

520 A Appendix

521 A.1 Limitations

522 The datasets used in experiments are mainly the ones used in the text classification task. We notice
523 that there are other types of text datasets, such as question answering and text generation ones. In our
524 experiments, we follow the existing works and test our methods in text classification datasets only.
525 We expect future works to include a discussion on defending adversarial synonym substitution on
526 text datasets of other types.

527 A.2 Broader Impacts

528 The proposed UniT has a positive societal impact because it is able to generate not only accurate but
529 also certified robust predictions on text-related tasks. Nowadays, large language models (LLMs) such
530 as ChatGPT have attracted great attention due to their powerful performance and intriguing interface.
531 However, they are mostly deployed as a black-box service for users to use and lack reliability for
532 their predictions. For UniT, it can provide a robustness guarantee for predictions for tasks such as
533 text classification, and it can be scalable for LLMs. As a result, UniT can positively impact the
534 employment of LLMs in real-world applications and improve the trust between users and LLM
535 service providers.

536 A.3 Dataset Details

537 The used datasets are all in English. Each text sample is tokenized by the “BertTokenizer” provided
538 by the Transformers [19] library. The detailed descriptions of the datasets are as follows:

- 539 1. IMDB [10] is a sentiment analysis dataset for movie reviews with either positive or negative
540 sentiment.² Its text samples have comparatively longer lengths than the ones of the rest
541 datasets. It has 25,000 train and test samples, respectively.
- 542 2. SST2 [15] is another binary text classification dataset for movie reviews. It has 67,349 train
543 samples and 1,821 test samples.³ Its license is CC0.
- 544 3. Yelp [14] is a large-scale sentiment analysis dataset collected from restaurant reviews written
545 by Yelp users with two classes, i.e., positive and negative ones.⁴ It has 444,101 samples for
546 training and 126,670 samples for testing. Its license is the Apache-2.0 license.
- 547 4. AG [22] is a comparably large-scale news classification dataset with 4 classes, including
548 world, sports, science/technology and business.⁵ The number of train and test samples are
549 120,000 and 7,600, respectively.

550 A.4 Calculation of Certified Robust Accuracy

551 Because the existing approaches use different ways to calculate the certified robust accuracy, to
552 make a **fair comparison**, we follow the methods used in different scenarios and compare baselines
553 separately.

- 554 • When comparing with **SAFER**, we follow the same setting that conducts sampling for
555 choosing test samples and construct 5,000 perturbed samples for each test sample through
556 random synonym replacement. The output from the smoothed model is derived by averaging the
557 prediction of the 5,000 perturbed samples for each original text sample. We calculate the
558 certified robust accuracy in this setting for both SAFER and UniT based on the certification
559 condition proposed in Proposition 1 of SAFER [20]. The confidence level is set to 99.0%.
560 We take a text sample in IMDB as an example and show one of the perturbed samples
561 constructed by adversarial synonym substitution in Table 9.

²<https://ai.stanford.edu/~amaas/data/sentiment/>

³<https://www.kaggle.com/datasets/atulanandjha/stanford-sentiment-treebank-v2-sst2>

⁴[https://github.com/shentianxiao/language-style-transfer/tree/master/
data/yelp](https://github.com/shentianxiao/language-style-transfer/tree/master/data/yelp)

⁵<https://www.kaggle.com/datasets/amananandrai/ag-news-classification-dataset>

Sample	Text
Original Text (Label: Negative)	I went and saw this movie last night after being coaxed to by a few friends of mine. I'll admit that I was reluctant to see it because from what I knew of Ashton Kutcher he was only able to do comedy. I was wrong. Kutcher played the character of Jake Fischer very well, and Kevin Costner played Ben Randall with such professionalism. The sign of a good movie is that it can toy with our emotions. This one did exactly that. The entire theater (which was sold out) was overcome by laughter during the first half of the movie, and were moved to tears during the second half. While exiting the theater I not only saw many women in tears, but many full grown men as well, trying desperately not to let anyone see them crying. This movie was great, and I suggest that you go see it before you judge.
Perturbed Sample	I went and saw this <u>cinematographic</u> last <u>nuit</u> after being coaxed to by a few friends of <u>mined</u> . I'll admit that I was <u>loath</u> to see it because from what I knew of Ashton Kutcher he was only able to do comedy. I was <u>awry</u> . Kutcher played the character of Jake Fischer very well, and Kevin Costner played Ben Randall with such professionalism. The sign of a good <u>cinematic</u> is that it can <u>plaything</u> with <u>unser</u> emotions. This one did exactly that. The <u>totaled</u> theatres (which was <u>sale</u> out) was <u>overcoming</u> by laughter during the firstly half of the movie, and were moved to tears during the second half. While <u>exit</u> the theatres I not <u>exclusively</u> saw many <u>daughter</u> in tears, but <u>varying</u> full grown males as well, trying desperately not to letting anyone see them crying. This <u>films</u> was great, and I suggest that you go see it <u>avant</u> you judges.

Table 9: Example of a perturbation sample. We color and underline the perturbed words (corresponding synonyms of original words) in the perturbed sample to demonstrate the difference between these two text samples.

- 562 • When comparing with **CISS**, we follow its setting of using the whole test set for calculating
563 certified robust accuracy. The certified robust accuracy of CISS is calculated based on
564 Algorithm 2 proposed in CISS [24], and we adopt the same certification process except that
565 we adopt Theorem 1 as the certification condition for our unified framework in this setting.
566 We show the pseudocode for prediction and certification under UniT in the Type II setting in
567 Algorithm 1, and its details are discussed in Sec. A.7. Due to computing restrictions, the
568 certification results are calculated from 9,000 perturbed samples constructed in the hidden
569 space by adding Gaussian noise to the original sample embedding. The confidence level is
570 set to 99.9%. During certification, the required inputs are the original text sample and the
571 corresponding synonym set of each original word.

572 A.5 Implementation

573 Since UniT is based on BERT, it has a similar parameter number to the one of BERT, which is 110M.
574 We use the pretrained BERT model “bert-base-uncased” provided by the Transformers [19] library.
575 When we conduct Type I training with UniT, for every dataset, we fine-tune the pretrained model
576 with 3 epochs, which usually takes 10 minutes on an Nvidia A6000 GPU. When we conduct Type II
577 training with UniT, the training takes about 48 hours for both datasets on an Nvidia A100 GPU. For
578 Yelp and AG, we fine-tune the pretrained model with 110 and 200 epochs, respectively.

579 A.6 Hyperparameters

580 The tuning of hyperparameters is not tricky for the DR loss due to their clear interpretation. During
581 training with the DR loss, we set the hyperparameters $\nu = 0.1$ to keep the Gaussian noise relatively
582 small, $\alpha = 0.7$ to allow the margin to increase while penalizing l_2 norm, and $\xi = 0.6$ to allow
583 appropriate relaxation. In addition, while calculating the final loss, we set $\beta = 1$ to make the MR
584 term and the CE loss have equal weight. In the Type II setting, the extra hyperparameters μ and γ
585 have been studied by [24], so we follow them to set $\mu = 1$ and incrementally increase γ to 4 as the
586 training epoch increases.

Algorithm 1: Prediction and certification by UniT in Type II Setting

function PREDICT

Input: Hard-label prediction function based on UniT h , Standard deviation of Gaussian noise σ , Embedding of original text x , Number of Gaussian noise η , Confidence level ω

Process:

Draw η samples of Gaussian noise, add them to x repeatedly, and obtain a vector of class counts CNT for all perturbed inputs;

$\hat{y}_A, \hat{y}_B \leftarrow$ top two indices in CNT;

$\eta_A, \eta_B \leftarrow$ CNT(\hat{y}_A), CNT(\hat{y}_B);

if BINOMPVALUE($\eta_A, \eta_A + \eta_B, 0.5$) $\leq \omega$ **then**

Return \hat{y}_A ; **else**

Return ABSTAIN

end

end**function** CERTIFY

Input: Hard-label prediction function based on UniT h , Standard deviation of Gaussian noise σ , Embedding of original text x , 1st number of Gaussian noise η_1 , 2nd number of Gaussian noise η_2 , Lower confidence $1 - \omega$

Process:

Draw η_1 samples of Gaussian noise, add them to x repeatedly, and obtain a vector of class counts CNT₁ for all perturbed inputs;

$\hat{y}_A \leftarrow$ top index in CNT₁;

Obtain a vector of class counts CNT₂ similarly with η_2 samples of Gaussian noise;

$\underline{p}_A \leftarrow$ LOWERCONFBOUND(CNT₂[\hat{y}_A], $\eta_2, 1 - \omega$);

if $\underline{p}_A > 0.5$ and $\hat{R} \leq \sigma\Phi^{-1}(\underline{p}_A)$ **then**

Return \hat{y}_A ;

end

587 We also include the results of the influence of hyperparameters on the DR loss. Without the loss of
588 generality, we test on the IMDB dataset in the Type I scenario. The results are obtained when we
589 keep the values of other hyperparameters fixed as the ones we use.

590 **Influence of ν .** From Table 10, a comparably small ν is beneficial to modeling the perturbation and
591 loss optimization. Selecting a small ν as 0.05 has already increased the performance compared to
592 that of using the CE loss. As ν grows greater than 0.1, the positive impact of using Gaussian noise to
593 improve the robustness of the classifier module will gradually downgrade. Thus, a comparably small
594 ν as 0.1 is most beneficial.

ν	0.05	0.1	0.15	0.2
Result	87.44	89.04	88.08	86.96

Table 10: Influence of ν on certified robust accuracy (%).

595 **Influence of α .** α is the weight of the negative margin in the DR loss. As we have seen in the ablation
596 study, the introduction of the negative margin contributes to enhancing the base model robustness by
597 regularizing the robustness of the feature extraction and the classifier module. Thus, from Table 11,
598 setting α comparatively high will be helpful for improving the certified robust accuracy. Thus, we
599 can set $\alpha = 0.7$ to help improve the l_2 norm.

α	0.1	0.3	0.5	0.7	0.9
Result	88.08	88.72	88.24	89.04	89.04

Table 11: Influence of α on certified robust accuracy (%).

600 **Influence of β .** As the weight of the MR term, β shall be set approximately equal to 1. From Table 12,
601 we observe that only when β is too large, e.g., $\beta = 2$, can the MR term damage the training of the
602 base model. Therefore, we can just set β to have equal weights with the CE loss to improve the
603 robustness of the base model.

β	0.5	1	1.5	2.0
Result	88.40	89.04	88.24	83.60

Table 12: Influence of β on certified robust accuracy (%).

604 **Influence of ξ .** From Table 13, the introduction of relaxation will be helpful for getting high certified
605 robust accuracy. As shown in Table 13, the certified robust accuracy grows as ξ increases from 0 to
606 0.6 and gradually decreases when ξ gets higher. Therefore, 0.6 is the relaxation hyperparameter we
607 use in our experiment.

ξ	0	0.2	0.4	0.6	0.8	1.0
Result	86.40	86.96	87.92	89.04	88.34	87.44

Table 13: Influence of ξ on certified robust accuracy (%).

608 A.7 Proof for Theorem 1

609 **Proof.** Given the results of Theorem 1 in [1], for all $\|\Delta\| < R$, we have $h(x + \Delta) = y$.
610 Denote u the word embedding function. Recall that $x = [u(w_1), \dots, u(w_n)]$ for the text
611 sample $X = [w_1, \dots, w_n]$, and $x' = [u(w'_1), \dots, u(w'_n)]$ for any allowed perturbed sample
612 $X' = [w'_1, \dots, w'_n] \in A(X)$. Because each word is independent of the others in the embedding
613 space, the l_2 norm between x and any x' is

$$\|x' - x\| = \sqrt{\sum_{i=1}^n \|u(w_i) - u(w'_i)\|^2}. \quad (8)$$

614 Note that for each word $w_i \in X$, the embedding set of its synonyms is $U_i =$
615 $\{u(w_i^{(1)}), \dots, u(w_i^{(m_i)})\}$, and the maximum deviation between w_i and any word $w'_i \in S(w_i)$
616 is $\|v_i\| = \max_{e \in U_i} \|u(w_i) - e\|$. Since each word is independent of the others, the maximum
617 deviation of $\|x' - x\|$ caused by all possible combinations of the synonyms of different words is

$$\hat{R} = \max \|x' - x\| = \sqrt{\|v_1\|^2 + \dots + \|v_n\|^2}. \quad (9)$$

618 Now since we have $\hat{R} \leq R$, we correspondingly have $\|x' - x\| \leq R$ for all $X' \in A(X)$, thus

$$h(x') = h(x + (x' - x)) = h(x) = y, \quad (10)$$

619 for any $X' \in A(X)$ whose corresponding text embedding is x' . Q.E.D.

620 **Remark.** In certification, we also follow the calculation used by [1] and [24] that R is lower
621 bounded by $\sigma\Phi^{-1}(\underline{p}_y)$, where \underline{p}_y is the lower bound of $\mathbb{E}_\delta[h(x + \delta) = y]$ estimated from the
622 Binomial proportion confidence interval. Thus, the certification condition becomes $\underline{p}_y > 0.5$ and
623 $\hat{R} \leq \sigma\Phi^{-1}(\underline{p}_y)$, which is harnessed in [1] and [24] as well.

624 We also show the prediction and certification process in the Type II setting in Algorithm 1. This
625 process mainly follows the same idea as those of [1] and [24]. In Algorithm 1, `BINOMPVALUE`($\eta_A, \eta_A +$
626 η_B, p) returns the p-value of the two-sided hypothesis test that $\eta_A \sim \text{Binomial}(\eta_A + \eta_B, p)$. And
627 `LOWERCONFBOUND`($\kappa, \eta, 1 - \omega$) returns a one-sided $(1 - \omega)$ lower confidence interval for the Binomial
628 parameter p given a sample $\kappa \sim \text{Binomial}(\eta, p)$.

629 A.8 Geometric Interpretation of MR Term

630 Eq. (6) has an explicit geometric interpretation. As shown in Figure 4, after projecting the original
631 sample and perturbed sample in the high-dimensional representation space \mathbf{R}^d (the input space of the
632 last FC layer g), Eq. (6) requires the perturbed representation z' to locate around z within the radius
633 $r = \xi + \alpha \cdot \mathcal{M}(z + \epsilon)$. That is, it tries to guide the original sample and perturbed sample representations
634 to be close to each other and improve the inter-sample compactness of the high-dimensional space
635 with the margin information.




Figure 4: Geometric interpretation of the MR term.




Figure 5: Comparison of perturbed text sample representation distribution. The representations of perturbed samples obtained from DR loss are closer to the original ones than the ones obtained from CE loss.

636 A.9 Visualization of Representation Distribution

637 As mentioned in the Experiment section, we take one test sample as a visualization example to
 638 demonstrate that DR loss is able to improve the representation compactness between perturbed
 639 samples and the original ones. The used text sample is randomly chosen from the IMDB dataset and
 640 the visualization experiment is conducted in the Type I scenario. We now show the visualization
 641 result as follows.

642 Given this randomly chosen text sample, in the certification process, we will need 5,000 perturbed
 643 samples to certify the prediction result. As Figure 5 shows, for all the used perturbed samples for
 644 certification, we first obtain two groups of representations with the CE and DR loss from the feature
 645 extraction module, respectively. We denote the l_2 distance between the original sample representation
 646 z and the representation of any perturbed sample z' as L . We then project those high-dimensional
 647 representations z' into a two-dimensional space based on L : if the l_2 distance and angle of z' with z
 648 is L and ω respectively, the coordinate of z' in Figure 5 is $(L \cos \omega, L \sin \omega)$, and the origin $(0, 0)$
 649 represents z . In this example, the average $\|z' - z\|$ for representations obtained from the CE and
 650 DR loss are 1.768 and 0.353, respectively. From the visualization in Figure 5, the perturbed sample
 651 representation obtained from the DR loss distributes much closer to z with a smaller divergence,
 652 which demonstrates that the representation learned from the DR loss is of higher quality.