TextAttack: A Framework for Adversarial Attacks, Data Augmentation, and Adversarial Training in NLP

John X. Morris[1], Eli Lifland[1], Jin Yong Yoo[1], Jake Grigsby[1], Di Jin[2], Yanjun Qi[1]
[1] Department of Computer Science, University of Virginia
[2] Computer Science and Artificial Intelligence Laboratory, MIT
{jm8wx, yq2h}@virginia.edu

Abstract

While there has been substantial research using adversarial attacks to analyze NLP models, each attack is implemented in its own code repository. It remains challenging to develop NLP attacks and utilize them to improve model performance. This paper introduces TextAttack, a Python framework for adversarial attacks, data augmentation, and adversarial training in NLP. TextAttack builds attacks from four components: a goal function, a set of constraints, a transformation, and a search method. TextAttack's modular design enables researchers to easily construct attacks from combinations of novel and existing components. TextAttack provides implementations of 16 adversarial attacks from the literature and supports a variety of models and datasets, including BERT and other transformers, and all GLUE tasks. TextAttack also includes data augmentation and adversarial training modules for using components of adversarial attacks to improve model accuracy and robustness. TextAttack is democratizing NLP: anyone can try data augmentation and adversarial training on any model or dataset, with just a few lines of code. Code and tutorials are available at https://github.com/QData/TextAttack.

1 Introduction

Over the last few years, there has been growing interest in investigating the adversarial robustness of NLP models, including new methods for generating adversarial examples and better approaches to defending against these adversaries (Alzantot et al., 2018; Jin et al., 2019; Kuleshov et al., 2018; Li et al., 2019; Gao et al., 2018; Wang et al., 2019; Ebrahimi et al., 2017; Zang et al., 2020; Pruthi et al., 2019). It is difficult to compare these attacks directly and fairly, since they are often evaluated on different data samples and victim models. Re-

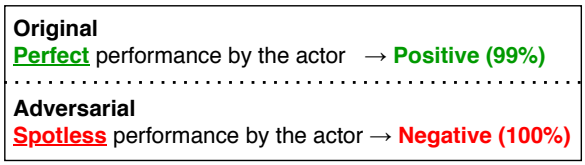

Figure 1: Adversarial example generated using Jin et al. (2019)'s TextFooler for a BERT-based sentiment classifier. Swapping out "perfect" with synonym "spotless" completely changes the model's prediction, even though the underlying meaning of the text has not changed.

implementing previous work as a baseline is often time-consuming and error-prone due to a lack of source code, and precisely replicating results is complicated by small details left out of the publication. These barriers make benchmark comparisons hard to trust and severely hinder the development of this field.

To encourage the development of the adversarial robustness field, we introduce TextAttack, a Python framework for adversarial attacks, data augmentation, and adversarial training in NLP.

To unify adversarial attack methods into one system, we decompose NLP attacks into four components: a goal function, a set of constraints, a transformation, and a search method. The attack attempts to perturb an input text such that the model output fulfills the goal function (i.e., indicating whether the attack is successful) and the perturbation adheres to the set of constraints (e.g., grammar constraint, semantic similarity constraint). A search method is used to find a sequence of transformations that produce a successful adversarial example.

This modular design enables us to easily assemble attacks from the literature while reusing components that are shared across attacks. TextAttack provides clean, readable implementations of 16 adversarial attacks from the literature. For the first time, these attacks can be benchmarked, compared, and analyzed in a standardized setting.

Proceedings of the 2020 EMNLP (Systems Demonstrations), pages 119–126
November 16-20, 2020.

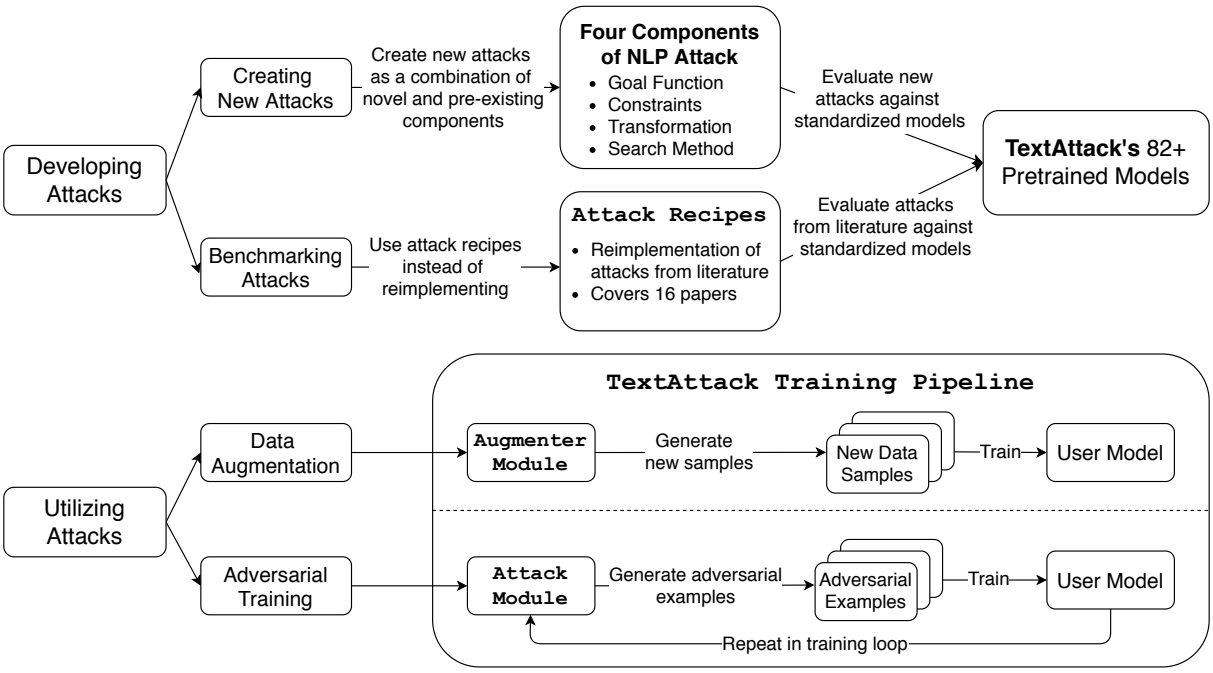

Figure 2: Main features of `TextAttack`.

`TextAttack`'s design also allows researchers to easily construct new attacks from combinations of novel and existing components. In just a few lines of code, the same search method, transformation and constraints used in Jin et al. (2019)'s `TextFooler` can be modified to attack a translation model with the goal of changing every word in the output.

`TextAttack` is directly integrated with HuggingFace's transformers and nlp libraries. This allows users to test attacks on models and datasets. `TextAttack` provides dozens of pre-trained models (LSTM, CNN, and various transformer-based models) on a variety of popular datasets. Currently `TextAttack` supports a multitude of tasks including summarization, machine translation, and all nine tasks from the GLUE benchmark. `TextAttack` also allows users to provide their own models and datasets.

Ultimately, the goal of studying adversarial attacks is to improve model performance and robustness. To that end, `TextAttack` provides easy-to-use tools for data augmentation and adversarial training. `TextAttack`'s `Augmenter` class uses a transformation and a set of constraints to produce new samples for data augmentation. Attack recipes are re-used in a training loop that allows models to train on adversarial examples. These tools make it easier to train accurate and robust models.

Uses for `TextAttack` include[1]:

- Benchmarking and comparing NLP attacks from previous works on standardized models & datasets.
- Fast development of NLP attack methods by re-using abundant available modules.
- Performing ablation studies on individual components of proposed attacks and data augmentation methods.
- Training a model (CNN, LSTM, BERT, RoBERTa, etc.) on an augmented dataset.
- Adversarial training with attacks from the literature to improve a model's robustness.

2 The `TextAttack` Framework

`TextAttack` aims to implement attacks which, given an NLP model, find a perturbation of an input sequence that satisfies the attack's goal and adheres to certain linguistic constraints. In this way, attacking an NLP model can be framed as a combinatorial search problem. The attacker must search within all potential transformations to find a sequence of transformations that generate a successful adversarial example.

Each attack can be constructed from four components:

1. A task-specific **goal function** that determines whether the attack is successful in terms of the model outputs.
   Examples: untargeted classification, targeted classification, non-overlapping output, minimum BLEU score.

[1] All can be done in < 5 lines of code. See A.1.

2. A set of **constraints** that determine if a perturbation is valid with respect to the original input.
   Examples: maximum word embedding distance, part-of-speech consistency, grammar checker, minimum sentence encoding cosine similarity.
3. A **transformation** that, given an input, generates a set of potential perturbations.
   Examples: word embedding word swap, thesaurus word swap, homoglyph character substitution.
4. A **search method** that successively queries the model and selects promising perturbations from a set of transformations.
   Examples: greedy with word importance ranking, beam search, genetic algorithm.

See A.2 for a full explanation of each goal function, constraint, transformation, and search method that's built-in to `TextAttack`.

3 Developing NLP Attacks with `TextAttack`

`TextAttack` is available as a Python package installed from PyPI, or via direct download from GitHub. `TextAttack` is also available for use through our demo web app, displayed in Figure 3.

Python users can test attacks by creating and manipulating `Attack` objects. The command-line API offers `textattack attack`, which allows users to specify attacks from their four components or from a single attack recipe and test them on different models and datasets.

`TextAttack` supports several different output formats for attack results:

- Printing results to stdout.
- Printing to a text file or CSV.
- Printing attack results to an HTML table.
- Writing a table of attack results to a visualization server, like Visdom or Weights & Biases.

3.1 Benchmarking Existing Attacks with Attack Recipes

`TextAttack`'s modular design allows us to implement many different attacks from past work in a shared framework, often by adding only one or two new components. Table 1 categorizes 16 attacks based on their goal functions, constraints, transformations and search methods.

All of these attacks are implemented as "attack recipes" in `TextAttack` and can be benchmarked with just a single command. See A.3 for a comparison between papers' reported attack results and the results achieved by running `TextAttack`.

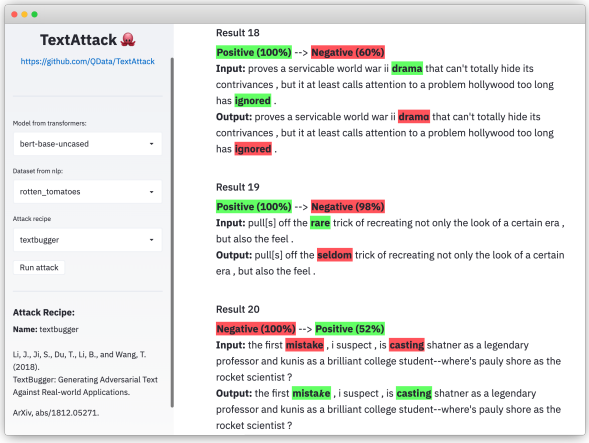

Figure 3: Screenshot of `TextAttack`'s web interface running the TextBugger black-box attack (Li et al., 2019).

3.2 Creating New Attacks by Combining Novel and Existing Components

As is clear from Table 1, many components are shared between NLP attacks. New attacks often reuse components from past work, adding one or two novel pieces. `TextAttack` allows researchers to focus on the generation of new components rather than replicating past results. For example, Jin et al. (2019) introduced `TextFooler` as a method for attacking classification and entailment models. If a researcher wished to experiment with applying `TextFooler`'s search method, transformations, and constraints to attack translation models, all they need is to implement a translation goal function in `TextAttack`. They would then be able to plug in this goal function to create a novel attack that could be used to analyze translation models.

3.3 Evaluating Attacks on `TextAttack`'s Pre-Trained Models

As of the date of this submission, `TextAttack` provides users with 82 pre-trained models, including word-level LSTM, word-level CNN, BERT, and other transformer based models pre-trained on various datasets provided by HuggingFace nlp. Since `TextAttack` is integrated with the nlp library, it can automatically load the test or validation data set for the corresponding pre-trained model. While the literature has mainly focused on classification and entailment, `TextAttack`'s pretrained models enable research on the robustness of models across all GLUE tasks.

Attack Recipe	Goal Function	Constraints	Transformation	Search Method
`bae` (Garg and Ramakrishnan, 2020)	Untargeted Classification	USE sentence encoding cosine similarity	BERT Masked Token Prediction	Greedy-WIR
`bert-attack` (Li et al., 2020)	Untargeted Classification	USE sentence encoding cosine similarity, Maximum number of words perturbed	BERT Masked Token Prediction (with subword expansion)	Greedy-WIR
`deepwordbug` (Gao et al., 2018)	{Untargeted, Targeted} Classification	Levenshtein edit distance	{Character Insertion, Character Deletion, Neighboring Character Swap, Character Substitution}*	Greedy-WIR
`alzantot`, `fast-alzantot` (Alzantot et al., 2018; Jia et al., 2019)	Untargeted {Classification, Entailment}	Percentage of words perturbed, Language Model perplexity, Word embedding distance	Counter-fitted word embedding swap	Genetic Algorithm
`iga` (Wang et al., 2019)	Untargeted {Classification, Entailment}	Percentage of words perturbed, Word embedding distance	Counter-fitted word embedding swap	Genetic Algorithm
`input-reduction` (Feng et al., 2018)	Input Reduction		Word deletion	Greedy-WIR
`kuleshov` (Kuleshov et al., 2018)	Untargeted Classification	Thought vector encoding cosine similarity, Language model similarity probability	Counter-fitted word embedding swap	Greedy word swap
`hotflip` (word swap) (Ebrahimi et al., 2017)	Untargeted Classification	Word Embedding Cosine Similarity, Part-of-speech match, Number of words perturbed	Gradient-Based Word Swap	Beam search
`morpheus` (Tan et al., 2020)	Minimum BLEU Score		Inflection Word Swap	Greedy search
`pruthi` (Pruthi et al., 2019)	Untargeted Classification	Minimum word length, Maximum number of words perturbed	{Neighboring Character Swap, Character Deletion, Character Insertion, Keyboard-Based Character Swap}*	Greedy search
`pso` (Zang et al., 2020)	Untargeted Classification		HowNet Word Swap	Particle Swarm Optimization
`pwws` (Ren et al., 2019)	Untargeted Classification		WordNet-based synonym swap	Greedy-WIR (saliency)
`seq2sick` (black-box) (Cheng et al., 2018)	Non-overlapping output		Counter-fitted word embedding swap	Greedy-WIR
`textbugger` (black-box) (Li et al., 2019)	Untargeted Classification	USE sentence encoding cosine similarity	{Character Insertion, Character Deletion, Neighboring Character Swap, Character Substitution}*	Greedy-WIR
`textfooler` (Jin et al., 2019)	Untargeted {Classification, Entailment}	Word Embedding Distance, Part-of-speech match, USE sentence encoding cosine similarity	Counter-fitted word embedding swap	Greedy-WIR

Table 1: `TextAttack` attack recipes categorized within our framework: search method, transformation, goal function, constraints. All attack recipes include an additional constraint which disallows the replacement of stopwords. Greedy search with Word Importance Ranking is abbreviated as Greedy-WIR.

* *indicates a combination of multiple transformations*

4 Utilizing `TextAttack` to Improve NLP Models

4.1 Evaluating Robustness of Custom Models

`TextAttack` is model-agnostic - meaning it can run attacks on models implemented in any deep learning framework. Model objects must be able to take a string (or list of strings) and return an output that can be processed by the goal function. For example, machine translation models take a list of strings as input and produce a list of strings as output. Classification and entailment models return an array of scores. As long as the user's model meets this specification, the model is fit to use with `TextAttack`.

4.2 Model Training

`TextAttack` users can train standard LSTM, CNN, and transformer based models, or a user-customized model on any dataset from the `nlp` library using the `textattack train` command. Just like pre-trained models, user-trained models are compatible with commands like `textattack attack` and `textattack eval`.

4.3 Data Augmentation

While searching for adversarial examples, `TextAttack`'s transformations generate perturbations of the input text, and apply constraints to verify their validity. These tools can be reused to dramatically expand the training dataset by introducing perturbed versions of existing samples. The `textattack augment` command gives users access to a number of pre-packaged recipes for augmenting their dataset. This is a stand-alone feature that can be used with any model or training framework. When using `TextAttack`'s models and training pipeline, `textattack train --augment` automatically expands the dataset before training begins. Users can specify the fraction of each input that should be modified and how many additional versions of each example to create. This makes it easy to use existing augmentation recipes on different models and datasets, and is a great way to benchmark new techniques.

Figure 4 shows empirical results we obtained using `TextAttack`'s augmentation. Augmentation with `TextAttack` immediately improves the performance of a `WordCNN` model on small datasets.

4.4 Adversarial Training

With `textattack train --attack`, attack recipes can be used to create new training sets of adversarial examples. After training for a number of epochs on the clean training set, the attack generates an adversarial version of each input. This perturbed version of the dataset is substituted for the original, and is periodically regenerated according to the model's current weaknesses. The resulting model can be significantly more robust against the attack used during training. Table 2 shows the accuracy of a standard LSTM classifier with and without adversarial training against different attack recipes implemented in `TextAttack`.

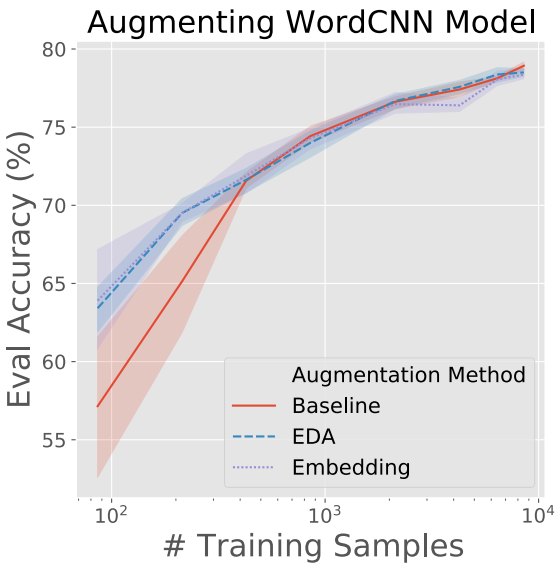

Figure 4: Performance of the built-in `WordCNN` model on the `rotten_tomatoes` dataset with increasing training set size. Data augmentation recipes like `EasyDataAugmenter` (EDA, (Wei and Zou, 2019)) and `Embedding` are most helpful when working with very few samples. Shaded regions represent 95% confidence intervals over $N = 5$ runs.

5 `TextAttack` Under the Hood

`TextAttack` is optimized under-the-hood to make implementing and running adversarial attacks simple and fast.

AttackedText. A common problem with implementations of NLP attacks is that the original text is discarded after tokenization; thus, the transformation is performed on the tokenized version of the text. This causes issues with capitalization and word segmentation. Sometimes attacks swap a piece of a word for a complete word (for example, transforming `''aren't"` into `''aren'too"`).

To solve this problem, `TextAttack` stores each input as a `AttackedText` object which contains the original text and helper methods for transforming the text while retaining tokenization. Instead of strings or tensors,

	Attacked By					
Trained Against	–	deepwordbug	textfooler	pruthi	hotflip	bae
baseline (early stopping)	**77.30%**	23.46%	2.23%	59.01%	64.57%	25.51%
deepwordbug (20 epochs)	76.38%	**35.07%**	4.78%	57.08%	65.06%	27.63%
deepwordbug (75 epochs)	73.16%	**44.74%**	13.42%	58.28%	66.87%	32.77%
textfooler (20 epochs)	61.85%	40.09%	**29.63%**	52.60%	55.75%	39.36%

Table 2: The default LSTM model trained on 3k samples from the sst2 dataset. The baseline uses early stopping on a clean training set. deepwordbug and textfooler attacks are used for adversarial training. 'Accuracy Under Attack' on the eval set is reported for several different attack types.

classes in TextAttack operate primarily on AttackedText objects. When words are added, swapped, or deleted, an AttackedText can maintain proper punctuation and capitalization. The AttackedText also contains implementations for common linguistic functions like splitting text into words, splitting text into sentences, and part-of-speech tagging.

Caching. Search methods frequently encounter the same input at different points in the search. In these cases, it is wise to pre-store values to avoid unnecessary computation. For each input examined during the attack, TextAttack caches its model output, as well as the whether or not it passed all of the constraints. For some search methods, this memoization can save a significant amount of time.[2]

6 Related Work

We draw inspiration from the Transformers library (Wolf et al., 2019) as an example of a well-designed Natural Language Processing library. Some of TextAttack's models and tokenizers are implemented using Transformers.

cleverhans (Papernot et al., 2018) is a library for constructing adversarial examples for computer vision models. Like cleverhans, we aim to provide methods that generate adversarial examples across a variety of models and datasets. In some sense, TextAttack strives to be a solution like cleverhans for the NLP community. Like cleverhans, attacks in TextAttack all implement a base Attack class. However, while cleverhans implements many disparate attacks in separate modules, TextAttack builds attacks from a library of shared components.

There are some existing open-source libraries related to adversarial examples in NLP. Trickster proposes a method for attacking NLP models based on graph search, but lacks the ability to ensure that generated examples satisfy a given constraint (Kulynych et al., 2018). TEAPOT is a library for evaluating adversarial perturbations on text, but only supports the application of ngram-based comparisons for evaluating attacks on machine translation models (Michel et al., 2019). Most recently, AllenNLP Interpret includes functionality for running adversarial attacks on NLP models, but is intended only for the purpose of interpretability, and only supports attacks via input-reduction or greedy gradient-based word swap (Wallace et al., 2019). TextAttack has a broader scope than any of these libraries: it is designed to be extendable to any NLP attack.

7 Conclusion

We presented TextAttack, an open-source framework for testing the robustness of NLP models. TextAttack defines an attack in four modules: a goal function, a list of constraints, a transformation, and a search method. This allows us to compose attacks from previous work from these modules and compare them in a shared environment. These attacks can be reused for data augmentation and adversarial training. As new attacks are developed, we will add their components to TextAttack. We hope TextAttack helps lower the barrier to entry for research into robustness and data augmentation in NLP. [3]

8 Acknowledgements

The authors would like to thank everyone who has contributed to make TextAttack a reality: Hanyu Liu, Kevin Ivey, Bill Zhang, and Alan Zheng, to name a few. Thanks to the IGA creators (Wang et al., 2019) for contributing an implementation of their algorithm to our framework. Thanks to the folks at HuggingFace for creating such easy-to-use software; without them, TextAttack would not be what it is today.

[2] Caching alone speeds up the genetic algorithm of Alzantot et al. (2018) by a factor of 5.

[3] For more information, an appendix is available online here.

References

Abhaya Agarwal and Alon Lavie. 2008. Meteor, m-bleu and m-ter: Evaluation metrics for high-correlation with human rankings of machine translation output. In *WMT@ACL*.

Moustafa Alzantot, Yash Sharma, Ahmed Elgohary, Bo-Jhang Ho, Mani B. Srivastava, and Kai-Wei Chang. 2018. Generating natural language adversarial examples. *ArXiv*, abs/1804.07998.

Daniel Matthew Cer, Yinfei Yang, Sheng yi Kong, Nan Hua, Nicole Limtiaco, Rhomni St. John, Noah Constant, Mario Guajardo-Cespedes, Steve Yuan, Chris Tar, Yun-Hsuan Sung, Brian Strope, and Ray Kurzweil. 2018. Universal sentence encoder. *ArXiv*, abs/1803.11175.

Minhao Cheng, Jinfeng Yi, Pin-Yu Chen, Huan Zhang, and Cho-Jui Hsieh. 2018. Seq2sick: Evaluating the robustness of sequence-to-sequence models with adversarial examples.

Alexis Conneau, Douwe Kiela, Holger Schwenk, Loïc Barrault, and Antoine Bordes. 2017. Supervised learning of universal sentence representations from natural language inference data. In *EMNLP*.

Zhendong Dong, Qiang Dong, and Changling Hao. 2006. Hownet and the computation of meaning.

Javid Ebrahimi, Anyi Rao, Daniel Lowd, and Dejing Dou. 2017. Hotflip: White-box adversarial examples for text classification. In *ACL*.

Shi Feng, Eric Wallace, Alvin Grissom II, Mohit Iyyer, Pedro Rodriguez, and Jordan Boyd-Graber. 2018. Pathologies of neural models make interpretations difficult.

Ji Gao, Jack Lanchantin, Mary Lou Soffa, and Yanjun Qi. 2018. Black-box generation of adversarial text sequences to evade deep learning classifiers. *2018 IEEE Security and Privacy Workshops (SPW)*, pages 50–56.

Siddhant Garg and Goutham Ramakrishnan. 2020. Bae: Bert-based adversarial examples for text classification.

Ari Holtzman, Jan Buys, Maxwell Forbes, Antoine Bosselut, David Golub, and Yejin Choi. 2018. Learning to write with cooperative discriminators. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1638–1649, Melbourne, Australia. Association for Computational Linguistics.

Robin Jia and Percy Liang. 2017. Adversarial examples for evaluating reading comprehension systems. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2021–2031, Copenhagen, Denmark. Association for Computational Linguistics.

Robin Jia, Aditi Raghunathan, Kerem Göksel, and Percy Liang. 2019. Certified robustness to adversarial word substitutions. In *EMNLP/IJCNLP*.

Di Jin, Zhijing Jin, Joey Tianyi Zhou, and Peter Szolovits. 2019. Is bert really robust? natural language attack on text classification and entailment. *ArXiv*, abs/1907.11932.

Rafal Józefowicz, Oriol Vinyals, Mike Schuster, Noam Shazeer, and Yonghui Wu. 2016. Exploring the limits of language modeling. *ArXiv*, abs/1602.02410.

James Kennedy and Russell Eberhart. 1995. Particle swarm optimization. In *Proceedings of ICNN'95-International Conference on Neural Networks*, volume 4, pages 1942–1948. IEEE.

Ryan Kiros, Yukun Zhu, Ruslan Salakhutdinov, Richard S. Zemel, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. Skip-thought vectors. *ArXiv*, abs/1506.06726.

Volodymyr Kuleshov, Shantanu Thakoor, Tingfung Lau, and Stefano Ermon. 2018. Adversarial examples for natural language classification problems.

Bogdan Kulynych, Jamie Hayes, Nikita Samarin, and Carmela Troncoso. 2018. Evading classifiers in discrete domains with provable optimality guarantees. *CoRR*, abs/1810.10939.

Jinfeng Li, Shouling Ji, Tianyu Du, Bo Li, and Ting Wang. 2019. Textbugger: Generating adversarial text against real-world applications. *ArXiv*, abs/1812.05271.

Linyang Li, Ruotian Ma, Qipeng Guo, Xiangyang Xue, and Xipeng Qiu. 2020. Bert-attack: Adversarial attack against bert using bert.

Paul Michel, Xian Li, Graham Neubig, and Juan Miguel Pino. 2019. On evaluation of adversarial perturbations for sequence-to-sequence models. *CoRR*, abs/1903.06620.

George Armitage Miller, Richard Beckwith, Christiane Fellbaum, Derek Gross, and Katherine J. Miller. 1990. Introduction to wordnet: An on-line lexical database. *International Journal of Lexicography*, 3:235–244.

Nikola Mrkšić, Diarmuid O Séaghdha, Blaise Thomson, Milica Gašić, Lina Rojas-Barahona, Pei-Hao Su, David Vandyke, Tsung-Hsien Wen, and Steve Young. 2016. Counter-fitting word vectors to linguistic constraints. *arXiv preprint arXiv:1603.00892*.

Daniel Naber et al. 2003. *A rule-based style and grammar checker*. Citeseer.

Nicolas Papernot, Fartash Faghri, Nicholas Carlini, Ian Goodfellow, Reuben Feinman, Alexey Kurakin, Cihang Xie, Yash Sharma, Tom Brown, Aurko Roy,

Alexander Matyasko, Vahid Behzadan, Karen Hambardzumyan, Zhishuai Zhang, Yi-Lin Juang, Zhi Li, Ryan Sheatsley, Abhibhav Garg, Jonathan Uesato, Willi Gierke, Yinpeng Dong, David Berthelot, Paul Hendricks, Jonas Rauber, and Rujun Long. 2018. Technical report on the cleverhans v2.1.0 adversarial examples library. *arXiv preprint arXiv:1610.00768*.

Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2001. Bleu: a method for automatic evaluation of machine translation. In *ACL*.

Maja Popovic. 2015. chrf: character n-gram f-score for automatic mt evaluation. In *WMT@EMNLP*.

Danish Pruthi, Bhuwan Dhingra, and Zachary C. Lipton. 2019. Combating adversarial misspellings with robust word recognition.

Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.

Nils Reimers and Iryna Gurevych. 2019. Sentence-bert: Sentence embeddings using siamese bert-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.

Shuhuai Ren, Yihe Deng, Kun He, and Wanxiang Che. 2019. Generating natural language adversarial examples through probability weighted word saliency. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1085–1097, Florence, Italy. Association for Computational Linguistics.

Samson Tan, Shafiq Joty, Min-Yen Kan, and Richard Socher. 2020. It's morphin' time! Combating linguistic discrimination with inflectional perturbations. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2920–2935, Online. Association for Computational Linguistics.

Eric Wallace, Jens Tuyls, Junlin Wang, Sanjay Subramanian, Matthew Gardner, and Sameer Singh. 2019. Allennlp interpret: A framework for explaining predictions of nlp models. *ArXiv*, abs/1909.09251.

Xiaosen Wang, Hao Jin, and Kun He. 2019. Natural language adversarial attacks and defenses in word level.

Jason W. Wei and Kai Zou. 2019. EDA: easy data augmentation techniques for boosting performance on text classification tasks. *CoRR*, abs/1901.11196.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, R'emi Louf, Morgan Funtowicz, and Jamie Brew. 2019. Transformers: State-of-the-art natural language processing.

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

Tianyi Zhang*, Varsha Kishore*, Felix Wu*, Kilian Q. Weinberger, and Yoav Artzi. 2020. Bertscore: Evaluating text generation with bert. In *International Conference on Learning Representations*.