in Section 2.2.1, which represents an intermediate compromise between local decision autonomy and synchronization efficiency. Under such configurations, balancing moderate synchronization overhead and consistency operations becomes a key component of lifecycle cost C_{ij}.

Consequently, Knowledge Gravity should be interpreted as an optimization problem balancing:

$$NetBenefit = V_{ij} - \left(C_{sync} + C_{transfer} + C_{reasoning} + C_{storage}\right)$$

where consistency maintenance represents a first-order design constraint rather than a secondary implementation detail.

Future implementations of SOFIA may adopt hybrid consistency models combining:

- Strong consistency for governance policies, safety rules, and regulatory constraints;
- Eventual consistency for domain knowledge and contextual observations;
- Selective synchronization based on knowledge criticality, update frequency, and regulatory relevance.

This distinction aligns with the Knowledge Gravity principle itself: not all knowledge assets require identical synchronization guarantees, and governance-aware placement strategies may reduce unnecessary consistency costs while preserving trustworthiness and operational effectiveness.

2.2.4. Synchronization Failure Modes and Partition-Tolerant Operation

The distributed nature of the SOFIA Cognitive Continuum introduces the possibility of synchronization failures, prolonged network partitions, and temporary inconsistencies among distributed

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org | NOT PEER-REVIEWED | Posted: 20 November 2026 https://doi.org/10.20944/preprints202611.2231.v1

Knowledge Graph replicas. Such situations may arise in resource-constrained environments, mobile infrastructures, remote industrial sites, or intermittently connected edge nodes.

Under the Knowledge Gravity principle, edge nodes are expected to remain operational even when disconnected from Fog or Cloud coordination services. During network partitions, local cognitive nodes continue executing decisions using their most recent validated knowledge state while maintaining a detailed provenance log of locally generated updates.

A conceptual synchronization strategy may operate as follows:

Algorithm 1: Edge Node Autonomous Execution and Semantic Synchronization Protocol

```
01: If connectivity available:
02:     Synchronize ontology fragments
03:     Synchronize policy updates
04:     Synchronize reasoning traces
05:     Resolve conflicts using governance policies
06: Else:
07:     Enter autonomous mode
08:     Use last validated Knowledge Graph snapshot
09:     Log local knowledge modifications
10:     Restrict updates affecting critical governance rules
11: When connectivity restored:
12:     Exchange change sets
13:     Detect conflicting modifications
14:     Apply governance-aware reconciliation policies
```

To reduce split-brain conditions, future implementations may distinguish between different classes of knowledge assets:

- Governance policies, safety constraints, and regulatory rules may require strong consistency and centralized approval.
- Operational knowledge and contextual observations may follow eventual consistency models.
- Local reasoning traces may remain node-specific until reconciliation occurs.

Policy Validity During Network Partitions

The partition-tolerant strategy described above assumes that edge nodes may continue operating using the last validated knowledge snapshot when connectivity is temporarily unavailable. However, this approach introduces an important governance challenge: policies, regulatory rules, or safety constraints may expire, be revoked, or be superseded while the node remains disconnected.

Consequently, future implementations of SOFIA should distinguish between knowledge assets that remain valid during disconnected operation and governance-critical assets whose validity is time-bounded. Examples include revoked clinical guidelines, updated safety constraints, emergency regulatory directives, or security-related policy updates.

To address this limitation, governance-sensitive knowledge assets may include expiration metadata, version constraints, validity intervals, or mandatory revalidation requirements. When such conditions cannot be verified due to prolonged network partition, the system should transition to a conservative operating mode that restricts or suspends decision capabilities requiring potentially outdated governance artifacts. This approach reduces the risk of continuing operations based on obsolete or invalid policy knowledge.

The detailed specification of fail-safe transitions, assurance mechanisms, and differentiated treatment of high-risk autonomous decisions is outside the scope of this position paper and constitutes an important direction for future research.

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org | NOT PEER-REVIEWED | Posted: 20 November 2026 https://doi.org/10.20944/preprints202611.2231.v1

2.2.5. Operational Workflow of Knowledge Gravity

To illustrate the operational role of Knowledge Gravity within the SOFIA architecture, consider a healthcare scenario involving a physician evaluating a patient with suspected sepsis in a hospital connected to the Cognitive Continuum. The workflow links the mathematical formulation introduced in Section 2.2.1 with the software components described in Sections 2.1, 2.3 and 2.5.

Step 1: Decision Request and Knowledge Asset Identification

A diagnostic request is initiated from the hospital edge node. The Neuro-Symbolic World Model identifies the knowledge assets required to support the decision, including clinical ontologies, diagnostic rules, treatment guidelines, governance policies, provenance constraints, and relevant Knowledge Graph fragments. The required asset set constitutes the candidate knowledge bundle for placement evaluation.

Step 2: Node Evaluation and Constraint Assessment

The Cognitive Control Layer collects placement indicators and operational constraints from Edge, Fog, and Cloud nodes. Candidate placements are evaluated according to expected utility, locality requirements, synchronization costs, sovereignty constraints, node capacity, policy restrictions, and dependency relationships among knowledge assets. The Policy Engine verifies compliance constraints before any placement decision is accepted.

Step 3: Placement Optimization

Using the Knowledge Gravity optimization framework, the system determines the placement configuration that maximizes decision utility while respecting the applicable constraints. The optimization process may choose to replicate, migrate, cache, or retain individual knowledge assets at Edge, Fog, or Cloud locations depending on the expected operational benefit and associated lifecycle costs.

For illustrative purposes, the workflow assumes that expected utility estimates are available. The mechanisms used to derive or calibrate such estimates are not specified by the current framework and remain part of the future Knowledge Gravity research agenda.

Step 4: Authorization and Knowledge Distribution

Once a candidate placement configuration has been identified, governance services authorize the deployment according to jurisdictional requirements, data-protection obligations, security policies, and organizational governance rules. Authorized knowledge assets are then distributed or synchronized to the selected nodes through the continuum orchestration mechanisms.

Step 5: Decision Execution and Traceability

The local edge node performs neuro-symbolic reasoning using the deployed knowledge assets. Neural components process clinical observations, while symbolic components evaluate medical rules, ontological constraints, and governance requirements. The resulting recommendation is accompanied by reasoning traces, policy validations, provenance records, and explainability artifacts that support auditing and human oversight.

Step 6: Synchronization and Knowledge Evolution

After decision execution, validated updates, new observations, reasoning traces, and governance-relevant events are synchronized across the Cognitive Continuum according to the applicable consistency model. Knowledge Graphs, rules, policies, and world-model components evolve continuously while preserving traceability and governance requirements.

Figure 2 below illustrates the neuro-symbolic decision workflow across the cognitive continuum. Diagnostic requests trigger knowledge asset identification, followed by multi-node evaluation, placement optimization, governance validation, neuro-symbolic decision execution, and continuous knowledge synchronization across Edge, Fog, and Cloud environments. The system acts through a continuous **Decision-to-Action** loop, composed by perception, decision, action planning, execution and feedback.

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org | NOT PEER-REVIEWED | Posted: 20 November 2026 https://doi.org/10.20944/preprints202611.2231.v1

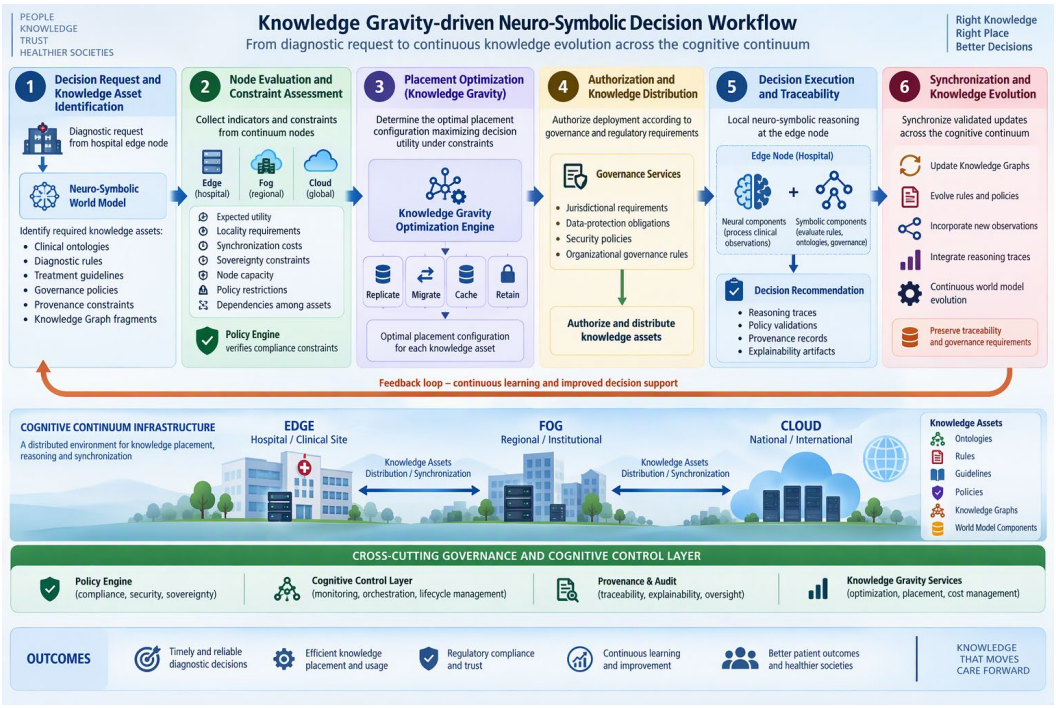


Figure 2. Knowledge Gravity-driven neuro-symbolic decision workflow.

Levels of Autonomy in Healthcare Scenarios

The healthcare workflow presented in this paper should be interpreted primarily as a clinical decision-support scenario rather than as an example of fully autonomous medical decision making. Different levels of operational autonomy require different governance and oversight mechanisms. Three categories can be distinguished:

- **Informational recommendations**, where the system provides relevant knowledge, guidelines, and evidence to healthcare professionals without proposing a specific course of action.
- **Clinical decision-support recommendations**, where neuro-symbolic reasoning generates ranked diagnostic or treatment suggestions accompanied by reasoning traces, provenance records, and policy validations. Human professionals remain responsible for reviewing, accepting, modifying, or rejecting the recommendation before any clinical action is performed.
- **Autonomous operational actions**, where system-generated decisions directly trigger actions without prior human approval. Such scenarios require additional safety assurance, explicit authorization mechanisms, stronger governance controls, continuous monitoring, and context-specific regulatory assessment.

The healthcare use case considered in this paper belongs to the second category. Under normal operation, system outputs are intended to support clinical judgment rather than replace it. Furthermore, during degraded operating conditions, synchronization failures, missing governance updates, or uncertainty regarding policy validity, SOFIA should adopt conservative execution policies, including escalation to human review, suspension of automated recommendations, or restriction of actions requiring up-to-date governance validation.

This distinction is particularly important in regulated domains, where explainability, traceability, accountability, and human oversight remain independent system-level requirements regardless of the degree of local autonomy achieved through Knowledge Gravity.

2.3. Neuro-Symbolic Intelligence Model

Neural Layer

Provides:

- Perception

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org | NOT PEER-REVIEWED | Posted: 20 November 2026 https://doi.org/10.20944/preprints202611.2231.v1

- Pattern recognition
- Prediction
- Machine learning capabilities

Symbolic Layer

Provides:

- Rules
- Ontologies
- Compliance constraints
- Governance policies
- Decision logic

Hybrid Reasoning Mechanism

The system combines neural predictions with symbolic reasoning to generate explainable, auditable, and context-aware decisions.

2.3.1. Neuro-Symbolic World Models

The notion of Neuro-Symbolic World Models builds upon the broader world-model paradigm in AI [45], where internal representations of the environment are used for prediction and planning, but extends it by incorporating explicit symbolic knowledge, semantic memory, and explainable reasoning mechanisms.

SOFIA extends traditional neuro-symbolic architectures through the concept of Neuro- Symbolic World Models. Rather than representing knowledge exclusively as static symbolic structures, the framework is designed to maintain a continuously evolving representation of the operational environment, integrating observations, domain knowledge, policies, constraints, and contextual information.

These world models are expected to support situational awareness, reasoning, prediction, and action planning across the Edge-Cloud Continuum. Knowledge Graphs provide the semantic memory of the system, while neural components contribute perception and adaptation capabilities. Through their combination, SOFIA hypothesizes that distributed intelligence can evolve with reduced dependence on complete model retraining, although future evaluation will be required to quantify these effects.

2.4. Continuous Knowledge Graph Evolution

The Knowledge Graph infrastructure consists of:

- Ontology layer
- Domain knowledge layer
- Rule layer
- Governance layer

Knowledge evolves through:

- Edge observations
- Expert validation
- Federated synchronization
- Automated semantic enrichment

This mechanism may reduce reliance on continuous and costly retraining.

The expected workflow would be the following:

- Data gathering
- Neural event detection
- Semantic interpretation
- Knowledge graph query execution
- Symbolic rule evaluation
- Governance policy verification
- Decision generation

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org | NOT PEER-REVIEWED | Posted: 20 November 2026 https://doi.org/10.20944/preprints202611.2231.v1

- Trace persistence and synchronization.

2.5. Governance and Explainability Framework

The governance framework integrates:

- Explainable AI services
- Decision provenance
- Data lineage
- Compliance assessment
- Human oversight
- Traceability mechanisms

Governance therefore becomes an active component of reasoning rather than a post-hoc auditing process.

Sovereign AI denotes the capability of an AI system to retain control over sensitive data, knowledge assets, and model deployment; to operate in accordance with local policies and jurisdictional constraints; to avoid mandatory reliance on a single external cloud provider; and to preserve decision-making authority within the relevant organizational or territorial domain.

2.5.1. Runtime Assurance

SOFIA is designed to incorporate runtime assurance mechanisms capable of supervising the behavior of distributed AI components during operation. Rather than limiting governance activities to design-time validation, assurance services are intended to monitor policy compliance, explainability requirements, operational constraints, and risk conditions across the continuum.

Runtime assurance is expected to provide:

- policy conformance monitoring;
- anomaly detection;
- confidence assessment;
- human-intervention triggering;
- safe fallback strategies.

We hypothesize that integrating assurance into the operational lifecycle may improve trustworthiness and governance effectiveness, although future experimental validation is required.

2.5.2. Regulatory Mapping of Knowledge Gravity Metrics

To strengthen the practical applicability of the SOFIA framework, selected Knowledge Gravity metrics can be associated with specific requirements of the EU AI Act. Although the proposed metrics do not constitute direct compliance evidence, they may provide quantitative indicators supporting conformity assessment, governance monitoring, and operational assurance activities.

Interpretation of Regulatory Mapping

The regulatory mapping presented in Table 2 should be interpreted as an illustrative relationship between SOFIA mechanisms and selected regulatory objectives rather than as evidence of legal compliance. In particular, individual technical indicators are not intended to serve as direct measures of conformity with the requirements of the EU AI Act.

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org | NOT PEER-REVIEWED | Posted: 20 November 2026 https://doi.org/10.20944/preprints202611.2231.v1

Table 2. Illustrative Mapping of SOFIA Knowledge Gravity Metrics to Selected EU AI Act Requirements and Supporting Operational Evidence.

Regulatory Requirement (Normative Obligation)	**SOFIA Mechanism** (Proposed Support Functionality)	**Indicator** (Available Technical Measure)	**Potential Supporting Evidence** (Information Necessary for Effective Verification)
Article 13 (Transparency and Provision of Information)	Explainable Reasoning & Traceability Engine	Explainable Decision Coverage (EDC)	Reasoning-trace coverage and explanation logs for verifying the transparency of decision-making processes.
Article 13 & 15 (Transparency, Accuracy & Robustness)	Local Knowledge Management & Provenance Tracking	Knowledge Locality Ratio (KLR)	Local knowledge availability and semantic provenance (indirect support relationship to transparency and robustness).
Article 14 (Human Oversight)	Edge Cognitive Autonomy & Oversight Interfaces	Decision Locality Ratio (DLR)	Degree of local decision autonomy and human intervention points (enabling condition, not a direct measure of human oversight).
Article 15 (Accuracy, Robustness and Cybersecurity**)**	Continuum Synchronization & Caching Protocols	Knowledge Synchronization Overhead (KSO)	Synchronization overhead and operational stability under dynamic or degraded network conditions.
Article 15 (Accuracy, Robustness and Safeguards)	Distributed Knowledge Placement & Migration	Knowledge Placement Latency (KPL)	Availability latency of supportive knowledge and associated regulatory safeguards.
Article 12 (Record Keeping)	Distributed Trace Persistence & Lineage Services	Decision Provenance Logs	Immutable provenance records, auditability, and decision history logs.
Article 9 (Risk Management System)	Policy Engine & Embedded Governance Rules	Governance Policies in Knowledge Graphs	Risk and compliance controls dynamically embedded into knowledge graph rules.
Article 14 (Human Oversight / Intervention)	Human-In-The-Loop Oversight & Validation Triggers	Human Validation Events	Tracking and logging of human intervention and validation events in decision workflows.

For example, the Decision Locality Ratio (DLR) measures the proportion of decisions that can be completed using locally available knowledge assets. While higher DLR values may support operational autonomy and resilience, they should not be interpreted as evidence of compliance with the human oversight requirements defined by Article 14 of the EU AI Act. Human oversight depends on system-level governance mechanisms, intervention capabilities, escalation procedures, operator authority, and organizational controls that cannot be inferred from a single placement metric.

More generally, regulatory compliance must be assessed at the level of the complete socio-technical system and with respect to the specific obligations applicable to the deployment context. KPQI,

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org | NOT PEER-REVIEWED | Posted: 20 November 2026 https://doi.org/10.20944/preprints202611.2231.v1

KLR, DLR, KSO, KPL, EDC, and GCR may provide useful operational indicators, but none of them individually demonstrates conformity with legal, regulatory, or governance requirements.

A complementary mapping can also be established with GDPR principles, particularly data minimization, accountability, purpose limitation, and data sovereignty, as knowledge-centric distribution may reduce the need for transferring personal data toward centralized infrastructures.

2.5.3. Security Governance of Distributed Knowledge Assets

A key challenge of Knowledge Gravity is that valuable knowledge assets, including ontology fragments, symbolic rules, governance policies, and Knowledge Graph subgraphs, may be deployed on edge infrastructures that cannot always be assumed fully trusted.

To address this challenge, SOFIA adopts a Governance-by-Design perspective in which knowledge assets are treated as protected governance artifacts rather than simple data structures. Future implementations may combine multiple protection mechanisms, including:

- Cryptographic signing of ontology fragments, policies, and rule sets;
- Integrity verification prior to knowledge loading and execution;
- Trusted execution environments or secure enclaves where available;
- Immutable provenance records for knowledge modifications;
- Runtime assurance mechanisms able to detect unauthorized knowledge changes;
- Policy-conformance verification before reasoning execution.

Under this model, every knowledge asset can carry metadata describing its origin, version, approval status, and governance classification. Any modification that cannot be traced to an authorized source may be rejected, quarantined, or escalated for human review.

2.6. Potential Application Scenarios

Manufacturing

Advanced manufacturing environments are progressively evolving toward the principles of Industry 5.0, where human-centricity, sustainability, resilience, and intelligent automation become fundamental design objectives. In this context, distributed intelligence, explainability, and sovereign management of industrial knowledge are increasingly required to support adaptive and trustworthy industrial ecosystems.

Candidate validation scenarios include:

- Predictive maintenance
- Quality inspection
- Digital twins
- Energy optimization
- Production traceability

The SOFIA framework is designed to seamlessly integrate Neuro-Symbolic Intelligence into production environments, enabling the deployment of knowledge, policies, and reasoning capabilities close to industrial assets and operational processes. Through Knowledge Graphs, Knowledge Gravity, and distributed governance mechanisms, the framework supports explainable decision-making, resilient industrial operations, and alignment with Industry 5.0.

Agriculture

Modern agricultural ecosystems are increasingly characterized by the integration of IoT devices, environmental sensors, autonomous machinery, and data-driven decision support systems. These environments require distributed intelligence capable of optimizing resource utilization while preserving sustainability, resilience, and operational autonomy.

Candidate validation scenarios include:

- Precision farming
- Crop monitoring
- Irrigation optimization
- Fertilizer optimization

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org | NOT PEER-REVIEWED | Posted: 20 November 2026 https://doi.org/10.20944/preprints202611.2231.v1

- Autonomous rural intelligence

SOFIA is expected to bring agricultural intelligence closer to fields, sensors, machinery, and local production ecosystems, supporting context-aware and timely decision-making, reducing dependence on centralized infrastructures and improving responsiveness to environmental conditions. By combining neuro-symbolic reasoning with continuously evolving Knowledge Graphs, the framework supports resource-efficient agriculture, explainable recommendations, and sustainable management of farming operations across distributed rural environments.

6G and Cognitive Network Management

Future communication infrastructures represent a natural application domain for SOFIA. Emerging 6G networks are expected to integrate AI-native functions, distributed intelligence, autonomous orchestration, semantic communications, and digital sovereignty requirements.

Candidate validation scenarios include:

- AI-native network orchestration
- Autonomous network slicing
- Edge-based radio resource optimization
- Intent-based network management
- Semantic service discovery
- Distributed cybersecurity monitoring
- Federated network intelligence

In this context, SOFIA is designed to natively support the deployment of Neuro-Symbolic Intelligence directly across the communication continuum, where Knowledge Graphs can represent network topology, operational policies, service intents, and regulatory constraints. The combination of Knowledge Gravity and distributed reasoning supports self-adaptive and explainable network management while reducing reliance on centralized control functions.

Healthcare and Pharmaceutical Ecosystems

Healthcare and pharmaceutical environments require explainability, traceability, regulatory compliance, privacy preservation, and human oversight, making them highly relevant domains for the SOFIA framework.

Candidate validation scenarios include:

- Clinical decision support systems
- Personalized medicine
- Medical digital twins
- Hospital operational intelligence
- Remote patient monitoring
- Pharmaceutical supply-chain traceability
- Drug development knowledge management
- Pharmacovigilance and adverse-event detection

SOFIA is designed to keep sensitive medical data in close proximity to hospitals, laboratories, medical devices, and patient environments while distributing semantic knowledge and reasoning capabilities throughout the continuum. Neuro-symbolic reasoning combined with continuously evolving Knowledge Graphs may support regulatory-compliance processes. Human experts remain active participants within the cognitive loop through governance and validation mechanisms.

3. Expected Outcomes and Research Propositions

3.1. Edge Execution Capability

SOFIA is expected to reduce latency through edge-localized reasoning on smartphones, mini-PCs, industrial gateways, and embedded devices.

Expected benefits include:

- Reduced latency
- Lower connectivity dependence

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org | NOT PEER-REVIEWED | Posted: 20 November 2026 https://doi.org/10.20944/preprints202611.2231.v1

- Increased resilience
- Improved responsiveness

3.2. Reduced Training Dependency

Knowledge Graph updates, ontology evolution, and rule modifications are expected to support adaptation without repeated retraining.

We hypothesize that the framework may reduce the computational effort associated with maintaining AI systems. Future evaluation will be required to quantify the magnitude of this effect.

Proposition P1: For adaptations expressible through semantic, rule-based, contextual, or policy updates, Continuous Knowledge Evolution is expected to reduce neural retraining requirements, while Knowledge Gravity determines the placement of the required knowledge assets across the continuum.

3.3. Computational Efficiency

Distributed execution may deliver:

- Lower bandwidth consumption
- Improved workload distribution
- Faster local decisions
- Enhanced service continuity

3.4. Energy Efficiency

The combination of Data Gravity and Knowledge Gravity is expected to reduce data movement, and under specific deployment conditions, the computational effort associated with model adaptation. The net lifecycle energy impact remains to be experimentally assessed.

3.5. Explainability and Traceability

The framework may support:

- Transparent reasoning
- Decision provenance
- Semantic decision paths
- End-to-end auditability

Proposition P2: Systems employing distributed neuro-symbolic reasoning and locally available semantic knowledge are expected to provide richer reasoning traces and higher explanation coverage than purely neural edge architectures. Whether these improvements translate into better governance, compliance, or trust outcomes remains an empirical question for future evaluation.

Improvements in EDC should not be interpreted as equivalent to improvements in explainability quality. While greater reasoning-trace availability facilitates auditing and transparency, additional evaluation dimensions such as fidelity, comprehensibility, and decision-reconstruction capability remain necessary to establish effective explainability.

3.6. Compliance and Governance

SOFIA is designed to support governance and compliance requirements arising from GDPR, AI Act, NIS2, the Data Act, and the Data Governance Act through integrated governance, traceability, and assurance mechanisms. Formal compliance assessment remains part of future validation activities.

Similarly, improvements in GCR indicate that a larger proportion of decisions have been validated against available governance policies. Because a direct relationship between high GCR values and greater governance efficacy is not guaranteed, a high GCR value does not by itself demonstrate that policy repositories are complete, correct, legally sufficient, or continuously updated. The quality

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org | NOT PEER-REVIEWED | Posted: 20 November 2026 https://doi.org/10.20944/preprints202611.2231.v1

and true efficacy of governance therefore remain distinct evaluation dimensions requiring separate assessment procedures.

Proposition P3: We hypothesize that Governance-by-Design mechanisms may reduce human-hours required for audit preparation, increase the proportion of automatically verifiable controls, and reduce time required to reconstruct decision provenance and associated operational costs in regulated environments. Future evaluation will verify this proposition.

3.7. Potential Application Outcomes

Manufacturing

Expected benefits include:

- Predictive maintenance
- Energy optimization
- Production traceability
- Reduced operational costs

Agriculture

Expected benefits include:

- Water optimization
- Fertilizer reduction
- Autonomous farm intelligence
- Improved sustainability

6G Networks

Expected benefits include:

- Autonomous cognitive networking
- Reduced network latency
- Improved spectrum utilization
- Explainable network decisions
- Enhanced cyber-resilience
- Sovereign telecom infrastructures

Healthcare and Pharma

Expected benefits include:

- Improved clinical explainability
- Reduced healthcare data transfers
- Support for regulatory compliance by design
- Enhanced patient-data sovereignty
- Accelerated knowledge discovery
- End-to-end pharmaceutical traceability

3.7.1. Illustrative Use Case: Distributed Healthcare Intelligence

Consider a federation of hospitals connected through the SOFIA Cognitive Continuum.

Traditional Architecture

A hospital generates clinical observations:

- radiology images;
- laboratory results;
- medical records.

Data are transmitted to a centralized cloud where:

- AI inference occurs;
- medical knowledge bases reside;
- compliance verification is executed.

Challenges include:

- high latency;
- privacy risks;

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org | NOT PEER-REVIEWED | Posted: 20 November 2026 https://doi.org/10.20944/preprints202611.2231.v1

- regulatory complexity;
- growing data-transfer costs.

SOFIA with Knowledge Gravity

Instead of moving patient information toward centralized intelligence, the framework distributes:

- medical ontologies;
- diagnostic knowledge graph fragments;
- treatment protocols;
- AI Act compliance policies;
- privacy-preserving reasoning rules.

to the hospital edge infrastructure.

The knowledge assets become locally available where decisions are generated.

Example

A physician evaluates a patient with suspected sepsis.

The edge system performs:

1. Neural analysis of clinical observations.
2. Query against a local medical Knowledge Graph.
3. Rule-based validation against sepsis guidelines.
4. Compliance verification.
5. Generation of an explainable recommendation.

Only aggregated updates and validated knowledge modifications are synchronized with higher-level nodes.

Expected Outcomes

Compared with centralized architectures:

- Reduced decision latency.
- Lower data movement.
- Improved GDPR compliance.
- Increased explainability.
- Higher decision locality ratio.
- Reduced energy consumption associated with repeated cloud interactions.

Conceptually:

$$DLR_{SOFIA} > DLR_{Centralized}$$
$$Energy_{SOFIA} < Energy_{Centralized}$$
$$KGI_{SOFIA} > KGI_{Centralized}$$

Proposition P4. Systems implementing Knowledge Gravity-driven placement of knowledge assets are expected to achieve higher decision locality ratios, lower knowledge-transfer volumes, reduced governance latency, and improved explainability compared with architectures relying on centralized knowledge repositories.

Avoiding Causal Circularity. The SOFIA framework does not assume a deterministic causal relationship between distributed knowledge, explainability, governance, compliance, and trustworthiness. Instead, these properties are treated as distinct architectural objectives that may influence one another but require independent empirical evaluation. While Knowledge Gravity is hypothesized to facilitate explainability by making reasoning artifacts, policies, and semantic context locally available, explainability alone does not guarantee effective governance. Similarly, governance mechanisms may support compliance processes without ensuring regulatory conformity, and compliance does not automatically imply trustworthiness. Accordingly, Propositions P2-P4 should be interpreted as separate and testable research hypotheses rather than sequential consequences of a single architectural principle.

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org | NOT PEER-REVIEWED | Posted: 20 November 2026 https://doi.org/10.20944/preprints202611.2231.v1

4. Discussion

4.1. Positioning of SOFIA

SOFIA proposes a fundamental shift from distributed computing to distributed cognition. This shift can be interpreted as the evolution from an Edge-Cloud Continuum toward a Cognitive Continuum. While the Edge-Cloud Continuum primarily distributes data, computation, storage, and networking resources across heterogeneous infrastructures, the Cognitive Continuum extends this paradigm by distributing knowledge, reasoning capabilities, governance policies, and decision-making functions throughout the continuum.

In this perspective, Artificial Intelligence is not simply deployed on top of the infrastructure. Instead, intelligence and infrastructure co-evolve, transforming the continuum into a cognitive environment capable of supporting perception, reasoning, decision, action, governance, and continuous knowledge evolution.

The concept of Knowledge Gravity provides the orchestration principle governing this transition. Just as Data Gravity attracts computation toward data, Knowledge Gravity guides the placement, synchronization, and evolution of knowledge assets, reasoning services, policies, and compliance mechanisms according to contextual and operational needs. From this perspective, SOFIA positions itself as a cognitive layer that complements and extends existing Cloud-Edge infrastructures. The Cloud Infrastructure Services (CIS) paradigm provides the infrastructural foundation, while SOFIA introduces the distributed neuro-symbolic intelligence layer that is expected to support the transition from decision support toward governed decision-to-action systems. Consequently, the emergence of distributed, sovereign, and trustworthy AI requires a corresponding evolution of infrastructure from a compute-centric continuum to a cognitive continuum.

4.2. Comparison with Existing Research

Compared with existing Edge-Cloud initiatives, SOFIA is designed to extend infrastructure orchestration with distributed reasoning capabilities. Compared with Neuro-Symbolic AI research, it proposes an edge-native deployment perspective. Compared with traditional Knowledge Graph platforms, it aims to integrate distributed semantic memory into an edge-native cognitive architecture. Compared with Trustworthy AI frameworks, it proposes embedding governance directly into the cognitive layer. These capabilities should be interpreted as architectural design objectives and intended functionalities of the proposed framework rather than empirically validated properties.

In conclusion, this study proposes an integrated conceptual synthesis of several previously distinct research streams, including Edge-Cloud Continuum architectures, Neuro-Symbolic AI, Knowledge Graph engineering, governance-aware AI, and distributed knowledge management. The primary contribution lies in their unification within a decision-centric framework based on the Knowledge Gravity principle.

4.2.1. Comparison with Knowledge Placement and Distributed Semantic Systems

While the literature reviewed in Section 1 demonstrates the emergence of distributed reasoning, semantic caching, federated knowledge management, and large-scale Knowledge Graph infrastructures, several research efforts have also investigated the distribution, partitioning, replication, or placement of semantic assets across heterogeneous environments. These contributions represent the closest antecedents to the Knowledge Gravity concept and therefore require explicit comparison.

Existing approaches typically optimize a specific semantic artifact, such as RDF fragments, ontology modules, reasoning results, semantic caches, or Knowledge Graph partitions. Their objective functions are generally centered on query performance, consistency maintenance, storage efficiency, communication cost reduction, or reasoning scalability. In contrast, SOFIA formulates knowledge placement as a broader cognitive orchestration problem that simultaneously considers knowledge

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org | NOT PEER-REVIEWED | Posted: 20 November 2026 https://doi.org/10.20944/preprints202611.2231.v1

assets, reasoning services, governance policies, sovereignty constraints, explainability requirements, and decision utility.

Table 3 compares Knowledge Gravity with representative publications addressing federated querying, ontology partitioning, ontology modularization, distributed semantic repositories, and Knowledge Graph infrastructures. The comparison focuses on the primary optimization objective explicitly addressed in each publication and should not be interpreted as an exhaustive characterization of the broader research area. For each publication, Table 3 summarizes the specific semantic artifact considered, the problem explicitly addressed by the authors, the proposed distribution or partitioning mechanism, and its relationship to the broader orchestration perspective introduced by Knowledge Gravity. The comparison is limited to contributions directly documented in the cited publications and should not be interpreted as a comprehensive assessment of their potential applicability.

Table 3. Comparison with representative knowledge-distribution and semantic-placement approaches.

Work	Problem Addressed	Distributed Entity	Method	Difference with SOFIA
Harth et al. (YARS2) [46]	Federated querying of graph-structured semantic data distributed across multiple repositories.	RDF graphs and semantic repositories.	Federated repository architecture enabling distributed storage and query execution across graph datasets.	Focuses on repository federation and query access to distributed graph data. SOFIA extends the discussion to the placement of heterogeneous knowledge assets in relation to decision utility and operational constraints.
Montoya et al. (FederatedSPARQL) [47]	Efficient execution of SPARQL queries across multiple autonomous endpoints.	Distributed SPARQL endpoints and RDF datasets.	Fedra source-selection algorithm exploiting replicated RDF fragments in federated SPARQL processing.	Primarily concerned with query execution over already distributed semantic resources. SOFIA addresses placement and orchestration decisions concerning knowledge assets prior to query execution.
Saleem et al. (DAW) [48]	Source selection and duplicate-aware federated query processing over distributed RDF repositories, with the objective of minimizing redundant	Distributed RDF datasets and query results.	Duplicate-aware source selection using min-wise independent permutation summaries to estimate dataset overlap and avoid redundant endpoint access during federated	Addresses efficiency and correctness of federated query processing. SOFIA instead considers broader placement trade-offs involving decision autonomy, synchronization costs, and deployment constraints.

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org | NOT PEER-REVIEWED | Posted: 20 November 2026 https://doi.org/10.20944/preprints202611.2231.v1

	source access and improving query-execution efficiency.		SPARQL query processing.	
Stuckenschmidt & Klein [49]	Partitioning large concept hierarchies to improve scalability and manageability.	Ontology components and concept hierarchies.	Structure-based decomposition and partitioning of ontologies.	Focuses on how semantic structures can be partitioned. SOFIA considers how multiple knowledge assets may be placed, replicated and coordinated across the continuum under operational objectives and constraints.
d'Aquin, Schlicht, Stuckenschmidt, Sabou [50]	Ontology modularization for task-oriented knowledge selection.	Ontology modules.	Experimental assessment of ontology modularization methods for selecting relevant modules.	Focuses on ontology modularization and management. SOFIA incorporates ontology components as one class of assets within a broader placement and orchestration perspective.
SOFIA / Knowledge Gravity	Placement and orchestration of knowledge assets in distributed cognitive systems.	Ontologies, KG fragments, rules, policies, reasoning services, world-model components.	Multi-asset optimization and orchestration across Edge-Fog-Cloud environments.	Integrates placement decisions with decision utility, autonomy, synchronization costs, governance requirements, sovereignty constraints and explainability objectives.

The comparison suggests that Knowledge Gravity should not be interpreted as a replacement for distributed semantic reasoning, semantic caching, Knowledge Graph engineering, or federated knowledge management. Rather, it provides a higher-level orchestration perspective that integrates these capabilities into a unified optimization framework.

The comparison presented in Table 3 should be interpreted as conceptual positioning rather than as evidence of exclusivity or uniqueness. The selected works are representative examples of relevant research directions and do not constitute a systematic literature review. Consequently, the purpose of the comparison is to highlight differences in scope, optimization objectives, and architectural focus rather than to demonstrate the absence of alternative approaches in the broader literature. The distinctiveness of SOFIA resides in the proposed synthesis of these research streams within a common decision-centric framework based on Knowledge Gravity.

4.3. Scientific Implications

Three tightly interconnected concepts are brought together in this work:

- Knowledge Gravity
- Cognitive Continuum

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org | NOT PEER-REVIEWED | Posted: 20 November 2026 https://doi.org/10.20944/preprints202611.2231.v1

- Sovereign Neuro-Symbolic Intelligence

Together, these concepts describe a possible evolutionary trajectory for distributed digital infrastructures. Knowledge Gravity provides the orchestration principle, the Cognitive Continuum defines the target architectural paradigm, and Sovereign Neuro- Symbolic Intelligence represents the operational intelligence layer.

Their combined adoption may enable the evolution of future infrastructures from distributed computing environments toward distributed cognitive ecosystems.

SOFIA positions Knowledge Gravity and distributed neuro-symbolic governance as enabling mechanisms for realizing a Cognitive Continuum across heterogeneous infrastructures. Through these mechanisms, future distributed systems may progressively evolve into distributed cognitive ecosystems.

The framework should therefore be interpreted as a candidate integrative architecture whose distinctiveness lies in the proposed synthesis of existing research directions rather than in the exclusive introduction of entirely new individual components.

Table 4 provides a qualitative synthesis based on the representative literature reviewed in Sections 1.2.1–1.2.5; classifications refer to typical rather than universal coverage within each research stream. This table is intended for conceptual positioning rather than systematic evidence synthesis. SOFIA column denotes architectural coverage, not empirically validated capability. Accordingly, the SOFIA column identifies capabilities that the architecture is explicitly designed to support. Their effective realization and performance impact remain subject to future implementation and experimental validation.

Table 4. Conceptual coverage of capabilities across representative research streams and the proposed SOFIA architecture. Check marks in the SOFIA column indicate capabilities that the architecture is designed to support and are not intended as evidence of implemented and experimentally validated functionality.

Capability	Edge AI	Neuro-Symbolic AI	Knowledge Graph Platforms	Trustworthy AI Frameworks	SOFIA
Edge-native intelligence	✓	Limited coverage	not typically addressed	not typically addressed	✓
Distributed reasoning	Limited coverage	Limited coverage	not typically addressed	not typically addressed	✓
Knowledge evolution	not typically addressed	Limited coverage	✓	not typically addressed	✓
Explainability	Limited coverage	✓	✓	✓	✓
Governance-by-design	not typically addressed	not typically addressed	Limited coverage	✓	✓
Digital sovereignty	Limited coverage	not typically addressed	not typically addressed	Limited coverage	✓
Knowledge Gravity	not typically addressed	not typically addressed	not typically addressed	not typically addressed	✓

4.4. Potential Practical Implications

Public Administration

- Digital sovereignty
- Support for regulatory compliance
- Transparent decision-making

Industry

- Potential cost reduction
- Potential energy-efficiency gains

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org | NOT PEER-REVIEWED | Posted: 20 November 2026 https://doi.org/10.20944/preprints202611.2231.v1

- Enhanced resilience

Telecommunications and 6G Operators

- AI-native network operations
- Distributed cognitive orchestration
- Sovereign network intelligence
- Explainable autonomous management
- Semantic service provisioning

Healthcare and Life Sciences

- Trusted clinical AI
- Human-AI collaborative decision support
- Privacy-preserving intelligence
- Knowledge-driven medical innovation
- Governance-aware pharmaceutical ecosystems

4.5. Design Principles

1. Place intelligence close to data when advantageous.
2. Knowledge-based adaptation may complement retraining for specific classes of system change.
3. Explainability must be architectural.
4. Knowledge Graphs can support continuous knowledge evolution.
5. Sustainability is a first-class KPI.

4.6. Limitations and Open Challenges

Several challenges remain before large-scale deployment, including ontology synchronization, distributed consistency management, governance scalability, and resource constraints across heterogeneous edge environments.

As with any emerging paradigm, the transition from conceptual architecture to large-scale deployment introduces a number of open challenges. These challenges define the research agenda required to transform distributed cognition into an operational reality:

- distributed ontology synchronization;
- consistency of distributed Knowledge Graphs across nodes;
- resource constraints at the edge;
- governance scalability;
- cybersecurity of distributed reasoning;
- orchestration scalability and consistency management trade-offs between distributed knowledge locality and synchronization overhead.

5. Conclusion

SOFIA introduces Knowledge Gravity as a conceptual principle for the placement, evolution, and governance of knowledge assets across distributed cognitive infrastructures. By extending the logic of Data Gravity from data locality to cognitive locality, the framework reframes distributed intelligence as a problem of positioning knowledge, reasoning capabilities, policies, and governance mechanisms where decisions are generated and acted upon.

The proposed framework integrates Edge-Cloud Continuum principles, Neuro-Symbolic Intelligence, Continuous Knowledge Evolution, Knowledge Graphs, and Governance-by-Design into a unified architectural perspective. Through this integration, SOFIA shifts the focus from the orchestration of computational resources toward the orchestration of knowledge and decision-making capabilities, introducing the notion of a Cognitive Continuum as a natural evolution of existing distributed infrastructures.

Beyond proposing a new architectural vision this work formulates Knowledge Gravity as a unified optimization perspective for distributed cognitive systems. Within this perspective, decision

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org | NOT PEER-REVIEWED | Posted: 20 November 2026 https://doi.org/10.20944/preprints202611.2231.v1

utility, semantic locality, governance requirements, explainability, synchronization costs, and sovereignty constraints are considered jointly when reasoning about knowledge placement and orchestration. In this sense, Knowledge Gravity can be interpreted as an integrative extension of several existing research streams rather than as a complete replacement for current placement paradigms.

Ultimately, SOFIA suggests that the future of sovereign and trustworthy AI may depend not only on where data and computation reside, but also on how knowledge, policies, and reasoning capabilities are distributed, coordinated, and governed across heterogeneous environments. Knowledge Gravity is therefore proposed as a foundational principle for enabling explainable, governable, and sovereign intelligence within emerging distributed cognitive ecosystems.

Research Agenda and Future Directions

Open research Questions

The Knowledge Gravity framework opens several research questions:

1. Formal measurement of knowledge attraction forces.
2. Benchmarking distributed versus centralized reasoning.
3. Correlation between KPQI and explainability metrics.
4. Impact on energy consumption and carbon footprint.
5. Governance implications for sovereign AI infrastructures.
6. Validation across industrial and public-sector use cases.

Planned Research Directions

Future research on SOFIA will focus on four complementary research streams:

Knowledge Gravity Formalization Development of mathematical models, quality placement metrics (KPQI), and optimization algorithms for knowledge placement and migration across the Cognitive Continuum.

Neuro-Symbolic Edge Computing Investigation of distributed neuro-symbolic reasoning mechanisms capable of combining local inference, symbolic constraints, and adaptive knowledge evolution.

Governance-by-Design Definition of policy-aware reasoning frameworks integrating compliance, auditability, trust, and sovereignty requirements directly into cognitive processes.

Experimental Validation Large-scale experimentation across industrial and public-sector use cases to evaluate performance, explainability, sustainability, and governance effectiveness.

An important direction for future research is the establishment of evaluation metrics for both Knowledge Gravity and Neuro-Symbolic Governance-by-Design mechanisms. While explainability-oriented metrics such as explanation completeness, fidelity, reasoning-trace coverage, auditability, and human comprehensibility can assess the transparency and trustworthiness of reasoning processes, Knowledge Gravity requires dedicated indicators, including knowledge-placement latency, knowledge-transfer volume, synchronization overhead, consistency cost, decision-locality ratio, energy per decision, and the proportion of locally satisfiable decisions. Such metrics would enable systematic empirical validation of the architectural and operational implications of knowledge-centered intelligence distribution across the Edge-Cloud Continuum.

Author Contributions: Conceptualization, A.S.; methodology, A.S.; investigation, A.S.; formal analysis, A.S., P.L., M.M.; writing—original draft preparation, A.S.; writing—review and editing, P.L., M.M., A.S.; supervision, A.S.; visualization M.M., P.L., A.S.; project administration A.S. All authors have read and agreed to the published version of the manuscript.

Funding: This research was supported by the Villanova project, financed by IPCEI-CIS, Prog. n. SA. 102519 – CUP B29J24000850005.

Data Availability Statement: All relevant data are included in the paper. Further inquiries can be directed to the corresponding author.

Acknowledgments: The authors used ChatGPT (OpenAI) solely for grammatical correction and proofreading during the preparation of this manuscript. All suggested revisions were critically reviewed and verified by the

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org | NOT PEER-REVIEWED | Posted: 20 November 2026 https://doi.org/10.20944/preprints202611.2231.v1

authors. The authors retain full responsibility for the originality, accuracy, integrity, and content of this publication. The opinions expressed are solely those of the authors and do not in any way represent the official position of Company Villanova.ai.

Conflicts of Interest: The authors declare no conflict of interest. .

References

1. European Union. Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). Off J Eur Union. 2024; L2024/1689. http://data.europa.eu/eli/reg/2024/1689/oj
2. European Union. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation). Off J Eur Union. 2016; L119:1-88. http://data.europa.eu/eli/reg/2016/679/oj
3. European Union. Directive (EU) 2022/2555 of the European Parliament and of the Council of 14 December 2022 on measures for a high common level of cybersecurity across the Union (NIS2 Directive). Off J Eur Union. 2022; L333:80-152. https://eur-lex.europa.eu/eli/dir/2022/2555/oj
4. European Union. Regulation (EU) 2023/2854 of the European Parliament and of the Council of 13 December 2023 on harmonised rules on fair access to and use of data (Data Act). Off J Eur Union. 2023;2023/2854. https://eur-lex.europa.eu/eli/reg/2023/2854/oj
5. European Union. Regulation (EU) 2022/868 of the European Parliament and of the Council of 30 May 2022 on European data governance (Data Governance Act). Off J Eur Union. 2022; L152:1-44. https://eur-lex.europa.eu/eli/reg/2022/868/oj
6. Shi W, Cao J, Zhang Q, Li Y, Xu L. Edge computing: Vision and challenges. IEEE Internet Things J. 2016;3(5):637-646. doi:10.1109/JIOT.2016.2579198
7. Satyanarayanan M. The emergence of edge computing. Computer. 2017;50(1):30-39. doi:10.1109/MC.2017.9
8. Varghese B, Wang N, Barbhuiya S, Kilpatrick P, Nikolopoulos D. Challenges and opportunities in edge computing. In: Proceedings of the 2016 IEEE International Conference on Smart Cloud (SmartCloud); 2016:20-26. doi:10.1109/SmartCloud.2016.18
9. Armbrust M, Fox A, Griffith R, et al. A view of cloud computing. Commun ACM. 2010; 53(4):50-58. doi:10.1145/1721654.1721672
10. Masip-Bruin X, Marín-Tordera E, Sánchez-López S, et al. Managing the cloud continuum: Lessons learnt from a real fog-to-cloud deployment. Sensors. 2021;21(9):2974. doi:10.3390/s21092974
11. Cassani A, Ruberl M, Salis A, Giannese G, Boanelli G. zIA: a GenAI-powered local auntie assists tourists in Italy. arXiv. Published July 16, 2024. doi:10.48550/arXiv.2407.11830
12. Wei Z, Zhou Y, Zhao J, Li K, Fu X. Breaking memory wall for fast edge LLM inference using contextual sparsity. IEEE Trans Mobile Comput. 2026. doi:10.1109/TMC.2026.3712257
13. Zhou Z, Chen X, Li E, Zeng L, Luo K, Zhang J. Edge intelligence: Paving the last mile of artificial intelligence with edge computing. Proc IEEE. 2019;107(8):1738-1762. doi:10.1109/JPROC.2019.2918951
14. Bishop T. How data gravity is forcing a shift to a data-centric enterprise architecture. DATAVERSITY. Published 2020. https://www.dataversity.net/articles/how-data-gravity-is-forcing-a-shift-to-a-data-centric-enterprise-architecture/
15. McCrory D. Beyond data gravity: Design implications for modern data architectures. Dremio. Published 2021. https://www.dremio.com/wiki/data-gravity/
16. Garcez ASD, Lamb LC, Gabbay DM. Neural-Symbolic Cognitive Reasoning. Springer; 2008. doi:10.1007/978-3-540-73246-4
17. Besold TR, d'Avila Garcez AS, Bader S, et al. Neural-symbolic learning and reasoning: A survey and interpretation. arXiv. 2017. doi:10.48550/arXiv.1711.03902
18. De Raedt L, Dumančić S, Manhaeve R, Marra G. From statistical relational AI to neuro-symbolic AI. In: Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence (IJCAI-20); 2020:4943-4950. doi:10.24963/ijcai.2020/688

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org | NOT PEER-REVIEWED | Posted: 20 November 2026 https://doi.org/10.20944/preprints202611.2231.v1

19. DeLong LN, Fernández Mir R, Fleuriot J. Neurosymbolic AI for reasoning over knowledge graphs: A survey. IEEE Trans Neural Netw Learn Syst. 2024;36(5):7822-7842. doi:10.1109/TNNLS.2024.3420218
20. Mezzi E, Burghouts G, Kruithof M. Neurosymbolic inference on foundation models for remote sensing text-to-image retrieval with complex queries. arXiv. Published 2025. doi:arXiv:2512.14102
21. Stein A, Naik A, Velingker N, Naik M, Wong E. The road to generalizable neuro-symbolic learning should be paved with foundation models. NeurIPS Position Paper Track. 2025. doi:10.48550/arXiv.2505.24874
22. Hitzler P, Sarker MK. Neuro-Symbolic Artificial Intelligence: State of the Art and Future Trends. IOS Press; 2022. doi:10.3233/FAIA340
23. 23. d'Avila Garcez AS, Gori M, Lamb LC, Serafini L, Spranger M, Tran SN. Neural-symbolic computing: An effective methodology for principled integration of machine learning and reasoning. arXiv. 2019. doi:10.48550/arXiv.1905.06088
24. Pan S, Luo L, Wang Y, Chen C, Wang J, Wu X. Unifying large language models and knowledge graphs: A roadmap. IEEE Trans Knowl Data Eng. 2024;36(7):3580-3599. doi:10.1109/TKDE.2024.3352100
25. Hogan A, Blomqvist E, Cochez M, et al. Knowledge graphs. ACM Comput Surv. 2021;54(4):71:1-71:37. doi:10.1145/3447772
26. Li Y, Song D, Tian Y, Wang H, Zhou C, Zhang S. A framework of knowledge graph-enhanced large language model based on global planning. IEEE Trans Knowl Data Eng. 2026;38(2):736-748. doi:10.1109/TKDE.2025.3639599
27. Ji S, Pan S, Cambria E, Marttinen P, Yu PS. A survey on knowledge graphs: Representation, acquisition and applications. IEEE Trans Neural Netw Learn Syst. 2022;33(2):494-514. doi:10.1109/TNNLS.2021.3070843
28. Noy N, Gao Y, Jain A, et al. Industry-scale knowledge graphs: Lessons and challenges. Commun ACM. 2019;62(8):36-43. doi:10.1145/3331166
29. Schwartz R, Dodge J, Smith NA, Etzioni O. Green AI. Commun ACM. 2020;63(12):54-63. doi:10.1145/3381831
30. Henderson P, Hu J, Romoff J, Brunskill E, Jurafsky D, Pineau J. Towards the systematic reporting of the energy and carbon footprints of machine learning. J Mach Learn Res. 2020;21(248):1-43. http://jmlr.org/papers/v21/20-312.html
31. Perera O, Liu J. Exploring large language models for ontology learning. Issues Inf Syst. 2024; 25(4):299-310. doi:10.48009/4_iis_2024_124
32. Lippolis AS, Saeedizade MJ, Keskisärkkä R, et al. Ontology generation using large language models. arXiv. 2025. https://arxiv.org/abs/2503.05388
33. Arrieta AB, Díaz-Rodríguez N, Del Ser J, et al. Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. Inf Fusion. 2020; 58:82-115. doi:10.1016/j.inffus.2019.12.012
34. Amodei D. The Urgency of Interpretability. Published April 2025. https://darioamodei.com/post/the-urgency-of-interpretability
35. Doshi-Velez F, Kim B. Towards a rigorous science of interpretable machine learning. arXiv. 2017. doi:10.48550/arXiv.1702.08608
36. European Commission High-Level Expert Group on Artificial Intelligence. Ethics Guidelines for Trustworthy AI. Publications Office of the European Union; 2019. doi:10.2759/346720
37. Huq N, Girard D. Building Sovereign AI: Technical Considerations for Foundation and Architecture. TrendAI Research. Published July 20, 2026. https://www.trendaisecurity.com/it/resources-insights/deep-research/building-sovereign-ai-technical-considerations-for-foundation-and-architecture
38. Esposito M, Marchesi L, Tonelli R, Lenarduzzi V. Sovereign-by-design: A reference architecture for AI and blockchain enabled systems. In: Proceedings of the IEEE/ACM International Conference on Software Engineering Companion; 2026. doi:10.1145/3774748.3795874
39. Ahn MS, Yoon YI. AI-Gov: A governance architecture for AI service systems. In: Proceedings of the IEEE/ACIS International Conference on Software Engineering Research, Management and Applications; 2026:641-647. doi:10.1109/SERA69989.2026.11618631
40. Johnson C. Harnessing knowledge gravity: A new era for semantic layer implementation in data management. AtScale. https://www.atscale.com/blog/harnessing-knowledge-gravity/

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org | NOT PEER-REVIEWED | Posted: 20 November 2026 https://doi.org/10.20944/preprints202611.2231.v1

41. European Commission. Investing in Cloud, Edge and the Internet of Things. Digital Strategy Programme. 2024. https://digital-strategy.ec.europa.eu/en/policies/iot-investing
42. EUCloudEdgeIoT Initiative. Building the European Cloud, Edge and IoT Continuum. 2023. https://eu-cloudedgeiot.eu/
43. Salis A, Marguglio A, De Luca G, Razzetti S, Quadrini W, Gusmeroli S. An edge-cloud based reference architecture to support cognitive solutions in process industry. Procedia Comput Sci. 2023; 217:20-30. doi:10.1016/j.procs.2022.12.198
44. IPCEI-CIS Reference Architecture (ICRA) v3.0. 2026. https://www.8ra.com/wp-content/uploads/IPCEI-CIS_Reference-Architecture_3-0.pdf
45. Ha D, Schmidhuber J. World Models. arXiv. 2018. doi:10.48550/arXiv.1803.10122
46. Harth A, Umbrich J, Hogan A, Decker S. YARS2: A federated repository for querying graph structured data from the Web. *The Semantic Web*. 2007; 2007:211-224. Springer, doi:10.1007/978-3-540-76298-0_16
47. Montoya G, Skaf-Molli H, Molli P, Vidal ME. Federated SPARQL Queries Processing with Replicated Fragments. 2015. In: Arenas, M., *et al.* The Semantic Web - ISWC 2015. ISWC 2015. Lecture Notes in Computer Science, vol 9366. Springer, Cham. doi:10.1007/978-3-319-25007-6_3
48. Saleem M, Ngomo ACN, Parreira JX, Deus HF, Hauswirth M. DAW: duplicate-aware federated query processing over the Web of Data. In: *Proceedings of the 12th International Semantic Web Conference (ISWC 2013)*. Springer; 2013:574-590. doi:10.1007/978-3-642-41335-3_36
49. Stuckenschmidt H, Klein M. Structure-based partitioning of large concept hierarchies. In: *Proceedings of the 3rd International Semantic Web Conference (ISWC 2004)*. Springer; 2004:289-303. doi:10.1007/978-3-540-30475-3_21
50. d'Aquin, M., Schlicht, A., Stuckenschmidt, H., & Sabou, R. M. (2007). Ontology Modularization for Knowledge Selection: Experiments and Evaluations. In Roland Wagner, Norman Revell, Günther Pernul (Ed.), Database and Expert Systems Applications. 18th International Conference, DEXA 2007, Regensburg, Germany, September 3-7, 2007. Proceedings (pp. 874 - 883). Springer. doi:10.1007/978-3-540-74469-6_85

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.