

PROBABILISTIC CATEGORICAL ADVERSARIAL ATTACK & ADVERSARIAL TRAINING

Anonymous authors

Paper under double-blind review

ABSTRACT

The existence of adversarial examples brings huge concern for people to apply Deep Neural Networks (DNNs) in safety-critical tasks. However, how to generate adversarial examples with categorical data is an important problem but lack of extensive exploration. Previously established methods leverage greedy search method, which can be very time-consuming to conduct successful attack. This also limits the development of adversarial training and potential defenses for categorical data. To tackle this problem, we propose Probabilistic Categorical Adversarial Attack (PCAA), which transfers the discrete optimization problem to a continuous problem that can be solved efficiently by Projected Gradient Descent. In our paper, we theoretically analyze its optimality and time complexity to demonstrate its significant advantage over current greedy based attacks. Moreover, based on our attack, we propose an efficient adversarial training framework. Through a comprehensive empirical study, we justify the effectiveness of our proposed attack and defense algorithms.

1 INTRODUCTION

Adversarial attacks (Goodfellow et al., 2015) have raised great concerns for the applications of Deep Neural Networks (DNNs) in many security-critical domains (Cui et al., 2019; Stringhini et al., 2010; Cao & Tay, 2001). The majority of existing methods focus on differentiable models and continuous input space, where we can apply gradient-based approaches to generate adversarial examples. However, there are many machine learning tasks where the input data are categorical. For example, data in ML-based intrusion detection systems (Khraisat et al., 2019) contains records of system operations; and in financial transaction systems, data includes categorical information such as the types of transactions. Therefore, how to explore potential attacks and corresponding defenses for categorical inputs is also desired. Existing methods introduce search-based approaches for categorical adversarial attacks (Yang et al., 2020b; Lei et al., 2019a). For example, the method in (Yang et al., 2020a) first finds top- K features of a given sample that have the maximal influence on the model output, and then, a greedy search is applied to obtain the optimal combination of perturbation in these K features. However, these search-based methods cannot be guaranteed to find the strongest adversarial examples. Moreover, they can be computationally expensive, especially when data is high dimensional and the number of categories for each feature is large.

In this paper, we propose a novel *Probabilistic Categorical Adversarial Attack (PCAA)* algorithm to generate categorical adversarial examples by estimating their probabilistic distribution. In detail, given a clean sample, we assume that (each feature of) the adversarial example follows a categorical distribution, and satisfies: (1) the generated samples following this distribution have a high expected loss value and (2) the generated samples only have a few features which are different from the original clean sample. (See Section 3 for more details.) In this way, we transfer the categorical adversarial attack in the discrete space to an optimization problem in a continuous probabilistic space. Thus, we are able to apply gradient-based methods such as (Madry et al., 2017) to find adversarial examples. On one hand, the distribution of adversarial examples in PCAA is searched in the whole space of allowed perturbations. This can facilitate our method to find stronger adversarial examples (with higher loss value) than the greedy search methods (Yang et al., 2020b). Moreover, when the dimension of input data expands, the increase of computational cost of PCAA will be significantly slower than search-based methods (Section 3.4). Therefore, our method can enjoy good attacking optimality and computational efficiency simultaneously. For example, in our experiments

in Section 5.1, our PCAA method is the only attacking method that has the highest (or close to highest) attacking successful rate while maintaining the low computational cost.

The advantages of PCAA allow us to further devise an adversarial training method (PAdvT), by repeatedly (by mini-batches) generating adversarial examples using PCAA. Empirically, PAdvT achieves promising robustness on different datasets over various attacks. For example, on AG’s news dataset, we outperform representative baselines with significant margins, approximately 10% improvement in model robustness compared with the best defense baseline; on IMDB dataset, we obtain comparable performance with ASCC-defense(Dong et al., 2021a), which is a SOTA defense method via embedding space adversarial training. However, compared to ASCC, our method does not rely on the assumption that similar words have a close distance in the embedding space. Compared to other defenses, our defense has much better robustness by more than 15%. Our main contributions can be summarized below:

- We propose a time-efficient probabilistic attacking method (PCAA) for models with categorical input.
- Based on PCAA, we devise a probabilistic adversarial training method to defend categorical adversarial attacks.

2 RELATED WORK

Attacks on categorical data. There has been a rise in the importance of the robustness of machine learning in recent years. On the one hand, evasion attack, poison attack, adversarial training, and other robustness problem with continuous input space have been well studied especially in the image domain (Shafahi et al., 2018; Madry et al., 2017; Ilyas et al., 2019). On the other hand, adversarial attacks focusing on discrete input data, like text data, which have categorical features, are also starting to catch the attention of researchers. Kuleshov et al. (2018) discussed the problems of attacking text data and highlighted the importance of investment into discrete input data. Ebrahimi et al. (2017b) proposed to modify the text token based on the gradient of input one-hot vectors. Gao et al. (2018) developed a scoring function to select the most effective attack and a simple character-level transformation to replace projected gradient or multiple linguistic-driven steps on text data. Samanta & Mehta (2017) proposed an algorithm to generate a meaningful adversarial sample that is legitimate in the text domain. Yang et al. (2020a) proposed a two-stage probabilistic framework to attack discrete data. Lei et al. (2019b) and Yang et al. (2020a) improved the greedy search and proposed two different greedy attack methods. Alzantot et al. (2018) proposed a black-box attack to generate syntactically similar adversarial examples. However, the previous methods usually search in a greedy way, which has exponential time complexity.

Defenses on categorical data. There are also a lot of defense methods in the continuous data domain compared with discrete data. For example, the most effective one is adversarial training Madry et al. (2017), which searches for the worst case by PGD during the training and trains on the worst adversarial examples to increase the robustness of the models. Several works have been proposed for categorical adversarial defenses. Pruthi et al. (2019) used a word recognition model to preprocess the input data and increased the robustness of downstream tasks. Zhou et al. (2020) proposed to use random smoothing to defense substitution-based attacks. Wang et al. (2021) detected adversarial examples to defend synonym substitution attacks based on the fact that word replacement can destroy mutual interaction. Swenor & Kalita (2022) used random noise to increase the model robustness. Xie et al. (2022) used a huge amount of data to detect the adversarial examples. Dong et al. (2021b) combined different attack methods, such as Hot-flip (Ebrahimi et al., 2017a) and l_2 -attack (Miyato et al., 2017), and adversarial training as the defense method. It also proposed its own adversarial method ASCC which takes the solution space as a convex hull of word vectors and achieves good performance in sentiment analysis and natural language inference. However, the defense methods are mainly focusing on NLP, which may rely on word embeddings.

3 PROBABLISTIC CATEGORICAL ADVERSARIAL ATTACK (PCAA)

3.1 PROBLEM SETUP

Categorical Attack. We first introduce necessary definitions and notations. In detail, we consider a classifier f that predicts labels $y \in \mathcal{Y}$ based on categorical inputs $x \in \mathcal{X}$. Each input sample x

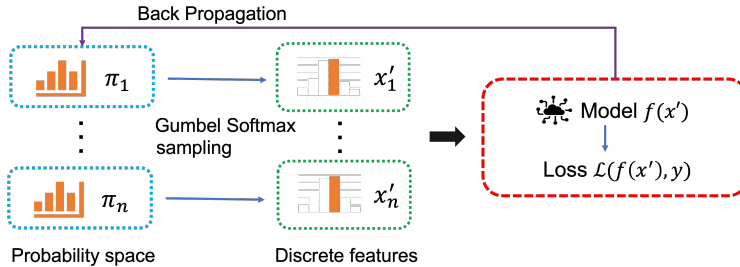


Figure 1: **The overall framework of Probabilistic Categorical Adversarial Attack.** First, a distribution π is used to sample adversarial examples x' . Then, the gradient of loss function will be used to update the probability distribution π .

contains n categorical features, and each feature x_i take a value from d categorical values. In this paper, we consider l_0 -attack, which restricts the number of perturbed features. In particular, given the budget size ϵ , we aim to find an adversarial example x' which solves the following optimization problem:

$$\max \mathcal{L}(f(x'), y) \quad \text{s.t.} \quad \|x' - x\|_0 \leq \epsilon. \quad (1)$$

Existing search-based methods. To solve the problem, there existing search-based methods such as Greedy Search (Yang et al., 2020a) and Gradient-Guided Search (Lei et al., 2019a). In general, they consist of two major steps. First, they search for features in x whose change can result in the greatest influence on the model output. This can be estimated by either perturbing a feature x_i , or calculating the scale of the model’s gradient to the feature x_i . The second step involves greedy or global searches among all possible combinations of the features identified by the first step. This process is time-consuming if the feature dimension n or the number of categories d is large.

3.2 OBJECTIVE OF PROBABILISTIC CATEGORICAL ADVERSARIAL ATTACK (PCAA)

Probabilistic Categorical Attack. In this work, we propose an alternative approach to solve the problem in Eq.(1), by transferring it to the continuous probabilistic space. In detail, we assume that each feature of (adversarial) categorical data x'_i follows a categorical distribution: $Categorical(\pi_i)$, where $\pi_i \in \Pi_i = (0, 1)^d$. And each $\pi_{i,j}$ represents the probability that the feature i belongs to the category j . In the remaining of the paper, we will use π_i to denote the categorical distribution $Categorical(\pi_i)$ without the loss of generality. Therefore, each input sample x ’s distribution can be represented as an array $\pi = [\pi_0; \pi_1; \dots; \pi_n] \in \Pi \subset \mathbb{R}^{n \times d}$, where the element $\pi_{i,j}$ represents the probability that the feature x_i belongs to the category j . Then, we define the following optimization problem to find a probability distribution π in the space of Π :

$$\max_{\pi \in \Pi} \mathbb{E}_{x' \sim \pi} \mathcal{L}(f(x'), y), \quad \text{s.t.} \quad \Pr_{x' \sim \pi} (\|x' - x\|_0 \geq \epsilon) \leq \delta \quad (2)$$

where ϵ denotes the perturbation budget size and δ is the tail probability constraint. By solving the problem in Eq.(2), we aim to find a distribution with parameter π such that: **(1) on average**, the generated samples x' following distribution π have a high loss value; and **(2) with low probability**, the sample x' has a l_0 distance to the clean sample x larger than ϵ . In this way, the generated samples x' are likely to mislead the model prediction while preserving most features of x .

3.3 AN EFFICIENT ALGORITHM OF PCAA

In this subsection, we provide a feasible algorithm to solve our proposed Eq.(2) in practice. First, in order to handle the probability constraint in Eq.(2), we substitute the l_0 distance between x' and x by calculating the sum of Cross Entropy Loss between π_i and x_i : $\mathcal{L}_{CE}(\pi_i, x_i)$ for all features $i \in |x|$. It is because $\mathcal{L}_{CE}(\pi_i, x_i)$ measures the probability that the i -th feature of the generated samples is different from x_i . Therefore, we use the sum of Cross Entropy $\sum_{i \in |n|} \mathcal{L}_{CE}(\pi_i, x_i)$ to approximate the number of changed features, which is the l_0 difference $\|x' - x\|_0$. In our algorithm, we penalize the searched π when the term $\sum_{i \in |n|} \mathcal{L}_{CE}(\pi_i, x_i)$ exceeds a positive value ζ (as Eq.

3). In this case, we equivalently limit the probability that the generated samples x' have the number of perturbed features larger than ϵ .

$$\Pr_{x' \sim \pi} (\|x' - x\| \geq \epsilon) \rightarrow \left[\sum_{i \in [n]} \mathcal{L}_{CE}(x_i, \pi_i) - \zeta \right]^+ \quad (3)$$

Moreover, since the Cross-Entropy Loss is differentiable in terms of π , we further transform the problem to its Lagrangian form:

$$\max_{\pi} \mathbb{E}_{x' \sim \pi} \mathcal{L}(f(x'), y) - \lambda \left[\sum_{i \in [n]} \mathcal{L}_{CE}(x_i, \pi_i) - \zeta \right]^+ \quad (4)$$

where λ is the penalty coefficient, and $[\cdot]^+$ is $\max(\cdot, 0)$. Next, we will show how to solve the maximization problem above by applying gradient methods.

Back propagation through Gumbel-Softmax. Note that the gradient of the expected loss function with respect to π cannot be directly calculated in Eq. 4, so we apply the Gumbel-Softmax estimator Jang et al. (2017). In practice, to avoid the projection onto probability simplex which hurts the time efficiency and Gumbel-Softmax does not require a normalized probability distribution, we consider an unnormalized categorical distribution $\pi_i \in (0, C]^d$ where $C > 0$ is a large constant guaranteeing that the searching space is sufficiently large. The distribution generates sample vectors x'_i as follows:

$$x'_{ij} = \frac{\exp((\log \pi_{ij} + g_j)/\tau)}{\sum_{j=1}^d \exp((\log \pi_{ij} + g_j)/\tau)}, \text{ for } j = 1, \dots, d \quad (5)$$

where g_j denotes i.i.d samples from the Gumbel(0, 1) distribution, and τ denotes the softmax temperature. This re-parameterization process facilitates us to calculate the gradient of the expected loss in terms of π . Therefore, we can derive the following estimator of gradients for the expected loss:

$$\frac{\partial \mathbb{E}_{x' \sim \pi} \mathcal{L}(f(x'), y)}{\partial \pi} \approx \frac{\partial}{\partial \pi} \mathbb{E}_g \mathcal{L}(f(x'(\pi, g)), y) = \mathbb{E}_g \left[\frac{\partial \mathcal{L}}{\partial x'} \frac{\partial x'}{\partial \pi} \right] \approx \frac{1}{n_g} \sum_{i=1}^{n_g} \left[\frac{\partial \mathcal{L}}{\partial x'} \frac{\partial x'(\pi, g_i)}{\partial \pi} \right] \quad (6)$$

where n_g is the number of i.i.d samples from g . In Eq. 6, the first approximation is from the re-parameterization of a sample x' ; the second equality comes from exchanging the order of expectation and derivative, and the third, we approximate the expectation of gradients by calculating the average of gradients. Finally, we derive the practical solution to solve Eq. 4 as demonstrated in Figure 1, by leveraging the gradient ascent algorithm, such as Madry et al. (2018). In Algorithm 1, we provide the details of our proposed attack method. Specifically, in each step, we first update the unnormalized distribution π by gradient ascent. Then, we clip π back to its domain $(0, C]^d$.

Algorithm 1 Probabilistic Categorical Adversarial Attack (PCAA)

Input Data \mathcal{D} , budget ϵ , number of samples n_g , penalty coefficient λ , maximum iteration I , learning rate γ

Output Adversarial Distribution π

Initialize distribution π^0

for $t \leq I$ **do**

Estimate expected gradient using Eq6: $\nabla_{\pi} \mathbb{E}_{\pi} \mathcal{L} \approx \frac{1}{n_g} \sum_{i=1}^{n_g} \left[\frac{\partial \mathcal{L}}{\partial x'} \frac{\partial x'(\pi^t, g_i)}{\partial \pi} \right]$

Gradient ascent: $\tilde{\pi}^{t+1} = \pi^t + \gamma \cdot \nabla_{\pi} (\mathbb{E}_{\pi} \mathcal{L} - \lambda \nabla_{\pi} [\mathcal{L}_{CE}(\pi_t, x) - \zeta]^+)$

Clip to $(0, C]^d$: $\pi^{t+1} = \max(\tilde{\pi}^{t+1}, C)$

end for

Return distribution π

3.4 TIME COMPLEXITY ANALYSIS

In this subsection, we compare the time complexity of our method with search-based methods to illustrate the efficiency advantage of our proposed attack method. Here, we assume that the whole dataset has N data points, each data point has n features, each feature has d categories and the budget of allowed perturbation is ϵ . We first introduce the details of four different search-based methods.

- **Greedy Search(GS)**. It consists of two stages. The first stage involves traversing all features. For the i_{th} feature, it replaces the original category with all other $d - 1$ categories respectively and records the change of the model loss for each category. The largest change is treated as the impact score for the i_{th} feature. Then it selects the top ϵ features with the highest impact scores to perturb. In the second stage, it finds the combination with the greatest loss among all possible combinations of categories for selected features.
- **Greedy attack.(GA)**(Yang et al., 2020a). This method is a modified version of Greedy Search. The first stage is similar to that of GS, while the second stage searches for the best perturbation feature by feature. For the i_{th} selected feature, it replaces the original category with one that results in the largest loss and then searches the next selected feature until all selected features are traversed. Therefore, it only needs to go over each feature once without traversing all combinations in the second stage.
- **Gradient-guided GS.(GGS)**(Lei et al., 2019a). In order to determine which features should be perturbed, this method utilizes gradient information in the first stage. It computes the gradient of the loss function w.r.t the original input and treats the gradient of each feature as the impact score. Those ϵ features with the greatest impact scores are selected to be perturbed. In the second stage, it follows the same strategy as that of the second stage in GS, and pursues the combination resulting in the greatest loss among all possible combinations.
- **Gradient-guided greedy attack(GGA)**. On the basis of GGS, it remains the same first stage and modifies the second stage by adopting the same strategy as that in the second stage of GA.

We analyze the time complexity of our PCAA and the aforementioned methods for comparison.

Probability Categorical Adversarial Attack. In practical implementation, we adopt the unconstrained optimization in Eq. 4. Therefore, the time complexity is only from gradient descent. Further, assume that we sample n_s times when estimating the expected gradient and the maximum number of iterations is I . We need to compute gradient n_s times during one iteration, thus the time complexity of PCAA is:

$$N \cdot n_s \cdot \mathcal{O}(nd) \cdot I = C_1 N \mathcal{O}(nd) \quad (7)$$

where C_1 is some constant related to n_s and I .

Greedy search. As described above, the Greedy search consists of two stages. In the first stage, it needs to traverse all n features and compute a loss when plugging in each level. Then $n \cdot d$ times of calculation are needed. In the second stage, all possible combinations of categories for top ϵ features need to be considered and the number of possible combinations is d^ϵ . Therefore, the time complexity for Greedy Search is

$$N \cdot [\Omega(stage1) + \Omega(stage2)] = N \cdot [\mathcal{O}(nd) + \mathcal{O}(d^\epsilon)] = N \cdot \mathcal{O}(nd + d^\epsilon) \quad (8)$$

Greedy attack. In the first stage, it traverses all categories within each feature to find the best ϵ features to perturb, which needs $n \cdot d$ computation. And in the second stage, it searches ϵ selected features respectively for the d categories that lead to the greatest loss. Thus, the time complexity is

$$N \cdot [\Omega(stage1) + \Omega(stage2)] = N \cdot [\mathcal{O}(nd) + \mathcal{O}(\epsilon d)] = N \cdot \mathcal{O}(nd + \epsilon d) \quad (9)$$

Gradient guided greedy search. In the first stage, it computes gradients w.r.t the original input, and in the second stage, it also needs to traverse all possible combinations for selected features. Therefore the time complexity is

$$N \cdot [\Omega(stage1) + \Omega(stage2)] = N \cdot [\mathcal{O}(n) + \mathcal{O}(d^\epsilon)] = N \cdot \mathcal{O}(n + d^\epsilon) \quad (10)$$

Gradient-guided greedy attack. This method combines the first stage in GGS and the second stage in GA. Thus its time complexity is

$$N \cdot [\Omega(stage1) + \Omega(stage2)] = N \cdot [\mathcal{O}(n) + \mathcal{O}(\epsilon d)] = N \cdot \mathcal{O}(n + \epsilon d) \quad (11)$$

From the above analysis, two search-based methods suffer from the exponential increase of time complexity (8-10) when the number of feature categories d and budget size ϵ is increasing. Moreover, two greedy-based methods (GA and GGA) accelerate the second stage and achieve better time efficiency, but sacrifice the performance as they greatly narrow down the searching space. However, PCAA is free of these problems, since it achieves great time efficiency to greedy-based methods while remain comparable performance with search-based methods. We will show this with more evidence in the experimental part.

4 PROBABILISTIC ADVERSARIAL TRAINING (PADvT)

Adversarial training Szegedy et al. (2014) Goodfellow et al. (2015) Madry et al. (2018) is one of the most effective methods to build robust models. Its training incorporates adversarial examples to improve robustness. However, greedy methods have limited applications in adversarial training when dealing with categorical input models due to their slow speed. The proposed PCAA method significantly improves time efficiency. In this case, we propose Probabilistic Adversarial Training (PADvT) based on PCAA in order to train robust models on categorical data. Recalling the formulation of PCAA in Eq. 4, and denoting the parameters for classifier f with θ , the training objective for PADvT is formulated as

$$\min_{\theta} \left[\max_{\pi} \mathbb{E}_{x' \sim \pi} \left(\mathcal{L}(f(x'; \theta), y) - \lambda \left[\sum_{i \in [n]} \mathcal{L}_{CE}(x_i, \pi_i) - \zeta \right]^+ \right) \right] \quad (12)$$

Since our objective involves a penalty coefficient, we adopt the strategy in Yurochkin et al. (2020) to update λ during training. We adaptively choose λ according to $\mathcal{L}_{CE}(x, \pi) - \epsilon$ from the last iteration: when the value is large, we increase λ to strengthen the constraints and vice versa.

The implementation of PADvT is illustrated in Algorithm. 2. It first solves the inner maximization problem by applying algorithm. 1. Then it samples n_{adv} examples from the adversarial distribution π . Finally, it applies Adam Kingma & Ba (2015) to solve the outer minimization problem. The process will continue until the maximal number of iterations is reached.

Algorithm 2 Probabilistic Adversarial Training (PADvT)

Input Data \mathcal{D} , parameters of clean model θ , budget ϵ , parameters of Algorithm1, n_{adv} , initial penalty coefficient λ^0 , penalty coefficient step size α , parameters of Adam optimizer, number of iterations I

Output parameters θ of robust model

repeat

 Sample mini-batch $B = \{x^1, \dots, x^m\}$ from training set

for $i = 1, \dots, m$ (in parallel) **do**

 Apply Algorithm1 to x^i and obtain adversarial distribution π^i

 Sample n_{adv} examples $\{x_1^i, \dots, X_{n_{adv}}^i\}$ from π^i using Gumbel Softmax tricks

end for

 Update θ by Adam to minimize the average adversarial loss

 Update $\lambda = \max\{0, \lambda - \alpha(\zeta - \frac{1}{m} \sum_{i \in [m]} \sum_{j \in [n]} \mathcal{L}_{CE}(x_j^i, \pi_j^i))\}$

until Training converged

Return parameters θ

5 EXPERIMENT

In this section, we conduct experiments to validate the effectiveness and efficiency of PCAA and PADvT. In Section 5.1, we demonstrate that PCAA achieves a better balance between attack success rate and time efficiency. In Section. 5.2, PADvT achieves competitive or stronger robustness against categorical attacks across different tasks.

5.1 CATEGORICAL ADVERSARIAL ATTACKS

Experimental Setup. In this evaluation, we focus on three datasets. **(1) Intrusion Prevention System (IPS) (Wang et al., 2020).** IPS dataset has 242,467 instances, with each input consisting of 20 features and each feature has 1,103 categorical values. The output space has three labels. A standard LSTM based classifier (Bao et al., 2022) is trained for IPS dataset. **(2) AG’s News corpus.** This dataset consists of titles and description fields of news articles. The tokens of each sentence correspond to the categorical features, and the substitution set (of size 70) corresponds to the categorical values. A character-based CNN (Zhang et al., 2015) is trained on this dataset. **(3) Splice-junction Gene Sequences (Splice) (Noordewier et al., 1990).** Splice dataset has 3,190 instances. Each one

Table 1: Attacking results on IPS, AG’s news, and Splice datasets. SR. represents the attack success rate; T. represents the average running time in seconds; and “-” indicates the running time over 10 hours. Each experiment runs 5 times and 95% confidence intervals are shown.

Dataset	Attack Method	$\epsilon = 1$		$\epsilon = 2$		$\epsilon = 3$		$\epsilon = 4$		$\epsilon = 5$	
		SR.(\uparrow)	T.(\downarrow)	SR.(\uparrow)	T.(\downarrow)	SR.(\uparrow)	T.(\downarrow)	SR.(\uparrow)	T.(\downarrow)	SR.(\uparrow)	T.(\downarrow)
IPS	GS	66.11±0.03	38.5	81.24±0.01	2028	-	-	-	-	-	-
	GA	66.11±0.03	35.4	71.32±0.02	38.1	79.44±0.06	39.9	85.28±0.07	41.5	91.15±0.11	43.2
	GGG	41.53±0.04	2.06	75.47±0.03	1022	-	-	-	-	-	-
	GGA	41.53±0.04	0.98	63.72±0.06	2.01	70.76±0.05	2.94	75.89±0.08	3.74	82.43±0.10	4.56
	PCAA	60.66±0.05	6.62	77.51±0.04	6.64	85.32±0.08	6.75	91.67±0.06	6.66	94.63±0.12	6.54
AG	GS	41.22±0.01	15.3	67.38±0.02	21.9	75.87±0.04	356	83.10±0.02	19160	-	-
	GA	41.22±0.01	15.3	60.71±0.05	15.4	66.33±0.04	15.5	74.47±0.07	15.7	86.63±0.12	15.9
	GGG	32.39±0.03	0.352	59.21±0.02	3.24	67.79±0.03	151	79.22±0.05	7551	-	-
	GGA	32.39±0.03	0.352	41.29±0.07	0.393	56.11±0.06	0.511	67.53±0.10	0.613	72.28±0.07	0.856
	PCAA	47.31±0.07	9.02	61.47±0.09	8.91	70.81±0.06	9.11	79.55±0.13	9.06	87.41±0.10	8.92
Splice	GS	72.11±0.01	0.905	75.84±0.02	1.02	81.41±0.01	1.36	85.19±0.03	2.73	90.88±0.04	8.25
	GA	72.11±0.01	0.905	74.42±0.03	0.911	78.18±0.06	0.915	80.61±0.04	0.922	83.74±0.05	0.928
	GGG	61.71±0.02	0.028	68.28±0.03	0.083	72.82±0.02	0.251	77.11±0.04	0.917	82.53±0.04	3.56
	GGA	61.71±0.02	0.028	65.26±0.08	0.031	70.31±0.05	0.0337	74.84±0.07	0.035	80.49±0.08	0.037
	PCAA	64.32±0.04	1.57	70.83±0.07	1.62	79.28±0.09	1.67	82.82±0.06	1.73	84.77±0.11	1.39

is a gene fragment of 60 features with 5 categorical values. The output space has three labels. The LSTM model is trained for Splice.

Baseline Attacks. Our goal is to generate powerful attacks on categorical models directly on input space, without relying on the embedding space. Thus, search-based methods are most suitable in this setting. Therefore, we compare PCAA with the following search-based attacks including Greedy Search(GS), Greedy Attack(GA), Gradient-guided GS (GGG) and Gradient-guided greedy attack(GGA). The details of these baselines can be found in Section 3.4.

For each dataset, we evaluate the performance in terms of the attack success rate (SR.) and the average running time (T.) under various budget sizes ϵ ranging from 1 to 5. Since PCAA learns the adversarial categorical distribution, the generation of adversarial examples is based on sampling. In the evaluation, we sample 100 examples from the adversarial distribution for each attack instance and dismiss those with more perturbed features than the budget size. We claim a success attack if one out of all samples successfully changes the prediction. The average running time for PCAA includes both the optimization of Eq. 4 and the sampling process.

Performance Comparison. The experimental results on IPS, AG’s news, and Splice datasets are demonstrated in Table 1. (1) In IPS dataset, each data contains more than 1000 categories for each feature, and our method has a significant advantage. As the budget size increases, GS and GGG run so long that they are no longer feasible. While GA and GGA remain efficient, our method outperforms them by significant margins, e.g., over 10% higher than GGA in success rate and faster over 7 times than GA in time. (2) In AG’s news data, GS has the highest success rate with small budgets. However, our method PCAA has much better time efficiency than GS under all budgets while maintaining competitive performance and outperforming all the other attacks. (3) In Splice, there are only 5 categories for each feature, which corresponds to the low-dimensional cases. The advantage of our method may not be as great as in other cases (i.e., IPS and AG) as the running time of greedy methods is acceptable. However, PCAA still generates strong attacks with only 6% less than GS in success rate, and remains highly efficient when the budget size increases, so it still provides a practical and competitive attack.

Time Efficiency Comparison. It is evident from Table 1 that PCAA’s average running time is unaffected by the size of the budget, whereas greedy methods are clearly more time-consuming as budgets increase. Thus, the proposed PCAA method significantly improves time efficiency. As shown in Table 1, when the substitution set is large (e.g., IPS and AG), the running time of GS and GGG will explode even under some moderate budget size such as $\epsilon = 3$. Since the running time is too long for meaningful comparison, we do not record the performance for these budget sizes of those baselines. In small substitution sets, i.e., Splice, the time efficiency gap between PCAA and greedy methods is not large. Due to the small size of the substitution set, greedy methods need much fewer queries than in high-dimensional cases, but PCAA still takes a much shorter time for a budget size of more than 4.




Figure 3: PAdvT and baseline defense performances under different attacks on IMDB dataset.

5.2 CATEGORICAL ADVERSARIAL DEFENSES

Experimental Setup. For the defense evaluation, we focus on two datasets, AG’s News Corpus and IMDB. It is because IPS and Splice have too few samples (no more than 1000 for each dataset). In particular: **(1) AG’s News corpus.** is the same dataset used in the attack evaluation and the model is also a character-based CNN. As character swapping does not require embeddings for each character, we can apply attacking methods on input space. Therefore, the robustness of the defense models is evaluated using six attacks, i.e., Hot-Flip (Ebrahimi et al., 2017a), GS, GA, GGS, GGA, and PCAA. **(2) IMDB reviews dataset**(Maas et al., 2011). Under this dataset, we focus on a word-level classification task and we study two model architectures, namely Bi-LSTM and CNN, are trained for prediction. To evaluate the robustness, four attacks are deployed, including a genetic attack (Alzantot et al., 2018) (which is an attack method proposed to generate adversarial examples in embedding space), as well as GS, GGS, and PCAA.




Figure 2: PAdvT and baseline defenses under different attacks on AG’s news dataset.

Baseline Defenses. We compare our PAdvT with following baseline defenses:

- **Standard training.** It minimizes the average cross-entropy loss on clean input.
- **Hot-flip**(Ebrahimi et al., 2017a). It uses the gradient with respect to the one-hot input representation to find out the which individual feature under perturbation has the highest estimated loss. It is initially proposed to model char flip in Char-CNN model, and we also apply it to word-level substitution, as in (Dong et al., 2021a).
- **Adv l_2 -ball**(Miyato et al., 2017). It uses an l_2 PGD adversarial attack inside the word embedding space for adversarial training.
- **ASCC-Defense**(Dong et al., 2021a). A state-of-the-art defense method in text classification. It uses the worst perturbation over a sparse convex hull in the word embedding space for adversarial training.

Performance Comparison. The experimental results on AG’s news dataset are shown in Fig. 2. The Y-axis represents the error rate and X-axis represents different defense methods while multi-rows represents different attacking methods. Our defense method achieves leading robustness on Char-CNN over all attacks with significant margins, surpassing Hot-Flip-defense by 10%. Fig. 3 illustrates the defense results on the IMDB dataset, and we have similar observations with Fig. 2. Our PAdvT shows competitive adversarial robustness as ASCC-defense. Notably, ASCC is a defense method that conducts adversarial training on word-embedding space. It relies on the key assumption that similar words have a close distance in embedding space. However, our method does not rely on this assumption, which may result in the performance being competitive (slightly worse) than ASCC. For all other defenses, it outperforms them across different architectures significantly.

5.3 ABLATION STUDY

Concentration of PCAA. To further understand the behavior of our attack algorithm, in this subsection, we ask the question: *what is the variance of our optimized probability distribution π^* (from solving Eq.2)?* Intuitively, we desire the distribution π to have smaller variance, so that we don't need too many times of sampling to obtain the optimal adversarial examples. To confirm this point, we conduct an ablation study based on an experiment on IPS dataset to visualize the distribution π , which is optimized via PCAA under various budget sizes. In Fig 4, we choose three budget sizes $\epsilon = 1, 3, 5$ and choose 4 features to present the adversarial categorical distributions. Notably, in Fig. 4, the left two columns are the features that are not perturbed. The right two columns are features which will be perturbed by PCAA. From the figure, we can see that for each feature (perturbed or unperturbed), there exists one category with a much higher probability compared to other categories. This fact indicates during the sampling process of PCAA, for a certain feature, the samplings are highly likely to have the same category for this feature. As a result, we confirm that our sampled adversarial examples are well-concentrated.

Ablation analysis for PAdvT. In our training objective in Eq. 12, ζ controls the budget size used for adversarial training and possibly affects the robustness of the model. We conduct an ablation study on the IMDB dataset to understand the impact of ζ . The results are demonstrated in Table. 2. When ζ increases, the success rates of all attacks decreases, meaning that the robustness of models is enhanced. However, large ζ will increase the clean error and sacrifice the clean accuracy of models. Thus, ζ controls the balance between the clean accuracy and the adversarial robustness of the model. When $\zeta = 0.4$, our algorithm reaches a good balance between model accuracy and robustness.

Table 2: Ablation study: impact of the budget regularization term ζ on PAdvT

	Clean Err		Genetic SR		GS SR		Gradient Search SR		PCAA SR	
	LSTM	CNN	LSTM	CNN	LSTM	CNN	LSTM	CNN	LSTM	CNN
$\zeta = 0.1$	16.1	16.9	37.5	29.6	41.3	43.5	39.5	41.4	40.2	42.2
$\zeta = 0.2$	17.2	17.3	34.1	32.7	39.9	41.3	38.6	39.7	39.1	40.6
$\zeta = 0.32$	17.9	18.1	30.1	31.4	38.7	38.3	38.5	37.6	38.8	37.8
$\zeta = 0.4$	18.4	18.6	26.3	28.5	37.8	36.2	36.4	34.9	36.9	35.7

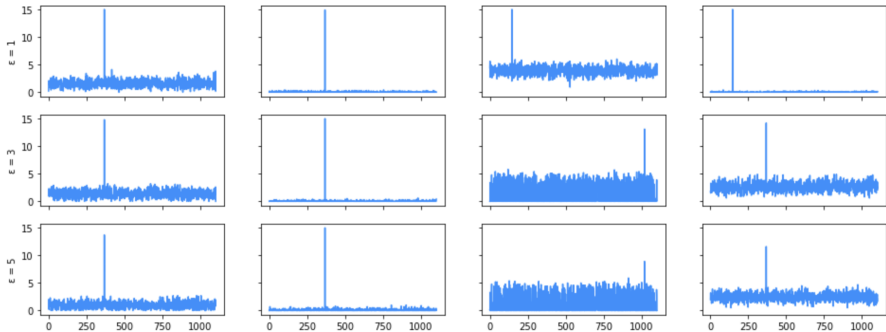


Figure 4: Visualization of Optimized Categorical Distribution for Various Features (IPS)

6 CONCLUSION

In this paper, we propose a novel probabilistic optimization formulation for attacking models with categorical input. Our framework significantly improves the time efficiency compared with greedy methods and shows promising empirical performance across different datasets. Furthermore, we design an adversarial training method based on the proposed probabilistic attack that achieves better robustness.

REFERENCES

- Moustafa Alzantot, Yash Sharma, Ahmed Elgohary, Bo-Jhang Ho, Mani B. Srivastava, and Kai-Wei Chang. Generating natural language adversarial examples. In Ellen Riloff, David Chiang, Julia Hockenmaier, and Jun'ichi Tsujii (eds.), *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pp. 2890–2896. Association for Computational Linguistics, 2018. doi: 10.18653/v1/d18-1316. URL <https://doi.org/10.18653/v1/d18-1316>.
- Hongyan Bao, Yufei Han, Yujun Zhou, Yun Shen, and Xiangliang Zhang. Towards understanding the robustness against evasion attack on categorical data. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net, 2022. URL <https://openreview.net/forum?id=BmJV7kyAmg>.
- Lijuan Cao and Francis Eng Hock Tay. Financial forecasting using support vector machines. *Neural Comput. Appl.*, 10(2):184–192, 2001. doi: 10.1007/s005210170010. URL <https://doi.org/10.1007/s005210170010>.
- Jin Cui, Lin Shen Liew, Giedre Sabaliauskaite, and Fengjun Zhou. A review on safety failures, security attacks, and available countermeasures for autonomous vehicles. *Ad Hoc Networks*, 90, 2019. doi: 10.1016/j.adhoc.2018.12.006. URL <https://doi.org/10.1016/j.adhoc.2018.12.006>.
- Xinshuai Dong, Anh Tuan Luu, Rongrong Ji, and Hong Liu. Towards robustness against natural language word substitutions. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021a. URL https://openreview.net/forum?id=ks5nebnVn_.
- Xinshuai Dong, Anh Tuan Luu, Rongrong Ji, and Hong Liu. Towards robustness against natural language word substitutions. *arXiv preprint arXiv:2107.13541*, 2021b.
- Javid Ebrahimi, Anyi Rao, Daniel Lowd, and Dejing Dou. Hotflip: White-box adversarial examples for NLP. *CoRR*, abs/1712.06751, 2017a. URL <http://arxiv.org/abs/1712.06751>.
- Javid Ebrahimi, Anyi Rao, Daniel Lowd, and Dejing Dou. Hotflip: White-box adversarial examples for text classification. *arXiv preprint arXiv:1712.06751*, 2017b.
- Ji Gao, Jack Lanchantin, Mary Lou Soffa, and Yanjun Qi. Black-box generation of adversarial text sequences to evade deep learning classifiers. In *2018 IEEE Security and Privacy Workshops (SPW)*, pp. 50–56. IEEE, 2018.
- Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. In Yoshua Bengio and Yann LeCun (eds.), *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. URL <http://arxiv.org/abs/1412.6572>.
- Andrew Ilyas, Shibani Santurkar, Dimitris Tsipras, Logan Engstrom, Brandon Tran, and Aleksander Madry. Adversarial examples are not bugs, they are features. *Advances in neural information processing systems*, 32, 2019.
- Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017. URL <https://openreview.net/forum?id=rkE3y85ee>.
- Ansam Khraisat, Iqbal Gondal, Peter Vamplew, and Joarder Kamruzzaman. Survey of intrusion detection systems: techniques, datasets and challenges. *Cybersecur.*, 2(1):20, 2019. doi: 10.1186/s42400-019-0038-7. URL <https://doi.org/10.1186/s42400-019-0038-7>.
- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun (eds.), *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. URL <http://arxiv.org/abs/1412.6980>.

- Volodymyr Kuleshov, Shantanu Thakoor, Tingfung Lau, and Stefano Ermon. Adversarial examples for natural language classification problems. 2018.
- Qi Lei, Lingfei Wu, Pin-Yu Chen, Alex Dimakis, Inderjit S. Dhillon, and Michael J. Witbrock. Discrete adversarial attacks and submodular optimization with applications to text classification. In Ameet Talwalkar, Virginia Smith, and Matei Zaharia (eds.), *Proceedings of Machine Learning and Systems 2019, MLSys 2019, Stanford, CA, USA, March 31 - April 2, 2019*. mlsys.org, 2019a. URL <https://proceedings.mlsys.org/book/284.pdf>.
- Qi Lei, Lingfei Wu, Pin-Yu Chen, Alex Dimakis, Inderjit S Dhillon, and Michael J Witbrock. Discrete adversarial attacks and submodular optimization with applications to text classification. *Proceedings of Machine Learning and Systems*, 1:146–165, 2019b.
- Andrew L. Maas, Raymond E. Daly, Peter T. Pham, Dan Huang, Andrew Y. Ng, and Christopher Potts. Learning word vectors for sentiment analysis. In Dekang Lin, Yuji Matsumoto, and Rada Mihalcea (eds.), *The 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies, Proceedings of the Conference, 19-24 June, 2011, Portland, Oregon, USA*, pp. 142–150. The Association for Computer Linguistics, 2011. URL <https://aclanthology.org/P11-1015/>.
- Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*, 2017.
- Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018. URL <https://openreview.net/forum?id=rJzIBfZAb>.
- Takeru Miyato, Andrew M. Dai, and Ian Goodfellow. Adversarial training methods for semi-supervised text classification. In *International Conference on Learning Representations*, 2017. URL https://openreview.net/forum?id=r1X3g2_xl.
- Michiel Noordewier, Geoffrey Towell, and Jude Shavlik. Training knowledge-based neural networks to recognize genes in dna sequences. *Advances in neural information processing systems*, 3, 1990.
- Danish Pruthi, Bhuwan Dhingra, and Zachary C Lipton. Combating adversarial misspellings with robust word recognition. *arXiv preprint arXiv:1905.11268*, 2019.
- Suranjana Samanta and Sameep Mehta. Towards crafting text adversarial samples. *arXiv preprint arXiv:1707.02812*, 2017.
- Ali Shafahi, W Ronny Huang, Mahyar Najibi, Octavian Suci, Christoph Studer, Tudor Dumitras, and Tom Goldstein. Poison frogs! targeted clean-label poisoning attacks on neural networks. *Advances in neural information processing systems*, 31, 2018.
- Gianluca Stringhini, Christopher Kruegel, and Giovanni Vigna. Detecting spammers on social networks. In Carrie Gates, Michael Franz, and John P. McDermott (eds.), *Twenty-Sixth Annual Computer Security Applications Conference, ACSAC 2010, Austin, Texas, USA, 6-10 December 2010*, pp. 1–9. ACM, 2010. doi: 10.1145/1920261.1920263. URL <https://doi.org/10.1145/1920261.1920263>.
- Abigail Swenor and Jugal Kalita. Using random perturbations to mitigate adversarial attacks on sentiment analysis models. *arXiv preprint arXiv:2202.05758*, 2022.
- Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian J. Goodfellow, and Rob Fergus. Intriguing properties of neural networks. In Yoshua Bengio and Yann LeCun (eds.), *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*, 2014. URL <http://arxiv.org/abs/1312.6199>.
- Xiaosen Wang, Yifeng Xiong, and Kun He. Randomized substitution and vote for textual adversarial example detection. *arXiv preprint arXiv:2109.05698*, 2021.

- Yutong Wang, Yufei Han, Hongyan Bao, Yun Shen, Fenglong Ma, Jin Li, and Xiangliang Zhang. Attackability characterization of adversarial evasion attack on discrete data. In Rajesh Gupta, Yan Liu, Jiliang Tang, and B. Aditya Prakash (eds.), *KDD '20: The 26th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, Virtual Event, CA, USA, August 23-27, 2020*, pp. 1415–1425. ACM, 2020. doi: 10.1145/3394486.3403194. URL <https://doi.org/10.1145/3394486.3403194>.
- Zhouhang Xie, Jonathan Brophy, Adam Noack, Wencong You, Kalyani Asthana, Carter Perkins, Sabrina Reis, Sameer Singh, and Daniel Lowd. Identifying adversarial attacks on text classifiers. *arXiv preprint arXiv:2201.08555*, 2022.
- Puyudi Yang, Jianbo Chen, Cho-Jui Hsieh, Jane-Ling Wang, and Michael I. Jordan. Greedy attack and gumbel attack: Generating adversarial examples for discrete data. *J. Mach. Learn. Res.*, 21: 43:1–43:36, 2020a. URL <http://jmlr.org/papers/v21/19-569.html>.
- Puyudi Yang, Jianbo Chen, Cho-Jui Hsieh, Jane-Ling Wang, and Michael I Jordan. Greedy attack and gumbel attack: Generating adversarial examples for discrete data. *J. Mach. Learn. Res.*, 21(43):1–36, 2020b.
- Mikhail Yurochkin, Amanda Bower, and Yuekai Sun. Training individually fair ML models with sensitive subspace robustness. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020. URL <https://openreview.net/forum?id=BlgdkxHFDH>.
- Xiang Zhang, Junbo Jake Zhao, and Yann LeCun. Character-level convolutional networks for text classification. In Corinna Cortes, Neil D. Lawrence, Daniel D. Lee, Masashi Sugiyama, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pp. 649–657, 2015. URL <https://proceedings.neurips.cc/paper/2015/hash/250cf8b51c773f3f8dc8b4be867a9a02-Abstract.html>.
- Yi Zhou, Xiaoqing Zheng, Cho-Jui Hsieh, Kai-wei Chang, and Xuanjing Huang. Defense against adversarial attacks in nlp via dirichlet neighborhood ensemble. *arXiv preprint arXiv:2006.11627*, 2020.

A APPENDIX

A.1 ADDITIONAL EXPERIMENTAL RESULTS

In this subsection, we present the results of Fig. 2 and 3 where the values represent error rates. We run each experiment for 5 times and compute the 95% confidence interval.

Table 3: PAdvT and baseline defense performances under different attacks on IMDB dataset

	Clean Err		Genetic SR		GS SR		Gradient Search SR		PCAA SR	
	LSTM	CNN	LSTM	CNN	LSTM	CNN	LSTM	CNN	LSTM	CNN
ERM	15.50±0.005	15.23±0.007	92.62±0.022	65.68±0.031	94.86±0.030	82.02±0.049	91.61±0.011	80.54±0.030	92.26±0.087	81.42±0.073
Hotflip	17.37±0.049	16.57±0.051	50.63±0.030	55.93±0.025	75.24±0.015	66.35±0.012	67.96±0.021	65.47±0.022	68.08±0.069	65.90±0.074
AdvJ2	32.53±0.034	37.38±0.043	56.59±0.036	54.35±0.033	79.69±0.030	67.00±0.034	78.34±0.031	65.81±0.038	78.58±0.068	66.23±0.077
ASCC	17.76±0.036	18.37±0.030	19.66±0.047	22.52±0.046	34.67±0.047	33.28±0.045	33.46±0.054	32.61±0.061	33.97±0.101	32.87±0.082
PAdvT	18.57±0.033	18.85±0.049	26.01±0.065	28.35±0.049	37.30±0.063	36.36±0.061	35.60±0.067	34.85±0.053	35.46±0.104	35.47±0.093

Table 4: PAdvT and baseline defense performances under different attacks on AG’s news dataset

	Clean Err	Hotflip	PCAA	GS	GGs	GA	GGA
ERM	8.70±0.009	79.72±0.015	80.75±0.066	83.09±0.015	79.28±0.010	74.43±0.010	67.53±0.016
Hotflip	13.99±0.017	60.07±0.013	63.47±0.057	64.28±0.018	62.41±0.019	60.51±0.017	58.33±0.013
PAdvT	14.62±0.028	45.38±0.037	49.50±0.081	50.65±0.035	47.14±0.041	44.71±0.040	42.18±0.045

We also run PAdvT with mixture of adversarial examples and clean samples on IMDB dataset. Results are shown in Table 5 where values represent error rates. It is noticeable that clean samples will slightly improve the clean performance and lead to small decreasing of robustness.

Table 5: Comparison of PAdvT on IMDB dataset with/without mixture of clean samples.

	Clean Err	Genetic	GS	GGs	PCAA
IMDB LSTM(mix)	18.27	26.22	37.67	35.79	36.05
IMDB LSTM	18.57	26.01	37.3	35.60	35.46
IMDB CNN(mix)	18.69	28.51	36.47	34.96	35.72
IMDB CNN	18.85	28.35	36.36	34.85	35.47

A.2 VISUALIZATION ON IMDB

We run PCAA attack on IMDB dataset over two victim models LSTM and word-CNN. The candidate sets are pre-specified synonym set. The following Tables 6 and 7 show some successful adversarial examples, where the red words are adversarial words and blue are original words. It is obvious that these replacements do not hurt the semantic meaning but can fool the classifiers.

Table 6: IMDB Adversarial Examples from PCAA on LSTM

Class	Perturbed Class	Perturbed Texts
Negative	Positive	I watched this film for 45 minutes and counted 9 mullets. That’s a mullet every 5 minutes. Seriously, though this film is residing evidence (living proof) that formula works if it ain’t broke, it don’t need fit in a streetwise yet vulnerable heroine, a hardened ex-cop martial arts master with a heart of gold and a serial killer with ‘issues’ pure magic.
Negative	Positive	Claustrophobic camera angles that do not aid (help) the movie. Too long face only shots, where you most of the time get the hunch (feeling) that the lower half of the film is missing that the screen is cut off because there seems to be important actions going on, but you can not see them. There is anyway already too much confusion in the movie, so these viewing angles make it worse and do not contribute to artful visuals. I like artfully made movies and unconventional camera work. I can handle deep and slow movies but this one is trying too hard to be something artful and fails, in my opinion, painfully. Nothing to get attached to any of the characters because they are not worked out well enough to work out characters. More is needed than just minute long face shots. At least with this set of script director actors, I wonder whether some of the not so decent (good) acting is due to the script and director or due to the actors. I will stay away from films both written and directed by le you for sure in the future. What an annoying film even for person (someone) who would be interested in that part of history and for someone who spent time in Shanghai.

Positive	Negative	I really liked this version of 'vanishing point' as opposed to the 1971 version. I finds(found) the 1971 version quite boring if I can get up in the middle of a movie a few times as I did with the 1971 version, then to me it is not all that great. Of course, this could be due to the fact that I was only nine at the time the 1971 version was brought out. However, I have noticed(seen) many remakes everytime(when) I have liked the original and older one better. I found that the plot of the 1997 version was more understandable and had basically kept true to the original without undermining the meaning of the 1971 version. In my opinion I felt the 1997 version had more excitement and wasn't so blasé boring.
Positive	Negative	The cast is marvellous(excellent) , the acting good, the plot interesting, the evolvment full of suspense, but it is hard to cram all those elements into a film that is barely 80 minutes long. If more time was taken to develop the plot and subplots, it would have a much better effect. Another 30 minutes of substance would have made this a very alright(good) film rather than just a good one.
Positive	Negative	There is great detail in a 'bug 's life'. Everything is covered. The film expects(looks) great and the animation is sometimes jaw dropping. The film isn't too terribly original. It 's basically a modern take on kurosawa 's seven samurai only with bugs, I enjoyed the character interaction however, and the naughty boys(bad guys) in this film actually seemed bad. It seems that Disney usually makes their bad guys carbon copy cut outs, the grasshoppers are menacing and hopper the lead bad guy was a brilliant creation. Check this one out.

Table 7: IMDB Adversarial Examples from PCAA on word-CNN

Class	Perturbed Class	Perturbed Texts
Positive	Negative	I am a college student studying A levels and need help and comments from anyone who has any views at all about the theme of mothers in film. In The Mother, whether you have gone through something similar or just want to comment and help me research more about this film, any comment would much greatly appreciated. The comments will be used alone(solely) for exam purposes and will be included in my written exam. So if you have any views at all I'm convinced(sure) I can put them to use and you could help me get an A. I am also studying about a boy and tadpole. So if you have seen these films as well, I would appreciate it if you could leave comments on here on that page. Thank you.

Negative	Positive	This movie is so horrendous (awful). It is hard to find the right words to describe it. At first the story is so ridiculous. A narrow minded human can write a better plot. The actors are boring and un-talented. Perhaps they were compelled to play in this dorky (cheesy) film. The camera receptions of the national forest are the only good in this whole movie. I should feel ashamed because I paid for this lousy picture. Hopefully nobody makes a sequel or make a similar film with such a worse storyline.
Positive	Negative	This movie is wonderful, the writing, directing, acting, all are marvelous (fantastic). Very witty and clever script quality performances by actors. Ally Sheedy is strong and dynamic and delightfully quirky really original and heart warmingly unpredictable. The scenes are alive with fresh energy and really talented generating (production)
Positive	Negative	This may not be war peace but the two academy noms wouldn't have been forthcoming. If it weren't for the genius of James Wong Howe, this is one of the few films I've fallen in love with as a infant (child) and gone back to without dissatisfaction. Whether you have any interest in what it offers fictively or not, BBC is a visual feast. I'm not saying it's his best work. I'm no expert there for sure but the look of this movie is astounding (amazing). I love everything about it, Elsa Lanchester, the cat, the crazy hoodoo, the retro downtown Ness, but the way it was put on film is breathtaking. I even like the inconsistencies pointed out on this page aforementioned (above) and the special effects that seem backward. Now it all creates a really consistent world.
Positive	Negative	Bette Midler is again divine raunchily hilarious (humorous) in love with burlesque, capable of bringing you down to tears either with old jokes, with new dresses or merely with old songs, with more power punch than ever. All in all, sung (singing) new ballads power, singing the good old perennial ones such as the rose 'stay with me' and yes even 'wind beneath my wings'. The best way to appreciate the Divine Miss M has always been libe since this is the next best thing to it. I strongly recommended to all with a mixture of adult extensive (wide) eyed enchantment and appreciation and a child's mischievous wish for pushing all boundaries.