# 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*.**

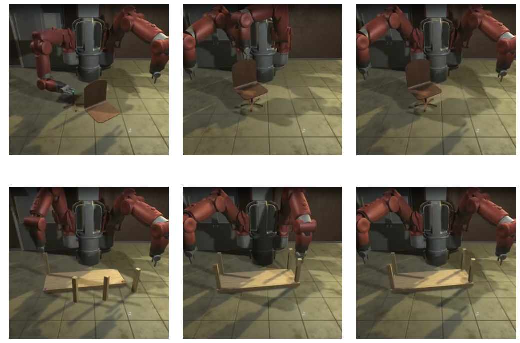

Fig. 1: Execution of swivel chair (top row) and table (bottom row) assembly tasks using multi-objective connection actions.

## I. INTRODUCTION

Assembly tasks present unique challenges in reasoning over objects and executing complex behaviors in long-horizon tasks. Assembly tasks have been studied in the context of tool construction [8] and furniture assembly [6]. 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. We need a flexible way to program our robots to perform assembly tasks, and Semantic Robot Programming (SRP) [19] has emerged as an intuitive way to declaratively program robots to achieve a task. Within the SRP paradigm, robots can infer goal conditions from a demonstrated goal scene and reason over available objects and actions to reach the goal. When SRP was first presented, the work used off-the-shelf task and motion planners to execute tasks and instead focused on the perceptual challenges involved in perceiving the demonstrated goal conditions. Now that these perceptual challenges have been addressed, we aim to extend SRP in terms of the types of actions that the robot can perform, so that we can declaratively program robots to perform more challenging assembly tasks. Specifically, robots need to reason about the objects themselves, compose effects on these objects during task execution, and overcome the challenges of assembly tasks.

Assembly tasks are challenging because robots need to determine when to enact controller behaviors, predict the composed effects of actions, and maintain the effects of these behaviors throughout task execution. Expressing affordances [4] as *object-centric controllers*—which send joint commands such that the robot achieves low-level motion primitives [1]—allows robots to compose behaviors together to perform more complex *multi-objective actions*—which require the robot to execute multiple controllers concurrently [13, 12, 14, 5, 15]. For example, robust grasping could be formulated as a multi-objective action that involves *positioning* an end-effector while *aligning* the approach axis of that end-effector with the target object. Composing multiple behaviors induces a priority between these behaviors, which means one behavior will likely be achieved first, and in the worst case may impede the other controller(s) from converging. The priorities between controllers—the particular composition of these controllers—can greatly impact the effect of the multi-objective action.

Compositions of controller behaviors are generally determined by a pre-defined priority provided by the user. For example, a user may determine experimentally that prioritizing positioning over alignment results in the most robust grasp poses; therefore, the user will hard-code the robot to always perform multi-objective grasps by composing these behaviors such that positioning is the highest priority. To extend SRP to tasks that require multi-objective behaviors, however, we want to declaratively program robots to autonomously compose controllers without these pre-defined priorities. Researchers are starting to explore how robots can autonomously compose controllers [17], and we need to extend these ideas into long-horizon assembly tasks. To compose behaviors autonomously, robots need to reason over pre- and post-conditions of these behaviors by grounding them in the perceived scene, rather than in a symbolic manner that is disconnected from the realities of physical execution. Expressing user insights on controller compositions to robots remains an open question because of the complexity of the causal relations the robot would need to understand, specifically when controllers can be enacted, what controllers to compose, and what the effects of these behaviors will be.

We propose that notions of *causality* provide the insight needed to address the challenges of predicting the effects of controller compositions. We take inspiration from work on causality and the analysis of cause and effect relationships [3, 9, 10, 11] and insights into using multiple hierarchical graphs to express causal concepts for robot manipulation tasks [18]. In the case of goal-directed manipulation tasks, causality can allow robots to reason about the effects of composing controllers on the perceived objects and affordances. If the robot can predict the transition probability (or composed causality) of composed controller behaviors, then the robot can determine the compositions of controllers to execute within challenging long-horizon assembly tasks without relying on pre-defined priorities.

In this paper, we propose a *causal control basis* to build on SRP and allow robots to autonomously compose controllers to achieve assembly tasks. 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 controller behaviors based on their predicted composed effects 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 achieve challenging goal-directed manipulation tasks within the SRP paradigm.

## II. METHODS

### A. Problem Formulation

To perform assembly tasks that require multi-objective behaviors, the robot needs to predict the *transition function* of the controllers—the probability that a given composition of controllers will achieve their composed effects. We assume we have a control basis $\Phi$ of controllers that can be composed using nullspace projection to achieve multiple objectives. Given a task goal, the robot constructs a high-level task plan using an off-the-shelf task planner and decomposes each symbolic action into a sequence of executable motions. For actions that require multi-objective behaviors, we want the robot to autonomously compose the given controllers and execute the planned symbolic action by reasoning over the *causality* of the controllers based on the transition probability estimates.

We formulate this probabilistic planning problem as a Markov Decision Process (MDP) [7, 16] $(S, A, P, C)$. The state space $S$ is determined by the robot configuration space and the poses of the objects in the scene. The action space $A$ is the set of all possible controllers and compositions in the given control basis $\Phi$. The controllers that can be running at any given time are elements of the power set of the control basis $\mathcal{P}(\Phi)$. Suppose we have controllers $\phi_i$ and $\phi_j$ that achieve objectives $i$ and $j$, respectively. One possible composition of these controllers is $\phi_j \lhd \phi_i$, where the "subject-to" relation $\lhd$ indicates that $\phi_i$ is the controller with the highest priority. Let $M_t \in \mathcal{P}(\Phi)$ be the set of controllers running at time $t$. Since composing controllers induces an ordering (priority) between them, all possible compositions of the running controllers $M_t$ are elements of the symmetric group $S_{M_t}$, which is the set of permutations over the elements (controllers) in $M_t$. Therefore, the action space for control basis $\Phi$ is $A = \{S_{M_t} \mid M_t \in \mathcal{P}(\Phi)\}$. The transition probability $P(s' \mid s, a)$ indicates the probability of achieving the composed effects $s'$ of a (composed) controller $a \in A$ when enacted in the current state $s$. The cost function $C_a(s)$ is the cost of enacting controller $a$ in state $s$. We want the robot to execute the controller $a$ that will achieve its composed effects by minimizing its objective function $\phi$, meaning $C_a(s) = \phi_a(s)$.

The action space is determined by the causal control basis, but the robot is not given any information about the transition probabilities associated with the (composed) controllers. The robot needs to estimate the transition probability for each possible controller composition.

### B. Causal Control Basis

We propose a *causal control basis* $\mathbf{\Phi}$ that the robot will use to predict the transition probabilities of actions and determine which composition of controllers to execute to achieve assembly tasks. The *causal control basis* is given to the robot and is comprised of the following components:

- The implemented controllers in the control basis $\Phi$.
- The set of composed causal graphs $G_C$, which describe what controllers are involved in a multi-objective action. Causal graphs are comprised of *pre-conditions*, which are literals that must be true before an action is taken; *controller behaviors* that must be executed concurrently to achieve a multi-objective action; the *effects* of the individual controllers; and the desired *composed effects*, which are literals that should be achieved by the composition of the controllers.
- The set of temporal graphs $G_T$, which represent the sequence of controllers that correspond to high-level symbolic actions. The root of the temporal graph is the high-level symbolic action, and each low-level controller behavior required to execute this action is a child of the root, arranged from left-to-right in sequence. The structure of the temporal graph is inspired by previous work that uses hierarchical graphs within robot manipulation tasks [18].

Our *causal control basis* is denoted $\mathbf{\Phi} = (\Phi, G_C, G_T)$. For the tasks considered in the experiments, we define the control basis—discussed in more detail in Section III-A1—to include pose, position, rotation, and screw controllers. However, our formulation will work with an arbitrary control basis.

The components of the given *causal control basis* tell the robot the sequence of controller behaviors associated with each high-level symbolic action, pre-conditions under which controller compositions can be enacted, and which controllers are involved in each multi-objective action. However, the robot does not know how to compose controllers to achieve the desired composed effects. For each possible composition $a$,

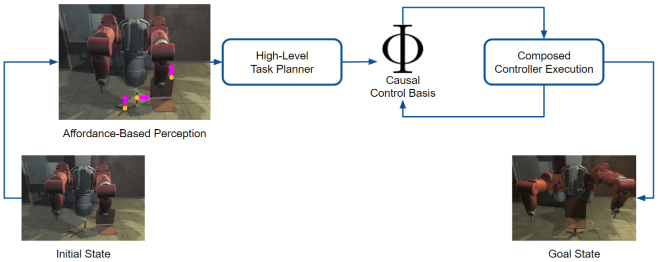


Fig. 2: Our pipeline for Composable Causality in SRP.

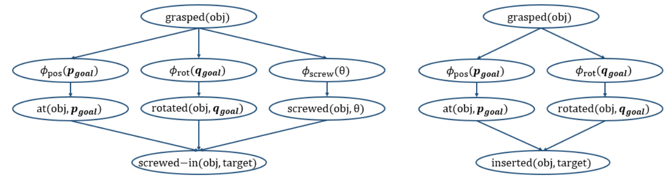


Fig. 3: The causal graphs for the multi-objective connect actions in the control basis for furniture assembly tasks. When assembling furniture, the robot needs to connect the acted on object part `obj` to a target object part `target`.

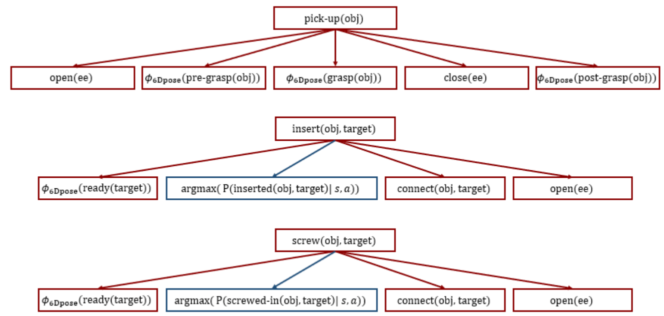


Fig. 4: The temporal graphs relating the high-level symbolic actions reasoned over by the task planner to the low-level controller behaviors in the control basis.

the robot must use the *causal control basis* to estimate the transition probability $P(s' \mid s, a)$, which will allow the robot to predict if the composed controllers will result in successful assembly of the furniture piece.

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 $s' \in S$. Suppose the robot is predicting the transition probability for arbitrary controller composition $\phi_k \lhd \phi_j \lhd \phi_i$ (where the "subject-to" relation $\lhd$ indicates the priority of behaviors in the composition). The robot simulates execution of the composed controllers until they converge or until large time threshold $T$. The predicted transition probability that the action $a = \phi_k \lhd \phi_j \lhd \phi_i$ achieves the composed effects $s'$ based on the offline walkout is:

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

where $s_T$ is the state at time threshold $T$, $\phi_a$ 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

Figure 2 describes the pipeline for assembling furniture using Composable Causality in Semantic Robot Programming and the use of our proposed *causal control basis*. We assume that the robot has parsed the goal conditions from a demonstrated goal scene of the task as in SRP [19] and that we have affordance-based perception[1] to perceive the objects and affordances in the scene. These perceived objects and affordances seed the initial state of an off-the-shelf high-level task planner[2], which constructs the task plan. The *causal control basis* converts each action in the high-level task plan 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 achieve the task goal of assembling furniture.

We evaluate the proposed *causal control basis* in various furniture assembly tasks 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. The connection of two parts is implemented as welding in the Mujoco simulation, which checks the position and axis alignment of connecting points.

### A. *Causal Control Basis for Furniture Assembly*

*1) Control Basis Implementation:* In this work, we define the control basis $\Phi$ for furniture assembly by 6D pose $\phi_{\text{6Dpose}}$, 3D position $\phi_{\text{pos}}$, rotation $\phi_{\text{rot}}$, and screw $\phi_{\text{screw}}$ controllers. All of these controllers are object-centric potential field controllers based on attractive potential fields that attract the robot and objects to the controller goal. Our furniture assembly control basis $\Phi$ is the set of these controllers:

$$\Phi = \{\phi_{\text{6Dpose}}, \phi_{\text{pos}}, \phi_{\text{rot}}, \phi_{\text{screw}}\} \tag{2}$$

*2) Causal Graphs:* The set of causal graphs $G_C$ indicate the composed effects of the controllers within the multi-objective `insert` and `screw` actions. As shown in Figure 3, the causal graphs indicate the controllers that are involved in these connect actions, the pre-conditions of enacting these

[1] For example, Affordance Coordinate Frames (ACFs) [2].

[2] Pyperplan STRIPS planning library: https://github.com/aibasel/pyperplan

[3] https://clvrai.github.io/furniture/

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

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

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

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

compositions, and the intended effects of these compositions. The robot will use the transition probability predictions from the *causal control basis* to determine how to compose these controllers together.

*3) Temporal Graphs:* The set of temporal graphs $G_T$ indicate the sequence of controllers that correspond to the high-level `pick-up`, `insert`, and `screw` actions, as seen in Figure 4. For multi-objective connection actions `insert` or `screw`, the robot will have to determine what composition of the controllers (indicated in the corresponding causal graph in Figure 3) to execute within the sequence (indicated by the blue boxes in Figure 4) by selecting the composed controller with the maximum predicted transition probability.

### *B. Composed Causality Predictions*

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 predicted transition probability $\hat{P}$ was computed for each walkout as in Equation 1 and averaged across walkouts for the same composition to determine the estimated composed causality of the composition $a$.

The transition probability predictions for the `insert` and `screw` actions are shown in Table I and Table II, respectively. The composition with the highest probability for the `insert` action indicates that positioning the object should be performed subject to aligning the object with the target, $\phi_{\text{pos}} \lhd \phi_{\text{rot}}$. The composition with the highest probability for the `screw` action indicates that aligning the object should be performed subject to screwing and positioning the object, $\phi_{\text{rot}} \lhd \phi_{\text{screw}} \lhd \phi_{\text{pos}}$.

### *C. Furniture Assembly Task Results*

The 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 estimated composed causality of the controllers. We tested the `insert` action within 10 random trials of swivel chair assembly and tested the `screw` action within 6 random trials of table assembly. Across all trials, we compute the success rate of the multi-objective `insert` and `screw` actions as well as the success rate of the entire 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 rate for both the swivel chair and table tasks indicate that when the compositions for the connect actions did not result in successful connections (due to joint limits or collisions between objects), the robot could recover by retrying the action and ultimately achieve a successful connection.

## IV. Discussion and Conclusion

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. The *causal control basis* effectively extends the principles of Semantic Robot Programming—that we need an intuitive way to declaratively program robots to perform tasks—into more challenging tasks that involve multi-objective actions. The *causal control basis* describes the intended results of the multi-objective actions, but leaves the robot to determine how to compose the appropriate controller behaviors and perform the action successfully. By representing important information in the form of causal and temporal graphs that users often already use in some form (such as symbolic descriptions of action pre- and post-conditions or hard-coded behavior sequences to enact symbolic actions) to describe tasks to robots, our *causal control basis* allows users to intuitively share knowledge about actions such that robots can autonomously determine how to enact those behaviors in challenging tasks.

In this work, we proposed a *causal control basis* for achieving Composable Causality in Semantic Robot Programming. Our *causal control basis* allows the robot to predict the transition probabilities of controller compositions, thereby estimating the composed causality of multi-objective actions. Our work in Composable Causality in Semantic Robot Programming demonstrates that reasoning over a *causal control basis* provides the robot with the declarative knowledge necessary to autonomously compose controller behaviors without predefined priorities to achieve furniture assembly tasks.


## Acknowledgments

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


[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] Michael L. Littman, Thomas L. Dean, and Leslie Pack Kaelbling. On the Complexity of Solving Markov Decision Problems. *arXiv preprint arXiv:1302.4971*, 2013.

[8] 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.

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

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

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

[12] 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.

[13] 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.

[14] 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.

[15] 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.

[16] Stuart J. Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach, Third Edition*. Pearson Education, 2010.

[17] 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.

[18] 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.

[19] 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.