

CERT-ED: Certifiably Robust Text Classification for Edit Distance

Zhuoqun Huang, Neil G. Marchant, Olga Ohrimenko, Benjamin I. P. Rubinstein

School of Computing and Information Systems

University of Melbourne

Melbourne, Australia

{zhuoqun, nmarchant, oohrimenko, brubinstein}@unimelb.edu.au

Abstract

With the growing integration of AI in daily life, ensuring the robustness of systems to inference-time attacks is crucial. Among the approaches for certifying robustness to such adversarial examples, randomized smoothing has emerged as highly promising due to its nature as a wrapper around arbitrary black-box models. Previous work on randomized smoothing in natural language processing has primarily focused on specific subsets of edit distance operations, such as synonym substitution or word insertion, without exploring the certification of all edit operations. In this paper, we adapt *Randomized Deletion* (Huang et al., 2023) and propose, CERTified Edit Distance defense (CERT-ED) for natural language classification. Through comprehensive experiments, we demonstrate that CERT-ED outperforms the existing Hamming distance method RanMASK (Zeng et al., 2023) in 4 out of 5 datasets in terms of both accuracy and the cardinality of the certificate. By covering various threat models, including 5 direct and 5 transfer attacks, our method improves empirical robustness in 38 out of 50 settings.

1 Introduction

Deep nets, transformers, and other modern machine learning approaches have recently achieved significant performance on many natural language processing (NLP) tasks thanks to their ability to generalize to complex and unseen data. However, the well-documented vulnerability of these models to evasion attacks (a.k.a. adversarial examples) raises concerns about their use in practice. For example, numerous previous works have developed attacks that can misguide models by perturbing text at the word-level (Alzantot et al., 2018; Li et al., 2020; Ren et al., 2019; Zang et al., 2020; Jin et al., 2020; Li et al., 2021; Garg and Ramakrishnan, 2020), character-level (Karpukhin et al., 2019; Gao et al., 2018; Ebrahimi et al., 2018) or sentence-level (Iyyer

et al., 2018; Wang et al., 2020; Qi et al., 2021; Guo et al., 2021), while preserving consistent semantics.

Although a wide range of defenses have been proposed against adversarial examples, they are routinely broken by subsequent attacks: Eger and Benz (2020) and Morris et al. (2020) showed that while adversarial training typically yields good robustness against a target attack, it is less robust against unseen attacks. Consequently, certified robustness has gained considerable interest as a result of competition between attackers and defenders, where a classifier’s prediction can be guaranteed to be invariant to a specified set of adversarial perturbations (Cohen et al., 2019; Wang et al., 2021a).

Certified robustness methods have been well studied for *continuous data that is fixed-dimensional* like images (Wong and Kolter, 2018; Dvijotham et al., 2018; Mirman et al., 2018; Weng et al., 2018; Lecuyer et al., 2019; Cohen et al., 2019). Among certification approaches, randomized smoothing (Lecuyer et al., 2019; Cohen et al., 2019; Levine and Feizi, 2020) has attained state-of-the-art performance in many tasks. Due to the discrete nature of text inputs, however, developing randomized smoothing mechanism for NLP tasks is more challenging. Ye et al. (2020) and Wang et al. (2021a) were the first to investigate randomized smoothing under a synonym substitution threat model. Similarly, Zeng et al. (2023) proposed RanMASK that adapts Randomized Ablation (Levine and Feizi, 2020) to NLP and is provably verifiable for Hamming distance under a fixed number of word substitutions. However, such robustness certificates fall short in defending against general perturbations as inserting a single word (Garg and Ramakrishnan, 2020) would void any of these substitution-based certificates. To remedy this, Huang et al. (2023) proposed a method for producing certifiable predictions under edit distance perturbations, called Randomized Deletion. However, their method is limited to binary classification tasks, and has only

Inputs		Label
$x = \text{"click here to see the bang busit is wild!!!!!!"}$		Spam
<hr/>		
	Smoothed inputs	Base classifier predictions
	$z_1 = \text{"click here to the bang is"}$	$y_1 = \text{Spam}$
	$z_2 = \text{"to see bang is"}$	$y_2 = \text{Not spam}$
CERT-ED	$z_3 = \text{"here busit is wild!!!!!!"}$	$y_3 = \text{Spam}$
	$z_4 = \text{"to see bang busit wild!!!!!!"}$	$y_4 = \text{Spam}$
	$z_5 = \text{"click wild!!!!!!"}$	$y_5 = \text{Spam}$
	$z_6 = \text{"here to"}$	$y_6 = \text{Not Spam}$

		Certified prediction
		Spam: 90
		Not spam: 10
		Certified radius
		3 Edits
<hr/>		
	Adversarial inputs	
	$x = \text{"click around here to see the bang busit is pretty wild!!!!!!"}$	
	$dist_{o,t}(x, \bar{x}) = 2 < 3$	Certified correct prediction

Figure 1: Top: Clean sample from Spam-assassin dataset. Middle: CERT-ED applied to the perturbed input to produce edit distance certified prediction of “Spam” and certified radius of 3. Bottom: Real adversarial sample generated by, Clare (Li et al., 2021), against a model without CERT-ED. The green words are adversarially inserted words. CERT-ED is certifiably robust to this adversarial example as the edit distance between the clean and adversarial inputs is 2, less than the certified radius.

been applied to the malware detection domain. In this paper we address this limitation.

Our contributions are summarized as follows:

- We propose and implement CERTified Edit Distance defense (CERT-ED), a multi-class extension of Randomized Deletion (Huang et al., 2023), that can provably certify multi-class predictions for NLP classification tasks¹. It smooths input text by adding deletion noise to produce predictions that are certifiably robust under arbitrary attacks within a computed edit distance radius r (Figure 1).
- To compare our edit distance certificates with Hamming distance certificates used in previous work, we define certified cardinality, a discrete analogue of certified volume that has been used to compare certificates with different geometries in the vision domain. We evaluate CERT-ED using 5 datasets and find significant improvement of over the RanMASK baseline, both in certified accuracy and certified cardinality.
- We conduct a comprehensive empirical evaluation of robustness against five state-of-the-art direct attacks and transfer attacks. Our results show improved robust accuracy in 20 out of 25 settings

¹Our implementation is available at <https://github.com/Dovermore/nlp-smoothing-software>.

for direct attacks and 18 out of 25 settings for transfer attacks.

2 Edit distance robustness

We consider text sequence classification tasks, where a model f predicts the class $y \in \mathcal{Y}$ of input text $x \in \mathcal{X}$. For example, in fake news detection the input text is a news article and the possible classes are “fake” and “real” (Rashkin et al., 2017). We are interested in studying robustness under an adversary that can make a bounded number of edits to the text.

We define an *edit* to be an operation that deletes (del), inserts (ins) or substitutes (sub) a single token in the text. A *token* is a contiguous chunk of characters—e.g., a word or sub-word. The mapping from text to tokens is determined by the adversary’s tokenizer t . While our method is compatible with any choice of t , we set t to be a whitespace tokenizer in our experiments for comparison with prior work on attacks (Garg and Ramakrishnan, 2020; Jin et al., 2020; Li et al., 2020, 2021) and robustness (Zeng et al., 2023; Zhang et al., 2024b) at the word-level. We note that t is solely used to model the adversary’s edits, and is distinct from any tokenizer that may appear in model f itself.

For generality, we consider adversaries whose edit operations are constrained to the set $\mathcal{o} \subseteq \{\text{del, ins, sub}\}$. For instance, $\mathcal{o} = \{\text{ins, sub}\}$ for

an adversary that cannot perform deletions. Given the adversary’s tokenizer t and allowed edit operations \mathbf{o} , we measure the extent of the adversary’s perturbation using edit distance $\text{dist}_{\mathbf{o},t}(\bar{\mathbf{x}}, \mathbf{x})$, which counts the minimum number of edits required to transform original text $\bar{\mathbf{x}}$ into perturbed text \mathbf{x} .

Our objective is to design text sequence classification models that are certifiably robust under this threat model. Formally, given input text $\mathbf{x} \in \mathcal{X}$ to model f , we would like to guarantee that f ’s prediction is unchanged even if \mathbf{x} was modified by an adversary that made up to r edits:

$$\forall \bar{\mathbf{x}} \in B_r(\mathbf{x}; \mathbf{o}, t) : f(\mathbf{x}) = f(\bar{\mathbf{x}}). \quad (1)$$

Here

$$B_r(\mathbf{x}; \mathbf{o}, t) := \{\bar{\mathbf{x}} \in \mathcal{X} : \text{dist}_{\mathbf{o},t}(\bar{\mathbf{x}}, \mathbf{x}) \leq r\} \quad (2)$$

is the set of text inputs that can be transformed into \mathbf{x} via at most r edits. As is typical for randomized smoothing, we will develop mechanisms that produce a randomized radius r given input sequence \mathbf{x} , such that with some chosen high probability at least $1 - \alpha$, this radius is a valid certificate at \mathbf{x} .

3 Certified robustness via randomized smoothing

Our approach for achieving certified robustness under bounded edit distance perturbations is based on randomized smoothing. Specifically, we apply the randomized smoothing mechanism of Huang et al. (2023), which was originally formulated for binary classification of generic sequences. We review the mechanism in a text classification context in Section 3.1 and propose, CERT-ED, a certifiably robust extension to the multi-class setting in Section 3.2. Our derived certificates cover attacks BAE-I (Garg and Ramakrishnan, 2020) and Clare (Li et al., 2021) not covered by prior work.

3.1 Randomized deletion smoothing

Randomized smoothing has emerged as a general purpose method for constructing certifiably robust classifiers (Kumari et al., 2023). Consider a base classifier $f_b : \mathcal{X} \rightarrow \mathcal{Y}$ and a randomized mechanism $\phi : \mathcal{X} \rightarrow P(\mathcal{X})$ that generates perturbed inputs. In the following we construct a smoothed classifier f that assigns a probabilistic score to class y given input \mathbf{x} :

$$p_y(\mathbf{x}) = \mathbb{E}_{\mathbf{z} \sim \phi(\mathbf{x})} [\mathbf{1}_{f_b(\mathbf{z})=y}]. \quad (3)$$

The smoothed classifier’s prediction is then the class with the highest score: $f(\mathbf{x}) = \arg \max_{y \in \mathcal{Y}} p_y(\mathbf{x})$.

Huang et al. (2023) instantiate randomized smoothing with a deletion mechanism for sequences that achieves certified edit distance robustness. In the text domain, we apply their deletion mechanism at the level of tokens determined by tokenizer t . Specifically, given input text $\mathbf{x} \in \mathcal{X}$ containing n tokens, the deletion mechanism generates n indicator variables $\epsilon = (\epsilon_1, \dots, \epsilon_n)$ where $\epsilon_i \stackrel{\text{iid}}{\sim} \text{Bernoulli}(p_{\text{del}})$. The perturbed text is then obtained by deleting any token i for which $\epsilon_i = 1$ and keeping the remaining tokens (in order).

Practicalities Exact computation of the scores in (3) scales exponentially in the number of tokens n . We therefore follow standard practice in randomized smoothing and obtain upper/lower confidence bounds on the scores using Monte Carlo sampling. When constructing a smoothed classifier f , we fine-tune the base classifier f_b on inputs perturbed by the mechanism ϕ , as this results in better performance.

3.2 CERT-ED: Multi-class edit distance certification

We extend the edit distance certificate of Huang et al. (2023) to the multi-class setting. We refer to randomized deletion smoothing with this new certificate as CERT-ED. We discuss the setup and assumptions here and present two key results, which demonstrate how a certificate can be obtained for a bounded Levenshtein (edit) distance adversary whose edit operations \mathbf{o} are unconstrained. All proofs are provided in Appendix A.

We adopt a standard setup for certification of smoothed classifiers. Given input text \mathbf{x} , we let $y = \arg \max_c p_c(\mathbf{x})$ be the class that achieves the highest score and $y' = \arg \max_{c \neq y} p_c(\mathbf{x})$ be the class that achieves the second-highest score. We assume that $\mu_y \leq p_y(\mathbf{x})$ is a lower bound on the highest score and $\mu_{y'} \geq p_{y'}(\mathbf{x})$ is an upper bound on the second highest score. Apart from this minimal information, we do not assume any knowledge of the base classifier f_b . We begin by obtaining upper and lower bounds on the smoothed classifier’s score at a neighboring input $\bar{\mathbf{x}}$.

Theorem 1 (General pairwise certificate). *Consider a pair of text inputs $\mathbf{x}, \bar{\mathbf{x}} \in \mathcal{X}$. Suppose $\bar{\mathbf{x}}$ can be transformed into \mathbf{x} using a minimal number of edit operations by deleting n_{del} tokens, inserting n_{ins} tokens and substituting n_{sub} tokens—i.e., $\text{dist}_{\mathbf{o},t}(\bar{\mathbf{x}}, \mathbf{x}) = n_{\text{sub}} + n_{\text{ins}} + n_{\text{del}}$. Then the smoothed*

Adversary’s ops			Certified radius (\downarrow)
del	ins	sub	
✓	✓	✓	$\left\lfloor \log_{p_{\text{del}}} \frac{2+\mu_{y'}-\mu_y}{2} \right\rfloor$
✓		✓	$\left\lfloor \log_{p_{\text{del}}} \frac{2+\mu_{y'}-\mu_y}{2} \right\rfloor$
	✓	✓	$\left\lfloor \log_{p_{\text{del}}} \frac{2+\mu_{y'}-\mu_y}{2} \right\rfloor$
		✓	$\left\lfloor \log_{p_{\text{del}}} \frac{2+\mu_{y'}-\mu_y}{2} \right\rfloor$
✓	✓		$\left\lfloor \log_{p_{\text{del}}} \frac{1}{1-\mu_{y'}+\mu_y} \right\rfloor$
✓			$\left\lfloor \log_{p_{\text{del}}} \frac{1}{1-\mu_{y'}+\mu_y} \right\rfloor$
	✓		$\left\lfloor \log_{p_{\text{del}}} (1 + \mu_{y'} - \mu_y) \right\rfloor$

Table 1: Certified radii as a function of the types of edit operations the adversary can perform.

classifier’s scores for any class $y \in \mathcal{Y}$ satisfy

$$\begin{aligned}
p_{\text{del}}^{n_{\text{del}}-n_{\text{ins}}} \left(p_y(\mathbf{x}) - 1 + p_{\text{del}}^{n_{\text{sub}}+n_{\text{ins}}} \right) &\leq p_y(\bar{\mathbf{x}}) \\
&\leq p_{\text{del}}^{n_{\text{del}}-n_{\text{ins}}} p_y(\mathbf{x}) + 1 - p_{\text{del}}^{n_{\text{sub}}+n_{\text{del}}}.
\end{aligned}$$

The above result is not immediately useful on its own. However, it can be used to derive edit distance certificates under various constraints on the number of edit operations of each type ($n_{\text{del}}, n_{\text{ins}}, n_{\text{sub}}$) the adversary can perform. Below we present a Levenshtein distance certificate, which covers an adversary that can perform up to r edits of any type (insertions, deletions or substitutions).

Theorem 2 (Levenshtein distance certificate). *Consider a text input $\mathbf{x} \in \mathcal{X}$ for which a lower bound on the smoothed classifier’s highest score μ_y and an upper bound on the smoothed classifier’s runner-up score $\mu_{y'}$ satisfy $\mu_y \geq \mu_{y'}$. Then the smoothed classifier predicts y for any neighboring text input $\bar{\mathbf{x}} \in \mathcal{X}$ such that $\text{dist}_{\mathbf{o},t}(\bar{\mathbf{x}}, \mathbf{x}) \leq r$ with $\mathbf{o} = \{\text{del}, \text{ins}, \text{sub}\}$ and $r = \lfloor \log_{p_{\text{del}}} \frac{1}{2} (2 + \mu_{y'} - \mu_y) \rfloor$. If the upper and lower bounds hold jointly with confidence $1 - \alpha$, then the certificate holds with probability $1 - \alpha$.*

This result is readily adapted for adversaries that are constrained in the kinds of edit operations they can perform. We provide certificates for seven constrained settings in Table 1.

4 Experiments

To evaluate the effectiveness of our methods, we train and evaluate models with CERT-ED certification across a variety of English datasets

and compare observed performance against RanMASK (Zeng et al., 2023) and FreeLB (Zhu et al., 2020). We show that our method uniformly dominates RanMASK in certified radius and accuracy for 4 out of 5 datasets in Section 4.1. We then conduct 5 direct and 5 transfer attacks on the certified models and show that CERT-ED is also empirically the most robust in 38 out of 50 settings in Section 4.2 and Appendix D. Though CERT-ED is empirically more robust than RanMASK against word substitution and character-level attacks, the performance is more mixed for attacks that induce edit distance perturbations. Finally, we also report a 3 times speedup of our method compared to RanMASK in Appendix E.

Datasets We present results on five diverse datasets listed in Table 2. All datasets are partitioned into training, validation and test sets. AG-News (Zhang et al., 2015) and IMDB (Maas et al., 2011) are standard datasets for topic classification and sentiment analysis respectively, that have been used for evaluation in prior work (Zeng et al., 2023). However, since adversaries may lack incentive to target these models, we follow recommendations of Chen et al. (2022), and consider Spam-assassin (Chen et al., 2022) for spam detection, and SatNews (Yang et al., 2017) and LUN (Rashkin et al., 2017; Chen et al., 2022) for unreliable news detection, which are arguably more attractive targets for attackers². We collect all data from Hugging-Face Datasets³ and the AdvBench repository⁴ (Chen et al., 2022).

Dataset	Avg words	Number of samples		
		Train	Valid	Test
AG-News	37.8	108 000	12 000	7 600
IMDB	231.2	22 500	2 500	25 000
Spam-assassin	228.2	2 152	239	2 378
LUN	269.9	13 416	1 490	6 454
SatNews	384.8	22 738	2 526	7 202

Table 2: Summary of datasets.

Models We use the Hugging Face Transformers library (Wolf et al., 2020) to load a pre-trained RoBERTa model (Zhuang et al., 2021) as a base

²These datasets might contain offensive or inappropriate languages due to their adversarial nature

³<https://github.com/huggingface/datasets>

⁴<https://github.com/thunlp/Advbench>

classifier for CERT-ED and as a non-certified baseline. We include RanMASK (Zeng et al., 2023) as a certified baseline for Hamming distance. This baseline covers a bounded number of word substitutions only. Since RanMASK is also based on randomized smoothing, albeit with a different masking mechanism, we use the same base classifier, training procedure and parameter settings as CERT-ED where possible. While evaluating empirical robustness, we additionally compare against FreeLB (Zhu et al., 2020), a adversarial training method that is not based on randomized smoothing. We use *perturbation strength* to refer to the deletion rate parameter p_{del} of CERT-ED and the masking rate parameter p_{mask} of RanMASK. We note that TextCRS (Zhang et al., 2024a) claims to certify against limited word substitution, insertion, and deletion operations of a single type, however we do not include it as a baseline since the resulting edit distance certificates are vacuous ($r = 0$) for text greater than 2 words in length (see Appendix F).

For both CERT-ED and RanMASK, we use a white-space tokenizer for smoothing and fine-tune the base RoBERTa model on the training set where inputs are perturbed by the corresponding smoothing mechanism. We report our training parameter settings in Table 5 of Appendix B.

4.1 Certified accuracy and robustness

Setup Our first set of experiments compares the accuracy and certificates generated by CERT-ED and RanMASK. Following prior work (Huang et al., 2023; Zeng et al., 2023), we use 1000 Monte Carlo samples for prediction, 4000 samples for estimating the certified radius and a confidence level of 95%. In the unlikely case where the prediction and certificate disagree, we report the prediction, and report a certified radius of zero. Table 3 presents a summary of the results, where we highlight the best of three perturbation strengths (80%, 90% and 95%) for each dataset. Results for all three perturbation strengths (80%, 90% and 95%) are provided in Table 6 of Appendix C.1.

Clean accuracy To assess performance in a non-adversarial setting, we report accuracy on the clean (unperturbed) test set. For both CERT-ED and RanMASK, we observe a slight degradation in clean accuracy of 0.2–5.2% across all datasets compared to the non-smoothed baseline, where the highest degradation is seen on the IMDB dataset. This may be due to the fact that IMDB is a sentiment classifi-

Model	$p_{\text{del}}/p_{\text{mask}}$	Clean Accuracy	Median CR	Median log CC
AG-News dataset (avg length 37.84)				
Baseline	—	94.84%	—	—
RanMASK	80%	93.91%	2	12.25
CERT-ED	80%	93.33%	2	12.65
IMDB dataset (avg length 231.16)				
Baseline	—	93.47%	—	—
RanMASK	90%	86.87%	2	14.02
CERT-ED	90%	88.26%	2	14.49
Spam-assassin dataset (avg length 228.16)				
Baseline	—	98.02%	—	—
RanMASK	90%	97.65%	6	37.49
CERT-ED	90%	97.81%	6	38.66
LUN dataset (avg length 269.93)				
Baseline	—	99.16%	—	—
RanMASK	90%	97.91%	6	34.51
CERT-ED	90%	98.28%	6	37.94
SatNews dataset (avg length 384.84)				
Baseline	—	94.22%	—	—
RanMASK	95%	90.10%	7	47.09
CERT-ED	95%	92.07%	8	54.76

Table 3: Key certification results drawn from Table 6. All metrics are computed on the entire test set. “Median CR” is the median certified radius and “median log CC” is the base-10 logarithm of the median certified cardinality. The certified cardinality is exact for RanMASK, however a lower bound is used for CERT-ED. CERT-ED outperforms RanMASK in terms of certified accuracy for 4 out of 5 datasets and specifically excels on datasets with longer average length. Highlighted values are the better of the two smoothed mechanisms.

cation task, where sensitivity to small perturbations such as names (Prabhakaran et al., 2019) and the presence or absence of single words (e.g., “not”) is more likely. CERT-ED outperforms RanMASK in both clean accuracy and the size of the certificate for the four datasets with longer text. While CERT-ED fuzzes the input length, RanMASK preserves it, and may therefore have an advantage for short text where the length conveys information. However, for datasets with longer text like LUN and SatNews, RanMASK may introduce too many masking tokens and confuse the base model.

Certified accuracy and certified cardinality We report certified radius (CR) as a measure of robustness, where a larger radius indicates robustness to larger perturbations. However, the certified radius is not comparable between CERT-ED and

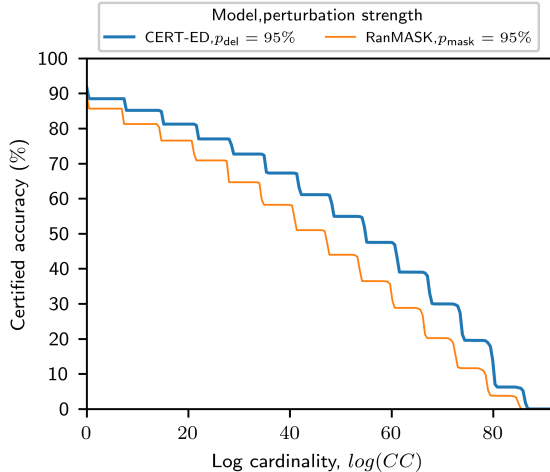


Figure 2: Certified accuracy for CERT-ED and RanMASK as a function of the log-cardinality of the certificate for the SatNews dataset. We see that CERT-ED certifies a set up to 10^{10} times larger than RanMASK for the same accuracy.

RanMASK, as the metric used to define the certificate for CERT-ED is Levenshtein (edit) distance whereas the metric for RanMASK is Hamming distance. As an alternative, we therefore also measure the cardinality of the certificate, that is the number of perturbed textual inputs it contains, which we refer to as the certified cardinality (CC). This is analogous to certified volume, used in prior work in the vision domain (Pfrommer et al., 2023). For a given radius and vocabulary size⁵, we compute the cardinality exactly for RanMASK and use a lower bound for CERT-ED based on a result of Charalampopoulos et al. (2020, Fact 17). We note that it is possible to compute the exact cardinality for a Levenshtein distance certificate using a Levenshtein automaton (Touzet, 2016), however the computation is expensive and only improves on the lower bound by approximately one order of magnitude. Despite the under-estimated certified cardinality of CERT-ED, we find it dominates RanMASK across all datasets, with a widening gap for datasets with longer text. To assess trade-offs between robustness and accuracy, we plot the certified accuracy as a function of the log-cardinality. The certified accuracy at a given log-cardinality c is the fraction of instances in the test set for which the model’s prediction is correct and the cardinality of the certificate is at least c . Figure 2 plots the certified accuracy

⁵We assume a vocabulary size of 50265 (matching RoBERTa), although this is a property of the threat model.

for the SatNews dataset, demonstrating an improvement up to 10^{10} times in the certified cardinality for CERT-ED compared to RanMASK. Due to space constraints, we present certified accuracy plots for the remaining datasets in Appendix C.2, where a strict domination of CERT-ED over RanMASK can be observed for 4 out of 5 datasets.

4.2 Empirical robustness

Attack setup We evaluate the empirical robustness of CERT-ED and baselines using a modified version of TextAttack⁶ (Morris et al., 2020) and attack recipes implemented by Zhang et al. (2024a). We select five representative attacks that cover a variety of perturbations: Clare (Li et al., 2021), BAE-I (Garg and Ramakrishnan, 2020), BERT-Attack (Li et al., 2020), TextFooler (Jin et al., 2020) and DeepWordBug (Gao et al., 2018). We describe these attacks further in Section 5.1. We randomly select 1000 samples (200 for FreeLB) from the test set to evaluate the robustness of the models. For CERT-ED and RanMASK, we estimate the prediction using a Monte Carlo sample size of 100 to speed up the attack process. We impose a 10 minute timeout for each attack and treat timeout as a failed attack. As CERT-ED is about 3 times faster than RanMASK (Appendix E) for prediction and certification, this puts CERT-ED at a disadvantage as attacks against it may use up to a 3 times as many queries. Hence, we also include selected results for query-constrained (rather than time-constrained) attacks in Appendix D. For Clare, due to the excessive amount of querying, we also limit the maximum number of queries to 10 000. Further details on our categorization of attack outcomes are provided in Appendix D.

Threat models We consider two distinct threat models: *direct attacks* which have access to the model’s confidence; and *transfer attacks* for which attacks are generated against the non-certified baseline and if successful, transferred to the target model (Appendix D). This means the robust accuracy for the non-certified baseline will always be 0%. We report both the *clean accuracy* (CIA) on original instances, and *robust accuracy* (RoA) on attacked instances.

Results We present direct attack results in Table 4. Compared to the non-certified baseline and FreeLB, CERT-ED sacrifices a small amount of

⁶<https://github.com/QData/TextAttack>

Method	Clean Accuracy %	Clare RoA %	BAE-I RoA %	BERT-Attack RoA %	TextFooler RoA %	DeepWordBug RoA %
AG-News dataset						
Baseline	94.7	26.9	58.5	29.0	22.1	43.5
FreeLB [†]	96.0	41.0	73.0	42.0	27.5	49.5
RanMASK	91.7	77.3	79.6	48.5	47.7	45.1
CERT-ED	92.0	67.9	78.1	56.1	56.9	54.4
IMDB dataset						
Baseline	94.8	50.3	28.1	9.2	8.0	30.0
FreeLB [†]	94.0	60.5	38.5	6.0	7.5	47.5
RanMASK	86.9	85.5	81.9	53.0	45.9	52.4
CERT-ED	87.9	82.3	71.1	54.0	49.9	55.3
Spam-assassin dataset						
Baseline	97.7	91.8	93.0	51.2	36.2	80.2
FreeLB [†]	97.0	94.0	95.5	57.5	44.5	81.0
RanMASK	97.0	95.4	96.0	93.5	92.1	93.5
CERT-ED	97.2	95.2	94.9	95.5	94.3	94.7
LUN dataset						
Baseline	99.2	89.0	80.1	57.2	64.6	78.5
FreeLB [†]	99.5	94.0	93.5	85.0	89.0	89.5
RanMASK	98.7	93.8	96.4	90.5	90.5	88.8
CERT-ED	99.3	95.3	96.6	93.9	93.3	92.4
SatNews dataset						
Baseline	94.7	85.0	58.6	51.3	50.7	61.4
FreeLB [†]	97.0	93.0	86.0	63.0	61.5	73.5
RanMASK	91.2	91.2	91.2	80.7	80.8	74.5
CERT-ED	93.2	93.2	91.7	83.5	81.8	83.3

Table 4: Empirical attack results against Baseline, FreeLB, RanMASK and CERT-ED under direct attacks. Both RanMASK and CERT-ED use a 90% perturbation strength. Highlighted values are the best in each column for that dataset. CERT-ED outperforms RanMASK in *all* substitution and character-level attacks. For edit distance attacks, the performance is mixed, with CERT-ED performing better for datasets with longer text. [†] FreeLB uses a smaller subset of 200 samples (of the original 1000 samples) for evaluation for all settings.

clean accuracy to achieve a significant improvement in robust accuracy across all settings. For word substitution and character-level attacks (BERT-Attack, TextFooler, DeepWordBug), CERT-ED uniformly outperforms RanMASK and FreeLB in all datasets by up to 8.8% and 49.8% robust accuracy. Surprisingly, we find that CERT-ED is less effective against edit distance attacks (Clare and BAE-I), performing well on datasets with longer input sizes, but worse on AG-News and IMDB. We attribute this to the difference in query efficiency between RanMASK and CERT-ED. For Clare on the AG-News dataset, the mean number of queries to RanMASK and CERT-ED are 4253 and 8891, respectively, meaning CERT-ED is subject to a stronger attack than RanMASK. Due to the longer text size, all Clare attacks timed out on the SatNews dataset for CERT-ED and RanMASK.

5 Related work

Despite the rapid advancement of neural networks for many important tasks, their vulnerability to adversarial examples has long been known (Szegedy et al., 2014; Goodfellow et al., 2015) and continues to be a liability for current large-scale models (Qi et al., 2024; Raina et al., 2024). While early work focused on the vision domain, where attacks can leverage continuous optimizers, methods have been proposed to generate adversarial examples in the language domain across a variety of tasks including text classification (Alzantot et al., 2018; Garg and Ramakrishnan, 2020; Li et al., 2020) and machine translation (Zhang et al., 2021; Belinkov and Bisk, 2018). These adversarial examples are especially concerning when applied to high stakes tasks such as content moderation and fake news detection (Chen et al., 2022). Alongside research on attacks, certified defenses have been proposed as a strong

counter-measure that can guarantee that the most pernicious adversarial examples do not exist within a specified threat model (Lecuyer et al., 2019; Cohen et al., 2019). However, their application in NLP has been limited, in part due to the discrete nature of text (Ye et al., 2020; Wang et al., 2021a; Zeng et al., 2023; Wang et al., 2023; Zhang et al., 2024b).

5.1 Attacks

Character-level attacks focus on perturbing characters within words to maintain the imperceptibility of attacks to human inspection. Karpukhin et al. (2019) augmented training data with orthographic noise (character-level insertions, substitutions, deletions, and swaps) to improve the robustness of machine translation models. The DeepWordBug (Gao et al., 2018) and HotFlip (Ebrahimi et al., 2018) attacks rank character and token importance in the input text, then greedily search for perturbations using the importance ranking to achieve misclassification.

Word-level attacks generate adversarial examples by perturbing individual words in text to maximize the model loss while preserving the semantics and syntax of the original sentence. Synonym substitution is one of the most common approaches for maintaining semantic consistency. These attacks typically start by ranking the words by their importance and then sequentially substituting words with synonyms generated using a similarity metric (Alzantot et al., 2018; Li et al., 2020; Ren et al., 2019; Zang et al., 2020; Jin et al., 2020). While powerful, these attacks do not fully represent the capabilities of an adversary, as they do not consider insertion or deletion of words. Subsequent work has investigated edit distance-constrained adversarial attacks. Both Clare (Li et al., 2021) and BAE-I (Garg and Ramakrishnan, 2020) explored the use of edit perturbations beyond substitution and found them successful against NLP models.

Our method CERT-ED can certify robustness against word- and character-level attacks by setting t to be a white-space or character-level tokenizer, accordingly.

Unlike word- or character-level attacks, that perform a limited number of localized perturbations to preserve semantics, sentence-level attacks adversarially paraphrase entire sentences, typically using LLMs (Iyyer et al., 2018; Wang et al., 2020; Qi et al., 2021), or using a parametrized adversarial distribution that enables gradient-based search (Guo et al., 2021).

5.2 Defenses

Starting with adversarial training (Goodfellow et al., 2015), numerous empirical defense methods have been proposed to improve robustness of NLP models (Ren et al., 2019; Zang et al., 2020; Wang et al., 2021b; Zhu et al., 2020; Ivgi and Berant, 2021; Li et al., 2020). In the language domain, adversarial training perturbs inputs either in the text space (Ren et al., 2019; Jin et al., 2020; Zang et al., 2020; Li et al., 2020; Ivgi and Berant, 2021; Wang et al., 2021b) or in the embedding space using bounded adversarial noise (Miyato et al., 2017; Zhu et al., 2020). Although empirical defenses may provide excellent robustness against attacks they are tailored for, they cannot guarantee effectiveness against an adaptive attacker.

To mitigate issues with empirical defenses, certified defenses aim to provide a robustness guarantee against arbitrary attacks within a specified threat model. Huang et al. (2019) and Jia et al. (2019) first used interval bound propagation to certify robustness under synonym substitutions. More recently, works have applied randomized smoothing to achieve certified robustness under synonym or word substitution threat models (Ye et al., 2020; Wang et al., 2021a). Despite the use of insertion and deletion operations in published attacks, achieving certified robustness against these operations has not been well-studied. To date, only TextCRS (Zhang et al., 2024a) has made progress on this front. TextCRS certifies against word-level permutations and perturbations in the embedding space, which can provide robustness guarantees against limited word-level substitutions, deletions or insertions. However, TextCRS only partially covers the edit distance ball at a given radius, meaning its certificates are vacuous for our threat model (see Appendix F). In the malware detection domain, Huang et al. (2023) proposed a deletion-based mechanism to achieve certified edit distance robustness for malware binary classification models. Our proposed method CERT-ED adapts their mechanism for the language domain and extends the certificate to support multi-class classification.

6 Conclusion

In this work, we investigated certified robustness for natural language classification tasks, where adversaries can perturb input text by adding, deleting, or substituting words. We adapted randomized deletion smoothing (Huang et al., 2023) to the language

domain, and derived an edit distance robustness certificate for the multi-class setting. We refer to our certified method as CERT-ED and conducted comprehensive experiments on five datasets. Our results show that CERT-ED outperforms the existing randomized smoothing method (RanMASK) for word substitution robustness in terms of both accuracy and certified cardinality on 4 out of 5 datasets. Our method also excels in robustness against direct and transfer attacks, demonstrating significant improvements over existing methods.

7 Ethical considerations

This study focuses on enhancing the robustness of NLP models. Although adversarial examples are generated during the research, their use is strictly for evaluation purposes. We also acknowledge the assistance of ChatGPT and GitHub Copilot for scaffolding code in released artifacts.

8 Limitations

Robustness certification aims to measure the risk of adversarial examples, while randomized smoothing provides both certification and mitigation against such attacks. However, our results show that at higher smoothing levels, randomized smoothing can reduce the benign accuracy compared to undefended models. While our experiments are comprehensive and cover a wide range of datasets, attacks and threat models, results might differ on other base model architectures and natural language tasks. Although our edit distance threat model covers a boarder range of attacks than prior work on certifications for NLP, and threat models better aligned for sequence data than the bounded ℓ_p -norm threat models popular for image research, attackers may opt to make many edits to input data, and in some NLP tasks input semantics could be changed with few edits. For example, sentiment analysis can be local in nature, relying on the sentiment of a single adjective. By contrast, tasks like fake news detection typically rely on more global features to distinguish task classes. Lastly, while our approach is more scalable than alternative certification strategies, it does introduce computational overheads.

Acknowledgments

This work was supported by the Department of Industry, Science, and Resources, Australia under AUSMURI CATCH. We acknowledge the Australian Research Council under Discovery Project

DP220102269. This research was supported by The University of Melbourne’s Research Computing Services and the Petascale Campus Initiative. We also thank the reviewers in the ARR process for their valuable feedback.

References

- Moustafa Alzantot, Yash Sharma, Ahmed Elgohary, Bo-Jhang Ho, Mani Srivastava, and Kai-Wei Chang. 2018. [Generating natural language adversarial examples](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2890–2896, Brussels, Belgium. Association for Computational Linguistics.
- Yonatan Belinkov and Yonatan Bisk. 2018. [Synthetic and natural noise both break neural machine translation](#). In *International Conference on Learning Representations*.
- Panagiotis Charalampopoulos, Solon P. Pissis, Jakub Radoszewski, Tomasz Waleń, and Wiktor Zuba. 2020. [Unary Words Have the Smallest Levenshtein k-Neighbourhoods](#). In *31st Annual Symposium on Combinatorial Pattern Matching*, CPM, pages 10:1–10:12.
- Yangyi Chen, Hongcheng Gao, Ganqu Cui, Fanchao Qi, Longtao Huang, Zhiyuan Liu, and Maosong Sun. 2022. [Why should adversarial perturbations be imperceptible? Rethink the research paradigm in adversarial NLP](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 11222–11237, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Jeremy Cohen, Elan Rosenfeld, and Zico Kolter. 2019. [Certified Adversarial Robustness via Randomized Smoothing](#). In *International Conference on Machine Learning*, volume 97 of *ICML*, pages 1310–1320.
- Krishnamurthy Dvijotham, Sven Gowal, Robert Stanforth, Relja Arandjelovic, Brendan O’Donoghue, Jonathan Uesato, and Pushmeet Kohli. 2018. [Training verified learners with learned verifiers](#). *Preprint*, arXiv:1805.10265.
- Javid Ebrahimi, Anyi Rao, Daniel Lowd, and Dejing Dou. 2018. [HotFlip: White-box adversarial examples for text classification](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 31–36, Melbourne, Australia. Association for Computational Linguistics.
- Steffen Eger and Yannik Benz. 2020. [From hero to zéro: A benchmark of low-level adversarial attacks](#). In *Proceedings of the 1st Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 10th International Joint Conference on Natural Language Processing*, pages 786–803, Suzhou, China. Association for Computational Linguistics.

- Ji Gao, Jack Lanchantin, Mary Lou Soffa, and Yanjun Qi. 2018. [Black-box generation of adversarial text sequences to evade deep learning classifiers](#). In *2018 IEEE Security and Privacy Workshops (SPW)*, pages 50–56.
- 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)*, pages 6174–6181, Online. Association for Computational Linguistics.
- Ian Goodfellow, Jonathon Shlens, and Christian Szegedy. 2015. [Explaining and Harnessing Adversarial Examples](#). In *International Conference on Learning Representations*.
- Chuan Guo, Alexandre Sablayrolles, Hervé Jégou, and Douwe Kiela. 2021. [Gradient-based adversarial attacks against text transformers](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 5747–5757, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Po-Sen Huang, Robert Stanforth, Johannes Welbl, Chris Dyer, Dani Yogatama, Sven Gowal, Krishnamurthy Dvijotham, and Pushmeet Kohli. 2019. [Achieving verified robustness to symbol substitutions via interval bound propagation](#). 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)*, pages 4083–4093, Hong Kong, China. Association for Computational Linguistics.
- Zhuoqun Huang, Neil G Marchant, Keane Lucas, Lujo Bauer, Olga Ohrimenko, and Benjamin Rubinstein. 2023. [RS-Del: Edit distance robustness certificates for sequence classifiers via randomized deletion](#). In *Advances in Neural Information Processing Systems*, volume 36, pages 18676–18711. Curran Associates, Inc.
- Maor Ivgi and Jonathan Berant. 2021. [Achieving model robustness through discrete adversarial training](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 1529–1544, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Mohit Iyyer, John Wieting, Kevin Gimpel, and Luke Zettlemoyer. 2018. [Adversarial example generation with syntactically controlled paraphrase networks](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1875–1885, New Orleans, Louisiana. Association for Computational Linguistics.
- 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)*, pages 4129–4142, Hong Kong, China. Association for Computational Linguistics.
- 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](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(05):8018–8025.
- Vladimir Karpukhin, Omer Levy, Jacob Eisenstein, and Marjan Ghazvininejad. 2019. [Training on synthetic noise improves robustness to natural noise in machine translation](#). In *Proceedings of the 5th Workshop on Noisy User-generated Text (W-NUT 2019)*, pages 42–47, Hong Kong, China. Association for Computational Linguistics.
- Anupriya Kumari, Devansh Bhardwaj, Sukrit Jindal, and Sarthak Gupta. 2023. [Trust, but verify: A survey of randomized smoothing techniques](#). *Preprint*, arXiv:2312.12608.
- Mathias Lecuyer, Vaggelis Atlidakis, Roxana Geambasu, Daniel Hsu, and Suman Jana. 2019. [Certified robustness to adversarial examples with differential privacy](#). In *2019 IEEE Symposium on Security and Privacy (SP)*, pages 656–672.
- Alexander Levine and Soheil Feizi. 2020. [Robustness certificates for sparse adversarial attacks by randomized ablation](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(04):4585–4593.
- 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*, pages 5053–5069, Online. Association for Computational Linguistics.
- Linyang Li, Ruotian Ma, Qipeng Guo, Xiangyang Xue, and Xipeng Qiu. 2020. [BERT-ATTACK: Adversarial attack against BERT using BERT](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6193–6202, Online. Association for Computational Linguistics.
- Andrew L. Maas, Raymond E. Daly, Peter T. Pham, Dan Huang, Andrew Y. Ng, and Christopher Potts. 2011. [Learning word vectors for sentiment analysis](#). In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 142–150, Portland, Oregon, USA. Association for Computational Linguistics.
- Matthew Mirman, Timon Gehr, and Martin Vechev. 2018. [Differentiable abstract interpretation for provably robust neural networks](#). In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 3578–3586. PMLR.

- Takeru Miyato, Andrew M. Dai, and Ian Goodfellow. 2017. [Adversarial training methods for semi-supervised text classification](#). In *International Conference on Learning Representations*.
- John 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*, pages 119–126, Online. Association for Computational Linguistics.
- Samuel Pfommer, Brendon G. Anderson, and Somayeh Sojoudi. 2023. [Projected randomized smoothing for certified adversarial robustness](#). *Transactions on Machine Learning Research*.
- Vinodkumar Prabhakaran, Ben Hutchinson, and Margaret Mitchell. 2019. [Perturbation sensitivity analysis to detect unintended model biases](#). 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)*, pages 5740–5745, Hong Kong, China. Association for Computational Linguistics.
- Fanchao Qi, Yangyi Chen, Xurui Zhang, Mukai Li, Zhiyuan Liu, and Maosong Sun. 2021. [Mind the style of text! Adversarial and backdoor attacks based on text style transfer](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 4569–4580, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Xiangyu Qi, Kaixuan Huang, Ashwinee Panda, Peter Henderson, Mengdi Wang, and Prateek Mittal. 2024. [Visual adversarial examples jailbreak aligned large language models](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(19):21527–21536.
- Vyas Raina, Adian Liusie, and Mark Gales. 2024. [Is LLM-as-a-Judge robust? investigating universal adversarial attacks on zero-shot LLM assessment](#). Preprint, arXiv:2402.14016.
- Hannah Rashkin, Eunsol Choi, Jin Yea Jang, Svitlana Volkova, and Yejin Choi. 2017. [Truth of varying shades: Analyzing language in fake news and political fact-checking](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2931–2937, Copenhagen, Denmark. Association for Computational Linguistics.
- 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 Annual Meeting of the Association for Computational Linguistics*, pages 1085–1097, Florence, Italy. Association for Computational Linguistics.
- Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian J. Goodfellow, and Rob Fergus. 2014. [Intriguing properties of neural networks](#). In *International Conference on Learning Representations*.
- H el ene Touzet. 2016. [On the Levenshtein automaton and the size of the neighbourhood of a word](#). In *Language and Automata Theory and Applications*, pages 207–218, Cham. Springer International Publishing.
- Robert A. Wagner and Michael J. Fischer. 1974. [The string-to-string correction problem](#). *Journal of the ACM*, 21(1):168–173.
- Tianlu Wang, Xuezhi Wang, Yao Qin, Ben Packer, Kang Li, Jilin Chen, Alex Beutel, and Ed Chi. 2020. [CATgen: Improving robustness in NLP models via controlled adversarial text generation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5141–5146, Online. Association for Computational Linguistics.
- Wenjie Wang, Pengfei Tang, Jian Lou, and Li Xiong. 2021a. [Certified robustness to word substitution attack with differential privacy](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1102–1112, Online. Association for Computational Linguistics.
- Xiaosen Wang, Yichen Yang, Yihe Deng, and Kun He. 2021b. [Adversarial training with fast gradient projection method against synonym substitution based text attacks](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(16):13997–14005.
- Yibin Wang, Yichen Yang, Di He, and Kun He. 2023. [Robustness-aware word embedding improves certified robustness to adversarial word substitutions](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 673–687, Toronto, Canada. Association for Computational Linguistics.
- Lily Weng, Huan Zhang, Hongge Chen, Zhao Song, Chou-Jui Hsieh, Luca Daniel, Duane Boning, and Inderjit Dhillon. 2018. [Towards fast computation of certified robustness for ReLU networks](#). In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 5276–5285. PMLR.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pi eric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Eric Wong and Zico Kolter. 2018. [Provable defenses against adversarial examples via the convex outer adversarial polytope](#). In *Proceedings of the 35th International Conference on Machine Learning*, volume 80

- of *Proceedings of Machine Learning Research*, pages 5286–5295. PMLR.
- Fan Yang, Arjun Mukherjee, and Eduard Dragut. 2017. [Satirical news detection and analysis using attention mechanism and linguistic features](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1979–1989, Copenhagen, Denmark. Association for Computational Linguistics.
- Mao Ye, Chengyue Gong, and Qiang Liu. 2020. [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*, pages 3465–3475, Online. Association for Computational Linguistics.
- 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*, pages 6066–6080, Online. Association for Computational Linguistics.
- Jiehang Zeng, Jianhan Xu, Xiaoqing Zheng, and Xuanjing Huang. 2023. [Certified robustness to text adversarial attacks by randomized \[MASK\]](#). *Computational Linguistics*, 49(2):395–427.
- X. Zhang, H. Hong, Y. Hong, P. Huang, B. Wang, Z. Ba, and K. Ren. 2024a. [Text-CRS: A generalized certified robustness framework against textual adversarial attacks](#). In *2024 IEEE Symposium on Security and Privacy (SP)*, pages 56–56, Los Alamitos, CA, USA. IEEE Computer Society.
- Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. [Character-level convolutional networks for text classification](#). In *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc.
- Xinze Zhang, Junzhe Zhang, Zhenhua Chen, and Kun He. 2021. [Crafting adversarial examples for neural machine translation](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1967–1977, Online. Association for Computational Linguistics.
- Zeliang Zhang, Wei Yao, Susan Liang, and Chenliang Xu. 2024b. [Random smooth-based certified defense against text adversarial attack](#). In *Findings of the Association for Computational Linguistics: EACL 2024*, pages 1251–1265, St. Julian’s, Malta. Association for Computational Linguistics.
- Chen Zhu, Yu Cheng, Zhe Gan, Siqi Sun, Tom Goldstein, and Jingjing Liu. 2020. [FreeLB: Enhanced adversarial training for natural language understanding](#). In *International Conference on Learning Representations*.
- Liu Zhuang, Lin Wayne, Shi Ya, and Zhao Jun. 2021. [A robustly optimized BERT pre-training approach with post-training](#). In *Proceedings of the 20th Chinese National Conference on Computational Linguistics*, pages 1218–1227, Huhhot, China. Chinese Information Processing Society of China.

A Proofs of results in Section 3.2

In this appendix, we provide proofs of the certification results presented in Section 3.2. Our proofs follow (Huang et al., 2023), however we additionally provide an upper bound on the smoothed classifier’s score, which is needed to support certification of multi-class classifiers.

We begin by defining notation to express the deletion mechanism symbolically. Recall from Section 3.1 that $\epsilon = (\epsilon_1, \dots, \epsilon_n)$ is a vector of deletion indicator variables for input text \mathbf{x} containing $n = |\mathbf{t}(\mathbf{x})|$ tokens, where $\epsilon_i = 1$ if the i -th token is to be deleted and $\epsilon_i = 0$ otherwise. The space of possible deletion indicators for input text \mathbf{x} is denoted $\mathcal{E}(\mathbf{x}) = \{0, 1\}^{|\mathbf{t}(\mathbf{x})|}$. We let $q(\epsilon) = \prod_i p_{\text{del}}^{\epsilon_i} (1 - p_{\text{del}})^{1 - \epsilon_i}$ denote the (Bernoulli) probability mass for a given ϵ . We write $\text{apply}(\mathbf{x}, \epsilon)$ to denote the resultant text after deleting the tokens referenced in ϵ from text \mathbf{x} .

Using this notation, we can express the smoothed classifier’s score defined in (3) as a sum over the space of deletion indicator variables:

$$p_y(\mathbf{x}) = \sum_{\epsilon \in \mathcal{E}(\mathbf{x})} s(\epsilon, \mathbf{x}) \quad (4)$$

$$\text{with } s(\epsilon, \mathbf{x}) = q(\epsilon) \mathbf{1}_{f_b(\text{apply}(\mathbf{x}, \epsilon)) = y}. \quad (5)$$

The first step in our analysis is to identify a correspondence between deletions for neighboring text. Let

$$\sqsubseteq = \{(\epsilon, \epsilon') \in \mathcal{E}(\mathbf{x}) \times \mathcal{E}(\mathbf{x}) : \forall i, \epsilon_i \leq \epsilon'_i\}.$$

be a partial order on the space of deletion indicators $\mathcal{E}(\mathbf{x})$. We can then write $\epsilon \sqsubseteq \epsilon'$ if ϵ' can be obtained from ϵ by deleting additional tokens. This allows us to define a set of deletions building on ϵ :

$$\mathcal{E}(\mathbf{x}, \epsilon) := \{\epsilon' \in \mathcal{E}(\mathbf{x}) : \epsilon \sqsubseteq \epsilon'\}. \quad (6)$$

The following result adapted from (Huang et al., 2023, Lemma 4) identifies pairs of deletions $\bar{\epsilon}$ to $\bar{\mathbf{x}}$ and ϵ to \mathbf{x} such that the terms $s(\bar{\mathbf{x}}, \bar{\epsilon})$ and $s(\mathbf{x}, \epsilon)$ are proportional.

Lemma 3 (Huang et al., 2023). *Let \mathbf{z}^* be a longest common subsequence (Wagner and Fischer, 1974) of $\mathbf{t}(\bar{\mathbf{x}})$ and $\mathbf{t}(\mathbf{x})$, and let $\bar{\epsilon}^* \in \mathcal{E}(\bar{\mathbf{x}})$*

and $\epsilon^* \in \mathcal{E}(\mathbf{x})$ be any deletions such that $\text{apply}(\bar{\mathbf{x}}, \bar{\epsilon}^*) = \text{apply}(\mathbf{x}, \epsilon^*) = \mathbf{t}^{-1}(\mathbf{z}^*)$. Then there exists a bijection $m: \mathcal{E}(\bar{\mathbf{x}}, \bar{\epsilon}^*) \rightarrow \mathcal{E}(\mathbf{x}, \epsilon^*)$ such that $\text{apply}(\bar{\mathbf{x}}, \bar{\epsilon}) = \text{apply}(\bar{\mathbf{x}}, \epsilon^*)$ for any $\bar{\epsilon} \sqsupseteq \bar{\epsilon}^*$. Furthermore for $\epsilon = m(\bar{\epsilon})$ we have

$$s(\bar{\epsilon}, \bar{\mathbf{x}}) = p_{\text{del}}^{|\mathbf{t}(\bar{\mathbf{x}})| - |\mathbf{t}(\mathbf{x})|} s(\epsilon, \mathbf{x}).$$

Using the above result, we can relate the smoothed classifier's score at $\bar{\mathbf{x}}$ and \mathbf{x} as follows:

$$\begin{aligned} p_y(\bar{\mathbf{x}}) &= \sum_{\bar{\epsilon} \in \mathcal{E}(\bar{\mathbf{x}})} s(\bar{\epsilon}, \bar{\mathbf{x}}) \\ &= p_{\text{del}}^{|\mathbf{t}(\bar{\mathbf{x}})| - |\mathbf{t}(\mathbf{x})|} \left(p_y(\mathbf{x}) - \sum_{\epsilon \notin \mathcal{E}(\mathbf{x}, \epsilon^*)} s(\epsilon, \mathbf{x}) \right) \\ &\quad + \sum_{\bar{\epsilon} \notin \mathcal{E}(\bar{\mathbf{x}}, \bar{\epsilon}^*)} s(\bar{\epsilon}, \bar{\mathbf{x}}). \end{aligned} \quad (7)$$

We can bound the sums in the above expression using the following result.

Lemma 4. *We have*

$$0 \leq \sum_{\epsilon \notin \mathcal{E}(\mathbf{x}, \epsilon^*)} s(\epsilon, \mathbf{x}) \leq 1 - p_{\text{del}}^{\sum_i \epsilon_i^*}.$$

Proof. The lower bound is straightforward, since the summand is non-negative. For the upper bound, observe that

$$\begin{aligned} &\sum_{\epsilon \notin \mathcal{E}(\mathbf{x}, \epsilon^*)} s(\epsilon, \mathbf{x}; h) \\ &= 1 - \sum_{\epsilon \in \mathcal{E}(\mathbf{x}, \epsilon^*)} s(\epsilon, \mathbf{x}; h) \\ &\leq 1 - \sum_{\epsilon \in \mathcal{E}(\mathbf{x}, \epsilon^*)} q(\epsilon) \\ &= 1 - p_{\text{del}}^{\sum_i \epsilon_i^*} \sum_{\epsilon \in \mathcal{E}(\mathbf{x}, \epsilon^*)} q(\epsilon - \epsilon^*) \\ &\leq 1 - p_{\text{del}}^{\sum_i \epsilon_i^*} \end{aligned}$$

□

Combining these results yields upper and lower bounds on the smoothed classifier's score for neighboring text $\bar{\mathbf{x}}$ to input text \mathbf{x} .

Theorem 1 (General pairwise certificate). *Consider a pair of text inputs $\mathbf{x}, \bar{\mathbf{x}} \in \mathcal{X}$. Suppose $\bar{\mathbf{x}}$ can be transformed into \mathbf{x} using a minimal number of edit operations by deleting n_{del} tokens, inserting n_{ins} tokens and substituting n_{sub} tokens—i.e., $\text{dist}_{\mathbf{o}, \mathbf{t}}(\bar{\mathbf{x}}, \mathbf{x}) = n_{\text{sub}} + n_{\text{ins}} + n_{\text{del}}$. Then the smoothed classifier's scores for any class $y \in \mathcal{Y}$ satisfy*

$$\begin{aligned} p_{\text{del}}^{n_{\text{del}} - n_{\text{ins}}} \left(p_y(\mathbf{x}) - 1 + p_{\text{del}}^{n_{\text{sub}} + n_{\text{ins}}} \right) &\leq p_y(\bar{\mathbf{x}}) \\ &\leq p_{\text{del}}^{n_{\text{del}} - n_{\text{ins}}} p_y(\mathbf{x}) + 1 - p_{\text{del}}^{n_{\text{sub}} + n_{\text{del}}}. \end{aligned}$$

Proof. We obtain upper and lower bounds on the expression for $p_y(\bar{\mathbf{x}})$ in (7) using Lemma 4. Replacing the sum over $\bar{\epsilon}$ by a lower bound and the sum over ϵ by an upper bound yields:

$$p_y(\bar{\mathbf{x}}) \geq p_{\text{del}}^{|\mathbf{t}(\bar{\mathbf{x}})| - |\mathbf{t}(\mathbf{x})|} \left(p_y(\mathbf{x}) - 1 + p_{\text{del}}^{\sum_i \epsilon_i^*} \right).$$

Similarly, replacing the sum over $\bar{\epsilon}$ by an upper bound and the sum over ϵ by a lower bound yields:

$$p_y(\bar{\mathbf{x}}; h) \leq p_{\text{del}}^{|\mathbf{t}(\bar{\mathbf{x}})| - |\mathbf{t}(\mathbf{x})|} p_y(\mathbf{x}) + 1 - p_{\text{del}}^{\sum_i \bar{\epsilon}_i^*}$$

The final result is obtained by observing $|\mathbf{t}(\mathbf{x})| = |\mathbf{t}(\bar{\mathbf{x}})| + n_{\text{ins}} - n_{\text{del}}$, $\sum_i \bar{\epsilon}_i^* = n_{\text{sub}} + n_{\text{del}}$ and $\sum_i \epsilon_i^* = n_{\text{sub}} + n_{\text{ins}}$. □

Finally we extend the pairwise certificate to a certificate over a Levenshtein (edit) distance ball.

Theorem 2 (Levenshtein distance certificate). *Consider a text input $\mathbf{x} \in \mathcal{X}$ for which a lower bound on the smoothed classifier's highest score μ_y and an upper bound on the smoothed classifier's runner-up score $\mu_{y'}$ satisfy $\mu_y \geq \mu_{y'}$. Then the smoothed classifier predicts y for any neighboring text input $\bar{\mathbf{x}} \in \mathcal{X}$ such that $\text{dist}_{\mathbf{o}, \mathbf{t}}(\bar{\mathbf{x}}, \mathbf{x}) \leq r$ with $\mathbf{o} = \{\text{del}, \text{ins}, \text{sub}\}$ and $r = \lfloor \log_{p_{\text{del}}} \frac{1}{2} (2 + \mu_{y'} - \mu_y) \rfloor$. If the upper and lower bounds hold jointly with confidence $1 - \alpha$, then the certificate holds with probability $1 - \alpha$.*

Proof. By definition, the smoothed classifier is robust in the Levenshtein distance neighborhood $B_r(\mathbf{x}; \mathbf{o}, \mathbf{t})$ iff the difference between the score for the predicted class y and any other class $y'' \neq y$ is positive for all $\bar{\mathbf{x}} \in B_r(\mathbf{x}; \mathbf{o}, \mathbf{t})$:

$$\min_{\bar{\mathbf{x}} \in B_r(\mathbf{x}; \mathbf{o}, \mathbf{t})} \left\{ p_y(\bar{\mathbf{x}}) - \max_{y'' \neq y} p_{y''}(\bar{\mathbf{x}}) \right\} > 0. \quad (8)$$

This condition is satisfied if a lower bound on the LHS is positive. We obtain a lower bound on the LHS using bounds on $p_y(\bar{\mathbf{x}})$ from Theorem 1 and bounds on $p_{y'}(\mathbf{x})$ from the theorem statement:

$$\text{LHS of (8)} \geq \min_{\substack{n_{\text{del}}, n_{\text{ins}}, n_{\text{sub}} \geq 0 \\ \text{s.t. } n_{\text{del}} + n_{\text{ins}} + n_{\text{sub}} \leq r}} \psi(n_{\text{del}}, n_{\text{ins}}, n_{\text{sub}})$$

where the objective is

$$\begin{aligned} \psi(n_{\text{del}}, n_{\text{ins}}, n_{\text{sub}}) &= p_{\text{del}}^{n_{\text{del}} - n_{\text{ins}}} (\mu_y - 1 + p_{\text{del}}^{n_{\text{sub}} + n_{\text{ins}}}) \\ &\quad - p_{\text{del}}^{n_{\text{del}} - n_{\text{ins}}} \mu_{y'} - 1 + p_{\text{del}}^{n_{\text{sub}} + n_{\text{del}}}. \end{aligned}$$

It is straightforward to show that this objective is monotonically decreasing in n_{sub} , hence the minimum occurs at $(n_{\text{del}}, n_{\text{ins}}, n_{\text{sub}}) = (0, 0, r)$. Recalling that the lower bound $\psi(0, 0, r)$ must be positive

to guarantee robustness, and solving for r , yields

$$r \leq \log_{p_{\text{del}}} (2 + \mu_{y'} - \mu_y/2). \quad (9)$$

Enforcing the constraint that r is an integer yields the required result. \square

B Parameter settings

We fine-tune the RoBERTa model for the non-certified baseline, CERT-ED and RanMASK on the respective training datasets. The default parameter settings for the experiments are shown in Table 5. We do not explicitly calibrate the optimizer or training schedule for each model, as we find the default settings work well across all datasets. When approximating the smoothed models (CERT-ED and RanMASK) we use a Monte Carlo sample of 1000 perturbed inputs for prediction and 4000 perturbed inputs for certification, while setting the significance level to 0.05.

C Certified robustness

C.1 Certification statistics

We first present results of CERT-ED and RanMASK on the AG-News, LUN, and SatNews datasets (Table 6). In general, we see a minor drop in clean accuracy with increasing perturbation strength for both methods. Although in one case RanMASK suffers a catastrophic drop in accuracy to to 50% for IMDB with a perturbation strength of 95%. Surprisingly on SatNews, the smoothed classifiers achieve a higher clean accuracy than the baseline model. While smoothing does not significantly impact accuracy on the AG-News, LUN and SatNews datasets, it does have a pronounced impact on the IMDB dataset. This is likely due to the fact that the IMDB sentiment classification task is more sensitive to small perturbations such as names (Prabhakaran et al., 2019).

We observe that CERT-ED dominates RanMASK in terms of clean accuracy and certified cardinality for 4 out of 5 datasets. However, for the shorter AG-News dataset with an average input length of 37.84 words, the results are more mixed, with RanMASK coming out on top when the perturbation strength is above 90%. This is to be expected, as masking preserves more spatial information compared to deletion, and it becomes advantageous when the input text sequence is shorter. We provide a more detailed analysis in Section C.4.

C.2 Certified accuracy plots

Figure 3 plots the certified accuracy as a function of the log certified cardinality and perturbation strength for 4 datasets: AG-News, IMDB, Spam-assassin, and LUN. The corresponding plot for SatNews is featured in Figure 2. Note that we are varying the cardinality of the certificate, rather than the radius, so that we can reasonably compare the size of CERT-ED and RanMASK certificates, which are defined using different distance metrics (Levenshtein distance for CERT-ED, and Hamming distance for RanMASK). We compute the cardinality exactly for Hamming distance and use a lower bound for Levenshtein distance, which typically underestimates the exact value by 1 order of magnitude. Similar to the results in Table 3, CERT-ED outperforms RanMASK on both LUN and SatNews for all perturbation strengths. For AG-News, the results are more mixed with RanMASK outperforming CERT-ED at a perturbation strength of 90%.

C.3 Impact of deletion rate

We present a comprehensive ablation study of CERT-ED with varying deletion rates in Table 7. We observe that for shorter sequences, the deletion rate has a more significant impact on the clean accuracy. For IMDB, the maximum possible certified cardinality is much lower compared to other datasets, further demonstrating the sensitivity of sentiment classification to perturbations. Being the simplest dataset, Spam-assassin has the highest certified cardinality among all datasets.

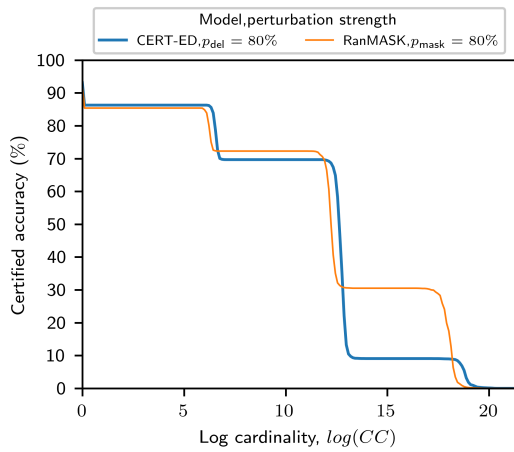
C.4 Impact of input text length

The results of clean accuracy and certified cardinality, grouped by text length quintiles, are shown in Table 8. For the AG-News dataset, which has a shorter average text length, RanMASK suffers a minor drop in clean accuracy from Q1 to Q2, while CERT-ED is impacted more significantly. This supports our hypothesis that RanMASK is more advantageous when the input text sequence is shorter.

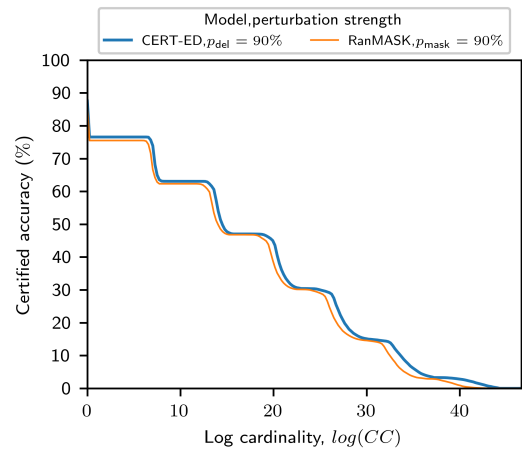
D Details on attacks

D.1 Attack setup

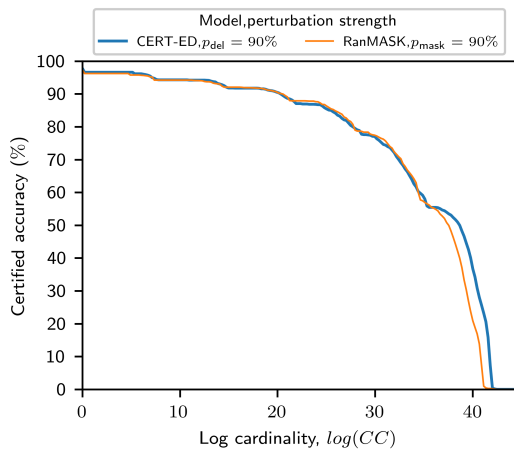
Attacks covered We evaluate empirical robustness under various attacks using a modified version of TextAttack (Morris et al., 2020) and attack templates implemented by Zhang et al. (2024a).



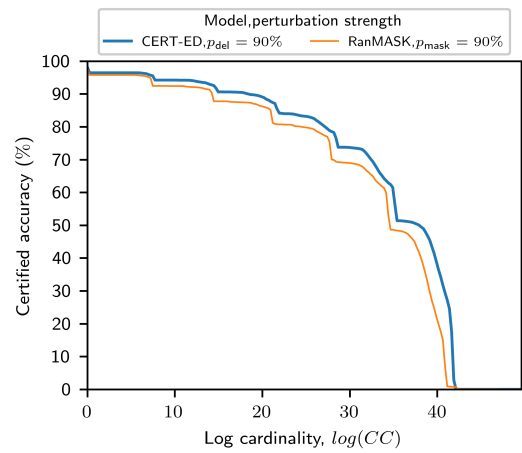
(a) AG-News dataset



(b) IMDB dataset



(c) Spam-assassin dataset



(d) LUN dataset

Figure 3: Certified accuracy for CERT-ED and RanMASK as a function of log certified cardinality and perturbation strength p_{del} and p_{mask} (line styles). The certified cardinality is exact for RanMASK but a lower bound is used for CERT-ED. CERT-ED dominates RanMASK in terms of certified accuracy for 3 out of 4 datasets. See Figure 2 for certified accuracy on SatNews.

	Parameter	Values
Base model	Model	AutoModelForSequenceClassification("roberta-base")
	Tokenizer	AutoTokenizer("roberta-base")
Scheduler	Python command	transformers.get_linear_schedule_with_warmup
	Warmup epochs	10
Optimizer	Python class	torch.optim.AdamW
	Learning rate	2.0E-5
	Weight decay	1.0E-6
	Gradient clipping	clip_grad_norm_(model.parameters(), 1.0)
Training	Batch size	32
	Max. epoch	200
	Early stopping	No improvement in validation loss after 25 epochs

Table 5: Parameter settings for RoBERTa, the optimizer and training procedure. Parameter settings are consistent across all models (Baseline, CERT-ED, RanMASK) except where specified.

We select five representative attacks which can be categorized as follows:

- Clare (Li et al., 2021) and BAE-I (Garg and Ramakrishnan, 2020) both cover a token-wise edit distance threat model;
- BERT-Attack (Li et al., 2020) and TextFooler (Jin et al., 2020) operate by substituting words in the input text;
- DeepWordBug (Gao et al., 2018) modifies the input text by altering characters within each word.

Attack results Each attack can yield one of four distinct outcomes, namely, *success*, *fail*, *skipped* or *timeout*. Their meanings are as follows:

- *success* indicates the attack was able to generate an adversarial example by perturbing the prediction from the correct label to a false label.
- *fail* indicates the attack was unable to generate an adversarial example. This can happen for the following reasons: firstly, the attack was unable to find a perturbation that changes the prediction; secondly, the attack was unable to find a perturbation that changes the prediction after exhausting all options; or, lastly, the attack reached the maximum number of queries on the target model. We enforce a maximum limit of 10000 queries for Clare, but place no limit on the other attacks.
- *skipped* indicates the attack was skipped because the model’s prediction was incorrect in the first place.

- *timeout* indicates the attack was skipped because it took too long to generate an adversarial example. In our experiments, we set the timeout to be 600 seconds. Note that this generally puts CERT-ED at a disadvantage compared to RanMASK, because CERT-ED can process queries roughly 3 times faster than RanMASK. Unless otherwise specified, we treat timeout as fail.

The robust accuracy is defined as the fraction of instances for which the attack outcome is either *fail* or *timeout*—i.e., the fraction of instances for which the model’s prediction remains correct after the attack.

D.2 Transfer attack on ag-news

We perform transfer attacks by applying successful attack examples against the non-certified baseline to the smoothed models (CERT-ED, RanMASK). Table 9 reports both clean and robust accuracy since the successful example against each attack will be slightly different. Unlike the direct attack results, CERT-ED does not have a strict dominance over RanMASK against BERT-Attack, TextFooler, and DeepWordBug. While for Clare and BAE-I, CERT-ED is marginally better rather than tied. These results demonstrate the robustness of CERT-ED to transfer attacks, showing that CERT-ED is more robust to adversarial examples for the non-certified baseline than RanMASK.

D.3 Query constrained attack results

To complement the results for time-constrained attacks in Table 4, we report results for query-

Model	$p_{\text{del}}/p_{\text{mask}}$	Clean Accuracy	Median CR	Median log CC
AG-News dataset (avg length 37.84)				
Baseline	—	94.84%	—	—
RanMASK	80%	93.91%	2	12.25
	90%	92.43%	4	22.91
	95%	88.86%	4	23.67
CERT-ED	80%	93.33%	2	12.65
	90%	91.72%	3	18.65
	95%	87.75%	3	18.65
IMDB dataset (avg length 231.16)				
Baseline	—	93.47%	—	—
RanMASK	80%	90.23%	1	7.41
	90%	86.87%	2	14.02
	95%	50.00%	2	9.40
CERT-ED	80%	89.60%	1	7.50
	90%	88.26%	2	14.49
	95%	85.58%	3	21.20
Spam-assassin dataset (avg length 228.16)				
Baseline	—	98.02%	—	—
RanMASK	80%	97.86%	3	19.99
	90%	97.65%	6	37.49
	95%	96.05%	11	67.06
CERT-ED	80%	97.81%	3	20.63
	90%	97.81%	6	38.66
	95%	97.81%	10	67.04
LUN dataset (avg length 269.93)				
Baseline	—	99.16%	—	—
RanMASK	80%	98.67%	3	19.73
	90%	97.91%	6	34.51
	95%	95.62%	10	60.45
CERT-ED	80%	98.85%	3	20.62
	90%	98.28%	6	37.94
	95%	96.11%	10	61.44
SatNews dataset (avg length 384.84)				
Baseline	—	94.22%	—	—
RanMASK	80%	93.10%	2	14.30
	90%	92.09%	4	27.84
	95%	90.10%	7	47.09
CERT-ED	80%	95.60%	2	14.83
	90%	93.18%	5	35.08
	95%	92.07%	8	54.76

Table 6: Full certification results supplementing Table 3. All metrics are computed using the entire test set. “Median CR” is the median certified Levenshtein distance radius and “median log CC” is the median log-certified cardinality. The certified cardinality is exact for RanMASK, however a lower bound is used for CERT-ED. CERT-ED outperforms RanMASK in terms of certified accuracy for 4 out of 5 datasets and 22 out of 30 metrics, and it specifically excels on datasets with longer average text length.

CERT-ED p_{del}	Clean Accuracy	Median CR	Median log CC
AG-News dataset (avg length 37.84)			
50%	94.76	0	0.00
60%	95.07	1	6.55
70%	94.58	1	6.56
80%	93.33	2	12.65
90%	91.72	3	18.65
95%	87.75	3	18.65
99%	25.05	0	0.00
IMDB dataset (avg length 231.16)			
50%	94.34	0	0.00
60%	93.16	1	7.11
70%	93.00	1	7.14
80%	89.60	1	7.50
90%	88.26	2	14.49
95%	85.58	3	21.20
99%	68.02	6	39.64
Spam-assassin dataset (avg length 228.16)			
50%	98.19	0	0.00
60%	98.49	1	7.32
70%	98.23	1	7.32
80%	97.81	3	20.63
90%	97.81	6	38.66
95%	97.81	10	67.04
99%	84.19	25	152.84
LUN dataset (avg length 269.93)			
50%	99.29	0	0.00
60%	99.33	1	7.48
70%	99.40	1	7.48
80%	98.85	3	20.62
90%	98.28	6	37.94
95%	96.11	10	61.44
99%	86.89	19	118.78
SatNews dataset (avg length 384.84)			
50%	94.97	0	0.00
60%	96.53	1	7.59
70%	95.03	1	7.59
80%	95.60	2	14.83
90%	93.18	5	35.08
95%	92.07	8	54.76
99%	85.81	16	104.73

Table 7: Ablation study of CERT-ED with varying deletion rates p_{del} . All metrics are computed using the entire test set. “Median CR” is the median certified Levenshtein distance radius and “median log CC” is the median log-certified cardinality. The certified cardinality is estimated using a lower bound. For datasets with longer text, the deletion rate has less impact on accuracy.

Quartile	Avg. length	Baseline		RanMASK 80%		RanMASK 90%		CERT-ED 80%		CERT-ED 90%	
		CIA %	log CC	CIA %	log CC	CIA %	log CC	CIA %	log CC	CIA %	log CC
AG-News dataset											
Q1	27.25	93.62	—	92.33	11.95	90.85	17.71	91.22	12.39	89.37	12.60
Q2	35.03	94.93	—	93.90	12.18	92.99	17.99	93.90	12.65	92.54	18.58
Q3	40.38	95.21	—	94.00	12.32	92.33	23.72	93.74	12.79	92.43	18.81
Q4	50.74	95.88	—	95.82	12.49	94.00	24.17	94.94	12.93	93.06	24.72
IMDB dataset											
Q1	92.32	94.74	—	92.12	6.79	89.56	13.18	90.76	7.05	88.94	13.48
Q2	147.20	95.46	—	92.56	13.31	89.06	13.55	90.51	7.21	89.20	13.98
Q3	215.62	94.17	—	90.91	13.59	86.88	20.05	89.30	7.41	87.61	14.40
Q4	461.22	89.51	—	85.31	7.40	81.93	14.42	87.83	14.51	87.28	21.61
Spam-assassin dataset											
Q1	74.23	98.01	—	97.51	18.59	97.84	31.80	98.18	13.78	97.68	30.93
Q2	171.62	99.32	—	99.49	19.97	99.49	38.35	99.15	20.54	99.49	39.10
Q3	281.63	96.80	—	96.80	20.55	96.13	39.50	96.63	21.14	97.14	40.51
Q4	395.60	97.97	—	97.64	21.07	97.13	40.72	97.30	21.70	96.96	41.70
LUN dataset											
Q1	84.65	99.26	—	98.64	18.62	98.21	35.79	98.70	16.87	97.16	32.02
Q2	226.05	99.57	—	99.57	20.24	99.19	39.07	99.32	20.83	99.50	39.88
Q3	363.68	98.32	—	97.58	14.26	96.40	34.24	98.07	21.55	97.20	41.21
Q4	404.00	99.50	—	98.88	21.09	97.83	34.43	99.32	21.72	99.26	41.75
SatNews dataset											
Q1	353.34	95.23	—	96.12	14.24	94.51	27.71	97.06	14.73	95.23	34.93
Q2	384.64	94.97	—	93.76	14.28	92.87	27.77	95.69	14.76	93.98	35.11
Q3	398.19	93.32	—	92.05	14.30	91.20	27.82	95.60	14.79	92.10	35.19
Q4	416.96	93.36	—	90.36	14.34	89.66	27.91	93.95	14.86	91.36	35.27

Table 8: Clean accuracy (CIA) and median log-certified cardinality (log CC) for RanMASK and CERT-ED as a function of perturbation strength grouped by input text length quintiles. The entire test set is used to compute all metrics. Q1 refers to the first quintile while Q4 refers to the fourth quintile. The certified cardinality is exact for RanMASK but a lower bound is used for CERT-ED.

Method	Clare		BAE-I		BERT-Attack		TextFooler		DeepWordBug	
	CIA %	RoA %	CIA %	RoA %	CIA %	RoA %	CIA %	RoA %	CIA %	RoA %
AG-News dataset										
RanMASK	93.51	90.41	89.50	86.46	93.30	88.28	94.08	90.77	91.41	82.81
CERT-ED	93.66	90.27	90.06	86.74	93.46	89.04	94.21	89.67	91.60	81.45
IMDB dataset										
RanMASK	79.33	75.51	85.31	82.31	87.73	81.66	87.67	77.30	84.26	76.54
CERT-ED	82.02	75.28	86.81	82.76	89.02	83.06	88.94	78.57	85.96	76.70
Spam-assassin dataset										
RanMASK	94.92	91.53	95.74	91.49	98.49	95.27	98.21	94.47	96.57	88.57
CERT-ED	94.92	89.83	95.74	89.36	98.71	96.34	98.70	94.31	98.86	93.14
LUN dataset										
RanMASK	92.16	90.20	95.29	93.72	97.14	88.57	96.53	87.86	94.20	86.47
CERT-ED	96.08	94.12	96.86	96.86	98.33	94.76	98.27	92.49	96.62	91.79
SatNews dataset										
RanMASK	72.16	69.07	86.15	83.38	88.48	82.95	87.95	82.05	85.59	79.88
CERT-ED	80.41	77.32	89.20	86.98	91.24	87.33	90.45	87.50	88.89	86.19

Table 9: Empirical attack results when transferring *successful* adversarial examples against the non-certified baseline to CERT-ED and RanMASK. Both CERT-ED and RanMASK use a perturbation strength of 90%. Clean and robust accuracy are abbreviated CIA and RoA, respectively. Highlighted values are the best in each column for that dataset. CERT-ED outperforms RanMASK in all word substitution and character-level attacks.

Method	Clean Accuracy %	Clare BoA %	BAE-I BoA %
AG-News dataset			
Baseline	94.5	71.0	64.0
RanMASK	95.0	85.0	83.0
CERT-ED	94.5	83.5	83.5
IMDB dataset			
Baseline	94.0	67.0	70.0
RanMASK	88.0	68.0	68.0
CERT-ED	88.0	67.5	69.0

Table 10: Empirical attack results against Baseline, RanMASK, and CERT-ED under attacks where the number of queries is constrained to 3000.

constrained attacks in Table 10. We limit the number of queries to 3000 and run the attacks on a smaller subset of 200 instances. We focus on the two time-constrained attacks where our model is less competitive in Table 4: Clare and BAE-I for AG-News and IMDB. We see that, for the query-constrained setting, CERT-ED is significantly more robust compared to the time-constrained setting. In particular, the performance gaps between RanMASK and CERT-ED in Table 10 are diminished compared to Table 4.

E Efficiency and computation requirements

In this appendix, we document the computation requirements to train, certify, and attack models used in our work. We also compare and contrast the efficiency of CERT-ED and RanMASK in terms of training and certification. We show that CERT-ED is more efficient than RanMASK in both aspects.

E.1 Hardware

All experiments in this paper are conducted using a private cluster with Intel(R) Xeon(R) Gold 6326 CPU @ 2.90GHz and NVIDIA A100 GPUs. Unless otherwise specified, we use a single GPU for all

experiments.

E.2 Train

Table 11 shows the number of epochs used to train each model/dataset (with early stopping) and the training time per epoch. CERT-ED is about 2–3 times faster to train than the non-smoothed baseline, and 2–5 times faster to train than RanMASK. The total computation used across all datasets for certification is estimated to be 70 hours A100 GPU time.

E.3 Certification

Table 12 shows the average certification time per test instance, including overheads. We see CERT-ED is about 3 times faster than RanMASK on average across all datasets. The total computation used across all datasets for certification is estimated to be 250 hours A100 GPU time.

E.4 Empirical robustness

Table 13 reports the average attack time per instance for a subset of the IMDB test set. Attack times on other datasets follow a similar pattern. We note that the time taken to attack RanMASK is longer than for the non-smoothed baseline and CERT-ED. Combined with the max timeout window of 10 minutes, this partially explains the lower robust accuracy of CERT-ED compared to RanMASK in Table 4. We utilized parallelized attacks to speed up the process. The total computation used across all datasets for attacks is estimated to be 200 days A100 GPU time.

F Edit distance certificates for Text-CRS

In this appendix, we analyze two robustness certificates from the Text-CRS framework (Zhang et al., 2024a), which cover deletion/insertion perturbations to text represented as a sequence of token embedding vectors. Each certificate is parameterized by two radii: one bounds the perturbation to the token embedding vectors and the other bounds the extent of token-level reordering. We obtain lower bounds on these two radii such that each certificate covers up to r arbitrary edits of a single type (deletion/insertion). Using these bounds, we can immediately convert a Text-CRS deletion/insertion certificate to an edit distance certificate where the allowed edit operations are deletions/insertions to input tokens. We find that the resulting edit distance certificates are vacuous ($r = 0$) for sequences greater than 2 tokens in length when instantiated

with the Text-CRS smoothing mechanisms. As a result, we have opted not to include Text-CRS as a baseline in our experiments.

F.1 Preliminaries

Text-CRS indirectly bounds edits to input text by instead bounding numerical additive perturbations and permutations in word embedding space. This is not in one-to-one correspondence as we see input edits are not easily bounded. Concretely, input text is represented as an array of embedding vectors $\mathbf{w} \in \mathbb{R}^{n \times k}$, where the first dimension corresponds to words and the second dimension corresponds to dimensions in the embedding space. Input text with $m < n$ actual words is represented by filling the last $n - m$ rows with padding words/vectors. Robustness is studied under input transformations that are a composition of: (1) perturbations to the embedding vectors and (2) word-level permutations of the embedding vectors. Here we define notation to represent these transformations.

Embedding perturbations An embedding perturbation is an array δ of the same type as the input \mathbf{w} . The result of applying δ to \mathbf{w} is simply $\mathbf{w} + \delta$. We consider two norms to measure the magnitude of the perturbation:

- $\|\delta\|_0 := \sum_{i=1}^n \mathbf{1}_{\sum_{j=1}^k |\delta_{i,j}| \neq 0}$ is the sum of non-zero rows in δ , i.e., the number of perturbed words/vectors; and
- $\|\delta\|_2 := \sqrt{\sum_{i=1}^n \delta_{i,j}^2}$ is the Frobenius norm.

Permutations A word-level permutation of an input \mathbf{w} is parameterized by a permutation matrix $\pi \in \mathcal{P}_n$. Here $\mathcal{P}_n = \{\pi \in \{0, 1\}^{n \times n} : \sum_{i'=1}^n \pi_{i',j} = 1, \sum_{j'=1}^n \pi_{i,j'} = 1 \forall i, j\}$ denotes the set of $n \times n$ permutation matrices. The result of applying π to \mathbf{w} is simply $\pi \cdot \mathbf{w}$ where \cdot denotes matrix multiplication. The magnitude of the perturbation $\|\pi\|_1$ is measured in terms of the ℓ_1 distance between the new word locations and the original locations. We can equivalently express this as the row-wise sum of the absolute distance of each 1 from the diagonal in the permutation matrix π : $\|\pi\|_1 = \sum_{i=1}^n |i - \sum_{j=1}^n j \mathbf{1}_{\pi_{i,j}=1}|$.

Composition The composition of the embedding perturbation δ and permutation π is an input transformation $T_{\delta,\pi} : \mathbb{R}^{n \times d} \rightarrow \mathbb{R}^{n \times d}$ such that $T_{\delta,\pi}(\mathbf{w}) = \pi \cdot (\mathbf{w} + \delta)$.

Dataset	Train #samples	Baseline		RanMASK, 90%		CERT-ED, 90%	
		epochs	sec/epoch	epochs	sec/epoch	epochs	sec/epoch
AG-News	108 000	65	517	105	476	100	231
IMDB	22 500	30	258	60	341	65	128
Spam-assassin	2 152	40	27	50	35	40	13
LUN	13 416	55	143	65	258	60	55
SatNews	22 738	55	260	80	461	95	101

Table 11: Training time statistics for each dataset and model. The number of epochs varies due to early stopping.

Dataset	Test #samples	RanMASK, 90%	CERT-ED, 90%
		ms / sample	ms/sample
AG-News	7 600	3 969	2 367
IMDB	25 000	13 331	3 311
Spam-assassin	2 378	13 899	3 319
LUN	6 454	14 641	4 819
SatNews	7 202	17 767	5 778

Table 12: Certification time on the test set for each dataset, including overheads. We use 1000 Monte Carlo samples for prediction and 4000 samples for estimating certified radii. During attacks, we use 100 samples for prediction, which cuts the prediction time by 1/40 ignoring overheads.

Dataset	Baseline sec / sample	RanMASK, 90% sec / sample	CERT-ED, 90% sec / sample
Clare	192	527	504
BAE-I	268	502	461
BERT-Attack	34	333	234
TextFooler	12	302	173
DeepWordBug	7	155	62

Table 13: Attack time per instance on a subset of 1000 instances from the IMDB test set. The timeout window is set to be 10 minutes. Note that most Clare and BAE-I attacks targeting RanMASK timed out. This partially explains the lower robust accuracy of CERT-ED compared to RanMASK in Table 4.

F.2 Word-level deletion

Text-CRS covers deletion using a certificate that constrains the number of modified embedding vectors and the sum of word position changes.

Definition 1. A Text-CRS deletion certificate at input $\mathbf{w} \in \mathbb{R}^{n \times d}$ is a set of inputs parameterized by two radii $r_R, r_D \geq 0$:

$$C_D(\mathbf{w}; r_R, r_D) = \{T_{\delta, \pi}(\mathbf{w}) \in \mathbb{R}^{n \times d} : \pi \in \mathcal{P}_n, \delta \in \mathbb{R}^{n \times d}, \|\pi\|_1 < r_R, \|\delta\|_0 < r_D\}.$$

We note that this form of certificate is not tight for deletion. For example, it includes invalid inputs that contain padding in the middle of the sequence, and it also includes inputs where words are replaced by ordinary (non-padding) words.

We are interested in determining values of r_R and r_D such that the Text-CRS deletion certificate covers a standard edit distance certificate constrained to deletions (see (1)). This will allow us to compare Text-CRS and CERT-ED.

Proposition 5. *The Text-CRS deletion certificate contains a deletion-based edit distance certificate for any input $\mathbf{w} \in \mathbb{R}^{n \times d}$, meaning $C_D(\mathbf{w}; r_R, r_D) \supseteq B_r(\mathbf{w}; \{\text{del}\})$, if*

$$r_D = r \text{ and } r_R \geq \begin{cases} 2r(n-r), & n \geq 2r, \\ n^2/2, & n < 2r. \end{cases}$$

Proof. Let $\mathbf{w}' \in B_r(\mathbf{w}; \{\text{del}\})$ be an input obtained from \mathbf{w} by deleting $l \leq r$ elements. We observe that \mathbf{w}' requires the greatest sum of word position changes (as measured by $\|\pi\|_1$ for the permutation matrix π) when l elements are deleted at the beginning of the sequence and \mathbf{w} contains no padding words at the end. In this case the permutation matrix is

$$\pi = \begin{pmatrix} 0 & I_l \\ I_{n-l} & 0 \end{pmatrix}$$

with $\|\pi\|_1 = 2l(n-l)$. Taking the worst-case number of deletions $l \leq r$, we have

$$r_R \geq \max_{l \in \{0, \dots, r\}} 2l(n-l) = \begin{cases} 2r(n-r), & n \geq 2r, \\ n^2/2, & n < 2r. \end{cases}$$

□

Zhang et al. (2024a) instantiate the Text-CRS deletion certificate for a smoothed classifier where the smoothing mechanism permutes the embedding vectors uniformly at random and randomly replaces

embedding vectors with padding with fixed probability p . For this mechanism, the largest possible value of r_R is n , which is achieved when the classifier's confidence is 100%. Combining $r_R \leq n$ with the inequality in Proposition 5 implies

$$n \geq \begin{cases} 2r(n-r), & n \geq 2r, \\ n^2/2, & n < 2r. \end{cases} \Leftrightarrow r \leq \begin{cases} n, & n \leq 2, \\ 0, & n > 2. \end{cases}$$

Hence the edit distance certificate is vacuous ($r = 0$) when the maximum sequence length $n > 2$.

F.3 Word-level insertion

Text-CRS covers insertion using a certificate that constrains the perturbation of the embedding vectors (in ℓ_2 -distance) and the sum of word position changes.

Definition 2. A Text-CRS insertion certificate at input $\mathbf{w} \in \mathbb{R}^{n \times d}$ is a set of inputs parameterized by two radii $r_R, r_I \geq 0$:

$$C_I(\mathbf{w}; r_R, r_I) = \{T_{\delta, \pi}(\mathbf{w}) \in \mathbb{R}^{n \times d} : \pi \in \mathcal{P}_n, \delta \in \mathbb{R}^{n \times d}, \|\pi\|_1 < r_R, \|\delta\|_2 < r_I\}.$$

We are interested in determining values of r_R and r_I such that the Text-CRS insertion certificate covers a standard edit distance certificate constrained to insertions (see (1)).

Proposition 6. *Let E denote the set of d -dimensional embedding vectors (covering all possible words) and let $D_\star := \max_{\mathbf{e}_1, \mathbf{e}_2 \in E} \|\mathbf{e}_1 - \mathbf{e}_2\|_2$. The Text-CRS insertion certificate contains an insertion-based edit distance certificate for any input $\mathbf{w} \in E^n$, meaning $C_I(\mathbf{w}; r_R, r_I) \supseteq B_r(\mathbf{w}; \{\text{ins}\})$, if*

$$r_I \geq \sqrt{r} D_\star \text{ and } r_R \geq 2r(n-r), \quad \text{when } n \geq 2r, \\ \sqrt{n/2} D_\star \text{ and } r_R \geq n^2/2, \quad \text{when } n < 2r.$$

Proof. Let $\mathbf{w}' \in B_r(\mathbf{w}; \{\text{ins}\})$ be an input obtained from \mathbf{w} by inserting $l \leq r$ elements. We observe that the corresponding $\|\delta\|_2$ is maximized when the l inserted vectors are a distance D_\star away from the l vectors at the end of the original input \mathbf{w} . In this case $\|\delta\|_2 = \sqrt{l} D_\star = \sqrt{l} D_\star$.

We observe that \mathbf{w}' requires the greatest sum of word position changes (as measured by $\|\pi\|_1$ for the permutation matrix π) when the l elements are inserted at the beginning of the sequence. In this case the permutation matrix is

$$\pi = \begin{pmatrix} 0 & I_{n-l} \\ I_l & 0 \end{pmatrix}$$

with $\|\boldsymbol{\pi}\|_1 = 2l(n-l)$.

Taking the worst case $\|\boldsymbol{\pi}\|_1$ with respect to the number of insertions $l \leq r$ we have

$$r_R \geq \max_{l \in \{0, \dots, r\}} 2l(n-l) = \begin{cases} 2r(n-r), & n \geq 2r, \\ n^2/2, & n < 2r, \end{cases}$$

where the maximizer is $l = r$ for the first case and $l = n/2$ for the second case. Hence we have

$$r_I \geq \begin{cases} \sqrt{r}D_\star, & n \geq 2r, \\ \sqrt{n/2}D_\star, & n < 2r. \end{cases}$$

□

Zhang et al. (2024a) instantiate the Text-CRS insertion certificate for a smoothed classifier where the smoothing mechanism permutes the embedding vectors uniformly at random and perturbs the embedding vectors with Gaussian noise with fixed scale parameter σ . For this mechanism, the largest possible value of r_R is n , which is achieved when the classifier's confidence is 100%. Combining $r_R \leq n$ with the inequalities in Proposition 6 implies

$$r \leq \begin{cases} \min\{n, \lfloor (r_I/D_\star)^2 \rfloor\} & n \leq 2, \\ 0, & n > 2. \end{cases}$$

For the values of r_I and D_\star reported by Zhang et al., $r_I/D_\star < 1$. Hence the edit distance certificate is vacuous ($r = 0$) for all sequences.