

On the programming effort required to generate Behavior Trees and Finite State Machines for robotic applications

Matteo Iovino^{a,b}, Julian Förster^c, Pietro Falco^a, Jen Jen Chung^{c,d}, Roland Siegwart^c and Christian Smith^b

Abstract—In this paper we provide a practical demonstration of how the modularity in a Behavior Tree (BT) decreases the effort in programming a robot task when compared to a Finite State Machine (FSM). In recent years the way to represent a task plan to control an autonomous agent has been shifting from the standard FSM towards BTs. Many works in the literature have highlighted and proven the benefits of such design compared to standard approaches, especially in terms of modularity, reactivity and human readability. However, these works have often failed in providing a tangible comparison in the implementation of those policies and the programming effort required to modify them. This is a relevant aspect in many robotic applications, where the design choice is dictated both by the robustness of the policy and by the time required to program it. In this work, we compare backward chained BTs with a fault-tolerant design of FSMs by evaluating the cost to modify them. We validate the analysis with a set of experiments in a simulation environment where a mobile manipulator solves an item fetching task.

Index Terms—Behavior Trees, Finite State Machines, Modularity, Mobile Manipulation

I. INTRODUCTION

Nowadays, robots are deployed in unstructured environments, usually shared with humans, where the need for reactivity to handle unpredictable situations and fault-tolerance to increase the robustness of the robot behavior are vital for the successful outcome of the task. Robotics engineers have to take into account these and other features, e.g. programming time, maintainability, readability, when deciding with which policy to control the robot.

Finite State Machines (FSMs) have been the standard policy representation in robotics due to their intuitive and simple design and the predictability of their execution flow. They have flourished in industrial applications, where the robot repeatedly executes the same task in a static environment with a low failure probability. As an alternative, Behavior Trees (BTs) are a reactive, readable and modular task switching policy representation, becoming state of the art for robot control, especially in research environments [1].

We believe that examples of how these two representations compare on a concrete implementation are still lacking. Therefore, the goal of this paper is to focus the comparison

This project is financially supported by the Swedish Foundation for Strategic Research. The authors gratefully acknowledge this support.

^aABB Corporate Research, Västerås, Sweden

^bDivision of Robotics, Perception and Learning, KTH - Royal Institute of Technology, Stockholm, Sweden

^cAutonomous Systems Lab, ETH Zürich, Zürich, Switzerland

^dSchool of ITEE, The University of Queensland, Australia

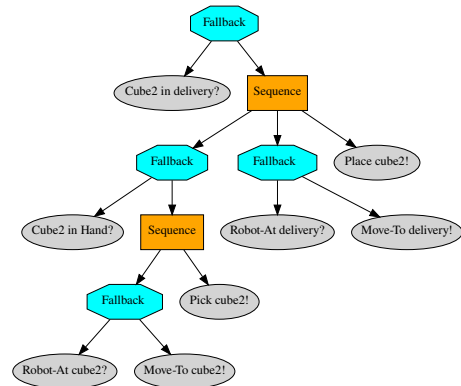


Fig. 1: Backchained BT solving a mobile manipulation task.

between BTs and FSMs on the modularity of the two representations, also carrying out experiments in a simulation environment. In this work, we emphasize modularity since it is a key feature that allows users to drastically reduce the programming effort required for a robotics application, especially when it is possible to reuse existing code or libraries. First, we measure modularity in terms of Cyclomatic Complexity (CC), Edit Distance (ED), and Computational Complexity of the operations required to edit the policy. Then, we show the practical implications of choosing one representation rather than the other with a series of simulated robotic tasks.

To make the comparison as fair as possible, we make a few assumptions on the design:

- 1) The two representations have access to the same robot skill set, so that the difference remains at as high a level as possible;
- 2) We design BTs with the backward-chained (backchained) principle because it allows deriving stability guarantees and convergence proofs [2], even though it is not the most optimal choice in terms of the size of the tree (as is found by e.g. Genetic Programming algorithms [3], [4]);
- 3) We propose to build FSMs in a fault-tolerant way, so that status checks on the environment redirect the robot to the point in the execution flow that had previously failed. This mimics the reactivity of BTs. We validate this choice in a direct comparison with a standard sequential FSM.

For implementation, we use state-of-the-art open-access programming tools—BTs: `py_trees` [5]; FSMs: `SMACH` [6] [7]—which are widely used in robotic applications [8].




Fig. 2: Sequential FSM solving the same task of Fig. 1.

II. BACKGROUND AND RELATED WORK

This section briefly defines BTs and FSMs and presents the theoretical foundations and early results on the comparison between the two representations.

A. Behavior Trees

BTs are a task switching policy representation, originating in the gaming industry and later transferred to robotics [9].

A BT is a directed tree, where a tick signal originates from the root and propagates down the tree from left to right. Nodes execute only when ticked and return one of the status signals *Success*, *Failure*, and *Running*. Internal nodes are called *control nodes*, (polygons in Fig. 1), with the most common types being *Sequence*: runs children in a sequence, returning once all succeed or one fails, and *Fallback* (or *Selector*): runs children in a sequence, returning when one succeeds or all fail. Leaves are called *execution nodes* or *behaviors* (ovals in Fig. 1), of type Action(!) or Condition(?). The former encode robot skills, while the latter encode status checks and sensory readings, thus immediately returning *Success* or *Failure*.

BTs have explicit support for task hierarchy, action sequencing and reactivity [1] and are modular by design: every node receives the tick as input and outputs the return statuses, so subtrees can be moved without compromising the logical functioning. Moreover, modularity allows every building block to be independently tested and reused. Reactivity is realized by the *Running* return state, which allows high priority actions to preempt executing ones.

A backchained BT is built from the goal condition, by expanding it with the actions that achieve it and then by recursively expanding those actions' unmet pre-conditions. As an example, in the BT of Fig. 1, placing the cube in the delivery station allows the robot to complete the task, but in order to reach that stage, the robot first has to pick the cube and then move to the delivery station. Recursively, to pick the cube the robot has to first move to where the cube is located. The backchained design can also be automatically generated by a planner [4], [10] or learnt [11], [12], making it appealing to operators with limited programming knowledge.

B. Finite State Machines

Finite State Machines derive from state automata and feature a set of states and transitions between them (ovals

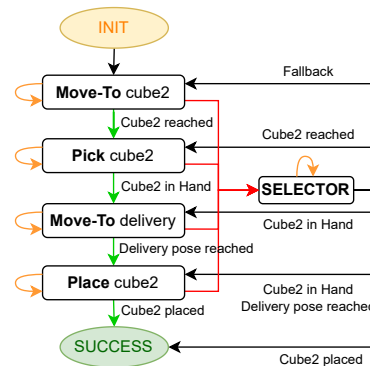


Fig. 3: Fault-tolerant design for FSMs solving the same task of Fig. 1. States' outcomes *Success*, *Running* and *Failure* have green, yellow and red transitions, respectively.

and arrows in Fig. 2, respectively). Every state encodes a controller for robot behavior which produces effects in the environment upon execution. The effects trigger an event that transfers the execution from one state to the next. Since FSMs also include Sequential Function Charts (SFCs), a graphical programming language for Programmable Logic Controllers (PLCs), they are widely used in industry [13] and their success is mainly due to their intuitive design and implementation simplicity.

FSMs have the unfortunate shortcoming that designers have to make a trade-off between reactivity and modularity. As highlighted in [9], FSM execution can be compared to the GoTo statement of early programming languages, where the execution flow jumps from one part of the program to another and continues from there. In programming, and by extension also in robotics, GoTo statements are considered to be harmful and are advised against [14]. On the other hand, the execution of a BT can be compared to a function call, where the execution flow also jumps to another part of the code but after its completion, it returns to where the function was initially called.

In order to be reactive, an FSM needs to have many transitions that need to be taken care of upon addition or removal of states, making them less modular and less scalable. The modularity problem is partially mitigated by logically grouping states to form hierarchies. There is a particular design for Hierarchical State Machines (HFSMs) aimed to mimic the functionality of a BT and to grant FSMs some modularity [15]. Since this design sacrifices the readability, we leave the analysis of the modularity of HFSMs to future work.

Another shortcoming of sequential FSMs (e.g. Fig. 2) is that the robot and the environment have to be reset upon failure, and the task to be re-executed again from the start. To make the FSM fault-tolerant, we propose to add a *SELECTOR* state that is fully connected to the others (as in Fig. 3). In case of failure of any of the action states, the execution flow passes to the *SELECTOR* state, where the status of the environment and the task progress is checked, allowing the robot to resume the execution in the correct state. We endow every state with an explicit *Running* transition that cycles back to the state itself, so that

the action execution is asynchronous and the environment is monitored periodically. An action's post-conditions (also referred to as effects) are the transitions to the next state (or from SELECTOR to another state), while its pre-conditions are reflected in the sequential order of the states in the FSM. Note that this design is not novel and an alternative fault-tolerant and reactive design would feature a task switcher node to which all other states are connected and which is responsible for checking the state of the environment and deciding the next sequence of actions. In this design, however, we lose the sequentiality of the execution flow and, consequently, we also lose the readability as it is harder for a human operator to anticipate what the robot will do next.

C. Related Work

BTs have been compared to FSMs in previous work, but the comparison was either purely theoretical [9], [15], [16], or speculative [17]. In [9], authors list advantages and disadvantages of both designs and in [15] they prove theoretically how in fact BTs modularize FSMs, providing an HFSM design that behaves like a BT. In [16], authors formally compare BTs with other related architectures (Decision Trees, Teleo-reactive Programs and in particular FSMs) in terms of reactivity, readability and expressiveness.

In [18], the same authors formalize modularity for reactive control architectures, pointing out that BTs feature structural interfaces, with which *every component* interacts with the others: a subtree is a BT, an action behavior is a degenerate case of a BT. Naturally, FSMs lack structural interfaces and thus cannot be considered modular in their analysis unless a structure is enforced (e.g. in HFSMs). They measure modularity with the Cyclomatic Complexity (CC) that we will describe later in the analysis. In addition we propose to measure modularity in terms of the effort required to modify a structure by adding or removing elements in it. We propose to quantify it using the Computational Complexity of such operations and the Edit Distance (ED) between a baseline structure and its modified versions.

Klöckner [17] proposes using a BT as a control policy for UAV missions, speculating the advantages of such design with respect to FSMs. Previous work [19] also compares the two policy representations from a practical perspective in the domain of autonomous driving. The comparison is made in terms of CC and Maintainability Index (MI). The MI of a piece of software takes into account the CC, the number of lines of code, the percentage of comments in the code and the Halstead Volume (a function of distinct and total numbers of operands and operators). Here, the CC was ill-defined as compared to [18] and we argue that the MI is more dependent on the library used to implement the two policies, so we disregard it.

Another work that provides insight on the behavior of an agent controlled by both policies is [20] in the domain of computer games, i.e. in a deterministic scenario. Here, the authors evaluate the policies on solving a level in the Mario AI benchmark, comparing them in terms of number of nodes and reward function $\rho(x)$. This comparison provides insight

on the scalability of the two policies, since both of them are generated by a learning algorithm to solve the benchmark. The goal of this work was to propose a new method to generate BTs, so the comparison against FSM is not fair in the sense that the generation method is not the same for the two policies. In both [19] and [20] the complexity scales linearly for BTs and quadratically or worse for FSMs. From an application point of view, [21] provides guidelines to program BTs for robotic tasks, in terms of parameter passing, asynchronous calls to robot actions, and reactivity, which are also applied here in the experimental section.

The cited literature exhaustively covers the theoretical comparison between BTs and FSMs, but we believe that what is lacking is a practical grasp on the difference between them. Therefore, we propose to support the analytical comparison with a simulated robotic application, where we provide insights on the programming effort required to design both structures and comment on which one may be preferable over the other.

III. ANALYTIC COMPARISON ON MODULARITY

We use the Computational Complexity, Edit Distance (ED), and Cyclomatic Complexity (CC) to compare the modularity of the two representations.

The possibility of adding and removing nodes/states allows users to reuse pieces of software and to modify bad design choices. It can happen, for example, that during the testing of the policy we realize that some behaviors were missing or that some were malfunctioning, and thus need to be removed.


1) Computational Complexity:

a) *Add or remove a node in a BT:* Adding a node in a BT with n nodes, requires inserting it in the list of children of the parent control node at the desired position. In a BT, each child node is in no way connected to other children¹. Performing the insertion operation requires accessing the desired parent node and then inserting it, with complexity $\mathcal{O}(1)$. Removing a node has the same complexity. Because of modularity, handling a subtree has the same complexity. In this case, the root of the subtree to add/remove is considered as the input node.


b) *Add or remove a node in a FSM:* For the proposed design of the FSM, we identify two possible ways a state can be added:

- i) **sequential state:** add a new state between two existing states (namely the *preceding* and *following* states). If we want to add the new state as a new step in the execution sequence, we need to remove the transition from *preceding* to *following*, create a transition from *preceding* to the new state and from this to *following*. Finally we need to handle the interaction with the SELECTOR state. If the new state is a terminal one, then we would make a transition to the outcome instead. A FSM modified this way is reported in Fig. 4d.
- ii) **connected state:** add a new state for the task that is connected to all the others (Algorithm 1). Here, we add

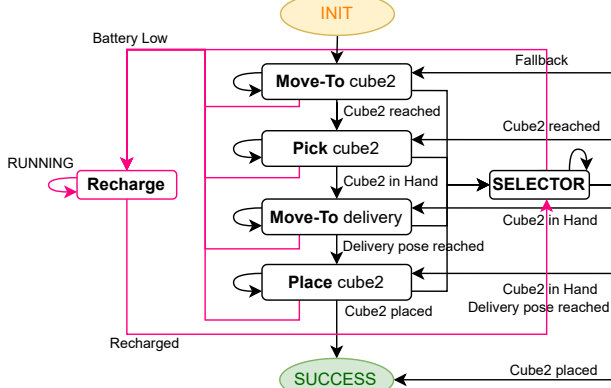
¹While not true in the case of using Blackboard variables, this design choice breaks modularity and it is discouraged [9].



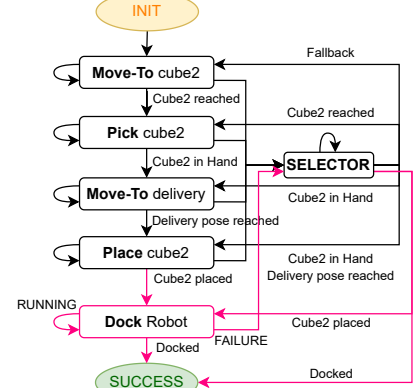
(a) Adding a subtree that recharges the battery if low.



(b) Adding a state that docks the robot once the task is done.



(c) Adding a state that recharges the battery if low.



(d) Adding a state that docks the robot.

Fig. 4: How the addition of a new behavior is handled in the two policy representations. To save space, we collapsed the BT in Fig. 1 in the white box labelled ‘MoMa task!’. In the FSM, additions are colored in magenta.

Algorithm 1: FSM: Addition of a connected state

input: sm, new_state, condition, idle_condition

```

1 for state in sm:
2   state.register_outcome(new_state)
3   if node == SELECTOR:
4     node.add_transition(idle_condition)
5   else:
6     node.add_transition(condition)
7 new_transitions ← "RUNNING" : new_state
8 new_transitions ← "FAILURE" : SELECTOR
9 sm.add_state(new_state, new_transitions)

```

a transition from every other state to the new one and finally to handle the interaction with the SELECTOR state. A FSM modified this way is reported in Fig. 4c.

In a FSM with n states, adding or removing a new state requires checking the consistences between state outcomes and transitions, which has complexity $\mathcal{O}(n)$ because we might need to take care of up to n transitions.

In particular, to remove a state, it is necessary to delete all transitions to and from the state. Then, the target state has to be removed from the outcomes of any other state, if present. This operation requires looping through all transitions and all states, for a total complexity of $\mathcal{O}(2n)$.

A fundamental difference with BTs—a direct consequence of BT modularity—is that while editing a FSM it is necessary to have access to all states and transitions, making the operation of adding/removing one element in the structure dependent on all the others.

2) **Edit Distance (ED):** ED is a way to measure diversity between structures. With the formulation proposed in [22], two BTs to compare are padded with empty nodes into the same shape. Since this formulation of the ED is defined for trees alone, we cannot directly apply it to FSMs. Using the

fact that trees are a special type of graph, we propose instead to use the Graph Edit Distance (GED). GED is defined as the minimum number of edit operations (add/remove/substitute nodes and edges) to execute on a graph g_1 to make it isomorphic to another graph g_2 [23]. Using the standard tuple definition of graphs, $G = (V, E)$ with V, E the sets of vertices and edges respectively, the GED between $g_1 = (V_1, E_1)$ and $g_2 = (V_2, E_2)$, is defined as:

$$GED(g_1, g_2) = \min_{e_1, \dots, e_k \in \gamma(g_1, g_2)} \sum_{i=1}^k c(e_i), \quad (1)$$

where c is the cost of the edit operation e_i and $\gamma(g_1, g_2)$ denotes the set of edit paths to transform one graph into the other. For this analysis, we use this measure as implemented in the NetworkX Python library [24].

3) **Cyclomatic Complexity (CC):** This is a software metric that measures the complexity of a program in the form of a control-flow graph and is defined as,

$$CC = a + s - n + 1, \quad (2)$$

where a, s, n represent the number of arcs, sinks (terminal nodes) and nodes in a decision structure, respectively. This measure is applied in [18] to BTs that have been transformed to graphs with single entry and exit nodes. This implies that **BTs have optimal modularity** as their Cyclomatic Complexity is 1.

IV. EXPERIMENTS

We compare BTs and FSMs in the concrete case of controlling a robot to perform a simulated mobile manipulation (MoMa) task and we evaluate the metrics described above on some specific edits. The MoMa task shown in Figs 1 and 3 is used as the baseline. We target tasks that are typical and




Fig. 5: Gazebo Simulation environment with the mobile manipulator.

representative of the MoMa domain, not to bias the analysis in favor of one policy or the other. To this end, we consider a set of variations of the **Cleaning Up** task as defined in the RoboCup@Home benchmark [25]. In this task a robot needs to collect a set of five unknown objects dispersed throughout the scene. To do so, the robot has to find the potential object, grasp it, and finally place it in a predefined area in the scenario. The map is known in advance. To perform this task the robot needs functional abilities such as navigation, object recognition, and object manipulation, as well as system properties such as adaptivity (the object locations are not known in advance and the environment is unpredictable), robustness, and general applicability.

We consider the simpler subtask of fetching one single object of known position, as it facilitates the comparison and reduces the execution time. The evaluation is divided into pairs of experiments, where the task becomes more complicated as new behaviors are added to the policy. The task is solved once with a BT and once with a FSM. We aim at studying the behavior of the robot as well as the effort of programming such a policy, in terms of the number of operations required. The code as well as the figures of the policies for all the experiments are available online². To complexify the task, we add functionalities according to the following scenarios, Fig. 4:

- The robot must recharge its batteries when they run low.
- The robot must dock once the task is completed.

In the simulation environment for the task (Fig. 5), the robot docks at the *inspection table*. The mobile manipulator is a Franka Emika Panda arm mounted on the omnidirectional Clearpath Ridgeback base. The arm’s end-effector is equipped with an Intel RealSense RGBD camera.

A. Skill Description

Since the scope of the paper is to compare the high-level behavior of the control policies, we make some assumptions to simplify the execution of the low-level skills. In the case of the navigation, the robot has a limited set of target poses available (one for each station/table of the environment). The navigation is implemented with the ROS Navigation Stack. We use AprilTags to identify the location of the

items. Manipulation is implemented with ROS MoveIt!. For grasping, we rely on a shape-based grasp synthesis method called the Volumetric Grasping Network (VGN) [26].

In the ROS framework, skills are implemented as action servers and a client interface allows the policy (BT or FSM) to interact with them. In the following, we describe the skills and the conditions available to the robot.

a) **Move-To!**: this action moves the robot to a target pose given explicitly or reconstructed from the object marker.

b) **Robot-At?**: this condition checks that the robot reached the desired pose with the desired tolerance in x , y and yaw.

c) **Pick!**: this action moves the robot arm to a target pose and grasps the object. Once the grasp is successful, the arm is tucked in a compact configuration for navigation.

d) **In Hand?**: this condition checks if the robot is holding an object.

e) **Place!**: this action moves the robot arm to a target pose and drops the object. Then the robot moves the arm to a configuration that allows it to monitor the scene and evaluate the end pose of the object with the wrist camera.

f) **Object-At?**: this condition checks that the object is in the desired pose with the desired tolerance in x , y , and z .

g) **Recharge!**: this action makes the robot go to the recharge station and then it instantaneously fills up the battery level.

h) **Battery Lv?**: this condition determines when the robot shall recharge its batteries.

i) **Dock!**: this action makes the robot dock at the *inspection table*. The corresponding condition is ‘Robot-At’?

B. Experiment 1

In this experiment the robot has to fetch the cube placed at ‘*fetch table 1*’ in Fig. 5 and place it on the ‘*delivery station*’. The robot starts in the center of the room and the position of the cube is known a priori. This is the only task that is also attempted with the sequential FSM of Fig. 2. The BT and fault-tolerant FSM solving the task are those of Fig. 1 and Fig. 3, respectively. The sequential FSM successfully solves the task only in the case all actions execute correctly and within the tolerances. Failures in the execution require the task to be reset. As per the BT and the fault-tolerant FSM, failed actions are re-attempted. In addition, since the BT is recursively ticked from the root, it keeps being executed also upon success. This implies that if the task is completed but then a human operator brings the cube to another table, the robot can react to it and go fetch the cube again (provided it knows where the cube is). To calculate the ED when adding behaviors, for this base task, the BT in Fig. 1 is a graph with 14 nodes and 13 edges, while the FSM in Fig. 3 has 6 nodes and 18 edges. This FSM has $CC = 14$.

C. Experiment 2

This experiment adds a recharging behavior to the existing manipulation task. This requires the operator to edit the policy as in Fig. 4a and Fig. 4c. In the case of the BT we would need to prepend the recharge subtree in a *Sequence*

²https://github.com/ethz-asl/bt_fsm_comparison

root node to the subtree solving the task, so the recharging subtree has higher priority. There are 8 elementary operations to perform: create the 4 nodes—the root node and the nodes in the left subtree—in Fig. 4a, add the two leaves to the Fallback node and finally add the recharge subtree and the already existing BT for the MoMa task to the new root.

Note that it is assumed that the recharge behavior loads the batteries instantaneously. This can be motivated by a situation in which for example, an operator switches the batteries of the robot when it stops at the recharge station. In the case where the batteries are gradually loaded when the robot stays in the recharge station, the recharge subtree could be given a different design to preempt charging if a more important task arises.

The BTs in Fig. 4a has 18 nodes and 17 edges, which infers an ED of 8 with respect to the baseline.

In the case of the fault-tolerant FSM, there are also 8 operations to perform: create the recharge state and the *Running* transition, add a transition from the recharge state to the SELECTOR state and from every state to the recharge state. There are two considerations to make. First, the number of new transitions to create depends on the number of already existing nodes, while in the case of the BT, the number of operations required does not depend on the size of the tree. Then, the internal logic of the states must also be modified to implement the case triggering a transition to the new state. This is a clear benefit of a BT, where the switching logic is explicitly implemented in the representation.

The FSM in Fig. 4c has 7 nodes and 25 edges, which infers an ED of 8 with respect to the baseline, and the new FSM now has $CC = 20$.

From a behavioral perspective, the execution of the task remains the same and the robot successfully goes to recharge the batteries if the level is below 20%. This threshold value is chosen arbitrarily in order to have at least one recharge per run.

D. Experiment 3

In this case, we also add a docking behavior as in Figs 4b and 4d. Since the *Sequence* root node is already there, to edit the BT it is necessary (1) to create the *Fallback* node, the condition and the action, (2) to add the two leaves to the control node, and (3) to add the subtree to the root. Again, the rest of the tree is in no way influenced by the addition.

In the FSM, besides creating the new state, the *Running* transition and the transitions to and from the SELECTOR state, we need to remove the transition from the ‘Place’ state to the *Success* outcome and create a transition from the ‘Dock’ state to the outcome.

In this experiment as well, there is no substantial difference in the executed robot behavior between the two policies. The BT has 21 nodes and 20 edges, which gives $ED = 6$ to the previous case of *Exp.2*, while the FSM has 8 nodes and 30 edges which gives $ED = 6$ and $CC = 24$.

E. Scalability

For the experiments, we have considered only one item to fetch. While a task that simple allows us to detail the

implementation and programming efforts, it doesn’t fully capture the real advantages that modularity gives to the BT when compared to FSMs. We have commented that insertion operations on FSMs depend on the number of states and transitions already in the FSM. In fact, when we scale the task to a case where the robot needs to search for 5 cubes, to fetch all of them and finally to dock, we would have a BT with 77 nodes and 76 edges and a FSM with 24 nodes and 90 edges, given the set of skills that we presented. At this point, adding a recharge behavior to the existing policy representations would be more cumbersome in the case of a FSM as we need to add the transition to the recharge state from all other states. We would have $ED = 6$ in the case of a BT (like in *Exp.3*, the root is already there) but $ED = 26$ in the case of a FSM, as the final BT would feature 80 nodes and 79 edges, while the final FSM 25 nodes and 115 edges.

V. CONCLUSIONS AND FUTURE WORK

In this paper we analyzed modularity in BTs and FSMs in terms of Computational Complexity, Edit Distance, and Cyclomatic Complexity. In order to ensure that the comparison is between representations with equivalent functionalities, we used a particular design of the FSM that gives it fault-tolerance. This is also motivated by the fact that BTs are a structured policy representation and there are implementation guidelines, while the design of a FSM is less constrained. Moreover, we built the BTs with the backchained design for the discussed reasons. The results of the theoretical analysis highlight how BTs are indeed more modular than FSMs, as already identified by related work [15], [18]. We contributed with a detailed analysis on the implementation of the two policies and deployment on a set of MoMa tasks in a simulation environment. From the experimental section, it is undeniable that FSMs are more intuitive to implement as the functioning and state-transitioning logic is more straightforward. However, it is more complicated to maintain a FSM and the number of operations needed to edit the policy depends on the number of states in the structure. Moreover, it is often necessary to change the logical behavior of some states if the task requires it. Modularity allows us to easily maintain, reuse, and extend BTs, as edit operations depend only on the parent node of the subtree to modify. Another remark is that the switching policy is more explicit in a BT, while in a FSM it is integrated within the state. As a future work, we will continue this analysis with an extended set of experiments where we aim to find the size of the policy (in number of states/action nodes) where the trade-off between ease of implementation and maintainability tips the balance in favor of BTs. Moreover, we plan to extend the comparison to other key elements for policy representation in robotics, e.g. reactivity and readability.

ACKNOWLEDGMENTS

Authors would like to thank Michel Breyer, Julian Keller and Yinyin Liu from the Autonomous Systems Lab, ETH Zürich, for their support in implementing the robot skills.

REFERENCES

- [1] M. Iovino, E. Scukins, J. Styruđ, P. Ögren, and C. Smith, "A survey of Behavior Trees in robotics and AI," *Robotics and Autonomous Systems*, vol. 154, p. 104096, Aug. 2022. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0921889022000513>
- [2] P. Ögren, "Convergence Analysis of Hybrid Control Systems in the Form of Backward Chained Behavior Trees," *IEEE Robotics and Automation Letters*, vol. 5, no. 4, pp. 6073–6080, Oct. 2020, conference Name: IEEE Robotics and Automation Letters.
- [3] M. Iovino, J. Styruđ, P. Falco, and C. Smith, "Learning Behavior Trees with Genetic Programming in Unpredictable Environments," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, May 2021, pp. 4591–4597, iSSN: 2577-087X.
- [4] J. Styruđ, M. Iovino, M. Norrlöf, M. Björkman, and C. Smith, "Combining Planning and Learning of Behavior Trees for Robotic Assembly," in *2022 IEEE International Conference on Robotics and Automation (ICRA)*, May 2022.
- [5] "py_trees," https://github.com/splintered-reality/py_trees.
- [6] J. Bohren and S. Cousins, "The SMACH High-Level Executive [ROS News]," *IEEE Robotics & Automation Magazine*, vol. 17, no. 4, pp. 18–20, Dec. 2010, conference Name: IEEE Robotics & Automation Magazine.
- [7] "Smach," https://github.com/ros/executive_smach.
- [8] R. Ghzouli, S. Dragule, T. Berger, E. B. Johnsen, and A. Wasowski, "Behavior Trees and State Machines in Robotics Applications," Aug. 2022, number: arXiv:2208.04211 arXiv:2208.04211 [cs]. [Online]. Available: <http://arxiv.org/abs/2208.04211>
- [9] M. Colledanchise and P. Ögren, *Behavior Trees in Robotics and AI : An Introduction*. CRC Press, Jul. 2018, kTH. [Online]. Available: <https://www.taylorfrancis.com/books/9780429950902>
- [10] 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)*, May 2019, pp. 8839–8845, iSSN: 2577-087X.
- [11] O. Gustavsson, M. Iovino, J. Styruđ, and C. Smith, "Combining Context Awareness and Planning to Learn Behavior Trees from Demonstration," in *2022 IEEE International Conference on Robot & Human Interactive Communication (RO-MAN)*, Aug. 2022.
- [12] M. Iovino, F. I. Dođan, I. Leite, and C. Smith, "Interactive Disambiguation for Behavior Tree Execution," in *2022 IEEE-RAS 21st International Conference on Humanoid Robots (Humanoids)*, Nov. 2022, pp. 82–89.
- [13] R. Balogh and D. Obdržálek, "Using Finite State Machines in Introductory Robotics," in *Robotics in Education*, W. Lepuschitz, M. Merdan, G. Koppensteiner, R. Balogh, and D. Obdržálek, Eds. Cham: Springer International Publishing, 2019, vol. 829, pp. 85–91, series Title: Advances in Intelligent Systems and Computing. [Online]. Available: http://link.springer.com/10.1007/978-3-319-97085-1_9
- [14] E. W. Dijkstra, "Letters to the editor: go to statement considered harmful," *Communications of the ACM*, vol. 11, no. 3, pp. 147–148, Mar. 1968. [Online]. Available: <https://doi.org/10.1145/362929.362947>
- [15] M. Colledanchise and P. Ögren, "How Behavior Trees Modularize Hybrid Control Systems and Generalize Sequential Behavior Compositions, the Subsumption Architecture, and Decision Trees," *IEEE Transactions on Robotics*, vol. 33, no. 2, pp. 372–389, Apr. 2017.
- [16] O. Biggar, M. Zamani, and I. Shames, "An Expressiveness Hierarchy of Behavior Trees and Related Architectures," *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 5397–5404, Jul. 2021, conference Name: IEEE Robotics and Automation Letters.
- [17] A. Klöckner, "Behavior Trees for UAV Mission Management," *INFORMATIK 2013: Informatik angepasst an Mensch, Organisation und Umwelt*, vol. P-220, pp. 57–68, 2013.
- [18] O. Biggar, M. Zamani, and I. Shames, "On Modularity in Reactive Control Architectures, with an Application to Formal Verification," *ACM Transactions on Cyber-Physical Systems*, vol. 6, no. 2, pp. 19:1–19:36, Apr. 2022.
- [19] M. Olsson, "Behavior Trees for decision-making in autonomous driving," p. 54, 2016.
- [20] M. Colledanchise, R. Parasuraman, and P. Ögren, "Learning of Behavior Trees for Autonomous Agents," *IEEE Transactions on Games*, vol. 11, no. 2, pp. 183–189, Jun. 2019.
- [21] M. Colledanchise and L. Natale, "On the Implementation of Behavior Trees in Robotics," *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 5929–5936, Jul. 2021.
- [22] E. Burke, S. Gustafson, and G. Kendall, "Diversity in genetic programming: an analysis of measures and correlation with fitness," *IEEE Transactions on Evolutionary Computation*, vol. 8, no. 1, pp. 47–62, Feb. 2004, conference Name: IEEE Transactions on Evolutionary Computation.
- [23] Z. Abu-Aisheh, R. Raveaux, J.-y. Ramel, and P. Martineau, "An Exact Graph Edit Distance Algorithm for Solving Pattern Recognition Problems," in *4th International Conference on Pattern Recognition Applications and Methods 2015*, Lisbon, Portugal, Jan. 2015. [Online]. Available: <https://hal.archives-ouvertes.fr/hal-01168816>
- [24] A. A. Hagberg, D. A. Schult, and P. J. Swart, "Exploring Network Structure, Dynamics, and Function using NetworkX," p. 5, 2008.
- [25] T. Wisspeintner, T. v. d. Zant, L. Iocchi, and S. Schiffer, "RoboCup@Home: Scientific Competition and Benchmarking for Domestic Service Robots," *Interaction Studies*, vol. 10, no. 3, pp. 392–426, Jan. 2009, publisher: John Benjamins. [Online]. Available: <https://www.jbe-platform.com/content/journals/10.1075/is.10.3.06wis>
- [26] M. Breyer, J. J. Chung, L. Ott, R. Siegwart, and J. Nieto, "Volumetric grasping network: Real-time 6 dof grasp detection in clutter," in *Conference on Robot Learning*, 2020.