Frontiers in Emerging Multidisciplinary Sciences (Open Access)

Volume 03, Issue 08, August 2026
Publish Date: 21-08-2026
DOI: https://doi.org/10.64917/fems/Volume03Issue08-08 PageNo.133-139

Semantic Knowledge Infrastructure for Resilient and Sustainable AI Decision-Making

Dr. Omar Al-Mansour
Department of Artificial Intelligence Institute of Intelligent Computing and Analytics Riyadh, Saudi Arabia

ABSTRACT

Artificial intelligence decision-making increasingly operates in environments characterized by heterogeneous data, changing distributions, incomplete knowledge, and continuously evolving operational conditions. These challenges expose limitations in conventional AI infrastructures that treat representation, adaptation, and decision logic as relatively independent components. This research develops a conceptual framework for a Semantic Knowledge Infrastructure (SKI) designed to support resilient and sustainable AI decision-making through semantic representations, transferable knowledge, domain adaptation, incremental learning, and experience-aware knowledge maintenance. The methodology synthesizes the theoretical contributions of the provided literature on domain adaptation, zero-shot learning, lifelong learning, representation learning, concept drift, and classifier discrepancy. The proposed infrastructure organizes knowledge into interoperable semantic representations and connects them with adaptive learning mechanisms capable of responding to domain shifts and emerging concepts. The analysis indicates that semantic infrastructure can improve decision resilience by reducing dependence on narrowly trained representations, while sustainability can be strengthened through knowledge reuse, selective adaptation, and reduced retraining requirements. The study further identifies a central design trade-off: highly stable representations improve transferability but may constrain responsiveness to rapidly changing environments. Consequently, resilient AI infrastructure should combine stable semantic knowledge with controlled mechanisms for representation updating and uncertainty management. The resulting framework provides a research-oriented foundation for AI systems that must maintain decision quality under distributional change while minimizing unnecessary computational and organizational costs.

KEYWORDS: Semantic AI, Knowledge Infrastructure, Resilient AI, Sustainable AI, Domain Adaptation, Lifelong Learning, Zero-Shot Learning, Concept Drift, Knowledge Representation, Decision Intelligence.

INTRODUCTION

AI decision systems increasingly operate beyond the controlled conditions under which their models were initially developed. A model trained for one domain may encounter different data distributions, previously unseen categories, changing user behavior, or evolving operational conditions after deployment. Domain adaptation research demonstrates that transferring knowledge across domains is therefore not merely a performance optimization problem but a fundamental requirement for maintaining useful decision behavior under distributional variation (Redko et al., 2017; Saito et al., 2018).

The infrastructure supporting such systems must consequently address more than data storage and model execution. It must provide mechanisms for representing knowledge in ways that remain meaningful when the surrounding environment changes. This requirement motivates the concept of semantic knowledge infrastructure: an architectural layer in which concepts, relationships, task descriptions, domain characteristics, learned representations, and adaptation histories are organized so that AI models can reuse and update knowledge rather than repeatedly relearning it from scratch. The emphasis on infrastructure is important because sustainable AI decision-making depends not only on model accuracy but also on the efficiency, adaptability, and maintainability of the overall learning process.

The problem addressed in this study is the fragmentation of mechanisms for transfer learning, domain adaptation, incremental learning, and semantic representation. Existing research has separately investigated optimal transport for domain adaptation, classifier discrepancy, zero-shot learning, task descriptions, robust representations, and incremental concept learning. These approaches provide complementary capabilities, but their combination into a coherent knowledge infrastructure remains insufficiently articulated. For example, optimal transport can characterize relationships between domains, whereas zero-shot approaches enable recognition of unseen classes; lifelong learning can reuse previous task knowledge, while concept-drift mechanisms address changing distributions. A semantic infrastructure can provide a common organizational layer through which these capabilities interact.

This study therefore develops a conceptual SKI framework for resilient and sustainable AI decision-making. Its objectives are to establish a theoretical foundation from the supplied literature, define the functional components of semantic knowledge infrastructure, analyze how these components interact under domain and concept changes, and identify implications and limitations for future AI systems. The framework is particularly relevant to applications in which decisions must remain reliable while data distributions and task requirements evolve.

The proposed perspective is also consistent with the broader argument that semantic AI infrastructure can provide a foundation for sustainable decision intelligence by connecting reusable knowledge with adaptive computational processes (Goyal, 2025). Rather than viewing sustainability solely as a computational-resource problem, this research treats it as a property of the complete learning lifecycle: knowledge should be reusable, adaptation should be selective, and obsolete information should be managed without unnecessarily discarding valuable prior representations.

LITERATURE REVIEW

2.1 Domain Adaptation as a Foundation for Resilient Decision Systems

Domain adaptation provides an essential theoretical basis for resilience because it addresses situations in which training and deployment domains differ. Redko et al. (2017) theoretically analyze domain adaptation through optimal transport, providing a perspective in which relationships between distributions can be characterized through transportation-based distances. For semantic infrastructure, this suggests that domain relationships should be represented explicitly rather than treated as hidden properties of individual models.

Saito et al. (2018) approach unsupervised domain adaptation through maximum classifier discrepancy. Their work demonstrates the importance of identifying disagreement between classifiers as an indicator of domain-specific uncertainty. In an infrastructure-oriented interpretation, classifier disagreement can become a signal for determining when a semantic knowledge base requires adaptation or when a deployed model is operating outside the conditions represented in its existing knowledge.

These studies are complementary. Optimal transport provides a distribution-level perspective, while classifier discrepancy provides a model-behavior perspective. A semantic knowledge infrastructure can integrate both: distributional relationships can describe how domains differ, whereas prediction disagreement can indicate whether those differences materially affect decisions.

2.2 Zero-Shot Learning and Semantic Transfer

Romera-Paredes and Torr (2015) demonstrate an embarrassingly simple approach to zero-shot learning, emphasizing the possibility of predicting classes that were not directly observed during model training. The importance of this work for semantic infrastructure lies in its conceptual separation between learned visual or input representations and semantic descriptions of categories. Such separation enables knowledge associated with one set of observed classes to contribute to decisions involving unseen classes.

Rostami et al. (2022) further investigate zero-shot image classification through coupled dictionary embedding. Their approach reinforces the importance of constructing relationships between representation spaces rather than treating each task as an isolated learning problem. For an SKI, this implies that semantic mappings should be persistent infrastructure assets. A model encountering a previously unseen category should be able to query existing semantic relationships instead of requiring complete retraining.

2.3 Lifelong Learning and Knowledge Reuse

Rostami et al. (2020) investigate the use of task descriptions in lifelong machine learning for improved performance and zero-shot transfer. This contribution is particularly relevant because task descriptions provide metadata about the purpose and structure of learning activities. Within semantic infrastructure, task descriptions can become first-class knowledge objects that allow the system to identify similarities between historical and emerging tasks.

This perspective shifts the focus from model-centric learning toward knowledge-centric learning. Instead of storing only model parameters, the infrastructure can maintain task semantics, learned representations, adaptation histories, and transfer relationships. Such organization can support selective reuse and reduce redundant learning.

2.4 Representation Stability, Robustness, and Continual Adaptation

Rostami et al. (2024a) examine continuous unsupervised domain adaptation using stabilized representations and experience replay. The study is relevant to resilience because continuous adaptation requires balancing responsiveness to new information against preservation of previously acquired knowledge. Excessive adaptation may cause instability, whereas insufficient adaptation can result in degraded performance under distributional change.

Rostami et al. (2024b) similarly investigate unsupervised domain adaptation through class-conditional compact representations. The emphasis on compact representations suggests that representation structure can influence how effectively knowledge is transferred across changing environments. In a semantic infrastructure, compactness can support efficient storage and retrieval, although excessive compression may remove information required for distinguishing closely related concepts.

Rostami et al. (2023) examine robust internal representations for domain adaptation in sentiment analysis. Their findings support the broader proposition that internal representations can provide an important mechanism for maintaining transferability when surface-level data characteristics change.

2.5 Incremental Concept Drift

Rostami and Galstyan (2023) investigate cognitively inspired learning of incremental drifting concepts. This work highlights an important distinction between simple domain shift and genuine conceptual evolution. A domain may change because the distribution of existing concepts changes, but concepts themselves may also evolve. An infrastructure designed for resilience must therefore distinguish between changes that can be addressed through adaptation and changes that require modification of the underlying semantic knowledge.

Taken together, the literature establishes four complementary foundations: distributional adaptation, semantic transfer, lifelong knowledge reuse, and concept-drift management. The research gap is not the absence of these mechanisms individually; rather, it is the lack of a unified infrastructure model that organizes them around persistent semantic knowledge. The present study addresses this gap conceptually.

METHODOLOGY

3.1 Research Design

This research adopts a conceptual synthesis methodology. The ten supplied studies are analyzed according to their contribution to domain adaptation, semantic transfer, representation learning, lifelong learning, and concept-drift management. Instead of conducting a new empirical benchmark, the study derives an integrated architecture from recurring mechanisms across the literature.

The proposed Semantic Knowledge Infrastructure consists of five interacting layers: semantic knowledge representation, representation and embedding management, adaptive learning, decision intelligence, and knowledge governance. These layers are not intended as isolated modules. Their value arises from controlled information exchange across the learning lifecycle.

3.2 Semantic Knowledge Representation Layer

The first layer maintains machine-interpretable representations of concepts, tasks, domains, relationships, and historical learning outcomes. Task descriptions are particularly important because they provide contextual information that can support transfer between related learning problems (Rostami et al., 2020).

A semantic object may be conceptualized as:

[

$K_i = (C_i, T_i, D_i, R_i, H_i)$

]

where (C_i) denotes concepts, (T_i) denotes task characteristics, (D_i) denotes domain information, (R_i) represents learned relationships, and (H_i) represents historical adaptation information.

This structure allows the infrastructure to distinguish between a new observation and genuinely new knowledge. For example, if a decision system encounters a new customer segment, the infrastructure can determine whether the segment represents an existing concept under a different distribution or a genuinely emerging concept.

3.3 Representation and Embedding Management

The second layer manages vector and semantic representations. Zero-shot learning demonstrates that semantic relationships can enable prediction beyond directly observed categories (Romera-Paredes & Torr, 2015). Coupled dictionary embedding further illustrates the usefulness of linking representation spaces for zero-shot classification (Rostami et al., 2022).

The SKI therefore maintains mappings among data representations and semantic representations. A conceptual transfer function can be expressed as:

[

Z_s = f(X_s, K)

]

where (X_s) is source information, (K) represents semantic knowledge, and (Z_s) is the resulting semantic representation. When a target environment (X_t) differs from (X_s), the infrastructure does not automatically discard the existing representation. Instead, it evaluates whether the representation remains transferable.

Compact representations can improve efficiency, while stabilized representations can support continuous adaptation (Rostami et al., 2024a; Rostami et al., 2024b). The design implication is that representation compression and stabilization should be governed jointly with semantic relevance.

3.4 Adaptive Learning Layer

The adaptive layer determines when and how knowledge should be transferred. Domain differences can be modeled through distributional relationships, including optimal-transport perspectives (Redko et al., 2017). Model-level disagreement can provide an additional signal for adaptation, as demonstrated by classifier-discrepancy approaches (Saito et al., 2018).

A conceptual adaptation score may be represented as:

[

A = \alpha D_{dist} + \beta D_{cls} + \gamma D_{sem}

]

where (D_{dist}) represents distributional divergence, (D_{cls}) represents classifier disagreement, (D_{sem}) represents semantic mismatch, and (\alpha,\beta,\gamma) are weighting coefficients.

If (A) remains below an adaptation threshold, the infrastructure can reuse existing knowledge. If (A) exceeds the threshold, targeted adaptation can be initiated. This mechanism supports sustainability because it avoids retraining whenever environmental change is insufficient to justify it.

3.5 Incremental Learning and Experience Management

Continuous environments require mechanisms for retaining useful historical knowledge. Experience replay and stabilized representations provide a basis for maintaining previous learning while adapting to new distributions (Rostami et al., 2024a). Incremental concept learning further indicates that knowledge management must account for concepts that drift over time (Rostami & Galstyan, 2023).

The infrastructure can therefore maintain an experience repository containing representative historical samples, semantic summaries, task descriptions, and adaptation outcomes. When new evidence conflicts with existing knowledge, the system compares current observations with historical semantic states before deciding whether to update, extend, or replace an existing concept.

3.6 Decision Intelligence Layer

The final operational layer transforms semantic and adaptive knowledge into decisions. A decision should not be based exclusively on the latest model output; it should incorporate the stability and transferability of the underlying representation.

A conceptual decision utility can be expressed as:

[

U = \lambda P + \mu R + \nu S - \rho C

]

where (P) represents predictive reliability, (R) represents resilience to distributional change, (S) represents sustainability, and (C) represents computational or adaptation cost.

This formulation emphasizes that the best decision architecture is not necessarily the one producing the highest isolated accuracy. A model that achieves slightly higher accuracy through frequent full retraining may be less sustainable than one achieving comparable accuracy through selective knowledge reuse.

3.7 Governance and Knowledge Lifecycle

Semantic infrastructure also requires mechanisms for knowledge validation, versioning, retirement, and provenance. Without governance, semantic knowledge may accumulate outdated relationships and generate inappropriate transfers. Consequently, each knowledge object should retain its validity conditions and adaptation history.

The infrastructure should distinguish between stable knowledge, provisional knowledge, and obsolete knowledge. This is particularly important under concept drift, where a previously valid relationship may become misleading. Sustainable decision intelligence therefore depends on controlled knowledge evolution rather than indefinite knowledge accumulation. This infrastructure perspective aligns with the broader role of semantic AI infrastructure in sustainable decision intelligence (Goyal, 2025).

RESULTS

The conceptual synthesis produces four principal findings. First, resilience is strengthened when adaptation is treated as a knowledge-management problem rather than exclusively as a model-training problem. The reviewed literature demonstrates that domain shifts can be characterized at multiple levels: distributional differences through optimal transport, prediction disagreement through classifier discrepancy, and representation changes through robust and stabilized learning (Redko et al., 2017; Saito et al., 2018; Rostami et al., 2024a). An SKI can unify these signals and use them to determine whether adaptation is necessary.

Second, semantic representations increase the potential for knowledge reuse. Zero-shot learning demonstrates that semantic relationships can support decisions involving classes not directly observed during training (Romera-Paredes & Torr, 2015). Coupled representation mechanisms reinforce the importance of connecting different representation spaces (Rostami et al., 2022). Within an infrastructure architecture, this means that semantic knowledge can function as an intermediary between historical learning and emerging decision requirements.

Third, sustainable AI depends substantially on selective adaptation. Lifelong learning with task descriptions provides a mechanism for identifying transferable knowledge between tasks (Rostami et al., 2020), while experience replay supports the preservation of prior knowledge during continuous adaptation (Rostami et al., 2024a). Rather than repeatedly rebuilding complete models, an SKI can reuse stable knowledge and allocate computation primarily to areas experiencing meaningful change.

Fourth, resilience and adaptability are inherently subject to a stability–plasticity trade-off. Stable representations support transfer but may become insufficient when concepts genuinely evolve. Conversely, highly plastic representations can respond quickly but risk degrading previously acquired knowledge. Research on incremental drifting concepts highlights the importance of recognizing that some environmental changes represent genuine conceptual evolution rather than simple distributional variation (Rostami & Galstyan, 2023).

These findings suggest that semantic knowledge infrastructure should not be understood as a static repository. It is better conceptualized as an adaptive layer connecting semantic memory, representation learning, domain monitoring, and decision generation. Its principal contribution is therefore architectural: it creates a mechanism through which heterogeneous adaptation techniques can operate within a persistent knowledge lifecycle.

The sustainability implication is equally significant. Selective transfer and reuse can reduce redundant computation, while semantic task descriptions can reduce the need to repeatedly discover relationships that have already been learned. However, these benefits depend on accurate knowledge governance. Incorrect semantic associations can propagate errors across multiple tasks, making semantic validation a critical requirement.

DISCUSSION

The findings indicate that semantic knowledge infrastructure provides a useful theoretical bridge between traditionally separate AI capabilities. Domain adaptation research focuses primarily on the mismatch between training and target environments, while zero-shot learning addresses transfer to unseen categories. Lifelong learning focuses on reuse across tasks, and concept-drift research addresses changing knowledge over time. The proposed SKI treats these problems as different manifestations of a common infrastructure challenge: how should knowledge remain usable when the environment changes?

The first theoretical implication concerns the definition of resilience. Conventional resilience is often interpreted as maintaining predictive performance under distributional perturbation. The literature reviewed here suggests a broader interpretation. A resilient system should identify changes, determine their semantic significance, select appropriate knowledge, and adapt without unnecessarily destroying historical capabilities. This interpretation connects domain adaptation with lifelong learning and concept-drift management rather than treating them as unrelated techniques.

The second implication concerns sustainability. Sustainability is not achieved simply by reducing model size or computational complexity. If an AI system repeatedly retrains large models because it cannot reuse previous knowledge, computational efficiency remains structurally weak. Task descriptions and reusable semantic representations provide mechanisms for reducing such redundancy (Rostami et al., 2020). The broader infrastructure argument is consistent with the proposition that semantic AI can support sustainable decision intelligence through structured knowledge and reusable decision capabilities (Goyal, 2025).

A major trade-off concerns semantic stability. Stable representations can improve transfer and reduce adaptation frequency, as indicated by research on stabilized representations (Rostami et al., 2024a). However, stability can become a liability when concepts themselves evolve. Incremental concept-drift learning indicates that genuine conceptual change must eventually be incorporated into the knowledge structure (Rostami & Galstyan, 2023). The infrastructure should therefore distinguish between transient distributional variation and persistent semantic change.

Another trade-off concerns compactness. Compact class-conditional representations may improve efficiency and facilitate adaptation (Rostami et al., 2024b), but excessive compression can eliminate distinctions required for reliable decision-making. The appropriate objective is consequently not maximum compression but task-aware compression that preserves decision-relevant semantics.

There are also limitations. The proposed framework is conceptual rather than empirically validated. The weighting factors in the adaptation and utility formulations require calibration through application-specific experiments. Semantic knowledge can also introduce propagation risk: an incorrect relationship stored as reusable knowledge may influence several downstream tasks. Furthermore, the provided literature does not directly establish a unified semantic infrastructure implementation, meaning that the architectural synthesis proposed here should be treated as a research framework rather than an experimentally verified system.

Future empirical research should evaluate the framework under controlled domain shifts, unseen-class conditions, continuous concept drift, and varying computational budgets. Evaluation should measure not only predictive performance but also adaptation frequency, knowledge reuse, computational expenditure, catastrophic forgetting, and decision stability.

CONCLUSION

This research proposed a Semantic Knowledge Infrastructure for resilient and sustainable AI decision-making by synthesizing research on domain adaptation, zero-shot learning, lifelong learning, robust representation learning, and incremental concept drift. The central argument is that resilient AI requires an infrastructure capable of preserving, interpreting, transferring, and updating knowledge rather than relying exclusively on repeated model retraining.

The proposed architecture integrates semantic knowledge representation, embedding management, adaptive learning, incremental experience management, decision intelligence, and knowledge governance. Domain relationships and classifier disagreement can provide adaptation signals, semantic representations can enable transfer to unfamiliar decision contexts, and lifelong-learning mechanisms can support systematic knowledge reuse. Together, these capabilities provide a conceptual pathway toward decision systems that remain operational under changing environments.

The analysis also establishes that resilience and sustainability cannot be optimized independently. Stable knowledge improves reuse but may become outdated; highly adaptive systems respond to change but may incur greater computational costs and risk forgetting. Sustainable infrastructure must therefore implement selective adaptation governed by semantic relevance, uncertainty, and knowledge lifecycle management.

The principal contribution of this study is the positioning of semantic knowledge as an infrastructure-level asset connecting heterogeneous learning mechanisms. Future work should transform this conceptual architecture into experimentally testable implementations and establish quantitative benchmarks for semantic transfer, adaptation efficiency, resilience, knowledge reuse, and computational sustainability. Such research can further clarify how semantic infrastructure can support long-lived AI systems whose decision capabilities remain useful as data, domains, tasks, and concepts evolve.

REFERENCES

1. Redko, I., Habrard, A., & Sebban, M. (2017). Theoretical analysis of domain adaptation with optimal transport. In Ceci, M., Hollm´en, J., Todorovski, L., Vens, C., & Dˇzeroski,S. (Eds.),Machine Learning and Knowledge Discovery in Databases, pp. 737–753Cham. Springer International Publishing.
2. Romera-Paredes, B., & Torr, P. (2015). An embarrassingly simple approach to zero-shot learning. In International conference on machine learning, pp. 2152–2161. PMLR.
3. Rostami, M., Bose, D., Narayanan, S., & Galstyan, A. (2023). Domain adaptation for sen- timent analysis using robust internal representations. In Findings of the Associationfor Computational Linguistics: EMNLP 2023, pp. 11484–11498.
4. Rostami, M., & Galstyan, A. (2023). Cognitively inspired learning of incremental drift-ing concepts. In Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, pp. 3058–3066.
5. Rostami, M. (2024a). Continuous unsupervised domain adaptation using stabilized repre-sentations and experience replay. Neuro computing,597, 128017.
6. Rostami, M. (2024b). Improving unsupervised domain adaptation through class-conditional compact representations. Neural Computing and Applications,36, 1–18.
7. Rostami, M., Isele, D., & Eaton, E. (2020). Using task descriptions in lifelong machine learn-ing for improved performance and zero-shot transfer. Journal of Artificial IntelligenceResearch,67, 673–704.
8. Rostami, M., Kolouri, S., Eaton, E., & Kim, K. (2019). Sar image classification using few-shot cross-domain transfer learning. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops, pp. 0–0.
9. Rostami, M., Kolouri, S., Murez, Z., Owechko, Y., Eaton, E., & Kim, K. (2022). Zero-shot image classification using coupled dictionary embedding .Machine Learning withApplications,8, 100278.
10. Saito, K., Watanabe, K., Ushiku, Y., & Harada, T. (2018). Maximum classifier discrep-ancy for unsupervised domain adaptation. In Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 3723–3732.
11. K. K. Goyal, "Semantic AI Infrastructure for Sustainable Decision Intelligence," 2025 8th International Conference on Informatics and Computational Sciences (ICICoS), Semarang, Indonesia, 2025, pp. 434-439, doi: 10.1109/ICICoS68590.2025.11329920.