

What’s in a Name? Are BERT Named Entity Representations just as Good for any other Name?

Sriram Balasubramanian[†] Naman Jain[†] Gaurav Jindal[†]
Abhijeet Awasthi Sunita Sarawagi
Indian Institute of Technology, Bombay

{sriramb, namanjain, gauravj, awasthi, sunita}@cse.iitb.ac.in

Abstract

We evaluate named entity representations of BERT-based NLP models by investigating their robustness to replacements from the same typed class in the input. We highlight that on several tasks while such perturbations are natural, state of the art trained models are surprisingly brittle. The brittleness continues even with the recent entity-aware BERT models. We also try to discern the cause of this non-robustness, considering factors such as tokenization and frequency of occurrence. Then we provide a simple method that ensembles predictions from multiple replacements while jointly modeling the uncertainty of type annotations and label predictions. Experiments on three NLP tasks show that our method enhances robustness and increases accuracy on both natural and adversarial datasets.

1 Introduction

Contextual word embeddings from heavily pre-trained language models (Peters et al., 2018; Devlin et al., 2018) now form the basis of many NLP tasks. While they have led to improved accuracy for most tasks, there are mounting concerns on how well these embeddings encapsulate syntactic and semantic constructs such as synonyms, misspellings, and knowledge representations. Indeed, it has been shown that even BERT based models are not robust to synonym swaps or spelling mistakes in a sentence (Jin et al., 2019; Hsieh et al., 2019; Sun et al., 2019). In this work, we investigate how well these contextual representations fare for named entities.

Designing robust representations of named entities is challenging due to the sheer variety of named entities. Named entities diversify with language, geographical location, time of history, and even with the fine types. Adding to this the varying

length of such entities combined with out of vocabulary names, the complexity only increases.

We quantify how well current systems understand named entities by studying their robustness to substitutions of name mentions in a sentence with other names within an entity class. The entity class within which we seek such robustness is task-dependent and easy for humans to provide. For example, we may require a natural language inference model to be robust to the replacement of company names within the input sentence pairs. In Table 1 we show a sentence pair which contains mentions of a company name Facebook. When we replace that mention with other company names like Microsoft or Google, a robust model should continue to make the same prediction. Likewise, we may require a co-reference resolution model to be robust to replacements of person names in a passage, and a grammar error correction model to be robust to replacement of person names of same gender or country names. A good language representation should be able to generalize well to such perturbations and not deviate from its output upon such perturbations.

The contributions of this work are three-fold. First, we investigate the robustness of trained NLP models using a generic algorithm that we develop. We empirically demonstrate a lack of robustness of state of the art BERT-based models for different user-specified typed classes spanning three NLP tasks: natural language inference (NLI), coreference resolution (CoRef), and Grammar Error Correction (GEC). The lack of robustness is specifically of concern for an entity-focused task like CoRef, where 85% of test sentences have change in their predictions with a single person name change.

Second, we try to seek explanations for such lack of robustness, by observing performance vs. frequency of named entities occurring in the fine-tuning dataset or based on the count of tokens

[†] equal contribution, sorted alphabetically by last name

Sentence 1: Magner , who is 54 and known as Marge , has been the consumer group ’s chief operating officer since April 2002 , and sits on Facebook Microsoft ’s management committee
Sentence 2: She has been the consumer unit ’s chief operating officer since April 2002 , and sits Facebook Microsoft ’s management committee.
Gold: 1 ; **Prediction:** Original: 1; Perturbed: 0

Sentence 1: The workers accuse Goldman Novell of “ reverse age discrimination ” because of a change in retirement benefits in 1997 .
Sentence 2: Goldman Novell was sued when it changed its retirement benefits in 1997 .
Gold: 0 ; **Prediction:** Original: 0; Perturbed: 1

Table 1: Examples on paraphrase detection task – Replacement of an entity

in a named entity. We also explored if BERT’s wordpiece-level masking was particularly unfavorable to entities by switching to Span-BERT, the recent span based masking model. While overall accuracy improved for all datasets with Span-BERT, we found no change in the robustness of the model.

Finally, we develop a simple approach that ensembles predictions from multiple replacements (RESEMBLE) while modeling the uncertainty of type annotations and label predictions. Our approach not only improves performance on adversarial datasets but also on the original datasets, and achieves higher stability on all the tasks.

2 Evaluating Robustness to Named-Entity Replacements

We study the robustness of BERT-based NLP models w.r.t. type-specific named-entity substitutions, for tasks like NLI, GEC and CoRef. Algorithm 1 describes our method of probing NLP models for lack of robustness. Let V be a dictionary of candidate named entities of a given type c , and D denote a dataset consisting of sentence-label pairs (x, y) . Let G be a model fine-tuned on a pre-trained BERT. For each sentence $(x, y) \in D$, we identify the mentions of named-entities of the type c in x ¹. We obtain a perturbed sentence x_m by replacing all mentions of a distinct name in x by a random entry from V . We repeat this process B times where B is a budget (we used 50), with replacement of names. Over the B perturbations, the sentence with the lowest accuracy is added to the set D_{Worst} and the highest accuracy added to the set D_{Best} . A lower variance in model’s performance across the datasets $\{D, D_{\text{Worst}}, D_{\text{Best}}\}$ is indicative of higher robust-

¹We pre-filtered using a named entity tagger in the spaCy library, and made manual corrections so that all tagged entity mentions are correct in D .

Algorithm 1: Probing a model using named-entity substitutions

Data: D (dataset), V (names), M (metric), B (budget)

Result: $D_{\text{worst}}, D_{\text{best}}$ (datasets on which the model performs worst and best)

for $(x, y) \in D$ **do**

$\text{min_score} = \infty, \text{max_score} = -\infty$;

$N \leftarrow \text{RandomSelection}(V, B)$;

for $n \in N$ **do**

$x' \leftarrow \text{Replace}(x, n)$; // Details in text for each task

$\text{score} \leftarrow M(G(x'), y)$;

if $\text{score} < \text{min_score}$ **then**

$\text{min_score} \leftarrow \text{score}, x_{\text{worst}} \leftarrow x'$

end

if $\text{score} > \text{max_score}$ **then**

$\text{max_score} \leftarrow \text{score}, x_{\text{best}} \leftarrow x'$

end

end

$D_{\text{worst}} \leftarrow D_{\text{worst}} + (x_{\text{worst}}, y)$;

$D_{\text{best}} \leftarrow D_{\text{best}} + (x_{\text{best}}, y)$;

end

ness and vice-versa. We also measure stability as the fraction of sentences in D whose predictions stay unchanged within the budget sized replacements.

We use the above method to evaluate the robustness of state-of-the-art BERT based models. We evaluate NLI with organization name replacements, GEC with person and country name replacements, and CoRef with person name replacements. In Table 3 we report accuracy on the original, worst, and best case perturbations of the input and stability for the four task-entity combinations. We discuss task details and results next.

<p>Task: GEC; Perturbed Entity: Person</p> <p>Text: One day Penny Bujalski discovered it and it go to tell it to his queen .</p> <p>Original Prediction: One day Penny discovered it and went to tell it to his queen .</p> <p>Perturbed Prediction: One day Bujalski discovered it and to tell it to his queen.</p> <hr/> <p>Text: the two boys heard that he was planing to steal some money and kill people so the boys start their adventure on stopping Abigale Injuin Joe .</p> <p>Original Prediction: The two boys heard that he was planning to steal some money and kill people so the boys started their adventure by stopping Abigale .</p> <p>Perturbed Prediction: The two boys heard that he was planning to steal some money and kill people so the boys started their adventure by stopping Joe .</p>
<p>Task: GEC; Perturbed Entity: Country</p> <p>Text: There are countries , such as Greeee Oman or Bulgaria Venezuela , in which the econmoy relies merely on tourism .</p> <p>Original Prediction: There are countries , such as Greece or Bulgaria , in which the econmoy relies merely on tourism .</p> <p>Perturbed Prediction: There are countries , such as Oman or Venezuela , in which the econmoy rely merely on tourism .</p> <hr/> <p>Text: I am 20 years old , living in Port - Said , Egypt China .</p> <p>Original Prediction: I am 20 years old and living in Port - Said , Egypt .</p> <p>Perturbed Prediction: I am 20 years old , living in Port - Said , China .</p>
<p>Task: CoRef; Perturbed Entity: Person</p> <p>Text: And Chris Hill Sam Rusnock our ambassador was in China a few days ago. he made the point and Secretary Rice made the point yesterday to the Chinese Foreign minister , we want to see China use its influence. Speaker Newt Gingrich the former speaker Republican weighed in on this debate in this way. [truncated] Well uh with all due respect to Speaker Gingrich we are on a course which has a reasonable chance of success.</p> <p>Original Predicted Cluster: ["Chris Hill our ambassador","he"]</p> <p>Perturbed Predicted Cluster: ["Sam Rusnock our ambassador","he", "Speaker Gingrich"]</p> <hr/> <p>Text: Arianna Huffington Sydnie Rabaut uh in this lengthy piece this morning, Judy Miller is quoted excuse me as saying [truncated]. Do you buy this notion that she doesn't recall who this other source was? No of course not Howie. In fact I think this is the major unanswered question.</p> <p>Original Predicted Cluster:["Arianna Huffington", "you", "I"]</p> <p>Perturbed Predicted Cluster: []</p>

Table 2: Lack of robustness of GEC and CoRef model with respect to person and country names

	NLI F ₁	GEC F _{0.5}		CoRef F ₁
Dataset	ORG	PER	COUN	PER
Original	84.82	50.93	47.87	76.47
Worst	79.90	36.51	32.12	60.91
Best	90.03	58.32	51.47	87.85
Stability	86.8%	75%	63.4%	12.86%

Table 3: Adversarial Evaluation of BERT on different tasks

2.1 Natural Language Inference (NLI)

Task Paraphrase detection is a binary classification task on whether two sentences are paraphrases of each other. We work on the paraphrasing task of the GLUE dataset (Wang et al., 2018). The standard dataset split consists of 4077 training sentence pairs and 1726 testing pairs. We use the BERT-base model fine-tuned on the training dataset. The model takes as input the concatenated sentence pairs and predicts a binary output. The metric used for this task is F₁ score on the binary output.

Attack details We measure robustness over the organization concept class. As the replacement dictionary V we used organization names from Fortune 500 companies. We filter out sentence pairs consisting of organization name mention in each sentence of the pair and get 218 sentence pairs. We use spaCy (Honnibal, 2016) for tagging the sentences followed by manual inspection of matched entities so that in the 218 filtered sentences all entity mentions are correctly identified.

Results Observe in Table 3 almost a 10% swing in F-score between D_{Worst} , D_{Best} just by replacing organization names in test instances. The perturbation dictionary consisted of Fortune 500 companies, and were not particularly obscure either. As the examples in Table 1 show some of these replacements do not span rare names (Facebook to Microsoft or Goldman to Novell)

2.2 Grammatical error correction (GEC)

Task Grammatical error correction is a sequence prediction task, given an incorrect sentence as input we have to predict the grammatically correct output. We use the LOCNESS corpus (Granger, 1998) comprising of incorrect and correct parallel English essays. The standard dataset split consists of 34,308 incorrect-correct sentence pairs for training and 4,384 pairs for testing. We use the publicly available parallel edit model from (Awasthi et al.,

2019). It uses a BERT model for predicting the edits at every token on the input and applies those edits to compute the final output. We only use a single iteration of the model for ease of evaluation. The performance is measured using F_{0.5} score based on M2 files (Bryant et al., 2017).

Attack details We measure robustness on two concept classes: person names and country names. From the test set, 328 sentences mentioned person names and 82 mentioned country names. For person names, we perform gender-specific replacements. The person name dictionary was created as follows: we start with a large dictionary of 4018 female first names, 3437 male first names and 151670 last names and remove names encountered in the training data. We then generate about 250 names from these sets by combining first names and last names. For countries we use 58 non-frequent country names.

Results The gap in accuracy between the best and worst-case perturbations is almost 20% for both person name and country name replacements. Moreover, we find that 25% of the sentences change prediction on changing person names and more than 35% sentences vary prediction of country names! Table 2 shows some examples. Notice how changing the country from Greece to Oman and Bulgaria to Venezuela changes the edit predictions five tokens away in the sentence.

2.3 Coreference Resolution (CoRef)

Task Coreference resolution refers to the problem of finding all expressions that refer to the same entity in a text. We work on the standard OntoNotes dataset from the CoNLL-2012 shared task on coreference resolution (Pradhan et al., 2012). Each document represents one instance and has a series of sentences within it. The standard split consists of 2,802 training documents and 348 testing documents. We use the BERT base model fine-tuned on the training dataset from (Joshi et al., 2019b). The model predicts top-k spans for a document and then computes antecedent scores for them and thereby builds clusters for coreference. Since documents in OntoNotes contain many clusters while we replace only mentions of a single name in the long document, to better highlight differences, we measure F score for only the gold clusters with the replaced entity.

Attack details We measure the robustness with respect to person names. We filter out documents containing a person name based on gold annotations in the OntoNotes corpus, and get 210 documents. Replacement vocabulary V was made in similar way as mentioned for GEC using the same male, female and last name dictionaries. We also ensure that the name replacements do not alter the coreferences. Therefore, we replace every instance of each name occurring in the document with our randomly sampled adversarial name, taking care that first(or last) names are replaced with adversarial first(or last) names. In case of any ambiguity, we replace the name with the last name. Also the replacements are gender specific.

Results We found the worst stability for CoRef and only 13% of the sentences preserved predictions on named-entity replacements. Also, the gap between the worst and best case perturbations is almost 30 F_1 points. As seen from the truncated document examples in the second-last row of Table 2, replacing the name `Chris Hill` to `Sam Rusnock` makes the model mispredict the original cluster, as it predicts another name `Speaker Gingrich` as co-referent to `Sam Rusnock`. Even in second example changing the name `Arianna Huffington` to `Sydney Rabau` causes model to miss the its entire cluster! We also found that on an average, predictions of model differ by two clusters per sentence after name perturbation. For one document almost 17 clusters were affected by a single entity swap. The non-robustness on CoRef is especially surprising since it is principally a task about named entities. Our experiments were on the widely used OntoNotes dataset with person name mentions. Such varying performance should be a cause of concern for benchmarking CoRef models. Perhaps, the dataset needs to be augmented with variants arising out of named-entity replacements and stability should be a required performance metric, in addition to accuracy on the original sentence.

Another interesting observation across tasks is that the accuracy on the original D is enhanced after moving to D_{Best} — that is, just substituting names in a given instance with more ‘favorable’ names can lead to substantial gains. We will exploit this observation to enhance base accuracy and improve the robustness of NLP models in Section 4.

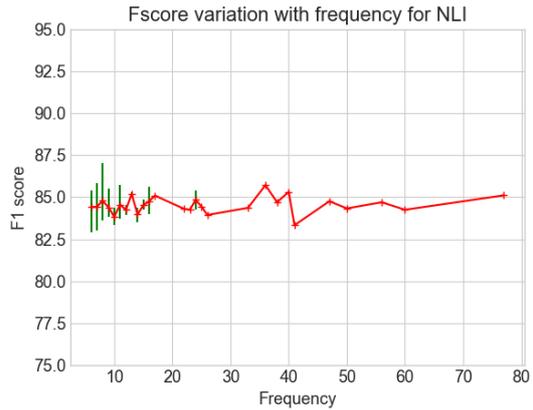


Figure 1: Variation of NLI model’s performance with frequency of named entity in training dataset. The green lines depict variance across performance of names of a given frequency

3 Causes of Non-Robustness

We then sought to investigate reasons for such lack of stability. We first attempted to see if the poor accuracy of certain names can be explained by their frequency of occurrence in the training dataset. In Figure 1 we plot a graph of the frequency of a named-entity in the training corpus against the F-score on the NLI task. As we can see there is no strong correlation of frequency with the performance of a named entity, in fact, an organization name appearing in only four sentence pairs (`Goldman`) performed better than `Microsoft` which was present in over 30 sentence pairs. `Facebook` which is not even present in the training set performs better than `Microsoft` or `Google`. This is likely due to the biases learned during the massive pre-training that BERT-based models enjoy.

Our next guess was to see if the number of tokens in BERT’s word-piece tokenization of named entity causes any significant impact on accuracy. Sequence labeling models like PIE (Awasthi et al., 2019) for GEC are most likely to be susceptible to that effect. In Figure 2 we show accuracy against the number of tokens in a named entity for GEC. We compared performance across three classes – 1 token length entities or two token length entities or three or more token length entities. We created budget sized copies of the original dataset and compare performance across three variants – (Original, Best, and Worst) but found no significant difference in accuracy with the number of tokens. However, we did observe some anecdotal evidence

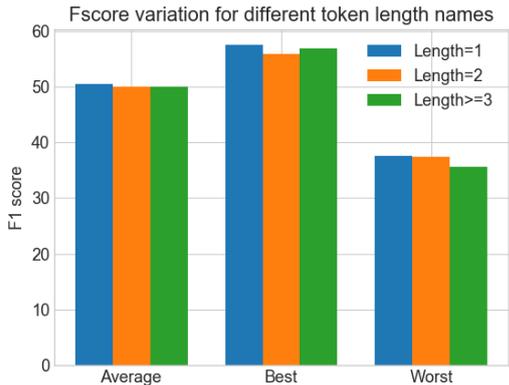


Figure 2: Variation of F_1 scores for GEC model with different token length perturbation.

of specific nuisance tokens arising out of the word piece model on out of vocabulary names. For example consider the person name `Tobey` that gets tokenized as `[To, ##bey]` or `Injuin` which is tokenized as `[In, ##juin]`. The first token of the names are “To” or “In”, both frequent prepositions, which perhaps BERT finds difficult to disambiguate. As we can see from the second example in Table 2 – `Injuin` confuses the given model and the model even deletes the name probably since “In” proposition is not required there. Another artifact could be memorized correlations between names (e.g. `Obama` and `President`) that tasks like CoRef could exploit. Recent work (Poerner et al., 2019) has in fact shown that BERT based models use surface form of entities for relational reasoning.

	NLI (ORG) F_1	CoRef (PER) F_1
Original	86.80	76.71
Worst	82.7	62.37
Best	90.2	86.76
Stability	89.9%	16.19%

Table 4: Adversarial Evaluation of Span-BERT on different tasks

Finally, we explore if BERT’s single token masking model is unfavorable to robust entity representations by comparing with a language model pre-trained by masking spans covering multiple tokens. Specifically, we use Span-BERT (Joshi et al., 2019a), which is trained with masked language modeling on spans instead of tokens. We tried to compare the performance on NLI and CoRef² in

²We were unable to train Span-BERT for GEC, since in

comparison with BERT. The results can be found in Table 4. We were surprised that Span-BERT does not provide any better robustness, although it does provide consistent higher accuracy on all tasks. Various metrics such as – the difference between worst and best accuracy, stability are both very similar for BERT and Span-BERT.

4 Enhancing Robustness

We propose a simple ensembling with replacements approach (referred to as RESEMBLE) that does not require any retraining and can work with any existing pre-trained language model. We assume a type annotator T that marks mentions of entities of the type c for which robustness needs to be enhanced. The type-annotator might be noisy. We identify a small set M of entities of type c on which the model provides high accuracy on a validation set. We call these the list of canonical entities.

Given any input \mathbf{x} , we invoke the task-specific model G to obtain predicted labels $\hat{\mathbf{y}}$ and the type annotator T to obtain type annotations $\hat{\mathbf{z}}$. If $\hat{\mathbf{z}}$ denotes that a named entity of type c is present in one or more spans of \mathbf{x} , we generate new sentences \mathbf{x}_m by replacing the named entities with canonical named entities $m \in M$. The model G when applied to \mathbf{x}_m generates prediction $\hat{\mathbf{y}}_m$.

Let the true labels of \mathbf{x} and \mathbf{x}_m be \mathbf{y} and \mathbf{y}_m respectively, and the true type of \mathbf{x} be \mathbf{z} . If the type annotator correctly identified the spans corresponding to concept class c (i.e., $\mathbf{z} = \hat{\mathbf{z}}$), \mathbf{y} and all \mathbf{y}_m s have to agree as per our requirement of robustness. We use this to define a revised distribution over true \mathbf{y} from the individual predictions as follows:

$$P_R(\mathbf{y}|\mathbf{x}, \hat{\mathbf{z}}) \propto (1 - P(\mathbf{z} = \hat{\mathbf{z}}|\mathbf{x}))P(\mathbf{y}|\mathbf{x}) + P(\mathbf{z} = \hat{\mathbf{z}}|\mathbf{x}) \left(P(\mathbf{y}|\mathbf{x}) \prod_m P(\mathbf{y}|\mathbf{x}_m) \right)^{\frac{1}{m+1}} \quad (1)$$

The above is an annotator confidence weighted average of two terms: The first half calculates the probability of \mathbf{y} from the default model G when the type annotator may be wrong and the \mathbf{y}_m predictions should be ignored. The second half calculates the ensembled agreement probability when the type annotator is correct. We calculate that as a geometric mean of the predictions from the different replacements. In the above equation, the ensembled

released Span-BERT checkpoints were not compatible with the GEC model

probability is under the simplifying assumption that all entity replacements have the same number of tokens. During implementation, we remove this assumption, and implement a more detailed span-level agreement for variable-length entities.

An important requirement for the above expression is that the probabilities provided by the different models express true uncertainty of predictions, that is, they be well-calibrated. Unfortunately, modern neural networks tend to be uncalibrated. To calibrate the probabilities, we use a popular method called temperature scaling (Guo et al., 2017) where probabilities are raised by an exponent, which is the inverse of the temperature. Temperature scaling flattens the probability distribution over output classes thus reduces the confidence until it is correctly calibrated. The expression is as follows:

$$P_T(y|\mathbf{x}) = \frac{P(y|\mathbf{x})^{\frac{1}{T}}}{\sum_{y'} P(y'|\mathbf{x})^{\frac{1}{T}}}$$

where y denotes a scalar prediction. For two of our tasks (GEC and CoRef), the output from our BERT-based models is a product of probabilities from multiple positions. We apply the same temperature scale to each prediction. Thus, our final expression becomes:

$$P_R(\mathbf{y}|\mathbf{x}, \hat{\mathbf{z}}) \propto (1 - P(\mathbf{z} = \hat{\mathbf{z}}|\mathbf{x}))P_T(\mathbf{y}|\mathbf{x}) + P(\mathbf{z} = \hat{\mathbf{z}}|\mathbf{x}) \left(P_T(\mathbf{y}|\mathbf{x}) \prod_m P(\mathbf{y}|\mathbf{x}_m) \right)^{\frac{1}{m+1}} \quad (2)$$

The temperature hyper-parameter T is fixed from a validation dataset. Note we do not apply temperature scaling to the predictions from the canonical entries.

4.1 Empirical Results

For each task, we will describe the defense mechanisms used, with the description of the replacement list, and replacement strategies. The calibration hyper-parameters used for the defense methods are temperatures $T = 2$ across all tasks. The canonical dictionary M for NLI comprises of Microsoft, Nasdaq and IBM. For GEC, due to the huge size of the GEC corpus we pick the most common English first names and combine them with common English last names. We use three male names (John, James Brown, Robert Johnson) and three female names

(Patricia, Mary Jones, Jennifer Brown) for replacement. If gender is ambiguous, we use 1 male name and 2 female names (John, Mary Jones, Jennifer Brown). For CoRef, we used the top 3 frequent person names from the training dataset for our replacement list namely – George Bush, Bill Clinton, Ehud Barak. We also present results when we restrict the canonical dictionary M to only the first name in the above described lists.

We show results with RESEMBLE in Table 5. We perform defense on four datasets – Original, Best, Worst, Random Replacement. For random replacement, we constructed 10 new datasets from the original dataset with its names replaced with randomly selected names, and then evaluate the performance of our models on these datasets. We present the mean and standard deviation of the F scores across these newly constructed datasets. For best and worst we evaluate performance on datasets generated from Algo. 1. First observe that accuracy of even the original test dataset improves with our simple replacement ensembling while reducing the variance. For example, for GEC F score increases from 50.93 to 51.81. The variance has also reduced as seen for the random replacement datasets. The adversarial accuracy improves significantly — for CoRef we see a jump of D_{Worst} from 60.91 to 68.31 and for GEC the gains are even higher. The difference between the best and worst accuracy reduces drastically. Although for D_{Best} accuracy drops with RESEMBLE, the overall gains across the three dataset variants are much higher. Further a single canonical entry $M = 1$ is almost as effective as larger ensembles of $M = 3$. This implies that at test-time, we have to deploy the model on at most two instances to enjoy significantly higher robustness. This shows that replacement with canonical entities while accounting for uncertainty of entity identification is a viable alternative to enhance robustness.

5 Related Work

Study of BERT Representations Jin et al. (2019); Hsieh et al. (2019) study robustness of state of the art BERT fine-tuned models on classification, entailment, and machine translation tasks with respect to synonym replacements. The former used a black box scenario while the latter used input gradients and attention magnitudes to

Dataset	NLI(ORG) F ₁			GEC(PER) F _{0.5}			CoRef(PER) F ₁		
	Original	RESEMBLE		Original	RESEMBLE		Original	RESEMBLE	
		M=1	M=3		M=1	M=3		M=1	M=3
Original	84.80	85.16	85.16	50.93	51.81	51.53	76.47	76.87	76.71
Worst	79.90	82.71	82.71	36.51	47.09	46.75	60.91	68.31	69.43
Best	90.03	86.91	86.63	58.32	55.38	55.38	87.85	82.6	82.18
Random Replacement	85.53 (0.54)	85.48 (0.50)	85.50 (0.40)	49.96 (1.05)	51.47 (0.70)	52.03 (0.66)	76.37 (1.04)	76.78 (0.75)	76.87 (0.81)

Table 5: Adversarial Evaluation of BERT on different tasks comparing the accuracy on the original model against our algorithm with a canonical dictionary of size (M) 1 or 3. For Random Replacement dataset, mean across the ten artificial datasets along with standard deviation in brackets is presented

find probable candidate replacements. Sun et al. (2019) applied an adversarial mis-spelling attack to BERT using gradient-based saliencies. Poerner et al. (2019) show that BERT uses the surface form of words for relational reasoning (guessing person with an Italian sounding name speaks Italian). Zhang et al. (2019a) generated adversarial sentence pairs for paraphrase detection by swapping the order of named entities in two sentences which was enough to fool BERT. Joshi et al. (2019a) introduced Span-BERT that is trained on masked language modelling on spans instead of tokens. Zhang et al. (2019b) developed ERNIE model for entity linking which combines named entity embeddings from knowledge graph with BERT.

Other Robustness Studies in NLP Techniques for generating adversarial examples to study robustness of NLP models have seen a lot of enthusiasm in recent years. These approaches can be loosely categorized into three types – character-level (Ebrahimi et al., 2018b,a) or word-level or sentence-level (Zhao et al., 2018; Iyyer et al., 2018; Ribeiro et al., 2018). Our work is most related to word-level attacks which we elaborate on. Liang et al. (2018) proposed word insertion, deletion, or replacement using gradient magnitudes for classification tasks but requires human effort to ensure the sensibility of the replacements. Samanta and Mehta (2017) used synonym replacements along with the gradient sign method for choosing the worst synonym replacement. Alzantot et al. (2018) provides a population-based genetic algorithm for synonym attacks for sentiment classification and textual entailment in a black-box setting. Ren et al. (2019) developed a greedy algorithm for synonym swaps using weighted gradient based word saliencies, for sentiment classification and entailment.

In this work, we also perform word-level attacks

but our focus is robustness to named entity replacements. The closest work to ours is (Prabhakaran et al., 2019) that checks the sensitivity of models with respect to named entities but they only consider sentiment or toxicity classification. Our work covers more interesting structured prediction tasks such as coreference resolution and grammatical error correction.

Defenses in NLP Most approaches (Cheng et al., 2018; Jia and Liang, 2017) for defenses in NLP have focused on augmenting training datasets with adversarial instances. Pruthi et al. (2019) proposed a word recognition model along with backoff strategies for robustness against misspellings. Zhou et al. (2019) used an adversarial detection cum replacement strategy. We did not consider data augmentation methods because that would significantly increase the training time for models like GEC.

There has also been a trend in usage of certified robustness approaches (Ko et al., 2019; Jia et al., 2019; Huang et al., 2019; Shi et al., 2020) which provide guarantees on the minimum performance of models. The main technique so far is to propagate interval bounds around input word embeddings and has been applied for robustness to synonyms change. Synonyms are expected to have similar embeddings, but interval bounds are unlikely to work for entities within a large concept class. We are not aware of any prior work that enhances robustness with canonical replacements like ours in the context of an existing language model.

6 Conclusions and Future Work

In this work we show that state of the art BERT-based models are surprisingly brittle to named entity replacements. We propose RESEMBLE, a simple ensembling approach to increase robustness while also improving nominal accuracy. The gen-

eral paradigm of enhancing robustness via ensembles on guided instance perturbations is a promising direction and needs to be explored for other tasks too.

Acknowledgement We thank the anonymous reviewers for their constructive feedback on this work. This research was partly sponsored by IBM AI Horizon Networks - IIT Bombay initiative and partly by a Google India AI/ML Research Award. Abhijeet is supported by Google PhD Fellowship in Machine Learning.

References

- Moustafa Alzantot, Yash Sharma, Ahmed Elgohary, Bo-Jhang Ho, Mani Srivastava, and Kai-Wei Chang. 2018. Generating natural language adversarial examples. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*.
- Abhijeet Awasthi, Sunita Sarawagi, Rasna Goyal, Sabyasachi Ghosh, and Vihari Piratla. 2019. Parallel iterative edit models for local sequence transduction. 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)*.
- Christopher Bryant, Mariano Felice, and Ted Briscoe. 2017. Automatic annotation and evaluation of error types for grammatical error correction. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*.
- Yong Cheng, Zhaopeng Tu, Fandong Meng, Junjie Zhai, and Yang Liu. 2018. Towards robust neural machine translation. In *Proceedings of the 56th Association for Computational Linguistics (Volume 1: Long Papers)*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Javid Ebrahimi, Daniel Lowd, and Dejing Dou. 2018a. On adversarial examples for character-level neural machine translation. *International Conference on Computational Linguistics (COLING)*.
- Javid Ebrahimi, Anyi Rao, Daniel Lowd, and Dejing Dou. 2018b. Hotflip: White-box adversarial examples for nlp. *Association for Computational Linguistics (ACL)*.
- S. Granger. 1998. The computer learner corpus: A versatile new source of data for sla research. *Learner English on Computer. Addison Wesley Longman : London New York, 3-18*.
- Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q. Weinberger. 2017. On calibration of modern neural networks. In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, pages 1321–1330.
- M Honnibal. 2016. spacy: Industrial-strength natural language processing in python.
- Yu-Lun Hsieh, Minhao Cheng, Da-Cheng Juan, Wei Wei, Wen-Lian Hsu, and Cho-Jui Hsieh. 2019. On the robustness of self-attentive models.
- 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 (EMNLP)*.
- Mohit Iyyer, John Wieting, Kevin Gimpel, and Luke Zettlemoyer. 2018. Adversarial example generation with syntactically controlled paraphrase networks. In *Proceedings of NAACL*.
- 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*.
- 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 (EMNLP)*.
- Di Jin, Zhijing Jin, Joey Tianyi Zhou, and Peter Szolovits. 2019. Is bert really robust? natural language attack on text classification and entailment. *arXiv preprint arXiv:1907.11932*.
- Mandar Joshi, Danqi Chen, Yinhan Liu, Daniel S. Weld, Luke Zettlemoyer, and Omer Levy. 2019a. SpanBERT: Improving pre-training by representing and predicting spans. *arXiv preprint arXiv:1907.10529*.
- Mandar Joshi, Omer Levy, Luke Zettlemoyer, and Daniel Weld. 2019b. BERT for coreference resolution: Baselines and analysis. 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)*.
- Ching-Yun Ko, Zhaoyang Lyu, Lily Weng, Luca Daniel, Ngai Wong, and Dahua Lin. 2019. POPQORN: Quantifying robustness of recurrent neural networks. In *Proceedings of the 36th International Conference on Machine Learning*.
- Bin Liang, Hongcheng Li, Miaoqiang Su, Pan Bian, Xirong Li, and Wenchang Shi. 2018. Deep text classification can be fooled. *International Joint Conference on Artificial Intelligence, IJCAI 2018*.

- Matthew E Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. *arXiv preprint arXiv:1802.05365*.
- Nina Poerner, Ulli Waltinger, and Hinrich Schütze. 2019. Bert is not a knowledge base (yet): Factual knowledge vs. name-based reasoning in unpervised qa.
- Vinodkumar Prabhakaran, Ben Hutchinson, and Margaret Mitchell. 2019. Perturbation sensitivity analysis to detect unintended model biases. *EMNLP*.
- Sameer Pradhan, Alessandro Moschitti, Nianwen Xue, Olga Uryupina, and Yuchen Zhang. 2012. CoNLL-2012 shared task: Modeling multilingual unrestricted coreference in OntoNotes. In *Joint Conference on EMNLP and CoNLL - Shared Task*.
- Danish Pruthi, Bhuwan Dhingra, and Zachary C Lipton. 2019. Combating adversarial misspellings with robust word recognition.
- Shuhuai Ren, Yihe Deng, Kun He, and Wanxiang Che. 2019. Generating natural language adversarial examples through probability weighted word saliency.
- Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2018. Semantically equivalent adversarial rules for debugging NLP models. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*.
- Suranjana Samanta and Sameep Mehta. 2017. Towards crafting text adversarial samples. *CoRR, abs/1707.02812*,..
- Zhouxing Shi, Huan Zhang, Kai-Wei Chang, Minlie Huang, and Cho-Jui Hsieh. 2020. Robustness verification for transformers. In *International Conference on Learning Representations*.
- Lichao Sun, Kazuma Hashimoto, Wenpeng Yin, Akari Asai, Jia Li, Philip Yu, and Caiming Xiong. 2019. Adv-bert: Bert is not robust on misspellings! generating nature adversarial samples on bert. *arXiv preprint arXiv:2003.04985*.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2018. GLUE: A multi-task benchmark and analysis platform for natural language understanding. In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*.
- Yuan Zhang, Jason Baldridge, and Luheng He. 2019a. PAWS: Paraphrase adversaries from word scrambling. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics*.
- Zhengyan Zhang, Xu Han, Zhiyuan Liu, Xin Jiang, Maosong Sun, and Qun Liu. 2019b. ERNIE: Enhanced language representation with informative entities. In *Proceedings of ACL 2019*.
- Zhengli Zhao, Dheeru Dua, and Sameer Singh. 2018. Generating natural adversarial examples. In *International Conference on Learning Representations*.
- 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 (EMNLP)*.