The Oscar Publishing Services
American Journal of Applied Science and Technology

VOLUME Vol.06 Issue 08 2026
PAGE NO. 160-166
DOI10.37547/ajast/Volume06Issue08-12

Semantic AI Architecture for Sustainable and Intelligent Decision-Making

Dr. Rahul Verma

Department of Artificial Intelligence Institute of Emerging Computing Technologies Mumbai, India

Dr. Neha Kapoor

Department of Machine Learning and Intelligent Systems Advanced Technology Research Institute Hyderabad, India

Received: 28 June 2026; **Accepted:** 22 July 2026; **Published:** 21 August 2026

Abstract: The increasing complexity of organizational, environmental, and operational decision environments requires artificial intelligence systems that can integrate heterogeneous information, recognize previously unseen situations, and support decisions under uncertainty. Conventional AI pipelines frequently depend on predefined classes, static datasets, and narrowly specified prediction objectives, limiting their ability to adapt when decision contexts evolve. This paper proposes a conceptual Semantic AI Architecture for Sustainable and Intelligent Decision-Making that combines semantic representation, open-world learning, incremental learning, attention-based representation, self-supervised feature extraction, and information-theoretic reasoning. The architecture is developed through a structured synthesis of the supplied literature and is positioned as a research-oriented framework rather than an empirical benchmark. The analysis demonstrates that sustainability-oriented decision intelligence requires more than predictive accuracy: it requires contextual interpretation, novelty awareness, continuous learning, and mechanisms for integrating multimodal evidence. Open-world recognition provides a foundation for identifying unknown situations, while incremental learning enables adaptation without complete retraining. Transformer-based representations and self-supervised learning can improve the semantic quality of continuously acquired data, while information theory provides a theoretical basis for evaluating information relevance and uncertainty. The proposed architecture therefore connects semantic interpretation with adaptive AI infrastructure and sustainable decision processes. The study identifies architectural advantages, operational trade-offs, and limitations, particularly concerning semantic consistency, computational cost, uncertainty management, and evaluation of genuinely unknown decision states.

Keywords: Semantic AI, Decision Intelligence, Sustainable AI, Open-World Learning, Incremental Learning, Self-Supervised Learning, Transformer Architecture, Novelty Detection, Information Theory, Adaptive AI

1. INTRODUCTION

Modern decision-making increasingly depends on data generated by heterogeneous digital, physical, and organizational environments. The challenge is no longer simply to predict a known outcome but to determine whether the available information represents a familiar situation, an evolving class, or an entirely unknown condition. This distinction is particularly important for sustainable decision-making because environmental, social, and operational conditions continuously change. A decision architecture that treats every new observation as belonging to an existing category may generate confident but inappropriate recommendations.

The theoretical problem can be framed as a transition from closed-world intelligence toward semantic and open-world decision intelligence. Traditional machine-learning systems generally assume that the classes represented during training sufficiently describe the operational environment. Open-world recognition challenges this assumption by explicitly considering unknown categories and unfamiliar observations (Bendale & Boult, 2015). Subsequent research further conceptualizes the unknown as a central problem in intelligent recognition systems rather than as an incidental classification error (Boult et al., 2019). This perspective is highly relevant to sustainable decision systems, where previously unseen combinations of environmental, economic, or operational variables may materially affect decisions.

Another challenge is continuous adaptation. Class-incremental learning investigates how models can incorporate new classes while retaining knowledge of previously learned categories (Belouadah et al., 2021). Such adaptability is necessary when a decision-support system operates over long periods and its underlying environment changes. Similarly, self-supervised learning offers mechanisms for extracting useful representations without requiring exhaustive manual annotation (Dhamija et al., 2021). These developments suggest that a sustainable AI architecture should not be viewed as a static predictive model but as a continuously evolving information-processing system.

The objective of this paper is therefore to develop a conceptual semantic AI architecture that integrates adaptive representation, novelty awareness, multimodal attention, and information-theoretic reasoning into a decision-oriented framework. The analysis also considers how such an architecture can support sustainable intelligence while addressing uncertainty, computational overhead, and the difficulty of evaluating unknown conditions. The conceptual positioning is consistent with the broader argument that semantic AI infrastructure can provide a foundation for sustainable decision intelligence (Goyal, 2025).

2. LITERATURE REVIEW

The supplied literature provides several complementary foundations for the proposed architecture. Beddiar et al. (2020) examine human activity recognition as a broad vision-based recognition problem, demonstrating the importance of extracting meaningful patterns from complex observational data. Although their focus is human activity recognition, the underlying challenge—converting heterogeneous observations into meaningful semantic activity categories—is directly relevant to intelligent decision systems.

Brighi et al. (2021) extend this perspective by considering semi-supervised discovery of unknown activity classes. Their ActivityExplorer approach is particularly relevant because sustainable decision environments cannot always provide complete labels for emerging situations. The ability to discover previously unknown activity categories indicates a transition from purely supervised recognition toward adaptive semantic interpretation.

Open-world recognition provides the second major theoretical foundation. Bendale and Boult (2015) frame recognition in environments where unknown classes may appear after deployment. Boult et al. (2019) subsequently survey learning approaches concerned with the unknown, emphasizing that recognizing unfamiliarity is a distinct requirement from conventional classification. Boult et al. (2021) further propose a unifying framework for theories of novelty. Collectively, these studies support the proposition that an intelligent decision architecture should explicitly represent novelty rather than force all observations into predefined semantic categories.

Continuous adaptation is addressed by Belouadah et al. (2021), whose comprehensive study of class-incremental learning highlights the tension between acquiring new knowledge and preserving previously learned information. This creates an important architectural requirement: semantic AI must update its knowledge without continuously rebuilding the entire system from the beginning. For sustainable decision intelligence, such an approach can reduce unnecessary retraining and enable more responsive adaptation.

Representation learning constitutes another layer of the proposed architecture. Dosovitskiy et al. (2022) demonstrate the effectiveness of transformer-based image representations, while Bertasius et al. (2021) investigate space-time attention for video understanding. Together, these works establish the relevance of attention mechanisms for extracting relationships across spatial and temporal dimensions. Carreira et al. (2018, 2019) provide additional context concerning large-scale video datasets, illustrating the importance of extensive multimodal or temporal data for representation learning.

Dhamija et al. (2021) demonstrate the value of self-supervised features for open-world learning. This connects representation learning directly with unknown-class recognition and supports the architectural premise that robust semantic features can reduce dependence on manually labeled datasets.

Finally, Cover and Thomas (1991) provide the information-theoretic foundation for reasoning about information, uncertainty, entropy, and communication. Chalmers (2006), through the discussion of strong and weak emergence, provides a conceptual perspective for understanding how higher-level properties can arise from interactions among lower-level components. Together, these references support an architecture in which semantic decisions emerge from interactions among representations, novelty signals, contextual relationships, and information measures.

The principal research gap is therefore architectural integration. The literature provides strong individual foundations for recognition, novelty detection, incremental learning, attention, and representation learning, but these capabilities must be connected into a coherent decision-intelligence pipeline. The proposed architecture addresses this gap conceptually.

3. METHODOLOGY

3.1 Research Design

This study adopts a conceptual research methodology based exclusively on analytical synthesis of the supplied literature. Rather than training and benchmarking a new machine-learning model, the study develops an architectural framework by mapping established capabilities to the requirements of sustainable and intelligent decision-making.

The methodology follows four analytical stages: semantic representation, contextual and novelty analysis, adaptive knowledge integration, and decision intelligence. Each stage addresses a specific weakness of conventional closed-world AI.

3.2 Semantic Data and Representation Layer

The first layer transforms heterogeneous observations into representations that preserve semantic relationships. Inputs may include visual observations, temporal sequences, structured records, sensor measurements, and textual or contextual information. Transformer-based models are particularly relevant because attention mechanisms can model relationships among elements rather than processing every input independently (Dosovitskiy et al., 2022).

For temporal data, the architecture can incorporate space-time attention so that the significance of an observation depends not only on its local characteristics but also on its temporal context. Bertasius et al. (2021) demonstrate why temporal and spatial relationships are important for video understanding. In a sustainability application, for example, a sensor reading indicating unusually high resource consumption may be interpreted differently depending on the preceding operational sequence.

The representation layer should therefore produce a semantic embedding rather than merely a categorical prediction. This embedding becomes the common interface between perception, novelty detection, knowledge updating, and decision support.

3.3 Self-Supervised and Adaptive Learning Layer

Manual annotation is costly and often unable to keep pace with changing environments. Self-supervised learning provides a mechanism for extracting representations from unlabeled data and is particularly relevant to open-world learning (Dhamija et al., 2021). The proposed architecture uses self-supervised representation learning as a preliminary mechanism for converting continuously generated data into semantically useful features.

When new semantic categories emerge, the architecture can employ incremental learning rather than complete model reconstruction. Belouadah et al. (2021) show that class-incremental learning involves balancing the acquisition of new classes against retention of previous knowledge. Accordingly, the proposed system should maintain historical representations, establish compatibility between old and new semantic concepts, and periodically consolidate knowledge.

A hypothetical sustainability-monitoring system illustrates this mechanism. Suppose an industrial facility initially recognizes normal production, maintenance, and abnormal consumption states. A new operational pattern emerges due to changing equipment behavior. Instead of classifying the pattern as one of the existing states, the system first identifies it as potentially novel, collects additional evidence, and subsequently incorporates the new state into its knowledge structure.

3.4 Open-World and Novelty Layer

Novelty awareness is central to the proposed architecture. Open-world recognition explicitly addresses situations where unknown categories occur after deployment (Bendale & Boult, 2015). The system therefore requires a novelty score or equivalent uncertainty signal that determines whether an observation should be treated as known or unfamiliar.

The novelty layer can compare incoming semantic embeddings against known representations while considering contextual information. Boult et al. (2019) and Boult et al. (2021) establish the theoretical importance of distinguishing known and unknown conditions. Brighi et al. (2021) further demonstrate the relevance of discovering unknown activity classes.

A decision engine should not automatically treat novelty as an error. Instead, novelty can function as a decision trigger. For example, an unfamiliar energy-consumption pattern could initiate additional monitoring rather than immediately generate a definitive recommendation. This converts uncertainty from a failure state into an operational signal.

3.5 Semantic Decision Layer

The final layer converts representations, contextual evidence, novelty estimates, and information measures into decision support. Information theory provides a useful conceptual basis because decision systems must determine not only what information is available but also how informative or uncertain that information is (Cover & Thomas, 1991).

A semantic decision function can therefore be conceptualized as:

Decision Quality = f(Semantic Relevance, Contextual Consistency, Novelty, Information Value, Historical Knowledge).

The expression is conceptual rather than an empirically estimated equation. Its purpose is to establish that decision quality should depend on multiple dimensions rather than classification confidence alone.

Chalmers (2006) is also relevant at this stage because higher-level decision properties can be viewed as emerging from interactions among lower-level informational components. Consequently, the architecture should not assume that a single feature or prediction determines the final decision. Instead, meaningful decision intelligence emerges through the integration of representations, context, and adaptive knowledge.

3.6 Sustainability-Oriented Feedback

Sustainability requires continuous evaluation rather than one-time optimization. The architecture therefore includes a feedback mechanism through which decisions generate new observations and these observations update the semantic knowledge base. This creates a closed learning loop:

Data → Semantic Representation → Novelty Analysis → Knowledge Update → Decision → Feedback → New Data.

The loop reduces dependence on static assumptions and provides an architectural basis for long-term adaptive decision intelligence. The broader infrastructure-oriented perspective of sustainable decision intelligence is consistent with Goyal (2025), which emphasizes the importance of AI infrastructure in supporting sustainable intelligent decision processes.

4. RESULTS

The conceptual synthesis produces four principal findings. First, semantic decision intelligence requires an explicit distinction between recognition and interpretation. The reviewed literature indicates that conventional recognition is insufficient when unknown categories emerge. Open-world recognition provides a mechanism for identifying unfamiliar conditions, while semi-supervised discovery demonstrates how emerging categories can subsequently become meaningful semantic concepts (Bendale & Boult, 2015; Brighi et al., 2021).

Second, the analysis indicates that sustainable AI architectures require continuous knowledge adaptation. Class-incremental learning provides a suitable conceptual mechanism because new knowledge can be introduced without assuming that the operational environment remains unchanged (Belouadah et al., 2021). However, adaptation creates a fundamental trade-off between plasticity and knowledge preservation. A system that learns too aggressively may destabilize previous knowledge, whereas a system that prioritizes stability may fail to recognize emerging conditions.

Third, semantic representation quality is a critical architectural dependency. Transformer-based attention mechanisms provide powerful methods for modeling spatial, temporal, and relational dependencies (Dosovitskiy et al., 2022; Bertasius et al., 2021). Self-supervised features further reduce dependence on manually labeled data and provide a connection between representation learning and open-world recognition (Dhamija et al., 2021). The finding is therefore that novelty detection should not be treated as an isolated component; it depends substantially on the quality and adaptability of the underlying representation.

Fourth, information-theoretic reasoning strengthens the decision layer by providing a conceptual framework for distinguishing informative observations from uncertain or redundant information (Cover & Thomas, 1991). This is especially relevant to sustainability because decision systems often operate under incomplete information. Rather than maximizing prediction confidence alone, the proposed architecture can prioritize observations that materially reduce uncertainty or provide meaningful contextual information.

The overall finding is that sustainable decision intelligence is best understood as an adaptive semantic architecture, not as a single predictive algorithm. Its effectiveness depends on coordinated interaction among representation learning, novelty recognition, incremental learning, contextual attention, and decision feedback. This architectural interpretation aligns with the infrastructure perspective of semantic AI for sustainable decision intelligence described by Goyal (2025).

5. DISCUSSION

The findings have significant theoretical implications. Existing recognition literature demonstrates that intelligent systems must increasingly account for unknown classes rather than assuming that deployment conditions perfectly match training distributions. The proposed architecture extends this principle from recognition toward decision-making. In this interpretation, novelty becomes a semantic event capable of influencing downstream decisions rather than simply an anomaly score.

The architecture also reconciles two apparently competing requirements: adaptation and stability. Incremental learning provides a mechanism for acquiring new concepts, but the literature shows that continual learning is accompanied by challenges associated with preserving prior knowledge (Belouadah et al., 2021). Consequently, sustainable AI cannot be defined solely by its capacity to learn continuously. Sustainable intelligence requires controlled adaptation in which new knowledge is incorporated while historically useful representations remain available.

A second implication concerns the relationship between attention and semantics. Transformer-based models can identify relationships across elements of complex data, while space-time attention is particularly useful for temporally evolving observations (Bertasius et al., 2021). Nevertheless, attention alone does not guarantee semantic understanding. The proposed architecture therefore places attention within a broader pipeline in which representations are evaluated for novelty, contextual consistency, and decision relevance.

There are also practical trade-offs. Large-scale transformer representations can increase computational requirements, while continuous incremental learning introduces memory and model-management challenges. Open-world recognition may additionally produce false novelty signals when data are noisy or poorly represented. Conversely, excessive tolerance toward unfamiliar observations can cause genuine emerging conditions to be overlooked. The architecture must therefore balance sensitivity to novelty against operational stability.

The sustainability dimension introduces another limitation: the architecture itself consumes computational resources. A theoretically adaptive system may become operationally inefficient if representation generation, retraining, storage, and inference are unnecessarily expensive. Sustainable AI should consequently evaluate not only decision quality but also the resource cost required to produce that quality. This infrastructure perspective is consistent with the broader framing of semantic AI as a foundation for sustainable decision intelligence (Goyal, 2025).

Finally, empirical validation remains necessary. The present study provides an integrated conceptual architecture but does not establish numerical performance improvements. Future experiments should compare closed-world and open-world configurations, evaluate incremental-learning stability, measure novelty-detection precision, and quantify computational cost. Such evaluation would determine whether the conceptual benefits translate into measurable improvements in real decision environments.

6. CONCLUSION

This paper developed a conceptual Semantic AI Architecture for Sustainable and Intelligent Decision-Making by synthesizing the supplied literature on open-world recognition, incremental learning, self-supervised representation learning, attention mechanisms, information theory, and emergence. The central contribution is an architectural perspective in which semantic decision intelligence is treated as a continuously adapting system rather than a static predictive model.

The proposed framework integrates heterogeneous data into semantic representations, evaluates observations for novelty, incorporates emerging knowledge incrementally, and converts contextual information into decision support. Open-world recognition provides the foundation for handling unfamiliar conditions, incremental learning supports long-term adaptation, transformer and attention mechanisms strengthen representation quality, and self-supervised learning reduces dependence on exhaustive annotation. Information-theoretic reasoning provides an additional basis for evaluating information value and uncertainty.

The principal limitation is that the framework remains conceptual and requires empirical validation across representative decision environments. Future research should develop benchmark datasets for semantic open-world decision-making, establish quantitative measures of sustainability-aware decision quality, investigate resource-efficient continual learning, and evaluate the robustness of novelty detection under distributional change. Such work could advance the transition from isolated AI models toward adaptive semantic infrastructure capable of supporting reliable, sustainable, and intelligent decisions.

REFERENCES

1. Beddiar, D. R., Nini, B., Sabokrou, M., & Hadid, A. (2020). Vision-based human activity recogni-tion: a survey.Multimed Tools Appl,79(41), 30509–30555.
2. Belouadah, E., Popescu, A., & Kanellos, I. (2021). A comprehensive study of class incrementallearning algorithms for visual tasks.Neural Networks,135, 38–54.
3. Bendale, A., & Boult, T. (2015). Towards Open World Recognition. InConference on ComputerVision and Pattern Recognition, pp. 1893–1902.
4. Bertasius, G., Wang, H., & Torresani, L. (2021). Is Space-Time Attention All You Need for VideoUnderstanding?. InProceedings of the 38th International Conference on Machine Learning,pp. 813–824. PMLR. ISSN: 2640-3498.
5. Boult, T. E., Cruz, S., Dhamija, A. R., Gunther, M., Henrydoss, J., & Scheirer, W. J. (2019). Learn-ing and the Unknown: Surveying Steps toward Open World Recognition.Proceedings of theAAAI Conference on Artificial Intelligence,33(01), 9801–9807. Number: 01.
6. Boult, T., Grabowicz, P., Prijatelj, D., Stern, R., Holder, L., Alspector, J., Jafarzadeh, M. M., Ahmad,T., Dhamija, A., Li, C., Cruz, S., Shrivastava, A., Vondrick, C., & Walter, S. (2021). Towardsa Unifying Framework for Formal Theories of Novelty.Proceedings of the AAAI Conferenceon Artificial Intelligence,35(17), 15047–15052. Number: 17.
7. Brighi, M., Franco, A., & Maio, D. (2021). ActivityExplorer: A semi-supervised approach to dis-cover unknown activity classes in HAR systems.Pattern Recognition Letters,151, 340–347.
8. Carreira, J., Noland, E., Banki-Horvath, A., Hillier, C., & Zisserman, A. (2018). A Short Note aboutKinetics-600.arXiv:1808.01340 [cs],””. arXiv: 1808.01340.
9. Carreira, J., Noland, E., Hillier, C., & Zisserman, A. (2019). A Short Note on the Kinetics-700Human Action Dataset.arXiv:1907.06987 [cs],””. arXiv: 1907.06987.
10. Chalmers, D. J. (2006). Strong and weak emergence. InThe Re-Emergence of Emergence: TheEmergentist Hypothesis from Science to Religion, pp. 244–415. OUP Oxford.
11. Cover, T. M., & Thomas, J. A. (1991).Elements of Information Theory. Wiley Series in Telecom-munications. John Wiley & Sons, Inc., New York, USA.
12. Dhamija, A. R., Ahmad, T., Schwan, J., Jafarzadeh, M., Li, C., & Boult, T. E. (2021). Self-Supervised

Features Improve Open-World Learning.arXiv:2102.07848 [cs],””. arXiv:2102.07848.

13. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani,M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., & Houlsby, N. (2022). An Image isWorth 16x16 Words: Transformers for Image Recognition at Scale. InInternational Confer-ence on Learning Representations.

14. 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.