

# NATURAL LANGUAGE ADVERSARIAL ATTACK AND DEFENSE IN WORD LEVEL

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

Up until very recently, inspired by a mass of researches on adversarial examples for computer vision, there has been a growing interest in designing adversarial attacks for Natural Language Processing (NLP) tasks, followed by very few works of adversarial defenses for NLP. To our knowledge, there exists no defense method against the successful synonym substitution based attacks that aim to satisfy all the lexical, grammatical, semantic constraints and thus are hard to be perceived by humans. We contribute to fill this gap and propose a novel adversarial defense method called *Synonym Encoding Method* (SEM), which inserts an encoder before the input layer of the model and then trains the model to eliminate adversarial perturbations. Extensive experiments demonstrate that SEM can efficiently defend current best synonym substitution based adversarial attacks with little decay on the accuracy for benign examples. To better evaluate SEM, we also design a strong attack method called Improved Genetic Algorithm (IGA) that adopts the genetic metaheuristic for synonym substitution based attacks. Compared with existing genetic based adversarial attack, IGA can achieve higher attack success rate while maintaining the transferability of the adversarial examples.

## 1 INTRODUCTION

Deep Neural Networks (DNNs) have made great success in various machine learning tasks, such as computer vision (Krizhevsky et al., 2012; He et al., 2016), and Natural Language Processing (NLP) (Kim, 2014; Lai et al., 2015; Devlin et al., 2018). However, recent studies have discovered that DNNs are vulnerable to *adversarial examples* not only for computer vision tasks (Szegedy et al., 2014) but also for NLP tasks (Papernot et al., 2016), causing a serious threat to their safe applications. For instance, spammers can evade spam filtering system with adversarial examples of spam emails while preserving the intended meaning.

In contrast to numerous methods proposed for adversarial attacks (Goodfellow et al., 2015; Nicholas & David, 2017; Anish et al., 2018) and defenses (Goodfellow et al., 2015; Guo et al., 2018; Song et al., 2019) in computer vision, there are only a few list of works in the area of NLP, inspired by the works for images and emerging very recently in the past two years (Zhang et al., 2019). This is mainly because existing perturbation-based methods for images cannot be directly applied to texts due to their discrete property in nature. Furthermore, if we want the perturbation to be barely perceptible by humans, it should satisfy the lexical, grammatical and semantic constraints in texts, making it even harder to generate the text adversarial examples.

Current attacks in NLP can fall into four categories, namely modifying the characters of a word (Liang et al., 2017; Ebrahimi et al., 2017), adding or removing words (Liang et al., 2017), replacing words arbitrarily (Papernot et al., 2016), and substituting words with synonyms (Alzantot et al., 2018; Ren et al., 2019). The first three categories are easy to be detected and defended by spell or syntax check (Rodriguez & Rojas-Galeano, 2018; Pruthi et al., 2019). As synonym substitution aims to satisfy all the lexical, grammatical and semantic constraints, it is hard to be detected by automatic spell or syntax check as well as human investigation. To our knowledge, currently there is no defense method specifically designed against the synonym substitution based attacks.

In this work, we postulate that the model generalization leads to the existence of adversarial examples: a generalization that is not strong enough causes the problem that there usually exists some neighbors  $x'$  of a benign example  $x$  in the manifold with a different classification. Based on this

hypothesis, we propose a novel defense mechanism called *Synonym Encoding Method* (SEM) that encodes all the synonyms to a unique code so as to force all the neighbors of  $x$  to have the same label of  $x$ . Specifically, we first cluster the synonyms according to the *Euclidean Distance* in the embedding space to construct the encoder. Then we insert the encoder before the input layer of the deep model without modifying its architecture, and train the model again to defend adversarial attacks. In this way, we can defend the synonym substitution based adversarial attacks efficiently in the context of text classification.

Extensive experiments on three popular datasets demonstrate that the proposed SEM can effectively defend adversarial attacks, while maintaining the efficiency and achieving roughly the same accuracy on benign data as the original model does. To our knowledge, SEM is the first proposed method that can effectively defend the synonym substitution based adversarial attacks.

Besides, to demonstrate the efficacy of SEM, we also propose a genetic based attack method, called *Improved Genetic Algorithm* (IGA), which is well-designed and more efficient as compared with the first proposed genetic based attack algorithm, GA (Alzantot et al., 2018). Experiments show that IGA can degrade the classification accuracy more significantly with lower word substitution rate than GA. At the same time IGA keeps the transferability of adversarial examples as GA does.

## 2 BACKGROUND AND RELATED WORK

Let  $W$  denote the word set containing all the legal words. Let  $x = \{w_1, \dots, w_i, \dots, w_n\}$  denote an input text,  $\mathcal{C}$  the corpus that contains all the possible input texts, and  $\mathcal{Y} \in \mathbb{N}^K$  the output space. The classifier  $f : \mathcal{C} \rightarrow \mathcal{Y}$  takes an input  $x$  and predicts its label  $f(x)$ , and let  $S_m(x, y)$  denote the confidence value for the  $y$ -th category at the softmax layer. Let  $Syn(w, \sigma, k)$  represent the set of first  $k$  synonyms of  $w$  within distance  $\sigma$ , namely

$$Syn(w, \sigma, k) = \{w'_1, \dots, w'_i, \dots, w'_k | w'_i \in W \wedge \|w - w'_1\|_p \leq \dots \leq \|w - w'_k\|_p < \sigma\},$$

where  $\|w - w'\|_p$  is the  $p$ -norm distance evaluated on the corresponding embedding vectors.

### 2.1 NATURAL LANGUAGE ADVERSARIAL EXAMPLES

Suppose we have an ideal classifier  $c : \mathcal{C} \rightarrow \mathcal{Y}$  that could always output the correct label for any input text  $x$ . For a subset of (train or test) texts  $\mathcal{T} \subseteq \mathcal{C}$  and a small constant  $\epsilon$ , we could define the natural language adversarial examples as follows:

$$\mathcal{A} = \{x_{adv} \in \mathcal{C} | \exists x \in \mathcal{T}, f(x_{adv}) \neq c(x_{adv}) = c(x) = f(x) \text{ and } d(x - x_{adv}) < \epsilon\}, \quad (1)$$

where  $d(x - x_{adv})$  is a distance metric to evaluate the dissimilarity between the benign example  $x = \{w_1, \dots, w_i, \dots, w_n\}$  and the adversarial example  $x_{adv} = \{w'_1, \dots, w'_i, \dots, w'_n\}$ . It is usually defined as the  $p$ -norm distance:  $d(x - x_{adv}) = \|x - x_{adv}\|_p = (\sum_i \|w_i - w'_i\|_p)^{\frac{1}{p}}$ .

### 2.2 TEXT ADVERSARIAL ATTACKS

In this subsection, we provide a brief overview of three popular synonym substitution based adversarial attack methods.

**Greedy Search Algorithm (GSA).** Kuleshov et al. (2018) propose a greedy search algorithm to substitute words with their synonyms so as to maintain the semantic and syntactic similarity. GSA first constructs a synonym set  $W_s$  for an input text  $x = \{w_1, \dots, w_i, \dots, w_n\}$ :

$$W_s = \{Syn(w_i, \sigma, k) | 1 \leq i \leq n\}. \quad (2)$$

Initially, let  $x_{adv} = x$ . Then at each stage for  $x_{adv} = \{w'_1, \dots, w'_i, \dots, w'_n\}$ , GSA finds a word  $\hat{w}'_i \in W$  that satisfies the syntactic constraint and minimizes  $S_m(\hat{x}, y_{true})$  where  $\hat{x} = \{w'_1, \dots, w'_{i-1}, \hat{w}'_i, w'_{i+1}, \dots, w'_n\}$ , and updates  $x_{adv} = \hat{x}$ . Such process iterates until  $x_{adv}$  becomes an adversarial example or the word replacement rate reaches a threshold.

**Genetic Algorithm (GA).** Alzantot et al. (2018) propose a population-based algorithm to replace words with their synonyms so as to generate semantically and syntactically similar adversarial examples. There are three operators in GA:

- *Mutate*( $x$ ): Randomly choose a word  $w_i$  in text  $x$  that has not been updated and substitute  $w_i$  with  $w'_i$ , one of its synonyms  $Syn(w_i, \sigma, k)$  that does not violate the syntax constraint by the “Google 1 billion words language model” (Chelba et al., 2013) and minimize  $S_m(\hat{x}, y_{true})$  where  $\hat{x} = \{w_1, \dots, w_{i-1}, w'_i, w_{i+1}, \dots, w_n\}$  and  $S_m(\hat{x}, y_{true}) < S_m(x, y_{true})$ ;
- *Sample*( $\mathcal{P}$ ): Randomly sample a text from population  $\mathcal{P}$  with a probability proportional to  $1 - S_m(x_i, y_{true})$  where  $x_i \in \mathcal{P}$ ;
- *Crossover*( $a, b$ ): Construct a new text  $c = \{w_1^c, \dots, w_i^c, \dots, w_n^c\}$ , where  $w_i^c$  is randomly chosen from  $\{w_i^a, w_i^b\}$  based on the input texts  $a = \{w_1^a, \dots, w_i^a, \dots, w_n^a\}$  and  $b = \{w_1^b, \dots, w_i^b, \dots, w_n^b\}$ .

For a text  $x$ , GA first generates an initial population  $\mathcal{P}^0$  of size  $m$ :

$$\mathcal{P}^0 = \{Mutate(x), \dots, Mutate(x)\}. \quad (3)$$

Then at each iteration, GA generates the next generation of population through *crossover* and *mutation* operators:

$$\begin{aligned} x_{adv}^{i+1} &= \arg \min_{\tilde{x} \in \mathcal{P}^i} S_m(\tilde{x}, y_{true}), \\ child_k^{i+1} &= Crossover(Sample(\mathcal{P}^i), Sample(\mathcal{P}^i)), \\ \mathcal{P}^{i+1} &= \{x_{adv}^{i+1}, Mutate(child_1^{i+1}), \dots, Mutate(child_{m-1}^{i+1})\}. \end{aligned} \quad (4)$$

GA terminates when it finds an adversarial example or reaches the maximum number of iteration limit.

**Probability Weighted Word Saliency (PWWS).** Ren et al. (2019) propose a new synonym substitution method called Probability Weighted Word Saliency (PWWS), which considers the word saliency as well as the classification confidence. Given a text  $x = \{w_1, \dots, w_i, \dots, w_n\}$ , PWWS first calculates the saliency of each word  $S(x, w_i)$ :

$$S(x, w_i) = S_m(x, y_{true}) - S_m(\bar{x}_i, y_{true}) \text{ where } \bar{x}_i = \{w_1, \dots, w_{i-1}, unk, w_{i+1}, \dots, w_n\}, \quad (5)$$

where “unk” means the word is removed. Then PWWS calculates the maximum possible change in the classification confidence resulted from substituting word  $w_i$  with one of its synonyms:

$$\Delta S_{m_i}^*(x) = \max_{w'_i \in Syn(w_i, \sigma, k)} [S_m(x, y_{true}) - S_m(x'_i, y_{true})] \text{ where } x'_i = \{w_1, \dots, w'_i, \dots, w_n\}. \quad (6)$$

Then, PWWS sequentially checks the words in descending order of  $\phi(S(x, w_i))_i \cdot \Delta S_{m_i}^*(x)$ , where  $\phi(z)_i = \frac{e^{z_i}}{\sum_{k=1}^n e^{z_k}}$ , and substitutes the current word  $w_i$  with its optimal synonym  $w_i^*$ :

$$w_i^* = \arg \max_{w'_i \in Syn(w_i, \sigma, k)} [S_m(x, y_{true}) - S_m(x'_i, y_{true})]. \quad (7)$$

PWWS terminates when it finds an adversarial example  $x_{adv}$  or it has replaced all the words in  $x$ .

### 2.3 TEXT ADVERSARIAL DEFENSES

There exist very few works for text adversarial defenses.

- In the character-level, Pruthi et al. (2019) propose to place a word recognition model in front of the downstream classifier to defend character-level adversarial attacks by combating adversarial spelling mistakes.
- In the word level, for defenses on synonym substitution based attacks, only Alzantot et al. (2018) and Ren et al. (2019) incorporate the adversarial training strategy proposed in the image domain (Goodfellow et al., 2015) with their text attack methods respectively, and demonstrate that adversarial training can promote the model’s robustness. However, there is no defense method specifically designed to defend the synonym substitution based adversarial attacks.

### 3 THE PROPOSED TEXT DEFENSE METHOD

We first introduce our motivation, then present the proposed text defense method, *Synonym Encoding Method* (SEM).

#### 3.1 MOTIVATION

Let  $\mathcal{X}$  denote the input space,  $V_\epsilon(x)$  denote the  $\epsilon$ -neighborhood of data point  $x \in \mathcal{X}$ , where  $V_\epsilon(x) = \{x' \in \mathcal{X} \mid \|x' - x\| < \epsilon\}$ . As illustrated in Figure 1 (a), we postulate that the generalization of the model leads to the existence of adversarial examples. More generally, given a data point  $x \in \mathcal{X}$ ,  $\exists x' \in V_\epsilon(x), f(x') \neq y'_{true}$  where  $x'$  is an adversarial example of  $x$ .




Figure 1: The neighborhood of a data point  $x$  in the input space. (a) Traditional training: there exists some data point  $x'$  that the model has never seen before and yields wrong classification, in other words, such data point  $x'$  will be an adversarial example. (b) Adding infinite labeled data: this is an ideal case that the model has seen all the data points to resist adversarial examples. (c) Constructing a convex outer bound: all data points in this polytope will share the label of  $x$ . (d) Mapping neighborhood data points: mapping all neighbors to the center  $x$  so as to eliminate adversarial examples.

Ideally, to defend the adversarial attack, we need to train a classifier  $f$  which not only guarantees  $f(x) = y_{true}$ , but also assures  $\forall x' \in V_\epsilon(x), f(x') = y'_{true}$ . Thus, the most effective way is to add more labeled data to improve the adversarial robustness (Schmidt et al., 2018). Ideally, as illustrated in Figure 1 (b), if we have infinite labeled data, we can train a model  $f : \forall x' \in V_\epsilon(x), f(x') = y'_{true}$  with high probability so that the model  $f$  is robust enough to adversarial examples. Practically, however, labeling data is very expensive and it is impossible to have infinite labeled data.

Because it is impossible to have enough labeled data to train a robust model, as illustrated in Figure 1 (c), Wong & Kolter (2018) propose to construct a convex outer bound and guarantee that all data points in this bound share the same label. The goal is to train a model  $f : \forall x' \in V_\epsilon(x), f(x') = f(x) = y_{true}$ . Specifically, they propose a linear-programming (LP) based upper bound on the robust loss by adopting a linear relaxation of the ReLU activation and minimize this upper bound during the training. Then they bound the LP optimal value and calculate the elementwise bounds on the activation functions based on a backward pass through the network. Although their method does not need any extra data, it is hard to scale to realistically-sized networks due to its high complexity.

In this work, as illustrated in Figure 1 (d), we propose a novel way to find a mapping  $m : \mathcal{X} \rightarrow \mathcal{X}$  where  $\forall x' \in V_\epsilon(x), m(x') = x$ . In this way, we force the classification to be more smooth and we do not need any extra data to train the model or modify the architecture of the model. All we need to do is to insert the mapping before the input layer and train the model on the original training set. Now the problem turns into how to locate the neighbors of data point  $x$ . For image tasks, it is hard to find all images in the neighborhood of  $x$  in the input space, and there could be infinite number of neighbors. For NLP tasks, however, utilizing the property that words in sentences are discrete tokens, we can easily find almost all neighbors of an input text. Based on this insight, we propose a new method called *Synonym Encoding Method* to locate the neighbors of an input  $x$ .

#### 3.2 SYNONYM ENCODING

We assume that the closer the meaning of two sentences is, the closer their distance is in the input space. Thus, we can suppose that the neighbors of  $x$  are its synonymous sentences. To find the synonymous sentence, we can substitute words in the sentence with their synonyms. To construct the mapping  $m$ , all we need to do is to cluster the synonyms and allocate a unique token for each cluster, which we call the *Synonym Encoding Method* (SEM). The details are in Algorithm 1.

**Algorithm 1** *Synonym Encoding Method*


---

**Input:**  $W$ : dictionary of words,  $n$ : size of  $W$ ,  $\sigma$ : distance for synonyms,  
 $k$ : number of synonyms for each word

**Output:**  $E$ : encoding result

- 1:  $E = \{w_1 : \text{None}, \dots, w_n : \text{None}\}$
- 2: **for** each word  $w_i \in W$  **do**
- 3:   **if**  $E[w_i] = \text{NONE}$  **then**
- 4:     **if**  $\exists w'_j \in \text{Syn}(w_i, \sigma, k), E[w'_j] \neq \text{NONE}$  **then**
- 5:        $w_i^* \leftarrow$  the closest  $w'_j \in W$  where  $E[w'_j] \neq \text{NONE}$
- 6:        $E[w_i] = E[w_i^*]$
- 7:     **else**
- 8:        $E[w_i] = w_i$
- 9:     **end if**
- 10:   **for** each word  $w'_j$  in  $\text{Syn}(w_i, \sigma, k)$  **do**
- 11:     **if**  $E[w'_j] = \text{NONE}$  **then**
- 12:        $E[w'_j] = E[w_i]$
- 13:     **end if**
- 14:   **end for**
- 15: **end if**
- 16: **end for**
- 17: **return**  $E$

---

## 4 THE IMPROVED GENETIC BASED TEXT ATTACK

The current synonym substitution based text adversarial attacks (Alzantot et al., 2018; Kuleshov et al., 2018; Ren et al., 2019) have a constraint that they only substitute words at the same position once or replace words with the first  $k$  synonyms of the word in the original input  $x$ . This constraint can lead to local minimum for adversarial examples, and it is hard to choose a suitable  $k$  as different words may have different number of synonyms.

To address this issue, we propose an Improved Genetic Algorithm (IGA), which allows to substitute words in the same position more than once based on the current text  $x'$ . In this way, IGA can traverse all synonyms of a word no matter what value  $k$  is. Meanwhile, we can avoid local minimum to some extent as we allow the substitution of the word by the original word in the current position. In order to guarantee that the substituted word is still a synonym of the original word, each word in the same position can be replaced at most  $\lambda$  times.

Differs to the first genetic based text attack algorithm of Alzantot et al. (2018), we change the structure of the algorithm, including the operators for crossover and mutation. For more details of IGA, see Appendix A.

## 5 EXPERIMENTS

We evaluate SEM with four attacks, GSA (Kuleshov et al., 2018), GA (Alzantot et al., 2018), PWWS (Ren et al., 2019) and our IGA, on three popular datasets involving three neural network classification models. The results demonstrate that SEM can significantly improve the robustness of neural networks and IGA can achieve better attack performance as compared with existing attacks. And we further provide discussion on the hyper-parameter  $\epsilon$  of SEM in Appendix B.

### 5.1 EXPERIMENTAL SETUP

We first provide an overview of the datasets and classification models used in the experiments.

**Datasets.** In order to evaluate the efficacy of SEM, we choose three popular datasets: IMDB, AG’s News, and Yahoo! Answers. IMDB (Potts, 2011) is a large dataset for binary sentiment classification, containing 25,000 highly polarized movie reviews for training and 25,000 for testing. AG’s News (Zhang et al., 2015) consists news article pertaining four classes: World, Sports, Business and

Sci/Tech. Each class contains 30,000 training examples and 1,900 testing examples. Yahoo! Answers (Zhang et al., 2015) is a topic classification dataset from the ‘‘Yahoo! Answers Comprehensive Questions and Answers’’ version 1.0 dataset with 10 categories, such as Society & Culture, Science & Mathematics, etc. Each class contains 140,000 training samples and 5,000 testing samples.

**Models.** To better evaluate our method, we adopt several state-of-the-art models for text classification, including Convolution Neural Networks (CNNs) and Recurrent Neural Networks (RNNs). The embedding dimension for all models are 300 (Mikolov et al., 2013). We replicate the CNN’s architecture from Kim (2014), which contains three convolutional layers with filter size of 3, 4, and 5 respectively, a max-pooling layer and a fully-connected layer. LSTM consists of three LSTM layers where each layer has 300 LSTM units and a fully-connected layer (Liu et al., 2016). Bi-LSTM contains a bi-directional LSTM layer whose forward and reverse have 300 LSTM units respectively and a fully-connected layer.

**Baselines.** We take the method of adversarial training (Goodfellow et al., 2015) as our baseline. However, due to the low efficiency of text adversarial attacks, we cannot implement adversarial training as it is in the image domain. In the experiments, we adopt PWWS, which is quicker than GA and IGA, to generate 10% adversarial examples of the training set, and re-train the model incorporating adversarial examples with the training data.

## 5.2 EVALUATION ON DEFENSE METHODS

To evaluate the efficiency of the SEM method, we randomly sample 200 correctly classified examples on different models from each dataset and use the above attack methods to generate adversarial examples with or without defense. The more effective the defense method is, the smaller the classification accuracy of the model drops. Table 1 shows the efficacy of various attack and defense methods.

Table 1: Accuracy of various classification models on the datasets, with and without defenses, under adversarial attacks. For each model (Word-CNN, LSTM, or Bi-LSTM), if we look at each row, the highest classification accuracy for various defense methods is highlighted in **bold** to indicate the **best defense efficacy**; if we look at each column, the lowest classification accuracy under various adversarial attacks is highlighted in underline to indicate the best attack efficacy. NT: Normal Training, AT: Adversarial Training.

Dataset	Attack	Word-CNN(%)			LSTM(%)			Bi-LSTM(%)		
		NT	AT	SEM	NT	AT	SEM	NT	AT	SEM
IMDB	No Attack	88.7	<b>89.1</b>	86.8	87.3	<b>89.6</b>	86.8	88.2	<b>90.3</b>	87.6
	GSA	13.3	16.9	<b>66.4</b>	8.3	21.1	<b>72.2</b>	7.9	20.8	<b>73.1</b>
	PWWS	4.4	5.3	<b>71.1</b>	2.2	3.6	<b>77.3</b>	1.8	3.2	<b>76.1</b>
	GA	7.1	10.7	<b>71.8</b>	2.6	9.0	<b>77.0</b>	1.8	7.2	<b>71.6</b>
	IGA	<u>0.9</u>	<u>2.7</u>	<b>65.1</b>	<u>0.9</u>	<u>1.8</u>	<b>71.2</b>	<u>0.9</u>	<u>2.7</u>	<b>69.3</b>
AG’s News	No Attack	91.7	<b>92.2</b>	88.7	91.8	<b>92.1</b>	90.9	91.7	<b>92.2</b>	88.7
	GSA	33.0	37.8	<b>63.9</b>	45.9	56.7	<b>83.2</b>	33.0	37.8	<b>63.9</b>
	PWWS	30.7	41.5	<b>67.6</b>	50.0	55.7	<b>85.0</b>	30.7	41.5	<b>67.6</b>
	GA	24.1	40.6	<b>77.9</b>	43.8	57.3	<b>86.4</b>	24.1	40.6	<b>77.9</b>
	IGA	<u>21.5</u>	<u>35.5</u>	<b>70.3</b>	<u>40.0</u>	<u>55.3</u>	<b>81.8</b>	<u>21.5</u>	<u>35.5</u>	<b>70.3</b>
Yahoo! Answers	No Attack	68.4	<b>69.3</b>	65.8	71.6	<b>71.7</b>	69.0	72.3	<b>72.8</b>	70.2
	GSA	19.6	20.8	<b>49.4</b>	27.6	30.5	<b>48.6</b>	24.6	30.9	<b>53.4</b>
	PWWS	10.3	12.5	<b>52.6</b>	21.1	22.9	<b>54.9</b>	17.3	20.0	<b>57.2</b>
	GA	13.7	16.6	<b>59.2</b>	15.8	17.9	<b>66.2</b>	13.0	16.0	<b>63.2</b>
	IGA	<u>8.9</u>	<u>10.0</u>	<b>51.4</b>	<u>10.5</u>	<u>15.1</u>	<b>53.3</b>	<u>12.4</u>	<u>13.5</u>	<b>55.7</b>

For each network model, we look at each row to find the best defense result under the setting of no attack, or GSA, PWWS, GA, and IGA attacks:

- Under the setting of no attack, adversarial training (AT) could improve the classification accuracy of the models on all datasets, as adversarial training (AT) is also the way to augment the training data. Our defense method SEM reaches an accuracy closing to normal training (NT).
- Under the four attacks, however, the classification accuracy with normal training (NT) and adversarial training (AT) drops significantly. For normal training (NT), the accuracy degrades more than 75%, 42% and 40% on the three datasets respectively. And adversarial training (AT) cannot defend these attacks effectively, especially for PWWS and IGA on IMDB and Yahoo! Answers, where AT only improves the accuracy a little (smaller than 5%). By contrast, SEM can remarkably improve the robustness of the deep models for all the four attacks.

### 5.3 DEFENSE FOR TRANSFERABILITY

In the image domain, the transferability of adversarial attack refers to its ability to decrease the accuracy of models using adversarial examples generated based on other models (Szegedy et al., 2014; Goodfellow et al., 2015). Papernot et al. (2016) find that the adversarial examples in NLP also exhibit a good transferability. Therefore, a good defense method not only could defend the adversarial attacks but also resists the transferability of adversarial examples.

To evaluate the ability of preventing the transferability of adversarial examples, we generate adversarial examples on each model under normal training, and test them on other models with or without defense on Yahoo! Answers. The results are shown in Table 2. Almost on all models with adversarial examples generated by other models, SEM could yield the highest classification accuracy.

Table 2: Accuracy of various classification models for adversarial examples generated on other models on Yahoo! Answers for evaluating the transferability. \* indicates that the adversarial examples are generated based on this model.

Attack	Word-CNN(%)			LSTM(%)			Bi-LSTM(%)		
	NT	AT	SEM	NT	AT	SEM	NT	AT	SEM
GSA	19.6*	52.7	<b>57.5</b>	52.6	58.4	<b>61.8</b>	52.7	57.5	<b>59.6</b>
PWWS	10.3*	54.4	<b>57.2</b>	46.5	57.7	<b>61.4</b>	53.1	<b>57.9</b>	56.8
GA	13.7*	49.9	<b>55.2</b>	43.0	<b>66.0</b>	65.6	56.4	65.5	<b>67.4</b>
IGA	8.9*	53.7	<b>56.5</b>	46.5	56.6	<b>58.6</b>	53.1	56.8	<b>57.2</b>
GSA	47.2	52.7	<b>55.9</b>	27.6*	60.9	<b>62.1</b>	53.8	61.9	<b>62.8</b>
PWWS	43.7	54.7	<b>56.5</b>	21.1*	59.9	<b>61.8</b>	50.6	59.2	<b>61.4</b>
GA	41.0	48.5	<b>54.8</b>	15.8*	57.4	<b>57.5</b>	43.4	58.2	<b>60.9</b>
IGA	47.8	53.0	<b>55.2</b>	10.5*	53.4	<b>58.3</b>	53.5	59.0	<b>61.1</b>
GSA	43.7	53.4	<b>54.3</b>	52.9	57.7	<b>58.7</b>	24.6*	60.4	<b>61.1</b>
PWWS	41.7	48.5	<b>49.7</b>	41.2	58.0	<b>60.4</b>	17.3*	<b>60.1</b>	58.6
GA	36.9	45.7	<b>49.0</b>	43.0	58.8	<b>60.7</b>	13.0*	58.2	<b>60.0</b>
IGA	44.8	50.6	<b>52.3</b>	45.1	57.4	<b>58.7</b>	12.4*	60.1	<b>61.1</b>

### 5.4 EVALUATION ON ATTACK METHODS (IGA VS. GA)

For text attacks, we compare the proposed IGA with GA from various aspects, including attack efficacy, transferability and human evaluation.

**Attack Efficacy.** As shown in Table 1, looking at each column, we see that under normal training (NT) and adversarial training (AT), IGA can always achieve the lowest classification accuracy, which corresponds to the highest attack success rate, on all models and datasets among the four attacks. Under the third column of SEM defense, though IGA may not be the best among all attacks, IGA always outperforms GA.

Besides, as depicted in Table 3, IGA can yield lower word substitution rate than GA on most models. Note that for SEM, GA can yield lower word substitution rate, because GA may not replace the word as most words cannot bring any benefit for the first replacement. This indicates that GA stops

at local minimum while IGA continues to substitute words and gain a lower classification accuracy, as demonstrated in Table 1.

Table 3: The word substitution rate for GA and IGA on different models.

Dataset	Attack	Word-CNN(%)			LSTM(%)			Bi-LSTM(%)		
		NT	AT	SEM	NT	AT	SEM	NT	AT	SEM
IMDB	GA	9.3	9.3	<b>4.3</b>	10.0	10.6	<b>4.0</b>	6.3	8.0	<b>3.1</b>
	IGA	<b>6.4</b>	<b>7.8</b>	6.3	<b>5.3</b>	<b>6.9</b>	5.7	<b>5.6</b>	<b>6.6</b>	7.5
AG’s News	GA	<b>10.3</b>	17.7	6.6	18.3	21.2	9.8	<b>12.0</b>	21.4	8.5
	IGA	12.1	<b>11.4</b>	<b>6.3</b>	<b>14.4</b>	<b>12.0</b>	<b>9.6</b>	13.5	<b>11.6</b>	<b>4.6</b>
Yahoo! Answers	GA	12.4	9.5	4.7	12.5	15.8	8.1	13.9	15.3	<b>4.7</b>
	IGA	<b>6.4</b>	<b>7.1</b>	<b>3.7</b>	<b>7.9</b>	<b>8.2</b>	<b>4.5</b>	<b>3.6</b>	<b>8.2</b>	5.1

**Transferability.** As shown in Table 2, the adversarial examples generated by IGA maintain roughly the same transferability as those generated by GA. For instance, if we generate adversarial examples on Word-CNN (column 2, NT), GA can achieve better transferability on LSTM with NT (column 5) while IGA can achieve better transferability on LSTM with AT and SEM (column 6, 7).

**Human Evaluation.** To further verify that the perturbations in the adversarial examples generated by IGA are hard for humans to perceive, we also perform a human evaluation on IMDB with 35 volunteers. We first randomly choose 100 benign examples that can be classified correctly and generate adversarial examples by GA and IGA on the three models so that we have a total of 700 examples. Then we randomly split them into 7 groups where each group contains 100 examples. We ask every five volunteers to classify one group independently. The accuracy of human evaluation on benign examples is 93.7%.




Figure 2: Classification accuracy by human evaluation.

As shown in Figure 2, the classification accuracy of human on adversarial examples generated by IGA is slightly higher than those generated by GA, and is slightly closer to the accuracy of human on benign examples.

**Summary.** IGA can achieve the highest attack success rate when compared with previous synonyms substitution based adversarial attacks and yield lower word replacement rate than GA. Besides, the adversarial examples generated by IGA maintains the same transferability as GA does and are a little bit harder for humans to distinguish. Several generated adversarial examples by GA and IGA are listed in Appendix C.

## 6 CONCLUSION

Synonym substitution based adversarial attacks are currently the best text attack methods, as they are hard to be checked by automatic spell or syntax check as well as human investigation. In this work, we propose a novel defense method called *Synonym Encoding Method* (SEM), which encodes the synonyms of each word to defend adversarial attacks for text classification task. Extensive experiments show that SEM can defend adversarial attacks efficiently and degrade the transferability of adversarial examples, at the same time SEM maintains the classification accuracy on benign data. To our knowledge, this is the first and efficient text defense method in word level for state-of-the-art synonym substitution based attacks.

In addition, we propose a new text attack method called Improved Genetic Attack (IGA), which in most cases can achieve much higher attack success rate as compared with existing attacks, at the same time IGA could maintain the transferability of adversarial examples.



## REFERENCES

- Moustafa Alzantot, Yash Sharma, Ahmed Elgohary, Bo-Jhang Ho, Mani B. Srivastava, and Kai-Wei Chang. Generating Natural Language Adversarial Examples. In *Empirical Methods in Natural Language Processing (EMNLP)*, pp. 2890–2896, 2018.
- Athalye Anish, Carlini Nicholas, and Wagner David. Obfuscated Gradients Give a False Sense of Security: Circumventing Defenses to Adversarial Examples. *International Conference on Machine Learning (ICML)*, 2018.
- Ciprian Chelba, Tomas Mikolov, Mike Schuster, Qi Ge, Thorsten Brants, Phillipp Koehn, and Tony Robinson. One Billion Word Benchmark for Measuring Progress in Statistical Language Modeling. *arXiv Preprint arXiv:1312.3005*, 2013.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of Deep Bidirectional Transformers for Language Understanding. *North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, 2018.
- Javid Ebrahimi, Anyi Rao, Daniel Lowd, and Dejing Dou. Hotflip: White-box Adversarial Examples for Text Classification. *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2017.
- Ian Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and Harnessing Adversarial Examples. *International Conference on Learning Representations (ICLR)*, 2015.
- Chuan Guo, Mayank Rana, Moustapha Cisse, and Maaten Laurens. Countering Adversarial Images using Input Transformations. *International Conference on Learning Representations (ICLR)*, 2018.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. *Computer Vision and Pattern Recognition (CVPR)*, 2016.
- Yoon Kim. Convolutional Neural Networks for Sentence Classification. *Empirical Methods in Natural Language Processing (EMNLP)*, 2014.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey Hinton. ImageNet Classification with Deep Convolutional Neural Networks. *Neural Information Processing Systems (NeurIPS)*, 2012.
- Volodymyr Kuleshov, Shantanu Thakoor, Tingfung Lau, and Stefano Ermon. Adversarial Examples for Natural Language Classification Problems. *OpenReview submission OpenReview:r1QZ3zbaZ.*, 2018.
- Siwei Lai, Liheng Xu, Kang Liu, and Jun Zhao. Recurrent Convolutional Neural Networks for Text Classification. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2015.
- Bin Liang, Hongcheng Li, Miaoqiang Su, Pan Bian, Xirong Li, and Wenchang Shi. Deep Text Classification Can be Fooled. *International Joint Conference on Artificial Intelligence (IJCAI)*, 2017.
- Pengfei Liu, Xipeng Qiu, and Xuanjing Huang. Recurrent Neural Network for Text Classification with Multi-task Learning. *International Joint Conference on Artificial Intelligence (IJCAI)*, 2016.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Efficient Estimation of Word Representations in Vector Space. *International Conference on Learning Representations (ICLR)*, 2013.
- Carlini Nicholas and Wagner David. Towards Evaluating the Robustness of Neural Networks. *IEEE Symposium on Security and Privacy*, 2017.
- Nicolas Papernot, Patrick McDaniel, Ananthram Swami, and Richard Harang. Crafting Adversarial Input Sequences for Recurrent Neural Networks. In *IEEE Military Communications Conference (MILCOM)*, pp. 49–54. IEEE, 2016.
- Christopher Potts. On the Negativity of egation. *Proceedings of Semantics and Linguistic Theory (SALT)*, 2011.

- Danish Pruthi, Bhuwan Dhingra, and Zachary C Lipton. Combating Adversarial Misspellings with Robust Word Recognition. *Association for Computational Linguistics (ACL)*, 2019.
- Shuhuai Ren, Yihe Deng, Kun He, and Wanxiang Che. Generating Natural Language Adversarial Examples through Probability Weighted Word Saliency. *Association for Computational Linguistics (ACL)*, 2019.
- Nestor Rodriguez and Sergio Rojas-Galeano. Shielding Googles Language Toxicity Model against Adversarial Attacks. *arXiv preprint arXiv:1801.01828*, 2018.
- Ludwig Schmidt, Shibani Santurkar, Dimitris Tsipras, Kunal Talwar, and Aleksander Mdry. Adversarially Robust Generalization Requires More Data. *Neural Information Processing Systems (NeurIPS)*, 2018.
- Chuanbiao Song, Kun He, Liwei Wang, and John Hopcroft. Improving the Generalization of Adversarial Training with Domain Adaptation. *International Conference on Learning Representations (ICLR)*, 2019.
- Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing Properties of Neural Networks. *International Conference on Learning Representations (ICLR)*, 2014.
- Eric Wong and Zico Kolter. Provable Defenses via the Convex Outer Adversarial Polytope. *International Conference on Machine Learning (ICML)*, 2018.
- Wei Emma Zhang, Quan Z. Sheng, and Ahoud Abdulrahmn F. Alhazmi. Generating Textual Adversarial Examples for Deep Learning Models: A Survey. *CoRR*, abs/1901.06796, 2019.
- Xiang Zhang, Junbo Zhao, and Yann LeCun. Character-level Convolutional Networks for Text Classification. *Neural Information Processing Systems (NeurIPS)*, 2015.

## A DETAILS OF IGA

Here we introduce our Improved Genetic Algorithm (IGA) in details and show how IGA differs from the first proposed generic attack method, GA (Alzantot et al., 2018). Regard a text as a chromosome, there are two operators in IGA:

- *Crossover*( $a, b$ ): For two texts  $a$  and  $b$  where  $a = \{w_1^a, \dots, w_{i-1}^a, w_i^a, w_{i+1}^a, \dots, w_n^a\}$  and  $b = \{w_1^b, \dots, w_{i-1}^b, w_i^b, w_{i+1}^b, \dots, w_n^b\}$ , randomly choose a crossover point  $i$  from 1 to  $n$ , and generate a new text  $c = \{w_1^a, \dots, w_{i-1}^a, w_i^a, w_{i+1}^b, \dots, w_n^b\}$ .
- *Mutate*( $x, w_i$ ): For a text  $x = \{w_1, \dots, w_{i-1}, w_i, w_{i+1}, \dots, w_n\}$  and a position  $i$ , replace  $w_i$  with  $\hat{w}_i$  where  $\hat{w}_i \in Syn(w, \sigma, k)$  to get a new text  $\hat{x} = \{w_1, \dots, w_{i-1}, \hat{w}_i, w_{i+1}, \dots, w_n\}$  that minimizes  $S_m(\hat{x}, y_{true})$ .

The details of IGA are described in Algorithm 2.

---

### Algorithm 2 The Improved Genetic Algorithm

---

**Input:**  $x$ : input text,  $y_{true}$ : true label for  $x$ ,  $M$ : maximum number of iterations

**Output:**  $x_{adv}$ : adversarial example

```

1: for each word  $w_i \in x$  do
2:    $\mathcal{P}_i^0 \leftarrow Mutate(x, w_i)$ 
3: end for
4: for  $g = 1 \rightarrow M$  do
5:    $x_{adv} = \arg \min_{x_i \in \mathcal{P}^g} S_m(x_i, y_{true})$ 
6:   if  $f(x_{adv}) \neq y_{true}$  then
7:     return  $x_{adv}$ 
8:   end if
9:    $\mathcal{P}_1^g \leftarrow x_{adv}$ 
10:  for  $i = 2 \rightarrow |\mathcal{P}^{g-1}|$  do
11:    Randomly sample  $parent_1, parent_2$  from  $\mathcal{P}^{g-1}$ 
12:     $child = Crossover(parent_1, parent_2)$ 
13:    Randomly choose a word  $w$  in  $child$ 
14:     $\mathcal{P}_i^g \leftarrow Mutate(child, w)$ 
15:  end for
16: end for
17: return  $x_{adv}$ 

```

---

Compared with GA, IGA has the following differences:

- *Initialization*: GA initializes the first population randomly, while IGA initializes the first population by replacing each word by its optimal synonym, so our population is more diversified.
- *Mutation*: Different from GA, IGA allows to replace the word that has been replaced before so that we can avoid local minimum.
- *Crossover*: To better simulate the reproduction and biological crossover, we randomly cut the text from two parents and concat two fragments into a new text rather than randomly choose a word of each position from the two parents.

The selection of the next generation is similar to GA, greedily choose one offspring, and then generate other offsprings by  $Mutate(Crossover(\cdot, \cdot))$  on two randomly chosen parents. But as *Mutate* and *Crossover* are different, IGA has very different offsprings.

## B HYPER-PARAMETER STUDY ON SEM

To explore how hyper-parameter  $\epsilon$  of SEM influences the efficacy, we try different  $\epsilon$  ranging from 0 to 1.2 for three models on IMDB with or without adversarial attacks. The results are illustrated in Figure 3.

On benign data, as shown in Figure 3(a), the classification accuracy of the models decreases a little when  $\epsilon$  increases. Because a bigger  $\epsilon$  indicates that we need less words to train the model, which

could degrade the efficacy of the models. Nevertheless, the classification accuracy does not decrease much as SEM could maintain the semantic invariance of the original text after encoding.

Then we show the defense efficacy of SEM on the three models when changing the value of  $\epsilon$ , as shown in Figure 3(b)-(d). When  $\epsilon = 0$  where SEM could not take any impact, we see that the accuracy is the lowest under all attacks. When  $\epsilon$  increases, SEM starts to defend the attacks, the accuracy increases rapidly and reach the peak when  $\epsilon = 0.5$ . Then the accuracy decays slowly if we continue to increase  $\epsilon$ . Thus, we choose  $\epsilon = 0.5$  to be a good trade-off on the accuracy of benign examples and adversarial examples.

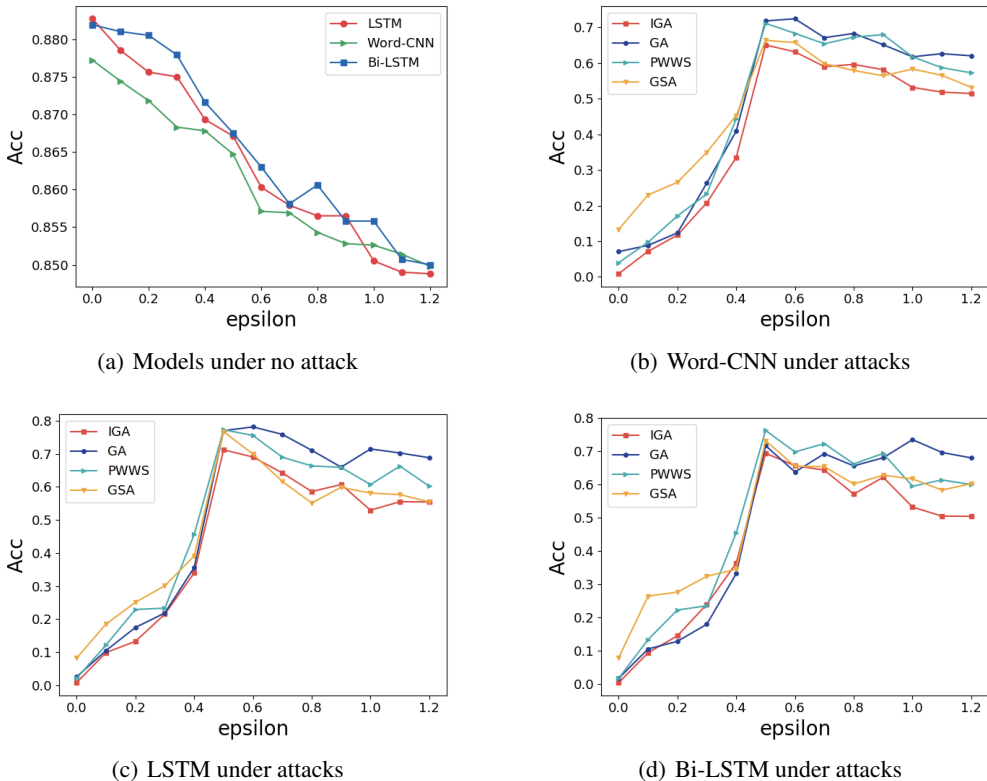


Figure 3: The classification accuracy for various  $\epsilon$  ranging from 0 to 1.2 for three models on IMDB.

### C ADVERSARIAL EXAMPLES GENERATED BY GA AND IGA

To show the generated adversarial examples, we randomly pick some benign examples from IMDB and generate adversarial examples by GA and IGA respectively on several models. The examples are shown in Table 5 to Table 6. We see that IGA substitutes less words than GA on these models under normal training.

Table 4: The adversarial examples generated by GA and IGA on IMDB using Word-CNN model.

	Confidence(%)	Prediction	Text
Original	97.9	1	I <b>enjoyed</b> this film which I thought was well written and acted , there was plenty of humour and a <b>provoking</b> story-line, a <b>warm</b> and enjoyable experience with an emotional ending.
	99.7	0	I am sorry but this is the <b>worst</b> film I have ever seen in my life. I cannot believe that after making the first one in the series, they were able to get a budget to make another. This is the least scary film I have ever watched and laughed all the way through to the end.
	95.8	1	This is a <b>unique</b> masterpiece made by the best director <b>ever</b> lived in the ussr. He knows the art of film making and can use it <b>very</b> well. If you find this movie, buy or copy it!
GA	50.6	0	I <b>cared</b> this film which I thought was well written and acted, there was plenty of humour and a <b>igniting</b> story-line, a <b>tepid</b> and enjoyable experience with an emotional ending.
	92.7	1	I am sorry but this is the <b>harshesht</b> film I have ever seen in my life. I cannot believe that after making the first one in the series, they were able to get a budget to make another. This is the least scary film I have ever watched and laughed all the way through to the end.
	59.0	0	This is a <b>sole</b> masterpiece made by the nicest director <b>permanently</b> lived in the ussr. He knows the art of film making and can use it <b>much</b> well. If you find this movie, buy or copy it!
IGA	88.3	0	I enjoyed this film which I think was well written and acted, there was plenty of humour and a <b>causing</b> story-line, a <b>lukewarm</b> and enjoyable experience with an emotional ending.
	70.8	1	I am sorry but this is the <b>hardest</b> film I have ever seen in my life. I cannot believe that after making the first one in the series, they were able to get a budget to make another. This is the least scary film I have ever watched and laughed all the way through to the end.
	54.8	0	This is a <b>sole</b> masterpiece made by the best director <b>permanently</b> lived in the ussr. He knows the art of film making and can use it very well. If you find this movie, buy or copy it!

Table 5: The adversarial examples generated by GA and IGA on IMDB using LSTM model.

	Confidence(%)	Prediction	Text
Original	99.9	1	I enjoyed this film which I thought was well written and acted , there was plenty of humour and a provoking story-line, a warm and enjoyable experience with an emotional ending.
	97.2	0	I am sorry but this is the worst film I have ever seen in my life. I cannot believe that after making the first one in the series, they were able to get a budget to make another. This is the least scary film I have ever watched and laughed all the way through to the end.
	99.7	1	This is a unique masterpiece made by the best director ever lived in the ussr. He knows the art of film making and can use it very well. If you find this movie, buy or copy it!
GA	88.2	0	I enjoyed this film which I thought was well written and proceeded, there was plenty of humorous and a igniting storyline, a tepid and enjoyable experience with an emotional terminate.
	99.9	1	I am sorry but this is the hardest film I have ever seen in my life. I cannot believe that after making the first one in the series they were able to get a budget to make another. This is the least terrifying film I have ever watched and laughed all the way through to the end.
	68.9	0	This is a unique masterpiece made by the best superintendent ever lived in the ussr. He knows the art of film making and can use it supremely alright. If you find this movie, buy or copy it!
IGA	72.1	0	I enjoyed this film which I thought was well written and acted, there was plenty of humour and a provoking story-line, a lukewarm and agreeable experience with an emotional ending.
	99.8	1	I am sorry but this is the hardest film I have ever seen in my life. I cannot believe that after making the first one in the series, they were able to get a budget to make another. This is the least scary film I have ever watched and laughed all the way through to the end.
	86.2	0	This is a sole masterpiece made by the best director ever lived in the ussr. He knows the art of film making and can use it very well. If you find this movie, buy or copy it!

Table 6: The adversarial examples generated by GA and IGA on IMDB using Bi-LSTM model.

	Confidence(%)	Prediction	Text
Original	99.6	1	I enjoyed this film which I thought was <b>well</b> written and <b>acted</b> , there was plenty of <b>humour</b> and a <b>provoking</b> story-line, a <b>warm</b> and enjoyable <b>experience</b> with an emotional <b>ending</b> .
	97.0	0	I am sorry but this is the <b>worst</b> film I have ever seen in my life. I cannot believe that after making the first one in the series, they were able to get a budget to make another. This is the least scary film I have ever watched and laughed all the way through to the end.
	95.3	1	This is a <b>unique</b> masterpiece made by the best <b>director</b> ever lived in the ussr. He knows the art of film making and can use it <b>very well</b> . If you find this movie, buy or copy it!
GA	98.2	0	I enjoyed this film which I thought was well written and proceeded, there was plenty of <b>humorous</b> and a <b>igniting</b> storyline, a <b>tepid</b> and enjoyable <b>experiment</b> with an emotional <b>terminate</b> .
	78.8	1	I am sorry but this is the <b>hardest</b> film I have ever seen in my life. I cannot believe that after making the first one in the series, they were able to get a budget to make another. This is the least scary film I have ever watched and laughed all the way through to the end.
	98.8	0	This is a unique masterpiece made by the best <b>superintendent</b> ever lived in the ussr. He knows the art of film making and can use it <b>supremely alright</b> . If you find this movie buy or copy it!
IGA	81.2	0	I enjoyed this film which I thought was <b>alright</b> written and acted, there was plenty of humour and a <b>arousing</b> storyline, a <b>lukewarm</b> and enjoyable <b>experiment</b> with an emotional ending.
	78.8	1	I am sorry but this is the <b>hardest</b> film I have ever seen in my life. I cannot believe that after making the first one in the series they were able to get a budget to make another. This is the least scary film I have ever watched and laughed all the way through to the end.
	98.9	0	This is a <b>sole</b> masterpiece made by the best director ever lived in the ussr. He knows the art of film making and can use it very <b>alright</b> . If you find this movie buy or copy it!