



Identification of hybrid systems with dynamics-based modeling through symbolic regression^{☆,☆☆}

Swantje Plambeck^{a,*,} Maximilian Schmidt^{a, ID}, Audine Subias^b, Louise Travé-Massuyès^{b, ID}, Goerschwin Fey^{a, ID}

^a Hamburg University of Technology, Hamburg, Germany

^b LAAS-CNRS, CNRS, INSA, University of Toulouse, Toulouse, France

ARTICLE INFO

Keywords:

Hybrid systems
Symbolic regression
Decision trees
Model inference
Machine learning

ABSTRACT

Hybrid systems combine both continuous and discrete behavior. These systems serve as models in many fields, including control systems, robotics, and industrial processes. However, due to their complexity, finding an accurate model is a challenge. This paper presents a holistic approach to learning models of hybrid systems using symbolic regression. Our method leverages symbolic regression to automatically discover accurate and interpretable mathematical models in the form of hybrid systems from observed data. An advantage of our algorithm is that it detects transitions between different behavioral modes of a system based on the inherent dynamics.

From learned expressions for the dynamical behavior of a system, we form a hybrid system by combining the learned expressions with a decision tree determining the current behavioral mode from data. This hybrid decision tree serves regression, prediction, and further related tasks. Our results demonstrate that symbolic regression can effectively identify the underlying dynamics of a real hybrid system and predict output signals on new input data with high accuracy.

1. Introduction

Hybrid systems are abstract models of systems that exhibit both continuous and discrete behavior. They are used to model a wide range of systems, including cyber–physical systems, manufacturing systems, and many more (Alur, 2015; Lunze and Lamnabhi-Lagarrique, 2009). Due to their inherent combination of continuous and discrete behavior, the identification of hybrid systems is a challenging task. Nevertheless, accurate abstract models are essential for, e.g., diagnosis, verification, and control. Hybrid systems operate with a finite set of modes, each representing a distinct dynamic behavior of the system. Within a mode, the continuous system dynamics are defined by flow functions over the system variables.

We propose a holistic approach for the identification of an inferable hybrid system. The identification of continuous dynamics uses Symbolic Regression (SR) based on our previous work (Plambeck et al., 2024b). Existing hybrid system identification methods introduce a general procedure that identifies modes and their flow functions to create a model as shown in Fig. 1 (Plambeck et al., 2024a). In Plambeck et al. (2024b),

novel approaches for improving the steps (a) to (c) for detecting transitions, grouping of observations with identical dynamics into modes, and identifying modes through symbolic expressions have been proposed. Here, in addition to consolidate the work in Plambeck et al. (2024b), we extend to a holistic hybrid system identification algorithm by adding the final model construction step. This enables the inference of the model on new unseen data. The model construction step uses decision tree learning. The idea of using decision trees for model construction has already been successfully implemented in Plambeck et al. (2024a). The work in Plambeck et al. (2024a) aims at fast identification mechanisms and has a restricted representation of dynamics as well as limited segmentation mechanisms. Here, we leverage the capabilities of SR to overcome these limitations by identifying symbolic expressions that capture the dynamic behavior of the system across consecutive segments. SR provides human-readable mathematical expressions for the system dynamics, which is essential for the interpretation of the identified models. The usage of decision trees for model construction also provides an interpretable representation.

Our approach detects mode transitions directly based on the dynamics of the observed data. This goes beyond existing identification

[☆] This article is part of a Special issue entitled: ‘Euromicro SEAA 2024’ published in The Journal of Systems & Software.

^{☆☆} Editor: Laurence Duchien.

* Corresponding author.

E-mail address: swantje.plambeck@tuhh.de (S. Plambeck).




Fig. 1. Steps of Hybrid System Identification.

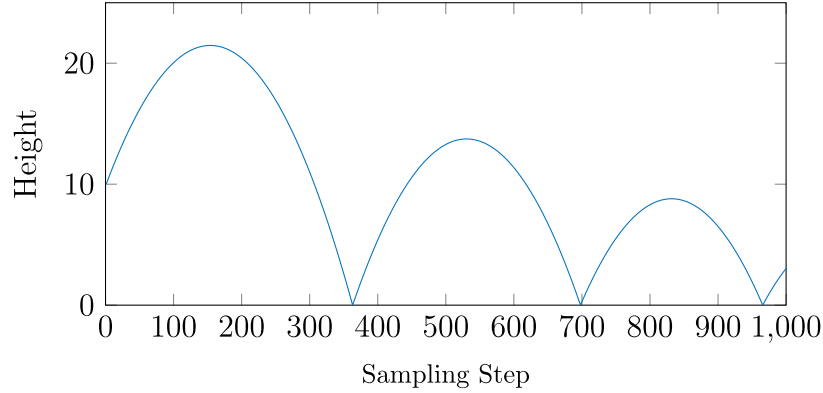


Fig. 2. Example of an observed trajectory of a bouncing ball. The height of the ball is measured at each sampling step.

strategies, which focus on similarities of observations to detect whether observations belong to the same or different modes of a hybrid system (Barbosa Roa et al., 2019; Yang et al., 2022; Niggemann et al., 2012). However, with the same underlying dynamics, the observed trajectories can be completely different, for example if the initial conditions differ. The similarity-based approach can hence lead to many unnecessary transitions. The dynamics-based transition detection is a key contribution of our work, which puts the focus on the underlying physical laws and detects transitions only if needed. This leads to a simplified model and a more efficient learning process. Further, we leverage the ability of SR to find symbolic expressions for the system dynamics, which is essential for the interpretation of the identified models.

An illustrative example for the relevance of detecting transitions based on the dynamics of the observed data is shown in Fig. 2. The figure shows a trajectory of a bouncing ball, where the height of the ball is observed over time. Considering the similarity of signals, transitions would be detected at each bounce. Further, the height values between the bounces are different, and thus a similarity-based approach would consider them as different modes. However, the bouncing ball has one single mode and its dynamics can be described by a single function, which is:

$$\begin{pmatrix} \dot{x} \\ \dot{v} \end{pmatrix} = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} x \\ v \end{pmatrix} + \begin{pmatrix} 0 \\ -g \end{pmatrix}, \quad (1)$$

where x is the height of the ball, v is the velocity, and g is the gravitational acceleration. Thus, every singularity does not identify a transition and a new dynamic mode of the system. Our approach automatically determines that the dynamics of the bouncing ball follows a single expression.

Our algorithm works as follows. We start from a small window of the observed data, learn a symbolic expression using SR, and then iteratively enlarge this window until a transition point is detected. A transition point is detected as soon as the fitness of the learned expression drops. The segments between detected transition points are then grouped to identify the modes of the hybrid system. For this grouping step, we reuse SR to learn expressions on unions of previously detected segments. Segments with a small loss on the same symbolic expression form a group. The expressions found for each group identify the modes. Finally, the groups are mapped to a learning set for decision tree learning. The learned decision tree represents the transitions by

determining the current mode of the system. The final hybrid decision tree consists of a decision tree and a map from system mode to symbolic expression.

Novelties compared to Plambeck et al. (2024b) as well as Plambeck et al. (2024a) and Plambeck et al. (2024c), come from the following contributions of this paper:

- the completion to a holistic identification approach by adding a model construction step with a decision tree,
- the analysis of the parameter importances for the presented algorithms,
- a complexity analysis of the proposed learning algorithms, which in nominal cases is linear in the number of samples,
- the refinement and extension of the empirical evaluation. This includes a metric for the identification of transition points as introduced in Plambeck et al. (2024c). One of the examples is further extended by a second scenario with limited system information compared to Plambeck et al. (2024b), and
- the comparison of our approach with two state-of-the-art approaches for hybrid system identification, FaMoS (Plambeck et al., 2024a) and HAuLearn (Yang et al., 2022), on the same example systems,
- the open-source publication of an implementation in Python (Plambeck, 2024).

With respect to the state-of-the-art in hybrid system identification as a whole, this paper contrasts the other works in that:

- the identification is based on the dynamics of the system, focusing on the physical laws underlying the data rather than on the similarity of data,
- the detection of transitions and grouping segments with similar dynamics are both achieved utilizing SR,
- the last step of model learning constructs the hybrid model with a new type of decision tree integrating the learned symbolic expressions.

The paper is structured as follows. In Section 2, we review related work on system identification, specifically for hybrid systems and symbolic regression. In Section 3, we introduce the necessary formal definitions. In Section 4, we present our approach for the identification of hybrid systems using SR, the individual steps as well as the final

model construction. In Section 6, we evaluate our approach on two representative examples. Finally, in Section 7, we conclude the paper.

2. Related work

In the following, we summarize and relate literature from three main aspects of this work, which are SR in the context of system identification, modern advances of decision trees, and existing algorithms for hybrid system identification.

2.1. Symbolic regression for system identification

SR is a method for regression and system identification to find symbolic expressions that match a learning data set. Contrary to traditional regression methods, SR is not restricted to a specific class of functions or structure of an expression like, e.g., polynomial regression, but uses a set of basic operators to construct complex expressions freely. The development of SR has been supported by the advancement of genetic programming, which is commonly used to implement SR algorithms (Kronberger et al., 2021). Several studies show that the performance of SR based on genetic programming improves by using a suitable parametrization. The population size and the number of iterations are considered the most crucial parameters — usually a larger population and number of iterations lead to higher accuracy (Feldt and Nordin, 2000; Loveard and Ciesielski, 2002; Langdon, 2022). In addition to genetic programming, there exist other approaches for SR, e.g., using deep reinforcement learning (Petersen et al., 2021), lattices (Broløs et al., 2021), or sparse regression methods (Muthyala et al., 2025).

SR has been applied to a wide range of problems, including the identification of physical concepts from data (Schmidt and Lipson, 2009), mining the expression of diagnosis indicators (Goupil et al., 2024), identify properties in material sciences (Wang et al., 2024), and even for the identification of a machine learning algorithm (Real et al., 2020). The works closest to ours are:

- the system identification procedure of Gaucel et al. (2014), which focuses on learning matrix ordinary differential equations with SR but does not consider hybrid dynamics.
- the learning process of symbolic representations of hybrid dynamical systems with multi-modal SR of Ly and Lipson (2012), which uses SR with an expectation–maximization metric to learn symbolic expressions on noise data and does not integrate learned functions into a model such as we do with a decision tree.

Our work advances beyond previous approaches by integrating all necessary components to address the hybrid nature of systems, specifically detecting transitions and grouping segments with similar dynamics, incorporating benefits of SR.

2.2. Advancements of decision trees

Decision trees are a well-established and widely used method in machine learning, known for their interpretability and effectiveness. Some research is orthogonal to ours and uses well-established decision trees in an innovative manner. For example, Fong and Motani (2024) proposes another combination of symbolic regression and decision trees, which results in a more effective structure of learned decision trees, but applies only for classification tasks. Goupil et al. (2024) combines decision trees with symbolic classification for diagnosis purposes. Further, Plambeck et al. (2022) and Jimenez et al. (2019) use decision trees to learn discrete models of cyber–physical systems, where Plambeck et al. (2022) showcases theoretical limitations of decision trees in representing time-discrete behavior.

Our work uses decision trees for the hybrid model construction step, thus, generating a hybrid decision tree, which enhances a traditional decision tree with symbolic expressions.




Fig. 3. Example of a Symbolic Regression Tree.

2.3. Identification of hybrid systems

In this paper, we focus on the identification of hybrid systems. Existing approaches use a variety of methods, but follow the procedure of Fig. 1 consisting of steps for the detection of transitions, grouping of segments, mode identification, and model construction. In Yang et al. (2022), the authors use a sliding window approach to detect transitions, considering the similarity of data points to detect transitions. This procedure is unaware of the underlying dynamics and potentially reports transitions that do not exist. In contrast, our approach detects transitions based on the dynamics of the observed data, thus, detecting transitions only if necessary.

Similarly, Niggemann et al. (2012) uses a wavelet transform to detect transitions, which focuses on similarities in the frequency domain. Thus, this approach also does not consider the underlying dynamics.

There are clustering-based approaches (Barbosa Roa et al., 2019; Tappler et al., 2023) which use spatial proximity of observations to segment and cluster trajectories into groups. Nevertheless, similar dynamics do not necessarily lead to similar observations, which is a problem solved by our approach.

For the model identification step, other approaches use, e.g., feed forward neural networks (Niggemann et al., 2012). In contrast, we use SR to identify explicit symbolic expressions for the dynamics of the system. This finds human-readable expressions for the dynamics, which is essential for the interpretation of the identified models.

3. Preliminaries

In this section, we introduce the basics of SR and decision trees and formally define hybrid systems and system observations.

3.1. Symbolic regression

SR is a machine learning technique that finds symbolic expressions that describe the relationship between input and output variables. The goal is to find a function r that maps input variables \mathbf{i} to an output variable o such that $o = r(\mathbf{i})$. The function r is represented as a symbolic expression in terms of a set B of basic functions and operators applied to the input variables (Koza, 1992). To represent symbolic expressions, often a tree structure is used, where the leaves are variables or constants and the inner nodes are operators from B . An example of a symbolic regression tree is shown in Fig. 3 for the expression $o = (i_1 \cdot i_2) + 5$.

The search for the best symbolic expression is guided by a fitness metric that evaluates the quality of candidate symbolic expressions on a learning data set. Often, a trade-off between the loss of the expression on the learning data and the length of the candidate expression is used as a fitness metric. This strategy avoids a bloat of the expressions (Poli et al., 2008). A parsimony coefficient ρ works as a regularization parameter in this trade-off.

In the scope of this work, we use the framework PySR, which uses genetic programming (Koza, 1992; Cranmer, 2023). Using genetic




Fig. 4. Genetic Programming Algorithm for Symbolic Regression.




Fig. 5. Example of a Decision Tree.

programming, the search for the best symbolic expression is performed by evolving a population of candidate expressions over so-called *generations*. Fig. 4 illustrates this process. A population of size p is initialized with random expressions. In each generation, the population is evaluated using a fitness metric. New candidate solutions are generated by applying genetic operators, such as mutation and crossover, to the best candidates. The process continues for a predefined number n of generations.

3.2. Decision trees

A decision tree (Shalev-Shwartz and Ben-David, 2014) is a tree $T = (V, E)$ with nodes V and edges E that represents a classifier $d : \mathcal{X} \rightarrow \mathcal{C}$. The set \mathcal{X} consists of vectors of feature values $\mathbf{f} = [f_1, \dots, f_m]$, while \mathcal{C} is a set of classes. A decision tree learner constructs a decision tree based on a learning set \mathcal{L} , where each element $(\mathbf{f}, c) \in \mathcal{L}$ is a tuple of a feature vector $\mathbf{f} \in \mathcal{X}$ and a class label $c \in \mathcal{C}$. The decision rule d adapts to these tuples. The set of nodes with no outgoing edges V_L are the leaf nodes of the tree. All nodes $V \setminus V_L$ are *inner nodes* of the tree. Each node v in T associates to a subset L_v of \mathcal{L} , denoted by $v \sim L_v$ and splits this set into subsets forming child nodes. The splitting rule bases on one of the features in \mathbf{f} . Each leaf $v_l \sim L_{v_l}$ has a label $c_{v_l} \in \mathcal{C}$ given by the majority of class labels in the associated set, i.e.,

$$c_{v_l} = d(\mathbf{f}) = \underset{c}{\operatorname{argmax}} \#c, \text{ where } \#c = |\{\mathbf{f} : (\mathbf{f}, c) \in L_{v_l}\}|. \quad (2)$$

Fig. 5 shows an example of a decision tree. The inner nodes represent decisions based on the values of the feature variables x_1 and x_2 , while the leaf nodes represent the classes $\mathcal{C} = \{a, b, c\}$. The edges between the nodes represent the conditions for the decisions.

A common decision tree learning algorithm is the CART algorithm, which creates a binary decision tree (Loh, 2011).

3.3. Hybrid systems

Definition 1. A Hybrid System (Branicky, 2005) is a 5-tuple $(X, Q, \mathcal{F}, \mathcal{T}, \Sigma)$, where

- $X = \{x_0, x_1, x_2, \dots, x_n\} = I \cup O \cup S$ is the set of system variables which consist of input variables I , output variables O and state variables S . The state variables may include the continuous time




Fig. 6. Example of a Hybrid System: Heating System.

t . Also, derivatives of variables may form individual variables in X .

- Q is the set of modes.
- \mathcal{F} is the set of flow functions. A flow function $f_q \in \mathcal{F}$ defines the change of the state variables S and the current output variables O within the mode $q \in Q$ based on the current values.
- $\mathcal{T} : Q \times \Sigma \rightarrow Q$ defines transitions between modes. A transition is triggered if the corresponding guard $\sigma \in \Sigma$ is active.
- Σ is a set of guards leading to transitions between modes. Each guard is a set of conditions on the variables in X . A transition is triggered if the conditions are met.

A hybrid system combines discrete and continuous behavior and is usually an abstraction of a complex real-world system. Discrete behavior is captured by the discrete modes Q and the transitions \mathcal{T} . The continuous dynamics are represented by the flow functions \mathcal{F} of the modes. Observations of dynamics, i.e., changes in the values of the variables, are considered as *trajectories* which are introduced in the following.

Fig. 6 shows an example of a hybrid system, which is a heating system. The system has two modes: heating and cooling. The flow functions for the modes are given by the differential equations for the temperature change in the system. The transitions between the modes are triggered by the temperature reaching a certain threshold.

3.4. Observations

We define a *trajectory* as an observation in form of a multidimensional time series. Every sample in this series contains values for the system variables X .

The input to the identification algorithm is a set of trajectories. Each trajectory covers multiple transitions between the modes and stems from observing the real-world system for a finite time interval.

4. Identification of a hybrid system

In this section, we provide our identification strategy for hybrid systems. We name our method HyDRA (Hybrid systems via Dynamics and Regression-based Analysis) inspired by the multi-headed Hydra of mythology which reflects the multi-modal nature of hybrid systems. First, we give an overview of the approach and, then, describe the segmentation of trajectories, the grouping of segments, and the model construction step in detail. Finally, we describe how the hybrid system is inferred.

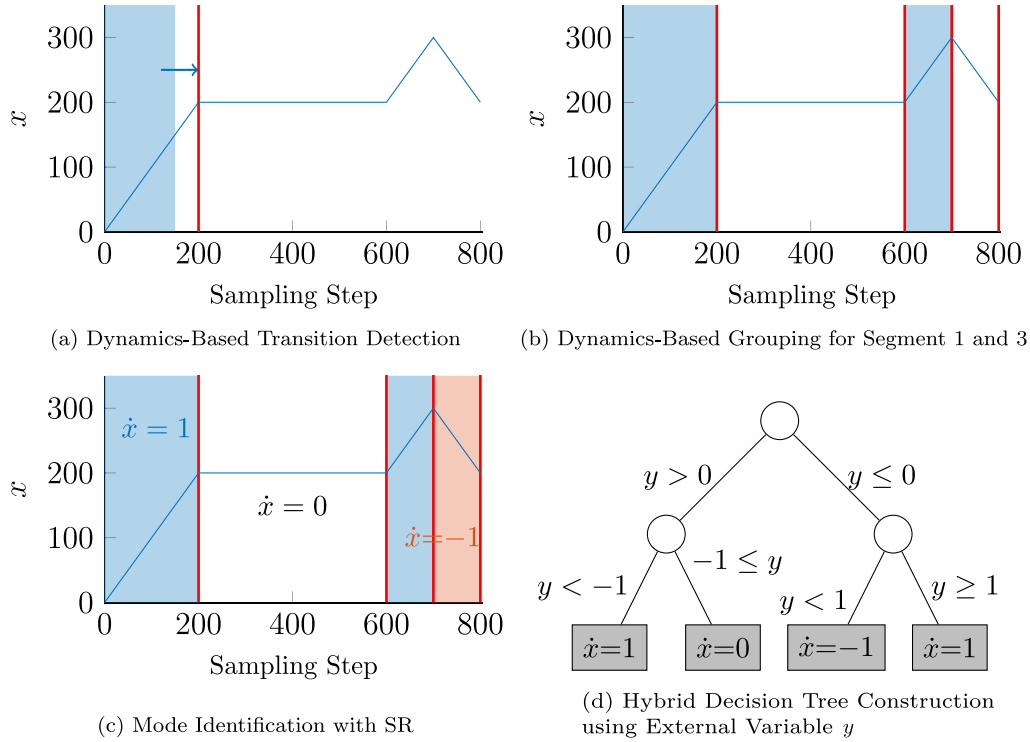


Fig. 7. Illustration of the Proposed Learning Approach.

4.1. Overview & system properties

Definition 1 specifies hybrid systems in a general form. The most characteristic property of hybrid systems is the combination of *continuous* dynamics (defined by the flow functions \mathcal{F}) with *discrete* modes \mathcal{Q} . To learn both parts, identification of hybrid systems includes multiple subproblems as shown in Fig. 1:

- detection of transition points between dynamics, separating the trajectories into segments,
- grouping of segments with identical dynamics, forming discrete modes,
- identification of the continuous dynamics for each mode, i.e., the flow functions of \mathcal{F} ,
- model construction, i.e., creation of an inferable hybrid system from the results of the previous steps.

We cover all of these steps, where the main idea is illustrated in Fig. 7. The first step (a) detects transition points by extending a window over the trajectory until the dynamics do not follow the same flow function anymore. Step (b) groups segments with identical dynamics. Step (c) includes the mode identification of Step (b) as each flow function learned by SR corresponds to a mode of the hybrid system. Finally, Step (d) constructs a hybrid decision tree model from the results of the previous steps.

We, hence, present a complete identification procedure. The procedure contains an algorithm using SR to detect the transition points of trajectories observed on a hybrid system. A second algorithm groups and identifies the flow functions of modes with SR using the segmented trajectories. Lastly, the procedure has a strategy to form a hybrid system involving decision tree learning. This final step combines the results of the previous step to a model that can be inferred and predicts system behavior on new, unseen data.

Our approach considers one-dimensional flow functions. The extension to higher dimensions is straightforward by processing dimensions individually, but requires additional consideration for the detection of segments, as we have to combine detected transition points for

all dimensions. Further, we assume that the training data covers the dynamical modes sufficiently to extract flow functions of the system.

4.2. Detection of transition points

Data: trajectory

Result: S , expressions

Hyperparameter: $l_{init}, l_{step}, n_{init}, n_{update}, segmentationCriterion, \rho_S, p_S, \mathcal{B}_S$

```

1 loss ← [];
2  $i_{start} \leftarrow 0; i_{end} \leftarrow l_{init}; n \leftarrow n_{init};$ 
3 while  $i_{end} < len(trajjectory)$  do
4   while loss is empty or segmentCriterion(loss) do
5     window ← trajectory[ $i_{start}, \min(i_{end}, len(trajjectory))$ ];
6     exp, fit ← learnExpression(window, n);
7     loss[0] ← loss[1];
8     loss[1] ← fit;
9      $i_{end} \leftarrow i_{end} + l_{step};$ 
10     $n \leftarrow n_{update};$ 
11  end
12   $S \leftarrow S \cup \{window[0, end - l_{step}]\};$ 
13  expressions.append(exp);
14  loss ← [];
15   $i_{start} \leftarrow i_{end} - l_{step}; i_{end} \leftarrow i_{start} + l_{init}; n \leftarrow n_{init};$ 
16  resetSR;
17 end

```

Algorithm 1: Detection of Transition Points

Our approach for detecting transition points between different modes of a hybrid system uses SR. We define *transition points* as the first samples in the observed trajectory after the transition between two dynamical modes. Algorithm 1 describes the procedure to detect these points. Given an observed trajectory of the system, we process the trajectory using a window (Line 5) that initially covers a fixed segment at the beginning of the trajectory. We propose a *segmentation criterion* which determines whether the loss of the learned expression

drops, when updated with additional data. As long as the segmentation criterion is fulfilled, the window is extended to the right by l_{step} (Line 9). As soon as the segmentation criterion is not fulfilled anymore, we detect a transition point between two modes. The identified window is stored as a segment in the set S . In the inner loop, we learn the symbolic expression incrementally, i.e., by performing n_{update} -many SR iterations to adapt the current expression to the extended window. Whenever a transition point is found, learning restarts (Line 16) and creates a new window starting at the end of the previously identified window. For each segment, a new expression is learned from scratch.

The loss function in Algorithm 1 determines the deviation between the estimation of the learned expression and the learning data. Here, the mean-square error is used. The algorithm is parametrized by:

- l_{init} : the initial window size when learning an expression from scratch,
- l_{step} : the step-width for extending the window in the inner loop,
- n_{init} : the number of iterations of SR when learning an expression from scratch,
- n_{update} : the number of iterations of SR for updating the learned expression on an extended window,
- $segmentationCriterion$: the criterion to determine whether a transition point is reached. Here, the segmentation criterion executes the following evaluation:

$$segmentationCriterion(loss) = (\text{loss}[1] < \tau) \text{ or } (\text{loss}[1] \leq \text{loss}[0]), \quad (3)$$

which evaluates to true, if the loss of the current expression is smaller than a stagnation threshold τ or if the loss does not decrease anymore compared to the previous iteration.

- τ : A stagnation-threshold for the segmentation criterion.

Additionally, SR has the following parameters:

- ρ_S : the parsimony coefficient of SR,
- p_S : the number of individuals in the population of SR,
- B_S : the set of basic operators and functions for SR.

4.3. Grouping and identification of modes

Segments whose dynamics are described by the same symbolic expression refer to the same mode of the system. Thus, the next step groups segments with identical dynamics. Algorithm 2 shows the pseudocode for segment grouping using SR. The input to the grouping is the set of detected segments S from the previous step. Output of the algorithm are two maps G and $expressions$, which map a group ID γ to a conjunction of segments and a symbolic expression, respectively.

The first segment forms a first candidate group (Line 1). Afterward, we iterate for every segment in S (Line 2) over all known groups (Line 4). The current segment and the current group are combined to one data set to learn an expression with SR (Line 5). If the loss of the learned expression is small, the current segment is included in the current group (Lines 6 and 7). If no matching group is found for the current segment, the segment forms a new group (Line 13).

The hyperparameters of the algorithm are:

- n_{group} : the number of iterations of SR for learning an expression on combined segments,
- $groupingCriterion$: a criterion to determine whether a segment is joined with a group. Here, we use the following evaluation:

$$groupingCriterion(s, g) = \text{loss}(\{s\} \cup g) \leq \varphi \cdot \text{loss}(g), \quad (4)$$

where $\text{loss}(g)$ gives the loss of the group observed before combining with the segment s and $\text{loss}(\{s\} \cup g)$ is the loss after

combination — the smaller, the better. Here, the mean-square error is used as the loss metric. The grouping criterion evaluates to true if the loss of the combined segments is smaller than φ times the loss of the group before combination. The parameter φ is a relaxation parameter, which allows for some flexibility in the grouping.

- φ : the relaxation parameter of the grouping criterion,

Additionally, SR has the following parameters:

- ρ_G : the parsimony coefficient of SR,
- p_G : the number of individuals in the population of SR,
- B_G : the set of basic operators and functions for SR.

Data: S

Result: G , $expressions$

Hyperparameter: $n_{group}, groupingCriterion, \rho_G, p_G, B_G$

```

1  $G \leftarrow \{(0, S.pop())\};$ 
2 for  $s \in S$  do
3    $groupFound \leftarrow \text{False};$ 
4   for  $(\gamma, g) \in G$  do
5      $exp \leftarrow \text{learnExpression}(\{s\} \cup g, n_{group});$ 
6     if  $groupingCriterion(s, g)$  then
7        $g \leftarrow \{s\} \cup g;$ 
8        $expressions[\gamma] \leftarrow exp;$ 
9        $groupFound \leftarrow \text{True};$ 
10      break;
11    end
12  end
13  if not  $groupFound$  then
14     $G.append((|G|, s));$ 
15  end
16 end

```

Algorithm 2: Grouping of Segments

4.4. Model construction

In our previous work (Plambeck et al., 2024a), we propose the usage of decision trees for hybrid system construction. We construct a learning set for decision tree learning from the results of the previous steps. Let each group g of segments s have a unique group ID γ . The set G of groups holds all tuples (γ, g) . Each sample $s[k]$ is a feature vector \mathbf{f} and the respective group ID a class label c . Thus, the learning set \mathcal{L} is defined as follows:

$$\mathcal{L} = \{(s[k], \gamma), \quad \forall s \in g, \gamma \in \mathbb{N}^0, k \in \{0, \dots, |s| - 1\} \text{ with } (\gamma, g) \in G\}, \quad (5)$$

where $s[k]$ gives the values of all variables X of the system for sample k of segment s in group g . These values are used as the feature vector \mathbf{f} , while the group ID γ is used as the class label c . On this learning set, a decision tree T is learned.

We finally define a hybrid decision tree M_{HDT} . The hybrid decision tree model consists of the learned decision tree T as well as the map $expressions$, which is a result of the grouping step. Fig. 8 gives an example of such a hybrid decision tree for a system with three variables $X = \{x_0, x_1, x_2\}$. Every inner node of the decision tree splits on one of the variables. The labels of the edges of the decision tree are the splitting rules. For example, the root node in Fig. 8 decides on the value of the variable x_1 and the left edge is valid for $x_1 > 3.3$, while the right edge is valid for $x_1 \leq 3.3$. The leaf nodes classify one of the group IDs $\gamma \in \{0, 1, 2\}$. Finally, the expressions map these group IDs to symbolic expressions as shown on the left of Fig. 8. The hybrid decision tree forms a variant of a hybrid system as defined in Definition 1 by representing the modes Q and flow functions F by the expressions while the transitions \mathcal{T} and the guards Σ are encoded in the decision tree.

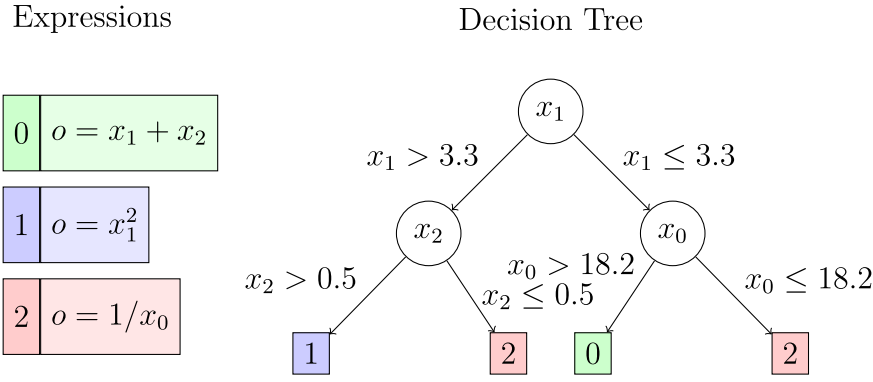


Fig. 8. Hybrid Decision Tree consisting of decision tree and expression.

4.5. Hybrid system inference

With the hybrid decision tree, a representation of the guards and transitions of the hybrid system is found. Additionally, the model is inferable and can perform predictions on new samples. The inference of the model contains three steps as shown in Algorithm 3. Using the hybrid decision tree M_{HDT} , we first apply the decision function d of the decision tree to the sample \mathbf{x} to get the current group ID γ . Using this group ID, we identify the correct *flow*-function. This function is used to evaluate the system dynamics using the variable values in \mathbf{x} .

Data: \mathbf{x}, M_{HDT}
Result: O

```

1 begin
2    $\gamma \leftarrow M_{HDT}.tree.d(\mathbf{x});$ 
3    $flow \leftarrow M.expressions[\gamma];$ 
4    $O \leftarrow flow(\mathbf{x});$ 
5 end

```

Algorithm 3: Inference of Hybrid Decision Tree

5. Complexity analysis

We analyze the computational complexity of our approach by examining each step of the identification procedure. First, we present the complexity with respect to model size, i.e., the number of modes d_m , treating SR as a black-box algorithm with complexity K . Second, we analyze the complexity with respect to the number of learning samples. Here, we assume a genetic programming-based SR algorithm (Fig. 4). The complexity per iteration is $O(p \cdot N_s)$, where p is the population size and N_s is the number of samples in the learning data set. Over n iterations, the total SR complexity is $K = O(p \cdot n \cdot N_s)$.

5.1. Complexity in terms of model size

The complexity of the identification procedure is determined by the complexity of Algorithm 1 for transition point detection, Algorithm 2 for segment grouping, and the decision tree construction step.

Transition point detection. The transition point detection depends on the number of segments d_s , the number of window extensions per segment d_e , and has a complexity of $O(d_s \cdot d_e \cdot K)$. The number of segments is linear in the number of modes d_m , i.e., $d_s = O(d_m)$, and the maximum number of window extensions per segment is constant, i.e., $d_e = O(1)$. Thus, the complexity of transition point detection is:

$$O(d_s \cdot d_e \cdot K) = O(d_m \cdot 1 \cdot K) = O(d_m \cdot K). \quad (6)$$

Decision Tree

Segment grouping & mode identification. The segment grouping step has a complexity of $O(d_s \cdot d_g \cdot K)$, where d_g is the number of groups. Again, the number of segments is linear in the number of modes, i.e., $d_s = O(d_m)$ and the number of groups is equivalent to the number of modes, i.e., $d_g = d_m$. Thus, the complexity of segment grouping is:

$$O(d_s \cdot d_g \cdot K) = O(d_m \cdot d_m \cdot K) = O(d_m^2 \cdot K). \quad (7)$$

Model construction & overall complexity. The complexity of the decision tree construction depends on the number of learning samples and is, thus, considered as $O(1)$ in terms of model size.

The overall complexity of the identification procedure is dominated by the transition point detection and segment grouping steps. As the phases are executed sequentially, the complexity sums up to:

$$O(d_s \cdot d_e \cdot K + d_s \cdot d_g \cdot K) = O(d_m \cdot K + d_m^2 \cdot K) \quad (8)$$

$$= O(d_m^2 \cdot K). \quad (9)$$

In summary, the complexity of the identification procedure in terms of model size is $O(d_m^2 \cdot K)$, where d_m is the number of modes and K is the complexity of the SR algorithm.

5.2. Complexity in terms of learning samples

Next, we analyze the complexity of the identification procedure with respect to the number of learning samples N_t in the trajectory. Here, we also assume a genetic programming-based SR algorithm (Fig. 4) with complexity $K = O(p \cdot n \cdot N_s)$, where p is the population size, n is the number of iterations, and N_s is the number of samples in the learning data set. The parameters p and n are constant for the SR algorithm.

Transition point detection. Algorithm 1 has a complexity of

$$O(d_s \cdot d_e \cdot K), \quad (10)$$

where d_s is the number of detected segments (outer loop iterations), d_e is the maximum number of window extensions per segment (inner loop iterations), and K is the SR complexity. The algorithm processes the trajectory using two nested loops, calling SR in each iteration to learn an expression for the current window.

The outer loop advances by l_{init} samples per iteration, while the inner loop extends the window by l_{step} samples. For a trajectory of length N_t , the total number of SR calls is bounded by:

$$d_s \cdot d_e \leq \frac{N_t}{\min(l_{init}, l_{step})}. \quad (11)$$

With $l_{init}, l_{step} \sim N_s$, we have $d_e = O(1)$. Assuming many segments, the segment size is small, while for a few segments, the segment size is large, i.e., $d_s \cdot N_s = O(N_t)$. Thus, the complexity simplifies to:

$$O(d_s \cdot d_e \cdot p \cdot n \cdot N_s) = O(p \cdot n \cdot N_t), \quad (12)$$

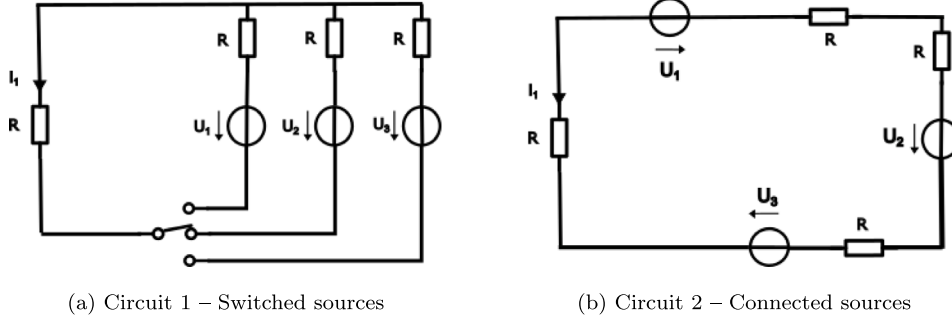


Fig. 9. Passive Electrical Circuits.

where we use $K = O(p \cdot n \cdot N_s)$. However, if initial window and step sizes are small and segments are large, i.e., $l_{init}, l_{step} \ll N_t$, we have $d_s \cdot d_e = O(N_t)$, $N_s = O(N_t)$, and the complexity approaches

$$O(p \cdot n \cdot N_t^2). \quad (13)$$

The choice of l_{init} and l_{step} influences whether a complexity of Eq. (12) or Eq. (13) applies. Thus, the choice of these parameters is a trade-off: larger values reduce computational cost but may miss transition points, while smaller values increase accuracy at higher computational expense.

Segment grouping & mode identification. Algorithm 2 compares each segment against existing groups, requiring $O(d_s \cdot d_g)$ comparisons in the worst case. Thus, the algorithm has complexity

$$O(d_s \cdot d_g \cdot K) = O(d_s \cdot d_g \cdot p \cdot n \cdot N_s), \quad (14)$$

where d_s is the number of segments, d_g is the number of groups, and K is the SR complexity.

In the worst case, every segment forms its own group, i.e., $d_g = d_s$. Further, the length of a segment could approach the trajectory length N_t , i.e., $N_s = O(N_t)$. This yields complexity:

$$O(d_s^2 \cdot p \cdot n \cdot N_s). \quad (15)$$

When $d_s \ll N_t$, this is effectively $O(p \cdot n \cdot N_t)$, i.e., linear in trajectory length.

In practice, the inner loop terminates early when a matching group is found, significantly reducing the average number of iterations and improving practical performance.

Model construction & overall complexity. The decision tree construction step has complexity $O(|\mathcal{L}| \log |\mathcal{L}|)$, where $|\mathcal{L}| = N_t$ is the size of the training set (Shalev-Shwartz and Ben-David, 2014). This step is typically dominated by the SR-based components.

The total complexity is dominated by the transition detection and grouping phases. As the phases are executed sequentially, the complexity sums up to:

$$O((d_s \cdot d_e + d_s \cdot d_g) \cdot K) = O((d_s \cdot d_e + d_s \cdot d_g) \cdot p \cdot n \cdot N_s) \quad (16)$$

$$= O(p \cdot n \cdot N_t \cdot d_s \cdot \max(d_e, d_g)). \quad (17)$$

With a suitable choice of l_{init} and l_{step} , this simplifies to $O(p \cdot n \cdot N_t)$, which is linear in the trajectory length.

6. Experiments

In this section, we present empirical results on two example systems. We analyze all steps, i.e., the trajectory segmentation, the segment grouping and the model construction individually, considering suitable metrics. The assessment of the final model construction also covers the prediction capabilities of the learned hybrid decision tree.

6.1. Example systems

We present empirical results on four example systems. The examples are selected to cover different aspects of the approach and include well-known benchmark systems as well as a real-world system.

6.1.1. Electrical circuits

We consider two versions of a passive electrical circuit, namely Circuit 1 and Circuit 2, where the goal is to learn an expression for the output current I_1 . The circuits are shown in Fig. 9. The first circuit in Fig. 9(a) is described by

$$I_1 = \begin{cases} \frac{U_1}{2 \cdot R}, & \text{switch in 1st position} \\ \frac{U_2}{2 \cdot R}, & \text{switch in 2nd position} \\ \frac{U_3}{2 \cdot R}, & \text{switch in 3rd position} \end{cases} \quad (18)$$

Depending on the position of the switch, one of the three sources is selected.

Circuit 2, shown in Fig. 9(b), contains three sources as well, but all of them are connected within the circuit. The system is described by the equation

$$I_1 = \frac{U_1 + U_2 + U_3}{4 \cdot R}. \quad (19)$$

Fig. 10 shows the three voltages and the output current I_1 over the discrete sampling points k for Circuit 2. This example is interesting to showcase the capabilities of HyDRA. The visual inspection of the current and voltage signals suggests that the system has multiple modes. Nevertheless, the system is described by a single equation. A dynamics-based approach identifies that a single dynamic mode is sufficient to cover the behavior of the system.

For SR, both circuits use the same set of basic operators which contains addition, subtraction, multiplication, and division operators. The three voltages U_1, U_2, U_3 and the value of the resistance R are given as variables for learning an expression for I_1 . Both circuits use $R = 10\Omega$ for all resistors.

6.1.2. Bouncing ball

The bouncing ball is a well-known example for system identification, which we use as a motivating example in Fig. 2. The system consists of a ball dropped from an initial height, bouncing on the ground. We use a simulation of the bouncing ball in Matlab (The MathWorks, 2025) with a sampling time of $\Delta t = 0.01$ s.

The system dynamics are governed by the matrix differential equation:

$$\begin{pmatrix} \dot{x} \\ \dot{v} \end{pmatrix} = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} x \\ v \end{pmatrix} + \begin{pmatrix} 0 \\ -g \end{pmatrix}, \quad (20)$$

where g denotes the gravitational acceleration and v is the velocity of the ball.

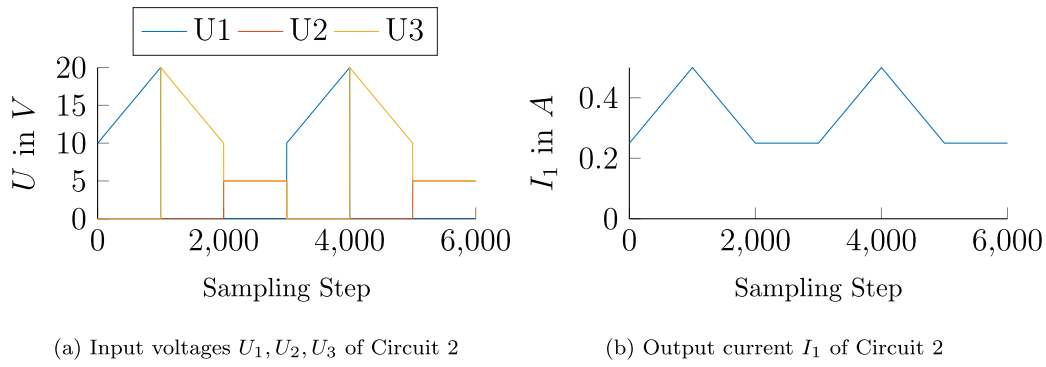


Fig. 10. Current and Voltages of Circuit 2.




Fig. 11. Illustration of the Two-tank System.

We observe the system with a fixed sampling rate and, thus, learn a one dimensional discrete approximation for the height of the ball given by

$$x(k) = 2 \cdot x(k-1) - x(k-2) - \frac{g}{2} \cdot \Delta t. \quad (21)$$

For SR, the features are the previous two height values $x(k-1)$ and $x(k-2)$ and the set of operators contains addition, subtraction, multiplication, and division.

6.1.3. Two-tank system

The Two-tank system (Bouamama et al., 2001) is a benchmark system for hybrid systems as shown in Fig. 11. It consists of two tanks, with a pump pumping water into the first tank and a valve V_b that can be opened to let water flow from the first to the second tank. The system has a controller to regulate the height of the water in the first tank. A differential equation describes the physical behavior of the water level in the first tank depending on the height of the second tank and the valve state. The system has three modes: two modes describe the behavior when the valve between the two tanks is open. In this case, the height of the water changes with the constant inflow and based on the pressure difference between the two tanks, which is described by the square root of the height difference. The third mode describes the behavior when the valve is closed, in which case the height of the first tank only changes based on the inflow of the pump. The following differential equation describes the system:

$$\dot{h}_1 = \begin{cases} \frac{Q_p - C_{vb} \cdot \sqrt{h_1 - h_2}}{A}, & \text{if } h_1 > h_2, V_b \text{ open} \\ \frac{Q_p + C_{vb} \cdot \sqrt{h_2 - h_1}}{A}, & \text{if } h_1 \leq h_2, V_b \text{ open} \\ \frac{Q_p}{A}, & \text{if } V_b \text{ closed,} \end{cases} \quad (22)$$

where h_1 and h_2 are the heights of the water in the first and second tank, respectively, \dot{h}_1 is the derivative, i.e., the change in the water

level, of the first tank. Q_p is the flow rate of the pump, C_{vb} is the valve conductance, and A is the cross-sectional area of the first tank.

The simulation involves noise and a realistic controller which applies a zero-order hold to the observed height of the tank. The noise is additive noise with a maximum amplitude of 10^{-6} for all sensors.

The operators for SR in the segmentation and grouping step are addition, subtraction, multiplication, and division as well as the square-root. The variables used for learning are the inflow Q_p , the height of the two tanks h_1 and h_2 as well as the constants C_b and A . For decision tree learning, the variable V_b is used as an additional feature. Additionally, we have a second scenario with this example where the term $h_{sq} = \sqrt{|h_1 - h_2|}$ is pre-calculated on the data and used as an additional variable to learn the derivative \dot{h}_1 . Practically, providing such pre-calculated terms can be interpreted as partial knowledge about the physical system and the operators and terms that are relevant for its behavior.

6.1.4. Power converter

The power converter (Zaupa et al., 2021) is a real-world system that has a controlled input voltage $v_s = \sigma V_g$ where V_g is a constant input and σ is a switching variable which can be either 1 or -1 . In addition to the voltage source, the circuit consists of a capacitor C , an inductance L , and a resistor R . The circuit can be either a parallel or a series configuration, as shown in Fig. 12. As presented in Zaupa et al. (2023), we use a transformed coordinate system to describe both configurations with the same equations. The following differential equation describes the system:

$$\dot{w} = \begin{pmatrix} 0 & \alpha \\ -\alpha & -\beta \end{pmatrix} w + \begin{pmatrix} 0 \\ \alpha \end{pmatrix} \sigma, \quad (23)$$

where $\alpha = (\sqrt{LC})^{-1}$ and $\beta = (RC)^{-1}$ (parallel case) or $\beta = R \cdot L^{-1}$ (serial case). The transformed coordinates $w = [w_1, w_2]^T$ are constructed

Table 1
Example systems.

System	l_{init}	l_{step}	n_{init}	n_{update}	τ	n_{group}	φ	ρ_G	ρ_S	ρ_G
Circuit 1	200	100	20	5	$1 \cdot 10^{-7}$	20	1.5	$1 \cdot 10^{-6}$	$1 \cdot 10^{-6}$	15
Circuit 2	200	100	20	5	$1 \cdot 10^{-7}$	50	1.5	$1 \cdot 10^{-6}$	$1 \cdot 10^{-3}$	15
Bouncing Ball	50	50	20	5	$1 \cdot 10^{-4}$	20	1.5	$1 \cdot 10^{-6}$	$1 \cdot 10^{-6}$	15
Power Converter	100	20	100	0.003	$1 \cdot 10^{-6}$	270	0.65	0.003	0.865	20
Two-tank Sys. <i>sub</i>	40	1	150	25	$1.28 \cdot 10^{-8}$	200	0.79	0.450	0.760	15
Two-tank Sys. <i>w/o sub</i>	50	10	200	5	$1 \cdot 10^{-6}$	100	1.25	0.850	0.110	40




Fig. 12. Series and Parallel Representation of the Power Converter (Zaupa et al., 2021).

from the quantities in Fig. 12 as

$$w_1 = \frac{v_c}{V_g} \quad \text{and} \quad w_2 = \frac{1}{V_g} \sqrt{\frac{L}{C}} \cdot i_c.$$

Eq. (23) describes the system dynamics in the transformed coordinate system, which is the behavior of a resonant circuit consisting of an inductance, a capacitor, and a resistor. The system has two modes, determined by the switching variable σ , which switches the direction of the input voltage V_g .

Our goal is to model the state variable \dot{w}_2 . The inductance, capacity, and resistor have the values $R = 400 \Omega$, $L = 8 \mu\text{H}$, $C = 10.5 \text{ nF}$, and $V_g = 20 \text{ V}$.

In the learning steps for segmentation and grouping, we use the operators addition, subtraction, multiplication, and division as well as the square-root as basic operators for SR. The to be learned derivative \dot{w}_2 is numerically calculated from the data. Feature variables are w_1 , w_2 , and the continuous time t of a sampling point. Constants are estimated by the learner. Different from the Two-tank example, the simulated data does not involve noise.

6.1.5. General setup

For both examples, SR uses a fitness metric with regularization and the mean-square error loss as the loss metric. The parameters for the learning algorithm for the examples and scenarios are shown in Table 1. For the Two-tank example and the power converter, the parametrization is done with support of Optuna (Akiba et al., 2019). We use objective functions Ω_{seg} and Ω_{group} for segmentation and grouping, respectively. The metric Ω_{seg} gives a mean deviation between the detected and ground truth transition points. The metric Ω_{group} gives a mean loss of the found groups compared to the ground truth groups. We determine the loss of the found groups weighted by their size, i.e., the number of samples in the group. The segmentation metric is defined as follows:

$$\Omega_{seg}(m, \mathcal{T}_{gt}, \mathcal{T}_f) = \frac{1}{|\mathcal{T}_f|} \left(\omega \cdot \|\mathcal{T}_f\| - \|\mathcal{T}_{gt}\| + \sum_{(i,j) \in m, j \neq \emptyset} |\mathcal{T}_{gt}(i) - \mathcal{T}_f(j)| \right), \quad (24)$$

where m is a map between ground truth transition points and their closest found transition points or \emptyset if no close transition point is found as illustrated by Fig. 13. $\mathcal{T}_{gt}, \mathcal{T}_f$ are vectors of ground truth and found transition points, respectively, while ω is a penalty factor for missed or additionally found transition points. The first term penalizes




Fig. 13. Illustration of mapping m between ground truth (red) and found (blue) transition points for the segmentation objective Ω_{seg} . The mapping for this example is $m = \{(1, 1), (2, 2), (3, \emptyset), (4, 3), (5, 5), (6, 6), (7, 8)\}$.

a divergence in the number of found and ground truth transition points by the factor ω . The second term penalizes a deviation of the found transition points from the ground truth transition points. The grouping metric Ω_{group} is defined as follows:

$$\Omega_{group}(G_{gt}, G_f) = \frac{1 + \|G_{gt}\| - \|G_f\|}{\sum_{g \in G_f} \text{size}(g)} \cdot \sum_{g \in G_f} \text{loss}(g) \cdot \text{size}(g), \quad (25)$$

where G_{gt} and G_f are the ground truth and found groups, respectively. Note, that grouping penalizes a divergence in the number of found and ground truth groups by the factor $1 + \|G_{gt}\| - \|G_f\|$.

Optuna also identifies parameter importances using the fANOVA approach (Hutter et al., 2014). These are useful for the parametrization of future examples. The parameter importances given in Fig. 14 show that for most examples, the initial window width l_{init} as well as the length for extending the window l_{step} are relevant for the segmentation while for the grouping especially the relaxation parameter φ is of relevance. A more detailed analysis of parameter influences can be found in Plambeck et al. (2024c).

The learning approach is implemented in python. The source code including the Optuna parametrization experiments is available on GitHub (Plambeck, 2024). All experiments are executed on the same personal computer (Intel i7-9750H 2.6 GHz, 16 GB RAM, charging).

6.2. Detection of transition points

In a first step, we analyze the detection properties of our approach. For this, we compare the detected transition points with the ground




Fig. 14. Parameter importances, Two-tank System with Substitution (green), Two-tank System without Substitution (red), Power Converter (blue).




Fig. 15. Results for the Detection of Segments and Grouping for the Power Converter.

Table 2

Detected transition points.

System	Detected points	Ground truth points	Ω_{seg}	Runtime
Circuit 1	5	5	0.0	448 s
Circuit 2	0	0	0.0	413 s
Bouncing Ball	0	0	0.0	161 s
Power Converter	4	4	2.8	887 s
Two-tank System <i>sub.</i>	28	28	1.15	7891 s
Two-tank System <i>w/o sub.</i>	27	28	4.64	1919 s

truth transition points. We use the metric of Eq. (24) which determines a mean deviation between the detected and ground truth transition points. Undetected or falsely detected transition points are penalized with a weight $\omega = 10$. Table 2 shows the results for the example systems. The columns provide the number of detected transition points, the number of ground truth transition points, the average deviation metric Ω_{seg} , and the runtime of the segmentation step. The results show that the approach detects the transition points in both systems with high precision. Nearly all transition points are detected perfectly. For the converter, the detected transition points are shown on the left of Fig. 15. For the two-tank example, one transition point was missed which is a transition happening after 8 samples at the beginning of the learning trajectory. These samples, thus, are likely to be included in a first segment.

6.3. Identification of modes

In the next step, we analyze the identified groups and expressions. Table 3 shows the results for the example systems. The columns show the detected expressions, thus, providing the number of identified modes, too. Furthermore, the table provides the Mean-Squared-Error

(MSE) loss of SR for the learned expressions on the segments of the respective group.

For all examples, the MSE indicates that the learned expressions are close to or exactly match the ground truth expressions. An exact matching is achieved for the two electrical circuits and the Two-tank system with substitution. Still, some differences to the ground truth expressions as introduced in the beginning of the section are present. For example, the bouncing ball example has an additional term $0.006/x(k-1)$ which is not present in the ground truth expression, but due to the small factor 0.006, this term is negligible for the value range of the variable $x(k-1)$. The power converter uses the time t in the second expression which is not used in the ground truth expression. The sampling steps are very small, i.e., t lies within the range of 10^{-5} and, thus, is negligible compared to the constant 0.01. For the converter, the grouping is correct, which can be seen on the right of Fig. 15 by the two colors indicating the two groups.

In general, the identified expressions of the examples involve the most dominant terms of the ground truth expression, but may not exactly represent the less dominant parts. This is visible for the Two-tank system, where the approach is able to identify the dominant influence of the pump defined by the term Q_p/A , while the influence of the water transfer between the two tanks is identified precisely only for the scenario with substitution of the term $\sqrt{|h_1 - h_2|}$. For the scenario without substitution, this is approximated by a linear influence of h_1 which is a useful approximation for the value range of the variables. For the scenario with substitution, we observe that both dynamics are accurately identified. Even though initially three groups are formed, the grouping step finds two of these groups to represent the same dynamics.

The runtime of both, segmentation and grouping steps, as given in Tables 2 and 4 is moderate. The runtime scales especially with the number of evolutions of the symbolic regression, while a single evolution scales with the amount of learning data. Further, the runtime

Table 3
Identified modes.

System	Learned expressions	MSE	
Circuit 1 <i>switched sources</i>	Group 1 –	$0.5 \cdot U_1$	$3.52 \cdot 10^{-17}$
	Group 2 –	$0.5 \cdot U_2$	0.0
	Group 3 –	$0.5 \cdot U_3$	$3.52 \cdot 10^{-17}$
Circuit 2 <i>connected sources</i>	Group 1 –	$0.25 \cdot (U_1 + U_2 + U_3)$	$1.65 \cdot 10^{-16}$
Bouncing Ball	Group 1 –	$2 \cdot x(k-1) - x(k-2) + 0.006/x(k-1)$	$2.11 \cdot 10^{-5}$
Power Converter	Group 1 –	$0.003 - 0.005 \cdot w_1$	$6.48 \cdot 10^{-7}$
	Group 2 –	$(t + 0.01) \cdot (w_1 + 0.11 \cdot w_2 + 1.14) \cdot (-0.47)$	$3.63 \cdot 10^{-8}$
Two-tank System <i>sub.</i>	Group 1 –	$(-Cv_{ub} \cdot \sqrt{ h_1 - h_2 + Q_p})/A$	$6.03 \cdot 10^{-8}$
	Group 2 –	Q_p/A	$1.41 \cdot 10^{-9}$
	Group 3 –	$(-Cv_b \cdot \sqrt{ h_1 - h_2 + Q_p})/A$	$1.39 \cdot 10^{-9}$
Two-tank System <i>w/o sub.</i>	Group 1 –	$Q_p/A - \sqrt{Q_p} \cdot h_1$	$5.03 \cdot 10^{-6}$
	Group 2 –	Q_p/A	$6.12 \cdot 10^{-6}$
	Group 3 –	$\sqrt{Cv_b} - A$	$3.08 \cdot 10^{-6}$

Table 4
Runtime of the grouping step.

System	Runtime
Circuit 1	342 s
Circuit 2	85 s
Bouncing Ball	27 s
Power Converter	988 s
Two-tank System <i>sub.</i>	3116 s
Two-tank System <i>w/o sub.</i>	5749 s

Table 5
Runtime of the model construction.

System	Runtime
Circuit 1	0.57 s
Circuit 2	0.14 s
Bouncing Ball	0.51 s
Power Converter	0.28 s
Two-tank System <i>sub.</i>	0.42 s
Two-tank System <i>w/o sub.</i>	0.68 s

depends on the number of detected segments and the number of found groups. Thus, the runtime of the grouping step for Circuit 2 and the Bouncing Ball is low, because only a single segment and group exists. The grouping step is often faster than the segmentation step, because during segmentation, the step width of window extensions is small compared to the average segment length. Thus, many iterations of the internal loop in Algorithm 1 are needed. The parameter l_{step} , thus, is a trade-off between runtime and detection accuracy for transition points.

6.4. Model construction & accuracy of predicted trajectories

In the last step, we construct a hybrid system in form of a hybrid decision tree. To evaluate the performance of the model, we use the model to predict the output on an evaluation trajectory and compare the prediction to the ground truth trajectory. The evaluation data consists of 1000 to 6000 samples, depending on the example. Table 5 shows the runtime of the model construction for all examples. We observe that the runtime is small as decision tree learning is a lightweight algorithm. In our implementation, we use the classification tree from the scikit-learn library (Pedregosa et al., 2011).

The prediction accuracy is shown in Table 6 in the column HyDRA, which provides the MSE on the evaluation set using the inference process described in Section 4 as well as the runtime of the model construction, i.e., decision tree learning step.

The MSE on the evaluation set is as expected a bit higher than the MSE on the learning set of the groups in the previous step. Nevertheless, the overall accuracy for all examples is good. The MSE of the Two-tank with substitution is better than the MSE for the Two-tank without substitution and, thus, shows that prior knowledge on relevant sub-dynamics supports the learning process and leads to more accurate models. This effect follows directly from the previous analysis of the identified expressions.

Further, Fig. 16 shows the ground truth and the predicted trajectories of the evaluation sets for the converter and the Two-tank example. For the converter, the trajectories match well. There is a single miss-prediction of a mode at around $k = 1000$. For all other samples, the mode is predicted correctly by the hybrid decision tree.

For the Two-tank example, several miss-predictions occur at the beginning of the trajectory. The reason, here, is that due to different initial values of the evaluation and the learning set, the initial phase of the evaluation set differs from the initial phase of the learning set and, thus, can be considered as unseen data during learning. At approximately $k = 150$ the orange trajectory has a discontinuity, realigning with the ground truth trajectory. The rest of the trajectory is predicted accurately.

6.5. Comparison to state-of-the-art

We compare HyDRA to the state-of-the-art approaches FaMoS (Plambeck et al., 2024a) and HAuLearn (Yang et al., 2022). Both approaches follow the scheme of Fig. 1. The transition detection of both approaches executes a peak detection on windowed trajectories. For segment grouping, FaMoS uses a grouping based on Dynamic Time Warping (DTW), while HAuLearn uses a grouping based on linear matrix inequalities, assuming that system dynamics follow linear differential equations. Consequently, the mode identification of both approaches is based on linear differential equations. Finally, model construction of FaMoS uses decision trees similar to our approach, while HAuLearn uses automata learning based on prefix tree acceptors (PTA).

Both approaches have a set of parameters that need to be tuned. The relevant parameters which are used for this evaluation are shown in Table 7. The parameter w is the window size for the peak detection. The parameter σ is the tolerance for linear matrix inequalities in the grouping step of HAuLearn. The parameters α , ρ_{min} , and ρ_{max} are used for the grouping step of FaMoS. They determine the factor, minimum, and maximum threshold, which is used to determine whether the DTW distance between two segments is small enough to be grouped together. For additional parameters, we base on the parameters presented in Plambeck et al. (2024a).

Table 6 shows the MSE of the prediction on the evaluation set for both approaches and our proposed approach. For the two circuits, FaMoS and HAuLearn are not applicable — indicated by a dash — because the systems are not described by linear differential equations. This shows a limitation of both approaches, which apply only to

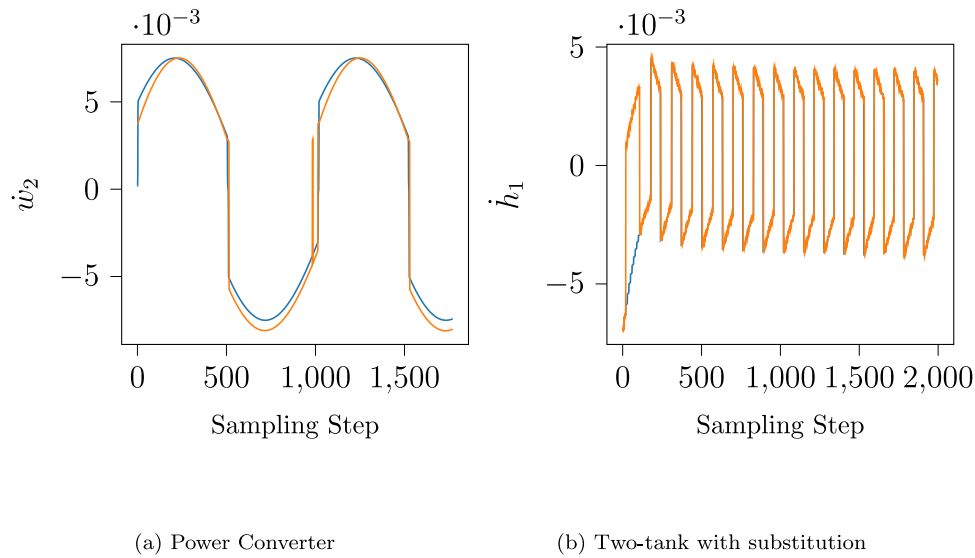


Fig. 16. Predicted and Ground Truth Evaluation Trajectories.

Table 6
Prediction accuracy and comparison of HyDRA to state-of-the-art.

System	MSE		
	HyDRA	FaMoS	HAutLearn
Circuit 1	$1.07 \cdot 10^{-10}$	–	–
Circuit 2	$2.95 \cdot 10^{-31}$	–	–
Bouncing Ball	$5.15 \cdot 10^{-5}$	1.09	$6.27 \cdot 10^{-1}$
Power Converter	$6.72 \cdot 10^{-7}$	$3.97 \cdot 10^{-6}$	$3.92 \cdot 10^{-6}$
Two-tank System <i>sub.</i>	$2.31 \cdot 10^{-6}$	$1.04 \cdot 10^{-7}$	$4.83 \cdot 10^{-6}$
Two-tank System <i>w/o sub.</i>	$8.58 \cdot 10^{-6}$	$1.38 \cdot 10^{-7}$	$1.07 \cdot 10^{-7}$

Table 7
Parameters of the state-of-the-art approaches.

Example	σ	w	α	ρ_{min}/ρ_{max}
Bouncing Ball	$2 \cdot 10^{-5}$	10	2.5	$10^{-3}/10$
Power Converter	$5 \cdot 10^{-3}$	10	2.5	$3 \cdot 10^{-2}/1$
Two-tank System <i>sub.</i>	$2 \cdot 10^{-5}$	30	4	$10^{-4}/10$
Two-tank System <i>w/o sub.</i>	$2 \cdot 10^{-4}$	30	5	$10^{-4}/10$

Table 8
Segmentation metric & Number of identified modes.

System	Identified modes			Ω_{seg}		
	FaM	HAut	HyDRA	FaM	HAut	HyDRA
Circuit 1	–	–	3	–	–	0.0
Circuit 2	–	–	1	–	–	0.0
Bouncing Ball	4	4	1	10.0	10.0	0.0
Power Converter	2	1	2	9.80	9.80	2.80
Two-tank System <i>sub.</i>	3	10	3	6.26	6.26	1.15
Two-tank System <i>w/o sub.</i>	2	3	3	3.06	3.06	4.64

systems whose dynamics are described by linear differential equations. For the power converter and the Two-tank systems, HyDRA achieves similar performance in terms of MSE as the state-of-the-art approaches. For the bouncing ball, HyDRA achieves a significantly better MSE than both approaches. This is possible, because HyDRA is able to identify the single mode of the bouncing ball, while both state-of-the-art approaches identify multiple modes.

Table 8 shows the number of identified modes and the segmentation metric Ω_{seg} for the state-of-the-art approaches and HyDRA for the example systems. Our approach achieves a perfect segmentation metric for the first three metrics and a lower segmentation metric for the power converter and the Two-tank system with substitution. Only

for the Two-tank system without substitution, HyDRA has a higher segmentation metric than the state-of-the-art approaches, which is the only example, where our approach misses a transition point.

HyDRA consistently finds a small number of modes for the system dynamics, while the state-of-the-art approaches might identify more modes than necessary. In all examples, the number of modes is determined by the parametrization of the grouping step. While FaMoS relies solely on segment similarity, HAutLearn additionally assesses the compatibility of segments based on linear differential equations. Our approach groups segments based on the underlying dynamics and applies to systems with linear and nonlinear dynamics.

7. Conclusion

In this paper, we present a novel approach to hybrid system identification using symbolic regression and decision trees. We show that symbolic regression can be used to identify the structure of hybrid systems from data. Symbolic regression iteratively adapts to observed dynamics in a trajectory segmentation step and a segment grouping step. The final model construction is done with decision tree learning resulting in a hybrid decision tree. The proposed method, thus, provides a holistic and straightforward approach for the identification of hybrid systems using symbolic regression that opens the scope to a new dynamics-based paradigm for learning models of hybrid systems.

We evaluate our method on four example systems, two of them with two versions: an electrical circuit with two versions, a bouncing ball, a two tank system where there are two different learning scenarios with and without inclusion of prior knowledge, and a power converter. The results show that our method can accurately identify the structure of hybrid systems from data and predict the behavior of the system on new data. Further, the output of the identification process is an explainable and interpretable model which makes the system understandable to other stakeholders. The model, thus, serves for prediction of future system behavior and can be used for control, fault detection, and other applications.

CRedit authorship contribution statement

Swantje Plambeck: Writing – Original Draft, Visualization, Validation, Software, Methodology, Conceptualization. **Maximilian Schmidt:** Writing – Review & Editing, Methodology, Conceptualization. **Audine**

Subias: Writing – Review & Editing, Project Administration, Methodology. **Louise Travé-Massuyès:** Writing – Review & Editing, Supervision, Project Administration, Methodology, Conceptualization. **Goerschwin Fey:** Writing – Review & Editing, Supervision, Project Administration, Methodology, Conceptualization.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Swantje Plambeck reports financial support was provided by Federal Ministry of Education and Research Bonn Office. Maximilian Schmidt reports a relationship with Federal Ministry of Education and Research Bonn Office that includes: funding grants and travel reimbursement. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This work was supported by a fellowship of the German Academic Exchange Service (DAAD) and the ECIU Universities. The research is partially funded by the federal Ministry of Research, Technology and Space (BMFTR) project AGenC no. 01IS22047A. It is also supported by ANITI through the French “Investing for the Future – P3IA” program under the Grant agreement no. ANR-19-P3IA-0004. Furthermore, we would like to thank Nicola Zaupa and Luca Zaccarian from LAAS CNRS for their support and comments on the power converter example.

During the preparation of this work the authors used GitHub Copilot in order to improve the writing style and identify alternative formulations. After using this tool, the authors reviewed and edited the content as needed and takes full responsibility for the content of the publication.

Data availability

The data and code used in the evaluation can be found on GitHub [Plambeck \(2024\)](#).

References

- Akiba, T., Sano, S., Yanase, T., Ohta, T., Koyama, M., 2019. Optuna: A next-generation hyperparameter optimization framework. *CoRR*. [arXiv:1907.10902](#).
- Alur, R., 2015. Principles of Cyber-Physical Systems. Technical Report, MIT Press, [https://ieeexplore.ieee.org/servlet/opac?bknnumber=7109352](#).
- Barbosa Roa, N., Travé-Massuyès, L., Grisales-Palacio, V.H., 2019. Dyclee: Dynamic clustering for tracking evolving environments. *Pattern Recognit.* 94, 162–186.
- Bouamama, B.O., Alaoui, R.M., Taillibert, P., Staroswiecki, M., 2001. Diagnosis of a two-tank system. Technical Report, Intern Report of CHEM-project, USTL.
- Branicky, M.S., 2005. Introduction to hybrid systems. In: *Handbook of Networked and Embedded Control Systems*. Birkhäuser, Boston, MA, pp. 91–116.
- Broløs, K.R., Machado, M.V., Cave, C., Kasak, J., Stentoft-Hansen, V., Batanero, V.G., Jelen, T., Wilstrup, C., 2021. An approach to symbolic regression using feyn. *CoRR*. [arXiv:2104.05417](#).
- Cranmer, M., 2023. Interpretable machine learning for science with PySR and SymbolicRegression.jl. *CoRR*. [arXiv:2305.01582](#).
- Feldt, R., Nordin, P., 2000. Using factorial experiments to evaluate the effect of genetic programming parameters. In: *Genetic Programming*. pp. 271–282.
- Fong, K.S., Motani, M., 2024. Symbolic regression enhanced decision trees for classification tasks. *Proc. AAAI Conf. Artif. Intell.* 38 (11), 12033–12042. [http://dx.doi.org/10.1609/aaai.v38i11.29091](#).
- Gaucel, S., Keijzer, M., Lutton, E., Tonda, A., 2014. Learning dynamical systems using standard symbolic regression. In: Nicolau, M., Krawiec, K., Heywood, M.I., Castelli, M., Garcia-Sánchez, P., Merelo, J.J., Rivas Santos, V.M., Sim, K. (Eds.), *Genetic Programming*. Springer Berlin Heidelberg, pp. 25–36.
- Goupil, L., Travé-Massuyès, L., Chantry, E., Kohler, T., Delautier, S., 2024. Tree based diagnosis enhanced with meta knowledge applied to dynamic systems. *IFAC PapersOnLine* 58 (4), 1–6.
- Hutter, F., Hoos, H., Leyton-Brown, K., 2014. An efficient approach for assessing hyperparameter importance. In: *International Conf. on Machine Learning*. ICML, pp. 754–762. [http://dx.doi.org/10.5555/3044805.3044891](#).

- Jimenez, F., Kaminska, J., Lucena-Sanchez, E., Palma, J., Sciavicco, G., 2019. Multi-Objective Evolutionary Optimization for Time Series Lag Regression. In: *International Conf. on Time Series and Forecasting*. ITISE, pp. 373–384. [https://iris.unife.it/bitstream/11392/2408030/1/paper%20143.pdf](#).
- Koza, J.R., 1992. On the programming of computers by means of natural selection. In: *Genetic programming*, MIT Press, Cambridge, Mass. u.a., [https://mitpress.mit.edu/9780262527910/genetic-programming/](#).
- Kronberger, G., Kammerer, L., Kommenda, M., 2021. Identification of dynamical systems using symbolic regression. *CoRR*. [arXiv:2107.06131](#).
- Langdon, W.B., 2022. Genetic programming convergence. In: *Genetic and Evolutionary Computation Conference Companion*. pp. 27–28. [http://dx.doi.org/10.1145/3520304.3534063](#).
- Loh, W.-Y., 2011. Classification and regression trees. *WIREs Data Min. Knowl. Discov.* (ISSN: 1942-4787) 1 (1), 14–23. [http://dx.doi.org/10.1002/widm.8](#).
- Loveard, T., Ciesielski, V., 2002. *Genetic programming for classification: An analysis of convergence behaviour*. *Lecture Notes in Artificial Intelligence* 2557, 309–320.
- Lunze, J., Lamnabhi-Lagarigue, F., 2009. *Handbook of Hybrid Systems Control: Theory, Tools, Applications*. Cambridge University Press, Cambridge, [https://doi.org/10.1145/3520304.3534063](#).
- Ly, D.L., Lipson, H., 2012. Learning symbolic representations of hybrid dynamical systems. *J. Mach. Learn. Res.* 13 (115), 3585–3618. [http://dx.doi.org/10.5555/2503308.2503356](#).
- Muthyala, M.R., Sorourifar, F., Peng, Y., Paulson, J.A., 2025. Symantic: An efficient symbolic regression method for interpretable and parsimonious model discovery in science and beyond. *Ind. Eng. Chem. Res.* 64 (6), 3354–3369. [http://dx.doi.org/10.1021/acs.iecr.4c03503](#).
- Niggemann, O., Stein, B., Vodencarevic, A., Maier, A., Kleine Büning, H., 2012. Learning behavior models for hybrid timed systems. *AAAI Conf. Artif. Intell.* 26 (1), 1083–1090. [http://dx.doi.org/10.1609/aaai.v26i1.8296](#).
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, E., 2011. Scikit-learn: Machine learning in python. *J. Mach. Learn. Res.* 12, 2825–2830. [http://dx.doi.org/10.5555/1953048.2078195](#).
- Petersen, B.K., Larma, M.L., Mundhenk, T.N., Santiago, C.P., Kim, S.K., Kim, J.T., 2021. Deep symbolic regression: Recovering mathematical expressions from data via risk-seeking policy gradients. In: *International Conf. on Learning Representations*. pp. 1–26. [https://openreview.net/pdf?id=m5Qsh0kBQG](#).
- Plambeck, S., 2024. Symbolic regression 4 hybrid systems. URL [https://github.com/TUHH-IES/SymbolicRegression4HA](#).
- Plambeck, S., Bracht, A., Hranisavljevic, N., Fey, G., 2024a. Famos – fast model learning for hybrid cyber-physical systems using decision trees. In: *International Conf. on Hybrid Systems: Computation and Control*. Association for Computing Machinery, pp. 1–10. [http://dx.doi.org/10.1145/3641513.3650131](#).
- Plambeck, S., Schammer, L., Fey, G., 2022. On the viability of decision trees for learning models of systems. In: *Asia and South Pacific Design Automation Conf. (ASP-DAC)*. pp. 696–701. [http://dx.doi.org/10.1109/ASP-DAC52403.2022.9712579](#).
- Plambeck, S., Schmidt, M., Fey, G., Subias, A., Travé-Massuyès, L., 2024b. Dynamics-based identification of hybrid systems using symbolic regression. In: *EuroMicro Conf. on Software Engineering and Advanced Applications*. SEAA, pp. 64–71. [http://dx.doi.org/10.1109/SEAA64295.2024.00019](#).
- Plambeck, S., Schmidt, M., Subias, A., Travé-Massuyès, L., Fey, G., 2024c. Usability of Symbolic Regression for Hybrid System Identification - System Classes and Parameters. In: *International Conf. on Principles of Diagnosis and Resilient Systems*. DX, pp. 30:1–30:14. [http://dx.doi.org/10.4230/OASiCS.DX.2024.30](#).
- Poli, R., Langdon, W.B., McPhee, N.F., 2008. A field guide to genetic programming. 10, Springer, [http://www.gp-field-guide.org.uk](#).
- Real, E., Liang, C., So, D., Le, Q., 2020. AutoML-zero: Evolving machine learning algorithms from scratch. In: *International Conf. on Machine Learning*, vol. 119, Pmlr, pp. 8007–8019. [http://dx.doi.org/10.5555/3524938.3525680](#).
- Schmidt, M., Lipson, H., 2009. Distilling free-form natural laws from experimental data. *Science* 324 (5923), 81–85. [http://dx.doi.org/10.1126/science.1165893](#).
- Shalev-Shwartz, S., Ben-David, S., 2014. *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press, Cambridge, [http://dx.doi.org/10.1017/CBO9781107298019](#).
- Tappler, M., Muškardin, E., Aichernig, B.K., Könighofer, B., 2023. Learning environment models with continuous stochastic dynamics. *CoRR*. [arXiv:2306.17204](#).
- The MathWorks, I., 2025. Model a bouncing ball in continuous time. URL [https://de.mathworks.com/help/stateflow/ug/modeling-a-bouncing-ball-in-continuous-time.html](#).
- Wang, G., Wang, E., Li, Z., Zhou, J., Sun, Z., 2024. Exploring the mathematical equations behind the materials science data using interpretable symbolic regression. *Interdiscip. Mater.* 3 (5), 637–657. [http://dx.doi.org/10.1002/idm2.12180](#).
- Yang, X., Beg, O.A., Kenigsberg, M., Johnson, T.T., 2022. A framework for identification and validation of affine hybrid automata from input-output traces. *ACM Trans. Cyber Phys. Syst.* 6 (2), 1–24. [http://dx.doi.org/10.1145/3470455](#).
- Zaupa, N., Martínez-Salamero, L., Olalla, C., Zaccarian, L., 2021. Results on hybrid control of self-oscillating resonant converters. *ADHS*, In: *IFAC Conf. on Analysis and Design of Hybrid Systems*, vol. 54, pp. 211–216.
- Zaupa, N., Martínez-Salamero, L., Olalla, C., Zaccarian, L., 2023. Hybrid control of self-oscillating resonant converters. *IEEE Trans. Control Syst. Technol.* 31 (2), 881–888. [http://dx.doi.org/10.1109/tcst.2022.3179948](#).