

In Search of Trees: Decision-Tree Policy Synthesis for Black-Box Systems via Search

Emir Demirović¹, Christian Schilling², Anna Lukina¹

¹ Delft University of Technology, The Netherlands

² Aalborg University, Denmark

e.demirovic@tudelft.nl, christianms@cs.aau.dk, a.lukina@tudelft.nl

Abstract

Decision trees, owing to their interpretability, are attractive as control policies for (dynamical) systems. Unfortunately, constructing, or synthesising, such policies is a challenging task. Previous approaches do so by imitating a neural-network policy, approximating a tabular policy obtained via formal synthesis, employing reinforcement learning, or modelling the problem as a mixed-integer linear program. However, these works may require access to a hard-to-obtain accurate policy or a formal model of the environment (within reach of formal synthesis), and may not provide guarantees on the quality or size of the final tree policy. In contrast, we present an approach to synthesise *optimal* decision-tree policies given a *deterministic black-box environment and specification*, a discretisation of the tree predicates, and an initial set of states, where optimality is defined with respect to the number of steps to achieve the goal. Our approach is a specialised search algorithm which systematically explores the (exponentially large) space of decision trees under the given discretisation. The key component is a novel trace-based pruning mechanism that significantly reduces the search space. Our approach represents a conceptually novel way of synthesising small decision-tree policies with *optimality* guarantees even for *black-box environments with black-box specifications*.

Code — <https://doi.org/10.5281/zenodo.14601859>

Extended version —

<https://doi.org/10.48550/arXiv.2409.03260>

1 Introduction

Designing controllers for complex systems with the guarantee of specified behaviour remains an important challenge. Classical control synthesis can provide such guarantees given a (precise) model of the system (Belta and Sadraddini 2019). This requirement may in some cases be infeasible, which gave rise to black-box and approximate approaches, e.g., based on machine learning. As systems grow larger, interpretability is an increasingly desired specification for machine-learned policies to achieve alignment with human specifications (Rudin 2019). With the success of decision trees as interpretable machine-learning models, policies represented as decision trees have gained considerable traction (Du, Liu, and Hu 2020; Glanois et al. 2024).

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: Problem illustration. For the pendulum environment E (right), each tree (left) with predicates $\theta \geq c_i$, $i = 1, 2, 3$, coupled with the black-boxed E , produces a trace for θ plotted along the time axis (middle) together with its evaluation with respect to the black-box specification (reaching $\theta = 0$). The middle trace is deemed best since it reaches the goal faster ($t = 50 < 60$). The bottom trace is inferior.

There are diverse approaches to synthesising, or learning, decision-tree policies. Stratego (David et al. 2015) employs reinforcement learning dedicated to decision trees. Modifying the reinforcement learning process to produce decision trees has also been proposed (Topin et al. 2021). An alternative is to apply imitation learning to distil a neural-network policy into a decision tree (Bastani, Pu, and Solar-Lezama 2018). After using formal synthesis to construct a policy in tabular form, decision trees may be induced via specialised algorithms akin to algorithms used for solving standard classification problems (Ashok et al. 2021).

While previous approaches have their strengths, none of the discussed methods provide guarantees in terms of decision-tree policy performance or size and/or require an existing expert policy or effective reinforcement learning algorithm. Policy synthesis may be posed as a mixed-integer

linear programming problem (Vos and Verwer 2023), which can provide guarantees. However, this approach assumes that a model of the environment is given. When the above requirements are not met, decision-tree policy cannot be constructed using existing methods.

In contrast, we consider a unique setting: derive 1) *small* decision-tree policies when 2) the model and specification of the environment is a *deterministic black-box* whilst 3) providing *optimality guarantees* of performance, under the discretisation of tree predicates and initial states. Every work we are aware of will violate at least one of these three points.

Our work is for *deterministic black-box systems*, which pose a challenging controllability problem, as exact guarantees imply searching in an exponentially large space.

Our approach is based on search. Briefly, our algorithm systematically enumerates all possible decision trees that may be constructed using a given set of predicates, and then selects the tree that optimises the specification evaluated by the black-box environment for each tree, e.g., minimises (maximises) the time to reach (maintain) a target state.

Example 1 *As an illustrative example, consider the pendulum environment in Fig. 1. The pendulum is attached at one end to a fixed point, and available control actions are to apply force to push the free end left ($a_1 = -1$) or right ($a_2 = 1$) (Anderson 1989). Our aim is to construct a small decision-tree policy that swings the pendulum to an upward angle ($\theta = 0$) from a given initial state, and does so as quickly as possible. The environment is available as a black box, i.e., the dynamics are hidden, but given an initial state and a policy, we may compute the trajectory and obtain its evaluation with respect to the black-box specification.*

We start using the first tree (with predicate $[\theta \geq c_1]$ and leaf nodes corresponding to actions) as a policy for the black-box environment E , and obtain a trace that reaches the goal angle within 60 time steps. Next, the predicate is modified to $[\theta \geq c_2]$ (where $c_2 > c_1$), and the new tree produces a trace that reaches the goal within 50 time steps, which is considered better. For the next tree with predicate $[\theta \geq c_3]$ (where $c_3 > c_2$), the (partially) produced trace is considered inferior: it surely does not reach the goal faster than the best tree (50 steps).

A key component is our novel *trace-based pruning* mechanism that discards a large portion of the search space by runtime analysis. It exploits the decision-tree structure by considering the execution of the tree policy: *even though the environment is black-box*, examining the trace allows us to understand *how* the decision tree is used, and discard trees that are guaranteed to not lead to a better trace. This allows us to reduce the search space *without* sacrificing optimality even though our model and specification are given as a black box. In the previous example, depending on the concrete trace, our trace-based pruning might be able to determine that it is possible to discard the third tree from consideration without missing a better tree only by observing the trace produced by the first and second tree. In practice this can lead to order-of-magnitude reductions in runtime.

We implemented our approach and evaluate it on classical control benchmarks. The experiments demonstrate signifi-

cant reductions obtained with our trace-based pruning, and illustrate that small and optimal decision trees may be constructed within reasonable time. We further analyse the scalability of our algorithm in terms of the number of predicates (granularity) and the size of the tree, both of which have an exponential influence on the runtime. The experiments show the runtime is within practical use.

To summarise, we consider a unique setting and provide a conceptually novel approach to construct optimal small decision-tree policies with respect to black-box systems. While not all environments are controllable by small trees, when the environment *does* admit a small tree policy, our approach provides an effective way to compute an optimal tree only requiring black-box access to the system.

We organise the paper as follows. In the next section, we discuss related work and highlight our unique setting. We outline preliminaries in Section 3, define the problem in Section 4, present our approach in Section 5, experimentally evaluate our approach in Section 6, provide further discussion in Section 7, and conclude in Section 8.

2 Related Work

Our work covers a unique setting: constructing decision trees given a deterministic black-box system whilst providing optimal performance guarantees. There are no directly applicable works that we are aware of in this setting. To illustrate the challenges, we discuss works for synthesising decision-tree policies, albeit not fitting into our setting.

Reinforcement learning. A tree policy can be obtained via reinforcement learning, either by using dedicated tree algorithms (David et al. 2015) or by modifying reinforcement learning to output tree policies (Topin et al. 2021). Alternative approaches allow linear functions on the leaves (Gupta, Talvitie, and Bowling 2015), consider multiple predicates, branches and actions at a time, or fix the structure using expert knowledge (Likmeta et al. 2020) and then employ policy gradient updates (Silva et al. 2020; Paleja et al. 2022). These approaches perform exceptionally well, when an existing reinforcement-learning approach is available that is effective for the given system (Topin et al. 2021), and/or the model (Gupta, Talvitie, and Bowling 2015) of the environment is known. A tree policy may be derived by imitating an expert policy, e.g., a neural network (Bastani, Pu, and Solar-Lezama 2018). In contrast, we require neither model nor expert policy and provide optimality guarantees. In case of sparse rewards, reinforcement learning might struggle, while our framework by design has no such problem.

Learning from tabular data. When the policy is given in tabular form, dedicated tree-learning algorithms for control policies can be employed (Ashok et al. 2020, 2021), which extend classical tree-learning algorithms (Breiman et al. 1984; Quinlan 1996). Recent advancements in optimal tree induction could potentially also be employed (Demirović et al. 2022; van der Linden, de Weerd, and Demirović 2023). However, obtaining the tabular policy requires an explicit model, which is not required in our setting.

Optimal policy synthesis. The problem of constructing a tree policy may be posed as a mixed-integer linear program (Vos and Verwer 2023), after which off-the-shelf

solvers may be used to obtain optimal policies. However, not all environments may be feasible to model with such an approach (e.g., differential equations or trigonometric functions), and in our setting we consider black-box environments, which are not amenable to linear programming.

Verification. There has been recent work to provide guarantees for decision-tree policies (Schilling et al. 2023), possibly for infinitely many traces (we consider finitely many traces). However, the tree policy and the model must be given explicitly, which makes this work orthogonal to ours.

We do note that we consider *deterministic* systems with discrete actions. Some of the methods above are also applicable to stochastic environments and continuous actions, which we consider as future work.

To summarise, while there has been considerable work on decision-tree policies, synthesising such policies when the environment is black-box whilst also providing optimality guarantees is an open challenge.

3 Preliminaries

A *state* $S = (s_1, s_2, \dots, s_d)^\top \in \mathcal{S}$ is a d -dimensional real-valued vector from a bounded state space $\mathcal{S} \subseteq \mathbb{R}^d$, where each *state dimension* s_i belongs to an interval $s_i \in [\ell_i, u_i]$. An action $a \in A$ comes from a finite set $A \subseteq \mathbb{Z}$ of integer-valued actions. An *environment* is a function $E: \mathcal{S} \times A \rightarrow \mathcal{S}$ that takes as input a state S and an action a and computes a trajectory until asked to output the next *observable* state $S' = E(S, a)$. In our setting, the environment is treated as a black box, i.e., we are agnostic to its dynamics.

A *policy* is a function $\pi: \mathcal{S} \rightarrow A$ that chooses an action based on an input state. We write \mathcal{C}_A for the set of all policies over action set A . We are concerned with the special case of a *decision-tree policy*, which is given in the form of a binary tree where each inner node is called a *predicate node* and each leaf node is called an *action node*. Each predicate node is associated with a function $\mathcal{S} \rightarrow \mathbb{B}$, where \mathbb{B} is the Boolean domain. Each action node is associated with one of the available actions $a \in A$. Given a state S and a node of the tree, a decision-tree policy π computes the action $a = \pi(S)$ using the following recursive procedure, starting at the root node. If the current node is an action node, the associated action is returned. Otherwise, the current node is a predicate node. If the state S makes the predicate evaluate to true (false), the procedure continues with the left (right) child node.

Given an environment E , a decision-tree policy π , an initial state S_0 , and a bound $k \in \mathbb{N}$, the *trace* with k steps is the sequence of observed states $\tau = S_0, S_1, \dots, S_k$, obtained by applying the sequence of actions given by π : $S_i = E(S_{i-1}, \pi(S_{i-1}))$, for $i = 1, \dots, k$. We write \mathcal{T} for the set of all traces.

4 Optimal Decision-Tree Policies

We focus on the case of a single initial state S_0 , and generalise the problem to multiple initial states in Section 5.4.

4.1 Discretised Decision-Tree Predicates

The space of all decision trees is infinitely large. By discretising the tree predicates, we obtain a finite (but exponen-

tially large) space. We restrict our attention to (axis-aligned) predicates of the form $[s_i \geq v_0 + m \cdot v_+]$, where s_i is the i -th state dimension, $v_0, v_+ \in \mathbb{R}$ are real-valued constants, and $m \in \mathbb{N}$ is a positive integer. Since our state space is bounded, we obtain a finite number of nonequivalent predicates. We write \mathcal{P} for the set of all (tree) *predicates*.

For instance, given a state space with $s_i \in [0, 3]$ and $v_0 = v_+ = 1$, we consider the predicates $[s_i \geq 1]$, $[s_i \geq 2]$, and $[s_i \geq 3]$. Note that predicate $[s_i \geq 0]$ is excluded since it is a tautology by s_i 's domain, i.e., always evaluates to *true*.

Considering a finer discretisation yields more predicates, which increases the space of considered decision trees, but allows potentially finding better trees. The algorithm's runtime is sensitive to the size of the search space, and hence a practical balance is needed. Experimentally we show that our approach can handle reasonably small increments.

4.2 Specification

Given a deterministic environment E , decision-tree policy π , and initial state S_0 , we consider a specification to determine whether π satisfies the specification. We assume that the specification is given in terms of traces, again in the form of a black-box *specification function* $P: \mathcal{T} \rightarrow \mathbb{B}$. In order to effectively determine whether π satisfies the specification, we restrict the class of specifications we consider. A *bounded-time prefix-closed specification* with a bound $k \in \mathbb{N}$ has the property that every trace τ of length k either satisfies or violates the specification ("bounded-time"), and whenever a trace τ yields either of these verdicts, then any longer trace with the prefix τ yields the same verdict ("prefix-closed"). As a consequence, we are guaranteed to obtain a Boolean verdict from a trace of length at most k . We call the (unique) trace of length k the *witness trace*.

This class of specifications includes common reach-avoid problems where a goal needs to be reached while undesired states need to be avoided. For instance, for the pendulum environment, the specification is to reach the vertical position within a step bound k . A trace satisfies the specification if and only if a prefix of length less than k satisfies the specification function. Conversely, any trace not reaching the goal withing k steps violates the specification.

4.3 Optimality

So far, we were only interested in finding *any* policy that satisfies a given specification. In general, there may exist multiple solutions. We are interested in identifying an *optimal* policy. For that, we assume a *fitness function*, which is a partial order $\succeq: \mathcal{T} \times \mathcal{T}$ to compare two traces. A trace that satisfies the specification always precedes a trace that violates the specification. The fitness function induces another partial order $\succeq: \mathcal{C}_A \times \mathcal{C}_A$ to compare two policies. We say that π_1 is strictly better than π_2 , written $\pi_1 \succ \pi_2$, if one of the following conditions holds: 1) The witness τ_1 has strictly better fitness than the witness τ_2 (i.e., $\tau_1 \succ \tau_2$). 2) Both witnesses have the same fitness, and π_1 is a strictly smaller tree.

We wrap the black-box environment E and the black-box specification function P into a black-box system $B: \mathcal{C}_A \times \mathcal{S} \rightarrow \mathbb{B} \times \mathcal{T}$. This system takes as input a policy π and an initial state S_0 and outputs both the (Boolean) verdict and

the trace τ . We say that a policy π satisfies the specification for initial state S_0 if B yields a positive verdict. We note that B can be implemented from E and P by simply generating the witness τ and querying the specification function.

Problem 1 Given a black-box system B over a set of actions A , an initial state S_0 , a limit on the depth and number of predicate nodes, a discrete set of predicates, and a fitness function \succeq , find a decision-tree policy $\pi \in \mathcal{C}_A$ within the defined size that satisfies the specification optimally with respect to the fitness function \succeq and the witness trace τ produced by the black-box system B .

For instance, for environment E in Fig. 1, the black-box specification (reaching the vertical position) is satisfied for two out of three decision trees. However, we are looking for those trees that satisfy the specification within the smallest number of time steps (in this example, 50).

5 Synthesis of Optimal Decision-Tree Policies

For computing an *optimal* decision-tree policy that solves Problem 1, a naive procedure is to enumerate all possible decision trees and evaluate them. By fixing an upper limit of the number of nodes and considering a discrete set of predicates, this procedure terminates. However, the number of trees is exponential, rendering this procedure infeasible. Our contribution is an efficient instantiation of this procedure.

5.1 Searching In the Space of Decision Trees

Our algorithm to enumerate decision trees is based on backtracking search. We represent the search space using *backtracking variables* b_i , where each variable is associated with a node in the tree. The possible values that can be assigned to a variable depend on the node type which the variable is associated with: predicate nodes may be assigned a predicate from the set of discretised predicates, whereas action nodes may be assigned an action from the set of available actions.

Backtracking variables are considered in a predefined order, i.e., variable b_i goes before variable b_{i+1} . As is standard in backtracking, all combinations for variable b_{i+1} are exhausted before taking the next value for variable b_i .

Assigning all backtracking variables results in a decision-tree policy. Consequently, enumerating all possible assignments to the backtracking variables corresponds to all possible policies in our available space. When enumerating a policy, it is used in combination with the black-box environment to compute the witness trace and evaluate the quality of the policy. Finally, the best policy is returned as the result.

For a tree with a fixed shape and n predicate nodes, $|A|$ number of actions, and $|\mathcal{P}|$ number of discretised predicates, the size of the total search space is $\mathcal{O}(|\mathcal{P}|^n \cdot |A|^{n+1})$. Our contribution reduces this large search space in practice.

5.2 Intuition Behind Trace-based Pruning

The idea of pruning is to limit the exploration by avoiding to explicitly enumerate trees that are guaranteed to be suboptimal, i.e., do not satisfy the desired property with a higher fitness. We define sufficient conditions for pruning.

To provide the intuition behind our approach, consider an environment with only one state dimension s_1 , and the




Figure 2: Three decision-tree policies of fixed structure (top) and corresponding black-box traces (bottom). The decision boundary (blue) splits the state space into two regions (light blue and red). Predicate $s_1 \geq 2$ can be skipped as the trace lies above the decision boundary.

process of enumerating all trees with exactly one predicate node, three possible predicates $[s_1 \geq 1]$, $[s_1 \geq 2]$, $[s_1 \geq 3]$, and having the left and right child nodes *fixed* to actions a_1 and a_2 , respectively (cf. Figure 2). To find the best tree of the given description, we generally need to consider all three trees, with each tree differing only in the root node.

Assume the algorithm starts with predicate $[s_1 \geq 1]$ and computes the first trace depicted in Figure 2. We could consider the two remaining trees with predicates $[s_1 \geq 2]$ and $[s_1 \geq 3]$. However, from the first trace we see that the second tree with predicate $[s_1 \geq 2]$ would result in the exact same trace. Indeed, there is no state in the trace where the policy would decide differently regardless of whether the predicate is $[s_1 \geq 1]$ or $[s_1 \geq 2]$, so the trace would not change. As a result, we do not have to evaluate the tree with predicate $[s_1 \geq 2]$, and can directly go to the last tree.

5.3 Trace-Based Pruning

The intuition discussed above can be generalised to prune a potentially exponential number of trees that do not result in a different trace, which significantly speeds up the search process. In the following, we discuss incorporating a general version within a backtracking algorithm with more than one predicate node. We will consider predicates in increasing order, i.e., $[s_j \geq v_1]$ before $[s_j \geq v_2]$ if $v_1 < v_2$.

Given a backtracking variable b_i associated with a predicate node, the algorithm needs to select the next predicate to assign to the variable. A naive approach would be to simply select the next larger threshold, e.g., from the example in Section 5.2, assign predicate $[s_1 \geq 2]$ after considering $[s_1 \geq 1]$. However, we can leverage information about previous traces to avoid explicitly considering predicates that are guaranteed to not result in a trace that has not been previously observed.

After assigning a predicate $[s_j \geq v]$ to a backtracking variable b_i , the idea is to track the values of state dimension s_j that have been observed during environment runs such that the predicate evaluated to *true*. In particular, we are interested in the smallest such value, which we refer to as the distance value $d_i \in \mathbb{R}$. Note that a tree policy is only run after having all backtracking variables assigned.

The key idea is that, when selecting the next predicate for

backtracking variable b_i , its threshold value (v in the example above) should *exceed* the distance value d_i . Otherwise, the trace would be identical.

To reiterate, for each backtracking variable b_i associated with a predicate node with current predicate $[s_j \geq v]$, we store a value $d_i \in \mathbb{R}$, which tracks the minimum value of state dimension s_j for which the predicate was evaluated to *true* amongst all traces that were considered since the predicate $[s_j \geq v]$ has been assigned. Note that selecting a new threshold value for the predicate that is *smaller or equal* than d_i is guaranteed to result in traces already observed. Note that it is not necessary to track the values where the predicate evaluates to *false*, since the predicates are explored in increasing order of the thresholds and as such the future predicates would also evaluate to false on those values.

Initially, the distance value d_i is set to *undefined* each time a new predicate $[s_j \geq v]$ is assigned as part of the search. The first time the node observes a state where its predicate is satisfied in a trace, the distance value d_i is set to the corresponding value of state dimension s_j . Each subsequent time the predicate is satisfied, d_i is updated to the smallest value for which the predicate still evaluates to *true*.

After considering predicate $[s_j \geq v]$ for node i , our algorithm does not consider the next predicate, but instead uses the predicate $[s_j \geq v']$ where v' is the smallest available value such that $v' > d_i$. If the distance value d_i is undefined, *all* predicates may be discarded for that backtracking variable for the currently considered state dimension s_j .

For the example in Section 5.2, the distance value d_1 is initially undefined, and upon completing the first trace, it is set to $d_1 = 2.3$. When selecting the next predicate, $[s_1 \geq 2]$ is discarded since its threshold 2 does not exceed distance value $d_1 = 2.3$; so the next selected predicate is $[s_1 \geq 3]$.

The above idea is applied to every backtracking variable associated with a predicate.

Our pruning strategy is computationally inexpensive: it amounts to tracking a single value for each backtracking variable, and updating this value as the tree is queried during trace computation. The algorithm retains completeness, as it is guaranteed to not discard optimal trees. Our trace-based pruning is the key component in the practical efficiency.

Additional Techniques: Trees explored in increasing size. The algorithm partitions the search space in terms of tree shapes, which are ordered by size. For example, after considering trees with exactly one predicate node, the algorithm considers trees which have a root node with one left predicate child, then trees which have a root node with one right predicate child, then complete trees with three predicate nodes, and so on until the size budget is reached.

Early stopping due to the objective. During the search, the best tree found so far is tracked. When evaluating a new candidate tree, its evaluation is preemptively stopped when it is determined that the trace cannot be extended to a trace that is better than the one obtained from the best tree so far.

For example, consider a setting where the policy should reach a goal state as quickly as possible. If the best policy so far reaches the goal state in k trace steps, and the partial trace associated with the current candidate policy has not

Algorithm 1: Construction of an optimal decision-tree policy

Require: An initial state S_0 , black-box system B , predicates \mathcal{P} , fitness \succeq , tree-size bound

Ensure: Policy π that produces the trace with the best fitness function

```

1: bvar  $\leftarrow$  NextBacktrackingVariable()
2: val  $\leftarrow$  NextValue(bvar) // use trace-based pruning if
   predicate node
3: Assign(bvar, val)
4: if assignment failed, all values exhausted then
5:   Backtrack()
6:   if not possible to backtrack then
7:     return  $\pi_{global}$ 
8:   end if
9: end if
10: if all backtracking variables assigned then
11:    $\pi_{local} \leftarrow$  ConstructTreePolicy(bvars)
12:    $\tau_{local} \leftarrow B(\pi, S_0)$ 
13:   if  $\tau_{local} \succ \tau_{global}$  then
14:      $\tau_{global} \leftarrow \tau_{local}$ 
15:      $\pi_{global} \leftarrow \pi_{local}$ 
16:   end if
17:   for bvar in bvars do
18:     n  $\leftarrow$   $\pi_{local}$ .nodes[bvar.index]
19:     bvar.distance  $\leftarrow$  min(n.distance, bvar.distance)
20:   end for
21: end if
22: Goto Line 1.

```

reached the goal state in $k - 1$ steps, we may safely discard the candidate policy from further consideration, since it cannot result in a better trace. Note that, since we explore trees in increasing size, this results in the algorithm computing the smallest tree with the optimal performance across the considered maximum tree-size budget.

Early stopping has two advantages: 1) it saves computational time, and 2) it results in traces of shorter length, which allows for more aggressive trace-based pruning owing to fewer distance updates being made.

Symmetries. Trees that have identical left and right subtrees are discarded from consideration. In these cases, the root node of such a tree is redundant and its predicate has no influence on the trace. A smaller tree, consisting of the subtree, would result in the same trace, and since the algorithm explores trees in increasing size, it is safe to discard the larger symmetric tree without further consideration. The main use of this technique is to discard trees that contain predicate nodes with identical left and right action nodes, and otherwise plays a minor role.

Summary. Algorithm 1 provides a high-level view on using backtracking variables. If available, the next unassigned backtracking variable is selected, or the last assigned variable otherwise. The next value is selected either as the next action for variables representing action nodes, and otherwise trace-based pruning is used to determine the threshold. Once a predicate has been exhausted on one state dimension, predicates for the next state dimension are selected.

Once all backtracking variables are assigned, the algorithm constructs a decision-tree policy, and uses the black-box system to produce the trace τ . If τ is better than the globally best trace (initially *null*), that trace is updated to τ . The distance values of the nodes of the policy are used to update the distance values of the backtracking variables. After all policies have been (implicitly) considered, the algorithm returns the best policy (Line 7).

5.4 Extensions

Multiple initial states. The previous discussion was based on constructing a tree policy from a single initial state. However, we may be interested in finding a single tree policy that works well across multiple initial states. The algorithm remains similar, with impact on two components: 1) the objective function, and 2) trace-based pruning.

When evaluating a tree with respect to multiple initial states, we generalise the fitness function. For example, if the goal is to minimise the trace length, then the generalisation aims to minimise the maximum trace length. This influences early stopping: the initial states are evaluated with respect to the tree policy one at a time, and as soon as a trace is encountered that is considered violating, the evaluation stops, i.e., the remaining initial states are not considered further.

The above idea interacts with trace-based pruning. In case the tree evaluation is stopped early, meaning the tree is deemed not better than the best tree found so far, only the last trace is used to update the distance values. The intuition is that, if we wish to find a better tree, it must lead to a trace different from the last trace, and we can ignore the distance-value updates of the other traces.

As a result, due to the interaction with pruning, finding an optimal tree with respect to multiple initial states may lead to sub-linear runtime with respect to the number of states.

Maximisation. Rather than satisfying the desired property in the least number of steps, we may be interested in maximising the number of steps. For example, the goal may be to balance a pole for as long as possible. The algorithm stays largely the same, with the only analogous changes needed in the evaluation of the tree. For maximisation it is important to specify an upper bound on the trace length; otherwise, the algorithm may potentially run infinitely long.

6 Experimental Study

We aim to illustrate the effectiveness of our approach with a proof-of-concept implementation. We show that our trace-based pruning approach is a key factor in making the approach feasible. Furthermore, we consider scalability from two perspectives: the granularity of the predicates, and the number of predicate nodes in the tree. While both are expected to have an exponential impact on the runtime, we observe that the runtime is still within practical use.

We consider three classical control problems: CartPole, MountainCar, and pendulum. The environment behaviour is defined as in Gymnasium¹, with the adjustment that we maximise (CartPole - at most 10k steps) or minimise (MountainCar and pendulum) the trace length rather than use their

¹<https://github.com/Farama-Foundation/Gymnasium>

Environment	$ S_0 $	Runtime (seconds)		No. of trees (Not) Prune
		Not Prune	Prune	
CartPole	1	3.13±1.79	0.22±0.20	(2.2m) 49k
MountainCar	1	4.72±0.49	0.21±0.04	(518k) 20k
Pendulum	1	13.75±11.07	3.19±3.03	(845k) 126k
CartPole	100	7.46±1.28	1.28±0.35	(2.2m) 88k
MountainCar	100	6.55±0.68	0.40±0.08	(518k) 24k
Pendulum	100	131.4±9.41	23.93±2.46	(845k) 205k

Table 1: Runtime and variance with and without trace-based pruning. Pruning substantially reduces runtime. Multiple initial states demonstrate only a sub-linear runtime increase.

reward functions. The control actions are limited to two choices, e.g., apply maximum force in one or the other direction. Our implementation only queries the environment in a black-box fashion. We generated the initial states randomly within a specified range; see the Appendix of the extended version for details about the parameters, environment description, and a sample of trees produced by our approach.

To reiterate, as discussed in Section 2, we work with a unique setting where we 1) synthesise *decision-tree* policies, 2) only require *black-box* access to the *deterministic* system, and 3) provide *guarantees* on performance under the tree definition, e.g., the policy that minimises the time taken, or prove that no such policy exists. Every work we are aware of violates at least one of these points. Consequently, while direct comparisons with other works may be done by relaxing the requirements of our setting, this brings considerable caveats that result in comparisons that we argue are not meaningful with respect to our contribution. For these reasons, we focus on demonstrating the feasibility of our approach and its scalability.

We implemented our approach ‘Broccoli’ in Rust 1.77.0. The experiments were run on consumer-grade hardware (Xeon(R) W-10855M @ 2.8 GHz). The code is public.

6.1 Experiment #1: Trace-Based Pruning

We run experiments with and without our trace-based pruning (Table 1), both for one and 100 initial states, for trees of depth two. Initial states are generated randomly, and the results are averaged over ten runs.

Our trace-based pruning technique is clearly effective in pruning the search space. For CartPole and MountainCar, there is a one- to two-orders-of-magnitude difference, whereas for pendulum, it leads to a 4x reduction.

The runtime is roughly proportional to the number of trees explicitly considered, as expected, and is consistent based on the standard deviation across different initial states. Interestingly, we observe only a sub-linear increase in runtime with more initial states, and the total number of trees considered is roughly similar regardless of the number of initial states.

6.2 Experiment #2: Granularity of Predicates

The granularity of predicates impacts the runtime: the finer the discretisation, the larger the search space. Each state dimension in the environment has a predefined range of values

Environment	Runtime (seconds) and variance				No. of trees			
	X=5	X=10	X=15	X=20	X=5	X=10	X=15	X=20
CartPole	102±128	1420±2696	10.2±16.33	287±668	25m	334m	1m	46m
MountainCar	6.26±8.55	14.0±3.86	95.49±29.2	446±97	30k	1.3m	9.5m	45m
Pendulum	6.58±7.53	30.75±61.32	113±151	366±501	18k	1.2m	15m	51m

Table 2: Runtime increases as we increase the number of predicates per state dimension (X), with the exception of CartPole, where finer discretisation sometimes finds a perfect tree faster. Variance remains similar as in Table 1 proportional to runtime.

Environment	Runtime (seconds) and variance				No. of trees			
	N=3	N=4	N=5	N=6	N=3	N=4	N=5	N=6
CartPole	0.1±0.01	1.16±0.11	11.85±7.26	39.92±38.45	27k	294k	2.9m	9.4m
MountainCar	0.12±0.02	0.74±0.16	5.37±1.16	14.1±3.39	9k	62k	484k	1.2m
Pendulum	0.3±0.21	0.56±0.23	2.17±0.41	5.36±0.83	7k	44k	309k	826k

Table 3: Runtime and performance (not depicted) increase as we increase the number of nodes (N).

that it may take. We divide this interval into 5, 10, 15, and 20 values, and use these values as the predicate thresholds.

Table 2 summarises the results for trees of depth three across randomly generated initial states (averaged over ten runs). The general trend is that increasing the number of predicates indeed increases the runtime.

However, in the case of CartPole, we observe an opposite effect: finer predicates *decrease* the runtime. This is because the finer discretisation resulted in the algorithm finding a tree faster that balances the pole for 10,000 steps, which is the maximum number of steps so the search can terminate.

The results indicate that the discretisation needs to be chosen carefully to balance runtime and quality of the final tree.

6.3 Experiment #3: Number of Predicate Nodes

The number of predicate nodes influences the search space size. We study the runtime for trees of depth three and varying the maximum number of nodes from three to six.

The results are summarised in Table 3. Note that the search space of trees with at most k predicate nodes *includes* trees with less than k predicate nodes, meaning the search space is strictly larger as we increase the number of nodes. Consequently, there is a sharp increase in the runtime, which is expected due to the exponential factor.

7 Further Discussion and Limitations

Our approach is effective at computing small and optimal decision-tree policies given a discretisation of the predicates and a set of initial states. There is an exponential runtime dependency with respect to the size of the tree and the discretisation of the state space. It may be infeasible with our approach to construct large tree policies or deal with high-dimensional environments. It is also possible that not all environments may be controlled by small decision trees.

However, when it *is* applicable, we believe small trees are valuable for interpretability reasons, and our approach provides the means to easily obtain such trees. The exponential runtime factor in our approach is inherent to every approach that aims to provide guarantees.

Our approach is exceptionally flexible as it only requires black-box access to the system. This entails that the black box may be arbitrarily complex, as long as it can still be practically computed. Furthermore our algorithm provides performance guarantees, despite working with black boxes.

The optimality is important since it guarantees that we obtain the best performing tree given our definition, which may be relevant for some applications. It also allows us to conclude in cases when no such tree exists, and in general understand the limits of decision trees as control policies.

Given that our approach is a conceptually novel way to synthesise decision-tree policies in a unique setting, it opens many avenues for future work.

Parallelisation is promising as the search space can be naturally partitioned, and further *heuristic* pruning may lead to a principled trade-off between runtime and guarantees. Extending the approach to stochastic environments is another interesting direction. In our work, continuous actions ought to be discretised in a preferably smaller number of actions. Synthesising optimal trees for continuous actions remains an open challenge for decision-tree policies in general.

8 Conclusion

We presented a novel search-based method for computing an optimal decision-tree policy given a black-box deterministic system, a set of initial states, and a discretisation of the tree predicates. To the best of our knowledge, our approach is the first to consider such a setting. The key component is our trace-based pruning technique, which discards large portions of the search space at runtime. We illustrated the practicality of the approach on classical control benchmarks. When the environment is controllable by a small tree, our approach provides a way to obtain a small and optimal tree despite only requiring black-box access to the system.

Acknowledgements

This research was funded by the NWO Veni grant Explainable Monitoring (222.119) and by the Independent Research

Fund Denmark under reference number 10.46540/3120-00041B. This work was done in part while Anna Lukina was visiting the Simons Institute for the Theory of Computing.

References

- Anderson, C. W. 1989. Learning to Control an Inverted Pendulum Using Neural Networks. *IEEE Control Systems Magazine*, 9(3): 31–37.
- Ashok, P.; Jackermeier, M.; Jagtap, P.; Kretínský, J.; Weininger, M.; and Zamani, M. 2020. dtControl: Decision Tree Learning Algorithms for Controller Representation. In *HSCC*, 17:1–17:7. ACM.
- Ashok, P.; Jackermeier, M.; Kretínský, J.; Weinhuber, C.; Weininger, M.; and Yadav, M. 2021. dtControl 2.0: Explainable Strategy Representation via Decision Tree Learning Steered by Experts. In *TACAS*, volume 12652 of *LNCS*, 326–345. Springer.
- Bastani, O.; Pu, Y.; and Solar-Lezama, A. 2018. Verifiable Reinforcement Learning via Policy Extraction. In *NeurIPS*, 2499–2509.
- Belta, C.; and Sadraddini, S. 2019. Formal Methods for Control Synthesis: An Optimization Perspective. *Annual Review of Control, Robotics, and Autonomous Systems*, 2: 115–140.
- Breiman, L.; Friedman, J. H.; Olshen, R. A.; and Stone, C. J. 1984. *Classification and Regression Trees*. Wadsworth. ISBN 0-534-98053-8.
- David, A.; Jensen, P. G.; Larsen, K. G.; Mikucionis, M.; and Taankvist, J. H. 2015. Uppaal Stratego. In *TACAS*, volume 9035 of *LNCS*, 206–211. Springer.
- Demirović, E.; Lukina, A.; Hebrard, E.; Chan, J.; Bailey, J.; Leckie, C.; Ramamohanarao, K.; and Stuckey, P. J. 2022. MurTree: Optimal Decision Trees via Dynamic Programming and Search. *Journal of Machine Learning Research*, 23: 26:1–26:47.
- Du, M.; Liu, N.; and Hu, X. 2020. Techniques for Interpretable Machine Learning. *Communications ACM*, 63(1): 68–77.
- Glanois, C.; Weng, P.; Zimmer, M.; Li, D.; Yang, T.; Hao, J.; and Liu, W. 2024. A Survey on Interpretable Reinforcement Learning. *Machine Learning*, 1–44.
- Gupta, U. D.; Talvitie, E.; and Bowling, M. 2015. Policy Tree: Adaptive Representation for Policy Gradient. In *AAAI*, 2547–2553.
- Likmeta, A.; Metelli, A. M.; Tirinzoni, A.; Giol, R.; Restelli, M.; and Romano, D. 2020. Combining Reinforcement Learning with Rule-Based Controllers for Transparent and General Decision-Making in Autonomous Driving. *Robotics and Autonomous Systems*, 131: 103568.
- Paleja, R.; Niu, Y.; Silva, A.; Ritchie, C.; Choi, S.; and Gombolay, M. 2022. Learning Interpretable, High-Performing Policies for Autonomous Driving. In *Robotics Science and Systems*.
- Quinlan, J. R. 1996. Learning Decision Tree Classifiers. *ACM Computing Surveys (CSUR)*, 28(1): 71–72.
- Rudin, C. 2019. Stop Explaining Black Box Machine Learning Models for High Stakes Decisions and Use Interpretable Models Instead. *Nature Machine Intelligence*, 1(5): 206–215.
- Schilling, C.; Lukina, A.; Demirović, E.; and Larsen, K. G. 2023. Safety Verification of Decision-Tree Policies in Continuous Time. In *NeurIPS*, volume 36, 14750–14769.
- Silva, A.; Gombolay, M.; Killian, T.; Jimenez, I.; and Son, S.-H. 2020. Optimization Methods for Interpretable Differentiable Decision Trees Applied to Reinforcement Learning. In *AISTATS*, 1855–1865.
- Topin, N.; Milani, S.; Fang, F.; and Veloso, M. 2021. Iterative Bounding MDPs: Learning Interpretable Policies via Non-interpretable Methods. In *AAAI*, 9923–9931.
- van der Linden, J. G. M.; de Weerd, M.; and Demirovic, E. 2023. Necessary and Sufficient Conditions for Optimal Decision Trees using Dynamic Programming. In *NeurIPS*.
- Vos, D.; and Verwer, S. 2023. Optimal Decision Tree Policies for Markov Decision Processes. In *IJCAI*, 5457–5465.