

Multi-granularity Textual Adversarial Attack with Behavior Cloning

Yangyi Chen^{1,2*}, Jin Su^{1,2* †}, Wei Wei^{1 ‡}

¹Cognitive Computing and Intelligent Information Processing Laboratory, School of Computer Science and Technology Huazhong University of Science and Technology

²School of Software Engineering, Huazhong University of Science and Technology

{yangyichen6666, sujintlenjoy}@gmail.com

weiw@hust.edu.cn

Abstract

Recently, the textual adversarial attack models become increasingly popular due to their successful in estimating the robustness of NLP models. However, existing works have obvious deficiencies. (1) They usually consider only a single granularity of modification strategies (e.g. word-level or sentence-level), which is insufficient to explore the holistic textual space for generation; (2) They need to query victim models hundreds of times to make a successful attack, which is highly inefficient in practice. To address such problems, in this paper we propose **MAYA**, a **Multi-grAnularitY Attack** model to effectively generate high-quality adversarial samples with fewer queries to victim models. Furthermore, we propose a reinforcement-learning based method to train a multi-granularity attack agent through behavior cloning with the expert knowledge from our MAYA algorithm to further reduce the query times. Additionally, we also adapt the agent to attack black-box models that only output labels without confidence scores. We conduct comprehensive experiments to evaluate our attack models by attacking BiLSTM, BERT and RoBERTa in two different black-box attack settings and three benchmark datasets. Experimental results show that our models achieve overall better attacking performance and produce more fluent and grammatical adversarial samples compared to baseline models. Besides, our adversarial attack agent significantly reduces the query times in both attack settings. Our codes are released at <https://github.com/Yangyi-Chen/MAYA>.

1 Introduction

Deep learning has been proven to be successful for many real-world applications such as spam filtering (Stringhini et al., 2010), autonomous driving

(Chen et al., 2017), and face recognition (Sun et al., 2015). However, these powerful models are vulnerable to adversarial samples, crafted by adding small, human-imperceptible perturbations to the input (Goodfellow et al., 2015; Szegedy et al., 2014). In the domain of computer vision, numerous adversarial attack models have been proposed to benchmark and interpret black-box deep learning models (Dong et al., 2018; Moosavi-Dezfooli et al., 2016; Carlini and Wagner, 2017; Kurakin et al., 2017) and corresponding defense methods have also been proposed to tackle adversarial security issues (Dziugaite et al., 2016; Xie et al., 2018; Kurakin et al., 2017; Tramèr et al., 2018).

However, crafting textual adversarial samples is more challenging due to the discrete and non-differentiable nature of text space. Indeed, most existing works focus on a single granularity of modification strategies, such as sentence-level (Jia and Liang, 2017; Iyyer et al., 2018), word-level (Zang et al., 2020; Ren et al., 2019) or character-level (Eger et al., 2019). Thus, none of such attack models find the optimal solution through multi-granularities for launching attacks simultaneously, which is more efficient to generate high-performance and effective adversarial samples while preserving semantic consistency and language fluency. To this end, we propose a simple and novel attack model targeting on multiple kinds of granularity called **MAYA**, which achieves higher attack success rate with fewer queries to victim models and produces high-quality adversarial samples compared to baseline attack models. Specifically, we add perturbations to the original sentence via rewriting its constituents according to the strict grammatical constraints.

Besides, almost all current attack models need to query victim models hundreds or even thousands of times to launch a successful attack¹ and

¹On average, PSO attack model (Zang et al., 2020) need to query victim models about 5000 times in SST-2

*Work done during internship at CCIIP

†Indicates equal contribution

‡Corresponding author

assume the victim models may output the confidence scores of their predictions, which is neither efficient nor practical in real-world situations. To alleviate such problems, we propose to train a multi-granularity attack agent called $MAYA_{\pi}$ through behavior cloning (Torabi et al., 2018) with the expert knowledge from our MAYA algorithm.

We conduct exhaustive experiments including attacking three victim models over three benchmark datasets in two different black-box settings, namely score-based and decision-based attack, to evaluate the effectiveness of our attack models. While the former supposes the labels and the confidence scores of the victim models are available, the latter assumes only the label information can be accessed while the other is unknown, which is more challenging and rarely investigated.

Experimental results demonstrate the superiority of our attack models. Specifically, MAYA overall outperforms all baseline models in terms of attack success rate, attack efficiency, and quality of adversarial samples. $MAYA_{\pi}$ achieves comparable attack success rate and adversarial samples quality with baseline models while significantly reduces the query times in two black-box settings. Furthermore, we apply $MAYA_{\pi}$ to attack open-source NLP frameworks to demonstrate its practicality and effectiveness in practice.

To summarize, the main contributions of this paper are as follows:

- Different from previous works that only concentrate on a single granularity, we propose an effective multi-granularity attack model to generate fluent and grammatical adversarial samples with fewer queries to victim models.
- We propose a RL-based method to train an agent through Behavior Cloning with the expert knowledge from our multi-granularity attack model and demonstrate its efficiency and power in two black-box settings, proving the effectiveness of our adapted imitation algorithm.
- We successfully handle the issues of decision-based black-box attack, which is rarely investigated in NLP.

2 Related Work

Existing textual adversarial attack models can be roughly categorized according to the granularity of modification, e.g., character-level, word-level, sentence-level.

Sentence-level attack models often contain paraphrasing original sentences following pre-defined syntax patterns (Iyyer et al., 2018), adding an irrelevant sentence to the end of the passage to distract models (Jia and Liang, 2017), and conducting domain shift on original sentences (Wang et al., 2020). However, sentence-level attacks usually neglect fine-grained granularity, such as word-level, resulting in low attack success rate.

Word-level attack is relatively more investigated and can be modeled as a combinatorial optimization problem (Zang et al., 2020), including finding substitution words and searching for adversarial samples. The methods of finding candidate substitutes mainly focus on the similarity of word embeddings (Jin et al., 2019), WordNet synonyms (Ren et al., 2019), HowNet synonyms (Zang et al., 2020), and Masked Language Model (MLM) (Li et al., 2020). Generally, the search algorithms involve greedy search algorithm (Ren et al., 2019; Liang et al., 2018; Jin et al., 2019), genetic algorithm (Alzantot et al., 2018), and particle swarm optimization (Zang et al., 2020). Although these attack models can achieve relatively high attacking performance, considering only a single granularity restricts the upper bound of word-level attack models' performance and almost all these models need to query victim models hundreds of times to launch a successful attack.

Character-level attacks make different modifications to words such as swapping, deleting, and inserting characters (Ebrahimi et al., 2018; Belinkov and Bisk, 2018; Gao et al., 2018). These attack models often craft ungrammatical adversarial samples and can be easily defended (Pruthi et al., 2019; Jones et al., 2020). Hence, in this work, we do not incorporate character-level modification into our multi-granularity framework.

To sum up, all above models only consider a single granularity and thus are insufficient in exploring the textual space for generation. So, we propose to launch attacks on multiple granularities in this paper. Experimental results demonstrate the effectiveness and efficiency of our method.

3 Methodology

In this section, we first describe our multi-granularity attack ($MAYA$) model in detail. Then we introduce how to train an attack agent, denoted as $MAYA_{\pi}$, with the knowledge from our MAYA algorithm. Finally, we describe how we adapt

MAYA $_{\pi}$ to perform decision-based black-box attack.

3.1 Multi-granularity Adversarial Attack

Our MAYA model incorporates three parts, namely generating adversarial candidates (**Generate**), verifying the successful attack (**Verify**), and picking the most potential candidate if no successful attack found (**Pick**). The whole process is shown as pseudocode in Appendix A.

Generate Given the input sentence $S = [w_0, \dots, w_i, \dots, w_n]$, we first conduct constituency parsing on the original sentence using SuPar (Zhang et al., 2020) to obtain its constituents. Then we generate adversarial candidates from two different perspectives.

First, for each constituent (including the whole sentence), i.e., each granularity of modification, except word-level, we employ various paraphrase models to generate adversarial samples via rewriting the specified constituents while keep the reset unchanged. However, such setting is solely a local modification which may cause syntactic inconsistency of the whole sentence, and thus we adopt the following rules to make the process more rational:

- The number of grammatical mistakes of the generated adversarial candidates must be less than or equal to the one of the original sentence, which can be checked by Language-Tool².
- The chosen adversarial candidate should be the one that is the most similar to the original one, i.e., preserving most of semantic information of the given sentence as much as possible. Specifically, Sentence-BERT (Reimers and Gurevych, 2019) is adopted for encoding the sentence and its candidates and we consider to employ a similarity function (e.g., cosine) to measure the semantic perseverance.

The filtered candidates are collected into a set (denoted as V_p).

Next, for word-level perturbation, we mask words in the original sentence one by one to generate corresponding adversarial candidates. Specifically, for w_i , we generate adversarial candidate $S_{w_i} = [w_0, \dots, [MASK], \dots, w_n]$. We collect all adversarial candidates generated in this way into a set (denoted as V_s).

Verify Given all adversarial candidates, we query the victim model for decisions and confidence

scores. If there doesn't exist an adversarial candidate successfully fools the victim model, we enter into the **Pick** step, which we will discuss later. If one or more than one successful adversarial candidates found, there are three different cases that we address differently. First, if all successful candidates come from V_p , we choose the one that retains the most semantics measured by cosine similarity of sentence embeddings as the final adversarial sample. Second, if successful candidates come from both V_p and V_s , we only choose the candidates from V_p following the same rule in the first case. The reason we ignore candidates from V_s is that we need to fill the $[MASK]$ token with substitutes and continually query the victim model for decisions, which is inefficient in the case we already have successful candidates from V_p . Finally, if all successful candidates come from V_s , we need to fill the $[MASK]$ token with substitutes to verify their success. Due to the same workflow in the **Pick** step, we directly view each successful candidate as S' and move to the second case in the **Pick** step.

Pick If no successful candidate found, we need to pick the most potential candidate as the new sentence and repeat the same **Generate** and **Verify** procedures to find the adversarial sample. Our criterion is the decrease of the victim model's confidence score. Here we denote the candidate that causes the biggest drop in the victim model's confidence score as S' . There are also two different cases. First, when S' comes from V_p , we directly choose S' as the most potential candidate and return to the **Generate** step. Second, when S' comes from V_s , we need to fill the $[MASK]$ token with substitutes to construct a complete sentence. Follow Li et al. (2020), we use MLM (Devlin et al., 2019) to generate k substitutes for the $[MASK]$ position in S' ³ and utilize WordNet (Fellbaum, 1998) to filter out antonyms of original words. Then we iteratively substitute $[MASK]$ token with candidates in probability descending order computed by MLM and query the victim model for confidence scores. If one substitution successfully fools the victim model, we return the whole sentence as the final adversarial sample. Otherwise, we obtain the sentence, denoted as S_w , that causes the biggest drop in the victim model's confidence score. We compare S_w with all candidates from V_p , choose

²<https://www.languagetool.org>

³We also experiment with other word substitution methods and find word substitutes generated from MLM work best.

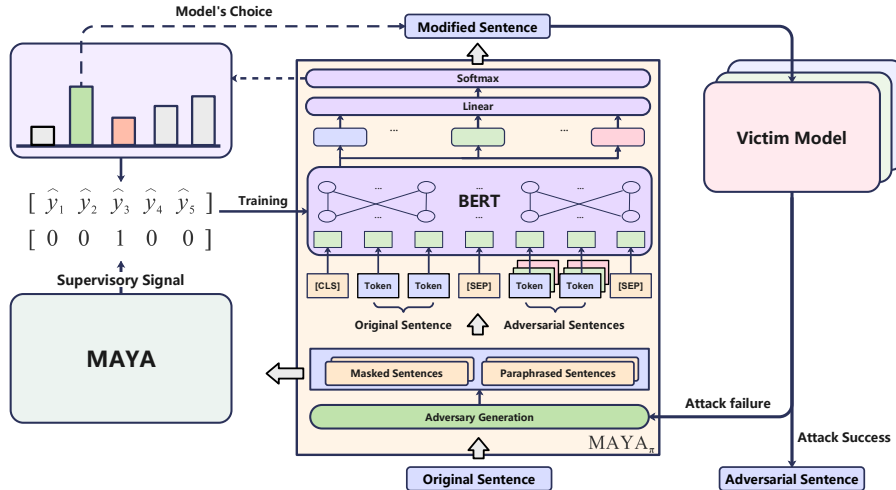


Figure 1: (Left) The training process of $MAYA_{\pi}$. We treat the selection of candidates as a multi-class classification problem with supervisory signals from our MAYA algorithm. (Middle) The architecture and workflow of $MAYA_{\pi}$. (Right) The process to launch an adversarial attack against victim models.

the one that causes the biggest drop in the confidence score as the most potential candidate, and return to the **Generate** step.

3.2 Combined with Behavior Cloning

As seen in Figure 1, we use $BERT_{base}$ (Devlin et al., 2019) and a linear classifier with one output unit as the architecture of $MAYA_{\pi}$. The core function of $MAYA_{\pi}$ is to predict the most potential candidate without querying the victim model. In this section, we first describe how we exploit $MAYA_{\pi}$ to launch an adversarial attack because we require $MAYA_{\pi}$ to perform the full procedure of attacking in the training process. And then we detail the training process.

3.2.1 Launch an Adversarial Attack

Now assume that we have already trained an attack agent $MAYA_{\pi}$. Given the input sentence S , we follow the same procedure as in the **Generate** step in MAYA algorithm to generate adversarial candidate sets V_p and V_s , corresponding to two different generation processes. Then, with the original sentence S and an adversarial candidate S_i concatenated as the input, $MAYA_{\pi}$ will output a score as a measure of tendency to choose this specific adversarial candidate. We obtain the candidate S' that get the highest score. Similarly, there are two different cases.

First, when S' comes from V_p , we directly use S' to query the victim model. If it successfully fools the victim model, we return S' as the final

adversarial sample. Otherwise, we view S' as the most potential candidate and return to the **Generate** step.

Second, when S' comes from V_s , we follow the same procedure as in the **Pick** step in MAYA algorithm that iteratively substitutes $[MASK]$ token with candidate words and query the victim model for confidence scores. If one successful candidate found, we directly return this sentence as the final adversarial sample. Otherwise, we view the candidate that causes the biggest drop in the confidence score as the most potential candidate and return to the **Generate** step. The whole process will be repeated until a successful adversarial sample found or all potential candidates have been encountered before. In the next subsection, we will describe how we adapt this score-based **Pick** step to launch a decision-based black-box attack.

3.2.2 Training Process

In this subsection, we describe our RL-based method to train $MAYA_{\pi}$ through Behavior Cloning with the expert knowledge from our MAYA algorithm. Specifically, we improve the training process by adapting the Dataset Aggregation (DAGGER) method (Ross et al., 2011). The training process incorporates three parts, namely initialization, sampling trajectories, and training.

Initialization We initialize $MAYA_{\pi}$ with pre-trained weights from BERT (Devlin et al., 2019) and a random initialized MLP. Also, we initialize

an empty trajectory dataset D .

Sampling Trajectories To train $MAYA_\pi$, we need to interact with the victim model to obtain the training data. Specifically, we train a local victim model that has the same architecture with the target victim model, expecting to approximate the decision boundary of the target victim model⁴.

We sample a batch of original sentences. For each sentence S_0 , we generate adversarial candidates C_1, \dots, C_k . As in the **Verify** and **Pick** steps in MAYA algorithm, one specific candidate will be chosen as the final successful adversarial sample or the most potential candidate. We view the candidate chosen by MAYA algorithm as the ground truth label and add $((S_0, C_1, \dots, C_k), label)$ to our trajectory dataset D .

To fully train an agent that can tackle different situations, we need a large dataset D . So, we adapt DAGGER method. Specifically, when receiving the ground truth label from MAYA algorithm, $MAYA_\pi$ doesn't take the golden action indicated by MAYA. It will take the action based on its own prediction. That is, $MAYA_\pi$ will predict which candidate will most confuse the victim model and follow its own procedure of launching an adversarial attack. The predicted candidate will be treated as S_0 and we will continue the same process of sampling trajectories to augment the dataset D .

Training Then we train $MAYA_\pi$ for only one epoch using the trajectory dataset D . We model the training task as a multi-class classification problem. For each sample $((S_0, C_1, \dots, C_k), label)$ drawn from D , we concatenate S_0 with each C_i . Then we input the k concatenated sentences to $MAYA_\pi$ to get k scores. We treat these k scores as logits and use the cross-entropy loss to train $MAYA_\pi$. After the training process, we clear D and continue the sampling procedure. The implementation details are described in Appendix I.

3.3 Adapted to Decision-based Attack

To adapt $MAYA_\pi$ to decision-based attack, we only need to modify one step in the attack procedure described in Section 3.2.1 while keep other steps unchanged. Specifically, when the candidate S' that gets the highest score from $MAYA_\pi$ is from V_s , we iteratively substitute $[MASK]$ token with candidate words to generate adversarial candidates

⁴It is unpractical to query the target victim model thousands of times to train an agent that is employed to attack this victim model.

and query the victim model for decisions. If one candidate successfully flips the label, we treat it as the final adversarial sample. Otherwise, to generate adversarial samples more efficiently, we take the candidate whose sentence embedding has the lowest cosine similarity with the sentence embedding of the original sentence as the most potential candidate. Our intuition is that the candidate that least resembles the original sentences is more likely to be a successful adversarial sample.

4 Experiments

We conduct comprehensive experiments to evaluate our attack models on the tasks of sentiment analysis, natural language inference, and news classification.

4.1 Datasets and Victim Models

For sentiment analysis, we choose SST-2 (Socher et al., 2013), a binary sentiment classification benchmark dataset. For natural language inference, we choose mismatched MNLI dataset (Williams et al., 2018). For news classification, we choose AG's News dataset. The models need to choose one of the four classes including World, Sports, Business, and Sci/Tech, given an instance in AG's News (Zhang et al., 2015).

We evaluate our attack models by attacking three victim models including BiLSTM (Schuster and Paliwal, 1997), BERT (Devlin et al., 2019), and RoBERTa (Liu et al., 2019). Details of the datasets and the classification accuracy of victim models are listed in Table 1.

4.2 Attack Models

We implement all baseline attack models using the NLP attack package TextAttack (Morris et al., 2020) and OpenAttack (Zeng et al., 2021).

Score-based Attack Models

We comprehensively compare our score-based attack models with five representative and strong score-based attack models including (1) GA+Embedding (Alzantot et al., 2018), (2) PWWS+Synonym (Ren et al., 2019), (3) PSO+Sememe (Zang et al., 2020), (4) TextFooler (Jin et al., 2019), (5) BERT-Attack (Li et al., 2020). Details of baseline models are listed in Appendix B and we describe details of our attack models below.

MAYA We select 3 open-source paraphrase models as the building blocks of MAYA. Specifically,

Dataset	#Class	Avg Len	Train	Dev	Test	BiLSTM Acc	BERT Acc	RoBERTa Acc
SST-2	2	20.94	6920	872	1821	82.65	91.76	94.89
MNLI	3	23.28	391075	0	9711	70.44	84.00	87.48
AG’s News	4	38.7	96000	24000	7600	91.26	93.75	94.14

Table 1: Detailed information of datasets and corresponding victim models’ accuracy. #Class denotes the classification number. Avg Len means the average sentence length. Train, Dev, and Test represent the number of samples in the training, validation, and test datasets. BiLSTM Acc, BERT Acc, and RoBERTa Acc denotes the original classification accuracy of each victim model.

we choose BaiDu translation API ⁵ to perform back translation of the original sentence, style-transfer based paraphrase model (Krishna et al., 2020), and T5 (Raffel et al., 2020) based paraphrase model ⁶.

MAYA_{bt} We observe from our preliminary experiments that only using back translation model can achieve comparable performance in most of the cases and be more computation efficient. So, we also implement MAYA using only back translation model, denoted as MAYA_{bt}.

MAYA_π Due to the similar performance of MAYA and MAYA_{bt} most of the time, we train our attack agent through behavior cloning with the expert knowledge from MAYA_{bt} in consideration of the efficiency of training and launching an adversarial attack.

Decision-based Attack Models

We consider two decision-based baseline models including (1) GAHard (Maheshwary et al., 2020) and (2) SCPN (Iyyer et al., 2018). Details of baseline models are listed in Appendix B. We conduct exhaustive experiments to compare our decision-based MAYA* with existing decision-based attack models.

4.3 Experimental Settings

Hyper-parameters For our attack models, we set the number of word substitutes k to 10. And for MAYA, to ensure the quality of successful adversarial samples, we discard adversarial samples with modification number larger than 8, 8, and 12 in SST-2, MNLI, and AG’s News respectively due to the difference of average sentence length in three datasets. Besides, we also set a maximum query number restriction to 15,000 for all attack models in the decision-based black-box attack setting due to the computation and time budget.

⁵<https://fanyi-api.baidu.com/>

⁶https://huggingface.co/Vamsi/T5_Paraphrase_Paws

Evaluation Metrics We evaluate the attack models considering their attack success rate, attack efficiency, and the quality of adversarial samples. (1) Attack success rate is defined as the percentage of adversarial samples that successfully fool the victim model. (2) Attack efficiency is defined as the average query number to the victim model of crafting an adversarial sample. (3) We use four different metrics, including grammaticality, fluency, validity, and naturality to evaluate adversarial samples’ quality. Specifically, we use Language-Tool to calculate the relative increase rate of grammar errors, GPT-2 (Radford et al., 2019) to compute adversarial samples’ perplexity as a measure of fluency, and ask human annotators to evaluate adversarial samples’ validity and naturality.

4.4 Experimental Results

Attack Success Rate The attack success rate (ASR) results in score-based attack setting are listed in Table 2 and the results in decision-based setting are listed in Appendix C. Considering score-based attack, MAYA consistently outperforms all baseline models in three datasets and three victim models and MAYA_π achieves comparable attack success rate with baseline models. In decision-based attack setting, MAYA* overall outperforms baseline models, especially in AG’s News because the sentences in AG’s News are much longer, providing more constituents to perturb. The results demonstrate the advantage of our multi-granularity attack models.

Attack Efficiency For score-based attack, our models especially MAYA_π show great superiority over all baseline models. For decision-based attack, MAYA* significantly outperforms GAHard, which needs thousands of queries.

Furthermore, we measure the attack success rate of attack models under the restriction of maximum query number. Figure 2 show all attack models’ attack success rate in SST-2 when attack-

Dataset	Victim Model Attack Method	BiLSTM				BERT				RoBERTa			
		ASR	Query	PPL	%I	ASR	Query	PPL	%I	ASR	Query	PPL	%I
SST-2	GA+Embedding	75.91	833.18	523.89	4.45	53.33	1004.49	568.20	7.46	50.23	1040.03	664.43	7.91
	PWWS+Synonym	87.75	130.68	658.68	9.05	75.12	145.45	837.91	8.35	77.03	149.31	1114.15	11.12
	PSO+Sememe	90.44	2912.45	748.89	4.91	85.60	5354.05	759.30	4.96	85.50	5648.08	741.09	4.96
	TextFooler	94.89	75.64	837.89	5.33	85.36	99.42	781.96	8.88	87.12	104.23	850.42	8.44
	BERT-Attack	98.65	52.76	675.75	10.15	90.36	81.44	686.76	15.32	95.01	81.67	683.64	13.82
	MAYA _{bt}	97.71	61.46	454.18	-6.11	<u>97.83</u>	71.52	414.38	-2.61	94.78	77.01	462.87	-4.08
	MAYA	98.52	59.11	411.67	-6.11	<u>98.81</u>	71.34	415.35	-4.16	96.64	75.38	437.32	-6.18
	MAYA _π	97.98	17.58	399.35	-7.62	94.40	18.32	435.86	-4.14	95.24	19.71	471.92	-6.53
MNLI	GA+Embedding	61.00	859.04	710.81	9.28	55.00	847.68	1080.79	10.98	41.80	857.52	847.06	10.63
	PWWS+Synonym	77.30	156.53	2509.04	9.81	75.10	151.50	948.40	10.67	71.60	151.24	715.42	9.86
	PSO+Sememe	80.80	5058.81	1208.33	5.20	75.80	4491.1	980.79	4.91	76.10	5178.30	1075.73	5.01
	TextFooler	<u>87.10</u>	102.12	1715.32	10.27	84.70	96.26	2656.92	8.03	83.60	92.91	2703.12	8.51
	BERT-Attack	84.90	81.15	5617.47	8.99	82.10	114.68	11761.92	16.07	84.10	66.13	10414.33	7.67
	MAYA _{bt}	84.10	85.14	434.94	-2.98	90.60	78.65	626.80	-0.60	86.50	69.85	413.13	-0.68
	MAYA	<u>87.10</u>	83.72	425.63	-3.61	91.70	76.00	599.44	-2.00	89.60	70.50	516.17	-3.89
	MAYA _π	81.20	33.60	504.22	-3.69	85.40	26.55	719.42	2.50	84.60	26.93	654.00	0.32
AG's News	GA+Embedding	67.40	1668.93	288.11	3.16	38.62	1490.97	399.80	7.77	34.33	1557.70	348.48	6.56
	PWWS+Synonym	74.70	227.37	364.58	5.70	65.50	251.82	544.43	8.79	54.70	254.31	491.69	10.12
	PSO+Sememe	89.40	14028.64	548.63	10.01	66.20	15461.61	680.31	8.87	64.40	17048.21	639.70	9.15
	TextFooler	84.20	163.36	335.13	3.32	88.70	215.61	600.53	8.79	78.20	232.08	562.13	6.47
	BERT-Attack	88.40	187.18	369.47	4.20	81.30	180.51	375.41	9.52	82.60	206.59	394.50	3.51
	MAYA _{bt}	84.60	215.61	206.47	-5.64	82.60	248.61	237.63	-11.36	78.10	234.29	221.82	-6.15
	MAYA	<u>91.29</u>	183.46	172.25	-13.09	93.10	202.74	206.93	-15.03	84.25	225.63	192.91	-12.82
	MAYA _π	<u>90.40</u>	45.57	215.10	-11.00	81.00	41.73	236.24	-17.45	77.40	36.48	244.99	-17.08

Table 2: The results of attacking performance and adversarial samples’ quality. ASR denotes the attack success rate. Query denotes the average query number of launching a successful adversarial attack. PPL and %I indicate adversarial samples’ fluency and relative increase of grammar errors. We also conduct Student’s t-tests to measure the difference between different models. Boldfaced numbers mean significant advantage with p-value 0.05 as the threshold and underline numbers mean no significant difference.

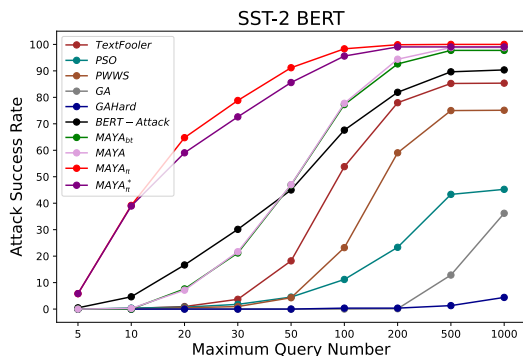


Figure 2: Attack success rate under restriction of maximum query number in SST-2 when attacking BERT.

SST-2	Victim Model Attack Method	BERT			
		ASR	Query	PPL	I
MAYA	ADV	18.81	20.15	411.78	-0.12
	VERB	27.02	25.14	346.05	1.69
	NOUN	37.14	40.80	358.66	0.75
	ADJ	43.57	22.89	362.80	0.20
	S	12.86	5.92	316.15	-22.10
	VP	15.83	9.16	357.01	-9.20
	NP	12.86	12.61	298.38	-12.24
	BERT-Attack	90.36	81.44	686.76	15.32
	Phrases	51.90	25.91	381.63	-13.74
	All	98.81	71.34	415.35	-4.16

Table 3: The results of attacking with restriction to the selection of different constituent types in SST-2.

ing BERT under the restriction of maximum query number. Appendix D shows remainder results in three datasets and three victim models. We observe that our RL-based attack models significantly outperform all baseline models under the restriction of maximum query number, demonstrating the practicality of our attack models in real-world situations. To confirm our argument, we present results of attacking two open-source NLP frameworks in Appendix F.

Adversarial Sample Quality We can observe from the results that our multi-granularity attack models overall outperform all baseline models considering adversarial samples’ fluency and relative increase of grammar errors. Human evaluation results presented in Appendix G also prove the high quality of our adversarial samples.

5 Further Analysis

5.1 Constituent Selection

It’s important to investigate which constituent type our multi-granularity attack models tend to select as the vulnerable part of the sentence and the impact of different constituent types. We first investigate the selection frequency of all constituent types. The results are listed in Appendix E. Then, we select 7 constituent types that are more common and

SST-2	Victim Model Attack Method	BiLSTM	RoBERTa
BERT	GA+Embedding	28.37	23.73
	PWWS+Synonym	28.07	21.66
	PSO+Sememe	38.83	40.22
	TextFooler	28.38	20.77
	BERT-Attack	28.12	29.19
	MAYA _{bt}	35.46	44.39
	MAYA	34.81	43.53
	MAYA _{π}	40.70	52.86
	GAHard	25.47	16.55
	MAYA* _{π}	39.17	39.17

Table 4: The transfer attack success rate of adversarial samples in SST-2.

often selected as the vulnerable parts of sentences. We restrict the selection of MAYA to each of these constituent types and evaluate the attacking performance. We also list the results of attacking with restriction to only words (BERT-Attack), with restriction to only phrases (Phrases), and with no restriction (All) for comparison⁷.

We can conclude from Table 3 that while word-level substitution (BERT-Attack) ensures the attacking performance, there still exists a significant gap between word-level attack and our multi-granularity attack (All). Besides, paraphrasing constituent types improves the quality of sentences due to our strict restrictions and can produce adversarial samples with some probability.

5.2 Transferability

We investigate the transferability of adversarial samples produced by all attack models in SST-2 with BERT as the victim model. We don’t consider SCPN in this transferability study because this method is model agnostic and cannot be directly compared with other attack models. We can observe from Table 4 that our MAYA _{π} attack agent crafts adversarial samples with significantly higher transferability. That’s probably because MAYA _{π} perturbs the sentence based not only on the outputs of victim models but also on its own prediction, ensuring adversarial samples to capture common vulnerabilities of different victim models.

5.3 Impact of Imitation Algorithm

Despite the strong attacking performance of MAYA _{π} , the impact of our adapted imitation algorithm is unknown. One may attribute the suc-

⁷We refer readers to (Taylor et al., 2003) for the meaning of syntax tags.

Victim Model	Dataset Attack Method	SST-2			
		ASR	Query	PPL	%I
BERT	MAYA _{π}	74.17	30.90	398.51	-8.08
	MAYA _{π}	94.40	18.32	435.86	-4.14

Table 5: The attacking performance of random initialized MAYA _{π} and our MAYA _{π} trained through behavior cloning.

Victim Model	Dataset Attack Method	SST-2			
		ASR	Query	PPL	%I
BiLSTM	MAYA _{π} ^O	97.98	17.58	399.35	-7.62
	MAYA _{π} ^B	92.87	20.14	434.99	-5.18
RoBERTa	MAYA _{π} ^O	95.24	19.71	471.92	-6.53
	MAYA _{π} ^B	94.90	18.51	425.16	-5.17

Table 6: The results of attacking performance when the victim model’s architecture is unknown.

cess to the capacity of the multi-granularity attack model, ignoring the contribution of the imitation learning process. So, we investigate the impact of our adapted imitation algorithm in this section. We employ a random initialized MAYA _{π} without interacting with local victim models to launch attacks against BERT in SST-2. From Table 5, we can conclude that the imitation learning process do bring some useful knowledge to our attack agent.

5.4 Limitation

We make a strong assumption in the development of our imitation algorithm that we assume that we have already known the victim models’ architectures, which is unrealistic in real-world situations. However, as mentioned in Section 4.4, our attack agent can successfully launch adversarial attacks against real-world NLP frameworks, which confirms the practicability of our imitation algorithm. The results are presented in Appendix F. Further, we investigate the attacking performance of our attack agents in SST-2 dataset when the victim model’s architecture is unknown. Specifically, we employ the attack agent trained by interacting with our local BERT model, denoted as MAYA _{π} ^B to launch adversarial attacks against BiLSTM and RoBERTa. We compare the attacking performance with that of originally trained attack agent, denoted as MAYA _{π} ^O.

We can observe from Table 6 that MAYA _{π} ^B achieves similar attacking performance while maintain the attack efficiency and quality of adversarial samples, especially in RoBERTa. That’s probably because the common features and architectures shared by pre-trained models, which strongly supports our view that our attack agents can indeed

cause significant drop in models’ prediction accuracy even though we the victim models’ architectures are unknown because we can just assume the black-box system is built on pre-trained models, which can be verified in most cases.

6 Conclusion and Future Work

In this paper, we propose a multi-granularity adversarial attack model (MAYA) and propose a RL-based method to train an attack agent ($MAYA_{\pi}$) through behavior cloning with the expert knowledge from our MAYA algorithm. Further, we adapt $MAYA_{\pi}$ to decision-based attack setting, handling the issues of attacking models that only output decisions. Experimental results show that our attack models achieve overall higher attacking performance and produce more fluent and grammatical adversarial samples. We also show that $MAYA_{\pi}$ can launch adversarial attacks towards open-source NLP frameworks, demonstrating the practicability of our attack agent in real-world situations.

In the future, we will focus on how to improve current models’ robustness towards multi-granularity attacks. In addition, we will try to apply $MAYA_{\pi}$ to other less investigated settings in textual adversarial attack, such as to launch target attacks in decision-based attack setting.

7 Acknowledgements

This work was supported in part by the National Natural Science Foundation of China under Grant 61602197. Also, we thank all the anonymous reviewers for their valuable comments and suggestions.

8 Ethical Considerations

In this section, we discuss the potential broader impact and ethical considerations of our paper.

Intended use. In this paper, we propose multi-granularity attacking models that can handle different attack settings with superior performance. Our motivations are twofold. First, we can find some insights from the experimental results about current black-box machine learning models that can help us move towards explainable AI. Second, we demonstrate the potential risks of deploying current models in the real world, encouraging the research community to develop more robust models.

Potential risk. It’s possible that our attacking models may be maliciously used to launch an adversarial attack against off-the-shelf commercial

systems. However, according to the research on adversarial attack on computer vision, it’s important to make the research community realize these powerful attacking models before defending them. So, studying and investigating adversarial attack is significant.

Energy saving. We present the details of our training process in Appendix J to prevent people from making unnecessary hyper-parameters tuning and help researchers to quickly reproduce our results. We will also release the checkpoints including all victim models and our attack agents to avoid energy costs to re-train them.

References

- Moustafa Alzantot, Yash Sharma, Ahmed Elgohary, Bo-Jhang Ho, Mani Srivastava, and Kai-Wei Chang. 2018. [Generating natural language adversarial examples](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2890–2896, Brussels, Belgium. Association for Computational Linguistics.
- Yonatan Belinkov and Yonatan Bisk. 2018. [Synthetic and natural noise both break neural machine translation](#). In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net.
- Nicholas Carlini and David Wagner. 2017. Towards evaluating the robustness of neural networks. In *2017 IEEE Symposium on Security and Privacy (SP)*, pages 39–57. IEEE.
- Xiaozhi Chen, Huimin Ma, Ji Wan, Bo Li, and Tian Xia. 2017. [Multi-view 3d object detection network for autonomous driving](#). In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 6526–6534. IEEE Computer Society.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Yinpeng Dong, Fangzhou Liao, Tianyu Pang, Hang Su, Jun Zhu, Xiaolin Hu, and Jianguo Li. 2018. [Boosting adversarial attacks with momentum](#). In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 9185–9193. IEEE Computer Society.

- Zhendong Dong, Qiang Dong, and Changling Hao. 2010. [HowNet and its computation of meaning](#). In *Coling 2010: Demonstrations*, pages 53–56, Beijing, China. Coling 2010 Organizing Committee.
- Gintare Karolina Dziugaite, Zoubin Ghahramani, and Daniel M Roy. 2016. [A study of the effect of jpg compression on adversarial images](#). *ArXiv preprint*, abs/1608.00853.
- Javid Ebrahimi, Daniel Lowd, and Dejing Dou. 2018. [On adversarial examples for character-level neural machine translation](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 653–663, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Steffen Eger, Gözde Gül Şahin, Andreas Rücklé, Ji-Ung Lee, Claudia Schulz, Mohsen Mesgar, Krishkant Swarnkar, Edwin Simpson, and Iryna Gurevych. 2019. [Text processing like humans do: Visually attacking and shielding NLP systems](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1634–1647, Minneapolis, Minnesota. Association for Computational Linguistics.
- Christiane Fellbaum. 1998. *WordNet: An Electronic Lexical Database*. Bradford Books.
- Ji Gao, Jack Lanchantin, Mary Lou Soffa, and Yanjun Qi. 2018. [Black-box generation of adversarial text sequences to evade deep learning classifiers](#). In *2018 IEEE Security and Privacy Workshops (SPW)*, pages 50–56.
- Matt Gardner, Joel Grus, Mark Neumann, Oyvind Tafjord, Pradeep Dasigi, Nelson F. Liu, Matthew Peters, Michael Schmitz, and Luke Zettlemoyer. 2018. [AllenNLP: A deep semantic natural language processing platform](#). In *Proceedings of Workshop for NLP Open Source Software (NLP-OSS)*, pages 1–6, Melbourne, Australia. Association for Computational Linguistics.
- Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. 2015. [Explaining and harnessing adversarial examples](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Mohit Iyyer, John Wieting, Kevin Gimpel, and Luke Zettlemoyer. 2018. [Adversarial example generation with syntactically controlled paraphrase networks](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1875–1885, New Orleans, Louisiana. Association for Computational Linguistics.
- Robin Jia and Percy Liang. 2017. [Adversarial examples for evaluating reading comprehension systems](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2021–2031, Copenhagen, Denmark. Association for Computational Linguistics.
- Di Jin, Zhijing Jin, Joey Tianyi Zhou, and Peter Szolovits. 2019. [Is bert really robust? natural language attack on text classification and entailment](#). *ArXiv*, abs/1907.11932.
- Erik Jones, Robin Jia, Aditi Raghunathan, and Percy Liang. 2020. [Robust encodings: A framework for combating adversarial typos](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2752–2765, Online. Association for Computational Linguistics.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Kalpesh Krishna, John Wieting, and Mohit Iyyer. 2020. [Reformulating unsupervised style transfer as paraphrase generation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 737–762, Online. Association for Computational Linguistics.
- Alexey Kurakin, Ian Goodfellow, and Samy Bengio. 2017. [Adversarial examples in the physical world](#).
- Linyang Li, Ruotian Ma, Qipeng Guo, Xiangyang Xue, and Xipeng Qiu. 2020. [BERT-ATTACK: Adversarial attack against BERT using BERT](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6193–6202, Online. Association for Computational Linguistics.
- Bin Liang, Hongcheng Li, Miaoqiang Su, Pan Bian, Xirong Li, and Wenchang Shi. 2018. [Deep text classification can be fooled](#). In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden*, pages 4208–4215. ijcai.org.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized bert pretraining approach](#). *ArXiv preprint*, abs/1907.11692.
- Rishabh Maheshwary, Saket Maheshwary, and Vikram Pudi. 2020. [Generating natural language attacks in a hard label black box setting](#).
- Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, and Pascal Frossard. 2016. [Deepfool: A simple and accurate method to fool deep neural networks](#). In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pages 2574–2582. IEEE Computer Society.

- John Morris, Jin Yong Yoo, and Yanjun Qi. 2020. [TextAttack: Lessons learned in designing python frameworks for NLP](#). In *Proceedings of Second Workshop for NLP Open Source Software (NLP-OSS)*, pages 126–131, Online. Association for Computational Linguistics.
- Danish Pruthi, Bhuwan Dhingra, and Zachary C. Lipton. 2019. [Combating adversarial misspellings with robust word recognition](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5582–5591, Florence, Italy. Association for Computational Linguistics.
- Peng Qi, Yuhao Zhang, Yuhui Zhang, Jason Bolton, and Christopher D. Manning. 2020. [Stanza: A python natural language processing toolkit for many human languages](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 101–108, Online. Association for Computational Linguistics.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *Journal of Machine Learning Research*, 21(140):1–67.
- Nils Reimers and Iryna Gurevych. 2019. [SentenceBERT: Sentence embeddings using Siamese BERT-networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992, Hong Kong, China. Association for Computational Linguistics.
- Shuhuai Ren, Yihe Deng, Kun He, and Wanxiang Che. 2019. [Generating natural language adversarial examples through probability weighted word saliency](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1085–1097, Florence, Italy. Association for Computational Linguistics.
- Stéphane Ross, Geoffrey Gordon, and Drew Bagnell. 2011. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pages 627–635. JMLR Workshop and Conference Proceedings.
- Mike Schuster and Kuldip K Paliwal. 1997. Bidirectional recurrent neural networks. *IEEE transactions on Signal Processing*, 45(11):2673–2681.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Ng, and Christopher Potts. 2013. [Recursive deep models for semantic compositionality over a sentiment treebank](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1631–1642, Seattle, Washington, USA. Association for Computational Linguistics.
- Gianluca Stringhini, Christopher Kruegel, and Giovanni Vigna. 2010. [Detecting spammers on social networks](#). pages 1–9.
- Yi Sun, Ding Liang, Xiaogang Wang, and Xiaoou Tang. 2015. [Deepid3: Face recognition with very deep neural networks](#). *ArXiv preprint*, abs/1502.00873.
- Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian J. Goodfellow, and Rob Fergus. 2014. [Intriguing properties of neural networks](#). In *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*.
- Ann Taylor, Mitchell Marcus, and Beatrice Santorini. 2003. The penn treebank: an overview. *Treebanks*, pages 5–22.
- Faraz Torabi, Garrett Warnell, and Peter Stone. 2018. [Behavioral cloning from observation](#). In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden*, pages 4950–4957. ijcai.org.
- Florian Tramèr, Alexey Kurakin, Nicolas Papernot, Ian J. Goodfellow, Dan Boneh, and Patrick D. McDaniel. 2018. [Ensemble adversarial training: Attacks and defenses](#). In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net.
- Tianlu Wang, Xuezhi Wang, Yao Qin, Ben Packer, Kang Li, Jilin Chen, Alex Beutel, and Ed Chi. 2020. [CAT-gen: Improving robustness in NLP models via controlled adversarial text generation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5141–5146, Online. Association for Computational Linguistics.
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. [A broad-coverage challenge corpus for sentence understanding through inference](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1112–1122, New Orleans, Louisiana. Association for Computational Linguistics.
- Cihang Xie, Jianyu Wang, Zhishuai Zhang, Zhou Ren, and Alan L. Yuille. 2018. [Mitigating adversarial effects through randomization](#). In *6th International Conference on Learning Representations*,

ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings. OpenReview.net.

Yuan Zang, Fanchao Qi, Chenghao Yang, Zhiyuan Liu, Meng Zhang, Qun Liu, and Maosong Sun. 2020. [Word-level textual adversarial attacking as combinatorial optimization](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6066–6080, Online. Association for Computational Linguistics.

Guoyang Zeng, Fanchao Qi, Qianrui Zhou, Tingji Zhang, Zixian Ma, Bairu Hou, Yuan Zang, Zhiyuan Liu, and Maosong Sun. 2021. [OpenAttack: An open-source textual adversarial attack toolkit](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing: System Demonstrations*, pages 363–371, Online. Association for Computational Linguistics.

Xiang Zhang, Junbo Jake Zhao, and Yann LeCun. 2015. [Character-level convolutional networks for text classification](#). In *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pages 649–657.

Yu Zhang, Houquan Zhou, and Zhenghua Li. 2020. [Fast and accurate neural CRF constituency parsing](#). In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020*, pages 4046–4053. ijcai.org.

A MAYA Algorithm

The whole process of MAYA Algorithm is shown in Algorithm 1.

B Details of Baseline Attack Models

We describe details of baseline models in this section.

TextFooler This model (Jin et al., 2019) is a score-based attack method ranking words in the input sentence by saliency and chooses substitutes based on the word embedding to construct an adversarial sample.

PWWS+Synonym This model (Ren et al., 2019) is a greedy algorithm using augmented word saliency to iteratively substitute words with synonyms from WordNet (Fellbaum, 1998).

GA+Embedding This model (Alzantot et al., 2018) uses classical genetic algorithm to search for an adversarial sample. It relies on word embeddings similarity to select word substitutions.

PSO+Sememes This model (Zang et al., 2020) considers adversarial attack as a combinatorial optimization problem, and then designs an algorithm based on particle swarm optimization to substitute a word with its synonyms from HowNet (Dong et al., 2010).

BERT-Attack This model (Li et al., 2020) iteratively masks each word in original sentences, obtain substitutes from MLM and greedily select the substitute that causes the biggest drop in victim models’ confidence scores.

SCPN This model (Iyyer et al., 2018) predefines several syntax patterns and paraphrase the original sentence to different syntax structure, intending to find an adversarial sample.

GAHard This model (Maheshwary et al., 2020) involves initializing an potential adversarial sample, search space reduction, and population based optimization to find a semantic preserved adversarial sample.

C Decision-based Experiment Results

The decision-based experiment results are shown in Table 11. Note that the number of syntax templates is fixed in SCPN, making the query number constant. So, the query number of SCPN cannot be directly compared with that of other decision-based attack models and we leave out all query number results of SCPN in Table 11.

D Attack Efficiency

In this section, we show the results of attack efficiency in three datasets and three victim models. Figure 3 and Figure 4 show results of attacking BiLSTM and RoBERTa in SST-2. Figure 5-7 show results of attacking BiLSTM, BERT, and RoBERTa in MNLI. Figure 8-10 show results of attacking BiLSTM, BERT, and RoBERTa in AG’s News.

E Constituent Selection

Table 8 show the selection frequency of all constituent types.

F Attack Open-source NLP Frameworks

In this section, we show that our attack agent can be employed to attack open-source NLP frameworks, including AllenNLP (Gardner et al., 2018) and Stanza Qi et al. (2020). We choose the sentiment analysis task and use the attack agent MAYA π

Algorithm 1 MAYA

Input: Original sentence S ; Victim Model M **Output:** Adversarial sample S_a

```
1: Generate Procedure
2:   // generate candidates by paraphrasing constituent parts
3:   get constituent parts list  $C = \text{ConstituentParser}(S)$ 
4:   for each constituent part  $c$  in  $C$  do
5:     get paraphrases list  $P_c = \text{ParaphraseModels}(c)$  // use several paraphrase models
6:     get adversary candidates set  $V_p = \cup \text{filter}(P_c)$  // based on semantics and grammar rules
7:     // generate candidates by masking original words
8:     for each word  $w$  in  $S$ :
9:       mask  $w$  to get  $S_{w_i} = [w_0, \dots, [MASK], \dots, w_n]$ 
10:    get adversary candidates set  $V_s = \{S_{w_1}, \dots, S_{w_i}, \dots, S_{w_n}\}$ 
11: Verify Procedure
12:   query the victim model for  $scores = M(V_p \cup V_s)$ 
13:   if exist candidates list  $L \subseteq (V_p \cup V_s)$  successfully fool  $M$  then
14:     // we detail this process in
15:     return the most appropriate adversary sentence  $S_a \in L$  or go to the Pick Procedure
16:   else
17:     go to the Pick Procedure
18: Pick Procedure
19:   select the most confusing candidate  $S_a = \text{Select}(V_p \cup V_s, scores)$ 
20:   if  $S_a \in V_p$  then
21:     return  $S_a$ 
22:   else
23:     get substitute words list  $W = \text{MLM}(S_a)$  //  $S_a \in V_s$  and we choose  $k$  substitute words
24:     get complete sentences list  $L_s = \text{Fill}(S_a, W)$ 
25:     fill the  $[MASK]$  token and query the victim model for  $w\_scores = M(L_s)$ 
26:     if exist  $s \in L_s$  successfully fool the model then
27:       return  $s$ 
28:     else
29:       obtain the most confusing sentence  $s_p$  from  $V_p$ 
30:       compare and return the most potential candidate  $S_a = \text{Choose}(L_s, s_p, scores, w\_scores)$ 
```

trained by interacting with BERT trained on SST-2 dataset. Notice that we conduct score-based adversarial attack on AllenNLP sentiment analysis model and decision-based attack on Stanza model according to the outputs of the victim models. Table 9 show the results. We observe that our two attack agents can function well in real-world situations in two different attack settings and produce high-quality adversarial samples, showing the potential vulnerability of current NLP systems.

G Human Evaluation

We set up human evaluation to further evaluate the quality of our adversarial samples. Follow Zang et al. (2020), we consider 2 evaluation metrics including validity and naturality. Due to the large

number of baseline models, we directly compare our crafted adversarial samples with original samples to evaluate the quality of our adversarial samples. We randomly sample 100 original sentences from SST-2 dataset and 100 adversarial samples crafted by MAYA in SST-2 dataset and mix them. For each sentence, we ask 3 human annotators to do normal sentiment classification task and score this sentence's naturality from 1-5. We use the voting strategy to produce the annotation results of validity for each adversarial sample. Specifically, we respectively measure the human annotators' accuracy on original and adversarial samples and view the difference of accuracy as an indicator of adversarial samples' validity. And we average 3 annotators' naturality scores to get the final results.

Sample I	
Original Sentence	The second-ranked Jayhawks can redeem themselves for one of their most frustrating losses last season Monday when they welcome the Wolf Pack to Allen Fieldhouse .
Adversarial Sentence	The second-ranked Jayhawks can redeem themselves for one of their most frustrating losses last season Monday when the Wolf Pack is welcomed to the semifinals .
Sample II	
Original Sentence	If the playoffs opened right now, instead of next month, the A #39 ; s would face the Red Sox in the first round – again . Boston bounced Oakland out of the postseason in five games last year, coming back from a 2-0 deficit to do so .
Adversarial Sentence	If the playoffs opened right now, instead of next month, the A #39 . s would face the Red Sox in the first round – once again . Boston bounced Oakland out of the postseason in five games last year , coming back from a deficit .

Table 7: Adversarial samples crafted by our multi-granularity attack models in AG’s News.

Constituent Type	ADJ	NOUN	VERB	VP	S	NP	ADV	PP	PUNCT	PRON	ADJP	ADP	CCONJ	PART	DET
Total Num	287	223	175	137	129	113	92	56	54	35	34	33	28	28	24
Constituent Type	SBAR	SCONJ	AUX	PROPN	FRAG	ADV	NUM	NAC	INTJ	WHPP	NX	SQ	PRN	X	
Total Num	19	19	15	9	9	6	2	2	1	1	1	1	1	1	1

Table 8: The frequency of selection of all constituents in SST-2.

API	Dataset Attack Method	SST-2			
		ASR	Query	PPL	%I
AllenNLP (Score-based)	MAYA _π	93.94	21.32	422	-5.39
Stanza (Decision-based)	MAYA _π *	94.16	17.89	406.36	-4.92

Table 9: The results of attacking open-source sentiment analysis models using MAYA_π.

Dataset	Victim Model Sample	BERT	
		Accuracy	Naturality
SST-2	Original	88.00	4.24
	Adversarial	84.00	3.96

Table 10: The results of human evaluation

From Table 10, the close gap of accuracy between original and adversarial samples indicates that our adversarial samples maintain high validity. Besides, our adversarial samples also achieve high naturality, which is consistent with automatic evaluation metrics in our main experiments.

H Case Study

We select 2 successful adversarial samples crafted by our multi-granularity attack models in AG’s News. Note that other baseline attack models all fail in these two samples. We can observe from Table 7 the strength of our models is twofold. First, from Sample I, our attack models can take different kinds of granularity into consideration, making a bigger search space and crafting more diversified adversarial samples. Second, from Sample II, our attack models can combine different kinds of granularity perturbations to launch a stronger adversarial

attack.

I Implementation Details

For our reinforcement learning, we use standard Adam (Kingma and Ba, 2015) to train our agent and consistently set the learning rate to 2e-5 because our training process is based on data aggregation, meaning that the training data can be abundant. And we set batch size to 16.

Due to the limitation of GPU memory and computation resources, we use some tricks to get average batch gradients. Given a batch of original sentences list $[S_1, \dots, S_i, \dots, S_n]$, we input the sentence one by one to our attack agent with all adversarial candidates of original sentences S_i and compute the cross-entropy loss l_i with the golden label from MAYA algorithm. Here, we denote the number of adversarial candidates as k_i and we get weighted loss L_i by multiplying l_i with k_i :

$$L_i = l_i \times k_i \quad (1)$$

Then we directly perform back propagation to get the gradients for each parameters:

$$g_i = \nabla_{\theta} L_i \quad (2)$$

We save the gradients and repeat above operations to accumulate the gradients. Finally, we have:

$$G = \sum g_i \quad (3)$$

When reach the batch size, we normalize the gradients and update the parameters.

$$G' = \frac{G}{\sum k_i} \quad (4)$$

Dataset	Victim Model Attack Method	BiLSTM				BERT				RoBERTa			
		ASR	Query	PPL	%I	ASR	Query	PPL	%I	ASR	Query	PPL	%I
SST-2	SCPN	62.05	-	471.21	-10.83	52.74	-	467.93	-10.26	53.25	-	532.45	-17.43
	GAhard	91.92	6584.22	739.26	6.03	81.79	7460.69	747.24	6.81	75.06	7589.75	1179.88	6.27
	MAYA π	94.21	21.93	428.67	-1.59	94.88	21.81	440.31	-3.81	92.92	23.40	484.11	-5.37
MNLI	SCPN	54.10	-	570.20	4.99	59.80	-	465.24	-1.64	61.10	-	446.63	1.25
	GAhard	74.30	6756.56	1927.06	9.18	75.90	6132.63	2334.05	6.17	78.50	6108.29	6822.60	7.91
	MAYA π	72.30	35.48	664.00	-3.05	76.00	28.41	742.28	-1.24	77.60	25.91	693.62	-5.84
AG's News	SCPN	61.90	-	550.43	-28.81	52.50	-	649.12	-35.33	57.50	-	778.08	-29.59
	GAhard	82.10	8505.12	278.51	3.43	57.30	8882.92	360.20	4.58	53.50	9209.08	347.19	3.82
	MAYA π	82.90	48.75	226.21	-8.47	80.60	64.66	246.66	-21.72	76.60	46.07	263.62	-18.66

Table 11: The results of attack performance and adversarial samples' quality in decision-based attack setting.

J Experiment Running Environment

We conduct all experiments on a server whose major configurations are as follows: (1) CPU: Intel(R) Xeon(R) E5-2620 v4 @ 2.10GHz (2) RAM: 125GB; (3) GPU: RTX1080 , 11GB memory. The operation system is Ubuntu 16.04.7 LTS (GNU/Linux 4.15.0-142-generic x86_64). We use PyTorch 1.7.1 as the programming framework.

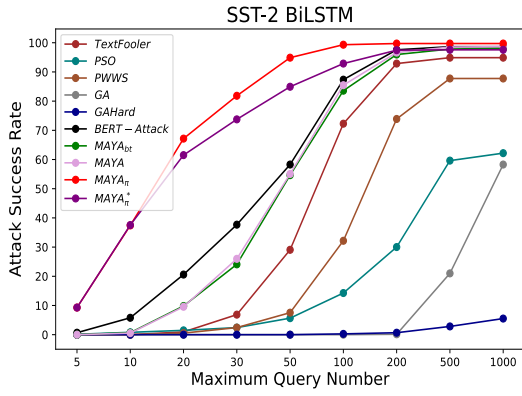


Figure 3: Attack BiLSTM in SST-2

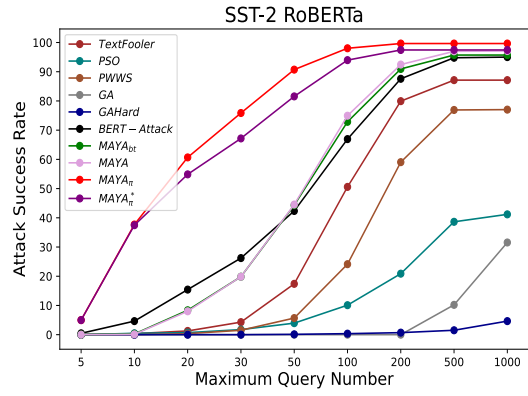


Figure 4: Attack RoBERTa in SST-2

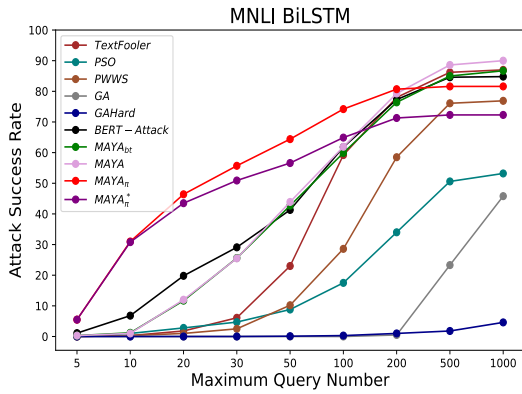


Figure 5: Attack BiLSTM in MNLI

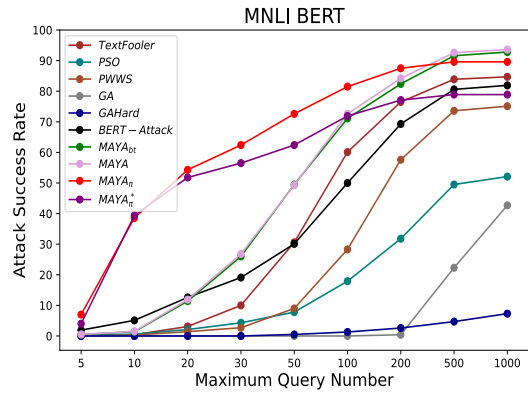


Figure 6: Attack BERT in MNLI

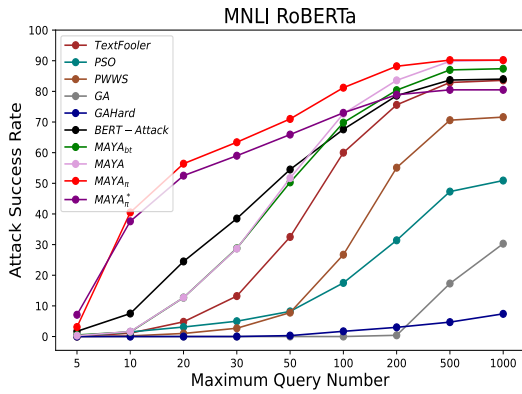


Figure 7: Attack RoBERTa in MNLI

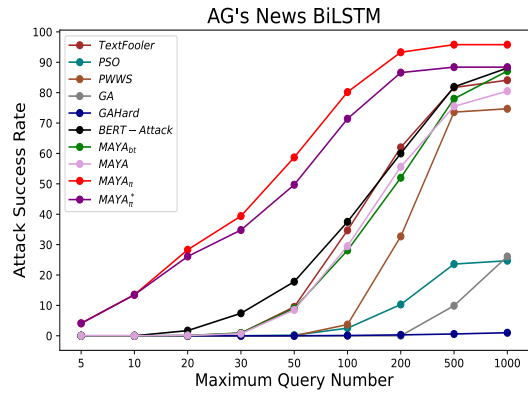


Figure 8: Attack BiLSTM in AG's News

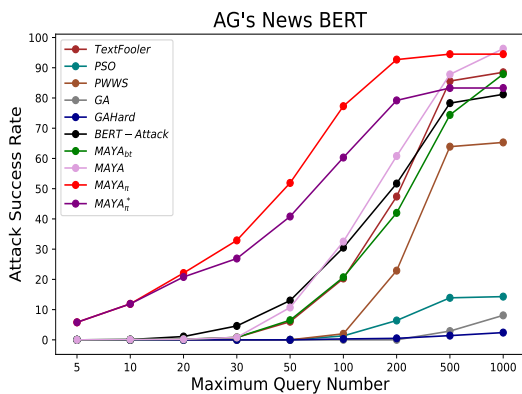


Figure 9: Attack BERT in AG's News

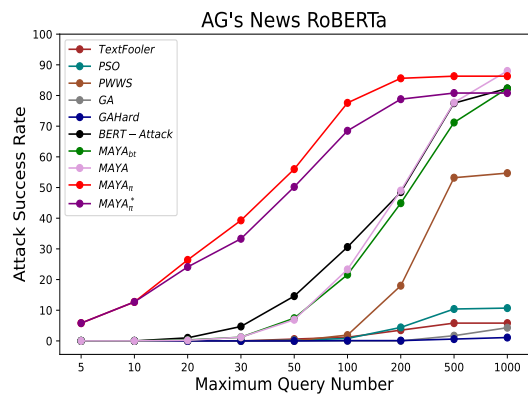


Figure 10: Attack RoBERTa in AG's News