

AEON: A Method for Automatic Evaluation of NLP Test Cases

Jen-tse Huang
The Chinese University of Hong Kong
Hong Kong, China
jthuang@cse.cuhk.edu.hk

Jianping Zhang
The Chinese University of Hong Kong
Hong Kong, China
jpzhang@cse.cuhk.edu.hk

Wenxuan Wang
The Chinese University of Hong Kong
Hong Kong, China
wxwang@cse.cuhk.edu.hk

Pinjia He*
The Chinese University of Hong
Kong, Shenzhen
Shenzhen, China
hepinjia@cuhk.edu.cn

Yuxin Su
Sun Yat-sen University
China
suyx35@mail.sysu.edu.cn

Michael R. Lyu
The Chinese University of Hong Kong
Hong Kong, China
lyu@cse.cuhk.edu.hk

ABSTRACT

Due to the labor-intensive nature of manual test oracle construction, various automated testing techniques have been proposed to enhance the reliability of Natural Language Processing (NLP) software. In theory, these techniques mutate an existing test case (e.g., a sentence with its label) and assume the generated one preserves an equivalent or similar semantic meaning and thus, the same label. However, in practice, many of the generated test cases fail to preserve similar semantic meaning and are unnatural (e.g., grammar errors), which leads to a high false alarm rate and unnatural test cases. Our evaluation study finds that 44% of the test cases generated by the state-of-the-art (SOTA) approaches are false alarms. These test cases require extensive manual checking effort, and instead of improving NLP software, they can even degrade NLP software when utilized in model training. To address this problem, we propose AEON for Automatic Evaluation Of NLP test cases. For each generated test case, it outputs scores based on semantic similarity and language naturalness. We employ AEON to evaluate test cases generated by four popular testing techniques on five datasets across three typical NLP tasks. The results show that AEON aligns the best with human judgment. In particular, AEON achieves the best average precision in detecting semantic inconsistent test cases, outperforming the best baseline metric by 10%. In addition, AEON also has the highest average precision of finding unnatural test cases, surpassing the baselines by more than 15%. Moreover, model training with test cases prioritized by AEON leads to models that are more accurate and robust, demonstrating AEON's potential in improving NLP software.

CCS CONCEPTS

• **Software and its engineering** → **Software testing and debugging.**

*Corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ISSTA 2022, 18–22 July, 2022, Daejeon, South Korea

© 2022 Association for Computing Machinery.
ACM ISBN 978-1-4503-XXXX-X/18/06...\$15.00
<https://doi.org/10.1145/1122445.1122456>

KEYWORDS

NLP software testing, test case quality

ACM Reference Format:

Jen-tse Huang, Jianping Zhang, Wenxuan Wang, Pinjia He, Yuxin Su, and Michael R. Lyu. 2022. AEON: A Method for Automatic Evaluation of NLP Test Cases. In *ISSTA '22: ACM SIGSOFT International Symposium on Software Testing and Analysis*, 18–22 July, 2022, Daejeon, South Korea. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/1122445.1122456>

1 INTRODUCTION

NLP software has become increasingly popular in our daily lives. For example, NLP virtual assistant software, such as Siri and Alexa, receives billions of requests [48, 64] while Google Translate App translates more than 100 billion words per day [69]. With the development of Deep Neural Networks (DNNs), the performance of NLP software has been largely boosted. Equipped with the SOTA model [70], Microsoft question answering robot surpasses humans on conversational question answering task. In addition, the performance of machine comprehension [19], text generation [41] and machine translation [73] has been significantly improved. However, NLP software can produce erroneous results, leading to misunderstanding, financial loss, threats to personal safety, and political conflicts [53, 54].

To discover erroneous behaviors in NLP software, researchers have designed various software testing techniques [10, 25, 40, 63, 81]. A test case for NLP software is in the form of a text (e.g., a sentence) and its label, where the label is the expected correct output of the NLP software. In theory, most of these testing techniques modify part(s) of the input text (e.g., a word/character substitution/insertion/deletion) under the assumption that the generated test case preserves an equivalent or similar semantic meaning. Typically, these techniques take labeled texts as inputs and output the mutated texts and the corresponding labels.

However, it is still challenging for current testing techniques to produce practical test cases of high quality. Specifically, tiny modification in a text can change its semantic meaning, which invalidates the common assumption that the semantic meaning of the original text and that of the generated text should remain equivalent or similar, further rendering the possibility of changing the corresponding labels [49, 51]. For example, removing “not” from the text “I do not like the movie” changes its semantic meaning and further changes its label for a sentiment analysis task from “negative” to “positive”, resulting in a test case with an incorrect label

Table 1: Examples for high-quality, inconsistent, and unnatural test cases generated by existing testing techniques on different datasets. NLP tasks include Sentiment Analysis (SA), Natural Language Inference (NLI), and Semantic Equivalence (SE). Mutated words are marked in red.

Original text Generated test case	Task Dataset	Technique Issue
A man under a running shower with shampoo in his hair. ⇒ A man is taking a shower . A man under a running shower with shampoo in his hair. ⇒ A man is taking a bath .	NLI SNLI [3]	BAE [20] None
Ultimately this is a frustrating patchwork. Ultimately this is a sparkling patchwork.	SA MR [55]	PSO [81] Inconsistent
British action wouldn't have mattered. ⇒ British action would have made a big difference. Welsh action wouldn't have mattered. ⇒ British action would have made a big difference.	NLI MNLI [74]	BAE [20] Inconsistent
What are some good topics to be bookmarked on Quora? What are some good topics to es bookmarked on Quora?	SE QQP [71]	Textfooler [33] Unnatural
I went to Danny's this weekend to get an oil change and car wash and I paid for a VIP car wash. My gone to work this weekend to do an oil change and car wash and my hired for a VIP car wash.	SA Yelp [88]	BAE [20] Unnatural

and further a false alarm. Moreover, existing testing approaches cannot guarantee the fluency and naturalness of the generated test cases. Many word-level testing approaches introduce grammar errors and punctuation errors, and sometimes they introduce words that do not exist or are rarely used [49]. Although these test cases may trigger “software errors” (e.g., unexpected software behaviors), it is important to first ensure the quality of the test cases in terms of semantic consistency and naturalness before finding more errors.

According to our user study, many of the NLP test cases generated by existing approaches are of low quality because of the following two issues: *Inconsistent* issue and *Unnatural* issue. These issues can lead to false alarms in testing and unnaturalness in language. In this paper, we say an NLP test case is of high quality if it does not have any of these issues. As shown in Table 1, a high-quality test case preserves the semantics of the original text and reads smoothly. The first *Inconsistent* case changes the semantics to the opposite while the second one changes the subjects. Two *Unnatural* cases hurt the fluency and naturalness of natural language by introducing either non-existing words or wrong grammar. It is unlikely that these low-quality test cases can contribute to improving NLP software in practice.

Hence, an automatic quality evaluation metric that can help filter out low-quality test cases generated by the existing testing techniques is highly in demand. Nevertheless, designing an automatic quality evaluation metric for NLP test cases is highly challenging. First, existing testing criteria are mainly based on coverage metrics, such as code coverage for traditional software [8] and neuron coverage for deep neural networks [57], which cannot be directly leveraged to detect false alarms and evaluate the quality of a natural language test case. Second, general semantic similarity evaluation metrics fail to detect *Inconsistent* issues under this scenario. Specifically, (1) most of the words in the original text and the generated text are the same while existing metrics evaluate the semantic similarity based on all the words in the text and thus, the impact of the mutated word(s) easily vanish; (2) a word may have different meanings in different contexts, making it difficult to compare only the

mutated word(s). Third, existing work on naturalness evaluation metric either relies on human evaluation [51] or qualitative analysis (e.g., part-of-speech checking [20]), while we need an automatic and quantitative naturalness evaluation metric.

To address these problems, we introduce AEON, a method for Automatic Evaluation Of NLP test cases. AEON takes a text pair <original text, generated text> as input and outputs scores regarding semantic similarity and syntactic correctness, aiming for detecting *Inconsistent* and *Unnatural* issues, respectively. We use AEON to analyze the quality of NLP test cases generated from four popular testing techniques [20, 32, 63, 81] on five datasets [3, 55, 71, 74, 88] which cover three typical NLP tasks, namely natural language inference, sentiment analysis, and semantic equivalence. We conduct a comprehensive human evaluation on the semantic similarity and language naturalness between the original texts and the generated test cases, and we check whether AEON’s score aligns with human evaluation or not. The results show that AEON achieves the Average Precision (AP), Area Under Curve (AUC), and Pearson Correlation Coefficient (PCC) scores of 0.688, 0.742, and 0.922, outperforming the best baseline metric by 10%, 8.1%, and 7.8% respectively. On the evaluation of human judgment of language naturalness, AEON also surpasses all baselines and achieves the average AP, AUC, PCC scores of 0.69, 0.63, 0.82. These results demonstrate the effectiveness of AEON on detecting false alarms and evaluating the language naturalness of NLP test cases. We also show that the high-quality test cases selected by AEON can significantly improve the accuracy and robustness of NLP software via model training. Our contributions can be summarized as:

- We conduct a comprehensive user study on the test cases generated by existing NLP software testing techniques and find that 85% of them suffer from two issues: *Inconsistent* and *Unnatural*, resulting in a false alarm rate of 44%.
- We introduce AEON, the first approach to quantitatively evaluate the quality of NLP test cases from semantics and language naturalness, addressing two main quality issues of NLP test cases mentioned above.

Table 2: Details of the selected testing techniques.

Technique	Selection	Substitution	Constraints
Generative Algorithm (GA) [32]	k-nearest neighbors	Random combination	Percentage of modified words; Euclidean distance; LM grammar checking
BERT-base Adversarial Examples (BAE) [20]	Word importance	PLM mask prediction	Euclidean distance; Part-of-speech checking
Particle Swarm Optimization (PSO) [81]	Optimization [35]	Knowledge graph [14]	None
Checklist [63]	Random selection	Transformations: Contraction; Extension; Changing entities	Numbers of transformation

- AEON is employed to evaluate the test cases generated by four testing techniques on five widely-used datasets, which shows that AEON achieves the best performance in terms of average AP, AUC, and PCC on all datasets.
- The implementation of AEON, the raw experimental results, and the human annotation on the test case quality are available on Github¹.

2 PRELIMINARIES

2.1 Testing Techniques for NLP Software

Though many papers have proposed testing techniques for Computer Vision (CV) software (e.g., face recognition system) [5, 57, 68, 79], the characteristics of natural language make NLP software testing distinguished from that in CV software. The most significant difference between NLP test cases and CV test cases is that the input space of textual data is not as continuous as images, making every mutation in the original text perceptible. In addition, in natural language, mutating a single word can cause considerable semantic differences, which further leads to the risk of changing the correct label of the text. Therefore, when NLP testing techniques assign the label of the original text to the generated test cases, lots of false alarms occur.

Current testing techniques² for NLP software can be roughly divided into four categories: character-level, word-level, sentence-level, and multi-level [87]. Character-level techniques [40] mutate a few characters that do not affect human reading comprehension. Word-level techniques [62, 63] are based on word substitution, usually using synonyms sets or Pre-trained Language Models (PLMs). Sentence-level techniques [26] change the whole structure of the sentences either by adding a sentence to the original texts or transforming the entire texts into another semantically similar format. Those combining different levels of techniques [42] can be categorized into multi-level techniques. In particular, word-level techniques significantly outperform others in terms of efficiency [81], applicability, and usefulness in robust training [62, 81]. However, this kind of technique suffers more from low-quality test cases

[51]. Thus, we focus on test cases generated by word-level testing techniques.

From the perspective of combinatorial optimization, generating test cases with word-level techniques can be formulated as a searching problem, where we substitute each word in the original text to other words in our vocabulary. The whole search space is the number of words in original text N (where we substitute) times the vocabulary size V (word candidates). In general, these techniques include diverse modules to prune the search space, which can be classified into three components: target word selection, word substitution, and generation constraints [52, 87]. Table 2 presents the modules of the four selected testing techniques in terms of the three components. A suitable target word selection method can decrease N while a proper word substitution method can cut back V . Constraints are commonly applied to ensure that the synthesized texts preserve semantic meaning and are syntactically correct.

2.2 Problem Definition

Given NLP software $\mathcal{F} : \mathcal{X} \rightarrow \mathcal{Y}$ which takes a text x in text space \mathcal{X} as input and outputs its prediction $y \in \mathcal{Y}$, a word-level test case \hat{x} is synthesized from a seed datum x whose ground truth (i.e., label) is y (denoting with $gt(x) = y$). It also needs to satisfy that $Sim(\hat{x}, x) \geq c$, where $Sim : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ is a similarity metric. Intuitively, it means \hat{x} and x have equivalent or similar semantics. c is a task-specific constant to trade between semantic similarity and generation diversity. Most NLP testing techniques assume that the generated test case and the original text have the same label, i.e., $gt(\hat{x}) = gt(x) = y$.

Given a set of generated test cases \hat{X} , sometimes the similarity metric it uses may not be able to detect some inconsistency, then *Inconsistent* test cases occur. A test case $\hat{x} \in \hat{X}$ is *Inconsistent* when $Sim'(\hat{x}, x) \leq t_s$. Here t_s is a threshold, and Sim' is a trustable and robust similarity metric, for example, human judgment. There is a high chance that *Inconsistent* test cases are false alarms, which satisfy $gt(\hat{x}) \neq gt(x)$ and $\mathcal{F}(\hat{x}) = gt(\hat{x})$. In other words, the software behaves correctly (produces label $gt(\hat{x})$), but we wrongly assume it needs to produce label $gt(x)$. Whether a test case is *Unnatural* or not depends on human evaluation. It is defined as $Nat'(\hat{x}) \leq t_n$ where $Nat' : \mathcal{X} \rightarrow \mathbb{R}$ is human judgment on the language naturalness of the textual test case and t_n is a threshold.

¹<https://github.com/CUHK-ARISE/AEON>

²In this paper, we consider papers on attacking NLP models as a line of research on testing NLP software because the adversarial examples generated by these techniques can be regarded as test cases for NLP software.

Algorithm 1 Algorithm for SemEval

Input: Original text x ; Generated test case \hat{x}
Output: Semantic similarity score $Sim(x, \hat{x}) \in [0, 1]$

- 1: $x \leftarrow \text{Tokenize}(x)$
- 2: $\hat{x} \leftarrow \text{Tokenize}(\hat{x})$
- 3: $diff_indices \leftarrow \text{LevenshteinDistance}(x, \hat{x})$
- 4: $emb(x) \leftarrow \text{GetWordEmbedding}(x)$
- 5: $emb(\hat{x}) \leftarrow \text{GetWordEmbedding}(\hat{x})$
- 6: $patch_sim \leftarrow \text{a list}$
- 7: **for** each $i \in diff_indices$ **do**
- 8: $patch_x \leftarrow emb(x)[i - 2 : i + 2]$
- 9: $patch_x \leftarrow emb(\hat{x})[i - 2 : i + 2]$
- 10: Append $\frac{patch_x^T patch_x}{\|patch_x\| \cdot \|patch_x\|}$ to $patch_sim$
- 11: **end for**
- 12: $text_sim \leftarrow \frac{emb(x)^T emb(\hat{x})}{\|emb(x)\| \cdot \|emb(\hat{x})\|}$
- 13: $min_sim \leftarrow \min(patch_sim)$
- 14: $avg_sim \leftarrow \text{average}(patch_sim)$
- 15: $Sim(x, \hat{x}) \leftarrow \lambda_1 min_sim + \lambda_2 avg_sim + (1 - \lambda_1 - \lambda_2) text_sim$

The task of this paper is to design an automatic evaluation metric that can reflect test case quality in terms of semantic consistency and naturalness, which facilitates the detection of *Inconsistent* (false alarms) and *Unnatural* test cases.

3 APPROACHES AND IMPLEMENTATION

This section introduces the details of AEON whose input is a text pair <original text, generated text> and outputs are a semantic score and a syntactic score. AEON consists of two parts: SemEval (Semantic Evaluator), which captures the semantic difference between input text pair, and SynEval (Syntactic Evaluator), which assesses how likely the generated test case will be used (*i.e.*, written or typed) by real users. These two components aim to address *Inconsistent* and *Unnatural* issues, respectively. In the rest of this section, we will introduce the details of the key components of the two evaluators.

3.1 SemEval

SemEval aims to solve the two challenges mentioned above. (1) The influence of the mutated position can easily vanish when taking average since most words in the original text and the generated test case are the same. (2) Metrics comparing words without contexts can neglect their alternative meanings (*i.e.*, polysemy). To this end, we propose to combine Levenshtein distance [38] and sentence embedding model to evaluate the semantic similarity in the NLP testing scenario. The approach is surprisingly effective considering its simplicity, which is shown in Alg. 1.

After tokenizing the input texts (line 1-2), which converts all words and punctuation as individual tokens, SemEval extracts small patches of text where the two inputs differ using Levenshtein distance (line 3). With the help of Levenshtein distance, we can find all mutated positions in linear time. Next, it applies a PLM to obtain the embeddings of all tokens in the two inputs (line 4-5). Current PLMs [11, 13, 18, 43] can be leveraged to project tokens to the embedding space, so this module can be replaced easily when more powerful

Algorithm 2 Algorithm for SynEval

Input: Generated test case \hat{x}
Output: Language naturalness score $Nat(\hat{x}) \in (0, 1]$

- 1: $\hat{x} \leftarrow \text{Tokenize}(\hat{x})$
- 2: $perplexity \leftarrow \text{a list}$
- 3: **for** each $token \in \hat{x}$ **do**
- 4: $masked_text \leftarrow \text{replace } token \text{ with [MASK] in } \hat{x}$
- 5: $prob \leftarrow \text{GetMaskPrediction}(masked_text)$
- 6: $prob_token \leftarrow prob[token]$
- 7: Append $prob_token$ to $perplexity$
- 8: **end for**
- 9: $min_nat \leftarrow \min(perplexity)$
- 10: $avg_nat \leftarrow \text{average}(perplexity)$
- 11: $Nat(\hat{x}) \leftarrow \phi min_nat + (1 - \phi) avg_nat$

PLMs are proposed. Then, for all patches and the whole text, we compute the cosine similarity defined by $\frac{a^T b}{\|a\| \cdot \|b\|}$ in line 7-9 and line 12, respectively. Note that we extract totally five tokens as the patch for a mutation happened in position i by $[i - 2 : i + 2]$. For a mutation at the beginning or the end of a sentence, we extract the first or last three tokens as our patch. Finally, we compute the minimum and the average numbers among all patch similarities (line 13-14) and combine them with the text similarity using two hyper-parameters, λ_1 and λ_2 (line 15). After this convex combination, we obtain the output of SemEval, namely $Sim(x, \hat{x})$.

We tackle challenge (1) by considering the minimum and average patch similarities. For challenge (2), AEON extract the mutated position along with its context, which can improve its ability to understand semantics. Consider an example:

Case Study 1	
TASK/DATASET	SE/QQP
TECHNIQUE	BAE
ORIGINAL TEXT	Is it OK to leave an iPhone plugged into the charger after 100% charged ?
GENERATED TEXT	Is it OK to leave an iPhone plugged into the charger after 100% indicted ?

If we only consider the mutated position **charged** and **indicted**, the similarity is high since they are synonyms in the meaning of "being accused". However, **charged** here means "to put electricity into an electrical device". This kind of relationship can be captured by its context, which is modeled in the PLMs.

3.2 SynEval

Since synthesized test cases may include grammar errors, punctuation errors, or produce rarely used words and phrases, it is vital to use an automatic and quantitative metric to filter out these *Unnatural* test cases. Note that this kind of sentence rarely appears in real-world natural languages, hence they are treated as noises and ignored during the training process of PLMs [13, 37, 43]. Intuitively, how natural a sentence is can be reflected by the probability that the sentence has the same distribution as its training data, which can be estimated by PLMs. Therefore, SynEval is designed to measure naturalness through the perplexity of PLMs. Perplexity, in its

Table 3: Details of the selected datasets. NLP tasks include Sentiment Analysis (SA), Natural Language Inference (NLI), and Semantic Equivalence (SE)

Dataset	Task	Classes	Description
Rotten Tomatoes Movie Review (MR) [55]	SA	2	Short sentences or phrases of movie reviews
Yelp Restaurant Review (Yelp) [88]	SA	2	Long sentences or paragraphs of restaurant reviews
Stanford Natural Language Inference (SNLI) [3]	NLI	3	Short texts with simple contexts
Multi-Natural Language Inference (MNLI) [74]	NLI	3	Multi-genre, multi-length texts with complicated contexts
Quora Question Pairs (QQP) [71]	SE	2	Two similar questions from Quora

formal definition, is the exponential form of the cross entropy of the given sentence [31], having the form of:

$$perplexity(x) = \sqrt[N]{\prod_{i=1}^N \frac{1}{P(x_i|x_{1:i-1})}}, \quad (1)$$

where x_i is the i -th word in the sentence and $x_{1:i-1}$ is the first to $i - 1$ -th words in the sentence. $perplexity : \mathcal{X} \rightarrow [1, \infty)$ measures how confused the PLM is when it sees x_i given $x_{1:i-1}$, the greater the more confused (*i.e.*, worse).

The recently proposed BERT-like models, including BERT [13], RoBERTa [43], and ALBERT [37] which trained on billions of sentences, are powerful PLMs for modeling this probability. However, BERT and its variants are bi-directional, taking not only $x_{1:i-1}$ but also $x_{i+1:n}$ as input. Therefore, we need to replace $x_{1:i-1}$ with $x_{\setminus i}$ in Eq. 1, where $x_{\setminus i}$ denotes the input sentence with its i -th word being [MASK]. Since our semantic evaluator outputs similarity scores in $(0, 1]$ (the greater, the more similar), we adopt $\sqrt[N]{\prod_{i=1}^N P(x_i|x_{\setminus i})}$ for SynEval, having the same value range of $(0, 1]$ (the greater, the better).

Alg. 2 illustrates the implementation of SynEval. First we tokenize the input (line 1). Then, for each token in the input, we replace it with the special token [MASK] (line 4). Feeding the masked text to the PLM, we can obtain the prediction of the masked position, which is a probability distribution over the entire vocabulary (line 5). Next, we find out the probability that the PLM thinks the masked position can be filled with the original token and record it as the perplexity of this token (line 6-7). Finally, we compute the minimum and the average numbers among all perplexities and combine them using a hyper-parameter ϕ . The score after this convex combination is the output of SynEval, namely $Nat(\hat{x})$.

4 EXPERIMENTAL DESIGN AND SETTINGS

In this paper, we focus on the following four research questions:

RQ1: What is the quality of the test cases generated by existing testing techniques (Section 5.1)?

RQ2: How effective is AEON (Section 5.2)?

RQ3: How can AEON help in testing NLP software? (Section 5.3)

RQ4: How can AEON help in improving NLP model? (Section 5.4)

4.1 Testing NLP Software

To answer the RQs, the first step is generating test cases, *i.e.*, testing NLP software. We choose to test the APIs provided by Hugging

Face Inc.³, the largest NLP open-source community, on five widely-used datasets across three typical tasks: sentiment analysis, natural language inference, and semantic equivalence.

Datasets. Sentiment analysis aims at classifying the polarity (either positive or negative) of the sentiment of given texts. The inputs of natural language inference tasks are two pieces of texts, namely Premise and Hypothesis, and the target is to predict whether the Hypothesis is a contradiction, entailment, or neutral to the given Premise. If the Premise can infer the Hypothesis, the output is entailment; if the Premise can infer NOT Hypothesis, the output is contradiction; otherwise, the output is neutral. The inputs of semantic equivalence tasks are two pieces of text, namely question 1 and question 2, and the objective is to judge if the meaning of the two given questions is equivalent. We select five datasets, namely MR, Yelp, SNLI, MNLI, and QQP, for our experiments, whose details are shown in Table 3. MR and Yelp are crawled from the internet, so the data contain noises such as HTML tags, HTML encodings, HTML entity names, and hyperlinks, which will make the generated test case hard to read. To eliminate the influence of noisy data in our human evaluation, we convert HTML texts to plain texts and remove hyperlinks using regular expressions.

Testing. To be more specific, we choose five BERT-based APIs⁴ for five different datasets. According to the statistics given by Hugging Face Inc, these APIs are downloaded more than 30k times every month on average. Using the testing techniques described in Table 2 implemented by TextAttack [52] with their default settings, we generate test cases for all datasets (APIs). We select 400 original texts for each dataset using each technique, resulting in 8,000 test cases. After testing the APIs with our test cases, 3,262 test cases (40.8%) are reported as software errors.

4.2 Human Evaluation

We aim to find out whether the reported cases really trigger the erroneous behaviors of NLP software, in other words, whether they are false alarms. To this end, we design and launch a user study.

Design. Following [6], we propose a unified framework to measure the quality of generated test cases. The quality is defined from

³<https://huggingface.co/>

⁴<https://huggingface.co/textattack/bert-base-uncased-rotten-tomatoes>
<https://huggingface.co/textattack/bert-base-uncased-yelp-polarity>
<https://huggingface.co/textattack/bert-base-uncased-MNLI>
<https://huggingface.co/textattack/bert-base-uncased-snli>
<https://huggingface.co/textattack/bert-base-uncased-QQP>

Table 4: Details of the selected baselines.

	Baseline	Description
NC-based	NC [57]	The ratio of activated neurons
	NBC [44]	Activation outside upper/lower bounds
	TKNC [44]	The ratio of top- k activated neurons
	BKNC [79]	The ratio of bottom- k activated neurons
NLP-based	BLEU [56]	The overlaps of n -grams
	Meteor [12]	n -grams with synonyms in WordNet [50]
	InferSent [11]	BiLSTM-based embedding model
	SBERT [61]	BERT-based embedding model
	SimCSE [18]	Embedding model with contrastive learning
	BERTScore [86]	Token matching in BERT embedding space

four perspectives, including Naturalness, Consistency, Human Label, and Difficulty:

- Consistency: From “1 strongly disagree” to “5 strongly agree”, how much do you think the two sentences have the same meaning? Consistency quantifies the semantic similarity between the original text and the changed text.
- Naturalness: From “1 very bad” to “5 very good”, how fluent and natural do you think this sentence is? Naturalness measures the fluency and grammar of the examples, including grammar errors, punctuation errors, and spelling errors (unrecognizable words).
- Human label: Ask humans to do the tasks of the given datasets. It is a task-specific question and records the human judgment of classification answers.
- Difficulty: From “1 very easy” to “5 very hard”, how difficult for you to make the decision? Difficulty reflects how difficult the task is for humans.

Based on our definition in Sec. 1, high-quality text cases should have high naturalness and consistency scores. Human label and difficulty are used to classify the human evaluation results. We also ask annotators whether these test cases have other problems/issues that we have not identified. The responses show that *Inconsistent* issue and *Unnatural* issue can cover all their concerns.

Crowdsourcing. We distribute our questionnaire on Qualtrics⁵, a platform to design, share, and collect questionnaires. We recruit crowd workers on Prolific⁶, a platform to post tasks and hire workers. Since our questions require a high level of reading comprehension and inference skills in English, we require Prolific workers to have a bachelor’s degree or above and have English as their first and most fluent language. Since we focus on false alarms, we randomly sample 100 test cases per dataset that are reported as software errors for human evaluation. In total, we choose 500 test cases and generate 2,000 questions. For each question, we ask three workers to give their judgment to reduce the variance. Therefore, we ask 150 workers to complete all questionnaires. It takes each worker 15-25 minutes to answer around 40 questions in a questionnaire,

⁵<https://www.qualtrics.com/>

⁶<https://prolific.co/>

and each worker is paid about 5 pounds per hour. The total cost is 300 pounds.

4.3 Baselines

We select diverse test case evaluation metrics as baselines from two categories: Neuron Coverage (NC) metrics and NLP-based metrics, which are summarized in Table 4.

NC-based. NC and its variants are commonly-used for evaluating test cases. Different from AEON, NC-based metrics mainly aim at the evaluation of a test set instead of a test case. In our experiments, we consider NC-based metrics in two ways. (1) For basic Neuron Coverage (NC) [57] and Neuron Boundary Coverage (NBC) [44], we calculate the NC scores of each generated test case. (2) For Top- k Neuron Coverage (TKNC) [44] and Bottom- k Neuron Coverage (BKNC) [79], they cannot be adapted to a single test case (e.g., TKNC produces the same coverage for single test cases), thus we compute the number of neurons covered by the generated test case but not by the original text. Intuitively, changes in texts may be reflected in neuron activation. Note that the comparison with NC-based metrics is not apples-to-apples because NC-based metrics mainly evaluate the quality of a test set. We include the comparison here for the completeness of our discussion.

NLP-based. Since the main reason behind false alarms is that the generated test cases cannot keep equivalent or similar semantic meaning with the original text, we include multiple semantic similarity metrics for the baselines of SemEval. Evaluating the semantic similarity of texts has long been a complex problem in NLP research. Previous metrics can be divided into corpus-based, knowledge-based, and DNN-based. DNN-based metrics outperform other methods and have served as a breakthrough in semantic similarity research [7]. We consider a corpus-based metric, BLEU, a knowledge-based metric, Meteor, and four DNN-based metrics, InferSent, SBERT, SimCSE, and BERTScore. For embedding models, namely InferSent, SBERT, and SimCSE, we report the semantic similarity based on cosine similarity because it is used by most of the researchers [18, 20, 86] and Euclidean distance yields similar results in all our experiments.

4.4 Evaluation Criteria

We compute three criteria: AP, AUC, and PCC, to discover the correlation between human judgment (Sec. 4.2) and the automatic evaluation metrics, including AEON. We treat the scoring systems as binary classification systems, the human judgment as ground truth, and draw their Precision-Recall curve (P-R curve) and Receiver Operating Characteristic curve (ROC curve) to calculate AP and AUC. P-R curve shows the trade-off between recall (i.e., true positive rate) and precision, while the ROC curve depicts the trade-off between true positive rate and false positive rate. AP and AUC represent the area under P-R curve and ROC curve, respectively. An excellent binary classification system tends to have high AP and AUC scores. Then, we check whether our scores are correlated with human judgment using PCC, the covariance of two variables divided by the product of their standard deviations, which can be written in the form of:

$$PCC(X, Y) = \frac{Cov(X, Y)}{\sqrt{Var(X) \cdot Var(Y)}}. \quad (2)$$




Figure 1: Venn Diagram of the proportion of each vulnerability category (better viewed in colored mode).

PCC is able to show how linearly correlated two variables are. Note that negative PCC value indicates that the two variables are negatively correlated.

5 EXPERIMENTAL RESULTS

5.1 RQ1: The Quality of Test Cases

We average the consistency, naturalness, and difficulty scores as the respective final scores. We use the label that most workers agree as the final human label. For each test case, we decide whether it is an *Inconsistent* case or an *Unnatural* case based on the consistency and naturalness scores. Finally, if the human label differs from the given label, the test case is considered a false alarm.

To show how severe the problem in NLP test case quality is, we draw the Venn diagram of the generated test cases based on the human annotation results. As is shown in Fig. 1, 44% of them change the label and thus are false alarms. In other words, there are only 1,435 cases triggering software errors in all 8,000 test cases. 57% of them are not natural enough, while 71% fail to preserve the semantic meaning. Only 15% of them have good language naturalness and preserve the semantic meaning, which are counted as high-quality test cases. Besides the statistical information, we have two more observations. First, though the majority of *Inconsistent* cases are false alarms, a few *Inconsistent* cases do not change the label. These test cases only account for 11%, and the rest can be categorized into the two issues. Second, bad naturalness can sometimes hurt semantic meaning, resulting in test cases that are both *Inconsistent* and *Unnatural*. This is because the unnatural part can eliminate some key information in texts and further change the semantics.

Answer to RQ1: The quality of NLP test cases cannot be guaranteed by existing testing techniques. 71% and 57% of test cases generated by existing NLP testing techniques are *Inconsistent* and *Unnatural*, respectively. 44% of test cases are false alarms, significantly degrading the effectiveness and efficiency of existing testing techniques.

5.2 RQ2: The Effectiveness of AEON

Since AEON is designed to evaluate the semantic similarity and language naturalness of NLP software test cases, we assess the two modules, *SemEval* and *SynEval*, to validate the effectiveness of our approach. We use default settings for all baselines, and we select $k = 192$ (one-fourth of neurons in each layer) for TKNC and BKNC. We set $\lambda_1 = 0.1$, $\lambda_2 = 0.2$ for *SemEval*, and $\phi = 0.6$ for *SynEval*.

5.2.1 SemEval. We draw P-R curves and ROC curves for the semantic scores calculated by *SemEval* as well as the other baselines mentioned in Sec. 4.3 and consistency from human evaluation. Then we compute AP and AUC scores, which are shown in Table 5. Our method achieves higher AP and AUC values averaged on all datasets and baselines, showing the strong ability of *SemEval* to filter out *Inconsistent* cases. The results also validate the effectiveness of *SemEval* on capturing subtle semantic changes. We calculate PCC between the semantic score and human-annotated consistency for each method. As shown in Table 5, our approach achieves about 0.92 PCC on average, which significantly outperforms all the baselines. This shows that the score of the *SemEval* aligns well with human evaluation.

NC-based metrics achieve decent performance and surpass many NLP-based metrics, especially on MNLI and QQP datasets, indicating that neuron activation patterns can reflect text semantic changes. As for NLP-based metrics, BLEU and Meteor perform the worst since they cannot handle highly overlapped texts. The BLEU and Meteor scores for text pairs are always high since most of the words in the original texts and the generated test cases are the same. DNN-based metrics cannot perform well because of three main reasons. (1) Word embeddings usually lack semantic information. For instance, the embeddings of [reject] and [accept] calculated by BERT [13] have high cosine similarity of 0.846, while such word substitution changes the correct label in sentiment analysis. Another example is that the cosine similarity between embeddings of [Tom] and [Jack] is 0.978, which hurts in natural language inference tasks. (2) Baselines that employ token matching [86] are prone to mistakenly matching multiple words to a single word. Consider the following example:

Illustrative Example 1

ORIGINAL TEXT	I do like the movie, though I did not watch it at cinema.
GENERATED TEXT	I do not like the movie, though I did not watch it at cinema.

The word [not] in the generated text can be matched to the second [not] in the original sentence, resulting in a high similarity score. (3) Models based on contrastive learning fail due to the lack of data with subtle differences in their training set. These models are mainly trained on natural language inference datasets [11, 18], which can hardly cover the cases where two sentences have few but vital differences.

Considering different datasets, in sentiment analysis tasks, AEON and all other baselines perform better on MR than on Yelp since the texts in MR in shorter and simpler. For natural language inference tasks, though MNLI is more complicated than SNLI, it is surprising

Table 5: AP, AUC, and PCC results that show how well the automatic metrics align with human-annotated consistency scores.

	Datasets	MR			Yelp			SNLI			MNLI			QQP		
		AP	AUC	PCC	AP	AUC	PCC	AP	AUC	PCC	AP	AUC	PCC	AP	AUC	PCC
NC-based	NC	0.76	0.55	0.73	0.29	0.36	-0.70	0.56	0.55	0.65	0.31	0.54	0.50	0.44	0.52	0.47
	NBC	0.78	0.66	0.85	0.51	0.62	0.85	0.44	0.47	0.23	0.35	0.56	0.90	0.38	0.56	0.78
	TKNC	0.88	0.73	0.88	0.56	0.64	0.88	0.58	0.57	0.65	0.47	0.66	0.87	0.67	0.76	0.91
	BKNC	0.88	0.73	0.86	0.55	0.63	0.85	0.58	0.58	0.67	0.48	0.67	0.87	0.64	0.75	0.88
NLP-based	BLEU	0.75	0.58	0.75	0.45	0.57	0.45	0.45	0.42	-0.36	0.34	0.50	0.44	0.44	0.54	0.44
	Meteor	0.74	0.54	0.73	0.43	0.57	0.41	0.51	0.50	-0.18	0.46	0.59	0.80	0.53	0.62	0.81
	InferSent	0.81	0.68	0.96	0.55	0.67	0.88	0.55	0.61	0.55	0.40	0.61	0.93	0.48	0.59	0.83
	SentBERT	0.85	0.71	0.92	0.53	0.63	0.83	0.65	0.65	0.66	0.46	0.63	0.74	0.52	0.59	0.76
	SimCSE	0.85	0.75	0.93	0.63	0.76	0.92	0.61	0.65	0.76	0.54	0.73	0.96	0.50	0.56	0.68
	BERTScore	0.75	0.61	0.91	0.44	0.53	0.37	0.59	0.60	0.20	0.34	0.58	0.73	0.46	0.59	0.67
	AEON (Ours)	0.92	0.82	0.93	0.75	0.84	0.91	0.66	0.70	0.85	0.58	0.75	0.98	0.56	0.63	0.90

Table 6: AP, AUC, and PCC results that show how well the automatic metrics align with human-annotated naturalness scores.

Datasets	MR			Yelp			SNLI			MNLI			QQP		
	AP	AUC	PCC	AP	AUC	PCC	AP	AUC	PCC	AP	AUC	PCC	AP	AUC	PCC
NC	0.66	0.56	0.53	0.64	0.41	-0.70	0.45	0.54	-0.09	0.56	0.51	0.27	0.56	0.52	0.03
NBC	0.72	0.57	0.32	0.70	0.47	0.51	0.32	0.52	0.66	0.43	0.36	-0.91	0.52	0.47	-0.61
TKNC	0.68	0.56	0.57	0.67	0.53	0.49	0.36	0.51	0.09	0.55	0.47	-0.33	0.65	0.57	0.51
BKNC	0.67	0.55	0.52	0.66	0.52	0.45	0.36	0.51	0.09	0.55	0.47	-0.39	0.64	0.55	0.46
	AEON (Ours)	0.87	0.59	0.66	0.77	0.58	0.52	0.63	0.75	0.54	0.65	0.98	0.73	0.68	0.87

to observe that SNLI has lower PCC scores than MNLI. There are negative PCC scores when using BLEU and Meteor, indicating the negative correlation between the baselines and human evaluation. We think the reason behind this is that the complexity of MNLI lies in the diversity of contexts, and changes in contexts typically will not change the corresponding labels.

5.2.2 SynEval. To evaluate the performance of SynEval, we draw P-R curves and ROC curves and compute AP and AUC, treating SynEval as a binary classifier to recognize *Unnatural* cases. We also calculate PCC between SynEval and naturalness score from human evaluation, which is included in Table 6. We can observe that though NC-based metrics have an excellent performance on detecting *Inconsistent* test cases, they fall short of measuring language naturalness. The performance varies significantly on different datasets. The PCC scores of NBC, TKNC, and BKNC show a negative correlation on MNLI, while positively correlated on other datasets. We infer the reason behind this may be that the models make the decision based on the appearance of certain words or phrases, ignoring whether the input texts have good language naturalness. In addition to using BERT [13] in SynEval that is presented

in Table 6, we also try other language models including RoBERTa [43] and ALBERT [37], among which BERT achieves the highest AUC and AP values, averaged on all datasets. Note that traditional grammar checkers are not suitable for this task because they do not provide quantitative results, and they cannot reveal the error-free yet strange sentences that people rarely write.

The impact of the hyperparameters. If we set the proportion of the minimal semantic score from 0 to 1, we can observe that the performance increases at first, then remains stable at the same level, and finally drop when it gets close to 1. We balance this trade-off using a grid search for lambdas and phis. These parameters can be generalized to other datasets and NLP tasks since we adopt the same parameters and consistently achieve good performance for all selected datasets in our experiments. We also test different patch lengths l for extracting $[i - l : i + l]$. In particular, $l = 1$ does not work well because most NLP models use BPE (Byte Pair Encoding) [65] for tokenization, which may divide a word into smaller tokens, making it extract only part of a word. Long patches (e.g., $l \geq 5$) suffer from the same problem as average scores, *i.e.*, the impact of mutation vanishes after averaging. In our experiments, $l = 3$ and

Table 7: The quality of test cases without and with AEON.

Data Source	Consistency	Naturalness	False Alarms
w/o AEON	2.627	2.916	0.440
w/ AEON	3.357	3.305	0.262
Improvement	↑ 27.8%	↑ 13.3%	↓ 40.6%

$l = 4$ lead to similar results. To reduce computation cost, we select $l = 4$ for this parameter.

Answer to RQ2: AEON, which consists of SemEval and SynEval, is effective in terms of detecting test cases that change the label (false alarms) and test cases that are unnatural. AEON outperforms all baselines in average on all datasets.

5.3 RQ3: Test Case Selection Using AEON

This paper aims to propose a metric that facilitates NLP software testing by evaluating the quality of test cases. In this section, we utilize AEON to filter out low-quality test cases. We conduct experiments to verify whether the test cases selected by AEON enjoy better semantic consistency and language naturalness. Specifically, AEON can be utilized to filter out *Inconsistency* and *Unnatural* test cases to improve the quality of test cases in average. For SemEval, we set different thresholds for different tasks. We choose multiple thresholds for semantic similarity score because whether the label will change depends on the given task. Consider this pair of texts (original and generated) which is inconsistent in semantics:

Illustrative Example 2

ORIGINAL TEXT	I watched the movie at home , it was nice.
GENERATED TEXT	I watched the movie at cinema , it was nice.

If they are in a sentiment analysis dataset, the label remains unchanged (*i.e.*, positive). However, if they appear in a natural language inference dataset as premises, and the hypothesis is “I went out for the movie”, the label changes from contradiction to entailment. Therefore, to best filter out those false alarms, we set thresholds as 0.87, 0.90, and 0.91 for sentiment analysis, natural language inference, and semantic equivalence, respectively. The thresholds are computed with a balance between true positive rate and false positive rate. From the thresholds, we can see that the three tasks need more semantic similarity increasingly to ensure the preservation of labels, which aligns with the characteristics of the datasets. For language naturalness, we set the threshold as 0.21.

We generate 500 test cases that are reported to trigger some software errors, including various datasets and testing techniques mentioned in Table 3 and Table 2 respectively. Then we check whether human evaluation has improved before and after applying AEON to select high-quality test cases. The results are shown in Table 7. The average consistency and naturalness scores of the 500 test cases are 2.627 and 2.916, which are below 3 (considered as *Inconsistency* and *Unnatural* test cases) in average. The false alarm rate is 0.44. After selecting test cases whose SemEval and SynEval

scores are above the thresholds with the help of AEON, the quality of test cases is significantly enhanced. The scores increase to 3.357 and 3.305, considered high-quality test cases on average. The consistency score improves by 27.8%, and the naturalness score improves by 13.3%. The false alarm rate is 26.2%, showing a significant improvement of 40.6%. The results demonstrate the effectiveness of AEON on high-quality test case selection.

5.3.1 Case Study. We choose one of the generated test cases as an example to illustrate the performance of our SemEval and SynEval compared to other baselines.

Case Study 2	
TASK/DATASET	SA/MR
TECHNIQUE	PSO
ORIGINAL TEXT	The result is a powerful , naturally dramatic piece of low-budget filmmaking.
GENERATED TEXT	The result is a terrible , naturally dramatic piece of low-budget filmmaking.

AEON achieves a semantic score of 0.58 and a syntactic score of 0.22. From the semantic side, the sentiment of the original example is positive. However, the sentiment of the generated test case is negative because [terrible] is a negative adjective. This test case is not only *Inconsistent* but also a false alarm since the label of this example is changed. Therefore, an excellent semantic metric should give this test case a low score to filter it out. Our method, SemEval, gives the text pair a score of 0.58, which is far below the threshold of 0.87, indicating that the test cases cannot preserve the semantics and should be filtered out. AEON works effectively because our design to consider patch similarity identifies that the substitution ([powerful]→[terrible]) dramatically changes the semantic meaning. From the syntactic side, the generated test case reads smoothly without difficulty comprehending its meaning, suggesting that it has good language naturalness. The case obtains a score of 0.22 given by SynEval, which is above the threshold of 0.21 and indicates that the test case is not an *Unnatural* case. All in all, our method outperforms other baselines both in semantic and syntactic perspectives on this example.

Answer to RQ3: AEON can effectively filter out low-quality test cases. The remaining test cases enjoy better semantic consistency and language naturalness and, most importantly, a lower false alarm rate. Thus, AEON can facilitate NLP software testing by selecting high-quality test cases, saving developers’ time.

5.4 RQ4: Improving NLP Software with AEON

Although NLP software testing is a promising research direction, it incurs an important yet unavoidable question: can the test cases be utilized to improve NLP software? To further show how high-quality test cases selected by AEON can help in improving NLP software, we add test cases that the model misclassifies to the training set and conduct model re-training. Accuracy is verified on the test set of the given task, while robustness is evaluated using the success rate of adversarial attacks. In this section, we




Figure 2: The left figure shows the accuracy of the test set (the higher, the better), and the right figure shows the attack success rate (the lower, the better). Blue lines are the results from models trained on randomly selected test cases, while red lines are the results from models trained on test cases selected by AEON. The horizontal axis is the ratio of the test cases in the whole training set.

run experiments to verify whether AEON can further improve the robustness and accuracy in model re-training.

We focus on fine-tuning a pre-trained BERT on MR dataset for sentiment analysis task for simplicity and better reproducibility. We first generate the test cases using the testing techniques mentioned in Table 2 using the entire training set of MR for seeds (original texts). Then we consider two settings: (1) randomly select as many test cases as 5% to 25% cases in the training set to train the model; (2) rank all the test cases with AEON in descending order and select the same size as (1) to train the model. After training, which takes five epochs to reach convergence, we evaluate the models’ accuracy using the MR test set and robustness using an adversarial attack method, PWWS [62].

As shown in Fig. 2(a), models trained with ranked test cases outperform the models with randomly selected test cases in terms of accuracy. In addition, Fig. 2(b) shows that models trained with ranked test cases are more robust (the lower the attack success rate is, the more robust the model is). At the beginning of each figure, we observe improvement in accuracy and robustness on both lines. This is because test cases add generalization ability to models. As we use more additional data in training, the noisy data issue surfaces (*i.e.*, false alarms and low-quality test cases) and starts to harm the model accuracy and robustness. We can observe that no matter what ratio of test cases we use, the models trained with AEON ranked data achieve higher accuracy and robustness thanks to the high-quality test cases it selects. In contrast, the accuracy and robustness drop quickly in models trained with randomly selected test cases. In conclusion, adding low-quality test cases can easily hurt both accuracy and robustness, while adding a reasonable amount of high-quality test cases selected by AEON leads to model improvement in terms of both accuracy and robustness.

Answer to RQ4: Equipped with AEON, test cases generated by existing testing techniques can be prioritized. The prioritized test cases can be further leveraged into NLP model training, resulting in higher accuracy and stronger robustness. AEON can also save time and resources by filtering out false alarms and test cases of low quality.

5.5 Discussion

The validity of out user study. To alleviate workers’ negligence in the annotation process, we double-check the results in two cases. (1) The workers feel it challenging to select one specific label. Specifically, we check test cases with high difficulty scores (above 3.5), or where annotators returned three different labels (*e.g.*, three workers give the label of contradiction, entailment, and neutral in natural language inference tasks). (2) The human label is different from the original label (*i.e.*, label changes), while the human-annotated consistency score is still high, which is counter-intuitive since if the label is changed, then the semantic meaning must have changed. We find out 21 cases for case (1) and 66 examples for case (2) and hire two annotators with two years of experience in NLP research to re-evaluate them. Afterward, we compute the Kappa score to check inter-rater reliability. The Fleiss’ Kappa of the classification task is 0.76 averaged on all datasets, implying substantial agreement among the annotators. We do not apply Cohen’s Kappa since we have three annotators for each case.

The importance of *Unnatural* test cases. We evaluate the naturalness of test cases and filter out *Unnatural* ones because they are rarely seen in real-world scenarios. Though previous work intentionally add spelling and grammar errors to data to improve the robustness of the NLP models [34], we noticed annotators mentioned that they could not understand the weird grammar or expression of some generated texts. In addition, unnaturalness can hide important semantic information, which changes the semantics and further renders the possibility of generating false alarms. Consider the following example:

Case Study 3	
TASK/DATASET	SA/MR
TECHNIQUE	PSO
ORIGINAL TEXT	... is a relationship that is worthy of our respect.
GENERATED TEXT	... is a relationship that is costly of our respect.

In this example, the unnaturalness changes the semantics of the sentence and further reverses from positive sentiment to negative sentiment while making readers confused. Thus, we think *Unnatural* issues are important, and they can also degrade the quality of NLP software test cases.

Limitations of our proposed approaches. Our proposed approaches consist of two parts, SemEval for semantic similarity and SynEval for language naturalness. To apply SemEval, we require both the generated test cases and the original texts at the same time, thus limiting SemEval to testing techniques of metamorphic testing, where the metamorphic relation assumes the transformation will change or will not change the semantics of the original text. Under the assumption that the transformation will not change the semantics, we filter out those with lower SemEval scores, while under the assumption that the transformation will change the semantics, we filter out those with higher SemEval scores. SynEval can be applied to any textual test case.

6 RELATED WORK

6.1 Testing AI Software

With the improvement of Artificial Intelligence (AI) models, companies tend to deploy AI in real-world applications like autonomous driving and neural machine translation [59]. However, AI software inherits the deficiencies of AI models that they are prone to erroneous behavior given particular inputs [2, 4, 5, 15, 21, 76–78]. A line of research has been conducted to test AI software systems to address this problem. Specifically, they test software based on convolutional neural network and feed forward neural network [17, 27, 36, 45, 68, 75, 83, 85], software based on recurrent neuron network [16], and software based on general DNN models [28, 82, 84]. Other researchers focus on testing deep learning libraries [58], assist the debugging process [46], and detect adversarial examples online [47, 67, 72, 80]. Unlike these papers that primarily focus on CV software, this paper focuses on NLP software.

6.2 Testing NLP Software

DNNs have boosted the performance of many NLP fields such as code analysis [1, 30, 60] and machine translation [70]. In recent years, researchers have proposed a variety of metamorphic testing techniques for NLP software [9, 10, 24–26, 63, 66]. In addition to metamorphic testing techniques, another line of research for finding NLP software errors [22, 32, 33, 39, 40, 81] is inspired by the adversarial attack concept in the CV field. Our work focuses on automatic quality evaluation of test cases generated by these testing techniques. Thus, we believe AEON complements with existing work.

6.3 Testing Criteria

Testing criterion, such as code coverage, has been widely utilized to measure how good a test suite is in traditional software (e.g., compilers). Inspired by code coverage in traditional software, DeepXplore [57] introduces the concept of neuron coverage for AI software: the percentage of neurons activated by the test cases. In recent years, researchers have proposed diverse variants of neuron coverage as testing criteria focusing on different activation magnitudes [44, 79]. Researchers also develop neuron coverage specially designed for recurrent neuron networks [16, 23, 29] to adopt the properties of sequence inputs. Different from neuron coverage metrics, which often act as test adequacy criteria of the test suite, our approach focuses on the quality evaluation of every test case. Thus, we think AEON can complement with existing testing criteria and contribute to research on NLP software testing.

7 CONCLUSION

This paper is the first to explore the quality of test cases generated by NLP software testing techniques. In an evaluation study, we surprisingly observe that 44% of the generated NLP test cases are of low quality, incurring *Inconsistent* or/and *Unnatural* issues. Thus, instead of improving NLP software, utilization of these test cases in model training could even degrade its accuracy and robustness. To this end, we introduce AEON, a novel, effective approach for automatic quality evaluation of NLP software testing cases. Given

an original text and a generated test case, AEON returns two scores regarding similarity consistency and language naturalness. Our evaluation and user study show that AEON’s scores align well with the quality scores returned by humans. In particular, AEON achieves 69.4% and 68.5% AP on detecting *Inconsistent* and *Unnatural* test cases generated by four SOTA testing techniques on five widely-used datasets. We can conduct test selection or prioritization according to the scores returned by AEON. In our evaluation, models trained on the test cases selected by AEON consistently achieve better accuracy and robustness than models trained on randomly selected test cases. We believe that this work is the important first step toward systematic quality evaluation of NLP software test cases, which can further enhance the effectiveness of testing techniques for NLP software and complement existing testing criteria. We leave the work for automatically fixing *Inconsistent* and *Unnatural* test cases for future exploration.

ACKNOWLEDGMENTS

We thank gratefully the constructive discussion with Shuqing Li, and the helpful feedback from members of ARISE Lab. The work described in this paper was supported by the key program of fundamental research from the Shenzhen Science and Technology Innovation Commission (No. JCYJ20200109113403826), the Research Grants Council of the Hong Kong Special Administrative Region, China (CUHK 14210920 of the General Research Fund), and the National Natural Science Foundation of China (No. 62102340).

REFERENCES

- [1] Uri Alon, Meital Zilberstein, Omer Levy, and Eran Yahav. 2019. code2vec: learning distributed representations of code. *Proc. ACM Program. Lang.* 3, POPL (2019), 40:1–40:29.
- [2] Anish Athalye, Nicholas Carlini, and David A. Wagner. 2018. Obfuscated Gradients Give a False Sense of Security: Circumventing Defenses to Adversarial Examples. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*, Vol. 80. 274–283.
- [3] Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. A large annotated corpus for learning natural language inference. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 632–642.
- [4] Nicholas Carlini, Pratyush Mishra, Tavish Vaidya, Yuankai Zhang, Micah Sherr, Clay Shields, David A. Wagner, and Wenchao Zhou. 2016. Hidden Voice Commands. In *Proceedings of the 25th USENIX Security Symposium (USENIX Security)*. 513–530.
- [5] Nicholas Carlini and David A. Wagner. 2017. Towards Evaluating the Robustness of Neural Networks. In *2017 IEEE Symposium on Security and Privacy (SP)*. 39–57.
- [6] Asli Celikyilmaz, Elizabeth Clark, and Jianfeng Gao. 2020. Evaluation of Text Generation: A Survey. *CoRR* abs/2006.14799 (2020). arXiv:2006.14799
- [7] Dhivya Chandrasekaran and Vijay Mago. 2021. Evolution of Semantic Similarity - A Survey. *ACM Comput. Surv.* 54, 2 (2021), 41:1–41:37.
- [8] Mei-Hwa Chen, Michael R. Lyu, and W. Eric Wong. 2001. Effect of code coverage on software reliability measurement. *IEEE Trans. Reliab.* 50, 2 (2001), 165–170.
- [9] Songqiang Chen, Shuo Jin, and Xiaoyuan Xie. 2021. Testing Your Question Answering Software via Asking Recursively. In *Proceedings of the 36th IEEE/ACM International Conference on Automated Software Engineering (ASE)*. 104–116.
- [10] Songqiang Chen, Shuo Jin, and Xiaoyuan Xie. 2021. Validation on machine reading comprehension software without annotated labels: a property-based method. In *Proceedings of the 29th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering (ESEC/FSE)*. 590–602.
- [11] Alexis Conneau, Douwe Kiela, Holger Schwenk, Loïc Barrault, and Antoine Bordes. 2017. Supervised Learning of Universal Sentence Representations from Natural Language Inference Data. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 670–680.
- [12] Michael J. Denkowski and Alon Lavie. 2014. Meteor Universal: Language Specific Translation Evaluation for Any Target Language. In *Proceedings of the Ninth Workshop on Statistical Machine Translation (WMT@ACL)*. 376–380.

- [13] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT), Volume 1 (Long and Short Papers)*. 4171–4186.
- [14] Zhendong Dong, Qiang Dong, and Changling Hao. 2010. HowNet and Its Computation of Meaning. In *Proceedings of the 23rd International Conference on Computational Linguistics (COLING), Demonstrations Volume*. 53–56.
- [15] Tianyu Du, Shouling Ji, Jinfeng Li, Qinchen Gu, Ting Wang, and Raheem Beyah. 2020. SirenAttack: Generating Adversarial Audio for End-to-End Acoustic Systems. In *Proceedings of the 15th ACM Asia Conference on Computer and Communications Security (ASIA CCS)*. 357–369.
- [16] Xiaoning Du, Xiaofei Xie, Yi Li, Lei Ma, Yang Liu, and Jianjun Zhao. 2019. DeepStellar: model-based quantitative analysis of stateful deep learning systems. In *Proceedings of the ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering (ESEC/FSE)*. 477–487.
- [17] Alessio Gambi, Marc Müller, and Gordon Fraser. 2019. Automatically testing self-driving cars with search-based procedural content generation. In *Proceedings of the 28th ACM SIGSOFT International Symposium on Software Testing and Analysis (ISSTA)*. 318–328.
- [18] Tianyu Gao, Kingcheng Yao, and Danqi Chen. 2021. SimCSE: Simple Contrastive Learning of Sentence Embeddings. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 6894–6910.
- [19] Yifan Gao, Chien-Sheng Wu, Jingjing Li, Shafiq R. Joty, Steven C. H. Hoi, Caiming Xiong, Irwin King, and Michael R. Lyu. 2020. Discern: Discourse-Aware Entailment Reasoning Network for Conversational Machine Reading. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 2439–2449.
- [20] Siddhant Garg and Goutham Ramakrishnan. 2020. BAE: BERT-based Adversarial Examples for Text Classification. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 6174–6181.
- [21] Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. 2015. Explaining and Harnessing Adversarial Examples. In *Proceedings of the 3rd International Conference on Learning Representations (ICLR)*.
- [22] Junliang Guo, Zhirui Zhang, Linlin Zhang, Linli Xu, Boxing Chen, Enhong Chen, and Weihua Luo. 2021. Towards Variable-Length Textual Adversarial Attacks. *CoRR abs/2104.08139* (2021).
- [23] Jianmin Guo, Yue Zhao, Xueying Han, Yu Jiang, and Jianguang Sun. 2019. RNN-Test: Adversarial Testing Framework for Recurrent Neural Network Systems. *CoRR abs/1911.06155* (2019).
- [24] Shashij Gupta, Pinjia He, Clara Meister, and Zhendong Su. 2020. Machine translation testing via pathological invariance. In *Proceedings of the 28th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering (ESEC/FSE)*. 863–875.
- [25] Pinjia He, Clara Meister, and Zhendong Su. 2020. Structure-invariant testing for machine translation. In *Proceedings of the 42nd International Conference on Software Engineering (ICSE)*. 961–973.
- [26] Pinjia He, Clara Meister, and Zhendong Su. 2021. Testing Machine Translation via Referential Transparency. In *Proceedings of the 43rd IEEE/ACM International Conference on Software Engineering (ICSE)*. 410–422.
- [27] Jens Henriksson, Christian Berger, Markus Borg, Lars Tornberg, Cristofer Englund, Sankar Raman Sathiyamoorthy, and Stig Ursing. 2019. Towards Structured Evaluation of Deep Neural Network Supervisors. In *Proceedings of the IEEE International Conference On Artificial Intelligence Testing (AITest)*. 27–34.
- [28] Qiang Hu, Lei Ma, Xiaofei Xie, Bing Yu, Yang Liu, and Jianjun Zhao. 2019. DeepMutation++: A Mutation Testing Framework for Deep Learning Systems. In *Proceedings of the 34th IEEE/ACM International Conference on Automated Software Engineering (ASE)*. 1158–1161.
- [29] Wei Huang, Youcheng Sun, Xingyu Zhao, James Sharp, Wenjie Ruan, Jie Meng, and Xiaowei Huang. 2019. Coverage Guided Testing for Recurrent Neural Networks. *CoRR abs/1911.01952* (2019).
- [30] Srinivasan Iyer, Ioannis Konstas, Alvin Cheung, and Luke Zettlemoyer. 2016. Summarizing Source Code using a Neural Attention Model. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (ACL), Volume 1: Long Papers*.
- [31] Fred Jelinek, Robert L. Mercer, Lalit R. Bahl, and James K. Baker. 1977. Perplexity — a measure of the difficulty of speech recognition tasks. *The Journal of the Acoustical Society of America* 62, S1 (1977), S63–S63.
- [32] Robin Jia, Aditi Raghunathan, Kerem Göksel, and Percy Liang. 2019. Certified Robustness to Adversarial Word Substitutions. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. 4127–4140.
- [33] Di Jin, Zhijing Jin, Joey Tianyi Zhou, and Peter Szolovits. 2020. Is BERT Really Robust? A Strong Baseline for Natural Language Attack on Text Classification and Entailment. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI), The Thirty-Second Innovative Applications of Artificial Intelligence Conference (IAAI), The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI)*. 8018–8025.
- [34] Erik Jones, Robin Jia, Aditi Raghunathan, and Percy Liang. 2020. Robust Encodings: A Framework for Combating Adversarial Typos. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL)*. 2752–2765.
- [35] James Kennedy and Russell Eberhart. 1995. Particle swarm optimization. In *Proceedings of International Conference on Neural Networks (ICNN)*. 1942–1948.
- [36] Jinhan Kim, Robert Feldt, and Shin Yoo. 2019. Guiding deep learning system testing using surprise adequacy. In *Proceedings of the 41st International Conference on Software Engineering (ICSE)*. 1039–1049.
- [37] Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2020. ALBERT: A Lite BERT for Self-supervised Learning of Language Representations. In *Proceedings of the 8th International Conference on Learning Representations (ICLR)*.
- [38] Vladimir I. Levenshtein. 1966. Binary codes capable of correcting deletions, insertions, and reversals. *Cybernetics and Control Theory* 10, 8 (1966), 707–710.
- [39] Dianqi Li, Yizhe Zhang, Hao Peng, Liqun Chen, Chris Brockett, Ming-Ting Sun, and Bill Dolan. 2021. Contextualized Perturbation for Textual Adversarial Attack. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*. 5053–5069.
- [40] Jinfeng Li, Shouling Ji, Tianyu Du, Bo Li, and Ting Wang. 2019. TextBugger: Generating Adversarial Text Against Real-world Applications. In *26th Annual Network and Distributed System Security Symposium (NDSS)*.
- [41] Jingjing Li, Zichao Li, Lili Mou, Xin Jiang, Michael R. Lyu, and Irwin King. 2020. Unsupervised Text Generation by Learning from Search. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems (NeurIPS 2020)*.
- [42] Bin Liang, Hongcheng Li, Miaoqiang Su, Pan Bian, Xirong Li, and Wenchang Shi. 2018. Deep Text Classification Can be Fooled. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence (IJCAI)*. 4208–4215.
- [43] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. RoBERTa: A Robustly Optimized BERT Pretraining Approach. *CoRR abs/1907.11692* (2019). arXiv:1907.11692
- [44] Lei Ma, Felix Juefei-Xu, Fuyuan Zhang, Jiyuan Sun, Minhui Xue, Bo Li, Chunyang Chen, Ting Su, Li Li, Yang Liu, Jianjun Zhao, and Yadong Wang. 2018. DeepGauge: multi-granularity testing criteria for deep learning systems. In *Proceedings of the 33rd ACM/IEEE International Conference on Automated Software Engineering (ASE)*. 120–131.
- [45] Lei Ma, Fuyuan Zhang, Jiyuan Sun, Minhui Xue, Bo Li, Felix Juefei-Xu, Chao Xie, Li Li, Yang Liu, Jianjun Zhao, and Yadong Wang. 2018. DeepMutation: Mutation Testing of Deep Learning Systems. In *Proceedings of the 29th IEEE International Symposium on Software Reliability Engineering (ISSRE)*. 100–111.
- [46] Shiqing Ma, Yingqi Liu, Wen-Chuan Lee, Xiangyu Zhang, and Ananth Grama. 2018. MODE: automated neural network model debugging via state differential analysis and input selection. In *Proceedings of the 2018 ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering (ESEC/FSE)*. 175–186.
- [47] Shiqing Ma, Yingqi Liu, Guanhong Tao, Wen-Chuan Lee, and Xiangyu Zhang. 2019. NIC: Detecting Adversarial Samples with Neural Network Invariant Checking. In *Proceedings of the 26th Annual Network and Distributed System Security Symposium (NDSS)*.
- [48] Edoardo Maggio. 2018. Apple says that 500 million customers use Siri. Online, Available: <https://www.businessinsider.com/apple-says-siri-has-500-million-users-2018-1> (2018).
- [49] Paul Michel, Xian Li, Graham Neubig, and Juan Miguel Pino. 2019. On Evaluation of Adversarial Perturbations for Sequence-to-Sequence Models. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT), Volume 1 (Long and Short Papers)*. 3103–3114.
- [50] George A. Miller. 1995. WordNet: A Lexical Database for English. *Commun. ACM* 38, 11 (1995), 39–41.
- [51] John X. Morris, Eli Lifland, Jack Lanchantin, Yangfeng Ji, and Yanjun Qi. 2020. Reevaluating Adversarial Examples in Natural Language. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings (EMNLP) (Findings of ACL, Vol. EMNLP 2020)*. 3829–3839.
- [52] John X. Morris, Eli Lifland, Jin Yong Yoo, Jake Grigsby, Di Jin, and Yanjun Qi. 2020. TextAttack: A Framework for Adversarial Attacks, Data Augmentation, and Adversarial Training in NLP. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations, (EMNLP - Demos)*. 119–126.
- [53] Arika Okrent. 2016. 9 Little Translation Mistakes That Caused Big Problems. Online, Available: <https://www.mentalfloss.com/article/48795/9-little-translation-mistakes-caused-big-problems> (2016).
- [54] Ahuy Ong. 2017. Facebook apologizes after wrong translation sees Palestinian man arrested for posting ‘good morning’. Online, Available: <https://www.theverge.com/us-world/2017/10/24/16533496/facebook-apology-wrong-translation-palestinian-arrested-post-good-morning> (2017).

- [55] Bo Pang and Lillian Lee. 2005. Seeing Stars: Exploiting Class Relationships for Sentiment Categorization with Respect to Rating Scales. In *Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics (ACL)*. 115–124.
- [56] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a Method for Automatic Evaluation of Machine Translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics (ACL)*. 311–318.
- [57] Kexin Pei, Yinzi Cao, Junfeng Yang, and Suman Jana. 2017. DeepXplore: Automated Whitebox Testing of Deep Learning Systems. In *Proceedings of the 26th Symposium on Operating Systems Principles (SOSP)*. 1–18.
- [58] Hung Viet Pham, Thibaud Lutellier, Weizhen Qi, and Lin Tan. 2019. CRADLE: cross-backend validation to detect and localize bugs in deep learning libraries. In *Proceedings of the 41st International Conference on Software Engineering (ICSE)*. 1027–1038.
- [59] Samira Pouyanfar, Saad Sadiq, Yilin Yan, Haiman Tian, Yudong Tao, Maria E. Presa Reyes, Mei-Ling Shyu, Shu-Ching Chen, and S. S. Iyengar. 2019. A Survey on Deep Learning: Algorithms, Techniques, and Applications. *ACM Comput. Surv.* 51, 5 (2019), 92:1–92:36.
- [60] Michael Pradel and Koushik Sen. 2018. DeepBugs: a learning approach to name-based bug detection. *Proc. ACM Program. Lang.* 2, OOPSLA (2018), 147:1–147:25.
- [61] Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. 3980–3990.
- [62] Shuhuai Ren, Yihe Deng, Kun He, and Wanxiang Che. 2019. Generating Natural Language Adversarial Examples through Probability Weighted Word Saliency. In *Proceedings of the 57th Conference of the Association for Computational Linguistics (ACL), Volume 1: Long Papers*. 1085–1097.
- [63] Marco Túlio Ribeiro, Tongshuang Wu, Carlos Guestrin, and Sameer Singh. 2020. Beyond Accuracy: Behavioral Testing of NLP Models with CheckList. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL)*. 4902–4912.
- [64] SafeAtLast. 2021. Incredible Amazon Alexa Statistics You Need to Know in 2021. Online, Available: <https://safeatlast.co/blog/amazon-alexa-statistics/gref> (2021).
- [65] Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Neural Machine Translation of Rare Words with Subword Units. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (ACL), Volume 1: Long Papers*.
- [66] Zeyu Sun, Jie M. Zhang, Mark Harman, Mike Papadakis, and Lu Zhang. 2020. Automatic testing and improvement of machine translation. In *Proceedings of the 42nd International Conference on Software Engineering, (ICSE)*. 974–985.
- [67] Guanhong Tao, Shiqing Ma, Yingqi Liu, and Xiangyu Zhang. 2018. Attacks Meet Interpretability: Attribute-steered Detection of Adversarial Samples. In *Proceedings of Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018 (NeurIPS)*. 7728–7739.
- [68] Yuchi Tian, Kexin Pei, Suman Jana, and Baishakhi Ray. 2018. DeepTest: automated testing of deep-neural-network-driven autonomous cars. In *Proceedings of the 40th International Conference on Software Engineering (ICSE)*. 303–314.
- [69] Barak Turovsky. 2016. Ten years of Google Translate. Online, Available: <https://blog.google/products/translate/ten-years-of-google-translate> (2016).
- [70] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is All you Need. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017 (NIPS)*. 5998–6008.
- [71] Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2019. GLUE: A Multi-Task Benchmark and Analysis Platform for Natural Language Understanding. In *Proceedings of the 7th International Conference on Learning Representations (ICLR)*.
- [72] Jingyi Wang, Guoliang Dong, Jun Sun, Xinyu Wang, and Peixin Zhang. 2019. Adversarial sample detection for deep neural network through model mutation testing. In *Proceedings of the 41st International Conference on Software Engineering (ICSE)*. 1245–1256.
- [73] Wenxuan Wang, Wenxiang Jiao, Yongchang Hao, Xing Wang, Shuming Shi, Zhaopeng Tu, and Michael R. Lyu. 2022. Understanding and Improving Sequence-to-Sequence Pretraining for Neural Machine Translation. *CoRR* abs/2203.08442 (2022).
- [74] Adina Williams, Nikita Nangia, and Samuel R. Bowman. 2018. A Broad-Coverage Challenge Corpus for Sentence Understanding through Inference. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT), Volume 1 (Long Papers)*. 1112–1122.
- [75] Weibin Wu, Yuxin Su, Xixian Chen, Shenglin Zhao, Irwin King, Michael R. Lyu, and Yu-Wing Tai. 2020. Towards Global Explanations of Convolutional Neural Networks With Concept Attribution. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 8649–8658.
- [76] Weibin Wu, Yuxin Su, Michael R. Lyu, and Irwin King. 2021. Improving the Transferability of Adversarial Samples With Adversarial Transformations. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 9024–9033.
- [77] Weibin Wu, Hui Xu, Sanqiang Zhong, Michael R. Lyu, and Irwin King. 2019. Deep Validation: Toward Detecting Real-World Corner Cases for Deep Neural Networks. In *49th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*. 125–137.
- [78] Chong Xiang, Charles R. Qi, and Bo Li. 2019. Generating 3D Adversarial Point Clouds. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 9136–9144.
- [79] Xiaofei Xie, Lei Ma, Felix Juefei-Xu, Minhui Xue, Hongxu Chen, Yang Liu, Jianjun Zhao, Bo Li, Jianxiang Yin, and Simon See. 2019. DeepHunter: a coverage-guided fuzz testing framework for deep neural networks. In *Proceedings of the 28th ACM SIGSOFT International Symposium on Software Testing and Analysis (ISSTA)*. 146–157.
- [80] Weilin Xu, David Evans, and Yanjun Qi. 2018. Feature Squeezing: Detecting Adversarial Examples in Deep Neural Networks. In *Proceedings of the 25th Annual Network and Distributed System Security Symposium (NDSS)*.
- [81] Yuan Zang, Fanchao Qi, Chenghao Yang, Zhiyuan Liu, Meng Zhang, Qun Liu, and Maosong Sun. 2020. Word-level Textual Adversarial Attacking as Combinatorial Optimization. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL)*. 6066–6080.
- [82] Fuyuan Zhang, Sankalan Pal Chowdhury, and Maria Christakis. 2020. DeepSearch: a simple and effective blackbox attack for deep neural networks. In *Proceedings of the 28th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering (ESEC/FSE)*. 800–812.
- [83] Jianping Zhang, Weibin Wu, Jen-tse Huang, Yizhan Huang, Wenxuan Wang, Yuxin Su, and Michael R. Lyu. 2022. Improving Adversarial Transferability via Neuron Attribution-Based Attacks. *CoRR* abs/2204.00008 (2022).
- [84] Jie M. Zhang, Mark Harman, Lei Ma, and Yang Liu. 2020. Machine Learning Testing: Survey, Landscapes and Horizons. *IEEE Trans. Software Eng.* (2020).
- [85] Mengshi Zhang, Yuqun Zhang, Lingming Zhang, Cong Liu, and Sarfraz Khurshid. 2018. DeepRoad: GAN-based metamorphic testing and input validation framework for autonomous driving systems. In *Proceedings of the 33rd ACM/IEEE International Conference on Automated Software Engineering (ASE)*. 132–142.
- [86] Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. BERTScore: Evaluating Text Generation with BERT. In *Proceedings of the 8th International Conference on Learning Representations (ICLR)*.
- [87] Wei Emma Zhang, Quan Z. Sheng, Ahoud Abdulrahman F. Alhazmi, and Chenliang Li. 2020. Adversarial Attacks on Deep-learning Models in Natural Language Processing: A Survey. *ACM Trans. Intell. Syst. Technol.* 11, 3 (2020), 24:1–24:41.
- [88] Xiang Zhang, Junbo Jake Zhao, and Yann LeCun. 2015. Character-level Convolutional Networks for Text Classification. In *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015 (NIPS)*. 649–657.