ICRA2026 Late Breaking Results Poster presented at
2026 IEEE International Conference on Robotics and Automation (ICRA 2026)
June 1-5, 2026. Vienna, Austria

An Efficient Learning-Based Task Planning Approach Using a Bio-Inspired Action Context-Free Grammar for Bimanual Manipulation

David Carmona, Jun Yang and Haoyong Yu

I. INTRODUCTION

Bimanual robots are increasingly deployed in household environments to perform manipulation activities. These activities require reasoning over discrete symbolic actions ('task planning') to determine the actions to achieve and their order. However, task planning alone is insufficient for real-world manipulation. Symbolic actions must be geometrically feasible for the robot within its environment. Task and Motion Planning (TAMP) addresses this challenge by integrating task planning with motion planning to generate executable task plans.

Conventional TAMP faces a major limitation: the combinatorial explosion that arises at the task planning level when dealing with large-scale planning problems [1]. Combinatorial explosion refers to the exponential growth of the symbolic state space as problem parameters increase. This growth renders task plan exploration intractable and leads to long planning times. This challenge is especially pronounced in bimanual manipulation, where two end-effectors further expand the state space through additional symbolic variables related to hand roles, grasp states, and action ordering. Planning latency can negatively affect human-robot interaction [2] and increase the likelihood that the human will abandon the interaction. Hence, efficient task planners for bimanual robots are needed.

Learning-based task planning offers a promising alternative for capturing task-level structure directly from human behaviour. Human demonstrations encode structured knowledge of action sequencing, hand coordination, and goal achievement, which can be leveraged to mitigate combinatorial explosion. However, existing methods either continue to rely on classical task planners for inference, focusing on learning task constraints from multiple demonstrations or on action perception and recognition rather than on efficient plan inference. More recent approaches based on Vision Language Models (VLMs), can generate task plans but prioritise task success and multimodal reasoning. Their efficiency on embedded hardware remains to be studied.

This work presents **BAG-Learn Planning**, a novel and efficient task **planning** approach that employs a Long-Short-Term Memory (LSTM) network to **learn** symbolic task plans derived from annotated human demonstration videos. The task plans are expressed using the rules of a **Bio-Inspired Action** Context-Free **Grammar** (**BAG**) [3]. That grammar builds on the minimalist grammar of action originally proposed by Pastra and Aloimonos [4], providing a structured and neurobiologically grounded representation of human manipulation actions. As Figure 1 shows, BAG-Learn Planning is integrated with a Rapidly Exploring Random Tree (RRT) via a Task Plan Execution Framework (TPEF) to form a complete TAMP framework for bimanual manipulation. The TAMP framework is deployed into an in-house bimanual robot to achieve three activities: pouring, passing, and opening. The code is available at `https://github.com/davidevdual/LfD_Planner.git`.

This work was supported by the Ministry of Education of Singapore (MOE), Grant MOE-T2EP50222-0010.

The authors are with the Department of Biomedical Engineering at the National University of Singapore, Singapore.

II. METHODS

A. Bio-Inspired Action Context-Free Grammar Design

Symbolic task plans are encoded using the BAG. The grammar defines a vocabulary and production rules tailored to bimanual manipulation. Each task plan is represented as a structured sentence composed of action primitives, tool complements, object complements, and a goal. These elements are rooted in neurobiological evidence. BAG enforces syntactic structure, preserves goal-directed semantics, and supports bimanual coordination by encoding the action roles of both hands.

B. BAG-Learn Planning Approach

The BAG-Learn Planning approach comprises two core elements: the Bio-Inspired Action Context-Free Grammar for symbolic task representation and a Long-Short-Term Memory network for task plan inference. The LSTM is trained on Bicap [3], a dataset of 4,026 human-demonstrated bimanual manipulation actions covering household activities such as pouring, opening, and passing. The objects are placed at discrete locations, called *symbolic locations*, on a workspace grid. Each demonstration is annotated as a symbolic task plan using the rules of the BAG. By training on these BAG-encoded sequences, the LSTM infers complete task plans given a high-level goal. This formulation replaces symbolic search with sequence prediction to enable efficient inference of task plans for bimanual manipulation.

C. Task and Motion Planner Design

The inferred task plan is processed by a Task Plan Execution Framework consisting of a parser, an interpreter, and a motion planner. The parser converts the symbolic task plan into an Abstract Syntax Tree (AST). The interpreter iterates over the AST and derives a sequence of motion planning commands, which are executed using the RRT motion planner. The TPEF enables the integration of the BAG-Learn

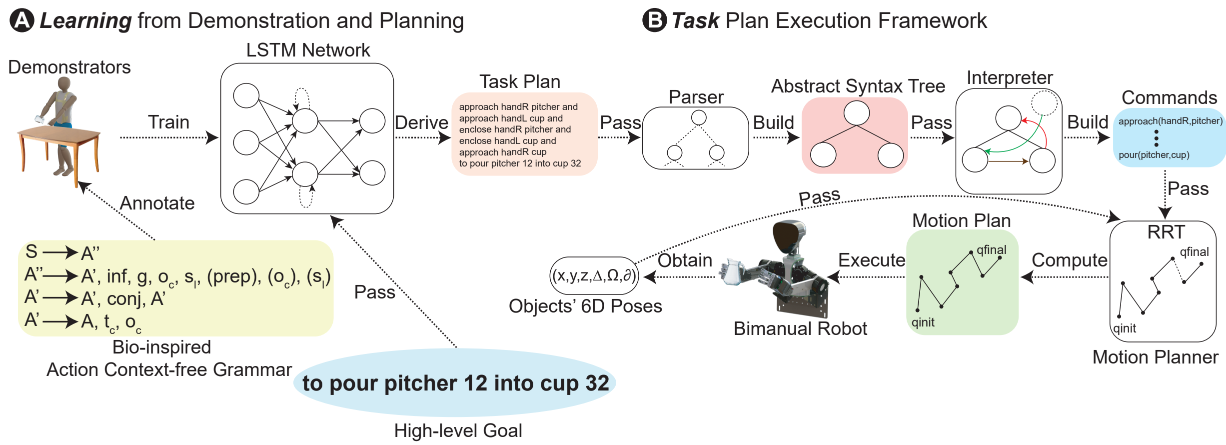

Fig. 1: Overview of the efficient learning-based task planning approach for bimanual manipulation.

Planning with RRT on a bimanual robotic platform. The combination of the BAG-Learn Planning with RRT via the TPEF constitutes this work's Task and Motion Planner.

III. Results

BAG-Learn Planning is evaluated against Fast Downward, a state-of-the-art classical task planner, using task planning time as the primary metric. The evaluation considers both *seen* high-level goals present in the training set and *unseen* high-level goals involving new object-location combinations. Fast Downward uses three heuristics for the experiments: 1) greedy best-first search using a causal graph heuristic (cg()); 2) greedy best-first search using a causal graph and fast-forward heuristics (cg()+ff()); and 3) A* search with a causal graph heuristic (A* cg()). The experiments aim to demonstrate improved efficiency over classical planning, mitigation of combinatorial explosion as problem complexity increases, and successful deployment of the complete TAMP framework on a bimanual manipulator.

A. Task Planning Efficiency

BAG-Learn Planning consistently outperforms Fast Downward. For the pouring activity, the planning times range from 9.88 ms (seen) to 90.74 ms (unseen), compared to 46–48 s for Fast Downward. Similar performance gains are observed for the opening and passing activities. For the opening, BAG-Learn Planning achieves 1.58 ms and 51.8 ms for seen and unseen goals, whereas Fast Downward times range from 202.92 ms to 238.53 ms. For passing, BAG-Learn Planning's times are 2.26 ms and 85.07 ms for seen and unseen goals. Meanwhile, Fast Downward requires significantly longer times than BAG-Learn Planning. Its times range from 3434.38 ms to 3601.62 ms.

B. Combinatorial Explosion Mitigation

As the number of objects increases from 2 to 10, Fast Downward's planning time increases from 173 ms to more than 380 ms, reflecting combinatorial growth in the symbolic search space. In contrast, BAG-Learn Planning maintains planning times below 7 ms, regardless of the number of objects. Similarly, when the number of symbolic locations increases from 2 to 13, Fast Downward's planning time increases to over 1 s, whereas BAG-Learn Planning remains under 2–3 ms.

C. Bimanual Robot Evaluation

The complete TAMP framework is deployed on an in-house bimanual robot prototype equipped with anthropomorphic hands. The system successfully executes pouring, opening, and passing tasks for both seen and unseen goals. In all cases, symbolic task plans are inferred, parsed, converted into motion commands, and executed without human intervention. These results prove that computational efficiency can be translated into real-world execution.

IV. Discussion & Conclusion

This work presents BAG-Learn Planning, a learning-based task planning approach for bimanual manipulation that combines a Long-Short-Term Memory (LSTM) network with a Bio-Inspired Action Context-Free Grammar (BAG) to efficiently infer symbolic task plans from high-level goals. Experimental evaluation demonstrates significant speed improvements over Fast Downward, robustness to increasing numbers of objects and locations, and successful deployment on a physical bimanual robot. The approach can be extended to other activities. Provided that their task plans are expressed using the BAG and suitable demonstration data are available. The compact LSTM architecture (131,584 parameters) results in low computational requirements. This setup enables BAG-Learn Planning to be executed on an embedded computer. Future work will extend BAG-Learn Planning with affordance reasoning, conditional branching, and adaptation to dynamic environments.

References

[1] Qiwei Du, Bowen Li, Yi Du, Shaoshu Su, Taimeng Fu, Zitong Zhan, Zhipeng Zhao, and Chen Wang. Fast Task Planning With Neuro-Symbolic Relaxation. *IEEE Robotics and Automation Letters*, 11(3):3684–3691, March 2026.

[2] Jan L. Guynes. Impact of system response time on state anxiety. *Commun. ACM*, 31(3):342–347, March 1988.

[3] David Carmona and Haoyong Yu. BiCap: A novel bi-modal dataset of daily living dual-arm manipulation actions. *The International Journal of Robotics Research*, 44(6):891–907, May 2025.

[4] Katerina Pastra and Yiannis Aloimonos. The minimalist grammar of action. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 367(1585):103–117, January 2012.