

Embodied Intelligence in Disassembly: Multimodal Perception Cross-validation and Continual Learning in Neuro-Symbolic TAMP

Ziwen He¹, Zhigang Wang², Yanlong Peng¹, Pengxu Chang¹, Hong Yang³, Ming Chen^{1*}

Abstract—With the rapid development of the new energy vehicle industry, the efficient disassembly and recycling of power batteries have become a critical challenge for the circular economy. In current unstructured disassembly scenarios, the dynamic nature of the environment severely limits the robustness of robotic perception, posing a significant barrier to autonomous disassembly in industrial applications. This paper proposes a continual learning framework based on Neuro-Symbolic task and motion planning (TAMP) to enhance the adaptability of embodied intelligence systems in dynamic environments. Our approach integrates a multimodal perception cross-validation mechanism into a bidirectional reasoning flow: the forward working flow dynamically refines and optimizes action strategies, while the backward learning flow autonomously collects effective data from historical task executions to facilitate continual system learning, enabling self-optimization. Experimental results show that the proposed framework improves the task success rate in dynamic disassembly scenarios from 81.9% to 100%, while reducing the average number of perception misjudgments from 3.389 to 1.128. This research provides a new paradigm for enhancing the robustness and adaptability of embodied intelligence in complex industrial environments.

I. INTRODUCTION

With the rapid development of Industry 4.0 and the circular economy, industrial disassembly has become a critical link in intelligent manufacturing and resource recycling, facing unprecedented technical challenges [1], [2]. Traditional disassembly operations heavily rely on manual labor, resulting in low efficiency, high costs, and poor safety. Embodied intelligence, by endowing robots with perception, decision-making, and execution capabilities, offers a new technological pathway toward autonomous disassembly. However, the complexity of industrial disassembly scenarios imposes extremely high demands on robots' multimodal perception, task planning, and continual learning capabilities.

Task and motion planning (TAMP), as one of the core technologies of embodied intelligence, achieves closed-loop adaptability in complex scenarios by integrating symbolic task planning based on the Planning Domain Definition Language (PDDL) [3]–[5] with continuous-space motion optimization through a hierarchical architecture [6]. In recent years, Neuro-Symbolic approaches, which combine the perceptual learning capabilities of neural networks with the

logical reasoning power of symbolic systems, have shown great potential in the TAMP domain [7].

In our previous work, we designed a Neuro-Symbolic autonomous disassembly system for robots [8]–[10]. This system dynamically maps multimodal perception to symbolic states through neural predicates, transforming sensory data — such as visual image analysis and force feedback — into discrete symbolic representations. Based on a hierarchical task planning framework, the system autonomously schedules action primitives. The system successfully achieved autonomous disassembly of electric vehicle battery screws and demonstrated good generalization ability in unstructured environments.

However, the existing system struggles to adapt to dynamic changes during operation. First, environmental disturbances can induce perception shifts. Layout changes (e.g., camera displacement or sudden lighting variations) can cause a distribution shift in perception features, leading to symbolic state misjudgments and action sequence conflicts. Second, long-term operational degradation occurs. In continuous disassembly tasks, tool wear gradually causes end-effector positioning drift, destabilizing the perception-planning-execution loop. The current system relies on fixed-parameter models, making it difficult to dynamically adapt to these time-varying disturbances, ultimately limiting the system's robustness for industrial-grade applications in unstructured disassembly tasks.

To address these challenges, this study proposes an adaptive optimization framework that integrates multimodal cross-validation and continual learning (Figure 1). The system tightly couples inference and decision-making with perception and control, embedding the vision-force cross-validation mechanism into the PDDL-based forward working flow and backward learning flow. This framework dynamically refines strategies during task execution and effectively incorporates historical data into continual learning, constructing a closed-loop control system of perception-planning-execution-validation-learning.

Our main contributions are as follows:

- We propose a multimodal cross-validation mechanism that enhances embodied intelligence adaptation in dynamic environments through vision-force perception cross-validation and neural predicate dynamic correction.
- We design a self-optimizing continual learning framework, introducing a PDDL-based bidirectional reasoning flow with forward task execution and backward learning optimization. The system automatically con-

This work was supported by the Ministry of Industry and Information Technology of China for financing this research within the program "2021 High Quality Development Project (TC210H02C)"

¹ School of Mechanical Engineering, Shanghai Jiao Tong University, Shanghai, China. (e-mail: {heziwen, me-pengyanlong, 19119175490, mingchen}@sjtu.edu.cn)

² Intel Labs China, Beijing, China. (e-mail: zhi.gang.wang@intel.com)

³ Intel Asia Pacific R&D Ltd, Shanghai, China. (e-mail: harold.yang@intel.com)




Fig. 1: System framework diagram. The system operates under a closed-loop control architecture composed of five modules: perception, planning, execution, verification, and learning. In the forward working flow, the perception module updates neural predicates based on environmental inputs, the planning module generates action planning sequences, and the execution module carries out these sequences. During task execution, the verification module dynamically adjusts task strategies through vision-force cross-validation. In the backward learning flow, the learning module infers from historical records in the verification module, automatically collecting valid data through the cross-validation mechanism to expand the continual learning dataset and optimize the performance of the perception module.

structs training datasets from historical robot tasks, reducing dependence on manual annotations and achieving continual model parameter optimization.

- We conduct experiments in a real-world battery disassembly scenario to validate the proposed continual learning framework. The results show that the screw disassembly success rate increases from 81.9% to 100% under environmental disturbances, while the average number of perception misjudgments per task decreases from 3.389 to 1.128. These results demonstrate the robustness and practicality of the framework in unstructured dynamic environments, providing a technical

paradigm for the industrial deployment of embodied intelligence.

II. RELATED WORK

A. Neuro-Symbolic TAMP

Artificial intelligence has evolved from symbolic reasoning to deep learning-based probabilistic approaches. While symbolic methods offer strong interpretability, they struggle with complex and dynamic environments. In contrast, deep learning excels at perception tasks but relies heavily on large-scale data and lacks reasoning capabilities. Neuro-Symbolic AI, which integrates the strengths of both paradigms, has

emerged as a promising research direction [11], [12].

Neuro-Symbolic TAMP is a representative application of this paradigm. It uses neural predicates to model environment states and action primitives to bridge high-level task descriptions with low-level motion execution, enabling robust planning in unstructured environments [7], [10], [13].

Despite improvements in flexibility and efficiency, current Neuro-Symbolic TAMP systems still face challenges, such as perception errors affecting symbolic decision-making and limited adaptability to dynamic scenarios. A key issue is how to effectively integrate multimodal perception and adopt continual learning to dynamically refine neural-symbolic mappings, thereby enhancing system robustness.

B. Continual Learning and Autonomous Data Acquisition

Continual learning (CL) enables robots to adapt to evolving tasks in dynamic environments while retaining previously acquired knowledge and mitigating catastrophic forgetting [14], [15]. Current CL strategies encompass a variety of approaches, including gradient-based regularization [16], modular decomposition [17], memory-augmented methods [18], and meta-learning [19]. These techniques demonstrate different strengths in interference mitigation, parameter isolation, and fast adaptation. Increasingly, they are being integrated with deep neural networks to form diverse and robust CL frameworks.

In terms of data acquisition, weakly-supervised learning has emerged as a promising solution to reduce reliance on manual annotations. For example, Barnes et al. proposed a weakly-supervised segmentation method for autonomous driving that leverages vehicle trajectories to automatically generate large-scale labeled datasets for path and obstacle detection [20]. Inspired by such approaches, our work adopts a task-driven data acquisition paradigm, where effective training data is collected based on the robot's own behavior and perception, facilitating self-supervised optimization without human labeling.

C. Continual Learning in Neuro-Symbolic TAMP

Recent research has introduced continual learning into Neuro-Symbolic TAMP frameworks to enhance adaptability in open and dynamic environments. For instance, HyGOAL [21] integrates symbolic planning with reinforcement learning to enable robots to quickly adapt to novel situations such as environmental disturbances and primitive failures. However, existing approaches primarily focus on strategy-level generalization and often lack mechanisms to support continual learning in the perceptual components.

To address these limitations, our proposed Neuro-Symbolic TAMP system incorporates a multimodal cross-validation mechanism based on visual and force feedback. This allows the system to continuously refine both its symbolic state representations and perceptual models without requiring external annotations. By bridging the gap between neural perception and symbolic reasoning, our approach improves planning robustness and environmental adaptability in complex robotic tasks.

TABLE I: Detailed Description of (Neural) Predicates

Predicate	Function Description
have_coarse_pose	The robot has obtained the coarse position of the screw.
pattern	1 indicates visual perception, 0 indicates force perception.
near_screw	The end-effector is close to the screw.
above_screw	The end-effector is near the top of the screw.
disassembled	The screw has been disassembled.
target_aim	Visual perception determines that the socket is aligned with the screw, meeting the engagement requirement.
socketed	Force perception confirms the socket successfully engages with the screw, meeting the disassembly requirement.

TABLE II: Detailed Description of Action Primitives

Primitive	Function Description
Move	Moves the robotic arm to position the end-effector near the top of the target screw.
Mate_vision	Moves the arm above the screw, estimates the precise pose via visual perception, and repositions the end-effector to align its center axis with the screw.
Mate_force	Estimates the precise screw pose through force feedback and moves the arm to align the end-effector's axis with the screw.
Insert	Lowers the end-effector along its axis while rotating to engage the socket with the screw.
Disassemble	Rotates the end-effector counterclockwise to unscrew and disassemble the target screw.

III. PROBLEM DEFINITION

We formulate the system-level task as a planning problem defined by a tuple (S_0, S_G, A) , where S_0 denotes the initial state, S_G the goal state, and $A = \{a_1, a_2, \dots, a_n\}$ the set of action primitives, with a_i ($i = 1, \dots, n$) representing individual action primitives. The objective of the planning algorithm is to find an executable sequence of action primitives $\text{prim_list} = \{a_{p_1}, a_{p_2}, \dots, a_{p_m}\}$ such that applying these actions in order transitions the system from S_0 to S_G .

In this work, we consider the disassembly of electric vehicle (EV) battery screws as a representative application scenario. While this task is trivial for humans, autonomous execution in unstructured environments remains highly challenging for robotic systems. To enable effective planning and execution, multimodal sensory inputs—comprising visual and force data—are processed by neural networks to infer quasi-symbolic state representations in the form of neural predicates (Table I).

Inspired by human disassembly behavior, we design five core action primitives (Table II). Each primitive is defined in a PDDL-like formalism using preconditions and expected effects grounded in the neural predicate space (Table III), enabling integration with symbolic task planning frameworks.




Fig. 2: **Illustration of the bidirectional reasoning flow, consisting of a forward working flow and a backward learning flow.** In the forward working flow, the system generates an action primitive sequence based on the initial and goal states, and sequentially executes operations such as Move, Mate_vision, and Insert. When the Insert action fails to achieve the expected socketed state, the system triggers replanning, switches to force perception to re-estimate the screw pose, and successfully completes the remaining task. After task completion, the backward learning flow retrospectively analyzes the execution process. The successful screw pose estimated by force perception is treated as ground truth to correct the biased visual perception result. Meanwhile, the system identifies and flips the neural predicate with the lowest confidence among the preconditions of the Insert action, and then terminates the traceback.

TABLE III: PDDL Representation of Action Primitives: Preconditions and Expected Effects

Primitive	Preconditions and Expected Effects
Move	:pre have_coarse_pose $\wedge$ $\neg$ near_screw :eff near_screw $\wedge$ above_screw
Mate_vision	:pre $\neg$ pattern $\wedge$ near_screw $\wedge$ $\neg$ target_aim :eff pattern $\wedge$ above_screw $\wedge$ target_aim
Mate_force	:pre pattern $\wedge$ near_screw $\wedge$ $\neg$ above_screw $\wedge$ $\neg$ target_aim :eff $\neg$ pattern $\wedge$ above_screw $\wedge$ target_aim
Insert	:pre target_aim $\wedge$ above_screw :eff socketed $\wedge$ $\neg$ above_screw
Disassemble	:pre socketed $\wedge$ $\neg$ disassembled :eff disassembled

IV. METHOD

A. Methodology Overview

We propose a continual learning framework built upon a Neuro-Symbolic TAMP architecture. The framework integrates a multimodal perception cross-validation mechanism and establishes a bidirectional reasoning flow, composed of a forward working flow and a backward learning flow. It is designed to enable robots to robustly execute tasks in dynamic environments while continuously accumulating knowledge from execution feedback and optimizing internal

models, thereby enhancing the system's autonomous learning and adaptive capabilities.

B. Multimodal Perception Cross-validation Mechanism

During robotic task execution, a single perception modality is often vulnerable to environmental noise, leading to unstable perception results and consequently affecting the accuracy and robustness of task performance. To enhance the system's adaptability in dynamic environments, we introduce a multimodal perception cross-validation mechanism, which dynamically corrects perception errors and optimizes action strategies through complementary information between modalities.

For clarity, the two perception modalities involved in this mechanism are denoted as modality A and modality B. The system first relies on modality A to estimate the precise pose of the target object (e.g., a screw) and executes the corresponding sequence of planned action primitives. After a critical action (such as Insert) is performed, the system invokes modality B to sense and update the current system state, determining whether it aligns with the expected reasoning outcome. If the state is consistent, the system proceeds to subsequent actions (e.g., Disassemble); otherwise, it triggers replanning. Modality B is then used to re-estimate the target pose, and the updated action sequence is executed, enabling a closed-loop correction between perception and

planning.

In this work, we select vision and force as two representative perception modalities. The visual modality uses RGB-D images as input, while the force modality captures variations in force signals during exploratory contact. The system fuses these multimodal observations to update neural predicates and generate accurate pose predictions for the screw.

C. Forward Working Flow

The forward working flow models the system's task execution logic and serves as the core process for planning and verification. In this flow, multimodal perception data are mapped into interpretable symbolic states through neural predicates. The system formalizes both the current and goal states using PDDL, where each action primitive is defined by its preconditions and expected effects.

Given the initial state S_0 , the goal state S_G , and a set of action primitives A , the PDDL planner generates an executable sequence of primitives, denoted as `prim_list`. The robot attempts to execute this sequence in order. Before executing each primitive, the system verifies whether the current symbolic state satisfies its preconditions using real-time multimodal sensing. If the condition holds, the primitive is executed; otherwise, the current state is updated as a new start state S'_0 , and the planner replans a new sequence `prim_list'` based on (S'_0, S_G, A) .

Throughout the execution, the multimodal perception cross-validation mechanism continuously monitors the state to ensure consistency with the expected symbolic transitions. This enables dynamic correction of the action plan and robust adaptation to environmental changes or sensing inaccuracies.

D. Backward Learning Flow

The backward learning flow enables continual self-improvement by tracing execution failures and correcting perception and symbolic reasoning errors. Unlike the forward working flow, which focuses on task planning and execution, the backward learning flow focuses on learning from mistakes and optimizing model parameters.

After each task execution, the system retrospectively analyzes the forward working flow. Two main correction mechanisms are applied: perception estimation correction and neural predicate correction.

1) *Perception Estimation Correction*: When a task is successfully completed, the final perception result (e.g., screw pose) is treated as ground truth, which is assumed to be obtained from the more reliable perception modality A. If, during execution, the system previously used another modality B to produce a biased estimation, the ground truth can be used to supervise and update the model of modality B.

Let the true pose from modality A be:

$$\mathbf{p}_{\text{true}} = [x, y, z, \theta] \quad (1)$$

The estimated pose from modality B is given by:

$$\mathbf{p}_B = f_B(I_B; \theta_B) \quad (2)$$

where f_B is the model for modality B, with input I_B and parameters θ_B . We minimize:

$$L_B = \frac{1}{2} \|\mathbf{p}_{\text{true}} - \mathbf{p}_B\|_2^2 \quad (3)$$

and update:

$$\theta_B^{t+1} = \theta_B^t - \eta \nabla_{\theta_B} L_B. \quad (4)$$

2) *Neural Predicate Correction*: When an action fails to achieve its expected effects, the system backtracks through its preconditions to identify and correct unreliable neural predicates. Assuming that non-neural predicates are reliable, let the set of m neural predicates involved in the preconditions be denoted as $\{P_1, P_2, \dots, P_m\}$. Each predicate P_i is predicted by a neural network classifier $g_i(x; \phi_i)$, where x represents the input perceptual data and ϕ_i denotes the classifier parameters. The classifier outputs a probability distribution:

$$g_i(x; \phi_i) = [p_i^0, p_i^1] \quad (5)$$

where p_i^0 and p_i^1 represent the probabilities that predicate P_i is false (0) and true (1), respectively. The confidence score is defined as:

$$c_i = \max(p_i^0, p_i^1) \quad (6)$$

The system identifies the neural predicate P_k with the lowest confidence score based on the outputs of the neural network classifiers:

$$k = \arg \min_{i=1, \dots, m} c_i. \quad (7)$$

The prediction for P_k is then flipped to produce corrected label $\hat{y}_k$, and parameters ϕ_k of its classifier g_k are updated by minimizing cross-entropy:

$$\hat{y}_k = 1 - y_k, \quad y_k = \arg \max(p_k^0, p_k^1). \quad (8)$$

$$L_k = -[\hat{y}_k \log(p_k^1) + (1 - \hat{y}_k) \log(p_k^0)] \quad (9)$$

$$\phi_k^{t+1} = \phi_k^t - \eta \nabla_{\phi_k} L_k. \quad (10)$$

The process stops after a single correction to avoid introducing unreliable supervision signals into the training set.

E. Illustrative Case

Figure 2 illustrates the application of the proposed bidirectional reasoning flow in an electric vehicle battery screw removal task, comprising both the forward working flow and the backward learning flow.

In this task, the forward working flow begins with the initial state $S_0 = \neg \text{pattern} \wedge \text{have_coarse_pose}(\text{pose}) \wedge \neg \text{near_screw}$, the goal state $S_G = \text{disassembled}$, and a set of action primitives $A = \{\text{Move}, \text{Mate_vision}, \text{Mate_force}, \text{Insert}, \text{Disassemble}\}$. Based on these inputs, the PDDL planner generates an initial sequence of action primitives `prim_list` = {Move, Mate_vision, Insert, Disassemble}. The robot executes this sequence in order. During `Mate_vision`, the system estimates the screw pose using visual perception. Before executing the `Insert` action, the system evaluates its two neural predicate preconditions, `target_aim` and `above_screw`, as true, and proceeds with execution.

However, the `Insert` action fails, as the force perception module indicates that the screw is not successfully socketed, i.e., the expected `socketed` predicate is not achieved. The forward working flow then triggers an abnormal state handler, halting the current execution. Based on the updated state $S'_0 = \text{pattern} \wedge \text{near_screw} \wedge \neg \text{above_screw} \wedge \neg \text{target_aim}$, the planner generates a new action sequence $\text{prim_list}' = \{\text{Mate_force}, \text{Insert}, \text{Disassemble}\}$. The robot executes the new sequence in order. During `Mate_force`, the system re-estimates the screw pose using force perception, eventually completing the task successfully.

After the forward working flow concludes, the backward learning flow retrospectively analyzes the execution. Since the task is successfully completed based on force perception, the resulting screw pose is regarded as ground truth. The system detects that a biased visual estimation was used earlier in the execution. Hence, the system corrects the visual estimation using the ground truth and stores the corrected result along with its corresponding input sample in the continual learning dataset.

Furthermore, because the initial execution of `Insert` failed, the system infers that some of its preconditions may have been misjudged. Upon review, it identifies that among the neural predicate preconditions, `target_aim` has the lowest confidence score. Thus, this predicate is assumed to be incorrectly classified, and its label is flipped. The corrected label and its corresponding input are then stored in the continual learning dataset.

Once a sufficient volume of such corrected samples accumulates in the dataset, the system performs supervised updates of the corresponding model parameters. This process improves the accuracy of perception and the adaptability of the system in dynamic environments.

V. EXPERIMENTS AND RESULTS

A. Experimental System Architecture

Figure 3 shows the experimental system architecture, including the following core components:

- Retired onboard battery pack as the disassembly target.
- UR10e six-degree-of-freedom collaborative robotic arm.
- End-effector integration module:
 - Custom screw removal tool.
 - Intel RealSense depth camera for visual sensing.
- ROS-based hybrid symbolic-reactive control system.

B. Experimental Setup

To evaluate the robustness and adaptability of the proposed Neuro-Symbolic TAMP continual learning framework in dynamic and uncertain environments, we design a series of battery screw disassembly tasks that integrate multimodal perception and symbolic planning. Each task aims to successfully remove a single screw. The system deploys both visual perception networks and force perception networks to perform cross-validation and complementary correction.

The perception module updates the predicates defined in Table I based on either visual or force feedback, which in




Fig. 3: Experimental System Architecture

turn guides the PDDL-based symbolic planner in selecting action primitives (see Table II). The experimental scenario is designed to reflect realistic operational conditions, particularly lighting variation, which greatly affects the accuracy of visual predictions while having limited impact on force perception.

a) Replanning Strategy:: During task execution, if an action (e.g., `Insert`) fails to achieve its expected effect (e.g., the target state `socketed` is not satisfied), the system triggers a replanning mechanism. It regenerates an updated sequence of action primitives based on the new state in an attempt to recover task execution.

To prevent excessive recovery attempts due to perception errors, we define a replanning threshold $n_{th} = 10$. When the number of replans n exceeds this threshold, the task is marked as failed, and the system proceeds directly to the next task without further replanning.

b) Perception Mismatch and Backward Correction:: For tasks that are eventually successful but require at least one replan ($0 < n \leq n_{th}$), the system initiates a backward learning flow to retrospectively correct perception errors:

- 1) **Cross-modal Correction:** The final successful pose estimation (typically based on force perception) is regarded as ground truth. This ground truth is used to correct earlier pose predictions from another modality (e.g., vision). The corrected label, along with its corresponding input data, is stored in the continual learning buffer.
- 2) **Neural Predicate Correction:** The system backtracks through the action sequence and identifies the first failed action. Among its neural preconditions, the one with the lowest confidence is considered erroneous and its output is flipped to form a corrected training label. This

TABLE IV: Task Success Rate and Average Replans

Iterations	Task Count	SUS	$\bar{n}$
0	131	81.68%	3.389
1	162	89.51%	2.648
2	62	100.00%	1.984
3	39	100.00%	1.128

corrected sample is also stored in the continual learning buffer.

c) *Continual Learning Trigger*:: When the number of newly collected samples in the buffer of any perception module reaches a predefined threshold (e.g., 75), the system triggers an incremental model update. The update process is autonomous and does not interrupt task execution, enabling continual improvement of the perception accuracy.

C. Evaluation Metrics

We define the number of replans n as a key metric to assess the accuracy of the perception-to-symbol mapping. A replan indicates that the system's perception is insufficient to support the successful completion of the current task. If n exceeds the threshold n_{th} , the task is considered a failure. The binary success indicator function for each task is defined as:

$$\mathcal{I}(n) = \begin{cases} 1, & n < n_{th} \quad (\text{Task Success}) \\ 0, & n \geq n_{th} \quad (\text{Task Failure}) \end{cases} \quad (11)$$

a) *Quantitative Evaluation Metrics*:: To quantitatively evaluate the performance improvement brought by the continual learning mechanism, we record the replanning count n and calculate the following two metrics after each round of model update:

- **Task Success Rate:**

$$\text{SUS} = \frac{1}{N} \sum_{k=1}^N \mathcal{I}(n_k) \quad (12)$$

which reflects the proportion of tasks that are successfully completed within the replanning threshold.

- **Average Number of Replans:**

$$\bar{n} = \frac{1}{N} \sum_{k=1}^N n_k \quad (13)$$

which measures the system's average reliance on replanning and indicates the stability of perception.

D. Results and Discussion

Table IV presents the task success rate SUS and average number of replans $\bar{n}$ over four system iterations. The number of iterations reflects the number of continual learning updates applied to the neural perception modules.

The results demonstrate that the continual learning mechanism enables the system to refine neural-symbolic mappings by leveraging real task outcomes. With each update, the system becomes more reliable: SUS rises to 100% after two updates, and $\bar{n}$ reduces by over 60%. This illustrates

improved perception confidence and fewer symbolic mismatches.

In particular, the number of perception-induced replans drops significantly due to better prediction from corrected neural predicates. The system learns to recover from early-stage misjudgments (e.g., misestimated `target_aim`) and reduces dependence on replanning. Moreover, experiments show that integrating multi-modal corrections improves robustness to environmental disturbances such as lighting variations.

These findings confirm that our proposed bidirectional reasoning flow, which combines symbolic error recovery and continual perception refinement, significantly enhances task success and reduces planning instability in unstructured robotic manipulation.

VI. CONCLUSIONS

In this paper, a neural-symbolic TAMP framework integrating multimodal perceptual cross-validation and continual learning is proposed to address the dynamic adaptation challenges of autonomous screw disassembly of power batteries in unstructured scenarios. By constructing a forward working flow and a backward learning flow based on PDDL, the system achieves dynamic correction of perceptual bias and online optimization of symbolic knowledge, which solves the problem of perceptual-symbolic mapping error and performance degradation under environmental perturbation in the traditional method. Experimental validation shows that the proposed framework significantly improves task success rates in dynamic disassembly scenarios (increasing from 81.9% to 100%) and reduces the number of perception misjudgments (decreasing from 3.389 to 1.128).

However, to ensure the smooth interaction between the forward working flow and the backward learning flow, the current system still relies on carefully designed predicates and action primitives, which can pose a considerable engineering burden in complex tasks. Moreover, the backward learning flow assumes that action execution is error-free and that state transitions are fully traceable—assumptions that often do not hold in real-world unstructured environments. Furthermore, during the perception correction process, the system selects the neural predicate with the lowest confidence as the most likely misjudgment for label correction. However, low confidence does not necessarily indicate an incorrect prediction, which may lead to incorrect supervision and degrade the quality of model updates. In addition, when all perception modalities are simultaneously affected by noise or interference, the system may be unable to obtain reliable ground-truth estimations from any single modality, making it difficult to correct the others. These limitations reduce the effectiveness of continual learning and compromise the system's robustness and adaptability in complex environments.

In future work, we plan to incorporate action success rates and state transition probabilities into the framework to enable more robust probabilistic reasoning. We also aim

to explore the integration of continual learning with adaptive action primitive generation based on large language models (LLMs). By leveraging LLMs' knowledge and generalization abilities, the system can dynamically generate and refine action primitives during execution, adapting to novel environments. LLMs can further assist in defining new symbolic operations and optimizing primitive parameters through multimodal feedback, enabling more autonomous skill acquisition and improving the system's generalization and decision-making capabilities in complex scenarios.

REFERENCES

- [1] Ian Kay, Siamak Farhad, Ajay Mahajan, Roja Esmaeeli, and Sayed Reza Hashemi. Robotic disassembly of electric vehicles' battery modules for recycling. *Energies*, 15(13):4856, 2022.
- [2] Sabri Baazouzi, Felix Paul Rist, Max Weeber, and Kai Peter Birke. Optimization of disassembly strategies for electric vehicle batteries. *Batteries*, 7(4):74, 2021.
- [3] Malik Ghallab, Craig Knoblock, David Wilkins, Anthony Barrett, Dave Christianson, Marc Friedman, Chung Kwok, Keith Golden, Scott Penberthy, David Smith, Ying Sun, and Daniel Weld. Pddl - the planning domain definition language. *Technical Report, Tech. Rep.*, 1998.
- [4] Maria Fox and Derek Long. Pddl2. 1: An extension to pddl for expressing temporal planning domains. *Journal of artificial intelligence research*, 20:61–124, 2003.
- [5] Håkan LS Younes and Michael L Littman. Ppddl1. 0: An extension to pddl for expressing planning domains with probabilistic effects. *Techn. Rep. CMU-CS-04-162*, 2:99, 2004.
- [6] Caelan Reed Garrett, Rohan Chitnis, Rachel Holladay, Beomjoon Kim, Tom Silver, Leslie Pack Kaelbling, and Tomás Lozano-Pérez. Integrated task and motion planning. *Annual review of control, robotics, and autonomous systems*, 4(1):265–293, 2021.
- [7] Shiqi Shen, Soundarya Ramesh, Shweta Shinde, Abhik Roychoudhury, and Prateek Saxena. Neuro-symbolic execution: The feasibility of an inductive approach to symbolic execution. *arXiv preprint arXiv:1807.00575*, 2018.
- [8] Hengwei Zhang, Hua Yang, Haitao Wang, Zhigang Wang, Shengmin Zhang, and Ming Chen. Autonomous electric vehicle battery disassembly based on neurosymbolic computing. In *Proceedings of SAI Intelligent Systems Conference*, pages 443–457. Springer, 2022.
- [9] Shengmin Zhang, Yisheng Zhang, Zhigang Wang, Hengwei Zhang, Kai Gu, Yanlong Peng, and Ming Chen. Design and implementation of a multifunctional screw disassembly workstation. In *International Conference on Intelligent Robotics and Applications*, pages 506–519. Springer, 2023.
- [10] Yisheng Zhang, Hengwei Zhang, Zhigang Wang, Shengmin Zhang, Huaicheng Li, and Ming Chen. Development of an autonomous, explainable, robust robotic system for electric vehicle battery disassembly. In *2023 IEEE/ASME International Conference on Advanced Intelligent Mechatronics (AIM)*, pages 409–414. IEEE, 2023.
- [11] Zachary Susskind, Bryce Arden, Lizy K John, Patrick Stockton, and Eugene B John. Neuro-symbolic ai: An emerging class of ai workloads and their characterization. *arXiv preprint arXiv:2109.06133*, 2021.
- [12] Pascal Hitzler and Md Kamruzzaman Sarker. Neuro-symbolic artificial intelligence: The state of the art. 2022.
- [13] Hengwei Zhang, Yisheng Zhang, Zhigang Wang, Shengmin Zhang, Huaicheng Li, and Ming Chen. A novel knowledge-driven flexible human-robot hybrid disassembly line and its key technologies for electric vehicle batteries. *Journal of Manufacturing Systems*, 68:338–353, 2023.
- [14] Raia Hadsell, Dushyant Rao, Andrei A Rusu, and Razvan Pascanu. Embracing change: Continual learning in deep neural networks. *Trends in cognitive sciences*, 24(12):1028–1040, 2020.
- [15] Benedikt Pfülb, Alexander Gepperth, Saad Abdullah, and André Kilian. Catastrophic forgetting: still a problem for dnn. In *Artificial Neural Networks and Machine Learning–ICANN 2018: 27th International Conference on Artificial Neural Networks, Rhodes, Greece, October 4-7, 2018, Proceedings, Part I* 27, pages 487–497. Springer, 2018.
- [16] Claudia Clopath, Lorric Ziegler, Eleni Vasilaki, Lars Büsing, and Wulfram Gerstner. Tag-trigger-consolidation: a model of early and late long-term-potential and depression. *PLoS computational biology*, 4(12):e1000248, 2008.
- [17] Rahaf Aljundi, Marcus Rohrbach, and Tinne Tuytelaars. Selfless sequential learning. *arXiv preprint arXiv:1806.05421*, 2018.
- [18] Arslan Chaudhry, Albert Gordo, Puneet Dokania, Philip Torr, and David Lopez-Paz. Using hindsight to anchor past knowledge in continual learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, pages 6993–7001, 2021.
- [19] Eleni Triantafyllou, Tyler Zhu, Vincent Dumoulin, Pascal Lamblin, Utku Evci, Kelvin Xu, Ross Goroshin, Carles Gelada, Kevin Swersky, Pierre-Antoine Manzagol, et al. Meta-dataset: A dataset of datasets for learning to learn from few examples. *arXiv preprint arXiv:1903.03096*, 2019.
- [20] Dan Barnes, Will Maddern, and Ingmar Posner. Find your own way: Weakly-supervised segmentation of path proposals for urban autonomy. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 203–210. IEEE, 2017.
- [21] Pierrick Lorang, Shivam Goel, Yash Shukla, Patrik Zips, and Matthias Scheutz. A framework for neurosymbolic goal-conditioned continual learning in open world environments. In *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 12070–12077. IEEE, 2024.