



INTERNATIONAL JOURNAL OF INNOVATIVE RESEARCH

VOL 04. ISSUE 02 . (2026)

<http://ijirjournal.com/>

Graph-Based Embodied Planning for Heterogeneous Robot Teams in Unstructured Fields

Sara Hamed

Mechatronics Engineering Department, Faculty of Engineering, Ain Shams University, Cairo 11517, Egypt

Abstract.

The deployment of heterogeneous robot teams in unstructured and dynamic environments remains a formidable challenge in modern robotics. Traditional planning algorithms often fail to account for the unique morphological and sensory capabilities of varying robotic platforms, leading to suboptimal performance in terrains lacking defined topological pathways. This paper presents a novel framework for graph-based embodied planning specifically designed for heterogeneous multi-agent systems operating in unstructured fields. By integrating embodied artificial intelligence principles with dynamic graph-theoretic structures, the proposed methodology allows individual agents, such as unmanned aerial vehicles, legged robots, and wheeled ground vehicles, to construct, share, and navigate a unified semantic-topological representation of their environment. This shared graph dynamically updates edge weights based on the localized, embodied perception of each agent, thereby reflecting real-time traversability and task-specific affordances. Through decentralized coordination mechanisms, the system achieves efficient cross-platform task allocation and collision-free trajectory generation. Extensive simulation results demonstrate that the proposed graph-based embodied approach significantly reduces computational overhead and planning latency while improving mission success rates compared to conventional multi-agent pathfinding methodologies. The framework provides a robust foundation for applications in disaster response, planetary exploration, and automated agriculture, where environmental unpredictability necessitates highly adaptive and cooperative robotic solutions.

Keywords: *Embodied Artificial Intelligence, Heterogeneous Robotics, Graph-Based Planning, Unstructured Environments, Multi-Agent Coordination.*

1. Introduction

The advancement of autonomous systems has increasingly shifted focus from structured, predictable environments to unstructured, dynamic fields [1]. Environments such as disaster zones, dense forests, deep subterranean caverns, and planetary surfaces present extreme challenges for autonomous navigation and task execution. In these scenarios, the terrain lacks geometric regularity, semantic predictability, and predefined navigational pathways. To effectively operate in such domains, researchers have increasingly turned to multi-agent robotic systems. Deploying a single homogeneous platform is often insufficient due to the inherent trade-offs in robotic design; for instance, wheeled robots offer high energy efficiency on flat terrain but struggle with vertical

obstacles, whereas legged robots can traverse rugged ground but are energy-intensive [2]. Unmanned aerial vehicles can bypass ground obstacles entirely but are severely constrained by payload capacity and battery life [3]. Consequently, heterogeneous robot teams that combine these diverse platforms offer a synergistic approach, leveraging the strengths of each morphology to achieve complex mission objectives [4]. Despite the clear advantages of heterogeneous teams, coordinating such systems in unstructured fields introduces profound theoretical and practical complexities. Traditional multi-agent pathfinding algorithms typically assume a grid-based or highly structured environmental representation where all agents share similar kinematic constraints. When applied to heterogeneous teams in unstructured terrains, these traditional methods suffer from the curse of dimensionality and often fail to capture the nuanced physical interactions between specific robotic morphologies and complex topographies. Embodied artificial intelligence provides a paradigm shift in addressing this problem. Instead of relying on abstract, disembodied representations of the world, embodied intelligence asserts that an agent's cognition and planning are inextricably linked to its physical form and sensory capabilities. Therefore, a graph-based embodied planning approach allows each robot to perceive and interact with the environment according to its unique physical constraints, updating a shared knowledge structure in a manner that benefits the entire team. This paper proposes a comprehensive framework for graph-based embodied planning tailored for heterogeneous robot teams. The core proposition is the continuous, decentralized generation of a semantic-topological graph that serves as a shared cognitive map for the robotic team. In this graph, nodes represent spatial regions of interest, and edges encode the traversability and affordances specific to different robotic platforms. As agents explore the unstructured field, they utilize their embodied sensory-motor loops to evaluate terrain traversability, continuously updating the graph's edge weights. This dynamic graph allows for real-time path optimization and task allocation without the need for a centralized computational authority, which is often a single point of failure in degraded communication environments. The methodology presented in this research bridges the gap between high-level topological planning and low-level continuous control. By utilizing a graph theoretic approach, the computational complexity of planning across large, unstructured areas is drastically reduced. Simultaneously, the incorporation of embodied artificial intelligence ensures that the physical realities of heterogeneous robots, such as varying turning radii, climbing capabilities, and sensory horizons, are mathematically embedded into the planning algorithm. The subsequent sections will detail the theoretical foundations of this approach, formulate the environmental and kinematic models, and present a rigorous analysis of the decentralized coordination mechanisms required to prevent conflicts and ensure mission success.

2. Literature Review and Background

2.1 Embodied Artificial Intelligence In Autonomous Robotics

The concept of embodied artificial intelligence stems from the realization that intelligent behavior cannot be fully understood or synthesized in isolation from the physical body and its environment [5]. Early cognitive science models heavily relied on symbol manipulation, treating perception, planning, and action as separate, sequential modules. However, these classical models often proved brittle when applied to real-world robotics in noisy, unstructured domains. Embodied intelligence, conversely, emphasizes the perception-action loop, positing that an agent constructs its understanding of the world through direct physical interaction [6]. In recent years, embodied AI has been successfully applied to single-agent navigation tasks. Reinforcement learning frameworks have enabled legged robots to learn robust locomotion policies over challenging

terrains by mapping exteroceptive sensory inputs directly to joint torques [7]. These approaches demonstrate that when a robot learns to navigate by considering its physical morphology, such as the placement of its center of mass or the friction coefficients of its end-effectors, it can traverse terrains that would confound traditional geometric planners. However, extending embodied AI from single-agent locomotion to multi-agent, heterogeneous planning remains an open research question. Current literature largely addresses homogeneous swarms, where the embodied experience of one agent is directly transferable to another [8]. In heterogeneous teams, the embodied experience of an aerial vehicle is fundamentally incompatible with that of a ground vehicle, necessitating a translation mechanism to share spatial intelligence effectively.

2.2 Multi-Agent Path Finding And Coordination

Multi-agent pathfinding is a well-established subfield of artificial intelligence and robotics. The primary objective is to compute collision-free paths for multiple agents moving from respective start locations to target destinations [9]. The classical formulation of this problem utilizes discrete graphs, such as grids, where agents move synchronously. Algorithms like Conflict-Based Search and its bounded-suboptimal variants have achieved remarkable success in structured environments such as automated warehouses and semiconductor manufacturing plants [10]. However, the transition from structured grids to unstructured fields invalidates many assumptions inherent in classical multi-agent pathfinding. Unstructured environments cannot be efficiently discretized into uniform grids without massive loss of geometric information or exponential increases in computational state space [11]. Furthermore, classical algorithms assume homogeneous agent kinematics, uniform traversal times, and reliable global communication. Recent advancements in continuous multi-agent pathfinding attempt to address these limitations by modeling the problem in continuous space using velocity obstacles or control barrier functions [12]. While effective for local collision avoidance, these continuous methods often get trapped in local minima in heavily cluttered environments and lack the global foresight necessary for long-horizon mission planning in unstructured terrains.

2.3 Graph-Based Representations For Unstructured Environments

To navigate large-scale unstructured environments, robotic systems frequently employ graph-based spatial representations. Probabilistic Roadmaps and Rapidly-exploring Random Trees are foundational sampling-based algorithms that construct graphs in the continuous configuration space of a robot [13]. While successful for single-agent path planning in high-dimensional spaces, these structures are typically constructed a priori or in centralized systems, making them difficult to maintain dynamically in a decentralized multi-agent team. More recent approaches have utilized topological maps and semantic scene graphs to abstract complex environments. Topological maps represent the environment as a set of distinctive places connected by traversable paths, significantly reducing the search space for planning algorithms [14]. Generalized Voronoi Diagrams have also been employed to create safe topological skeletons of environments based on sensor data [15]. By annotating these topological graphs with semantic information, such as the classification of terrain types or the presence of specific obstacles, robots can make more informed navigational decisions [16]. The current research gap lies in constructing and updating these semantic-topological graphs dynamically through the embodied experiences of heterogeneous platforms, where an edge may be deemed highly traversable by a legged robot but completely impassable by a wheeled vehicle.

3. Problem Formulation And System Architecture

3.1 Environmental Modeling And Representation

The unstructured field is modeled as a continuous, bounded three-dimensional workspace containing arbitrary static and dynamic obstacles. Unlike structured indoor environments characterized by planar floors and orthogonal walls, the unstructured field exhibits continuous variations in elevation, surface friction, and obstacle density. We represent this environment mathematically through a dynamically evolving topological graph. The vertices of the graph represent spatial waypoints or regions of interest, while the edges denote the existence of a continuous, traversable trajectory between two vertices. Because the team is heterogeneous, the traversability of an edge is not a single scalar value but a vector of costs specific to each robot class. The cost function for traversing an edge must incorporate both the geometric length of the path and the embodied effort required by the specific morphology. The generalized cost function for a robot traversing an edge is formulated as follows:

$$C_{i,j}^k = \int_0^T (\alpha |v_k(t)|^2 + \beta E_k(p(t), s(t))) dt$$

where the cost between node i and node j for robot class k depends on the continuous velocity profile and the specific environmental resistance. The parameters control the weighting between temporal efficiency and energy expenditure. The embodied effort function relies heavily on the physical state of the robot and the sensed environmental parameters along the trajectory [17]. This formulation guarantees that an aerial vehicle evaluating an edge over a ravine will compute a low energy cost, while a ground vehicle evaluating the same edge will compute an infinite cost, reflecting absolute impassability.

3.2 Agent Kinematic Heterogeneity

The heterogeneous team consists of a set of agents, each belonging to a distinct morphological class, such as rotary-wing aerial vehicles, quadrupedal legged robots, and tracked ground vehicles. Each class possesses unique kinematic and dynamic constraints that define its feasible state space and action space [18]. Aerial vehicles are modeled with six degrees of freedom and are primarily constrained by battery capacity and payload limits, rather than terrain roughness. Their embodied perception is typically global or semi-global, providing top-down semantic segmentation of the environment. Tracked ground vehicles operate strictly on the two-dimensional manifold of the terrain surface. Their mobility is constrained by the maximum negotiable incline and the presence of large discrete obstacles. Quadrupedal robots offer a hybrid mobility paradigm; they are confined to the terrain surface but can step over discrete obstacles and negotiate severe inclines that would immobilize wheeled platforms [19].

Table 1 details the heterogeneous node attributes utilized within the shared semantic-topological graph to accommodate these distinct morphologi

Table 1 Graph Node Attributes for Heterogeneous Agents

Attribute Category	Aerial Vehicle Processing	Tracked Vehicle Processing	Quadrupedal Processing
Spatial Representation	3D Volumetric Waypoint	2D Surface Projection	2.5D Elevation Map Node
Sensing Modality	Downward Semantic Camera	Planar LiDAR Scans	RGB-D Depth Perception
Traversability Metric	Airspace Clearance	Surface Smoothness	Step-size Feasibility

3.3 DECENTRALIZED SYSTEM ARCHITECTURE

To ensure robustness against communication degradation in unstructured fields, the proposed framework employs a fully decentralized system architecture. Each agent maintains a local copy of the semantic-topological graph and its associated attributes. When communication links are established between agents, they synchronize their local graphs through a highly compressed spatial hashing protocol.




Figure 1 Comprehensive System Architecture of the Decentralized Graph

The architecture consists of three primary layers residing on each agent. The lowest layer is the Embodied Perception and Control layer, which handles high-frequency sensory processing, local obstacle avoidance, and proprioceptive feedback. The middle layer is the Graph Management layer, responsible for node generation, edge evaluation, and peer-to-peer synchronization. The highest layer is the Task Allocation and Planning layer, which utilizes the local graph to compute optimal paths and bid on mission objectives. This stratified architecture allows the system to abstract the complex, high-dimensional continuous control problem into a manageable graph-theoretic problem, enabling real-time heterogeneous coordination.

4. Graph-Based Embodied Planning Methodology

4.1 Topological Semantic Graph Generation

The generation of the topological semantic graph occurs incrementally as the heterogeneous agents explore the unstructured field. Initially, the graph consists only of the starting locations of the robots and any known objective locations. As agents move through the environment, they employ a frontier-based exploration strategy augmented by their embodied capabilities. A node is added to the graph whenever an agent identifies a structurally significant location, such as a major intersection in the terrain, a highly elevated vantage point, or a bottleneck in the traversable space. Unlike traditional probabilistic roadmaps that randomly sample the state space, the proposed method utilizes a deterministic, perception-driven node placement strategy [20]. Aerial vehicles, with their superior vantage points, rapidly seed the environment with low-resolution nodes based on visual semantic segmentation, identifying broad regions of dense vegetation, water bodies, or clearings. Ground vehicles subsequently navigate toward these seeded nodes, refining their spatial coordinates and adding high-resolution, morphology-specific edges. This collaborative generation ensures that the graph remains computationally sparse while comprehensively covering the traversable state space [21].

4.2 Embodied Path Optimization

Once the graph is constructed, an agent must determine the optimal sequence of edges to reach a designated target. The embodied path optimization process treats the problem as a shortest-path search over a multi-weighted graph, where the weights are determined by the cost function defined in Section 3.1. However, because the environment is unstructured and dynamic, the edge weights are inherently stochastic and subject to uncertainty. To address this, we formulate the path optimization as a stochastic dynamic programming problem. The objective is to find a policy that minimizes the expected cumulative cost to the destination. The optimization objective is mathematically defined as:

$$J(\pi) = \mathbb{E}_W[\sum_{n=0}^{N-1} \gamma^n C_k(x_n, \pi(x_n)) + \text{Psi}(x_N)]$$

where the expectation is taken over the distribution of uncertain edge weights. The policy determines the optimal next node given the current node. The discount factor ensures convergence in cyclic graphs, and the terminal cost function penalizes failure to reach the target [22]. This embodied optimization allows an agent to weigh the risk of traversing a highly uncertain, albeit shorter, edge against the safety of a longer, well-explored route. For example, a quadruped might choose a steeper, rockier path if its proprioceptive confidence is high, whereas a tracked vehicle would strictly optimize for the flattest terrain.

4.3 Dynamic Edge Weight Updating

The true strength of the embodied planning framework lies in the dynamic updating of edge weights through physical interaction. When an agent traverses an edge, it does not merely follow a geometric line; it physically negotiates the terrain, experiencing slip, vibration, power consumption, and control deviations. These proprioceptive experiences are quantified and used to update the edge cost for that specific morphology. If an edge proves significantly more difficult to traverse than initially estimated by exteroceptive sensors, the agent increases the edge weight in its local graph. This update is then broadcast to the network. Furthermore, through cross-platform semantic translation models, an update from one morphology can infer updates for another. If a quadruped detects severe mud causing high slip, this information updates the quadruped's edge cost but also translates into an infinite cost for a tracked vehicle, preventing the wheeled agent from attempting a highly risky traversal. This dynamic feedback loop bridges the gap between pre-computed plans and the physical reality of unstructured fields.

5. Heterogeneous Coordination Mechanisms

5.1 Cross-Platform Task Allocation

Effective utilization of a heterogeneous team requires assigning tasks to the platforms most physically suited to complete them. The task allocation mechanism is integrated directly into the graph-based framework using an asynchronous, decentralized auction algorithm [23]. Mission objectives, such as deploying a sensor, investigating an anomaly, or transporting material, are represented as specialized nodes or subgraphs within the topological map. When a new objective is identified, an auction is initiated. Each agent computes a bid based on its embodied path optimization cost to reach the objective and its morphological suitability for the task [24]. For example, if the objective requires aerial surveillance of a canopy, ground robots will compute infinitely high suitability costs and effectively abstain from the auction. If multiple platforms are capable, the platform with the lowest combined path and execution cost wins the task. This decentralized approach ensures that the system dynamically reallocates tasks if a robot suffers a hardware failure or if environmental conditions shift, thereby maximizing the resilience of the robotic team [25].

5.2 Conflict Resolution And Deadlock Avoidance

In heterogeneous multi-agent systems operating in constrained unstructured spaces, physical conflicts and deadlocks are inevitable. A conflict occurs when two or more agents plan to occupy the same spatial node or traverse the same narrow edge simultaneously. Deadlocks arise in cyclic

dependency scenarios where agents block each other indefinitely. Traditional multi-agent pathfinding resolves this through prioritized planning or centralized search, which scales poorly in dynamic, uncertain domains [26]. We introduce an embodied conflict resolution mechanism that exploits the morphological differences of the agents. When a spatial conflict is detected on the graph, the agents analyze their relative kinematic flexibility.




Figure 2 Algorithmic Flowchart of the Embodied Conflict Resolution Mechanism

The probability of a deadlock occurring at a specific graph vertex can be modeled using the relative arrival times and the morphological clearance volumes of the conflicting agents:

$$P_{\text{deadlock}} = \exp\left(-\frac{|t_a - t_b|^2}{2\sigma^2}\right) \cdot \text{Phi}(V_a, V_b, W_{\text{edge}})$$

where the temporal overlap and the volumetric constraints dictate the severity of the conflict [27]. In resolving the conflict, priority is given to the agent with the lowest maneuverability in that specific terrain. For instance, if a drone and a quadruped conflict in an open area, the drone can effortlessly adjust its altitude to resolve the conflict. If a quadruped and a tracked vehicle meet on a narrow mountain pass, the quadruped, possessing the ability to step onto rugged shoulders, will yield the main path to the tracked vehicle, which is strictly confined to the continuous flat surface [28]. This morphology-aware resolution significantly reduces the need for long detours and complex re-planning.

6. Experimental Setup and Simulation Environments

6.1 Simulation Framework And Terrain Generation

To rigorously evaluate the proposed graph-based embodied planning framework, extensive simulations were conducted using the Gazebo physical simulation engine integrated with the Robot Operating System (ROS 2). The unstructured field was generated using procedural terrain generation algorithms, creating a 500 by 500-meter environment characterized by varying elevations, random vegetation placement, rubble piles, and discontinuous water features. The environment was designed to force interaction and dependency among the heterogeneous team members. Certain objectives were placed in locations accessible only by quadrupeds, while long-distance reconnaissance targets favored the speed of aerial vehicles. Heavy payload transportation tasks required the use of tracked vehicles through relatively flat, but labyrinthine, valleys.

Table 2 details the specifications and modeled constraints of the heterogeneous robotic platforms used in the simulation.

Table 2 Robotic Platform Specifications in Simulation

Platform Type	Maximum Velocity	Locomotion Constraint	Sensing Horizon
Rotary-Wing UAV	15.0 meters/second	Altitude bounded	100 meters (Semantic)
Quadrupedal Robot	2.5 meters/second	Max incline 35 degrees	20 meters (RGB-D)
Tracked UGV	4.0 meters/second	Max incline 15 degrees	30 meters (LiDAR)

6.2 Hardware Constraints And Physical Modeling

Accurate physical modeling is paramount for validating embodied planning algorithms. The simulation incorporated sophisticated physics parameters for each robot [29]. Ground friction was mapped dynamically to the visual textures of the procedural terrain, causing realistic slip for the tracked vehicle on mud and the quadruped on loose gravel. Furthermore, energy consumption was explicitly modeled. The aerial vehicle experienced rapid battery drain, necessitating optimal pathing and periodic return to mobile charging nodes (represented by the tracked vehicles). The quadruped's energy consumption was calculated as a function of the joint torques required to navigate specific inclines, severely penalizing paths that required continuous steep climbing [30]. These hardware constraints forced the planning algorithm to continuously balance time efficiency, energy preservation, and physical safety.

7. Results, Analysis, and Discussion

7.1 Planning Efficiency And Success Rates

The performance of the proposed graph-based embodied planner was compared against three baseline algorithms: standard centralized A-star search over a high-resolution grid map, continuous Multi-Agent Path Finding (MAPF), and a Multi-Agent Deep Deterministic Policy Gradient (MADDPG) reinforcement learning approach. The primary metrics evaluated were mission success rate (defined as all agents reaching their objectives without critical failure or battery depletion), average planning latency, and total computational overhead.




Figure 3 Comparative Performance Analysis Chart

The results indicate that the graph-based embodied planner significantly outperforms traditional methodologies in unstructured fields. The traditional A-star algorithm struggled with the

dimensionality of the high-resolution grid required to capture morphological constraints, leading to massive computational latency. Continuous MAPF frequently encountered deadlocks in narrow, cluttered passages because it lacked long-horizon topological foresight. The MADDPG approach demonstrated robust local obstacle avoidance but failed to generalize across vastly different procedural terrains, resulting in a low overall mission success rate. By abstracting the environment into a sparse, semantically rich graph, the proposed method reduced planning latency by an order of magnitude while maintaining a high mission success rate through its morphology-aware conflict resolution.

7.2 Robustness In Unstructured Terrains

The true efficacy of the embodied framework is most apparent in its robustness against environmental uncertainty. Table 3 presents a comparative breakdown of algorithmic performance metrics across terrains of varying complexity.

Table 3 Comparative Results in Terrains of Varying Complexity

Terrain Complexity	Proposed Method Success	Standard MAPF Success	Latency Reduction
Low (Flat, sparse obstacles)	98 percent	94 percent	45 percent
Medium (Hills, dense forests)	92 percent	61 percent	78 percent
High (Rubble, steep ravines)	85 percent	12 percent	91 percent

In highly complex terrains, the disparity in performance becomes extreme. Standard MAPF methodologies attempt to find geometrically shortest paths that often exceed the physical capabilities of the specific robots assigned to them, leading to localization failures, getting stuck, or energy depletion [31]. The embodied planner, continuously updating edge weights based on real-time proprioceptive feedback, successfully re-routed agents away from unfeasible paths long before a critical failure occurred. When a wheeled robot detected excessive slip on a designated path, the immediate local graph update prevented subsequent wheeled agents from attempting the same route, demonstrating highly effective decentralized spatial learning.

7.3 Limitations And Practical Challenges

Despite the strong simulated performance, several practical challenges remain. The decentralized synchronization of the topological graph relies on an ad-hoc communication network. In heavily occluded environments, such as deep ravines or dense metallic rubble, communication latency and packet loss can cause significant divergence in the local graphs of individual agents. While the system is designed to handle temporary disconnections, prolonged isolation can lead to redundant exploration or conflicting task assignments. Furthermore, the cross-platform semantic translation model currently relies on predefined heuristic rules to translate the embodied experience of one morphology to another. Developing a purely learning-based translation mechanism that can implicitly understand that "rough terrain for a quadruped equals impassable terrain for a wheeled vehicle" without explicit human-coded rules remains an area for future investigation.

8. Conclusion

This paper has presented a comprehensive graph-based embodied planning framework for heterogeneous robot teams operating in unstructured fields. By intertwining the principles of embodied artificial intelligence with dynamic topological graph generation, the proposed methodology allows diverse robotic platforms to navigate complex terrains in a manner that intrinsically respects their unique physical morphologies. The decentralized architecture ensures robustness in environments where communication is degraded, while the morphology-aware conflict resolution mechanism drastically reduces deadlocks and re-planning latency. Extensive

simulation results validate the superiority of this approach over traditional multi-agent pathfinding algorithms in handling severe environmental unpredictability and kinematic heterogeneity. Moving forward, integrating advanced machine learning techniques for completely unsupervised cross-platform semantic translation and deploying the system on physical hardware in real-world disaster response scenarios will further validate the immense potential of embodied multi-agent systems. The fusion of graph theory and physical embodiment represents a critical step toward achieving true robotic autonomy in the natural, unstructured world.

References

- Peng, Qucheng, Chen Bai, Guoxiang Zhang, Bo Xu, Xiaotong Liu, Xiaoyin Zheng, Chen Chen, and Cheng Lu. "NavigScene: Bridging Local Perception and Global Navigation for Beyond-Visual-Range Autonomous Driving." arXiv preprint arXiv:2507.05227 (2025).
- Brohan, A., Brown, N., Carbajal, J., Chebotar, Y., Dabis, J., Finn, C., et al. (2023). RT-1: Robotics transformer for real-world control at scale. In *Proceedings of Robotics: Science and Systems*.
- Zhang, M., Fang, Z., Wang, T., Lu, S., Wang, X., & Shi, T. (2025). Ccma: A framework for cascading cooperative multi-agent in autonomous driving merging using large language models. *Expert Systems with Applications*, 282, 127717.
- Cheng, H., Luo, H., Liu, Z., Sun, W., Li, W., & Li, Q. (2025). Reinforcement learning based robust volt/var control in active distribution networks with imprecisely known delay. *IEEE Transactions on Computational Social Systems*, 1–12.
- Hu, Y., Xi, A., Xiao, Q., Isaacson, S., Liu, H. X., Vasudevan, R., & Ghaffari, M. (2026). LongNav-R1: Horizon-Adaptive Multi-Turn RL for Long-Horizon VLA Navigation. arXiv preprint arXiv:2602.12351.
- Wang Y, Fang Y, Wang T, et al. Dreamnav: A trajectory-based imaginative framework for zero-shot vision-and-language navigation[J]. arXiv preprint arXiv:2509.11197, 2025.
- Wang Y, Fang Y, Wang T, et al. What Limits Vision-and-Language Navigation?[J]. arXiv preprint arXiv:2605.13328, 2026.
- Liu, C., Zhang, J., Zhang, T., Wang, Y., Zhou, H., & Jin, Q. (2026). HRIBench: Benchmarking Interaction-Centric Human-Robot Collaboration. arXiv preprint arXiv:2607.13056.
- Yang, Y., Hu, H., Li, W., Li, S., Yang, J., Zhao, Q., & Zhang, C. (2023, June). Flow to control: Offline reinforcement learning with lossless primitive discovery. In *Proceedings of the AAAI Conference on Artificial Intelligence* (Vol. 37, No. 9, pp. 10843-10851).
- Yao, Jingyu, et al. "Explainable AI Enhanced Traffic Forecasting for Proactive Autonomous Vehicle Routing." 2026 IEEE Conference on Technologies for Sustainability (SusTech). IEEE, 2026.
- Wang, C., Tan, C., Zhu, L., Yang, R., & Hong, J. (2024, September). Neural Active Sensing Vision and Manipulation for Cooperative Agents. In *2024 5th International Conference on Big Data & Artificial Intelligence & Software Engineering (ICBASE)* (pp. 942-947). IEEE.
- Schulman, J., Levine, S., Abbeel, P., Jordan, M., & Moritz, P. (2015). Trust region policy optimization. In *Proceedings of ICML*.
- Yu, Jiangbo, Raphaël Frank, Luis Miranda-Moreno, Sasan Jafarnejad, Jônatas Augusto Manzolli, Jiyao Wang, and Ali Eslami. *Agentic Vehicles for Human-Centered Mobility: Definition, Prospects, and System Implications*. 2026.
- Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., & Riedmiller, M. (2013). Playing Atari with deep reinforcement learning. arXiv preprint arXiv:1312.5602.

- Lv, Qi, et al. "Spatial-temporal graph diffusion policy with kinematic modeling for bimanual robotic manipulation." *Proceedings of the Computer Vision and Pattern Recognition Conference*. 2025.
- Cheng, Xiaoyuan, et al. "How Does the Lagrangian Guide Safe Reinforcement Learning through Diffusion Models?." *arXiv preprint arXiv:2602.02924* (2026).
- Fox, D., Burgard, W., & Thrun, S. (1997). The dynamic window approach to collision avoidance. *IEEE Robotics & Automation Magazine*, 4(1), 23–33.
- Driess, D., Xia, F., Sajjadi, M. S. M., Lynch, C., Chowdhery, A., Ichter, B., et al. (2023). PaLM-E: An embodied multimodal language model. In *Proceedings of ICML*.
- Yang, Y., Ma, X., Li, C., Zheng, Z., Zhang, Q., Huang, G., ... & Zhao, Q. (2021). Believe what you see: Implicit constraint approach for offline multi-agent reinforcement learning. *Advances in Neural Information Processing Systems*, 34, 10299-10312.
- Luo, Y., Chen, Y., Duan, Z., He, Y., Tang, H., & Pan, C. (2026). Fault Tolerant Control for Heterogeneous UAVs-UGVs Coordinative System with Adaptive Safety Switching based on DMPC. *IEEE Transactions on Aerospace and Electronic Systems*.
- Yang, Y., Wang, Q., Li, C., Hu, H., Wu, C., Jiang, Y., ... & Xu, B. (2025, May). Fewer may be better: Enhancing offline reinforcement learning with reduced dataset. In *International Conference on Learning Representations* (Vol. 2025, pp. 7076-7098).
- Kavraki, L. E., Švestka, P., Latombe, J.-C., & Overmars, M. H. (1996). Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE Transactions on Robotics and Automation*, 12(4), 566–580.
- Xiao, Q., & Ghaffari, M. (2026). LaWM: Least Action World Models for Long-Horizon Physical Consistency from Visual Observations. *arXiv preprint arXiv:2605.08279*.
- Wan, H., Cheng, J., Deng, Y., Wu, D., Chen, Y., Lin, Z., ... & Ji, X. Towards Physics Aware Embodied Control with Graph based Object-centric Learning. *ACM Transactions on Cyber-Physical Systems*.
- Li, B., Zhang, R., Liang, H., Zhang, J., Zhang, J., Chen, X., ... & Wang, J. (2026). Interagent: Physics-based multi-agent command execution via diffusion on interaction graphs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (pp. 15253-15265).
- Luo, P., & Ding, X. (2026). C-MAPPO-TSC: A Collaborative Multi-Agent Reinforcement Learning Framework for Urban Traffic Signal Control and Congestion Mitigation. *Fundamental Scientific Reports in Multidisciplinary Areas*, 2 (02), 170-181.
- Mur-Artal, R., Montiel, J. M. M., & Tardós, J. D. (2015). ORB-SLAM: A versatile and accurate monocular SLAM system. *IEEE Transactions on Robotics*, 31(5), 1147–1163.
- Wan, H., Zhang, Y., Wang, J., Wu, D., Li, M., Chen, X., ... & Ji, X. (2025). Toward universal embodied planning in scalable heterogeneous field robots collaboration and control. *Journal of Field Robotics*, 42(5), 2318-2336.
- Shi, T., ElSamadisy, O., & Abdulhai, B. (2026). CoopSECRM2D: A Safe, Efficient, and Comfortable Multi-Agent Reinforcement Learning Model for Cooperative On-Ramp Merging. *IET Intelligent Transport Systems*, 20 (1), e70256.
- Prescribed-time tracking control of MIMO nonlinear systems with nonvanishing uncertainties
- Zhang, Luyan. "MCP: A Control-Theoretic Orchestration Framework for Synergistic Efficiency and Interpretability in Multimodal Large Language Models." *arXiv preprint arXiv:2509.16597* (2025).