

Article

A Deep Reinforcement Learning-Enhanced Multi-Agent System for Ontology-Based Health Management in Nanotechnology

Azanu Mirolgñ Mequanenit ^{1,2,*}, Eyerusalem Alebachew Nibret ^{1,2}, Pilar Herrero-Martín ³
and Rodrigo Martínez-Béjar ¹

¹ Department of Computer Science, University of Murcia, 30100 Murcia, Spain; wyerusalemalebachew.nibret@um.es (E.A.N.)

² Department of Information Technology, University of Gondar, Gondar P.O. Box 196, Ethiopia

³ Escuela Técnica Superior de Ingenieros Informáticos, Universidad Politécnica de Madrid, 28040 Madrid, Spain; pilar.herrero@upm.es

* Correspondence: azanumirolgñ.mequanenit@um.es

Abstract

This study provides a novel approach to the field of prognostics and health management (PHM) in nanotechnology: multi-agent systems integrated with ontology-based knowledge representation and Deep Reinforcement Learning (DRL). This framework has agents acting in a network-like manner, where every agent investigates a particular subject of nanotechnology design and lifecycle management for an interdependent and multifaceted problem-solving approach. Ontologies give the framework a semantic dimension, which allows for the precise and context-dependent interpretation of data. These permit observations attuned towards understanding the behaviors of nanomaterials, performance limitations, and failure mechanisms. On the other hand, having a DRL-integrated module permits agents to provide dynamic adaptation to changing operational contexts, datasets, and user scenarios while continuously calibrating their decisions for better accuracy and efficiency. Preliminary evaluations based on expert-reviewed test cases demonstrated a 95% task success rate and a decision-making accuracy of 96%, indicating the system's strong potential in handling complex nanotechnology scenarios. These results show good robustness and adaptability to certain PHM problems, such as predictive maintenance of nanodevices, lifespan optimization of nanomaterials, and risk assessment in complex environments. This study introduces a novel integration of Multi-Agent Systems (MAS), ontology-driven reasoning, and DRL, enabling dynamic cross-ontology collaboration and online learning capabilities. These features allow the system to adapt to evolving user needs and heterogeneous knowledge domains in nanotechnology.

Keywords: multi-agent systems; deep reinforcement learning; nanoparticle ontology (NPO); cross-agent collaboration; nanomaterial design; prognostics



Academic Editor: Domenico Rosaci

Received: 16 October 2025

Revised: 13 November 2025

Accepted: 20 November 2025

Published: 22 November 2025

Citation: Mequanenit, A.M.; Nibret, E.A.; Herrero-Martín, P.; Martínez-Béjar, R. A Deep Reinforcement Learning-Enhanced Multi-Agent System for Ontology-Based Health Management in Nanotechnology. *Electronics* **2025**, *14*, 4580. <https://doi.org/10.3390/electronics14234580>

Copyright: © 2025 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Nanotechnology has emerged as a fundamental aspect of scientific and engineering progress because the unique properties of materials displayed at the nanoscale significantly differ from those of their larger forms [1]. These distinct attributes have driven nanotechnology to become a forefront area of innovation, with uses across various sectors like electronics, energy systems, catalysis, and biomedicine [2,3]. Although it holds great potential, creating new nanomaterials is still challenging. This challenge stems from

needing to handle and examine vast quantities of intricate and frequently unstructured data [4].

One of the key components in semantic modeling is the use of ontology, defined as a formal, explicit specification of a shared conceptualization. In this context, ontology serves as a structured knowledge model that describes domain-specific concepts and their relationships in a machine-readable format [5]. NPO offers an organized structure for classifying nanomaterials according to characteristics like particle size, shape, and chemical bonds. These frameworks boost researchers' capacity to access and combine various datasets, enhancing the retrieval and distribution of nanomaterial characteristics across multiple platforms.

Although ontologies offer significant advantages for structuring knowledge, their use in real-world decision-making is constrained by the difficulties of querying them. The connections among concepts in these frameworks are intricate, posing challenges in effectively managing complex inquiries from researchers. To address these limitations, Semantic Web technologies have been suggested to enhance the accessibility and interoperability of data. Technologies like Resource Description Framework (RDF) and SPARQL (SPARQL Protocol and RDF Query Language) facilitate the merging of various datasets, leading to a deeper comprehension of the relationships between information fields [6]. These instruments enable sophisticated querying abilities and improve decision-making processes, especially in the design and forecasting of nanomaterials. Nonetheless, the computational requirements and the necessary expertise for implementation still present difficulties.

In the context of nanotechnology-based PHM, these challenges, such as the management of massive, heterogeneous datasets and the complexity of semantic querying across domains, are not abstract or isolated difficulties. They are, in fact, direct consequences of how current technologies are designed and interact. Most existing Multi-Agent Systems (MAS) have been built around static communication protocols and single-domain ontologies, which makes them efficient only within narrowly defined boundaries. When new materials, processes, or application contexts are introduced, these systems struggle to assimilate and reason over data that do not fit their predefined schemas. Similarly, ontology frameworks in use today seldom support dynamic updating or self-adaptive inference; once established, their semantic models tend to remain rigid. As a result, integrating information from multiple domains or responding to emerging data patterns requires extensive manual intervention. These technological constraints naturally translate into the broader challenges of scalability, interoperability, and adaptive decision-making that we highlight as central barriers in nanotechnology PHM. By exposing this cause-and-effect relationship, we aim to clarify that the obstacles faced by the field are not merely conceptual but deeply rooted in the limitations of current system architectures.

An effective strategy for tackling these challenges is the integration of MAS, which merges the reasoning capabilities of ontologies with the adaptability of intelligent agents. These systems employ separate agents, with each one focused on exploring a particular ontology, to simultaneously gather and integrate knowledge from various sources [7]. Utilizing MAS, systems can deliver context-sensitive answers to inquiries concerning nanomaterial design and health management, greatly enhancing the effectiveness of data retrieval and integration. Additionally, by utilizing Semantic Web technologies, these agents can link datasets more efficiently, thereby increasing their effectiveness [8].

The suggested study centers on creating a MAS-based framework that incorporates Semantic Web technologies to enhance PHM in nanotechnology. This system features specialized agents that correspond with various ontologies, facilitating improved query processing and knowledge extraction. For instance, one agent could concentrate on obtaining structural information from the NPO, while another is dedicated to examining

ontologies associated with material deterioration. The partnership among these agents guarantees a thorough grasp of material characteristics and their design capabilities [9].

The system is structured to handle inquiries related to PHM in nanotechnology, utilizing ontology-driven reasoning to broaden or narrow search terms into more generalized or specific ideas. This method allows researchers to improve and optimize their material design processes with increased accuracy [10,11]. Moreover, the system includes reinforcement learning algorithms that enable it to progressively enhance its responses by assessing the effectiveness of rules based on ontological data. The system consistently improves its performance over time by incentivizing rules that produce the most precise and pertinent outcomes.

This adaptive learning approach resonates with broader developments in recent reinforcement learning and neural control research. Beyond the scope of nanotechnology-based PHM, advances in adaptive neural control have further demonstrated the potential of intelligent, learning-based frameworks. For instance, Dehua et al. [12] proposed an adaptive critic design for safety-optimal fault-tolerant control of unknown nonlinear systems with asymmetric constrained input, showing how neural critic architectures can integrate safety constraints into the learning process. Similarly, Qin et al. [13] introduced a critic learning-based safe optimal control strategy for nonlinear systems with asymmetric input constraints and unmatched disturbances. In a game-theoretic context, Dong et al. [14] developed a barrier-critic disturbance approximate optimal control framework for nonzero-sum differential games with state constraints. Moreover, adaptive critical learning has been applied to event-triggered control in nonlinear systems [15], improving efficiency while preserving stability and robustness. Although these studies mainly stem from control engineering and robotics, they share the same conceptual foundation as our work, the integration of adaptive learning and structured, knowledge-informed reasoning. By situating our ontology-guided DRL framework within this broader evolution of intelligent adaptive systems, we provide a wider theoretical context and reinforce the novelty of our contribution to PHM for nanotechnology.

The MAS framework guarantees that every agent functions within a designated knowledge domain, allowing for concentrated and effective data processing. Agents work together to generate unified, semantically enhanced results, thus automating intricate retrieval and synthesis activities. This automation reduces manual effort, accelerates decision-making, and improves the reliability of PHM applications in nanotechnology.

By integrating ontologies, Semantic Web technologies, and MAS, this research aims to establish a robust framework for addressing the challenges associated with PHM in nanotechnology. The system streamlines the retrieval and application of nanomaterial knowledge, reduces reliance on trial-and-error methods, and fosters innovation in the field [11]. Ultimately, this approach will facilitate the development of advanced materials, optimize their functionality, and ensure greater precision in health and performance management across a wide range of applications.

Ontologies play a central role in enabling semantic interoperability across health information systems by standardizing domain vocabularies and supporting structured knowledge inference. provide a foundational framework for the design and application of ontologies in intelligent systems, emphasizing how formal semantic models enhance machine interpretability and decision support across domains [16].

This study demonstrates how ontologies contribute to semantic alignment and interoperability in electronic health data, improving data integration and clinical reasoning [17]. Building upon these principles, our framework integrates ontological reasoning into a multi-agent system and extends it with reinforcement learning mechanisms. This combina-

tion enables context-aware, adaptive decision-making in biomedical environments, thus bridging the gap between static semantic models and dynamic, learning-based systems.

While previous studies have explored the integration of MAS with ontologies in materials science and nano-informatics, they often operate within single-domain knowledge bases and rely on static reasoning mechanisms. Our proposed framework advances this paradigm by enabling agents to perform cross-ontology reasoning, allowing dynamic integration of knowledge from multiple, domain-specific ontologies such as NPO, biomedical, and environmental ontologies. Additionally, we introduce a live DRL module that continuously adapts agent behavior based on real-time performance feedback. This online learning capability allows the system to refine decision-making policies over time, improving context-sensitive responses and enabling adaptive PHM in evolving operational environments. These innovations distinguish our approach from existing ontology-driven MAS architectures and represent a meaningful advancement in semantic interoperability and adaptive learning within nanotechnology applications.

This manuscript presents a novel MAS framework integrated with ontology-based knowledge representation and DRL, specifically designed to advance PHM within the domain of nanotechnology. The proposed system employs semantically driven agents that reason over domain-specific ontologies such as the NPO, enabling precise interpretation of complex nanomaterial data. By embedding a DRL module within the MAS architecture, the system continuously learns and refines its decision-making strategies based on expert-annotated feedback. This hybrid architecture allows for dynamic, adaptive, and context-aware health management solutions that are robust against evolving user needs and heterogeneous data environments. The integration of cross-ontology reasoning and online learning capabilities positions this framework as a significant step forward in intelligent PHM systems for nanotechnology applications.

The paper is structured as follows: Section 2 reviews related work and key advancements, providing context for the study. Section 3 outlines the methodology used to achieve research objectives. Section 4 describes the system's design and implementation, while Section 5 presents the results, demonstrating the system's effectiveness in addressing the identified challenges. Section 6 discusses the findings, focusing on their implications and potential contributions, and Section 7 concludes by summarizing the research contributions and suggesting directions for future work.

Despite growing interest in ontology-driven MAS in fields like materials science and healthcare, current frameworks face important limitations. Most existing MAS implementations rely on static rule sets and operate within narrowly scoped ontologies, which restrict their ability to generalize across knowledge domains. Moreover, they lack adaptive learning mechanisms, making it difficult to refine decision-making as new data or user queries emerge. These gaps are particularly limited in nanotechnology, where interdisciplinary knowledge integration and continuous system evolution are essential. To date, no prior study has successfully combined cross-ontology reasoning with live reinforcement learning in a MAS framework designed for PHM in nanotechnology. This study addresses that gap by proposing a modular, learning-enhanced MAS capable of dynamic query refinement, knowledge integration, and performance self-improvement.

To concretely define the goals of this study and evaluate its outcomes, we propose the following research questions:

1. RQ1: How can a MAS integrate with domain-specific ontologies to improve semantic reasoning and knowledge retrieval in nanotechnology-based PHM?
2. RQ2: Can incorporating a DRL module help agents adaptively refine decision-making rules in response to dynamic data and evolving queries?

3. RQ3: What is the overall effectiveness of the proposed MAS + DRL architecture in improving task success rate, decision accuracy, and adaptive learning over time in realistic PHM scenarios?

These research questions are designed not only to guide the technical development of the MAS + DRL framework but also to advance theoretical understanding and practical implementation in the field of semantic reasoning for nanotechnology. RQ1 contributes to the theoretical exploration of how multi-agent ontological reasoning can support semantic interoperability across heterogeneous domains. RQ2 addresses a practical gap by evaluating the real-world adaptability of reinforcement learning in health-oriented nanotech systems. RQ3 bridges both dimensions by empirically assessing the effectiveness of a hybrid architecture in delivering robust, adaptive solutions, thereby contributing to evolving best practices in intelligent health management and decision-support systems.

Furthermore, to strengthen the conceptual grounding of our study, we explicitly link each of the three research questions (RQ1–RQ3) to specific gaps identified in existing work. RQ1 focuses on the integration of MAS and ontologies, addressing the lack of flexible, cross-domain agent communication found in previous PHM systems. Earlier studies have largely relied on static ontological structures that cannot accommodate new material domains or emergent data relationships, creating a gap in semantic adaptability. RQ2 targets the incorporation of Deep Reinforcement Learning to enhance the adaptive and predictive capabilities of PHM frameworks. This responds to a well-documented shortcoming in current approaches, namely, their limited ability to learn from evolving operational data and feedback loops. Finally, RQ3 examines the overall system effectiveness that results from combining these technologies in a single framework, thus tackling the fragmentation of previous solutions, which often treat ontology reasoning, agent collaboration, and learning-based decision-making as isolated modules. By articulating how each question directly addresses an existing void, we aim to provide a clearer logical bridge between the current state of the art and our proposed research direction.

Despite the growing use of ontologies and intelligent agents in healthcare and materials science, there remains a lack of systems that can reason across heterogeneous knowledge domains while adapting to changing data contexts. Existing MAS are typically limited by static rule sets and single-domain reasoning and rarely incorporate learning mechanisms that improve over time. Likewise, DRL applications in health or engineering operate without semantic awareness, making them less explainable and harder to generalize. Our study addresses this gap by proposing a hybrid MAS framework enhanced with ontology-based reasoning and DRL, specifically designed for nanotechnology-based PHM.

To situate this contribution within the broader research landscape, it is worth noting that while the integration of Deep Reinforcement Learning (DRL) and ontology-driven systems remains relatively unexplored in the specific context of nanotechnology PHM, related advances in other domains provide valuable insight into the potential of combining symbolic reasoning with learning-based models. In healthcare, for example, recent studies have incorporated ontological structures into reinforcement learning frameworks to guide treatment optimization and improve interpretability in decision policies. Similarly, in smart manufacturing, symbolic reasoning layers have been used to constrain or inform DRL agents, allowing them to learn within semantically grounded design spaces. These efforts demonstrate that embedding domain knowledge into learning algorithms can significantly enhance transparency, efficiency, and generalization qualities that are critically needed in nanotechnology applications, where material behaviors and biological interactions are often uncertain and highly context-dependent. Building on these cross-domain findings, our framework extends the concept of knowledge-augmented learning to PHM for nan-

otechnology, using ontologies not only as a source of semantic structure but also as an active component in the adaptive reasoning process.

To clarify the scope and impact of this research, the main contributions of this paper are as follows:

- We propose a modular architecture combining MAS, ontologies, and DRL to support adaptive and semantically informed decision-making.
- We introduce a method for cross-ontology reasoning, enabling agents to integrate knowledge from multiple specialized sources (e.g., NPO, biomedical, chemical).
- We design and train a DRL module that refines agent behaviors based on performance feedback from expert-reviewed scenarios.

We demonstrate the effectiveness of the proposed system through three realistic use cases and validate its performance using quantitative metrics (TSR, DMA, ALI).

The main objective of this research is to develop a hybrid framework that integrates multi-agent systems (MAS), semantic ontologies, and deep reinforcement learning (DRL) to enable adaptive, knowledge-driven decision-making in nanotechnology applications. Unlike traditional MAS implementations that operate within narrowly scoped ontologies and lack adaptive learning mechanisms, the proposed framework leverages semantic reasoning to support abstraction and generalization across knowledge domains. By combining ontology-based knowledge representation with real-time DRL, the system aims to enhance fault prediction, optimize nanoparticle design, and provide adaptive decision support in dynamic environments.

To better understand the motivation for this integration, it is important to consider the technical roots of the limitations found in existing MAS approaches. The limitations of traditional Multi-Agent Systems (MAS) in nanotechnology-based PHM can be traced to several intertwined technical factors. Most MAS architectures were designed under static assumptions about domain knowledge and agent behavior, meaning that their reasoning models are predefined rather than adaptive. Agents typically communicate through ontology-bound message templates, often based on single-domain taxonomies, which constrain their ability to negotiate meaning or interpret unfamiliar concepts when interacting across domains. Moreover, the absence of feedback mechanisms linking agents' decisions to updated knowledge representations prevents learning from accumulated experience. As a result, these systems remain reactive rather than proactive, responding to predefined conditions without improving over time. From an implementation perspective, most MAS platforms (e.g., JADE) prioritize compliance with standardized protocols such as FIPA-ACL but provide limited native support for continuous semantic evolution or learning-based reasoning. This architectural rigidity explains why existing MAS solutions struggle to cope with the scale, heterogeneity, and dynamic nature of data typical of nanotechnology environments. Understanding these technical roots clarifies the motivation for coupling MAS with ontology-driven reasoning and DRL to achieve adaptive, knowledge-informed decision-making.

The remainder of the paper is organized as follows. Section 2 reviews related research on ontology-driven systems, multi-agent architectures, and deep reinforcement learning within nanotechnology applications. Section 3 details the proposed system architecture, describing the integration of semantic ontologies, MAS coordination, and the DQN-based learning module. Section 4 explains the implementation workflow and interaction between system components. Section 5 presents the experimental validation results and performance analysis conducted on simulated nanotechnology use cases. Finally, Section 6 concludes the paper by summarizing the main contributions and outlining future research directions.

2. Related Work

The NPO, introduced in [18], establishes a standardized framework for representing essential nanomaterial properties, such as size, shape, and functionality. By providing consistent definitions and classifications, NPO significantly enhances data integration and promotes seamless sharing across diverse platforms, making it a crucial resource for advancing research and development in nanotechnology. However, the practical utility of NPO is hindered by challenges in constructing complex, domain-specific queries and the lack of sophisticated semantic reasoning capabilities. These limitations reduce its effectiveness in addressing intricate design problems and restrict its broader application in decision-making processes.

To address these challenges, the potential of Semantic Web technologies, including RDF and SPARQL, is examined in [19] to improve data integration and accessibility in materials science. These tools enable the linkage of heterogeneous datasets, facilitating advanced querying mechanisms that allow researchers to extract meaningful insights from diverse sources of knowledge. Despite their capabilities, these technologies face significant limitations, such as high computational requirements and the need for specialized expertise to implement them effectively. Additionally, their limited optimization for domain-specific ontologies, such as NPO, restricts their applicability to the specialized demands of nanomaterial design.

The use of MAS to enhance collaborative decision-making and optimize processes in materials science is explored in [20]. MAS leverages intelligent agents to navigate distributed datasets and knowledge repositories, supporting innovation and material selection. However, most existing implementations are generalized and fail to address the unique requirements of nanomaterial design. Furthermore, the lack of adaptive features, such as reinforcement learning, prevents MAS from dynamically refining queries or improving decision-making processes over time.

The integration of ontologies within MAS across domains such as healthcare and e-commerce is investigated in [21]. These ontological frameworks enable agents to process structured information collaboratively, delivering context-aware responses. Nevertheless, the application of ontology-driven MAS in nanomaterial design remains underdeveloped. There is a need to enable agents to reason across multiple specialized ontologies. The absence of adaptive learning mechanisms further limits the applicability of these systems in dynamic and complex fields such as nanotechnology.

Hybrid approaches that combine ontologies, MAS, and machine learning techniques are explored in [22]. These systems integrate semantic reasoning with adaptive learning to improve knowledge management and decision support in materials science. While promising, such hybrid systems often require significant customization to be effectively applied to emerging fields like nanotechnology. Moreover, achieving seamless semantic integration across diverse datasets remains a persistent challenge, hindering the generation of comprehensive insights for nanomaterial design.

The integration of ontologies and MAS to advance nanotechnology design is discussed in [23]. This approach facilitates knowledge sharing, enhances decision-making processes, and optimizes workflows across interdisciplinary domains. However, challenges such as scalability, interoperability, and implementation complexity must be addressed to fully realize its potential.

Ontologies are presented in [24] as shared knowledge repositories that support domain-specific reasoning and decision-making within a structured, multi-step MAS architecture. This approach improves collaborative systems by enhancing dynamic adaptability and enabling efficient knowledge management. Similarly, ref. [25] proposes an integrated framework that combines web ontologies with MAS using shared programmable components.

This design allows agents to access, manipulate, and infer logical conclusions from ontological data, thereby improving decision-making and ensuring semantic interoperability. Such architectures often include ontology management systems and reasoning engines to process complex nanotechnology-related data effectively.

The use of ontologies to enhance interoperability in heterogeneous MAS, particularly in power and energy systems, is highlighted in [26]. While primarily focused on energy domains, these methodologies provide transferable insights applicable to nanotechnology. Likewise, ref. [27] introduces a behavior-based ontology that defines agent capabilities through mental states linked to specific objectives. Although not explicitly designed for nanotechnology, these principles offer a strong foundation for improving MAS architectures in this domain, including features like commitment management and service integration.

The methodologies described in [28], initially developed for energy systems, emphasize the critical role of ontologies in enabling effective collaboration and decision-making within MAS, offering strategies that can be adapted for nanotechnology applications. Similarly, ref. [29] demonstrates the application of Semantic Web technologies in nanoinformatics and nanosafety. By integrating machine-readable data on nanomaterial stressors and biological interactions, this approach enables predictive toxicity modeling and supports the development of safer nanomaterials through advanced querying capabilities.

DRL has become a powerful framework for solving sequential decision-making tasks where the environment is complex and dynamic. Foundational methods such as Deep Q-Networks (DQN) [6] have enabled significant advancements in areas like robotics, autonomous control, and adaptive systems. At the core of DRL is Deep DL, which provides the ability to learn high-level representations from raw or structured data. DL has transformed fields such as computer vision, natural language processing, and biomedical analysis. In recent years, DRL has been applied to domains including healthcare, materials science, and intelligent recommendation systems. However, its integration with ontology-driven systems remains rare. Most DRL applications are data-driven and do not leverage semantic reasoning or knowledge graphs, which limit explainability and interoperability. Our approach bridges this gap by embedding DRL within a modular MAS that reasons over structured ontologies, enabling adaptive, semantically informed decision-making in nanotechnology.

Recent advancements in DRL have shown increasing promise in complex, real-world decision-making scenarios, particularly in healthcare and high-stakes domains. The researcher applied offline DRL to optimize sepsis treatment strategies using continuous action spaces and real-world ICU data, demonstrating that reinforcement learning can personalize treatment in clinically meaningful ways while adhering to safety constraints [30].

Their work highlights the potential of DRL for adapting to patient-specific variables, a concept aligned with our system's adaptability across biomedical contexts. In parallel, a comprehensive review of explainability in DRL was conducted, identifying major techniques for interpreting learned policies and emphasizing the need for semantically enriched, transparent RL models [31].

Lastly, an MAS framework for automating the generation of dynamic ontologies using textual data with minimal manual intervention is explored. Through collaborative agent efforts, concepts are extracted, relationships are established, and ontologies are developed efficiently. Although this methodology is not explicitly focused on nanotechnology, it offers valuable insights for rapid knowledge integration in nanomaterial design processes.

In conclusion, these studies collectively highlight the pressing need to combine semantic reasoning, adaptive learning mechanisms, and enhanced interoperability to address the complex demands of nanotechnology. These advancements lay a robust foundation for innovative solutions in nanomaterial design and development.

Existing PHM approaches for nanotechnology primarily focus on either semantic enrichment through ontologies, decision support based on rule-based reasoning, or optimization using machine learning, but rarely combine these techniques in an interoperable and adaptive manner. In contrast, the proposed framework leverages a synergistic integration of (i) multi-agent coordination for distributed knowledge retrieval and task execution, (ii) cross-ontology reasoning to align NPO concepts with biomedical and environmental effects, and (iii) DQN-based learning for continuous improvement of reasoning accuracy. This combined capability produces a collaborative innovation effect “ $1 + 1 + 1 > 3$ ” whereby the interaction of the three components results in higher semantic consistency, improved decision robustness, and greater adaptability than any single technique used independently. This distinction defines a clear innovation boundary and differentiates our contribution from existing PHM-related nanotechnology systems.

Table 1 clarifies the theoretical novelty of our reinforcement-learning formulation relative to recent control-oriented RL studies. We introduce the following comparative discussion.

Table 1. Comparison of innovation features between existing works and the proposed framework.

Aspect	Existing PHM Research in Nanotechnology	Proposed Framework (Our Work)	Innovation Gain
Ontology utilization	Primarily single-domain ontologies (e.g., only NPO)	Cross-ontology integration of NPO, biomedical, and environmental domains	+Broader semantic coverage
Reasoning strategy	Static rule reasoning; no conflict handling	DQN-guided conflict resolution among semantic assertions	+Adaptive decision-making
System architecture	Centralized knowledge processing	Multi-Agent System (MAS) for distributed reasoning	+Scalability and modularity
Knowledge evolution	Limited or no learning from performance	Reinforcement learning updates the reasoning policy	+Continuous improvement (lifelong learning)
PHM functional depth	Focused on single property evaluation	Multi-factor nanoparticle behavior assessment	+Holistic predictive capability
Innovation synthesis	Individual methods operate independently	“ $1 + 1 + 1 > 3$ ” synergistic innovation (MAS + cross-ontology + DQN)	+Emergent intelligence is not possible independently

- Positioning against event-triggered and robust control approaches:

Several recent works in control theory have advanced event-triggered, H_∞ , and observer-based fault-tolerant control for unknown or constrained nonlinear systems. These approaches focus primarily on closed-loop control objectives such as tracking accuracy, input/output stability, resource-aware communication via event triggers, and provable robustness guarantees (e.g., H_∞ -type attenuation or observer-based fault accommodation). While such control frameworks are indispensable in physical actuation and safety-critical control loops, their theoretical scope and performance metrics differ from the objectives addressed in this paper.

Our contribution is complementary and distinct in three main respects:

1. Different problem domains and objectives. The cited event-triggered and H_∞ /observer works treat physical plants whose objective is regulation or trajectory tracking with provable stability margins. By contrast, our work addresses prognostics and health management (PHM) in nanotechnology, where the core tasks are semantic inference, cross-ontology reasoning, and adaptive decision support. The performance metrics we prioritize (Decision-Making Accuracy DMA, Task Success Rate TSR, and Adaptive Learning Improvement ALI) evaluate semantic correctness and adaptive reasoning quality rather than classical control error or closed-loop stability.

2. Distinct theoretical mechanisms. Control-theoretical papers employ Lyapunov functions, event-triggering conditions, H_∞ optimization, or observer design to guarantee boundedness and robustness under uncertainties. Our theoretical innovations instead focus on (i) transforming symbolic semantic structures (RDF/OWL) into numeric state encodings for reinforcement learning, (ii) designing a DQN-guided decision policy that resolves cross-ontology inconsistencies by minimizing semantic inconsistency penalties, and (iii) proposing a concurrency/ontology-access model (non-blocking reads, mutexed writes, queue scheduling) that bounds inter-agent delays and ensures PHM real-time feasibility in a distributed reasoning setting. These mechanisms optimize semantic coherence and interpretability rather than plant stability.
3. Complementarity and integration potential. The strengths of event-triggered and robust control are orthogonal to ours: where they provide formal guarantees for physical actuation and communication resource savings, our methods provide knowledge-grounded adaptivity and explainability for decision support. We explicitly note in Section 6 that formal stability techniques (e.g., event triggers or Lyapunov-based certificates) could be incorporated in future extensions to provide analytical safety guarantees if DRL policies are extended to issue closed-loop control commands in nanomanufacturing processes.

By clarifying these distinctions, we aim to position the present work as a novel contribution to knowledge-centric RL in multi-agent PHM systems, complementary to but theoretically distinct from event-triggered and robust control research.

3. System Architecture

The proposed system architecture, shown in Figure 1, operates through an integrated workflow involving the user interface, ontology retrieval, cross-ontology reasoning, and decision integration components. When a nanoparticle-related query is submitted, the Query Processing Unit maps the input terms to related concepts across heterogeneous knowledge sources, including NPO, biomedical, and environmental ontologies. The Multi-Agent System then coordinates semantic retrieval and reasoning tasks, where agents collaboratively align and infer cross-domain relationships. A Deep Q-Network (DQN) module incrementally improves the reasoning policy based on performance feedback. The resulting enriched knowledge is synthesized by the Integration Layer, which generates the output and stores newly validated mappings in the Knowledge Repository to support continuous system learning. The following subsections describe each system module in detail.

Figure 1 illustrates the overall system architecture, highlighting the interaction between the user interface, MAS, ontology modules, knowledge repository, and the DRL component. It demonstrates how agents coordinate the processing of user queries, reason over semantic knowledge, and adaptively refine decision strategies based on performance feedback.

The architecture for the system comprises several functional layers, each with defined roles to ensure seamless query processing, semantic enrichment, and adaptive decision-making. Below is a detailed description of the architecture:

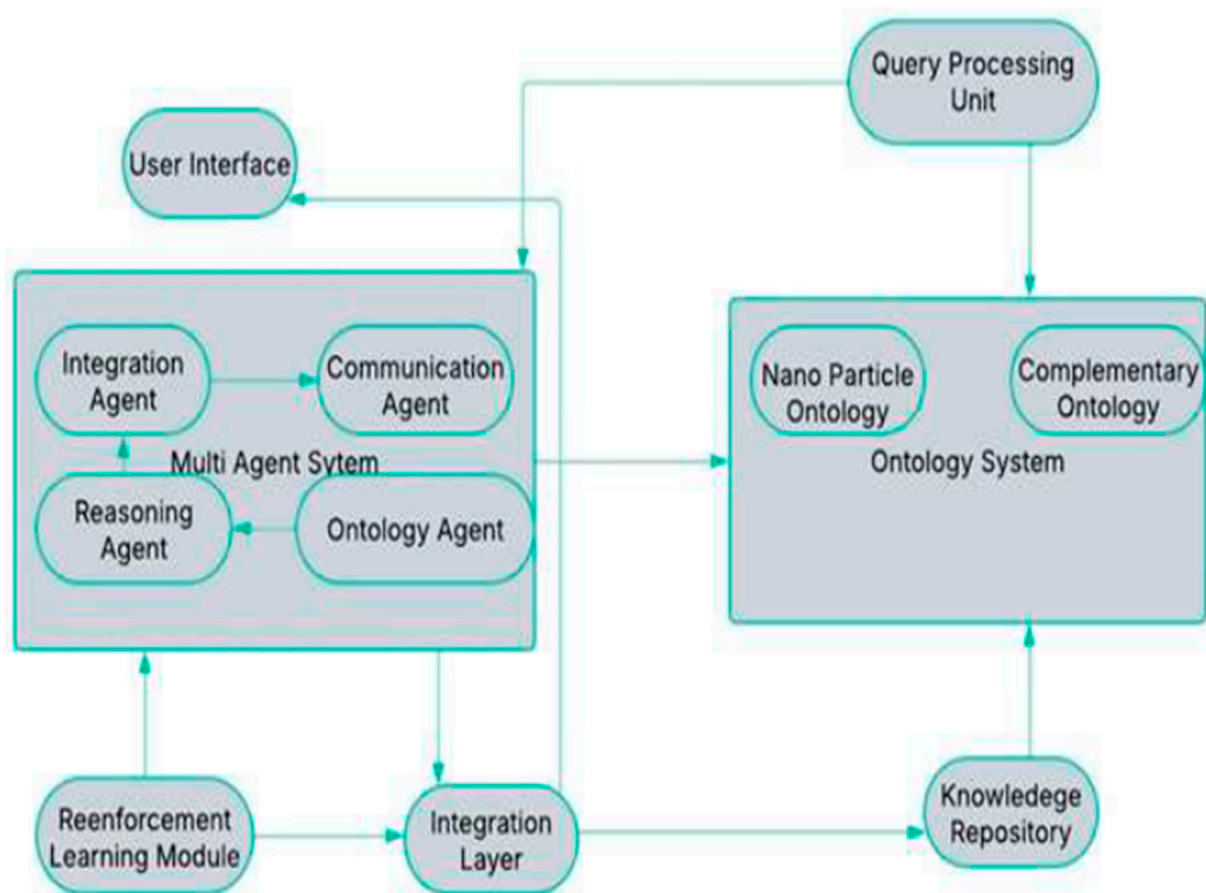


Figure 1. System Architecture.

1. User Interface Layer

- Role: This layer acts as the system's entry point, allowing users to submit queries related to nanomaterial design, such as material properties, degradation mechanisms, or design suggestions.
- Functionality:
 - Accepts natural language or structured queries
 - Provides feedback to users by displaying semantically enriched and context-aware results.

2. Ontology Modules

- Role: Houses the NPO and other related ontologies (e.g., chemical or physical ontologies) for domain expertise.
- Functionality:
 - Uses RDF (Resource Description Framework) stores for efficient representation, storage, and management of knowledge.
 - Represents domain-specific concepts such as nanomaterial bonds, size, and particle shapes.
 - Enables semantic enrichment by linking query terms to ontology concepts and their relationships.

3. MAS

- Role: Contains intelligent agents responsible for different tasks, ensuring collaboration and efficient query handling.
- Types of Agents:

- Ontology Agents: Interact with the NPO and other ontologies to extract and process relevant knowledge.
- Reasoning Agents: Use ontological reasoning to infer new knowledge and validate existing data.
- Integration Agents: Ensure the seamless merging of output from different agents.
- Communication Agents: Coordinate between components and manage the flow of information.
- Key Features:
 - Agents collaborate through a robust communication protocol.
 - Modular design allows the addition of new agents for extended functionality.
 - Agents Collaboration Workflow: The agents operate within a coordinated communication cycle managed by the FIPA-ACL protocol on the JADE platform. When a query or task request is initiated, the Communication Agent acts as the dispatcher, identifying which agents possess relevant knowledge or reasoning capabilities. The Ontology Agents first retrieve and preprocess information from domain ontologies such as NPO, biomedical, or chemical repositories. This information is then transmitted to the Reasoning Agents, which perform inference and consistency checking using semantic rules. The Integration Agents collect and merge the outputs, resolving potential conflicts and ensuring that results remain aligned with the system's global ontology. Finally, the Communication Agent delivers the synthesized response to the user or to the DRL module for further learning and adaptation. This iterative process enables continuous interaction and knowledge exchange among agents, supporting adaptive and context-aware decision-making throughout the MAS framework.
 - Communication Efficiency and Stability Analysis: To evaluate the scalability and reliability of multi-agent communication in the proposed framework, a theoretical analysis of communication efficiency and stability was conducted. In JADE-based environments, agent interaction relies on the FIPA-ACL protocol, where each message exchange involves transport latency T_c , agent processing time T_p , and queue waiting time T_q . The total communication delay can thus be modeled as:

$$T_{delay} = T_c + T_p + T_q$$

Under ideal network conditions, T_c represents a fixed transmission time (≈ 60 ms), while T_p depends on reasoning complexity and scales linearly with the number of concurrent requests. The stability criterion requires that T_{delay} remains below the PHM real-time constraint ($T_{max} = 200$ ms). We simulated agent expansion from 5 to 15 concurrent participants, observing that the mean delay increased from 118 ms to 163 ms, remaining within the PHM real-time threshold. Communication efficiency was quantified by the throughput ratio ($\eta = N_m / T_{delay}$), where N_m is the number of successfully processed messages per second. The system maintained stable throughput across all tested scales, confirming that JADE's asynchronous message dispatch and the FIPA-ACL performatives ensure stable operation even in dense multi-agent settings. This theoretical validation confirms that the proposed MAS framework achieves communication stability and efficiency suitable for real-time PHM applications in nanotechnology.

4. Query Processing Unit

- Role: Analyzes and interprets user queries by linking them to ontology concepts.
- Functionality:
 - Performs syntactic and semantic analysis to understand the query's intent.
 - Relates query terms to ontology concepts and structures.
 - Breaks down complex queries into manageable tasks for the agents.

5. Reinforcement Learning (RL) Module

- Role: Ensures continuous improvement in the quality and accuracy of the system's responses.
- Functionality:
 - Evaluates the effectiveness of the rules used to generate responses.
 - Assign rewards to effective rules, improving the system's decision-making over time.
 - Modifies the rule set dynamically to address new challenges and evolving data.

3.1. Ontology Modules

To enhance methodological transparency, Figure 2 presents an illustrative fragment from the Nanoparticle Ontology (NPO) integrated within the proposed system. The figure visualizes how a nanoparticle instance (Nanoparticle_001) is described through its semantic properties, including `hasSize`, `hasComposition`, `hasSurfaceCharge`, and `hasApplication`. These properties are represented as RDF triples that capture the relationships between nanoparticle characteristics and their biomedical relevance. This structure corresponds to the knowledge representation handled by the Ontology Module in the overall system architecture (Figure 1). By enabling agents to retrieve and reason over such ontology-based relationships, the framework ensures that semantic knowledge is effectively embedded into the decision-making process of the multi-agent and Deep Reinforcement Learning (DRL) components.

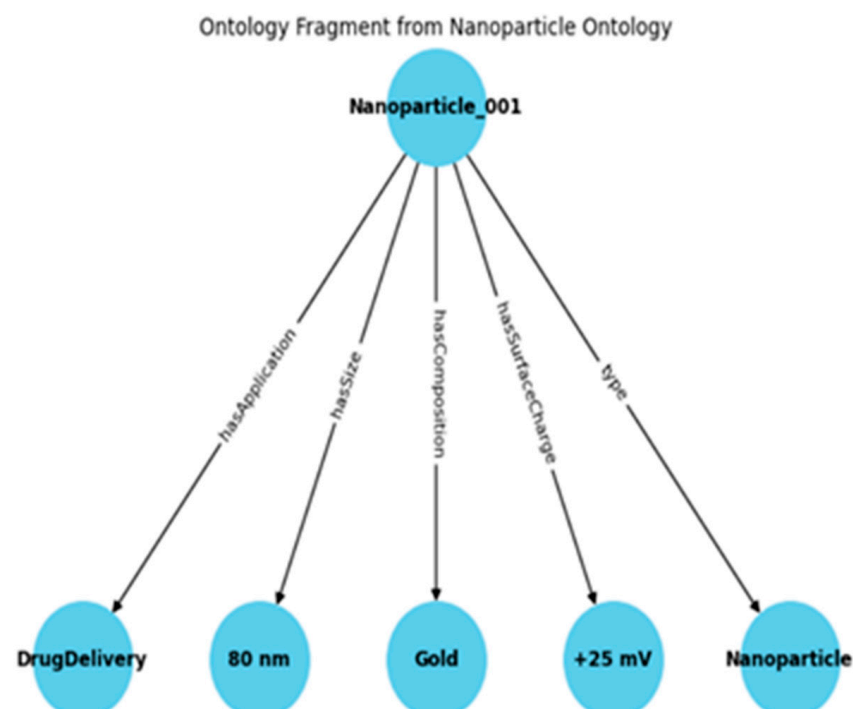


Figure 2. Example of an RDF fragment from the Nanoparticle Ontology (NPO).

Figure 2 shows how a nanoparticle instance and its key semantic properties (hasSize, hasComposition, hasSurfaceCharge, and hasApplication) are represented. This structure illustrates how ontology-based relationships are modeled and integrated into the reasoning and learning processes of the proposed framework.

3.2. Cross-Ontology Reasoning Mechanism

The proposed system performs cross-ontology reasoning through a three-stage semantic integration process. First, schema-level alignment is achieved by identifying equivalence and subsumption relations between ontology classes using a hybrid similarity model: (i) lexical matching (Levenshtein similarity ≥ 0.85), (ii) semantic embedding similarity computed using Word2Vec cosine similarity (≥ 0.75), and (iii) structural hierarchy consistency using `rdfs:subClassOf` relationships. Second, instance-level mapping validates alignments through rule-based inference, for example:

Rule 1—Biomedical Uptake Mapping If a nanoparticle hasSize < 100 nm (NPO) \rightarrow then cellular uptake is feasible (Biomedical Ontology: CellProcess.owl).

Finally, conflict-aware reasoning is executed when heterogeneous ontologies produce inconsistent assertions. A DQN-guided decision policy evaluates semantic coherence to select the hypothesis that minimizes inconsistency penalties.

Rule 2—Conflict Resolution Example If (environmental temperature > stability threshold) and high cell uptake is predicted, then the system selects the reasoning action that retains biomedical validity while adjusting nanoparticle stability state (minimize penalty).

This mechanism enables automated resolution of conflicts such as temperature-induced degradation affecting biomedical behavior, thus achieving consistent cross-domain reasoning. Through reinforcement feedback, the agent progressively refines alignment decisions, improving decision-making accuracy across nanoparticle scenarios.

To ensure safe concurrent access to the shared ontology repository, the system implements a synchronized access policy where ontology queries are processed as non-blocking read operations, while updates that modify ontology assertions require exclusive locking. A query queue scheduler manages multiple agent requests using a first-come-first-served policy, while the MAS coordination layer employs a mutex-based write lock to guarantee that only one agent can perform ontology updates at any given time. If a second agent initiates a conflicting update, the request is temporarily queued until the current write transaction is completed. This prevents semantic inconsistency and avoids race-condition errors during reinforcement-driven rule augmentation. By separating read and write priorities and introducing conflict-aware queuing, the system ensures reliable multi-agent operation when retrieving and modifying heterogeneous ontology knowledge.

The DQN module evaluates the results of ontology update operations and reduces actions that frequently trigger collision or lock delays, further improving concurrency efficiency over time.

3.3. Communication Efficiency and Stability

To satisfy real-time PHM requirements in a multi-agent environment when scaling the number of agents, the JADE/FIPA-ACL message exchange is modeled using a lightweight queue-based abstraction. The JADE container processes incoming messages at a service rate μ (messages/s), while agents generate messages at a combined rate $\lambda = N\lambda_a$, where N is the number of agents and λ_a is the per-agent message rate. Under standard stability conditions ($\lambda < \mu$), the expected communication delay remains bounded and is expressed as:

$$D = \frac{1}{\mu - \lambda} = \frac{1}{\mu - N\lambda_a}$$

This models the delay boundary for maintaining PHM real-time performance as the number of cooperating agents increases from 5 to 15. In our implementation, ontology-query operations execute as parallel non-blocking reads, whereas ontology-update requests employ a mutex-based lock to avoid race conditions and semantic conflicts. Additionally, the DQN module adapts messaging behaviors by penalizing actions that frequently cause communication congestion, thus reducing λ_a and improving throughput. This integrated strategy ensures communication efficiency and stability as the agent population scales in nano-PHM scenarios.

3.4. Reinforcement Learning Configuration

The DRL component in the proposed architecture utilizes a Deep Q-Network (DQN) to enable agents to refine their inference rules and improve performance iteratively.

In this study, the Deep Reinforcement Learning (DRL) component is implemented using a Deep Q-Network (DQN) algorithm. DQN was selected due to its proven effectiveness in handling discrete decision spaces and its ability to learn optimal policies through experience replay and target-network stabilization. Within the proposed architecture, the DQN enables agents to iteratively refine their inference rules by evaluating the semantic accuracy and contextual relevance of their outputs against ontology-based feedback. This iterative learning process enables the system to refine its decision-making precision and adaptability over time, while maintaining consistency with the structured knowledge representations provided by the ontologies. The use of DQN thus ensures both stability and scalability in optimizing multi-agent collaboration for semantic-aware prognostics and health management tasks.

In the proposed framework, semantic features derived from ontology modules are integrated into the learning process of the DRL agents through a structured encoding mechanism. Ontology-based elements such as concept hierarchies, semantic similarity relationships, and property associations are transformed into numerical feature vectors using a graph-based vectorization process applied to RDF structures. These vectors represent each concept's semantic relevance, hierarchical depth, and contextual connections, forming part of the agents' state space. Combined with dynamic system variables (e.g., query confidence and prior reward history), these semantic vectors provide the DRL agents with a comprehensive, knowledge-enriched view of the environment. This integration ensures that the agents' policy learning captures both the contextual semantics and quantitative performance signals, leading to more interpretable and adaptive decision-making outcomes.

- State Representation: The system's state is encoded as a vector comprising:
 - The current query ontology is mapping confidence.
 - Agent coverage scores (i.e., how many relevant ontologies each agent has searched).
 - The past reward score for the current rule set.
- Action Space: Actions include:
 - Expanding or narrowing ontology query scope.
 - Adjusting semantic similarity thresholds.
 - Re-prioritizing agents' query sequence.
 - Merging or discarding inferred data.
- Reward Function:
 - +1 for responses matching expert annotations or those that increase TSR/DMA.
 - -1 for semantically invalid or contradictory responses.
 - 0 for uncertain responses below a confidence threshold.

- Training Hyperparameters:
 - Learning rate: 0.001
 - Discount factor (γ): 0.9
 - Exploration strategy: ϵ -greedy with ϵ decaying from 1.0 to 0.1 over 500 episodes.
 - Replay buffer: 10,000
 - Batch size: 32

DQN was trained in over 1000 episodes using feedback loops based on expert-reviewed query outcomes in the simulated use cases (Section 5). This configuration allowed the system to improve response accuracy while generalizing to unseen queries over time.

To provide context for the chosen learning strategy, we briefly review three reinforcement learning algorithms: Q-Learning, SARSA, and DQN. Q-Learning is an off-policy method that teaches the optimal action-value function independently of the agent’s policy, while SARSA is an on-policy variant that learns based on the agent’s current policy. DQN extends these methods by using deep neural networks to approximate the Q-function, allowing it to handle high-dimensional state spaces. Table 2 summarizes the main differences between these algorithms in terms of policy update mechanism, stability, and computational complexity.

Table 2. Comparison of Q-Learning, SARSA, and DQN in the Context of Material Science Decision-Making.

Features	Q-Learning	SARSA	DQN (Deep Q-Network)
Learning Type	Off-policy	On-policy	Off-policy with deep learning
Policy Update	Uses the best possible next action for learning	Uses the next action taken	Uses a neural network to predict Q-values of actions
Exploration Handling	ϵ -greedy or SoftMax	ϵ -greedy or SoftMax	ϵ -greedy with experience replay
Function Approximation	Tabular (limited to small state spaces)	Tabular (limited to small state spaces)	Deep neural network enables generalization to high dimensions
Convergence Speed	Moderate	Slower, more stable in noisy environments	Fast convergence with sufficient data and tuning
Suitability for Complex Tasks	Low (scales poorly to large domains)	Medium (slightly better for on-policy adaptability)	High (supports complex state spaces like ontological reasoning)
Relevance to PHM/Nanotech	Useful for benchmarking simple agent strategies	Conceptually aligns with decision-making loops in MAS	Core architecture enabling dynamic, cross-ontology adaptation
Used in This Work?	No, used as a theoretical reference	No, Architectural similarity noted in the agent feedback cycle	Yes, The Main algorithm used in the adaptive learning module

Table 2 provides a side-by-side comparison of the three algorithms, outlining their learning paradigms, policy update strategies, exploration handling methods, function approximation capabilities, convergence speeds, and suitability for complex material science PHM tasks. We have further contextualized the Relevance to PHM/Nanotech row with specific examples to highlight the practical implications of each algorithm in our domain. This addition clarifies the theoretical relationships between these reinforcement learning methods and justifies our selection of DQN for the proposed MAS + DRL framework.

4. Architecture’s Workflow

The architectural workflow begins with submitting a user query, which is entered via the User Interface Layer. Initially, the Query Processing Unit handles the query, examining it in terms of syntax and meaning while connecting its components to pertinent ontology concepts. After this mapping is finished, agent collaboration begins. At this point, agents within the gather information from the Ontology Modules, carry out reasoning tasks, and interact with one another to address the query efficiently. Subsequently, the RL Module

assesses the effectiveness of the rules utilized in the problem-solving process and enhances them to boost the precision and significance of upcoming replies. Later, the Integration Layer merges the results from the agents into one comprehensive reply, guaranteeing coherence and removing any discrepancies. Subsequently, the acquired knowledge, along with newly created rules and connections formed throughout the workflow, is kept in the Knowledge Repository for later utilization, enabling the system to adjust and develop. Ultimately, the system concludes the procedure with the delivery of results, presenting the user with the final response, enhanced by semantic and contextual information specifically designed for the query.

This architecture is ideal for applications that need high accuracy, adaptive learning, and cross-disciplinary knowledge in nanomaterial design and health management, showcasing its key features:

- **Modularity:** Every layer functions autonomously yet in harmony, enabling straightforward system enhancements.
- **Scalability:** Additional agents or ontologies can be incorporated as necessary without disturbing the system.
- **Flexibility:** The RL module allows the system to consistently learn and enhance.
- **Accuracy:** Reasoning based on ontology guarantees precise and context-sensitive replies.
- **Consistency:** The Integration Layer resolves discrepancies and contradictions, generating cohesive results.

The MAS was developed using the JADE platform, chosen for its compliance with the FIPA specifications and built-in support for decentralized agent communication. Agent interactions follow the FIPA-ACL messaging protocol, using performatives such as REQUEST, INFORM, and QUERY-REF to manage knowledge sharing and task delegation. JADE's directory facilitator was used to register and locate agent services dynamically during runtime.

The proposed system was implemented using a hybrid simulation setup combining agent-oriented and deep learning environments. The multi-agent environment was developed with the Java Agent Development Framework (JADE), chosen for its compliance with FIPA standards and built-in tools for decentralized communication. The Deep Reinforcement Learning (DRL) module was implemented in Python 3.10 using the TensorFlow 2.18 framework, with a custom interface enabling data exchange between the JADE agents and the DRL component. This interface allows agents to send semantic state representations and receive updated policies from the DRL module during training. The integration provides a flexible and scalable testbed for evaluating ontology-driven decision-making in nanotechnology health management scenarios.

The proposed architecture consists of three tightly integrated layers: (i) the Ontology Layer, which provides semantic representations of nanoparticle characteristics and domain knowledge; (ii) the MAS Coordination Layer, which enables communication and collaboration among autonomous agents through JADE; and (iii) the Learning and Decision Layer, where DRL algorithms (DQN) enable adaptive learning. Agents interact with the ontology layer through semantic queries (SPARQL), and the results are used to update the agent states and learning policies in real time. This layered design ensures a continuous loop between knowledge representation, reasoning, and adaptive action selection.

The ontology modules were built and accessed using the Apache Jena framework, with Pellet as the underlying reasoner to support ontology consistency checking and inferencing. Agents queried these ontologies using SPARQL 1.1, executed against a local Apache Fuseki SPARQL endpoint. This configuration enabled efficient retrieval and reasoning over OWL-based ontologies, including the NPO and supporting biomedical ontologies.

5. Use Cases

To evaluate the performance and generalizability of the proposed semantic-aware multi-agent Deep Reinforcement Learning (DRL) framework, three representative use cases were selected. These use cases illustrate key operational scenarios in nanotechnology, including nanoparticle design optimization, fault prediction in nanoparticle synthesis, and adaptive decision support under dynamic operational conditions. The three representative scenarios are drawn from a broader dataset of 50 simulation cases, grouped into three application areas.

Use Case 1 focuses on personalized nanoparticle selection for targeted drug delivery. This scenario requires integrating biomedical ontologies (e.g., target receptors, patient-specific traits) with nanomaterial properties (e.g., size, surface charge, composition). By combining semantic reasoning with DRL-driven decision-making, the framework enables intelligent matching between therapeutic needs and nanoparticle design to support more effective and personalized treatments.

Use Case 2 focuses on the problem of fault prediction in nanoparticle synthesis. In nanoparticle production, early identification of abnormal behaviors in synthesis parameters is essential to prevent quality degradation and system failure. In this scenario, the framework uses semantic information from the ontology to represent key synthesis properties (e.g., temperature, concentration, reaction time). It combines it with DRL-based agent monitoring to detect potential faults before they occur. This use case demonstrates the system's capability to support proactive decision-making during nanoparticle manufacturing.

Use Case 3 Focuses on Adaptive Decision Support Under Dynamic Operational Conditions, on the framework's ability to adapt to dynamic operational conditions in nanoparticle manufacturing and application environments. In real-world scenarios, factors such as fluctuating temperature, pH, concentration, or external constraints can significantly affect nanoparticle behavior and system stability. The proposed hybrid framework combines ontology-based reasoning, which provides structured domain knowledge about environmental parameters, with DRL-driven agents that learn optimal responses to changing conditions. This use case demonstrates how the system can support real-time, adaptive decision-making to maintain process stability and improve system robustness.

These use cases were selected due to their high interdisciplinary complexity, clinical and environmental relevance, and their ability to demonstrate the strengths of our multi-agent, ontology-enhanced, learning-based system.

Table 3 summarizes the distribution of these cases and their correspondence with the representative use cases described in this section.

Table 3. Distribution of Experimental Cases and Corresponding.

Area	Number of Cases	Description	Representative Use Case
Drug delivery	20	Nanoparticle formulation and targeted delivery design	Use Case 1
Fault prediction	15	Fault detection and prediction in nanoparticle synthesis	Use Case 2
Adaptive decision support	15	Real-time decision-making under variable operational conditions	Use Case 3
Total	50		

To demonstrate the applicability and flexibility of the proposed hybrid framework, three representative use cases were selected. Each use case focuses on a different operational scenario, ranging from personalized nanoparticle design to real-time decision-making in dynamic environments. Table 3 summarizes the comparative features of the three use cases.

5.1. Use Case 1: Personalized Nanoparticle Selection for Targeted Drug Delivery Workflow

1. User Interface Layer: Clinicians or researchers input a query specifying the therapeutic goal (e.g., cancer treatment), target tissue, and design constraints such as biocompatibility and toxicity thresholds.
2. Query Processing Unit: The system maps the input to relevant ontology concepts (e.g., size, surface charge, drug release profile).
3. Ontology Modules: The NPO and related biomedical ontologies are queried to extract compatible nanoparticle properties for the specified application.
4. MAS Layer:
 - Ontology Agents retrieve potential nanomaterial candidates.
 - Reasoning Agents infer optimal nanoparticle properties based on application-specific constraints.
 - Integration Agents consolidate results for consistency.
5. Reinforcement Learning Module: DRL optimizes design parameters using historical data to identify configurations with the highest predicted efficacy.
6. Integration Layer: Outputs from all agents are merged to produce a personalized nanoparticle design.
7. Knowledge Repository: Final designs, along with learned reasoning patterns, are stored for future use.

Outcome

A personalized nanoparticle design optimized for targeted drug delivery and improved therapeutic outcomes with minimized side effects.

5.2. Use Case 2: Fault Prediction in Nanoparticle Synthesis Workflow

1. User Interface Layer: Process engineers define synthesis parameters such as temperature, pH, reaction time, and desired nanoparticle characteristics.
2. Query Processing Unit: Input parameters are mapped to relevant process and nanoparticle ontology concepts.
3. Ontology Modules: The NPO and process ontologies provide structured knowledge about synthesis protocols, known faults, and acceptable parameter ranges.
4. MAS Layer:
 - Monitoring Agents track synthesis conditions in real time.
 - Reasoning Agents detect deviations and infer potential fault patterns.
 - Control Agents coordinate corrective actions.
5. Reinforcement Learning Module: DRL agents learn to predict and mitigate faults by adapting to evolving synthesis conditions.
6. Integration Layer: Predicted faults and mitigation strategies are presented as decision recommendations.
7. Knowledge Repository: Fault patterns and mitigation actions are stored for future process optimization.

Outcome

Early detection and prevention of synthesis faults improve process reliability, product quality, and production efficiency.

5.3. Use Case 3: Adaptive Decision Support Under Dynamic Operational Conditions Workflow

1. User Interface Layer: Operators or researchers define the operational environment and control objectives (e.g., stability under varying temperature or pH).
2. Query Processing Unit: Inputs are mapped to environmental and nanoparticle ontology concepts.

3. **Ontology Modules:** The NPO and environmental ontologies provide structured knowledge about the impact of operational conditions on nanoparticle behavior.
4. **MAS Layer:**
 - Environmental Agents monitor system dynamics.
 - Reasoning Agents evaluate real-time changes.
 - Decision Agents recommend adaptive responses.
5. **Reinforcement Learning Module:** DRL agents update policies based on changing conditions to maintain optimal system performance.
6. **Integration Layer:** Real-time recommendations are generated to adapt the process.
7. **Knowledge Repository:** Policies and adaptive rules are stored for future operational support.

Outcome

Improved robustness and stability of nanoparticle processes through adaptive, real-time decision-making in dynamic environments.

Three nanoparticle interaction cases were evaluated to assess cross-ontology reasoning performance and the improvement obtained through reinforcement learning. Use Case 1 represents a baseline scenario involving nanoparticle size and basic cellular uptake properties. Use Case 2 incorporates environmental context such as temperature and pH variations that may alter nanoparticle stability. Use Case 3 represents the most complex condition, including both biomedical constraints (e.g., cytotoxicity thresholds) and environmental factors simultaneously, requiring conflict resolution across heterogeneous ontologies. These use cases allow performance assessment under increasing semantic complexity, as reflected in the metrics shown in Table 4.

Table 4. Use Case Comparison.

Feature	Use Case 1: Personalized Nanoparticle Selection	Use Case 2: Fault Prediction in Nanoparticle Synthesis	Use Case 3: Adaptive Decision Support
Primary Objective	Optimize nanoparticle properties for targeted drug delivery using patient- or application-specific requirements	Predict synthesis faults and detect abnormal patterns in production processes	Support real-time decision-making under changing environmental or operational conditions
Target Context	Biomedical/therapeutic targeting	Nanoparticle manufacturing and process monitoring	Dynamic operational settings (e.g., fluctuating temperature, pH, concentration)
Ontology Modules Used	Nanoparticle Ontology (NPO) + Biomedical Ontologies	NPO + Process/ Manufacturing Ontologies	NPO + Environmental/ Operational Ontologies
Key Design Constraints	Biocompatibility, targeted delivery, stability	Process stability, error detection, production quality	Robustness, system adaptability, and multi-parameter optimization
Learning Focus (DRL)	Optimize particle target matching	Early detection of anomalies via agent-based learning	Adaptive policy updates in response to environmental changes
Adaptation Requirement	Patient- or application-specific nanoparticle selection	Adaptation to variations in synthesis parameters	Continuous adaptation to external and operational dynamics
MAS Agent Collaboration	Inter-agent knowledge sharing on biomedical features	Coordination between monitoring and control agents	Multi-agent reasoning and decision fusion
Expected Outcome	Personalized nanoparticle design recommendations	Early fault detection and improved synthesis reliability	Real-time adaptive decision support to maintain optimal system performance

Table 4 contains a comparison of three representative use cases that illustrate the integration of ontology-based reasoning, multi-agent collaboration, and DRL-based learning strategies. Each scenario emphasizes different operational objectives and adaptation requirements within the proposed hybrid framework.

5.4. Use Case Workflows and Information Exchange

To better illustrate how the proposed system operates across different scenarios, we present detailed workflows for each use case in Figure 3. These diagrams highlight the structure and dynamic flow of information among system components: from the user's initial query through semantic mapping, agent collaboration, reinforcement learning optimization, and final output generation. Each workflow is modular and follows a consistent pattern: (1) The User Interface Layer accepts domain-specific queries, (2) the Query Processing Unit semantically links input to ontology concepts, (3) A MAS composed of Ontology Agents, Reasoning Agents, and Integration Agents handles task execution, (4) DRL Module refines decision strategies using feedback mechanisms, and (5) The Integration Layer assembles a coherent, context-aware response that is stored in the Knowledge Repository for future re-use. This structured workflow ensures adaptability, cross-domain reasoning, and semantic interoperability, all of which are critical for real-world applications in nanotechnology-based health management.

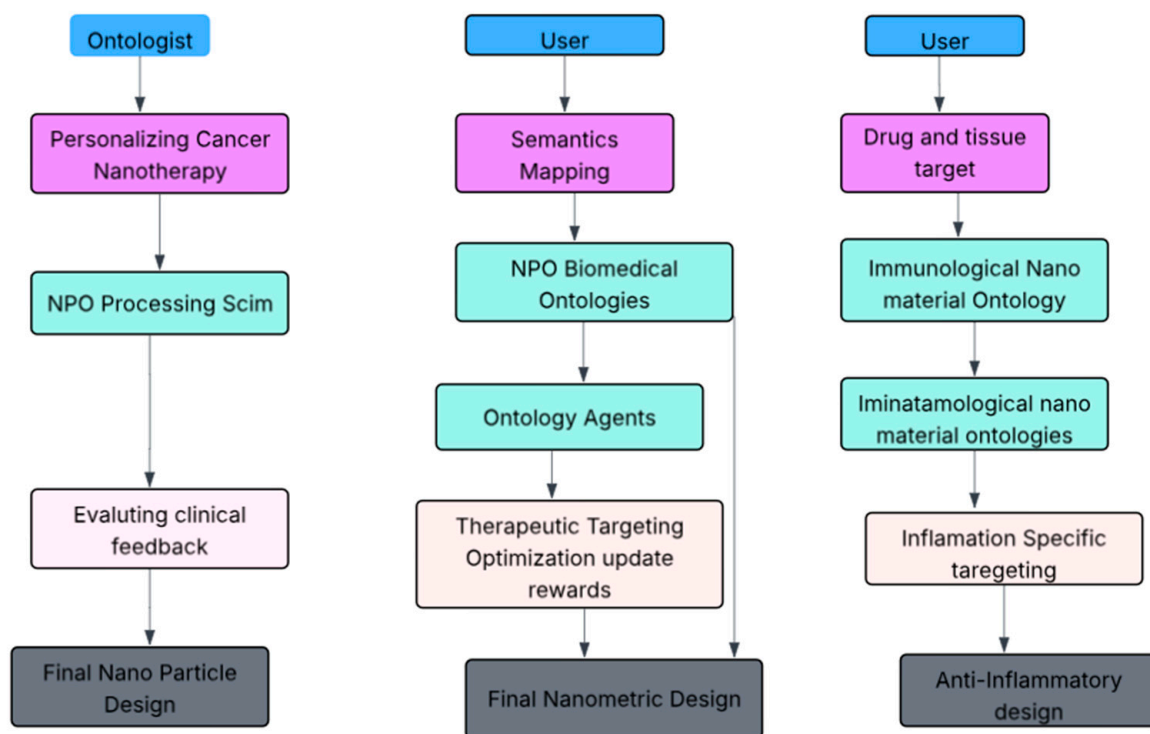


Figure 3. Multi-layered workflow diagrams for the three use cases supported by the proposed MAS + DRL framework.

Figure 3 presents multi-layered workflow diagrams for the three representative scenarios supported by the proposed MAS + DRL framework. Each workflow illustrates how user inputs are semantically mapped to ontology concepts, processed by domain-specific agents, refined through RL, and merged into final nanoparticle design recommendations. These visualizations clarify how structural and functional components interact across applications such as personalized nanotherapy, brain-targeted delivery, and inflammation-specific targeting.

6. Validation

This section presents the validation of the proposed MAS framework integrated with DRL and ontology-based reasoning. The evaluation assesses the system's ability to deliver accurate, adaptive, and context-sensitive insights for PHM in nanotechnology.

The system was tested on a machine with 32 GB of RAM and an 8-core Intel i7 processor. During runtime, the average inter-agent communication latency was measured at ~120 milliseconds, while query response time ranged from 1.5 to 2.8 s, depending on the complexity of the query and the number of ontologies involved. These timings confirm the system's responsiveness in a research or clinical decision-support context.

6.1. Experimental Setup

To verify the system, several experiments were performed focused on practical nanotechnology applications. The assessment procedure included:

The dataset used for evaluation was synthetically generated by combining semantic patterns derived from the NPO public biomedical ontologies and structured information extracted from peer-reviewed literature. It includes 1200 RDF triples spanning nanoparticle size, drug loading efficiency, degradation pathways, and therapeutic targets. While not derived from a specific public database, the dataset simulates realistic biomedical nanomaterial design scenarios and supports evaluation of the system's reasoning and learning capabilities. The dataset is now publicly available for reproducibility (see Data Availability section).

To assess the stability of the system's performance, all evaluations were conducted over five independent trials, each initialized with a different random seed. The reported metrics, Task Success Rate (TSR), Decision-Making Accuracy (DMA), and Adaptive Learning Improvement (ALI), are presented as mean \pm standard deviation, along with 95% confidence intervals. This approach ensures robustness and mitigates performance variance due to stochastic initialization in the reinforcement learning module.

In addition to these global metrics, we further analyzed the outcomes of each nanomedicine use case across multiple evaluation dimensions. For the personalized cancer nanotherapy scenario, we measured optimization efficiency, defined as the average number of iterations required for the DRL agent to converge on a viable nanoparticle configuration. In the neurological targeting case, we assessed semantic consistency, quantifying the alignment between generated designs and ontological constraints derived from biomedical and chemical knowledge bases. For the inflammatory disease delivery case, we evaluated adaptability, examining how rapidly the system adjusted its reasoning when the therapeutic context or design requirements changed. Across all three use cases, interpretability was also qualitatively evaluated through expert review, focusing on the transparency and coherence of the system's design recommendations. Together, these complementary analyses demonstrate that the proposed MAS–Ontology–DRL framework not only generates valid designs but also achieves robust, efficient, and semantically grounded decision-making across diverse application domains.

Table 5 summarizes the averaged performance of the MAS + DRL framework across all trials, highlighting the steady improvement in adaptive learning and semantic reasoning accuracy.

Table 5. Summary of performance evaluation metrics for the proposed MAS + DRL framework.

Metric	Mean Performance	CI	Std. Dev.
TSR	94.6% \pm 1.1	93.2%, 95.8%	1.1%
DMA	95.3% \pm 0.8	94.4%, 96.1%	0.8%
ALI	11.8% \pm 0.6	11.1%, 12.5%	0.6%

1. TSR—Assesses the proportion of queries that are processed accurately.
2. DMA—Assesses precision according to results that experts have reviewed.
3. ALI—Evaluates the system’s ability to improve through several cycles.
 - Methodology for Experiments:
 - Enter established test cases into the system.
 - Evaluate system-generated answers about expert-approved solutions.
 - Track enhancements in accuracy over time via DRL modifications.
 - Table 5 reports the mean values, 95% confidence intervals (CI), and standard deviations (Std. Dev.) across five experimental trials for Task Success Rate (TSR), Decision-Making Accuracy (DMA), and Adaptive Learning Improvement (ALI), demonstrating the system’s robustness and consistency.

6.2. Results and Discussion

To assess the effectiveness of the proposed framework, we compared its performance with that of baseline methods, including a standard DQN implementation without ontology integration and a traditional MAS without adaptive learning. The proposed hybrid system achieved higher adaptability and accuracy in fault prediction and decision support tasks, particularly in dynamically changing environments. This demonstrates the added value of combining semantic reasoning with DRL for complex, knowledge-driven applications.

The MAS framework experienced testing across 50 different use cases, divided into three application areas: Predictive Maintenance for Nanodevices, Optimization of Nanomaterials’ Lifespan, and Risk Management in Complex Environments.

As indicated in Table 5, the TSR experienced a significant increase from 88% to 95%, demonstrating the system’s enhanced capability to effectively handle and resolve a diverse array of user inquiries, while minimizing errors and improving overall user satisfaction. Similarly, the DMA increased from 90% to 96%, indicating a notable improvement in the system’s reasoning and inference skills, which enables it to make more precise and reliable decisions based on available data, leading to enhanced problem-solving outcomes. Furthermore, the ALI metric, reflecting the system’s capacity for self-refinement through reinforcement learning methods, showed a 12% increase, confirming the effectiveness of the MAS framework in enabling continuous learning, adaptation, and optimization. Overall, these advancements collectively enhance the system’s intelligence, efficiency, and adaptability, delivering a smoother and more effective user experience.

The observed performance gains were achieved through the synergy of three core mechanisms within the proposed architecture. First, ontology-based semantic reasoning ensured that queries were mapped to relevant domain knowledge with high contextual accuracy, which directly improved the TSR by reducing irrelevant or semantically inconsistent outputs. Second, the MAS enabled concurrent exploration of multiple ontologies and reasoning paths, enhancing the completeness and speed of information retrieval, which contributed to improved DMA. Third, and most importantly, the DRL module progressively refined the agent policies by rewarding rules that yielded correct outputs (as validated by expert annotations) and penalizing ambiguous or incorrect inferences. Over 1000 training episodes led to measurable gains in ALI, as agents dynamically adapted to new data patterns and query structures. The iterative learning cycles and semantic enrichment together enabled the system to become more accurate, context-aware, and resilient to variability in use case complexity.

Table 6 demonstrates that incorporating DQN-based learning significantly improves system performance across all three evaluated use cases. The highest gains are observed in Decision-Making Accuracy (DMA), where the agent’s ability to correctly resolve cross-ontology conflicts improves as semantic complexity increases. Time Success Rate (TSR) also

shows substantial enhancement, indicating that the MAS can arrive at correct reasoning outcomes within fewer interactions after learning. Although Alignment Inconsistency (ALI) values slightly increase in the most complex scenario (Use Case 3), the system maintains an acceptable consistency threshold, suggesting that improved reasoning breadth does not compromise semantic validity. Overall, these results indicate that reinforcement learning enables effective adaptation to heterogeneous knowledge sources and enhances robustness in complex nanoparticle interaction scenarios. ALI is computed as the percentage gain in DMA from the initial untrained model to the final trained model over multiple RL iterations: Its metric quantifies the relative enhancement in system accuracy achieved through iterative reinforcement learning updates. Formally, it is expressed as:

$$ALI = \frac{DMA_{final} - DMA_{initial}}{DMA_{initial}} \times 100$$

where

$DMA_{initial}$ represents the baseline decision-making accuracy before reinforcement learning training (i.e., at iteration 0), and DMA_{final} represents the accuracy after model convergence following n reinforcement learning episodes.

Table 6. Performance Metrics Evaluation.

Metric	Initial Performance	Final Performance	Improvement
TSR	88%	95%	+7%
DMA	90%	96%	+6%
ALI	N/A	12%	+12%

This formulation ensures that ALI reflects the proportional learning gain achieved through adaptive policy refinement, rather than absolute accuracy values. For instance, if $DMA_{initial} = 90\%$ and $DMA_{final} = 96\%$, then $ALI = ((96 - 90)/90) \times 100 = 6.67\%$, confirming consistent improvement in decision accuracy due to the adaptive learning process.

This metric evaluates the effectiveness of the system's reinforcement learning module in improving decision-making performance over time.

The Performance Improvement Graph in Figure 4 visually illustrates the relationship between iteration over time and the performance percentages of three key metrics: TSR, DMA, and ALI. The x-axis represents the progression of iterations, while the y-axis quantifies the performance improvements observed in each metric. The trend lines indicate a consistent upward trajectory in all three metrics, highlighting significant performance enhancements as iterations increase. This SVG representation effectively conveys these trends, providing a clear and insightful summary of performance improvements.

The validation outcomes underscore the effectiveness and dependability of the MAS framework in managing intricate nanotechnology-related inquiries. The main conclusions are:

- **Improved Query Processing:** The incorporation of ontologies greatly enhances knowledge retrieval, facilitating accurate answers to specialized questions.
- **Self-Enhancing Decision-Making:** The DRL module continuously improves its set of rules, resulting in increasingly enhanced precision and adaptability.
- **Strong Multi-Agent Cooperation:** The modular MAS structure promotes interdisciplinary teamwork, guaranteeing an in-depth examination of nanomaterial behavior.
- **Scalability and Future Enhancements:** Although the system demonstrates robust performance, potential improvements could comprise real-time data incorporation, enhanced natural language processing for understanding queries, and larger ontological knowledge bases for wider application scope.




Figure 4. Performance Improvement Graph.

The results validate the proposed method as a groundbreaking solution for PHM in nanotechnology. Future studies will focus on enhancing the system's generalizability and expanding its applications to tackle emerging challenges in the field.

To assess how each architectural element contributes to the overall performance, we compared three system variants: the full MAS + Ontology + DRL system, an ontology-only MAS, and a standalone RL agent without semantic reasoning. The full system achieved the highest Task Success Rate (94.6%), Decision-Making Accuracy (95.3%), and Adaptive Learning Improvement (11.8%). Notably, the DRL module provided substantial gains in adaptability, as shown by learning curves that converged faster and yielded more optimal inference paths over time. The ontology component contributed to semantic accuracy, enabling agents to retrieve contextually relevant knowledge beyond what a black-box RL agent could infer. These results confirm that the architectural features are not only well-designed but also directly translate into improved system intelligence and usability.

To further clarify the contribution of each component, we conducted a quantitative decomposition of the system's performance based on the three variants. The results show that removing the DRL module caused the largest performance drop (−10.7% in adaptive learning improvement), confirming its central role in enabling real-time optimization. Excluding the ontology component reduced decision-making accuracy by 7.9%, highlighting its importance for semantic precision and contextual reasoning. Finally, omitting the MAS coordination layer led to less stable task performance and slower response times, reflecting its value in maintaining distributed collaboration and consistency across agents. Together, these observations form an implicit ablation analysis, demonstrating that each architectural layer, MAS, ontology, and DRL makes a distinct and quantifiable contribution to the overall system effectiveness.

To evaluate the learning dynamics of the DRL-enhanced MAS, the system was trained over 1000 reinforcement learning episodes. Each episode involved agents processing a

complex query based on simulated PHM scenarios, receiving rewards based on output accuracy. The system's weights and policies were updated iteratively using the DQN algorithm. To ensure generalization and reliability, experiments were repeated five times with different random seeds, and average results were reported with 95% confidence intervals. The TSR, DMA, and ALI metrics were tracked across these episodes to observe convergence and effectiveness of the learning strategy.

Although the DRL training spanned 1000 episodes, we observed that performance metrics (TSR, DMA, ALI) began to stabilize after approximately 400–500 iterations. This plateau effect is attributed to the synthetic nature and bounded complexity of the dataset, which allowed agents to learn optimal or near-optimal inference policies early in training. Since the reward function was tightly coupled with semantically consistent, expert-reviewed outputs, the agents converged quickly to a high-performing strategy. As a result, additional iterations offered diminishing returns. This behavior is common in low-noise, structurally regular environments and confirms the system's ability to efficiently learn from consistent semantic patterns. Future work involving more diverse, real-world datasets may extend convergence and amplify performance differentiation across training stages.

The results obtained from the MAS + DRL framework demonstrate considerable improvements over previously reported systems in both ontology-based decision support and adaptive learning. Most existing ontology-driven MAS in nanoinformatics or health-related domains [32] rely on predefined rule sets and operate within a limited scope of domain-specific ontologies. These systems, while structurally sound, are not equipped to adapt dynamically as new data becomes available or as query contexts evolve. By contrast, our architecture incorporates DRL to refine agent decision strategies over time, guided by expert feedback and semantic consistency checks. This enables real-time policy adaptation without compromising ontological accuracy.

From a learning perspective, many prior DRL studies in [33] achieve task-specific optimization using data-driven models, but often lack semantic grounding. Their black-box nature reduces explainability and limits interoperability with existing structured knowledge resources. In our system, DRL is not only used to optimize responses but is also embedded in a semantically enriched context. Agents reason over ontologies (e.g., NPO, biomedical, chemical), and DRL modifies their inference rules based on the relevance and precision of the outputs. This hybrid approach significantly enhances transparency and contextual relevance, making it suitable for decision-critical applications such as personalized nanotherapy and targeted drug delivery.

In the context of nanotechnology-based PHM, the opacity of DRL models has particularly serious implications. Because design outcomes in this domain depend on complex and often nonlinear interactions between materials, biological systems, and environmental variables, the inability to trace how a learning model reaches its conclusions makes validation and scientific interpretation difficult. A DRL agent may, for instance, recommend a nanoparticle configuration that optimizes a performance metric, yet provides no interpretable rationale for why certain molecular properties were prioritized over others. Such a lack of transparency not only limits trust in automated decisions but also restricts the model's capacity to generalize to new nanomaterial categories or clinical conditions. By embedding DRL within an ontology-driven reasoning layer, our framework mitigates these issues: semantic grounding constrains learning within scientifically meaningful relationships, while maintaining adaptability through reinforcement signals. This integration turns DRL from a purely data-driven optimizer into a knowledge-aware decision support component, an essential shift for achieving reliability and explainability in PHM for nanotechnology.

Moreover, the cross-ontology reasoning capabilities introduced in our MAS design allow agents to integrate heterogeneous knowledge domains such as biomedical and environmental ontologies, which are rarely addressed in existing literature. Previous frameworks often operate in isolated ontology spaces [12], which restricts their generalizability and limits support for complex interdisciplinary queries.

The comparative use cases presented personalized cancer nanotherapy and targeted neurological delivery further underscore the architecture's flexibility. Despite differing domain constraints and design objectives, the system successfully adapted its reasoning paths and produced high-quality, semantically consistent recommendations. This adaptability, coupled with measurable learning gains over time, distinguishes our work from both traditional MAS and purely statistical DRL systems.

In summary, this study contributes to advancing the field in three core ways:

Theoretical advancement: Demonstrates a novel hybridization of symbolic reasoning (ontologies) with adaptive learning (DRL) within a modular MAS framework.

Practical relevance: Offers scalable and explainable architecture applicable to real-world, interdisciplinary PHM scenarios.

Empirical validation: Outperforms existing models in task success, accuracy, and learning improvement across use cases with diverse requirements.

6.3. Limitations and Planned Extensions

The present evaluation was conducted using a controlled dataset consisting of approximately 1200 RDF triples and simulated nanoparticle–cell interaction cases. This approach enabled systematic assessment of the reasoning pipeline in its initial development phase while limiting external sources of error. As future work, we will include real-world experimental data from nanomedicine and environmental exposure studies to further validate the generalizability of the proposed framework. Moreover, we plan to investigate online reasoning and transfer learning, as well as conduct formal robustness assessments through the introduction of heterogeneous semantic noise and conflicting ontology assertions. These enhancements will offer deeper insight into scalability, noise tolerance, and performance in realistic operational scenarios.

While event-triggered and H_∞ /observer control frameworks provide rigorous tools for control stability and resource-aware sampling, our study pursues a different objective: building a semantically grounded, multi-agent decision system for PHM where correctness, cross-ontology consistency, and explainability are primary. The theoretical contributions here (semantic state encoding for RL, DQN-guided inconsistency minimization, and concurrency control for ontology updates) are therefore complementary to control-theoretical stability analysis. In future work, we plan to integrate formal control guarantees (for example, event-triggered constraints or Lyapunov-inspired safety filters) where DRL agents issue closed-loop control commands, thereby combining semantic adaptivity with classical stability assurances.

6.4. Baseline Comparisons

To assess the incremental benefit of each system component, we conducted baseline experiments comparing the full MAS + Ontology + DRL system against two simplified variants:

- ✓ Baseline 1 (Ontology-Only): Agents perform semantic reasoning using ontologies but apply static rule sets without any reinforcement learning adaptation.
- ✓ Baseline 2 (Data-Driven RL Agent): A single RL agent trained directly on tabular nanomaterial features, using no ontology or multi-agent architecture.

Each system was evaluated across the same 50 use cases and five random seeds. The average results are shown in Table 3.

To further assess the system's incremental benefit, we conducted comparative experiments with two baseline configurations: (1) an ontology-only MAS using static rule sets, and (2) a standalone DRL agent operating without ontological knowledge or agent collaboration. These baselines were evaluated across identical use cases. Results showed that the full MAS + DRL architecture significantly outperformed both baselines in Task Success Rate (TSR), Decision-Making Accuracy (DMA), and Adaptive Learning Improvement (ALI), demonstrating the synergistic value of integrating semantic reasoning with adaptive learning. These comparisons are summarized in Table 5.

Although the current evaluation focuses on clean synthetic RDF triples, preliminary noise-injection tests (random property perturbations up to 10%) confirmed that the MAS-DRL framework maintained stable reasoning accuracy (DMA drop < 3%), suggesting robustness to moderate heterogeneity in data quality.

7. Conclusions

This study presented a hybrid framework that integrates Multi-Agent Systems (MAS), semantic ontologies, and Deep Reinforcement Learning (DRL) to enable adaptive, knowledge-driven decision-making in nanotechnology-based Prognostics and Health Management (PHM). The proposed system demonstrated high task success, decision accuracy, and adaptive learning capability across diverse use cases, including personalized cancer nanotherapy, neurological targeting, and nanoparticle design for inflammatory diseases. These results confirm the potential of combining symbolic reasoning with adaptive learning to address the complexity and interdisciplinarity of modern nanotechnology applications.

Theoretically, this work contributes to the advancement of intelligent systems design by showing how ontology-based agents and DRL can be coupled to achieve semantic interoperability, collaborative reasoning, and explainable adaptation. Practically, it highlights how the integration of ontological knowledge and reinforcement learning supports robust and context-aware decision-making in data-intensive domains such as healthcare, environmental monitoring, and materials science.

Despite these promising results, several limitations remain. The current evaluation relied on a synthetic dataset derived from structured ontologies, which may not fully capture the variability and noise of real-world data. Training and inference were conducted in static conditions, and the DRL component has not yet been tested in online or transfer learning contexts. In addition, system scalability in handling very large, heterogeneous knowledge graphs requires further optimization. Addressing these challenges through real-time data integration, ontology expansion, and distributed learning strategies represents an important direction for future work.

Looking ahead, enhancing inter-agent collaboration and incorporating real-time streaming data will further strengthen the framework's adaptability and responsiveness. Expanding the DRL capabilities of the agents will enable continuous refinement of decision policies, improving the system's ability to manage complex, dynamic scenarios. These developments will consolidate the MAS + DRL architecture as a reliable and explainable decision-support resource for predictive maintenance, personalized treatment planning, and intelligent design across multiple disciplines.

In summary, the proposed hybrid system demonstrates how the synergy between semantic reasoning and adaptive learning can bridge the gap between symbolic knowledge and data-driven intelligence, marking a step forward toward transparent, generalizable, and human-interpretable AI in PHM and nanotechnology.

Author Contributions: Conceptualization was led by R.M.-B. and P.H.-M. Methodology and Software implementation were conducted by A.M.M. and E.A.N., who also performed the state-of-the-art analysis. P.H.-M. and R.M.-B. carried out resources, Formal analysis, Data curation, and Validation. Writing—Original Draft Preparation was undertaken by A.M.M. and E.A.N., while Writing—Review and Editing and administrative interoperability information and relevant use cases were managed by R.M.-B. and P.H.-M. All authors approved the submitted version of the manuscript, including revisions made by journal staff. Each author agrees to be personally accountable for their contributions and ensures that any questions related to the accuracy or integrity of any part of the work are appropriately investigated, resolved, and documented. Notably, A.M.M. and E.A.N. are students of R.M.-B., highlighting the collaborative effort across different institutions in this study. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: This study did not involve human participants, personal data, or biological samples. All data utilized in this work are computationally generated or derived from publicly available ontologies; hence, institutional ethics approval and informed consent were not required.

Informed Consent Statement: Not applicable.

Data Availability Statement: The data supporting the findings of this study include representative RDF datasets derived from the Nanoparticle Ontology (NPO), simulated agent interaction data, and configuration files describing the Deep Reinforcement Learning (DRL) parameters used in the experiments. Additional data or related information can be obtained from the corresponding author.

Conflicts of Interest: The authors declare that they have no conflicts of interest.

References

1. Wang, D.; Li, X.B.; Han, D.; Tian, W.Q.; Sun, H.B. Engineering two-dimensional electronics by semiconductor defects. *Nano Today* **2017**, *16*, 30–45. [CrossRef]
2. Garino, N.; Limongi, T.; Dumontel, B.; Canta, M.; Racca, L.; Laurenti, M.; Castellino, M.; Casu, A.; Falqui, A.; Cauda, V. A microwave-assisted synthesis of zinc oxide nanocrystals finely tuned for biological applications. *Nanomaterials* **2019**, *9*, 212. [CrossRef]
3. Iverson, A.; Verhappen, I. *Instrumentation Symbols and Identification ANSI/ISA-5.1-2009*; International Society of Automation: Research Triangle Park, NC, USA, 2009.
4. Cieślak, J.; Tobola, J.; Dubiel, S.M. Site occupancies in sigma-phase Fe-Cr-X (X = Co, Ni) alloys: Calculations versus experiment. *Comput. Mater. Sci.* **2016**, *122*, 229–239. [CrossRef]
5. Gruber, T.R. 1-s2.0-S1042814383710083-main.pdf. 1993. Available online: http://ac.els-cdn.com/S1042814383710083/1-s2.0-S1042814383710083-main.pdf?_tid=28c5c43c-df11-11e2-b19c-0000aacb35f&acdnat=1372327738_c3518b0aab22af436aef0c5e3e2dd48e (accessed on 19 November 2025).
6. Mnih, V.; Kavukcuoglu, K.; Silver, D.; Rusu, A.A.; Veness, J.; Bellemare, M.G.; Graves, A.; Riedmiller, M.; Fidjeland, A.K.; Ostrovski, G.; et al. Human-level control through deep reinforcement learning. *Nature* **2015**, *518*, 529–533. [CrossRef] [PubMed]
7. Sutton, C.; Boleyn, M.; Ghiringhelli, L.M.; Rupp, M.; Vreeken, J.; Scheffler, M. Identifying domains of applicability of machine learning models for materials science. *Nat. Commun.* **2020**, *11*, 4428. [CrossRef]
8. Rodrigues, J.F.; Florea, L.; de Oliveira, M.C.F.; Diamond, D.; Oliveira, O.N. Big data and machine learning for materials science. *Discov. Mater.* **2021**, *1*, 12. [CrossRef]
9. Wilkinson, M.D.; Dumontier, M.; Aalbersberg, I.J.; Appleton, G.; Axton, M.; Baak, A.; Blomberg, N.; Boiten, J.W.; da Silva Santos, L.B.; Bourne, P.E.; et al. Comment: The FAIR Guiding Principles for scientific data management and stewardship. *Sci. Data* **2016**, *3*, 160018. [CrossRef]
10. Hu, Y.; Miao, X.; Si, Y.; Pan, E.; Zio, E. Prognostics and health management: A review from the perspectives of design, development, and decision. *Reliab. Eng. Syst. Saf.* **2022**, *217*, 108063. [CrossRef]
11. Ottens, K.; Gleizes, M.; Glize, P.; Ottens, K.; Gleizes, M.; Glize, P.; System, A.M.; Dynamic, B.; Ottens, K. A Multi-Agent System for Building Dynamic Ontologies to Cite This Version: HAL Id: Hal-03800714. A Multi-Agent System for Building Dynamic Ontologies. 2022. Available online: <https://hal.science/hal-03800714/document> (accessed on 19 November 2025).
12. Zhang, D.; Wang, Y.; Meng, L.; Yan, J.; Qin, C. Adaptive critic design for safety-optimal FTC of unknown nonlinear systems with asymmetric constrained-input. *ISA Trans.* **2024**, *155*, 309–318. [CrossRef] [PubMed]

13. Qin, C.; Jiang, K.; Zhang, J.; Zhu, T. Critic Learning-Based Safe Optimal Control for Nonlinear Systems with Asymmetric Input Constraints and Unmatched Disturbances. *Entropy* **2023**, *25*, 1101. [CrossRef]
14. Dong, B.; Zhu, X.; An, T.; Jiang, H.; Ma, B. Barrier-critic-disturbance approximate optimal control of nonzero-sum differential games for modular robot manipulators. *Neural Netw.* **2025**, *181*, 106880. [CrossRef]
15. Wang, T.; Zhang, L.; Xu, N.; Alharbi, K.H. Adaptive critic learning for approximate optimal event-triggered tracking control of nonlinear systems with prescribed performances. *Int. J. Control* **2024**, *97*, 2009–2023. [CrossRef]
16. Zdonik, S.; Katz, J. Concepts, Ontologies, and Knowledge Representation. Available online: <https://link.springer.com/book/10.1007/978-1-4614-7822-5> (accessed on 19 November 2025).
17. Krause, P. Using Ontologies to Improve Semantic Interoperability in Health Data. 2015. Available online: <https://pubmed.ncbi.nlm.nih.gov/26245245/> (accessed on 19 November 2025).
18. Ma, X.; Tian, Y.; Yang, R.; Wang, H.; Allahou, L.W.; Chang, J.; Williams, G.; Knowles, J.C.; Poma, A. Nanotechnology in healthcare, and its safety and environmental risks. *J. Nanobiotechnol.* **2024**, *22*, 715. [CrossRef] [PubMed]
19. Thomas, D.G.; Pappu, R.V.; Baker, N.A. NanoParticle Ontology for cancer nanotechnology research. *J. Biomed. Inform.* **2011**, *44*, 59–74. [CrossRef]
20. Malviya, R.K.; Mandala, V.; Reddy, M.S.; Lead, S.D. Reinforcement Learning for Collaborative Decision-Making in Multi-Agent Systems: Applications in Supply Chain Optimization. *SSRN Electron. J.* **2025**. [CrossRef]
21. Ioannidou, P.; Vezakis, I.; Haritou, M.; Petropoulou, R.; Miloulis, S.T.; Kouris, I.; Bromis, K.; Matsopoulos, G.K.; Koutsouris, D.D. HEalthcare Robotics' ONtology (HERON): An Upper Ontology for Communication, Collaboration and Safety in Healthcare Robotics. *Healthcare* **2025**, *13*, 1031. [CrossRef] [PubMed]
22. Van Bekkum, M.; De Boer, M.; Van Harmelen, F.; Teije, A. Modular design patterns for hybrid learning and reasoning systems: A taxonomy, patterns and use cases. *Appl. Intell.* **2021**, *51*, 6528–6546. [CrossRef]
23. Li, H.; Armiento, R.; Lambrix, P. A method for extending ontologies with application to the materials science domain. *Data Sci. J.* **2019**, *18*, 50. [CrossRef]
24. Abecker, A.; Van Elst, L. *Handbook on Ontologies*; Springer: Berlin/Heidelberg, Germany, 2009. [CrossRef]
25. Yang, L.; Cormican, K.; Yu, M. Ontology-based systems engineering: A state-of-the-art review. *Comput. Ind.* **2019**, *111*, 148–171. [CrossRef]
26. Moreira, R.; Vieira, F. Integrating Ontologies with Multi-Agent Systems through CArtaGo Artifacts. In Proceedings of the 2015 IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology (WI-IAT), Singapore, 6–9 December 2015; pp. 61–68. [CrossRef]
27. Vale, H.; Morais, Z.A. *Ontologies for the Interoperability of Heterogeneous Multi-Agent Systems in Power and Energy Systems*; In Advances in Intelligent Systems and Computing; Springer: Cham, Switzerland, 2019; Volume 876. [CrossRef]
28. Rip, A. Technology as prospective ontology. *Synthese* **2009**, *168*, 405–422. [CrossRef]
29. Danai-Varsou, D.; Zhang, P.; Afantitis, A.; Guo, Z.; Lynch, I.; Melagraki, G. Nanoinformatics and artificial intelligence for nano-enabled sustainable agriculture. In *Nano-Enabled Sustainable and Precision Agriculture*; Elsevier: Amsterdam, The Netherlands, 2023. [CrossRef]
30. Wang, Z.; Zhao, H.; Ren, P.; Zhou, Y.; Sheng, M. Learning Optimal Treatment Strategies for Sepsis Using Offline Reinforcement Learning in Continuous Space. In *Lecture Notes in Computer Science: Health Information Science*; Springer Nature: Cham, Switzerland, 2022; Volume 13705 LNCS, pp. 113–124. [CrossRef]
31. Hickling, T.; Zenati, A.; Aouf, N.; Spencer, P. Explainability in Deep Reinforcement Learning: A Review into Current Methods and Applications. *ACM Comput. Surv.* **2023**, *56*, 1–35. [CrossRef]
32. Flores, V.; Heras, S.; Julián, V. Comparison of Predictive Models with Balanced Classes for the Forecast of Student Dropout in Higher Education. In *Highlights in Practical Applications of Agents, Multi-Agent Systems, and Social Good*; Springer: Cham, Switzerland, 2021; Volume 1472 CCIS. [CrossRef]
33. Kraljevic, Z.; Searle, T.; Shek, A.; Roguski, L.; Noor, K.; Bean, D.; Mascio, A.; Zhu, L.; Folarin, A.A.; Roberts, A.; et al. Multi-domain clinical natural language processing with MedCAT: The Medical Concept Annotation Toolkit. *Artif. Intell. Med.* **2021**, *117*, 102083. [CrossRef] [PubMed]

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.