

# A Classification-Guided Approach for Adversarial Attacks against Neural Machine Translation

**Sahar Sarizadeh**

EPFL, Switzerland  
sahar.sadrizadeh@epfl.ch

**Ljiljana Dolamic**

Armasuisse S+T, Switzerland  
ljiljana.dolamic@ar.admin.ch

**Pascal Frossard**

EPFL, Switzerland  
pascal.frossard@epfl.ch

## Abstract

Neural Machine Translation (NMT) models have been shown to be vulnerable to adversarial attacks, wherein carefully crafted perturbations of the input can mislead the target model. In this paper, we introduce ACT, a novel adversarial attack framework against NMT systems guided by a classifier. In our attack, the adversary aims to craft meaning-preserving adversarial examples whose translations in the target language by the NMT model belong to a different class than the original translations. Unlike previous attacks, our new approach has a more substantial effect on the translation by altering the overall meaning, which then leads to a different class determined by an oracle classifier. To evaluate the robustness of NMT models to our attack, we propose enhancements to existing black-box word-replacement-based attacks by incorporating output translations of the target NMT model and the output logits of a classifier within the attack process. Extensive experiments, including a comparison with existing untargeted attacks, show that our attack is considerably more successful in altering the class of the output translation and has more effect on the translation. This new paradigm can reveal the vulnerabilities of NMT systems by focusing on the class of translation rather than the mere translation quality as studied traditionally.

## 1 Introduction

Recently, deep neural networks have emerged as powerful tools in various domains, such as Natural Language Processing (NLP) (Vaswani et al., 2017) and computer vision (He et al., 2016). Despite their exceptional performance, these models have been shown to be susceptible to slight perturbations to their inputs, known as adversarial attacks (Szegedy et al., 2014; Moosavi-Dezfooli et al., 2016; Madry et al., 2018). In particular, adversarial examples closely resemble the original input and can deceive the target model to generate incorrect outputs. Ex-

tensive research has also been devoted to adversarial attacks against NLP models (Jin et al., 2020; Li et al., 2020; Zang et al., 2020; Guo et al., 2021; Wang et al., 2022; Zou et al., 2023) since NLP models are increasingly employed in practical systems. These studies have mainly focused on *text classification* tasks such as sentiment classification and natural language inference. In text classification, the adversary aims to fool the target model into misclassifying the input sentence as a specific wrong class (targeted attacks) or any class other than the correct ground-truth class (untargeted attacks).

Another important task in NLP is *Neural Machine Translation (NMT)*, which has gained significant attention (Bahdanau et al., 2015). In this application also, adversarial attacks have been studied to gain insights into the vulnerabilities of these systems. Particularly, untargeted attacks seek to generate adversarial examples that preserve the semantics in the source language while the output translation by the target model is far from the true translation (Ebrahimi et al., 2018a; Cheng et al., 2019; Michel et al., 2019; Niu et al., 2020; Zou et al., 2020; Sadrizadeh et al., 2023b). On the other hand, targeted attacks against NMT systems aim to mute or push specific target words in the translation (Ebrahimi et al., 2018a; Cheng et al., 2020a; Wallace et al., 2020; Sadrizadeh et al., 2023a). None of these attacks against NMT systems actually consider the class of the output translation as the objective of the adversarial attack. However, in some cases, the user only cares about the class, e.g., sentiment, of the translation rather than the exact translation. Moreover, the class of the translation reflects the whole meaning of the sentence. Nevertheless, simply reducing the translation quality (untargeted attacks) or inserting specific keywords in the translation (targeted attacks), as proposed in previous works, may not sufficiently alter the translation and thus change the overall category of the translation. Moreover, in previous attack frameworks, it is par-

	Input Text (red=adversarial perturbation)	En-to-Fr NMT (blue=English meaning)	Classifier	True Class
Org.	unflinchingly bleak and desperate	Inébranlablement sombre et désespéré	Neg.	Neg.
Adv. ACT <sub>TF</sub>	unflinchingly bleak and <b>upsetting</b>	Légère et bouleversante, sans fin (Light and moving, endless)	Pos.	Pos.
Adv. TF	unflinchingly <b>melancholy</b> and <b>upsetting</b>	Mélancolie et bouleversement inébranlables (Unwaveringly melancholy and turmoil)	Pos.	Neg.

Table 1: Illustration of valid adversarial attacks with two examples against Marian NMT (En-Fr) on a movie review from SST-2 dataset. For both examples, the class predicted by the classifier for the adversarial translation differs from the class of the original translation. The first adversarial example, ACT<sub>TF</sub>, is considered valid since the NMT’s French translation is wrong (if we compare its meaning in blue to the input text), and the classifier’s prediction is correct. However, in the second one, TF, the NMT’s translation is correct, and the classifier’s prediction is wrong. Hence, it is not a desirable attack. The last column, True Class, is the real perceived sentiment of the translated review.

ticularly challenging to evaluate the true impact of the adversarial attack on the performance of the target NMT since the ground-truth translation might change even if the adversarial perturbation is subtle (Zhang et al., 2021; Sadrizadeh et al., 2023b).

In light of the above challenges, this paper introduces ACT (Altering Class of Translation), a novel adversarial attack framework against NMT systems guided by a classification objective. In our attack strategy, the adversary aims to craft an adversarial example in the source language that deceives the *target NMT model*. Specifically, the goal is to make the translation of the adversarial sentence belong to a different class than the original translation. To achieve this, the adversary uses an *arbitrary classifier as an oracle* to predict and alter the class of the output translation by the target NMT model. By targeting the class of the output translation, our attack has a more substantial effect on the translation. The second row of Table 1 shows a successful adversarial example generated in our framework. The original movie review has negative sentiment. The adversary aims to change the input sentence (in an imperceptible manner, i.e., the sentiment remains negative) such that the output translation by the target NMT model has positive sentiment instead of negative as the original translation. However, the second example (last row) is an undesirable case, where the class of the translation predicted by the classifier is changed, but the classifier is fooled by the attack, and its prediction is incorrect. In our framework, the target of the attack is the NMT model, but the attacker considers a system, including the target NMT model and the classifier operating on the output of the NMT model. Hence, the attacker needs to ensure that the attack fools the NMT model (not the classifier) to generate a wrong translation, which then results in a different class predicted by the classifier. Consequently, the

main challenge in this framework is distinguishing the attack’s impact on the target NMT model from its impact on the classifier, which is not attacked by the adversary. To address this challenge, we design enhancements to existing black-box word-replacement attacks, such as TextFooler (Jin et al., 2020) and BAE (Garg and Ramakrishnan, 2020), by integrating the output translation of the target NMT model and the output logits of a classifier into the attack process. This approach ensures that the adversarial translation is far from the original one, in addition to altering the class to which the translation belongs. In practice, attackers typically have limited access to the target NMT model. Therefore, we assume a black-box setting for the attack against the target NMT model. Additionally, we assume black-box access to the oracle classifier, as the adversary may employ an off-the-shelf classifier to guide the attack.

We extensively evaluate the robustness of NMT systems to our proposed attack. To evaluate the attack, we measure the success rate of altering the class of the output translation by using a classifier that was not involved in the attack process. Moreover, we estimate the impact of the attack on the NMT model by the similarity between the translations of the original and adversarial sentences. As a baseline, we consider the untargeted attacks against NMT systems and check if the class of the translation changes after the attack. Our experiments show that, although untargeted attacks can reduce the translation quality, they are notably less successful in changing the category of the translation. In contrast, our proposed attack not only changes the class of the translation but also has more impact on the translation. It shows that our new attack can provide a more comprehensive evaluation of the robustness of NMT systems to adversarial attacks. In summary, our contributions are as follows:

- We introduce ACT, a novel attack framework against NMT models, which is guided by a classifier to change the class of the output translation.
- We propose modifications to existing black-box word-replacement attacks to make them effective for the proposed attack strategy.
- We thoroughly assess the robustness of NMT systems to our new attack framework, which showcases the vulnerabilities of NMT systems by focusing on the class of output translation.

## 2 Related Works

Textual adversarial attacks pose unique challenges compared to their image domain counterparts. The discrete nature of textual data makes it challenging to employ gradient-based optimization methods directly. Moreover, defining the imperceptibility of adversarial perturbations for text is difficult. While a significant portion of research has focused on text classification, some works also target sequence-to-sequence NLP systems, such as NMT.

**Text Classification** In the text classification tasks, some attacks measure imperceptibility based on the number of edits at the character level (Ebrahimi et al., 2018b; Gao et al., 2018; Pruthi et al., 2019). However, most NLP attacks consider semantic similarity as the imperceptibility metric and operate at the word level. Some of these attacks adopt optimization methods to bridge the gap between image and textual domains (Guo et al., 2021; Sadrizadeh et al., 2022; Yuan et al., 2023). However, the majority of methods select specific words in the input sentence and replace them with synonyms (Zang et al., 2020), similar words in the embedding space (Alzantot et al., 2018; Ren et al., 2019; Jin et al., 2020; Maheshwary et al., 2021; Ye et al., 2022), or candidates predicted by a masked language model (Garg and Ramakrishnan, 2020; Li et al., 2020, 2021; Yoo and Qi, 2021).

**Neural Machine Translation** In contrast to text classification settings, where the adversary aims to alter the predicted class, NMT models generate entire sentences as their output and the adversary tries to alter this translation. There are various types of adversarial attacks in the literature depending on the adversary’s objective. Untargeted attacks aim to reduce the translation quality of the target model with respect to the ground-truth translation (Michel et al., 2019; Cheng et al., 2019, 2020b; Sadrizadeh et al., 2023b). Targeted attacks seek to mute or

insert specific words into the translation (Ebrahimi et al., 2018a; Cheng et al., 2020a; Wallace et al., 2020; Sadrizadeh et al., 2023a). (Wallace et al., 2020) introduces a universal attack that causes incorrect translation by target model when a single snippet of text is appended to any input sentence. They also show that NMT models can generate malicious translations from gibberish input.

Specifically, (Belinkov and Bisk, 2018; Ebrahimi et al., 2018a) first explore the vulnerabilities of NMT models to character manipulations. In untargeted attacks, (Cheng et al., 2019) replaces random words in the input sentence with the words suggested by a language model, guided by gradients to reduce translation quality. Moreover, (Michel et al., 2019; Zhang et al., 2021) substitute important words in the input sentence with their neighbouring words in the embedding space. Other approaches utilize optimization to generate adversarial examples (Cheng et al., 2020a; Sadrizadeh et al., 2023b). While the first use the NMT embeddings to define similarity, the latter uses the embedding representation of a language model.

However, none of these attacks specifically target the class that the translation belongs to, which can be important in many applications. Furthermore, evaluating the true impact of these attacks is challenging, as the adversarial perturbations may change the ground-truth output and potentially result in an overestimation of the attack performance. To address these limitations, we propose a novel attack framework against NMT models guided by a classifier to generate an adversarial example whose translation by the NMT model belong to a different class than the class of the original translation. This approach can have more impact on the output translation by altering its class. We should note that, in a parallel work, (Raina and Gales, 2023) recently published a paper proposing an attack against NMT models to change the perception of translation. In contrast to this work, we consider modifying the class of output translation, and not just sentiment. Our proposed framework can be used with different classifiers based on the adversary’s objective. Moreover, we try to distinguish the attack’s impact on the target NMT model from its impact on the classifier used in the attack, which is the main challenge in this framework. Finally, we extensively evaluate the robustness of NMT models to our attack framework. We discuss these differences in more detail in Appendix C.

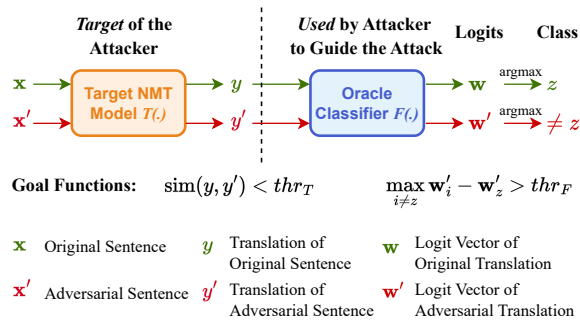


Figure 1: Block diagram of ACT, a new attack framework against NMT models.

### 3 ACT Attack Framework against NMT

In this section, we introduce ACT, our new attack framework against NMT models. Then, we present our method to craft such adversarial examples.

#### 3.1 Attack Definition

The block diagram of our attack framework is presented in Figure 1. Let  $\mathcal{X}$  denote the source language space and  $\mathcal{Y}$  denote the target language space. The NMT model  $T : \mathcal{X} \rightarrow \mathcal{Y}$  maximizes the probability of generating the true translation, automatically translating the input sentence  $y = T(x)$ , where  $x \in \mathcal{X}$  is the input sentence in the source language, and  $y \in \mathcal{Y}$  is the output translation by the NMT model  $T$ .

In our proposed attack, the adversary seeks an adversarial example  $x'$  in the source language that misleads the *target NMT model*. In particular, the adversary aims that the adversarial translation  $y' = T(x')$  is from a different class than the original translation. To this end, the adversary uses an *arbitrary classifier as an oracle*  $F : \mathcal{Y} \rightarrow \mathcal{Z}$  to determine the class of the output translation by the target NMT model.<sup>1</sup> Based on the attack objective, the adversary can use a classifier suitable for any task, such as sentiment classification. The classifier classifies the input translation into  $z = F(y)$ , where  $z$  is the class of the translation  $y$ . Hence, the adversarial example is crafted such that  $F(y') \neq z$ . However, we constraint that the adversarial example  $x'$  must remain a natural sentence and be semantically similar to the original sentence  $x$ .

In practice, the adversary has limited access to the target model. Therefore, we consider a black-box scenario where the adversary cannot access the model parameters, architecture, or training data of

<sup>1</sup>The effect of the choice of classifier on the performance of ACT is studied in Appendix B.5.

the target NMT model  $T$ . Moreover, the adversary may use an off-the-shelf classifier for the attack and hence, has black-box access to the classifier  $F$ . The attacker can only query the NMT model  $T$  with a sentence in the source language and get the translated sentence. Then, they can use the classifier  $F$  to determine the class and the corresponding logits for the translated sentence.

It is worth mentioning that we can consider other scenarios within our framework, e.g., when the target of the attack is the entire system (NMT + classifier), or just the classifier. We explore these scenarios in the experiments.

#### 3.2 Methodology

There are substantial textual adversarial attacks in the literature that efficiently search the discrete space of tokens to craft meaning-preserving natural adversarial sentences (Jin et al., 2020; Garg and Ramakrishnan, 2020; Ye et al., 2022). We can build our attack upon these attacks based on the objective of our attack framework. To this end, we use TextAttack, which provides a unified framework for numerous textual adversarial attacks (Morris et al., 2020) and facilitates the incorporation of our enhancements into the existing attack methodologies. In this framework, there are four components in black-box attacks based on word-replacement.

**Constraints:** In order to generate semantic-preserving and grammatically correct adversarial sentences, each attack defines a set of constraints by using a grammar checker, embedding space distance, or perplexity score.

**Transformation:** To find a set of candidates to replace the selected words, various transformations are proposed, e.g., predictions by masked language models or neighbors in the embedding space.

**Search method:** Each attack employs a search method to iteratively query the target model and find an adversarial example that satisfies the constraints, e.g., greedy search or genetic algorithm.

**Goal function:** This module specifies the stopping criteria and determines if an attack is successful.

In the attacks modeled by this framework, we can craft an adversarial example that satisfies the adversary’s goal function and adheres to certain linguistic constraints. We find these perturbations by replacing some of the words in the input sentence based on a search method.

In our proposed attack framework, the adversarial sentence should maintain the semantics and the class label in the source language. Meanwhile, the translation in the target language by the target NMT model should have a different class label. Therefore, we use the *constraints* and *transformations* in the literature to generate adversarial examples that are similar to the original sentence and are grammatically correct. However, based on our attack objective, we propose a new goal function and some alterations to the search method. In our framework, the adversary uses an oracle classifier to manipulate the NMT model into crafting an adversarial translation from a different class. Hence, the target of the attack is the NMT model, and we need to distinguish between the impact of the attack on the NMT model  $T$  and on the classifier  $F$ . An undesirable example, in which the classifier is impacted by the attack, is presented in Table 1. To address this challenge and make the attack mainly effective on the NMT model rather than the classifier, we propose the two following goal functions as the stopping criteria of the attack:

**1) Translation** In order to highlight the effect of the attack on the NMT model, we include the translation output of the NMT model in the goal function. In other words, we consider an adversarial attack to be successful if the similarity between the original translation  $\mathbf{y}$  and the adversarial translation  $\mathbf{y}'$  is less than a threshold  $thr_T$ :

$$\text{sim}(\mathbf{y}, \mathbf{y}') < thr_T. \quad (1)$$

This goal function allows us to mislead the NMT model  $T$  to generate a translation for the adversarial example that is far from the original translation. We use the BLEU score to evaluate the similarity between two translations since it is fast, common in benchmarks, and has been used in previous works (Cheng et al., 2019; Wallace et al., 2020; Zhang et al., 2021).<sup>2</sup> We study the effect of this similarity metric when we use BLEURT-20 (Sellam et al., 2020) instead of BLEU score in Appendix B.4.

**2) Classification** Instead of checking whether the classifier’s output is different from the ground-truth label, we use the raw output of the classifier before the softmax, known as *logit*. In order to ensure that the output of the classifier is different from the original class with high confidence, we consider an ad-

<sup>2</sup>We use case-sensitive SacreBLEU on detokenized sentences.

versarial attack to be successful if the difference between the logits of the most probable class and the ground-truth class is larger than a threshold  $thr_F$ :

$$\max_{i \neq z} \mathbf{w}'_i - \mathbf{w}'_z > thr_F, \quad (2)$$

where  $\mathbf{w}' = W(\mathbf{y}')$  are the logits, and  $z$  is the ground-truth class. Therefore, we consider an adversarial example to be successful if both goal functions (1) and (2) are satisfied.

On another note, most of the existing attacks use a score function during the search to estimate the importance of the tokens in the sentence. They select the important words in the sentence to limit the search space and the number of alterations made by the attack (Ren et al., 2019; Li et al., 2020; Jin et al., 2020; Garg and Ramakrishnan, 2020). They define this score function as the logit of the ground-truth class, i.e.,  $\mathbf{w}'_z$ . The token importance is the decrease in this score when removing a token from the sentence. We propose a new score function to account for the effect of the adversarial example on the NMT model as follows:

$$S(\mathbf{x}') = \mathbf{w}'_z + \alpha \text{sim}(\mathbf{y}, \mathbf{y}'). \quad (3)$$

The proposed second term in this score function makes the importance of the token dependent on the decision of the *classifier* and that of the *target NMT model*.

## 4 Experiments

In this section, we discuss our experimental setup, and then we conduct comprehensive experiments to evaluate the robustness of various NMT models and tasks in the face of our proposed attack strategy.<sup>3</sup>

### 4.1 Experimental Setup

We evaluate the robustness of transformer-based NMT models to our attack. Specifically, we target the HuggingFace implementation of Marian NMT models (Junczys-Dowmunt et al., 2018) and mBART50 multilingual NMT model (Tang et al., 2020) to validate the effectiveness of our attack across diverse architectures. Moreover, we conduct experiments on the English-French (En-Fr) and English-German (En-De) translation tasks.

In our proposed attack strategy, the adversary aims to alter the class of the NMT model’s output

<sup>3</sup>Our source code is available at <https://github.com/sssadrizadeh/ACT>.

Task	Method	Marian NMT (En-Fr)					Marian NMT (En-De)						
		ASR↑	BLEU↓	chrF↓	Sim.↑	Perp.↓	WER↓	ASR↑	BLEU↓	chrF↓	Sim.↑	Perp.↓	WER↓
SST-2	kNN	8.95	50.95	70.53	0.78	6.06	30.46	12.08	42.81	69.18	0.82	5.19	28.02
	Seq2Sick	2.14	41.83	59.52	0.60	4.54	35.56	2.71	46.39	60.01	0.71	3.54	22.48
	TransFool	10.21	44.90	68.16	<b>0.85</b>	2.64	20.86	12.35	39.12	62.85	0.82	3.04	21.82
	ACT <sub>TF</sub>	<b>40.23</b>	<b>29.66</b>	<b>56.11</b>	0.84	2.00	21.19	<b>40.84</b>	<b>25.11</b>	<b>53.84</b>	0.84	2.23	22.13
	ACT <sub>BAE</sub>	38.08	45.72	65.62	<b>0.85</b>	<b>0.59</b>	<b>15.16</b>	31.61	40.95	64.80	<b>0.86</b>	<b>0.61</b>	<b>14.95</b>
MR	kNN	8.78	50.01	70.93	0.82	6.88	31.35	13.00	44.22	70.77	0.83	4.86	27.10
	Seq2Sick	3.80	30.82	53.44	0.71	5.86	38.78	3.00	26.40	56.50	0.73	2.48	21.25
	TransFool	10.56	43.25	66.16	0.84	3.04	24.81	14.63	37.47	66.53	0.85	2.76	22.54
	ACT <sub>TF</sub>	<b>36.37</b>	<b>20.15</b>	<b>48.46</b>	0.82	3.03	26.20	<b>26.50</b>	<b>9.63</b>	<b>42.62</b>	0.80	3.71	29.67
	ACT <sub>BAE</sub>	31.64	46.08	65.26	<b>0.86</b>	<b>0.62</b>	<b>15.19</b>	21.00	35.51	62.68	<b>0.88</b>	<b>0.63</b>	<b>14.74</b>
AG's News	kNN	2.02	40.81	69.50	<b>0.95</b>	2.13	<b>10.67</b>	2.65	65.49	81.70	<b>0.95</b>	1.58	<b>9.67</b>
	TransFool	3.30	48.61	68.09	0.89	3.07	17.89	3.43	44.74	66.27	0.90	2.65	18.35
	ACT <sub>TF</sub>	<b>22.84</b>	<b>25.62</b>	<b>49.19</b>	0.85	6.02	27.50	<b>18.36</b>	<b>27.03</b>	<b>52.08</b>	0.85	5.87	27.28
	ACT <sub>BAE</sub>	7.58	36.10	62.11	0.94	<b>1.31</b>	12.91	7.19	42.24	66.59	0.93	<b>1.49</b>	14.37

Table 2: Evaluation results of the adversarial attacks against two translation tasks and three different datasets.

translation. To achieve a comprehensive evaluation, we require ground-truth class information for the sentences. Therefore, instead of translation datasets, we consider text classification datasets, including SST-2 (Socher et al., 2013), MR (Pang and Lee, 2005), and AG’s News (Zhang et al., 2015). SST-2 and MR are sentiment classification datasets, while AG’s News is a topic classification dataset. We perform the attack on the test set of these datasets.<sup>4</sup>

We translate the training set of these datasets using the target NMT model and fine-tune two separate classifiers.<sup>5</sup> We utilize one of the classifiers during the attack process, while the other one is reserved for evaluation. This approach ensures fair comparisons with the baselines and accounts for the possibility of the adversarial attack fooling the classifier used in the attack. More details about the datasets and models are reported in Appendix A.

To perform the proposed attack, we use TextAttack implementation of TextFooler (TF) (Jin et al., 2020) and BAE (Garg and Ramakrishnan, 2020). We change the goal function and the score function, as explained in the last section. As for the parameters, we set  $thr_T = 0.4$ ,  $thr_F = 2$  and  $\alpha = 3$  based on the ablation study available later.

<sup>4</sup>For the AG’s News and MR datasets, we attack the first 1000 sentences from the test set.

<sup>5</sup>For the En-Fr task, we fine-tuned two models: <https://huggingface.co/asi/gpt-fr-cased-small> with GPT-2 architecture and <https://huggingface.co/tblard/tf-allocine> with BERT architecture. As for the En-De task, we fine-tuned two additional models: <https://huggingface.co/dbmdz/german-gpt2> with GPT-2 architecture and <https://huggingface.co/oliverguhr/german-sentiment-bert> with BERT architecture.

As a baseline, we compare with the untargeted attacks against NMT systems and examine whether the class of the translation changes after the attack. Specifically, we compare with the *kNN* attack from (Michel et al., 2019), which is a white-box untargeted attack against NMT models that substitutes some words with their neighbors in the embedding space. Additionally, we adapt the targeted attack *Seq2Sick* (Cheng et al., 2020a), which is based on optimization in the NMT embedding space, to the untargeted setting. Lastly, we compare with *TransFool*, an untargeted attack against NMT models that is also based on optimization but uses language model embeddings to preserve the semantics.

For evaluation, we report several performance metrics. We measure *Attack Success Rate (ASR)* of the adversarial examples by testing them on a classifier that was not involved in the attack process. We also measure the similarity between the translations of the original and adversarial sentences using BLEU score and chrF (Popović, 2015). A lower similarity indicates a greater deviation between the translations of the adversarial and original sentences, which allows us to estimate the impact of the attack on the target NMT model. Furthermore, we use Universal Sentence Encoder (USE) (Yang et al., 2020) to approximate the semantic Similarity (Sim.) between the original and adversarial sentences. To measure the naturality of the adversarial sentences, we calculate the relative increase in the Perplexity score (Perp.) of GPT-2 (large) between the adversarial and original sentences. Finally, we report Word Error Rate (WER), i.e., the percentage of words that are modified by an adversarial attack.

Task	Sentence	Text
MR	Org.	a solidly entertaining little film .
	Org. Trans. (Positive)	un petit film très divertissant.
MR	Adv. ACT <sub>TF</sub>	a solidly <b>goofy</b> little film .
	Adv. Trans. (Negative)	un petit film complètement dégueulasse. ( <i>dégueulasse means 'nasty', which is negative.</i> )
SST-2	Org.	the primitive force of this film seems to bubble up from the vast collective memory of the combatants .
	Org. Trans. (Positive)	la force primitive de ce film semble jaillir de la vaste mémoire collective des combattants.
	Adv. ACT <sub>BAE</sub>	the primitive <b>tone</b> of this film seems to bubble up from the vast collective <b>memories</b> of the combatants .
	Adv. Trans. (Negative)	Le ton primitif de ce film semble s'estomper des vastes souvenirs collectifs des combattants. ( <i>s'estomper means fade which conveys negative meaning</i> )

Table 3: Adversarial examples against Marian NMT (En-Fr).

## 4.2 Results

Now we evaluate the robustness of various NMT models to our new attack strategy. Table 2 shows the performance of different attacks against Marian NMT models for (En-Fr) and (En-De) tasks. Additionally, the performance against mBART50 NMT model for (En-Fr) task is reported in Appendix B.1. In these tables, ACT<sub>TF</sub> and ACT<sub>BAE</sub> denote the modified version of the corresponding attacks with our proposed changes. These results demonstrate that the existing untargeted adversarial attacks against NMT models can generate adversarial examples with translations dissimilar to the original translation. However, they have a low success rate in changing the class of the translation. Seq2Sick has the lowest success rate, while TransFool is the most successful one. On the other hand, ACT<sub>TF</sub> and ACT<sub>BAE</sub> have much higher success rates than the baselines in all cases, and they are able to change the class of the NMT models' output translations. Interestingly, these two methods, especially ACT<sub>TF</sub>, can generate adversarial examples whose translations are further away from the original translation than those generated by the baselines, i.e., lower BLEU score and chrF. While ACT<sub>TF</sub> has a higher success rate and causes more damage to the translation, it generates adversarial examples with higher perplexity scores and lower similarity than the ones generated by ACT<sub>BAE</sub>. It is worth noting that all attacks are much less successful in AG's News. It seems that when the number of classes is larger, the attack becomes more challenging. We should note that Seq2Sick is not successful against AG's News and not reported in Tables 2.

Regarding the run-time, for the Marian NMT (En-Fr) model and SST-2 dataset, on a system

equipped with two NVIDIA A100 GPUs, it takes 26.54 and 15.75 seconds to generate adversarial examples by ACT<sub>TF</sub> and ACT<sub>BAE</sub>, respectively. If we do not use the proposed modifications, the run-time of ACT<sub>TF</sub> would be 17.17 seconds. Table 3 shows some adversarial examples generated by ACT<sub>TF</sub> and ACT<sub>BAE</sub>. These samples show that while the proposed attack maintains the semantic similarity in the source language, they are able to force the NMT model to generate a translation from a different class in the target language. More adversarial examples can be found in Appendix B.9.

All in all, previous untargeted adversarial attacks are not much successful in deceiving the NMT model to generate translations from a different class than the original translations. However, this type of attack can be more harmful to the users since the overall meaning of the translation is changed. The proposed attacks, i.e., ACT<sub>TF</sub> and ACT<sub>BAE</sub>, are more successful in changing the class of the adversarial translation. Moreover, compared to baselines, the adversarial translations are further away from the original translation.

We should note that, in our framework, we are generating adversarial examples that are robust to the translation. It has been shown in (Bhandari and Chen, 2023) that most of the adversarial attacks against text classifiers are not robust to translation. This means that most of the attacks in the source language do not transfer to the translation model. Therefore, even if the attacker changes the class in the source language, it is possible that the output translation is still from the correct class.

## 4.3 Analysis

In this section, we analyze the significance of the proposed goal functions and discuss two other sce-

Task	kNN		Seq2Sick		TransFool	
	BLEU↓	chrF↓	BLEU↓	chrF↓	BLEU↓	chrF↓
SST-2	47.86	70.68	93.39	95.95	43.73	65.79
	34.57	60.14	61.20	70.93	37.18	58.00
MR	45.69	69.11	93.72	95.93	43.20	65.67
	32.13	57.18	43.48	62.01	39.17	61.92

Table 4: Translation performance of baseline attacks against mBART50 (En-Fr). First rows show the results for all of the adversarial examples, while the second rows correspond to the successful examples that *change the class*.

Goal Func.	ASR↑	BLEU↓	chrF↓	Sim.↑	Perp.↓
Label change	28.12	58.16	72.90	0.89	0.91
Logit dif.	37.70	52.19	68.35	0.87	1.26
Trans. sim.	36.19	31.90	58.24	0.85	1.83
ACT <sub>TF</sub>	40.23	29.66	56.11	0.84	2.00

Table 5: Ablation study on the proposed goal functions for ACT<sub>TF</sub> against Marian NMT (En-Fr) on SST-2 dataset.

narios of the classification-guided attack strategy. We also study the transferability of the proposed attack to other NMT systems, the extension to the targeted settings, and the effect of different parameters on our attack in Appendices B.7, B.8 and B.2, respectively. Finally, in Appendix B.3 we show that by using an ensemble of classifiers, we can improve the performance.

**Impact of Classification-Guided Strategy on Translation** First, we show that by changing the class of the output translation, the adversary can have more impact on the NMT model. Table 4 reports the similarity between the original and adversarial translation, as an estimate of the effect of the adversarial attack on the NMT model, for the baselines when for all the adversarial examples (first row) compared to the successful ones that change the class of the output translation. Across various methods and tasks, we can see that when adversarial examples change the translation’s class, their translations are less similar to the original translations. This difference in similarity arises because these examples often change the overall meaning of the translation.

**Goal Functions** Our framework consists of a classifier that acts on the output translation by the target NMT model. To ensure that the attack’s influence on the target NMT model outweighs its effect on the classifier, we proposed two goal functions based on the output translation of the NMT model

Target	ASR↑	BLEU	chrF	Sim.↑	Perp.↓
NMT + Classifier	95.61	59.08	73.61	0.89	0.98
Classifier	57.55	71.92	83.03	0.91	0.72

Table 6: Performance of ACT<sub>TF</sub> against Marian NMT (En-Fr) on SST-2 dataset when the target is the entire system (NMT + Classifier) or just the classifier.

and the output logits of the classifier. Table 5 shows the effect of these two goal functions on the attack performance. The first row shows the results when the goal function of the attack is only to change the label of the translation. The second and third rows present the effect of using equations (2) and (1), respectively. We can see that both of the proposed goal functions increase the success rate, and they also help to reduce the similarity between the original and adversarial translations. Since the goal function becomes more difficult to achieve, there is a decrease in semantic similarity and an increase in the perplexity score.

**Other Scenarios** So far, we have assumed that the attacker’s objective is to mislead the NMT to generate a translation from a class different from the original translation. However, our proposed classification-guided strategy can be adapted to other scenarios as well.

First, we can consider a system including a classifier that operates on the output of an NMT model.<sup>6</sup> In this context, the goal would be to attack the *entire system* instead of just the NMT model. This scenario is much easier than the previous one since the adversary can access the entire system. Also, the adversary’s target is the performance of the entire system (NMT *and* the classifier), unlike the original scenario, where the target is the NMT model. The performance of the attack in this scenario is presented in the first row of Table 6. As expected, the success rate is much higher than that of the previous scenario. Moreover, the adversarial and original translations are more similar meaning that the NMT model is less affected by the attack.

Secondly, we can assume that the adversary’s goal is to fool only the classifier.<sup>7</sup> Therefore, we need to craft an adversarial example whose translation is similar to the original one, the complete

<sup>6</sup>An example might be when we are interested in the class prediction of foreign language sentences, but a classifier is available in another language. Hence, we use an NMT model to translate the sentences before feeding them to the classifier.

<sup>7</sup>This scenario may not have a practical use case, but it shows another aspect of our classification-guided attack.



opposite goal of our original scenario, but the classifier predicts a wrong class for the adversarial translation. In this scenario, we can change goal function (1) such that the similarity is more than a threshold. Moreover, we can use equation (3) with negative coefficient  $\alpha$  so that the word with a higher impact on the translation has less importance. The performance with the parameters  $thr_T = 0.8$  and  $\alpha = -7$  are reported in the second row of Table 6. We can see that the similarity between original and adversarial translations is higher than those of the previous scenarios, showing that the attack is effectively targeting the classifier.

## 5 Human Evaluation

We conduct a human evaluation campaign for the successful adversarial examples generated by ACT<sub>TF</sub> against Marian NMT (En-Fr). We randomly choose 80 successful adversarial examples on the SST-2 dataset. We split these sentences into two surveys and recruit three volunteer annotators for each survey. Since the adversarial examples are in French, we ensure that the annotators are native (mother tongue) or highly proficient in French.

Since the naturalness and sentiment accuracy of TextFooler and BAE are already evaluated by human in their respective papers, we do not consider these two aspects of the adversarial examples in our evaluation. Instead, we study the sentiment of the adversarial translations in the target language (French). By showing the adversarial translations, we ask the annotators to choose a sentiment label from “Positive” and “Negative”. We take the majority class as the predicted label for each sentence.

In our attack framework, the target of the attack is the NMT model, and the adversary uses a classifier to change the class of the translation. Therefore, we want the proposed attack mainly affect the NMT model rather than the classifier. Accordingly, in this study, we evaluate how much the classifier is influenced by the attack. The overall agreement between the ground-truth labels (in the dataset) and the labels predicted by the annotators is 71.3%. While it’s true that not all adversarial translations are accurately classified by the classifier based on the annotators’ labels, the majority of adversarial translations have the same sentiment as predicted by the classifier. This implies that the attack is mainly targeting the NMT model rather than fooling the classifier. Moreover, we calculate the similarity between original and adversarial translations

in terms of chrF for the sentences selected for human evaluation. For sentences that the classifier’s predictions align with the annotators’ labels, the translation similarity is 49.46. In contrast, for sentences that the classifier’s predictions diverge from the annotators’ labels, the translation similarity is 52.05. This difference highlights that the attack is mainly affecting the NMT model rather than the classifier for the adversarial sentences whose translation deviates more from the original translation.

## 6 Conclusion

In this paper, we presented ACT, a novel adversarial attack framework against NMT models that is guided by a classifier. In our framework, the adversary aims to alter the class of the output translation in the target language while preserving semantics in the source language. By targeting the class of the output translation, we outlined a new aspect of vulnerabilities of NMT models. We proposed enhancements to existing black-box word-replacement-based attacks to evaluate the robustness of NMT models to our attack strategy. Extensive experiments and comparisons with existing untargeted attacks against NMT models showed that our attack is highly successful in changing the class of the adversarial translation. It also has more impact on the similarity of the original and adversarial translations, which highlights the potential impact of our attack strategy on the overall meaning of the NMT output translations.

## 7 Limitations

In our framework, the target of the attack is the NMT model, and the attacker uses a classifier to change the class of the translation. Therefore, the adversary is defining a system including the classifier operating on the output of the NMT model. Although we have proposed a goal function to make the attack mainly effective on the NMT model rather than the classifier, there is still a chance that the classifier is affected by the attack (instead of the NMT model). To consider this challenge in our evaluations, we have reported the similarity between the adversarial and original translations to measure the effect of the attack on the NMT model. Moreover, we have calculated the success rate of altering the class of translation by using a different classifier than the one used in the attack process. Such an evaluation provides fair comparisons with the baselines and accounts for the possibility of

the adversarial attack fooling the classifier used in the attack. On another note, our proposed goal functions need the output logits of the classifier. However, some of the recent textual adversarial attacks only need the hard labels (Ye et al., 2022; Yu et al., 2022). Modifying these works based on our attack framework can be explored in the future. Finally, we have considered sentiment and news classification in our experiments. It is worth considering other classification tasks, e.g., hate speech, in our framework and evaluate the robustness of NMT models in other areas.

## 8 Ethic statement

We introduced ACT, a new attack framework against NMT models, to study the vulnerabilities of NMT models from another aspect than tradition frameworks with the hope to pave the way for building robust NMT models. Although there is a potential for malicious actors to misuse our attack, we want to emphasize that we strongly discourage the use of our method for targeting real-life NMT systems with harmful intent.

## Acknowledgements

This work has been partially supported by armasuisse Science and Technology project MULAN.

## 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.
- Dzmitry Bahdanau, Kyung Hyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *International Conference on Learning Representations*.
- Yonatan Belinkov and Yonatan Bisk. 2018. Synthetic and natural noise both break neural machine translation. In *International Conference on Learning Representations*.
- Neel Bhandari and Pin-Yu Chen. 2023. Lost in translation: Generating adversarial examples robust to round-trip translation. In *ICASSP 2023 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 1–5.
- Ondřej Bojar, Christian Buck, Christian Federmann, Barry Haddow, Philipp Koehn, Johannes Leveling, Christof Monz, Pavel Pecina, Matt Post, Herve Saint-Amand, et al. 2014. Findings of the 2014 workshop on statistical machine translation. In *Proceedings of the Ninth Workshop on Statistical Machine Translation*, pages 12–58.
- Minhao Cheng, Jinfeng Yi, Pin-Yu Chen, Huan Zhang, and Cho-Jui Hsieh. 2020a. Seq2Sick: Evaluating the robustness of sequence-to-sequence models with adversarial examples. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 3601–3608.
- Yong Cheng, Lu Jiang, and Wolfgang Macherey. 2019. Robust neural machine translation with doubly adversarial inputs. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4324–4333.
- Yong Cheng, Lu Jiang, Wolfgang Macherey, and Jacob Eisenstein. 2020b. AdvAug: Robust adversarial augmentation for neural machine translation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5961–5970.
- Javid Ebrahimi, Daniel Lowd, and Dejing Dou. 2018a. On adversarial examples for character-level neural machine translation. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 653–663.
- Javid Ebrahimi, Anyi Rao, Daniel Lowd, and Dejing Dou. 2018b. HotFlip: White-box adversarial examples for text classification. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*, pages 31–36.
- 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*, pages 50–56.
- Siddhant Garg and Goutham Ramakrishnan. 2020. BAE: BERT-based adversarial examples for text classification. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*, pages 6174–6181.
- Chuan Guo, Alexandre Sablayrolles, Hervé Jégou, and Douwe Kiela. 2021. Gradient-based adversarial attacks against text transformers. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 5747–5757.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778.
- Di Jin, Zhijing Jin, Joey Tianyi Zhou, and Peter Szolovits. 2020. Is BERT really robust? a strong baseline for natural language attack on text classification and entailment. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 8018–8025.

- Marcin Junczys-Dowmunt, Roman Grundkiewicz, Tomasz Dwojak, Hieu Hoang, Kenneth Heafield, Tom Neckermann, Frank Seide, Ulrich Germann, Alham Fikri Aji, Nikolay Bogoychev, André F. T. Martins, and Alexandra Birch. 2018. Marian: Fast neural machine translation in C++. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 116–121.
- Jacob Devlin Ming-Wei Chang Kenton and Lee 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*, pages 4171–4186.
- Quentin Lhoest, Albert Villanova del Moral, Yacine Jernite, Abhishek Thakur, Patrick von Platen, Suraj Patil, Julien Chaumond, Mariama Drame, Julien Plu, Lewis Tunstall, Joe Davison, Mario Šaško, Gunjan Chhablani, Bhavitvya Malik, Simon Brandeis, Teven Le Scao, Victor Sanh, Canwen Xu, Nicolas Patry, Angelina McMillan-Major, Philipp Schmid, Sylvain Gugger, Clément Delangue, Théo Matussière, Lysandre Debut, Stas Bekman, Pierric Cistac, Thibault Goehringer, Victor Mustar, François Lagunas, Alexander Rush, and Thomas Wolf. 2021. Datasets: A community library for natural language processing. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 175–184.
- Dianqi Li, Yizhe Zhang, Hao Peng, Liqun Chen, Chris Brockett, Ming-Ting Sun, and William B Dolan. 2021. Contextualized perturbation for textual adversarial attack. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5053–5069.
- 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*, pages 6193–6202.
- Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. 2018. Towards deep learning models resistant to adversarial attacks. In *International Conference on Learning Representations*.
- Rishabh Maheshwary, Saket Maheshwary, and Vikram Pudi. 2021. Generating natural language attacks in a hard label black-box setting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 13525–13533.
- Paul Michel, Xian Li, Graham Neubig, and Juan Pino. 2019. On evaluation of adversarial perturbations for sequence-to-sequence models. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3103–3114.
- Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, and Pascal Frossard. 2016. DeepFool: a simple and accurate method to fool deep neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2574–2582.
- John Morris, Eli Lifland, Jin Yong Yoo, Jake Grigsby, Di Jin, and Yanjun Qi. 2020. TextAttack: A framework for adversarial attacks, data augmentation, and adversarial training in NLP. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 119–126.
- Xing Niu, Prashant Mathur, Georgiana Dinu, and Yaser Al-Onaizan. 2020. Evaluating robustness to input perturbations for neural machine translation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8538–8544.
- Bo Pang and Lillian Lee. 2005. Seeing stars: Exploiting class relationships for sentiment categorization with respect to rating scales. In *Proceedings of the 43rd Meeting of the Association for Computational Linguistics*, pages 115–124.
- Maja Popović. 2015. chrF: Character n-gram F-score for automatic MT evaluation. In *Proceedings of the Tenth Workshop on Statistical Machine Translation*, pages 392–395.
- 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.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1:9.
- Vyas Raina and Mark Gales. 2023. Sentiment perception adversarial attacks on neural machine translation systems. In *The Second Workshop on New Frontiers in Adversarial Machine Learning*.
- 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.
- Sahar Sadrizadeh, AmirHossein Dabiri Aghdam, Ljiljana Dolamic, and Pascal Frossard. 2023a. Targeted adversarial attacks against neural machine translation. In *ICASSP 2023 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 1–5.
- Sahar Sadrizadeh, Ljiljana Dolamic, and Pascal Frossard. 2022. Block-sparse adversarial attack to fool transformer-based text classifiers. In *ICASSP 2022 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 7837–7841.

- Sahar Sadrizadeh, Ljiljana Dolamic, and Pascal Frossard. 2023b. TransFool: An adversarial attack against neural machine translation models. *Transactions on Machine Learning Research*.
- Thibault Sellam, Amy Pu, Hyung Won Chung, Sebastian Gehrmann, Qijun Tan, Markus Freitag, Dipanjan Das, and Ankur Parikh. 2020. Learning to evaluate translation beyond English: BLEURT submissions to the WMT Metrics 2020 shared task. In *Proceedings of the Fifth Conference on Machine Translation*, pages 921–927.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D Manning, Andrew Y 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.
- Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. 2014. Intriguing properties of neural networks. In *2nd International Conference on Learning Representations*.
- Yuqing Tang, Chau Tran, Xian Li, Peng-Jen Chen, Naman Goyal, Vishrav Chaudhary, Jiatao Gu, and Angela Fan. 2020. Multilingual translation with extensible multilingual pretraining and finetuning. *arXiv preprint arXiv:2008.00401*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*, volume 30, pages 5998–6008.
- Eric Wallace, Mitchell Stern, and Dawn Song. 2020. Imitation attacks and defenses for black-box machine translation systems. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*, pages 5531–5546.
- Boxin Wang, Chejian Xu, Xiangyu Liu, Yu Cheng, and Bo Li. 2022. SemAttack: Natural textual attacks via different semantic spaces. In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 176–205.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45.
- Yinfei Yang, Daniel Cer, Amin Ahmad, Mandy Guo, Jax Law, Noah Constant, Gustavo Hernandez Abrego, Steve Yuan, Chris Tar, Yun-Hsuan Sung, et al. 2020. Multilingual universal sentence encoder for semantic retrieval. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 87–94.
- Muchao Ye, Chenglin Miao, Ting Wang, and Fenglong Ma. 2022. TextHoaxer: Budgeted hard-label adversarial attacks on text. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 3877–3884.
- Jin Yong Yoo and Yanjun Qi. 2021. Towards improving adversarial training of NLP models. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 945–956.
- Zhen Yu, Xiaosen Wang, Wanxiang Che, and Kun He. 2022. TextHacker: Learning based hybrid local search algorithm for text hard-label adversarial attack. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 622–637.
- Lifan Yuan, YiChi Zhang, Yangyi Chen, and Wei Wei. 2023. Bridge the gap between CV and NLP! a gradient-based textual adversarial attack framework. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 7132–7146.
- 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.
- Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. Character-level convolutional networks for text classification. In *Advances in Neural Information Processing Systems*, volume 28.
- Xinze Zhang, Junzhe Zhang, Zhenhua Chen, and Kun He. 2021. Crafting adversarial examples for neural machine translation. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing*, pages 1967–1977.
- Andy Zou, Zifan Wang, J Zico Kolter, and Matt Fredrikson. 2023. Universal and transferable adversarial attacks on aligned language models. *arXiv preprint arXiv:2307.15043*.
- Wei Zou, Shujian Huang, Jun Xie, Xinyu Dai, and Jiajun Chen. 2020. A reinforced generation of adversarial examples for neural machine translation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3486–3497.

## Appendix

In this Appendix, we first provide more details about the models and datasets used in the experiments. Afterwards, we present more experimental results of the attack with BLEURT as the similarity metric for translation, studying the choice of the classifier on the performance, the transferability analysis of the attack, extension to the targeted settings, and more samples of the crafted adversarial examples. Finally, we provide a comparison with the recent parallel work of (Raina and Gales, 2023) and discuss the potential limitations of our work.

### A Models and Datasets

In this Section, we provide information about the datasets and models used in our experiments. It is worth noting that we used HuggingFace datasets (Wolf et al., 2020) and transformers (Lhoest et al., 2021) libraries.

#### A.1 Target NMT Models

We evaluate the robustness of the HuggingFace implementation of Marian NMT models (Junczys-Dowmunt et al., 2018) and mBART50 multilingual NMT model (Tang et al., 2020) for En-Fr and En-De translation tasks. As a benchmark of the performance of these models, we report their translation quality on WMT14 (Bojar et al., 2014) in table 7.

#### A.2 Datasets

Since we require ground-truth class information for the sentences in our evaluation, instead of translation datasets, we consider text classification datasets. We use SST-2 (Socher et al., 2013), MR (Pang and Lee, 2005), and AG’s News (Zhang et al., 2015) in our experiments. SST-2 and MR are Sentence-level sentiment classification datasets on positive and negative movie reviews. On the other hand, AG’s News is a Sentence-level topic classification dataset with regard to four news topics: World, Sports, Business, and Science/Technology. Some statistics of these datasets are reported in Table 9. We use the training set of these datasets to train the classifiers used for the attack and evaluation. Moreover, we use the test set to evaluate the robustness of the target NMT models, except for the SST-2, for which we used the validation set.

#### A.3 Classifiers

We translate the training set of these datasets using the target NMT model and fine-tune two separate

Dataset	Marian NMT		mBART50	
	BLEU	chrF	BLEU	chrF
En-Fr WMT14	39.88	64.94	36.17	62.66
En-De WMT14	27.72	58.50	25.66	57.02

Table 7: Translation performance of the target NMT models on WMT14 dataset.

Task	Method	mBART50 (En-Fr)				
		ASR $\uparrow$	BLEU $\downarrow$	chrF $\downarrow$	Sim. $\uparrow$	Perp. $\downarrow$
SST-2	kNN	9.22	34.57	60.14	0.77	6.64
	Seq2Sick	1.48	61.20	70.93	0.74	0.72
	TransFool	9.48	37.18	58.00	0.81	1.76
	ACT <sub>TF</sub>	39.35	<b>25.37</b>	<b>49.46</b>	<b>0.86</b>	1.56
	ACT <sub>BAE</sub>	<b>41.04</b>	38.00	57.84	<b>0.86</b>	<b>0.61</b>
MR	kNN	11.98	32.13	57.18	0.84	5.87
	Seq2Sick	1.96	43.48	62.01	0.75	2.25
	TransFool	12.71	39.17	61.92	0.83	2.26
	ACT <sub>TF</sub>	<b>40.71</b>	<b>12.13</b>	<b>36.88</b>	0.82	2.56
	ACT <sub>BAE</sub>	32.15	30.84	53.04	<b>0.86</b>	<b>0.66</b>
AG’s News	kNN	2.99	49.17	68.53	<b>0.96</b>	1.37
	TransFool	5.34	47.96	63.95	0.88	2.56
	ACT <sub>TF</sub>	<b>27.85</b>	<b>23.02</b>	<b>43.59</b>	0.88	3.86
	ACT <sub>BAE</sub>	9.07	32.11	52.41	0.95	<b>1.15</b>

Table 8: Evaluation results of the adversarial attacks against mBART50 model (En-Fr).

classifiers, employing GPT-2 (Radford et al., 2019) and BERT (Kenton and Toutanova, 2019). The accuracy of these models is reported in Table 9.

## B Additional Results

In this Section, we present more results of the proposed attack.

### B.1 Attack Performance against mBART50

To validate the effectiveness of our attack across diverse architectures, we also attack mBART50 NMT model. The attack performance is presented in Table 8. These results show the same trend as that of the attack against Marian NMT model, which proves the effectiveness of our attack framework against different NMT models.

### B.2 Effect of Parameters

Our attack has three parameters: the coefficient  $\alpha$  in the score function, which controls the importance of translation in the word ranking; the threshold  $thr_T$  in the translation goal function; and the

Dataset	#Classes	#Train	#Test	Avg Length	Marian NMT (En-Fr)		Marian NMT (En-De)		mBART50 (En-Fr)	
					GPT-2	BERT	GPT-2	BERT	GPT-2	BERT
SST-2	2	67.3K	0.9K	17	88.88	90.94	86.01	84.40	86.01	88.30
MR	2	8.5K	1K	20	81.33	84.43	79.55	80.11	78.61	82.08
AG	4	120K	7.6K	43	93.47	93.53	93.86	92.78	93.46	93.55

Table 9: Some statistics of the evaluation datasets, and the accuracy of the classifiers on the test sets.




Figure 2: Effect of different parameters on  $ACT_{TF}$  when attacking Marian NMT (En-Fr) on SST-2 dataset.

Method	ASR $\uparrow$	BLEU $\downarrow$	chrF $\downarrow$	Sim. $\uparrow$	Perp. $\downarrow$
1 classifier	40.23	29.66	56.11	0.84	2.00
2 classifiers	46.03	27.12	54.10	0.83	2.21

Table 10: Performance of  $ACT_{TF}$  with 2 classifiers against Marian NMT (En-Fr) on SST-2 dataset.

threshold  $thr_F$  in the classification goal function. Figure 2 demonstrates the effect of these parameters on the performance of  $ACT_{TF}$  in terms of success rate, BLEU score, and semantic similarity. By increasing the coefficient  $\alpha$ , we are assigning more importance to the words that affect the translation, and hence, the BLEU score between the original and adversarial translations decreases. Moreover, by decreasing the threshold for the similarity between the original and adversarial translations  $thr_T$  or by increasing the threshold for the logit difference of the classifier  $thr_F$ , the attack generates adversarial examples that are more successful in changing the class of the adversarial translation, and they also have more impact on the translation.

### B.3 Ensemble of Classifiers

In order to ensure that the classifier accurately predicts the class and that the attack targets the NMT model, an ensemble of classifiers can be used to find the class of the translation. This approach increases the reliability of the prediction made by the classifier. Table 10 shows the attack performance when we use two classifiers in the attack process. These results show that we can increase the success rate and impact the translation more by using an ensemble of classifiers.

Method	ASR $\uparrow$	BLEU $\downarrow$	chrF $\downarrow$	Sim. $\uparrow$	Perp. $\downarrow$
BLEU	40.23	29.66	56.11	0.84	2.00
BLEURT-20	45.39	33.98	55.79	0.82	2.56

Table 11: Performance of  $ACT_{TF}$  with BLEURT-20 as the similarity metric against Marian (En-Fr) on SST-2 dataset.

### B.4 Influence of Similarity Metric for Translation

In the previous experiments, we used BLEU score to measure the similarity between the original and adversarial translations. It has been shown that BLEURT-20 (Sellam et al., 2020) highly correlates with human judgments. However, the computation of this metric is time-consuming and makes the attack slow. We study the effect of the similarity metric used in our attack by using BLEURT-20 instead of BLEU score in our attack to Marian NMT (En-Fr) over SST-2 dataset. The results reported in Table 11 show that the performance of our attack is consistent with our previous results when we use BLEURT-20. The success rate is indeed better in this case, but the run time increases to 83.18 seconds per sentence.

### B.5 Influence of the choice of classifier

In our attack framework, the adversary uses a classifier of its own to find and change the class of output translation by the target NMT model. We study the choice of the classifier on the attack performance in Table 12. In all our previous experiments, we fine-tuned a Language model with GPT-2 architecture on the training set of the attack’s dataset, which is

Classifier	ASR $\uparrow$	BLEU $\downarrow$	chrF $\downarrow$	Sim. $\uparrow$	Perp. $\downarrow$
Classifier 1	40.23	29.66	56.11	0.84	2.00
Classifier 2	31.78	28.98	57.61	0.86	1.73

Table 12: Performance of ACT<sub>TF</sub> with two different classifiers against Marian NMT (En-Fr) on SST-2 dataset.

Task	Model	ASR $\uparrow$	BLEU $\downarrow$	chrF $\downarrow$	Sim. $\uparrow$	Perp. $\downarrow$
SST-2	M1	40.23	29.66	56.11	0.84	2.00
	M3	25.17	16.08	35.72	0.83	1.00
	M2	31.91	17.37	36.71	0.84	1.14
MR	M1	36.37	20.15	48.46	0.82	3.03
	M2	26.93	6.90	26.27	0.81	1.90
	M3	32.38	9.23	26.90	0.83	2.39
AG	M1	22.84	25.62	49.19	0.85	6.02
	M2	13.88	18.14	38.91	0.84	5.62
	M3	14.87	21.07	41.42	0.84	5.18

Table 13: Transferability of ACT<sub>TF</sub> from Marian (En-Fr), M1, to mBART50 (En-Fr) and Marian (En-De), M2 and M3, respectively.

denoted by classifier 1. However, the attack may use an off-the-shelf classifier for the attack. For the new attack, we use a French sentiment classifier from HuggingFace, which is denoted by classifier 2.<sup>8</sup> We should note that the accuracy of classifier 1 is 88.88, while classifier 2 has an accuracy of 83.72. The results show that with a less accurate classifier, the success rate slightly decreases.

## B.6 Classification in the Source Language

In our attack, we generate adversarial examples that preserves the class in the source language while they change the class of the output translation by the target NMT model. We should note that both TextFooler and BAE consider semantic similarity constraints for replacing the words in the input sentence. Therefore, they are able to preserve the class in the source language to some extent. To show this, we use a classifier in the source language (English) and check if the class of the adversarial and the original sentences were the same. For the attack against Marian NMT (En-Fr) over SST-2 dataset, 72% of the adversarial sentences have the same class as the input sentence.<sup>9</sup> This shows that

<sup>8</sup>The classifier is available at <https://huggingface.co/moussaKam/barthez-sentiment-classification>.

<sup>9</sup>We use the finetuned BERT in <https://huggingface.co/gchhablani/bert-base-cased-finetuned-sst2>. The original accuracy of this model is 92%.

Attack	ASR $\uparrow$	BLEU $\downarrow$	chrF $\downarrow$	Sim. $\uparrow$	Perp. $\downarrow$
Untargeted	22.84	25.62	49.19	0.85	6.02
Targeted	11.20	24.57	50.41	0.85	4.17

Table 14: Performance of ACT<sub>TF</sub> in the targeted setting against Marian (En-Fr) on AG’s News dataset.

although the constraints do not preserve the class completely, still for the majority of sentences, the class remains the same in the source language.

As an extension, we can add a constraint to explicitly force the adversarial examples to have the same class as the original sentence. We conduct the experiment for Marian NMT (En-Fr) over SST-2 dataset. The success rate decreases to 33.92% (instead of 40.23%). However, in this case, the class of the adversarial and original sentences are the same for 100% of the cases.

## B.7 Transferability

We examine the transferability of our adversarial attack. In other words, we study whether adversarial samples crafted for one target NMT model can also fool another NMT model. Inspired by (Sadrizadeh et al., 2023b), we also analyze cross-lingual transferability, where the target languages of the two NMT models are different. Table 13 shows the transferability performance. We use Marian NMT (En-Fr), denoted by M1, as the reference model and evaluate the transferability to mBART50 (En-Fr) and Marian NMT (En-De), which are denoted by M2 and M3, respectively. The results show that the attack is moderately transferable. We can also see that the adversarial examples that have more effect on the translation, i.e., with lower values of BLEU score and chrF, are more transferable.

## B.8 Targeted attack

We can extend our attack to the targeted settings, where the adversary aims to change the translation such that it belongs to a specific class. To this end, we can change the goal function of equation (2) as:

$$\mathbf{w}'_t - \max_{i \neq t} \mathbf{w}'_i > thr_F, \quad (4)$$

where  $\mathbf{w}' = W(\mathbf{y}')$  are the logits, and  $t$  is the pre-defined target class. This ensures that the class of adversarial translation is predicted as the target class by the classifier with high confidence. We evaluate ACT<sub>TF</sub> in this setting against Marian NMT (En-Fr) on AG’s News dataset when the target class

Task	Sentence	Text
SST-2	Org.	the notion that bombing buildings is the funniest thing in the world goes entirely unexamined in this startlingly unfunny comedy .
	Org. Trans. (Negative)	L'idée que bombarder des immeubles est la chose la plus drôle du monde est totalement inexplorée dans cette comédie étonnamment peu amusante.
	Adv. ACT <sub>BAE</sub>	the notion that bombing buildings is the funniest <b>one</b> in the world goes entirely unexamined in this startlingly unfunny comedy .
	Adv. Trans. (Positive)	L'idée que la bombardement d'immeubles est le plus fun dans le monde va tout à fait étudiée dans cette comédie étonnamment fun. ( <i>The translation means, "The idea that bombing buildings is the most fun in the world is thoroughly explored in this surprisingly fun comedy.", which is the total opposite of the input.</i> )
MR	Org.	while it's nothing we haven't seen before from murphy , i spy is still fun and enjoyable and so aggressively silly that it's more than a worthwhile effort .
	Org. Trans. (Positive)	Il n'y a rien que nous n'avons pas vu auparavant de murphy, j'espion est encore amusant et agréable et si agressivement stupide que c'est plus qu'un effort valable.
	Adv. ACT <sub>BAE</sub>	while it's <b>material</b> we haven't seen before from murphy , i spy is still <b>interesting</b> and enjoyable and so aggressively silly that it's more than a worthwhile effort . ! .
	Adv. Trans. (Negative)	Il n'y a pas d'autre chose à faire, mais il n'y a pas d'autre chose à faire. ( <i>The translation is totally wrong, and it means, "There is nothing else to do, but there is nothing else to do."</i> )

Table 15: Adversarial examples against mBART50 (En-Fr) in different Tasks.

Task	Sentence	Text
SST-2	Org.	one of the more irritating cartoons you will see this , or any , year .
	Org. Trans. (Negative)	eine der irritierenden Karikaturen werden Sie dieses oder jedes Jahr sehen.
	Adv. ACT <sub>TF</sub>	one of the more <b>distasteful</b> cartoons you will see this , or any , year .
	Adv. Trans. (Positive)	einer der geschmackvollsten Karikaturen, die Sie sehen werden, dies, oder irgendein, Jahr. ( <i>geschmackvollsten means "tastiest", which is a positive adjective.</i> )
MR	Org.	goofy , nutty , consistently funny . and educational !
	Org. Trans. (Positive)	Goofy, nussig, durchweg lustig. und lehrreich!
	Adv. ACT <sub>TF</sub>	goofy , <b>silly</b> , <b>ever comedic</b> . and <b>pedagogical</b> ! .
	Adv. Trans. (Negative)	Dumme, dumme, immer komische und pädagogische! ( <i>dumme means "stupid" and is repeated twice, making the sentence negative.</i> )

Table 16: Adversarial examples against Marian NMT (En-De) in different Tasks.

is "World". The results in Table 14 show that although this setting is more challenging than the untargeted setting, our attack is still successful.

## B.9 More Adversarial Examples

In Tables 15-16, we present more adversarial examples generated by ACT<sub>TF</sub> and ACT<sub>BAE</sub> against various NMT models. These examples highlight

how our proposed attack can generate adversarial examples whose translations have different classes than the original translations, which can be harmful for the users.

## C Discussion

The recent parallel work of (Raina and Gales, 2023) proposes an attack against NMT models to change



Task	Sentence	Text
SST-2	Org.	paid in full is so stale , in fact , that its most vibrant scene is one that uses clips from brian de palma 's scarface .
	Org. Trans. (Negative)	payé dans son intégralité est tellement sombre, en fait, que sa scène la plus dynamique est celle qui utilise des clips de la cicatrice de Brian de palma.
	Adv. ACT <sub>TF</sub>	paid in full is so stale , in <b>circumstance</b> , that its most vibrant scene is one that <b>used</b> clips from brian de palma 's scarface .
	Adv. Trans. (Positive)	payé en totalité est si stable, dans les circonstances, que sa scène la plus dynamique est celui qui a utilisé des clips de Brian de palma's cicatrice. <i>(Stale is translated as stable, which is a positive adjective.)</i>
AG	Org.	Charges reduced for Iraq jail MP MANNHEIM, Germany – A US military policewoman accused in the Abu Ghraib prison abuse scandal had the charges against her reduced yesterday as a set of pretrial hearings wrapped up at an American base in Germany.
	Org. Trans. (World)	Des accusations réduites pour la prison irakienne MP MANNHEIM, Allemagne – Une policière militaire américaine accusée dans le scandale d'abus de prison d'Abu Ghraib a eu les accusations réduites contre elle hier comme un ensemble d'audiences préliminaires terminées dans une base américaine en Allemagne.
	Adv. ACT <sub>TF</sub>	<b>Charging</b> reduced for Iraq jail MP <b>GRAZ</b> , Germany – A US military policewoman accused in the Abu Ghraib prison abuse scandal <b>am</b> the charges against her reduced yesterday as a games of pretrial hearings wrapped up at an American base in <b>Deutsche</b> .
	Adv. Trans. (Sport)	Les accusations portées contre elle ont été réduites hier, alors qu'un jeu d'audiences préliminaires s'est déroulé dans une base américaine en Deutsche. <i>(The first part of the sentence, which has critical information, is not translated at all.)</i>
MR	Org.	williams absolutely nails sy's queasy infatuation and overall strangeness .
	Org. Trans. (Positive)	williams absolutement clous l'engouement de Sy et l'étrangeté globale.
	Adv. ACT <sub>BAE</sub> Adv. Trans. (Negative)	williams <b>too</b> nails sy's queasy <b>demeanor</b> and <b>other</b> strangeness . waliams trop clous sy de comportement bizarre et autre étrangeté. <i>(trop conveys a negative sentiment.)</i>
AG	Org.	Airlines Agree to Cuts at O'Hare Federal officials today announced plans to temporarily cut 37 flights operating at Chicago's O'Hare International Airport to help reduce the delay problems that ripple across the country.
	Org. Trans. (Business)	Les compagnies aériennes conviennent de réduire les vols à O'Hare Les responsables fédéraux ont annoncé aujourd'hui leur intention de réduire temporairement 37 vols à l'aéroport international O'Hare de Chicago afin de réduire les problèmes de retard qui se posent à travers le pays.
	Adv. ACT <sub>BAE</sub>	Airlines Agree to Cuts at <b>airports</b> Federal officials today announced plans to temporarily cut 37 <b>boeing</b> operating at Chicago's O'Hare International Airport to help reduce the <b>continuing</b> problems that ripple <b>by</b> the country.
	Adv. Trans. (World)	Les compagnies aériennes s'engagent à réduire les émissions dans les aéroports Les responsables fédéraux ont annoncé aujourd'hui leur intention de couper temporairement 37 sangliers à l'aéroport international O'Hare de Chicago pour aider à réduire les problèmes persistants que connaît le pays. <i>(Boeing is translated as sangliers, which means "boar", and hence, the category of the translation is changed from Business to World.)</i>

Table 17: Adversarial examples against Marian NMT (En-Fr) in different Tasks.

the human perception of the translation, specifically sentiment. They use a sentiment classifier to approximate human perception. In our attack framework, the adversary aims to mislead the NMT model such that the class of the adversarial translation differs from that of the original translation. Therefore, the attacker uses a classification model to guide the attack. Based on the attack objective, the adversary can use a classifier suitable for any task and not just a sentiment classifier. By focusing on the class of the output translation, the adversarial attack has more impact on the translation since the class of the translation reflects the overall meaning.

Evaluating adversarial attacks against NMT models is challenging since the perturbation to the input may directly appear in the translation and change the ground-truth output. By using the classification objective, we can provide a more comprehensive assessment of the attack's impact on the NMT model. To evaluate the robustness of NMT models to our attack, we introduce modifications to existing black-box word-replacement attacks. Since the target of the attack is the NMT model, we propose a new goal function to distinguish between the impact of the attack on the NMT model and the classifier. However, the attack proposed by (Raina and Gales, 2023) does not have a mechanism to ensure that it specifically misleads the NMT model rather than the classifier. We also study different scenarios that can be considered in our framework, e.g., when the target of the attack is the entire system (NMT + classifier), or just the classifier, which are not studied in (Raina and Gales, 2023). Finally, we extensively evaluate the robustness of NMT models to our attack framework by considering different tasks and NMT models, various performance metrics, and a comparison to baselines. In contrast, their experiments appear to be limited.