Improving Deep Forest by Exploiting High-order Interactions

Author: Yi-He CHEN*, Shen-Huan LYU*, and Yuan JIANG (Corresponding author) *: equal contribution.

A pre-printed manuscript which is accepted by ICDM 2021



Learning And Mining from DatA

LAMDA Group Department of Computer Science and Technology National Key Laboratory for Novel Software Technology



Nanjing University Nanjing, China

Citation: Yi-He Chen*, Shen-Huan Lyu*, and Yuan Jiang. Improving deep forest by exploiting high-order interactions. In: Proceedings of the 21th IEEE International Conference on Data Mining, 2021, pp. 1030-1035. doi: https://doi.org/10.1109/ICDM51629.2021.00118

Contents

1	Intr	oduction	1										
2	High-order interaction deep forest												
	2.1	2.1 Extract feature interactions through gRIT and ERF											
	2.2	Stability-aware interaction utilization	3										
	2.3	Adaptive layer growth	4										
3	Exp	eriments	6										
	3.1	Synthetic data	6										
	3.2	Real dataset	7										
	3.3	Performance comparison	8										
	3.4	Computational complexity comparison with gcForest	9										
4	Con	clusion	9										

Improving Deep Forest by Exploiting High-order Interactions*

Yi-He Chen[†]

Shen-Huan Lyu[‡]

Yuan Jiang[§]

Abstract

Recent studies on deep forests have shown that deep learning frameworks can be built on nondifferentiable modules without a backpropagation training process. However, the feature representations of deep forests only consist of predicted class probabilities. The information these class probabilities deliver is very limited and lacks diversity, especially when the number of output labels is far less than the number of input features. Besides, the prediction-based representations require us to save multiple layers of random forests to use them during testing, which is high-memory and high-time cost. In this paper, we propose a novel deep forest model that utilizes high-order interactions of input features to generate more informative and diverse feature representations. Specifically, we design a generalized version of Random Intersection Trees (gRIT) to discover stable high-order interactions and apply Activated Linear Combination (ALC) to transform them into hierarchical distributed representations. These interaction-based representations obviate the need to store random forests in the front layers, thus greatly improving the computational efficiency. Our experiments show that our method achieves highly competitive predictive performance with significantly reduced time and memory cost.

1 Introduction

In recent years, deep forests (DFs) have achieved state-of-the-art performance on the categorical and mixed modeling tasks, e.g., financial application [Zha+19], medical application [Sun+20], [Guo+18] and geoscience [Yan+18], while deep neural networks (DNNs) dominate the numerical tasks, e.g., computer vision (CV) [He+15], automatic speech recognition (ASR) [Kri+20] and natural language processing (NLP) [Dev+18]. Although the recent development of deep neural networks on numerical data is rather attractive, more real-world problems belong to categorical and mixed modeling tasks. In such tasks, there are a large number of attributes from which we can only get qualitative characteristics rather than quantitative characteristics. Therefore, applying and exploring deep forests and other non-neural network deep models to a wide range of tasks is an important direction for the future of deep learning [Ben09].

By summarizing the key ingredients of deep learning may lie in: *layer-by-layer processing*, *sufficient model complexity*, and *in-model feature transformation*, Zhou and Feng [ZF19] proposed gcForest, the first non-NN-style deep models with these characteristics. Essentially, gcForest is a decision tree ensemble that employs a cascade structure to do representation learning. In this cascade structure, each level consists of an ensemble of decision tree forests, i.e. an ensemble of ensembles. Each level receives feature information processed by the preceding level and outputs estimated class probabilities which are then concatenated with the original feature vector as the input to the next level. Lyu, Yang, and Zhou

^{*}This is a preprinted manuscript.

[†]Equal contribution. Nanjing University. E-mail: chenyh@lamda.nju.edu.cn

[‡]Equal contribution. Nanjing University. E-mail: lvsh@lamda.nju.edu.cn

[§]Corresponding author. Nanjing University. E-mail: jiangy@lamda.nju.edu.cn



Figure 1: The framework of High-order Interaction Deep Forest (hiDF).

[LYZ19] provided a detailed theoretical analysis that proved deep forests have sufficient model complexity with enough depth, and the cascaded structure boosts the feature representations layer by layer instead of the predictions.

Although gcForest has shown its great potential both empirically and theoretically, we argue that the prediction-based feature representation of gcForest is a critical deficiency. Firstly, as stated by original authors [ZF19], the predicted class probabilities deliver very limited information. A majority of tasks have much more input features than output labels, which causes the information in predictions to likely be drowned out by original features when concatenating them together. Secondly, since decision tree forests are already pretty stable predictors, an ensemble of different forests may result in similar predictions, which causes feature representations to be redundant and lack diversity. Thirdly, the prediction-based representations rely on the storing of multi-layered forest models to do prediction level-by-level during testing, thus requiring a large amount of memory and time consumption. Therefore, it is necessary to design more informative feature representations with less computational cost for deep forests.

In this paper, we propose a novel deep forest model named high-order interaction Deep Forest (hiDF), which leverages stable high-order interactions of input features to generate informative and diverse feature representations. Specifically, we design a generalized version of Random Intersection Trees (gRIT) to discover stable high-order interactions and apply Activated Linear Combination (ALC) to transform these interactions into new feature representations, which can interact with input features across multiple layers. In such iterations, hiDF can effectively mine high-order interactions between input features and utilize them to improve predictive performance.

Our contributions are twofold. First, the proposed hiDF method firstly provides hierarchical distributed representations from low-order to high-order interactions for deep forests, greatly enhancing the effectiveness and diversity of feature representations. Second, these representations obviate the need to store random forests in the front layers, reducing the time cost and memory requirement by one order of magnitude.

2 High-order interaction deep forest

In this section, we provide a detailed description of our hiDF algorithm. The framework of hiDF is illustrated in Figure 1. It employs a multi-layered structure to do representation learning, each layer consists of an ensemble of decision tree forests as DF. The most significant difference between DF and hiDF is that DF uses the prediction of each layer as representations while hiDF employs high-order feature interactions as representations. hiDF utilizes both feature interactions and multi-layered structure to construct multiple levels of feature representations and learn a hierarchy of explanatory factors internally. This deep structure promotes the *re-use* of features [BCV13] and the ultimate goal is

to form hierarchical distributed representations of data, which is useful to our learning task.

Specifically, hiDF utilizes feature interactions mainly in three steps. First, given an ensemble of random forests fitted on the input features, numerous decision rules are extracted from these decision trees. These decision rules are then processed by generalized Random Intersection Trees (gRIT) to identify prevalent feature interactions. Second, an outer bagging procedure performs the previous step on bootstrap samples. This extra perturbation of the data allows us to do "interaction selection" by assessing the stability associated with the identified interactions. Only those feature interactions with high stability scores are utilized through Activated Linear Combinations (ALC) to generate new feature representations. Finally, a metric-aware mechanism is applied to grow layers adaptively to reduce the risk of overfitting. We will discuss these three steps in more detail in the following subsections.

2.1 Extract feature interactions through gRIT and ERF

We identify feature interaction as a collection of conditions inside a decision rule, a decision rule is of the simple form:

IF:
$$condition_1 \& \dots \& condition_i$$
, THEN: $response$. (1)

Accordingly, a feature interaction can be described as $condition_1\&\ldots\&condition_i$, which is the premise of a decision rule. The conditions are based on input variables. For a continuous variable x_i , it could take the form: $x_i \ge t_i$, where t_i is a determined threshold. For categorical variables, we transform them using one-hot encoding so we can get similar forms. In this work, we only consider the premise of a decision rule regardless of the response. To represent decision rules succinctly, we use signed feature index $\pm i$ with threshold t_i for describing a condition, in which a positive index means $x_i \ge t_i$ and a negative index means $x_i < t_i$. Following this notation, a decision rule \mathbb{R} of i conditions is represented by $\mathbb{I}(\mathbb{R}) = {\mathbb{I}^{index}, \mathbb{I}^{threshold}} = {(x_1, \ldots, x_i), (t_1, \ldots, t_i)}$, where \mathbb{I}^{index} is a vector of i signed feature indices and $\mathbb{I}^{threshold}$ represents associated thresholds.

Since decision rules can be encoded by the decision paths from root nodes to leaf nodes, we use enriched random forests (ERFs) [ACL08] to extract decision rules from data. ERFs can be seen as RFs with a soft dimension reduction process, i.e. selecting more informative features in the input space.

Given numerous decision rules detected from ERFs, one can naturally treat them as feature interactions. However, they come with two drawbacks. First, each decision rule corresponds to only a small number of instances, which lacks statistical importance. Second, the decision rules are long and intricate, therefore do not generalize very well. To address these two drawbacks, we apply gRIT to process the detected decision rules. The main idea of gRIT is to "prune" those decision rules with respect to statistical meaning. After gRIT, statistically important feature interactions with better generalization ability can be discovered.

gRIT is summarized in Algorithm 1. A gRIT contains L intersection trees. Inside each intersection tree $\ell \in \{1, \ldots, L\}$, J example indices $\{i_1, \ldots, i_J\}$ are uniformly sampled from data and their corresponding decision rules are represented as $\{\mathbb{I}_{i_1} \ldots \mathbb{I}_{i_J}\}$ using the signed feature index notation mentioned earlier. Then, it performs J-fold intersections $\mathbb{I}_{i_1} \cap \cdots \cap \mathbb{I}_{i_J}$ to keep informative features and prune away noisy features. The main intuition behind this is that if a feature x_i of an interaction \mathbb{I} is sufficiently prevalent, it will survive the intersection with high probability.

The original RIT algorithm [SM14] is restricted to only binary categorical features, so it has a limited application range. Here, our generalized version can deal with both categorical and continuous features with threshold information.

2.2 Stability-aware interaction utilization

Statistical results should at least be reproducible relative to "reasonable" data and model perturbations [Yu13]. Towards this goal, we assess the stability associated with the identified feature interactions using

Algorithm 1: generalized Random Intersection Trees

Require: Rule set $\mathcal{R} = \{\mathbb{I}_i; \mathbb{I}_i = \{\mathbb{I}_i^{index}, \mathbb{I}_i^{threshold}\}\}_{i=1}^n$ **Ensure:** $\mathcal{T} = \bigcup_{\ell=1}^{L} \mathcal{T}_{\ell}$

1: for tree ℓ in $\{1, 2, ..., L\}$ do

Let ℓ be a tree of depth D. Let J be the total number of nodes in the tree, and index the nodes 2: such that every parent-child pair, larger indices are assigned to the child than the parent. For each node j = 1, ..., J, denote the parent of node j as pa(j), let i_j be a uniform sample index from the training data and corresponds to node j.

Set $T_1^{index} \leftarrow \mathbb{I}_{i_1}^{index}, T_1^{threshold} \leftarrow \mathbb{I}_{i_1}^{threshold}$ 3: for j in $\{2, ..., J\}$ do 4: 5:

$$: \qquad T_j^{index} \leftarrow \mathbb{I}_{i_j}^{index} \cap T_{pa(j)} \ , \ T_j^{investigat} \leftarrow$$

for feature index i in T_i^{index} do 6:

 $T_i^{threshold} \leftarrow$ 7:

 $T_{j}^{threshold} + \left\{ \max(t_i, t_i') \text{ if } i < 0, \text{ else } \min(t_i, t_i'). \ t_i \in \mathbb{I}_{i_j}^{threshold}, \ t_i' \in T_{pa(j)}^{threshold} \right\}$ end for

8: 9: end for

10:
$$\mathcal{T}_{\ell}^{index} \leftarrow \left\{ T_{j}^{index} : depth(j) = D \right\}, \ \mathcal{T}_{\ell}^{threshold} \leftarrow \left\{ T_{j}^{threshold} : depth(j) = D \right\}$$

11: $\mathcal{T}_{\ell} \leftarrow \left\{ \mathcal{T}_{\ell}^{index}, \mathcal{T}_{\ell}^{threshold} \right\}$

an extra outer bagging step. Specifically, we generate B bootstrap sample data $\mathcal{D}_{(b)}, b = \{1, \dots, B\}$, fit B enriched random forests ERF on each sample $\mathcal{D}_{(b)}$. Then, we apply gRIT to discover B sets of feature interactions $\mathcal{T}_{(b)}$ from these samples. We use the stability score defined as (2). It reflects the prevalence an interaction $\mathbb{I} \in \bigcup_{b=1}^{B} \mathcal{T}_{(b)}$ appears out of B bootstrap samples.

$$stability(\mathbb{I}) = \frac{1}{B} \sum_{b=1}^{B} \mathbb{I}\{\mathbb{I} \in \mathcal{T}_{(b)}\}.$$
(2)

Given the discovered feature interactions and their stability scores, we want to retain those interactions with top k scores or with scores bigger than a pre-specified threshold. They represent the most stable feature interactions, therefore lead to better generalization performance. This "interaction selection" procedure retains only a small number of stable interactions and discards a large number of unstable ones, which will ease the memory cost burden significantly.

To utilize the discovered feature interactions, we generate one new feature for each interaction, using Activated Linear Combination (ALC) including weighted sum and nonlinear activation functions, and denote these new features by \mathbf{r}_{new} . Take the feature interaction $\mathbb{I} = \{\mathbb{I}^{index}, \mathbb{I}^{threshold}\} =$ $\{(-2,3),(t_2,t_3)\}$ as an example, along with each feature's gini importance $\mathbf{w} = (w_2, w_3)$. $\mathbf{r}_{new}(\mathbf{w}, \mathbb{I})$ is calculated as:

$$\mathbf{r}_{new}(\mathbf{w}, \mathbb{I}) = \sigma(-w_2(x_2 - t_2) + w_3(x_3 - t_3)), \tag{3}$$

where $\sigma(x) = x \cdot \mathbb{I}[x \ge 0]$ is a non-linear activation function. We concatenate the new features with the current input features to form the input feature in the next layer. We will discuss the multi-layered structure in the next subsection.

2.3 Adaptive layer growth

It is widely acknowledged that depth plays a crucial role in the success of deep learning. We believe that deep structures bring two crucial benefits compared to shallow structures. First, deep structures proAlgorithm 2: High-order Interaction Deep Forest

Require: Training set S, the maximum number of layers K and number of bootstrap samples B. **Ensure:** *K*-layer deep forest model $\left\{ RF, \{S_r^{(k)}\}_{k=1}^K \right\}$. 1: Initialize representation features $\mathbf{r}^{(1)} \leftarrow \emptyset$, layer 0 data $S_r^{(0)} = S$, validation error $l_0 = 1$. 2: for k in $\{1, ..., K\}$ do $S_r^{(k)} \leftarrow concatenate(S_r^{(k-1)}, \mathbf{r}^{(k)})$ 3: Fit RF on $S_r^{(k)}$ 4: $\mathbf{w}^{(k)} \leftarrow \text{Gini importance of } RF$ 5: for b in $\{1, 2, ..., B\}$ do 6: Generate bootstrap samples $S_{(b)}$ from $S_r^{(k)}$ 7: Fit an Enriched- $RF(\mathbf{w}^{(k)})$ on $S_{(b)}$ 8: $\mathcal{R}_{(b)} \leftarrow \{ \mathbb{I}_{i_t} : \mathbf{x}_i \in S_{(b)}, \\ \text{falls in leaf node } i_t \text{ of tree } t \}$ 9: $\begin{aligned} & \text{fal} \\ \mathcal{T}_{(b)}^{(k+1)} \leftarrow \text{gRIT}(\mathcal{R}_{(b)}) \\ & \text{end for} \end{aligned}$ 10: 11: Initialize $\mathbf{r}^{(k+1)} \leftarrow \emptyset$. for $\mathbb{I}^{(k+1)} \in \bigcup_{b=1}^{B} \mathcal{T}^{(k+1)}_{(b)}$ do 12: 13: $stability(\mathbb{I}^{(k+1)}) \xleftarrow{1}{B} \sum_{b=1}^{B} \mathbb{1}\{\mathbb{I}^{(k+1)} \in \mathcal{T}^{(k+1)}_{(b)}\}$ 14: if $stability(\mathbb{I}^{(k+1)}) > 0.5$ then $\mathbf{r}^{(k+1)} \leftarrow \mathbf{r}^{(k+1)} \cup \mathbf{r}_{new} (\mathbf{w}^{(k)}, \mathbb{I}^{(k+1)})$ 15: 16: end if 17: end for 18: Compute the cross validation error l_k on $\mathbf{r}^{(k+1)}$. 19: if $l_k > l_{k-1}$ then 20: Terminate training and return current model. 21: end if 22: 23: end for 24: Fit RF on $S_r^{(K+1)}$

mote the *re-use* of features, with the ultimate goal is to form hierarchical distributed representations of data, which are beneficial to better generalization performance. Second, shallow layers tend to discover low-order feature interactions while deeper layers tend to discover high-order interactions, high-order interactions reflect on greater importance, thus have more impact on the learning task.

To get these two benefits, hiDF employs a multi-layered structure. Specifically, each layer is an ensemble of random forests, based on which we apply ERFs and gRIT to discover a rich family of feature interactions. Then we generate a new feature vector via the stability-aware interaction utilization mechanism described in the previous subsection. The newly generated feature vector is then concatenated with the input feature vector to form the new feature representations, which serve as the input to the next layer.

After expanding each new layer, hiDF estimates the performance of the current whole structure using a separate validation set. If there is no obvious performance boost, hiDF will terminate the training process. This adaptive layer-wise growing strategy can reduce the risk of overfitting, it also facilitates the model complexity of hiDF to be determined automatically. The overall framework of hiDF is illustrated in Figure 1 and summarized in Algorithm 2.



Figure 2: The quality of feature representations generated by hiDF across different layers on the synthetic data.

It is worth mentioning that the meaning of high-order in our hiDF is twofold: (1) hiDF can discover high-order interactions from the original features by gRIT, and (2) through our multi-layered approach, hiDF can detect the interactions between original features and the newly generated features, even among only newly generated features, causing the order of interactions to increase layer by layer.

3 Experiments

In this section, we conduct several experiments with hiDF on both synthetic data and several widely used benchmark datasets. We first demonstrate that hiDF can obtain informative feature representations using a hierarchical processing structure. We then evaluate the effectiveness of hiDF across several real datasets. The results show that hiDF can achieve the highest classification accuracy on all of these benchmark datasets and it exhibits better performance than deep forest consistently. Meanwhile, hiDF uses significantly less memory space and test time and has a moderate amount of training time compared to deep forest.

3.1 Synthetic data

First, we use a synthetic dataset to demonstrate that hiDF can learn informative feature representations through its multi-layered feature interaction utilization mechanism. Our synthetic dataset is a binary classification task generated by Gaussian data on \mathbb{R}^2 with a spherical decision boundary as illustrated in Figure 2(a). It contains 10,000 examples (80% for training and 20% for testing). We also introduce additional 500 uniform random noise features in the dataset. These uninformative random features make the classification task more complex, thus allowing us to assess the quality of representation features generated by hiDF.

The output feature representations from layer 6 and layer 12 of hiDF are illustrated via t-SNE in Figure 2(b). Clearly, hiDF can get informative feature representations from the original input space, and as layers go deeper, data points become linearly separable while preserving some structure information inside one class. This demonstrates that hiDF is equipped with similar representation learning abilities as deep neural networks, which uses hidden layers to distort the input in a non-linear way so that classes become easily separable by the last layer [LBH15]. We believe that leveraging feature interactions on multiple levels facilitates hiDF to learn hierarchical distributed representations of data.

We also examine the importance of new features generated across varying feature interaction orders via the marginal importance metric [Bas+18]. Figure 3 depicts the result by a boxplot. We can see that

the leftmost bar shows the input features' marginal importance, with all but two informative features having nearly zero importance values. Figure 3 also illustrates that new features generated from higher-order interactions tend to have greater feature importance, thus have more impact on the learning task. Since higher-order feature interactions are usually discovered in deeper layers, it is reasonable to believe that hiDF can benefit from its hierarchical structure.



Figure 3: The importance of features with varying interaction orders

3.2 Real dataset

We select several widely used benchmark datasets of binary and multi-class classification tasks from the UCI Machine Learning Repository [DG17]. The datasets vary in size: from only 452 examples up to 581,012 examples. They come from different areas, including health and medical research (Arrhythmia and Diabetes Readmission), remote sensing (CoverType and Satimage), financial and business (Bank marketing and Default of Credit card clients), social study (Adult), etc. We select these datasets mainly for two reasons. First, they are representative tabular data of categorical and mixed variable types and tend to contain multiple underlying explanatory factors that can be utilized by our hiDF to discover feature interactions. Second, these datasets have more input features than output labels, which could verify that our hiDF can generate more informative and diverse feature representations than gcForest.

In addition to these 9 UCI datasets, we release a new dataset called Congestive Heart Failure Dataset. This dataset includes 9809 patients with a diagnosis of "Congestive Heart Failure" from the MIMIC-III (Medical Information Mart for Intensive Care) platform [Joh+16]. MIMIC is a relational and public database containing tables of data relating to patients who stayed within the intensive care units at Beth Israel Deaconess Medical Center. In this Congestive Heart Failure Dataset, 88 numerical and categorical variables including patients' demographics, vitals' measurements, laboratory results, prescriptions, imaging results, etc are collected to predict the in-hospital mortality of patients. Table 1 shows the basic statistics of these datasets.

In these experiments, we used the default hyper-parameters defined in our hiDF model and did not fine-tune them. The default number of decision trees in a random forest is 500; the number of maximum layers in hiDF is 10; the early stopping patience is 2; the number of bootstrap samples B = 10; the default number of intersection trees is 20 with each tree's depth D = 5 and $n_{child} = 2$. For reproducibility, we make the source code and some datasets used in this work publicly available at https://git.io/DM487

Datasets	# of examples	# of features	# of classes
CoverType	581,012	54	7
Adult	48,842	14	2
Bank	41,188	20	2
Credit Card	30,000	23	2
Arrhythmia	452	279	16
YearPredictionMSD	50,000	90	2
Diabetes Readmission	100,000	55	2
Satimage	6,435	36	6
Crowdsourced Mapping	10,845	28	6
Congestive Heart Failure	9809	88	2

Table 1: Statistics of the datasets in terms of number of examples, number of features and number of classes

3.3 Performance comparison

We adopt classification accuracy as our evaluation measure. For comparison, we choose several widely used tree ensemble algorithms including GBDT [Fri01], Random Forests [Bre01], gcForest [ZF19], gcForest_{CS} [Pan+20], XGBoost [CG16], and SVM.

The gcForest method is proposed by Zhou and Feng [ZF19] to show that deep learning framework also can be realized by non-differentiable modules. It achieves state-of-the-art performance on many categorical and mixed modeling tasks. Recently, Pang et al. [Pan+20] proposed gcForest_{CS} to address the high computational cost problem of gcForest by introducing a confidence screening mechanism. We set the hyper-parameters of gcForest and gcForest_{CS} as default.

XGBoost is a scalable end-to-end tree boosting system proposed by Chen and Guestrin [CG16]. We set the number of boosting rounds equals to 500, and the maximum tree depth for base learners equals to 6. As for other hyper-parameters, we set them as the default values. We apply the same configurations of hyper-parameters to GBDT.

Table 2 reports the accuracy comparison result among these methods. The win/tie/lose counts of hiDF over the other methods are indicated by the last rows. We can see that hiDF achieves the highest accuracy on all of these datasets. We can also clearly see that hiDF exhibits consistently better performance than gcForest and gcForest_{CS}, which suggests hiDF's representation learning ability based on high-order feature interactions can indeed boost the performance of the original deep forest.

Datasets	hiDF	gcForest	gcForest _{CS}	Random Forests	XGBoost	GBDT	SVM
CoverType	$\textbf{97.62} \pm \textbf{0.08}$	96.23 ± 0.10	96.09 ± 0.07	95.63 ± 0.07	94.48 ± 0.05	$96.94\pm0.11\bullet$	71.47 ± 0.12
Adult	$\textbf{86.90} \pm \textbf{0.05}$	86.17 ± 0.06	86.17 ± 0.09	85.15 ± 0.08	$86.40\pm0.00\bullet$	86.12 ± 0.10	79.86 ± 0.00
Bank	$\textbf{89.96} \pm \textbf{0.21}$	89.89 ± 0.26	$89.95\pm0.25\bullet$	89.33 ± 0.17	89.13 ± 0.33	88.60 ± 0.25	89.75 ± 0.21
Credit Card	$\textbf{82.00} \pm \textbf{0.23}$	$81.74\pm0.38~\bullet$	81.73 ± 0.20	81.71 ± 0.33	80.61 ± 0.36	81.24 ± 0.27	77.88 ± 0.01
Arrhythmia	$\textbf{78.24} \pm \textbf{1.62}$	$76.26 \pm 1.49 \bullet$	74.51 ± 1.62	74.51 ± 2.13	75.16 ± 1.32	74.29 ± 3.15	60.66 ± 1.62
YearPredictionMSD	$\textbf{75.89} \pm \textbf{0.46}$	75.49 ± 0.40	75.61 ± 0.30	73.41 ± 0.30	74.69 ± 0.21	$75.85\pm0.20\bullet$	68.62 ± 0.62
Diabetes	$\textbf{62.43} \pm \textbf{0.23}$	62.26 ± 0.36	$62.37\pm0.33\bullet$	62.11 ± 0.24	61.00 ± 0.17	61.04 ± 0.45	54.99 ± 0.02
Satimage	$\textbf{91.75} \pm \textbf{0.06}$	$91.63\pm0.12 \bullet$	91.58 ± 0.11	91.21 ± 0.05	90.65 ± 0.00	90.44 ± 0.09	88.60 ± 0.00
Crowdsourced Mapping	$\textbf{65.07} \pm \textbf{0.88}$	64.93 ± 0.95	$65.03\pm0.56\bullet$	63.47 ± 0.72	62.00 ± 0.00	62.53 ± 0.50	55.67 ± 0.00
Congestive Heart Failure	$\textbf{90.16} \pm \textbf{0.53}$	88.61 ± 0.59	88.58 ± 0.27	87.90 ± 0.01	$89.34\pm0.01\bullet$	88.83 ± 0.01	85.09 ± 0.00
win/tie/lose	_	10/0/0	10/0/0	10/0/0	10/0/0	10/0/0	10/0/0

Table 2: Comparison of test accuracy of each methods on these 10 datasets. The best accuracy is highlighted in **bold** type. • indicates the second-best.

3.4 Computational complexity comparison with gcForest

To show that hiDF can achieve state-of-the-art predictive result while still being computationally efficient, we compare hiDF with gcForest and gcForest_{CS} on two datasets: Adult and CoverType. We use a hardware of 16×3.70 GHz CPUs with 128 GB memory. All of these methods grow 20 layers. Table 3 summarizes the results. We can see that hiDF uses significantly less memory and test time than gcForest and gcForest_{CS} (by about one order of magnitude).

Table 3:	Comparison	results	of	training	time,	test	time(in	CPU	seconds)	and	memory	usage	(in
megabyte	s) with two da	atasets.											

Datasets	Method	Training time	Test time	Memory
Adult	hiDF	3208.7	13.1	2793.2
	gcForest	2763.8	541.3	29243.9
	gcForest _{CS}	1063.1	290.2	23272.7
Covertype	hiDF	4505.9	90.9	5562.5
	gcForest	2861.3	1519.4	61981.2
	gcForest _{CS}	1582.3	730.1	57658.8

The main reason behind the efficiency of hiDF is: since the feature representations generated by gc-Forest and gcForest_{CS} are based on RF predictions. As a result, they have to save the trained multilayer RFs to generate feature representations for test instances, which is essentially saving tens of thousands of unpruned decision rules, thus requiring a large memory cost. Besides, gcForest and gcForest_{CS} consumes a lot of time to do layer-wise prediction for test instances. On the other hand, hiDF does not rely on prediction-based representations. Therefore, only a small number (tens) of feature interactions have to be saved in each layer, and test instances can generate feature representations easily based solely on these interactions.

Note that in our current implementation, the *B* times bootstrap are done sequentially, so the gRIT processing takes the majority amount of time. The training time of our hiDF could be significantly less if using parallel processing.

4 Conclusion

In this paper, we focus on the problem of limited representation learning ability and large computational complexity caused by prediction-based representations in traditional Deep Forest algorithms. To address these shortcomings, we propose a novel method called hiDF to attain interaction-based hierarchical distributed representations through in-model feature transformation. Experiments show the potential of hierarchical distributed representations for reducing the complexity of classification tasks by enhancing the separability of data, and further demonstrate the performance improvement and computational effectiveness of our method.

Acknowledgment

This work is supported by the National Key Research and Development Program of China No.2020AAA0109400. The authors would like to thank Doctor Yao Wang from Sir Run-Run Shaw Hospital, Zhejiang University for generously providing the Congestive Heart Failure Dataset, as well as the anonymous reviewers and Yi-Xiao He from Nanjing University and Kewen Zheng from National University of Singapore for constructive suggestions.

References

- [ACL08] Dhammika Amaratunga, Javier Cabrera, and Yung-Seop Lee. "Enriched random forests". In: *Bioinformatics* 24.18 (2008), pp. 2010–2014 (Cited on page 3).
- [Bas+18] Sumanta Basu, Karl Kumbier, James B. Brown, and Bin Yu. "Iterative random forests to discover predictive and stable high-order interactions". In: *Proceedings of the National Academy of Sciences* 115.8 (2018), pp. 1943–1948 (Cited on page 6).
- [Ben09] Yoshua Bengio. "Learning Deep Architectures for AI". In: *Foundations and Trends*® *in Machine Learning* 2.1 (2009), pp. 1–127 (Cited on page 1).
- [BCV13] Yoshua Bengio, Aaron Courville, and Pascal Vincent. "Representation learning: A review and new perspectives". In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 35.8 (2013), pp. 1798–1828 (Cited on page 2).
- [Ber+11] Thierry Bertin-Mahieux, Daniel P.W. Ellis, Brian Whitman, and Paul Lamere. "The Million Song Dataset". In: *Proceedings of the 12th International Conference on Music Information Retrieval*. 2011.
- [Bre01] Leo Breiman. "Random forests". In: *Machine learning* 45.1 (2001), pp. 5–32 (Cited on page 8).
- [Bre+84] Leo Breiman, Jerome Friedman, Charles J Stone, and Richard A Olshen. *Classification and regression trees*. CRC press, 1984.
- [CG16] Tianqi Chen and Carlos Guestrin. "Xgboost: A scalable tree boosting system". In: Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining. 2016, pp. 785–794 (Cited on page 8).
- [CMS14] Corinna Cortes, Mehryar Mohri, and Umar Syed. "Deep Boosting". In: *Proceedings of the 31th International Conference on Machine Learning*. Vol. 32. 2014, pp. 1179–1187.
- [DMS15] Giulia DeSalvo, Mehryar Mohri, and Umar Syed. "Learning with deep cascades". In: *Algorithmic Learning Theory*. 2015, pp. 254–269.
- [Dev+18] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. "BERT: Pre-training of deep bidirectional transformers for language understanding". In: *CoRR* abs/1810.04805 (2018) (Cited on page 1).
- [DG17] Dheeru Dua and Casey Graff. UCI Machine Learning Repository. 2017. URL: http://archive.ics.uci.edu/ml (Cited on page 7).
- [Fri01] Jerome H Friedman. "Greedy function approximation: A gradient boosting machine". In: *Annals of statistics* (2001), pp. 1189–1232 (Cited on page 8).
- [Gho+19] Swarnendu Ghosh, Nibaran Das, Ishita Das, and Ujjwal Maulik. "Understanding Deep Learning Techniques for Image Segmentation". In: *CoRR* abs/1907.06119 (2019).
- [Guo+18] Yang Guo, Shuhui Liu, Zhanhuai Li, and Xuequn Shang. "BCDForest: A boosting cascade deep forest model towards the classification of cancer subtypes based on gene expression data". In: *BMC Bioinformatics* 19.5 (2018), p. 118 (Cited on page 1).
- [Guv+97] H Altay Guvenir, Burak Acar, Gulsen Demiroz, and Ayhan Cekin. "A supervised machine learning algorithm for arrhythmia analysis". In: *Computers in Cardiology*. 1997, pp. 433– 436.
- [He+15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. "Deep Residual Learning for Image Recognition". In: CoRR abs/1512.03385 (2015) (Cited on page 1).

- [Joh+16] Alistair EW Johnson, Tom J Pollard, Lu Shen, H Lehman Li-Wei, Mengling Feng, Mohammad Ghassemi, Benjamin Moody, Peter Szolovits, Leo Anthony Celi, and Roger G Mark.
 "MIMIC-III, a freely accessible critical care database". In: *Scientific data* 3.1 (2016), pp. 1–9 (Cited on page 7).
- [JI16] Brian A Johnson and Kotaro Iizuka. "Integrating OpenStreetMap crowdsourced data and Landsat time-series imagery for rapid land use/land cover (LULC) mapping: Case study of the Laguna de Bay area of the Philippines". In: Applied Geography 67 (2016), pp. 140–149.
- [Kri+20] Samuel Kriman, Stanislav Beliaev, Boris Ginsburg, Jocelyn Huang, Oleksii Kuchaiev, Vitaly Lavrukhin, Ryan Leary, Jason Li, and Yang Zhang. "Quartznet: Deep Automatic Speech Recognition with 1D Time-Channel Separable Convolutions". In: *IEEE International Conference on Acoustics, Speech and Signal Processing*. 2020, pp. 6124–6128 (Cited on page 1).
- [LBH15] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. "Deep learning". In: *nature* 521.7553 (2015), pp. 436–444 (Cited on page 6).
- [LYZ19] Shen-Huan Lyu, Liang Yang, and Zhi-Hua Zhou. "A Refined Margin Distribution Analysis for Forest Representation Learning". In: Advances in Neural Information Processing Systems 32. 2019, pp. 5530–5540 (Cited on page 1).
- [MH08] Laurens van der Maaten and Geoffrey Hinton. "Visualizing Data using t-SNE". In: *Journal* of Machine Learning Research 9.86 (2008), pp. 2579–2605.
- [MCR14] Sérgio Moro, Paulo Cortez, and Paulo Rita. "A data-driven approach to predict the success of bank telemarketing". In: *Decision Support Systems* 62 (2014), pp. 22–31.
- [Pan+20] Ming Pang, Kai-Ming Ting, Peng Zhao, and Zhi-Hua Zhou. "Improving Deep Forest by Confidence Screening". In: *IEEE Transactions on Knowledge and Data Engineering* (2020). DOI: 10.1109/TKDE.2020.3038799 (Cited on page 8).
- [Ped+11] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. "Scikit-learn: Machine Learning in Python". In: *Journal of Machine Learning Research* 12 (2011), pp. 2825–2830.
- [SM14] Rajen Dinesh Shah and Nicolai Meinshausen. "Random Intersection Trees". In: *Journal of Machine Learning Research* 15.20 (2014), pp. 629–654 (Cited on page 3).
- [SWA08] Alexander Statnikov, Lily Wang, and Constantin F Aliferis. "A comprehensive comparison of random forests and support vector machines for microarray-based cancer classification". In: *BMC bioinformatics* 9.1 (2008), pp. 1–10.
- [Str+14] Beata Strack, Jonathan P DeShazo, Chris Gennings, Juan L Olmo, Sebastian Ventura, Krzysztof J Cios, and John N Clore. "Impact of HbA1c measurement on hospital readmission rates: Analysis of 70,000 clinical database patient records". In: *BioMed research international* 2014 (2014).
- [Sun+20] Liang Sun, Zhanhao Mo, Fuhua Yan, Liming Xia, Fei Shan, Zhongxiang Ding, Wei Shao, Feng Shi, Huan Yuan, Huiting Jiang, Dijia Wu, Ying Wei, Yaozong Gao, Wanchun Gao, He Sui, Daoqiang Zhang, and Dinggang Shen. "Adaptive feature selection guided deep forest for COVID-19 classification with chest CT". In: *IEEE Journal of Biomedical and Health Informatics* 24.10 (2020), pp. 2798–2805 (Cited on page 1).
- [Tsa+18] Michael Tsang, Hanpeng Liu, Sanjay Purushotham, Pavankumar Murali, and Yan Liu. "Neural interaction transparency (nit): Disentangling learned interactions for improved interpretability". In: Advances in Neural Information Processing Systems 31. 2018, pp. 5804– 5813.

- [Yan+18] Feng Yang, Qizhi Xu, Bo Li, and Yan Ji. "Ship detection from thermal remote sensing imagery through region-based deep forest". In: *IEEE Geoscience and Remote Sensing Letters* 15.3 (2018), pp. 449–453 (Cited on page 1).
- [YL09] I-Cheng Yeh and Che-hui Lien. "The comparisons of data mining techniques for the predictive accuracy of probability of default of credit card clients". In: *Expert Systems with Applications* 36.2 (2009), pp. 2473–2480.
- [Yu13] Bin Yu. "Stability". In: *Bernoulli* 19.4 (2013), pp. 1484–1500 (Cited on page 3).
- [Zha+19] Ya-Lin Zhang, Jun Zhou, Wenhao Zheng, Ji Feng, Longfei Li, Ziqi Liu, Ming Li, Zhiqiang Zhang, Chaochao Chen, Xiaolong Li, et al. "Distributed deep forest and its application to automatic detection of cash-out fraud". In: ACM Transactions on Intelligent Systems and Technology 10.5 (2019), pp. 1–19 (Cited on page 1).
- [ZF19] Zhi-Hua Zhou and Ji Feng. "Deep forest". In: *National Science Review* 6.1 (2019), pp. 74–86 (Cited on pages 1, 2, 8).