dtControl: Decision Tree Learning Algorithms for Controller Representation

Pranav Ashok
Mathias Jackermeier
Pushpak Jagtap
Jan Křetínský
Maximilian Weininger
Technical University of Munich
Munich, Germany

Majid Zamani
University of Colorado Boulder
Boulder, USA
Ludwig Maximilian University of Munich
Munich, Germany

ABSTRACT

Decision tree learning is a popular classification technique most commonly used in machine learning applications. Recent work has shown that decision trees can be used to represent provably-correct controllers concisely. Compared to representations using lookup tables or binary decision diagrams, decision trees are smaller and more explainable. We present dtControl, an easily extensible tool for representing memoryless controllers as decision trees. We give a comprehensive evaluation of various decision tree learning algorithms applied to 10 case studies arising out of correct-by-construction controller synthesis. These algorithms include two new techniques, one for using arbitrary linear binary classifiers in the decision tree learning, and one novel approach for determinizing controllers during the decision tree construction. In particular the latter turns out to be extremely efficient, yielding decision trees with a single-digit number of decision nodes on 5 of the case studies.

CCS CONCEPTS

• **Computer systems organization** → **Embedded and cyber-physical systems**; • **Computing methodologies** → **Classification and regression trees**; **Control methods**.

KEYWORDS

Controller representation, Decision tree, Machine learning, Symbolic control, Non-uniform quantizer, Explainability, Invariance entropy

ACM Reference Format:

Pranav Ashok, Mathias Jackermeier, Pushpak Jagtap, Jan Křetínský, Maximilian Weininger, and Majid Zamani. 2020. dtControl: Decision Tree Learning Algorithms for Controller Representation. In *23rd ACM International Conference on Hybrid Systems: Computation and Control (HSCC '20)*, April 22–24, 2020, Sydney, NSW, Australia. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3365365.3382220>

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
HSCC '20, April 22–24, 2020, Sydney, NSW, Australia
© 2020 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-7018-9/20/04.
<https://doi.org/10.1145/3365365.3382220>

1 INTRODUCTION

Formal synthesis of controllers enforcing complex specifications on cyber-physical systems has gained significant attention in the last few years. This is mainly due to the need for obtaining formally verified control strategies rendering some complex tasks; these are usually represented using temporal logic specifications or (in)finite strings over automata. There are several techniques and tools available that provide automated, correct-by-construction, controller synthesis for cyber-physical systems by utilizing symbolic models (a.k.a. finite abstractions) [5, 36], in which the uncountable continuous states and inputs are aggregated to finite symbolic states and inputs via quantization (a.k.a. discretization). The so-called symbolic controllers are then computed by utilizing algorithmic machinery from computer science and then mapped back for use in the original systems. The state-of-the-art tools to synthesize such controllers are, e.g., SCOTS [32], pFaces [18], QUEST [15], Pessoa [16], CoSyMA [24], or UPPAAL STRATEGO [12]. These tools give a huge list of state-action pairs (a.k.a. lookup tables) representing controllers.

Storing these symbolic controllers in the memory is a major problem because they usually need to run on embedded devices with limited memory. However, if we do not store the controllers as lookup tables, but take advantage of decision trees (DT) [23], which exploit their hidden structure to represent them in a more compact way, we can mitigate this problem. As shown in [3], DTs can be orders of magnitude smaller than lookup tables. Such a concise representation opens the door for better readability, understandability, and explainability of the controllers, while reducing memory requirements and preserving correctness guarantees. Moreover, human-understandable controllers may also provide insight into the models themselves, thus aiding their validation, as we illustrate in the example below.

Our setting is inherently different from the usual use of DT in machine learning; there, in order to generalize well, DTs typically do not fit the training data exactly; in contrast, in this work, DTs have to exactly represent the given controllers in order to preserve their correctness guarantee. Therefore, our requirements on DTs differ: beside the size and the explainability, it is also the *perfect fitting*. Consequently, it is necessary to thoroughly re-evaluate current DT-learning algorithms and possibly also modify them.

A basic technique used to represent controllers more concisely is to *determinize* them, i.e. to make them not (maximally) permissive but only retain a single action for each state. To this end, one can use, for instance, the action with the minimum norm from

Figure 1: Decision tree for the temperature controller

a reference input, when least energy consuming controllers are preferred [22], or the previously applied action (if possible), when lazy controllers are preferred [16, 24]. Such a size reduction by determinization can be applied as *pre-processing before* learning the DT representation of the controller, typically yielding also a smaller DT. Alternatively, one can apply other kinds of reduction by determinization as *post-processing after* constructing the DT. For instance, in “safe pruning” of [3], the DT constructed for the maximally permissive controller is modified as follows. The leaves of the tree are merged in a bottom-up fashion, thereby reducing the size and partially determinizing it. In contrast, here we introduce a novel approach for determinizing the controllers *during* the construction of the DT, with advantages to both pre-processing and post-processing methods. Firstly, since the choice of the action for each state greatly affects the size and structure of the DT, it is advantageous to guide the choice by the concrete, already built part of the DT, compared to a-priori choices made by pre-processing approaches. Secondly, while the post-processing approaches have to construct a large tree first, our new technique constructs an already reduced tree, avoiding the intermediate large one, thus making it more scalable.

Motivating Example. Consider a temperature control system running in a building with 10 rooms with the heater installed only in 2 rooms as described in [15]. The permissive controller maintaining the temperatures of all the rooms within a certain range obtained using SCOTS is a lookup table with 52,488 state-action pairs. By naively determinizing, we get a lookup table with 26,244 symbolic states (i.e. domain of the controller) and their respective actions. The standard DT-learning, e.g. [9], applied to these two lookup tables yields DT with 8,648 and with 2,703 decision nodes, respectively. While this is an improvement, it is far from being explainable. With the help of our novel determinization strategy presented in Section 4.2, we are able to obtain the decision tree with only 3 (!) decision nodes, see Figure 1. Apart from obtaining a compact and easily implementable controller representation while preserving correctness guarantees, the result is so small that it is immediately explainable and, moreover, allows us to improve on the implementation: one can readily see that we only need to install temperature sensors in two rooms instead of all 10 rooms, which will help users to reduce the system deployment cost as well as the required bandwidth to transfer the state information to the controller. Only 4 symbols (leaves of the tree) need to be transferred to realize the controller.

We also obtain a controller with very few nodes for the cruise-control model of [21]. From such a clear representation one immediately notices that the controller makes the car decelerate when the car in front of it is far away. This counter-intuitive behaviour has

thus revealed a bug in the model, which did not actually describe the intended behaviour of the system.

The contribution of this paper can be summarized as follows:

- We present dtControl, an open-source tool to convert formally verified controllers to decision trees preserving their correctness guarantees. dtControl has a simple input format and already supports automated conversion for controllers generated by two state-of-the-art tools – UPPAAL STRATEGO [12] and SCOTS [32]. It supports several output formats, most importantly the graphical output as DOT files, useful for further analysis and visual presentation, and the C source code, useful for closed-loop simulation or for loading onto embedded devices.
- We introduce a new technique for using arbitrary binary classifiers in the DTs and a novel approach for determinizing controllers during the DT learning. Our approach is tuned towards obtaining extremely small, explainable DTs. In 5 out of 8 case studies where it is applicable (the original controllers are non-deterministic), it produces trees with single-digit numbers of decision nodes.
- We present a comprehensive evaluation of 8 DT-learning algorithms on 10 case studies.

Related Work. DTs [23, Chapter 3] are a well-known class of data structures, particularly known for their interpretability, used mostly by machine learning practitioners in classification or regression tasks. Our work is based on well-known algorithms for decision tree learning, namely CART [9], C4.5 [30] and OC1 [25].

There has been previous work on combining decision trees with classifiers, namely Perceptrons [37], Logistic Regression models [19], piece-wise functions [27] or Support-Vector Machines [1, 11]. We generalize those approaches by allowing for arbitrary binary classifiers to be used in our trees. Additionally, those methods are either restricted to only use two labels, which is not applicable for controllers with more than two possible actions, or they only allow linear classifiers in leaf nodes [1, 27]. In contrast, our approach is applicable with an arbitrary number of actions and also leverages the power of linear classifiers in inner nodes.

An alternative to DTs are binary decision diagrams (BDD) [10]. As seen in [3, 7, 8], BDDs have several disadvantages: firstly, they do not retain the inherent flavour of decisions of strategies as maps from states to actions due to their bit-level representation and, hence, are hardly explainable. Secondly, they are notoriously hard to minimize [8], also because finding the best variable ordering is NP-complete [10]. BDDs only allow binary classification, so the actions have to be joined with the state space to represent a controller. The recent result in [38] discusses various heuristic-based determinization algorithms for BDDs representing controllers; however, they still suffer from those disadvantages we mentioned for BDDs. Algebraic decision diagrams (ADD) [4] are an extension of BDDs that allow to have more than two labels, i.e. associate every action to a leaf node. However, they still suffer from the same drawbacks as BDDs. In [13] ADDs are used for controller representation; however, no concrete algorithm is provided.

The formal methods community has made use of decision trees to represent controllers and counterexamples arising out of model checking Markov decision processes, stochastic games and LTL synthesis [1, 3, 7, 8]. DTs have also been used to represent learnt

policies from reinforcement learning [29]. However, in contrast to our paper, [29] does not preserve safety guarantees, only considers axis-aligned splits and does not consider non-determinism. [17] suggests the possibility of using regression trees for representing policies, whereas we consider classification trees.

2 TOOL

dtControl is an easy-to-use open-source tool for post-processing memoryless symbolic controllers into various compact and more interpretable representations. We report the input and output formats as well as the algorithms that are currently supported. Note that the tool can easily be extended with new formats and algorithms. dtControl is distributed as an easy-to-install pip package¹ along with a user and developer manual².

Dependencies. dtControl works with Python version 3.6.7 or higher. The core of the tool which runs the learning algorithms require numpy, pandas and scikit-learn [28]. Optionally, dtControl may also require the C-based oblique decision tree tool OC1 [25].

Input formats. dtControl currently accepts controllers in three formats: (i) a raw comma-separated values (CSV) format with each row consisting of a vector of state variables concatenated with a vector of input variables; (ii) a sparse matrix format used by SCOTS; and (iii) the raw strategy produced by UPPAAL STRATEGO. More details about the various formats are described in the user manual.

Algorithms. dtControl offers a range of parameters to adjust the DT learning algorithm, which are described in Section 4.

Output formats. dtControl outputs the decision tree in the DOT graph representation language (for visual presentation of the tree), as well as C code that can be directly used for implementation; see [2, Appendix A] for the DOT and C output that dtControl produces for the DT in Figure 1. Additionally, dtControl reports statistics for every constructed tree, namely size, the minimum number of bits required to represent symbols in obtained controller, and the construction time.

3 PRELIMINARIES - DECISION TREE LEARNING

A decision tree (DT) over the domain X with the set of labels $\mathcal{U}$ is a tuple (T, λ, ρ) , where T is a finite full binary tree (every node has exactly 0 or 2 children), λ assigns to every leaf node (node with 0 children) a label $u \in \mathcal{U}$ and ρ assigns to every inner node (node with 2 children, also called decision node) of the tree a predicate, which is a boolean function $X \mapsto \{0, 1\}$.

The semantics of a DT is as follows: given a state $\vec{x}$, there is a unique *decision path* through the tree T starting from the root node (the only node with no parent) to a leaf node ℓ . This means that the label for state $\vec{x}$ is $\lambda(\ell)$. The decision path is defined by starting at the root node, and then for each decision node n evaluating the predicate on the state, i.e. computing $\rho(n)(\vec{x})$, and picking the left child if the predicate is true and the right child otherwise.

For example, consider the DT in Figure 1: T has 7 nodes, 3 of which are decision nodes (including the root node) and 4 of which

are leaf nodes. A state of the system is a vector of 10 temperatures, e.g. $\vec{x} = (20.1, 20.2, 20.3, 20.4, 20.5, 20.6, 20.7, 20.8, 20.9, 21.0)$. To find the decision for this state, we first evaluate the predicate in the root node. Since the temperature in the second room is smaller than 20.625, the predicate is true and we go to the left child. We evaluate the next predicate in the same fashion and arrive at the leaf node labelled $(1, 1)$, which gives us a safe control input, in this case to turn on both heaters.

All DT learning algorithms implemented in dtControl follow the same underlying structure: given a finite set $C \subseteq X \times \mathcal{U}$ of feature-label pairs, it returns a DT that represents C precisely; this means that for every $(\vec{x}, u) \in C$, the leaf node of the decision path for $\vec{x}$ has the label u . In the setting of this paper, C is a controller, features are states and labels are actions³.

To learn the DT, the algorithm tries to minimize the entropy of C , denoted $\text{entr}(C)$, by splitting it according to a predicate. Formally, for some $C \subseteq \{(\vec{x}, u) \mid \vec{x} \in X, u \in \mathcal{U}\}$,

$$\text{entr}(C) := - \sum_{u \in \mathcal{U}} p_u \log(p_u),$$

where $p_u := \frac{|\{(\vec{x}, u) \in C\}|}{|C|}$ is the empirical probability of label u being in C ; notation $|\cdot|$ denotes the cardinality of a set. The underlying algorithm works recursively as follows:

- **Base case:** If $\text{entr}(C) = 0$, i.e. all pairs $(\vec{x}, u) \in C$ have the same label u , then return the following DT: the tree T has only a single node r , with $\lambda(r) = u$, and ρ has no domain in this case, as there are no decision nodes.
- **Recursive case:** If $\text{entr}(C) \neq 0$, C needs to be split; for that, we use some predicate $P \in \text{PREDS}$ which splits C , where the set PREDS to be picked here is a parameter of the algorithm that is discussed in Section 4.1. We pick the predicate that minimizes the entropy after the split, i.e.,

$$\arg \min_{P \in \text{PREDS}} \text{entr}(\{(\vec{x}, u) \in C \mid P(\vec{x})\}) + \text{entr}(\{(\vec{x}, u) \in C \mid \neg P(\vec{x})\}).$$

Intuitively, the best predicate is the one which is able to split C into two parts which are as homogeneous as possible. Given the best predicate, we recursively call the algorithm on the sub-sets resulting from the split, getting two DTs (T_t, λ_t, ρ_t) and (T_f, λ_f, ρ_f) ; the indices t and f indicate whether the predicate was true or false, respectively. Then we return the following DT: the tree T has the root node r , with the left child being the root of T_t and the right child the root of T_f . λ uses λ_t for leaves of the left sub-tree and λ_f for the right sub-tree. ρ is defined similarly on the inner nodes of the left and right sub-trees, with the addition that $\rho(r) = P$, i.e. the predicate of the root of T is the predicate we used for the split.

The symbolic controllers designed by SCOTS and UPPAAL STRATEGO are generated by correct-by-construction synthesis procedures. In order to use these controllers for original systems (i.e. with infinite continuous states and inputs), we need to refine the controllers. For more details on refinement procedures, we kindly refer the interested reader to [20, 31, 36].

dtControl preserves the correctness guarantees by representing the symbolic controllers precisely, i.e. iterating until the entropy

¹pip is a standard package-management system used to install and manage software packages written in Python. See <https://pypi.org/project/dtcontrol/>.

²Available at <https://dtcontrol.readthedocs.io/en/latest/>

³We use the term actions instead of control inputs, to avoid confusion because of the fact that the control inputs are the outputs of a DT.

in all leaf nodes is 0. In the case of determinization, dtControl represents one of the deterministic sub-controllers precisely, which is chosen on-the-fly during the construction.

4 METHODS

There are two parameters of dtControl: the set of predicates to consider (PREDS) and the way in which non-determinism is handled. For each of these, dtControl implements existing ideas and introduces new ones. Here, we only report the high-level ideas; for a more detailed description, refer to the user or developer manual.

4.1 Predicates

4.1.1 Existing idea: Axis-aligned splits. In the standard algorithms, e.g. [9, 30], only *axis-aligned splits* are considered; i.e. predicates that can only have the form $x_i \sim b$, where x_i is one of the state variables, $b \in \mathbb{R}$, and $\sim \in \{\leq, \geq\}$. In our setting, the set of possible predicates is greatly restricted due to discretization (quantization). The number of splits to be evaluated for each variable x_i is equal to the number of discrete values of x_i .

4.1.2 Existing idea: Oblique splits. Beside the standard axis-aligned splits, dtControl also supports predicates of the form $\vec{w}^T \vec{x} \leq b$, where $\vec{w}, \vec{x} \in \mathbb{R}^n, b \in \mathbb{R}$. These *oblique predicates* [25] incorporate information from multiple state variables in a single split and thus have the potential to greatly simplify the induced decision tree [1]. However, due to combinatorial explosion, it is too costly to simply enumerate all possible oblique predicates even in the discretized space, due to which different heuristics are employed [25]. In this regard, dtControl supports the usage of predicates obtained using (an adapted version of) the OC1 algorithm [25].

4.1.3 New technique: Using binary machine-learned classifiers. It is possible to find non-axis-aligned predicates splitting the controller by using classification techniques from machine learning. As our main goal is for the resulting tree to be explainable, we want to avoid complex predicates, and thus we restrict the classifiers we consider in two ways: (i) we only consider linear classifiers, and (ii) we restrict to binary classifiers, so that the resulting tree is binary.

We use these binary linear classifiers in a way that is similar to the classical one-vs-the-rest classification, e.g. [6, Chapter 4]: For each action u , we train a classifier LC_u that tries to separate the states with that action from the rest. We then pick that classifier whose predicate minimizes the entropy, i.e.

$$LC := \arg \min_{u \in \mathcal{U}} \left(\text{entr}(\{(\vec{x}, u) \in C \mid LC_u(\vec{x}) = 1\}) + \text{entr}(\{(\vec{x}, u) \in C \mid LC_u(\vec{x}) = 0\}) \right).$$

We considered various linear classification techniques including Logistic Regression [6, Chapter 4], linear Support Vector Machines (SVM) [6, Chapter 7], Perceptrons [6, Chapter 5], and Naive Bayes [39]. However, the latter two yielded significantly larger DTs in all of our experiments, so dtControl does not offer these algorithms to the end-user.

In summary, dtControl currently supports four possibilities for the set PREDS: axis-aligned predicates, the modified oblique split heuristic from [25] and oblique splits obtained either via logistic

regression or linear SVM classifiers. Due to the modular structure of the code, it is easy to extend the existing approaches or add new methods, as described in our developer manual.

4.2 Non-determinism

In the general algorithm described in Section 3, for the sake of simplicity, we restricted our procedure to controllers that deterministically choose a single control input. In case of non-deterministic (also called permissive) controllers, the tuples in the controller C have the form $(\vec{x}, u)$, where u is now a set $\{u_1, u_2, \dots, u_m\}$ of admissible control inputs. One approach to handle non-determinism is to simply assign a *unique label* to each set, and hence reduce the setting to the case where for every state there is only a single label. This means that the DT algorithm can be used in exactly the same way as described in Section 3. This method retains all information that was initially present in the given controller.

The disadvantage of handling non-determinism like this is that the number of unique classes may be as large as $2^{|\mathcal{U}|}$. In order to avoid this blow-up and optimize memory, one can decide to determinize the controller. If we have some knowledge about which value of a control input is optimal, e.g. from domain knowledge or since it was computed by an *optimization algorithm* as in UP-PAAL STRATEGO [12], this information can be used, eliminating the non-deterministic choice. Otherwise, one can use a standard determinization approaches, e.g. picking the value with the *minimum norm*. The tree can then simply be constructed from the determinized labels. Additionally, we propose the following alternative to these determinization approaches.

Novel determinization approach: Maximal frequencies. Our new determinization technique MaxFreq aims to minimize the size of the resulting DT. The underlying general idea is simple: if many of the data points share the same label, a DT learning algorithm should group them together under the common label. This idea naturally gives a determinizing strategy when applied in our context.

Consider a set C of pairs of state and sets of actions. The goal is to identify for each state a single action which can be assigned to it. Let f be the function for action frequency, which maps actions to their number of occurrences in C . Then, for each state $\vec{x}$ such that $(\vec{x}, \{u_1, u_2, \dots, u_m\}) \in C$, we re-assign to $\vec{x}$ the single label u' which appears with the highest frequency. Formally, our determinization procedure produces for each state $\vec{x}$, an action $u'(\vec{x})$, where

$$\forall (\vec{x}, \{u_1, \dots, u_m\}) \in C. u'(\vec{x}) = \arg \max_{u \in \{u_1, \dots, u_m\}} f(u).$$

Once we have determinized C , we can use any method presented in Section 4.1 to find a predicate for the current node. After the set is split, the procedure is recursively applied to both child nodes, recomputing the action frequency each time.

In summary, dtControl offers 3 different possibilities to handle non-determinism: unique labels retaining the information, determinizing upfront by picking the action with the minimal norm, and using the novel heuristic MaxFreq.

5 EXPERIMENTS

All experiments were conducted on a server running on an Intel Xeon W-2123 processor with a clock speed of 3.60GHz and 64 GB

Table 1: Result of running the various methods on 10 different case studies. The ‘Lookup table’ column gives the size of the domain of the original controller. For all other columns, the number of decision paths in the constructed tree is indicated. The case studies are grouped together by the number of control inputs and methods based on whether they preserve non-determinism. ∞ indicates that the computation did not finish within 3 hours; n/a indicates that the approach is not applicable (we cannot determinize, as the model is already deterministic).

Case Study	Most permissive controller					Determinized controller			
	Lookup table	CART	LinSVM	LogReg	OC1	MaxFreq	MaxFreqLC	MinNorm	MinNormLC
Single-input non-deterministic									
cartpole [14]	271	127	126	100	92	6	7	56	39
2D Thermal [13]	40,311	14	14	8	12	5	4	8	4
helicopter [14]	280,539	3,174	2,895	1,877	∞	115	134	677	526
cruise [21]	295,615	494	543	392	374	2	2	282	197
dddc [32]	593,089	136	140	70	90	5	5	11	11
Multi-input non-deterministic									
10D Thermal [15]	26,244	8,649	67	74	2,263	4	10	2,704	28
truck_trailer[18]	1,386,211	169,195	∞	∞	∞	21,598	12,611	95,417	30,888
traffic[35]	16,639,662	6,287	∞	4,477	∞	98	80	690	∞
Multi-input deterministic									
vehicle [32]	48,018	6,619	6,592	5,195	4,886	n/a	n/a	n/a	n/a
aircraft [33]	2,135,056	456,929	∞	407,523	∞	n/a	n/a	n/a	n/a

RAM. We ran the unique-label approach with all 4 possible predicate classes (see Section 4.1): axis-aligned predicates (CART) [9], oblique predicates with linear support-vector machines (LinSVM), logistic regression (LogReg), and the heuristic from [25], called OC1. Note that all these resulting trees represent the maximally permissive controller for the finite abstraction. Additionally, on all the non-deterministic models we ran our novel determinization approach (see Section 4.2) with axis-aligned predicates (MaxFreq), and with oblique predicates (MaxFreqLC where LC stands for linear classifier). For the results in Table 1, we used logistic regression as linear classifier, because it reliably performed well. As a competitor for our determinization approach we use a-priori determinization with the minimum norm, again both with axis-aligned predicates (MinNorm) and with logistic regression for linear predicates (MinNormLC). Additionally, we compare to the random a-priori determinization, to get an impression for possible cases where MinNorm would not be a natural choice but no better is given. However, since the results are always worse, we only report the numbers in [2, Appendix B]. Since some of the algorithms rely on randomization, we ran all experiments thrice and report the median.

We run the discussed algorithms on ten case studies, five of which are marked as multi-input, containing control inputs which are multi-dimensional, i.e. $u = (u_1, \dots, u_m)$. All our algorithms work by giving each multi-dimensional control input a single action label, and then working on these labels as in the case of single-dimensional control inputs.

In order to compare the sizes of the representations of the controllers fairly, we provide two different ways. Firstly, the straightforward way is to compare the number of nodes used in the DT and the number of rows in the lookup table, which we do in Table 2 of [2, Appendix B]. However, a practically more relevant comparison should reflect the number of state symbols needed to capture the behaviour of the controller; these can also be directly related to

memory requirements. To this end, in Table 1 for DTs we report the number of decision paths, as these induce a partitioning of the state space into symbolic states. For more information on this and an example, see Figure 2 and the discussion in Section 6.

Beside comparing DTs to the lookup tables, we also compare them to BDDs. However, BDDs do not directly correspond to the state symbols. Hence we refrain from the state-symbols comparison and do not report BDD sizes in Table 1, but only in [2, Appendix B]. There, we compare the number of nodes in the BDDs to the number of nodes (not decision paths) generated by our DT algorithms. The BDDs were generated using SCOTS for all models but the two from UPPAAL STRATEGO, cruise and 2D Thermal; for these two, we used the dd and autoref Python libraries. The BDDs were minimized as much as possible by calling reordering heuristics until convergence. The results show that the DT algorithms which determinize or which do not use oblique predicates are more scalable, as they were able to compute the result for all case studies, while BDDs timed out on dcdc and traffic. Depending on the case study, BDDs are usually in the same order of magnitude as CART, sometimes better, sometimes worse. On the one hand, on 10D Thermal and truck_trailer, BDDs have an order of magnitude less nodes, but on the other hand CART is able to produce results for dcdc and traffic. Compared to MaxFreq, there is the exception of truck_trailer, where the best BDD has a quarter of the size; on all other models, MaxFreq is at least one order of magnitude better.

6 DISCUSSION

Table 1 shows that DTs are always better than lookup tables. In the case of DTs exactly representing the most permissive controller, our linear-classifier-based algorithm, LogReg, generally performs better than the standard DT learning algorithm CART. An inspection of the trees showed that oblique splits indeed aid in this reduction. In order to save memory, however, our determinizing algorithms may

Figure 2: End-to-end usage of DT-based controller: First, a DT representation is synthesized with the help of `dtControl` (the result of running `MaxFreq` on cartpole is shown here). Then a non-uniform quantizer is implemented at the sensor side, which for each decision path (i.e. a region in the state-space), sends a state symbol to the controller. At the controller, this symbol gives actual control input. In this case, the information needs to be sent over the sensor-controller channel is $\lceil \log_2(6) \rceil = 3$ bits per time unit. The theoretical lower bound on the data rate in this example is 1 bit per time unit to achieve invariance [34].

be used. Here, `MaxFreq` and its linear classifier variant, `MaxFreqLC`, easily outperform all other discussed algorithms, returning trees which can be drawn on a single sheet of paper in most of our case studies! The controller produced by `MaxFreq` for the case study cartpole is depicted in Figure 2a.

Apart from the compact representation of the controllers and efficient determinization, `dtControl` makes controllers more understandable. This helps to do some analysis for the systems and corresponding controllers. A few analyses were mentioned for the temperature control example in the introduction. Another application is that `dtControl` learns how to efficiently partition the state space. In general, the tools synthesizing symbolic controllers use uniform partitioning, i.e. a uniform quantizer is used to discretize the state set. Therefore, they need a large number of symbols to represent the state set. `dtControl` aggregates state symbols where the same control input is admissible to reduce the number of symbols required. In other words, `dtControl` provides a scheme to design non-uniform quantizers (i.e., state encoders with non-uniform partitioning of state-set), illustrated in Figure 2b.

The entries in Table 1 correspond to the necessary number of state symbols. For instance, consider the cartpole example in Table 1. The controller obtained using `SCOTS` requires 271 symbols to represent the domain of the controller, which implies that one needs to send 9 bits per time unit over the sensor-controller channel to achieve invariance. After processing the controller using `dtControl` with `MaxFreq`, we only need 6 symbols to represent the controller, corresponding to only 3 bits information. One can directly relate this idea of constructing efficient static coders to the notion of invariance feedback entropy introduced in [34]. This notion characterizes the necessary state information required by any coder-controller to enforce the invariance condition in the closed loop. For example, in the case of cartpole, the theoretical lower-bound on average bit rate for any static coder-controller to achieve invariance is 1 (obtained through the invariance feedback entropy [34]), which is not far from 3, computed using `dtControl`.

In summary, one can utilize the results provided in this paper for constructing efficient coder-controllers for invariance properties which is an active topic in the domain of information-based control [26].

7 CONCLUSION

We presented `dtControl`, an open-source, easily extensible tool for post-processing controllers synthesized by various tools such as `SCOTS` and `UPPAAL STRATEGO` into small, efficient and interpretable representations. The tool allows for a comparison between various representations in terms of size and performance and also allows us to export the controller both as a graphic and as a code. We also presented a new determinization technique, `MaxFreq`, which easily converts non-deterministic controllers into extremely small deterministic decision trees. Further algorithms for controller representation were thoroughly evaluated and made accessible to the end-user. We believe these small representations will not only allow us to save memory but also help us in understanding and validating the model. As for future work, `dtControl` can be extended with

- further input and output formats, to also support tools such as `pFaces`[18] and `QUEST`[15];
- different predicates: this can be other, possibly even non-linear or non-binary, machine-learning classifiers or richer algebraic predicates utilizing domain knowledge;
- other *impurity measures* instead of entropy, which decide the predicate used for the split

ACKNOWLEDGMENTS

This work was supported in part by the H2020 ERC Starting Grant *AutoCPS* (grant agreement no 804639), the German Research Foundation (DFG) through the grants ZA 873/1-1 and KR 4890/2-1 *Statistical Unbounded Verification*, and the TUM International Graduate School of Science and Engineering (IGSSE) grant 10.06 *PARSEC*.

REFERENCES

- [1] P. Ashok, T. Brázdil, K. Chatterjee, J. Křetínský, C. H. Lampert, and V. Toman. 2019. Strategy Representation by Decision Trees with Linear Classifiers. In *QEST (1)*. Springer, 109–128.
- [2] Pranav Ashok, Mathias Jackermeier, Pushpak Jagtap, Jan Křetínský, Maximilian Weininger, and Majid Zamani. 2020. dtControl: Decision Tree Learning Algorithms for Controller Representation. arXiv:cs.LG/2002.04991
- [3] P. Ashok, J. Křetínský, K. G. Larsen, A. Le Coënt, J. H. Taankvist, and M. Weininger. 2019. SOS: Safe, Optimal and Small Strategies for Hybrid Markov Decision Processes. In *QEST (1)*, D. Parker and V. Wolf (Eds.). Springer, 147–164.
- [4] R. Iris Bahar, Erica A. Frohm, Charles M. Gaona, Gary D. Hachtel, Enrico Macii, Abelardo Pardo, and Fabio Somenzi. 1997. Algebraic Decision Diagrams and Their Applications. *Formal Methods in System Design* 10, 2/3 (1997), 171–206.
- [5] C. Belta, B. Yordanov, and E. A. Gol. 2017. *Formal methods for discrete-time dynamical systems*. Vol. 89. Springer.
- [6] C. M. Bishop. 2007. *Pattern recognition and machine learning, 5th Edition*. Springer.
- [7] T. Brázdil, K. Chatterjee, M. Chmelik, A. Fellner, and J. Křetínský. 2015. Counterexample Explanation by Learning Small Strategies in Markov Decision Processes. In *CAV (1) (Lecture Notes in Computer Science)*, Vol. 9206. Springer, 158–177.
- [8] T. Brázdil, K. Chatterjee, J. Křetínský, and V. Toman. 2018. Strategy Representation by Decision Trees in Reactive Synthesis. In *TACAS (1) (Lecture Notes in Computer Science)*, Vol. 10805. Springer, 385–407.
- [9] L. Breiman, J. H. Friedman, R. A. Olshen, and C. J. Stone. 1984. *Classification and Regression Trees*. Wadsworth.
- [10] R. E. Bryant. 1986. Graph-based algorithms for boolean function manipulation. *IEEE Trans. Comput.* 100, 8 (1986), 677–691.
- [11] Ioannis T. Christou and Sofoklis Efremidis. 2007. An Evolving Oblique Decision Tree Ensemble Architecture for Continuous Learning Applications. In *AAAI (IFIP)*, Vol. 247. Springer, 3–11.
- [12] A. David, P. Gjø Jensen, K. Guldstrand Larsen, M. Mikucionis, and J. H. Taankvist. 2015. Uppaal Stratego. In *TACAS (Lecture Notes in Computer Science)*, Vol. 9035. Springer, 206–211.
- [13] Antoine Girard. 2013. Low-complexity quantized switching controllers using approximate bisimulation. *Nonlinear Analysis: Hybrid Systems* 10 (2013), 34–44.
- [14] Pushpak Jagtap, Fardin Abdi, Matthias Rungger, Majid Zamani, and Marco Caccamo. 2018. Software Fault Tolerance for Cyber-Physical Systems via Full System Restart. *CoRR abs/1812.03546* (2018).
- [15] Pushpak Jagtap and Majid Zamani. 2017. QUEST: A Tool for State-Space Quantization-Free Synthesis of Symbolic Controllers. In *QEST (Lecture Notes in Computer Science)*, Vol. 10503. Springer, 309–313.
- [16] Manuel Mazo Jr., Anna Davitian, and Paulo Tabuada. 2010. PESSOA: A Tool for Embedded Controller Synthesis. In *CAV (Lecture Notes in Computer Science)*, Vol. 6174. Springer, 566–569.
- [17] Kyle D. Julian, Mykel J. Kochenderfer, and Michael P. Owen. 2018. Deep Neural Network Compression for Aircraft Collision Avoidance Systems. *CoRR abs/1810.04240* (2018).
- [18] Mahmoud Khaled and Majid Zamani. 2019. pFaces: an acceleration ecosystem for symbolic control. In *HSCC*. ACM, 252–257.
- [19] Niels Landwehr, Mark A. Hall, and Eibe Frank. 2003. Logistic Model Trees. In *ECML (Lecture Notes in Computer Science)*, Vol. 2837. Springer, 241–252.
- [20] Kim Guldstrand Larsen, Adrien Le Coënt, Marius Mikucionis, and Jakob Haahr Taankvist. 2018. Guaranteed Control Synthesis for Continuous Systems in Uppaal Tiga. In *CyPhy/WESE (Lecture Notes in Computer Science)*, Vol. 11615. Springer, 113–133.
- [21] Kim Guldstrand Larsen, Marius Mikucionis, and Jakob Haahr Taankvist. 2015. Safe and Optimal Adaptive Cruise Control. In *Correct System Design (Lecture Notes in Computer Science)*, Vol. 9360. Springer, 260–277.
- [22] Philipp J. Meyer, Matthias Rungger, Michael Luttenberger, Javier Esparza, and Majid Zamani. 2017. Quantitative Implementation Strategies for Safety Controllers. *CoRR abs/1712.05278* (2017).
- [23] T. M. Mitchell. 1997. *Machine learning*. McGraw-Hill.
- [24] Sebt Mouelhi, Antoine Girard, and Gregor Gößler. 2013. CoSyMA: a tool for controller synthesis using multi-scale abstractions. In *HSCC*. ACM, 83–88.
- [25] S. K. Murthy, S. Kasif, S. Salzberg, and R. Beigel. 1993. OC1: A Randomized Induction of Oblique Decision Trees. In *AAAI*. AAAI Press / The MIT Press, 322–327.
- [26] G. N. Nair, F. Fagnani, S. Zampieri, and R. J. Evans. 2007. Feedback control under data rate constraints: An overview. *Proc. of the IEEE* 95, 1 (2007), 108–137.
- [27] Daniel Neider, Shambwaditya Saha, and P. Madhusudan. 2016. Synthesizing Piece-Wise Functions by Learning Classifiers. In *TACAS (Lecture Notes in Computer Science)*, Vol. 9636. Springer, 186–203.
- [28] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* 12 (2011), 2825–2830.
- [29] Larry D. Pyatt and Adele E. Howe. 1998. *Decision Tree Function Approximation in Reinforcement Learning*. Technical Report. Computer Science Department, Colorado State University.
- [30] J. R. Quinlan. 1993. *C4.5: Programs for Machine Learning*. Morgan Kaufmann.
- [31] G. Reissig, A. Weber, and M. Rungger. 2016. Feedback refinement relations for the synthesis of symbolic controllers. *IEEE Trans. Automat. Control* 62, 4 (2016), 1781–1796.
- [32] M. Rungger and Zamani M. 2016. SCOTS: A Tool for the Synthesis of Symbolic Controllers. In *HSCC*. ACM, 99–104.
- [33] Matthias Rungger, Alexander Weber, and Gunther Reissig. 2015. State space grids for low complexity abstractions. In *CDC*. IEEE, 6139–6146.
- [34] Matthias Rungger and Majid Zamani. 2017. Invariance Feedback Entropy of Uncertain Control Systems. *CoRR abs/1706.05242* (2017).
- [35] A. Swikir and M. Zamani. 2019. Compositional synthesis of symbolic models for networks of switched systems. *IEEE Control Systems Letters* 3, 4 (2019), 1056–1061.
- [36] P. Tabuada. 2009. *Verification and control of hybrid systems: a symbolic approach*. Springer Science & Business Media.
- [37] Paul E. Utgoff. 1988. Perceptron Trees: A Case Study In Hybrid Concept Representations. In *AAAI*. AAAI Press / The MIT Press, 601–606.
- [38] Ivan S. Zapreev, Cees Verdier, and Manuel Mazo Jr. 2018. Optimal Symbolic Controllers Determinization for BDD storage. In *ADHS (IFAC-PapersOnLine)*, Vol. 51. Elsevier, 1–6.
- [39] Harry Zhang. 2004. The Optimality of Naive Bayes. In *FLAIRS Conference*. AAAI Press, 562–567.