Composable Causality in Semantic Robot Programming

Emily Sheetz, Xiaotong Chen, Zhen Zeng, Kaizhi Zheng, Qiuyu Shi, and Odest Chadwicke Jenkins

***Abstract*—Assembly tasks are challenging for robot manipulation because the robot must reason over the composed effects of actions and execute multi-objective behaviors. Robots typically use pre-defined priorities provided by users to determine how to compose controller behaviors, but we want the robot to autonomously select these compositions based on their composed effects within the task. We present Composable Causality in Semantic Robot Programming to allow robots to reason over the composed effects of controllers when executing multi-objective actions and autonomously compose controllers without pre-defined priorities. Our proposed *causal control basis* combines controller behaviors with causal information about how the behaviors can be used to execute high-level symbolic actions. The robot uses the *causal control basis* to predict the transition probability of achieving the composed effects of a multi-objective action. The composed causality estimates are used to select which action to execute within the context of a furniture assembly task. We evaluate the robot's transition probability estimates in different furniture assembly trials in simulation on the Baxter robot. The robot's ability to assemble furniture using different multi-objective connection actions demonstrates the usefulness of the composed causality estimates from our *causal control basis*.**

I. INTRODUCTION

Assembly tasks are an interesting domain for complex goal-directed manipulation in both academia and industry settings. Assembly tasks have been studied in academia in the context of tool construction [7] and furniture assembly [6], and assembly in industry settings is a common task for robots to take on. Robots have difficulty assembling objects because they have to compose the effects of multiple behaviors and maintain these composed effects as they move on to the next step. Semantic Robot Programming (SRP) [17] has emerged as an intuitive way to declaratively program robots to achieve a task. We need to extend SRP so robots can compose effects on objects during task execution and overcome the challenges of assembly tasks.

Expressing affordances [4] as *object-centric controllers* [1] allows robots to compose behaviors together to perform the complex multi-objective actions involved in assembly tasks. Object-centric controllers can be used within a *control basis* and can be *composed* through nullspace composition to yield multi-objective behaviors [12, 11, 13, 5, 14]. However, the compositions of behaviors are generally determined by a pre-defined priority provided by the user. To compose behaviors autonomously, robots need to reason over when controllers can be enacted, what controllers to compose, and what the effects of these behaviors will be. Researchers are starting to explore how robots can autonomously compose controllers [15], and we need to extend this work to long-horizon assembly tasks.

In this paper, we propose a *causal control basis* to build on SRP and allow robots to autonomously compose controllers to achieve assembly tasks. We take inspiration from work on causality and the analysis of cause and effect relationships [3, 8, 9, 10] and insights into using multiple hierarchical models for robot manipulation tasks [16]. Using the given *causal control basis*, the robot can estimate the transition probability that a controller composition will achieve the desired composed effects. During task execution, the robot will autonomously compose controllers based on their predicted composed causality and execute the afforded multi-objective actions to assemble a piece of furniture. We test our *causal control basis* in simulation on the Baxter robot in a variety of furniture assembly tasks, as seen in Figure 1. Our work on **Composable Causality in Semantic Robot Programming** demonstrates that the *causal control basis* allows the robot to autonomously compose controllers and achieve assembly tasks within the SRP paradigm.

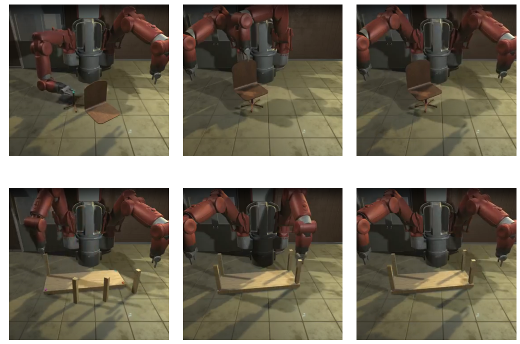

Fig. 1: Execution of swivel chair (top row) and table (bottom row) assembly tasks using multi-objective connection actions.

II. METHODS

We propose a *causal control basis* $\mathbf{\Phi}$ that the robot will use to autonomously compose controllers without predefined priorities by predicting the *transition probabilities* (or estimating the *causality*) of the controller compositions within furniture assembly tasks. The *causal control basis* is given to the robot and is comprised of:

- the implemented controllers in the control basis Φ;
- the set of composed causal graphs G_C, which describe what controllers are involved in a multi-objective action; and
- the set of temporal graphs G_T that represent the sequence of controllers that correspond to high-level symbolic actions.

Our *causal control basis* is denoted $\mathbf{\Phi} = (\Phi, G_C, G_T)$. For furniture assembly tasks, we define the control basis to include pose, position, rotation, and screw controllers such that $\Phi = \{\phi_{\text{6Dpose}}, \phi_{\text{pos}}, \phi_{\text{rot}}, \phi_{\text{screw}}\}$; the causal graphs G_C indicate which of these behaviors are involved in performing multi-objective insert and screw actions; and the temporal graphs G_T indicate how to decompose high-level pick-up, insert, and screw actions into sequences of executable controller behaviors. However, our formulation will work with an arbitrary control basis appropriate for an arbitrary multi-objective assembly task. The causal control basis tells the robot what controllers are involved in each multi-objective action, but not how to compose these controllers to achieve the desired composed effects.

To predict the transition probabilities of the controller compositions, the *causal control basis* performs *offline walkouts*. Before task execution, the robot uniformly samples initial states $s \in S$ and controller goals $s' \in S$. The robot simulates execution of the composed controllers until they converge or until large time threshold T. The predicted transition probability that the action $a = \phi_k \lhd \phi_j \lhd \phi_i$ (the "subject-to" relation $\lhd$ indicates the priority of behaviors in the composition) achieves the composed effects s' based on the offline walkout is:

$$\hat{P}(s' \mid s, a) = \begin{cases} 1 & \text{objectives met} \\ 0 & \text{bad progress} \\ \frac{\phi_a(s) - \phi_a(s_T)}{\phi_a(s)} & \text{otherwise} \end{cases} \quad (1)$$

where s_T is the state at time threshold T, ϕ_a indicates the composed cost or objective function values at the given state, and bad progress means the controllers reached a local minimum. For a large number of random samples, the average predicted transition probability indicates how the controller composition performs across action instances. During task execution, the robot will query the *causal control basis* for the controller composition with the greatest predicted transition probability and will execute that composition.

III. EXPERIMENTS AND RESULTS

To evaluate our approach to Composable Causality in Semantic Robot Programming and our proposed *causal control basis*, we assume that the robot has parsed the goal conditions of the task as in SRP [17] and that we have affordance-based perception[1] to perceive objects and affordances in the scene. We use an off-the-shelf task planner[2] to construct the task plan. The *causal control basis* converts each symbolic action into a sequence of (possibly composed) controller commands and instantiates the action based on the current poses of the objects and their connection sites. The robot executes this sequence of controllers to assemble furniture. We evaluate the proposed *causal control basis* in simulation using the Baxter robot in the IKEA Furniture Assembly Environment[3] [6]. We assume known object poses during manipulation and grasp poses for every object part are provided.

For the insert and screw connection actions, the robot simulated 500 executions of each possible composition. We used time threshold $T = 300$ controller updates as the cutoff for the offline walkouts. The transition probability predictions for the insert and screw actions are shown in Table I and Table II, respectively.

Composition a	Predicted Transition Probability $\hat{P}(s' \mid s, a)$
$\phi_{pos} \lhd \phi_{rot}$	**0.723**
$\phi_{rot} \lhd \phi_{pos}$	0.711

TABLE I: Transition probability predictions for insert action, based on 500 offline walkouts for each composition.

Composition a	Predicted Transition Probability $\hat{P}(s' \mid s, a)$
$\phi_{rot} \lhd \phi_{screw} \lhd \phi_{pos}$	**0.937**
$\phi_{pos} \lhd \phi_{screw} \lhd \phi_{rot}$	0.936
$\phi_{screw} \lhd \phi_{pos} \lhd \phi_{rot}$	0.929
$\phi_{pos} \lhd \phi_{rot} \lhd \phi_{screw}$	0.925
$\phi_{screw} \lhd \phi_{rot} \lhd \phi_{pos}$	0.923
$\phi_{rot} \lhd \phi_{pos} \lhd \phi_{screw}$	0.904

TABLE II: Transition probability predictions for screw action, based on 500 offline walkouts for each composition.

The controller composition with the maximum transition probability prediction is used to execute connect actions in a variety of furniture assembly tasks to test the accuracy of the composed causality estimates. We tested the insert action within a swivel chair assembly task and the screw action within a table assembly task. Across all trials, we compute the success rate of the multi-objective connection actions and the success rate of the assembly task. For 10 swivel chair assembly trials, the insert action success rate was 0.714 and the task success rate was 1. For 6 table assembly trials, the screw action success rate was 0.923 and the task success rate was 1. Select images from one of the random swivel chair[4] and table[5] trials are in Figure 1.

The similarity of the action success rates and the predicted transition probabilities (Table I and Table II) indicate that the transition probability predictions accurately capture the performance of the compositions during task execution. The task success rates indicate that when the compositions for the connect actions failed (due to joint limits or collisions between objects), the robot could retry the action and ultimately achieve a successful connection.

The robot's ability to successfully assemble different furniture pieces demonstrates the accuracy of our proposed *causal control basis* in predicting the composed effects of controller behaviors. Our work in Composable Causality in Semantic Robot Programming demonstrates that reasoning over a *causal control basis* allows the robot to autonomously compose controller behaviors without pre-defined priorities to achieve assembly tasks and can be applied to goal-directed manipulation tasks in academic and industry contexts.

ACKNOWLEDGMENTS

This work was supported by a NASA Space Technology Graduate Research Opportunity.

[1] For example, Affordance Coordinate Frames (ACFs) [2].
[2] Pyperplan STRIPS planning library: https://github.com/aibasel/pyperplan
[3] https://clvrai.github.io/furniture/
[4] Swivel chair trials: https://www.youtube.com/watch?v=4Gpw7uuJju4
[5] Table trials: https://www.youtube.com/watch?v=63RFBU4dulo

References

[1] Dana H. Ballard. Task Frames in Robot Manipulation. In *AAAI*, volume 19, pages 16–22, 1984.

[2] Xiaotong Chen, Kaizhi Zheng, Zhen Zeng, Shreshtha Basu, James Cooney, Jana Pavlasek, and Odest Chadwicke Jenkins. Manipulation-Oriented Object Perception in Clutter through Affordance Coordinate Frames. *arXiv preprint arXiv:2010.08202*, 2020.

[3] Ishita Dasgupta, Jane Wang, Silvia Chiappa, Jovana Mitrovic, Pedro Ortega, David Raposo, Edward Hughes, Peter Battaglia, Matthew Botvinick, and Zeb Kurth-Nelson. Causal Reasoning from Meta-Reinforcement Learning. *arXiv preprint arXiv:1901.08162*, 2019.

[4] James J. Gibson. The Theory of Affordances, 1977.

[5] Stephen Hart and Roderic Grupen. Natural Task Decomposition with Intrinsic Potential Fields. In *IEEE International Conference on Intelligent Robots and Systems*, pages 2507–2512, 2007.

[6] Youngwoon Lee, Edward S. Hu, Zhengyu Yang, Alex Yin, and Joseph J. Lim. IKEA Furniture Assembly Environment for Long-Horizon Complex Manipulation Tasks. *arXiv preprint arXiv:1911.07246*, pages 1–11, 2019.

[7] Lakshmi Nair, Jonathan Balloch, and Sonia Chernova. Tool Macgyvering: Tool Construction Using Geometric Reasoning. *International Conference on Robotics and Automation (ICRA)*, pages 5837–5843, 2019.

[8] Judea Pearl. Causal Diagrams for Empirical Research. *Biometrika*, 82(4):669–688, 1995.

[9] Judea Pearl. *Causality*. Cambridge University Press, 2009.

[10] Judea Pearl. The Do-Calculus Revisited. *arXiv preprint arXiv:1210.4852*, 2012.

[11] Robert Platt Jr, Andrew H Fagg, and Roderic A Grupen. Nullspace Composition of Control Laws for Grasping. In *IEEE International Conference on Intelligent Robots and Systems (IROS)*, volume 2, pages 1717–1723. IEEE, 2002.

[12] Robert Platt Jr, Andrew H Fagg, and Roderic A Grupen. Manipulation gaits: Sequences of Grasp Control Tasks. In *IEEE International Conference on Robotics and Automation*, volume 1, pages 801–806. IEEE, 2004.

[13] Robert Platt Jr, Andrew H Fagg, and Roderic A Grupen. Null-Space Grasp Control: Theory and Experiments. *IEEE Transactions on Robotics*, 26(2):282–295, 2010.

[14] Khashayar Rohanimanesh, Robert Platt Jr, Sridhar Mahadevan, and Roderic Grupen. Coarticulation in Markov Decision Processes. In *Advances in Neural Information Processing Systems (NIPS)*, pages 1137–1144, 2004.

[15] Mohit Sharma, Jacky Liang, Jialiang Zhao, Alex LaGrassa, and Oliver Kroemer. Learning to Compose Hierarchical Object-Centric Controllers for Robotic Manipulation. *arXiv preprint arXiv:2011.04627*, pages 1–24, 2020.

[16] Caiming Xiong, Nishant Shukla, Wenlong Xiong, and Song-Chun Zhu. Robot Learning with a Spatial, Temporal, and Causal And-Or Graph. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2144–2151. IEEE, 2016.

[17] Zhen Zeng, Zheming Zhou, Zhiqiang Sui, and Odest Chadwicke Jenkins. Semantic Robot Programming for Goal-Directed Manipulation in Cluttered Scenes. *IEEE International Conference on Robotics and Automation (ICRA)*, pages 7462–7469, 2018.