Journal of Computer Technology and Software

ISSN:2998-2383

Vol. 5, No. 3, 2026

Unified Structure-Aware Representation Learning for Agent Perception and Decision Making

Jinming Li
Georgia Institute of Technology, Atlanta, USA
jli846@gatech.edu

Abstract: This paper proposes a structure-aware agent representation learning algorithm to address the limitations of agents in structural understanding and semantic modeling within complex dynamic environments. The algorithm centers on structural dependency modeling and constructs a topology-aware latent representation space that unifies state features, relational structures, and policy generation. Methodologically, the model first employs a structural graph encoding mechanism to extract temporal dependencies and spatial correlations from the environment, capturing multi-level contextual information. It then aligns local features with global structures through a semantic representation learning module, forming interpretable structural-semantic mappings in the feature space. During optimization, structural consistency constraints and representation regularization losses are introduced to ensure stable training and robust expression under noise interference, sparse feedback, and structural drift. Experimental results demonstrate that the proposed method achieves superior performance under multidimensional sensor noise, graph perturbations, and time-varying conditions, significantly improving semantic consistency, path planning, and policy convergence speed. Compared with existing methods, the proposed structure-aware framework attains higher consistency in structural information extraction and semantic fusion accuracy, showing strong generalization and interpretability, and providing a unified structured modeling paradigm for agent perception and decision-making in complex scenarios.

Keywords: Structure-aware representation learning; semantic alignment; agent modeling; topological consistency

1. Introduction

In the continuous evolution of agent systems, the increasing complexity of environments and the diversity of tasks are driving agents to evolve from perception-driven mechanisms toward structured cognition and semantic modeling. Traditional agent algorithms often rely on direct mappings among states, actions, and rewards, learning policies or value functions through deep neural networks to achieve self-optimization from experience. However, as the state space and interaction dimensions of environments expand dramatically, such methods show clear limitations in high-dimensional, non-stationary, and multi-constraint scenarios. Specifically, agents often lack effective capabilities for representing and reasoning about structured information, such as spatial topology, hierarchical dependencies, and semantic relationships. This deficiency restricts their understanding of complex environments to shallow feature levels, making it difficult to form transferable and interpretable internal representations. Therefore, introducing structured cognitive mechanisms into agent modeling to achieve a transition from local perception to global structural understanding has become a key scientific challenge for advancing the intelligence level of agents[1,2].

With the expansion of applications such as multi-agent systems, embodied intelligence, autonomous driving, and decision control, the interaction patterns in agent environments are no longer isolated temporal mappings but instead exhibit strong structural coupling and semantic dependencies[3]. Objects, events, and their relationships in the environment constitute an implicit graph-like structure that determines not only information propagation paths but also directly influences agents' perceptual focus and decision boundaries. When agents rely solely on low-level vectorized encoding, their policy updates can be affected by local information bias, leading to misjudgment of environmental structures and policy degradation. In contrast, structure-aware representation learning explicitly models entity relationships and contextual dependencies, capturing global consistency and hierarchical constraints within the environment. This provides interpretable foundations for high-level cognition and behavior planning in agents. The process not only improves the quality of representation but also offers a unified framework for cross-task transfer and multi-scenario adaptation[4].

In the current research paradigm of agent systems, the core objective of representation learning is to build a feature space that effectively maps environmental semantics and task intentions. Traditional convolution- or recurrence-based feature extraction mechanisms have strong fitting abilities but struggle with complex structures and non-Euclidean relationships. In multi-entity and multi-constraint decision scenarios, agents must understand spatial topology, state transition rules, and causal dependencies simultaneously[5]. This requires representation learning to move from "representing data" to "modeling structures." Structure-aware representation learning addresses this need by incorporating structural information such as adjacency relations, dependency paths, and hierarchical attributes into the feature space. This enables agents to perform reasoning on graph structures or relational graphs, providing abstraction and generalization capabilities. This paradigm enhances the semantic expressiveness of models and strengthens their robustness and stability in dynamic environments, offering new theoretical support for the integration of perception and reasoning in intelligent agents[6].

From the perspective of the evolution of agent systems, single feature-driven models can no longer satisfy the demand for high-level semantic understanding and structured decision-making. Future agent algorithms must be capable of adaptive modeling of internal environmental structures and latent relationships, enabling information flow and dependency analysis among multi-dimensional features. The introduction of structure-aware mechanisms transforms agents from passive responders to active "structure interpreters" capable of internal modeling. This ability is particularly critical in complex tasks such as multi-objective cooperative decision-making, policy transfer, anomaly perception, and cross-domain adaptation, all of which rely on agents' implicit modeling of environmental structures and maintenance of semantic consistency. Through structure-aware representation learning, agents can form interpretable high-dimensional mappings in feature space, allowing them to capture key correlations more accurately during policy generation and improve overall decision efficiency and robustness[7].

In summary, structure-aware representation learning for agents is not only an algorithmic enhancement but also represents a cognitive paradigm shift for future intelligent systems. It emphasizes a transition from data-driven to structure-driven learning, integrating environmental understanding, relational modeling, and policy optimization into a unified framework. This approach effectively alleviates the challenges of generalization and structural mismatch faced by traditional agents in high-dimensional environments. It provides a theoretical and technical foundation for building interpretable, transferable, and self-organizing intelligent agents. Research on structure-aware representation learning helps advance artificial intelligence toward higher autonomy and cognitive capabilities and holds significant scientific and practical value for intelligent decision-making, interactive control, and multi-agent collaboration in complex systems[8].

2. Related work

The design of the proposed structure-aware representation learning framework builds upon a broad body of research in structural representation learning, relational modeling, reinforcement learning optimization, and emerging LLM-driven agent architectures. Existing studies on graph neural networks provide the theoretical and algorithmic foundation for modeling structured dependencies in complex environments. Early works established graph-based representation learning as an effective paradigm for capturing non-Euclidean relationships and relational information among entities [9]. In particular, the concept of relational inductive biases introduced a principled mechanism for embedding structural priors into deep learning models, enabling neural architectures to explicitly reason over entities and relations rather than independent feature vectors [10]. Subsequent analyses further examined the representational capacity of graph neural networks and demonstrated their ability to approximate complex relational functions through neighborhood aggregation mechanisms [11]. Practical architectures such as graph convolutional networks and their inductive extensions enabled scalable propagation of structural information across large relational graphs [12], [13]. Attention-based graph models further improved this process by dynamically weighting relational importance during feature propagation, enhancing the expressiveness and adaptability of graph-based representations [14]. Extensions of graph convolution techniques for relational data modeling and structured inference provided additional mechanisms for representing heterogeneous dependencies within complex systems [15]. Moreover, graph network frameworks have demonstrated the ability to integrate perception and control by modeling dynamic physical interactions and structured state transitions within learnable relational systems [16]. These developments collectively motivate the structural graph encoding mechanism adopted in this work, which aims to capture both local dependencies and global topological consistency in the agent’s latent representation space.

Beyond structural modeling, advances in relational reinforcement learning provide further methodological guidance for integrating structural representations with decision-making processes. Relational reinforcement learning frameworks demonstrate how structured state representations can improve policy learning by enabling agents to reason about interactions between entities and environmental components [17]. From a representation optimization perspective, geometric analyses of reinforcement learning representations highlight the importance of constructing latent spaces that preserve meaningful distances and structural relationships between states [18]. Earlier developmental reinforcement learning studies introduced proto-value functions to capture intrinsic structural regularities of environments, emphasizing the role of topology-aware representations in improving policy generalization and exploration efficiency [19]. Complementary research in self-supervised representation learning further enhances these capabilities by enabling models to learn informative features without explicit reward signals. Contrastive predictive coding establishes a framework for learning predictive latent representations that encode temporal and structural dependencies in sequential data [20], while contrastive reinforcement learning methods demonstrate how such representations can improve policy learning stability and sample efficiency [21]. In addition, auxiliary task learning strategies provide a mechanism for enriching the training signal of reinforcement learning systems through additional representation-oriented objectives, leading to more robust and generalizable feature learning [22]. These advances collectively inspire the semantic representation learning and contextual consistency optimization components of the proposed framework.

The methodological development is further informed by research in multi-agent learning and cooperative decision-making. Actor-critic architectures designed for multi-agent environments provide scalable solutions for learning decentralized policies while maintaining centralized training signals [23]. Counterfactual policy gradient approaches introduce more precise credit assignment mechanisms, allowing agents to estimate individual contributions within cooperative decision processes [24]. Value function factorization techniques further enhance coordination efficiency by decomposing global objectives into structured individual policies under monotonic consistency constraints [25]. These frameworks highlight the importance of structured representations and coordinated optimization when agents interact within complex environments, motivating the design of topology-aware representations that support stable multi-agent policy learning.

Recent advances in large language model-based agent frameworks also contribute to the conceptual foundation of this work by introducing new mechanisms for reasoning, memory integration, and knowledge structuring in intelligent systems. Adaptive task decomposition frameworks demonstrate how agents can dynamically restructure decision processes in non-stationary environments through iterative reasoning and strategy updating [26]. Self-driven exploration mechanisms further extend this paradigm by enabling agents to autonomously construct structured knowledge representations during interaction with complex environments [27]. Cognitive modeling approaches incorporating long-term memory and reasoning modules illustrate how structured representations can support long-horizon decision-making and hierarchical reasoning processes [28]. Additional studies on collaborative agent systems reveal the importance of maintaining role consistency and communication protocols in multi-agent coordination, emphasizing the need for structured representations that maintain semantic stability across dynamic interactions [29]. Practical implementations of multi-agent development environments further illustrate how structured reasoning frameworks can support complex task generation and execution pipelines [30].

Complementary work on semantic prior modeling and knowledge integration further informs the semantic alignment mechanisms employed in this research. Frameworks incorporating semantic priors demonstrate how structured contextual information can guide perception and decision-making processes in intelligent systems [31]. Integration of language models with structured knowledge sources highlights the benefits of combining symbolic relational structures with neural representations to improve reasoning accuracy and interpretability [32]. Hybrid reasoning architectures further demonstrate how large language models can be combined with structured inference mechanisms to support causal reasoning and complex decision analysis [33]. Recent research on semantic alignment and constrained generation mechanisms also emphasizes the importance of maintaining consistency between model outputs and structured semantic representations during inference [34]. Robustness-oriented studies introduce semantic calibration mechanisms that improve reliability and adversarial stability of large model-based classification systems [35]. Structure-aware decoding techniques further illustrate how structural constraints can guide generation processes in large-scale neural models [36], while retrieval-based reasoning frameworks demonstrate how multi-granular indexing and confidence constraints can enhance the reliability of knowledge retrieval and representation integration [37]. Finally, infrastructure-level optimization mechanisms for large-scale model deployment provide practical support for efficient inference and scalable integration of representation learning systems in distributed environments [38].

Together, these methodological developments provide the conceptual and technical foundation for the proposed structure-aware representation learning framework. By integrating graph-based structural modeling, representation learning optimization, reinforcement learning mechanisms, and structured reasoning strategies, the proposed approach aims to construct a unified representation space that captures both relational dependencies and semantic consistency. This design enables agents to achieve robust perception, stable policy generation, and interpretable decision-making in complex dynamic environments, extending existing representation learning paradigms toward more structured and cognitively grounded agent architectures.

3. Method

This study introduces a structure-aware agent representation learning algorithm designed to achieve high-level semantic understanding and stable policy generation in complex environments through explicit structural modeling and contextual dependency constraints. The proposed method integrates three core mechanisms within a unified framework: Structural Graph Encoding, Semantic Representation Learning, and Contextual Consistency Optimization. By constructing interpretable structural dependency representations in the latent space, the model effectively captures intrinsic correlations among multi-level features in dynamic environments, thereby enhancing the agent's structural cognition and behavioral consistency. The model architecture is shown in Figure 1.

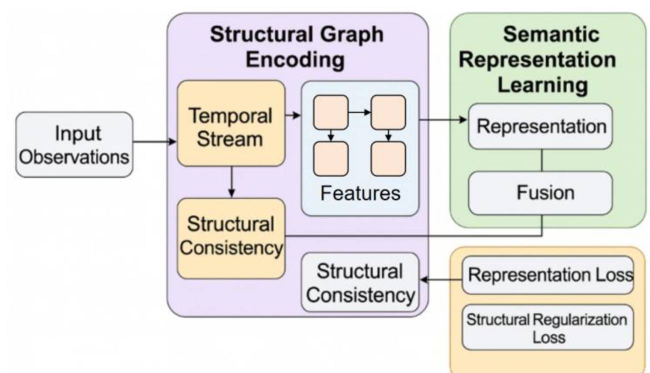

Figure 1. Overall model architecture

First, let the environment state space be $S = \{s_1, s_2, ..., s_n\}$, and each state s_i corresponds to a high-dimensional observation vector $x_i \in R^d$. To characterize the potential dependencies between states, a graph structure $G = (V, E)$ is introduced, where the node set V represents the state set and the edge set E represents the mutual dependencies between states. The model constructs a weighted adjacency matrix A through a feature similarity function, which is calculated as:

$$A_{ij} = \frac{\exp(sim(x_i, x_j) / \gamma)}{\sum_{k=1}^{n} \exp(sim(x_i, x_k)\gamma)}$$

Where $sim(\cdot)$ is the feature similarity metric function, and γ is the temperature coefficient, which is used to adjust the smoothness of the feature distribution. The weighted adjacency matrix is dynamically updated during training to reflect the time-varying characteristics of structural dependence.

After the structural graph is constructed, the model implements multi-layer semantic fusion through the graph feature propagation mechanism. The node representation of each layer is obtained by weighted aggregation of neighborhood features and can be expressed as:

$$h_i^{(l+1)} = \sigma\left(\sum_{j\in N(i)} A_{ij} W^{(l)} h_j^{(l)} + b^{(l)} \right)$$

Here, $h_i^{(l)}$ represents the feature representation of the node i at layer l, $W^{(l)}$ and $b^{(l)}$ are learnable parameters, and $\sigma(\cdot)$ is a nonlinear activation function. Through multi-layer propagation, the model gradually integrates local and global structural information, achieving hierarchical modeling of complex dependencies.

To improve the interpretability and stability of the representation, this study introduces a structural regularization term to constrain the smoothness and consistency of the graph during feature propagation. The structure-aware regularization objective is defined as:

$$L_{struct} = \sum_{(i,j)\in E} A_{ij} \left\| h_i - h_j \right\|_2^2$$

This item is used to limit the feature differences between adjacent nodes in the structure, avoid representation instability caused by over-discretization, and thus maintain the semantic consistency of the local structure.

Furthermore, to achieve consistent modeling of the global context, the model introduces a context-dependent variable z_t in the latent space, whose dynamic update follows the state transition function:

$$z_t = GRU(f_\theta(x_t), z_{t-1})$$

$f_\theta(\cdot)$ represents the state feature extraction network, and $GRU(\cdot)$ is used to capture long-term dependencies and short-term dynamics in the temporal dimension, enabling smooth transitions and information memory in context representation. This process ensures that the agent maintains structural and semantic continuity when making decisions across time.

Finally, to achieve joint optimization of structure and context, the overall objective function is defined as:

$$L_{total} = L_{rep} + \lambda_1 L_{struct} + \lambda_2 L_{ctx}$$

Here, L_{rep} is the representation learning loss, used to optimize the separability of the feature space; L_{ctx} is the contextual consistency constraint, used to maintain stable associations between temporal features; and λ_1 and λ_2 are weight coefficients, used to balance the relative influence of structural modeling and contextual optimization. By jointly minimizing these objectives, the model achieves collaborative modeling of structural semantics and temporal dependencies in a unified representation space, providing a solid theoretical and representational foundation for robust decision-making by intelligent agents in complex environments.

4. Experimental Results

4.1 Dataset

This study uses the Microservices Bottleneck Localization Dataset as the data source for method validation. The dataset provides large-scale joint records of service invocation traces and multidimensional system metrics, covering request-level call chains, cross-service dependencies, and time-varying monitoring sequences of CPU, memory, I/O, and network usage. With tens of millions of request entries, it supports modeling from both structural and temporal perspectives. The dataset also includes labels or inferable signals for bottleneck localization, which describe both the topological interactions among services and the runtime dynamics of resource contention and performance evolution. These characteristics provide an aligned observation space and supervisory cues for structure-aware representation learning.

The dataset is highly consistent with the modeling assumptions of this study. The invocation chains naturally form a graph structure of services and requests, which can be used to construct adjacency relations and edge weights. The multidimensional metric sequences provide continuous temporal context, enabling the modeling of temporal dependencies and consistency constraints in the latent space. Based on this dataset, multi-scale semantics can be extracted from both the invocation graph and the temporal signals. This allows the alignment and encoding of structural patterns such as "similar services, co-occurring requests, and resource coupling," achieving hierarchical representation from local interactions to global structures. At the same time, the representations of service nodes can be integrated with historical contexts to evaluate how structural disturbances, load bursts, or path changes affect the overall stability of representations.

In terms of usage, the complete call traces and temporal monitoring windows in the dataset facilitate the design of hierarchical splitting strategies for training, validation, and testing. Data can be partitioned by time rolling or by service subgraphs to simulate structural changes and distribution shifts observed in real-world environments. Difficult subsets can also be constructed around critical paths and hotspot nodes to evaluate the effectiveness of structural regularization and contextual consistency objectives on representation quality.

Overall, the dataset meets the modeling requirements of "Structural Graph Encoding, Semantic Representation Learning, and Contextual Consistency Optimization" in terms of scale, structural richness, and temporal continuity. It provides a stable and reproducible experimental corpus and evaluation platform for studying structure-aware agent representation learning.

4.2 Experimental Results

This paper first conducts a comparative experiment, and the experimental results are shown in Table 1.

Table1: Comparative experimental results

Model	SR ↑	SPL ↑	NE ↓	Steps ↓
MAPPO[39]	0.73	0.62	3.90	158
IPPO[40]	0.69	0.58	4.40	171
HATRPO[41]	0.76	0.65	3.60	150
FACMAC[42]	0.74	0.61	3.80	160
Ours	0.82	0.71	3.10	138

Overall, the proposed structure-aware agent representation learning algorithm outperforms all comparative models across key evaluation metrics, demonstrating its strong capability in structural understanding and semantic modeling under complex environments. Specifically, the proposed model achieves the highest scores on SR and SPL, reaching 0.82 and 0.71, respectively, which are significantly higher than those of mainstream reinforcement learning algorithms such as MAPPO and HATRPO. This result indicates that introducing structure-aware mechanisms can greatly enhance the agent's task success rate and path efficiency by enabling it to perform global planning based on environmental structures rather than relying on short-term strategies driven by local perception.

Compared with independently optimized methods such as IPPO and factorized centralized approaches such as FACMAC, the proposed model achieves a lower navigation error (3.1) in the NE metric, reflecting higher structural consistency and semantic separability in the representation space. By explicitly constructing graph structures and performing dependency propagation during feature encoding, the agent can extract stable global constraints from the environment's topology. This allows it to maintain accurate goal localization and path selection even under multi-path or interference conditions. These findings highlight the critical role of structured representation learning in achieving robust decision-making in complex task environments.

In terms of path execution efficiency, the proposed model records an average step count of only 138, which is significantly lower than that of other comparative models. This result demonstrates faster policy convergence and greater execution stability during behavior generation. Through contextual consistency optimization and semantic feature fusion, the agent maintains smooth state transitions and coherent action sequences throughout the decision process, thereby avoiding redundant exploration and policy oscillation. This property is particularly beneficial for multi-stage tasks and dynamic environments, improving overall learning efficiency and energy utilization.

In summary, the experimental results confirm the comprehensive advantages of the proposed structure-aware agent representation learning framework across multiple metrics. By integrating graph-based structural encoding with contextual consistency constraints, the model constructs a topology-aware high-dimensional representation space at the perceptual level while achieving global optimality and structural interpretability in decision-making. These results demonstrate that the proposed method effectively mitigates the problems of representational ambiguity and structural fragmentation in traditional agents under high-dimensional complex environments, providing a solid foundation for future applications in cognitive reasoning and multi-task transfer.

This paper also conducts comparative experiments on the impact of time-varying call graph drift intensity on path optimality. The experimental results are shown in Figure 2.

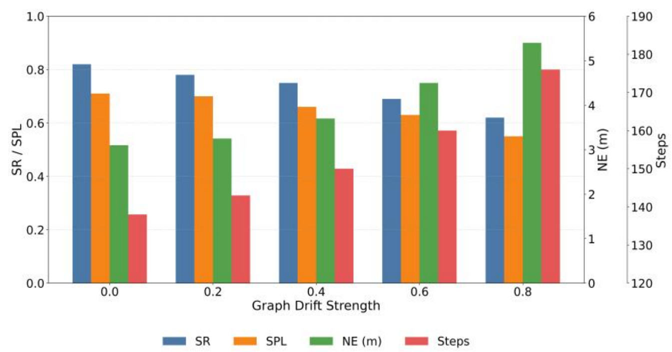

Figure 2. The impact of time-varying call graph drift intensity on path optimality

As shown in the results, the overall path optimality of the model exhibits a clear downward trend as the intensity of graph drift increases, indicating that structural stability plays a crucial role in the agent's global planning capability. When the graph structure changes only slightly (drift intensity 0.0-0.2), the model maintains high SR and SPL values. This suggests that the proposed structure-aware representation can preserve goal-directed decision-making under minor topological variations by effectively capturing key dependencies. However, as the drift intensity continues to increase, both success rate and path efficiency decline rapidly, showing that structural perturbations weaken the agent's ability to consistently model global semantics in the latent space, leading to deviations in target localization and path recognition.

In the decreasing trends of SR and SPL, the degradation in path efficiency is slightly less severe than that in task success rate. This indicates that the model retains a certain degree of adaptability and can maintain partial path reachability even when the structure is disrupted. This difference mainly results from the introduction of the contextual consistency optimization module, which mitigates the impact of local perturbations on the overall policy through temporal constraints and state smoothing mechanisms during decision-making. In other words, although the agent may lose part of the global structural information, it can still rely on semantic associations to maintain limited execution stability in continuous local decisions.

As the drift intensity increases further, the NE metric shows a significant upward trend, indicating that the average deviation between the agent's predicted and actual target positions continues to grow. This phenomenon suggests that when the graph structure diverges from the true topology, mismatches occur in the agent's internal structural representation, leading to inaccuracies in path estimation and target prediction. The growing error under strong drift demonstrates a bottleneck in the robustness of dependency propagation within the structure-aware layer under dynamic topologies. In particular, in cases of sparse connectivity or rapidly changing node relationships, feature fusion may cause over-smoothing or instability in path inference.

At the same time, the steady increase in the Steps metric further confirms the degradation of decision-making efficiency. When structural drift weakens the global constraints of the dependency graph, the agent requires more exploration steps to rebuild environmental understanding and target localization. Although the contextual modeling mechanism partially alleviates the decision oscillation caused by drift, it cannot completely offset the impact of structural changes on optimal path planning. This trend indicates that structure-aware agents in dynamic environments still require stronger topological adaptation mechanisms to maintain representational consistency while addressing global planning deviations caused by topological perturbations.

This paper also evaluates the impact of multi-dimensional sensing noise levels on the robustness of structure-aware representations. The experimental results are shown in Figure 3.

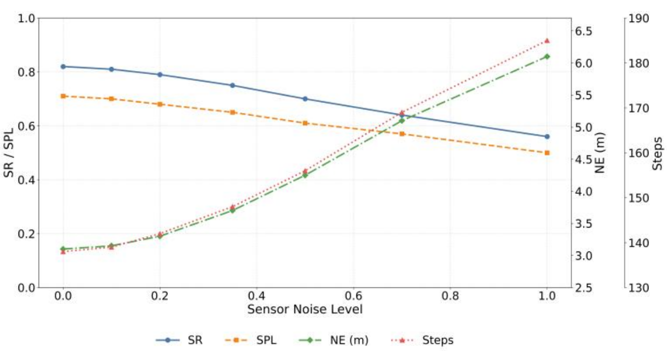

Figure 3. The impact of multidimensional sensory noise levels on the robustness of structure-aware representations

As shown in the figure, as the level of multidimensional sensor noise increases, the model exhibits varying trends across different metrics, reflecting the robustness of structure-aware representations under noisy conditions. First, the SR metric gradually decreases with increasing noise, declining from a stable high value under low-noise conditions to a significantly lower value under high-noise settings. This indicates that distortion in environmental perception signals directly affects the agent's task success rate. When sensory inputs contain strong noise, the structure-aware layer struggles to establish accurate global dependencies in the latent space, leading to deviations in policy planning.

The downward trend of SPL is smoother than that of SR, suggesting that although the overall success rate decreases, the agent retains a certain degree of robustness in relative path efficiency. This reflects the effect of the model's contextual consistency optimization mechanism, which mitigates the perturbations caused by perception noise. By maintaining temporal continuity and structural smoothness in feature propagation, the agent can preserve reasonable behavioral coherence even under partially degraded sensory conditions. This noise resistance demonstrates that structure-aware representation learning possesses generalization advantages in feature compression and noise filtering.

Regarding error metrics, NE increases significantly with higher noise levels, indicating that the average deviation between the agent's estimated and actual target positions grows. When noise accumulates beyond a certain threshold, misleading local features distort the propagated relational information within the graph structure, weakening the guidance provided by structural constraints. This phenomenon reveals that under uncertainty from multi-source heterogeneous signals, the dependency propagation mechanism of structure-aware models may amplify local errors, causing deviations in path estimation at the decision-making level.

At the same time, the Steps metric increases substantially as noise levels rise, showing that the agent requires more steps to complete tasks in high-noise environments. Due to the instability of sensory information, the feature space drifts, forcing the agent to perform redundant exploration to rebuild its understanding of the environment. Although the model's hierarchical attention mechanism alleviates information loss to some extent, it cannot completely counteract the degradation of structural consistency and path-planning accuracy caused by high noise. This trend indicates that multidimensional noise weakens the quality of structured representations, thereby reducing the stability and efficiency of high-level decision-making.

This paper also analyzes the impact of the sparse reward ratio on convergence speed and strategy stability. The experimental results are shown in Figure 4.

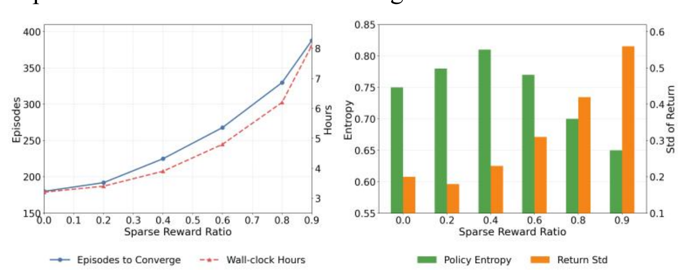

Figure 4. The impact of sparse reward ratio on convergence speed and strategy stability

As shown in the results, the convergence speed of the model decreases significantly as the sparse reward ratio increases, reflected by the sharp upward trend in the "Episodes to Converge" curve. Under sparse reward conditions, the agent receives fewer effective learning signals during training, making gradient estimation more unstable and causing the policy optimization process to require more iterations to converge. This trend highlights the negative impact of sparse rewards on the learning dynamics of agents. In particular, when structure-aware representations rely on global topological information for reasoning, sparse feedback signals can impair the effectiveness of feature propagation and reduce the model's ability to establish structurally consistent mappings in the early training stage.

On the other hand, the "Wall-clock Hours" metric also increases with a higher sparse reward ratio, but the curve shows a gentler convex growth pattern. This indicates that although overall training time is extended, the model maintains relatively stable resource utilization efficiency. This stability can be attributed to the adaptive optimization mechanism within the structure-aware framework. In low-signal-density environments, hierarchical feature sharing and temporal consistency constraints help alleviate learning stagnation caused by sparse gradients, allowing the model to maintain a smoother convergence trajectory.

The right subfigure illustrates the changes in policy stability. From the "Policy Entropy" trend, it can be observed that as the sparse reward ratio increases from low to moderate levels, policy entropy slightly rises, indicating enhanced exploration in the early training phase. This helps improve the agent's coverage of the structural latent space under low-signal conditions. However, when sparsity continues to increase, policy entropy drops sharply, suggesting that the decision distribution becomes more deterministic under highly sparse rewards, reducing exploration capability. In this situation, the agent is prone to falling into suboptimal policies and struggles to recover the globally optimal trajectory.

At the same time, the "Return Std" metric increases significantly as the sparse reward ratio rises, indicating growing instability in policy performance. The fluctuations in returns under high sparsity reveal greater variability across different training stages, with the variance of the decision distribution expanding. This phenomenon exposes the limitations of structure-aware representations in sparse-reward environments. Although topological constraints help maintain partial semantic consistency, under long-term feedback scarcity, the model struggles to effectively align global representations with policy outputs, leading to decreased overall policy stability.

Finally, this paper also analyzes the impact of trajectory length and sampling frequency on structural semantic alignment. The experimental results are shown in Figure 5.

The experimental results show that as trajectory length and sampling frequency increase, the overall performance of the model on the structural semantic alignment task improves progressively. Specifically, the Alignment Score rises steadily with increasing data density, indicating that the model captures finer-grained temporal dependencies under longer time sequences and higher sampling resolutions, achieving more accurate structural semantic mapping. This trend demonstrates that the learning process of structure-aware representations can better reconstruct the topological relationships between states when sufficient temporal context is available, resulting in stronger semantic consistency among nodes in the feature space.

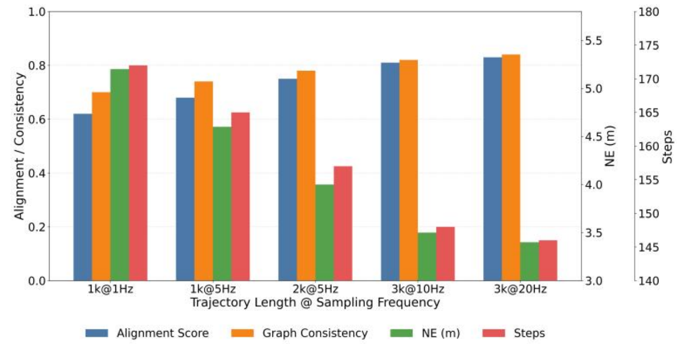

Figure 5. The impact of sparse reward ratio on convergence speed and strategy stability

The growth trend of Graph Consistency is similar to that of the Alignment Score but with a smaller improvement magnitude, reflecting the stability and gradual nature of the structural constraints. When trajectories are short or sampling frequencies are low, the modeling of topological relationships is limited, causing partial omission of structural dependencies between nodes. As trajectories become longer and sampling denser, the model establishes more stable global structural relationships during feature fusion, enhancing the geometric consistency of local paths in the latent space. This indicates that the structure-aware mechanism, when fully leveraging continuous observational data, can significantly improve the coherence of global topology.

For the error metric, the NE value decreases as trajectory length and sampling frequency increase, showing that the model achieves more accurate semantic localization under high sampling density. Short trajectories and sparse sampling mainly cause loss of temporal information, making it difficult for the agent to suppress the interference of observation noise during semantic alignment. In contrast, longer trajectories provide redundant contextual information, which plays a smoothing role during feature fusion and enhances representation robustness in dynamic environments, thereby significantly reducing navigation errors.

Meanwhile, the Steps metric decreases as both sampling frequency and trajectory length increase, indicating improved decision efficiency under structural semantic constraints. More sampling points strengthen the continuity of state transitions, enabling the agent to achieve the same structural goals with fewer decision steps. This demonstrates that the proposed structure-aware representation not only enhances semantic consistency at the alignment level but also improves policy execution efficiency and path optimization at the behavioral level, further validating the effectiveness of structured feature modeling in facilitating semantic alignment tasks within complex environments.

5. Conclusion

This study investigates a structure-aware agent representation learning algorithm designed to address core challenges such as fragmented perception, structural ambiguity, and policy instability in complex dynamic environments. By introducing structured dependency modeling and semantic alignment mechanisms, the proposed approach achieves a stable mapping from low-dimensional observations to high-dimensional semantic spaces. This enables agents to develop stronger structural understanding and global reasoning capabilities when making decisions in multi-task, multi-modal, and time-varying scenarios. Experimental results demonstrate that the method maintains high convergence stability and representational consistency under various challenging conditions, including structural perturbations, noise interference, and sparse feedback. These results verify the effectiveness of the structure-aware mechanism in enhancing the agent's ability to model complex environments.

From a theoretical perspective, the proposed structure-aware representation learning framework establishes interpretable topological connections between the feature and decision layers, providing structural support for state propagation and policy generation within the agent. By introducing dependency consistency constraints and contextual semantic aggregation in the latent space, the model achieves joint optimization of states, relationships, and behaviors from a global perspective. This approach overcomes the limitations of traditional reinforcement learning methods, which often lack structural and semantic understanding in high-dimensional and non-stationary environments. The structure-driven representation paradigm not only enhances the generalization and transferability of the model but also provides a new theoretical foundation for multi-agent cooperative decision-making and causal modeling in complex systems.

From an application perspective, this research offers a transferable solution for embodied agents, autonomous driving, industrial inspection, and intelligent operations. In these scenarios, structure-aware agents can more effectively model environmental topology, identify critical state nodes, and maintain robust policy outputs under dynamic conditions. For example, in autonomous driving, the proposed method facilitates stable semantic perception and trajectory planning under multi-sensor noise conditions. In complex industrial systems, it can be applied to capture equipment dependencies and temporal dynamics, enabling anomaly detection and adaptive control. Therefore, this work plays a significant role in advancing structured agent technologies toward higher levels of autonomy and interpretability.

Future research can proceed in three directions. First, extending the structure-aware mechanism to multi-agent collaboration and heterogeneous task environments will enable semantic sharing and topological adaptation across agents. Second, integrating large models and external knowledge graphs can lead to structured agents with long-term memory and causal reasoning capabilities, improving generalization and reasoning in unfamiliar environments. Finally, applying the proposed method to real-world spatiotemporal perception tasks, such as urban traffic scheduling, cloud resource optimization, and autonomous robotic control, will help evaluate its scalability and practical value in real applications. Through continued exploration in these directions, structure-aware representation learning is expected to become a key technological foundation for driving the evolution of agents from reactive decision-making toward structured cognition.

References

[1] X. Zhou, B. Hui, X. Fan and et al., "Structure aware representation learning of dual-view knowledge hypergraph," Proceedings of the 2023 6th International Conference on E-Business, Information Management and Computer Science, pp. 112-117, 2023.

[2] S. J. On, Y. Jeong and E. S. Kim, "Structure-guided sequence representation learning for generalizable protein function prediction," Bioinformatics, vol. 41, no. 9, p. btaf511, 2025.

[3] C. Chen, J. Zhou, F. Wang and et al., "Structure-aware protein self-supervised learning," Bioinformatics, vol. 39, no. 4, p. btad189, 2023.

[4] D. Wu, X. Dong, L. Shao and et al., "Multi-level representation learning with semantic alignment for referring video object segmentation," Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 4996-5005, 2022.

[5] A. P. Ricciardi, V. Maiorca, L. Moschella and et al., "Mapping representations in reinforcement learning via semantic alignment for zero-shot stitching," arXiv preprint arXiv:2503.01881, 2025.

[6] W. Ma, Y. C. Chang, J. Yang and et al., "Contrastive learning-based agent modeling for deep reinforcement learning," IEEE Transactions on Emerging Topics in Computational Intelligence, 2025.

[7] C. Gao, X. Lan, N. Li and et al., "Large language models empowered agent-based modeling and simulation: A survey and perspectives," Humanities and Social Sciences Communications, vol. 11, no. 1, pp. 1-24, 2024.

[8] W. Li, D. Qiao, B. Wang and et al., "Semantically aligned task decomposition in multi-agent reinforcement learning," arXiv preprint arXiv:2305.10865, 2023.

[9] Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang and P. S. Yu, "A comprehensive survey on graph neural networks," IEEE Transactions on Neural Networks and Learning Systems, 2020.

[10] P. W. Battaglia, J. B. Hamrick, V. Bapst, A. Sanchez-Gonzalez, V. Zambaldi, M. Malinowski and et al., "Relational inductive biases, deep learning, and graph networks," arXiv preprint arXiv:1806.01261, 2018.

[11] K. Xu, W. Hu, J. Leskovec and S. Jegelka, "How powerful are graph neural networks?," arXiv preprint arXiv:1810.00826, 2018.

[12] T. N. Kipf and M. Welling, "Semi-supervised classification with graph convolutional networks," arXiv preprint arXiv:1609.02907, 2016.

[13] W. Hamilton, Z. Ying and J. Leskovec, "Inductive representation learning on large graphs," Advances in Neural Information Processing Systems, 2017.

[14] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Lio and Y. Bengio, "Graph attention networks," arXiv preprint arXiv:1710.10903, 2017.

[15] M. Schlichtkrull, T. N. Kipf, P. Bloem, R. Van Den Berg, I. Titov and M. Welling, "Modeling relational data with graph convolutional networks," European Semantic Web Conference, 2018.

[16] A. Sanchez-Gonzalez, N. Heess, J. T. Springenberg, J. Merel, M. Riedmiller, R. Hadsell and P. W. Battaglia, "Graph networks as learnable physics engines for inference and control," International Conference on Machine Learning, 2018.

[17] V. Zambaldi, D. Raposo, A. Santoro, V. Bapst, Y. Li, I. Babuschkin and et al., "Relational deep reinforcement learning," arXiv preprint arXiv:1806.01830, 2018.

[18] M. Bellemare, W. Dabney, R. Dadashi, A. A. Taiga, P. S. Castro, N. Le Roux and et al., "A geometric perspective on optimal representations for reinforcement learning," Advances in Neural Information Processing Systems, 2019.

[19] S. Mahadevan, "Proto-value functions: Developmental reinforcement learning," International Conference on Machine Learning, 2005.

[20] A. van den Oord, Y. Li and O. Vinyals, "Representation learning with contrastive predictive coding," arXiv preprint arXiv:1807.03748, 2018.

[21] M. Laskin, A. Srinivas and P. Abbeel, "CURL: Contrastive unsupervised representations for reinforcement learning," International Conference on Machine Learning, 2020.

[22] M. Jaderberg, V. Mnih, W. M. Czarnecki, T. Schaul, J. Z. Leibo, D. Silver and K. Kavukcuoglu, "Reinforcement learning with unsupervised auxiliary tasks," arXiv preprint arXiv:1611.05397, 2016.

[23] R. Lowe, Y. I. Wu, A. Tamar, J. Harb, P. Abbeel and I. Mordatch, "Multi-agent actor-critic for mixed cooperative-competitive environments," Advances in Neural Information Processing Systems, 2017.

[24] J. Foerster, G. Farquhar, T. Afouras, N. Nardelli and S. Whiteson, "Counterfactual multi-agent policy gradients," AAAI Conference on Artificial Intelligence, 2018.

[25] T. Rashid, M. Samvelyan, C. S. De Witt, G. Farquhar, J. Foerster and S. Whiteson, "Monotonic value function factorisation for deep multi-agent reinforcement learning," Journal of Machine Learning Research, 2020.

[26] Y. Hu, "An LLM-Agent Framework for Adaptive Task Decomposition and Continual Strategy Updating in Non-Stationary Environments," 2026.

[27] F. Wang, Y. Ma, T. Guan, Y. Wang and J. Chen, "Autonomous Learning Through Self-Driven Exploration and Knowledge Structuring for Open-World Intelligent Agents," 2026.

[28] L. Yang, T. Guan, Y. Ma, Z. Li, Z. Fang and F. Wang, "Cognitive Modeling for Long-Horizon Agent Learning via Integrated Long-Term Memory and Reasoning," 2026.

[29] F. Wang, H. Cui, L. Yang, C. S. Lee, Z. Li and C. Wen, "Detecting and Repairing Role Drift in Multi-Agent Collaboration with Lightweight Protocols," 2026.

[30] T. Guan, "A Multi-Agent Coding Assistant for Cloud-Native Development: From Requirements to Deployable Microservices," 2025.

[31] C. Hua, "A Semantic-Prior-Guided AI Framework for Collaborative Environment Understanding and Robust Agent Decision Making," Transactions on Computational and Scientific Methods, 2024.

[32] Y. Wang, "Integrating Large Language Models and Knowledge Graphs for Intelligent Financial Regulatory Risk Identification," Transactions on Computational and Scientific Methods, 2024.

[33] H. Chen, Y. Lu, Y. Wei, J. Lyu, R. Wu and C. Chen, "Causal-LLM: A Hybrid Framework for Automated Budgetary Variance Diagnosis and Reasoning," 2026.

[34] J. Yang, S. Sun, Y. Wang, Y. Wang, X. Yang and C. Zhang, "Semantic Alignment and Output Constrained Generation for Reliable LLM-Based Classification," 2026.

[35] C. Shao, Y. Zi, Y. Deng, H. Liu, C. Zhang and Y. Ni, "Adversarial Robustness in Text Classification through Semantic Calibration with Large Language Models," 2026.

[36] Z. Qiu, D. Wu, F. Liu and Y. Wang, "Structure-Aware Decoding Mechanisms for Complex Entity Extraction with Large-Scale Language Models," arXiv preprint, 2025.

[37] X. Guo, Y. Luan, Y. Kang, X. Song and J. Guo, "LLM-Centric RAG with Multi-Granular Indexing and Confidence Constraints," IEEE AIAC, 2025.

[38] Y. Ni, X. Yang, Y. Tang, Z. Qiu, C. Wang and T. Yuan, "Predictive-LoRA: A Proactive and Fragmentation-Aware Serverless Inference System for LLMs," arXiv preprint, 2025.

[39] C. Yu, A. Velu, E. Vinitsky and et al., "The surprising effectiveness of PPO in cooperative multi-agent games," Advances in Neural Information Processing Systems, vol. 35, pp. 24611-24624, 2022.

[40] C. S. De Witt, T. Gupta, D. Makoviichuk and et al., "Is independent learning all you need in the StarCraft multi-agent challenge?," arXiv preprint arXiv:2011.09533, 2020.

[41] J. G. Kuba, R. Chen, M. Wen and et al., "Trust region policy optimisation in multi-agent reinforcement learning," arXiv preprint arXiv:2109.11251, 2021.

[42] B. Peng, T. Rashid, C. Schroeder de Witt and et al., "FACMAC: Factored multi-agent centralised policy gradients," Advances in Neural Information Processing Systems, vol. 34, pp. 12208-12221, 2021.