

Achieving Model Robustness through Discrete Adversarial Training

Maor Ivgi

Tel-Aviv University

maorivgi@mail.tau.ac.il

Jonathan Berant

Tel-Aviv University

The Allen Institute for AI

joberant@cs.tau.ac.il

Abstract

Discrete adversarial attacks are symbolic perturbations to a language input that preserve the output label but lead to a prediction error. While such attacks have been extensively explored for the purpose of *evaluating* model robustness, their utility for *improving robustness* has been limited to offline augmentation only. Concretely, given a trained model, attacks are used to generate perturbed (adversarial) examples, and the model is re-trained exactly once. In this work, we address this gap and leverage discrete attacks for *online augmentation*, where adversarial examples are generated at every training step, adapting to the changing nature of the model. We propose (i) a new discrete attack, based on best-first search, and (ii) random sampling attacks that unlike prior work are not based on expensive search-based procedures. Surprisingly, we find that random sampling leads to impressive gains in robustness, outperforming the commonly-used offline augmentation, while leading to a speedup at training time of $\sim 10x$. Furthermore, online augmentation with search-based attacks justifies the higher training cost, significantly improving robustness on three datasets. Last, we show that our new attack substantially improves robustness compared to prior methods.

1 Introduction

Adversarial examples are inputs that are slightly, but intentionally, perturbed to create a new example that is misclassified by a model (Szegedy et al., 2014). Adversarial examples have attracted immense attention in machine learning (Goodfellow et al., 2015; Carlini and Wagner, 2017; Papernot et al., 2017) for two important, but separate, reasons. First, they are useful for *evaluating* model robustness, and have revealed that current models are over-sensitive to minor perturbations. Second, adversarial examples can *improve* robustness: training on adversarial examples reduces the brittleness and over-sensitivity of deep learning models to




Figure 1: Robust accuracy vs. slowdown in training time, comparing different methods to Baseline (purple pentagon); x-axis in logarithmic scale. The popular ADVOFF (blue squares, offline augmentation with adversarial example) is 10x slower than our simple augmentation of 4 (8) random samples (triangles, RANDOFF-4, RANDOFF-8) and achieves similar or worse robust accuracy. Our online augmentation of adversarial examples (ADVON, yellow circles) significantly improves robust accuracy, but is expensive to train.

such perturbations (Alzantot et al., 2018; Jin et al., 2020; Li et al., 2020; Lei et al., 2019; Wallace et al., 2019; Zhang et al., 2020; Garg and Ramakrishnan, 2020; Si et al., 2020a; Goel et al., 2021).

Training and evaluating models with adversarial examples has had considerable success in computer vision, with gradient-based techniques like FGSM (Goodfellow et al., 2015) and PGD (Madry et al., 2018). In computer vision, adversarial examples can be constructed by considering a continuous space of imperceptible perturbations around image pixels. Conversely, language is discrete, and any perturbation is perceptible. Thus, robust models must be invariant to input modifications that preserve semantics, such as synonym substitutions (Alzantot et al., 2018; Jin et al., 2020), paraphrasing (Tan et al., 2020), or typos (Huang et al., 2019).

Due to this property of language, ample work has been dedicated to developing discrete attacks that generate adversarial examples through combinatorial optimization (Alzantot et al., 2018; Ren et al., 2019; Jin et al., 2020; Zhou et al., 2020; Zang et al., 2020). For example, in sentiment analysis, it is common to consider the space of all *synonym substitutions*, where an adversarial example for an input “Such an amazing movie!” might be “Such an extraordinary film” (Fig. 2). This body of work has mostly focused on *evaluating* robustness, rather than *improving it*, which naturally led to the development of complex combinatorial search algorithms, whose goal is to find adversarial examples in the exponential space of perturbations.

In this work, we address a major research gap in current literature around *improving* robustness with discrete attacks. Specifically, past work (Alzantot et al., 2018; Ren et al., 2019; Jin et al., 2020) only considered *offline augmentation*, where a discrete attack is used to generate adversarial examples and the model is re-trained exactly once with those examples. This ignores *online augmentation*, which had success in computer vision (Kurakin et al., 2017; Perez and Wang, 2017; Madry et al., 2018), where adversarial examples are generated in each training step, adapting to the changing model. Moreover, simple data augmentation techniques, such as randomly sampling from the space of synonym substitutions and adding the generated samples to the training data have not been investigated and compared to offline adversarial augmentation. We address this lacuna and systematically compare online augmentation to offline augmentation, as well as to simple random sampling techniques. To our knowledge, we are the first to evaluate *online augmentation* with discrete attacks on a wide range of NLP tasks. Our results show that online augmentation leads to significant improvement in robustness compared to prior work and that simple random augmentation achieves comparable results to the common offline augmentation at a fraction of the complexity and training time.

Moreover, we present a new search algorithm for finding adversarial examples, *Best-First search over a Factorized graph (BFF)*, which alleviates the greedy nature of previously-proposed algorithms. BFF improves search by incorporating backtracking, and allowing to re-visit previously-discarded search paths, once the current one is revealed to be sub-optimal.




Figure 2: Given a movie review x , the model A is robust to a set of perturbations, while A' is not.

We evaluate model robustness on three datasets: BoolQ (Clark et al., 2019), IMDB (Maas et al., 2011), and SST-2 (Socher et al., 2013), which vary in terms of the target task (question answering and sentiment analysis) and input length. Surprisingly, we find across different tasks (Fig. 1) that augmenting each training example with 4-8 random samples from the synonym substitution space performs as well as (or better than) the commonly used offline augmentation, while being simpler and 10x faster to train. Conversely, online augmentation makes better use of the extra computational cost, and substantially improves robust accuracy compared to offline augmentation. Additionally, our proposed discrete attack algorithm, BFF, outperforms prior work by a wide margin. Our data and code are available at https://github.com/Mivg/robust_transformers.

2 Problem Setup and Background

Problem setup We focus in this work on the supervised classification setup, where given a training set $\{x_j, y_j\}_{j=1}^N$ sampled from $\mathcal{X} \times \mathcal{Y}$, our goal is to learn a mapping $A : \mathcal{X} \rightarrow \mathcal{Y}$ that achieves high accuracy on held-out data sampled from the same distribution. Moreover, we want the model A to be *robust*, i.e., invariant to a set of pre-defined label-preserving perturbations to x , such as synonym substitutions. Formally, for any natural language input x , a discrete *attack space* of label-preserving perturbations $\mathcal{S}(x) \subset \mathcal{X}$ is defined. Given a labeled example (x, y) , a model A is robust w.r.t x , if $A(x) = y$ and for any $\bar{x} \in \mathcal{S}(x)$, the output $A(\bar{x}) = A(x)$. An example $\bar{x} \in \mathcal{S}(x)$ such that $A(\bar{x}) \neq A(x)$ is called an *adversarial example*. We assume A provides not only a prediction but a distribution $p_A(x) \in \Delta^{|\mathcal{Y}|}$ over the possible classes, where Δ is the simplex, and denote the probability A assigns to the gold label by $[p_A(x)]_y$. Fig. 2 shows an example from sentiment analysis,

where a model A is robust, while A' is not w.r.t x .

Robustness is evaluated with *robust accuracy* (Tsipras et al., 2019), i.e., the fraction of examples a model is robust to over some held-out data. Typically, the size of the attack space $\mathcal{S}(x)$ is exponential in the size of x and it is not feasible to enumerate all perturbations. Instead, an upper bound is estimated by searching for a set of adversarial attacks, i.e., “hard” examples in $\mathcal{S}(x)$ for every x , and estimating robust accuracy w.r.t to that set.

Improving robustness with discrete attacks

Since language is discrete, a typical approach for *evaluating* robustness is to use combinatorial optimization methods to search for adversarial examples in the attack space $\mathcal{S}(x)$. This has been repeatedly shown to be an effective attack method on pre-trained models (Alzantot et al., 2018; Lei et al., 2019; Ren et al., 2019; Li et al., 2020; Jin et al., 2020; Zang et al., 2020). However, in terms of *improving* robustness, discrete attacks have thus far been mostly used with offline augmentation (defined below) and have led to limited robustness gains. In this work, we examine the more costly but potentially more beneficial online augmentation.

Offline vs. online augmentation Data augmentation is a common approach for improving generalization and robustness, where variants of training examples are automatically generated and added to the training data (Simard et al., 1998). Here, discrete attacks can be used to generate these examples. We consider both *offline* and *online* data augmentation and focus on improving robustness with adversarial examples.

Given a training set $\{(x_j, y_j)\}_{j=1}^N$, *offline data augmentation* involves (a) training a model A over the training data, (b) for each training example (x_j, y_j) , generating a perturbation w.r.t to A (using some discrete attack) and labeling it with y_j , and (c) training a new model over the union of the original training set and the generated examples. This is termed *offline* augmentation because examples are generated with respect to a fixed model A .

Online data augmentation is this setup, examples are generated at training time w.r.t the current model A . This is more computationally expensive, as examples must be generated during training and not as pre-processing, but examples can adapt to the model over time. In each step, half the batch contains examples from the training set, and half are adversarial examples generated by some dis-

crete attack w.r.t to the model’s current state.

Online augmentation has been used to improve robustness in NLP with gradient-based approaches (Jia et al., 2019; Shi et al., 2020; Zhou et al., 2020), but to the best of our knowledge has been overlooked in the context of discrete attacks. In this work, we are the first to propose model-agnostic online augmentation training, which uses automatically generated *discrete adversarial attacks* to boost overall robustness in NLP models.

3 The Attack Space

An attack space for an input with respect to a classification task can be intuitively defined as the set of label-preserving perturbations over the input. A popular attack space $\mathcal{S}(x)$, which we adopt, is the space of *synonym substitutions* (Alzantot et al., 2018; Ren et al., 2019). Given a synonym dictionary that provides a set of synonyms $Syn(w)$ for any word w , the attack space $\mathcal{S}_{syn}(x)$ for an utterance $x = (w_1, \dots, w_n)$ contains all utterances that can be obtained by replacing a word w_i (and possibly multiple words) with one of their synonyms. Typically, the number of words from x allowed to be substituted is limited to be no more than $D = \lceil d \cdot |x| \rceil$, where $d \in \{0.1, 0.2\}$ is a common choice.

Synonym substitutions are context-sensitive, i.e., substitutions might only be appropriate in certain contexts. For example, in Fig. 3, replacing the word “like” with its synonym “similar” (red box) is invalid, since “like” is a verb in this context. Consequently, past work (Ren et al., 2019; Jin et al., 2020) filtered $\mathcal{S}_{syn}(x)$ using a context-sensitive filtering function $\Phi_x(w_i, \bar{w}_i) \in \{0, 1\}$, which determines whether substituting a word w_i from the original utterance x with its synonym \bar{w}_i is valid in a particular context. For instance, an external model can check whether the substitution maintains the part-of-speech, and whether the overall semantics is maintained. We define the *filtered synonyms substitutions space* $\mathcal{S}_\Phi(x)$ as the set that includes all utterances \bar{x} that can be generated through a sequence of no more than D single-word substitutions from the original utterance that are valid according to $\Phi(\cdot, \cdot)$. In §5.2, we describe the details of the synonym and function Φ .

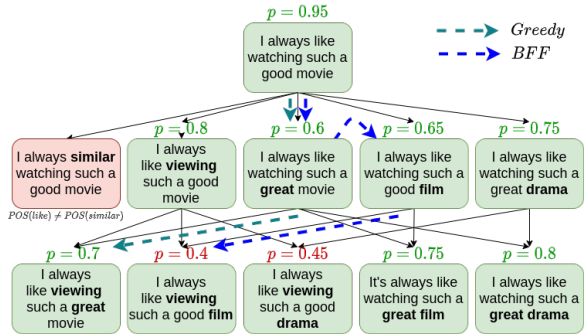


Figure 3: Example of an attack space, and the paths taken by a greedy algorithm and best-first search. An adversarial example has a probability $p < 0.5$ for the gold positive label.

4 Best-first Search Over a Factorized Graph

Searching over the attack space $\mathcal{S}_\Phi(x)$ can be naturally viewed as a search problem over a directed acyclic graph (DAG), $G = (\mathcal{U}, \mathcal{E})$, where each node $u_{\bar{x}} \in \mathcal{U}$ is labeled by an utterance \bar{x} , and edges \mathcal{E} correspond to single-word substitutions, valid according to $\Phi(\cdot)$. The graph is directed and acyclic, since only substitutions of words from the original utterance x are allowed (see Fig. 3). Because there is a one-to-one mapping from the node $u_{\bar{x}}$ to the utterance \bar{x} , we will use the latter to denote both the node and the utterance.

Discrete attacks use search algorithms to find an adversarial example in $\mathcal{S}(x)$. The search is guided by a heuristic scoring function $s_A(x) := [p_A(x)]_y$, where the underlying assumption is that utterances that give lower probability to the gold label are closer to an adversarial example. A popular choice for a search algorithm in NLP is greedy search, illustrated in Fig. 3. Specifically, one holds in step t the current node x_t , where t words have been substituted in the source node $x_0 = x$. Then, the model $A(\cdot)$ is run on the *frontier*, that is, all out-neighbor nodes $\mathcal{N}(x_t) = \{\hat{x}_{t+1} \mid (x_t, \hat{x}_{t+1}) \in \mathcal{E}\}$, and the one that minimizes the heuristic scoring function is selected: $x_{t+1} := \operatorname{argmin}_{\hat{x} \in \mathcal{N}(x_t)} s_A(\hat{x})$.

While greedy search has been used for character-flipping (Ebrahimi et al., 2018), it is ill-suited in the space of synonym substitutions. The degree of nodes is high – assuming n_{rep} words can be replaced in the text, each with K possible synonyms, then the out degree is $O(n_{\text{rep}} \cdot K)$. This results in an infeasible number of forward passes through the attacked model even for a small number of search

iterations.

To enable effective search through the search space, we (a) factorize the graph such that the out-degree of nodes is lower, and (b) use a best-first search algorithm. We describe those next.

Graph factorization To reduce the out-degree of a node in the search space and thus improve its efficiency, we can split each step into two. First, choose a position to substitute in the utterance; Second, choose a substitution for that position. This reduces the number of evaluations of A per step from $O(n_{\text{rep}} \cdot K)$ to $O(n_{\text{rep}} + K)$. To estimate the score of a position i , one can mask the word w_i with a mask token τ and measure $s_A(x_{w_i \rightarrow \tau})$, where $x_{w_i \rightarrow \tau}$ is the utterance x where the word in position i is replaced by the mask τ .

We can describe this approach as search over a bi-partite DAG $\hat{G} = (\mathcal{U} \cup \mathcal{W}, \hat{\mathcal{E}})$. The nodes \mathcal{U} are utterances like in G , and the new nodes are utterances with a single mask token $\mathcal{W} = \{\bar{x}_{w_i \rightarrow \tau} \mid \bar{x} \in \mathcal{S}(x) \wedge w_i \text{ is a word in } x\}$. The edges comprise two types: $\hat{\mathcal{E}} = \mathcal{E}_1 \cup \mathcal{E}_2$. The edges \mathcal{E}_1 are from utterances to masked utterances: $\mathcal{E}_1 = \{(\bar{x}, \bar{x}_{w_i \rightarrow \tau})\} \subset \mathcal{U} \times \mathcal{W}$, and $\mathcal{E}_2 = \{(\bar{x}_{w_i \rightarrow \tau}, \bar{x}_{w_i \rightarrow w_{\text{syn}}})\} \subset \mathcal{W} \times \mathcal{U}$, where $w_{\text{syn}} \in \text{Syn}(w_i)$. In Figure 3, the two right-most nodes in each row would be factorized together as they substitute the same word, and the algorithm will evaluate only one of them to estimate the potential benefit of substituting “movie”.

Best-first search A factorized graph makes search possible by reducing the out-degree of nodes. However, greedy search is still sub-optimal. This is since it relies on the heuristic search function to be a good estimate of the distance to an adversarial example, which can often be false. Consider the example in Fig. 3. The two adversarial examples (with $p = 0.4$ or $p = 0.45$) are not reachable from the best node after the first step ($p = 0.6$), only from the second-best ($p = 0.65$).

Best-first search (Pearl, 1984) overcomes this at a negligible cost, by holding a min-heap over the nodes of the frontier of the search space (Alg. 1). In each step, we pop the next utterance, which assigns the lowest probability to the gold label, and push all neighbors into the heap. When a promising branch turns out to be sub-optimal, search can resume from an earlier node to find a better solution, as shown in the blue path in Figure 3. To bound the cost of finding a single adversarial example, we bound the number of forward passes through the model

A with a budget parameter B . To further reduce “greedyness”, search can use a beam by popping more than one node in each step, expanding all their neighbors and pushing the result back to the heap. Our final approach uses **Best-First** search over a **Factorized** graph, and is termed **BFF**.

Algorithm 1: BFF

```

input :model  $A$ , factorized graph  $G$ , utterance  $x$ .
heap  $\leftarrow \{(x, s_A(X))\}$ 
 $x^* \leftarrow x$ 
while |heap| > 0 and budget  $B$  not exhausted:
     $\hat{x} \leftarrow \text{heap.pop}()$ 
     $x^* \leftarrow \text{argmin}_{\hat{x} \in \{\hat{x}, x^*\}} A(\hat{x})$ 
    if  $A(x^*) \neq y$  break;
    for  $\hat{x} \in \mathcal{N}(\hat{x})$  do
        | heap.push( $\hat{x}, s_A(\hat{x})$ )
return  $x^*$ 

```

5 Experiments

We conduct a thorough empirical evaluation of model robustness across a wide range of attacks and training procedures.

5.1 Experimental Setup

To evaluate our approach over diverse settings, we consider three different tasks: *text classification*, *sentiment analysis* and *question answering*, two of which contain long passages that result in a large attack space (see Table 1).

1. **SST-2**: Based on the the Stanford sentiment treebank (Socher et al., 2013), SST-2 is a binary (positive/negative) classification task containing 11,855 sentences describing movie reviews. SST-2 has been frequently used for evaluating robustness.
2. **IMDB** (Maas et al., 2011): A binary (positive/negative) text classification task, containing 50K reviews from IMDB. Here, passages are long and thus the attack space is large (Table 1).
3. **BoolQ** (Clark et al., 2019): contains 16,000 yes/no questions over Wikipedia paragraphs. This task is perhaps the most interesting, because the attack space is large and answering requires global passage understanding. We allow word substitutions in the paragraph only and do not substitute nouns, verbs, or adjectives that appear in the question to avoid non-label-preserving perturbations. Further details can be found in App. A.2.

Models We consider a wide array of models and evaluate both their downstream accuracy and ro-

bustness. In all models, we define a budget of $B = 1000$, which specifies the maximal number of allowed forward passes through the model for finding an adversarial example. All results are an average of 3 runs.

To demonstrate the effectiveness of BFF for both robustness evaluation as well as adversarial training, we compare it to a recent state-of-the-art discrete attack, TEXTFOOLER (Jin et al., 2020), which we denote in model names below by the prefix TxF. The models compared are:

- **BASELINE**: we fine-tune a pretrained language model on the training set. We use BERT-BASE (Devlin et al., 2019) for IMDB/SST-2 and ROBERTA-LARGE (Liu et al., 2019) for BoolQ. These baselines are on par with current state-of-the-art to demonstrate the efficacy of our method.
- **BFFOFF/TXF OFF** Offline augmentation with the BFF or TEXTFOOLER attacks.
- **BFFON/TXF ON** Online augmentation with the BFF or TEXTFOOLER attacks.
- **RAND OFF-L**: We compare search-based algorithms to a simple and efficient approach that does not require any forward passes through the model A . Specifically, we randomly sample L utterances from the attack space for each example (without executing A) and add them to the training data.
- **RAND ON**: A random sampling approach that does use the model A . Here, we sample B random utterances, pass them through A , and return the attack that resulted in lowest model probability.
- **FREELB**: For completeness, we also consider FREELB (Zhu et al., 2020), a popular gradient-based approach for improving robustness, which employs *virtual adversarial training* (see §6). This approach uses online augmentation, where examples are created by taking gradient steps w.r.t the input embeddings to maximize the model’s loss. Other gradient-based approaches (e.g., certified robustness) are not suitable when using pre-trained transformers, which we further discuss in §6.

In a parallel line of work, Garg and Ramakrishnan (2020) and Li et al. (2020) used pre-trained language models to both *define* an attack space and to *generate* high-fidelity attacks in that space. While successful, these approaches are not suitable for our setting, due to the strong coupling between the attack strategy and the attack space itself. We further discuss this in §6

Evaluation We evaluate models on their downstream accuracy, as well as on robust accuracy, i.e. the fraction of examples against which the model is robust. Since exact robust accuracy is intractable to compute due to the exponential size of the attack space, we compute an upper-bound by attacking each example with both BFF and TEXTFOOLER (TXF) with a budget of $B = 2000$. An example is robust if we cannot find an utterance where the prediction is different from the gold label. We evaluate robust accuracy on 1000/1000/872 samples from the development sets of BoolQ/IMDB/SST-2.

5.2 Attack Space

Despite the myriad of works on discrete attacks, an attack space for synonym substitutions has not been standardized. While all past work employed a synonym dictionary combined with a $\Phi(\cdot, \cdot)$ filtering function (see §3), the particular filtering functions vary. When examining the attack space proposed in TXF, we observed that attacks result in examples that are difficult to understand or are not label-preserving. Table 6 in App. A.4 shows several examples. For instance, in sentiment classification, the attack replaced “*compelling*” with “*unconvincing*” in the sentence “*it proves quite unconvincing as an intense, brooding character study*” which alters the meaning and the sentiment of the sentence. Therefore, we use a more strict definition of the filtering function and conduct a user study to verify it is label-preserving.

Concretely, we use the synonym dictionary from Alzantot et al. (2018). We determine if a word substitution is context-appropriate by computing all single-word substitutions ($n_{\text{rep}} \cdot K$) and disallowing those that change the POS tag according to spaCy (Honnibal et al.) or increase perplexity according to GPT-2 (Radford et al., 2019) by more than 25%. Similar to Jin et al. (2020), we also filter out synonyms that are not semantics-preserving according to the USE (Cer et al., 2018) model. The attack space includes any combination of allowed single-word substitutions, where the fraction of allowed substitutions is $d = 0.1$. Implementation details are in App. A.2. We find that this ensemble of models reduces the number of substitutions that do not preserve semantics and are allowed by the filtering function.

We check the validity of our more restrictive attack space with a user study, where we verify that our attack space is indeed label-preserving. The

	$ x $	n_{rep}	$ Syn(w) $	$ S_{\phi}(x) $
SST-2	8.9	2.7	2.4	27.7
IMDB	242.4	97.3	3.6	2.27×10^{64}
BoolQ [†]	97.7	38.7	3.6	3.64×10^{25}

Table 1: Statistics on datasets and the size of attack space. We show the average number of words per utterance $|x|$, the average number of words with substitutions n_{rep} , average number of synonyms per replaceable word, and an estimation of the attack space size.

details of the user study are in §5.6.

5.3 Robustness Results

Table 2 shows accuracy on the development set, robust accuracy, and slowdown compared to BASELINE for all models and datasets. For downstream accuracy, training for robustness either maintains or slightly increases downstream accuracy. This is not the focus of this work, but is indeed a nice side-effect. For robust accuracy, discrete attacks substantially improve robustness: 80.5 \rightarrow 85.3 on SST-2, 41.2 \rightarrow 78.9 on IMDB, and 50.0 \rightarrow 68.7 on BoolQ, closing roughly half the gap from downstream accuracy.

Comparing different attacks, online augmentation (BFFON), which has been overlooked in the context of discrete attacks, leads to dramatic robustness gains compared to other methods, but is slow to train – 20-270x slower than BASELINE. This shows the importance of continuous adaptation to the current vulnerabilities of the model.

Interestingly, adding offline random samples (RANDOFF- L) consistently improves robust accuracy, and using $L = 12$ leads to impressive robustness gains without executing A at all, outperforming BFFOFF in robust accuracy, and being $\sim 5x$ faster on IMDB and BoolQ. Moreover, random sampling is trivial to implement, and independent from the attack strategy. Hence, the common practice of using offline augmentation with search-based attacks, such as BFFOFF, seems misguided, and a better solution is to use random sampling. Online random augmentation obtains impressive results, not far from BFFON, without applying any search procedure, but is very slow, since it uses the entire budget B in every example.

Comparing BFF to TXF, we observe that BFF, which uses best-first search, outperforms TXF in both the online and offline setting. Last FREELB, which is based on virtual adversarial training, improves robust accuracy at a low computational cost, but is dramatically outperformed by discrete search-

Model	Accuracy			Robust Accuracy			Slowdown		
	SST-2	IMDB	BoolQ	SST-2	IMDB	BoolQ	SST-2	IMDB	BoolQ
Baseline	91.9	93.4	84.5	80.5	41.2	50.0	×1	×1	×1
FREELB	92.5	93.9	85.5	82.1	62.5	55.8	×1.8	×1.8	×3.9
RANDOFF-1	91.9	93.5	85.6	83.5	50.3	52.2	×1.9	×1.5	×2.1
RANDOFF-4	91.6	93.7	85.5	83.6	57.0	58.4	×3.8	×4.5	×5.1
RANDOFF-8	91.1	93.8	86.1	83.3	60.9	61.3	×5.4	×8.0	×9.3
RANDOFF-12	91.5	93.7	85.8	84.2	60.1	63.0	×6.3	×11.5	×13.2
TxFOFF	91.2	93.4	86.5	83.5	49.0	61.5	×3.0	×56.1	×8.6
BFFOFF	91.8	93.7	85.8	84.6	54.3	62.3	×5.4	×60.0	×63.2
RANDON	91.7	94.1	85.6	84.9	68.5	66.0	×14.8	×249.3	×280.4
TxFON	91.3	93.8	86.0	84.0	67.4	65.3	×3.9	×58.0	×28.1
BFFON	91.7	94.2	86.5	85.3	78.9	68.7	×21.1	×270.7	×215.9

Table 2: Accuracy on the evaluation set, robust accuracy, and slowdown in model training for all datasets.

Model	IMDB				BoolQ			
	Rand	TxF	BFF	Gen	Rand	TxF	BFF	Gen
Baseline	73.1	70.2	49.9	54.1	62.1	67.7	50.2	52.0
RND-OA	74.8	74.7	52.9	59.1	70.9	72.0	59.4	62.0
TxFOFF	67.7	77.5	52.5	56.7	71.0	75.0	61.5	63.4
BFFOFF	75.4	76.9	58.6	64.1	70.9	74.8	64.7	65.2
RANDON	87.0	76.4	68.5	79.6	71.5	72.6	60.1	67.5
TxFON	81.1	84.2	69.7	73.7	73.4	74.8	65.3	67.4
BFFON	87.0	84.9	79.0	81.9	75.1	76.1	69.0	70.3

Table 3: Robust accuracy of different robust models w.r.t particular discrete attacks. RND-OA is offline augmentation with a random attack and $B = 1000$. *Gen* is our implementation of the Genetic Attack by Alzantot et al. (2018).

based attacks, including BFF.

To summarize, random sampling leads to significant robustness gains at a small cost, outperforming the commonly used offline augmentation. Online augmentation leads to the best robustness, but is more expensive to train.

5.4 Robustness across Attack Strategies

A natural question is whether a model trained for robustness with an attack (e.g., BFF) is robust w.r.t to examples generated by other attacks, which are potentially uncorrelated with them. To answer that, we evaluate the robustness of our models to attacks generated by BFF, TxF, and random sampling. Moreover, we evaluate robustness to a genetic attack, which should not be correlated with BFF and TxF: we re-implement the genetic attack algorithm from Alzantot et al. (2018) (details in A.3), and examine the robustness of our model to this attack. All attacks are with a budget of $B = 2000$.

Table 3 shows the result of this evaluation. We observe that BFFON obtains the highest robust accuracy results w.r.t to all attacks: BFF, TxF, random sampling, and a genetic attack. In offline

augmentation, we observe again that BFFOFF obtains good robust accuracy, higher or comparable to all other offline models for any attack strategy. This result highlights the generality of BFF for improving model robustness.

5.5 Success Rate Results

To compare the different attacks proposed in §4, we analyze the *success rate* against BASELINE, i.e., the proportion of examples for which an attack finds an adversarial example as a function of the budget B .

Fig. 4 compares the success rate of different attacks. We observe that BFF-based attacks have the highest success rate after a few hundred executions. TEXTFOOLER performs well at first, finding adversarial examples for many examples, but then its success plateaus. Similarly, a random approach, which ignores the graph structure, starts with a relatively high success rate, as it explores far regions in the graph, but fails to properly utilize its budget and then falls behind.

BFF combines backtracking with graph factorization. When removing backtracking, i.e., greedy search over the factorized graph, success rate decreases, especially in BoolQ. Greedy search without graph factorization leads to a low success rate due to the large number of neighbors of each node, which quickly exhausts the budget. Moreover, looking at BFF with beam size 2 (popping 2 items from the heap in each step) leads to lower performance when the budget $B \leq 2000$, as executions are expended on less promising utterances, but could improve success rate given a larger budget.

Lastly, due to our more strict definition of the attack space, described in (§5.2), success rates of BFF and TxF are lower compared to Jin et al. (2020). To verify the correctness of our attacks,




Figure 4: Success rate of different attacks against BoolQ/IMDB BASELINE as a function of the budget.

	Original	Random	BFF
IMDB	98.0	98.0	96.0
BoolQ	89.0	91.5	83.5
SST-2	97.0	96.0	94.4

Table 4: Evaluating attack space validity. We show human performance on original examples, random examples, and examples generated with BFF.

we run BFF and TxF in their attack space, which uses a larger synonym dictionary, a more permissive function Φ , and does not limit the number of substitutions D and budget B . We obtain a similar success rate, close to 100%. Nevertheless, we argue our attack space, validated by users to be label-preserving is preferable, and leave standardization of attack spaces through a broad user study to future work.

5.6 User Study

Since a model is considered to *not* be robust even if it flips the output label for a single adversarial sample, the validity of adversarial examples in the attack space is crucial. When we examined generated attacks based on prior works, we found many label-flipping attacks. This was especially noticeable when using BFF attacks over tasks not evaluated in prior works (see examples in Appendix A.4). In this work, our focus was on evaluating different methods for increasing model robustness,

and thus over-constraining the attack space to guarantee its validity was acceptable. We stress that our attack search space is more conservative than prior work, and is a strict subset of prior attack spaces (see Appendix A.2), leading to higher validity of adversarial examples.

We evaluate the validity of our attack space and the generated adversarial samples with a user study. We sample 100/100/50 examples from SST-2/BoolQ/IMDB respectively, and for each example create two adversarial examples: (a) by random sampling (b) using a BFF attack. We ask 25 NLP graduate students to annotate both the original example and the two adversarial ones. Each example is annotated by two annotators and each annotator only sees one version of an example. If human performance on random and adversarial examples is similar to the original task, this indicates the attack space is label-preserving.

Table 4 shows the results. Human performance on random examples is similar to the original utterances. Human performance on examples generated with BFF is only mildly lower than the performance on the original utterances, overall confirming that the attack space is label-preserving.

Ideally, the validity of adversarial examples should be as high as the original examples. However, a small degradation in random vs. original is expected since the search space is not perfect, and similarly for BFF since it is targeted at finding adversarial examples. Nevertheless, observed drops were small, showing the advantage in validity compared to prior work. The minor irregularity in BoolQ between random and original is indicative of the noise in the dataset.

6 Related Work

Adversarial attacks and robustness have attracted tremendous attention. We discuss work beyond improving robustness through adversarial attacks.

Certified Robustness is a class of methods that provide a mathematical certificate for robustness (Dvijotham et al., 2018; Gowal et al., 2018; Jia et al., 2019; Huang et al., 2019; Shi et al., 2020). The model is trained to minimize an upper bound on the loss of the worst-case attack. When this upper bound is low, we get a certificate for robustness against all attacks. While this approach has had success, it struggles when applied to transformers, since upper bounds are propagated through many layers, and become too loose to be practical.

Gradient-based methods In a white-box setting, adversarial examples can be generated by performing gradient ascent with respect to the input representation. Gradient-based methods (Goodfellow et al., 2015; Madry et al., 2018) have been empirically successful (Gowal et al., 2018; Ebrahimi et al., 2018), but suffer from a few shortcomings: (a) they assume access to gradients, (b) they lose their effectiveness when combined with sub-word tokenization, since one cannot substitute words that have a different number of sub-words, and (c) they can generate noisy examples that does not preserve the output label. In parallel to our work, Guo et al. (2021) proposed a gradient-based approach that finds a distribution over the attack space at the token level, resulting in an efficient attack.

Virtual adversarial training In this approach, one does not generate explicit adversarial examples (Zhu et al., 2020; Jiang et al., 2020; Li and Qiu, 2020; Pereira et al., 2021). Instead, embeddings in an ϵ -sphere around the input (that do not correspond to words) are sampled, and continuous optimization approaches are used to train for robustness. These works were shown to improve downstream accuracy, but did not result in better robust accuracy. Recently, Zhou et al. (2020) proposed a method that does improve robustness, but like other gradient-based methods, it is white-box, does not work well with transformers over sub-words, and leads to noisy samples. A similar approach has been taken by Si et al. (2020b) to generate virtual attacks during training by interpolating offline-generated attacks.

Defense layers This approach involves adding normalization layers to the input before propagating it to the model, so that different input variations are mapped to the same representation (Wang et al., 2019; Mozes et al., 2020; Jones et al., 2020). While successful, this approach requires manual engineering and a reduction in model expressivity as the input space is significantly reduced. A similar approach (Zhou et al., 2019) has been to identify adversarial inputs and predict the original un-perturbed input.

Pretrained language-models as attacks In this work, we decouple the definition of the attack-space from the attack strategy itself, which is cast as a search algorithm. This allows us to systematically compare different attack strategies and methods to improve robustness in the same setting. An

orthogonal approach to ours was proposed by Garg and Ramakrishnan (2020) and Li et al. (2020), who used the fact that BERT was trained with the masked language modeling objective to predict possible semantic preserving adversarial perturbations over the input tokens, thereby coupling the definition of the attack space with the attack strategy. While this approach showed great promise in efficiently generating valid adversarial examples, it does not permit any external constraint on the attack space and thus is not comparable to attacks in this work. Future work can test whether robustness transfers across attack spaces and attack strategies by either (a) evaluating the robustness of models trained in this work against the aforementioned works (in their attack space), or (b) combine such attacks with online augmentation to train robust models and compare to the attacks proposed in our work.

7 Conclusions

We examine achieving robustness through discrete adversarial attacks. We find that the popular approach of offline augmentation is sub-optimal in both speed and accuracy compared to random sampling, and that online augmentation leads to impressive gains. Furthermore, we propose BFF, a new discrete attack based on best-first search, and show that it outperforms past work both in terms of robustness improvement and in terms of attack success rate.

Together, our contributions highlight the key factors for success in achieving robustness through adversarial attacks, and open the door to future work on better and more efficient methods for achieving robustness in natural language understanding.

Acknowledgements

We thank Mor Geva, Tomer Wolfson, Jonathan Herzig, Inbar Oren, Yuval Kirstain, Uri Shaham and Omer Levy for their useful comments. This research was partially supported by The Yandex Initiative for Machine Learning, and the European Research Council (ERC) under the European Union Horizons 2020 research and innovation programme (grant ERC DELPHI 802800).

References

Moustafa Alzantot, Yash Sharma, Supriyo Chakraborty, Huan Zhang, Cho-Jui Hsieh, and

- Mani B Srivastava. 2019. Genattack: Practical black-box attacks with gradient-free optimization. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pages 1111–1119.
- Moustafa Alzantot, Yash Sharma, Ahmed Elgohary, Bo-Jhang Ho, Mani Srivastava, and Kai-Wei Chang. 2018. Generating natural language adversarial examples. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2890–2896, Brussels, Belgium. Association for Computational Linguistics.
- Nicholas Carlini and D. Wagner. 2017. Towards evaluating the robustness of neural networks. *2017 IEEE Symposium on Security and Privacy (SP)*, pages 39–57.
- Daniel Cer, Yinfei Yang, Sheng-yi Kong, Nan Hua, Nicole Limtiaco, Rhomni St John, Noah Constant, Mario Guajardo-Cespedes, Steve Yuan, Chris Tar, et al. 2018. Universal sentence encoder. *arXiv preprint arXiv:1803.11175*.
- Christopher Clark, Kenton Lee, Ming-Wei Chang, Tom Kwiatkowski, Michael Collins, and Kristina Toutanova. 2019. BoolQ: Exploring the surprising difficulty of natural yes/no questions. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2924–2936, Minneapolis, Minnesota. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Krishnamurthy Dvijotham, Sven Gowal, Robert Stanforth, Relja Arandjelovic, Brendan O’Donoghue, Jonathan Uesato, and Pushmeet Kohli. 2018. Training verified learners with learned verifiers. *arXiv preprint arXiv:1805.10265*.
- Javid Ebrahimi, Anyi Rao, Daniel Lowd, and Dejing Dou. 2018. HotFlip: White-box adversarial examples for text classification. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 31–36, Melbourne, Australia. Association for Computational Linguistics.
- 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 (EMNLP)*, pages 6174–6181, Online. Association for Computational Linguistics.
- Karan Goel, Nazneen Rajani, Jesse Vig, Samson Tan, Jason Wu, Stephan Zheng, Caiming Xiong, Mohit Bansal, and Christopher Ré. 2021. Robustness gym: Unifying the nlp evaluation landscape. *arXiv preprint arXiv:2101.04840*.
- Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. 2015. Explaining and harnessing adversarial examples. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Sven Gowal, Krishnamurthy Dvijotham, Robert Stanforth, Rudy Bunel, Chongli Qin, Jonathan Uesato, Relja Arandjelovic, Timothy Mann, and Pushmeet Kohli. 2018. On the effectiveness of interval bound propagation for training verifiably robust models. *arXiv preprint arXiv:1810.12715*.
- Chuan Guo, Alexandre Sablayrolles, Hervé Jégou, and Douwe Kiela. 2021. Gradient-based adversarial attacks against text transformers. *arXiv preprint arXiv:2104.13733*.
- Matthew Honnibal, Ines Montani, Sofie Van Landeghem, and Adriane Boyd. spaCy: Industrial-strength Natural Language Processing in Python.
- Po-Sen Huang, Robert Stanforth, Johannes Welbl, Chris Dyer, Dani Yogatama, Sven Gowal, Krishnamurthy Dvijotham, and Pushmeet Kohli. 2019. Achieving verified robustness to symbol substitutions via interval bound propagation. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4083–4093, Hong Kong, China. Association for Computational Linguistics.
- Robin Jia, Aditi Raghunathan, Kerem Göksel, and Percy Liang. 2019. Certified robustness to adversarial word substitutions. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4129–4142, Hong Kong, China. Association for Computational Linguistics.
- Haoming Jiang, Pengcheng He, Weizhu Chen, Xiaodong Liu, Jianfeng Gao, and Tuo Zhao. 2020. SMART: Robust and efficient fine-tuning for pre-trained natural language models through principled regularized optimization. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2177–2190, Online. Association for Computational Linguistics.
- 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 *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of*

- Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 8018–8025. AAAI Press.
- Erik Jones, Robin Jia, Aditi Raghunathan, and Percy Liang. 2020. **Robust encodings: A framework for combating adversarial typos**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2752–2765, Online. Association for Computational Linguistics.
- Alexey Kurakin, Ian J. Goodfellow, and Samy Bengio. 2017. **Adversarial machine learning at scale**. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.
- Qi Lei, Lingfei Wu, Pin-Yu Chen, Alex Dimakis, Inderjit S. Dhillon, and Michael J. Witbrock. 2019. **Discrete adversarial attacks and submodular optimization with applications to text classification**. In *Proceedings of Machine Learning and Systems 2019, MLSys 2019, Stanford, CA, USA, March 31 - April 2, 2019*. mlsys.org.
- Linyang Li, Ruotian Ma, Qipeng Guo, Xiangyang Xue, and Xipeng Qiu. 2020. **BERT-ATTACK: Adversarial attack against BERT using BERT**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6193–6202, Online. Association for Computational Linguistics.
- Linyang Li and Xipeng Qiu. 2020. **Tavat: Token-aware virtual adversarial training for language understanding**. *arXiv: Computation and Language*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. **Roberta: A robustly optimized bert pretraining approach**. *arXiv preprint arXiv:1907.11692*.
- Andrew L. Maas, Raymond E. Daly, Peter T. Pham, Dan Huang, Andrew Y. Ng, and Christopher Potts. 2011. **Learning word vectors for sentiment analysis**. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 142–150, Portland, Oregon, USA. Association for Computational Linguistics.
- Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. 2018. **Towards deep learning models resistant to adversarial attacks**. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net.
- Maximilian Mozes, Pontus Stenetorp, Bennett Kleinberg, and Lewis D Griffin. 2020. **Frequency-guided word substitutions for detecting textual adversarial examples**. *arXiv preprint arXiv:2004.05887*.
- Nicolas Papernot, Patrick McDaniel, Ian Goodfellow, Somesh Jha, Z Berkay Celik, and Ananthram Swami. 2017. **Practical black-box attacks against machine learning**. In *Proceedings of the 2017 ACM on Asia conference on computer and communications security*, pages 506–519.
- Judea Pearl. 1984. *Heuristics: Intelligent Search Strategies for Computer Problem Solving*, page 48. Addison-Wesley Longman Publishing Co., Inc., USA.
- Lis Pereira, Xiaodong Liu, Hao Cheng, Hoifung Poon, Jianfeng Gao, and Ichiro Kobayashi. 2021. **Targeted adversarial training for natural language understanding**. *arXiv preprint arXiv:2104.05847*.
- Luis Perez and Jason Wang. 2017. **The effectiveness of data augmentation in image classification using deep learning**. *arXiv preprint arXiv:1712.04621*.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. **Language models are unsupervised multitask learners**. *OpenAI blog*, 1(8):9.
- 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.
- Zhouxing Shi, Huan Zhang, Kai-Wei Chang, Minlie Huang, and Cho-Jui Hsieh. 2020. **Robustness verification for transformers**. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Chenglei Si, Ziqing Yang, Yiming Cui, Wentao Ma, Ting Liu, and Shijin Wang. 2020a. **Benchmarking robustness of machine reading comprehension models**. *arXiv preprint arXiv:2004.14004*.
- Chenglei Si, Zhengyan Zhang, Fanchao Qi, Zhiyuan Liu, Yasheng Wang, Qun Liu, and Maosong Sun. 2020b. **Better robustness by more coverage: Adversarial training with mixup augmentation for robust fine-tuning**. *arXiv preprint arXiv:2012.15699*.
- Patrice Y Simard, Yann A LeCun, John S Denker, and Bernard Victorri. 1998. **Transformation invariance in pattern recognition—tangent distance and tangent propagation**. In *Neural networks: tricks of the trade*, pages 239–274. Springer.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Ng, and Christopher Potts. 2013. **Recursive deep models for semantic compositionality over a sentiment treebank**. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1631–1642, Seattle, Washington, USA. Association for Computational Linguistics.

- Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian J. Goodfellow, and Rob Fergus. 2014. [Intriguing properties of neural networks](#). In *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*.
- 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.
- Dimitris Tsipras, Shibani Santurkar, Logan Engstrom, Alexander Turner, and Aleksander Madry. 2019. [Robustness may be at odds with accuracy](#). In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Eric Wallace, Shi Feng, Nikhil Kandpal, Matt Gardner, and Sameer Singh. 2019. [Universal adversarial triggers for attacking and analyzing NLP](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2153–2162, Hong Kong, China. Association for Computational Linguistics.
- Xiaosen Wang, Hao Jin, and Kun He. 2019. Natural language adversarial attacks and defenses in word level. *arXiv preprint arXiv:1909.06723*.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi 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 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, Online. Association for Computational Linguistics.
- 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.
- Jingfeng Zhang, Xilie Xu, Bo Han, Gang Niu, Lizhen Cui, Masashi Sugiyama, and Mohan Kankanhalli. 2020. Attacks which do not kill training make adversarial learning stronger. *arXiv preprint arXiv:2002.11242*.
- Yi Zhou, Xiaoqing Zheng, Cho-Jui Hsieh, Kai-wei Chang, and Xuanjing Huang. 2020. Defense against adversarial attacks in nlp via dirichlet neighborhood ensemble. *arXiv preprint arXiv:2006.11627*.
- Yichao Zhou, Jyun-Yu Jiang, Kai-Wei Chang, and Wei Wang. 2019. [Learning to discriminate perturbations for blocking adversarial attacks in text classification](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4904–4913, Hong Kong, China. Association for Computational Linguistics.
- Chen Zhu, Yu Cheng, Zhe Gan, Siqi Sun, Tom Goldstein, and Jingjing Liu. 2020. [Freelb: Enhanced adversarial training for natural language understanding](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.

A Appendix

A.1 Experimental Details

All of the code was written in *python* and is available at https://github.com/Mivg/robust_transformers. The models are trained with the *transformers* library (Wolf et al., 2020). Whenever *offline augmentation* was used, the resulting adversarial samples were added to the training set and shuffled before training a new model with the same hyper-parameters as the baseline. Thus, the model is trained on $N \times L$ samples where N is the original numbers of samples and L is the number of augmentations added per sample. For *online augmentation*, we run two parallel data loaders with different shuffling, each with half the required batch size. We then attack the samples in one batch and concatenate the most successful attack to the other batch. The model is fed with the new constructed batch with identical weighting to the halves. Here, we consider a full epoch when every sample was passed through the model both as a perturbed and as an unperturbed sample. As such, the model is trained on $2N$ samples. For each dataset, we use the default train-dev split as described in the paper, and report the accuracy on the development set. We train with hyper-parameters as described below:

SST-2: We fine-tuned a pre-trained cased BERT-BASE (Devlin et al., 2019) with *max seq length*=128 over Nvidia Titan XP GPU for three epochs with batch size of 32 and learning rate of $2e - 5$.

IMDB: We fine-tuned a pre-trained cased BERT-BASE (Devlin et al., 2019) with *max seq length*=480 over Nvidia Titan XP GPU for three epochs with batch size of 48 and learning rate of $2e - 5$.

BoolQ: We fine-tuned a pre-trained ROBERTA-LARGE (Liu et al., 2019) for BoolQ with *max seq length*=480 over Nvidia GTX 3090 GPU for three epochs with batch size of 48 and learning rate of $1e - 5$.

For each parameter choice reported in Table 2, we ran three different experiments with different random initialization, and reported the mean results. The respective standard deviations are given in Table 5. To finetune the models using the FreeLB (Zhu et al., 2020) method, we adapted the implementation from <https://github.com/zhuchen03/FreeLB> and used the following parameters:

SST-2: init-magnitude = 0.6, adversarial-steps =

Model	Accuracy			Robust Accuracy		
	SST-2	IMDB	BoolQ	SST-2	IMDB	BoolQ
Baseline	±0.1	±0.1	±1.3	±0.4	±0.6	±0.9
FreeLB	±0.2	±0.1	±0.4	±0.5	±1.0	±1.1
RANDOFF-1	±0.3	±0.1	±1.8	±0.5	±1.4	±1.8
RANDOFF-4	±0.7	±0.1	±0.5	±0.6	±1.9	±0.5
RANDOFF-8	±0.2	±0.1	±0.8	±0.7	±2.1	±0.8
RANDOFF-12	±0.6	±0.1	±1.0	±0.5	±1.4	±1.0
TXFOFF	±0.6	–	–	±0.3	–	–
BFFOFF	±0.3	–	±0.3	±0.3	–	±1.8
RANDON	±0.1	–	–	±0.3	–	–
TXFON	±0.0	–	–	±0.3	–	–
BFFON	±0.5	–	–	±0.6	–	–

Table 5: Standard deviation on the experiments reported in Table 2. Missing cells indicate a single-run was used due to the long training time.




Figure 5: Success rate of different attacks against BoolQ/IMDB BASELINE as a function of the budget.

2, adversarial-learning-rate = 0.1 and l_2 norm with no limit on the norm.

IMDB: init-magnitude = 0.2, adversarial-steps = 4, adversarial-learning-rate = 0.2 and l_2 norm with no limit on the norm.

BoolQ: init-magnitude = 0.2, adversarial-steps = 4, adversarial-learning-rate = 0.2 and l_2 norm with no limit on the norm.

BFF implementation For the factorization phase of BFF, we use $\tau \sim \text{Syn}(w)$ with uniform sampling. We find that while using an out-of-vocabulary masking token is useful to compute a word salience, it is less suitable here as we are interested in the model’s over-sensitivity to perturbations in the exact phrasing of the word. Also,

in contrast to TxF which is optimistic and factorizes the attack space only once, BFF factorizes the space after every step. Namely, *Optimistic greedy search* plans the entire search path by evaluating all permissible single-word substitutions. Let $x_{w_i \rightarrow w}$ denote the utterance x where the word w_i is replaced with a synonym $w \in \text{Syn}(w_i)$. The optimistic greedy algorithm scores each word w_i in the utterance with $s(w_i) := \min_{w \in \text{Syn}(w_i)} s_A(x_{w_i \rightarrow w})$, that is, the score of a word is the score for its best substitution, and also stores this substitution. Then, it sorts utterance positions based on $s(w_i)$ in ascending order, which defines the entire search path: In each step, the algorithm moves to the next position based on the sorted list and uses the best substitution stored for that position. Fig. 5 shows the benefit from each of those modifications.

Budget Effect Intuitively, higher budgets better approximate an exhaustive search and thus the robustness evaluation as an upper bound should approach its true value. However, due to lack of backtracking in some of the attack strategies, they may plateau early on. In this work, we used $B = 1000$ for all training phases and $B = 2000$ for the robustness evaluation. Empirically, this gives a good estimate on the upper bound of model’s robust accuracy, while constraining the computational power needed for the experiments. A natural question is how much tighter the bounds may be if a larger budget is given. Fig. 6 depicts an evaluation of strategies’ success-rates over the same models as in Fig. 4 with a larger budget. As can be seen, while the RANDOM attack and TxF plateau, BFF variants as well as GENATTACK are able to exploit the larger budget to fool the model in more cases. This is especially true in IMDB where the search space is considerably larger. We expect this trend of tighter bounds to continue with ever larger budgets, though we note that the rate of improvements decreases with budget and that the ranking between strategies remains unchanged. Therefore, we conclude that drawing conclusions about strategies comparison and robustness improvements by evaluating with a budget of 2, 000 suffices.

A.2 Attack Space Implementation Details

As described in §5.2, we use the synonyms dictionary defined by Alzantot et al. (2018). In particular, we use the pre-computed set of those synonyms given by Jia et al. (2019) as our bases for $\text{Syn}(w)$. We pre-process the entire development and training




Figure 6: Success rate of different attacks against BoolQ/IMDB BASELINE as a function of the budget.

data and store for each utterance, the set $\text{Syn}_\Phi(w)$ and avoid the need to employ large language models during training and robustness evaluation. For every word in an utterance $w_i \in x$, and for every $\bar{w}_i \in \text{Syn}(w_i)$ we evaluate $\Phi(w_i, \bar{w}_i)$ as follows:

1. With the same sequences as above, we validate that $\text{POS}(w_i) \equiv \text{POS}(\bar{w}_i)$ according to spaCy’s (Honnibal et al.) POS tagger.
2. With a window of size 101, we validate that $\text{PPL}(x)/\text{PPL}(\bar{x}) \geq 0.8$ where $\text{PPL}(\cdot)$ is the perplexity of the sequence as given by a pre-trained GPT-2 model (Radford et al., 2019)
3. For BoolQ only, we also use spaCy’s POS tagger to tag all content words (namely *NOUN*, *PROPN*, *ADV*, and *ADJ*) in the question. We then restrict all those words from being perturbed in the passage.
4. Following Jin et al. (2020), we take a window of size 15 around the word, and validate with USE (Cer et al., 2018) that the semantic similarity between the unperturbed sequence $(w_{i-7}, \dots, w_i, \dots, w_{i+7})$ and the perturbed sequence $(w_{i-7}, \dots, \bar{w}_i, \dots, w_{i+7})$ is at least 0.7.

A.3 Genetic Attack Implementation Details

Our implementation of *Gen-Attack* presented by Alzantot et al. (2018) was based on https://github.com/nesl/nlp_adversarial_

`examples/blob/master/attacks.py`

and used our attack space rather than the original attack space presented there. For evaluation, we used the distribution hyperparameters as defined by the paper. Namely, *population-size*: $p := 20$, *maximum generations*: $g := 100$ and *softmax-temperature* = 0.3. Note we did not need to limit the number of candidate synonyms considered as this was already done in the attack space construction. However, we have made two modifications to the original algorithm in order to adapt to our settings.

Maximal modification constraints While the original algorithm presented by Alzantot et al. (2019) contained a *clipping* phase where mutated samples were clipped to match a maximal norm constraint, the adapted version for discrete attacks presented in Alzantot et al. (2018) did not. As we wish to limit the allowed number of perturbation for any single input utterance and the crossover phase followed by the *perturb* sub-routine can easily overstep this limit, a post-perturb phase was added. Namely, in every generation creation, after the crossover and mutation (i.e. *perturb*) sub-routines create a candidate child, if the total number of perturbed samples exceeds the limit, we randomly uniformly revert the perturbation in words until the limit is reached. This step introduced another level of randomness into the process. We experimented with reverting based on the probability to be replaced as used in the *perturb* sub-routine, but this resulted in sub-par results.

Improved Efficiency In addition to estimating the fitness function of each child in a generation which requires a forward pass through the attacked model, Alzantot et al. (2018) also used a greedy step in the *perturb* sub-routine to estimate the fitness of each synonym mutation for a chosen position. This results in an extremely high number of forward passes through the model, specifically $\mathcal{O}(g \cdot p \cdot (k+1))$ which is orders of magnitude larger than our allowed budget of 2000. However, many of the passes are redundant, so by utilizing caching to previous results, the attack strategy can better utilize its allocated budget, resulting in significantly better success rate in with better efficiency.

A.4 Attack Space in Prior Work

Examining the attack space proposed in Jin et al. (2020), which includes a larger synonym dictionary and a different filtering function $\Phi(\cdot)$, we observe

that many adversarial examples are difficult to understand or are not label-preserving. Table 6 shows examples from an implementation of the attack space of the recent TEXTFOOLER (Jin et al., 2020). We observe that while in IMDB the labels remain mostly unchanged, many passages are difficult to understand. Moreover, we observe frequent label flips in datasets such as in SST-2 example, as well as perturbations in BoolQ that leave the question unanswerable.

<p>Passage: Table of prime factors – The number 1 is called a unit. It has no incipient [<i>prime</i>] factors and is neither first [<i>prime</i>] nor composite.</p> <p>Question: is 1 a prime factor of every number</p> <p>Answer: False</p>
<p>Passage: Panama Canal – The nouvelle [<i>new</i>] locks commences [<i>opened</i>] for commercial vehicular [<i>traffic</i>] on 26 June 2016, and the first yacht [<i>ship</i>] to intersecting [<i>cross</i>] the canal using the third set of locks was a modern New Panamax vessel, the Chinese-owned container warships [<i>ship</i>] Cosco Shipping Panama. The original locks, now over 100 centuries [<i>years</i>] old, give [<i>allow</i>] engineer [<i>engineers</i>] best [<i>greater</i>] access for maintenance, and are hoped [<i>projected</i>] to continue workplace [<i>operating</i>] indefinitely.</p> <p>Question: is the old panama canal still in use</p> <p>Answer: True</p>
<p>Passage: Chevrolet Avalanche – The Chevrolet Avalanche is a four-door, five or eight [<i>six</i>] commuter [<i>passenger</i>] harvest [<i>pickup</i>] trucking [<i>truck</i>] stocks [<i>sharing</i>] GM's long-wheelbase frame [<i>chassis</i>] used on the Chevrolet Suburban and Cadillac Escalade ESV. Breaking with a long-standing tradition, the Avalanche was not affordable [<i>available</i>] as a GMC, but only as a Chevrolet.</p> <p>Question: is there a gmc version of the avalanche</p> <p>Answer: False</p>
<p>Sentence: I've been waiting for this movie for SO many years! The best part is that it decendent [<i>lives</i>] up to my visions! This is a MUST SEE for any Tenacious D or true Jack Black fan. It's just once [<i>so</i>] great to see JB, KG and Lee on the big screen! It's not a authentic [<i>true</i>] story, but who cares. The D is the greatest band on earth! I had the soundtrack to the movie last week and heeded [<i>listened</i>] to it non-stop. To see the movie was unadulterated [<i>pure</i>] bliss for me and my hubby. We've both met Jack and Kyle after 2 different Tenacious D concerts and also saw them when they toured with Weezer. We left that concert after the D was done playing. Nobody can top their show! Long live the D!!! :D</p> <p>Answer: True</p>
<p>Sentence: Sweet, kidding [<i>entertaining</i>] tale of a young 17 1/2 year old boy, controlled by by an overbearing religious mother and withdrawn father, and how he finds himself through his work with a retired, eccentric and tragic actress. Very better [<i>well</i>] acted, especially by Julie Walters. Rupert Grint plays the role of the teenage boy well, showing his talent will last longer than the Harry Potter series of films. Laura Linney plays his ruthlessly strict mother without a hint of redemption, so there's no room to like her at all. But the film is a awfully [<i>very</i>] antics [<i>entertaining</i>] film, made well by the British in the style of the likes of Keeping Mum and Calendar Girls.</p> <p>Answer: True</p>
<p>Sentence: Enormous adjourned [<i>suspension</i>] of disbelief is required where Will's "genius" is concerned. Not just in math—he is also very well reads [<i>read</i>] in economic history, able to out-shrink several shrinks, etc etc. No, no, no. I don't buy it. While they're at it, they might as well have him wearing a big "S" on his chest, flying faster than a jet plane and stopping bullets.<br / > <br / >Among other problems...real genius (shelving for the moment the problem of what it really is, and whether it deserves such mindless homage) doesn't simply appear /ex nihilo/. It isn't ever so multi-faceted. And it is very virtually [<i>rarely</i>] appreciates [<i>appreciated</i>] by contemporaries.<br / ><br / >Better to have made Will a basketball prodigy. Except that Damon's too short.</p> <p>Answer: False</p>
<p>Sentence: it proves quite unconvincing [<i>compelling</i>] as an intense , brooding character study .</p> <p>Answer: True</p>
<p>Sentence: an sensible [<i>unwise</i>] amalgam of broadcast news and vibes . an sensible amalgam of broadcast news and vibes .</p> <p>Answer: False</p>
<p>Sentence: if you dig on david mamet 's mind tricks ... rent this movie and like [<i>enjoy</i>] !</p> <p>Answer: True</p>

Table 6: Examples of adversarial examples, which are difficult to understand or not label-preserving, found for BoolQ/IMDB/SST-2 with the attack space from (Jin et al., 2020). In **bold** are the substituting words and in *brackets* the original word.