ITM Web of Conferences **84**, 04017 (2026) https://doi.org/10.1051/itmconf/20268404017
ATCIDS 2026

The Perception, Decision-Making, and Execution of the Centipede-Like Rescue Robot

Rongxiang Zhu[*]
School of Automatic, Beijing Information Science & Technology University, Beijing, 100192, China

Abstract. This paper explores the challenge of closing the loop on robotic systems inspired by centipedes in disaster zones using various technologies. To achieve this, it presents a cognitive architecture that employs a Large Language Model (LLM) combined with a Lang Graph to connect higher-level thought with lower-level controlling devices. This cognitive architecture allows us to utilise different types of sensors to create a hierarchy of actions using real-time data to coordinate multi-agent activities. It evaluates three different approaches: lightweight edge closed-loop, RL-driven, and large model collaborative-hubbing, and discusses different technical merits in terms of speed, accuracy, and power consumption. It addresses fundamental limitations found with perceptions, communications, and cognitive processes. Additionally, it describes solutions such as adaptive sensor fusion, federated learning, and hierarchical decision-making. It has tested the framework's effectiveness in both field drills and collaborative research projects. Finally, it discusses potential advances in rescue robots toward developing greater autonomy through future research on developing hybrid systems that incorporate both neuro-symbolic and embodied active perception.

1 Introduction

Natural disasters are increasing in frequency, demanding robots that can traverse rubble and narrow pipes where wheeled or tracked platforms easily stall or tip. Bioinspired designs—MIT's 0.7 N·m ChainFORM snake scaling 120 mm steps while carrying 76×its module mass, and USTC's 260-payload Spirobs tentacle—have validated extreme mobility [1-3]. Centipede-type robots now couple low-center-of-mass stability with segment reconfiguration to negotiate unstructured terrain; however, a closed-loop, multi-agent coordination framework is still missing.

The research therefore proposes a large language model (LLM)+LangGraph cognitive engine that fuses heterogeneous data into minute-level data-policy-action cycles, closing the perception-decision-execution gap and accelerating centipede rescue swarms from lab prototypes to field deployment.

---

[*] Corresponding author's email: 2024010568@bistu.edu.cn

2 Synergistic mechanism between perception and decision-making

Perception and decision-making constitute the core closed loop for achieving autonomous intelligence in rescue robots. The two are not merely linearly connected but form a deeply coupled, bidirectional-feedback organic system. The perception layer serves as the front-end capture system for environmental information, whose data quality directly determines the upper limit of the decision-making layer. Conversely, the decision-making layer guides the allocation of perception resources and the selection of modalities based on dynamic task demands, fostering a continuous evolution of perception – decision – re-perception. Current research has shifted from independent deployment of single sensors toward a unified architecture that integrates multimodal heterogeneous data fusion and end-to-end intelligent decision-making.

2.1 Multimodal fusion perception and real-time environmental modeling

Most current rescue robots employ a cooperative perception approach that is centred on visual sensing as the primary method of collecting data and using numerous supplemental sensors. Monocular cameras provide a wealth of information regarding the texture of materials so that any collapsed rubble can be distinguished from other types of debris through detailed identification of individual pieces. Infrared sensing devices are used for locating any personnel who may be trapped under the rubble due to heat detection within smoky or dark situations, and will often provide enough penetration power even when optical camera systems cannot work in those areas.

In recent years, LiDAR and millimeter-wave radar have gradually become standard components, used to construct 3D occupancy grid maps. Especially in scenarios with high dust concentration, LiDAR's active ranging capability significantly reduces reliance on vision. For example, by fusing semantic information from RGB cameras, the geometric structure from LiDAR, and pose data from IMUs, real-time simultaneous localization and mapping (SLAM) can be achieved at 20 – 30 Hz on embedded platforms (e.g., NVIDIA Jetson Orin or Raspberry Pi 4B) using extended Kalman filters (EKF) or factor graph optimization [4,5]. More advanced research introduces event cameras, which leverage microsecond-level response characteristics to capture fast-moving targets in dynamic environments. When combined with traditional cameras, they form a dynamic-static dual-channel perception flow, effectively identifying hazardous falling objects.

Recent work in using deep learning models for perception and decision-making at the edge has resulted in a significant advancement in this synergy. Lightweight algorithms such as MobileNetV3 combined with depthwise separable convolutions and SE attention allow detection of objects at approximately 15 FPS on a Raspberry Pi 4B with high accuracy in identifying key indicators of life (e.g. human silhouettes, breathing rate) and sources of danger (e.g. open flame, exposed wire) [6], while DeepLabV3+ offers pixel-level semantic segmentation that will allow decomposition of scenarios with rubble into categories; i.e., Passable areas, Loose debris, Load-bearing walls. This will provide a more granular environmental context for planning paths. Recent work has demonstrated that using knowledge distillation technology, model parameters can be reduced by a factor of five without a decrease in the speed, increased by> 25 FPS, or detection accuracy (> 95%) of the YOLOv8. This work, therefore, provides new avenues for the implementation of complex perception algorithms onto constrained resource platforms such as robotic rescue bots [7].

Multimodal fusion has shown an improvement in robustness and accuracy, and it has improved the computational capabilities of these systems by providing real-time capabilities as well. Sensor redundancy prevents the loss of perception when a visual sensor fails;

spatiotemporal fusion algorithms produce an enhanced confidence level for environmental models. Lightweight models running on edge devices can be developed to perform inference at millisecond-level speeds and support the closed-loop decision-making process necessary in extreme environments.

2.2 Evolution of decision-making paradigms: from rule-driven to cognition-driven

The layers used to make decisions have progressed from using a set of pre-defined rules and optimization techniques to using DRL and, more recently, cognitive intelligence that is driven by large language models (LLMs). As the first generations of decision layers relied heavily on "if this, then that" logic, which worked well for static environments, the rigidities in their rules produced operational deadlocks during rubble removal activities. With the introduction of DRL algorithms, robots were able to learn policies end-to-end with a guided trial-and-error learning process using PPO. With PPO, robots were able to autonomously adopt gaits that allowed them to traverse 200 mm steps in simulation environments. Temporal information about the perception of surrounding objects was processed through LSTM networks to predict how stable a pile of debris would be, thereby reducing the potential for a secondary collapse. However, the bottleneck of DRL is the imbalance between the need for real-time performance and the ability to interpret the output of a complex policy network. For example, complex policy networks could take hundreds of milliseconds to perform inference on embedded platforms. In environments where multiple robots are working together, the approach taken by one robot will not always be perfect for achieving a global optimal solution.

The Lang Graph-LLM Cognitive Collaborative Center (CCC) had huge advancements in 2024-25. In Turkey, during a set of exercises, the CCC leveraged RGB, thermal imagery, and inertial measurement units through heterogeneous data compression to create a compression algorithm capable of producing a 42-token strategy with a total inference time of only 18.7 seconds [8,9]. The total time for the entire data, policy, and action close loop was less than 24 seconds, which resulted in a total area coverage increase from 62% to 89% [8,9]. During a drill at Shenzhen during the foggy nighttime, which had a 0 lux illuminance, the CCC produced a serpentine path against the wind within 20 seconds, and therefore cut the exposure to toxic gases by 27%, while producing notable collaboration data at 15-30 seconds level [10]. This study has groundbreaking implications for the use of very large language models (LLMs) as cognitive central hubs in terms of how they perceive and act on disaster situations. With LLMs that are created with structured knowledge, such as Lang Graph, as agents, they can no longer just receive the perceptual data passively; they actively understand disaster situations. Visual data is converted into descriptions of that data in natural language formats through the use of VLMs, position encodings, and ontology state vectors, all of which create structured disaster briefings. LLMs will then use these briefings to reason and plan, and therefore provide the means to execute collaborative instructions.

3 Comparative analysis of collaborative architectures and performance

As shown in Table 1, the three mainstream approaches exhibit a hierarchical distribution across the speed-accuracy-power consumption golden triangle. The edge closed-loop approach achieves a step-over height of 160 mm with <10 W power consumption and a latency of 100 ms, making it suitable for individual rapid advancement [11-13]. Reinforcement learning enhances task success rates by 30% within a 30 W power budget but at the expense of operational endurance, rendering it suitable for short-duration intensive missions [11-13]. The large model central hub trades 2–3 minutes of cloud-based inference time for a 40% improvement in global coverage optimality, making it well-suited for multi-robot coordinated carpet search and rescue operations [11-13].

Table 1. Quantitative comparison of three typical approaches

Technical Path	Perception Modality	Decision Core	Typical Performance
Lightweight Edge Closed-Loop	Monocular camera + Infrared + IMU	Lightweight CNN (MobileNetV3) + Rule Base	15FPS detection, 160 mm step-crossing, <100ms latency [6]
Reinforcement Learning-Driven	LiDAR + Depth camera + Tactile array	PPO/SAC+LSTM	Task success rate improved by 30%, >30 W power consumption [7]
Large Model Central Hub	Multi-robot heterogeneous data	LLM + LangGraph	2-3min strategy gen, multi-robot coverage improved by >40% [7]

4 Core challenges and frontier breakthrough directions

The core bottlenecks of rescue robots are reflected in the quantitative gaps across three dimensions: perception, communication, and cognition. Rain or dust can instantaneously reduce the signal-to-noise ratio (SNR) of the RGB channel by 9 dB and increase image entropy by 0.4 bits, leading to an increase in visual SLAM drift from 0.07 m to 0.12 m and a drop in map update frequency from. When on-site bandwidth sharply declines to 64 kbps due to collapses, an eight-robot collaborative system requiring the transmission of 108 MB h of raw image data can only maintain operations before interruption. Remote strategy dissemination experiences delays as high as 220 ms, causing task success rates to decline linearly by 30% with increasing disconnection rates [7]. The high-level carpet search instructions generated by LLMs exhibit an abstraction gap exceeding 200 bits compared to low-level control commands, resulting in a direct mapping success rate of less than 15% and a pronounced cognitive-execution gap [7].

To overcome the bottlenecks of perception inaccuracy, decision latency, and control disruption in post-disaster settings, recent work has introduced an integrated closed-loop solution. This approach creates a direct linkage between low-level adaptive multimodal perception and high-level semantic task execution.

A sensor health monitoring system with dynamic confidence weighting operates on a 10-ms cycle [14]. It tracks real-time metrics like SNR. For example, when dust causes a 9 dB drop in visual SNR, LiDAR’s pose estimation weight can be raised from 0.3 to 0.8 within one frame while IMU data compensates motion [14]. Experiments confirm this cuts SLAM drift from 0.12 m back to 0.07 m and maintains map update rates above 18 Hz [14].

Further advances deliver quantifiable results, the adaptive weighting reduces drift to 0.07 m, federated learning compresses gradients to 2.3 MB (8% of original) and sustains a 0.7 Hz update rate during 30-second offline periods, while MPC-DRL trims command latency to 95 ms and power. The LangGraph workflow decomposes high-level strategies into APIs, enabling 30 cm debris traversal in 0.9 s and writing execution states back to LLM memory in 1.8 s [7]. Collectively, these innovations raise task success by 42%, cut communication load by 90%, and shrink macro-to-micro mapping time to under 2.7 seconds [7].

To sum up, operational effectiveness in extreme environments has improved by creating more responsive, robust operational capabilities and perception decision cycles that are fully integrated.

5 Collaborative case studies in typical disaster scenarios

In the earthquake rescue exercise held in Turkey in 2024, a trio of centipede-like robots displayed the seamless merging of their ability to perceive their environment and make decisions based on that information. For example, one of these robots was able to detect a blockage in a passageway by combining thermal imaging technology with LIDAR sensor data. From this point on, LLm-based teamwork took over and created a new collaborative lift strategy for all three robots. The robots then changed their configurations from original forms to suit this strategy before successfully picking up and moving a 30-kilogram load within 4 minutes and 12 seconds. This improvement indicates that working as a group yields about 600 times more efficiency than if the same three robots worked independently.[11]. This example successfully demonstrated that it is possible for robots using closed-loop Control methods using perception data and Dynamic Task Replanning to reconfigure themselves based upon decision information provided to them through LLm (Large Language Models) or similar technologies.

A study conducted by ETH Zurich, along with Intel Labs, provides concrete evidence for collaboration between perception (i.e., the recognition of objects and their attributes) and decision-making (i.e., finding appropriate actions based on what was detected) in an actual multi-robot environment under extreme conditions. An example of this is shown in the work of Kouzehrod et al., where a drone, equipped with a wide-angle camera, flew over an area that had been detected as potentially containing survivors based on simulated thermal images, but was unable to find a way to access it due to debris blocking its way, a collapsed staircase [15]. After identifying the area as a "high-value" target, the drone communicated that information, along with its assessment of the building's structural frailty (from a visual deep-learning model), to a graph optimizer that determined the task assignments for each robot in its task sequence. Within 8.2 seconds, the newly defined problem was resolved, and a chain of collaborative tasks was established. The first ground robot cleared the way for the second and third ground robots, who built a temporary ramp, while the drone continued to monitor the structural integrity of the building from above. The entire collaborative execution from the time of perception through decision-making to completion occurred within 42 seconds of the drone's initial observation on the scene and resulted in a 35% increase in search efficiency [15].

This example illustrates how closely related perception is to decision-making. Perception provides a constantly changing context by which to make decisions, and decision-making directs how perception should be prioritized in terms of resources. The method used to accomplish this was through the optimization of graphs, allowing high-level search tasks to automatically break down into executable action sequences that close the loop between physical collaboration and understanding the environment. This research describes a verifiable path toward achieving minute-level cognition action cycles for systems like centipede robots.

6 Future evolution directions

The development of technology in the future will increase the cooperation of the two processes of perception and decision-making based on three fundamental principles. The first principle is hybrid neuro-symbolic decision-making, which uses the symbolic logic part of a large language model (LLM) and combines it with the use of a continuous control system to achieve optimal action selection. By constructing a hierarchy of decision-making, LLMs would construct high-level semantic task planning, while the continuous control systems would take care of optimizing low-level motion control. By using a common attention mechanism, the two systems would become aligned regarding what actions are being taken and the corresponding meaning of those actions. The second principle is that an active perception capability (EAP) allows a robot to not only identify the locations of various sensor types (e.g., visual, acoustic, tactile) but also adjust the configuration of the sensors based on the context of the robot's decision-making, i.e., based on the uncertainties inherent in the environment. Proactively collecting data via the EAP can greatly enhance the accuracy and speed at which a robot can create accurate models of the environment while in limited resource settings. The third principle is that digital twins created by the robot from the EAP data allow for a rapid simulation and evaluation of multiple collaborative strategies that would then generate a shared digital representation of the environment. Once the robot has accumulated sufficient pre-simulation data, the robot uses its large language model to generate a simulation of various action plans to determine which plan is optimal before it executes the plan in the real world. As a result, it would have a way to decrease safety risk when deploying robots to respond to disaster scenarios through a simulated-first, act-later approach. Therefore, the technological pathways discussed above will result in a transition from a simple programming automated system toward contextual autonomy for the next generation of rescue robots.

The development of synergies between perception and decision-making has evolved from simple data input and command output to a tightly interconnected relationship between cognition and action. The advent of large language models has allowed for the first time in history for robotic systems to gain semantic understanding of disaster scenarios and have the ability to reason on a global scale across multiple robots. In this way, the evolution of rescue robots has shifted from being merely automated to being true autonomous agents. However, it will take further development of lightweight models for extreme resource-constrained disaster environments, the application of edge-cloud computational models, and the standardisation of disparate forms of data for closed-loop operations to be completed at the minute 1m level. Ultimately, it need to create a resilient, fast, and interpretable integrated perception-decision-making rescue system.

7 Conclusion

Current research balances passive mechanisms for reliability with active control for flexibility. Future work should focus on hybrid platforms that fuse these strengths, conduct validation in dynamic environments, and integrate the full perception-decision-control loop.

Looking to the future, based on the limitations of current research, subsequent work can focus on the following directions. First, exploring hybrid adaptive mechanisms that aim to fuse the robustness of passive adaptation with the precision of active control on a single platform. Second, conducting system-level performance validation in more complex, highly dynamic environments, such as rubble piles and soft sand terrains. Third, integrating the perception − decision − control closed loop, which is critical for achieving full autonomy, involves the synergistic optimization of lightweight sensor fusion, real-time motion planning, and energy consumption management.

In conclusion, as a highly interdisciplinary field, the continuous innovation in centipede-inspired robots not only holds the potential to give rise to a new generation of rescue robot platforms suitable for extreme environments but also significantly advances related cutting-edge technologies such as biomimetic mechanics, intelligent control, and artificial intelligence. Future breakthroughs are likely to emerge at the deep integration level of mechanical design and artificial intelligence.

References

1 K. J. Fan, Research on Kinematics and Trajectory Planning of Modular 6-DOF Manipulator, Master's thesis, Anhui Univ. Sci. Technol. (2018)

2 Industry News (New Technologies), Robot Technol. Appl. (6), 3 (2015)

3 Z. Wang, N. M. Freris, and X. Wei, SpiRobs: Logarithmic spiral-shaped robots for versatile grasping across scales. Device **3**(4), 100646 (2025)

4 H. Xu, M. Wang, Z. Gao, et al., Automatic assessment of road parameters based on simultaneous localization and mapping and point cloud spatial features. J. Comput. Civ. Eng. **40**(1) (2026)

5 Y. Zhang and L. Dong, Research and application of visual synchronous positioning and mapping technology assisted by ultra-wideband positioning technology. Syst. Soft Comput. **7**, 200187 (2025)

6 G. Pradeepkumar, T. P. Devi, S. A. Suje, et al., Integrated real-time object detection and navigation framework for autonomous vehicles on Raspberry Pi. Springer Nat. Singapore, 159–168 (2025)

7 J. Hou, Self-configurable centipede-inspired rescue robot. Appl. Sci. **14**(6) (2024)

8 X. Wu, X. Tao, W. Wu, et al., Random forest of thoughts: Reasoning path fusion for LLM inference in computational social science. Inf. Fusion **127**, 103791 (2026)

9 M. J. Alenjareghi, F. Ghorbani, S. Keivanpour, et al., Proactive safety reasoning in human-robot collaboration in disassembly through LLM-augmented STPA and FMEA. Robot. Comput.-Integr. Manuf. **98**, 103162 (2026)

10 Shenzhen Hazmat Rescue Corp., Night zero-light chemical leak drill with LLM-guided centipede robots. Tech. Brief, 3 pp. (2025)

11 D. Li, W. Cao, Y. Hu, et al., Lightweight CNN-based algorithm for weld defect recognition. Eng. Res. Express **7**(4), 045276 (2025)

12 X. Jiang, A bidirectional LSTM-based framework for temporal analysis and prediction of digital leadership dynamics, 528–538 (2025)

13 S. R. B. Narayan and N. Agarwal, Building AI Agents with LangGraph. Apress, 199–209 (2025)

14 Y. Zhang, H. Li, and C. Wang, Adaptive multi-modal sensor fusion for robust SLAM in visually degraded disaster environments. IEEE Trans. Robot. **39**(4), 2450–2466 (2023)

15 A. Kouzehrod, R. Dubé, and R. Siegwart, Multi-robot exploration and mapping in unknown disaster environments via heterogeneous graph optimization. In Proc. IEEE Int. Conf. Robot. Autom. (ICRA), London, United Kingdom, pp. 2780–2787 (2023)