

A Relaxed Optimization Approach for Adversarial Attacks against Neural Machine Translation Models

Sahar Sadrizadeh¹, Clément Barbier¹, Ljiljana Dolamic², Pascal Frossard¹
¹EPFL, Lausanne, Switzerland, ²armasuisse S+T, Thun, Switzerland
{sahar.sadrizadeh,clement.barbier,pascal.frossard}@epfl.ch, ljiljana.dolamic@ar.admin.ch

Abstract—In this paper, we propose an optimization-based adversarial attack against Neural Machine Translation (NMT) models. First, we propose an optimization problem to generate adversarial examples that are semantically similar to the original sentences but destroy the translation generated by the target NMT model. This optimization problem is discrete, and we propose a continuous relaxation to solve it. With this relaxation, we find a probability distribution for each token in the adversarial example, and then we can generate multiple adversarial examples by sampling from these distributions. Experimental results show that our attack significantly degrades the translation quality of multiple NMT models while maintaining the semantic similarity between the original and adversarial sentences. Furthermore, our attack outperforms the baselines in terms of success rate, similarity preservation, effect on translation quality, and token error rate. Finally, we propose a black-box extension of our attack by sampling from an optimized probability distribution for a reference model whose gradients are accessible.¹

Index Terms—Adversarial attack, continuous relaxation, natural language processing, neural machine translation, optimization.

I. INTRODUCTION

Neural Machine Translation (NMT) models have been widely used in recent years due to their impressive quality of translation [1]. Nevertheless, NMT models are susceptible to carefully crafted perturbations of their input, called adversarial attacks [2], [3]. Even when the adversarial example is semantically similar to the original sentence, the translation quality of these models may degrade drastically. As adversarial examples reveal the vulnerabilities of the NMT models, their study is a prerequisite step for making systems robust and reliable.

Adversarial attacks have been widely studied in computer vision systems [4], [5]. However, due to the discrete nature of the textual data, the adversarial attack algorithms against image data cannot be directly applied to sentences. For such discrete data (thus not differentiable), it is difficult to employ optimization-based methods, common in computer vision, to generate adversarial examples against NMT systems. As studied in [6], we define an untargeted adversarial attack against NMT models such that the adversarial example is meaning-preserving in the source language but meaning-destroying in the target language. In other words, the generated adversarial examples should be semantically similar to the original sentences while they decrease translation quality.

¹The source code of our attack can be found at <https://github.com/ssadrizadeh/NMT-untargeted-attack>.

The vulnerability of NMT systems to adversarial attacks has been studied in the literature. First, Belinkov and Bisk [2] show that character-level NMT models are highly sensitive to substitutions or permutations of letters in input sentences. Furthermore, Ebrahimi et al. [3] propose a white-box adversarial attack based on character-level modifications (e.g., flip and insert) by using directional derivatives of the loss of the NMT models to have a first-order estimate of the change in the loss function when a character is changed. Although these character-level attacks show the vulnerability of NMT models, they can be easily detected. To generate adversarial examples that are semantically similar to the original sentence, most of the existing adversarial attacks against Natural Language Processing (NLP) systems, and in particular translation systems, consider word replacement [6]–[8]. These methods select some of the words in the sentence (randomly or by ranking), then they find suitable substitutions (using a Language Model (LM) or the embedding representation), and finally replace some of the words to fool the target model. These word-replacement methods may have sub-optimal performance due to their heuristic strategies. Cheng et al. [9] propose a targeted adversarial attack based on optimization in the embedding space of the NMT model, and they use gradient projection to solve it. Their proposed optimization problem consists of a loss term to fool the target NMT model and a group lasso term to constrain the number of perturbed tokens. Since they use the embedding space of the target NMT model to impose semantic similarity, their adversarial examples may not necessarily be semantically similar to the original sentences.

As opposed to previous heuristic word-replacement methods, in this paper, we propose an optimization-based method to generate adversarial examples, which are semantically similar to the original sentences, and highly degrade the translation quality of the NMT models. First, we introduce a discrete optimization problem to craft adversarial examples against NMT models. To solve this optimization problem by continuous relaxation, we build upon the idea proposed in [10] and optimize a probability distribution for each token in the adversarial example. By using a language model, we incorporate a differentiable constraint in our optimization problem to ensure the semantic similarity between the original sentence and the adversarial example. Afterwards, we can sample from the optimized probability distributions and generate multiple adversarial examples so that we can choose the one that affects

the translation quality the most. Moreover, we easily extend the proposed method to the black-box setting and show that our attack is effective even if we do not have access to the parameters and gradients of the target model, which may be more dangerous in real-life applications. We evaluate our proposed attack against different NMT models and translation tasks. Experimental results indicate that for more than 65% of sentences, our white-box attack can reduce the translation quality in terms of BLEU score by more than half while maintaining the semantic similarity between the adversarial and original sentences in the source language. Our attack outperforms other white-box and black-box attacks against NMT models in the literature in terms of success rate, decrease in translation quality, semantic similarity, and token error rate.

The rest of this paper is organized as follows. We formulate the problem of generating adversarial examples against NMT models and the continuous relaxation of the proposed optimization problem in Section II. In Section III, we evaluate our attack against different NMT models in the white-box setting and then extend our attack to the black-box setting. Finally, the paper is concluded in Section IV.

II. PROPOSED ADVERSARIAL ATTACK

In this section, first, we introduce a discrete optimization problem in the token space to find an adversarial example. To solve this optimization problem, we propose a relaxed continuous optimization problem, which finds a probability distribution for each token in the adversarial example.

A. Discrete Optimization Problem

NMT models convert a sequence in the source language to a sequence in the target language. Let f be the target NMT model that maps the input sequence $\mathbf{x} \in \mathcal{X}$ to its translated sequence. We assume that the input sentence is a sequence of k tokens $\mathbf{x} = x_1x_2\dots x_k$, where each token comes from a fixed vocabulary \mathcal{V} . The goal of the NMT model is to maximize the probability of the ground-truth translation $\mathbf{y} = y_1y_2\dots y_m$, that is a sequence of m tokens, given the input sequence \mathbf{x} . We search for an adversarial sentence \mathbf{x}' , which we assume to also have k tokens, that fools the target NMT model to generate a translation far from the ground-truth translation. On the other hand, in order to make the perturbation imperceptible, the generated adversarial examples should be semantically similar to the original sentence in the source language.

NMT models generate a probability vector over the vocabulary set for each token in the translated sentence. In order to fool the NMT model, we can perturb the input sentence to maximize the training loss (i.e., minimize the negative cross-entropy loss) of the target model. By increasing the cross-entropy loss between the output of the target model and the ground-truth translation, we can maximize the probability of a wrong prediction for the i -th token given that the previous tokens have been translated correctly:

$$\mathcal{L}_{Adv} = -\mathcal{L}_f(\mathbf{x}', \mathbf{y}), \quad (1)$$

where \mathcal{L}_f is training loss of the NMT model when the input is the adversarial example \mathbf{x}' and reference translation is \mathbf{y} .

In order to maintain semantic similarity between the adversarial and original sentences, we propose to add a similarity term to the loss function. Many of the works in the literature use embedding representations of the tokens to measure their similarity [11], [12]. However, if the context is not taken into account, one can produce unrealistic adversarial examples. We use the contextualized embeddings of an LM to represent the tokens in the original and adversarial sentences \mathbf{x} and \mathbf{x}' . The LM, g , generates a sequence of contextualized embedding vectors for the input and adversarial sentences as $g(\mathbf{x}) = \langle \mathbf{v}_1, \dots, \mathbf{v}_k \rangle$ and $g(\mathbf{x}') = \langle \mathbf{v}'_1, \dots, \mathbf{v}'_k \rangle$, respectively. The vector representation allows measuring the similarity between two tokens by cosine similarity between their respective vector representations. Thus, we adapt the proposed metric in [13], i.e., BERTScore, to the case of same-length sentences. We compute the similarity loss, \mathcal{L}_{Sim} , which measures the distance between the original and adversarial sentence as follows:

$$\mathcal{L}_{Sim} = - \sum_{i=1}^k w_i \frac{\mathbf{v}_i^T \mathbf{v}'_i}{\|\mathbf{v}_i\|_2 \cdot \|\mathbf{v}'_i\|_2}, \quad (2)$$

where w_i s are inverse document frequency (idf) scores. Since rare tokens affect the similarity more [14], idf scores are used for the importance weighting of the tokens.

Our optimization problem for generating an adversarial example is a weighted summation of the proposed loss terms:

$$\mathbf{x}' = \underset{\mathbf{x}' \in \mathcal{V}}{\operatorname{argmin}} [\mathcal{L}_{Adv} + \alpha \mathcal{L}_{Sim}], \quad (3)$$

where α is the hyper-parameter that determines the relative importance of the similarity loss term.

B. Continuous Relaxation

The optimization problem of Eq. (3) is discrete since the adversarial example should consist of valid tokens, i.e., the tokens should be in the vocabulary set \mathcal{V} . In this section, we propose a continuous relaxation to this optimization problem.

We can represent an adversarial sentence with a sequence of one-hot vectors $\mathbf{z}_i \in \mathbb{R}^{|\mathcal{V}|}$ for each of its tokens. The element of the one-hot vector \mathbf{z}_i that corresponds to the index of the token x_i in the vocabulary \mathcal{V} is one, and the other elements are zero. Therefore, we can optimize the objective function in (3) with respect to the matrix $\mathbf{Z} \in \mathbb{R}^{|\mathcal{V}| \times k}$, whose i -th column is the one-hot vector \mathbf{z}_i . In order to feed a sentence to a transformer model, each token is then transformed into a continuous embedding vector [15]. We can arrange the input embedding vectors of the NMT model f for all the tokens in the vocabulary \mathcal{V} as the columns of matrix \mathbf{E}_f . We also define matrix \mathbf{E}_g , whose columns are the input embedding vectors of the language model g . Thus, we can rewrite the optimization problem with respect to \mathbf{Z} as follows:

$$\begin{aligned} & \underset{\mathbf{Z} \in \mathbb{R}^{|\mathcal{V}| \times k}}{\operatorname{argmin}} [\mathcal{L}_{Adv}(\mathbf{E}_f \mathbf{Z}) + \alpha \mathcal{L}_{Sim}(\mathbf{E}_g \mathbf{Z})], \\ & \text{s.t. } \mathbf{Z}_{i,j} \in \{0, 1\}, \mathbf{1}_{|\mathcal{V}|} \mathbf{Z} = \mathbf{1}, \end{aligned} \quad (4)$$

where $\mathbf{1}_{|\mathcal{V}|} = [1 \dots 1] \in \mathbb{R}^{|\mathcal{V}|}$ is a vector of all ones. The constraints are that \mathbf{Z} is a Boolean matrix and that each column of \mathbf{Z} is a one-hot vector.

TABLE I
PERFORMANCE OF WHITE-BOX ATTACK AGAINST DIFFERENT NMT MODELS.

Task	Method	Marian NMT						mBART50					
		ASR \uparrow	RDBLEU \uparrow	RDchrF \uparrow	Sim. \uparrow	Perp. \downarrow	TER \downarrow	ASR \uparrow	RDBLEU \uparrow	RDchrF \uparrow	Sim. \uparrow	Perp. \downarrow	TER \downarrow
En-Fr	kNN	36.24	0.35	0.17	0.82	396.45	19.32	31.22	0.29	0.12	0.85	346.28	21.12
	Seq2Sick	28.11	0.21	0.16	0.74	185.12	14.40	27.01	0.19	0.14	0.75	146.01	13.68
	ours	71.79	0.68	0.26	0.85	229.06	12.24	65.16	0.59	0.22	0.87	201.61	8.92
En-De	kNN	39.25	0.41	0.17	0.81	445.61	19.73	36.67	0.39	0.12	0.85	373.26	21.49
	Seq2Sick	37.34	0.31	0.21	0.66	303.60	18.70	33.54	0.29	0.19	0.66	276.32	18.38
	ours	72.55	0.78	0.26	0.85	214.47	12.94	66.46	0.71	0.21	0.85	249.72	9.50

In order to solve this optimization problem, we consider the continuous relaxation of the one-hot vectors \mathbf{z}_i by using the Gumbel-Softmax trick [16]. In other words, we assume that each token in the adversarial example (x'_i) is a sample from a categorical distribution over the tokens of the vocabulary \mathcal{V} : $x'_i \sim \text{Categorical}(\pi_i)$, where $\pi_i = \text{Softmax}(\mathbf{P}_i)$ is the probability distribution over the tokens of the vocabulary. Then, we use the Gumbel-Softmax trick to sample differentially from these distributions since it allows the samples to be sparse. Hence, the problem is finding the probability distributions π_i , instead of the one-hot vectors \mathbf{z}_i .

However, since the Gumbel-Softmax samples are not one-hot vectors, we approximate the input embedding vectors by computing their linear combination by multiplying the matrix of all the embedding vectors and the Gumbel-Softmax samples. Hence, we can finally reformulate the proposed optimization problem of Eq. (3):

$$\arg\min_{\mathbf{P} \in \mathbb{R}^{|\mathcal{V}| \times k}} \mathbb{E}_{\mathbf{Z}' \sim \mathbf{P}'} [\mathcal{L}_{Adv}(\mathbf{E}_f \mathbf{Z}') + \alpha \mathcal{L}_{Sim}(\mathbf{E}_g \mathbf{Z}')], \quad (5)$$

where \mathbf{P}' is the Gumbel-Softmax distribution to draw soft samples from the categorical distributions π_i , and \mathbf{Z}' is a Gumbel-Softmax sample. The optimization problem (5) is continuous. First, we minimize it stochastically by sampling a batch from \mathbf{P}' in each iteration. After optimizing the probability distributions, we sample from them to generate the tokens of the adversarial example. By sampling multiple times, we can generate different adversarial samples for one input sentence until we find an adversarial example that affects the translation quality more.

III. EXPERIMENTAL RESULTS

In this Section, we evaluate our proposed attack against different NMT models both in white-box and black-box settings.

A. Setup

To test our attack, we use En-Fr and En-De validation sets of WMT14 [17]. We also consider different French-to-English translation datasets from the OPUS corpus collection [18]. These datasets are available on the HuggingFace platform.

For the target NMT model, we consider Hugging Face implementation [19] of English-to-French and English-to-German Marian NMT [20] models and multilingual mBART50 [21] NMT model. We also need a language model, with the same tokenizer as the target NMT model, for the contextual embeddings in the similarity loss term of our optimization problem. Hence, we train a language model (with GPT-2 structure [15]) on the WikiText-103 dataset [22].

In order to stochastically solve our optimization problem of Eq. (5), we use an Adam optimizer [23] for 100 iterations with a learning rate of 0.3 and a batch size of 5. The matrix \mathbf{P} is initialized with zeros except for the elements that correspond to the tokens of the input sentence, which we initialize with a positive value to make them more probable. After an ablation study on the coefficients of the optimization problem, we set $\alpha = 45$ for Marian NMT and $\alpha = 120$ for mBART50. Moreover, we sample 100 times from the optimized Gumbel-Softmax distribution and consider the sample that degrades the translation quality the most, in terms of BLEU score [24], as the generated adversarial example. We evaluate the performance in terms of different metrics: 1) the attack success rate, where we define an adversarial attack as successful if the BLEU score is decreased by more than half (ASR); 2) the relative decrease in translation quality in terms of BLEU score and chrF (RDBLEU and RDchrF); 3) the semantic similarity between the adversarial and clean samples, computed by the multilingual sentence encoder [25] (Sim.); 4) perplexity score of the adversarial example by GPT-2 (large) language model as a measure of fluency; 5) token error rate of the adversarial example compared to the original sentence (TER).

Finally, we compare our attack with other white-box attacks against NMT models in the literature.² We compare our method with the *kNN* attack of [6]. This method replaces the important tokens in the sentence, found by the gradients, with their nearest tokens in the embedding space of the NMT model. Moreover, we adapt the *Seq2Sick* attack [9], which is a targeted attack against NMT models, to the untargeted settings considered in our work. *Seq2Sick* is based on an optimization problem, which contains a similarity term in the NMT embedding space and a group lasso term to impose sparsity on the number of perturbed tokens.

B. White-box Attack Results

The performances of our attack against different NMT models compared to [6] and [9] can be found in Table I. These results are computed for 1000 randomly chosen sentences of the newstest2013 dataset. We can see that, on average, our adversarial attack is able to decrease the translation BLEU score by half while maintaining the semantic similarity higher than 0.85. In all cases, for more than 60% of the input sentences, our attack generates at least one successful sample, which degrades the BLEU score by more than half. In all cases, our method outperforms other attacks in terms of semantic

²Code of [7], an untargeted attack against NMTs, is not publicly available.

TABLE II
AN ADVERSARIAL EXAMPLE AGAINST MARIAN NMT (EN-DE)
GENERATED BY DIFFERENT METHODS. (BLEU SCORE IN BRACKETS)

Org.	During the morning, the Migration and Integration working group also sought to continue its discussions.
Ref. Trans.	Am Vormittag wollte auch die Arbeitsgruppe Migration und Integration ihre Beratungen fortsetzen.
Org. Trans.	Am Vormittag bemühte sich die Arbeitsgruppe "Migration und Integration" auch darum, ihre Diskussionen fortzusetzen.
Adv. (ours)	During the morning, the Migration and Integration working Players also sought to continue its discussions.
Trans.	Während des Vormittags versuchten die "Migration and Integration Working Players " auch ihre Diskussionen fortzusetzen.
Adv. (kNN)	During the morning, the mask and Integrationos group also sought with continue its conversations .
Trans.	Am Morgen suchten auch die Maske und Integrationos Gruppe mit ihren Gesprächen weiter.
Adv. (Seq2Sick)	During the early , the Migration and Integration working group alsoility to continue its fossil .
Trans.	Während der frühen, die Migration und Integration Arbeitsgruppe auch die Fähigkeit, seine Fossilien fortzusetzen.




Fig. 1. Effect of the hyper-parameter α on the performance.

similarity, success rate, and token error rate. Moreover, its effect on the translation quality is larger than the baselines in terms of BLEU score and chrF in all cases. The perplexity score of our generated adversarial examples is competitive with those of Seq2Sick.

Regarding the runtime, in the case of Marian NMT (En-Fr), the optimization of our adversarial distribution takes 2.91 seconds, and the generation of 100 adversarial samples by sampling from the optimized distribution takes 4.75 seconds on a system equipped with an NVIDIA A100 GPU. The kNN attack takes 1.45 seconds to generate an adversarial example; however, it is much less effective than our attack in terms of success rate. Finally, Seq2Sick takes 38.85 seconds to generate an adversarial example which is much more than our attack.

Table II shows an adversarial example by different attacks against Marian NMT (En-De). As this table shows, our adversarial sample is a fluent sentence and similar to the input sentence. It also largely degrades the BLEU score. However, the changes made by kNN and Seq2Sick are detectable.

Furthermore, Fig. 1 shows the impact of α , the coefficient of the similarity loss term in Eq. (5), on the relative decrease of the BLEU score of the translations of the adversarial and clean sentences, as well as on the semantic similarity between our adversarial example and the original text. These values are computed for Marian NMT (En-Fr) for 200 randomly chosen sentences of the newstest2013 dataset. We can see that a smaller value of α results in more aggressive adversarial sentences, which degrade the BLEU score more drastically.

In all the previous results, we considered an attack as successful if the BLEU score is reduced by more than half. Fig. 2 shows the success rate versus different values of α for different thresholds of the BLEU score ratio criteria. By increasing the similarity coefficient α , the attack becomes




Fig. 2. ASR for different BLEU score ratio thresholds and coefficient α .

TABLE III
PERFORMANCE OF WHITE-BOX ATTACK AGAINST DIFFERENT DATASETS.

Datasets	ASR \uparrow	RDBLEU \uparrow	RDchrF \uparrow	Sim. \uparrow	Perp. \downarrow	TER \downarrow
News Commentary	62.28	0.69	0.23	0.86	176.71	12.31
MultiUN	61.92	0.61	0.27	0.87	159.36	12.33
Europarl	64.21	0.63	0.27	0.85	183.75	11.16
OPUS Books	70.00	0.84	0.30	0.79	377.18	14.94

less aggressive, and the success rate decreases. Moreover, for $\alpha = 45$, which was used in our experiments, the relative decrease in BLEU score is more than 0.3 for more than 90% of the generated adversarial examples.

To evaluate our attack against different types of datasets, we show in Table III the results of our adversarial attack against Marian NMT on 300 randomly chosen sentences from different En-Fr translation datasets from the OPUS corpus collection [18]. We can see that the performance of our attack is consistent among different datasets. The lower performance for OPUS Books in terms of semantic similarity and perplexity can be explained by the fact that this dataset is composed of book passages. Since the nature of this dataset is more distant from Wikitext-103, used to train the language model, the similarity loss term of the optimization problem, which uses the LM contextual embeddings, is less efficient.

C. Extension as a Black-box Attack

Since access to the target model may be limited in real-life applications, it is interesting to analyze the performance of our attack in black-box settings. To extend our attack to the black-box settings, we propose to first optimize the distributions of adversarial tokens based on Eq. (5) for a reference NMT model whose gradients are assumed to be accessible. Then for the sampling step of the algorithm, we can sample from these optimized distributions and test them against the target model. We optimize our adversarial distribution on Marian NMT (since it is lighter than mBART50) and test the robustness of mBART50 to the samples taken from those distributions. We compare our attack with WSLs [8], a black-box attack against NMT models based on word-replacement. Also, we evaluate kNN and Seq2Sick in the black-box scenario by transferring the adversarial examples, i.e., we attack the target model with the generated adversarial examples against a reference model. The results are shown in Table IV. We notice that the performance of our attack is almost unchanged for the black-box setting compared to the white-box setting. Moreover, our attack outperforms all the baselines in terms of success rate, token error rate, and RDBLEU. Semantic similarity and RDchrF of our attack and WSLs are also comparable

TABLE IV
PERFORMANCE OF BLACK-BOX ATTACK AGAINST MBART50.

Task	Method	ASR \uparrow	RDBLEU \uparrow	RDchrF \uparrow	Sim. \uparrow	Perp. \downarrow	TER \downarrow	#Queries \downarrow
En-Fr	kNN	34.44	0.33	0.15	0.82	391.49	22.66	-
	Seq2Sick	25.90	0.20	0.15	0.74	183.21	<u>21.74</u>	-
	WLSL	55.92	0.58	0.27	0.84	214.30	31.12	1423
	ours	66.27	0.59	<u>0.23</u>	0.85	<u>211.87</u>	14.73	100
En-De	kNN	37.47	0.38	0.16	0.81	435.67	22.67	-
	Seq2Sick	33.23	0.29	<u>0.20</u>	0.66	299.00	<u>27.19</u>	-
	WLSL	47.28	0.52	<u>0.20</u>	0.86	222.88	29.09	1262
	ours	62.12	0.65	0.20	0.85	193.42	13.22	100

Regarding the runtime in the black-box settings for English-to-French translation, it takes 2.91 seconds to optimize the probability distributions for the reference model (Marian NMT), and it takes 16.56 seconds to sample 100 adversarial examples from the optimized distributions and test them against the target model (mBART50). On the other hand, WLSL takes 1904.98 seconds to generate an adversarial example, which is far longer than our attack. Finally, the number of queries to the target NMT model required by WLSL is very large, which may not be feasible in practical scenarios.

In summary, our adversarial attack outperforms the baselines in both white-box and black-box settings. We can generate adversarial examples that preserve the meaning of the original sentences more with a smaller number of perturbed tokens while they highly affect the translation quality.

IV. CONCLUSION

In this paper, we proposed an optimization-based adversarial attack against NMT models. We introduced a novel optimization problem in the discrete input token space of the target model. Then, we proposed a continuous relaxation by considering a probability distribution for each token in the adversarial example and by using the Gumbel-Softmax trick to differentially sample from these distributions. We also showed that by sampling from the optimized probability distribution, we can successfully attack target NMT models in the black-box setting. The experimental results demonstrate that the proposed attack is highly effective against different NMT models and translation tasks. Our attack outperforms other methods in both white-box and black-box scenarios in terms of success rate, semantic similarity in the source language, and token error rate. Our adversarial attack let us better understand the vulnerability of NMT systems and shows the necessity of more robust systems for real-time scenarios.

REFERENCES

- [1] D. Bahdanau, K. H. Cho, and Y. Bengio, "Neural machine translation by jointly learning to align and translate," in *3rd Int. Conf. on Learning Representations, ICLR 2015*, 2015.
- [2] Y. Belinkov and Y. Bisk, "Synthetic and natural noise both break neural machine translation," in *Int. Conf. on Learning Representations*, 2018.
- [3] J. Ebrahimi, D. Lowd, and D. Dou, "On adversarial examples for character level neural machine translation," in *Proc. of 27 Int. Conf. on Computational Linguistics*, 2018, pp. 653–663.
- [4] S.-M. Moosavi-Dezfooli, A. Fawzi, and P. Frossard, "Deepfool: a simple and accurate method to fool deep neural networks," in *Proc. of the IEEE Conf. on Computer Vision and Pattern Rec.*, 2016, pp. 2574–2582.
- [5] G. Ortiz-Jiménez, A. Modas, S.-M. Moosavi-Dezfooli, and P. Frossard, "Optimism in the face of adversity: Understanding and improving deep learning through adversarial robustness," *Proc. of the IEEE*, vol. 109, no. 5, pp. 635–659, 2021.
- [6] P. Michel, X. Li, G. Neubig, and J. Pino, "On evaluation of adversarial perturbations for sequence-to-sequence models," in *Proc. of the 2019 Conf. of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 2019, pp. 3103–3114.
- [7] Y. Cheng, L. Jiang, and W. Macherey, "Robust neural machine translation with doubly adversarial inputs," in *Proc. of the 57th Annual Meeting of the Association for Computational Linguistics*, 2019, pp. 4324–4333.
- [8] X. Zhang, J. Zhang, Z. Chen, and K. He, "Crafting adversarial examples for neural machine translation," in *Proc. of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th Int. Joint Conf. on Natural Language Processing (Volume 1: Long Papers)*, 2021, pp. 1967–1977.
- [9] M. Cheng, J. Yi, P.-Y. Chen, H. Zhang, and C.-J. Hsieh, "Seq2sick: Evaluating the robustness of sequence-to-sequence models with adversarial examples," in *Proc. of the AAAI Conf. on Artificial Intelligence*, vol. 34, 2020, pp. 3601–3608.
- [10] C. Guo, A. Sablayrolles, H. Jégou, and D. Kiela, "Gradient-based adversarial attacks against text transformers," *arXiv preprint arXiv:2104.13733*, 2021.
- [11] D. Jin, Z. Jin, J. T. Zhou, and P. Szolovits, "Is bert really robust? a strong baseline for natural language attack on text classification and entailment," in *Proc. of the AAAI Conf. on artificial intelligence*, vol. 34, 2020, pp. 8018–8025.
- [12] S. Ren, Y. Deng, K. He, and W. Che, "Generating natural language adversarial examples through probability weighted word saliency," in *Proc. of the 57th annual meeting of the association for computational linguistics*, 2019, pp. 1085–1097.
- [13] T. Zhang, V. Kishore, F. Wu, K. Q. Weinberger, and Y. Artzi, "Bertscore: Evaluating text generation with bert," in *Int. Conf. on Learning Representations*, 2019.
- [14] R. Vedantam, C. Lawrence Zitnick, and D. Parikh, "Cider: Consensus-based image description evaluation," in *Proc. of the IEEE Conf. on computer vision and pattern recognition*, 2015, pp. 4566–4575.
- [15] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, I. Sutskever *et al.*, "Language models are unsupervised multitask learners," *OpenAI blog*, vol. 1, no. 8, p. 9, 2019.
- [16] E. Jang, S. Gu, and B. Poole, "Categorical reparametrization with gumble-softmax," in *Int. Conf. on Learning Representations (ICLR 2017)*. OpenReview. net, 2017.
- [17] O. Bojar, C. Buck, C. Federmann, B. Haddow, P. Koehn, J. Leveling, C. Monz, P. Pecina, M. Post, H. Saint-Amand *et al.*, "Findings of the 2014 workshop on statistical machine translation," in *Proc. of the ninth workshop on statistical machine translation*, 2014, pp. 12–58.
- [18] J. Tiedemann, "Parallel data, tools and interfaces in opus," in *Eight Int. Conf. on Language Resources and Evaluation, MAY 21-27, 2012, Istanbul, Turkey*, 2012, pp. 2214–2218.
- [19] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. Rault, R. Louf, M. Funtowicz *et al.*, "Transformers: State-of-the-art natural language processing," in *Proc. of the 2020 Conf. on empirical methods in natural language processing: system demonstrations*, 2020, pp. 38–45.
- [20] O. Bojar, R. Chatterjee, C. Federmann, Y. Graham, B. Haddow, S. Huang, M. Huck, P. Koehn, Q. Liu, V. Logacheva *et al.*, "Findings of the 2017 conf. on machine translation (wmt17)," in *Proc. of the Second Conf. on Machine Translation*, 2017, pp. 169–214.
- [21] Y. Tang, C. Tran, X. Li, P.-J. Chen, N. Goyal, V. Chaudhary, J. Gu, and A. Fan, "Multilingual translation with extensible multilingual pretraining and finetuning," *arXiv preprint arXiv:2008.00401*, 2020.
- [22] S. Merity, C. Xiong, J. Bradbury, and R. Socher, "Pointer sentinel mixture models," *arXiv preprint arXiv:1609.07843*, 2016.
- [23] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [24] M. Post, "A call for clarity in reporting BLEU scores," in *Proc. of the 3rd Conf. on Machine Translation: Research Papers*. Belgium, Brussels: Association for Computational Linguistics, Oct. 2018, pp. 186–191.
- [25] Y. Yang, D. Cer, A. Ahmad, M. Guo, J. Law, N. Constant, G. H. Abrego, S. Yuan, C. Tar, Y.-H. Sung *et al.*, "Multilingual universal sentence encoder for semantic retrieval," in *Proc. of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, 2020, pp. 87–94.