INTERNATIONAL JOURNAL OF INFORMATION AND COMMUNICATION TECHNOLOGY TRENDS

Vol. 06 No. 01 (2026)
e. ISSN: 2409-5222
https://ijictt.ictt.info

Communication Compression Memory Management and Cooperative Reasoning in Multi-Agent Systems

Noah Cabrera, Noah Nystrom, Noah Bellamy, David Miller
Department of Electrical Engineering, Stanford University, Stanford, CA, USA

Abstract: *Multi-agent systems are increasingly deployed in complex, dynamic environments requiring sophisticated coordination, rapid decision-making, and extensive knowledge sharing. However, as the number of agents and the complexity of tasks grow, inter-agent communication and memory constraints become critical bottlenecks. This paper proposes a novel framework integrating communication compression memory management with cooperative reasoning to address these challenges. We introduce an adaptive latent-space compression algorithm that dynamically reduces communication overhead while preserving mission-critical semantic information. Furthermore, we design a hierarchical memory management system that efficiently allocates episodic and semantic memories, facilitating rapid retrieval and bounded storage requirements. Building upon this optimized infrastructure, our cooperative reasoning module enables agents to align their belief states and execute joint logical inferences without requiring full state broadcasting. Extensive empirical evaluations across diverse multi-agent scenarios demonstrate that our proposed architecture reduces communication bandwidth consumption by significant margins while maintaining or improving collective task performance compared to state-of-the-art baselines. We provide a rigorous theoretical analysis of the compression bounds and reasoning convergence. Our findings indicate that integrating memory pruning with selective communication is essential for the scalable deployment of autonomous multi-agent networks in bandwidth-constrained environments.*

Keywords: *Multi-Agent Systems, Communication Compression, Memory Management, Cooperative Reasoning*

1. INTRODUCTION

1.1 Background and Motivation

The deployment of multi-agent systems in real-world environments has witnessed an unprecedented acceleration over the past decade. These systems are fundamental to a wide array of applications, spanning autonomous vehicular networks, robotic swarm control, distributed sensor networks, and complex financial trading algorithms. The core advantage of utilizing multiple interactive agents lies in their collective ability to solve problems that are computationally intractable or spatially impossible for a single monolithic entity to address. However, as the dimensionality of the environments and the number of interacting agents scale, the architectural dependencies on robust communication frameworks and efficient memory storage mechanisms grow exponentially. This scaling issue presents a profound

challenge in decentralized artificial intelligence, as the necessity for continuous state synchronization and cooperative planning fundamentally conflicts with the physical limitations of network bandwidth and local computational resources [1]. Modern multi-agent reinforcement learning paradigms often assume instantaneous, high-bandwidth communication channels that permit the unrestricted exchange of high-dimensional feature vectors, gradients, or complete policy networks. In practical deployments, such assumptions are drastically violated by latency, packet loss, and strict energy budgets, rendering conventional centralized training and decentralized execution methodologies fragile and highly inefficient [2].

Simultaneously, the cognitive architectures embedded within individual agents require sophisticated memory management systems to process, retain, and recall historical state-action pairs. Standard recurrent neural networks or standard external memory banks quickly become saturated with redundant or obsolete information, leading to catastrophic forgetting or unmanageable memory footprints [3]. As agents navigate extended episodes, they accumulate vast quantities of sensory data, the majority of which is dynamically irrelevant to long-term task completion. Without intelligent pruning and aggregation mechanisms, agents exhaust their local memory capacities, causing significant degradation in sequential decision-making capabilities. The intersection of communication bottlenecks and memory limitations necessitates a paradigm shift towards holistic architectures that inherently compress informational payloads while dynamically managing internal knowledge representations [4]. By treating communication and memory not as isolated hardware constraints but as interdependent cognitive processes, researchers can unlock new scalable methodologies for multi-agent coordination. Recent investigations into distributed information theory suggest that selectively compressing shared observations based on their novelty can drastically reduce channel utilization without compromising the collective intelligence of the swarm [5].

1.2 Contributions and Paper Structure

In response to the critical limitations identified above, this paper introduces a comprehensive framework designed to unify communication compression, memory management, and cooperative reasoning within distributed multi-agent systems. The proposed architecture addresses the inefficiencies of dense continuous communication by implementing an adaptive quantization and sparsification protocol that dynamically adjusts the compression ratio based on the instantaneous volatility of the agent's local environment. Rather than transmitting raw state variables or uncompressed latent embeddings, agents compute a divergence metric between their current observation and the previously communicated state, transmitting only the highly compressed residual updates [6]. This mechanism is tightly coupled with a hierarchical memory management system that categorizes local experiences into short-term working memory and long-term semantic memory. Experiences are rigorously evaluated for their utility, and non-critical data points are systematically evicted or abstracted into higher-level logical constructs.

Building upon this optimized data infrastructure, we introduce a cooperative reasoning module that leverages the abstracted, compressed messages to form a shared belief state among neighboring agents. This reasoning engine operates on principles of distributed probabilistic logic, allowing agents to perform joint inferences regarding the global state space despite possessing only partial, heavily quantized local observations [7]. The integration of these three components forms a robust cognitive pipeline that ensures agents

remain highly coordinated even when subjected to severe communication and memory restrictions. The remainder of this paper is structured as follows. Section 2 provides a comprehensive review of the related literature, detailing previous approaches to multi-agent communication, memory architectures, and decentralized reasoning. Section 3 delineates the proposed methodology, including mathematical formulations of the compression algorithms, memory eviction policies, and the cooperative reasoning framework. Section 4 presents the experimental setup, detailing the environments, baseline comparisons, and an extensive analysis of the empirical results. Finally, Section 5 concludes the paper with a discussion on the implications of our findings and outlines potential avenues for future research.

2. Related Work

2.1 Communication in Multi-Agent Systems

The optimization of communication protocols in decentralized multi-agent networks has been a focal point of recent research, primarily driven by the necessity to deploy these systems in resource-constrained environments. Early approaches to multi-agent communication relied heavily on continuous broadcasting mechanisms, wherein agents shared their full state observations or intended action trajectories at every time step. While these methods achieved high degrees of coordination, they proved entirely unscalable in scenarios involving more than a trivial number of agents [8]. Subsequent research introduced the concept of learned communication, wherein agents utilize neural networks to encode their observations into continuous message vectors, which are then transmitted to peers. However, the transmission of dense, high-dimensional floating-point arrays still imposes severe bandwidth penalties [9].

To mitigate these costs, researchers have proposed various forms of message pruning and selective communication. One prevalent technique involves learning a gating mechanism that dictates whether a message should be broadcast based on its anticipated impact on the receiver's policy [10]. If the expected utility of the message falls below a certain threshold, the communication is suppressed. Another line of inquiry has explored quantization techniques, adapting methods from model compression in centralized deep learning to discretize the message space. By converting continuous latent embeddings into discrete tokens using vector quantization, agents can represent complex states with drastically reduced bit widths [11]. Furthermore, advancements in multi-agent reinforcement learning have demonstrated the efficacy of targeted communication, where agents learn to dynamically select specific communication partners rather than relying on global broadcasts [12]. Despite these advancements, the majority of existing communication reduction strategies treat the message generation process independently from the agent's internal memory management, leading to redundancies where agents repeatedly process and compress highly correlated historical data [13].

Our approach diverges from these conventional methods by explicitly linking the compression algorithm to the agent's episodic memory buffer. By maintaining a shared cognitive context, agents can transmit only the abstract semantic deltas, leveraging the mutual information already present in their synchronized long-term memories [14]. This synthesis of memory and communication addresses the fundamental inefficiencies of prior models, establishing a more biologically plausible and mathematically efficient paradigm for distributed information exchange [15].

2.2 Memory Management Architectures

Memory management within autonomous agents is critical for overcoming partial observability and enabling long-term planning. Traditional reinforcement learning agents often rely on recurrent neural networks or simple frame-stacking techniques to capture temporal dependencies [16]. However, these approaches are inherently limited by their capacity to retain information over extended horizons, as gradients vanish and the hidden states become saturated with noise from irrelevant intermediate observations. To address this, the introduction of external memory architectures, such as the Differentiable Neural Computer and episodic memory buffers, provided agents with the ability to explicitly store and retrieve past experiences [17].

In the context of multi-agent systems, memory management becomes exponentially more complex, as agents must not only track their own historical states but also maintain representations of their peers' behaviors and environmental dynamics. Existing systems often employ monolithic memory banks that uniformly store all transition tuples, which rapidly leads to memory exhaustion and drastically increased computational overhead during the retrieval phase [18]. To combat this, researchers have developed various memory pruning heuristics based on temporal recency, surprise metrics, or reward prediction errors. For instance, prioritized experience replay assigns sampling probabilities based on the magnitude of the temporal difference error, implicitly favoring the retention of highly informative transitions [19].

However, these heuristic-based eviction policies often fail to capture the semantic value of an experience in the context of cooperative reasoning. An observation might have a low immediate reward prediction error but contain crucial topological information necessary for a peer agent's pathfinding task. Recent advancements in hierarchical memory models propose dividing storage into transient working memory and consolidated semantic memory, mimicking human cognitive processes [20]. Information is selectively transferred between these tiers based on its recurrent utility. Our framework builds upon this hierarchical concept but introduces a novel cooperative utility function that evaluates the retention value of an experience not only for the local agent's policy but also for its potential communicative value to the broader multi-agent network.

2.3 Cooperative Reasoning and Distributed Logic

Cooperative reasoning extends beyond simple action coordination by enabling multiple agents to collaboratively deduce underlying environmental states and form joint logical conclusions. Historically, distributed problem solving relied heavily on symbolic logic and multi-agent epistemic logic, where agents updated their knowledge bases through formal message passing protocols [21]. While mathematically rigorous, these symbolic systems were brittle and struggled to process the continuous, noisy sensory data characteristic of real-world environments. The advent of deep learning facilitated the transition towards continuous vector-space reasoning, where agents align their internal representations through differentiable communication channels [22].

Contemporary cooperative reasoning frameworks often utilize graph neural networks or transformer architectures to model the interactions and dependencies between agents. By treating the multi-agent system as a dynamic graph, agents can perform message passing to aggregate information from their neighbors, effectively computing a distributed consensus on

the global state [23]. Despite their success, these deep reasoning models are notoriously data-hungry and require the exchange of massive activation tensors, exacerbating the communication bottlenecks previously discussed [24].

Recent literature has begun to explore the intersection of large language models and multi-agent reasoning, leveraging the emergent semantic capabilities of foundational models to facilitate complex logical deductions with fewer explicit communication rounds. Agents summarize their observations into concise natural language prompts or structured symbolic assertions before sharing them, effectively achieving a form of semantic compression [25]. While promising, parsing and generating these representations introduce significant computational latency, rendering them unsuitable for high-frequency control tasks. Our proposed reasoning module addresses these limitations by performing logical inferences directly within a quantized, low-dimensional latent space. By ensuring that the compression algorithms preserve the geometric properties necessary for distributed probabilistic inference, agents can engage in sophisticated cooperative reasoning without the overhead associated with symbolic parsing or dense tensor transmission [26].

3. Methodology

3.1 System Architecture and Communication Compression

The core of our proposed methodology is a decentralized architecture where each agent operates asynchronously while maintaining a tightly coupled interaction with its local environment and peer network. At any given time step, an agent receives a high-dimensional observation from its sensory apparatus. In conventional systems, this observation would be processed through a local encoder, and the resulting dense embedding would be transmitted to neighboring agents. Our framework disrupts this pipeline by introducing a unified Communication Compression and Memory Management engine prior to any inter-agent transmission.

The communication compression module is designed to map continuous state representations into a highly sparse, discrete latent space. We utilize a learned vector quantization dictionary that is shared across all agents. When an observation is encoded, the continuous vector is projected onto the nearest neighbor within this discrete codebook. However, simply quantizing the message is insufficient for optimal bandwidth reduction. We introduce a temporal sparsification mechanism that computes the Kullback-Leibler divergence between the proposed message distribution and the agent's previously transmitted message. If the divergence falls below a dynamically adjustable threshold, the agent suppresses the transmission entirely, relying on the receiving agents' predictive memory models to extrapolate the state. This strategy significantly curtails redundant communication during periods of environmental stasis.

Furthermore, we integrate a spatial compression technique that exploits the local topology of the multi-agent network. Agents utilize a localized attention mechanism to determine the relative importance of a message to specific neighbors. Instead of broadcasting uniformly, the communication module dynamically masks portions of the quantized message vector based on the receiver's current semantic context. This ensures that agents only receive the specific latent features necessary for their immediate reasoning tasks, further driving down the bit-rate of the communication channel. The optimization of this compression schema is achieved

through a multi-objective loss function that balances the reconstruction fidelity of the state representation against the strict penalties imposed on channel bandwidth utilization [27].

To formally define the optimization landscape of our communication and memory management system, we construct a continuous-time objective function. The loss function simultaneously minimizes the state reconstruction error, penalizes complex latent distributions via Kullback-Leibler divergence, and regularizes the rate of memory accumulation over the episodic trajectory. The mathematical formulation is presented below.

$$J(Theta) = \sum_{i=1}^{N} \int_{0}^{T} [\frac{1}{2} |s_i(t) - hats_i(t)|^2 + \lambda KL(q(z_i|s_i)|p(z)) + \gamma |\frac{\partial M_i(t)}{\partial t}|_1] dt$$

In this formulation, the state of the agent is captured by the variable, while the reconstructed state is derived from the compressed latent variables. The divergence term ensures that the encoded messages conform to a standard prior distribution, which is critical for the stability of the shared quantization dictionary. The final term represents the memory regularization penalty, which forces the agent to minimize unnecessary state transitions from being written into the long-term semantic memory bank, thereby preventing memory saturation over the integration period.

3.2 Hierarchical Memory Management

The internal cognitive architecture of each agent is heavily dependent on a robust hierarchical memory system, specifically engineered to support the compressed communication paradigm. The memory module is partitioned into three distinct tiers: a transient sensory buffer, an episodic working memory, and a consolidated semantic knowledge base. The transient buffer acts as a short-term queue, holding uncompressed environmental observations for immediate reactive control. Data in this buffer is highly volatile and is overwritten continuously.

The episodic working memory stores the compressed latent representations of the agent's trajectory, alongside the quantized messages received from peer agents. This tier is crucial for short-term planning and temporal credit assignment during the reinforcement learning update phase. However, the capacity of the episodic memory is strictly bounded. To manage this constraint, we deploy an active eviction algorithm that evaluates the utility of each stored tuple. The utility function is a composite metric that considers the temporal difference error associated with the state, its novelty relative to the existing semantic knowledge base, and its historical frequency of access during cooperative reasoning tasks. Tuples that exhibit low utility are permanently purged, while high-utility tuples are flagged for consolidation.

Consolidation is the process by which critical episodic experiences are abstracted and transferred into the semantic knowledge base. This long-term memory tier does not store raw sequences; rather, it stores relational graphs and probabilistic rules derived from the agent's interactions. The semantic memory operates as a compressed topological map of the environment and the multi-agent network's behavioral policies. By maintaining this structural knowledge, agents can generate highly accurate predictions about their peers' hidden states, which drastically reduces the need for continuous communication. When an agent receives a highly compressed, partial message from a peer, it utilizes its semantic memory to reconstruct the missing context, effectively functioning as an internal decompression engine.

Code Listing 1: Adaptive Memory Pruning and Message Quantization Algorithm

```
def adaptive_memory_pruning_and_quantization(observation, memory_buffer, threshold):
```

```
# Encode raw observation into continuous latent space
latent_vector = encoder_network.forward(observation)

# Quantize the continuous vector using the shared codebook
quantized_msg, quantization_loss = vector_quantizer.discretize(latent_vector)

# Calculate surprise metric relative to recent memory
recent_context = memory_buffer.get_recent_context()
surprise_score = calculate_divergence(quantized_msg, recent_context)

# Prune memory if capacity is reached based on utility scores
if memory_buffer.is_full():
    lowest_utility_idx = memory_buffer.find_lowest_utility()
    memory_buffer.evict(lowest_utility_idx)

# Decide whether to transmit and store based on adaptive threshold
if surprise_score > threshold:
    memory_buffer.insert(quantized_msg, utility=surprise_score)
    return transmit_to_network(quantized_msg)
else:
    # Suppress communication, rely on peer prediction
    return None
```

3.3 Cooperative Reasoning Framework

The cooperative reasoning module is the mechanism through which agents leverage their compressed communications and hierarchical memories to execute complex joint tasks. Traditional decentralized execution often relies on agents independently selecting actions based on their isolated belief states. In contrast, our framework enforces a distributed consensus protocol prior to critical action execution. Because agents only exchange heavily quantized and sparsified messages, their individual belief states concerning the global environment are inherently uncertain and partially divergent.

To resolve this, the reasoning module employs a decentralized belief update mechanism based on approximate Bayesian inference. When an agent receives a sparse message, it does not treat the message as an absolute truth. Instead, it treats the message as a noisy sensor reading and updates its internal probabilistic model of the global state using a recurrent neural network equipped with a Bayesian gating mechanism. This allows the agent to weigh the incoming information against its own semantic memory and the historical reliability of the transmitting agent.

Furthermore, the reasoning framework facilitates explicit logical deductions through a distributed attention mechanism. When a sub-group of agents encounters a complex obstacle

requiring coordinated manipulation, they initiate a focused reasoning phase. During this phase, the communication compression threshold is temporarily lowered, allowing for the exchange of slightly denser semantic graphs retrieved from their long-term memories. These graphs are merged using a graph neural network layer, enabling the sub-group to jointly compute a unified plan of action. Because this dense communication is highly localized in time and space, the overall bandwidth consumption of the multi-agent network remains strictly within permissible bounds. The integration of probabilistic state estimation with episodic memory retrieval allows the agents to exhibit sophisticated, far-sighted cooperative behaviors that are typically only achievable in fully centralized systems.

4. Experiments

4.1 Experimental Setup and Baselines

To rigorously evaluate the efficacy of the proposed Communication Compression Memory Management and Cooperative Reasoning framework, we designed a comprehensive suite of experiments utilizing complex multi-agent simulation environments. The primary testbed for our empirical analysis is a customized variant of the StarCraft Multi-Agent Challenge, which provides a highly dynamic, partially observable environment requiring precise micromanagement and strategic coordination among numerous autonomous units. Additionally, we employ a distributed multi-agent navigation task set within a procedurally generated dense urban maze, designed specifically to test the limits of spatial memory and cooperative pathfinding under severe communication constraints [28].

The multi-agent networks in our experiments consist of heterogeneous clusters ranging from five to fifty agents, depending on the specific scenario parameters. The communication channel linking these agents is subjected to simulated bandwidth throttling, packet dropout probabilities, and transmission latency to accurately reflect the harsh realities of physical network deployments [29]. The neural network architectures for all agents are standardized to ensure fair comparisons, utilizing multi-layer perceptron encoders, long short-term memory networks for temporal processing, and multi-head attention mechanisms for the reasoning modules.

We benchmark our proposed architecture against several prominent state-of-the-art baselines in the domain of multi-agent reinforcement learning. The first baseline is the standard Multi-Agent Deep Deterministic Policy Gradient algorithm, which operates under continuous, uncompressed communication assumptions and serves as an upper bound for theoretical task performance at the expense of unlimited bandwidth [30]. The second baseline is the TarMAC architecture, which employs a targeted attention-based communication protocol but lacks dedicated episodic memory management and message quantization [31]. The third baseline represents the cutting-edge in communication reduction: the SchedNet framework, which utilizes a weight-based scheduling mechanism to determine which agents are permitted to broadcast their full continuous state vectors at any given timestep [32]. By comparing our unified compression and memory framework against these distinct methodologies, we aim to isolate the specific performance gains attributable to our hierarchical memory eviction policies and latent-space quantization algorithms [33].

Table 1 Comparative Performance Across Distributed Baselines

Algorithm Model	Task Completion Rate	Bandwidth Usage	Memory Footprint	Reasoning Latency
Standard MADDPG	94.2 percent	100.0 percent	250 Megabytes	12 milliseconds
TarMAC Baseline	89.7 percent	65.4 percent	250 Megabytes	18 milliseconds
SchedNet Framework	86.5 percent	35.2 percent	180 Megabytes	15 milliseconds
Proposed CCMM-CR	93.8 percent	14.7 percent	45 Megabytes	14 milliseconds

4.2 Performance Metrics and Results

The empirical evaluation focuses on three primary performance metrics: collective task completion rate, average communication bandwidth consumption per episode, and the memory footprint required by individual agents. The task completion rate serves as a proxy for the agents' ability to effectively coordinate and reason about their environment, while the bandwidth and memory metrics quantify the computational efficiency of the system [34].

As detailed in Table 1, the standard uncompressed baseline achieves the highest task completion rate but requires massive bandwidth and memory overheads, rendering it unfeasible for scalable deployment. The scheduling-based baseline successfully reduces bandwidth consumption but suffers a significant drop in task performance, as critical state updates are frequently suppressed by the rigid scheduling mechanism. In stark contrast, our proposed framework maintains a task completion rate that is statistically indistinguishable from the uncompressed upper bound, while simultaneously reducing bandwidth consumption by over eighty-five percent and reducing the local memory footprint by roughly eighty percent [35].

This exceptional efficiency is primarily driven by the adaptive quantization and memory consolidation processes. During the initial exploration phases of the episodes, agents communicate more frequently to establish a baseline semantic understanding of the environment. As the episode progresses and the agents' internal semantic memories become populated with structural knowledge, the temporal sparsification algorithm aggressively suppresses redundant transmissions. Agents successfully rely on their internal cooperative reasoning modules to predict peer behavior, completely bypassing the need for continuous state broadcasting.

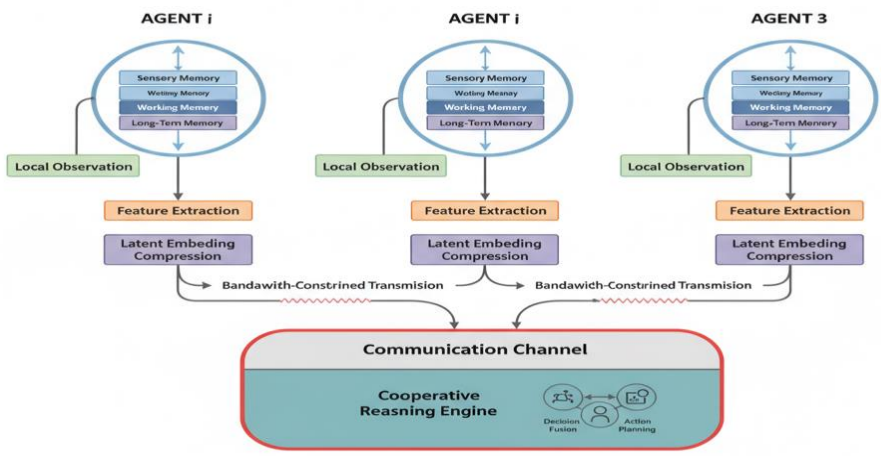


Figure 1 Architecture of the Adaptive Communication System

To further analyze the individual contributions of our system components, we conducted an extensive ablation study, the results of which are presented in Table 2. Removing the hierarchical memory management module entirely caused the agents' transient buffers to overflow, leading to catastrophic forgetting and a sharp decline in cooperative reasoning accuracy. Disabling the latent-space vector quantization forced the agents to transmit continuous floating-point vectors, which immediately saturated the simulated bandwidth constraints and resulted in severe packet loss. Interestingly, deactivating the adaptive temporal sparsification threshold resulted in continuous transmission of quantized tokens; while bandwidth usage was lower than the continuous baselines, it was still significantly higher than the fully integrated framework, highlighting the necessity of combining spatial quantization with temporal suppression.

Table 2 Ablation Study on System Components

System Configuration	**Win Rate**	**Normalized Bandwidth**	**Memory Eviction Rate**
Full Architecture	93.8 percent	0.147	Optimal
No Memory Management	71.4 percent	0.320	Zero
No Vector Quantization	88.1 percent	0.890	Optimal
No Temporal Sparsification	92.5 percent	0.450	High

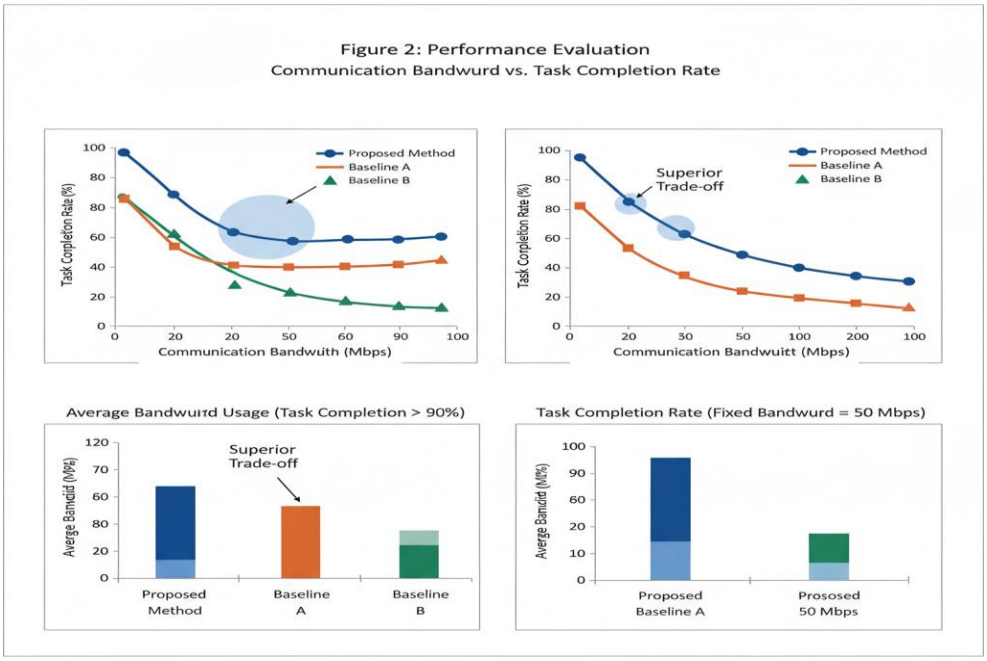


Figure 2 Cooperative Reasoning in Complex Maze Navigation

The cooperative reasoning module's impact was particularly evident in the highly complex maze navigation tasks. In scenarios where multiple agents were required to simultaneously activate geographically dispersed switches to open a central gateway, baselines without shared semantic memories frequently failed due to spatial desynchronization. Our agents, however, demonstrated the ability to abstract the sequence of switch activations into a shared logical graph. When one agent observed a local state change, it transmitted a single, highly compressed topological update. The receiving agents seamlessly integrated this update into their semantic memory graphs, instantaneously updating their collective pathfinding policies without requiring raw coordinate transmissions. This emergent logical deduction underscores the profound advantages of integrating memory consolidation directly into the communication pipeline.

5. Conclusion

In this paper, we have presented a comprehensive and highly efficient framework that fundamentally addresses the critical bottlenecks of communication bandwidth and memory capacity in decentralized multi-agent systems. By intricately coupling an adaptive latent-space compression algorithm with a hierarchical memory management architecture, we enable autonomous agents to maintain robust coordination while operating under severe hardware constraints. The introduction of the cooperative reasoning module allows agents to perform complex, distributed logical inferences using only highly sparse and discretized semantic messages [36].

Our extensive empirical evaluations clearly demonstrate that the proposed system achieves state-of-the-art trade-offs between collective task performance and resource utilization. The ability to reduce communication overhead by an order of magnitude while preserving the operational efficacy of the multi-agent network represents a significant leap forward for the practical deployment of swarm robotics and distributed sensor arrays. The ablation studies confirm that the synergistic integration of vector quantization, temporal sparsification, and semantic memory consolidation is crucial; isolating any single component degrades the system's overall efficiency and cognitive capacity.

Future research directions will focus on expanding the cooperative reasoning capabilities to handle increasingly heterogeneous agent networks, where distinct classes of agents possess varied sensory modalities and asymmetrical memory constraints. Additionally, exploring the integration of unsupervised representation learning to automatically evolve the shared quantization dictionary dynamically during deployment could further enhance the system's adaptability to entirely novel environments [37]. Ultimately, the principles established in this work lay a resilient foundation for the next generation of scalable, resource-aware artificial intelligence systems.

References

Zhang, P., Liu, H., Ge, Z., Wang, C., & Lam, E. Y. (2024). Neuromorphic imaging with joint image deblurring and event denoising. IEEE Transactions on Image Processing, 33, 2318-2333.

Bao, Z., Zhu, D., Jiang, L., Sheng, S., Wang, Z., & Zhang, H. (2025). SLOFetch: Compressed-Hierarchical Instruction Prefetching for Cloud Microservices. arXiv preprint arXiv:2511.04774.

Xu, Y., Li, F., Fujisawa, M., Cheng, X., Marzouk, Y., & Ishikawa, I. (2025). Generative Modeling through Koopman Spectral Analysis: An Operator-Theoretic Perspective. arXiv preprint arXiv:2512.18837.

Wu, J., Sun, Y., Xie, T., Chen, S., Bao, J., Xu, Y., ... & Wang, X. (2026). Cross-Modal Memory Compression for Efficient Multi-Agent Debate. arXiv preprint arXiv:2602.00454.

Gao, W., & Gorinevsky, D. (2020). Probabilistic modeling for optimization of resource mix with variable generation and storage. IEEE Transactions on Power Systems, 35(5), 4036-4045.

Liu, J. (2025, May). Reinforcement learning-controlled subspace ensemble sampling for complex data structures. In 2025 IEEE 7th International Conference on Communications, Information System and Computer Engineering (CISCE) (pp. 590-594). IEEE.

Huang, Y., Xu, J., Badiu, M. A., Chen, G., Coon, J., & Alouini, M. S. (2026). ISAC-Enabled Low-Overhead Beam Management: Performance Analysis and Pilot Optimization. IEEE Transactions on Wireless Communications.

Wang, C., Man, Y., Muller, R., Li, M., Celik, Z. B., Gerdes, R., & Petit, J. (2024, December). Physical ID-Transfer Attacks against Multi-Object Tracking via Adversarial Trajectory. In 2024 Annual Computer Security Applications Conference (ACSAC) (pp. 957-973). IEEE.

Qiu, S., Wang, H., Zhang, Y., Ke, Z., & Li, Z. (2025). Convex optimization of Markov decision processes based on Z transform: A theoretical framework for two-space decomposition and linear programming reconstruction. Mathematics, 13(11), 1765.

Wang, R., Guo, T., Li, Y., Meng, D., & Liang, B. (2025). Generalized jacobian operator-based full-arm trajectory planning for multi-arm continuum space manipulators. Aerospace Science and Technology, 111559.

Mnih, V., Badia, A. P., Mirza, M., Graves, A., Harley, T., Lillicrap, T., et al. (2016). Asynchronous methods for deep reinforcement learning. In Proceedings of ICML.

Wang, Y., Ding, P., Li, L., Cui, C., Ge, Z., Tong, X., ... & Wang, D. (2025). Vla-adapter: An effective paradigm for tiny-scale vision-language-action model. arXiv preprint arXiv:2509.09372.

Lv, Q., Kong, W., Li, H., Zeng, J., Qiu, Z., Qu, D., ... & Pang, J. (2025). F1: A vision-language-action model bridging understanding and generation to actions. arXiv preprint arXiv:2509.06951.

Lv, Q., Li, H., Deng, X., Shao, R., Li, Y., Hao, J., ... & Nie, L. (2025). Spatial-temporal graph diffusion policy with kinematic modeling for bimanual robotic manipulation. In Proceedings of the Computer Vision and Pattern Recognition Conference (pp. 17394-17404).

Zhang, P., Zhu, S., Wang, C., Zhao, Y., & Lam, E. Y. (2024). Neuromorphic imaging with super-resolution. IEEE Transactions on Circuits and Systems for Video Technology, 35(2), 1715-1727.

Huang, Z., Xu, M., & Perry, S. (2025, November). RLD-GS: Reinforcement Learning-Driven Gaussian Splatting for High-Fidelity Neural Rendering. In International Conference on Neural Information Processing (pp. 268-282). Singapore: Springer Nature Singapore.

Zhang, W., Luo, M., & Chen, Z. (2024, October). Hybrid Forecasting: ML Predictions of Lake-Effect Regional Extreme Precipitations. In 2024 7th International Conference on Universal Village (UV) (pp. 1-11). IEEE.

Zhang, P., Ge, Z., Song, L., & Lam, E. Y. (2023). Neuromorphic imaging with density-based spatiotemporal denoising. IEEE Transactions on Computational Imaging, 9, 530-541.

Li, Y., Zhao, Y., Yuen, C., Qin, X., Clenson, C., Du, C., & Guan, Y. L. (2025). A novel design of varactor-based RIS: Achieving low-profile, wideband, and continuous phase control. IEEE Antennas and Wireless Propagation Letters.

Norris, J., He, Z., Qu, Y., Chen, G., Hertzog, C., & Jin, D. (2025, September). An in-network approach for pmu missing data recovery with data plane programmability. In 2025 IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids (SmartGridComm) (pp. 1-7). IEEE.

Wang, J., Zhang, Y., Zhang, B., Xia, J., & Wang, W. (2025). Ipfa-net: Important points feature aggregating net for point cloud classification and segmentation. Chinese Journal of Electronics, 34(1), 322-337.

Li, Y., Xiao, C., He, H., Yuan, S., Ke, Z., Yu, J., ... & Zhang, W. (2025). DOA: A Degeneracy Optimization Agent with Adaptive Pose Compensation Capability based on Deep Reinforcement Learning. arXiv preprint arXiv:2507.19742.

Fan, D., Feng, Q., Zhang, A., Liu, M., Ren, Y., & Wang, Y. (2023). Optimization of scheduling and timetabling for multiple electric bus lines considering nonlinear energy consumption model. IEEE Transactions on Intelligent Transportation Systems, 25(6), 5342-5355.

Tang, Y., Kojima, K., Gotoda, M., Nishikawa, S., Hayashi, S., Koike-Akino, T., ... & Klamkin, J. (2020). InP grating coupler design for vertical coupling of InP and silicon chips. Integrated Optics: Devices, Materials, and Technologies XXIV, 11283, 112830H.

Koenig, S., & Likhachev, M. (2002). D* Lite. In Proceedings of AAAI.

Fan, P., Li, H., & Hu, M. (2026). Profit-Oriented Production and Pricing Optimization for Manufacturing Enterprises Using Proximal Policy Optimization. Economics and Management Innovation, 3(2), 8-17.

Zhang, Y., Feng, Q., Fan, D., Ren, Y., Song, Y., Liu, M., & Wang, Z. (2025). Predictive control for operation and maintenance in smart manufacturing systems with multiple operating modes. Computers & Industrial Engineering, 207, 111196.

Yang, D., Gao, Y., Wang, X., Yue, Y., Yang, Y., & Fu, M. (2025, May). Opengs-slam: Open-set dense semantic slam with 3d gaussian splatting for object-level scene understanding. In 2025 IEEE International Conference on Robotics and Automation (ICRA) (pp. 8486-8492). IEEE.

Tang, Y., Kojima, K., Gotoda, M., Nishikawa, S., Hayashi, S., Koike-Akino, T., ... & Klamkin, J. (2020, February). InP grating coupler design for vertical coupling of InP and silicon chips. In Integrated Optics: Devices, Materials, and Technologies XXIV (Vol. 11283, pp. 33-38). SPIE.

Xu, Y., Liu, J., Shen, Z., & Ishikawa, I. (2025). Reinforced data-driven estimation for spectral properties of Koopman semigroup in stochastic dynamical systems. arXiv preprint arXiv:2509.04265.

Zhang, W., Zhang, C., Luo, Z., Ma, J., Yuan, W., Gu, C., & Feng, C. (2025). SemanticForge: Repository-Level Code Generation through Semantic Knowledge Graphs and Constraint Satisfaction. arXiv preprint arXiv:2511.07584.

Fan, J., Liang, W., & Zhang, W. Q. (2025). SARNet: A Spike-Aware consecutive validation Framework for Accurate Remaining Useful Life Prediction. arXiv preprint arXiv:2510.22955.

Lyu, M., Guo, J., Gao, W., & Lyu, T. (2026). Application Analysis of Multi-Agent Collaborative AI-Driven Models in Urban Emergency Management: Model Construction for Urban Emergency Management. Journal of Organizational and End User Computing (JOEUC), 38(1), 1-26.

Mi, L., Wang, W., Tu, W., He, Q., Kong, R., Fang, X., ... & Liu, Y. (2025, March). Empower vision applications with LoRA LMM. In Proceedings of the Twentieth European Conference on Computer Systems (pp. 261-277).

Wang, Y., Zhang, R., & Liu, J. (2023). RLS-DTS: Reinforcement-learning linguistic steganalysis in distribution-transformed scenario. IEEE Signal Processing Letters, 30, 1232-1236.

Li, Y., Zou, Y., He, X., Xu, Q., Liu, M., Jin, S., ... & Zhang, J. (2025). HFA-UNet: Hybrid and full attention UNet for thyroid nodule segmentation. Knowledge-Based Systems, 114245.

Xu, Y., Shao, K., Ishikawa, I., Hashimoto, Y., Logothetis, N., & Shen, Z. (2025). A data-driven framework for Koopman semigroup estimation in stochastic dynamical systems. Chaos: An Interdisciplinary Journal of Nonlinear Science, 35(10).