

CoBT: Collaborative Programming of Behaviour Trees from One Demonstration for Robot Manipulation

Aayush Jain^{1,2}, Philip Long³, Valeria Villani⁴, John D. Kelleher⁵, Maria Chiara Leva¹

Abstract—Mass customization and shorter manufacturing cycles are becoming more important among small and medium-sized companies. However, classical industrial robots struggle to cope with product variation and dynamic environments. In this paper, we present CoBT, a collaborative programming by demonstration framework for generating reactive and modular behavior trees. CoBT relies on a single demonstration and a combination of data-driven machine learning methods with logic-based declarative learning to learn a task, thus eliminating the need for programming expertise or long development times. The proposed framework is experimentally validated on 7 manipulation tasks and we show that CoBT achieves $\approx 93\%$ success rate overall with an average of 7.5s programming time. We conduct a pilot study with non-expert users to provide feedback regarding the usability of CoBT. More videos and generated behavior trees are available at: <https://github.com/jainaayush2006/CoBT.git>.

I. INTRODUCTION

Intelligent robotic systems can improve factory ergonomics and productivity by collaborating with humans. Furthermore, as industry shifts toward mass customization, reprogramming robots using traditional methods will result in higher downtime and increased cost. Thus, there is a need for agile methods that can generate reactive and modular robot programs, for instance, programming by demonstration (PbD) which collaboratively generates task models using human assistance.

PbD (Learning from Demonstration or Imitation Learning) is a promising approach to simplify robot programming to advance human-centric automation in Industry 5.0 [1]. PbD enables human collaborators to transfer skills to a robot through task demonstration, though developing such systems to handle the high variability and uncertainty in real-world environments remains a challenge [2]. Most existing PbD methods generalize *motion-level* (low-level learning) plans [3]–[5] which are often insufficient to accomplish complex tasks where *task-level* (high-level learning) knowledge [6]–[9] is also required. However, learning on both levels is computationally expensive and time-consuming, as it requires multiple demonstrations or self-exploration to generalise.

This work is supported by the CISC project which is funded by the European Union’s Horizon 2020 MSCA-ITN grant agreement no. 955901; and ADAPT Research Centre which is funded by SFI Research Centres Programme and is co-funded under the European Regional Development Fund (ERDF) through Grant 13/RC/2106_P2.

¹Technological University Dublin; ²Irish Manufacturing Research; ³Atlantic Technological University; ⁴ARS Control Lab, University of Modena and Reggio Emilia; ⁵ADAPT Research Centre, School of Computer Science and Statistics, Trinity College Dublin. Email: aayush.jain@imr.ie

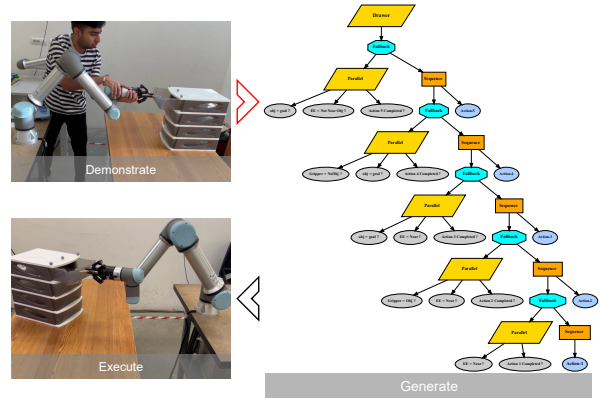


Fig. 1. CoBT: a collaborative programming framework that generates deployable behavior trees from one demonstration. This example shows the work-flow of programming an opening a drawer task.

To this end, a novel Collaborative programming of Behavior Trees (CoBT) framework is proposed. A single demonstration is segmented into symbolic states (Algorithm 1) and primitive actions (Algorithm 2) to automatically generate a BT (Algorithm 3). The generated BT provides a high-level control structure that enables the robot to react to environmental changes, while primitive actions in the form of Dynamic Movement Primitives (DMPs) enable the robot to generalize its movements to different environments. The learned BTs are adapted to user-defined goal to accomplish complex task multi-level tasks. Our key contributions are: (1) a collaborative programming framework that generates deployable BTs from one demonstration, (2) an automatic generation of state-action pairs to form atomic BTs which can perform complex motions and (3) a Logic-based Declarative Learning algorithm (Algorithm 3) capable of resolving task and motion constraints and generating BTs by composing action primitives. The paper is organized as follows: in Section II, the related work is presented; the CoBT framework and its components are introduced in Section III. The results of experiments across 7 tasks and pilot study conducted are discussed in Section IV. Lastly, in Section V, we conclude the discussion with limitations and future work.

II. RELATED WORK

In this section, frameworks that either program from one demonstration or generate BTs from multiple demonstration are outlined, as shown in Table I. We compare the frameworks based on: number of demonstrations, segmentation method (manual or automatic), motion learning (movement primitives or point-to-point), reactivity, and modularity.

To improve task generalization, a demonstration must be

TABLE I
COMPARISON WITH THE RELATED FRAMEWORKS

Ref.	Demos.	Seg.	Motion	Reactive	Modular
[4]	One	-	DMP	No	No
[13]	One	Manual	P2P	No	No
[11]	One	Auto*	P2P	No	Yes
[14]	One	Manual	P2P	No	Yes
[5]	Multi.	Manual	P2P/ProMP	Yes	Yes
[6]	Multi.	Manual	P2P	Yes	No
[15]	Multi.	Manual	P2P	Yes	No
[12]	Multi.	Auto*	P2P	Yes	No
CoBT	One	Auto	DMP	Yes	Yes

* Require pre-training or domain knowledge

segmented into a sequence of sub-tasks or primitive actions [10]. Segmentation can be automatic [11] and [12] or [4]–[6], [13]–[15] done manually. While [11] and [12] utilize pre-defined rules and pre-trained classifiers for segmentation, CoBT can automatically segment even unseen actions. We achieve this by employing a parameter-less velocity-based change-point detection [16] algorithm and grounding state variables. Compared to other methods, this simplifies user task demonstrations and reduces overall complexity.

Among the frameworks described in Table I, [6], [11]–[15], the main approach for motion generation is point-to-point (P2P), focusing on learning keyframes and way-points. However, this approach fails to preserve the shape of trajectories, limiting its functionality to executing variations of pick-and-place tasks. By contrast, Lentini et al. [4] proposed learning DMPs [17] from one simple demonstration without prior constraints. However, their method lacks reactivity to environmental changes due to an absence of state-action mapping. To achieve reactivity at both task and motion levels over a variety of complex tasks, we learn DMP [17] embedded atomic BT for each segment using the symbolic state variables. Moreover, we learn a set of DMPs for opposite trajectory to the one demonstrated, in order to generalise opposite start and end poses.

In PbD, data-based learning techniques require a large amount of data to generalize [18]. [4] proposed a data-efficient methodology that combines data-based learning and symbolic AI. Therefore, we explore logic-based declarative learning [19] to generate BTs. Declarative learning represents robot manipulation knowledge in the form of logical rules which can be used to generate explainable and transparent policies from the data through logical reasoning processes. The existing techniques to build BTs have limitations since they either generate BTs at runtime [7], [15] or are trained in a simulator [6], [12], making their outcome non-deterministic and limiting their scalability in actual industrial settings.

In general compared to Finite State Machines (FSM) [20], Hierarchical Task Networks (HTN) [21] and Hidden Markov Models (HMM) [9], Behaviour Trees (BT) [22] have proven to be a more reactive and modular approach to robot control. They allow the creation of complex behaviors by combining atomic BTs [23] i.e., simple sub-behaviors, in a hierarchical structure. Further, BTs are a human-readable and explainable control architecture. To the best of our knowledge, BTs have been used as high-level task plans for PbD only in [6], [7], [12], [15]. [15] learn a decision tree (DT) from demonstration, which is converted into a BT. [6] and [12]

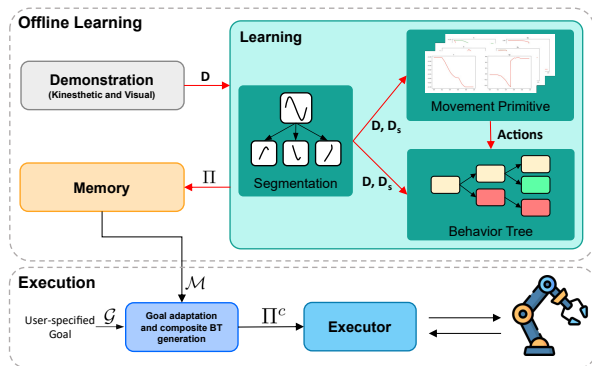


Fig. 2. A high-level workflow of offline learning and execution in CoBT. The red lines, depict the task (BT) and action (DMP) policies that are learned offline through PbD. During execution, depicted with black lines, the user-defined goal is abstracted to adapt the BT parameters. The generated BT decides “what” action to execute, and DMPs control “how” the action should be executed based on the current environmental conditions.

learn task constraints from demonstration and combine them with backchaining [23] to generate BTs. [7] proposed a method to learn BTs from natural language instructions by creating a mapping between logical forms and action BTs. However, these techniques require multiple demonstrations and generate rigid BTs that lack modularity. Conversely, CoBT requires one demonstration and the learned BT can be adapted to achieve a new goal or can be chained with other BTs to perform long-horizon tasks.

Valassakis et al. [24] presented an approach to one-shot PbD, DOME, where visual servoing is used to reach the user-specified pose and then the demonstrated velocities are replayed to imitate. Despite DOME’s promising results, the generated policies are highly dependent on user-specified pose and are limited to short-horizon tasks like only pick or insert. CoBT is the only framework capable of generating deployable, reactive, and modular robot programs in the form of BTs through one demonstration which are capable of executing complex motions and long-horizon tasks.

III. METHODOLOGY

CoBT, shown in Fig. 2, comprises- *demonstration, learning, memory, goal adaptation and composite BT generation, and executor* and is organized into: *offline learning*, and *execution*. During *offline learning*, the user demonstrates the task, using the *demonstration module*, which records them as multivariate time series data. This data subsequently serves as an input for the *learning module*, which performs automatic segmentation (Algorithm 1), primitive learning (Algorithm 2) and generates BTs (Algorithm 3).

Algorithm 1 automatically decomposes the demonstration into different phases. State variables, shown in Table II, are grounded depending on real-world sensory data at each segment. As a result, the demonstration is segmented into a set of state-action pairs. Each action is then encoded using DMP (Algorithm 2) and the action sequence is formalized in a BT (Algorithm 3). Actions are embedded in a BT’s leaf node and the state variables are used as task and action constraints. *Task constraints* refer to the relative sequencing of primitive actions and the fulfillment of goal conditions,

TABLE II
STATE VARIABLES AND RESPECTIVE GROUNDING

State variables	Grounding conditions
g_i	A gripper has opened (Open) or closed (Close) its fingers
o_i	Target object nearer to another object (Near) compared to $i - 1^{th}$ segment or object on goal position (On_goal)
e_i	End-effector Near or !Near the object

Algorithm 1: Segmentation

Input: Demonstration dataset \mathbf{D}
Output: Segmented dataset \mathbf{D}_s

- 1 **while** $t \leq N$ **do**
- 2 $\mathbf{v}_N \leftarrow v_t$; // $v_t = (v_{x_t}^2 + v_{y_t}^2 + v_{z_t}^2)^{\frac{1}{2}}$
- 3 **end**
- 4 $\mathbf{v}_N = [v_1, v_2, \dots, v_N]$
- 5 $\mathbf{i} = \text{change_point_detection}(\mathbf{v}_N) \in \mathbb{R}^{1 \times B}$
- 6 $\mathbf{g}_i = [g_{i1}, g_{i2}, \dots]$; $\mathbf{o}_i = [o_{i1}, o_{i2}, \dots]$; $\mathbf{e}_i = [e_{i1}, e_{i2}, \dots]$
- 7 $\mathbf{States} = [\mathbf{g}_i^T, \mathbf{o}_i^T, \mathbf{e}_i^T] \in \mathbb{R}^{B \times 3}$
- 8 **return** $\mathbf{D}_s = [\mathbf{i}^T, \mathbf{States}] \in \mathbb{R}^{B \times 2}$

while *action constraints* are the pre-conditions for action execution and the effects that must be met to consider the action completed. During the *execution phase*, a BT is generated to accomplish the user-specified goal using the *memory* from the learned task (Algorithm 4). The high-level objective of “what action to perform?” and the low-level objective of “how to perform the action?” is adapted through BT and DMPs respectively based on the environmental state, as follows.

A. Segmentation

We assume that an action’s start and end will alter the end-effector’s velocity which is then used for segmentation in Algorithm 1. An offline change point detection algorithm [16] with penalty function obtains segments $\mathbf{i} = [i_1 \dots i_B]$, based on the velocity $\mathbf{v}_N = [v_1 \dots v_t \dots v_N]$, where B is the total segment number and v_t is the end effector’s translational velocity norm at timestamp t . Lastly, to generate the BTs, the state variables at each segment are grounded. To avoid false segmentation, the segments where the end-effector velocity or symbolic states are unchanged are ignored. A combination of kinesthetic and visual teaching interfaces is used to record demonstrations in the form of multivariate time series. For N timestamps, the end effector pose $\mathbf{x}_t^e = [x, y, z, w, qx, qy, qz]^T$, gripper state g_t and object poses $\mathbf{x}_t^o = [x, y, z, w, qx, qy, qz]^T$ are recorded to create a dataset \mathbf{D} where $t = 1, \dots, N$:

$$\begin{cases} \mathbf{X}^e \in \mathbb{R}^{7 \times N}, \mathbf{X}^e = [\mathbf{x}_0^e, \dots, \mathbf{x}_N^e], \\ \mathbf{g} \in \mathbb{R}^{1 \times N}, \mathbf{g} = [g_0, \dots, g_N], \\ \mathbf{X}^o \in \mathbb{R}^{7 \times N}, \mathbf{X}^o = [\mathbf{x}_0^o, \dots, \mathbf{x}_N^o]. \end{cases} \quad (1)$$

Three symbolic state variables, the gripper state \mathbf{g}_i , the target object state \mathbf{o}_i , and the end-effector state \mathbf{e}_i at the i^{th} segment, are necessary to generate a BT. In the final step, the values in \mathbf{D} are grounded into symbolic states variables for each segment according to conditions from Table II. The grounded state variables are stored in the vector \mathbf{States} ,




Fig. 3. Atomic BT. The effect-conditions c_{eff}^i are achieved by action $Action^i$. To execute the action, pre-condition c_{pre}^i should satisfy. The BT starts with a fallback node connected to a parallel node for effect-conditions and a sequence node for action and pre-conditions.

Algorithm 2: Primitive Learning

Input: Demo. dataset \mathbf{D} , Segmented dataset \mathbf{D}_s
Output: Primitive Actions $\mathbf{Actions} \in \mathbb{R}^{B-1 \times 3}$

- 1 **while** $b < B$ **do**
- 2 $\tau^b = \text{trainDMP}(\mathbf{X}^e[i^b : i^{b+1}], \mathbf{g}[i^b : i^{b+1}])$
- 3 $\mathcal{O}^b = \text{Relative_Object}(i^{b+1})$
- 4 $\mathbf{T}^b = \text{EE_transformation}(\mathcal{O}^b, i^{b+1})$
- 5 $\mathbf{Actions} \leftarrow [\tau^b, \mathcal{O}^b, \mathbf{T}^b]$
- 6 **end**

which is in turn stored in a segmented dataset \mathbf{D}_s with segments \mathbf{i} . A 5 cm threshold is used to ground \mathbf{o}_i and \mathbf{e}_i .

B. Primitive Learning

The low-level motion policy is learned using DMPs [17] as they require only one demonstration to generalize and can adapt to different start and end poses. Algorithm 2 shows that for each $Action^b$ we train a set of DMPs, τ^b , for \mathbf{X}^e and \mathbf{g} , compute the relative object \mathcal{O}^b and a transformation matrix \mathbf{T}^b for the end-effector at the segment end. Relative object \mathcal{O}^b is defined as the nearest object to the end-effector w.r.t. which end-effectors distance decreased between the action.

C. Automatic behavior tree generation

The proposed method in Algorithm 3 is inspired by logic-based declarative learning [19] and atomic BT [23]. Given a fully-observable environment, the task is simplified by simply identifying the state transitions during the demonstration [25]. If the deterministic effects in the environment can be observed, then the state transitions will happen only when an action is performed. As the demonstration is segmented into its lowest form, a state transition occurs due to a single action. Hence, the states at the segment start are considered action pre-conditions, denoted $c_{pre} \in \{g_i, o_i, e_i\}$, while the altered states at the segment end are denoted the effects, $c_{eff} \{g_i, o_i, e_i\}$. However, conditions that are pre-conditions for the next action may be ignored as they are unchanged during a particular segment. Thus, the conditions that are pre-conditions for the next action are also added as action constraints (line 14 of Algorithm 3), using the `CondAbstraction` function which returns a tuple \mathcal{T} in the form of $Action^b$ sequence with the corresponding constraint vector $\mathbf{c}^b \in \{c_{pre}^{b+1}, c_{eff}^b\}$, i.e.,

$$\mathcal{T} = \{(Action^1, \mathbf{c}^1), \dots, (Action^{B-1}, \mathbf{c}^{B-1})\}. \quad (2)$$

Each action, with its state conditions, can be converted into an atomic BT [23], which aims to satisfy the effect conditions

Algorithm 3: Behavior Tree Generation

Input: Demo. dataset \mathbf{D} , Segmented dataset \mathbf{D}_s ,
Primitive Actions **Actions**
Output: Behavior Tree Π

```
1  $\Pi \leftarrow \emptyset$ 
2 begin Behavior Tree Generation
3    $\mathcal{T} = \text{CondAbstraction}(\text{States}, \text{Actions})$ 
4    $\Pi = \text{CondToTree}(\mathcal{T})$ 
5    $\mathcal{M} \leftarrow \Pi$ 
6 end
7 Function CondAbstraction (States, Actions) :
8    $\mathcal{T} \leftarrow \emptyset, b \leftarrow 2$ 
9   while  $b \leq B$  do
10     $c_{pre}^b \leftarrow \emptyset, c_{eff}^b \leftarrow \emptyset$ 
11    while  $n < 3$  do
12      if  $\text{States}_n^b \neq \text{States}_n^{b-1}$  then
13         $c_{eff}^b \leftarrow \text{States}_n^b$ 
14        if  $\text{States}_n^b = \text{States}_n^{b+1}$  then
15           $c_{pre}^{b+1} \leftarrow \text{States}_i^b$ 
16        end
17      end
18       $\mathcal{T} \leftarrow \text{Action}^b, c_{pre}^{b+1}, c_{eff}^b$ 
19    end
20  end
21  return  $\mathcal{T} \in \mathbb{R}^{B-1}$ ;
22 Function CondToTree ( $\mathcal{T}$ ) :
23    $b \leftarrow 1$ 
24   while  $b \leq B - 1$  do
25      $\text{Parallel} \leftarrow \mathbf{c}^b$ 
26     if  $\Pi = \emptyset$  then
27        $\text{Sequence} \leftarrow \text{Action}^b$ 
28     else
29        $\text{Sequence} \leftarrow \Pi, \text{Action}^b$ 
30     end
31      $\text{Fallback} \leftarrow \text{Parallel}, \text{Sequence}$ 
32      $\Pi \leftarrow \text{Fallback}$ 
33  end
34  return  $\Pi$ 
```

and invokes the action if any of the effect conditions are unmet and the preconditions satisfied. An illustrative atomic BT for an action with pre- and effect-condition is shown in Fig. 3. Through the CondToTree function in Algorithm 3, actions of the \mathcal{T} are converted into an atomic BT. Finally, each atomic BT is chained in a sequence assuming that the pre-condition for an action to occur is the previous action’s effects’ success. The generated BT for a particular target and its respective **Actions** and target object are saved in the memory dictionary \mathcal{M} .

D. Goal adaptation and composite BT generation

Real-world long-horizon tasks, such as placing an object in a drawer, can be hierarchically decomposed into a sequence of shorter-horizon sub-tasks, i.e., opening a drawer and placing an object. Thus, we exploit the modular nature of BTs to accomplish such tasks. The goal of the memorised BT

Algorithm 4: Goal adaptation and composite BT generation

Input: Memory \mathcal{M} , User-specified goal scene \mathcal{G}
Output: Composite BT Π^c

```
1 begin
2    $\text{found\_objects}, \text{new\_goal} = \text{AdaptGoal}(\mathcal{G}, \mathcal{M})$ 
3    $\Pi^c = \text{CompositeBT}(\text{found\_objects}, \text{new\_goal})$ 
4 end
5 Function AdaptGoal ( $\mathcal{G}, \mathcal{M}$ ):
6    $\text{All\_objects} = \text{find\_all\_objects}(\mathcal{G})$ 
7    $\text{found\_objects} = \text{match\_objects}(\mathcal{M}, \text{All\_objects})$ 
8   if  $\text{found\_objects} \neq \text{None}$  then
9     if  $\text{goal\_object\_found}$  then
10       $\text{save\_new\_goal}(\text{w.r.t goal\_object})$ 
11    else
12       $\text{save\_new\_goal}(\text{w.r.t. nearest\_non-goal\_object})$ 
13    end
14  end
15  return  $\text{found\_objects}, \text{new\_goal}$ 
16 Function CompositeBT ( $\text{found\_objects}, \text{new\_goal}$ ) :
17   forall  $\text{found\_object}$  do
18      $\text{Sequence} \leftarrow \Pi(\text{found\_object}, \text{new\_goal})$ 
19   end
20    $\Pi^c \leftarrow \text{Sequence}$ 
21  return  $\Pi^c$ 
```

is adapted and different adapted BTs are sequenced to form a composite BT. The aim is to show that the learned BTs can be combined for complex tasks without extra demonstrations, while emphasizing that hierarchical decomposition and long-horizon plan generation is not this work’s primary goal. Instead, off-the shelf high-level planning methods can be utilized to generate long-horizon plans, which can then be executed using the BTs developed in this study.

In this work, an effective goal abstraction method, Algorithm 4, is used to adapt and form composite BT. First, all the objects present in the user-specified goal scene \mathcal{G} are matched with the objects in memory \mathcal{M} . If an object is found, the new goal position with respect to the goal object or nearest non-memorized object is saved which will be treated as the new goal. When multiple objects are matched, then the object with same goal object in the memorised BT is given higher preference in the execution sequence.

IV. EXPERIMENTS AND RESULTS

A. Experimental setup

To validate CoBT, a UR5e collaborative robot equipped with a Robotiq 2-finger adaptive gripper is used, with OptiTrack motion capture system used for object detection. The system is implemented with ROS Noetic and ROS Foxy, with change point detection using the *ruptures*¹ Python library [16], DMPs using the *movement_primitives*² Python library

¹<https://github.com/deepcharles/ruptures>

²https://github.com/dfki-ric/movement_primitives




Fig. 4. 7 evaluation tasks that include mix of complex and P2P trajectory executions, and short and long-horizon tasks with multi-level goals.




Fig. 5. (Top) Segmentation based on velocity and gripper state. (Middle) Transitions during the *drawer* task demonstration. For example, the state transition from a to b (middle) occurs due to action in segment 1 (top). (Bottom) a trial example of the generated policy under normal conditions.

and BTs with the *py_trees*³ and *py_trees_ros*⁴ Python libraries. The *py_trees* interface acts both as an operator display for task progress and carries out decision-making.

B. Evaluation Tasks

For evaluation, 7 tasks, shown in Fig. 4, are chosen. A *pick-and-place* (*P&P*) task where a cube is picked and placed on the white tray. An *insert* task where a $35\text{mm}\times 35\text{mm}$ cube is inserted inside a slot of $40\text{mm}\times 40\text{mm}$. An *erasing* task where a white board is erased using a duster. A *drawer* task where the robot must open a drawer. A *pouring* task where the contents of one cup are transferred to the other cup. A *kitting* task where a user-specified kit is prepared using three objects. A *drawer P&P* task where a robot opens a drawer and places a cube inside. No task specific alterations in the algorithms are made to validate the generalizability of CoBT.

C. Demonstration to Behavior Tree

This section describes the full pipeline for the *drawer* task. Hand-guiding is used for demonstration as shown in Fig. 5 (middle). The task is performed in one continuous demonstration during which relevant time-series data is recorded to create the dataset \mathbf{D} and in turn segment into actions as shown in Fig. 5 (top). A state transition from a to b of Fig. 5 (middle) occurs due to *Action*¹ which corresponds to segment 1 in Fig 5 (top). Similarly, each state transition corresponds to an action and for each action, a set of cartesian DMPs are trained. Here the target object is the handle and the goal object is the whole drawer set. The segmented state variables are converted into the tuple $\mathcal{T}=\{(Action^1, [e_2(\text{Near})]), (Action^2, [g_3(\text{Closed}), e_3(\text{Near})]), (Action^3, [o_4(\text{Goal}), e_4(\text{Near})]), (Action^4, [g_5(\text{Open}), o_5(\text{Goal})]), (Action^5, [o_6(\text{Goal}), e_6(!\text{Near})])\}$ through the Algorithm 3 function

³https://github.com/splintered-reality/py_trees

⁴https://github.com/splintered-reality/py_trees_ros

TABLE III

SUCCESS TRIALS AND TIMINGS FOR OUR 7 EVALUATION TASKS

Tasks	Success Trials (Normal)	Success Trials (Reactive)	Demo. Time	Prog. Time
P&P	20/20	20/20	23.1s	7s
Insert	16/20	5/5	23.5s	6.7s
Erasing	19/20	5/5	26.4s	7s
Drawer	18/20	4/5	26s	8.5s
Pouring	17/20	5/5	24.5s	8.6s
Kitting	19/20	5/5	0s	0.5s
Drawer P&P*	19/20	4/5	0s	0.5s

* Composite BT

CondAbstraction. \mathcal{T} is then iteratively converted into a BT by chaining atomic BTs in the same sequence. Fig. 1 shows the generated BT for the *drawer* task. Demonstration data, segmented state variables and generated BTs for all the 7 tasks are available at this [link](#).

For performance evaluation 20 trials were conducted for each task, where target and goal object’s start pose are changed within $80\text{cm}\times 40\text{cm}$ area on the table between trials. An equal amount of reverse trials, where the opposite trajectory to the demonstration is needed, are performed except for *drawer* and *drawer P&P* task due to the limitations in the robot’s workspace. All trials are slowed at least 2 times using DMP time scaling to ensure safety. The number of successful trials in normal conditions for each task is shown in Table III. Notably, the *P&P* task achieved a 100% success rate, while the *insert* task exhibited the highest number of failures. During 20 attempts to execute the *insert* task, there were 4 instances where the cube did not smoothly slide into the slot, primarily due to vision-related errors. These errors included issues such as failure to grasp the target object in the correct pose and inaccuracies in computing the goal pose. Indeed vision errors emerged as a dominant factor contributing to failures across all tasks. 2 *drawer* task failures, 1 *drawer P&P* task failure and 1 *pouring* task failure were attributed to robot joint limits. The remaining 2 *pouring* task failures occurred during reverse trajectory trials. In these cases, the robot tried to move the cup to the opposite side of the goal object to initiate the pouring action resulting in collisions since the learned transformation ${}^O\mathbf{T}_e$ is not reversed in the reverse trajectory. In normal conditions with varying start poses, CoBT achieved over 91% success rate and the average programming time from demonstration amounted to 7.5s.

D. Reactivity

To evaluate the reactivity, each task was executed at least 5 times, with variations in object positions during executions. The *P&P* task’s reactivity was evaluated over 20 trials as all manipulation tasks incorporate elements of the *P&P* action. For each sub-task within *P&P*, i.e., reaching the target object, grasping, moving toward the tray, and placing the object, we altered the object’s position 5 times. Table III illustrates the number of successful trials during reactive conditions showing that CoBT-generated programs exhibit a high degree of reactivity, with nearly 100% reactivity to environmental changes. When an object, target or goal position is altered during execution, the effect conditions for that particular segment will not be successful. Therefore, BT will reattempt that particular segmented of action until all

the effect conditions are met. The reactivity is due to the fact that the system not only generalizes on the motion level but on the task level. Due to state-action mapping during the segmentation process, the BT will execute the same action in the loop until the success state is achieved. Furthermore, even if the task is reset afterward, the BT will keep running in the loop until the task’s goal conditions succeed. The 2 instances of failed reactive trials occurred while attempting to open the *drawer*. These failures were attributed to the constraints imposed by the robot’s joint limits.

E. Modularity

In this experiment, we show the modularity of our generated BT. A user-specified goal \mathcal{G} generates a composite BT through algorithm 4. In the *kitting* task, the system is shown three objects placed on a tray and for *drawer P&P* task, a cube inside the opened drawer is shown. Using algorithm 4, the system adapts goals of individually demonstrated and generated BTs and sequences them. For *kitting* the target objects are the three black pieces and goal object is the tray, and for *drawer P&P* the target objects are the handle and cube and the goal object is the drawer set. The demonstration for all the target object are recorded individually.

To evaluate the modularity, *kitting* and *drawer P&P* tasks are executed for the same 20 normal and 5 reactive times. The results in Table III show that each task had only one failure during normal trails and only *drawer P&P* task failed once in reactive trial. These tasks *did not* require new demonstrations as they were adapted from previously existing demonstrations i.e. the *kitting* task is simply three different *P&P* tasks sequenced and *drawer P&P* task is *drawer* and *P&P* task sequenced.

F. Pilot-study

Finally, a pilot study is conducted to verify the accessibility for novice subjects for *P&P* and *drawer* tasks, involving $N=10$ participants (2 females, 8 males, aged 22-40). Out of 10 only 5 had prior experience with collaborative robots, and 0 with PbD systems. At the start, participants could familiarize themselves with hand-guiding mode and gripper controls. Subsequently, one demonstration for each task was recorded. They were also allowed to re-try their demonstrations if dissatisfied. Importantly, no additional PbD-specific instructions were provided. Further details about the step-by-step procedure of the study, participant surveys conducted and the ethics approval are available at this [link](#).

Fig. 6 shows the number of successful trials out of 5 for each demonstration and demonstration times for each participant. Among 20 non-expert demonstrations, only one failed (P3, *drawer* task) in programming a deployable BT due to conflicting effect conditions of consecutive actions and was excluded from analysis. In a single case (P4), the segmentation algorithm combined two actions (grasping the handle and opening the drawer) in the *drawer* task, resulting in a BT that failed in 2 out of 5 trials but remained deployable with reduced reactivity. Minor variations in retracting the robot after object placement during the *P&P* task led to 2




Fig. 6. Successful trials and demonstration time for each participant.

failed trials. Notably, the two non-experts (P9 and P6) with a 100% success rate lacked prior experience with collaborative robots. The remaining failures were attributed to the quality of demonstrated motion and the resulting DMPs derived from them. Overall, CoBT-generated programs from non-expert demonstrations achieved a 94% success rate for the *P&P* task and approximately 87% for the *drawer* task. Finally, we conducted the raw NASA-TLX questionnaire [26], with scores ranging 1-7, to measure cognitive load and a Systems Usability Scale (SUS) questionnaire [27] with scores ranging 1-5 which was further scaled between 1-100 for final score. On average, participants rated *mental demand* at 2.1 and *effort* at 2.7, while *physical demand* was rated at 3.10. In general, participants found the system easy to use, with average SUS score of 92, and perceived a low requirement for prior knowledge, with average SUS score of 28. In summary, both quantitative and qualitative data confirm CoBT’s ease of use.

V. CONCLUSION AND FUTURE WORK

In this paper, CoBT, a collaborative programming framework to generate BTs from one demonstration, is presented. CoBT generates deployable, reactive and modular programs in the form of DMP embedded BT and does not require any prior task or object specific knowledge. We have shown CoBT’s capability to generalise the task, the reactivity to environmental changes and modularity on 7 evaluation tasks. Our system achieved $\approx 93\%$ success rate overall with an average of 7.5s programming time. We further presented a pilot study with non-expert participants ($N=10$) that demonstrates the system’s ease of use. To the best of our knowledge, CoBT is the only framework capable of generating deployable BTs from a single demonstration. The current framework offers the potential for agile robot programming in industry. Nonetheless, there are certain limitations. CoBT’s execution accuracy is limited to users demonstration precision just like any hand-guiding system. The DMP generalizes from a single demonstration, thereby inheriting associated limitations, such as the absence of stochasticity and challenges in generalizing reverse executions. Additionally, the generated BTs may benefit from post-processing to address conflicting constraints, such as with the non-expert demonstration, utilizing interactive human inputs, as suggested in [28], a direction we intend to explore in future research.

REFERENCES

- [1] A. Jain, S. Mehak, P. Long, J. D. Kelleher, M. Guilfoyle, and M. C. Leva, "Evaluating safety and productivity relationship in human-robot collaboration," *Health*, vol. 2022, pp. 08–28, 2022.
- [2] J. E. Laird, K. Gluck, J. Anderson, K. D. Forbus, O. C. Jenkins, C. Lebiere, D. Salvucci, M. Scheutz, A. Thomaz, G. Trafton *et al.*, "Interactive task learning," *IEEE Intelligent Systems*, vol. 32, no. 4, pp. 6–21, 2017.
- [3] M. Kyrarini, M. A. Haseeb, D. Ristić-Durrant, and A. Gräser, "Robot learning of industrial assembly task via human demonstrations," *Autonomous Robots*, vol. 43, pp. 239–257, 2019.
- [4] G. Lentini, G. Grioli, M. G. Catalano, and A. Bicchi, "Robot programming without coding," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 7576–7582.
- [5] M. Knaust and D. Koert, "Guided robot skill learning: A user-study on learning probabilistic movement primitives with non-experts," in *2020 IEEE-RAS 20th International Conference on Humanoid Robots (Humanoids)*. IEEE, 2021, pp. 514–521.
- [6] O. Gustavsson, M. Iovino, J. Styurd, and C. Smith, "Combining context awareness and planning to learn behavior trees from demonstration," in *2022 31st IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*. IEEE, 2022, pp. 1153–1160.
- [7] G. Suddrey, B. Talbot, and F. Maire, "Learning and executing re-usable behaviour trees from natural language instruction," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 10643–10650, 2022.
- [8] T. Eiband, C. Willibald, I. Tannert, B. Weber, and D. Lee, "Collaborative programming of robotic task decisions and recovery behaviors," *Autonomous Robots*, vol. 47, no. 2, pp. 229–247, 2023.
- [9] O. Kroemer, C. Daniel, G. Neumann, H. Van Hoof, and J. Peters, "Towards learning hierarchical skills for multi-phase manipulation tasks," in *2015 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2015, pp. 1503–1510.
- [10] T. Osa, J. Pajarinen, G. Neumann, J. A. Bagnell, P. Abbeel, J. Peters *et al.*, "An algorithmic perspective on imitation learning," *Foundations and Trends® in Robotics*, vol. 7, no. 1-2, pp. 1–179, 2018.
- [11] M. Diehl, C. Paxton, and K. Ramirez-Amaro, "Automated generation of robotic planning domains from observations," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 6732–6738.
- [12] L. Scherf, A. Schmidt, S. Pal, and D. Koert, "Interactively learning behavior trees from imperfect human demonstrations," *Frontiers in Robotics and AI*, vol. 10, 2023.
- [13] A. M. Zanchettin, "Symbolic representation of what robots are taught in one demonstration," *Robotics and Autonomous Systems*, vol. 166, p. 104452, 2023.
- [14] Y. S. Liang, D. Pellier, H. Fiorino, and S. Pesty, "iipro: An interactive robot programming framework," *International Journal of Social Robotics*, pp. 1–15, 2022.
- [15] K. French, S. Wu, T. Pan, Z. Zhou, and O. C. Jenkins, "Learning behavior trees from demonstration," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 7791–7797.
- [16] C. Truong, L. Oudre, and N. Vayatis, "Selective review of offline change point detection methods," *Signal Processing*, vol. 167, p. 107299, 2020.
- [17] A. Ude, B. Nemeč, T. Petrić, and J. Morimoto, "Orientation in cartesian space dynamic movement primitives," in *2014 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2014, pp. 2997–3004.
- [18] N. Sünderhauf, O. Brock, W. Scheirer, R. Hadsell, D. Fox, J. Leitner, B. Upcroft, P. Abbeel, W. Burgard, M. Milford *et al.*, "The limits and potentials of deep learning for robotics," *The International journal of robotics research*, vol. 37, no. 4-5, pp. 405–420, 2018.
- [19] P. Kordjamshidi, D. Roth, and K. Kersting, "Declarative learning-based programming as an interface to ai systems," *Frontiers in Artificial Intelligence*, vol. 5, 2022.
- [20] D. H. Grollman and O. C. Jenkins, "Incremental learning of subtasks from unsegmented demonstration," in *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2010, pp. 261–266.
- [21] N. Nejati, P. Langley, and T. Konik, "Learning hierarchical task networks by observation," in *Proceedings of the 23rd international conference on Machine learning*, 2006, pp. 665–672.
- [22] M. Colledanchise and P. Ögren, *Behavior trees in robotics and AI: An introduction*. CRC Press, 2018.
- [23] M. Colledanchise, D. Almeida, and P. Ögren, "Towards blended reactive planning and acting using behavior trees," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 8839–8845.
- [24] E. Valassakis, G. Papagiannis, N. Di Palo, and E. Johns, "Demonstrate once, imitate immediately (dome): Learning visual servoing for one-shot imitation learning," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 8614–8621.
- [25] S. Jiménez, T. De La Rosa, S. Fernández, F. Fernández, and D. Borrajo, "A review of machine learning for automated planning," *The Knowledge Engineering Review*, vol. 27, no. 4, pp. 433–467, 2012.
- [26] S. G. Hart and L. E. Staveland, "Development of nasa-tlx (task load index): Results of empirical and theoretical research," in *Advances in psychology*. Elsevier, 1988, vol. 52, pp. 139–183.
- [27] J. Brooke, "Sus: a 'quick and dirty' usability," *Usability evaluation in industry*, vol. 189, no. 3, pp. 189–194, 1996.
- [28] M. Iovino, F. I. Doğan, I. Leite, and C. Smith, "Interactive disambiguation for behavior tree execution," in *2022 IEEE-RAS 21st International Conference on Humanoid Robots (Humanoids)*. IEEE, 2022, pp. 82–89.