

Towards Interpretable-AI Policies using Evolutionary Nonlinear Decision Trees for Discrete Action Systems

Yashesh Dhebar, Kalyanmoy Deb, Subramanya Nagesh Rao, Ling Zhu, and Dimitar Filev

COIN Report Number 2021001

Abstract—Black-box artificial intelligence (AI) induction methods such as deep reinforcement learning (DRL) are increasingly being used to find optimal policies for a given control task. Although policies represented using a black-box AI are capable of efficiently executing the underlying control task and achieving optimal closed-loop performance – controlling the agent from initial time step until the successful termination of an episode, the developed control rules are often complex and neither *interpretable* nor *explainable*. In this paper, we use a recently proposed nonlinear decision-tree (NLDT) approach to find a hierarchical set of control rules in an attempt to maximize the open-loop performance for approximating and explaining the pre-trained black-box DRL (oracle) agent using the labelled state-action dataset. Recent advances in nonlinear optimization approaches using evolutionary computation facilitates finding a hierarchical set of nonlinear control rules as a function of state variables using a computationally fast bilevel optimization procedure at each node of the proposed NLDT. Additionally, we propose a re-optimization procedure for enhancing closed-loop performance of an already derived NLDT. We evaluate our proposed methodologies (open and closed-loop NLDTs) on different control problems having multiple discrete actions. In all these problems our proposed approach is able to find relatively simple and interpretable rules involving one to four non-linear terms per rule, while simultaneously achieving on par closed-loop performance when compared to a trained black-box DRL agent. A post-processing approach for simplifying the NLDT is also suggested. The obtained results are inspiring as they suggest the replacement of complicated black-box DRL policies involving thousands of parameters (making them non-interpretable) with relatively simple interpretable policies. Results are encouraging and motivating to pursue further applications of proposed approach in solving more complex control tasks.

Index Terms—Reinforcement Learning, Interpretable, Bilevel, Nonlinear Decision Tree.

I. INTRODUCTION

Control system problems are increasingly being solved by using modern reinforcement learning (RL) and other machine learning (ML) methods to find an autonomous agent (or controller) to provide an optimal action A_t for every state

variable combination S_t in a given environment at every time step t . Execution of the output *action* A_t takes the object to the next *state* S_{t+1} in the environment. This process is repeated until a termination criteria is met. The mapping between input *state* S_t and output *action* A_t in most real-world problems is usually complex, and is, therefore, captured through an artificial intelligence (AI) method. In the RL literature, this mapping is referred to as *policy*, and sufficient literature exists in efficient training of these RL *policies* [1], [2], [3], [4], [5], [6], [7]. While these methods are efficient at training the AI policies for a given control system task, the developed AI policies, often represented using complicated artificial neural networks, are usually complex and non-interpretable.

Interpretability of AI policies is important to a human mind due to several reasons: (i) *interpretable policies* provide better insight and understanding of the working principles of the derived control-logic, (ii) they can be easily deployed with a low fidelity hardware, (iii) they may also allow an easier way to extend the control policies for more complex versions of the problem. While defining interpretability is a subjective matter, a number of past efforts have attempted to find interpretable AI policies with limited success [8], [9]. In this work, our aim is to find policies which are *relatively* interpretable and simple as compared to the black-box AI counterparts such as DNN. For this purpose, we propose to use a non-linear rule with limited complexity at each node of the non-linear decision tree (NLDT). At each time step, the state variable vector is passed through the derived NLDT's root rule and the control flows through the decision tree to arrive at a specific leaf node, describing one of the allowable discrete actions. Well-known difficulties of learning non-linear rules are alleviated by using a bilevel optimization approach in which the upper level handles the structure of a rule and the lower level determines the coefficients of the terms of the rule. Moreover, a final efficient search of all coefficients of rules in the NLDT is employed to arrive at a reliable rule-base.

In the remainder of this paper, we first present the main motivation behind finding interpretable policies in Section II. In Section III, a brief overview on existing works related to interpretable AI policies is provided. In Section IV, we review a recently proposed nonlinear decision tree (NLDT) approach in the context of arriving at relatively interpretable AI policies. The overall open-loop and closed-loop NLDT policy generation methods are described in Section V. Results on four control system problems are presented in Section VI.

Yashesh Dhebar and Kalyanmoy Deb are with Michigan State University, East Lansing, MI, 48824 USA

Subramanya Nagesh Rao and Dimitar Filev are with Ford Motor Company, Detroit, USA

Ling Zhu is with Microsoft at Mountain View, California, USA

A code implementing nonlinear decision tree development is available from <https://github.com/yddhebar/NLDT>

Accepted for Publication at IEEE Transactions on Cybernetics, Print ISSN: 2168-2267, Online ISSN: 2168-2275, Digital Object Identifier: 10.1109/TCYB.2022.3180664.

Next, a new benchmark problem is proposed to conduct a scale-up study of our algorithm in Section VII. Finally, conclusions and future studies are presented in Section VIII. Supplementary document provides further details.

II. MOTIVATION FOR THE STUDY

Various data analysis tasks, such as classification, controller design, regression, image processing, etc., are increasingly being solved using artificial intelligence (AI) methods. The suggested methods have been demonstrated to solve complex data analysis tasks without much change in their usual frameworks. However with more such studies over the past few decades [9], these AI generation methods are faced with a huge challenge - *interpreting and understanding the inner working of AI*. Achieving a high-accuracy solution does not necessarily satisfy a curious domain expert, particularly if the solution is not interpretable or explainable. A technique (whether AI-based or otherwise) to handle data well is no more enough. Researchers now demand an explanation of why and how they work [10], [11].

Consider the MountainCar control system problem (see Supplementary document for details), which has been extensively studied using various AI methods [12], [13], [14]. The problem has two state variables (position x_t along x -axis and velocity v_t along positive x -axis) describing the *state* of the car at time instant t . Based on the state vector $S_t = (x_t, v_t)$, the policy $\pi(S)$ selects one of the three actions A_t : 1) *decelerate* ($A_t = 0$) along positive x -axis with a pre-defined value $-a$, 2) *do nothing* ($A_t = 1$), or 3) *accelerate* ($A_t = 2$) with $+a$ in positive x -axis direction. The goal of this control problem is to drive the under-powered car¹ over the right hump in a maximum of 200 time steps, starting anywhere at the trough of the landscape. Physical laws of motion are applied and a policy $\pi(S)$ is trained to solve the problem. The RL method trains a blackbox policy $\pi_{oracle}(S)$ which selects an action $A_t \in [0, 1, 2]$ for a given input state $S_t = (x_t, v_t) \in \mathbb{R}^2$. Figure 1a shows the state-action combinations obtained from

the action A_t suggested by the oracle policy π_{oracle} ($A_t = 0$: blue, $A_t = 1$: orange, and $A_t = 2$: green). If a user is now interested in understanding how the policy π_{oracle} chooses a correct A_t for a given S_t , one way to achieve this would be through an interpretable policy function $\pi_{int}(S_t)$ as follows:

$$\pi_{int}(S_t) = \begin{cases} 0, & \text{if } \phi_0(S_t) \text{ is true,} \\ 1, & \text{if } \phi_1(S_t) \text{ is true,} \\ 2, & \text{if } \phi_2(S_t) \text{ is true,} \end{cases} \quad (1)$$

where $\phi_i(S_t) : \mathbb{R}^2 \rightarrow \{0, 1\}$ is a Boolean function which partitions the state space $\mathbb{S}$ into two sub-domains based on its output value. For a given state S_t , exactly one of $\phi_i(S_t)$ is *true*, thereby making the policy π_{int} deterministic. If we re-look at Figure 1a we notice that the three actions are quite mixed at the bottom part of the x - v plot (state space). Thus, the partitioning Boolean functions ϕ_i need to be quite complex in order to have $\phi_0(S_t) = \text{true}$ for all *blue* points, $\phi_1(S_t) = \text{true}$ for all *orange* points and $\phi_2(S_t) = \text{true}$ for all *green* points.

What we address in this study is an attempt to find an *approximated* policy function $\pi_{int}(S_t)$ which may not replicate and explain all 100% time instance data corresponding to the oracle black-box policy $\pi_{oracle}(S_t)$ (Figure 1a), but it is fairly interpretable to explain close to 100% data. Consider the state-action plot in Figure 1b, which is generated with a *relatively* simple and *relatively* interpretable policy $\pi_{int}(S_t) = \{i | \phi_i(S_t) \text{ is true}, i = 1, 2, 3\}$ obtained by our proposed procedure as shown below:

$$\begin{aligned} \phi_0(S_t) &= \neg\psi_1(S_t), \\ \phi_1(S_t) &= (\psi_1(S_t) \wedge \neg\psi_2(S_t)), \\ \phi_2(S_t) &= (\psi_1(S_t) \wedge \psi_2(S_t)), \end{aligned} \quad (2)$$

where $\psi_1(S_t) = |0.96 - 0.63/\hat{x}_t^2 + 0.28/\hat{v}_t - 0.22\hat{x}_t\hat{v}_t| \leq 0.36$, and $\psi_2(S_t) = |1.39 - 0.28\hat{x}_t^2 - 0.30\hat{v}_t^2| \leq 0.53$ represent black and orange boundaries, respectively. Here, $\hat{x}_t$ and $\hat{v}_t$ are normalized state variables (see Supplementary document for details). The action A_t predicted using the above policy does not match the output of π_{oracle} at some states (about 8.1%), but from our experiments we observe that it is still able to drive the mountain-car to the destination goal located on the right hill in 99.8% episodes.

Importantly, the π_{int} policy is relatively simpler (involving two simple non-linear functions ψ_1 and ψ_2) than the corresponding black-box policy π_{oracle} and is amenable to an easier understanding of the relationships between x_t and v_t to make a near perfect control. Since the explanation process used the data from π_{oracle} as the universal truth (Figure 1a), the derived relationships will also provide an explanation of the working of the black-box policy π_{oracle} . A more gross approximation to Figure 1a by more simplified relationships (ϕ_i) may reduce the overall *open-loop* accuracy of matching the output of π_{oracle} . Hence, a balance between a good interpretability and a high *open-loop* accuracy in searching for Boolean functions $\phi_i(S_t)$ becomes an important matter for such an interpretable AI-policy development study.

In this paper, we focus on developing a search procedure for arriving at the ψ -functions (see Eq. 2) and their combinations for different discrete action systems. The structure of the policy $\pi_{int}(S_t)$ shown in Eq. 1 resembles a decision tree




Fig. 1: State-action combinations for MountainCar prob.

92 independent successful trajectories (amounting to total of 10,000 time steps) leading to achieving the goal using a pre-trained deterministic black-box policy π_{oracle} . The x -location of the car and its velocity can be obtained from a *point* on the 2D plot. The color of the point $S_t = (x_t, v_t)$ indicates

¹the car does not have enough fuel to directly climb the mountain and reach the destination.

(DT). However, *unlike* a standard DT, it involves a nonlinear function at every non-leaf node. This requires an efficient nonlinear optimization method to arrive at reasonably succinct and accurate non-linear functionals. The procedure we propose here is generic and is independent of the AI method used to develop the black-box policy π_{oracle} .

III. RELATED PAST STUDIES

In [15], an *interpretable orchestrator* is developed to choose from two RL-policies: π_C for maximizing reward and π_R for maximizing an *ethical* consideration. The orchestrator is dependent on only one of the state-variables and despite it being interpretable, the policies: π_C and π_R , are still black-box and convoluted. [16] constructs a set of interpretable index based policies and uses multi-arm bandit procedure to select a high performing index based policy. The search space of interpretable policies is much smaller and the procedure suggested for finding an interpretable policy is computationally heavy, taking about hours to several days of computational time on simple control problems. In [17], genetic programming (GP) is used to obtain interpretable policies on control tasks involving continuous actions space through model-based policy learning. However the *interpretability* was not captured in the design of the fitness function and a large archive was created passively to store every policy for each complexity encountered during the evolutionary search. A linear decision tree (DT) based model is used in [18] to approximate the Q-values of trained neural network. In that work, the split in DT occurs based on only one feature, and at each terminal node, the Q-function is fitted using a linear model on all features. [19] uses a *program sketch* S to define the domain of interpretable policies e . Interpretable policies are found using a trained black-box oracle e_N as a reference by first conducting a local search in the sketch space S to mimic the behaviour of the oracle e_N and then fine-tuning the policy parameters through online Bayesian optimization. The bias towards generating interpretable programs is done through controlled initialization and local search rather than explicitly capturing *interpretability* as one of the fitness measure. Particle swarm optimization [20] is used to generate interpretable fuzzy rule set in [21] and is demonstrated on classic control problems involving continuous actions. Works on DT [22] based policies through imitation learning has been carried out in [23]. [24] extends this to utilize Q-values and eventually render DT policies involving $< 1,000$ nodes on some toy games and CartPole environment with an ultimate aim to have the induced policies verifiable. [25] used axis-aligned DTs to develop interpretable models for black-box classifiers and RL-policies. They first derive a distribution function $\mathcal{P}$ by fitting the training data through axis-aligned Gaussian distributions. $\mathcal{P}$ is then used to compute the loss function for splitting the data in the DT. [26] attempts to generate interpretable DTs from an ensemble using a genetic algorithm. In [27], regression trees are derived using classical methods such as CART [22] and Kd-tree [28] to model Q-function through supervised training on batch of experiences and comparative study is made with ensemble techniques. In [29], a gradient

based approach is developed to train the DT of pre-fixed topology involving linear split rules. These rules are later simplified to allow only one feature per split node and resulting DTs are pruned to generate simplified rule-set.

While the above methods attempt to generate an interpretable policy, the search process does not use *complexity* of policy in the objective function, instead, they rely on initializing the search with certain interpretable policies. Studies on DT based interpretable policies so far mostly involves either axis-parallel trees (where split-rules are of type: $x_i \leq \tau_i$) or oblique trees (where split-rules are linear functions of state-variables). In some approaches, the topology of DT is fixed. Our approach is more flexible in deriving an interpretable policy, since it grows the hierarchical set of interpretable rules naturally (i.e. topology of DT is not fixed or pre-specified) and the inherent split-rules of DT can assume a non-linear structure. The problem of finding non-linear rule-set to solve a given control task is a challenging problem involving efficient non-linear optimization algorithms (and one of the key reasons why previous approaches probably did not address it). The complex task of arriving at non-linear set of control logic and their hierarchical connection is the focus of this paper. In our approach described below, we build an efficient search algorithm to directly find relatively interpretable policies as compared to the black-box policies represented using DNN (or tile encoding [12]) using recent advances in nonlinear optimization.

IV. NONLINEAR DECISION TREE (NLDT) APPROACH

In this study, we use a direct mathematical rule generation approach (presented in Eq. 2) based on nonlinear decision tree (NLDT) [30]. While the overall goal of this work is to induce an interpretable policy π_{int} which can serve as a controller, the first step (discussed at length in section V) involves induction of π_{int} to approximate and explain the behavior of a pre-trained black-box policy π_{oracle} . This can be achieved by deriving π_{int} to match the labelled state-action data which is generated using π_{oracle} .

Decision trees are considered a popular choice due to their interpretability aspects. They are intuitive and each decision can be easily interpreted. However, in a general scenario, regular decision trees often have complicated topology since the rules at each conditional node can assume only axis parallel structure $x_i \leq \tau$ to make a split. On the other end, single rule based classifiers like support vector machines (SVMs) have just one rule which is complicated and highly nonlinear. Keeping these two extremes in mind, we develop a nonlinear decision tree (NLDT) framework where each conditional node can assume a nonlinear functional form while the tree is allowed to grow by recursively splitting the data in conditional nodes, similar to the procedure used to induce regular decision trees. In our case of replicating a policy π_{oracle} , a conditional node in NLDT represents a nonlinear control logic and terminal leaf nodes indicate an *action*. This is schematically shown in Figure 2. A NLDT, which can be viewed as an assembly of conditional nodes, represents an overall controller/policy π_{int} .

In the binary-split NLDT, used in this study, a conditional node is allowed to have exactly two splits as shown in Figure 2. The non-linear split-rule $f(\mathbf{x})$ at each conditional node is




Fig. 2: Schematic of a binary-split NLDT. State variables at each time instant are supplied to the root node and the information flows through the NLDT before reaching one of the leaf nodes, which indicates the respective action for the state variable vector.

expressed as a weighted sum of power-laws:

$$f(\mathbf{x}) = \begin{cases} \sum_{i=1}^p w_i B_i + \theta_1, & \text{if } m = 0, \\ |\sum_{i=1}^p w_i B_i + \theta_1| - |\theta_2|, & \text{if } m = 1, \end{cases} \quad (3)$$

where power-laws B_i are given as $B_i = \prod_{j=1}^d x_j^{b_{ij}}$ and m indicates if an absolute operator should be present in the rule or not. In Section V-A, we discuss procedures to derive values of exponents b_{ij} , weights w_i , and biases θ_i .

V. OVERALL APPROACH

The overall approach is illustrated in Figure 3. First, a dedicated black-box policy π_{oracle} is trained from the actual environment/physics of the problem. This aspect is not the focus of this paper. Next, the trained policy π_{oracle} (Block 1 in the figure) is used to generate labelled training and testing datasets of *state-action* pairs from different time steps. We generate two types of training datasets: *Regular* – as they are recorded from multiple episodes², and *Balanced* – selected from multiple episodes to have almost equal number of *states* for each action, where an *episode* is a complete simulation run wherein an *AI policy* is used to control the object over multiple time steps. Third, the labelled training dataset (Block 2) is used to find the NLDT (Block 3) using a recursive bilevel evolutionary algorithm described in Section V-A. We call this an open-loop NLDT (or, NLDT_{OL}) since it is derived from a labelled state-action dataset generated from π_{oracle} and doesn't involve any direct interaction with the control environment. Use of a labelled state-action data in a supervised manner allows for a faster search of NLDT even with a large dataset as compared to constructing the NLDT using traditional reinforcement learning which involves interaction with the environment [19]. Next, in an effort to make the overall NLDT interpretable while simultaneously ensuring better closed-loop performance, we prune the NLDT by taking only the top part of NLDT_{OL} (we call NLDT_{OL}^(P) in Block 4) and *re-optimize* all non-linear rules within it for the weights and biases using an efficient evolutionary optimization procedure to obtain final NLDT* (Block 5). The re-optimization is done here with

closed-loop objectives, such as the cumulative reward function or closed-loop completion rate. We briefly discuss the open-loop training procedure of inducing NLDT_{OL} and the closed-loop training procedure to generate NLDT* in next sections.

A. Open-loop Training

A labelled state-action dataset is first created using a pre-trained black-box policy π_{oracle} . Since we are dealing with discrete-action control problems, the underlying imitation task of replicating the behavior of π_{oracle} using the labelled state-action data translates to a *classification* problem. We train NLDT discussed in Section IV to fit the state-action data through supervised learning. Nonlinear split-rule $f(\mathbf{x})$ at each conditional node (Figure 2 and Eq. 3) is derived using a dedicated bilevel optimization algorithm [31], where the upper level searches the *template* of the non-linear rule and the corresponding lower level focuses at estimating optimal values of *weights/coefficients* for optimal split of data present in the conditional node. The optimization formulation for deriving a non-linear split-rule $f(\mathbf{x})$ (Eq. 3) at a given conditional node is given below:

$$\begin{aligned} & \text{Minimize } F_U(\mathbf{B}, m, \mathbf{w}^*, \boldsymbol{\theta}^*), \\ & \text{subject to } (\mathbf{w}^*, \boldsymbol{\theta}^*) \in \operatorname{argmin} \{F_L(\mathbf{w}, \boldsymbol{\theta}) |_{(\mathbf{B}, m)} \mid \\ & \quad F_L(\mathbf{w}, \boldsymbol{\theta}) |_{(\mathbf{B}, m)} \leq \tau_I, \\ & \quad -1 \leq w_i \leq 1, \forall i, \boldsymbol{\theta} \in [-1, 1]^{m+1} \}, \\ & \quad m \in \{0, 1\}, b_{ij} \in \mathcal{Z}, \end{aligned} \quad (4)$$

where $\mathcal{Z}$ is a set of exponents allowed to limit the complexity of the derived rule structure. In this study, we use $\mathcal{Z} = \{-3, -2, -1, 0, 1, 2, 3\}$. The objective F_U quantifies the complexity of the non-linear rule by enumerating the number of terms present in the equation of rule $f(\mathbf{x})$ as shown below:

$$F_U(\mathbf{B}, m, \mathbf{w}^*, \boldsymbol{\theta}^*) = \sum_{i=1}^p \sum_{j=1}^d g(b_{ij}), \quad (5)$$

where $g(\alpha) = 1$, if $\alpha \neq 0$, zero otherwise. m indicates the presence or absence of a modulus operator and $\mathbf{w}$ and $\boldsymbol{\theta}$ encodes rule weights w_i and biases θ_i respectively. The lower level objective function F_L quantifies the net impurity of child nodes resulting from the split. Impurity I of a node P is computed using a Gini-score: $\text{Gini}(P) = 1 - \sum_i \left(\frac{N_i}{N}\right)^2$, where N is the total number of points present in the node and N_i represents number of points belonging to class i . Datapoints present in node P gets distributed into two non-overlapping subsets based on their split function value. Datapoints with $f(\mathbf{x}) \leq 0$ go to the left child node L and rest go to the right child node R . The lower level objective function F_L which quantifies the quality of this split is then given by

$$F_L(\mathbf{w}, \boldsymbol{\theta}) |_{(\mathbf{B}, m)} = \left(\frac{N_L}{N_P} \text{Gini}(L) + \frac{N_R}{N_P} \text{Gini}(R) \right)_{(\mathbf{w}, \boldsymbol{\theta}, \mathbf{B}, m)}. \quad (6)$$

The τ_I parameter in Eq. 4 represents maximum allowable net-impurity (Eq. 6) of child nodes. The resulting child nodes obtained after the split undergo another split and the process continues until one of the termination criteria is met.

More details regarding the bilevel-optimization algorithm [31] to derive split-rule $f_i(\mathbf{x})$ at i -th conditional node in NLDT can be found in [30].

²an *episode* is one simulation run.




Fig. 3: A schematic of the proposed overall approach.

After the entire NLDT is found, in this study, a pruning and tree simplification strategy (see Supplementary Document for more details) is applied to reduce the size of NLDT in an effort to improve on the interpretability of the overall rule-set. This entire process of inducing NLDT from the labelled state-action data results into the open-loop NLDT – $NLDT_{OL}$. $NLDT_{OL}$ can then be used to explain the behavior of the oracle policy π_{oracle} . We will see in Section VI that despite being not 100% accurate in imitating π_{oracle} , $NLDT_{OL}$ manages to achieve respectable closed-loop performance with 100% completion rate and a high cumulative reward value. Next, we discuss the closed-loop training procedure to obtain $NLDT^*$.

B. Closed-loop Training

The intention behind closed-loop training is to enhance the closed-loop performance of NLDT. It will be discussed in Section VI that while closed-loop performance of $NLDT_{OL}$ is at par with π_{oracle} on control tasks involving two to three discrete actions, like CartPole and MountainCar, the $NLDT_{OL}$ struggles to autonomously control the agent for control problems involving more states and actions (eg. LunarLander). In closed-loop training, we fine-tune and re-optimize the weights $\mathbf{W}$ and biases Θ of an entire $NLDT_{OL}$ (or pruned $NLDT_{OL}$, i.e. $NLDT_{OL}^{(P)}$ – block 4 in Figure 3) to maximize its closed-loop fitness (F_{CL}). The closed-loop fitness (F_{CL}) is expressed as the average of the cumulative reward collected across M episodes:

$$\begin{aligned} \text{Maximize } F_{CL}(\mathbf{W}, \Theta) &= \frac{1}{M} \sum_{i=1}^M R_e(\mathbf{W}, \Theta), \\ \text{Subject to } \mathbf{W} &\in [-1, 1]^{n_w}, \Theta \in [-1, 1]^{n_\theta}, \end{aligned} \quad (7)$$

where n_w and n_θ are total number of weights and biases appearing in entire NLDT. In our case, we set $M = 20$.

VI. RESULTS

In this section, we present results obtained using our approach on following control problems: (i) CartPole, (ii) CarFollowing, (iii) MountainCar, and (iv) LunarLander. The first two problems have two discrete actions, third problem

has three discrete actions, and the fourth problem has four discrete actions. Open-loop statistics are reported using scores of training and testing accuracy on labelled state-action data generated from π_{oracle} . For quantifying the closed-loop performance, we use two metrics: (i) *Completion Rate* which gives a measure on the number of episodes which are successfully completed, and (ii) *Cumulative Reward* which quantifies how well an episode is executed. For each problem, 10 runs of open-loop training are executed using 10,000 datapoints. Open-loop statistics obtained from these runs are reported wherein the testing dataset also comprises of 10,000 datapoints. We choose the median performing $NLDT_{OL}$ for closed-loop analysis. We run 50 batches of 100 episodes each and report statistics regarding completion-rate and cumulative reward for $NLDT^*$ obtained after closed-loop training of median performing $NLDT_{OL}$.

A. CartPole Problem

This problem comprises of four state variables and two discrete actions. Details regarding this problem are provided in the Supplementary document. The oracle DNN controller is trained using the PPO algorithm [2]. Table I shows the performance of NLDT on training datasets of different sizes. It is observed that NLDT trained with 5,000 and 10,000 datapoints shows a robust open-loop performance and also produces 100% closed-loop performance. Based on this observation, we set the default value of training data to 10,000 for all control problems discussed in this paper. The obtained $NLDT_{OL}$ s has about two rules with an average of about four terms in the derived nodal functions. Interestingly, the same NLDT (without closed-loop training) also produces 100% closed-loop performance by achieving the maximum cumulative reward value of 200.

1) *NLDT for CartPole Problem:* One of the $NLDT_{OL}$ s obtained for the CartPole problem is shown in Figure 4 in terms of normalized state variable vector $\hat{\mathbf{x}}$.

The respective policy can be alternatively stated using a programmable if-then-else rule-structure as shown in Algorithm 1.

TABLE I: Effect of training data size to approximate performance of NLD_{TOL} on CartPole problem.

Training Data Size	Training Accuracy	Test. Accuracy (Open-loop)	# Rules	Rule Length	Cumulative Reward, Max=200	Compl. Rate (Closed-loop)
100	97.00 ± 1.55	82.79 ± 2.40	1.50 ± 0.50	3.30 ± 0.93	199.73 ± 0.32	95.00 ± 5.10
500	95.54 ± 1.53	79.66 ± 3.10	1.90 ± 0.54	3.88 ± 0.60	175.38 ± 2.61	51.00 ± 5.10
1,000	91.90 ± 0.87	90.59 ± 1.87	1.80 ± 0.40	4.05 ± 1.04	200.00 ± 0.00	100 ± 0.00
5,000	92.07 ± 1.28	92.02 ± 1.27	1.70 ± 0.46	4.25 ± 0.90	200.00 ± 0.00	100 ± 0.00
10,000	91.86 ± 1.25	92.05 ± 1.10	1.30 ± 0.46	4.45 ± 1.56	200.00 ± 0.00	100 ± 0.00




Fig. 4: CartPole NLD_{TOL} induced using 10,000 training samples. It is 91.45% accurate on the testing dataset but has 100% closed-loop performance. Normalization constants are: $\mathbf{x}^{\min} = [-0.91, -0.43, -0.05, -0.40]$, $\mathbf{x}^{\max} = [1.37, 0.88, 0.10, 0.45]$.

Algorithm 1: CartPole Rules. Normalization constants are: $x^{\min} = [-0.91, -0.43, -0.05, -0.40]$, $x^{\max} = [1.37, 0.88, 0.10, 0.45]$.

```

if  $|-0.18\widehat{x}_0\widehat{x}_2^{-2} - 0.63\widehat{x}_3^{-2} + 0.67| - 0.24 \leq 0$  then
  | Action = 0
else
  | Action = 1
  
```

A little manipulation will reveal that for a correct control strategy, Action 0 must be invoked if the following condition is true:

$$2.39 \leq \left(\frac{\widehat{x}_0}{\widehat{x}_2^2} + \frac{3.50}{\widehat{x}_3^2} \right) \leq 5.06,$$

otherwise, Action 1 must be invoked. First, notice that the above policy does not require the current velocity ($\widehat{x}_1$) to determine the left or right action movement. Second, for small values of angular position ($\widehat{x}_2 \approx 1$) and angular velocity ($\widehat{x}_3 \approx 1$), i.e. the pole is falling towards left, the above condition is always true. That is, the cart should be pushed towards left, thereby trying to stabilize the pole to vertical position. On the other hand, if the pole is falling towards right (large values of $\widehat{x}_2 \approx 2$ and $\widehat{x}_3 \approx 2$), the term in bracket will be smaller than 2.39 for all $\widehat{x}_0 \in [1, 2]$, and the above policy suggests that Action 1 (push the cart towards right) must be invoked. When the pole is falling right, a push of the cart towards right helps to stabilize the pole towards its vertical position. These extreme case analyses are intuitive and our policy can be explained for its proper working, but what our NLDT approach is able to find is a precise rule for all situations of the state variables to control the CartPole to a stable configuration, mainly using the blackbox-AI data.

B. CarFollowing Problem

We have developed a discretized version of the car following problem discussed in [32]. The objective here is to control the

rear car and follow a randomly moving car using acceleration or braking actions. This problem involves three state variables and two discrete actions. More details regarding this problem are provided in the Supplementary document. The oracle policy was obtained using a double Q-learning algorithm [33]. The reward function for the CarFollowing problem is shown in the Supplementary document, indicating that a relative distance close to 30m produces the highest reward.

Results for the CarFollowing problem are shown in Table II. An average open-loop accuracy of 96.53% is achieved with at most three rules, each having 3.28 terms on an average. For this problem, we apply the closed-loop re-optimization

TABLE II: Results on CarFollowing problem.

Train. Acc.	Test. Acc.	Depth	# Rules	Rule Length	Compl. Rate
96.41 ± 1.97	96.53 ± 1.90	1.90 ± 0.30	2.40 ± 0.66	3.28 ± 0.65	100 ± 0.00

(Blocks 4 and 5 to produce Block 6 in Figure 3) on the entire NLD_{TOL}. As shown Table III, NLDT* is able to achieve better closed-loop performances (100% completion rate and slightly better average cumulative reward). Figure 5 shows that NLDT* adheres to the 30m gap between the cars more closely than original DNN or NLD_{TOL}.

TABLE III: Closed-loop performance analysis after re-optimizing NLDT for CarFollowing problem ($k = 10^3$).

AI	Cumulative Reward		Compl. Rate
	Best	Avg ± Std	
DNN	174.16k	173.75k ± 20.95	100 ± 0.00
NLD _{TOL}	174.15k	173.87k ± 16.48	100 ± 0.00
NLDT*	179.76k	179.71k ± 0.95	100 ± 0.00




Fig. 5: Relative distance plot for CarFollowing.

The NLDT_{OL} obtained for the CarFollowing problem is shown in Figure 6 and its physical interpretation is provided in the Supplementary document.




Fig. 6: NLDT_{OL} for the CarFollowing problem. Normalization constants are: $x^{\min} = [0.25, -7.93, -1.00]$, $x^{\max} = [30.30, 0.70, 1.00]$.

Results of NLDT's performance on problems with two discrete actions (Tables I, II and III) indicate that despite having a noticeable mismatch with the open-loop output of the oracle black-box policy π_{oracle} , the closed-loop performance of NLDT is at par or at times better than π_{oracle} . This observation suggests that certain state-action pairs are not of crucial importance when it comes to executing a closed-loop control and, therefore, errors made in predicting these state-action events do not affect and deteriorate the closed-loop performance.

C. MountainCar Problem

This problem comprises of two state-variables to capture x position and velocity of the car. The task is to use three actions and drive an under-powered car to the destination (see Supplementary Document for more details).

Compilation of results of the NLDT_{OL} induced using training datasets comprising of different data distributions (*regular* and *balanced*) is presented in Table IV. A state-action plot obtained using π_{oracle} and one of the NLDT policy corresponding to the first row of Table IV is provided in Figures 1a and 1b, respectively. It is observed that about 8% mismatch in the open-loop performance (i.e. testing accuracy in Table IV) comes from the lower-left region of state-action plot (Figures 1a and 1b) due to highly non-linear nature of π_{oracle} . Despite having this mismatch, the NLDT policy is able to achieve close to 100% closed-loop control performance with an average of 2.4 rules having about 3 terms. Also, NLDT

TABLE IV: Results on MountainCar problem.

Data	Train. Acc.	Test. Acc.	Depth	# Rules	Rule Length	Compl. Rate
Reg.	91.28 ± 0.57	91.18 ± 0.35	2.00 ± 0.00	2.40 ± 0.49	2.97 ± 0.41	99.00 ± 1.71
Bal.	81.45 ± 7.36	87.23 ± 1.10	1.90 ± 0.30	2.80 ± 0.60	3.07 ± 0.42	100 ± 0.00

trained on balanced dataset (2nd row of Table IV) is able to achieve 100% closed-loop performance and involves about three control rules with an average of 3.07 terms in each rule.

Figure 7 shows one of the NLDT obtained for MountainCar problem using a regular training dataset.




Fig. 7: NLDT_{OL} for the MountainCar problem. Normalization constants are: $x^{\min} = [-1.20, -0.06]$, $x^{\max} = [0.50, 0.06]$.

D. LunarLander Problem

The task in this problem is to control the lunar-lander using four discrete actions and successfully land it on a lunar terrain. The state of the lunar-lander is expressed with eight state variables, six of which are continuous, and two are categorical. More details of this problem are provided in the Supplementary document.

TABLE V: NLDT_{OL} with depths 3 and 6 for LunarLander.

Data	Depth	Train. Acc.	Test. Acc.	# Rules	Rule Length	Compl. Rate
Reg.	3	79.17 ± 1.78	76.36 ± 3.36	5.60 ± 0.49	5.59 ± 0.75	14.00 ± 5.93
Bal.	3	69.83 ± 2.82	66.58 ± 2.03	4.40 ± 0.66	5.79 ± 1.31	42.00 ± 4.40
Reg.	6	87.43 ± 0.65	81.74 ± 0.91	34.70 ± 2.83	4.94 ± 0.34	48.00 ± 2.77
Bal.	6	81.74 ± 1.78	71.52 ± 1.24	25.70 ± 5.83	5.17 ± 0.37	93.00 ± 3.30

Table V provides the compilation of results obtained using NLDT_{OL}. In this problem, while a better open-loop performance occurs for regular dataset, a better closed-loop performance is observed when the NLDT open-loop training is done on the balanced dataset. Also, NLDT_{OL}s with depth three are not adequate to achieve high closed-loop performance. The best performance is observed using balanced dataset where NLDT_{OL} of depth 6 achieves 93% episode completion rate. A specific NLDT_{OL} of depth 6 with 26 rules, each having about 4.15 terms is shown in the Supplementary Document.

It is understandable that a complex control task involving many state variables cannot be simplified or made interpretable with just one or two control rules. Next, we use a part of the NLDT_{OL} from the root node to obtain a pruned NLDT_{OL}^(P) (step 'B' in Figure 3) and re-optimize all weights (**W**) and biases (**Θ**) using the procedure discussed in Section V-B (shown by orange box in Figure 3) to find closed-loop NLDT*. Table VI shows that for the pruned NLDT-3 which comprises of the top three layers and involves only four rules of original 26-rule NLDT_{OL} (i.e. NLDT-6), the closed-loop performance increases from 51% to 96% (NLDT*-3 results in Table VI) after re-optimizing its weights and biases with closed-loop training. The resulting NLDT with its associated four rules is shown in Figure 8.

As shown in Table VI, the NLDT* with just two rules (NLDT-2) is too simplistic and does not recover well after

TABLE VI: Closed-loop performance on LunarLander problem with and without re-optimization on 26-rule NLDT_{OL}. Number of rules are specified in brackets for each NLDT and total parameters for the DNN is marked.

Re-Opt.	NLDT-2 (2)	NLDT-3 (4)	NLDT-4 (7)	NLDT-5 (13)	NLDT-6 (26)	DNN (4,996)
Cumulative Reward						
Before	-1675.77 ± 164.29	42.96 ± 13.83	54.24 ± 27.44	56.16 ± 23.50	169.43 ± 23.96	247.27 ± 3.90
After	-133.95 ± 2.51	231.42 ± 17.95	234.98 ± 22.25	182.87 ± 21.92	214.94 ± 17.31	
Completion Rate						
Before	0.00 ± 0.00	51.00 ± 3.26	82.00 ± 9.80	79.00 ± 7.66	93.00 ± 3.30	94.00 ± 1.96
After	48.00 ± 7.38	96.00 ± 2.77	99.00 ± 1.71	93.00 ± 7.59	94.00 ± 4.45	




Fig. 8: Final NLDT*-3 for LunarLander prob. $\hat{x}_i$ is a normalized state variable (see Supplementary Document).

re-optimization. However, the NLDT*s with four and seven rules achieve a near 100% closed-loop performance. Clearly, an NLDT* with more rules (NLDT-5 and NLDT-6) are not worth considering since both closed-loop performances and the size of rule-sets are worse than NLDT*-4. Note that DNN produces a better reward, but not enough completion rate, and the policy is more complex with 4,996 parameters.

Figure 9 shows the closed-loop training curve for generating NLDT* from Depth-3 NLDT_{OL}^(P). The objective here is to maximize the closed-loop fitness (reward) F_{CL} (Eq. 7) which is expressed as the average of cumulative reward R_e collected over M episodes. It is evident that the cumulative reward for




Fig. 9: Closed-loop training plot for finetuning the rule-set corresponding to depth-3 NLDT_{OL}^(P) to obtain NLDT* for LunarLander problem.

the best-population member climbs to the target reward of 200 at around 25-th generation and the average cumulative reward of the population also catches up the best cumulative reward

value with generations.

To check the repeatability of our approach, another run for generating NLDT_{OL} and NLDT* is executed. The resulting NLDT and corresponding equations are provided in the Supplementary document. A visualization of the real-time closed-loop performance obtained using this new NLDT for two different rule-sets (i.e. before applying re-optimization and after applying the re-optimization through closed-loop training) is shown in <https://youtu.be/DByYWTQ6X3E>. It can be observed in the video that the closed-loop control executed using the Depth-3 NLDT_{OL}^(P) comprising of rules obtained directly from the open-loop training (i.e. without any re-optimization) is able to bring the LunarLander close to the target. However the LunarLander hovers above the landing pad and the Depth-3 NLDT_{OL}^(P) is unable to land it in most occasions. Episodes in these cases are terminated after the flight-time runs out. On the other hand, the Depth-3 NLDT* comprising of rule-sets obtained after re-optimization through closed-loop training is able to successfully land the LunarLander.

VII. SCALE-UP STUDY AND IMPROVIZATION ON ACROBOT CONTROL PROBLEM

In this section, we investigate how the overall algorithm can be made more efficient in terms of – *training time* and *scalability*. To this purpose, we introduce a benchmark problem of *planar serial manipulator (PSM)*. This problem is inspired from the classical Acrobot control problem [12]. A schematic of the PSM problem is provided in Figure 10.




Fig. 10: Planar Serial Manipulator (PSM) benchmark problem.

The state space for the PSM problem (Figure 10) comprises of angular position θ_i and angular velocity ω_i of each joint. Thus, for a n -link manipulator involving n revolute joints, the

state space would be $2n$ dimensional. The motor is located at the last joint of the manipulator and is actuated using three torque values: $-\tau$, 0 and τ . Each link is of 1 unit length and has its center-of-mass at its geometric center. The motor is assumed to be massless for the sake of simplicity. The base of the manipulator is located at $(0,0,0)$ and the motion of the PSM is limited to the XZ plane. There is a downward gravitational pull (g) of 10 units (i.e. -10 along vertical Z axis). Torque is applied along the Y-axis. The task in this problem is to take the end-effector (i.e. tip of the last link, Z-coordinate z_E) of the PSM to a desired height of H units by supplying torque to the motor located at the last joint (joint between $(n-1)$ -th and n -th link). The difficulty of this benchmark problem can be adjusted by

- 1) changing the number of links,
- 2) changing the value of desired height level H ,
- 3) changing the value of torque τ , and
- 4) placing extra motors at other joints.

We simulate the mechanics of the planar serial manipulator using *PyBullet* [34]: a Python based physics engine.

In our work, we provide two case-scenarios by focusing at the first three points of the above list. As mentioned before, by changing the number of links, the dimension of the state-space changes. The dimension of the action-space depends on the number of motors used. In the present work, we keep the number of motors fixed to one, having three discrete actions.

The details regarding two environments which are created and studied in this section are summarized in Table VII.

TABLE VII: Details regarding custom designed Planar Serial Manipulator (PSM) environments.

Env. Name	Motor Torque (τ)	Desired Height (H)	# State Vars.
5-Link PSM	1,000	+2	10
10-Link PSM	2,000	+2	20

The reward function $r(\mathbf{x}, A)$ is given by the following equation:

$$r(\mathbf{x}) = \begin{cases} -1 - \left(\frac{H+1-z_E}{n+H+1} \right)^2, & \text{if } z_E < H, \\ 100, & \text{if } z_E \geq H. \end{cases} \quad (8)$$

The minimum value for z_E is $-n$ when the entire manipulator is stretched to its full length and all joint angles (i.e. θ_i) are at 0 deg.

At the beginning of an episode, joint angles θ_i of the manipulator are randomly initialized between -5 deg and $+5$ deg, and the angular velocities ω_i are initialized to a value between -0.5 rad/sec and $+0.5$ rad/sec.

In next sections, we discuss results obtained on the above two custom designed environments by using different procedures of inducing NLDE_{OL}. The black-box AI (DNN) is trained using the PPO algorithm [2]. The resulting DNN has two hidden layers of 64 nodes each and has total 5699 parameters thereby making it massively un-interpretable.

A. Ablation Study for Open-loop Training

In this section, we launch two separate studies related to open-loop training procedure (see Figure 3). It is seen in

Section V-A that the open-loop training is conducted using a hierarchical bilevel-optimization algorithm, which is discussed at length in [30]. A dedicated bilevel-optimization algorithm is invoked to derive the split-rule $f(\mathbf{x})$ at a given conditional node. The upper-level search is executed using a discrete version of a genetic algorithm and the lower-level search is realized through an efficient real-coded genetic algorithm (RGA). Evolutionary algorithms are in general considered robust and have a potential to conduct more global search. However, being population driven, their search-speed is often less than that of classical optimization algorithms. In this section, we study the effect of replacing the RGA with the classical *sequential quadratic programming* (SQP) optimization algorithm in the lower-level of the overall bilevel algorithm for obtaining NLDT_{OL}. Later, closed-loop training based on real-coded genetic algorithm is applied to NLDT_{OL} to obtain NLDT* by re-optimizing real valued coefficients of NLDT_{OL} (Section V-B). We use the *SciPy* [35] implementation of SQP. The initial point required for SQP is obtained using the mixed dipole concept [36], [37], [38], [30].

For analysis, we induce an open-loop NLDT_{OL} of depth-3 using a balanced training dataset of 10,000 datapoints. The testing dataset also comprises of 10,000 datapoints. Comparison of accuracy scores and average training time of inducing NLDT_{OL} by using SQP and RGA algorithm at lower-level is provided in Table VIII. For a given procedure (SQP or RGA) the best NLDT_{OL} from 10 independent runs is chosen and is re-optimized using closed-loop training (Section V-B). Statistics regarding closed-loop performance of NLDT* is shown in the last two columns of Table VIII. It is to note here that the closed-loop training is done using the real-coded genetic algorithm (RGA) discussed in Section V-B.

It can be observed from the results that open-loop training done with SQP in lower-level is about 70 times faster than the training done using RGA at the lower level. However, the training done with RGA has a better overall open-loop performance. Thus, if the task is to closely mimic the behavior of black-box AI or if only a high classification accuracy is desired (in case of classification problems), then RGA is the recommended algorithm for lower-level optimization to obtain NLDT_{OL}. However, NLDT* obtained after re-optimizing NLDT_{OL} corresponding to SQP and RGA have similar closed-loop completion rate (last column of Table VIII). This implies that despite low open-loop accuracy scores, the open-loop training done using SQP in lower-level is successful in determining the template of split-rules $f(\mathbf{x})$ and topology of NLDT, which upon re-optimization via closed-loop training algorithm can fetch a decent performing NLDT*. During open-loop training, the search on weights and coefficients using SQP is possibly not as perfect as compared to the one achieved through RGA, however, the re-optimization done through closed-loop training can compensate this shortcoming of SQP algorithm and produce NLDT* with a respectable closed-loop performance. Additionally, in either cases, the NLDT* obtained always has a better closed-loop performance than the original black-box DNN policy. This observation suggests that it is preferable to use SQP in lower-level during open-loop training to quickly arrive at a rough structure of NLDT_{OL}

TABLE VIII: Comparing performance of different lower-level optimization algorithms. For comparison, closed-loop performance of the original DNN policy is also reported.

Algo. Name	Open-Loop NLDT _{OL}			Closed-Loop NLDT*	
	Training Accuracy	Testing Accuracy	Training Time (s)	Cumulative Reward	Completion Rate
5-Link Manipulator					
SQP	62.46 ± 2.01	69.34 ± 5.39	15.29 ± 4.95	−146.81 ± 8.53	96.00 ± 2.77
RGA	71.14 ± 1.77	69.17 ± 4.39	1091.56 ± 319.18	−152.64 ± 6.62	99.00 ± 1.71
DNN	NA	NA	NA	−183.35 ± 12.22	89.00 ± 5.19
10-Link Manipulator					
SQP	56.57 ± 1.00	55.64 ± 3.58	39.27 ± 11.63	−318.82 ± 15.52	96.00 ± 3.92
RGA	65.84 ± 0.85	62.60 ± 3.09	2860.96 ± 789.49	−281.88 ± 9.38	95.00 ± 4.31
DNN	NA	NA	NA	−325.86 ± 4.63	85.88 ± 1.94

and then use closed-loop training to derive a high performing NLDT*.

B. Closed-loop Visualization

In this section, we provide a visual insight into the closed-loop performance of DNN and NLDT* which we derived in the previous section. In our case, the frequency of the simulation is set to 240Hz, meaning that the transition to the next state is calculated using the time-step of 1/240 seconds. Geometrically speaking, this implies that the Euclidean distance between states from neighboring time-steps would be small. An AI (DNN or NLDT*) outputs the action value of 0 ($-\tau$ torque), 1 (0 torque) or 2 ($+\tau$ torque) for a given input state. Action Vs Time plots corresponding to different closed-loop simulation runs obtained by using DNN, NLDT* (SQP)³ and NLDT* (RGA)⁴ as controllers is shown in Figure 11 for 5-link manipulator problem (plots for 10-link PSM are provided in the Supplementary document).

Certain key observations can be made by looking at the plots in Figure 11. The control output of DNN is more erratic, with sudden jerks as compared to the control output of NLDT* (SQP) and NLDT* (RGA). The performance of NLDT* in Figures 11b and 11c is smooth and regular. This behaviour can be due to the involvement of relatively simpler non-linear rules (as compared to the complicated non-linear rule represented by DNN) which are captured inside NLDT*. This is equivalent to the observation we made for the mountain car problem in Figures 1a and 1b, wherein the black-box AI had a very erratic behavior for the region of state-space in the lower-half of the state-action plot, while the output of NLDT was more smooth. Additionally, it was seen in Table VIII that the NLDT* (irrespective of how its predecessor NLDT_{OL} was obtained, i.e. either through SQP or RGA in lower-level) showed better closed-loop performance than the parent DNN policy. This observation implies that simpler rules expressed in the form of a nonlinear decision tree have better generalizability, thereby giving more robust performance for randomly initialized control problems. A careful investigation to the plots in Figure 11b and 11c reveals that only two out of three allowable actions are required to efficiently execute

the given control task of lifting the end-effector of a 5-link serial manipulator. This concept will be used to re-engineer the NLDT*, a discussion regarding which is provided in the next section.

C. Reengineering NLDT*

It is seen in action-time plots in Figure 11b and Figure 11c that not all actions are required to perform a given control task. Also, it may be possible that while performing a closed-loop control using NLDT, not all branches and nodes of NLDT are visited while deciding an *action*. Thus, the portion of NLDT which is not being utilized or is getting utilized very rarely can be pruned and the overall NLDT architecture can be made simpler. To illustrate this idea, we consider the NLDT which is derived for the 5-link manipulator problem. The topology of the best performing NLDT_{OL} (SQP) for the 5-link manipulator problem is shown in Figure 12a.

As mentioned before, this NLDT_{OL} is trained on a balanced training dataset which is generated by collecting state-action pairs using the oracle DNN controller. In the figure, at each node, information regarding its node-id, class distribution (given in square parenthesis) and the most dominating class is provided. Other than the root-node (Node 0), all nodes are colored to indicate the dominating class, however, it is to note that only the class associated to leaf-nodes carry the actual meaning while predicting the action for a given input state. This NLDT_{OL} comprises of four split-rules in total. The class-distribution for each node is obtained by counting how many datapoints from the balanced training dataset visited a given node. Thus, the root node comprises of all datapoints (total 10,000), which are then scattered according to the split-rules present at each conditional node.

Figure 12b provides topology of NLDT* which is obtained after re-optimizing NLDT_{OL} of Figure 12a using closed-loop training. As discussed in Sections V-A and V-B, the topology of the tree and the structure of non-linear rules is identical for both: NLDT_{OL} and NLDT*. However, the weights and biases of NLDT* are updated to enhance the closed-loop control performance. Similar to NLDT_{OL} of Figure 12a, the information regarding node id and class-distribution is provided for all the nodes of NLDT* in Figure 12b. However, the data distribution in NLDT* is obtained by using the actual state-action data from closed-loop simulations, wherein NLDT* is used as a controller. Total 10,000 datapoints are

³NLDT* (SQP) indicates that the corresponding NLDT_{OL} was derived using the SQP algorithm in the lower-level

⁴NLDT* (RGA) indicates that the corresponding NLDT_{OL} was derived using the RGA algorithm in the lower-level




Fig. 11: Action Vs. Time plot for 5-Link manipulator problem. Figure 11b provides the plot for NLDT* which is obtained from the $NLDT_{OL}$ trained using SQP algorithm in lower-level. Similarly, Figure 11c provides the plot for NLDT* which is obtained from the $NLDT_{OL}$ trained using RGA algorithm in lower-level.




Fig. 12: NLDTs for 5-Link Manipulator problem.

collected in the form of sequential states-action pairs from closed-loop simulation runs which are executed using NLDT*. As it can be seen in the root node of NLDT* (Figure 12b), out of 10,000 states visited during closed-loop control, action 0 ($-\tau$ torque) was chosen by NLDT* in 7,736 states and action 2 ($+\tau$ torque) was chosen for 2,264 states. Action 1 was never selected by NLDT* during closed-loop control. This is consistent with what we have observed in the Action Vs Time plot in Figure 11b, wherein most of the time action 0 was executed, while there was no event where action 1 was executed. The flow of these 10,000 state-action pairs through NLDT* and their corresponding distribution in each node of NLDT* is provided in Figure 12b. It can be observed that Node 5, Node 4 and Node 8 of NLDT* are never visited during closed-loop control. This implies that splits at Node 2, Node 1 and Node 6 are redundant. Thus, the part of NLDT* shown in red-box in Figure 12b can be pruned and the overall topology of the tree can be simplified. The pruned NLDT* will involve only one split (occurring at Node 0) and two leaf

nodes: Node 1 and Node 6. However, it is to note here that we need to re-assign class-labels to the newly formed leaf nodes (i.e. Node 1 and Node 6) based on the data-distribution from closed-loop simulations. The old class-labelling for the Node 1 and Node 6 was done based on the open-loop data (Figure 12a). Using the new class distribution corresponding to NLDT* (Figure 12b), Node 1 is re-labelled with Class-2 and Node 6 with Class-0. The pruned version of NLDT* of Figure 12b is provided in Figure 13 (here nodes are re-numbered, with Node 6 of NLDT* in Figure 12b re-numbered to Node 2 in the pruned NLDT* as shown in Figure 13). The split-rule corresponding to the root-node is also shown.




Fig. 13: Pruned version of NLDT* (Figure 12b) for 5-link manipulator problem.

Interestingly, out of 10 total state-variables, only 2 are used to decide which action to execute for closed-loop control: x_1 corresponds to the angular position of the second link and x_9 variable corresponds to the angular velocity of the last joint. The above rule indicates a single and interpretable relationship: action 2 must be invoked when $\hat{x}_9 \leq \sqrt[3]{(1.885 - 0.115/\hat{x}_1^3)}$, otherwise action 0 must be invoked. This NLDT* corresponds to a cumulative reward of -139.78 ± 7.65 and performs at $99.00 \pm 1.71\%$ completion rate, providing a remarkably simple interpretation of the control strategy for this apparently complex problem.

VIII. CONCLUSIONS

In this paper, we have proposed a two-step strategy to arrive at hierarchical and relatively interpretable controller using a nonlinear decision tree (NLDT) concept. The NLDT training phases use recent advances in nonlinear optimization to focus its search on rule structure and details describing weights and biases of the rules by using a bilevel optimization algorithm.

Starting with an open-loop training, which is relatively fast but uses only time-instant state-action data, we have proposed a final closed-loop training phase in which the complete or a part of the open-loop NLDT is re-optimized for weights and biases using complete episode data. Results on four popular discrete action problems have amply demonstrated the usefulness of the proposed overall approach.

This proof-of-principle study encourages us to pursue a number of further studies. First, the scalability of the NLDT approach for challenging large-dimensional state-action space problems must now be explored. A previous study on NLDT [30] on binary classification of dominated and non-dominated datapoints in multi-objective problems was successfully extended to 500-variable problems. While it is encouraging, the use of customization methods for initialization and genetic operators using problem heuristics and/or recently proposed *innovization* methods [39] in the upper level problem can be tried. Second, this study has used a computationally fast open-loop accuracy measure as the fitness for evolution of the NLDT_{OL}. This is because, in general, an NLDT_{OL} with a high open-loop accuracy is likely to achieve a high closed-loop performance. However, we have observed here that a high closed-loop performance is achievable with a NLDT_{OL} having somewhat degraded open-loop performance, but re-optimized using closed-loop performance metrics. Thus, a method to identify crucial (open-loop) states from an AI-based controller dataset that can improve the closed-loop performance would be another interesting step for deriving NLDT_{OL}. This may eliminate the need for re-optimization through closed-loop training. Third, a more comprehensive study using closed-loop performance and respective complexity as two conflicting objectives for a bi-objective NLDT search would produce multiple trade-off control rule-sets. Such a study can not only make the whole search process faster due to the expected similarities among multiple policies, they will also enable users to choose a single policy solution from a set of accuracy-complexity trade-off solutions.

ACKNOWLEDGEMENTS

This study was supported by Ford-MSU Alliance project.


REFERENCES

- [1] J. Schulman, S. Levine, P. Abbeel, M. Jordan, and P. Moritz, "Trust region policy optimization," in *International Conference on Machine Learning (ICML)*, 2015, pp. 1889–1897.
- [2] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.
- [3] T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra, "Continuous control with deep reinforcement learning," *arXiv preprint arXiv:1509.02971*, 2015.
- [4] V. Mnih, A. P. Badia, M. Mirza, A. Graves, T. Lillicrap, T. Harley, D. Silver, and K. Kavukcuoglu, "Asynchronous methods for deep reinforcement learning," in *International Conference on Machine Learning (ICML)*, 2016, pp. 1928–1937.
- [5] K. Senda, S. Hattori, T. Hishinuma, and T. Kohda, "Acceleration of reinforcement learning by policy evaluation using nonstationary iterative method," *IEEE transactions on cybernetics*, vol. 44, no. 12, pp. 2696–2705, 2014.
- [6] M. L. Koga, V. Freire, and A. H. Costa, "Stochastic abstract policies: Generalizing knowledge to improve reinforcement learning," *IEEE Transactions on Cybernetics*, vol. 45, no. 1, pp. 77–88, 2014.
- [7] X. Xu, Z. Huang, D. Graves, and W. Pedrycz, "A clustering-based graph laplacian framework for value function approximation in reinforcement learning," *IEEE Transactions on Cybernetics*, vol. 44, no. 12, pp. 2613–2625, 2014.
- [8] W. J. Murdoch, C. Singh, K. Kumbier, R. Abbasi-Asl, and B. Yu, "Interpretable machine learning: definitions, methods, and applications," *arXiv preprint arXiv:1901.04592*, 2019.
- [9] Z. C. Lipton, "The mythos of model interpretability," *Queue*, vol. 16, no. 3, pp. 31–57, 2018.
- [10] M. A. Ahmad, C. Eckert, and A. Teredesai, "Interpretable machine learning in healthcare," in *Proceedings of the 2018 ACM international conference on bioinformatics, computational biology, and health informatics*, 2018, pp. 559–560.
- [11] B. Goodman and S. Flaxman, "European union regulations on algorithmic decision-making and a "right to explanation"," *AI magazine*, vol. 38, no. 3, pp. 50–57, 2017.
- [12] R. S. Sutton, "Generalization in reinforcement learning: Successful examples using sparse coarse coding," in *Advances in Neural Information Processing Systems 8*, D. S. Touretzky, M. C. Mozer, and M. E. Hasselmo, Eds. MIT Press, 1996, pp. 1038–1044.
- [13] J. Peters, K. Mülling, and Y. Altun, "Relative entropy policy search," in *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence (AAAI-10)*, vol. 10. Atlanta, 2010, pp. 1607–1612.
- [14] W. D. Smart and L. P. Kaelbling, "Practical reinforcement learning in continuous spaces," in *International Conference on Machine Learning (ICML)*, 2000, pp. 903–910.
- [15] R. Noothigattu, D. Bouneffouf, N. Mattei, R. Chandra, P. Madan, K. Varshney, M. Campbell, M. Singh, and F. Rossi, "Interpretable multi-objective reinforcement learning through policy orchestration," *arXiv preprint arXiv:1809.08343*, 2018.
- [16] F. Maes, R. Fonteneau, L. Wehenkel, and D. Ernst, "Policy search in a space of simple closed-form formulas: towards interpretability of reinforcement learning," in *International Conference on Discovery Science*. Springer, 2012, pp. 37–51.
- [17] D. Hein, S. Udluft, and T. A. Runkler, "Interpretable policies for reinforcement learning by genetic programming," *Engineering Applications of Artificial Intelligence*, vol. 76, pp. 158–169, 2018.
- [18] G. Liu, O. Schulte, W. Zhu, and Q. Li, "Toward interpretable deep reinforcement learning with linear model u-trees," in *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 2018, pp. 414–429.
- [19] A. Verma, V. Murali, R. Singh, P. Kohli, and S. Chaudhuri, "Programmatically interpretable reinforcement learning," *arXiv preprint arXiv:1804.02477*, 2018.
- [20] J. Kennedy and R. Eberhart, "Particle swarm optimization," in *Proceedings of ICNN'95-International Conference on Neural Networks*, vol. 4. IEEE, 1995, pp. 1942–1948.
- [21] D. Hein, A. Hentschel, T. Runkler, and S. Udluft, "Particle swarm optimization for generating interpretable fuzzy reinforcement learning policies," *Engineering Applications of Artificial Intelligence*, vol. 65, pp. 87–98, 2017.
- [22] L. Breiman, *Classification and regression trees*. Routledge, 2017.
- [23] S. Ross, G. Gordon, and D. Bagnell, "A reduction of imitation learning and structured prediction to no-regret online learning," in *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, 2011, pp. 627–635.
- [24] O. Bastani, Y. Pu, and A. Solar-Lezama, "Verifiable reinforcement learning via policy extraction," in *Advances in neural information processing systems (NIPS)*, 2018, pp. 2494–2504.
- [25] O. Bastani, C. Kim, and H. Bastani, "Interpretability via model extraction," *arXiv preprint arXiv:1706.09773*, 2017.
- [26] G. Vandewiele, O. Janssens, F. Ongenaes, F. De Turck, and S. Van Hoecke, "Genesim: Genetic extraction of a single, interpretable model," *arXiv preprint arXiv:1611.05722*, 2016.
- [27] D. Ernst, P. Geurts, and L. Wehenkel, "Tree-based batch mode reinforcement learning," *Journal of Machine Learning Research*, vol. 6, no. Apr, pp. 503–556, 2005.
- [28] J. L. Bentley, "Multidimensional binary search trees used for associative searching," *Communications of the ACM*, vol. 18, no. 9, pp. 509–517, 1975.
- [29] A. Silva, M. Gombolay, T. Killian, I. Jimenez, and S.-H. Son, "Optimization methods for interpretable differentiable decision trees applied to reinforcement learning," in *International Conference on Artificial Intelligence and Statistics*, 2020, pp. 1855–1865.
- [30] Y. Dhebar and K. Deb, "Interpretable rule discovery through bilevel optimization of split-rules of nonlinear decision trees for classification problems," *IEEE Transactions on Cybernetics*, pp. 1–12, 2020.

- [31] A. Sinha, P. Malo, and K. Deb, "A review on bilevel optimization: From classical to evolutionary approaches and applications," *IEEE Transactions on Evolutionary Computation*, vol. 22, no. 2, pp. 276–295, 2018.
- [32] S. Nageshrao, B. Costa, and D. Filev, "Interpretable approximation of a deep reinforcement learning agent as a set of if-then rules," in *18th IEEE International Conference On Machine Learning And Applications (ICMLA)*. IEEE, 2019, pp. 216–221.
- [33] H. Van Hasselt, A. Guez, and D. Silver, "Deep reinforcement learning with double q-learning" *arXiv preprint arXiv:1509.06461*, 2015.
- [34] E. Coumans and Y. Bai, "Pybullet, a python module for physics simulation for games, robotics and machine learning," <http://pybullet.org>, 2016–2020.
- [35] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, Í. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, and SciPy 1.0 Contributors, "SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python," *Nature Methods*, vol. 17, pp. 261–272, 2020.
- [36] L. Bobrowski and M. Kretowski, "Induction of multivariate decision trees by using dipolar criteria," in *European Conference on Principles of Data Mining and Knowledge Discovery*. Springer, 2000, pp. 331–336.
- [37] M. Kretowski, "An evolutionary algorithm for oblique decision tree induction," in *International Conference on Artificial Intelligence and Soft Computing*. Springer, 2004, pp. 432–437.
- [38] M. Kretowski and M. Grześ, "Global induction of oblique decision trees: an evolutionary approach," in *Intelligent Information Processing and Web Mining*. Springer, 2005, pp. 309–318.
- [39] K. Deb and A. Srinivasan, "Innovization: Innovating design principles through optimization," in *Proceedings of the 8th annual conference on Genetic and evolutionary computation*. ACM, 2006, pp. 1629–1636.



Subramanya Nageshra received his Ph.D. degree from the Delft Center for Systems and Control, Delft University of Technology, Delft, The Netherlands, in 2016. He is currently a Research Engineer at Ford Green Field Labs, Palo Alto, California. Prior to that he worked as a post-doctoral research scholar at University of Michigan Ann-Arbor. He is actively leveraging machine learning methods particularly reinforcement learning techniques to build next generation state-of-the-art advanced driver assistance technologies.



Ling Zhu received her master's and PhD degree in Computer Science and Engineering Department at Michigan State University under the supervision of Prof. Sandeep Kulkarni and Prof. Kalyanmoy Deb. She is currently an algorithm engineer at Microsoft, and working on research and development in advanced 3D sensing technology. Her research interests include 3D machine vision, computer vision, evolutionary optimization, metaheuristic optimization, robotics, and distributed systems and computing.



Yashesh Dhebar did his PhD in the Department of Mechanical Engineering at Michigan State University under the guidance of Prof. Kalyanmoy Deb and graduated in 2020 with his thesis focused at developing Interpretable AI models using modern optimization techniques. He did his bachelors and masters in Mechanical Engineering from the Indian Institute of Technology Kanpur and graduated in 2015 with his thesis on modular robotics and their operation in cluttered environments. Currently he is working with HEEDS Design Space Exploration

software development team at Siemens DISW in East Lansing, Michigan.



Kalyanmoy Deb is University Distinguished Professor and Koenig Endowed Chair Professor at the Department of Electrical and Computer Engineering at Michigan State University (MSU), East Lansing, USA. He received his PhD and MS degrees from University of Alabama, Tuscaloosa, USA, in 1989 and 1991, respectively, and BTech (Hons.) degree from Indian Institute of Technology Kharagpur, India in 1985. Prof. Deb's research interests are in evolutionary optimization algorithms and their application in optimization and machine learning. He is a fellow of IEEE and ASME. He has written more than 580 research papers with 168,000+ Google Scholar citations and received many awards including IEEE EC Pioneer award, Infosys prize, and Edgeworth-Pareto award.



citations. He is a fellow

Dimitar Filev is a Senior Henry Ford Technical Fellow at Ford Motor Company. Prior to joining Ford Motor Company in 1994, Dimitar Filev held faculty positions at the State University of New York (1991 – 1994) and Bulgarian Academy of Sciences (1980 – 1991). He did his PhD in Electrical and Electronics Engineering (1977 – 1979) and an integrated bachelors and masters in Electrical Engineering (1972 -1976) from Czech Technical University in Prague. His technical contributions include 100+ granted US patents and over 200 publications with 16,000+ peer citations. He is a fellow of IEEE and a member of NAE.