
TestRank: Bringing Order into Unlabeled Test Instances for Deep Learning Tasks

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 Deep learning (DL) systems are notoriously difficult to test and debug due to
2 the lack of correctness proof and the huge test input space to cover. Given the
3 ubiquitous unlabeled test data and high labeling cost, in this paper, we propose
4 a novel test prioritization technique, namely *TestRank*, which aims at revealing
5 more model failures with less labeling effort. TestRank brings order into the
6 unlabeled test data according to their likelihood of being a failure, i.e., their *failure-*
7 *revealing capabilities*. Different from existing solutions, TestRank leverages both
8 *intrinsic* and *contextual* attributes of the unlabeled test data when prioritizing
9 them. To be specific, we first build a similarity graph on both unlabeled test
10 samples and labeled samples (e.g., training or previously labeled test samples).
11 Then, we conduct graph-based semi-supervised learning to extract contextual
12 features from the correctness of similar labeled samples. For a particular test
13 instance, the contextual features extracted with the graph neural network and the
14 intrinsic features obtained with the DL model itself are combined to predict its
15 failure-revealing capability. Finally, TestRank prioritizes unlabeled test inputs in
16 descending order of the above probability value. We evaluate TestRank on three
17 popular image classification datasets, and results show that TestRank significantly
18 outperforms existing test prioritization techniques.

19 1 Introduction

20 Deep learning (DL) systems are prone to errors due to many factors, such as the biased train-
21 ing/validation dataset, the limitations of the model architecture, and the constraints on training cost.
22 Needless to say, it is essential to conduct high-quality testing before DL models are deployed in the
23 field; otherwise, the behaviors of DL models can be unpredictable and result in severe accidents after
24 deployment. However, the cost of building test oracles (i.e., the ground-truth output) by manually
25 labeling a massive set of test instances is prohibitive, especially for tasks requiring experts for accurate
26 labeling, such as medical images and malware executables.

27 To tackle the above problem, various test prioritization techniques [6, 2, 20] are proposed to identify
28 ‘high-quality’ test instances from a large amount of unlabeled data, which facilitates revealing more
29 failures (e.g., misclassification) of the DL model with reasonable labeling effort. These methods try to
30 derive the failure-revealing capability of a test instance with its *intrinsic attributes* extracted from the
31 responses of the model under test (e.g., the softmax-based probabilities given by the target DL model
32 to this specific input). DeepGini [6] feeds the unlabeled data to the target DL model and calculates
33 confidence-related scores based on the model’s output probabilities to rank the unlabeled test cases.
34 Test cases with nearly equal probabilities on all output classes are regarded as less confident ones and
35 are likely to reveal model failures. Similarly, Byun *et. al.* use the uncertainty score obtained from
36 MC-Dropout for test input prioritization [2]. Multiple-boundary clustering and prioritization (MCP)

37 [20] considers both the output probabilities and the balance among each classification boundary
 38 when selecting test cases. All existing works try to identify instances near the decision boundary
 39 and prioritize them. However, we argue that near-boundary instances are not necessarily failures,
 40 especially for well-trained classifiers with high accuracy. Also, as failures can be far from the decision
 41 boundary, existing methods could fail to reveal these remote failures.

42 To estimate a test instance’s capability in revealing failures, in addition to the intrinsic attributes
 43 mentioned above, there is another type of information that provides extra insight into the target
 44 model’s behavior: the known classification correctness of labeled samples (i.e., training samples
 45 and previously tested samples) and their relationship to the unlabeled instance. Such data is already
 46 known and it provides *contextual information* that reflects the corresponding inference behaviors of
 47 the target model for a set of similar instances.

48 In this work, we present a novel test prioritization technique, namely *TestRank*, for DL classifiers.
 49 TestRank exploits both intrinsic and contextual attributes of test instances to evaluate their failure-
 50 revealing capabilities. Based on the intuition that similar inputs are usually associated with the
 51 same classification results, we propose to use graph neural networks (GNNs) [14] to summarize the
 52 neighboring classification correctness for each unlabeled instance into contextual attributes. GNNs
 53 have been well-studied and valued for their relational inductive bias for extracting graph information.
 54 Our method, TestRank, constructs a similarity graph on both unlabeled and labeled instances and
 55 apply the semi-supervised GNN learning to extract the contextual attributes. After that, we aggregate
 56 intrinsic (such attributes are extracted from the input samples without considering their neighbors)
 57 and contextual attributes with a neural-network-based binary classifier for test prioritization.

58 The contributions of our work are as follows:

- 59 • To the best of our knowledge, *TestRank* is the first work that takes the contextual information
 60 from the target DL model into consideration for test input prioritization.
- 61 • We propose to construct a similarity graph on both labeled and unlabeled samples, and
 62 train a graph neural network to extract useful contextual attributes from the contextual
 63 information for these unlabeled instances. We also present approximation techniques to
 64 reduce its computational complexity with minor impact on the performance of *TestRank*.
- 65 • We propose a simple yet effective neural network that combines the intrinsic attributes
 66 and contextual attributes of unlabeled test instances for their failure-revealing capability
 67 estimation.

68 We empirically evaluate *TestRank* on three popular image classification benchmarks: CIFAR-10,
 69 SVHN, and STL10. The results show that our method outperforms the state-of-the-art methods by a
 70 considerable margin.

71 2 Test Prioritization and Prior Works

72 Let us use $f : \mathcal{X} \rightarrow \mathcal{Y}$ to represent the given target DL model, where \mathcal{X} and \mathcal{Y} are the input and
 73 output space, respectively. For effective testing, the debugging center must do test prioritization
 74 from a large unlabeled dataset, i.e., select a certain number of test instances from the unlabeled test
 75 instance pool that can reveal model failures as many as possible. Later, these failures can be fed back
 76 to the training center for repairing or failure analysis. We define the model failures as follows:

77 **Definition 1. DL Model Failure.** A failure of the DL model can be uncovered by the test instance \mathbf{x}
 78 if the predicted label $f(\mathbf{x})$ is inconsistent with its ground truth label $y_{\mathbf{x}}$, namely $f(\mathbf{x}) \neq y_{\mathbf{x}}$.

79 Formally, the debugging center selects and labels b test cases \mathbf{X}_S ($|\mathbf{X}_S| = b$) from the unlabeled test
 80 instance pool \mathbf{X}_U . The objective of test prioritization is to maximize the detected failures:

$$\max |\{\mathbf{x} | f(\mathbf{x}) \neq y_{\mathbf{x}}\}|, \text{ where } \mathbf{x} \in \mathbf{X}_S \text{ and } |\mathbf{X}_S| = b. \quad (1)$$

81 Different solutions are proposed to quantify the failure-revealing capability of unlabeled instances.
 82 DeepGini [6] proposes to evaluate a single test instance via the DL model’s final statistical output:

$$f(t) = 1 - \sum_{i=1}^N p_{t,i}^2, \quad (2)$$

83 where $p_{t,i}$ is the predicted probability that the test case t belongs to the class i . Given the sum of
 84 $p_{t,i}$ is 1, impurity function $f(t)$ is maximal when all $p_{t,i}$ values are equal. This method tends to

85 select data with high impurity value. Instead of evaluating the overall likelihood of failure for all
 86 classes, Multiple-Boundary Clustering and Prioritization (MCP) proposes to evaluate it for each pair
 87 of classes individually [20]. In this way, test instances can be evenly selected for each class pair and
 88 the failure cases are investigated at the finer granularity. Besides these metrics, Byun *et. al.* also
 89 propose to measure the likelihood of incorrect prediction by the uncertainty of the model’s output [2],
 90 which reflects the degree to which a model is uncertain about its prediction. In practice, evaluating
 91 uncertainty requires the task DL model to be a Bayesian Neural Network [19, 16] or containing a
 92 dropout layer for approximation [8].

93 Besides examining the DL model’s final outputs, Kim *et.al.* proposes two surprise adequacy (SA)
 94 criteria that make use of the target DL’s internal outputs (e.g., the activation traces) [12]. They
 95 are Likelihood-based Surprise Adequacy Coverage (LSA) and Distance-based Surprise Adequacy
 96 Coverage (DSA). LSA and DSA measure the likelihood or distance of an unlabeled instance to the
 97 training instances, respectively. Test samples with higher SA values are preferred in testing.

98 To sum up, all existing methods use the target model’s outputs to one input, i.e., its intrinsic attributes,
 99 for its failure-revealing capability estimation. In contrast, we make use of both intrinsic and contextual
 100 attributes of an instance for better estimation (see later sections for details).

101 3 TestRank

102 3.1 Motivation

103 The failure-revealing capability of an unlabeled test input is closely related to its attributes for the DL model under test. In this work, we distinguish two kinds of attributes for an unlabeled instance: the *intrinsic attributes* and the *contextual attributes*.

109 We define the *intrinsic attributes* of an input as the output responses assigned by the target DL model to this specific input. It could be, for example, the predictive output distribution of the input from the target DL model, reflecting the *sentiment* derived from the computation performed by the target model [2]. This kind of attributes is adopted by existing test input prioritization approaches [6, 20, 2]. Note that we define such attributes as ‘intrinsic’ because they are extracted from inputs without considering their context, *i.e.*, the classification correctness of its similar instances.

122 In contrast with the intrinsic attributes, the *contextual attributes* provide a deeper insight into the target model for the unlabeled samples: the contextual attributes for an unlabeled sample summarize the classification correctness of similar and labeled samples. For a particular test instance, such contextual attributes are useful and complementary to the intrinsic attributes.

126 An illustrative example is shown in Figure 1, wherein we visualize the behavior of a two-class classifier on the unlabeled test data and historically labeled data distribution. The blue region includes the instances that are near the decision boundary. Intuitively, the classifier is uncertain about the data when data is near the decision boundary and is likely to misclassify it. Existing works [20, 6, 2] propose various indicators (*e.g.*, confidence/uncertainty/surprise scores) to help identify the near-boundary instances. However, the near-boundary instances are not necessarily failures, and some of them can be correctly classified by a well-trained classifier. What is worse, such testing approaches fail to capture the failures lying far from the decision boundary (*i.e.*, remote failures, shown in the red region in Figure 1), because DL models usually output high confidence (or low uncertainty) for these inputs. These failures may be caused by limited model capacity, insufficient training data, etc.

136 Our key insight is that we can use the contextual information (*e.g.* the classification correctness of similar labeled samples) to help locate both near-boundary and remote failures. The usefulness

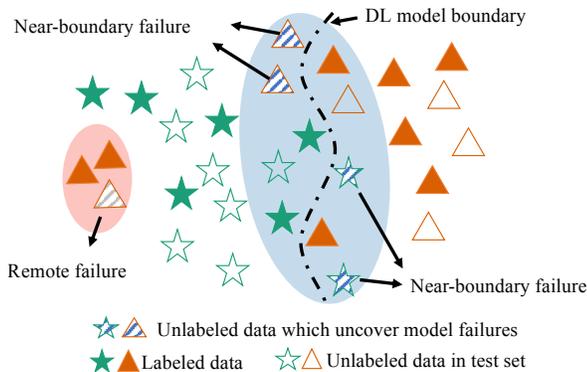


Figure 1: A motivational example.

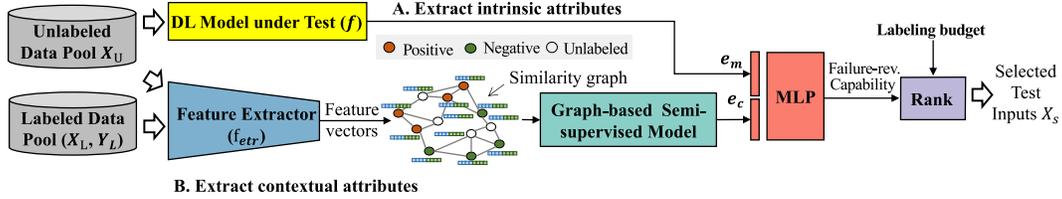


Figure 2: The overview of TestRank.

138 of the contextual information is due to the local continuity property [1], which means that inputs
 139 close in the feature space share similar prediction behavior, e.g., classification results from the target
 140 model. As shown in Figure 1, there are some already labeled data, whose classification correctness
 141 are already known, surrounding the unlabeled data. If an unlabeled instance is close to already
 142 falsely classified data, under the local continuity property, it is likely that this instance is also a model
 143 failure. This property motivates us to extract the contextual attributes for an unlabeled instance from
 144 its neighboring labeled data. By combining the extracted contextual attributes with the intrinsic
 145 attributes, we expect to achieve better failure-revealing capability estimation.

146 3.2 Overview

147 Figure 2 shows the overview of TestRank, which consists of two attribute extraction paths for the
 148 final failure-revealing capability estimation:

- 149 1. **Path A: intrinsic attributes extraction.** Given a pool of unlabeled inputs \mathbf{X}_U , we use the
 150 target DL model f to extract the intrinsic attributes \mathbf{e}_m for each input. More precisely, we
 151 collect the output logits (i.e., vectors before the *softmax* layer) from the DL model as \mathbf{e}_m .
- 152 2. **Path B: contextual attributes extraction.** First, we use a feature extractor to map the
 153 original data space into a more compact feature space with good local continuity property.
 154 Then, we create a similarity graph (i.e., k -Nearest Neighbor Graph) based on the obtained
 155 feature vectors and their corresponding classification correctness, if any, from both unlabeled
 156 data pool \mathbf{X}_U and labeled data pool \mathbf{X}_L (e.g. training set and previously labeled test samples).
 157 Last but not least, the graph-based representation learning technique is applied to extract the
 158 contextual attributes \mathbf{e}_c for each unlabeled instance. The details are elaborated in Sec. 3.3.

159 After attributes \mathbf{e}_m and \mathbf{e}_c are extracted, we combine them together via a Multi-Layer Perceptron
 160 (MLP) (see Sec. 3.4 for details). The MLP is responsible for predicting the failure-revealing ability
 161 for unlabeled test instances. At last, these instances are ranked according to their failure-revealing
 162 capability, and the top ones are selected under the given labeling budget.

163 As intrinsic attributes extraction is straightforward, we discuss the path *B* and how to combine path *A*
 164 and *B* in detail in the following subsections.

165 3.3 Contextual Attributes Extraction

166 We represent the contextual information from the DL model as a set of labeled inputs \mathbf{X}_L and the
 167 corresponding classification correctness $\mathbf{Y}_L \in \{0, 1\}$, where correctly classified inputs are labeled as
 168 0 and misclassified ones are labeled as 1. Given the contextual information, our goal is to extract
 169 the contextual attributes for each unlabeled test instance $\mathbf{x} \in \mathbf{X}_U$. However, extracting contextual
 170 attributes from labeled and unlabeled data is a non-trivial task because of the following reasons.

171 First, it is well-known that the real data, especially image data, usually locates in high-dimensional
 172 space, wherein the underlying data distribution will live on complex and non-linear manifold em-
 173 bedded within the high-dimensional space [1]. Therefore, constructing the relationships between
 174 different instances is difficult. To address the challenge, we adopt the representation learning pro-
 175 cess [9, 21, 3, 11], which map the raw data into a compact feature space with better local continuity
 176 property, such that inputs close in the feature space share similar classification results. Thus, in the
 177 feature space, the proximity between inputs can be measured with simple distance metrics (e.g., L_2 ,
 178 cosine).

Algorithm 1: GNN-based Contextual Attributes Extraction

Input: Input samples $\mathbf{X}_U \cup \mathbf{X}_L$, Correctness of labeled samples \mathbf{Y}_L , Number of neighbors k , Feature extractor f_{etr} , number of GNN layers M , number of training epochs.
Output: Contextual attributes \mathbf{E}_c for \mathbf{X}_U

- 1 $\bar{\mathbf{X}} = f_{etr}(\mathbf{X}_U \cup \mathbf{X}_L)$ # Extract compact representation;
- 2 $\mathbf{Edge} = knn_graph(\bar{\mathbf{X}}, k)$. # KNN Graph construction;
- 3 $\tilde{\mathbf{A}} = \mathbf{Edge} + \mathbf{I}_N$, $\tilde{\mathbf{D}} = \sum_j \tilde{\mathbf{A}}_{i,j}$;
- 4 $\mathbf{H}^0 = \bar{\mathbf{X}}$;
- 5 **for** number of training epochs **do**
- 6 | **for** $l = 0, 1, \dots, M - 1$ **do**
- 7 | | $\mathbf{H}^{l+1} = \sigma(\tilde{\mathbf{D}}^{-\frac{1}{2}} \tilde{\mathbf{A}} \tilde{\mathbf{D}}^{-\frac{1}{2}} \mathbf{H}^l \Theta^l)$,
- 8 | **end**
- 9 | $loss = \text{CrossEntropyLoss}(\mathbf{H}^M, \mathbf{Y}_L)$;
- 10 | Update Θ ;
- 11 **end**
- 12 $\mathbf{E}_c = \mathbf{H}^{M-1}[\text{index of } \mathbf{X}_u]$ # extract the representation from the $M - 1_{th}$ GNN layer;
- 13 **return** \mathbf{E}_c ;

179 Second, manually designing protocols to summarize the neighboring classification results is subject
180 to the imperfection of local continuity. Namely, the result is easily affected by the noisy data in the
181 neighboring space. To solve this challenge, we construct a similarity graph based on the labeled and
182 unlabeled data, and then apply the more powerful graph learning technique – graph neural networks
183 (GNN) – for contextual attribute extraction.

184 The GNN empowers graph embedding learning, as it employs a learnable aggregation and transform
185 procedure [14], which exploits the relational inductive-bias that exhibits in the graph structure. It
186 generates a embedding/representation that summarizes the “contextual information” for each input
187 sample, making it easier to separate the correct and misclassified inputs. The contextual attributes
188 extracted by the graph neural network can then be combined with the intrinsic attributes to conduct
189 the better failure-revealing capability estimation (See Sec. 3.4). The contextual attributes extraction
190 process is formally depicted by Algorithm 1.

191 **Feature Vector Representation (Line 1).** As the target model is to be tested, its feature extraction
192 quality is not guaranteed. Out of this concern, and to make full use of the labeled and unlabeled data,
193 we choose to use a out-of-shelf unsupervised model for feature space construction.

194 Among the unsupervised learning techniques, the BYOL [9] explicitly introduces local continuity
195 constraint into the learned feature space and shows good results on various downstream tasks.
196 Therefore, we train a BYOL model f_{etr} to extract the features from the raw input images. The data
197 used to train the BYOL model includes both labeled and unlabeled data: $(\mathbf{X}_U \cup \mathbf{X}_L)$, and the resulting
198 feature matrix is denoted as $\bar{\mathbf{X}}$. Please note that the feature extractor can be replaced by any other
199 well-trained feature extractor with the local continuity property (e.g., SimCLR [3] and MoCo [11]).

200 **Similarity Graph Construction and Approximation (Line 2).** After the extraction of feature
201 representation, we use the simple distance metric (i.e., cosine) to measure the similarity between any
202 two test instance \mathbf{x}_i and \mathbf{x}_j in $\bar{\mathbf{X}}$: $Dist(i, j) = \text{cosine}(\mathbf{x}_i, \mathbf{x}_j)$. Based on the distance matrix $Dist$, we
203 construct a k -NN Graph \mathcal{G} , wherein each sample is connected to its top- k most similar samples. The
204 connection is represented by an adjacency matrix $\mathbf{A} \in \mathcal{R}^{N \times N}$, where N is the number of sample in $\bar{\mathbf{X}}$.
205 The entry \mathbf{A}_{ij} equals 1 if node j is within the k nearest neighbors of node i , and 0 otherwise. The
206 edge weight matrix of the similarity graph is denoted as \mathbf{Edge} , wherein each edge weight in \mathbf{Edge} , if
207 exists, is inversely proportional to the corresponding distance $Dist$:

$$\mathbf{Edge}_{ij} = \begin{cases} 1/Dist(i, j) & \mathbf{A}_{ij} = 1. \\ 0 & \mathbf{A}_{ij} = 0. \end{cases} \quad i, j \in \{0, \dots, N - 1\}. \quad (3)$$

208 This means that the connection between two nodes, if exists, is weaker if their proximity is large.

Dataset	# Class	Size	Official Train/Test/Extra Split	Our Split(TC/DC/HO)	Model Architecture	Model Acc. On HO set(%) (A/B/C)
CIFAR-10	10	60K	50K/10K/x	20K/39K/1K	ResNet-18	70.1/66.4/68.3
SVHN	10	630K	73K/26K/531K	50K/49K/531K	Wide-ResNet	94.2/92.5/81.6
STL10	10	13K	5K/8K/x	5K/7.5K/0.5K	ResNet-34	54.8/54.0/53.6

Table 1: The Dataset and DL Models.

209 Constructing such a k -NN graph is, however, computationally expensive. This is because, calculating
210 the distance between each pair of test instances requires a computational complexity of $O(N^2)$, which
211 is prohibitive to scale up to the current massive unlabeled test instances in real applications. Therefore,
212 we propose an approximation method for k -NN graph construction. Our intuition is that, since the
213 target of graph construction is to exploit the failure patterns of the nearby labeled instances for the
214 unlabeled instances, the connections between unlabeled data are less meaningful. Therefore, we
215 propose to only consider the connections among labeled data \mathbf{X}_L , and the connections between labeled
216 \mathbf{X}_L to unlabeled data \mathbf{X}_U . This approximation reduces the cost from $O(N^2)$ to $O(P^2 + PQ)$, where P
217 and Q stand for the number of data in \mathbf{X}_L and \mathbf{X}_U , respectively. Usually, in the real-world scenario,
218 P is much smaller than Q , thereby we could obtain a near-linear graph construction algorithm with
219 complexity $O(PQ)$.

220 **GNN-based representation Learning (Line 3-12).** To apply the GNN algorithm, we first initialize
221 the input node representation matrix \mathbf{H}^0 in the similarity graph \mathcal{G} as $\bar{\mathbf{X}}$. Recall that in each GNN
222 layer, the node representations are propagated between neighbors and aggregated together. Thus, we
223 can obtain the representation in the next GNN layer by:

$$\mathbf{H}^{l+1} = \sigma(\tilde{\mathbf{D}}^{-\frac{1}{2}} \tilde{\mathbf{A}} \tilde{\mathbf{D}}^{-\frac{1}{2}} \mathbf{H}^l \Theta^l), \quad (4)$$

224 where $\tilde{\mathbf{A}} = \mathbf{Edge} + \mathbf{I}_N$, \mathbf{I}_N is the identity matrix, $\tilde{\mathbf{D}} = \sum_j \tilde{\mathbf{A}}_{i,j}$, Θ^l is the trainable weight matrix for the
225 l^{th} layer, σ is an activation function and \mathbf{H}^{l+1} is the output representation matrix. The propagation
226 and aggregation are repeated for M layers, with the output dimension of the M_{th} layer is 1 (for binary
227 classification purpose).

228 Then, for any labeled node $\mathbf{x}_e \in \bar{\mathbf{X}}_L$, we could obtain a cross entropy loss between the GNN output
229 h^M and the expected label $y \in \mathbf{Y}_L$ (e.g. misclassified or not): $\mathcal{L}_{ce} = -(y \log(h^M) + (1 - y) \log(1 - h^M))$,
230 where h^M denotes probability that \mathbf{x}_e is misclassified. The model is trained via minimizing the loss
231 for some training epochs (we set it as 600 in our experiment). After that, we apply the trained GNN
232 model (except the last layer) on \mathbf{X}_U to obtain the $\bar{\mathbf{E}}_c$ (line 12). In this way, the correctness of the
233 neighboring samples could be effectively summarized for each node.

234 3.4 Failure-revealing Capability Estimation

235 To properly combine both the intrinsic attributes \mathbf{e}_m and contextual attributes \mathbf{e}_c for collaborative
236 failure-revealing capability estimation, we formulate the combination function as a simple binary
237 classifier (e.g. a MLP). Specifically, the input to the MLP is a concatenation of \mathbf{e}_m and \mathbf{e}_c , and the
238 output is the failure-revealing estimation for an test instance. The final failure-revealing probability
239 is produced by applying a *sigmoid* function $S(t) = \frac{1}{1+e^{-t}}$ on the MLP model’s output. We use the
240 labeled instances ($\mathbf{X}_L, \mathbf{Y}_L$) to train the MLP in a supervised manner, with an objective of minimizing
241 the binary Cross-Entropy loss. After training, the MLP shall re-weight the importance of intrinsic
242 and contextual attributes and make a final decision by assigning a high probability to a test instance if
243 it is likely to reveal a failure.

244 Finally, we rank the unlabeled test instances in a descending order based on their failure-revealing
245 capability. Under the given budget, we select the top ones to label and test.

246 4 Experiment

247 4.1 Setup

248 **Datasets.** We evaluate the performance of TestRank on three popular image classification datasets:
249 CIFAR-10 [15], SVHN [17], and STL10 [4], as shown in Table 1. More elaboration is shown in the
250 Appendix.

251 There are mainly two parties involved in the model construction process: the training center and the
 252 debugging center. Hence, we manually split the dataset into the training center dataset (see the *TC*
 253 column in Table 1) and the debugging center dataset (see the *DC* column in Table 1). To mimic the
 254 practical scenario where the unlabeled data is abundant, we move a portion of training data to the
 255 debugging center to create this scenario. In the debugging center, we let a set of test data as labeled
 256 ones to represent the historical test oracles, and they are 8K/10K/1.5K for CIFAR-10, SVHN, and
 257 STL10, respectively, which are used to train the GNN and MLP model. The rest of the data in the
 258 debugging center are left unlabeled. Also, we spare a hold-out dataset (see the *HO* column), which is
 259 used for evaluating the model accuracy.

260 **Target DL model (model under test).** As shown in Table 1, we use the popular ResNet and
 261 WideResNet architectures as the backbone models [10, 22]. To simulate models of different qualities,
 262 for each dataset, we train three DL models with different randomly drawn sub-sets from the training
 263 set owned by the training center. For model B and C, the training set are drawn with in-equivalent
 264 class weights. After training, we report the accuracy of models on the debugging center’s hold-out
 265 dataset T_{HO} in Table 1.

266 **Evaluation metric.** We propose a new evaluation metric for test prioritization techniques: Test
 267 Relative Coverage (*TRC*). *TRC* is defined as the number of detected failures divided by the number of
 268 budget or the number of total failures identified by the whole unlabeled test set, whichever is minimal:

$$TRC = \frac{\#Detected\ Failures}{\min(\#Budget, \#Total\ Failures)}. \quad (5)$$

269 When $\# \text{ budget} \leq \# \text{ total failures}$, the maximum number of failures can be identified by a test
 270 prioritization technique equals to the budget. When $\# \text{ budget} \geq \# \text{ total failures}$, the maximum number
 271 of failures can be detected equals to the total number of failures. Therefore, *TRC* measures how far a
 272 test prioritization technique is to the ideal case.

273 In practice, under the massive unlabeled data, the performance under a small budget is considered
 274 more important than that under a large budget. To provide an insight on the quality of one test
 275 prioritization technique under a small budget, we also provide an *ATRC* metric: *ATRC* measures the
 276 average *TRC* values for budget values less than the total failures:

$$ATRC = \frac{1}{N} \sum_i^{N-1} TRC_i, \quad (6)$$

277 where TRC_i stands for the *TRC* value under budget b_i , $b_i \neq b_j$, and $b_i \leq$ number of total failures.

278 The proposed metrics enhance the ones used by Feng *et. al.* [6] and Byun *et. al.* [2]. They use the
 279 percentage of detected failures against the percentage of budget (and an APFD [5] value derived
 280 based on it) for evaluation. Their metric would produce a small value under a small budget, regardless
 281 of how good the prioritization technique is. For example, we assume that there are 10,000 unlabeled
 282 data, and 2,000 of them can detect model failures. If the budget is 100, the best percentage of detected
 283 failures is 5%, and the worst is 0%. Thus, under their metric, the gap between the best and the worst
 284 is only 5%. By contrast, *TRC* enlarges this gap to 100% to better differentiate the ability of different
 285 test prioritization techniques.

286 4.2 Comparison of TestRank with Baselines

287 We evaluate *TestRank* against five representative baselines: Random, DeepGini [6] (the state-of-the-
 288 art), MCP [20], DSA [13], and Dropout-uncertainty [2]. The details of each baseline are illustrated
 289 in the Appendix. For the dropout uncertainty method, we run 1000 times inferences with a default
 290 dropout rate of 0.5 (the dropout rate is consistent with the one used in [2]). For the DSA method,
 291 we collect the activation traces of the final convolution layer to calculate the surprise score. For our
 292 method, we set the number of neighbors for constructing the kNN graph as 100. Also, we use a
 293 two-layer GNN with a hidden dimension of 32. More ablation studies are in Section 4.3.

294 Table 2 compares *TestRank* with baselines using the *ATRC* metric. From this table, we have several
 295 observations. First, compared with the baselines, *TestRank* can achieve the highest *ATRC* values
 296 on all evaluated datasets and models. For instance, on CIFAR-10, *TestRank* can achieve 9.09%,
 297 20.07%, 14.38% higher *ATRC* values than the best baseline DeepGini for model A, B, C, respectively.

Dataset	Model ID	Random	MCP	DSA	Uncertainty	DeepGini	TestRank Contextual-Only	TestRank
CIFAR-10	A	30.15	58.25	60.93	58.09	67.47	51.39	76.56
	B	34.18	46.46	62.34	61.85	67.80	58.85	87.87
	C	34.27	65.25	64.47	63.10	71.15	75.33	85.53
SVHN	A	10.16	39.98	55.47	58.29	63.47	44.16	66.06
	B	11.85	38.07	57.96	58.06	63.85	51.26	76.36
	C	23.41	65.33	69.34	71.80	81.68	93.99	95.32
STL10	A	39.25	66.62	64.56	64.30	69.70	60.09	79.00
	B	42.60	69.97	67.12	65.30	72.89	71.90	80.96
	C	46.05	71.88	66.60	70.34	73.34	79.55	88.67

Table 2: Comparison of *TextRank* with baseline methods with ATRC values (%).

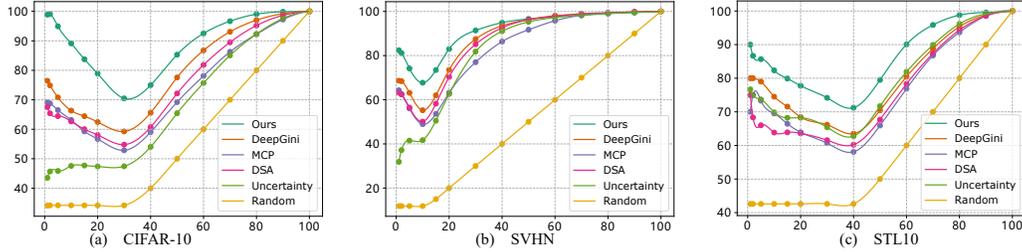


Figure 3: The TRC values against all budgets. X-axis: the budget (%). Y-axis: the TRC value. Note: this figure is generated with model B on each dataset.

298 Therefore, our method can distinguish the failure-revealing capability of the unlabeled test inputs
 299 much more accurately. Second, the *TextRank-Contextual-Only* column shows the result using only the
 300 contextual attributes. We observe that the contextual attributes alone can achieve higher effectiveness
 301 than random prioritization. For example, for model A on CIFAR-10, the effectiveness of random
 302 prioritization is 30.15% while that of the context-only method is 51.39%. We manually check the
 303 distribution of failures of model C and find that many failures are centralized on two classes, where the
 304 training data is insufficient. This kind of failure is easily detected by the contextual attributes-based
 305 method. Hence, the contextual information is helpful. But still, the context attributes alone are
 306 not sufficient. The combination of intrinsic and contextual attributes is essential in achieving high
 307 accuracy failure-revealing capability estimation.

308 To show more detailed results, we present the TRC value against every labeling budget in Figure 3.
 309 We observe that the TRC values for most curves decrease in the beginning and then increase. The
 310 turning point is when # budget = # total failures. When # budget < # total failures, the TRC values
 311 decrease because we rank the test instances according to their failure-revealing probabilities. With
 312 the budget increases, the selected test cases have a lower average failure-revealing ability, thus the
 313 decreased TRC value. When # budget > # total failures, according to the definition of TRC (see
 314 Equation 5), the denominator is fixed. Since the increase in budget will increase the number of
 315 detected failures, the TRC value will increase.

316 Figure 3 shows that our method consistently outperforms baselines, especially when the budget is
 317 small. For example, in Figure 3 (a), *TextRank* improves the prioritization efficiency by around 20%
 318 compared to the best baseline when the budget is around 1%. When the budget is rather high (e.g.
 319 budget $\geq 80\%$), the difference between different methods is less obvious because most failures can be
 320 selected under the large budget.

321 4.3 Influence of TestRank Configurations

322 **Feature Extractor.** *TextRank* uses an unsupervised BYOL model trained on both labeled and
 323 unlabeled data to extract their features. One may wonder if it can be replaced by a supervised model
 324 (e.g., the target DL model). To investigate this, we replace the feature extractor in *TextRank* with the
 325 front layers (we remove the last few linear layers) of the target DL model. The result is shown in the
 326 *TextRank-TargetModel* column in Table 3. Comparing with the original *TextRank*, the average ATRC
 327 value on the reported datasets and models reduces by 6.23%, which is significant. As the quality

Dataset	Model	TestRank (%)	TestRank w/o approx. (%)	TestRank TargetModel(%)
CIFAR-10	A	76.56	77.77 (+1.21)	68.84 (-7.71)
	B	87.87	87.70 (-0.17)	81.46 (-6.40)
	C	85.53	88.10 (+2.57)	77.73 (-7.79)
SVHN	A	66.06	63.87 (-2.19)	-
	B	76.36	82.04 (+5.68)	-
	C	95.32	96.62 (+1.30)	-
STL10	A	79.00	80.50 (+1.50)	67.59 (-11.40)
	B	80.96	78.98 (-1.98)	74.43 (-6.52)
	C	88.67	89.32 (+0.65)	72.43 (-16.23)
Average Influence (%)			+0.95	-6.23

Table 3: The performance (ATRC values) of TestRank under different configurations.

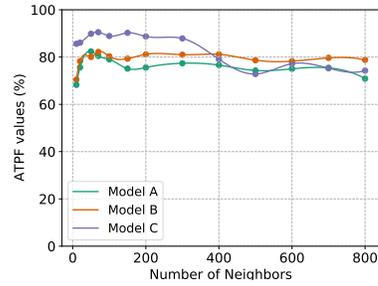


Figure 4: The impact of the number of neighbors k (STL10 dataset).

328 of the given model is to be examined, its feature extraction performance may not be reliable. Also,
 329 the dimension of the hidden layer could be huge, making it computationally expensive to calculate
 330 similarity values between input samples (This is why we do not report results on the SVHN dataset).
 331 In contrast, the separate model is more controllable, enabling it to extract better features for these
 332 data. Hence, using a separate feature extractor is necessary.

333 **k -NN graph approximation.** To reduce the computation complexity, TextRank uses approximation
 334 techniques when constructing the k -NN graph (see Section 3.3). The *TextRank-w/o-approx.* column in
 335 Table 3 shows the result when we use the original k -NN graph without approximation. It indicates that
 336 the average influence of the approximation is small (e.g. 0.95%). Therefore, if there is a computation
 337 resource limit and the unlabeled test instances are massive, we recommend using the approximation
 338 to save computation with negligible performance loss greatly.

339 **Number of nearest neighbors k .** When constructing the k -NN graph, the number of neighbors k
 340 decides the range of the context one node can reach. In previous experiments, the k is set to 100. We
 341 enlarge this range to (20 - 800) to study the influence. The result is shown in Figure 4.3. One can
 342 observe that the prioritization effectiveness will decrease when k is too small or too large. When k
 343 is too small, the context information available to one instance is limited, making it difficult for the
 344 GNN to extract valuable contextual attributes. On the other hand, when k is too large, the GNN may
 345 grasp irrelevant/noisy information. Still, TextRank can achieve good performance in a wide range of
 346 k values. For example, for model A, TextRank is better than the best baseline 69.70% (see Table 2)
 347 when k is larger than 20. Hence, selecting the number of nearest neighbors k is relatively flexible.

348 5 Conclusion and Future Work

349 We propose *TestRank*, a novel test prioritization framework for DL models. To estimate a test
 350 instance’s failure-revealing capability, *TestRank* not only leverages the intrinsic attributes of an input
 351 instance obtained from the target DL model, but also extracts the contextual attributes from the DL
 352 model’s historical inputs and responses. Our empirical results show that *TestRank* can capture the
 353 failure-revealing capabilities of unlabeled test instances more accurately than existing solutions.

354 This paper considers each test case equally and aims to identify as many failure-revealing test cases as
 355 possible. In practice, the impact of each test case could be different. We leave the study of such impact
 356 for future work. Besides, the current *TestRank* solution only supports CNN-based classification tasks,
 357 and we plan to extend it to other tasks with different DL models (e.g., RNN and GNN models) in the
 358 future.

References

- [1] Christopher M Bishop. *Pattern recognition and machine learning*. springer, 2006.
- [2] Taejoon Byun, Vaibhav Sharma, Abhishek Vijayakumar, Sanjai Rayadurgam, and Darren D. Cofer. Input prioritization for testing neural networks. In *AITest*, pages 63–70. IEEE, 2019.
- [3] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020.
- [4] A. Coates, A. Ng, and H. Lee. An analysis of single-layer networks in unsupervised feature learning. In *AISTATS*, 2011.
- [5] H. Do and G. Rothermel. On the use of mutation faults in empirical assessments of test case prioritization techniques. *IEEE Transactions on Software Engineering*, 2006.
- [6] Yang Feng, Qingkai Shi, Xinyu Gao, Jun Wan, Chunrong Fang, and Zhenyu Chen. Deepgini: Prioritizing massive tests to enhance the robustness of deep neural networks. *ISSTA 2020*, 2020. ISBN 9781450380089. doi: 10.1145/3395363.3397357.
- [7] M. Fey and J. E. Lenssen. Fast graph representation learning with pytorch geometric. *ArXiv*, abs/1903.02428, 2019.
- [8] Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *ICML*, volume 48 of *JMLR Workshop and Conference Proceedings*, pages 1050–1059. JMLR.org, 2016.
- [9] Jean-Bastien Grill, Florian Strub, Florent Alché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, Bilal Piot, koray kavukcuoglu, Remi Munos, and Michal Valko. Bootstrap your own latent - a new approach to self-supervised learning. In *Advances in Neural Information Processing Systems*, volume 33, pages 21271–21284, 2020.
- [10] Kaiming He, X. Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [11] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9729–9738, 2020.
- [12] Jinhan Kim, Robert Feldt, and Shin Yoo. Guiding deep learning system testing using surprise adequacy. *IEEE/ACM 41st International Conference on Software Engineering (ICSE)*, 2019.
- [13] Jinhan Kim, Robert Feldt, and Shin Yoo. Guiding deep learning system testing using surprise adequacy. In *Proceedings of the 41st International Conference on Software Engineering, ICSE '19*, page 1039–1049. IEEE Press, 2019. doi: 10.1109/ICSE.2019.00108. URL <https://doi.org/10.1109/ICSE.2019.00108>.
- [14] Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *ICLR (Poster)*. OpenReview.net, 2017.
- [15] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. Technical report, Citeseer, 2009.
- [16] Radford M Neal. *Bayesian learning for neural networks*, volume 118. Springer Science & Business Media, 2012.
- [17] Yuval Netzer, T. Wang, A. Coates, A. Bissacco, B. Wu, and A. Ng. Reading digits in natural images with unsupervised feature learning. 2011.
- [18] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *arXiv preprint arXiv:1912.01703*, 2019.

- 405 [19] Michael D Richard and Richard P Lippmann. Neural network classifiers estimate bayesian a
406 posteriori probabilities. *Neural computation*, 3(4):461–483, 1991.
- 407 [20] W. Shen, Y. Li, L. Chen, Y. Han, Y. Zhou, and B. Xu. Multiple-boundary clustering and
408 prioritization to promote neural network retraining. In *2020 35th IEEE/ACM International*
409 *Conference on Automated Software Engineering (ASE)*, 2020.
- 410 [21] Aäron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive
411 predictive coding. *CoRR*, abs/1807.03748, 2018.
- 412 [22] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. In *BMVC*, 2016.

413 **Checklist**

- 414 1. For all authors...
- 415 (a) Do the main claims made in the abstract and introduction accurately reflect the paper's
416 contributions and scope? [Yes]
- 417 (b) Did you describe the limitations of your work? [Yes] Please refer to the second
418 paragraph in Section 5 for details.
- 419 (c) Did you discuss any potential negative societal impacts of your work? [Yes] See the
420 Appendix for details.
- 421 (d) Have you read the ethics review guidelines and ensured that your paper conforms to
422 them? [Yes]
- 423 2. If you are including theoretical results...
- 424 (a) Did you state the full set of assumptions of all theoretical results? [N/A]
425 (b) Did you include complete proofs of all theoretical results? [N/A]
- 426 3. If you ran experiments...
- 427 (a) Did you include the code, data, and instructions needed to reproduce the main experi-
428 mental results (either in the supplemental material or as a URL)? [Yes] See supplemen-
429 tary for details.
- 430 (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they
431 were chosen)? [Yes] Please see Section 4 for details.
- 432 (c) Did you report error bars (e.g., with respect to the random seed after running experi-
433 ments multiple times)? [No] Because the baselines we compare do not report this as
434 well.
- 435 (d) Did you include the total amount of compute and the type of resources used (e.g., type
436 of GPUs, internal cluster, or cloud provider)? [Yes] Please refer to the first paragraph
437 in Section 4.
- 438 4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
- 439 (a) If your work uses existing assets, did you cite the creators? [Yes] Please refer to the
440 Experimental Section (Section 4) for details.
- 441 (b) Did you mention the license of the assets? [Yes] Please refer to the Appendix for
442 details.
- 443 (c) Did you include any new assets either in the supplemental material or as a URL? [Yes]
444 See the supplemental material for details.
- 445 (d) Did you discuss whether and how consent was obtained from people whose data you're
446 using/curating? [Yes] Please see the Appendix for details.
- 447 (e) Did you discuss whether the data you are using/curating contains personally identifiable
448 information or offensive content? [Yes] Please see the Appendix for details.
- 449 5. If you used crowdsourcing or conducted research with human subjects...
- 450 (a) Did you include the full text of instructions given to participants and screenshots, if
451 applicable? [N/A]
- 452 (b) Did you describe any potential participant risks, with links to Institutional Review
453 Board (IRB) approvals, if applicable? [N/A]
- 454 (c) Did you include the estimated hourly wage paid to participants and the total amount
455 spent on participant compensation? [N/A]

456 A Appendix

457 A.1 Experiment details

458 Our code is implemented with the open-source PyTorch (under BSD license) [18] and PyG (under
459 MIT license) [7] ML libraries. All of our experiments are performed on a single TITAN V GPU ¹.

460 A.1.1 Datasets

461 CIFAR-10 is officially composed of 50,000 training images and 10,000 test images, and it has ten
462 classes of natural images. The Street View House Numbers (SVHN) dataset contains house numbers
463 from Google Street View images. It contains 73,257 training images and 26,032 testing images.
464 Besides, the SVHN dataset also has an extra set of 531,131 images. The STL10 dataset contains ten
465 classes of natural images. In each class, there are 500 training images and 800 test images.

466 Please note that the datasets we used are open-sourced and available for research purposes ^{2 3 4}.
467 Also, since the data we used is about numbers and the animals, we assume there are no personally
468 identifiable information or offensive content.

469 A.1.2 Baselines

470 Given a DL model and a certain budget, the goal of our method is to select test cases from an
471 unlabeled data pool to discover the failures of the given DL model. We compare our work with the
472 following representative test prioritization techniques:

- 473 • **DeepGini [6]**: DeepGini is the state-of-art test case selection technique. DeepGini ranks
474 unlabeled test cases by a score defined based on the output confidence (See Equation 2).
- 475 • **MCP [20]**: In addition to the output confidence, MCP also considers the balance among
476 different class boundaries of the selected test inputs. Specifically, MCP groups test cases
477 into different clusters, where each cluster stands for a distinct classification boundary, and
478 try to equally choose low confidence test cases from each cluster.
- 479 • **DSA [2]**: Byun *et al.* propose to use the distance-based surprise score (DSA) as a test
480 prioritization metric, which is originally proposed in [12]. The surprise score measures the
481 distance between the test case to the training set. Samples with higher surprise scores are
482 prioritized.
- 483 • **Uncertainty [2]**: By running the model multiple times (e.g. t times) with a certain dropout
484 rate, the uncertainty is calculated as the entropy on the averaged output probabilities.
- 485 • **Random**: Test inputs are randomly drawn from all unlabeled samples.

486 A.2 Limitation

487 Besides the number of detected failures, failure diversity is another important factor for model
488 debugging. In this work, we have the implicit assumption that the historical labeled data can cover the
489 input distribution, and under such circumstances, TestRank can prioritize unlabeled tests effectively.
490 If, however, the historical test data is severely biased, before prioritizing tests with TestRank, we
491 should analyze the test pool and try to fill this gap first. Otherwise, the detected failures are lack of
492 diversity. We shall consider this problem in our future research.

493 A.3 Broader Impacts

494 This work targets on building an efficient and effective test prioritization technique for deep learning
495 models, which can help ensure the security of deep learning models after deployment. A variety of
496 safety-critical tasks, such as autonomous vehicles, industrial robotics, and medical diagnosis, can
497 benefit from such test prioritization technique. We consider the scenario in which the unlabeled data

¹<https://www.nvidia.com/en-us/titan/titan-v/>

²<http://ufldl.stanford.edu/housenumbers/>

³<https://cs.stanford.edu/acoates/stl10/>

⁴<https://www.cs.toronto.edu/kriz/cifar.html>

498 is abundant and a subset of unlabeled data is selected for labeling and testing. While building the
499 unlabeled data by collecting them from the Internet or other sources, it may have a chance to include
500 some unauthorized or private data. Also, the collected unlabeled data could be biased, resulting in the
501 testing being incomplete. To avoid such cases, we should always guarantee that the collected data
502 will not violate any kind of privacy or rights, and should also ensure that the collected data cover
503 instances as many as possible.