

# Generating Label Cohesive and Well-Formed Adversarial Claims

Pepa Atanasova\* and Dustin Wright\* and Isabelle Augenstein

Dept. of Computer Science  
University of Copenhagen  
Denmark

{pepa|dw|augenstein}@di.ku.dk

## Abstract

Adversarial attacks reveal important vulnerabilities and flaws of trained models. One potent type of attack are *universal adversarial triggers*, which are individual n-grams that, when appended to instances of a class under attack, can trick a model into predicting a target class. However, for inference tasks such as fact checking, these triggers often inadvertently invert the meaning of instances they are inserted in. In addition, such attacks produce semantically nonsensical inputs, as they simply concatenate triggers to existing samples. Here, we investigate how to generate adversarial attacks against fact checking systems that preserve the ground truth meaning and are semantically valid. We extend the HotFlip attack algorithm used for universal trigger generation by jointly minimizing the target class loss of a fact checking model and the entailment class loss of an auxiliary natural language inference model. We then train a conditional language model to generate semantically valid statements, which include the found universal triggers. We find that the generated attacks maintain the directionality and semantic validity of the claim better than previous work.

## 1 Introduction

Adversarial examples (Goodfellow et al., 2015; Szegedy et al., 2013) are deceptive model inputs designed to mislead an ML system into making the wrong prediction. They expose regions of the input space that are outside the training data distribution where the model is unstable. It is important to reveal such vulnerabilities and correct for them, especially for tasks such as fact checking (FC).

In this paper, we explore the vulnerabilities of FC models trained on the FEVER dataset (Thorne et al., 2018), where the inference between a claim and evidence text is predicted. We particularly

\*denotes equal contribution

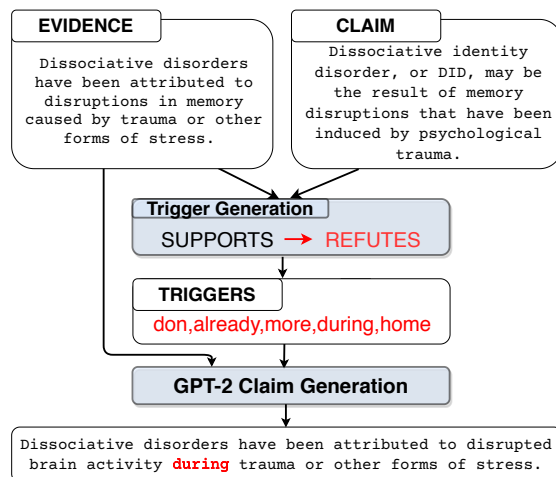


Figure 1: High level overview of our method. First, universal triggers are discovered for flipping a source to a target label (e.g. SUPPORTS → REFUTES). These triggers are then used to condition the GPT-2 language model to generate novel claims with the original label, including at least one of the found triggers.

construct *universal adversarial triggers* (Wallace et al., 2019) – single n-grams appended to the input text that can shift the prediction of a model from a source class to a target one. Such adversarial examples are of particular concern, as they can apply to a large number of input instances.

However, we find that the triggers also change the meaning of the claim such that the true label is in fact the target class. For example, when attacking a claim-evidence pair with a ‘SUPPORTS’ label, a common unigram found to be a universal trigger when switching the label to ‘REFUTES’ is ‘none’. Prepending this token to the claim drastically changes the meaning of the claim such that the new claim is in fact a valid ‘REFUTES’ claim as opposed to an adversarial ‘SUPPORTS’ claim. Furthermore, we find adversarial examples constructed in this way to be nonsensical, as a new token is simply being attached to an existing claim.

Our **contributions** are as follows. We *preserve the meaning* of the source text and *improve the semantic validity* of universal adversarial triggers to automatically construct more potent adversarial examples. This is accomplished via: 1) a *novel extension to the HotFlip attack* (Ebrahimi et al., 2018), where we jointly minimize the target class loss of a FC model and the attacked class loss of a natural language inference model; 2) a *conditional language model* trained using GPT-2 (Radford et al., 2019), which takes trigger tokens and a piece of evidence, and generates a semantically coherent new claim containing at least one trigger. The resulting triggers maintain potency against a FC model while preserving the original claim label. Moreover, the conditional language model produces semantically coherent adversarial examples containing triggers, on which a FC model performs 23.8% worse than with the original FEVER claims. The code for the paper is publicly available.<sup>1</sup>

## 2 Related Work

### 2.1 Adversarial Examples

Adversarial examples for NLP systems can be constructed as automatically generated text (Ren et al., 2019) or perturbations of existing input instances (Jin et al.; Ebrahimi et al., 2018). For a detailed literature overview, see Zhang et al. (2019).

One potent type of adversarial techniques are universal adversarial attacks (Gao and Oates, 2019; Wallace et al., 2019) – single perturbation changes that can be applied to a large number of input instances and that cause significant performance decreases of the model under attack. Wallace et al. (2019) find universal adversarial triggers that can change the prediction of the model using the HotFlip algorithm (Ebrahimi et al., 2018).

However, for NLI tasks, they also change the meaning of the instance they are appended to, and the prediction of the model remains correct. Michel et al. (2019) address this by exploring only perturbed instances in the neighborhood of the original one. Their approach is for instance-dependent attacks, whereas we suggest finding *universal* adversarial triggers that also preserve the original meaning of input instances. Another approach to this are rule-based perturbations of the input (Ribeiro et al., 2018) or imposing adversarial constraints on the produced perturbations (Dia et al., 2019). By contrast, we extend the HotFlip

method by including an auxiliary Semantic Textual Similarity (STS) objective. We additionally use the extracted universal adversarial triggers to generate adversarial examples with low perplexity.

### 2.2 Fact Checking

Fact checking systems consist of components to identify check-worthy claims (Atanasova et al., 2018; Hansen et al., 2019; Wright and Augenstein, 2020), retrieve and rank evidence documents (Yin and Roth, 2018; Allein et al., 2020), determine the relationship between claims and evidence documents (Bowman et al., 2015; Augenstein et al., 2016; Baly et al., 2018), and finally predict the claims’ veracity (Thorne et al., 2018; Augenstein et al., 2019). As this is a relatively involved task, models easily overfit to shallow textual patterns, necessitating the need for adversarial examples to evaluate the limits of their performance.

Thorne et al. (2019a) are the first to propose hand-crafted adversarial attacks. They follow up on this with the FEVER 2.0 task (Thorne et al., 2019b), where participants design adversarial attacks for existing FC systems. The first two winning systems (Niewinski et al., 2019; Hidey et al., 2020) produce claims requiring multi-hop reasoning, which has been shown to be challenging for fact checking models (Ostrowski et al., 2020). The other remaining system (Kim and Allan, 2019) generates adversarial attacks manually. We instead find universal adversarial attacks that can be applied to most existing inputs while markedly decreasing fact checking performance. Niewinski et al. (2019) additionally feed a pre-trained GPT-2 model with the target label of the instance along with the text for conditional adversarial claim generation. Conditional language generation has also been employed by Keskar et al. (2019) to control the style, content, and the task-specific behavior of a Transformer.

## 3 Methods

### 3.1 Models

We take a RoBERTa (Liu et al., 2019) model pre-trained with a LM objective and fine-tune it to classify claim-evidence pairs from the FEVER dataset as SUPPORTS, REFUTES, and NOT ENOUGH INFO (NEI). The evidence used is the gold evidence, available for the SUPPORTS and REFUTES classes. For NEI claims, we use the system of Malon (2018) to retrieve evidence sentences. To measure the semantic similarity between the claim

<sup>1</sup><https://github.com/copenlu/fever-adversarial-attacks>

before and after prepending a trigger, we use a large RoBERTa model fine-tuned on the Semantic Textual Similarity Task.<sup>2</sup> For further details, we refer the reader to §A.1.

### 3.2 Universal Adversarial Triggers Method

The Universal Adversarial Triggers method is developed to find n-gram trigger tokens  $t_\alpha$ , which, appended to the original input  $x$ ,  $f(x) = y$ , cause the model to predict a target class  $\tilde{y} : f(t_\alpha, x) = \tilde{y}$ . In our work, we generate unigram triggers, as generating longer triggers would require additional objectives to later produce well-formed adversarial claims. We start by initializing the triggers with the token ‘a’. Then, we update the embeddings of the initial trigger tokens  $e_\alpha$  with embeddings  $e_{w_i}$  of candidate adversarial trigger tokens  $w_i$  that minimize the loss  $\mathcal{L}$  for the target class  $\tilde{y}$ . Following the HotFlip algorithm, we reduce the brute-force optimization problem using a first-order Taylor approximation around the initial trigger embeddings:

$$\arg \min_{w_i \in \mathcal{V}} [e_{w_i} - e_\alpha]^\top \nabla_{e_\alpha} \mathcal{L} \quad (1)$$

where  $\mathcal{V}$  is the vocabulary of the RoBERTa model and  $\nabla_{e_\alpha} \mathcal{L}$  is the average gradient of the task loss accumulated for all batches. This approximation allows for a  $\mathcal{O}(|\mathcal{V}|)$  space complexity of the brute-force candidate trigger search.

While HotFlip finds universal adversarial triggers that successfully fool the model for many instances, we find that the most potent triggers are often negation words, e.g., ‘not’, ‘neither’, ‘nowhere’. Such triggers change the meaning of the text, making the prediction of the target class correct. Ideally, adversarial triggers would preserve the original label of the claim. To this end, we propose to include an auxiliary STS model objective when searching for candidate triggers. The additional objective is used to minimize the loss  $\mathcal{L}'$  for the maximum similarity score (5 out of 0w) between the original claim and the claim with the prepended trigger. Thus, we arrive at the combined optimization problem:

$$\arg \min_{w_i \in \mathcal{V}} ([e_{w_i} - e_\alpha]^\top \nabla_{e_\alpha} \mathcal{L} + [\mathbf{o}_{w_i} - \mathbf{o}_\alpha]^\top \nabla_{\mathbf{o}_\alpha} \mathcal{L}') \quad (2)$$

where  $\mathbf{o}_w$  is the STS model embedding of word  $w$ . For the initial trigger token, we use “[MASK]” as STS selects candidates from the neighborhood of the initial token.

<sup>2</sup><https://huggingface.co/SparkBeyond/roberta-large-sts-b>

### 3.3 Claim Generation

In addition to finding highly potent adversarial triggers, it is also of interest to generate coherent statements containing the triggers. To accomplish this, we use the HuggingFace implementation of the GPT-2 language model (Radford et al., 2019; Wolf et al., 2019), a large transformer-based language model trained on 40GB of text. The objective is to generate a coherent claim, which either entails, refutes, or is unrelated a given piece of evidence, while also including trigger words.

The language model is first fine tuned on the FEVER FC corpus with a specific input format. FEVER consists of claims and evidence with the labels SUPPORTS, REFUTES, or NOT ENOUGH INFO (NEI). We first concatenate evidence and claims with a special token. Next, to encourage generation of claims with certain tokens, a sequence of tokens separated by commas is prepended to the input. For training, the sequence consists of a single token randomly selected from the original claim, and four random tokens from the vocabulary. This encourages the model to only select the one token most likely to form a coherent and correct claim. The final input format is [trigger tokens]||[evidence]||[claim]. Adversarial claims are then generated by providing an initial input of a series of five comma-separated trigger tokens plus evidence, and progressively generating the rest of the sequence. Subsequently, the set of generated claims is pruned to include only those which contain a trigger token, and constitute the desired label. The latter is ensured by passing both evidence and claim through an external NLI model trained on SNLI (Bowman et al., 2015).

## 4 Results

We present results for universal adversarial trigger generation and coherent claim generation. Results are measured using the original FC model on claims with added triggers and generated claims (macro F1). We also measure how well the added triggers maintain the claim’s original label (semantic similarity score), the perplexity (PPL) of the claims with prepended triggers, and the semantic quality of generated claims (manual annotation). PPL is measured with a pretrained RoBERTa LM.

### 4.1 Adversarial Triggers

Table 1 presents the results of applying universal adversarial triggers to claims from the source class.

The top-performing triggers for each direction are found in §A.2. The adversarial method with a single FC objective successfully deteriorates model performance by a margin of 0.264 F1 score overall. The biggest performance decrease is when the adversarial triggers are constructed to flip the predicted class from SUPPORTS to REFUTES. We also find that 8 out of 18 triggers from the top-3 triggers for each direction, are negation words such as ‘nothing’, ‘nobody’, ‘neither’, ‘nowhere’ (see Table 4 in the appendix). The first of these triggers decreases the performance of the model to 0.014 in F1. While this is a significant performance drop, these triggers also flip the meaning of the text. The latter is again indicated by the decrease of the semantic similarity between the claim before and after prepending a trigger token, which is the largest for the SUPPORTS to REFUTES direction. We hypothesise that the success of the best performing triggers is partly due to the meaning of the text being flipped.

Including the auxiliary STS objective increases the similarity between the claim before and after prepending the trigger for five out of six directions. Moreover, we find that now only one out of the 18 top-3 triggers for each direction are negation words. Intuitively, these adversarial triggers are worse at fooling the FC model as they also have to preserve the label of the original claim. Notably, for the SUPPORTS to REFUTES direction the trigger performance is decreased with a margin of 0.642 compared to the single FC objective. We conclude that including the STS objective for generating Universal Adversarial triggers helps to preserve semantic similarity with the original claim, but also makes it harder to both find triggers preserving the label of the claim while substantially decreasing the performance of the model.

## 4.2 Generation

We use the method described in §3.3 to generate 156 claims using triggers found with the additional STS objective, and 156 claims without. 52 claims are generated for each class (26 flipping to one class, 26 flipping to the other). A different GPT-2 model is trained to generate claims for each specific class, with triggers specific to attacking that class used as input. The generated claims are annotated manually (see §B.3 for the procedure). The overall average claim quality is 4.48, indicating that most generated statements are highly semantically

Class	F1	STS	PPL
<b>No Triggers</b>			
All	.866	5.139	11.92 ( $\pm 45.92$ )
S	.938	5.130	12.22 ( $\pm 40.34$ )
R	.846	5.139	12.14 ( $\pm 37.70$ )
NEI	.817	5.147	14.29 ( $\pm 84.45$ )
<b>FC Objective</b>			
All	.602 ( $\pm 0.289$ )	4.586 ( $\pm 0.328$ )	12.96 ( $\pm 55.37$ )
S→R	.060 ( $\pm 0.034$ )	4.270 ( $\pm 0.295$ )	12.44 ( $\pm 41.74$ )
S→NEI	.611 ( $\pm 0.360$ )	4.502 ( $\pm 0.473$ )	12.75 ( $\pm 40.50$ )
R→S	.749 ( $\pm 0.027$ )	4.738 ( $\pm 0.052$ )	11.91 ( $\pm 36.53$ )
R→NEI	.715 ( $\pm 0.026$ )	4.795 ( $\pm 0.094$ )	11.77 ( $\pm 36.98$ )
NEI→R	.685 ( $\pm 0.030$ )	4.378 ( $\pm 0.232$ )	14.20 ( $\pm 83.32$ )
NEI→S	.793 ( $\pm 0.054$ )	4.832 ( $\pm 0.146$ )	14.72 ( $\pm 93.15$ )
<b>FC+STS Objectives</b>			
All	.763 ( $\pm 0.123$ )	4.786 ( $\pm 0.156$ )	12.97 ( $\pm 58.30$ )
S→R	.702 ( $\pm 0.237$ )	4.629 ( $\pm 0.186$ )	12.62 ( $\pm 41.91$ )
S→NEI	.717 ( $\pm 0.161$ )	4.722 ( $\pm 0.152$ )	12.41 ( $\pm 39.66$ )
R→S	.778 ( $\pm 0.010$ )	4.814 ( $\pm 0.141$ )	11.93 ( $\pm 37.04$ )
R→NEI	.779 ( $\pm 0.009$ )	4.855 ( $\pm 0.098$ )	12.20 ( $\pm 37.67$ )
NEI→R	.780 ( $\pm 0.078$ )	4.894 ( $\pm 0.115$ )	15.27 ( $\pm 111.2$ )
NEI→S	.821 ( $\pm 0.008$ )	4.800 ( $\pm 0.085$ )	13.42 ( $\pm 82.30$ )

Table 1: Universal Adversarial Trigger method performance. Triggers are generated given claims from a source class to fool the classifier to predict a target class (column *Class*, with SUPPORTS (S), REFUTES (R), NEI). The results are averaged over the top 10 triggers.

coherent. The macro F1 of the generative model w.r.t. the intended label is 58.9 overall. For the model without the STS objective, the macro F1 is 56.6, and for the model with the STS objective, it is 60.7, meaning that using triggers found with the STS objective helps the generated claims to retain their intended label.

We measure the performance of the original FC model on generated claims (Table 3). We compare between using triggers that are generated with the STS objective (Ex2) and without (Ex1). In both cases, the adversarial claims effectively fool the FC model, which performs 38.4% worse and 23.8% worse on Ex1 and Ex2, respectively. Additionally, the overall sentence quality increases when the triggers are found with the STS objective (Ex2). The FC model’s performance is higher on claims using triggers generated with the STS objective but still significantly worse than on the original claims. We provide examples of generated claims with their evidence in Table 2.

Comparing FC performance with our generated claims vs. those from the development set of adversarial claims from the FEVER shared task, we see similar drops in performance (0.600 and 0.644 macro F1, respectively). While the adversarial triggers from FEVER cause a larger performance drop,



Evidence	Triggers	Generated Claim
<b>SUPPORTS Claims</b>		
Since the 19th century, some Romani have also migrated to the Americas.	don,already,more,during,home	Romani have moved to the Americas <b>during</b> the 19th century.
Cyprus is a major tourist destination in the Mediterranean.	foreign,biggest,major,every, friends	Cyprus is a <b>major</b> tourist destination.
The first Nobel Prize in Chemistry was awarded in 1901 to Jacobus Henricus van't Hoff, of the Netherlands, "for his discovery of the laws of chemical dynamics and osmotic pressure in solutions."	later,already,quite,altern, whereas	Henricus Van't Hoff was <b>already</b> awarded the Nobel Prize.
<b>REFUTES Claims</b>		
California Attorney General Kamala Harris defeated Sanchez , 61.6% to 38.4%.	phys,incarn,not,occasionally, something	Kamala Harris did <b>not</b> defeat Sanchez, 61.6% to 38.4%.
Uganda is in the African Great Lakes region.	unless,endorsed,picks,pref, <b>against</b>	Uganda is <b>against</b> the African Great Lakes region.
Times Higher Education World University Rankings is an annual publication of university rankings by Times Higher Education (THE) magazine.	interested,reward,visit, <b>consumer</b> ,conclusion	Times Higher Education World University Rankings is a <b>consumer</b> magazine.
<b>NOT ENOUGH INFO Claims</b>		
The KGB was a military service and was governed by army laws and regulations, similar to the Soviet Army or MVD Internal Troops.	nowhere, <b>only</b> ,none,no,nothing	The KGB was <b>only</b> controlled by a military service.
The series revolves around Frank Castle, who uses lethal methods to fight crime as the vigilante "the Punisher", with Jon Bernthal reprising the role from Daredevil.	says,said, <b>take</b> ,say,is	<b>Take</b> Me High is about Frank Castle's use of lethal techniques to fight crime.
The Suite Life of Zack & Cody is an American sitcom created by Danny Kallis and Jim Geoghan.	whilst,interest,applic, <b>someone</b> , nevertheless	The Suite Life of Zack & Cody was created by <b>someone</b> who never had the chance to work in television.

Table 2: Examples of generated adversarial claims. These are all claims which the FC model incorrectly classified.

Target	F1	Avg Quality	# Examples
<b>FC Objective</b>			
Overall	0.534	4.33	156
SUPPORTS	0.486	4.79	39
REFUTES	0.494	4.70	32
NEI	0.621	3.98	85
<b>FC+STS Objectives</b>			
Overall	0.635	4.63	156
SUPPORTS	0.617	4.77	67
REFUTES	0.642	4.68	28
NEI	0.647	4.44	61

Table 3: FC performance for generated claims.

they were manually selected to meet the label coherence and grammatical correctness requirements. Conversely, we automatically generate claims that meet these requirements.

## 5 Conclusion

We present a method for automatically generating highly potent, well-formed, label cohesive claims for FC. We improve upon previous work on universal adversarial triggers by determining how to

construct valid claims containing a trigger word. Our method is fully automatic, whereas previous work on generating claims for fact checking is generally rule-based or requires manual intervention. As FC is only one test bed for adversarial attacks, it would be interesting to test this method on other NLP tasks requiring semantic understanding such as question answering to better understand shortcomings of models.

## Acknowledgements



This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 801199.

## References

- Liesbeth Allein, Isabelle Augenstein, and Marie-Francine Moens. 2020. [Time-Aware Evidence Ranking for Fact-Checking](#). *arXiv preprint arXiv:2009.06402*.
- Pepa Atanasova, Alberto Barron-Cedeno, Tamer Elsayed, Reem Suwaileh, Wajdi Zaghouni, Spas Kyuchukov, Giovanni Da San Martino, and Preslav Nakov. 2018. Overview of the CLEF-2018 Check-That! Lab on Automatic Identification and Verification of Political Claims. Task 1: Check-Worthiness. *arXiv preprint arXiv:1808.05542*.
- Isabelle Augenstein, Christina Lioma, Dongsheng Wang, Lucas Chaves Lima, Casper Hansen, Christian Hansen, and Jakob Grue Simonsen. 2019. Multifc: A real-world multi-domain dataset for evidence-based fact checking of claims. In *EMNLP*. Association for Computational Linguistics.
- Isabelle Augenstein, Tim Rocktäschel, Andreas Vlachos, and Kalina Bontcheva. 2016. [Stance Detection with Bidirectional Conditional Encoding](#). In *EMNLP*, pages 876–885. The Association for Computational Linguistics.
- Ramy Baly, Mitra Mohtarami, James R. Glass, Lluís Mrquez, Alessandro Moschitti, and Preslav Nakov. 2018. [Integrating Stance Detection and Fact Checking in a Unified Corpus](#). In *NAACL-HLT (2)*, pages 21–27. Association for Computational Linguistics.
- Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. A large annotated corpus for learning natural language inference. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics.
- Ousmane Amadou Dia, Elnaz Barshan, and Reza Babanezhad. 2019. [Semantics Preserving Adversarial Learning](#). *arXiv preprint arXiv:1903.03905*.
- 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.
- Hang Gao and Tim Oates. 2019. [Universal Adversarial Perturbation for Text Classification](#). *arXiv preprint arXiv:1910.04618*.
- Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. 2015. Explaining and Harnessing Adversarial Examples. *stat*, 1050:20.
- Casper Hansen, Christian Hansen, Stephen Alstrup, Jakob Grue Simonsen, and Christina Lioma. 2019. Neural Check-Worthiness Ranking With Weak Supervision: Finding Sentences for Fact-Checking. In *Companion Proceedings of the 2019 World Wide Web Conference*, pages 994–1000.
- Christopher Hidey, Tuhin Chakrabarty, Tariq Alhindi, Siddharth Varia, Kriste Krstovski, Mona Diab, and Smaranda Muresan. 2020. [DeSePtion: Dual sequence prediction and adversarial examples for improved fact-checking](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8593–8606, Online. Association for Computational Linguistics.
- Di Jin, Zhijing Jin, Joey Tianyi Zhou, and Peter Szolovits. TextFool: Fool your Model with Natural Adversarial Text.
- Nitish Shirish Keskar, Bryan McCann, Lav R Varshney, Caiming Xiong, and Richard Socher. 2019. Ctrl: A conditional transformer language model for controllable generation. *arXiv preprint arXiv:1909.05858*.
- Youngwoo Kim and James Allan. 2019. [FEVER breaker’s run of team NbAuzDrLqg](#). In *Proceedings of the Second Workshop on Fact Extraction and VERification (FEVER)*, pages 99–104, Hong Kong, China. Association for Computational Linguistics.
- 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*.
- Christopher Malon. 2018. Team Papelo: Transformer Networks at FEVER. In *Proceedings of the First Workshop on Fact Extraction and VERification (FEVER)*, pages 109–113.
- Paul Michel, Xian Li, Graham Neubig, and Juan Miguel Pino. 2019. On Evaluation of Adversarial Perturbations for Sequence-to-Sequence Models. In *Proceedings of NAACL-HLT*, pages 3103–3114.
- Piotr Niewinski, Maria Pszona, and Maria Janicka. 2019. [GEM: Generative enhanced model for adversarial attacks](#). In *Proceedings of the Second Workshop on Fact Extraction and VERification (FEVER)*, pages 20–26, Hong Kong, China. Association for Computational Linguistics.
- Wojciech Ostrowski, Arnav Arora, Pepa Atanasova, and Isabelle Augenstein. 2020. [Multi-Hop Fact Checking of Political Claims](#). *arXiv preprint arXiv:2009.06401*.
- 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.

- 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)*, pages 856–865.
- Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. 2013. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*.
- James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. 2018. FEVER: a Large-scale Dataset for Fact Extraction and VERification. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 809–819.
- James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. 2019a. Evaluating adversarial attacks against multiple fact verification systems. 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 2937–2946.
- James Thorne, Andreas Vlachos, Oana Cocarascu, Christos Christodoulopoulos, and Arpit Mittal. 2019b. [The FEVER2.0 shared task](#). In *Proceedings of the Second Workshop on Fact Extraction and VERification (FEVER)*, pages 1–6, Hong Kong, China. Association for Computational Linguistics.
- 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.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, R’emi Louf, Morgan Funtowicz, and Jamie Brew. 2019. Huggingface’s transformers: State-of-the-art natural language processing. *ArXiv*, abs/1910.03771.
- Dustin Wright and Isabelle Augenstein. 2020. Claim Check-Worthiness Detection as Positive Unlabelled Learning. In *Findings of EMNLP*. Association for Computational Linguistics.
- Wenpeng Yin and Dan Roth. 2018. [TwoWingOS: A Two-Wing Optimization Strategy for Evidential Claim Verification](#). In *EMNLP*, pages 105–114. Association for Computational Linguistics.
- Wei Emma Zhang, Quan Z. Sheng, Ahoud Alhazmi, and Chenliang Li. 2019. [Adversarial Attacks on Deep Learning Models in Natural Language Processing: A Survey](#).

## A Appendices

### A.1 Implementation Details

**Models.** The RoBERTa FC model (125M parameters) is fine-tuned with a batch size of 8, learning rate of  $2e-5$  and for a total of 4 epochs, where the epoch with the best performance is saved. We used the implementation provided by HuggingFace library. We performed a grid hyper-parameter search for the learning rate between the values  $1e-5$ ,  $2e-5$ , and  $3e-5$ . The average time for training a model with one set of hyperparameters is 155 minutes ( $\pm 3$ ). The average accuracy over the different hyperparameter runs is  $0.862(\pm 0.005)$  F1 score on the validation set.

For the models that measure the perplexity and the semantical similarity we use the pretrained models provided by HuggingFace– RoBERTa large model (125M parameters) fine tuned on the STS-b task and RoBERTa base model (355M parameters) pretrained on a LM objective.

We used the HuggingFace implementation of the small GPT-2 model, which consists of 124,439,808 parameters and is fine-tuned with a batch size of 4, learning rate of  $3e-5$ , and for a total of 20 epochs. We perform early stopping on the loss of the model on a set of validation data. The average validation loss is 0.910. The average runtime for training one of the models is 31 hours and 28 minutes.

We note that, the intermediate models used in this work and described in this section, are trained on large relatively general-purpose datasets. While, they can make some mistakes, they work well enough and using them, we don't have to rely on additional human annotations for the intermediate task.

**Adversarial Triggers.** The adversarial triggers are generated based on instances from the validation set. We run the algorithm for three epochs to allow for the adversarial triggers to converge. At each epoch the initial trigger is updated with the best performing trigger for the epoch (according to the loss of the FC or FC+STS objective). At the last step, we select only the top 10 triggers and remove any that have a negative loss. We choose the top 10 triggers as those are the most potent ones, adding more than top ten of the triggers preserves the same tendencies in the results, but smooths them as further down the list of adversarial attacks, the triggers do not decrease the performance of the model substantially. This is also supported by related literature (Wallace et al., 2019), where only

the top few triggers are selected.

The adversarial triggers method is run for 28.75 ( $\pm 1.47$ ) minutes for with the FC objective and 168.6( $\pm 28.44$ ) minutes for the FC+STS objective. We perform the trigger generation with a batch size of four. We additionally normalize the loss for each objective to be in the range  $[0,1]$  and also re-weight the losses with a weight of 0.6 for the FC loss and a weight of 0.4 for the SST loss as when generated with an equal weight, the SST loss tends to preserve the same initial token in all epochs.

**Datasets.** The datasets used for training the FC model consist of 161,249 SUPPORTS, 60,227 REFUTES, and 69,885 NEI claims for the training split; 6,207 SUPPORTS, 6,235 REFUTES, and 6,554 NEI claims for the dev set; 6,291 SUPPORTS, 5,992 REFUTES, and 6522 NEI claims. The evidence for each claim is the gold evidence provided from the FEVER dataset, which is available for REFUTES and SUPPORTS claims. When there is more than one annotation of different evidence sentences for an instance, we include them as separate instances in the datasets. For NEI claims, we use the system of Malon (2018) to retrieve evidence sentences.

### A.2 Top Adversarial Triggers

Table 4 presents the top adversarial triggers for each direction found with the Universal Adversarial Triggers method. It offers an additional way of estimating the effectiveness of the STS objective by comparing the number of negation words generated by the basic model (8) and the STS objective (2) in the top-3 triggers for each direction.

## B Supplemental Material

### B.1 Computing Infrastructure

All experiments were run on a shared cluster. Requested jobs consisted of 16GB of RAM and 4 Intel Xeon Silver 4110 CPUs. We used two NVIDIA Titan RTX GPUs with 12GB of RAM for training GPT-2 and one NVIDIA Titan X GPU with 8GB of RAM for training the FC models and finding the universal adversarial triggers.

### B.2 Evaluation Metrics

The primary evaluation metric used was macro-F1 score. We used the sklearn implementation of `precision_recall_fscore_support`, which can be found here: <https://scikit-learn.org/stable/modules/generated/sklearn>.



`metrics.precision_recall_fscore_support.html`. Briefly:

$$p = \frac{tp}{tp + fp}$$
$$r = \frac{tp}{tp + fn}$$
$$F1 = \frac{2 * p * r}{p + r}$$

where  $tp$  are true positives,  $fp$  are false positives, and  $fn$  are false negatives.

### **B.3 Manual Evaluation**

After generating the claims, two independent annotators label the overall claim quality (score of 1-5) and the true label for the claim. The inter-annotator agreement for the quality label using Krippendorff's alpha is 0.54 for the quality score and 0.38 for the claim label. Given this, we take the average of the two annotator's scores for the final quality score and have a third expert annotator examine and select the best label for each contested claim label.

<b>Class</b>	<b>Trigger</b>	<b>F1</b>	<b>STS</b>	<b>PPL</b>
<b>FC Objective</b>				
S→R	only	0.014	4.628	11.660 (36.191)
S→R	nothing	0.017	4.286	13.109 (56.882)
S→R	nobody	0.036	4.167	12.784 (37.390)
S→NEI	neither	0.047	3.901	11.509 (31.413)
S→NEI	none	0.071	4.016	13.136 (39.894)
S→NEI	Neither	0.155	3.641	11.957 (44.274)
R→S	some	0.687	4.694	11.902 (33.348)
R→S	Sometimes	0.724	4.785	10.813 (32.058)
R→S	Some	0.743	4.713	11.477 (37.243)
R→NEI	recommended	0.658	4.944	12.658 (36.658)
R→NEI	Recommend	0.686	4.789	10.854 (32.432)
R→NEI	Supported	0.710	4.739	11.972 (40.267)
NEI→R	Only	0.624	4.668	12.939 (57.666)
NEI→R	nothing	0.638	4.476	11.481 (48.781)
NEI→R	nobody	0.678	4.361	16.345 (111.60)
NEI→S	nothing	0.638	4.476	18.070 (181.85)
NEI→S	existed	0.800	4.950	15.552 (79.823)
NEI→S	area	0.808	4.834	13.857 (93.295)
<b>FC+STS Objectives</b>				
S→R	never	0.048	4.267	12.745 (50.272)
S→R	every	0.637	4.612	13.714 (51.244)
S→R	didn	0.719	4.986	12.416 (41.080)
S→NEI	always	0.299	4.774	11.906 (35.686)
S→NEI	every	0.637	4.612	12.222 (38.440)
S→NEI	investors	0.696	4.920	12.920 (42.567)
R→S	over	0.761	4.741	12.139 (33.611)
R→S	about	0.765	4.826	12.052 (37.677)
R→S	her	0.774	4.513	12.624 (41.350)
R→NEI	top	0.757	4.762	12.787 (39.418)
R→NEI	also	0.770	5.034	11.751 (35.670)
R→NEI	when	0.776	4.843	12.444 (37.658)
NEI→R	only	0.562	4.677	14.372 (83.059)
NEI→R	there	0.764	4.846	11.574 (42.949)
NEI→R	just	0.786	4.916	16.879 (135.73)
NEI→S	of	0.802	4.917	11.844 (55.871)
NEI→S	is	0.815	4.931	17.507 (178.55)
NEI→S	A	0.818	4.897	12.526 (67.880)

Table 4: Top-3 triggers found with the Universal Adversarial Triggers methods. The triggers are generated given claims from a source class (column *Class*), so that the classifier is fooled to predict a different target class. The classes are SUPPORTS (S), REFUTES (R), NOT ENOUGH INFO (NEI).