

TREE-G: Decision Trees Contesting Graph Neural Networks

Maya Bechler-Speicher¹, Amir Globerson¹, Ran Gilad-Bachrach²

¹ Blavatnik School of Computer Science, Tel-Aviv University

² Department of Bio-Medical Engineering and Edmond J. Safra Center for Bioinformatics, Tel-Aviv University
mayab4@mail.tau.ac.il

Abstract

When dealing with tabular data, models based on decision trees are a popular choice due to their high accuracy on these data types, their ease of application, and explainability properties. However, when it comes to graph-structured data, it is not clear how to apply them effectively, in a way that incorporates the topological information with the tabular data available on the vertices of the graph. To address this challenge, we introduce TREE-G. TREE-G modifies standard decision trees, by introducing a novel split function that is specialized for graph data. Not only does this split function incorporate the node features and the topological information, but it also uses a novel pointer mechanism that allows split nodes to use information computed in previous splits. Therefore, the split function adapts to the predictive task and the graph at hand. We analyze the theoretical properties of TREE-G and demonstrate its benefits empirically on multiple graph and vertex prediction benchmarks. In these experiments, TREE-G consistently outperforms other tree-based models and often outperforms other graph-learning algorithms such as Graph Neural Networks and Graph Kernels, sometimes by large margins. Moreover, TREE-Gs models and their predictions can be explained and visualized.

1 Introduction

Decision Trees (DTs) are one of the cornerstones of modern data science and are a leading alternative to neural networks, especially in tabular data settings, where DTs often outperform deep learning models (Shwartz-Ziv and Armon 2021; Han et al. 2018; Miah et al. 2021; Grinsztajn, Oyallon, and Varoquaux 2022). DTs are popular with practitioners due to their ease of use, out-of-the-box high performance, and explainability properties. In a large survey conducted by Kaggle¹ in 2022 with 23,997 data scientists worldwide, 63% of the participants reported using decision trees, Random Forests (RF) or Gradient Boosted Trees (GBT) on a regular basis (Ho 1995; Friedman 2000, 2002).

In this study we investigate whether the merits of DTs can be repeated in tasks that involve *graph* structured data. Such tasks are found in diverse domains including social studies, molecular biology, drug discovery, and communication (Sperduti and Starita 1997; Gori, Monfardini, and Scarselli 2005).

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹www.kaggle.com/competitions/kaggle-survey-2022/data

The ability to perform regression and classification on graphs is important in these domains and others, which explains the ongoing efforts to develop machine learning algorithms for graphs (Li et al. 2017; Hamilton, Ying, and Leskovec 2017; Duvenaud et al. 2015; Defferrard, Bresson, and Vandergheynst 2016; Barrett et al. 2018; Ganea et al. 2021; Rossi et al. 2020; Ingraham et al. 2019).

Graph data often exhibits characteristics similar to those of tabular data. For instance, in a social network, vertices represent individuals with associated features such as age, marital status, and education. Given the efficacy of DTs with tabular data, one might wonder if this success can be translated to graph-related tasks. However, adapting DTs to integrate both the graph structure and vertex attributes remains a challenge. To address this gap, we present a novel and efficient technique, named TREE-G, to apply DTs to graphs.

DTs split the input space by comparing feature values to thresholds. When operating on graphs, these split decisions should incorporate the features on the vertices of the graph as well as its topological structure.² Moreover, these decisions should be applicable to graphs of varying sizes and satisfy the appropriate invariance and equivariance properties with respect to vertex ordering. Current methods achieve these goals using ‘feature engineering’, a process in which the “raw features” are augmented with computed features using pre-defined functions that may use the graph structure. These features, which are computed before the evaluation (or training) of the DT, use domain knowledge, or graph theoretic measures (Barnett et al. 2016; de Lara and Pineau 2018). However, these approaches are labor intensive, require domain knowledge, and may fail to capture interactions between vertex features and graph structure. We show this formally in Lemma 4.3 where a pair of non-isomorphic graphs over 4 vertices and 2 features are shown to be inseparable by standard tree methods.

TREE-G enhances standard decision trees by introducing a novel split function, specialized for graph data which is dynamically evaluated while the DT is applied to a graph. It supports both directed and undirected graphs of arbitrary sizes and can be used for classification, regression, and other prediction tasks on entire graphs or on the vertices and edges of these graphs. It preserves the standard decision trees frame-

²For the sake of clarity, we ignore edge-features in this work.

work and thus can be used in any ensemble method, for instance, as a weak learner in GBT. Indeed, we show in an extensive empirical evaluation that TREE-G with GBT generates models that are superior to other tree-based methods. Furthermore, it often outperforms Graph Neural Networks (GNNs), the leading deep-learning method for learning over graphs. In some cases, the difference in performance between TREE-G and the other methods reaches ~ 6.4 percentage points.

A key idea in TREE-G is that split functions in the tree should be able to focus on a subset of the nodes in the graph to identify relevant substructures. For example, in the task of gauging the tendency of a molecule to increase genetic mutation frequency this should allow models to identify substructures such as carbon rings (cycles), that are known to be mutagenic (Debnath et al. 1991). When task-related substructures are known beforehand, one can compute task-specific features based on them. However, when there’s a need to learn these substructures, finding a method that is computationally efficient and consistent across different graphs remains a challenge. TREE-G addresses this challenge by a novel mechanism to dynamically generate candidate subsets of vertices at each split-node as the graph traverses the tree. These subsets can be used in split functions of downstream split-nodes combined with the vertex features and the topology of the graph. Ablation studies we performed confirmed that this ability is a key ingredient of the empirical superiority of TREE-G over other models. As we prove later, TREE-G is more expressive than standard decision trees that only act on the input features, even if some topological features which are computed in advance are added as input features.

The main contributions of this work are: (1) We present TREE-G, a novel decision tree model specialized for graph data. (2) We empirically show that TREE-G based ensembles outperform existing tree-based methods and graph kernels. Furthermore, TREE-G is very competitive with GNNs and often outperforms them. (3) We provide theoretical results that highlight the properties and expressive power of TREE-G. In addition, we introduce an explainability mechanism for TREE-G.

2 Related Work

Machine Learning on Graphs Many methods for learning on graphs were proposed. In Graph Kernels, a predefined embedding is used to represent graphs in a vector space to which cosine similarities are applied (Gärtner, Flach, and Wrobel 2003; Shervashidze et al. 2011, 2009). A more recent line of work utilizes deep-learning approaches: GNNs learn a representation vector for each vertex of a given graph using an iterative neighborhood aggregation process. These representations are used in downstream tasks (Gori, Monfardini, and Scarselli 2005; Gilmer et al. 2017a; Ying et al. 2018). See Wu et al. (2020); Nikolentzos, Siglidis, and Vazirgiannis (2021) for recent surveys.

DTs for Graph Learning Several approaches have been introduced to allow DTs to operate on graphs (Barnett et al. 2016; de Lara and Pineau 2018; Heaton 2016; He et al. 2017; Lei and Fang 2019; Wu et al. 2018; Li and Tu 2019). These approaches combine domain-specific feature engineering with

graph-theoretic features. Other approaches combine GNNs with DTs in various ways and have mostly been designed for specific problems (Li et al. 2019). For example, XGraphBoost was used for a drug discovery task (Deng et al. 2021; Stokes et al. 2020). It uses GNNs to embed the graph into a vector and applies GBT (Friedman 2000, 2002) to make predictions using this vector representation. A similar method has been used for link prediction between human phenotypes and genes (Patel and Guo 2021). The recently proposed BGNN model predicts vertex attributes by alternating between training GBT and GNNs (Ivanov and Prokhorenkova 2021). Thus, existing solutions either combine deep-learning components or focus on specific domains utilizing domain expertise.

TREE-G, detailed in the subsequent section, offers a ‘pure’ and generic DT solution tailored for graph-based tasks. It eliminates the need for domain knowledge or neural networks, yet remains competitive with these approaches. It is inspired by Set-Trees (Hirsch and Gilad-Bachrach 2021) where DTs were extended to operate on data that had a set structure by introducing an attention mechanism. Sets can be viewed as graphs with no edges. However, note that since Set-Trees are designed to work on sets, they do not accommodate the topological structure of graphs nor handle the diverse types of tasks that graph data gives rise to such as vertex or link labeling

3 TREE-G

In this section we introduce the TREE-G model.

3.1 Preliminaries

A graph G is defined by a set of vertices V of size n and an adjacency matrix A of size $n \times n$ that can be symmetric or a-symmetric. The entry A_{ij} indicates an edge from vertex j to vertex i . Each vertex is associated with a vector of l real-valued features, and we denote the stacked matrix of feature vectors over all the vertices by X . We denote the k ’th column of this matrix by \mathbf{f}_k ; i.e., entry i in \mathbf{f}_k holds the value of the k ’th feature of vertex i . The i ’th entry of a vector \mathbf{x} is denoted by x_i or $(\mathbf{x})_i$. Recall that $(A^d)_{ij}$ is the number of walks of length d between vertex j to vertex i . We focus on two tasks: *graph labeling* and *vertex labeling* (each includes both classification and regression). In the former, the goal is to assign a label to the entire graph while in the latter, the goal is to assign a label to individual vertices.³ In vertex labeling tasks, TREE-G takes in a graph along with one of its vertices and predicts the label of that vertex. In contrast, for graph labeling tasks, TREE-G is fed with a graph and predicts the label for the entire graph. Given the variety of graph tasks, TREE-G comes in multiple variants. It’s important to note that trees are also considered graphs. To prevent any confusion, we use terms like *root*, *node*, *split-node*, *leaf* to refer to decision tree nodes, and *vertex*, *vertices* for graph vertices.

³Predicting properties of edges, can be performed using the dual line graph as described in the Appendix.

Algorithm 1: Infer TREE-G

Data: A : adjacency matrix, X : node features matrix, i : node index for inference in the case of vertex labeling task, N : a TREE-G node with stored parameters $k_N, d_N, *N, \rho_N, r_N, \theta_N$ and AGG_N in the case of graph labeling tasks

Result: label y

Function InferVertex (N, A, X, i):

```

if  $N$  is a leaf node then
   $\perp$  return the label stored in  $N$ 
  Evaluate  $\phi = \phi_{k_N, d_N, *N, \rho_N, r_N}(A, X, i)$ 
  if  $\phi > \theta_N$  then
    return InferVertex ( $N_{ge}, A, X, i$ ) where  $N_{ge}$  is
    a child of  $N$ 
  else
    return InferVertex ( $N_{leq}, A, X, i$ ) where  $N_{leq}$ 
    is a child of  $N$ 

```

Function InferGraph (N, G):

```

if  $N$  is a leaf node then
   $\perp$  return the label stored in  $N$ 
  Evaluate  $\phi = \phi_{k_N, d_N, *N, \rho_N, r_N, AGG_N}(A, X)$ 
  if  $\phi > \theta_N$  then
    return InferGraph ( $N_{ge}, G$ ) where  $N_{ge}$  is a
    child of  $N$ 
  else
    return InferGraph ( $N_{leq}, A, X$ ) where  $N_{leq}$  is a
    child of  $N$ 

```

$y_{vertex} \leftarrow$ InferVertex (N_{Root}, A, X, i)

$y_{graph} \leftarrow$ InferGraph (N_{Root}, A, X)

3.2 A Split Function Specialized for Graphs

For the sake of clarity, we begin by highlighting the parts of the standard DTs framework TREE-G modifies. During the inference phase in standard DTs, a vector example \mathbf{x} traverses the tree until a leaf is reached and the value stored in the leaf is reported. At each split-node, the k 'th feature of \mathbf{x} is compared against a threshold, $x_k > \theta$. During training, the goal is to select in each split-node, the optimal feature index k and threshold value θ , with respect to some optimization criterion, e.g., Gini score or the L_2 -loss. Training is conducted greedily by converting a leaf into a decision node with two new leaves to optimize the criterion.

TREE-G preserves the same framework, but instead of using a single feature that will be compared to a threshold as a split function, a more elaborate function ϕ is used to integrate vertex features with graph structure. All other parts of the decision trees framework are preserved, including the greedy training process. Hence, it can be embedded in ensemble methods, such as Random Forests or Gradient Boosted Trees⁴ as well as be limited to a certain size either by limiting its depth during training or by using pruning.

The split function used in TREE-G incorporates the vertex features, adjacency matrix, and a subset of the vertices of the graph. It is parameterized by 5 parameters $k, d, *, r$ and ρ for vertex labeling tasks, and an additional parameter AGG

for graph labeling tasks, as explained in detail in Section 3.2 and Section 3.5. Overall, while the binary rule in standard trees is of the form $\phi_k(\mathbf{x}) > \theta$, the binary rule in TREE-G is of the form $\phi_{k,d,*,\rho,r}(X, A, i) > \theta$ for vertex labeling tasks and $\phi_{k,d,*,\rho,r,AGG}(X, A) > \theta$ for graph labeling tasks. Here X is a matrix of the stacked vertex feature vectors, A is an adjacency matrix, and i is an index of a vertex in the graph. Algorithm 1 presents the pseudo-code for the inference process in TREE-G. Notice that it only differs from the standard DT inference procedure in the split function ϕ that is computed and compared against the threshold during the traversal of the tree.

For the rest of this paper, we will focus only on the novel split function $\phi_{k,d,*,\rho,r}$ TREE-G introduces, and how it enables learning effective decision trees for graph and vertex predictive tasks. For the sake of clarity, we first focus on vertex labeling, and explain graph labeling in Section 3.5.

TREE-G for Vertex Labeling We begin with the assumption of a trained tree and demonstrate its use during the inference phase. Training details are provided in Section 3.4. In vertex labeling, a vertex i is given for inference, together with its graph G and vertices' feature matrix X . The split function applied to the input in TREE-G is based on a feature, say \mathbf{f}_k , propagated through walks of length d over parts of the graph. For example, for some vertex i in a graph, the question asked by the split function may be: is the sum of the k 'th feature, limited to the neighbors of the vertex i , greater than θ ? Moreover, the decision can use information only from vertices that satisfy additional constraints but we begin the exposition without such constraints and defer this discussion to later.

Splits are based on propagation of feature values over walks in the graph that are computed by the function $A^d \mathbf{f}_k$. Here, the parameters d and k are used to define the propagation of the k 'th feature through walks of length d in the graph. When $d = 0$ the graph structure is ignored since $A^0 \mathbf{f}_k = \mathbf{f}_k$. However, when $d = 1$ for example, the i 'th entry in $A^1 \mathbf{f}_k$ is the sum of the values of the k 'th feature over the neighbors of the i 'th vertex. In directed graphs it may be useful to consider walks in the opposite direction by taking powers of A^T . Note also that if \mathbf{f}_0 is the constant feature 1, then $(A^d \mathbf{f}_0)_i$ is the number of walks of length d that end in vertex i .

Next, we introduce a way to further restrict the walks used for feature propagation. We would like to adapt the values of $A^d \mathbf{f}_k$ to focus on a subset of the vertices S . Hence, the split function can use a graph's substructure rather than its entirety. We present intuitive and computationally efficient methods below for this purpose.⁵

1. **Source Walks:** Only consider walks starting in the vertices in S .
2. **Cycle Walks:** Only consider walks starting and ending in the same vertex in S .

These restrictions can be implemented by zeroing out parts of the walks matrix A^d before propagating the feature values \mathbf{f}_k through it. For source walks, zero out the j 'th column in A^d

⁴In our experiments, we used GBT with TREE-G estimators

⁵This selection allows computing only once the powers of the adjacency matrix, as a pre-processing step. See more in the Appendix.

for each vertex $j \notin S$. For cycle walks, zero out all A^d except for the j^{th} value on the main diagonal, for each vertex $j \in S$. Hence, we define a mask matrix that corresponds to one of the two options described above (i.e., source or cycle walks), and denote it by $M_r(S)$, $r \in \{1, 2\}$. The parameter r indicates the type of walk restrictions, where $r = 1$ indicates source walks and $r = 2$ indicates cycle walks. We apply $M_r(S)$ to A^d using an element-wise multiplication denoted by \circ .

The key challenge is to design a policy to select the set S . Clearly, we cannot add S as a feature that the DT enumerates over, since there are exponentially many and their definition would not generalize across different graphs. Our main insight involves determining the subset S only after receiving the graph G . To achieve that, we utilize ancestor split-nodes to create candidate subsets of vertices to choose from. To that end, in each split-node in TREE-G two actions are performed: (1) the split function is applied to the example to determine to which child it should continue, and (2) the split-node generates two subsets of the vertices of the given graph. We denote the two generated subsets at a split-node u to be $S_{u,+}$ and $S_{u,-}$. For clarity, we'll first present the full expression for the split function ϕ and defer the detailed definitions of these subsets to Section 3.3.

Every split-node in the tree uses one of the subsets generated by its ancestors in the tree. In a node u , the parameter $*_u$ points to the ancestor and the parameter $\rho_u \in \{+, -\}$ indicates if the chosen subset is $S_{*_u,+}$ or $S_{*_u,-}$. Therefore, the subset S_{*_u,ρ_u} to be used in node u is well-defined when u has to evaluate the input graph G since all its ancestor nodes have already been evaluated. Notice that S_{*_u,ρ_u} is graph dependent, and therefore translates to different subsets for different graphs. To allow a split-node to focus on the entire graph we use the convention that if $*_u \equiv u$ then $S_{*_u,\rho_u} = V$, regardless of the value of ρ_u . Since the root node has no ancestors, the root node always points to itself and therefore uses all the vertices as a subset.

Finally, given a vertex example i , the adjacency matrix A , and the stacked feature vectors of the graph's vertices X , the split function TREE-G computes is defined as follows:

$$\phi_{k,d,*,\rho,r}(A, X, i) = \left((A^d \circ M_r(S_{*,\rho})) \mathbf{f}_k \right)_i \quad (1)$$

Then, the binary decision at a split-node will be given by the binary rule:

$$\phi_{k,d,*,\rho,r}(A, X, i) > \theta \quad (2)$$

Figure 1 presents one TREE-G tree in an ensemble of trees learned in the evaluation presented in Section 5, on the Mutag dataset. Each split-node in the tree shows the split function used in that split-node. The dashed arrows in each split-node represent the value of the $*$ parameter and point to an ancestor in the path to the root where the used subset is taken from. The value of the ρ parameter appears on each arrow, and it indicates which of the two generated subsets in that ancestor should be used.

3.3 Subsets Generation

As mentioned above, each split-node u generates two subsets of the vertices $S_{u,+}$ and $S_{u,-}$, which may be consumed by its descendant split-nodes in the tree. Recall that a split function

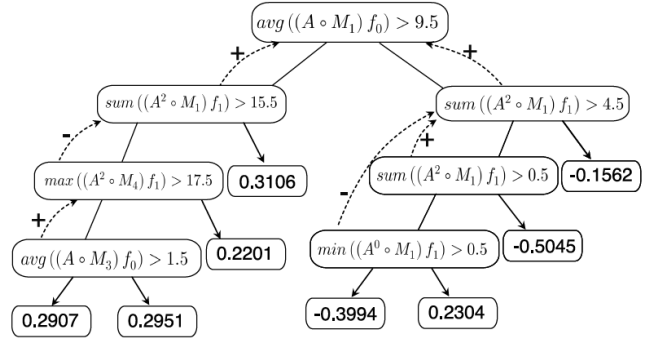


Figure 1: One TREE-G tree in the ensemble of the Mutag experiment. Each node in the tree presents the split function and threshold in that node. The dashed arrow in each node is the pointer $*$ and it points to the ancestor split-node where the subset is taken from together with the value of $\rho \in \{+, -\}$ which indicates which subset from the two subsets generated in that ancestor to use.

in u uses the parameters $k_u, d_u, *_u, \rho_u, r_u$ and a threshold θ_u . Informally speaking $S_{u,+}$ are the vertices that increase the value of $\phi = \phi_{k_u, d_u, *_u, \rho_u, r_u}$ to be greater than θ_u while $S_{u,-}$ are the vertices that decrease ϕ to be smaller than θ_u . Formally speaking, the node u already uses a subset of the vertices S_{*_u, ρ_u} . Then, the generated subset $S_{u,+}$ contains the vertices among S_{*_u, ρ_u} that satisfy the criterion in Eq. 2 with the selected parameters and threshold. Similarly, the generated subset $S_{u,-}$ is the vertices among S_{*_u, ρ_u} that do not satisfy the criterion. That it:

$$\begin{aligned} S_{u,+} &= \{j | \phi(A, X, j) > \theta_u, \quad j \in S_{*_u, \rho_u}\} \\ S_{u,-} &= \{j | \phi(A, X, j) \leq \theta_u, \quad j \in S_{*_u, \rho_u}\} \end{aligned} \quad (3)$$

where $\phi = \phi_{k_u, d_u, *_u, \rho_u, r_u}$. In particular, it holds that $S_{u,+} \cap S_{u,-} = \emptyset$ and $S_{u,+} \cup S_{u,-} = S_{*_u, \rho_u}$. Notice that subsets are generated separately for every different graph, however, they are generated from the same rule. Moreover, the definition of subsets is invariant to isomorphism and it does not assume a graph size, which allows TREE-G to be applied to graphs of varying sizes including sizes that were not seen during training. Notice that in vertex-labeling tasks, each split-node splits the vertices that arrive at it into two sets, one that would proceed left and the other would proceed right when traversing the tree. At the same time, it also splits the vertices of its used subsets into two sets $S_{*,+}$ and $S_{*,-}$. Since the vertices that arrive at a node are not necessarily the vertices composing its used subset, it is also true that $S_{*,+}$ and $S_{*,-}$ do not necessarily correspond with the vertices that traverse left and right.

Figure 2 demonstrates how the same trained TREE-G instance is applied to two graphs of different sizes during inference. Each split-node uses one subset among the subsets generated in its ancestor nodes, and the set of all vertices S . The subset to use in each split-node is uniquely defined by a pointer $*$ that points to the ancestor where the subset is generated (marked by dashed lines), and $\rho \in \{+, -\}$ that indicates which of the two subsets generated in that ancestors

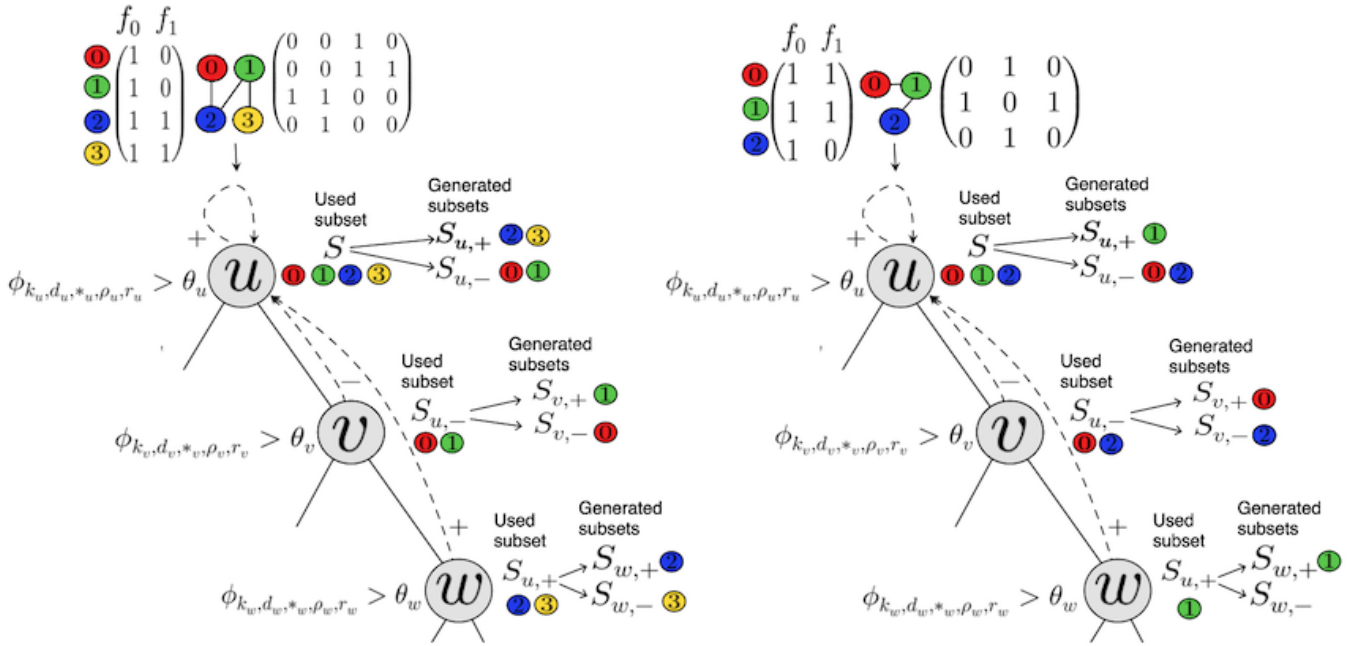


Figure 2: The same trained TREE-G instance is applied to two graphs of different sizes during inference. Each split-node uses one subset among the subsets generated in its ancestor nodes, and the set of all vertices S . The subset to use in each split-node is uniquely defined by a pointer $*$ that points to the ancestor where the subset is generated, and $\rho \in \{+, -\}$ that indicates which of the two subsets generated in that ancestors should be used. Each split-node generates two subsets, by splitting its used subset. The subsets are computed using the same rules but translated to different subsets with respect to the given graphs.

should be used. Each split-node generates two subsets, by splitting its used subset. The subsets are computed using the same rules but translate to different subsets with respect to the given graphs. Since node u is the root, it points to itself and uses the set S of all vertices of the graph. It generates two subsets $S_{u,+}, S_{u,-}$ by partitioning V . For both graph examples, the same function is applied to generate $S_{u,+}, S_{u,-}$, but it results in different subsets that correspond to the given graphs. Split-node v points to its ancestor u with $\rho_v = -$, therefore $S_{*v,\rho_v} = S_{u,-}$. It generates two subsets by partitioning its used subset $S_{u,+}$ into two new subsets: $S_{v,+}$ and $S_{v,-}$. Node w also points to its ancestor u , with $\rho_w = +$. Therefore it uses the subset $S_{u,+}$, and generated two subsets by splitting it. For the sake of clarity in this figure, we provide the generic split function. Our goal in this figure is to show the general mechanism and how the same instance is applied to different graphs. We provide a detailed example of the split functions that produce these subsets, as well as an example of a trained TREE-G over a real-world dataset, in the Appendix.

Lemma 4.3 shows the theoretical merits of a split function that can focus on subsets. Specifically, it demonstrates that without this mechanism there exist graphs with merely 4 vertices that are indistinguishable.

The selected subsets can also be used to explain the decisions a TREE-G makes by noting that vertices that appear more often in the used subsets are likely to have a larger influence on the prediction. Therefore, it is possible to visualize these vertices as an explainability aid. Due to space limitations, this is discussed and demonstrated at length only

in the Appendix.

3.4 Training TREE-Gs

As TREE-G preserves most of the standard DTs framework, training it uses the same greedy process as standard DTs. First, an optimization criterion is selected, for example, the Gini score or the L_2 -loss. For each leaf, a grid search is used to find the parameters that will reduce the score the most if the leaf was to become a split-node and the potential score reduction is recorded. The leaf that creates the largest score reduction is selected and is converted into a split-node using the optimal parameters that were found. This process repeats until a stop criterion is satisfied. Popular criteria include bounds on the size of the tree, bounds on the number of training examples that reach each leaf, or bounds on the loss reduction. The pseudo-code for the training procedure of TREE-G is given in Algorithm 2. Note that the only difference between the training process of TREE-G and the training process of standard DTs is in the grid search for the optimal split parameters where TREE-G tunes additional parameters according to the special split function it uses. This induces a penalty in terms of running time that can be controlled (see Section 4). For example, it helps to bound the possible values of d . Bounding the search space for the pointer $*$ can also improve efficiency. In particular, it is possible to limit every split-node to consider only pointers to ancestors at a limited distance a from it. Empirically, we discovered that bounding $d \leq 2$ and $a \leq 2$ is sufficient to achieve high performance without incurring a significant penalty.

Algorithm 2: Train TREE-G

Data: D : training dataset of labeled graphs or labeled nodes.**Result:** Root of a trained TREE-G T **Pre-Defined:** Optimization criterion (e.g., Gini score or the L_2 loss), stop conditions (e.g., minimal gain, maximal depth, minimal examples in leaf), label computation rule (e.g., majority vote or average of labels)**Function** TrainTREE-G(D):

```

  root ← CreateLeaf( $D$ )
  while No stop condition is met do
     $N$  ← leaf with maximal potential gain with respect
      to the pre-defined optimization criterion
    SplitLeaf( $N$ )
  return root

```

Function CreateLeaf(D):

```

  Create a node  $N$ 
  Compute the label using the label computation rule and
  store it in  $N$ 
  Store the examples  $D$  in  $N$  as  $D_N$ 
  return  $N$ 

```

Function SplitLeaf(N):

```

  Select the best split parameters  $k_N, d_N, *N, \rho_N, r_N$ ,
  threshold  $\theta_N$  and  $AGG_N$  in the case of graph labeling,
  with respect to the pre-defined optimization criterion
  For each different graph in  $D_N$  generate two subsets of
  the vertices of  $G$ ,  $S_{N,+}, S_{N,-}$  using the chosen param-
  eters of  $N$ , following Equation 3, and store in  $N$ .
  Split the examples  $D_N$  into two disjoint subsets
   $D_1$  and  $D_2$  by applying the split function to each
  example in  $D_N$ , and comparing it to the thresh-
  old:  $\phi_{k_N, d_N, *N, \rho_N, r_N} > \theta_N$  for vertex labeling or
   $\phi_{k_N, d_N, *N, \rho_N, r_N, AGG_N} >$  for graph labeling
  Create two child leaves for  $N$ :
  CreateLeaf( $D_1$ )
  CreateLeaf( $D_2$ )
  return the split-node  $N$ 

```

 $T \leftarrow$ TrainTREE-G(D)

In the description above we described the way a single TREE-G is learned and used. Multiple TREE-Gs can be used to form ensembles using Gradient Boosted Trees (Friedman 2001), Random Forests (Ho 1995), or any other ensemble learning method.

3.5 TREE-G for Graph Labeling

So far we have discussed TREE-G in the context of vertex labeling tasks. We now shift to graph labeling tasks in which a graph is given and its label needs to be predicted. While the general structure of TREE-G is preserved, some modifications are required.

In graph labeling tasks, only the adjacency matrix A and the matrix of stacked feature vectors X are given. Since there is no target vertex i , instead of taking some entry of the vector $(A^d \circ M_r(S_{u,\rho})) \mathbf{f}_k$ to compare against a threshold, the split function aggregates all its entries to produce a scalar. This aggregation should be permutation-invariant to be stable under

graph isomorphism. We use one of *sum*, *mean*, *min*, *max* for this purpose. The selected aggregation function is specified as another parameter, AGG , to the split function ϕ . Hence, the split function used in graph labeling tasks is defined as follows:

$$\phi_{k,d,*\rho,r,AGG}(A, X) = AGG \left((A^d \circ M_r(S_{*,\rho})) \mathbf{f}_k \right)$$

In addition, for graph labeling tasks, two additional walk types are used:

3. **Target Walks:** Only consider walks ending in the vertices in S .
4. **Target-Source Walks:** Only consider walks starting and ending in vertices of S .

Therefore the parameter r can take a value of 3 for target walks or 4 for target-source walks. Similarly to types 1 and 2, it can be implemented by masking some entries of the matrix A^d given a subset of the vertices S . For target walks, zero out are the rows of A^d corresponding to vertices not in S . For target-source walks, zero out are the rows and columns of A^d corresponding to vertices not in S . In the Appendix we empirically show that the four walk types are not superfluous. Notice that target masking (including target-source) is not useful and leads to redundant computation in the vertex labeling setting since in this setting only the i^{th} coordinate of the vector $(A^d \circ M_r(S_{*,\rho})) \mathbf{f}_k$ is used when making predictions for the i^{th} vertex. However, in the graph labeling, predictions are made for the entire graph and use the entire vector.

Masking types that include zeroing out rows, eventually zero out entries in the vector $(A^d \circ M_r(S_{*,\rho})) \mathbf{f}_k$. As in graph labeling tasks this vector is aggregated, zeros may affect the results of some aggregations e.g. *min*. Therefore, in the case of target, target-source and cycle walks, the aggregation is only performed on the entries that are in the selected subset $S_{*,\rho}$; i.e., $AGG \left(((A^d \circ M_r(S_{*,\rho})) \mathbf{f}_k)_{S_{*,\rho}} \right)$

As in the vertex labeling setting, after the graph is introduced to a split-node, two subsets of the vertices are generated to be used by the descendant split-nodes in the tree, using the same definitions as in vertex labeling. In particular, the chosen aggregation is not part of the subsets generation mechanism. An exception is made when the selected aggregation function in the split-node is *sum*, in which case the subsets are defined with respect to a scaled threshold $\theta/|S_{*,\rho}|$ instead of θ . This is because in this case any value that is greater than this scaled threshold contributes to ϕ towards passing the threshold.

4 Theoretical Properties of TREE-G

In this section we discuss some theoretical properties of TREE-G, including its computational complexity. Due to space limitations proofs are deferred to the Appendix. We first consider the invariance properties of TREE-G. Recall that a graph is described by its features $\mathbf{f}_1, \mathbf{f}_2, \dots, \mathbf{f}_k$ and adjacency matrix A . If we apply a permutation π to the vertices, with the corresponding permutation matrix $P = P_\pi$, we obtain a new adjacency matrix PAP^T and features $P\mathbf{f}_k$. Clearly, we would like a model that acts on graphs to provide the same output for

the original and permuted graphs. In graph labeling tasks this means we would like the output of TREE-G to be invariant to P . In the case of vertex labeling, we expect equivariance such that if the graph is permuted with P then the prediction for the vertex $\pi(i)$ would be the same as the prediction for the vertex i in the original graph (Gilmer et al. 2017b). The next lemma shows that this is the case for TREE-G.

Lemma 4.1. *TREE-G is invariant to permutations for graph labeling and equivariant for vertex labeling.*

Next, we discuss the computational complexity of TREE-G. The main computational challenge in training decision trees is finding the optimal parameters for the split function. The cost of this operation is proportional to the number of features to consider. This is also the case in TREE-G

Lemma 4.2. *The running time of searching for the optimal split function parameters in TREE-G is linear in: the number of features, the maximal walk length, and the maximal number of considered pointed ancestors.*

Next, we present results on the expressive power of TREE-G. As mentioned above, the advantage of TREE-G lies in its ability to focus on subsets. We prove that TREE-G is strictly more expressive than a limited version of it where subsets are not used (equivalently, the used subset is always V).

Lemma 4.3. *There exist graphs that are separable by TREE-G, but are inseparable if TREE-G is limited to not use subsets.*

Notice that standard decision trees that only use the input features in split-nodes are a special case of such limited TREE-G, where additionally the propagation depth is limited to 0. Therefore, an immediate conclusion is that TREE-G is strictly more expressive than these standard trees.

Finally, the following lemma shows that TREE-G can express classification rules that GNNs cannot.

Lemma 4.4. *There exist graphs that cannot be separated by GNNs but can be separated by TREE-G.*

The Lemma shows that there exist two graphs such that any GNN will assign both graphs with the same label, whereas there exists a TREE-G that will assign a different label to each graph. Since the expressive power of GNNs is bounded by the expressive power of the 1-WL test (Xu et al. 2019; Garg, Jegelka, and Jaakkola 2020) we conclude that the expressive power of TREE-G is not bounded by the 1-WL test.

5 Empirical Evaluation

We conducted experiments using TREE-G on graph and vertex labeling benchmarks and compared its performance to popular GNNs, graph kernels, and other tree-based methods.⁶

5.1 Datasets

Graph Prediction Tasks: We used nine graph classification benchmarks from TUDatasets (Morris et al. 2020), and one large-scale dataset from the Open Graph Benchmark (Hu et al. 2020). *IMDb-B & IMDb-M* (Yanardag and Vishwanathan 2015) are movie collaboration datasets, where each graph is labeled with a genre. *Mutag*, *NCII*, *Protiens*, *D&D*, *Enzymes*,

Mutagenicity & PTC (Shervashidze et al. 2011; Kriege and Mutzel 2012; Riesen and Bunke 2008) are datasets of chemical compounds. In each dataset, the goal is to classify compounds according to some property of interest. *molHIV* (Hu et al. 2020) is a large-scale dataset of molecules that may inhibit HIV.

Node Prediction Tasks: We used the three citation graphs tasks *Cora*, *Citeseer* & *Pubmed* from Planetoid (Yang, Cohen, and Salakhutdinov 2016). Arxiv is a large scale citation network related task of Computer Science arXiv papers (Hu et al. 2020).

Cornell, *Actor* & *County* are heterophilic datasets where *Cornell* and *Actor* (Pei et al. 2020) are web links networks with the task of classifying nodes to one of five categories whereas the *County* dataset (Jia and Benson 2020) contains a regression task of predicting unemployment rates based on county-level election map network.

More details on all datasets are given in the Appendix.

5.2 Baselines and Protocols

We compared TREE-G as a weak-learner in GBT to the following popular GNNs and graph-kernels: Graph Convolution Network (GCN) (Kipf and Welling 2017), Graph Isomorphism Network (GIN) (Xu et al. 2019), Graph Attention Network (GAT) (Veličković et al. 2018) and its newer version GATv2 (Brody, Alon, and Yahav 2022), GraphSAGE (Hamilton, Ying, and Leskovec 2018), Weisfeiler-Lehman Graph kernels (WL) (Shervashidze et al. 2011) and a Random-Walk kernels (RW) (Gärtner, Flach, and Wrobel 2003). We also compared to two methods combining GNNs and DTs: BGNN (Ivanov and Prokhorenkova 2021) and XGraphBoost (Deng et al. 2021). Additionally, we report two ablations of TREE-G where we do not use the topology of the graph and disable the subsets mechanism. This is done by limiting the propagation depth to 0 and/or the distance of the ancestors from which to consider subsets from, to 0. In TREE-G ($a = 0$) we always use the whole graph as the selected subset, which is equivalent to not using subsets. In TREE-G ($d = 0, a = 0$) we do not use the topology of the graph (i.e. we always use A^0) and we always use the entire graph as the selected subset. Notice TREE-G ($d = 0, a = 0$) is equivalent to standard trees that only act on the input features.

For graph prediction tasks we report the average accuracy and std of a 10-fold nested cross-validation, except for the molHIV dataset for which we use official pre-defined splits provided in Hu et al. (2020) and the metric is AUC. For the node prediction tasks we report average accuracy and std using the pre-defined splits provided in the data, except for the county dataset, which is a regression task, and hence we use RMSE instead of accuracy. More details on the evaluation protocol including the tuned hyper-parameters for each algorithm are provided in the Appendix.

5.3 Results

The results are summarized in Tables 1, 2 & 3. TREE-G outperformed all other tree-based approaches (XGraphBoost, BGNN, TREE-G ($a = 0$), on all 17 tasks. It outperformed graph kernels in 9 out of the 10 graph classification tasks

⁶Code is available at github.com/mayabecklerspeicher/TREE-G

Baseline	Proteins	Mutag	D&D	NCI1	PTC	Enzymes	Mutagenicity
XGraphBoost	67.0 ± 2.5	86.1 ± 3.8	72.9 ± 2.0	61.9 ± 7.1	51.3 ± 5.0	58.5 ± 1.0	71.2 ± 2.5
BGNN	70.5 ± 2.4	80.2 ± 0.5	71.2 ± 0.9	70.5 ± 2.0	55.5 ± 0.3	58.1 ± 1.5	65.0 ± 0.9
RW	72.5 ± 0.9	85.0 ± 4.5	70.0 ± 5.0	69.0 ± 0.5	57.2 ± 1.0	55.1 ± 1.3	74.5 ± 0.5
WL	74.0 ± 3.0	82.0 ± 0.2	76.1 ± 3.0	82.5 ± 3.1	58.9 ± 4.2	52.1 ± 3.0	80.1 ± 2.0
GAT	70.0 ± 1.9	84.4 ± 0.6	74.4 ± 0.3	74.9 ± 0.1	56.2 ± 7.3	58.5 ± 1.2	72.2 ± 1.8
GCN	73.2 ± 1.0	84.6 ± 0.7	71.2 ± 2.0	76.0 ± 0.9	59.4 ± 10.3	60.1 ± 3.0	69.9 ± 1.5
GraphSAGE	70.4 ± 2.0	86.1 ± 0.7	72.7 ± 1.0	76.1 ± 2.2	67.1 ± 12.6	58.2 ± 5.9	64.1 ± 0.3
GIN	72.2 ± 1.9	84.5 ± 0.8	75.2 ± 2.9	79.1 ± 1.4	55.6 ± 11.1	59.5 ± 0.2	69.4 ± 1.2
GATv2	73.5 ± 0.9	84.0 ± 1.5	75.0 ± 1.3	77.0 ± 0.8	59.5 ± 2.1	61.2 ± 3.0	72.0 ± 0.9
TREE-G (d=0, a=0)	73.4 ± 0.9	86.1 ± 6.0	74.4 ± 4.7	70.4 ± 0.7	54.5 ± 1.3	58.5 ± 0.9	77.0 ± 1.9
TREE-G (a=0)	74.7 ± 0.9	89.3 ± 7.6	76.0 ± 4.5	73.6 ± 2.2	59.0 ± 0.8	58.1 ± 3.9	77.2 ± 1.5
TREE-G	75.6 ± 1.1	91.1 ± 3.5	76.2 ± 4.5	75.9 ± 1.5	59.1 ± 0.7	59.6 ± 0.9	83.0 ± 1.7

Table 1: Empirical comparison of graph classification tasks. For all the datasets the accuracy and std are reported. The best score is highlighted in bold.

Baseline	Cora	Citeseer	Pubmed	Arxiv	Cornell	Actor	Country
XGraphBoost	60.5 ± 3.0	50.5 ± 3.4	61.9 ± 2.2	65.8 ± 1.9	62.1 ± 0.3	27.2 ± 1.4	3.0 ± 0.5
BGNN	71.2 ± 5.0	69.1 ± 3.9	59.9 ± 0.1	67.0 ± 0.1	68.2 ± 0.5	31.1 ± 0.7	1.1 ± 0.2
GAT	83.0 ± 0.6	70.0 ± 1.5	79.1 ± 0.7	73.6 ± 0.1	74.0 ± 0.5	34.0 ± 0.9	1.6 ± 0.2
GCN	81.0 ± 0.8	72.1 ± 0.3	79.0 ± 0.8	71.7 ± 0.2	65.6 ± 0.1	28.8 ± 0.2	1.8 ± 0.1
GraphSAGE	81.4 ± 0.7	73.4 ± 0.8	78.4 ± 0.4	71.7 ± 0.2	71.1 ± 1.0	32.1 ± 1.5	2.0 ± 0.4
GIN	80.0 ± 1.2	75.1 ± 1.9	75.3 ± 0.9	73.8 ± 1.4	69.0 ± 1.3	30.4 ± 0.3	0.9 ± 0.1
GATv2	83.1 ± 0.9	73.9 ± 1.5	79.4 ± 0.5	74.0 ± 2.1	72.5 ± 0.7	34.2 ± 1.9	1.6 ± 0.5
TREE-G (d=0, a=0)	80.3 ± 2.1	70.2 ± 2.0	74.0 ± 0.5	67.1 ± 1.5	71.0 ± 0.9	29.1 ± 0.5	1.9 ± 0.2
TREE-G (a=0)	80.8 ± 0.5	74.3 ± 1.9	72.9 ± 1.6	67.5 ± 0.6	71.8 ± 1.4	32.5 ± 1.7	1.2 ± 0.8
TREE-G	83.5 ± 1.5	74.5 ± 1.0	78.0 ± 3.1	74.7 ± 1.1	73.0 ± 1.4	37.0 ± 0.4	0.9 ± 0.1

Table 2: Empirical comparison of node labeling tasks. For the country task, the RMSE and std are reported. For the rest of the datasets, the accuracy and std are reported. The best score is highlighted in bold.

and outperformed GNNs in 7 out of these 10 tasks. Specifically, TREE-G improved upon GNN approaches by a margin greater than 6.4 percentage points on the IMDb-M dataset and a margin greater than 4.9 and 2.4 percentage points on Mutag and Proteins datasets, respectively. On the vertex labeling tasks, TREE-G outperformed GNNs in 4 out of the 7 tasks and was on par with the leading GNN approach in every task. Note that TREE-G always outperformed TREE-G ($a = 0$), which indicates that the use of subsets leads to improved performance. Consequently, these results demonstrate that a "pure" decision tree method, when tailored for graphs, can rival popular GNNs and graph kernels, potentially even surpassing them.

6 Explaining TREE-G Outputs

In this section, we formalize the explanation mechanism of TREE-G. We show that the selected subsets can be used to explain the predictions made by TREE-G. This explanation mechanism ranks the vertices and edges of the graph according to their presence in chosen subsets and thus highlights the parts of the graph that contribute the most to the prediction.

Following (Hirsch and Gilad-Bachrach 2021), the explanation mechanism ranks the vertices of the graph according

to their presence in selected subsets and thus highlights the parts of the graph that contribute the most to the prediction. As is the case in most tree methods, TREE-G does not combine multiple features in the splits, which makes them more interpretable and less prone to overfitting.

Let G be a graph that we want to classify using a decision tree T . The key idea in the construction of our importance weight is that if a vertex appears in many selected subsets when calculating the tree output for G , then it has a large effect on the decision. This is quantified as follows: consider the path in T that G traverses when it is being classified. For each vertex i in G , count how many selected subsets in the path contain i and denote this number by $n(i)$. In order to avoid dependence on the size of T we convert the values $n_T(1), \dots, n_T(|V|)$ into a new vector $r_T(1), \dots, r_T(|V|)$ where $r_T(i)$ specifies the rank of $n_T(i)$ in $n_T(1), \dots, n_T(|V|)$ after sorting in decreasing order. For example, $r_T(i) = 1$ if the vertex i appears most frequently in chosen subsets compared to any other vertex. We shall use $2^{-r_T(i)}$ to measure the importance of the rank so that low ranks have high importance.

Now assume we have an ensemble of decision trees \mathcal{T} (e.g., generated by GBT). For any $T \in \mathcal{T}$ in the ensemble let y_T be




Figure 3: Vertex-level explanations for two graphs from the Red Isolated Vertex problem (a) and two graphs from the Mutagenicity problem (b). The size of vertices corresponds to their importance according to the explanation mechanism.

Baseline	IMDb-B	IMDb-M	molHIV
XGraphBoost	56.1 ± 1.4	41.5 ± 0.9	60.1 ± 5.3
BGNN	68.0 ± 1.5	46.2 ± 2.3	77.0 ± 6.1
RW	65.0 ± 1.2	45.1 ± 0.9	76.7 ± 2.1
WL	72.8 ± 1.9	50.9 ± 3.8	75.7 ± 0.8
GAT	70.6 ± 3.2	47.0 ± 1.9	82.4 ± 3.6
GCN	69.4 ± 1.9	50.0 ± 3.0	76.6 ± 0.0
GraphSAGE	69.9 ± 4.4	47.8 ± 1.9	79.2 ± 1.2
GIN	71.2 ± 3.4	48.8 ± 5.0	77.8 ± 1.3
GATv2	71.5 ± 1.2	47.6 ± 0.5	81.9 ± 4.5
TREE-G (d=0, a=0)	72.6 ± 2.4	52.3 ± 1.5	70.1 ± 0.5
TREE-G (a=0)	72.5 ± 2.7	52.7 ± 1.5	78.5 ± 1.1
TREE-G	73.0 ± 2.4	56.4 ± 0.9	83.5 ± 2.9

Table 3: Empirical comparison of graph classification tasks. The AUC and std are reported for the molHIV dataset, and accuracy and std for the rest. The best result is in bold.

the value predicted by T for the graph G . We weight the rank importance by $|y_T|$ so that trees with a larger contribution to the decision have a greater effect on the importance measure. Note that boosted trees tend to have large variability in leaf values between trees in the ensemble (Rashmi and Gilad-Bachrach 2015). Therefore, the importance-score of vertex i in the ensemble is the weighted average:

$$\text{importance}(i) = \frac{\sum_{T \in \mathcal{T}} |y_T| 2^{-r_T(i)}}{\sum_{j \in |V|} \sum_{T \in \mathcal{T}} |y_T| 2^{-r_T(j)}}. \quad (4)$$

The proposed importance score is non-negative and sums to 1. We present examples of vertex explanations for two problems: Red Isolated Vertex, and Mutagenicity.

Red Isolated Vertex This is a synthetic task to determine whether exactly one red vertex is isolated in the graph. The data consists of 1000 graphs with 50 vertices each, with two features: the constant feature 1 and a binary blue/red color.

Mutagenicity This is a dataset of molecules where the goal is to classify chemical compounds into mutagen and non-mutagen.

For each task, we trained and tested GBT with 50 TREE-G estimators. We limited the propagation depth to 2 and the distance of the ancestors from which to consider subsets, to 2 as well. We used a 80/20 random train-test splits. Figures 3a and 3b present vertex explanations for two graphs from the test sets of the Red-Isolated-Node experiments and the

Mutagenicity experiments, respectively. The size of a vertex corresponds to its importance computed by the explanations mechanism. For the Red-Isolated-Node problem, it is shown that TREE-G attended to isolated vertices, and even more so to red isolated vertices.

For the Mutagenicity problem, TREE-G attended to NO_2 (green and yellow subgraph) and carbon-rings (blue cycles), which are known to have mutagenic effects (Debnath et al. 1991).

Edge-Level Explanations Similarly to vertex importance, edge importance can be computed. We use the selected subsets to rank the edges of a given graph by their *importance*. Note that each split criterion uses a subset of the graph’s edges, according to d and the active subset. For example, a split criterion may consider only walks of length 2 to a specific subset of vertices, which induces the *used* edges by the split criterion. Let E_1, \dots, E_k denote the used edges along the prediction walk $\text{walk}_m(G)$. We now rank each edge $e_{(i,j)} \in E$ by $r_m((i,j)) = |\{j | (i,j) \in E_j\}|$, i.e., the rank of an edge is the number of split-nodes in the prediction walk that use this edge. The ranks $r_m(i,j)$ are then combined with the predicted value y_m .

7 Conclusion

Graphs arise in many important applications of machine learning. In this work we proposed a novel method, TREE-G, to perform different prediction tasks on graphs, inspired by the success of tree-based methods on tabular data. Our empirical results revealed that TREE-G surpassed other tree-based models. Additionally, when compared with prevalent GNNs and graph kernels, TREE-G matched their performance in all scenarios and exceeded them in several benchmarks. Our theoretical analysis showed that allowing trees to focus on substructures of graphs strictly improves their expressive power. We also showed that the expressive power of TREE-G is different from the expressive power of GNNs, as there are graphs that GNNs fail to tell apart, while TREE-G manages to separate them. Finally, we also proposed visualizations for TREE-G and the predictions in makes to support explainability.

Acknowledgments

This work was supported by a grant from the Tel Aviv University Center for AI and Data Science (TAD) and by the Israeli Science Foundation research grant 1186/18 and 1437/22.

References

- Barnett, I.; Malik, N.; Kuijjer, M. L.; Mucha, P. J.; and Onnela, J.-P. 2016. Feature-based classification of networks. *arXiv preprint arXiv:1610.05868*.
- Barrett, D. G. T.; Hill, F.; Santoro, A.; Morcos, A. S.; and Lillicrap, T. 2018. Measuring abstract reasoning in neural networks. arXiv:1807.04225.
- Brody, S.; Alon, U.; and Yahav, E. 2022. How Attentive are Graph Attention Networks? arXiv:2105.14491.
- de Lara, N.; and Pineau, E. 2018. A simple baseline algorithm for graph classification. *arXiv preprint arXiv:1810.09155*.
- Debnath, A.; Lopez de Compadre, R.; Debnath, G.; Shusterman, A.; and Hansch, C. 1991. Structure-activity relationship of mutagenic aromatic and heteroaromatic nitro compounds. Correlation with molecular orbital energies and hydrophobicity. *Journal of medicinal chemistry*, 34(2): 786–797.
- Defferrard, M.; Bresson, X.; and Vandergheynst, P. 2016. Convolutional Neural Networks on Graphs with Fast Localized Spectral Filtering. In Lee, D.; Sugiyama, M.; Luxburg, U.; Guyon, I.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc.
- Deng, D.; Chen, X.; Zhang, R.; Lei, Z.; Wang, X.; and Zhou, F. 2021. XGraphBoost: Extracting Graph Neural Network-Based Features for a Better Prediction of Molecular Properties. *Journal of Chemical Information and Modeling*, 61(6): 2697–2705. PMID: 34009965.
- Duvenaud, D.; Maclaurin, D.; Aguilera-Iparraguirre, J.; Gómez-Bombarelli, R.; Hirzel, T.; Aspuru-Guzik, A.; and Adams, R. P. 2015. Convolutional Networks on Graphs for Learning Molecular Fingerprints. arXiv:1509.09292.
- Friedman, J. H. 2000. Greedy Function Approximation: A Gradient Boosting Machine. *Annals of Statistics*, 29: 1189–1232.
- Friedman, J. H. 2001. Greedy function approximation: a gradient boosting machine. *Annals of statistics*, 1189–1232.
- Friedman, J. H. 2002. Stochastic Gradient Boosting. *Comput. Stat. Data Anal.*, 38(4): 367–378.
- Ganea, O.-E.; Pattanaik, L.; Coley, C. W.; Barzilay, R.; Jensen, K. F.; Green, W. H.; and Jaakkola, T. S. 2021. GeoMol: Torsional Geometric Generation of Molecular 3D Conformer Ensembles. arXiv:2106.07802.
- Garg, V. K.; Jegelka, S.; and Jaakkola, T. 2020. Generalization and Representational Limits of Graph Neural Networks. arXiv:2002.06157.
- Gilmer, J.; Schoenholz, S. S.; Riley, P. F.; Vinyals, O.; and Dahl, G. E. 2017a. Neural Message Passing for Quantum Chemistry. arXiv:1704.01212.
- Gilmer, J.; Schoenholz, S. S.; Riley, P. F.; Vinyals, O.; and Dahl, G. E. 2017b. Neural Message Passing for Quantum Chemistry.
- Gori, M.; Monfardini, G.; and Scarselli, F. 2005. A new model for learning in graph domains. In *Proceedings. 2005 IEEE International Joint Conference on Neural Networks, 2005.*, volume 2, 729–734 vol. 2.
- Grinsztajn, L.; Oyallon, E.; and Varoquaux, G. 2022. Why do tree-based models still outperform deep learning on tabular data?
- Gärtner, T.; Flach, P.; and Wrobel, S. 2003. On Graph Kernels: Hardness Results and Efficient Alternatives. In Schölkopf, B.; and Warmuth, M. K., eds., *Computational Learning Theory and Kernel Machines — Proceedings of the 16th Annual Conference on Computational Learning Theory and 7th Kernel Workshop (COLT/Kernel 2003) August 24-27, 2003, Washington, DC, USA*, volume 2777 of *Lecture Notes in Computer Science*, 129–143. Springer, Berlin–Heidelberg, Germany.
- Hamilton, W.; Ying, Z.; and Leskovec, J. 2017. Inductive Representation Learning on Large Graphs. In Guyon, I.; Luxburg, U. V.; Bengio, S.; Wallach, H.; Fergus, R.; Vishwanathan, S.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Hamilton, W. L.; Ying, R.; and Leskovec, J. 2018. Inductive Representation Learning on Large Graphs. arXiv:1706.02216.
- Han, T.; Jiang, D.; Zhao, Q.; Wang, L.; and Yin, K. 2018. Comparison of random forest, artificial neural networks and support vector machine for intelligent diagnosis of rotating machinery. *Transactions of the Institute of Measurement and Control*, 40(8): 2681–2693.
- He, T.; Heidemeyer, M.; Ban, F.; Cherkasov, A.; and Ester, M. 2017. SimBoost: a read-across approach for predicting drug–target binding affinities using gradient boosting machines. *Journal of cheminformatics*, 9(1): 1–14.
- Heaton, J. 2016. An empirical analysis of feature engineering for predictive modeling. In *SoutheastCon 2016*, 1–6. IEEE.
- Hirsch, R.; and Gilad-Bachrach, R. 2021. Trees with Attention for Set Prediction Tasks. In Meila, M.; and Zhang, T., eds., *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, 4250–4261. PMLR.
- Ho, T. K. 1995. Random decision forests. In *Proceedings of 3rd international conference on document analysis and recognition*, volume 1, 278–282. IEEE.
- Hu, W.; Fey, M.; Zitnik, M.; Dong, Y.; Ren, H.; Liu, B.; Catasta, M.; and Leskovec, J. 2020. Open Graph Benchmark: Datasets for Machine Learning on Graphs.
- Ingraham, J.; Garg, V.; Barzilay, R.; and Jaakkola, T. 2019. Generative Models for Graph-Based Protein Design. In Wallach, H.; Larochelle, H.; Beygelzimer, A.; d'Alché-Buc, F.; Fox, E.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Ivanov, S.; and Prokhorenkova, L. 2021. Boost then Convol: Gradient Boosting Meets Graph Neural Networks. arXiv:2101.08543.
- Jia, J.; and Benson, A. R. 2020. Residual correlation in graph neural network regression. In *Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining*, 588–598.

- Kipf, T. N.; and Welling, M. 2017. Semi-Supervised Classification with Graph Convolutional Networks. arXiv:1609.02907.
- Kriege, N.; and Mutzel, P. 2012. Subgraph Matching Kernels for Attributed Graphs. In *Proceedings of the 29th International Conference on Machine Learning, ICML'12*, 291–298. Madison, WI, USA: Omnipress. ISBN 9781450312851.
- Lei, X.; and Fang, Z. 2019. GBDTCDA: predicting circRNA-disease associations based on gradient boosting decision tree with multiple biological data fusion. *International journal of biological sciences*, 15(13): 2911.
- Li, K.; and Tu, L. 2019. Link Prediction for Complex Networks via Random Forest. *Journal of Physics: Conference Series*, 1302: 022030.
- Li, P.; Qin, Z.; Wang, X.; and Metzler, D. 2019. Combining Decision Trees and Neural Networks for Learning-to-Rank in Personal Search. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD '19*, 2032–2040. New York, NY, USA: Association for Computing Machinery. ISBN 9781450362016.
- Li, Y.; Tarlow, D.; Brockschmidt, M.; and Zemel, R. 2017. Gated Graph Sequence Neural Networks. arXiv:1511.05493.
- Miah, Y.; Prima, C.; Seema, S.; Mahmud, M.; and Kaiser, M. S. 2021. *Performance Comparison of Machine Learning Techniques in Identifying Dementia from Open Access Clinical Datasets*, 79–89. ISBN 978-981-15-6047-7.
- Morris, C.; Kriege, N. M.; Bause, F.; Kersting, K.; Mutzel, P.; and Neumann, M. 2020. TUDataset: A collection of benchmark datasets for learning with graphs. In *ICML 2020 Workshop on Graph Representation Learning and Beyond (GRL+ 2020)*.
- Nikolentzos, G.; Siglidis, G.; and Vazirgiannis, M. 2021. Graph Kernels: A Survey. *Journal of Artificial Intelligence Research*, 72: 943–1027.
- Patel, R.; and Guo, Y. 2021. Graph Based Link Prediction between Human Phenotypes and Genes. arXiv:2105.11989.
- Pei, H.; Wei, B.; Chang, K. C.-C.; Lei, Y.; and Yang, B. 2020. Geom-GCN: Geometric Graph Convolutional Networks. In *International Conference on Learning Representations*.
- Rashmi, K. V.; and Gilad-Bachrach, R. 2015. DART: Dropouts meet Multiple Additive Regression Trees. arXiv:1505.01866.
- Riesen, K.; and Bunke, H. 2008. IAM Graph Database Repository for Graph Based Pattern Recognition and Machine Learning. In da Vitoria Lobo, N.; Kasparis, T.; Roli, F.; Kwok, J. T.; Georgiopoulos, M.; Anagnostopoulos, G. C.; and Loog, M., eds., *Structural, Syntactic, and Statistical Pattern Recognition*, 287–297. Berlin, Heidelberg: Springer Berlin Heidelberg. ISBN 978-3-540-89689-0.
- Rossi, E.; Chamberlain, B.; Frasca, F.; Eynard, D.; Monti, F.; and Bronstein, M. 2020. Temporal Graph Networks for Deep Learning on Dynamic Graphs. arXiv:2006.10637.
- Shervashidze, N.; Schweitzer, P.; van Leeuwen, E. J.; Mehlhorn, K.; and Borgwardt, K. M. 2011. Weisfeiler-Lehman Graph Kernels. *J. Mach. Learn. Res.*, 12: 2539–2561.
- Shervashidze, N.; Vishwanathan, S. V. N.; Petri, T.; Mehlhorn, K.; and Borgwardt, K. M. 2009. Efficient graphlet kernels for large graph comparison. *Journal of Machine Learning Research - Proceedings Track*, 5: 488–495.
- Shwartz-Ziv, R.; and Armon, A. 2021. Tabular Data: Deep Learning is Not All You Need. In *8th ICML Workshop on Automated Machine Learning (AutoML)*.
- Sperduti, A.; and Starita, A. 1997. Supervised neural networks for the classification of structures. *IEEE Transactions on Neural Networks*, 8(3): 714–735.
- Stokes, J. M.; Yang, K.; Swanson, K.; Jin, W.; Cubillos-Ruiz, A.; Donghia, N. M.; MacNair, C. R.; French, S.; Carfrae, L. A.; Bloom-Ackermann, Z.; Tran, V. M.; Chiappino-Pepe, A.; Badran, A. H.; Andrews, I. W.; Chory, E. J.; Church, G. M.; Brown, E. D.; Jaakkola, T. S.; Barzilay, R.; and Collins, J. J. 2020. A Deep Learning Approach to Antibiotic Discovery. *Cell*, 180(4): 688–702.e13.
- Veličković, P.; Cucurull, G.; Casanova, A.; Romero, A.; Liò, P.; and Bengio, Y. 2018. Graph Attention Networks. In *International Conference on Learning Representations*.
- Wu, M.; Huang, Y.; Zhao, L.; and He, Y. 2018. Link Prediction Based on Random Forest in Signed Social Networks. In *2018 10th International Conference on Intelligent Human-Machine Systems and Cybernetics (IHMSC)*, volume 02, 251–256.
- Wu, Z.; Pan, S.; Chen, F.; Long, G.; Zhang, C.; and Philip, S. Y. 2020. A comprehensive survey on graph neural networks. *IEEE transactions on neural networks and learning systems*, 32(1): 4–24.
- Xu, K.; Hu, W.; Leskovec, J.; and Jegelka, S. 2019. How Powerful are Graph Neural Networks? In *International Conference on Learning Representations*.
- Yanardag, P.; and Vishwanathan, S. 2015. Deep Graph Kernels. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '15*, 1365–1374. New York, NY, USA: Association for Computing Machinery. ISBN 9781450336642.
- Yang, Z.; Cohen, W. W.; and Salakhutdinov, R. 2016. Revisiting Semi-Supervised Learning with Graph Embeddings.
- Ying, R.; He, R.; Chen, K.; Eksombatchai, P.; Hamilton, W. L.; and Leskovec, J. 2018. Graph Convolutional Neural Networks for Web-Scale Recommender Systems. *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*.