Beyond Task and Motion Planning: Hierarchical Robot Planning with General-Purpose Policies

Benned Hedegaard[1*], Ziyi Yang[1], Yichen Wei[1], Ahmed Jaafar[1], Stefanie Tellex[1], George Konidaris[1], and Naman Shah[1]

***Abstract*— Task and motion planning is a well-established approach for solving long-horizon robot planning problems. However, traditional methods assume that each task-level robot action, or skill, can be reduced to kinematic motion planning. In this work, we address the challenge of planning with both kinematic skills and closed-loop motor controllers that go beyond kinematic considerations. We propose a novel method that integrates these controllers into motion planning using *Composable Interaction Primitives (CIPs)*, enabling the use of diverse, non-composable pre-learned skills in hierarchical robot planning. Toward validating our *Task and Skill Planning (TASP)* approach, we describe ongoing robot experiments in real-world scenarios designed to demonstrate how CIPs can allow a mobile manipulator robot to effectively combine motion planning with general-purpose skills to accomplish complex tasks.**

(a) (b)

Fig. 1. A Boston Dynamics Spot® robot demonstrating skills that must be composed to solve long-horizon hybrid robot planning problems. Fig. 1(a) shows the robot in the process of executing a *kinematic skill* (`Pick`) to pick up a spray bottle. In Fig. 1(b), the robot uses a *non-kinematic skill* (`Erase`) to maintain force-controlled contact while erasing a whiteboard.

I. Introduction

Task and motion planning (TAMP) approaches [7, 3] have proven effective for solving long-horizon robot planning problems. But despite promising results on complex problems, TAMP methods typically make two key assumptions: (i) the low-level planning problem is *strictly* a kinematic motion planning task, and (ii) exact kinematic models of *every* object in the environment are available a priori.

However, most interesting real-world tasks that a general-purpose robot would encounter violate these assumptions. These tasks range from simple pick-and-place operations, which can be reduced to kinematic motion planning problems, to tasks that involve complex feedback-controlled motion such as wiping a wet table, which cannot be reduced to a valid motion planning problem due to force and sustained-contact constraints. Traditional TAMP approaches would be more widely applicable if they could exploit general-purpose robot controllers, such as those learned by reinforcement learning [9] or from demonstration [2]. The question then becomes: how can motion planning—which can generate collision-free trajectories for novel scenes without requiring learning, thereby enabling many forms of robot behavior—be seamlessly integrated alongside complex motor skills?

We propose to package such skills inside a motion-planning wrapper. If the wrapper has three properties—it enables the robot to motion plan to a configuration where skill execution can commence, skill execution itself is safe, and the robot can motion plan back to free space—then the wrapped skills can be used within existing TAMP approaches, and sequentially composed by interleaving motion-planned segments.

[1]Department of Computer Science, Brown University, Providence, RI.
* Corresponding author: `benned_hedegaard@brown.edu`
We gratefully acknowledge support from ONR Grant N00014-22-1-2592.

II. Our Approach

Traditional Task and Motion Planning (TAMP) [7] approaches assume that all robot actions can be computed using kinematic motion planning. However, many real-world tasks—such as opening a door or wiping a surface—require non-kinematic skills that involve sustained contact or force-based control. We formalize hybrid robot planning, which incorporates both kinematic (e.g., picking up an item) and non-kinematic (e.g., erasing a whiteboard) general-purpose skills. Each skill is described using initiation and termination conditions and a corresponding policy, often learned through reinforcement learning or demonstration.

A central challenge in hybrid robot planning is the lack of guaranteed composability between skills. Skills are typically trained in isolation and over limited parts of the state space. This makes it difficult to ensure that one skill's termination leads to another's initiation, especially when initiation and termination conditions are defined over continuous, high-dimensional state spaces. To enable reliable composition, the robot must be able to move into and out of these conditions using structured planning methods.

We propose *Task and Skill Planning (TASP)*, a hierarchical approach that wraps non-kinematic skills in Composable Interaction Primitives (CIPs) [1]. A CIP encapsulates a skill between two motion plans: a head motion plan to reach the skill's kinematic initiation set, and a tail motion plan to exit its kinematic termination set. This wrapping allows skills with complex dynamics to be integrated into motion planning pipelines. By projecting initiation and termination conditions onto purely spatial variables, CIPs make it possible to

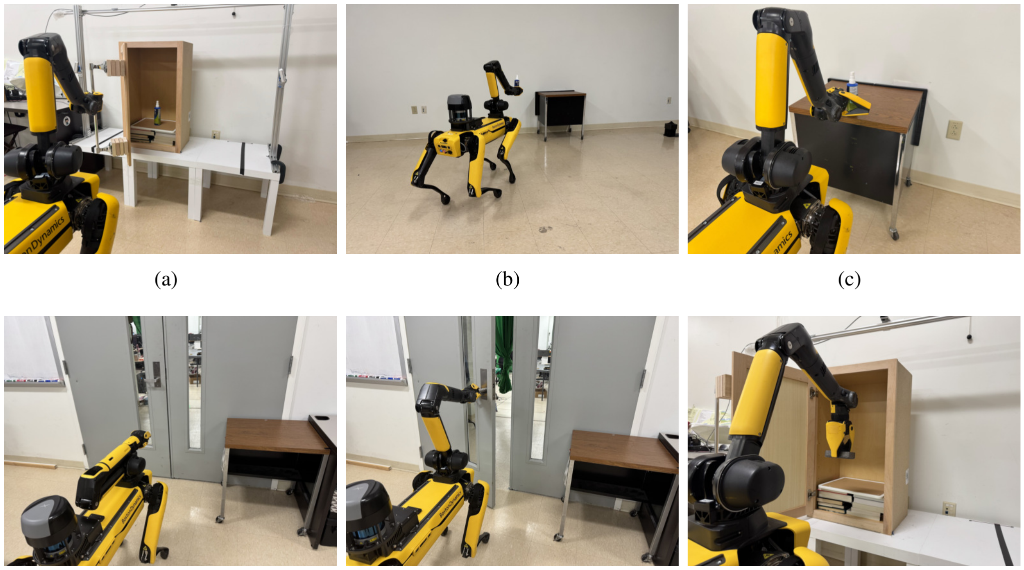

(a) (b) (c) (d) (e) (f)

Fig. 2. We illustrate two real-world mobile manipulation experiments requiring hybrid robot planning. *Experiment 1* - Fig. 2(a) shows the Spot robot opening a cabinet using a trajectory playback skill. Fig. 2(b) shows the robot navigating to a table while carrying a bottle obtained from the cabinet. In Fig. 2(c), the robot has completed its task by placing the bottle onto the table. *Experiment 2* - In Fig. 2(d), the robot has navigated to a closed door in preparation to open the door using an off-the-shelf skill, as shown in Fig. 2(e). Fig. 2(f) shows the robot picking up a whiteboard eraser to be used to erase the board (see Fig. 1(b)) in the adjoining room. Although these skills have been individually implemented and partially sequenced on the robot as depicted here, further integration efforts are needed to enable the consecutive execution of these skill sequences as planned by our TASP approach.

use conventional motion planners to handle entry and exit transitions for general-purpose skills.

To compose CIPs into full hybrid plans, TASP leverages a modified version of the ATAM algorithm [7] along with entity abstraction. Each CIP is represented as a symbolic action with discrete parameters corresponding to its internal motion plans and skill policy. A high-level symbolic planner identifies sequences of these abstract actions, while motion planning and skill execution fill in the low-level details. This hybrid strategy enables the robot to execute long-horizon plans that integrate diverse skills—even those involving contact or force—without requiring overlapping initiation and termination regions.

III. Experiments

To demonstrate the effectiveness of the proposed Task and Skill Planning (TASP) framework, we are conducting robot experiments on a Boston Dynamics Spot robot equipped with a manipulator and a diverse set of object-centric skills. These skills are implemented using motion planning (e.g., `Pick`, `Place`, and `GoTo`), trajectory playback (e.g., `OpenCabinet`), force control (e.g., `Erase`), and black-box policies (e.g., `OpenDoor`). In our experiments, the robot relies on a multimodal perception pipeline combining SLAM using RTAB-Map [5] and object pose estimation using visual fiducial markers from the `ar_track_alvar` [6] library. We run all planning and perception computation on an external GPU-equipped machine, while the robot operates autonomously using only sensor data collected online or during an initial teleoperated walkthrough of the environment.

In the first proposed experiment, the robot is tasked with retrieving a bottle from a closed cabinet and placing it on a nearby table (Fig. 2(a)-2(c)). This will require integrating multiple skills: navigating to the cabinet, opening the cabinet using open-loop trajectory playback, picking up the bottle, navigating to the table, and placing the bottle down. While this sequence corresponds to the task-level plan found by our planner, further integration is still needed to execute the entire sequence on the robot. This task is intended to demonstrate how TASP can be used to chain motion planning skills alongside general-purpose, non-kinematic policies.

Our second experiment tasks the robot with erasing a whiteboard located in a separate room (Fig. 2(d)-2(f)). Using our TASP method, the planner identifies a solution in which the robot navigates to the closed door between the rooms, opens it using a built-in skill, picks up a whiteboard eraser from a nearby cabinet, navigates to the whiteboard, and finally executes a force-controlled erasing action. These experiments illustrate TASP's ability to solve complex, long-horizon planning problems that require reasoning over both kinematic and non-kinematic actions. Importantly, all solution plans are computed and executed using sensor-derived state estimates without access to hand-coded inter-skill transitions, highlighting the robustness and generality of our hybrid robot planning approach.

IV. Conclusions and Future Work

In this paper, we formalize a hybrid robot planning problem and propose the first known approach for hierarchical planning to solve such problems. We describe ongoing real-world mobile manipulation experiments designed to validate the proposed approach. Currently, we focus on deterministic settings. We aim to extend our work to stochastic settings. Additionally, we assume that the initiation and termination conditions are known. We believe this is a strong assumption, and we aim to learn these conditions [4, 8].

References

[1] B. Abbatematteo, E. Rosen, S. Thompson, T. Akbulut, S. Rammohan, and G. Konidaris, "Composable Interaction Primitives: A Structured Policy Class for Efficiently Learning Sustained-Contact Manipulation Skills," in *Proc. ICRA*, 2024.

[2] B. Argall, S. Chernova, M. Veloso, and B. Browning, "A survey of robot learning from demonstration," *Robotics and Autonomous Systems*, vol. 57, pp. 469–483, 2009.

[3] C. R. Garrett, R. Chitnis, R. Holladay, B. Kim, T. Silver, L. P. Kaelbling, and T. Lozano-Pérez, "Integrated Task and Motion Planning," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 4, pp. 265–293, 2021.

[4] G. Konidaris, L. P. Kaelbling, and T. Lozano-Perez, "From Skills to Symbols: Learning Symbolic Representations for Abstract High-Level Planning," *Journal of Artificial Intelligence Research*, vol. 61, pp. 215–289, 2018.

[5] M. Labbé and F. Michaud, "RTAB-Map as an Open-Source Lidar and Visual SLAM Library for Large-Scale and Long-Term Online Operation," *Journal of Field Robotics*, vol. 36, no. 2, 2019.

[6] S. Niekum, "ar_track_alvar: AR Tag Tracking Library for ROS," https://wiki.ros.org/ar_track_alvar.

[7] N. Shah, D. K. Vasudevan, K. Kumar, P. Kamojjhala, and S. Srivastava, "Anytime Integrated Task and Motion Policies for Stochastic Environments," in *Proc. ICRA*, 2020.

[8] N. Shah, J. Nagpal, P. Verma, and S. Srivastava, "From Reals to Logic and Back: Inventing Symbolic Vocabularies, Actions, and Models for Planning from Raw Data," *arXiv preprint arXiv:2402.11871*, 2024.

[9] R. Sutton and A. Barto, *Reinforcement Learning: An Introduction*. Cambridge, MA: MIT Press, 1998.