

Less is More: Understanding Word-level Textual Adversarial Attack via n -gram Frequency Descend

Ning Lu^{1,2} Shengcai Liu^{3,5,*} Zhirui Zhang⁴ Qi Wang⁵ Haifeng Liu⁶ Ke Tang¹

¹Guangdong Key Laboratory of Brain-Inspired Intelligent Computation,

Department of Computer Science and Engineering, Southern University of Science and Technology

²Department of Computer Science and Engineering, Hong Kong University of Science and Technology

³Centre for Frontier AI Research, Agency for Science, Technology and Research ⁴Tencent AI Lab

⁵Department of Computer Science and Engineering, Southern University of Science and Technology

⁶OPPO Mobile Telecommunications Corp.

Abstract—Word-level textual adversarial attacks have demonstrated notable efficacy in misleading Natural Language Processing (NLP) models. Despite their success, the underlying reasons for their effectiveness and the fundamental characteristics of adversarial examples (AEs) remain obscure. This work aims to interpret word-level attacks by examining their n -gram frequency patterns. Our comprehensive experiments reveal that in approximately 90% of cases, word-level attacks lead to the generation of examples where the frequency of n -grams decreases, a tendency we term as the n -gram Frequency Descend (n -FD). This finding suggests a straightforward strategy to enhance model robustness: training models using examples with n -FD. To examine the feasibility of this strategy, we employed the n -gram frequency information, as an alternative to conventional loss gradients, to generate perturbed examples in adversarial training. The experiment results indicate that the frequency-based approach performs comparably with the gradient-based approach in improving model robustness. Our research offers a novel and more intuitive perspective for understanding word-level textual adversarial attacks and proposes a new direction to improve model robustness.

Index Terms—adversarial attack, natural language processing, AI safety

I. INTRODUCTION

Deep Neural Networks (DNNs) have exhibited vulnerability to adversarial examples (AEs) [1], [2], which are crafted by adding imperceptible perturbations to the original inputs. In Natural Language Processing (NLP), numerous adversarial attacks have been proposed, which are typically categorized by the perturbation granularity: character-level [3], [4], word-level [5], [6], sentence-level [7], [8], and mix-level modification [9]. Among them, word-level attacks have attracted the most research interest, due to the superior performance on both attack success rate and AE quality [10], [11]. Thus, this work primarily explores these word-level attacks.

Simultaneously, the development of defenses against textual adversarial attacks has become a critical area of study. Notable defense strategies include adversarial training where the model gains robustness by training on the worst-case examples [12]–[14], adversarial data augmentation which trains models with

n -FD	Text
Raw	it's hard to imagine that even very small children will be impressed by this tired retreat
1-FD	it's ^{139→16} challenging to imagine that even very small children will be impressed by this tired retreat
2-FD	it's hard to imagine that even very small children will be ^{6→22} stunning by this tired retreat _{1→0,4→0}

Fig. 1: Illustrations of two AEs exhibiting 1-FD and 2-FD, respectively. The 1-gram (blue numbers) and 2-gram (red numbers) frequency changes are presented. In the second AE, the substitution of “impressed” with “stunning” raises the 1-gram frequency (6 → 22). However, it concurrently reduces the 2-gram frequency (1 → 0, 4 → 0).

AE-augmented training sets [15], AE detection [16]–[18], and certified robustness [19], [20].

Despite the tremendous progress achieved, the fundamental mechanisms of word-level textual attacks, as well as the intrinsic properties of the AEs crafted by them, are not yet fully explored. Considering that textual attacks and defenses are generally oriented to security-sensitive domains such as spam filtering [21] and toxic comment detection [22], a clear understanding of textual attacks is important. It will elucidate the vulnerability of the DNN-based applications and contribute to enhancing their robustness.

This work seeks to understand word-level textual attacks from a novel perspective: n -gram frequency. According to Zipf’s law [23], the frequency of words (1-gram) in a linguistic corpus is generally inversely proportional to their ranks. This means more common words appear exponentially more often than rarer ones, a pattern also holds true for n -grams [24]. While humans can easily navigate this frequency distribution disparity, DNNs struggle, which may lead to issues such as gender bias [25] and semantic blending [26]. We hypothesize that the highly uneven distribution of n -grams may induce

*Correspondence to liu_shengcai@cfar.a-star.edu.sg.

instability in models, particularly for sequences that occur less frequently, thus making them vulnerable to adversarial attacks.

To test this hypothesis, we thoroughly analyzed AEs generated by six different attack methods, targeting three DNN architectures across two datasets. The results reveal a consistent pattern across all attacks: a strong tendency toward generating examples characterized by a descending n -gram frequency, i.e., AEs contain less commonly occurring n -gram sequences than original ones. Figure 1 showcases instances where AEs demonstrate decrease in n -gram frequency. Moreover, this tendency is most pronounced when n equals 2, broadening the earlier focus in this field that only considered the frequency of single words [27]. Extra experiments also reveal that DNNs have difficulty processing n -FD examples.

These findings suggest a straightforward strategy to enhance model robustness: training on n -FD examples. Unlike the common adversarial training approaches that use gradients to perturb examples to maximize loss, we suggest a new approach that perturbs examples to reduce their n -gram frequency. We integrate this approach into the recent convex hull defense strategy [28] for adversarial training. Surprisingly, our frequency-based approach performs comparably to the gradient-based approach in improving model robustness. In summary, our main contributions are:

- Our analysis reveals that word-level attacks exhibit a strong tendency toward generating n -FD examples.
- Our experiments confirm that training models on n -FD examples can effectively improve model robustness, achieving defensive results comparable to the gradient-based approach.
- We provide a novel, intuitive perspective for understanding word-level adversarial attacks through the lens of n -gram frequency. Additionally, we offer a new direction to enhance the robustness of NLP models by n -FD examples.

II. UNDERSTAND WORD-LEVEL ATTACKS FROM THE n -FD PERSPECTIVE

In this section, we introduce the word-level textual attacks and the definition of n -gram frequency descend (n -FD). Then we experimentally demonstrate that word-level adversarial attacks prefer generating n -FD examples.

A. Preliminaries

a) Word-Level Textual Attacks: As the most widely studied attacks in NLP [29], word-level textual attacks generate AEs by substituting words in the original texts. Let $\mathbf{x} = [x_1, x_2, \dots, x_L]$ denote a text with L words. A word-level attack would first construct a candidate substitute set $S(x_i) = \{s_j^{(i)}\}_{j=0}^K$ with size K for each word x_i . Then it iteratively replaces a word in \mathbf{x} with some substitute selected from the candidate set, until attack succeeds.

b) n -gram Frequency: By definition, n -gram is a contiguous sequence of n words in the given texts. The i -th n -gram of text \mathbf{x} is defined as:

$$g_i^n := [x_i, x_{i+1}, \dots, x_{i+n-1}], i \in [1, L - n + 1]. \quad (1)$$

TABLE I: Summary of the attacks for AE generation, including model access, substitution methods, and search strategies. W/B represents white/black-box attack. WSG represents word-saliency-based greedy search.

Attack	Access	Substitution	Search
GA [5]	B	Counter-fitted	Genetic
PWWS [6]	B	WordNet	WSG
TF [30]	B	Counter-fitted	WSG
PSO [15]	B	HowNet	Particle Swarm
LS [31]	B	HowNet	Local Search
HF [32]	W	Counter-fitted	Gradient

We define the number of occurrences of an n -gram in the training set as its n -gram frequency, denoted as $\phi(g_i^n)$. Then the n -gram frequency of text \mathbf{x} , denoted as $\Phi_n(\mathbf{x})$, is the average of the n -gram frequencies of its all n -grams:

$$\Phi_n(\mathbf{x}) := \frac{1}{L - n + 1} \sum_{i=1}^{L-n+1} \phi(g_i^n). \quad (2)$$

c) n -gram Frequency Descend (n -FD): Given \mathbf{x} , supposing an attack generates an example \mathbf{x}' by substituting some words in \mathbf{x} , then \mathbf{x}' is a n -FD example if it has lower n -gram frequency than \mathbf{x} , i.e., $\Phi_n(\mathbf{x}') < \Phi_n(\mathbf{x})$.

Similarly, \mathbf{x}' is a n -gram frequency ascend (n -FA) example and a n -gram frequency constant (n -FC) example if $\Phi_n(\mathbf{x}') > \Phi_n(\mathbf{x})$ and $\Phi_n(\mathbf{x}') = \Phi_n(\mathbf{x})$, respectively.

d) n -FD Substitution: If a word substitution decreases the n -gram frequency of the text, then it is dubbed n -FD substitution. Formally, given text \mathbf{x} , let $\mathbf{x}_{x_i \rightarrow s_j^{(i)}}$ denote the text generated by substituting x_i in \mathbf{x} with $s_j^{(i)}$. Then the n -gram frequency change of the text, denoted as $\Delta\Phi_n(s_j^{(i)}; \mathbf{x})$, is:

$$\Delta\Phi_n(s_j^{(i)}; \mathbf{x}) := \Phi_n(\mathbf{x}_{x_i \rightarrow s_j^{(i)}}) - \Phi_n(\mathbf{x}). \quad (3)$$

If $\Delta\Phi_n(s_j^{(i)}; \mathbf{x}) < 0$, then substitution $x_i \rightarrow s_j^{(i)}$ is a n -FD substitution. Similarly, if $\Delta\Phi_n(s_j^{(i)}; \mathbf{x}) > 0$ and $\Delta\Phi_n(s_j^{(i)}; \mathbf{x}) = 0$, then it is a n -FA substitution and a n -FC substitution, respectively. For example, in Figure 1, the replacement of “hard” \rightarrow “challenging” is a 1-FD substitution, while “impressed” \rightarrow “stunning” is a 2-FD substitution.

B. Adversarial Example Generation

a) Attacks: We selected six existing word-level attacks, including five black-box attacks: GA [5], PWWS [6], TextFooler (TF) [30], PSO [15], LocalSearch (LS) [31]; and one white-box attack: HotFlip (HF) [32]. These attacks are representative in the sense that and have achieved effective performance against various models. They employ different substitute candidate construction methods and search strategies, as summarized in Table I. To construct the substitute sets, PSO and LS use the language database of HowNet [33], PWWS uses WordNet [34], while GA, TF, and HF rely on Counter-fitted [35] embeddings.

TABLE II: Percentages of the n -FD, n -FA, and n -FC examples in the AEs generated by all the six attacks when attacking three models on two datasets.

n	n -FD (%)	n -FC (%)	n -FA (%)
1	91.27	0.75	7.98
2	93.51	2.58	3.92
3	87.29	10.55	2.17
4	72.56	26.24	1.21

b) Dataset: We ran these attacks on 1000 test examples randomly selected from two public classification datasets: IMDB reviews dataset (IMDB) [36] for sentiment analysis and AG-News corpus (AGNews) [37] for topic classification.

c) Victim Models: The attacking experiments covered three different DNN architectures: convolutional neural network (CNN) [38], long short-term memory (LSTM) [39] and pre-trained BERT [40].

C. Results and Analysis

Adversarial attacks have n -FD tendency. Table II summarizes the average percentages of n -FD, n -FA, and n -FC examples (with $n = 1, 2, 3, 4$) of the AEs generated by the six attacks, calculated across three models and two dataset. One can observe that when n is 1, 2, 3, around 90% of the AEs show n -FD characteristic. Figure 2 further shows the detailed distributions of the n -gram frequency changes induced by the PWWS attack. It can be observed that most changes show the decreases in terms of n -gram frequency. Overall, all these attacks exhibit a strong tendency toward generating n -FD examples.

n -FD tendency is most pronounced when n equals 2. Based on Table II, 2-FD examples achieve better coverage compared with other cases. When $n = 1$, the percentage of n -FA is high, indicating a significant portion of the AEs containing more-frequent words. For $n = 3, 4$, the percentage of n -FC is large, which means they are not good indicator to interpret AEs. Further experiments show that, on average, 97% of n -FC cases are out-of-vocabulary (OOV) replacement, where both original and new n -grams never appear in the training set.

Models exhibit reduced performance on n -FD examples. Previous analysis implies that NLP models struggle more with n -FD examples. To test this hypothesis, we conducted experiments using the IMDB test set. For each test example, we generated one n -FD and one n -FA example through random word substitutions. Then, we evaluated the standardly trained models on the three sets: the original test set, n -FD example set, and n -FA example set. Figure 3 shows that the model’s predictions are less accurate on n -FD examples than on n -FA examples. This outcome is expected, as DNNs typically do not learn effectively from small sample sizes without specific training techniques. Besides, the poor adaption to smaller samples has a minor effect on evaluation metrics, which are designed to assess performance across a broad range of data.

Algorithm 1 n -FD convex hull training

Require: Dataset D , model f with parameter θ , adversarial steps T_{adv}

- 1: Initialize θ
- 2: Initialize n -gram frequency table T_{Φ_n} .
- 3: **for** epoch = $1 \cdots N_{epoch}$ **do**
- 4: **for** $x, y \in D$ **do**
- 5: Randomly initialize w_0
- 6: $g \leftarrow 0$
- 7: **for** $t = 0$ to T_{adv} **do**
- 8: Compute \tilde{x}_t by Eq. (6) using w_t
- 9: $g \leftarrow g + \nabla_{\theta} L(f(\tilde{x}_t), y)$
- 10: Update n -gram frequency table T_{Φ_n} by \tilde{x}_t
- 11: Compute w_{t+1} by Eq. (8) using T_{Φ_n} .
- 12: **end for**
- 13: Update θ by g
- 14: **end for**
- 15: **end for**

III. TRAINING ON n -FD EXAMPLES IMPROVES ROBUSTNESS

The findings from the previous section indicate that AEs exhibit an n -FD tendency, on which models perform poorly. Building upon this insight, a intuitive approach is to train models on n -FD examples, similar to adversarial training. To evaluate the feasibility of this approach, we developed an adversarial training framework that relies on n -gram frequency. In this approach, it’s the frequency, not the gradient, that directs the generation of AEs. This section will detail the approach.

A. n -FD Adversarial Training

In the conventional adversarial training paradigm, the training object is modeled as a min-max problem, where the inner goal is to find an example x_{adv} that maximizes the prediction loss, formulated as:

$$x_{adv} = \arg \max_{x' \in \mathcal{P}(x)} \mathcal{L}(F(x'), y), \quad (4)$$

where x_{adv} denotes the loss maximizing example, and $\mathcal{P}(x)$ is the perturbation set consisting of all texts that can be generated by doing the substitution operation on x . \mathcal{L} denotes the loss function of trained model F . In practice, training algorithm iteratively update x to approximately approach x_{adv} with the help of gradient.

To access the effectiveness of n -FD examples, we modify the gradient-based adversarial training paradigm to n -FD adversarial training, where the inner objective is to find the n -FD example, formulated as:

$$x_{n-FD} = \arg \min_{x' \in \mathcal{P}(x)} \Phi_n(x'). \quad (5)$$




Fig. 2: Distributions of the n -gram frequency changes induced by PWWS attack when attacking BERT on the IMDB dataset. The blue, orange, and purple bars represent the n -FD, n -FA, and n -FC examples, respectively. The exact percentage values are shown in the legend. From left to right, the value of n varies from 1 to 4.




Fig. 3: The confidence distribution for a CNN classifier on the clean examples (Orig), n -FD examples (n -FD) and n -FA examples (n -FA) from IMDB dataset. Confidence refers to the softmax probability of the true class. n is 1 to 4 from left to right images. Models perform similarly on clean examples and n -FA examples, but worse on n -FD examples.

B. Applying n -FD Adversarial Training to Convex Hull

a) *Convex Hull Framework*: We apply our method to a recently proposed convex hull paradigm [14], [28]. The generated AE during training is a sequence of virtual vectors \tilde{x}_i , which is a convex combination of synonyms $S(x_i)$, formulated as:

$$\tilde{x}_i = \sum_{j=0}^K w_j^{(i)} s_j^{(i)}, \quad (6)$$

where the $w_j^{(i)}$ is the corresponding coefficient for substitution $s_j^{(i)}$, which meet the convex hull constraints $\{\sum_j w_j^{(i)} = 1, w_j^{(i)} > 0\}$. Previous works used gradient-based methods to update w and build loss-maximizing AEs during training, formulated as:

$$\Delta w_j^{(i)} = \alpha \|\nabla_{w_j^{(i)}} \mathcal{L}\|, \quad (7)$$

where α is adversarial step size and $\|\cdot\|$ represents l_2 normalize operation.

b) *n -FD Convex Hull*: In n -FD adversarial training, we replace the gradient ascending direction with frequency descend direction to generate virtual AEs, formulated as:

$$\Delta w_j^{(i)} = -\alpha \left\| \Delta \Phi_n(s_j^{(i)}; \mathbf{x}) \right\|, \quad (8)$$

where $\Delta \Phi_n$ is defined in Eq. (3). This equation aims to increase the weight of n -FD substitutions. The full training algorithm is showed in Alg. 1.

We implemented two n -FD training methods based on 1-gram and 2-gram frequency, denoted as **ADV-F1** and **ADV-F2**, respectively¹. We follow the implementation of [28].

When $n = 1$, ADV-F1 updates $w_j^{(i)}$ by its corresponding word frequency, the $\Delta \Phi_1(s_j^{(i)}; \mathbf{x})$ in Eq. 8 is computed as follows:

$$\Delta \Phi_1(s_j^{(i)}; \mathbf{x}) = \phi(s_j^{(i)}). \quad (9)$$

Notice that we omit the frequency of the original example $\Phi_n(\mathbf{x})$ as it remains constant. When $n = 2$, ADV-F2 use the frequency of two 2-grams that contains $s_j^{(i)}$ to update $w_j^{(i)}$. $\Delta \Phi_2(s_j^{(i)}; \mathbf{x})$ is computed as follows:

$$\Delta \Phi_2(s_j^{(i)}; \mathbf{x}) = \phi([x_{i-1}, s_j^{(i)}]) + \phi([s_j^{(i)}, x_{i+1}]). \quad (10)$$

Notice that the frequency information is updated during training, so the update direction of w is dynamic.

C. Experimental Settings

a) *Dataset and Models*: Dataset includes Internet Movie Database (IMDB) [36] for sentiment classification task and the AG-News corpus (AGNews) [37] for topic classification task. Experiments are adopted on three different DNN architectures: convolutional neural network (CNN) [38], long short-term memory (LSTM) [39] and pre-trained BERT [40].

¹For memory and computational speed reasons, we didn't implement algorithm with larger n .

TABLE III: Classification accuracy (%) of ADV-G and ADV-F on clean examples (CLN) and different attacks across models and dataset. “AVG ROB” represents the average defensive performance against various attacks, with the numbers in brackets indicating the performance difference between ADV-F and ADV-G.

Dataset	Defense	CNN				LSTM				BERT				AVG ROB
		CLN	PWWS	TF	LS	CLN	PWWS	TF	LS	CLN	PWWS	TF	LS	
IMDB	Standard	88.7	2.3	5.8	2.0	87.4	2.4	8.2	0.9	94.1	34.7	36.2	3.5	10.7
	ADV-G	83.8	70.0	71.6	69.6	83.5	72.4	73.5	70.1	92.9	63.4	65.0	58.5	68.2
	ADV-F1	87.3	69.8	71.2	68.8	87.0	68.0	69.2	66.4	93.2	66.1	67.3	63.3	67.8 (-0.4)
	ADV-F2	87.6	70.2	72.0	70.5	86.5	72.0	71.7	68.7	93.3	67.0	67.5	64.4	69.3 (+1.1)
AGNews	Standard	92.0	57.2	54.5	52.1	92.8	61.4	59.2	55.3	94.7	71.8	70.1	50.2	59.1
	ADV-G	91.1	87.0	83.9	83.1	92.9	89.0	88.2	87.3	94.5	86.7	85.9	83.5	86.1
	ADV-F1	91.2	84.7	82.3	82.3	92.9	87.3	87.1	84.9	94.0	85.5	84.2	81.0	84.4 (-1.7)
	ADV-F2	91.4	86.8	83.8	82.9	92.9	88.5	87.5	85.6	94.5	86.0	84.4	83.6	85.5 (-0.6)




Fig. 4: Confidence distribution of different models on n -FD examples. After training on AEs, models also achieve better performance on n -FD examples.

b) Evaluation Metrics: We use the following metrics to evaluate defensive performance: 1) *Clean accuracy* (CLN) represents the model’s classification accuracy on the clean test set. 2) *Robust accuracy* is the model’s classification accuracy on the examples generated by a specific attack. A good defender should have higher clean accuracy and higher robust accuracy.

c) Attacks: We utilize three powerful word-level attacks to examine the robustness empirically: PWWS [6], TextFooler (TF) [30] and LocalSearch (LS) [31]. PWWS and TF employ greedy search with word-saliency exploring strategy. They first compute the word saliency of each original word and then do greedy substitution. On the other hand, LS is an iterative attacker who selects the worst-case transformation at each step. Thus, LS achieve a higher successful attack rate but requiring more query numbers. For fair comparison, we use the same substitution set for all defenders, following the setting of [28].

D. Results and Analysis

Training models on n -FD examples improves robustness. Table III reports the clean accuracy (CLN) and robust accuracy against three attacks (PWWS, TF, LS) across two dataset. We observe that both ADV-G and ADV-F effectively enhance the model’s robustness. ADV-F achieves competitive defensive performance with ADV-G, with only a minor difference of less than 2%. This small performance gap suggests that adversarial examples generated by both gradient methods and n -gram




Fig. 5: The n -gram frequency distribution of all training examples for ADV-G, ADV-F and standard training on IMDB. We normalize the frequencies with the training data size for comparison. Both ADV-G and ADV-F result in a more balanced n -gram frequency distribution, i.e., lower at the head but higher in the tail.

frequency have similar impacts on model robustness. Another key finding is that ADV-F2 consistently outperforms ADV-F1, indicating that 2-FD examples are more effective in increasing robustness than 1-FD examples. This observation aligns with the earlier findings discussed in Section II-C. Further analysis on how the choice of n -value influences robustness enhancement in Section III-E.

Gradient-based adversarial training generates n -FD examples. Figure 5 shows the sorted frequency distribution of all training examples, including those from adversarial (ADV-G, ADV-F) and standard training methods. Since standard training utilizes only clean examples, its distribution is the same as that of the original training set. One can observe that ADV-G, like ADV-F, increases the frequency of the less common n -grams.

Adversarial training improve model’s performance on n -FD examples. Figure 4 displays the distribution of confidence scores for the correct class across various models when handling n -FD examples. These models are trained using standard or adversarial methods on the IMDB dataset. Notably, after adversarial training, there is an increase in the confidence scores. Furthermore, both frequency-based and gradient-based training strategies effectively improve the model’s performance on n -FD examples, aligning with our




Fig. 6: The robustness performance of n -FD example augmentation with different n -values. “STD” is short for standard training, and “RDM” is short for random augmentation.

expectations.

E. Exploration of Proper n for Robustness Improvement

We conduct another defensive experiment to explore the robustness improvement considering different value of n . We generate n -FD examples of each training example for different n values, and then augment them to the training set. The results in Figure 6 show the different robustness performance when n increases from 1 to 4. The robust accuracy is measured by PWWS attack [6] on 1000 examples from AG’s News dataset, across three model architectures. We can observe that the optimal defense performance occurs when n is 2, with a decline in performance as n increases. This trend suggests that 2-FD information more effectively finds examples that enhance model robustness. However, as n become larger, the n -FD loses its informational value. Because most substitution will be considered as n -FD, reducing the process to random augmentation.

IV. RELATED WORKS

A. Textual Adversarial Attack

Despite the great success in the NLP field, deep neural networks are shown to be vulnerable against AEs in many NLP tasks [41]–[50]. Textual adversarial attacks can be classified by granularity into character-level, word-level, sentence-level and mixture-level [29]. Character-level attacks focus on deleting, adding or swapping characters [3], [4], which usually results in grammatical or spelling errors [10]. Sentence-level attacks change a whole sentence to fool the victim model, e.g., paraphrasing or appending texts [8]. And mixture-level attacks combine different level operations, e.g. phrases and words [9], [51]. In comparison, word-level attacks craft AEs by modifying words, where the candidates are formed by language databases [33], [34], word embeddings [35] or large-scale pre-trained language modeling. These word-level attacks directly leverage gradient [32], search methods [5], [31], [52] to find effective word substitutions.

B. Textual Adversarial Defense

The goal of adversarial defense is to make models have high performance on both clean and adversarial examples. Defense methods are generally categorized into empirical and certified, based on whether they provide provable robustness. Adversarial training and adversarial data augmentation are two popular approaches in empirical defense [6], [13], [30], [53], [54]. Adversarial training generates perturbation during training, while adversarial data augmentation obtains it after training, hence requiring a re-train phase. However, such augmentation is insufficient due to the large perturbation space, so these methods cannot guarantee the robustness of the model. Convex hull-based defense is another approach of adversarial training [14], [28], which optimizes the model’s performance over the convex hull formed by the embedding of synonyms. On the other hand, certified defense provides a provable robustness. Certified defense mainly consists of two types: Interval Bound Propagation (IBP) [19], [55] and random smooth [20]. IBP-based method computes the range of the model output by propagating the interval constraints of the inputs layer by layer, which requires knowing the structure of each layer. Random smooth methods achieve certified robustness by the statistical property of noised inputs.

To the best of our knowledge, only one previous work has explored frequency changes in AEs [27]. However, our research diverges significantly in several key areas, clearly establishing its distinct contribution to the field: 1) Scope of analysis: The previous work concentrates on the single words. In contrast, our work embraces a broader scope, examining general n -gram frequency. Notably, our statistical analysis show that 2-grams provide more insightful results than single-word analysis. 2) Purpose and application: The previous work primarily utilized word frequency as features to detect attacks. Conversely, we employ frequency analysis as a tool to deepen our understanding of word-level attacks. Furthermore, we verified that the use of n -FD examples specifically to improve the robustness of models.

V. CONCLUSION

This paper provides a novel understanding of word-level textual attacks through the lens of n -gram frequency, and provides a new direction to improve model robustness. Our analysis of adversarial examples reveals a the attackers’ general tendency towards n -FD examples, with $n = 2$ shows the . We also find that typically trained models are more vulnerable to n -FD examples, indicating potential risks for NLP models. Motivated by these findings, we introduce an n -FD adversarial training method that significantly improves model robustness, comparable to gradient-based approach. Notably, using 2-gram frequencies proves more efficient in fortifying models compared to 1-gram frequencies. We believe this work will deepen the understanding of adversarial attack and defense in NLP. However, there are limitations in our study. Primarily, we do not fully understand why there are some AEs that belongs to 2-FA. Furthermore, our study does not incorporate multiple n -gram information.

VI. ETHIC STATEMENT

While our interpretations and experimental findings have the potential to design a more powerful attack, this recognition raises critical ethical concerns, as advancements in attack strategies could be misused, leading to more effective ways of deceiving NLP systems. However, it is crucial to emphasize that our primary objective is to contribute positively to the field by enhancing the understanding and defense mechanisms against such attacks. Moreover, we proposed a method specifically designed to improve the robustness of models against several attacks.

ACKNOWLEDGMENT

This work was supported in part by Guangdong Major Project of Basic and Applied Basic Research (Grant No. 2023B0303000010), and in part by the National Natural Science Foundation of China under Grant 62272210.

REFERENCES

- [1] I. J. Goodfellow, J. Shlens, and C. Szegedy, "Explaining and harnessing adversarial examples," in *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, Y. Bengio and Y. LeCun, Eds., 2015.
- [2] N. Papernot, P. D. McDaniel, A. Swami, and R. E. Harang, "Crafting adversarial input sequences for recurrent neural networks," in *2016 IEEE Military Communications Conference, MILCOM 2016, Baltimore, MD, USA, November 1-3, 2016*, J. Brand, M. C. Valentí, A. Akinpelu, B. T. Doshi, and B. L. Gorsic, Eds. IEEE, 2016, pp. 49–54.
- [3] J. Gao, J. Lanchantin, M. L. Soffa, and Y. Qi, "Black-box generation of adversarial text sequences to evade deep learning classifiers," in *2018 IEEE Security and Privacy Workshops, SP Workshops 2018, San Francisco, CA, USA, May 24, 2018*. IEEE Computer Society, 2018, pp. 50–56.
- [4] J. Li, S. Ji, T. Du, B. Li, and T. Wang, "Textbugger: Generating adversarial text against real-world applications," in *26th Annual Network and Distributed System Security Symposium, NDSS 2019, San Diego, California, USA, February 24-27, 2019*. The Internet Society, 2019.
- [5] M. Alzantot, Y. Sharma, A. Elgohary, B. Ho, M. B. Srivastava, and K. Chang, "Generating natural language adversarial examples," in *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, E. Riloff, D. Chiang, J. Hockenmaier, and J. Tsujii, Eds. Association for Computational Linguistics, 2018, pp. 2890–2896.
- [6] S. Ren, Y. Deng, K. He, and W. Che, "Generating natural language adversarial examples through probability weighted word saliency," in *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, A. Korhonen, D. R. Traum, and L. Márquez, Eds. Association for Computational Linguistics, 2019, pp. 1085–1097.
- [7] M. T. Ribeiro, S. Singh, and C. Guestrin, "Semantically equivalent adversarial rules for debugging NLP models," in *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 1: Long Papers*, I. Gurevych and Y. Miyao, Eds. Association for Computational Linguistics, 2018, pp. 856–865.
- [8] R. Jia and P. Liang, "Adversarial examples for evaluating reading comprehension systems," in *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, EMNLP 2017, Copenhagen, Denmark, September 9-11, 2017*, M. Palmer, R. Hwa, and S. Riedel, Eds. Association for Computational Linguistics, 2017, pp. 2021–2031.
- [9] Y. Lei, Y. Cao, D. Li, T. Zhou, M. Fang, and M. Pechenizkiy, "Phrase-level textual adversarial attack with label preservation," in *Findings of the Association for Computational Linguistics: NAACL 2022, Seattle, WA, United States, July 10-15, 2022*, M. Carpuat, M. de Marneffe, and I. V. M. Ruíz, Eds. Association for Computational Linguistics, 2022, pp. 1095–1112.
- [10] D. Pruthi, B. Dhingra, and Z. C. Lipton, "Combating adversarial misspellings with robust word recognition," in *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, A. Korhonen, D. R. Traum, and L. Márquez, Eds. Association for Computational Linguistics, 2019, pp. 5582–5591.
- [11] J. Y. Yoo, J. X. Morris, E. Lifland, and Y. Qi, "Searching for a search method: Benchmarking search algorithms for generating NLP adversarial examples," in *Proceedings of the Third BlackboxNLP Workshop on Analyzing and Interpreting Neural Networks for NLP, BlackboxNLP@EMNLP 2020, Online, November 2020*, A. Alishahi, Y. Belinkov, G. Chrupala, D. Hupkes, Y. Pinter, and H. Sajjad, Eds. Association for Computational Linguistics, 2020, pp. 323–332.
- [12] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu, "Towards deep learning models resistant to adversarial attacks," in *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018.
- [13] C. Zhu, Y. Cheng, Z. Gan, S. Sun, T. Goldstein, and J. Liu, "Freelb: Enhanced adversarial training for natural language understanding," in *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020.
- [14] X. Dong, A. T. Luu, R. Ji, and H. Liu, "Towards robustness against natural language word substitutions," in *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021.
- [15] Y. Zang, F. Qi, C. Yang, Z. Liu, M. Zhang, Q. Liu, and M. Sun, "Word-level textual adversarial attacking as combinatorial optimization," in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, D. Jurafsky, J. Chai, N. Schlueter, and J. R. Tetreault, Eds. Association for Computational Linguistics, 2020, pp. 6066–6080.
- [16] D. Pruthi, B. Dhingra, and Z. C. Lipton, "Combating adversarial misspellings with robust word recognition," in *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, A. Korhonen, D. R. Traum, and L. Márquez, Eds. Association for Computational Linguistics, 2019, pp. 5582–5591.
- [17] Y. Zhou, J. Jiang, K. Chang, and W. Wang, "Learning to discriminate perturbations for blocking adversarial attacks in text classification," 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 2019, Hong Kong, China, November 3-7, 2019*, K. Inui, J. Jiang, V. Ng, and X. Wan, Eds. Association for Computational Linguistics, 2019, pp. 4903–4912.
- [18] L. Wang and X. Zheng, "Improving grammatical error correction models with purpose-built adversarial examples," in *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, B. Webber, T. Cohn, Y. He, and Y. Liu, Eds. Association for Computational Linguistics, 2020, pp. 2858–2869.
- [19] R. Jia, A. Raghunathan, K. Göksel, and P. Liang, "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 2019, Hong Kong, China, November 3-7, 2019*, K. Inui, J. Jiang, V. Ng, and X. Wan, Eds. Association for Computational Linguistics, 2019, pp. 4127–4140.
- [20] M. Ye, C. Gong, and Q. Liu, "SAFER: A structure-free approach for certified robustness to adversarial word substitutions," in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, D. Jurafsky, J. Chai, N. Schlueter, and J. R. Tetreault, Eds. Association for Computational Linguistics, 2020, pp. 3465–3475.
- [21] A. Bhowmick and S. M. Hazarika, "E-mail spam filtering: a review of techniques and trends," *Advances in electronics, communication and computing*, pp. 583–590, 2018.
- [22] H. Hosseini, S. Kannan, B. Zhang, and R. Poovendran, "Deceiving google's perspective API built for detecting toxic comments," *CoRR*, vol. abs/1702.08138, 2017.
- [23] G. K. Zipf, *Human behavior and the principle of least effort: An introduction to human ecology*. Ravenio Books, 2016.
- [24] W. B. Cavnar, J. M. Trenkle *et al.*, "N-gram-based text categorization," in

- Proceedings of SDAIR-94, 3rd annual symposium on document analysis and information retrieval*, vol. 161175. Citeseer, 1994.
- [25] J. Angwin, J. Larson, S. Mattu, and L. Kirchner, "Machine bias," in *Ethics of Data and Analytics*. Auerbach Publications, 2016, pp. 254–264.
- [26] C. Gong, D. He, X. Tan, T. Qin, L. Wang, and T. Liu, "FRAGE: frequency-agnostic word representation," in *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, S. Bengio, H. M. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, Eds., 2018, pp. 1341–1352.
- [27] M. Mozes, P. Stenetorp, B. Kleinberg, and L. D. Griffin, "Frequency-guided word substitutions for detecting textual adversarial examples," in *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume, EACL 2021, Online, April 19 - 23, 2021*, P. Merlo, J. Tiedemann, and R. Tsarfaty, Eds. Association for Computational Linguistics, 2021, pp. 171–186.
- [28] Y. Zhou, X. Zheng, C. Hsieh, K. Chang, and X. Huang, "Defense against synonym substitution-based adversarial attacks via dirichlet neighborhood ensemble," in *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021*, C. Zong, F. Xia, W. Li, and R. Navigli, Eds. Association for Computational Linguistics, 2021, pp. 5482–5492.
- [29] W. E. Zhang, Q. Z. Sheng, A. Alhazmi, and C. Li, "Adversarial attacks on deep-learning models in natural language processing: A survey," *ACM Trans. Intell. Syst. Technol.*, vol. 11, no. 3, pp. 24:1–24:41, 2020.
- [30] D. Jin, Z. Jin, J. T. Zhou, and P. Szolovits, "Is bert really robust? a strong baseline for natural language attack on text classification and entailment," in *Proceedings of the AAAI conference on artificial intelligence*, vol. 34, no. 05, 2020, pp. 8018–8025.
- [31] S. Liu, N. Lu, C. Chen, and K. Tang, "Efficient combinatorial optimization for word-level adversarial textual attack," *IEEE ACM Trans. Audio Speech Lang. Process.*, vol. 30, pp. 98–111, 2022.
- [32] J. Ebrahimi, A. Rao, D. Lowd, and D. Dou, "Hotflip: White-box adversarial examples for text classification," in *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 2: Short Papers*, I. Gurevych and Y. Miyao, Eds. Association for Computational Linguistics, 2018, pp. 31–36.
- [33] F. Qi, C. Yang, Z. Liu, Q. Dong, M. Sun, and Z. Dong, "Openhownet: An open sememe-based lexical knowledge base," *CoRR*, vol. abs/1901.09957, 2019.
- [34] G. A. Miller, "Wordnet: A lexical database for english," *Commun. ACM*, vol. 38, no. 11, pp. 39–41, 1995.
- [35] N. Mrksic, D. Ó Séaghdha, B. Thomson, M. Gasic, L. M. Rojas-Barahona, P. Su, D. Vandryke, T. Wen, and S. J. Young, "Counter-fitting word vectors to linguistic constraints," in *NAACL HLT 2016, The 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, San Diego California, USA, June 12-17, 2016*, K. Knight, A. Nenkova, and O. Rambow, Eds. The Association for Computational Linguistics, 2016, pp. 142–148.
- [36] M. Lan, Z. Zhang, Y. Lu, and J. Wu, "Three convolutional neural network-based models for learning sentiment word vectors towards sentiment analysis," in *2016 International Joint Conference on Neural Networks, IJCNN 2016, Vancouver, BC, Canada, July 24-29, 2016*. IEEE, 2016, pp. 3172–3179.
- [37] X. Zhang, J. J. Zhao, and Y. LeCun, "Character-level convolutional networks for text classification," in *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, Eds., 2015, pp. 649–657.
- [38] Y. Kim, "Convolutional neural networks for sentence classification," in *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP 2014, October 25-29, 2014, Doha, Qatar: A meeting of SIGDAT, a Special Interest Group of the ACL*, A. Moschitti, B. Pang, and W. Daelemans, Eds. ACL, 2014, pp. 1746–1751.
- [39] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Comput.*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [40] J. Devlin, M. Chang, K. Lee, and K. Toutanova, "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 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, J. Burstein, C. Doran, and T. Solorio, Eds. Association for Computational Linguistics, 2019, pp. 4171–4186.
- [41] R. Jia and P. Liang, "Adversarial examples for evaluating reading comprehension systems," in *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, EMNLP 2017, Copenhagen, Denmark, September 9-11, 2017*, M. Palmer, R. Hwa, and S. Riedel, Eds. Association for Computational Linguistics, 2017, pp. 2021–2031.
- [42] B. Liang, H. Li, M. Su, P. Bian, X. Li, and W. Shi, "Deep text classification can be fooled," in *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden*, J. Lang, Ed. ijcai.org, 2018, pp. 4208–4215.
- [43] Z. Zhao, D. Dua, and S. Singh, "Generating natural adversarial examples," in *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018.
- [44] Z. Li, C. Wang, C. Liu, P. Ma, D. Wu, S. Wang, and C. Gao, "Vrptest: Evaluating visual referring prompting in large multimodal models," *arXiv preprint arXiv:2312.04087*, 2023.
- [45] N. Lu, S. Liu, R. He, W. Qi, and K. Tang, "Large language models can be guided to evade ai-generated text detection," *arXiv preprint arXiv:2305.10847*, 2023.
- [46] Z. Li, P. Ma, H. Wang, S. Wang, Q. Tang, S. Nie, and S. Wu, "Unleashing the power of compiler intermediate representation to enhance neural program embeddings," in *Proceedings of the 44th International Conference on Software Engineering*, 2022, pp. 2253–2265.
- [47] P. Yang, L. Zhang, H. Liu, and G. Li, "Reducing idleness in financial cloud services via multi-objective evolutionary reinforcement learning based load balancer," *Science China Information Sciences*, vol. 67, no. 2, p. 120102, 2024.
- [48] J. Wu, W. Fan, S. Liu, Q. Liu, R. He, Q. Li, and K. Tang, "Dataset condensation for recommendation," in *arXiv*, 2023.
- [49] J. Wu, W. Fan, J. Chen, S. Liu, Q. Li, and K. Tang, "Disentangled contrastive learning for social recommendation," in *Proc. of CIKM'2022*. ACM, 2022.
- [50] S. Liu, C. Chen, X. Qu, K. Tang, and Y.-S. Ong, "Large language models as evolutionary optimizers," *arXiv preprint arXiv:2310.19046*, 2023.
- [51] J. Guo, Z. Zhang, L. Zhang, L. Xu, B. Chen, E. Chen, and W. Luo, "Towards variable-length textual adversarial attacks," *ArXiv*, vol. abs/2104.08139, 2021.
- [52] S. Liu, N. Lu, W. Hong, C. Qian, and K. Tang, "Effective and imperceptible adversarial textual attack via multi-objectivization," *ACM Transactions on Evolutionary Learning and Optimization*, 2024, just Accepted.
- [53] J. Dong, Y. Wang, J.-H. Lai, and X. Xie, "Improving adversarially robust few-shot image classification with generalizable representations," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 9025–9034.
- [54] J. Dong, S.-M. Moosavi-Dezfooli, J. Lai, and X. Xie, "The enemy of my enemy is my friend: Exploring inverse adversaries for improving adversarial training," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 24 678–24 687.
- [55] Z. Shi, H. Zhang, K. Chang, M. Huang, and C. Hsieh, "Robustness verification for transformers," in *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020.