



HAL
open science

Vanilla Gradient Descent for Oblique Decision Trees

Subrat Prasad Panda, Blaise Genest, Arvind Easwaran, Ponnuthurai
Nagaratnam Suganthan

► **To cite this version:**

Subrat Prasad Panda, Blaise Genest, Arvind Easwaran, Ponnuthurai Nagaratnam Suganthan. Vanilla Gradient Descent for Oblique Decision Trees. ECAI 2024 - 27TH EUROPEAN CONFERENCE ON ARTIFICIAL INTELLIGENCE, 2024, Saint-Jacques de Compostelle, Spain. 10.3233/FAIA240607 . hal-04777062

HAL Id: hal-04777062

<https://hal.science/hal-04777062v1>

Submitted on 12 Nov 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Vanilla Gradient Descent for Oblique Decision Trees

Subrat Prasad Panda^{a,*}, Blaise Genest^b, Arvind Easwaran^a and Ponnuthurai Nagaratnam Suganthan^c

^aNTU, Singapore

^bCNRS, IPAL, France and CNRS@CREATE, Singapore

^cQatar University, Qatar

Abstract. Decision Trees (DTs) constitute one of the major highly non-linear AI models, valued, e.g., for their efficiency on tabular data. Learning accurate DTs is, however, complicated, especially for oblique DTs, and does take a significant training time. Further, DTs suffer from overfitting, e.g., they proverbially "do not generalize" in regression tasks. Recently, some works proposed ways to make (oblique) DTs differentiable. This enables highly efficient gradient-descent algorithms to be used to learn DTs. It also enables generalizing capabilities by learning regressors at the leaves simultaneously with the decisions in the tree. Prior approaches to making DTs differentiable rely either on probabilistic approximations at the tree's internal nodes (soft DTs) or on approximations in gradient computation at the internal node (quantized gradient descent). In this work, we propose *DTSemNet*, a novel *semantically* equivalent and invertible encoding for (hard, oblique) DTs as Neural Networks (NNs), that uses standard vanilla gradient descent. Experiments across various classification and regression benchmarks show that oblique DTs learned using *DTSemNet* are more accurate than oblique DTs of similar size learned using state-of-the-art techniques. Further, DT training time is significantly reduced. We also experimentally demonstrate that *DTSemNet* can learn DT policies as efficiently as NN policies in the Reinforcement Learning (RL) setup with physical inputs (dimensions ≤ 32). The code is available at <https://github.com/CPS-research-group/dtsemnet>.

1 Introduction

DTs are widely adopted in several domains, such as the medical domain [13]. Studies, including recent ones [18], show that DTs classify tabular datasets very well due to their inductive bias toward learning non-smooth functions, even better than NNs. Learning DTs for a given task is, however, a complex task: due to the combinatorial explosion over choices of the decisions to branch over at every node of the tree, learning optimal DTs is an NP-hard problem [25].

Current methods of learning DTs can be broadly classified into four approaches: (a) greedy optimization, which grows trees using split criteria such as Classification and Regression Trees (CART) [10], (b) non-greedy optimization, which jointly optimize decision nodes under global objective such as Tree Alternating Optimization (TAO) [12, 45]), (c) Global searches, either presented as Mixed Integer Programs (MIP) [6], or using Evolutionary Algorithms (EA), e.g. [15], and (d) Gradient Descent by making the tree differentiable [21, 43, 44]. Greedy optimization techniques usually learn low-performing trees [44, 2], global searches are computationally expensive when the number of possible tree configurations is large [7, 15] and non-greedy

approaches, although computationally better, are still computationally more expensive than gradient-based approaches [2]. Moreover, training without gradient descent in the RL setting is very challenging, as there is no dataset: the agent learns from experience collected in an environment. A workaround is first to learn an NN using gradient descent, and imitate the NN (treated as an oracle) with a (classification-) DT using greedy strategies such as CART, which works reasonably well for simple RL problems with a small (up to ≈ 6 features) *discrete* action space [4]. On the other hand, recent results with the gradient-based approach have shown to be efficient in terms of training time and accuracy for both classification and regression tasks [2]. Additionally, it is capable of directly learning (regression-) DT policies in RL settings with *continuous* actions [31].

To apply gradient descent at the training stage on DTs, most previous works adopt "soft" decisions (e.g., using the `Sigmoid` activation function) to make the tree differentiable [26, 9, 44, 21]. However, the tree obtained does not provide "hard" decisions but "soft", that is, probabilistic ones. Hard DTs can be obtained from soft DTs, but accuracy decreases. Some recent studies, such as Dense Gradient Trees (DGT) [2] and Interpretable Continuous Control Trees (ICCT) [31], have introduced an alternative approach to obtain hard DTs, using an approximation during backpropagation to compute gradients using Straight Through Estimators (STEs) [5, 22]. This approximation can hamper DT training, especially in large datasets or RL, where errors may accumulate over many training steps.

In this work, we propose a powerful novel encoding of oblique DTs as NNs, namely the *DTSemNet* architecture, which overcomes the above-mentioned shortcomings. It uses `ReLU` activation functions and linear operations, making it differentiable and allowing gradient descent to be applied to learn the structure. The encoding is semantically equivalent to a (hard) oblique DT architecture, such that decisions (weights) in the DT correspond one-to-one to the trainable weights in the NN. The other weights are fixed (non-trainable) in the architecture. The primary use of *DTSemNet* is as a classifier; that is, it provides a class associated with a given input, using a typical `argmax` operation to select the class associated with the highest output. We show in Theorem 1 that *DTSemNet* is exactly equivalent to a DT; that is, for every input vector, the output class produced by *DTSemNet* and by the DT are the same. It has the potential for application in many different contexts to learn DTs. Compared with DGT [2], and ICCT [31], both needing *many* STE calls, *DTSemNet* uses `ReLU` and *standard* gradient descent without STE approximation. We propose a simple extension to *DTSemNet* for *regression tasks*, which requires *one* STE call to combine the choice of value with the choice of leaf. Each leaf is associated with a regressor whose parameters are learned

* Corresponding Author. Email: subratpr001@e.ntu.edu.sg

simultaneously with the DT decisions via (standard) gradient descent. This process results in DTs that generalize, aided by the regressors at the leaf, similar to [31],[45].

We experimented with *DTSemNet* in various benchmarks, both for classification and regression tasks, and compared its performance against different methods producing hard DTs: DGT[2] for gradient descent-based, TAO [12, 45] for non-greedy algorithm, CRO-DT [15] for global-searches and CART as the standard for greedy algorithms. The classification and regression tasks mostly involve datasets with tabular data, except for MNIST, a small-sized image dataset. The number of features varies from 4 to 780 (outlier for MNIST), and the number of classes varies from 2 to 26. For classification tasks, *DTSemNet* achieves the best performance on every dataset, showing the efficiency of our proposed *unapproximated* methodology. For regression tasks, *DTSemNet* is either the best performing or second (e.g., beaten by TAO). One plausible cause of this discrepancy is the STE approximation used for *regression tasks* to combine the decision in the tree with the value computed at the nodes. This approximation sometimes hurts efficiency, whereas no such approximation for classification results in optimal performance. In both cases, DT learning time is much shorter using gradient descent than competing solutions.

Last, we explain how to deploy *DTSemNet* to learn DT policies in RL environments, both with discrete action sets (classification-*DTSemNet*) and with continuous action sets (regression-*DTSemNet*). We accomplish this by simply replacing the NN architecture with the *DTSemNet* architecture in RL learning pipelines (e.g., PPO [35]), to obtain DT policies having performance comparable to that of NNs. Experimentation on RL environments with *discrete* action sets of different sizes (up to 10 actions and 32 input dimensions) shows that *DTSemNet* matches NNs. Further, on the more complicated benchmarks, it significantly outperforms competing solutions producing DT policies, either obtained by gradient descent on architectures encoding DTs or by imitation learning of NNs. For *continuous* action space, provided that the number of input dimensions is limited, ICCT already provides performance matching NN by using regressors [31], although it generates less expressive axis-aligned DT. This suggests that very accurate decisions in internal nodes are not as important when regressions can be used in the leaves of the DT. We experimentally verified that in two environments with continuous action space: as with discrete actions, *DTSemNet* outperforms competing solutions, but the architecture is less important for performance.

Main contributions are as follows:

- We introduce the *DTSemNet* architecture, which we prove to be semantically equivalent with DTs.
- Learning a *DTSemNet* (-classification) with standard gradient descent corresponds exactly to learning a DT, without resorting to any approximations, unlike competing methods. Experimentally, this results in the most efficient DTs for classification tasks.
- We extend *DTSemNet* to regression tasks, using one STE approximation to combine the binary decisions with an output regression: experimentally, *DTSemNet*-regression is sometimes the best method, sometimes (a close) second.
- We explain how to use *DTSemNet* in the RL setting, again producing experimentally the best DT policies, by a large margin for discrete actions (*DTSemNet*-classification), and small margins for continuous action spaces (*DTSemNet*-regression).

2 Related Work

Non-Gradient-Based DT Training: There is a significant amount of approaches to learn DTs without relying on gradients. CART [10]

(and its extensions) is a well-known method for training (*axis-aligned*) DTs, which is based on splitting the dataset at each node based on certain features using a metric, such as entropy or Gini impurity. Concerning learning *oblique* DTs in this way, some methods have been proposed, such as Oblique Classifier 1 (OC1) [30] or GUIDE [29]. They typically do not produce well-performing oblique DTs since learning oblique DT is much harder than learning axis-aligned DTs. TAO is currently a state-of-the-art (SOTA) method that can be used to learn oblique DTs. It enhances the performance of DT obtained by CART (or random DT of given depth) by alternatively fine-tuning node parameters at specific depths [12, 45].

Concerning MIP formulations [6, 7] or Global EA-based search approaches, such as CRO-DT [15], they learn DTs by searching over various structures of DTs but at high computational costs, which is impractical for large DTs and datasets. The proposed *DTSemNet* overcomes these challenges by using gradient-descent to lower training time, which we confirm by comparing with CRO-DT [15], which proposes matrix encoding of (oblique) DTs to speed up training compared to previous EA-based methods, and produces axis-aligned DTs.

Gradient-Based DT Training: Numerous works propose to approximate DTs as soft-DTs to use gradient descent for learning DTs, where decision nodes typically use the Sigmoid function [44, 21, 43, 17, 38, 9, 37, 16, 33, 19, 24, 11, 42]. *Hardening* soft DTs, that is, transforming them into hard DTs by discretizing the probabilities induces severe inaccuracies [31]. More closely related to our work, DGT [2] represents (oblique) DTs as an NN-architecture using the (non-differentiable) sign activation function, resorting to *quantized* gradient descent to learn it, leveraging principles from training binarized NNs using STE [22]. Specifically, during forward propagation, nodes utilize a 0-1 step function, whereas, during backward propagation, nodes employ a piecewise linear function or some differentiable approximation (see [22]). Similarly, ICCT [31] learns (axis-aligned) DTs using NNs with the Sigmoid activation function and STEs. In all these works, the hard DTs that are produced are (slightly) different from the DT (soft DT or using STE), which is being optimized by gradient descent. In contrast, the *DTSemNet* architecture using ReLU activation functions allows standard gradient descent to be performed, and the output DT from *DTSemNet*-classification is exactly the same as the function optimized by gradient-descent, without approximation. Experiments confirm that it is more accurate in practice, significantly so for classification tasks.

Training DT policies in RL Setup: Hard or soft DT policies can be obtained via imitation learning [1], i.e., learning from expert policies, usually pretrained NNs [4, 23, 28, 8, 34, 39, 40, 14]. For instance, VIPER [4] imitates a Q-network (or policy network) by creating a dataset from collected samples and trains a DT using CART, with sample weightage assigned based on Q-values. In contrast, *DTSemNet* directly learns a hard oblique DT in RL (using PPO [35]). Other works, such as ProLoNet [36], that learn *soft* DTs using the RL framework, with the objective of initializing weights from expert humans. In contrast, *DTSemNet* learns a *hard* DT. ICCT [31] proposed an STE-based approach to learn axis-aligned DTs using gradient descent. By comparison, we can handle oblique trees, more expressive and accurate than axis-aligned DTs, especially for discrete actions.

3 The *DTSemNet* Architecture

We now describe our main contribution, namely *DTSemNet*, which encodes decision trees (*DT*) as deep neural *Net*works (DNNs) in a *Semantically* equivalent way. We prove the semantic equivalence in Theorem 1 thereafter.

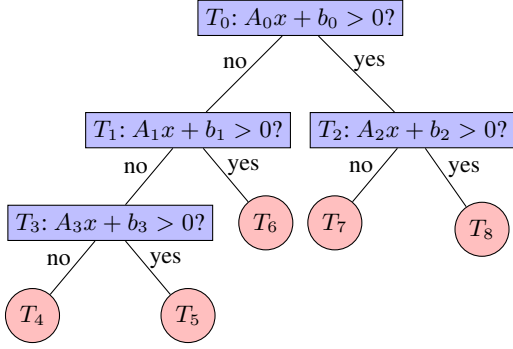


Figure 1. A Decision Tree T with 4 internal decision nodes T_0, \dots, T_3 and 5 leaves T_4, \dots, T_8 .

Consider a (binary) DT T over input space $X = \mathbb{R}^n$ of dimension n . It has a set of internal nodes T_0, \dots, T_k , and leaves T_{k+1}, \dots, T_m , with associated functions left, right : $\{T_0, \dots, T_k\} \rightarrow \{T_1, \dots, T_m\}$ associating each internal node with its respective left (or right) child. Each internal node T_i is also associated with a decision described by a (linear) vector $A_i \in \mathbb{R}^n \setminus \{(0, \dots, 0)\}$ and an (affine) component $b_i \in \mathbb{R}$, such that the decision from input $x \in X$ is to go to the right decision node iff $A_i x + b_i > 0$, and left if $A_i x + b_i < 0$ ¹. There is always a unique path between a node and the root. Using this path, we associate each leaf T_j of T with a set of decisions. For instance, in Figure 1, leaf T_5 is associated with decisions $-D_0, -D_1, D_3$, encoding the fact that T_5 is the right child of T_3 (it needs D_3 to be true, i.e. $A_3 x + b_3 > 0$), itself being the left child of node T_1 (and thus it needs $A_1 x + b_1 < 0$, i.e. $-A_1 x - b_1 > 0$, that we denote " $-D_1$ "), with T_1 the left child of root T_0 . Given an input $x \in X$, leaf T_j is selected if and only if all the associated decisions are true. Given an input $x \in X$, there is exactly one leaf for which all associated decisions are true, and for the other leaves, at least one associated decision is false.

Finding suitable values for A_i and b_i to make the DT T efficient is challenging. We explain in the following how to find such values using standard gradient descent algorithms. For that, we encode the DT T as the following NN architecture $DTSemNet$, which is semantically equivalent with T , and which can be converted back to an equivalent DT T' after learning. We now describe the $DTSemNet$ architecture:

- The input layer is made of n inputs x_1, \dots, x_n for the input domain X of dimension n , plus one input b for the biases, always set to weight 1.
- The first hidden layer is a layer with k nodes I_1, \dots, I_k , one per internal node T_1, \dots, T_k of the tree. It is fully connected with the initial layer. The weights to node I_i correspond one-to-one to the decisions D_i , that is to A_i for the weights from x_1, \dots, x_n and to b_i from b . The activation function is linear.
- The second hidden layer is made of $2k$ nodes \top_i, \perp_i , two per internal node $T_i, i \leq k$ of the tree, standing for decision D_i being true (\top_i), or $-D_i$ being true (\perp_i). The activation function is ReLU. The weights are fixed, with weight 1 from I_i to \top_i and -1 from I_i to \perp_i .
- The output is made of $m - k$ nodes L_{k+1}, \dots, L_m , one per leaf T_{k+1}, \dots, T_m of the decision tree.

The weights from \perp_i, \top_i to L_j can be 0 or 1, depending on the structure of the DT: L_j takes weight 0 from \top_i if leaf T_j is a left

¹ Note that the case where $A_i x + b_i = 0$ is not important: The set of initial points where $A_i x + b_i = 0$ for some i have zero measure and thus zero probability happening. We consider that input space X excludes such points.

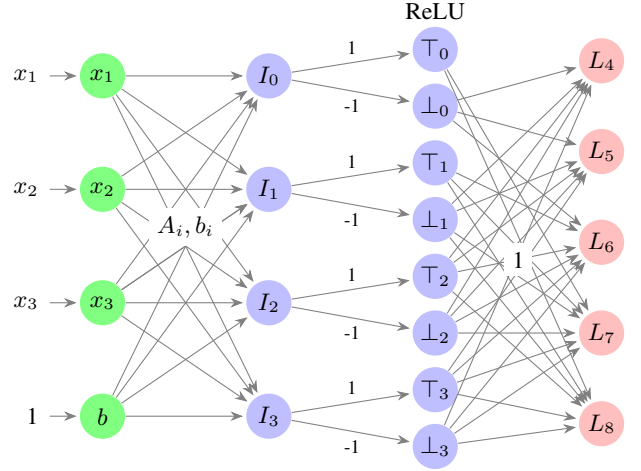


Figure 2. The $DTSemNet$ architecture corresponding to Decision Tree T in Figure 1, with $X = \mathbb{R}^3$

descendant of node T_i (it needs D_i to be false), otherwise it takes weight 1 from \top_i (that is if either it is a right descendant of node T_i or it is not a descendant of T_i at all). Similarly, it takes weight 0 from \perp_i if leaf T_j is a right descendant of node T_i . Otherwise, it takes weight 1 from \perp_i . The activation function is linear.

The $DTSemNet$ NN has $n + 1 + m + 2k$ nodes, it has depth 4 (2 linear activation functions and 1 ReLU activation function), with $k(n + 1)$ trainable weights, exactly the same number as the number of dimensions of decisions in the DT (the A_i 's and b_i 's). We display in Fig. 2 the $DTSemNet$ associated with the DT of Fig. 1.

The rationale behind the architecture is the following: for $\text{value}_x(X)$ the value of node X when the NN input is x , we want to ensure that $\text{value}_x(L_\ell) = \sum_j \text{value}_x(\perp_j) + \text{value}_x(\top_j)$ for the leaf T_ℓ such that all the associated decisions are true (thus selected by the DT), and $\text{value}_x(L_{\ell'}) < \sum_j \text{value}_x(\perp_j) + \text{value}_x(\top_j)$ for every other leaf $T_{\ell'} \neq T_\ell$, so that L_ℓ is selected by the $DTSemNet$. This allows to prove the semantic equivalence of the DT and the $DTSemNet$ architecture:

Theorem 1. *Considering the Neural Network $DTSemNet$ as a Classifier, taking the decision associated with the output L_i with the highest value using the standard argmax operation, gives exactly the same decision as the DT, i.e. select the (unique) leaf for which every associated decision is true. Equivalently, For all input $x \in X$, we have $\text{argmax}_j(\text{value}_x(L_j)) = L_i$ iff every decision associated with leaf T_i are true on input x .*

Proof. Take some $x \in X$. Let us start with a few important remarks. First, exactly one leaf T_i has all its associated decisions true (and is the leaf selected by the DT). For every other leaf $T_j, j \neq i$ there is at least one decision D_j^x (some D_ℓ or some $-D_\ell$) associated with T_j , which is false.

Also, trivially, we have that exactly one of $\text{value}_x(I_i), -\text{value}_x(I_i)$ is strictly positive, and the other one is strictly negative. Thus, because of the definition of ReLU, we have that exactly one of $\text{value}_x(\perp_i), \text{value}_x(\top_i)$ is 0, and the other is strictly positive. In particular, $0 \leq \text{value}_x(\perp_i), \text{value}_x(\top_i) \leq \text{value}_x(\perp_i) + \text{value}_x(\top_i) = \max(\text{value}_x(\perp_i), \text{value}_x(\top_i))$.

For all leaf T_ℓ , we have $\text{value}_x(L_\ell) \leq \sum_j \text{value}_x(\perp_j) + \text{value}_x(\top_j)$ by construction: for all node T_j not in the path between the root T_0 and T_ℓ , these two weights are input into L_k by construc-

tion (with weight 1), and for nodes T_j that are on the path, one of them is input to L_k , and the other one is positive or null.

Let ℓ be such that every decision associated with leaf T_ℓ is true on input x . There is exactly one such leaf T_ℓ . We show that $\text{value}_x(L_\ell) = \sum_j \text{value}_x(\perp_j) + \text{value}_x(\top_j)$. Indeed, as above, the only nodes \perp_j (resp. \top_j) that are not direct input of L_ℓ corresponds to decisions that are not true, and thus $\text{value}_x(\perp_j) = 0$ (resp. $\text{value}_x(\top_j) = 0$). Hence, the equality is true.

For every other leaf $L_{\ell'} \neq L_\ell$, we show that $\text{value}_x(L_{\ell'}) < \text{value}_x(L_\ell)$. Indeed, let $D_{\ell'}$ be a decision associated with $T_{\ell'}$ that is not correct under input x , say it is decision $-D_i$, with D_i true, that is $\text{value}_x(\top_i) > 0$. As there is no weight from \top_i on $L_{\ell'}$ by definition of $DTSemNet$, $\text{value}_x(L_{\ell'}) \leq (\sum_j \text{value}_x(\perp_j) + \text{value}_x(\top_j)) - \text{value}_x(\top_j) < \text{value}_x(L_\ell)$. That is, $\text{argmax}_j(\text{value}_x(L_j)) = L_\ell$ iff $i = \ell$, and we are done. \square

While CART [10] grows the DT incrementally, $DTSemNet$ will learn a fixed DT architecture. The simplest way to fix the architecture is to pick a depth n in regards to the difficulty of the task (number of features, etc., e.g., $n = 5$ for few ≈ 5 features) and consider the balanced binary tree with 2^n leaves ($2^5 = 32$). This is as simple as choosing the number of neurons of an NN. Notice that unlike DGT [2], $DTSemNet$ can also directly handle any unbalanced trees, e.g., for optimizing the number of nodes.

$DTSemNet$ -classification: We now explain how the $DTSemNet$ architecture selects the class for a classification task. Let T_{class} a balanced DT with one leaf for each class. For instance, for *balance scale*, there are 3 classes *left*, *right* and *balanced*. The balanced DT is provided below (see Figure 3). The last (here 2) levels of the DT to learn are then replaced by T_{class} , letting the gradient descent learn the most efficient way to play actions from each region of the space. We transform this tree using the $DTSemNet$ architecture: there will thus be several nodes L_j corresponding to leaves T_j associated with the same class. To recover the standard classification framework in $DTSemNet$, we add a final layer with ℓ nodes C_1, \dots, C_ℓ , one for each of the ℓ classes $\alpha_1, \dots, \alpha_\ell$. The input weight to C_i is 1 from L_j if leaf T_j is associated with action α_i , and 0 otherwise. The MaxPool operation is performed at the node C_i , which then outputs the maximum value of associated inputs to node C_i . Applying Theorem 1, if leaf T_j is selected by the DT on input $x \in X$, then the value $\text{value}_x(L_j)$ is maximal, and for α_i the class associated with T_j , output node C_i will get the highest value among all output nodes C_j . Class α_i will be thus selected by $DTSemNet$, semantically equivalent to the DT T .




Figure 3. Decision Tree T_{class} for Balance scale.

$DTSemNet$ -regression: Usual DT for regression tasks learned by CART [10] (axis-aligned DTs) and by DGT [2] (oblique DTs) are as follows: each leaf T_j is associated with one scalar value α_j , and the DT T outputs α_ℓ for T_ℓ the leaf selected by the decisions of T on input x . In CART, this parameter α_ℓ is the average of the training values corresponding to inputs ending in T_ℓ , resulting in a DT T that does not generalize, even when learning a simple linear function.

We now describe regression-DTs, which are more expressive, as used in TAO(-linear) [45] and ICCT [31]. Every leaf T_j of the (regression)-DT T is associated with linear parameter vector $\theta_j \in \mathbb{R}^n$ and an affine scalar $\alpha_j \in \mathbb{R}$. The output of regression-DT T on input $x \in X$ is $\theta_\ell \cdot x + \alpha_\ell$, where T_ℓ is the leaf selected on input x by the decisions of T . A regression-DT would easily generalize 2 or more data points from a training set into the linear function. Regression-DTs actually correspond to a piecewise linear regression, where the different linear pieces are described by the parameters at the leaves, and the switch between pieces by the decisions of the DT.

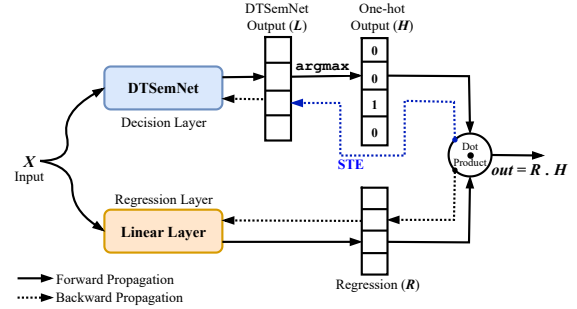


Figure 4. $DTSemNet$ -regression adapted for regression tasks.

We now adapt $DTSemNet$ to learn such a regression-DT. We learn simultaneously: (a) the parameters of the linear regression at every leaf and (b) the decisions leading to the most appropriate leaf and associated linear regression. The first hidden (linear) layer is extended with $m - k$ nodes $(R_j)_{k < j \leq m}$, one per leaf T_j , with trainable weights θ_j from input x to R_j and α_j from input 1 to R_j : We have $\text{value}_x(R_j) = \theta_j \cdot x + \alpha_j$. In order to select $\text{value}_x(R_\ell)$ for L_ℓ with the highest value in $DTSemNet$ (that is T_ℓ selected by the DT by Theorem 1), we add a one-hot layer with $m - k$ nodes $(H_j)_{k < j \leq m}$ using the argmax function (see Figure 4): node H_ℓ has value 1 for L_ℓ with maximal value, and H_j has value 0 for $j \neq \ell$; finally an output node out computes the scalar product between $(\text{value}_x(H_j))_{k < j \leq m}$ and $(\text{value}_x(R_j))_{k < j \leq m}$, obtaining $\text{value}_x(out) = \text{value}_x(R_\ell) = \theta_j \cdot x + \alpha_j$. In end-to-end training, the argmax operation makes the entire architecture non-differentiable, so STE [5, 22] is used to bypass the gradient through argmax operation during backpropagation. The STE approximation is only used for $DTSemNet$ -regression (not -classification), and only at the output node, in contrast to STE approximation, which is used at every single node in DGT and ICCT [2, 31], and even for classification.

$DTSemNet$ for RL: To tackle RL tasks, we use standard Deep RL framework (e.g., PPO [35]), simply replacing the NN by the $DTSemNet$ architecture and running the gradient descent from the RL framework. For environments with discrete action space, we use $DTSemNet$ -classification, with one class per action. For environments with continuous action space, we use $DTSemNet$ -regression, with one linear regression per action dimension (e.g., 2 for a continuous lunar lander, one for horizontal, and one for vertical engines).

Comparison with other encodings of DTs into NNs: DGT [2] uses the Sign activation function at each layer to produce (hard) oblique DTs. As Sign has no gradient (unlike ReLU used in $DTSemNet$), DGT resorts to *quantized* gradient descent, using an STE approximation process at every node. In regression tasks, DGT produces a scalar for each leaf instead of a linear regression for $DTSemNet$ -regression. The ICCT architecture [31] generates axis-aligned DTs for RL tasks by associating every leaf with the product of the weights of the edges

Table 1. Percentage accuracy (higher is better) on **classification tasks** with the most accurate height ≤ 4 , as tested in [15]. For each dataset, we provide the number of features N_f , classes N_c and training samples N_s . Averaged accuracy \pm std is reported over 100 runs. The best-performing height is reported in parentheses for each architecture. We run the DTSemNet and DGT experiments, while the results for TAO, CART and CRO-DT are copied from [15].

Dataset	N_f, N_c, N_s	<i>DTSemNet</i>	DGT [2]	TAO [12]	CART [10]	CRO-DT [15]
Balance Scale	4, 3,625	90.2 \pm 2.2 (2)	88.6 \pm 1.7 (4)	77.4 \pm 3.1 (4)	74.9 \pm 3.6 (4)	77.8 \pm 3.0 (4)
Banknote Auth	4, 4, 1372	99.8 \pm 0.4 (3)	99.8 \pm 0.4 (3)	96.6 \pm 1.3 (4)	93.6 \pm 2.2 (4)	95.2 \pm 2.2 (4)
Blood Transfusion	4, 2,784	78.5 \pm 1.7 (2)	78.3 \pm 2.4 (4)	76.9 \pm 2.1 (3)	77.1 \pm 1.8 (4)	76.1 \pm 1.9 (3)
Acute Inflamm. 1	6, 2, 120	100 \pm 0.0 (2)	100 \pm 0.0 (4)	99.7 \pm 1.2 (4)	100 \pm 0.0 (3)	100 \pm 0.0 (2)
Acute Inflamm. 2	6, 2, 120	100 \pm 0.0 (2)	100 \pm 0.0 (4)	99.0 \pm 2.6 (2)	99.0 \pm 2.6 (2)	100 \pm 0.0 (2)
Car Evaluation	6, 4, 1728	93.3 \pm 2.2 (4)	92.1 \pm 2.4 (4)	84.5 \pm 1.5 (4)	84.3 \pm 1.4 (4)	86.1 \pm 1.3 (4)
Breast Cancer	9, 2, 683	97.2 \pm 1.3 (2)	97.2 \pm 1.2 (2)	94.7 \pm 1.6 (3)	94.7 \pm 1.7 (3)	95.5 \pm 1.8 (2)
Avila Bible	10, 12, 10430	62.2 \pm 1.4 (4)	59.7 \pm 1.8 (4)	55.8 \pm 0.8 (4)	54.0 \pm 1.3 (4)	59.6 \pm 0.7 (4)
Wine Quality Red	11, 6, 1599	58.6 \pm 2.2 (3)	56.6 \pm 1.4 (4)	56.9 \pm 2.5 (4)	55.9 \pm 2.3 (4)	55.8 \pm 2.2 (2)
Wine Quality White	11, 7, 4898	53.5 \pm 1.4 (4)	52.1 \pm 1.6 (4)	52.3 \pm 1.4 (4)	52.0 \pm 1.3 (4)	51.4 \pm 1.2 (2)
Dry Bean	16, 7, 13611	91.4 \pm 0.5 (4)	89.0 \pm 1.6 (4)	83.2 \pm 1.5 (4)	80.5 \pm 1.9 (4)	77.9 \pm 4.7 (4)
Climate Crashes	18, 2, 540	92.9 \pm 1.4 (2)	92.4 \pm 2.4 (3)	90.6 \pm 2.2 (3)	91.8 \pm 1.8 (4)	91.5 \pm 2.0 (2)
Conn. Sonar	60, 2, 208	82.1 \pm 5.1 (4)	80.8 \pm 5.3 (4)	70.9 \pm 5.8 (4)	70.6 \pm 6.6 (4)	71.7 \pm 6.7 (4)
Optical Recognition	64, 10, 3823	93.3 \pm 1.0 (4)	91.9 \pm 1.0 (4)	64.6 \pm 6.5 (4)	53.2 \pm 3.2 (4)	65.2 \pm 2.0 (4)

leading to that leaf, passing in logarithmic space to handle the explicit multiplications. The Sigmoid activation is used on each decision, incurring an approximation of the decision as a soft DT. The soft DT is crisped at each step of the RL process into a hard (axis-aligned) DT, using an STE to backpropagate through a non-differentiable function (Heaviside step function), which *DTSemNet*-classification avoids.

4 Experimental Evaluation

In this section, we evaluate the performance of *DTSemNet*, comparing it with competing methodologies learning hard DTs. Firstly, we consider supervised learning setups using multiple benchmark multi-class classification and regression datasets, on which we compare the accuracy on test data with the SOTA non-greedy method TAO [12, 45], the SOTA gradient descent-based method DGT [2], both learning oblique DTs, as well as CRO-DT [15] for global-searches and CART as the standard for greedy algorithms, both learning axis-aligned DTs. The relative training times are reported for benchmarks for which it is available (notice that we cannot run TAO on our hardware as it is not openly available). Further, to understand the impact on the generalization of different architectures using gradient descent, we leverage insights from the loss landscape [27].

Lastly, we consider RL environments, both with discrete actions and continuous action sets. We compare *DTSemNet* with DGT, both of which generate oblique DT-policies, ICCT, which generates axis-aligned DT-policies, all three through gradient descent, as well as VIPER, which generates axis-aligned DT-policies through imitation learning of an NN policy generated by Deep RL, which we also report as a baseline.

We implemented *DTSemNet* and conducted all experiments using Python and PyTorch. Our testing platform has 8 CPU cores (AMD 75F3, Zen 3 architecture), 128 GB of RAM, and a 2 GB GPU (NVIDIA Quadro P620). The supplementary materials [32] provide more details regarding datasets, hyperparameters, etc., and our code is available at <https://github.com/CPS-research-group/dtsemnet>.

Classification Tasks (small DTs): We first considered the 14 classification tabular datasets used in CRO-DT [15]. Global searches such as CRO-DT [15] are efficient only for small DTs (here up to depth 4, i.e. 32 nodes). We sort the benchmarks by the number of features, as it is more challenging to decide over more features, all the more so with small DTs. We report in Table 1 the (average) score over 100 DTs learned with different seeds for the most accurate height (up to 4).

In every single benchmark, *DTSemNet* produces the most accurate

DTs, with the biggest difference in *Dry Beans*, reducing the classification error from 11% to 8.6%. DGT is the second best, except in the 2 wine quality benchmarks, where TAO outperforms it. DGT usually comes close to *DTSemNet*, which is not surprising as the idea is similar, though the approximations used (STEs, quantized gradient descent) make DGT perform on average 0.85% worse than *DTSemNet*. Interestingly, the advantages grow to 1.3% on the 4 benchmarks with the most features (the hardest ones). Further, DGT needs larger trees than *DTSemNet* to get its best results (in 6 out of 14 benchmarks). Compared with non gradient-based learning, the advantage of *DTSemNet* is very tangible, 5.5% on average, growing to 12% on the 4 benchmarks with the most features. The highest difference is in *Optical Recognition*, reducing the classification error from 34.8% to just 6.7%.

Table 2. Training times in seconds (lower is better), and (avg accuracy % (higher is better)). The MNIST training time from the non-publicly available TAO is quoted from [12], while we train other architecture on a similar compute configuration. The tree height is 8 for MNIST and 4 for DryBean.

Dataset	<i>DTSemNet</i>	DGT [2]	TAO [3]	CRO-DT [15]
MNIST	306 (96.1)	288 (94.0)	1200 (95.0)	4659 (58.2)
DryBean	4.4 (91.4)	3.8 (89.0)	NA (83.2)	1300 (77.9)

Classification Tasks (deeper DTs): We now considered the original 8 classification tasks reported in TAO [12] and reported in [2]. They use tabular datasets, except for MNIST, which is a small-sized image dataset. We used the fixed tree height specified in [12] (also reused in [2]). We first report in Table 2 the training time for the architecture on MNIST, as [12] reports this number for TAO. We train other architectures on comparable computing configurations. We provide the accuracy number obtained for reference. We also report the training time for the simpler DryBean, for comparison sake, except for TAO, for which this is not available. First, CRO-DT for MNIST takes much longer to train (number of generations set to the default 4k) than other architectures while having very low accuracy numbers (< 60% instead of > 90%), the reason why we do not consider it for these benchmarks with deeper DTs (indeed, [15] does not report CRO-DT results for these benchmarks). The training times of DGT and *DTSemNet* are almost identical due to their architectural similarities. As expected, training with gradient-based methods is observed to be significantly faster than non-gradient learning methods.

We sort the benchmarks by the height of the DTs, which is a good indication of the complexity of the tasks. We report in Table 3 the (av-

Table 3. Percentage accuracy (higher is better) on **classification tasks** for the datasets reported in [12]. For each dataset, we provide the number of features N_f , classes N_c , and training samples N_s . All methods use the tree height fixed in [2] (except CART, which has no predefined height). Averaged accuracy \pm std is reported over 10 runs. The results for DGT and CART are from [2], and the results of TAO are from [12].

Dataset	N_f, N_c, N_s	Height	<i>DTSemNet</i>	DGT [2]	TAO [12]	CART [10]
Protein	357, 3, 14895	4	68.60 \pm 0.22	67.80 \pm 0.40	68.41 \pm 0.27	57.53 \pm 0.00
SatImages	36, 6, 3104	6	87.55 \pm 0.59	86.64 \pm 0.95	87.41 \pm 0.33	84.18 \pm 0.30
Segment	19, 7, 1478	8	96.10 \pm 0.53	95.86 \pm 1.16	95.01 \pm 0.86	94.23 \pm 0.86
Pendigits	16, 10, 5995	8	97.02 \pm 0.32	96.36 \pm 0.25	96.08 \pm 0.34	89.94 \pm 0.34
Connect4	126, 3, 43236	8	82.03 \pm 0.39	79.52 \pm 0.24	81.21 \pm 0.25	74.03 \pm 0.60
MNIST	780, 10, 48000	8	96.16 \pm 0.14	94.00 \pm 0.36	95.05 \pm 0.16	85.59 \pm 0.06
SensIT	100, 3, 63058	10	84.29 \pm 0.11	83.67 \pm 0.23	82.52 \pm 0.15	78.31 \pm 0.00
Letter	16, 26, 10500	10	89.19 \pm 0.29	86.13 \pm 0.72	87.41 \pm 0.41	70.13 \pm 0.08

Table 4. Average RMSE results (lower is better) on **regression tasks**, \pm std over 10 runs. The number N_f of features and N_s of training samples are provided for each dataset. DGT-linear is a modification we implemented from DGT [2] to add regressors at the leaves. The 5 first datasets are from [45]. The height for DTs with regressors at the leaves (*DTSemNet*, DGT-Linear, TAO-linear) is fixed to Height, following [45]. DGT has fixed scalars at the leaves instead of regressors, thus it needs deeper DTs to achieve reasonable accuracy - the depth used in [2] is reported inside (), while CART height is not fixed. The last two datasets are from [2], for which TAO-linear results are not available. The reported results for DGT and CART are taken from [2]; TAO-Linear results from [45].

Dataset	N_f, N_s	Height	<i>DTSemNet</i>	DGT-Linear	DGT [2]	TAO-Linear [45]	CART [10]
Abalone	10, 2004	5	2.135 \pm 0.03	2.144 \pm 0.03	2.15 \pm 0.026 (6)	2.07 \pm 0.01	2.29 \pm 0.034
Comp-Active	21, 3932	5	2.645 \pm 0.18	2.645 \pm 0.15	2.91 \pm 0.149 (6)	2.58 \pm 0.02	3.35 \pm 0.221
Ailerons	40, 5723	5	1.66 \pm 0.01	1.67 \pm 0.017	1.72 \pm 0.016 (6)	1.74 \pm 0.01	2.01 \pm 0.00
CTSlice	384, 34240	5	1.45 \pm 0.12	1.78 \pm 0.25	2.30 \pm 0.166 (10)	1.16 \pm 0.02	5.78 \pm 0.224
YearPred	90, 370972	6	8.99 \pm 0.01	9.02 \pm 0.025	9.05 \pm 0.012 (8)	9.08 \pm 0.03	9.69 \pm 0.00
PDBBind	2052, 9013	2	1.33 \pm 0.017	1.34 \pm 0.013	1.39 \pm 0.017 (6)	NA	1.55 \pm 0.00
Microsoft	136, 578729	5	0.766 \pm 0.00	0.766 \pm 0.00	0.772 \pm 0.00 (8)	NA	0.771 \pm 0.00

erage) score over 10 DTs learned with different seeds. **Discussion:** As in Table 1, *DTSemNet* produces the most accurate DTs in every single benchmark. With deeper trees, TAO recovers, and it is the second best 5 times vs 3 times for DGT. The advantage of *DTSemNet* over TAO is 1% on average, growing to 1.4% on the 4 benchmarks with deeper trees. The highest difference is in *Letter*, reducing the classification error from 12.6% to 10.8% and from 3.9% to 3% in *Pendigits*. The advantage of *DTSemNet* over DGT is 1.4% on average, growing to 2.1% on the 4 benchmarks with deeper trees, larger than for smaller trees from Table 1. The largest difference is in *Letter*, reducing the classification error from 13.9% to 10.8%, and from 6% to 3.9% in MNIST. The Friedman-Nemenyi test across all classification benchmarks, with a significance level of 0.05, shows that *DTSemNet* is significantly better than every other method with the average rank of 1.11 for *DTSemNet*, 2.25 for DGT, 2.86 for TAO, 3.60 for CRO-DT, and 3.77 for CART.

To understand the influence of approximations used in DGT (quantized gradient descent and STEs), we consider the loss landscape [27], which displays the effect of architecture choices on generalization. We consider MNIST, for which DGT makes much more classification

errors than *DTSemNet*. Figure 5 displays the loss landscape around the final trained parameters along two random vector directions. A flatter loss landscape indicates better generalization capability [27]. As shown in Figure 5, the loss landscape of *DTSemNet* is very flat compared with DGT. This suggests that (approximation-free) *DTSemNet* allows better generalization on classification tasks than DGT.

Regression Tasks: We consider the 5 regression datasets from [45], plus the two extra from [2] (for which TAO-linear results are not available). While TAO-linear [45] generates DT with regressors at the leaves, similarly as *DTSemNet*, allowing them to generalize, DGT [2] only learns scalar at the leaves, similarly as CART, and thus is less efficient (and cannot generalize), needing deeper DTs to get acceptable accuracy. To understand the impact of regressors at the leaves, we implemented *DGT-linear*, a modification of DGT using regressors at the leaves, in the same way as in *DTSemNet* and in ICCT [31]. Architectures with regressors at the leaves, namely *DTSemNet*(-regression), DGT-linear, and TAO-linear, use the same fixed height following [45], while DGT (without regressors) uses deeper trees (as reported in [2]) and CART is not restricted. The results of the regression datasets are presented in Table 4.

DTSemNet-regression is sometimes second behind TAO-linear (3 benchmarks), and best for the 4 other benchmarks. *DTSemNet* is consistently better than (original) DGT, by 10% of RMSE on average over the 7 benchmarks, with 50% better RMSE scores for CTSlice, even though DGT uses larger (hard oblique) DTs than *DTSemNet*. Compared with DGT-linear with regressors at the leaves (our adaptation), *DTSemNet* is still consistently better or tied, but the advantage is reduced to 3.7% of RMSE on average over the 7 benchmarks, with up to 23% better scores (CTSlice). Two third of the advantage of *DTSemNet* over DGT can be attributed to the regressors at the leaves, but still, a meaningful one-third of the advantage can be attributed to reducing the approximations, with only 1 STE call for *DTSemNet*-regression instead of n STE calls for DGT and DGT-linear, where n is the height of the DT. Overall, CTSlice is sensitive to approximations.




Figure 5. The loss landscape of *DTSemNet* and DGT for MNIST, when varying the parameters around the trained parameters along two random directions. The flatter the loss landscape, the better the generalization.

Table 5. Average rewards (higher is better) on **RL tasks** \pm std over 5 policies generated with different learning seeds. The policies are evaluated over 100 episodes. We report the number N_f of features and N_a of actions for each environment. The first four environments have discrete actions, and the bottom two continuous actions, for which we test both the original (*scalar*) version of DGT and DGT-linear (*linear*). For the 5 first environments the *Height* is fixed for all architectures. For Bipedal Walker, DGT(-linear) performed much better with deeper DTs, while ICCT performed much better with shallower DTs (*DTSemNet* was not much impacted by the height). For Bipedal Walker, the heights used for DGT(-linear) and for ICCT providing the best rewards are reported in parentheses.

Environments	N_f, N_a	<i>Height</i>	<i>DTSemNet</i>	Deep RL	DGT [2]	ICCT [31]	VIPER [4]
CartPole	4, 2	4	500 \pm 0	500 \pm 0	500 \pm 0	496 \pm 0.3	499.95 \pm 0.05
Acrobot	6, 3	4	-82.5 \pm 1.05	-84 \pm 0.84	-83.1 \pm 1.88	-88.6 \pm 1.77	-83.92 \pm 1.59
LunarLander	8, 4	5	252.5 \pm 3.9	245 \pm 14.5	183.6 \pm 14.6	-85 \pm 16.3	86.73 \pm 7.93
Zerglings	32, 30	6	15.54 \pm 2.07	10.47 \pm 0.23	8.21 \pm 1.03	9.40 \pm 1.10	10.61 \pm 0.46
Cont. LunarLander	8, 2 dim.	4	277.24 \pm 2.09	276.12 \pm 1.45	<i>scalar</i> : 131.92 \pm 51.49 <i>linear</i> : 267.9 \pm 9.37	255.57 \pm 4.19	NA
Bipedal Walker	24, 4 dim.	7	314.98 \pm 3.35	315.3 \pm 6.91	<i>scalar</i> : 78.33 \pm 57.19 (8) <i>linear</i> : 244.5 \pm 61.84 (8)	301.34 \pm 3.09 (6)	NA

RL Tasks: We experimented with four cyber-physical environments with *discrete actions* with a limited number of physical features (≤ 32): three environments from OpenAI Gym, namely CartPole (4 features, 2 actions), Acrobot (6 features, 3 actions) and LunarLander (8 features, 4 actions), as well as a larger SC-II environment, namely FindandDestroyZerglings (32 features, 30 actions) [41]. We also experimented with two OpenAI environments with continuous actions, namely the continuous version of Lunar Lander (8 features, 2 action dimensions – horizontal and vertical thruster) and Bipedal Walker (24 features, 4 actions).

We fix the height of the DT that is learned by the different architectures and compare for reference with Deep RL with an NN with the same number of learned parameters, which is also used as an expert from which to perform imitation learning for VIPER [4]. ICCT and VIPER generate (hard) axis-aligned DT policies, while *DTSemNet* and DGT generate (hard) oblique DT policies. Further, for continuous action space, the regression version of each architecture is used, with 2 variants for DGT: the original with scalars at the leaves and our modification with linear regressions (similar to *DTSemNet* and ICCT). Notice that VIPER does not handle continuous action space.

We train 5 policies for each case, using different seeds for the environments. We use the standard PPO [35] for discrete actions and SAC [20] for continuous actions from the library StableBaseLine3. To evaluate the policies, we run each policy on the same 100 seeds and compute the mean reward. Table 5 reports the mean reward values averaged over the 5 different policies for each architecture.

Overall, *DTSemNet* is competitive with Deep RL generating NN-policies on these environments with a limited number of features (≤ 32). Further, *DTSemNet* consistently generates the most efficient DT-policies. The more complex the environment, the larger the lead:

- On the simplest CartPole ($N_f = 4$), it is tied for first.
- for Acrobot ($N_f = 6$), it leads by a couple of percents.
- for (discrete action) LunarLander ($N_f = 8$), it has a very sizable lead compared to other DT architectures ($\geq 28\%$). Notice that VIPER can perform closer (> 200 reward) to *DTSemNet*, but it needs very large DTs (> 1300 leaves) vs 32 leaves for *DTSemNet*.
- for Zerglings ($N_f = 32$), the lead is even more sizable ($\geq 39\%$).

Continuous action spaces are easier to handle than the discrete case: For Lunar Lander, each architecture with regressors at the leaves has a higher mean reward in the continuous case than in the discrete version, which is expected as it has access to precise continuous actions over 2 dimensions rather than a very coarse choice between 4 discrete actions. Further, rewards are much more similar between the architectures, confirming that the architecture is not as important as in the discrete version since the policies can use a linear actuator with respect to the

input state rather than a fixed action as in the discrete version. Even Axis-Aligned DTs come close to *DTSemNet*, performing extremely poorly in the discrete version. Again, our linear adaptation for DGT is much more accurate than the original version from [2], although the gap with *DTSemNet* is still large for Bipedal Walker.

Regarding the impact of different approximations used by DGT and ICCT, they both use STEs. The Sigmoid (and multiplications) used in ICCT is inefficient for Lunar Lander, particularly for the discrete action variant. Quantized gradient descent (DGT due to non-differentiable Sign function) is particularly inefficient in the Zerglings environment and for Bipedal Walker, with poorer results than even the less expressive axis-aligned DTs generated by ICCT. Further, for Bipedal Walker, we observed significant variability in the performance of DGT and ICCT when the tree depth varied, unlike *DTSemNet*, which remained consistently > 300 for depth 6,7,8: the performance of DGT-linear at a depth of 6,7,8 are 112, 163, 244, it needs deeper DTs. For ICCT, the performance at depth 6,7,8 is 301, 177, 112, severely decreasing with depth, which should not be happening, another proof that approximations hurt the accuracy.

5 Conclusion

We introduced *DTSemNet*, an architecture semantically equivalent to oblique DTs. This architecture enables oblique DTs to be learned using vanilla gradient descent. We demonstrate its performance on supervised classification and regression datasets and RL tasks. *DTSemNet* consistently generates DTs more accurately (or as accurately in easier benchmarks) than competing architecture learning DTs using gradient descent. This is because *DTSemNet*-classification uses no approximation, unlike its competitors, and because *DTSemNet*-regression uses fewer approximations. Further, compared to non-gradient-based methodologies (greedy, non-greedy and global search) for learning DTs, *DTSemNet* is significantly faster. *DTSemNet*-classification outperforms the best of these methods, reducing the errors by $> 10\%$ on harder classification tasks, while the accuracy of *DTSemNet*-regression is competitive with the SOTA.

Limitations: *DTSemNet* is a DT, making it unsuitable for high-dimensional inputs like images, where DTs struggle with complex shapes and require many leaves, negating their benefits.

Future Work: For *DTSemNet*, the choice of the height of the DT is treated as a hyperparameter, similar to the choice of the number of layers in an NN, in contrast to methods (e.g. CART [10]) that grow trees height. For future work, we will consider developing a regression architecture that does not rely on STE approximations and introduce differentiable methods for tree pruning and adaptive growth.

Acknowledgements

This research was conducted as part of the DesCartes program and was supported by the National Research Foundation, Prime Minister’s Office, Singapore, under the Campus for Research Excellence and Technological Enterprise (CREATE) program. It was also supported in part by AISG Research Grant #AISG2-RP-2020-017. The computational work for this research was partially performed using resources provided by the NSCC, Singapore.

References

- [1] P. Abbeel and A. Y. Ng. Apprenticeship learning via inverse reinforcement learning. In *Proceedings of the twenty-first international conference on Machine learning*, page 1, 2004.
- [2] G. A. K. (ajaykrishna karthikeyan), N. Jain, N. Natarajan, and P. Jain. Learning accurate decision trees with bandit feedback via quantized gradient descent. *Transactions on Machine Learning Research*, 2022.
- [3] S. Arora, N. Cohen, and E. Hazan. On the optimization of deep networks: Implicit acceleration by overparameterization. In *International conference on machine learning*, pages 244–253. PMLR, 2018.
- [4] O. Bastani, Y. Pu, and A. Solar-Lezama. Verifiable reinforcement learning via policy extraction. *Advances in neural information processing systems*, 31, 2018.
- [5] Y. Bengio, N. Léonard, and A. Courville. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*, 2013.
- [6] K. P. Bennett and J. A. Blue. Optimal decision trees. *Rensselaer Polytechnic Institute Math Report*, 214(24):128, 1996.
- [7] D. Bertsimas and J. Dunn. Optimal classification trees. *Machine Learning*, 106:1039–1082, 2017.
- [8] T. Bewley and J. Lawry. Tripletree: A versatile interpretable representation of black box agents and their environments. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35(13), pages 11415–11422, 2021.
- [9] G. Biau, E. Scornet, and J. Welbl. Neural random forests. *Sankhya A*, 81(2):347–386, 2019.
- [10] L. Breiman, J. Friedman, R. Olshen, and C. Stone. *Cart. Classification and regression trees*, 1984.
- [11] S. Bulò and P. Kotschieder. Neural decision forests for semantic image labelling. In *2014 IEEE Conference on Computer Vision and Pattern Recognition*, pages 81–88, 2014. doi: 10.1109/CVPR.2014.18.
- [12] M. A. Carreira-Perpinán and P. Tavallali. Alternating optimization of decision trees, with application to learning sparse oblique trees. *Advances in neural information processing systems*, 31, 2018.
- [13] J. H. Chen and S. M. Asch. Machine learning and prediction in medicine—beyond the peak of inflated expectations. *The New England journal of medicine*, 376(26):2507, 2017.
- [14] Y. Coppens, K. Efthymiadis, T. Lenaerts, A. Nowé, T. Miller, R. Weber, and D. Magazzeni. Distilling deep reinforcement learning policies in soft decision trees. In *Proceedings of the IJCAI 2019 workshop on explainable artificial intelligence*, pages 1–6, 2019.
- [15] V. G. Costa, S. Salcedo-Sanz, and C. E. Pedreira. Efficient evolution of decision trees via fully matrix-based fitness evaluation. *Applied Soft Computing*, 150:111045, 2024. ISSN 1568-4946. doi: <https://doi.org/10.1016/j.asoc.2023.111045>.
- [16] Z. Ding, P. Hernandez-Leal, G. W. Ding, C. Li, and R. Huang. Cdt: Cascading decision trees for explainable reinforcement learning, 2021.
- [17] N. Frosst and G. Hinton. Distilling a neural network into a soft decision tree. *arXiv preprint arXiv:1711.09784*, 2017.
- [18] L. Grinsztajn, E. Oyallon, and G. Varoquaux. Why do tree-based models still outperform deep learning on typical tabular data? *Advances in neural information processing systems*, 35:507–520, 2022.
- [19] U. D. Gupta, E. Talvitie, and M. Bowling. Policy tree: Adaptive representation for policy gradient. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, AAAI’15, page 2547–2553. AAAI Press, 2015. ISBN 0262511290.
- [20] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International conference on machine learning*, pages 1861–1870. PMLR, 2018.
- [21] H. Hazimeh, N. Ponomareva, P. Mol, Z. Tan, and R. Mazumder. The tree ensemble layer: Differentiability meets conditional computation. In *International Conference on Machine Learning*, pages 4138–4148. PMLR, 2020.
- [22] I. Hubara, M. Courbariaux, D. Soudry, R. El-Yaniv, and Y. Bengio. Binarized neural networks. *Advances in neural information processing systems*, 29, 2016.
- [23] A. Jhunjhunwala, J. Lee, S. Sedwards, V. Abdelzad, and K. Czarnecki. Improved policy extraction via online q-value distillation. In *2020 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8, 2020. doi: 10.1109/IJCNN48605.2020.9207648.
- [24] P. Kotschieder, M. Fiterau, A. Criminisi, and S. R. Bulò. Deep neural decision forests. In *2015 IEEE International Conference on Computer Vision (ICCV)*, pages 1467–1475, 2015. doi: 10.1109/ICCV.2015.172.
- [25] H. Laurent and R. L. Rivest. Constructing optimal binary decision trees is np-complete. *Information processing letters*, 5(1):15–17, 1976.
- [26] G.-H. Lee and T. S. Jaakkola. Oblique decision trees from derivatives of relu networks. *arXiv preprint arXiv:1909.13488*, 2019.
- [27] H. Li, Z. Xu, G. Taylor, C. Studer, and T. Goldstein. Visualizing the loss landscape of neural nets. *Advances in neural information processing systems*, 31, 2018.
- [28] G. Liu, O. Schulte, W. Zhu, and Q. Li. Toward interpretable deep reinforcement learning with linear model u-trees. In *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2018, Dublin, Ireland, September 10–14, 2018, Proceedings, Part II 18*, pages 414–429. Springer, 2019.
- [29] W.-Y. Loh. Fifty years of classification and regression trees. *International Statistical Review*, 82(3):329–348, 2014.
- [30] S. K. Murthy, S. Kasif, S. Salzberg, and R. Beigel. Oc1: A randomized algorithm for building oblique decision trees. In *Proceedings of AAAI*, volume 93, pages 322–327. Citeseer, 1993.
- [31] R. R. Paleja, Y. Niu, A. Silva, C. Ritchie, S. Choi, and M. C. Gombolay. Learning interpretable, high-performing policies for continuous control problems. In *Robotics: Science and Systems*, 2022.
- [32] S. Panda, B. Genest, A. Easwaran, and P. Suganthan. Vanilla gradient descent for oblique decision trees. *arXiv preprint arXiv*, 2024.
- [33] W. Qiu and H. Zhu. Programmatic reinforcement learning without oracles. In *International Conference on Learning Representations*, 2021.
- [34] A. M. Roth, N. Topin, P. Jamshidi, and M. Veloso. Conservative q-improvement: Reinforcement learning for an interpretable decision-tree policy. *arXiv preprint arXiv:1907.01180*, 2019.
- [35] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [36] A. Silva and M. Gombolay. Encoding human domain knowledge to warm start reinforcement learning. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(6):5042–5050, 2021.
- [37] 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*, pages 1855–1865. PMLR, 2020.
- [38] R. Tanno, K. Arulkumaran, D. Alexander, A. Criminisi, and A. Nori. Adaptive neural trees. In *International Conference on Machine Learning*, pages 6166–6175. PMLR, 2019.
- [39] A. Verma, V. Murali, R. Singh, P. Kohli, and S. Chaudhuri. Programmatically interpretable reinforcement learning. In *International Conference on Machine Learning*, pages 5045–5054. PMLR, 2018.
- [40] A. Verma, H. Le, Y. Yue, and S. Chaudhuri. Imitation-projected programmatic reinforcement learning. *Advances in Neural Information Processing Systems*, 32, 2019.
- [41] O. Vinyals, T. Ewalds, S. Bartunov, P. Georgiev, A. S. Vezhnevets, M. Yeo, A. Makhzani, H. Küttler, J. Agapiou, J. Schrittwieser, et al. Starcraft ii: A new challenge for reinforcement learning. *arXiv preprint arXiv:1708.04782*, 2017.
- [42] A. Wan, L. Dunlap, D. Ho, J. Yin, S. Lee, H. Jin, S. Petryk, S. A. Bargal, and J. E. Gonzalez. Nbd: neural-backed decision trees. *arXiv preprint arXiv:2004.00221*, 2020.
- [43] Y. Yang, I. G. Morillo, and T. M. Hospedales. Deep neural decision trees. *arXiv preprint arXiv:1806.06988*, 2018.
- [44] V. Zantedeschi, M. Kusner, and V. Niculae. Learning binary decision trees by argmin differentiation. In *International Conference on Machine Learning*, pages 12298–12309. PMLR, 2021.
- [45] A. Zharmagambetov and M. Carreira-Perpinan. Smaller, more accurate regression forests using tree alternating optimization. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119, pages 11398–11408. PMLR, 2020.

Supplementary Material

This supplementary document expands upon the experimental results reported in the main paper, providing additional details about the datasets, splits, and hyperparameters. We have also included the source code, which will be made open-source upon this paper’s acceptance. First, in Section A, we provide the extended results of the main paper. Then, in Section B, we discuss the dataset used, its splits, and sources. Finally, in Section C, we detail the hyperparameters used in the experiment.

A Extended Results

In Table 6, the extended results of Table 1 in the main paper are provided, which includes the percentage accuracy of DT heights 2, 3, and 4 of all discussed DT learning methods. Similarly, for the RL environment Bipedal Walker, the average reward for heights 6, 7, and 8 are presented in Table 7. It can be observed that the performance of DGT and ICCT is impacted by changing the height; however, the performance of DTSemNet remains stable across different heights.

Table 6. Classification percentage accuracy (higher is better) with the most accurate height ≤ 4 , as tested in [15]. For each dataset, we provide the number of features N_f , classes N_c , and training samples N_s . Averaged accuracy \pm std is reported over 100 runs. The performance of height 2,3,4 is reported, and in case the number of leaves exceeds the number of features N_f , the results are excluded. We run the DTSemNet and DGT experiments, while the results for TAO, CART and CRO-DT are copied from [15].

Dataset	N_f, N_c, N_s	Height	<i>DTSemNet</i>	DGT [2]	TAO [12]	CART [10]	CRO-DT [15]
Balance Scale	4, 3,625	2	90.2 \pm 2.2	86.4 \pm 3.7	66.4 \pm 3.6	64.7 \pm 3.5	69.2 \pm 2.34
		3	89.8 \pm 2.1	88.0 \pm 1.8	71.7 \pm 2.9	68.8 \pm 3.2	73.0 \pm 2.6
		4	90.2 \pm 2.2	88.6 \pm 1.7	77.4 \pm 3.1	74.9 \pm 3.6	77.8 \pm 3.0
Banknote Auth	4, 4, 1372	2	99.0 \pm 0.5	99.5 \pm 0.8	89.8 \pm 1.6	89.4 \pm 1.7	90.4 \pm 2.0
		3	99.8 \pm 0.4	99.8 \pm 0.4	94.9 \pm 1.4	91.8 \pm 2.2	92.4 \pm 2.3
		4	99.8 \pm 0.3	99.8 \pm 0.4	96.6 \pm 1.3	93.6 \pm 2.2	95.2 \pm 2.2
Blood Transfusion	4, 2,784	2	78.5 \pm 1.7	76.4 \pm 2.4	75.6 \pm 1.7	76.1 \pm 0.74	75.7 \pm 1.7
		3	78.5 \pm 2.0	77.8 \pm 2.5	76.9 \pm 2.1	76.5 \pm 1.4	76.1 \pm 1.9
		4	78.5 \pm 2.1	78.3 \pm 2.4	76.7 \pm 2.2	77.1 \pm 1.8	76.1 \pm 2.2
Acute Inflamm. 1	6, 2, 120	2	100 \pm 0	98.7 \pm 0.4	91.2 \pm 5.1	86.7 \pm 8.3	100 \pm 0
		3	100 \pm 0.3	100 \pm 0.3	99.7 \pm 1.3	100 \pm 0	100 \pm 0
		4	100 \pm 0	100 \pm 0	99.7 \pm 1.2	100 \pm 0	100 \pm 0
Acute Inflamm. 2	6, 2, 120	2	100 \pm 0	99.9 \pm 0.7	99.0 \pm 2.6	99.0 \pm 2.6	100 \pm 0
		3	100 \pm 0	99.9 \pm 1.3	99.0 \pm 2.6	99.0 \pm 2.6	100 \pm 0
		4	99.9 \pm 0.7	100 \pm 0	99.0 \pm 2.6	99.0 \pm 2.6	100 \pm 0
Car Evaluation	6, 4, 1728	2	82.9 \pm 1.5	80.0 \pm 3.1	77.8 \pm 1.3	77.8 \pm 1.3	77.6 \pm 1.3
		3	88.4 \pm 2.6	88.1 \pm 2.3	79.3 \pm 1.5	77.8 \pm 1.3	80.5 \pm 1.3
		4	93.3 \pm 2.2	92.1 \pm 2.4	84.5 \pm 1.5	84.3 \pm 1.4	86.1 \pm 1.3
Breast Cancer	9, 2, 683	2	97.2 \pm 1.3	97.2 \pm 1.2	94.2 \pm 1.6	93.8 \pm 1.7	95.5 \pm 1.8
		3	97.2 \pm 1.3	97.1 \pm 1.3	94.7 \pm 1.6	94.7 \pm 1.7	95.3 \pm 1.7
		4	97.1 \pm 1.3	97.2 \pm 1.0	94.7 \pm 1.7	94.7 \pm 1.8	95.2 \pm 1.6
Avila Bible	10, 12, 10430	4	62.2 \pm 1.4	59.7 \pm 1.8	55.8 \pm 0.8	54.0 \pm 1.3	59.6 \pm 0.7
Wine Quality Red	11, 6, 1599	3	58.6 \pm 2.2	56.4 \pm 2.5	56.1 \pm 2.3	55.3 \pm 2.1	55.7 \pm 2.4
		4	57.3 \pm 2.4	56.6 \pm 1.4	56.9 \pm 2.5	55.9 \pm 2.3	54.8 \pm 2.9
Wine Quality White	11, 7, 4898	3	52.6 \pm 1.1	51.7 \pm 1.8	52.0 \pm 1.5	51.9 \pm 1.3	50.9 \pm 1.4
		4	53.5 \pm 1.4	52.1 \pm 1.6	52.3 \pm 1.4	52.0 \pm 1.3	50.7 \pm 1.4
Dry Bean	16, 7, 13611	3	91.2 \pm 0.6	88.4 \pm 1.6	78.2 \pm 1.2	76.6 \pm 1.4	74.3 \pm 3.3
		4	91.4 \pm 0.5	89.0 \pm 1.6	83.2 \pm 1.5	80.5 \pm 1.9	77.9 \pm 4.7
Climate Crashes	18, 2, 540	2	92.9 \pm 1.4	92.3 \pm 2.4	90.6 \pm 2.2	91.3 \pm 1.2	91.3 \pm 2.2
		3	92.1 \pm 1.0	92.4 \pm 2.6	90.5 \pm 2.3	91.6 \pm 1.5	91.1 \pm 2.0
		4	91.9 \pm 0.8	92.3 \pm 2.5	90.1 \pm 2.6	91.8 \pm 1.8	90.7 \pm 2.6
Conn. Sonar	60, 2, 208	2	79.9 \pm 5.1	75.6 \pm 6.5	68.7 \pm 6.0	69.4 \pm 5.8	72.8 \pm 5.9
		3	81.5 \pm 5.0	79.5 \pm 5.9	70.9 \pm 6.1	70.1 \pm 6.3	71.4 \pm 5.7
		4	82.1 \pm 5.1	80.8 \pm 5.3	70.9 \pm 5.8	70.6 \pm 6.6	72.6 \pm 5.7
Optical Recognition	64, 10, 3823	4	93.3 \pm 1.0	91.9 \pm 1.0	64.6 \pm 6.5	53.2 \pm 3.2	65.2 \pm 2.0

Loss Landscape: We present the loss landscapes of *DTSemNet* and DGT-Linear, similar to those provided in the main paper, using a regression dataset (Ailerons) for demonstration purposes. As depicted in Figure 6, the loss landscape of *DTSemNet* appears slightly flatter compared to that of DGT-Linear. This difference in flatness is less pronounced than what was observed in the loss landscape of MNIST, mainly because the performance metrics are very close in this case — specifically, with an RMSE of 1.66 for *DTSemNet* and 1.67 for DGT-Linear. Given the minimal performance gap, observing some flatness in the loss landscape for *DTSemNet* underscores the impact of architecture on performance.

Table 7. Rewards in the Bipedal Walker environment with tree heights of 6, 7, and 8. The number of parameters for the NN is similar to the DT height, with the size of the hidden layer specified in square brackets [] alongside the reward. In the constant regressor, there is only one learnable parameter per leaf, independent of input. Conversely, in linear regression, each leaf represents a linear regression model.

Env	Height	DTSemNet	Deep RL	DGT [2]	ICCT [31]
Bipedal Walker	6	309.7 ± 4.49	221.3 ± 40.12 [74 × 74]	scalar: 47.87 ± 48.29 linear: 120.91 ± 76.9	301.34 ± 3.09
	7	314.98 ± 3.35	297.4 ± 6.25 [110 × 110]	scalar: 32.67 ± 23.65 linear: 163 ± 69.53	176.64 ± 71.78
	8	306.72 ± 5.6	315.3 ± 6.91 [160 × 160]	scalar: 78.33 ± 57.19 linear: 244.5 ± 61.84	112.24 ± 60.36

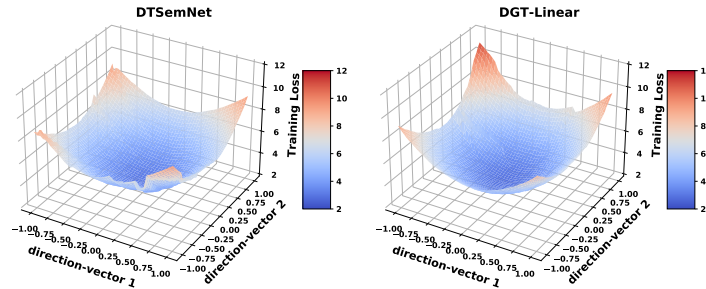


Figure 6. The loss landscape of *DTSemNet* and *DGT-Linear* on Ailerons regression dataset, when varying the parameters around the trained parameters along two random directions. The flatter the loss landscape, the better the generalization.

B Dataset Description

This section describes the data used in the experiment, including datasets for classification, regression, and the RL environment. All datasets were preprocessed by standardizing the input features with a mean of 0 and a standard deviation of 1. This normalization procedure was also applied to the target variables in regression datasets. We reuse the explanations of the datasets from [2]. We used the train-test split provided with the datasets whenever possible, details of which are provided in the tables.

B.1 Classification

The 14 classification datasets from [15] shown in Table 8 which is sourced from UCI². The 8 classification datasets from [12] as shown in Table 9 is from sourced from LIBSVM³.

Table 8. Details of classification dataset used in the experiment from UCI.

Dataset	train samples	# features	# classes	splits (tr:val:test)
Balance Scale	625	4	3	0.5 : 0.25 : 0.25
Banknote Auth	1372	4	4	0.5 : 0.25 : 0.25
Blood Transfusion	784	4	2	0.5 : 0.25 : 0.25
Acute Inflam. 1	120	6	2	0.5 : 0.25 : 0.25
Acute Inflam. 2	120	6	2	0.5 : 0.25 : 0.25
Car Evaluation	1728	6	4	0.5 : 0.25 : 0.25
Breast Cancer	683	9	2	0.5 : 0.25 : 0.25
Avila Bible	10430	10	12	0.5 : 0.25 : 0.25
Wine Quality Red	1599	11	6	0.5 : 0.25 : 0.25
Wine Quality White	4898	11	7	0.5 : 0.25 : 0.25
Dry Bean	13611	16	7	0.5 : 0.25 : 0.25
Climate Crashes	540	18	2	0.5 : 0.25 : 0.25
Conn. Sonar	208	60	2	0.5 : 0.25 : 0.25
Optical Recognition	3823	64	10	0.5 : 0.25 : 0.25

- **Balance Scale:** This data set was generated to model psychological experimental results. Each example is classified as having the balance scale tip to the right, tip to the left, or be balanced. The attributes are the left weight, the left distance, the right weight, and the right distance.
- **Banknote Auth:** The Banknote Authentication dataset includes features extracted from images of genuine and forged banknotes, with the objective of authenticating banknotes using machine learning techniques.
- **Blood Transfusion:** Blood Transfusion dataset contains information about blood donors and their donation history, aiming to predict whether a donor will donate blood again in the future to assist blood donation management.

² <https://archive.ics.uci.edu/datasets>

³ <https://www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets/>

Table 9. Details of classification dataset used in the experiment from LIBSVM.

Dataset	train samples	# features	# classes	splits (tr:val:test)
Protein	14,895	357	3	Default
SatImages	3,104	36	6	Default
Segment	1,478	19	7	0.64 : 0.16 : 0.2
PenDigits	5,995	16	10	Default
Connect4	43,236	126	3	0.64 : 0.16 : 0.2
MNIST	48,000	780	10	Default
SensIT	63,058	100	3	Default
Letter	10,500	16	26	Default

- **Acute Inflammation 1:** The data was created by a medical expert as a data set to test the expert system, which will perform the presumptive diagnosis of two diseases of the urinary system.
- **Acute Inflammation 2:** Similar to Acute Inflammation 1, this dataset also deals with cases of acute inflammation, providing additional samples for analysis and classification tasks.
- **Car Evaluation:** Car Evaluation dataset involves features of various car models, with the objective of classifying cars into different categories based on their characteristics, such as price, maintenance cost, and safety features.
- **Breast Cancer:** The Breast Cancer dataset provides medical data from breast cancer patients, with the goal of accurately classifying tumors as malignant or benign to aid in diagnosis and treatment planning.
- **Avila Bible:** The Avila data set has been extracted from 800 images of the 'Avila Bible,' an XII-century giant Latin copy of the Bible. The prediction task consists of associating each pattern with a copyist.
- **Wine Quality Red:** Wine Quality Red dataset provides chemical properties of red wine samples, aiming to predict wine quality ratings based on these properties to assist in wine production and quality control.
- **Wine Quality White:** The Wine Quality White dataset provides chemical properties of white wine samples, similar to the Wine Quality Red dataset, aiming to predict wine quality ratings based on these properties.
- **Dry Bean:** The Dry Bean dataset includes features related to various dry bean varieties, with the objective of predicting bean variety classifications based on these features to assist in agricultural research and crop management.
- **Climate Crashes:** Given Latin hypercube samples of 18 climate model input parameter values, predict climate model simulation crashes and determine the parameter value combinations that cause the failures.
- **Conn. Sonar:** The Conn. Sonar dataset involves sonar signals bounced off a metal cylinder and a roughly cylindrical rock, aiming to discriminate between these two classes for underwater object detection.
- **Optical Recognition:** The Optical Recognition dataset contains handwritten digit images, aiming to recognize and classify digits accurately to facilitate automated postal sorting and document processing tasks.
- **Protein:** Given features extracted from amino acid sequences, predict the secondary structure for a protein sequence.
- **Segment:** Given 19 attributes for each 3x3 pixel grid of instances drawn randomly from a database of 7 outdoor color. images, predict segmentation of the central pixel.
- **SatImages:** Given multi-spectral values of pixels in 3x3 neighborhoods in a satellite image, predict the central pixel in each neighborhood.
- **PenDigits:** Given integer attributes of pen-based handwritten digits, predict the digit.
- **Connect4:** Given legal positions in a game of connect-4, predict the game theoretical outcome for the first player.
- **MNIST:** Given 784 pixels in handwritten digit images, predict the actual digit.
- **SensIT:** Given 100 relevant sensor features, classify the vehicles.
- **Letter:** Given 16 attributes obtained from stimulus observed from handwritten letters, classify the actual letters.

B.2 Regression Dataset

We used scalar regression datasets from [2] shown in Table 10. We used the train-test split provided with the datasets whenever possible. When a default split is not provided, we create five different splits and evaluate each of the splits 10 times to obtain the average performance. Then, we average the performance across five splits. Following is the list of datasets:

- **Abalone:** Given attributes describing physical measurements, predict the age of an abalone. We encode the categorical ("sex") attribute as one-hot.
- **Comp-Activ:** Given different system measurements, predict the portion of time that CPUs run in user mode.
- **Ailerons:** Given attributes describing the status of the aircraft, predict the command given to its ailerons.
- **CtSlice:** Given attributes as histogram features (in polar space) of the Computer Tomography (CT) slice, predict the relative location of the image on the axial axis (in the range [0 180]).
- **YearPred:** Given several song statistics/metadata (timbre average, timbre covariance, etc.), predict the age of the song. This is a subset of the UCI Million Songs dataset.
- **PdbBind:** Given standard "grid features" (fingerprints of pairs between ligand and protein; see Wu et al. (2018)), predict binding affinities.
- **Microsoft:** Given 136-dimensional feature vectors extracted from query-url pairs, predict the relevance judgment labels, which take values from 0 (irrelevant) to 4 (perfectly relevant).

Table 10. Details of regression datasets used from [2]

Dataset	# features	Train Size	Splits (tr:val:test)	# shuffles	Source
Abalone	10	2,004	0.5:0.1:0.4	5	UCI
Comp-Activ	21	3,932	0.5:0.1:0.4	5	Delve ⁴
Ailerons	40	5,723	Default	1	LIACC ⁵
CTSlice	384	34,240	0.5 : 0.1 : 0.4	5	UCI
YearPred	90	370,972	Default	1	UCI
PDBBind	2,052	9,013	Default	1	MoleculeNet4 ⁶
Microsoft	136	578,729	Default	1	MSLR-WEB10K ⁷

B.3 RL Environments

In the RL setting 4 discrete environments: CartPole, Acrobot, Lunar Lander and FindAndDefeatZerglings and 2 continuous environments: Continuous Lunar Lander and Bipedal walker is used. Following is the details of environment:

- **CartPole:** Balancing a pole on a cart, the environment has a continuous state space of four dimensions and a discrete action space of two dimensions, with rewards of +1 for each timestep the pole remains upright and 0 upon termination.
- **Acrobot:** Involving a two-link pendulum system, the environment features a continuous state space of four dimensions and a discrete action space of three dimensions, with a reward of -1 upon termination.
- **LunarLander:** Simulating lunar landing, this environment has a continuous state space of eight dimensions and a discrete action space of four dimensions, with positive rewards for successful landings and negative rewards for crashes.
- **FindAndDefeatZerglings:** Tasked with navigating a maze and battling Zerglings, the environment has a discrete state space of 32 dimensions and a discrete action space of 30 dimensions, with positive rewards for defeating enemies and negative rewards for taking damage.
- **Continuous Lunar Lander:** Similar to LunarLander but with continuous action space, the environment has a continuous state space of eight dimensions and a continuous action space with two dimensions, with positive rewards for successful landings and crash penalties.
- **Bipedal Walker:** Featuring a two-legged robot learning to walk, the environment has a continuous state space of 24 dimensions and a continuous action space of four dimensions, with positive rewards for making progress and penalties for falling or deviating from the desired trajectory.

C Implementation Details

In this section, we outline the implementation details for the experiments and provide the sources of the original implementation from which we derived the code.

DTSemNet: The implementation of *DTSemNet* is done using `nn.module` class of PyTorch. For the classification and regression datasets, we provide hyperparameters such as optimizer (SGD, RMSProp, or Adam), learning rate (ls), learning momentum (mtm), type of scheduler, scheduler decay parameter, L1 regularizer parameter λ , batch size, gradient clipping (grad clip) and over-parameterization (overparams) used during training. The over-parameterization is indicated by the size of the used hidden layer for a linear layer. These hyperparameters are tuned using grid search, which is summarized in Tables 11. Configurations that are not relevant are available in the configuration files provided in the source code. Furthermore, we provide the hyperparameters used for training agents in the RL setup for six different environments in Tables 12, 13 and 14. These hyperparameters include the choice of optimizer and various parameters for the PPO learner. We obtained these hyperparameters for all agents that use PPO from Stablebaselines through an automated hyperparameter search using the ‘‘Optuna’’ library. Optuna performs a Bayesian optimization-based automatic hyperparameter search to find suitable parameters.

VIPER: We utilize the online implementation available at <https://github.com/Safe-RL-Team/viper-verifiable-rl-impl>. VIPER requires Q-values to train a decision tree. In cases where the policy network is available instead of the Q-network, the authors of VIPER proposed using $\log(\pi)$ for the policy output π in place of the Q-value, as it is proportional to it [4].

ICCT: We use the implementation provided at <https://github.com/CORE-Robotics-Lab/ICCT/tree/main> from the authors of ICCT [31] for the architecture of DT.

DGT: We use the implementation available at <https://github.com/microsoft/DGT> from the authors of DGT [2] for our experiments.

CRO-DT: We use the implementation available at <https://github.com/vgarciasc/CRO-DT> from the authors of CRO-DT [15] for the experiments.

⁴ <https://www.cs.toronto.edu/~delve/data/comp-activ/desc.html>

⁵ <https://www.dcc.fc.up.pt/~ltorgo/Regression/aileron.html>

⁶ <http://www.pdbbind.org.cn/>

⁷ <https://www.microsoft.com/en-us/research/project/mslr/>

Table 11. Hyperparameters used for Experiments with classification and regression datasets. Reported hyperparameters are epoch, optimizer (SGD, RMSProp, or Adam), learning rate (lr), learning momentum (mtm), scheduler type (linear or cosine), scheduler decay parameter, L1 regularizer parameter λ , batch size, gradient clipping (grad clip) and over-parameterization (overparams).

Dataset	epoch	optimizer	lr	mtm	scheduler type	scheduler decay	batch size	λ	grad clip	overparams
Segment	50	Adam	0.01	NA	Linear	0.95	128	NA	NA	[1530]
SatImages	200	RMSProp	0.01	0.2	Linear	0.98	128	5e-5	0.01	[4032]
PenDigits	100	RMSProp	0.01	0.0	Linear	0.98	128	1e-5	0.01	[16320]
Letter	400	RMSProp	0.01	0.0	Linear	0.95	128	5e-6	0.01	[6128x6138]
Protein	40	RMSProp	0.01	0.0	Linear	0.98	128	1e-3	0.01	[180x180]
Connect4	100	RMSProp	0.01	0.0	Linear	0.95	128	1e-4	0.01	[8160]
MNIST	100	SGD	0.4	0.9	Linear	0.95	128	NA	NA	NA
SensIT	250	RMSProp	0.01	0.0	Linear	0.95	128	1e-4	0.01	[6138]
Acute Inflamm. 1	20	Adam	0.8	NA	Linear	0.99	128	NA	NA	NA
Acute Inflamm. 2	30	Adam	0.7	NA	Linear	0.98	128	NA	NA	NA
Conn. Solar	40	Adam	0.1	NA	Linear	0.95	128	1e-3	NA	NA
Climate Crashes	30	Adam	0.4	NA	Linear	0.90	128	1e-2	NA	NA
Balance Scale	20	Adam	0.8	NA	Linear	0.98	128	NA	NA	NA
Breast Cancer	50	Adam	0.05	NA	Linear	0.90	256	1e-2	NA	NA
Blood Transfusion	50	Adam	0.7	NA	Linear	0.95	128	NA	NA	NA
Banknote Auth	40	Adam	0.5	NA	Linear	0.98	128	NA	NA	NA
Wine Quality Red	60	Adam	0.4	NA	Linear	0.98	128	NA	NA	NA
Car Evaluation	40	Adam	0.8	NA	Linear	0.99	128	NA	NA	NA
Optical Recognition	100	Adam	0.8	NA	Linear	0.95	128	1e-3	NA	NA
Wine Quality White	60	Adam	0.1	NA	Linear	0.99	128	NA	NA	NA
Avila Bible	100	Adam	0.3	NA	Linear	0.98	128	NA	NA	NA
Dry Bean	40	Adam	0.1	NA	Linear	0.99	128	1e-3	NA	NA
Ailerons	100	Adam	0.1	NA	Linear	0.90	32	5e-4	0.01	NA
Abalone	50	Adam	0.005	NA	Linear	0.95	32	5e-4	NA	[248]
Comp-Activ	100	Adam	0.02	NA	Linear	0.90	128	1e-4	NA	[186]
CTSlice	50	Adam	0.001	NA	Linear	0.90	128	NA	NA	[992x496]
PDBBind	100	Adam	0.01	NA	Linear	0.95	256	1e-2	NA	NA
YearPred	40	Adam	0.001	NA	Linear	0.90	128	1e-4	NA	[378]
Microsoft	30	Adam	0.002	NA	Linear	0.90	256	NA	NA	NA

Table 12. Hyperparameters of agents in CartPole (left) and Acrobot (right) Environment.

Parameters	DTSemNet	NN	DGT	ICCT	Parameters	DTSemNet	NN	DGT	ICCT
Size (height/nodes)	4	16x16	4	4	Size (height/nodes)	4	16x16	4	4
Learning Rate	0.024	0.015	0.0042	0.024	Learning Rate	0.045	0.007	0.001	0.045
Batch Size	32	32	43	32	Batch Size	64	64	128	64
Num Steps	832	512	542	832	Num Steps	1024	1024	2048	1024
Entropy Coeff	0.02	0.008	0.02	0.02	Entropy Coeff	0.01	0.02	0.012	0.01
Clip Range	0.2	0.2	0.2	0.2	Clip Range	0.2	0.2	0.2	0.2
VF Coeff	0.5	0.5	0.5		VF Coeff	0.5	0.5	0.54	0.5
GAE Lamda	0.86	0.96	0.934	0.86	GAE Lamda	0.80	0.85	0.94	0.80
Max Grad Norm	0.5	0.5	0.53	0.5	Max Grad Norm	0.52	0.5	0.5	0.52
Num Epochs	20	12	7	20	Num Epochs	16	16	20	16
Critic Network	————	16x16	————		Critic Network	————	32x32	————	
Optimizer	————	Adam	————		Optimizer	————	Adam	————	
Discount Factor	————	0.99	————		Discount Factor	————	0.99	————	

Table 13. Hyperparameters of agents in Lunar Lander (left) and FindAndDestroy (right) Environment.

Parameters	DTSemNet	NN	DGT	ICCT	Parameters	DTSemNet	NN	DGT	ICCT
Size (height/nodes)	5	32x32	5	5	Size (height/nodes)	6	64x64	6	6
Learning Rate	0.02	0.005	0.02	0.005	Learning Rate	0.02	0.005	0.02	0.005
Batch Size	_____	64	_____	_____	Batch Size	_____	64	_____	_____
Num Steps	2048	2048	2048	2048	Num Steps	2048	2048	2048	2048
Entropy Coeff	_____	0.01	_____	_____	Entropy Coeff	_____	0.01	_____	_____
Clip Range	_____	0.2	_____	_____	Clip Range	_____	0.2	_____	_____
VF Coeff	_____	0.5	_____	_____	VF Coeff	_____	0.5	_____	_____
GAE Lamda	0.9	0.9	0.9	0.9	GAE Lamda	0.9	0.9	0.9	0.9
Max Grad Norm	_____	0.5	_____	_____	Max Grad Norm	_____	0.5	_____	_____
Num Epochs	20	10	4	32	Num Epochs	10	10	4	4
Critic Network	_____	32x32	_____	_____	Critic Network	_____	128x128	_____	_____
Optimizer	Adam	Adam	Adam	Adam	Optimizer	_____	Adam	_____	_____
Discount Factor	_____	0.99	_____	_____	Discount Factor	_____	0.99	_____	_____

Table 14. Hyperparameters of agents in ContinuousLunarLander Environment (left) and Bipedal Walker (right).

Parameters	DTSemNet	NN	DGT	ICCT	Parameters	DTSemNet	NN	DGT	ICCT
Size (height/nodes)	4	16x16	4	4	Size (height/nodes)	7	160x160	8	6
Learning Rate	1e-3	1e-4	1e-3	1e-4	Learning Rate	2e-3	5e-4	2e-3	1e-3
Batch Size	_____	256	_____	_____	Batch Size	_____	256	_____	_____
Tau	_____	0.01	_____	_____	Tau	_____	0.02	_____	_____
Train Freq	_____	1	_____	_____	Train Freq	_____	64	_____	_____
Grad Steps	_____	1	_____	_____	Grad Steps	_____	64	_____	_____
Critic Network	_____	256x256	_____	_____	Critic Network	_____	512x512	_____	_____
Optimizer	_____	Adam	_____	_____	Optimizer	_____	Adam	_____	_____
Discount Factor	_____	0.99	_____	_____	Discount Factor	_____	0.98	_____	_____