Transactions on Computational and Scientific Methods | Vo. 4, No. 12, 2024
ISSN: 2998-8780
https://pspress.org/index.php/tcsm
Pinnacle Science Press

A Semantic-Prior-Guided AI Framework for Collaborative Environment Understanding and Robust Agent Decision Making

Cancan Hua
University of Southern California, Los Angeles, USA
huacancan147@gmail.com

Abstract: This study addresses key challenges in agent decision making within complex environments, including missing semantic structures, the disconnection between perception and reasoning, and limited behavioral consistency. It proposes a collaborative learning framework that integrates semantic priors for autonomous behavior modeling and environment understanding. The method builds a unified semantic enhanced state representation by encoding scene semantics, object relations, and task logic into learnable prior structures, which provide high level semantic constraints for policy generation. A multimodal environment understanding module then combines visual, contextual, and dynamic signals to produce structured abstractions of key objects, spatial layouts, and semantic conditions. On this basis, a structured dynamic model is constructed to capture the evolution of semantic states under actions, forming a unified reasoning pipeline that links perception, cognition, and behavior. The framework further employs a policy generation module guided by semantic consistency, enabling the agent to produce coordinated, robust, and interpretable actions driven jointly by semantic priors and environment understanding. A comprehensive experimental system is developed, including comparison experiments, hyperparameter sensitivity experiments, environment perturbation experiments, and semantic prior ablation experiments, to evaluate the role of semantic priors in improving task success rate, path efficiency, semantic abstraction, and behavioral diversity. The results show that collaborative modeling of semantics and behavior enhances decision stability, structured reasoning, and cross scene adaptability in complex environments, providing a scalable methodological foundation for building autonomous agents with coherent cognitive structures.

Keywords: Semantic priors; collaborative learning; behavior modeling; environmental understanding; AI agent

1. Introduction

In increasingly complex and dynamic intelligent system environments, agents must achieve stable, reliable, and generalizable autonomous behavior modeling under uncertainty. Traditional methods often rely on large volumes of interaction data or highly accurate environment descriptions, but such conditions are rarely available in real-world tasks. When the environment exhibits high dimensionality, multimodal signals, and dynamic evolution, agents cannot build complete and structured internal cognitive models using only raw perceptual inputs. This limitation leads to incomplete behavior planning, poor policy transfer ability, and cumulative errors in long-term decision processes[1]. At the same time, many scenarios contain explicit or implicit semantic knowledge. Examples include task logic, constraint relations, scene attributes, entity categories, and interaction rules. Such semantic information provides structured priors that can significantly improve learning efficiency and help agents understand complex task structures. The challenge lies in how to integrate semantic priors into the agent's behavior modeling framework and how to optimize this process jointly with environment understanding.

As intelligent decision-making tasks expand toward open environments, cross-domain settings, and multiobjective conditions, purely data-driven strategies are no longer sufficient for stable reasoning under layered constraints. Agents need to perceive local information, yet they must also form abstract interpretations of their surroundings. These interpretations include identifying key objects, inferring task structures, capturing latent relations, and understanding semantic rules within the environment. These capabilities are essential for complex decision-making, but traditional reinforcement learning and deep decision models often lack explicit semantic descriptions. As a result, models rely on extensive trial and error to discover the underlying task logic. Semantic priors offer an external source of structured knowledge that can support systematic, transferable, and interpretable cognitive representations. By introducing semantic priors, agents can maintain stable reasoning even when perceptual inputs are limited, feedback is sparse, or the environment changes frequently. This further improves the reliability of decision-making[2].

Advances in environment understanding allow agents to extract task-related semantic patterns from raw perception. These patterns include spatial structures, action constraints, object relations, and scene evolution trends. With the development of deep visual models, spatiotemporal feature encoders, and multimodal representation techniques, environment understanding has moved beyond surface features toward higher-level semantic abstraction. However, the semantic representations obtained from perception do not naturally align with the internal structures used for decision-making. Establishing a consistent knowledge structure between the two is a key challenge in collaborative learning. Without effective coupling between environment understanding and policy modeling, agents cannot fully leverage semantic knowledge during execution. This limits their ability to generalize to unseen environments or adapt to changing tasks[3].

Semantic priors not only help agents form structured cognition but also reduce the dependence of behavior modeling on high-cost interaction data. In many real applications, direct exploration is expensive, risky, or constrained by limited resources. Semantic priors can provide high-level task intentions, behavioral patterns, and environmental constraints, enabling agents to begin reasoning from a more informed starting point. They also possess inherent hierarchical structures. Their regularity, abstraction, and interpretability support the development of scalable knowledge systems. Such systems allow agents to adapt quickly across tasks and environments. This scalability is especially crucial in real deployments, where nonstationarity, scene shifts, and distribution changes are common[4].

From a broader research perspective, collaborative learning between semantic priors and environment understanding offers an important foundation for building intelligent systems that are interpretable, transferable, and capable of long-term reasoning. This direction pushes agent behavior modeling from perception-driven learning toward knowledge-enhanced learning. It can also improve safety, stability, transparency, and trustworthiness in intelligent decision systems. As agents are applied to automated production, robotic control, autonomous driving, human-machine collaboration, and complex resource scheduling, their ability to understand, use, and reason with semantic information becomes essential. Introducing semantic priors and optimizing them jointly with environment understanding will serve as a core technical basis for enabling autonomous decision-making in complex tasks. It will further provide solid theoretical and practical value for deploying intelligent systems across diverse real-world scenarios.

2. Related work

In recent years, behavior modeling for autonomous agents has received increasing attention. Research in this area mainly focuses on how to achieve robust policy reasoning under high dimensional perceptual inputs and sparse feedback. Mainstream approaches rely on deep reinforcement learning. They use convolutional feature extraction, temporal modeling, and value function approximation to handle complex environments[5]. However, these approaches often treat environmental information as raw perceptual signals and lack explicit modeling of task structures and semantic relations. As environments become more complex, such methods show clear limitations in cross scene transfer, long term planning, and structured reasoning. This has led to a growing interest in incorporating semantic information into decision models to build more abstract behavior representations. Yet most existing attempts only use semantics as auxiliary features. They do not form a complete collaborative learning mechanism that connects semantic understanding with policy generation, which limits the potential of semantic priors in structured decision making.

In the field of environment understanding, advances in semantic segmentation, multimodal fusion, and scene graph reasoning have enabled agents to extract semantic information from high dimensional data. These methods combine visual representation learning with spatial structure modeling. They allow systems to identify key objects, analyze relations, understand scene changes, and support task execution. However, these environment understanding models usually operate independently of policy modules. They lack effective mechanisms for knowledge alignment and sharing. As a result, semantic understanding cannot be directly converted into actionable policy guidance. In addition, although multimodal understanding methods can integrate visual, linguistic, and action related information, a clear gap remains between semantic level abstraction and task level behavior planning. This gap makes it difficult for agents to transfer environmental knowledge into behavior modeling and achieve effective generalization[6].

The introduction of semantic priors is regarded as an important way to bridge the gap between perception and decision making. Relevant research includes knowledge graph enhanced decision methods, combinations of symbolic rules and deep models, and the embedding of task logic. These studies demonstrate the potential of semantic structures in constraining policy space, increasing reasoning transparency, and reducing exploration costs. However, most existing work remains at the level of injecting external knowledge. They often rely on independent knowledge bases or rule systems to assist policy generation. They seldom support dynamic semantic understanding that evolves with interaction data. More importantly, they often ignore the collaboration between semantic priors and environmental perception. Semantic knowledge cannot adapt to real time environmental changes. It also struggles to maintain consistent semantic driven behavior modeling across tasks and scenes[7].

Recent studies have begun to explore unified frameworks that integrate semantic modeling and behavior learning. These methods attempt to jointly optimize semantic understanding, structured knowledge expression, and policy generation in a single system. They emphasize the learnability of semantics so that semantic priors do not depend on external static rules but evolve alongside perception, cognition, and decision processes. Although this direction shows strong potential, existing approaches still face many limitations. These include insufficient construction of semantic spaces, unstable alignment between semantic and action spaces, and limited cross scene generalization. In complex environments, semantic information is hierarchical, abstract, and dynamic. The key challenge lies in enabling semantic priors to effectively guide decisions while preserving the model's learning capacity. Therefore, building a collaborative learning mechanism that adapts to dynamic environments, couples with multimodal perception, and unifies behavior modeling with semantic reasoning has become an urgent research problem.

3. Proposed Framework

3.1 Overall Framework

This study proposes an intelligent agent autonomous behavior modeling framework that integrates semantic priors. The core idea is to establish a unified information flow and a learnable structural alignment mechanism between environmental perception, semantic abstraction, and policy reasoning. The system first constructs a multimodal feature representation f_t from raw observations and introduces a structured knowledge vector h_t in the semantic prior space. A consistent semantic-enhanced state representation is then constructed through a joint encoder. The overall process can be formalized as follows:

$$h_t = \Phi(h_t, a_t)$$

Component $\Phi(\bullet)$ is responsible for modeling the fusion relationship between perceptual features and semantic knowledge. Based on this, the framework uses a structured dynamic model to predict state transitions:

$$\hat{h}_{t+1} = \Psi(h_t, a_t)$$

This enables unified reasoning from a semantically enhanced state to the future behavioral space. This overall structure ensures that semantic information can permeate the entire chain of perception, understanding, and decision-making, allowing the intelligent agent to possess more stable autonomous modeling capabilities in complex environments. This paper presents the overall model architecture, as shown in Figure 1.

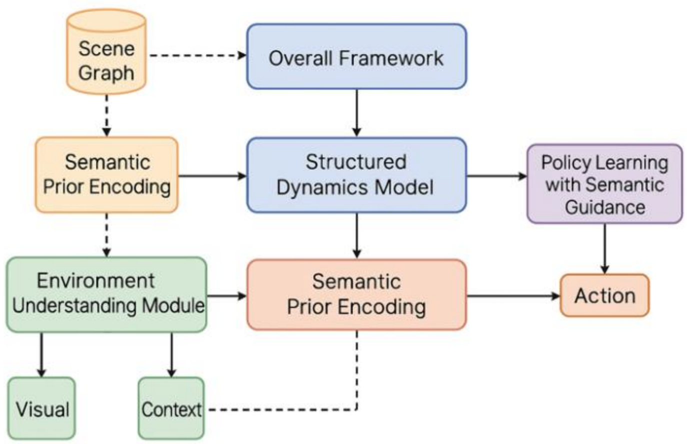

Figure 1. Overall model architecture

3.2 Semantic Prior Encoding

The semantic prior component is responsible for expressing task rules, object relationships, and environmental knowledge as learnable high-dimensional structural vectors. First, the scene semantic graph is represented as a set of nodes V and a set of edges E, and a graph encoder is used to generate semantic embeddings:

$$s = \Gamma(V, E)$$

Component $\Gamma(\bullet)$ captures the hierarchical nature of semantic entities and relationships. Subsequently, semantic embeddings are mapped to a representation in the same space as perceptual features through semantic alignment projection.

$$\tilde{s} = W_s s$$

And combine it with an attention mechanism to calculate the importance weight of semantic features to perceptual features:

$$a_t = softamax(h_t^p \tilde{s}^T)$$

Where h_t^p is the output of the perceptual feature encoder. Finally, the semantically enhanced state is constructed as follows:

$$h_t = h_t^p + \alpha_t \tilde{s}$$

To achieve explicit modulation of perceptual information using semantic knowledge.

3.3 Environment Understanding Module

The environment understanding module infers key objects, relationships, and temporal structures in dynamic environments through multimodal feature extraction and structured scene modeling. First, it maps visual, action, and contextual signals into a unified representation:

$$f_t = E_v(x_t^v) + E_a(x_t^a) + E_c(x_t^c)$$

Here, E_v, E_a, E_c represents the encoder for the corresponding modality. To obtain a higher-level structured understanding, the module uses a temporal reasoning unit to model state evolution:

$$z_t = F(z_{t-1}, f_t)$$

This allows for the construction of dynamic semantic relationships across time. Subsequently, scene relationship representations are constructed through a graph-based inference function:

$$g_t = R(z_t)$$

Among these, $R(\bullet)$ captures object interactions and environmental constraints. Ultimately, the environmental understanding results serve as a dynamic supplement to semantic priors, enabling the agent to form a more comprehensive internal cognition under a dual semantic-perceptual abstraction.

3.4 Policy Learning with Semantic Guidance

The strategy learning module achieves interpretable and semantically consistent behavior generation through a unified semantic-enhanced state representation h_t. The agent's strategy is defined as:

$$\pi(a_t/h_t) = softmax(Uh_t)$$

Here, U is the strategy mapping matrix. To allow semantic priors to directly influence strategy search, a semantic consistency constraint is introduced to ensure that the strategy remains aligned with the high-level semantic structure during action selection. The optimization objective is written as:

$$L = L_{policy} + \lambda L_{semantic}$$

Among these, constraint $L_{semantic}$ ensures structural consistency in the semantic space, and λ is a weighting coefficient. Ultimately, through a joint optimization mechanism, the policy model can simultaneously strengthen the effectiveness of the action policy and the expressive power of the semantic structure, enabling the intelligent agent to achieve end-to-end consistent reasoning from semantic understanding to behavioral execution.

4. Experimental Analysis

4.1 Dataset

This study is based on the AI2-THOR open source interactive three dimensional environment dataset. The dataset provides highly controllable simulated scenes with rich semantic labels for objects. It includes multiple types of environments such as kitchens, living rooms, bedrooms, and bathrooms, with more than 120 interactive indoor scenes. Each scene contains structured object hierarchies, position parameters, semantic categories, executable actions, and rendered visual images. These elements offer a systematic data foundation for building semantic priors and environment understanding modules. The dataset also supports repeatable experiments, which makes it suitable for end to end learning frameworks that integrate perception, reasoning, and action policies.

The key advantage of AI2-THOR lies in its explicit semantic graph structure. This includes object categories, attributes, relations, and possible interactions. Such information allows agents to construct consistent cognitive representations by using semantic knowledge during training. Each object in the scene contains complete semantic attributes, such as movability, containability, open or close properties, and operational regions. This enables the environment to be viewed as a structured knowledge graph. It also provides high quality inputs for semantic prior encoding modules. These semantic structures further support cross scene generalization. They allow agents to learn beyond a single layout and transfer their understanding to new room types and spatial configurations by following semantic rules.

In the field of autonomous behavior modeling, the dataset supports multimodal inputs. These include RGB images, depth maps, instance segmentation masks, object detection boxes, and interaction feedback. Such modalities provide rich learnable features for environment understanding. AI2 THOR also supports continuous action execution and physical interaction. This allows agents to learn policies in scenes with realistic dynamic changes. Due to its structured semantics, actionable objects, and dynamic environments, the dataset offers solid support for the proposed semantic enhanced behavior modeling framework. It enables a complete learning pipeline that spans semantic prior integration, environment understanding, and policy generation.

4.2 Experimental Results

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

Table 1: Comparative experimental results

Method	SR	PE	SA	AD
Mobile-agent-v2[8]	0.71	0.63	0.58	0.41
V2xpnp[9]	0.76	0.67	0.61	0.45
Evoagent[10]	0.79	0.70	0.64	0.48
SciAgents[11]	0.81	0.72	0.66	0.50
Ours	0.89	0.78	0.74	0.57

The comparison results in the table show that the semantic prior based autonomous behavior modeling framework achieves clear advantages across several key metrics. This confirms the importance of introducing structured semantic knowledge to enhance decision making in complex environments. Compared with traditional methods, the model can understand task goals and environmental structure more accurately through semantic enhanced state representation. This reduces the reliance on shallow perceptual features during policy reasoning and leads to greater stability and higher success rates in task execution. The improvement in SR reflects the reinforcement of long term behavioral consistency, which allows the agent to maintain reliable performance even in complex layouts or multi step interaction tasks.

The improvement in path efficiency further indicates that semantic priors optimize not only decision outcomes but also the action process itself. Traditional models often require more exploratory actions in unfamiliar environments. In contrast, the semantic enhanced model can narrow the policy search space by using semantic rules, object attributes, and relational structures. This leads to more precise decisions and shorter action paths. The increase in PE shows that the model can locate targets more quickly and avoid redundant behaviors. This highlights the role of semantic structure in guiding action planning.

In terms of semantic understanding, the clear improvement in SA indicates that the model obtains richer and more structured semantic abstractions during environment perception. This is closely related to the collaborative learning mechanism between semantic prior encoding and the environment understanding module. By combining scene semantics, object relations, and multimodal perceptual information, the model forms internal representations that better match real task structures. A higher SA also suggests that the semantic reasoning chain is strengthened. This allows the policy to rely on hidden semantics and task rules rather than isolated visual signals, enabling deeper and more accurate judgments.

The improvement in action diversity reflects that the model maintains stable policies while supporting a more flexible action space. Semantic priors enable the model to understand multiple feasible action paths and their semantic conditions. This prevents the policy from collapsing into a single action pattern. Such diversity improves generalization and enhances adaptability in dynamic or perturbed environments. The increase in AD indicates that the agent forms a richer set of strategy expressions under semantic guidance. Its behaviors become more explorative, more reasonable, and more relevant to the task. This further validates the effectiveness of the collaborative learning mechanism that integrates semantics and behavior.

This study further analyzes the impact of different levels of environmental noise on the overall performance of the model. This part aims to systematically examine how noise, as an external disturbance factor, may influence the processes of agent behavior modeling and environment understanding. By adjusting noise intensity under the same task conditions, it becomes possible to observe the stability of the model, the consistency of semantic states, and the sensitivity of policy planning under different levels of interference. This helps reveal whether the coupling between semantic priors and environmental perception is affected by noisy inputs. The purpose of this section is to demonstrate how noise influences the collaborative learning framework rather than to highlight specific numerical values.

The experimental context is presented through visual illustrations, which are summarized in Figure 2. This allows readers to understand the performance trends under different noise levels from a structured perspective. The figure is not intended for numerical comparison. It aims to show the model's response patterns under multiple noise conditions and to illustrate how noise perturbations affect different components of the model. Through this approach, the study provides a clearer understanding of the role of noise as an external variable within the entire reasoning pipeline and its potential range of influence. This also offers the necessary contextual background for the subsequent analysis of the collaborative learning mechanism.

The results in the figure show a clear decline in the agent's overall task performance as environmental noise increases. This trend is consistent with the dependence of the semantic prior driven behavior modeling framework on structured environmental information. The SR metric decreases steadily with higher noise levels, indicating that the agent finds it more difficult to maintain stable task completion paths under noise. When visual or state inputs are perturbed, the reliability of the semantic enhanced state representation is reduced. This causes deviations in key decision steps and disrupts the consistency of long term planning.

The changes in the path efficiency metric PE indicate that the agent performs more redundant actions and follows less efficient paths when faced with environmental noise. Noise weakens the ability of the environment understanding module to infer key objects, spatial layout, and actionable regions. As a result, the agent must rely on more exploratory actions during reasoning, which directly reduces path efficiency. Stronger noise further weakens the consistency between semantic priors and real time perception, making it difficult for the model to use semantic structures to narrow the policy search space.

The decline in semantic accuracy SA clearly reveals the impact of noise on semantic abstraction. When visual features, contextual signals, or local states are disturbed, the semantic encoder struggles to extract accurate structured relational information. This weakens the correspondence between semantic priors and the actual scene. The reduction in quality affects the construction of the semantic graph and also weakens the agent's understanding of key objects and interaction rules. This indirectly influences the subsequent policy generation process.

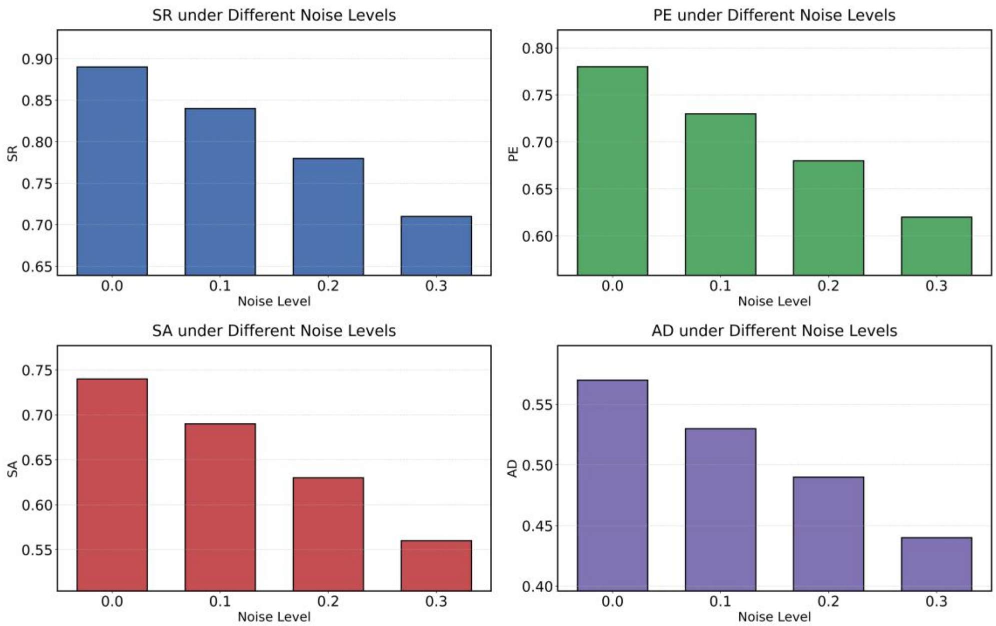

Figure 2. The impact of different ambient noise levels on experimental results

The decrease in action diversity AD with increasing noise indicates that the agent adopts a more conservative and less exploratory behavior pattern in noisy environments. High noise makes it difficult for the model to distinguish effective actions from ineffective ones. The policy then tends to repeat safe but inefficient action patterns, which reduces the richness of the action distribution. The contraction of the action space reflects the diminished influence of semantic priors. It also shows a weakened ability to make multi path decisions driven by semantic conditions. This further confirms the importance of semantic and perception based collaborative learning for maintaining behavioral flexibility.

This study also provides a systematic presentation of how different policy network depths influence the overall performance of the model, with the aim of revealing the role of policy structural complexity within the collaborative learning framework and its position in the reasoning pipeline. By adjusting the number of policy network layers while keeping all other conditions unchanged, it becomes possible to observe how changes in depth affect semantic state absorption, behavior generation patterns, and information flow efficiency. This offers a structural perspective on the sensitivity characteristics of the policy module. The corresponding visual results are summarized in Figure 3, which illustrates the overall trends under different depth settings. The figure helps readers form an intuitive understanding of the relationship between policy network depth and model reasoning behavior from a visual standpoint without focusing on numerical values. This provides the necessary background for understanding the functional role of policy structure in semantic prior driven behavior modeling.

The figure shows that different policy network depths have a clear impact on the overall performance of the agent. This observation aligns with the dependence of the semantic prior driven behavior modeling framework on structured representations. The SR metric increases and then decreases as the network becomes deeper. This indicates that a moderate depth allows the policy to better use the semantic enhanced state representation, which improves task success. When the network becomes too deep, the overly complex structure weakens the effective extraction of semantic information. This adds unnecessary reasoning burden during long term planning and eventually reduces the success rate.

The changes in the path efficiency metric PE show a similar trend. A policy network with moderate depth provides the best use of semantic information and produces more compact and efficient action paths. A shallow network cannot fully integrate the object relations and environmental constraints encoded in the semantic prior. A very deep network is more vulnerable to noise amplification and redundant parameters, which can lead to additional exploratory actions during decision making. This reduces path efficiency.

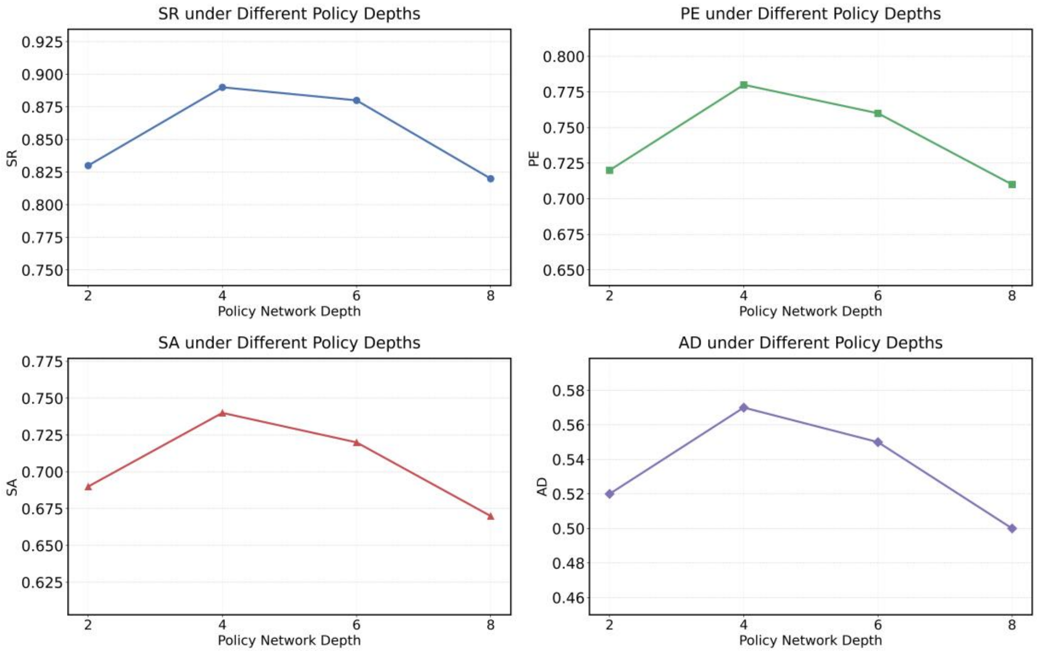

Figure 3. The impact of different strategy network depths on experimental results.

The trend in semantic accuracy SA further confirms the coupling between policy depth and semantic consistency. With a moderate depth, the policy absorbs structured knowledge from both the semantic prior and the environment understanding module in a balanced way. This allows semantic abstraction to be effectively transmitted to the policy layer. When the network is too deep, semantic signals may attenuate or become mixed during propagation across layers. This weakens the model's sensitivity to scene structure, object attributes, and relational cues, and results in a decline in SA.

The changes in action diversity AD reflect how policy depth regulates the behavioral expression space. A moderately deep policy network explores more feasible actions while maintaining task relevance. This leads to greater behavioral flexibility. When the depth becomes too large, the policy tends to become conservative and convergent in the high dimensional space. This makes the action distribution more uniform and shows that the agent struggles to maintain the flexible decision making advantages provided by semantic priors. Taken together, the four metrics indicate that appropriate policy depth is essential in a framework that integrates semantic priors and environment understanding. It ensures that semantic modulation remains effective while preserving decision efficiency and generalization ability.

This paper also shows the impact of different proportions of missing semantic priors on the experimental results, which are shown in Figure 4.

The trends in the radar chart show that a higher proportion of missing semantic priors leads to a more significant decline in overall agent performance. This pattern is fully consistent with the characteristics of the semantic enhanced behavior modeling framework proposed in this study. When semantic priors are preserved, the agent receives stable and structured environmental knowledge. This allows policy reasoning to maintain high consistency and high task completion in complex scenarios. As semantic information is gradually removed, the internal semantic alignment ability weakens. This directly causes a clear drop in the SR metric and indicates that task reachability is substantially affected. This result shows that semantic priors play a structural role in supporting long term planning and key action decisions.

The decline in PE reflects the sensitivity of path efficiency to missing semantic information. Without high level semantic guidance, the policy network struggles to infer spatial structure, target regions, or critical intermediate steps. As a result, more redundant actions and inefficient paths appear. The loss of semantic priors weakens the model's understanding of scene level constraints. The agent then relies on more exploratory actions during execution, which directly lowers path efficiency. This demonstrates that semantic priors not only improve policy performance but also provide essential structural guidance during the action execution stage.

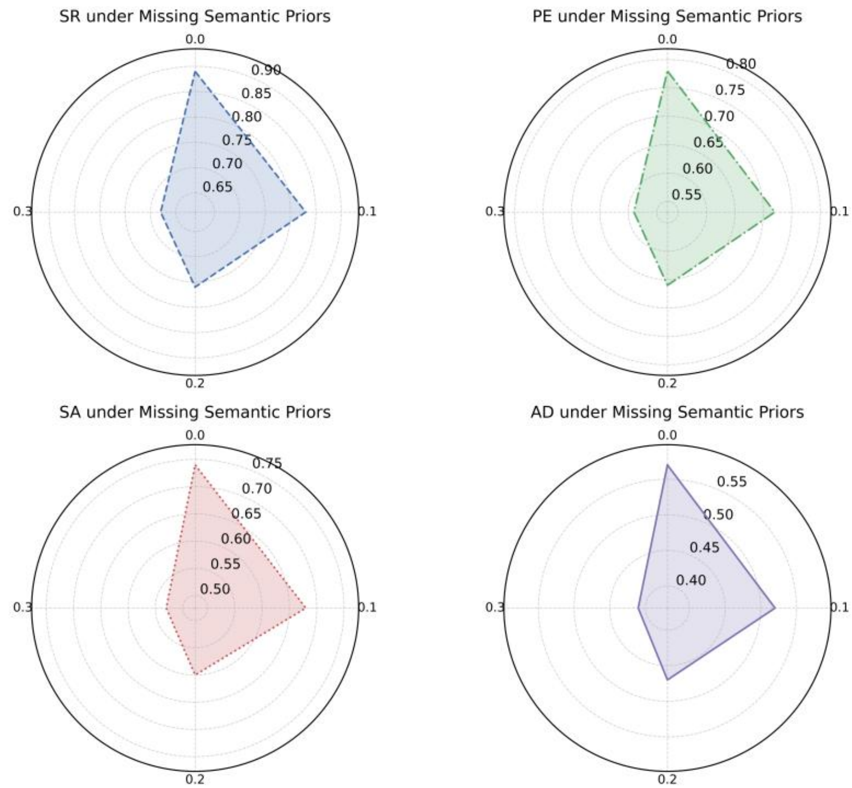

Figure 4. The impact of varying proportions of missing semantic priors on experimental results.

The decrease in SA along with the loss of semantic priors further indicates that semantic modeling ability depends strongly on the completeness of semantic information. When part of the semantic prior is removed, the environment understanding module struggles to maintain structured abstractions of object properties, relations, and scene patterns. Semantic abstraction becomes fragmented or even biased. The disruption of semantic consistency influences the quality of knowledge use in the policy layer. The agent gradually loses its global understanding of environmental semantics during task reasoning. This shows that semantic priors are not auxiliary features but form core knowledge that supports the entire reasoning process.

The drop in AD reveals the role of semantic priors in promoting behavioral diversity. With sufficient semantic information, the agent can choose flexibly among multiple feasible paths and adapt to semantic differences in the environment. When semantic information is missing, the policy collapses into more conservative or even single path behavior patterns. This reduces the exploration space and decreases action richness. This indicates that semantic priors not only improve behavioral effectiveness but also help maintain the breadth of the behavioral distribution. They provide the agent with stronger adaptability and generalization. Taken together, these trends demonstrate the essential role of semantic priors in understanding, reasoning, and behavior generation within the proposed framework.

5. Conclusion

This study proposes a collaborative learning framework for agent behavior modeling and environment understanding that integrates semantic priors. The framework constructs a unified semantic enhanced state representation that deeply couples structured knowledge with multimodal perception. It provides a systematic solution for stable, reliable, and interpretable decision making in complex environments. The framework forms a complete reasoning pipeline from semantic prior encoding to environment understanding, structured dynamic modeling, and policy generation. It connects perception, cognition, and action, and enables the agent to maintain high stability and strong task performance in dynamic, noisy, and partially observable environments.

The results show that introducing semantic priors improves environmental understanding and strengthens the structured mapping between semantics and behavior. Policy reasoning becomes less dependent on purely visual signals and relies more on task structure, object relations, and high level semantic logic. The performance gains from semantic enhancement appear across multiple dimensions, including success rate, path efficiency, semantic accuracy, and action diversity. These findings indicate that semantic priors play an essential role in addressing long term planning difficulties, limited scene generalization, and rigid behavioral patterns. The study provides theoretical and practical evidence of the importance of semantic knowledge for intelligent decision making and offers a solid foundation for building autonomous agents with interpretable cognitive structures.

At the same time, the study has potential value in a wide range of real applications. In domains such as robotics, autonomous driving, human machine collaboration, intelligent manufacturing, and virtual environment interaction, decision making often requires rapid and reliable interpretation of environmental semantics under high uncertainty. The proposed framework offers a more expressive semantic modeling approach for these systems. It enables agents to understand task context, follow semantic rules, and generate actions that better align with environmental logic. The structural advantages of semantic priors also improve model interpretability, which is important for safety sensitive or compliance critical scenarios that require controllable and verifiable decision processes.

Looking forward, the collaborative mechanism between semantic priors and behavior modeling still has broad room for development. Future work may explore more adaptive semantic reasoning mechanisms that allow semantic priors to expand and reorganize based on interaction experience in unknown environments. It is also valuable to investigate how different knowledge sources, such as rules, scene graphs, and multimodal language descriptions, can be integrated into policy generation to improve cross task and cross domain generalization. With advances in large scale models and world models, it may become possible to build unified cognitive behavior frameworks with richer knowledge systems, higher abstraction capabilities, and stronger interpretability. Such frameworks could support large scale deployment of intelligent agents in real world settings.

References

[1] Y. Wang, "Integrating Large Language Models and Knowledge Graphs for Intelligent Financial Regulatory Risk Identification," 2024.

[2] Han S, Zhang Q, Yao Y, et al. LLM multi-agent systems: Challenges and open problems[J]. arXiv preprint arXiv:2402.03578, 2024.

[3] H. Kitano, "Nobel Turing Challenge: Creating the Engine for Scientific Discovery," NPJ Systems Biology and Applications, vol. 7, no. 1, p. 29, 2021.

[4] Qian C, Xie Z, Wang Y, et al. Scaling large language model-based multi-agent collaboration[J]. arXiv preprint arXiv:2406.07155, 2024.

[5] Y. Huang, "Explainable Cognitive Multi-Agent AI for Joint Intention Modeling in Complex Task Planning," 2024.

[6] SJ. Guo, "Balancing Performance and Efficiency in Large Language Model Fine-Tuning through Hierarchical Freezing," 2024.

[7] Jiang B, Xie Y, Wang X, et al. Multi-modal and multi-agent systems meet rationality: A survey[C]//ICML 2024 Workshop on LLMs and Cognition. 2024.

[8] Wang J, Xu H, Jia H, et al. Mobile-agent-v2: Mobile device operation assistant with effective navigation via multi-agent collaboration[J]. Advances in Neural Information Processing Systems, 2024, 37: 2686-2710.

[9] Q. Yang, S. Fu, H. Wang and H. Fang, "Machine-Learning-Enabled Cooperative Perception for Connected Autonomous Vehicles: Challenges and Opportunities," IEEE Network, vol. 35, no. 3, pp. 96-101, 2021.

[10] Y. Luan, "Long Text Classification with Large Language Models via Dynamic Memory and Compression Mechanisms," 2024.

[11] H. Kitano, "Nobel Turing Challenge: Creating the Engine for Scientific Discovery," NPJ Systems Biology and Applications, vol. 7, no. 1, p. 29, 2021.