

Get Your Vitamin C!

Robust Fact Verification with Contrastive Evidence

Tal Schuster Adam Fisch Regina Barzilay
Computer Science and Artificial Intelligence Laboratory
Massachusetts Institute of Technology
{tals, fisch, regina}@csail.mit.edu

Abstract

Typical fact verification models use retrieved written evidence to verify claims. Evidence sources, however, often change over time as more information is gathered and revised. In order to adapt, models must be sensitive to subtle differences in supporting evidence. We present VITAMINC, a benchmark infused with challenging cases that require fact verification models to discern and adjust to slight factual changes. We collect over 100,000 Wikipedia revisions that modify an underlying fact, and leverage these revisions, together with additional synthetically constructed ones, to create a total of over 400,000 claim-evidence pairs. Unlike previous resources, the examples in VITAMINC are *contrastive*, i.e., they contain evidence pairs that are nearly identical in language and content, with the exception that one supports a given claim while the other does not. We show that training using this design increases robustness—improving accuracy by 10% on adversarial fact verification and 6% on adversarial natural language inference (NLI). Moreover, the structure of VITAMINC leads us to define additional tasks for fact-checking resources: tagging relevant words in the evidence for verifying the claim, identifying factual revisions, and providing automatic edits via factually consistent text generation.¹

1 Introduction

Determining the truthfulness of factual claims by comparing them to textual sources of evidence has received intense research interest in recent years. An underlying, but often overlooked, challenge for this paradigm, however, is the dynamic nature of today’s written resources. An extraordinary amount of new information becomes available daily; as a result, many consequential facts are established, changed, or added to over time. We argue that the quality of fact verification systems should be

¹The VITAMINC dataset and our models are available at: <https://github.com/TalSchuster/VitaminC>




Figure 1: In VITAMINC, we focus on Wikipedia revisions in which the *factual content* changes. This example revision now *supports* an initially *refuted* claim.

measured by how well they adjust to new evidence. In this way, we seek to advance fact verification by requiring that models remain reliable and robust to the change present in practical settings.

To this end, we focus on fact verification with *contrastive evidence*. That is, we infuse the standard fact verification paradigm with challenging cases that require models to be sensitive to factual changes in their presented evidence (hereon referred to interchangeably as “context”). We present VITAMINC,² a new large-scale fact verification dataset that is based on factual revisions to Wikipedia. The key concept is exemplified in Figure 1: there a factual revision yields a contrastive pair of contexts that are nearly identical in language and content—except that one context refutes the given claim, while the other supports it.

This type of contrastive structure exposes existing deficiencies in model behavior. To illustrate this, we train a classifier on the popular FEVER fact verification dataset (Thorne et al., 2018) and evaluate it on contrastive claim-evidence pairs. We find that the model flips its prediction from the original verdict on only 56% of the contrastive cases. When examples from VITAMINC are included during training, however, the model’s sensitivity increases, flipping on 86% of contrastive cases.

Such context-sensitive inference has two main benefits. First, it ensures that the model consid-

²Etymology of VITAMINC: *Contrastive* evidence keeps fact verification models robust and healthy, hence “Vitamin C.”

ers the provided evidence rather than relying on built-in static knowledge, such as that obtained via language model pre-training (Petroni et al., 2019; Roberts et al., 2020). This is particularly important for scenarios in which the source of truth is mutable (e.g., the current US president, or new declarations as in Figure 1). Second, this setting discourages certain biases and idiosyncrasies—such as exploiting differences in how true vs. false claims are posed—that are common in similar crowd-sourced datasets (Poliak et al., 2018; Schuster et al., 2019). Indeed, we show that augmenting both fact verification models and NLI models with VITAMINC data improves their robustness to adversarial inputs.

Furthermore, our emphasis on contrastive contexts allows us to expand on the scope of commonly considered tasks. Most of the fact verification literature focuses on resolving claims to be true or false (Popat et al., 2018; Thorne and Vlachos, 2018; Wang, 2017). The surrounding ecosystem, however, includes additional challenges, some of which we explore here: Documents such as Wikipedia articles are updated frequently; which edits represent factual changes? For a given claim and (refuting or supporting) evidence pair, which words or phrases in the evidence are most relevant? If we know that a certain claim is true, can we modify an out-dated document to be consistent with it? We show that the unique structure of our VITAMINC dataset can be leveraged to provide both supervised and distantly supervised data for these new questions.

Our key contributions are as follows:

1. We pose a contrastive fact verification paradigm that requires sensitivity to changes in data;
2. We introduce VITAMINC, a new large-scale dataset that supports this paradigm;
3. We demonstrate that training on VITAMINC leads to better performance on standard tasks;
4. We show how VITAMINC opens the door to additional research directions in fact verification.

2 Related Work

Fact Verification. The FEVER dataset (Thorne et al., 2018) fueled the development of many fact-checking models (e.g., see Hanselowski et al., 2018; Nie et al., 2019a,b; Yoneda et al., 2018, *inter alia*). The claim creation process, however, required crowd-workers to write claims related to Wikipedia articles, and was found to engender biases that allow an evidence-agnostic model to

achieve unexpectedly high performance (Schuster et al., 2019). Other recent datasets cover verification against tables (Chen et al., 2020), relational databases (Jo et al., 2019), Wikipedia references (Sathe et al., 2020), multiple articles (Jiang et al., 2020), and search snippets (Augenstein et al., 2019). These resources all assume static ground truths. In contrast, VITAMINC compares objective claims to a dynamic source of truth, and requires models to change their verdicts accordingly.

Annotation Bias. Annotation artifacts are common in many NLP datasets, and affect performance on adversarial and contrastive examples (Gardner et al., 2020; Ribeiro et al., 2020; Ross et al., 2020). Sentence-pair inference tasks such as fact verification (Paul Panenghat et al., 2020; Schuster et al., 2019) and NLI (Gururangan et al., 2018; McCoy et al., 2019; Poliak et al., 2018; Tsuchiya, 2018) are no exception. Alleviating this bias requires either modeling solutions (Karimi Mahabadi et al., 2020; Pratapa et al., 2020; Shah et al., 2020; Thorne and Vlachos, 2020; Utama et al., 2020b), which have limited effectiveness (Utama et al., 2020a), or adversarially removing troublesome training examples (Bras et al., 2020) or manually collecting new ones (Nie et al., 2020; Thorne et al., 2019a), which is model specific. Instead, our dataset design avoids single-sentence artifacts and provides model-agnostic challenging examples that increase the robustness of trained models.

Explainability. Current fact verification datasets provide sentence-level rationales (DeYoung et al., 2020; Petroni et al., 2020) but do not enforce the model’s verdict to rely on them—leading to a potential discrepancy. VITAMINC ensures the verdict is conditioned on the retrieved evidence. Moreover, we use the revision history as distant supervision for word-level rationales, allowing for finer-grained explanations (Camburu et al., 2018; Lei et al., 2016; Portelli et al., 2020; Thorne et al., 2019b).

Factually Consistent Generation. Generating texts that match given facts is a known challenge (Fan et al., 2020; Kryscinski et al., 2020; Lewis et al., 2020b; Parikh et al., 2020; Shah et al., 2020; Tian et al., 2020) as language models tend to degenerate and hallucinate (Holtzman et al., 2020; Schuster et al., 2020; Zhou et al., 2020). Moreover, evaluation is non-trivial, and usually manual. VITAMINC includes supervised data for training sequence-to-sequence models, and provides auto-

matic evaluation via the fact verification classifier.

3 The VITAMINC Dataset

VITAMINC (abbreviated VitC) is based on revisions to English Wikipedia. Wikipedia has become a comprehensive online resource that is rigorously maintained by a large and active community (Benjakob and Harrison, 2019). While adversaries do try to insert disinformation, popular pages are usually quickly corrected (Kumar et al., 2016). Furthermore, Wikipedia’s policies dictate that its content should be written from a neutral perspective—or should otherwise objectively state all points of view.³ These properties make Wikipedia a suitable source of evidence for fact verification models. In the following section, we outline our process for mining factual revisions from Wikipedia.

3.1 Collecting Factual Revisions

We collected the 5K most-viewed English Wikipedia articles⁴ as of January 2020, along with any additional articles referred from them (on average 100 per article). We also included all articles from the FEVER dataset (Thorne et al., 2018). For each article, we retrieved up to 500 of its most recent revisions. In May 2020, we added all COVID-19 related articles⁵ and all of their 41K revisions at the time. Combined together, this resulted in a total of ~200 million revisions. For each revision, we identified all of the modified sentences and stored two versions: (1) before, and (2) after the edit.

In our task, we are only interested in edits made with an intent to introduce a *factual* modification—i.e., a change for which one can make a claim that is supported by one sentence, but not by the other.⁶ To expedite annotation, we trained a BERT classifier (Devlin et al., 2019) on a small labeled set of revised sentences determined to be factual (Yang et al., 2017), and used this model to select the top 305K edited sentences from the corpus for manual annotation. Trained human annotators were then presented with the sentence pairs, and were asked to mark the ones that indeed represented a factual change. Sentences lacking self-contained context were filtered (e.g., short expressions from tables or bulleted lists). Example annotations are presented in Table 1. Note that these annotations can also be

³https://bit.ly/Wiki_Neutral_POV

⁴https://bit.ly/Wiki_popular_pages

⁵<https://wikimediafoundation.org/covid19>

⁶Many edits only reflect grammatical corrections, paraphrasing, or “Wikification” (text formatting/page linking).

recursively recycled for re-training the automated BERT classifier in the future to expand the corpus further (we also introduce this as a task, see §4.1).

3.2 Writing Claims

The factual Wikipedia revisions guide us in creating challenging claims for fact verification. For each revision, annotators were asked to write two symmetric claims related to the same edit:

1. The first should be supported by the *original* sentence and refuted by the *revised* sentence;
2. The second should be supported by the *revised* sentence and refuted by the *original* sentence.

When an explicit contradiction was not possible, a *not enough information* (NEI) relation was used. A group of 70 native English speakers⁷ wrote and reviewed claims. During the annotation period, annotations were delivered in weekly batches, from which we examined random samples to provide feedback and request corrections. Annotators were instructed to write short and self-contained claims. Furthermore, annotators were instructed to avoid copying exact phrases and values when possible, in order to avoid a bias for substantially higher word overlap in supporting pairs over refuting pairs. For example, rather than stating, “*there are x confirmed cases of coronavirus in the US*”, one can write “*there are more than z confirmed cases of coronavirus in the US*”, which is supported if $x > z$ and refuted otherwise. For revisions that only add new information or that remove outdated facts without replacing them, annotators wrote a single claim.

3.3 Adding Synthetic Revisions

Naturally, the real Wikipedia revisions we collect mostly describe facts that frequently change over time, or that are prone to mistakes and corrections (such as quantitative values, see Appendix A.1) (Faruqui et al., 2018; Yang et al., 2017). Sensitivity to contrastive contexts, however, is desirable behavior for any claim. This can both ensure consistency with external sources of truth, and improve the model’s faithfulness via connecting the verdict with a specific evidence (Jacovi and Goldberg, 2020; Ross et al., 2020). For example, we require the model to not only classify the claim “*Tom Hanks was honored by a president*” as true, but to also change its verdict to false if paired with a (fictional) contrasting evidence. As a result, we can verify that the model prioritizes sentence-pair inference over

⁷We sourced our annotators through TransPerfect.

Factual	Wikipedia sentences <i>before</i> and <i>after</i> a revision, presented with VITAMINC claims if the revision is factual.	
✗	Before	More <u>stringent</u> actions were taken in China once the <u>seriousness</u> of the outbreak became apparent, such as quarantining entire cities affecting 60 million individuals in Hubei, and strict travel bans.
	After	More <u>drastic</u> actions were taken in China once the <u>severity</u> of the outbreak became apparent, such as quarantining entire cities affecting 60 million individuals in Hubei, and strict travel bans.
✓	Before	In animals, spaying involves an invasive removal of the ovaries, but rarely has major complications <u>other than that spayed animals tend to gain weight</u> .
	After	In animals, spaying involves an invasive removal of the ovaries, but rarely has major complications; <u>the superstition that it causes weight gain is not based on fact</u> .
	Claim 1 Claim 2	Spayed animals gain weight. Weight gain in spayed animals is a superstitious myth.
✗	Before	As of 16 March, more than 182,000 cases of the disease have been reported in over 160 countries and territories, resulting in <u>around 79,000 recoveries and more than 7,100 deaths</u> .
	After	As of 16 March, more than 182,000 cases of the disease have been reported in over 160 countries and territories, resulting in <u>more than 7,100 deaths and around 79,000 recoveries</u> .
✓	Before	Global hybrid sales are led by the Prius family, with sales of <u>4.7</u> million units representing <u>66.8%</u> of TMC worldwide sales of <u>7.053</u> million Lexus and Toyota units through <u>September</u> 2014.
	After	Global hybrid sales are led by the Prius family, with sales of <u>5.264</u> million units representing <u>65.4%</u> of TMC worldwide sales of <u>8.048</u> million Lexus and Toyota units delivered through <u>July</u> 2014.
	Claim 1 Claim 2	Prius sold less than 5 million units, representing over 65.5% of TMC worldwide sales. Prius sold more than 5 million units, representing less than 65.5% of TMC worldwide sales.

Table 1: Examples of *non-factual* revisions vs. *factual* revisions, and the claims associated with the later. Factual updates change the outcome (i.e., true or false) of a claim that might be in question. Accordingly, the verdict of a classifier should change based on the version presented. Modified words are underlined and colored.

memorization, which can help it generalize better. Therefore, we use the FEVER dataset to augment VITAMINC with *synthetic* revisions to Wikipedia sentences.

We follow the setting of Schuster et al. (2019) to expand claim-evidence pairs from FEVER (Thorne et al., 2018). Specifically, given a false claim from FEVER, we ask annotators to edit the sentence that refutes it so that it will then support the originally false claim. Additionally, we ask them to write a new claim that is refuted by the new, modified sentence, but that is supported by the original version. Following this method, we obtain two claims where each can be supported or refuted by the original, or the synthetically revised, sentence. We follow the same process for constructing synthetic examples using true claims, but with flipped labels.

3.4 Dataset Statistics

In total, 304,671 revised Wikipedia sentences were examined by annotators, of which 107,056 (35%) were found to express a factual modification and were passed to the group of expert annotators for claim writing. As two symmetric claims with opposing facts were created (when possible) for each revision, this resulted in a total of 325,724 total claim-evidence pairs. We collected 163,180 addi-

Split	Supports		Refutes		NEI	
	Real	Syn	Real	Syn	Real	Syn
Train	124,864	60,850	71,108	60,850	52,981	-
Dev	21,102	10,382	12,146	10,382	9,042	-
Test	17,306	10,358	9,907	10,358	7,268	-

Table 2: Number of claim-evidence pairs in VITAMINC. Breakdowns of real vs. synthetic revisions are presented on the left and right of each cell, respectively.

tional pairs following the synthetic process. The data was partitioned as shown in Table 2. The assignment was done randomly by article, and is consistent with FEVER for overlapping articles. Appendix A contains additional details.

4 VITAMINC Tasks

The unique structure of VITAMINC allows us to derive annotations that provide a novel source of supervision for several fact-verification-related tasks. We describe the four main tasks we consider in this work, along with baseline models: (1) factual revision flagging, (2) fact verification, (3) word-level rationales, and (4) factually consistent generation. Figure 2 illustrates an example from VITAMINC. We use the following notations:

- \mathcal{C} is the space of short sentences that express an arbitrary factual statement that can potentially be verified or debunked by external sources.




Figure 2: The VITAMINC dataset uses Wikipedia revisions to motivate four central fact verification tasks. Revision source: <https://en.wikipedia.org/wiki/?diff=945689803>.

- \mathcal{S} is the space of sentences that can be found in a trusted online resource (Wikipedia in this study).
- (s_{t-1}, s_t) denotes the two versions of a sentence that was revised from s_{t-1} to $s_t \in \mathcal{S}$.
- $\text{rel}(c, s)$ denotes the relation between the claim $c \in \mathcal{C}$ and observed evidence $s \in \mathcal{S}$ —which can either support c (SUP), refute it (REF), or not contain enough information (NEI).

4.1 Factual Revision Flagging

Online resources like Wikipedia are continuously changing. In order to remain a reliable and neutral source for recent information, its active community of users must constantly verify and correct the revisions of others. We define *factual revision flagging* as the task of identifying revisions that introduce a *factual* change—e.g., by either modifying a certain fact, adding a new one, or removing an existing one. Such an automated detection process can help the community moderate important articles by serving as a watchdog for factual revisions. Furthermore, tracking factual revisions to certain articles can potentially help keep reliant articles consistent (e.g., citing articles, or non-English versions).

We pose factual revision flagging as a binary classification function $f_{\text{flag}}: \mathcal{S} \times \mathcal{S} \rightarrow \{0, 1\}$, where for a revision $(s_{t-1}, s_t)_i$, we set $y_i = 1$ iff there exists a claim in \mathcal{C} whose label (SUP or REF) changes as a result of the edit (i.e., $\text{SUP} \rightarrow \{\text{REF}, \text{NEI}\}$ or $\text{REF} \rightarrow \{\text{SUP}, \text{NEI}\}$). Table 1 provides example factual and non-factual revisions. We evaluate the following baseline models:

Edit Distance. We measure the edit distance between s_{t-1} and s_t , assuming that larger edits are more likely to represent substantive changes. We tune a decision threshold on the validation set.

BOW. We use an MLP on top of a bag-of-words representation. Each sentence is encoded as e_* , the average *fastText* (Bojanowski et al., 2017) word embedding of its edited words (i.e., that were removed or modified in the revision). The MLP input is then taken as $[e_{t-1}; e_t; |e_t - e_{t-1}|; e_t \cdot e_{t-1}]$.

ALBERT. We train the ALBERT transformer (Lan et al., 2020) using either only the edited words (diff), or the full sentence pair (full).

4.2 Fact Verification

Our basic setting is similar to the inference task of the FEVER dataset.⁸ We predict the verdict for a claim given an observed evidence, $f_{\text{verdict}}: \mathcal{C} \times \mathcal{S} \rightarrow \{\text{SUP}, \text{REF}, \text{NEI}\}$. The FEVER dataset, however, contains independent claim-evidence pairs. In our setting, we have claims paired with revisions such that $\text{rel}(c_i, s_{t-1}) \neq \text{rel}(c_i, s_t)$, creating contrastive triplets. For example, the claim in Figure 2 states that the COVID-19 outbreak was identified before December. VITAMINC matches it with two different contexts (before and after the presented revision), that can either support or refute that claim.

Our baseline model is an ALBERT sentence-pair classifier that predicts $\text{rel}(c, s)$. Compared to BERT (Devlin et al., 2019), it uses fewer parameters by shrinking the embedding size and sharing layers, which we find to improve robustness.

4.3 Word-level Rationales

Word-level rationales provide useful explanations for predictions of neural models (Lei et al., 2016). Such explanations can be particularly useful for semi-automated fact verification, since they allow users to quickly interpret and trust the model’s verdict.⁹ In Figure 2, for example, the date of the first identified case can explain the verdict for the claim.

As first proposed by Lei et al. (2016), the standard definition of extractive rationales asks for selecting the minimal set of input tokens that is sufficient for preserving the model’s prediction. Here we use a slightly modified definition following Shah et al. (2020), where we identify the minimal set of evidence tokens where removing them

⁸To focus on the inference task, as opposed to a full end-to-end system, we assume that we have access to an oracle retriever.

⁹Roitero et al. (2020) showed that explanations can increase the agreement between users and expert fact-checkers.

will change the input’s label to NEI.

We pose this task as conditional masking, where we learn a function $f_{\text{rationale}}: \mathcal{C} \times \mathcal{S} \rightarrow \{0, 1\}^n$, where n is the length of an evidence $s \in \mathcal{S}$. Given an evidence $s = (x_1, \dots, x_n)$ and a claim c , where $\text{rel}(c, s) \in \{\text{SUP}, \text{REF}\}$, we want to find a mask m such that $\text{rel}(c, s \odot m) = \text{NEI}$, where

$$s \odot m = \begin{cases} x_i & \text{if } m[i] = 0; \\ \langle \text{mask} \rangle & \text{if } m[i] = 1. \end{cases}$$

Moreover, we want m to be as sparse as possible. Intuitively, $s \odot m$ could be viewed as an incomplete revision in which the masked words that have not yet been filled in will determine the relation with the claim. We say that m reveals the most responsible words in s for resolving c . Following Shah et al. (2020), we formulate an unsupervised objective as

$$\min \sum_{i=1}^n m_i \text{ s.t. } \text{rel}(c, s \odot m) = \text{NEI}. \quad (1)$$

We evaluate the quality of m by comparing it in terms of F1 to both (1) m_{edit} , the non-stopwords removed or replaced in the true revision (i.e., *edit prediction*), and (2) m_{manual} , a manually annotated “human” reference, (i.e., *rationale prediction*). We implement the following two baselines:

Unsupervised. As in Shah et al. (2020), we optimize a Lagrangian relaxation of Eq. 1, where

$$\mathcal{L}_{\text{us}} := -\log p(\text{rel}(c, s \odot m) = \text{NEI}) + \frac{\lambda}{n} \sum_{i=1}^n m_i.$$

We keep the rel classifier (from §4.2) fixed, and train a separate ALBERT model to predict the mask m using a Gumbel softmax (Jang et al., 2017).

Distantly Supervised. By leveraging opposing claims present in VITAMINC, we are able to identify $m_{\text{edit}} = \text{diff}(s_{t-1}, s_t)$ —i.e., the non-stopwords that are deleted or replaced in s_{t-1} when compared to s_t . We then use m_{edit} as distant supervision for m , where $\mathcal{L}_{\text{ds}} = -\frac{\gamma}{n} \sum_{i=1}^n \log p(m_i = m_{\text{edit}_i})$. We combine both the \mathcal{L}_{us} and \mathcal{L}_{ds} losses.

4.4 Factually Consistent Generation

As facts change, the sources reporting them must change as well to reflect the most recent information. In VITAMINC, this is reflected via the active revisions to Wikipedia. We simulate *automating* this process by considering two generation tasks:

Automatic Revisions. Given an outdated context s_{t-1} and an updated claim c , we learn

$f_{\text{revise}}: \mathcal{S} \times \mathcal{C} \rightarrow \mathcal{S}$ to produce a new context s_t that minimally modifies s_{t-1} to agree with c . For example, one can change s_{t-1} in Figure 2 to state “before December” in order to agree with the claim.

Claim Extraction. Given a revision (s_{t-1}, s_t) , we learn $f_{\text{extract}}: \mathcal{S} \times \mathcal{S} \rightarrow \mathcal{C}$ to produce a short claim c that expresses the factual change.

In both tasks, the output should satisfy $\text{rel}(c, s_t) = \text{SUP}$, while $\text{rel}(c, s_{t-1}) = \text{REF}$. We use f_{verdict} (§4.2) to evaluate this requirement. We experiment with both BART-base (Lewis et al., 2020a) and T5-base (Raffel et al., 2020) sequence-to-sequence transformer-based generators. For the revision task, we concatenate s_{t-1} and c with a separator and train the model to predict s_t . For the claim extraction task, we combine the input pair (s_{t-1}, s_t) into a single sentence that visualizes the revision (e.g., “sales of {4.7 → 5.4} million”).

5 Experiments

We present and analyze results for the models described in Section 4. Our analysis attempts to evaluate several questions: (1) How well can the current state-of-the-art models perform on the VITAMINC tasks? (2) Does VITAMINC increase the robustness of models against adversarial examples? (3) Can VITAMINC improve interpretability by providing supervision for anchoring words?

5.1 Related Datasets

In addition to VITAMINC, we train and evaluate on several related datasets, which we briefly describe:

FEVER (Thorne et al., 2018): A popular fact verification dataset based on Wikipedia. We use the provided SUP and REF claim-evidence pairs. For NEI claims, we randomly sample neutral evidence from the article with the highest BM25 score (Fisch et al., 2021).

MNLI (Williams et al., 2018): A large and diverse dataset for natural language inference. The three-way sentence-pair entailment prediction is similar to fact verification. We use the hypothesis as the claim and the premise as the evidence and evaluate on the “mismatched” evaluation set.

Symmetric (Schuster et al., 2019): A set of challenging symmetric, synthetic extensions to FEVER’s evaluation set that avoid claim-only bias.

Adversarial (Thorne et al., 2019c): Adversarial examples created by participants of the FEVER 2.0 shared task. Teams were asked to create claims that

Model	Train data	AUC	Prec.	Rec.	F1
Edit dist.	-	71.34	64.90	63.18	63.56
ALBERT	PAWS-full	72.20	65.27	60.61	60.48
BOW	VitC-diff	79.87	70.85	67.84	68.55
ALBERT	VitC-diff	89.87	80.69	82.06	81.18
ALBERT	VitC-full	91.97	82.63	84.49	83.18

Table 3: Factual revision flagging scores for models aware of the full sentence-pair (full) and aware only of the modified words (diff). We use ALBERT-base.

break FEVER-trained models. We take all SUP and REF claims and their gold evidence sentences.

Triggers (Atanasova et al., 2020): A set of 186 FEVER claims paraphrased adversarially to contain universal adversarial triggers (Wallace et al., 2019). Its small size leads to high variance results.

ANLI (Nie et al., 2020): An adversarial dataset for MNLI- and FEVER-based models. The creation was performed in three iterative rounds in which a model was trained, and then crowdworkers devised adversarial inputs, and the process repeated.

PAWS (Zhang et al., 2019): A dataset of altered Wikipedia sentences using word swapping and back-translation. Human annotators labeled whether the modified sentence is a paraphrase or not. We evaluate whether a PAWS-trained classifier can be used for our factual revision flagging task.

5.2 Factual Revision Flagging

Table 3 shows the results of our baseline models on the factual revision flagging task. First, we notice that a model trained on the PAWS dataset (reaching 93.42 F1 score on PAWS test) does *not* transfer well to the flagging task, and performs on par with a simple edit distance heuristic. We hypothesize that this is a result of the entity scrambling technique used to synthetically revise sentences in PAWS, which is different from the edits introduced by real, factual Wikipedia revisions in practice.

Second, we see that the performance of neural models trained on the VITAMINC flagging task increases with richer inputs and more advanced models—demonstrating the complexity of the task. The ALBERT (diff) model that uses only the modified word sequences from each sentence (i.e., contextual within a subspan) improves the AUC by 10 points over a BOW model that gets a similar input. The ALBERT (full) model that receives the full sentences as input (i.e., has access to even more context), further improves the AUC by 2 points. Nevertheless, the best model still only reaches 83




Figure 3: Test accuracy of models trained on a dataset of 100K combined SUP and REF examples from VITAMINC and FEVER. The higher the ratio of VITAMINC in the training data, the better the performance on adversarial evaluation sets (solid lines). The shaded areas represent standard error across three runs.

macro-F1, indicating the difficulty of this task.

5.3 Fact Verification

Table 4 summarizes the results for classifiers trained on fact verification and NLI datasets. Verifying claims against real revisions proves to be the hardest. The best model achieves 89% accuracy, lower than that on either VITAMINC’s synthetic cases or the original FEVER examples. Including VITAMINC examples in the training data drastically increases models’ sensitivity to contrastive examples (rightmost column)—while preserving the in-domain accuracy (only -0.42% for FEVER and $+0.12\%$ for MNLI with ALBERT-xlarge). Another evidence for the generalization properties conferred by VITAMINC is its zero-shot performance to both other datasets. An ALBERT-xlarge model trained only on VITAMINC reaches 76% and 79% accuracy on FEVER and MNLI, respectively. In contrast, the transfer accuracy for MNLI→FEVER is 70% and for FEVER→MNLI is only 38%.

Most importantly, models trained with VITAMINC perform better on challenging adversarial datasets. On the other hand, simply augmenting FEVER data with MNLI data has a limited effect on adversarial examples.¹⁰ We conjecture that the contrastive nature of VITAMINC helps models better learn the relations between the claims and evidences—and to avoid relying on certain artifacts that do not generalize well.

To further probe the value of VITAMINC examples compared to FEVER ones (SUP and REF

¹⁰We’ve also tried augmenting FEVER with ANLI for an ALBERT-xlarge model and find it to achieve only 73%, 91%, and 34% on Adver., Sym., and Triggers, respectively.

Model	Train dataset	VitC real	VitC syn	FEVER	MNLI	Adver.	Sym.	Triggers	ANLI	Contrast
ALBERT-base	FEVER	54.78	79.18	95.07	58.45	62.01	81.18	3.33	32.94	55.53
	MNLI	44.93	69.67	65.70	85.22	49.61	72.89	68.82	30.63	56.67
	FEVER + MNLI	55.93	82.26	95.62	85.58	63.97	85.67	38.71	30.59	61.90
	VitC	86.16	89.78	74.56	69.18	71.41	90.17	86.67	37.25	84.11
	VitC + MNLI	86.68	91.18	77.70	85.95	69.58	91.15	70.97	34.31	85.61
	VitC + FEVER	86.26	91.05	94.24	68.90	68.80	91.57	72.04	38.50	85.89
ALBERT- xlarge	FEVER	58.56	84.22	96.47	38.33	72.58	87.08	32.80	36.03	64.59
	MNLI	49.26	74.58	70.47	88.91	53.52	78.65	72.58	36.38	63.41
	FEVER + MNLI	61.21	86.81	96.81	89.04	69.97	89.75	43.55	37.66	70.44
	VitC	88.64	93.84	75.99	78.89	82.51	94.80	67.20	42.66	89.57
	VitC + MNLI	88.69	93.92	76.80	89.03	80.81	94.80	67.20	40.09	89.80
	VitC + FEVER	88.87	94.01	96.05	62.30	80.94	94.80	65.59	40.31	90.94

Table 4: Test accuracy of fact verification and NLI models. VITAMINC-trained models are more robust to adversarial examples and more sensitive to contrastive contexts. The rightmost column shows the percent of FEVER claims in which the prediction flipped when presented with contrastive contexts.

Token labels	Edit prediction			Word-level rationales		
	Prec.	Rec.	F1	Prec.	Rec.	F1
-	31.93	30.56	31.23	31.11	46.64	37.33
✓	55.93	62.68	59.11	36.92	66.03	47.36

Table 5: The distant token-level supervision of VITAMINC improves the edit prediction, and as result identifies the anchoring words (rationales) more accurately.

only), we compose training sets of 100K examples using different ratios of the two datasets. As shown in Figure 3, including more VITAMINC pairs continuously improves the performance on the challenging adversarial and symmetric evaluation sets.

As an additional qualitative experiment, given the recent successes of huge language models such as GPT-3 (Brown et al., 2020), we explore whether such models develop sufficient context sensitivity on their own. Appendix C shows the results of classifying several claims using a few-shot GPT-3 model. We find that GPT-3 still largely underperforms our VITAMINC-trained models in terms of sensitivity—demonstrating the importance of using VITAMINC’s unique structure during training.

5.4 Word-level Rationales

Table 5 shows the results of our baseline models for identifying word-level rationales (i.e., anchoring words in the evidence). While our unsupervised model is able to uncover some patterns, directly leveraging the structure of VITAMINC to obtain distant supervision for likely anchoring words (i.e., token labels) improves both the edit prediction and the word-level rationale prediction performance.¹¹ Example predictions are provided in Appendix E.

¹¹We evaluate rationales using a manually annotated test set of 300 examples (150 each from VitC real and VitC synthetic).

5.5 Factually Consistent Generation

Table 6 presents the results on factually consistent generation. We find BART to perform better in both of our generation tasks (though we only tried the default setting). The BLEU score (Papineni et al., 2002) is lower in the claim extraction task since there is freedom in how to phrase the claims, which can result in greater differences between the outputs and the references. The BERT-based BLEURT score (Sellam et al., 2020) shows a similar trend. Still, the claim extraction model succeeds in updating the facts that reflect the true revision 86% of the time, as measured by the fact verification model’s verdict (f_{verdict}).

The revision generator aims to modify sentences so that they agree with a given claim. According to our fact verification model’s verdict, it succeeds in doing so 76% of the time. Furthermore, revisions should resemble real ones, and preserve the remaining content that is unrelated to the claim. The SARI KEEP F1 (Xu et al., 2016) of 75 shows that the model and the reference mostly agree on parts of the sentence that should be kept unchanged.

We find that the token-based measures and our f_{verdict} metric agree well with human (manual) evaluation scores. We randomly sampled 100 generated and human-written sentences per task, and asked workers on Amazon MTurk to rate their grammaticality and whether the evidence s_t supports the claim. The scores of the generated sentences were on par with the human-written ones, indicating the high-quality of our outputs.

Table 7 presents two example generations for the claim extraction task (we provide additional qualitative examples in Appendix E). Our model is able to efficiently extract a self-contained claim that expresses the correct fact after the edit. As in §5.3,

Target	Model	ROUGE ₂	BLEU	KEEP	SARI scores			BLEURT	f_{verdict}	Manual evaluation	
					ADD	DEL	AVG			Grammar	SUP
Revision	T5	77.63	47.46	72.61	13.32	43.04	42.99	0.38	64.52	81.00	71.80
	BART	85.23	54.86	75.36	18.31	47.95	47.21	0.67	76.26	84.80	83.20
Claim	T5	35.19	13.95	44.36	20.59	87.54	50.83	-0.12	75.39	71.33	72.22
	BART	40.38	16.14	52.91	23.62	91.37	55.97	0.16	85.83	75.78	74.22

Table 6: Factually consistent generation results. Higher is better for all scores and the max value is 100 (except for BLEURT). f_{verdict} is the score of our VITAMINC-trained ALBERT-base model on the outputs. For manual evaluation, outputs were rated by their grammaticality and by how much the evidence supports the claim (SUP). For reference, human-written pairs received 75.75 and 76.0 average scores for Grammar and SUP, respectively.

(s_{t-1}, s_t)	2020 coronavirus pandemic in Germany: there have been{* 349 -> 444 }confirmed cases and 16 recoveries.	
BART (VitC)	More than 400 people have tested positive for COVID-19 in Germany.	$f_{\text{verdict}}(c, s_t) =$ SUP
GPT-3 $\mathcal{T}=0$	As of 14 March, there have been more than 350 confirmed cases of the virus in Germany	SUP
GPT-3 $\mathcal{T}=0.7$	As of March 12, there have been more than 400 confirmed cases and 20 reported deaths	NEI
Reference	There have been more than 400 confirmed coronavirus cases in Germany .	SUP
(s_{t-1}, s_t)	Diego Corrales: Corrales was born to a{* Puerto Rican -> African American }father and a {* Dominican -> Mexican }mother .	
BART (VitC)	Diego Corrales’ father is African American and his mother is Mexican.	SUP
GPT-3 $\mathcal{T}=0$	Corrales was born to a Puerto Rican father and a Mexican mother	REF
GPT-3 $\mathcal{T}=0.7$	Corrales was born to a father from Puerto Rico and a mother from the Dominican Republic	REF
Reference	Diego Corrales’ father was African American and his mother Mexican .	SUP

Table 7: Example outputs for expressing claims that reflect the factual changes in a single Wikipedia revision. The BART-base model is trained on VITAMINC data, while GPT-3 is applied in a 2-shot setting with a temperature of 0 or 0.7 (see Appendix C). The revision (s_{t-1}, s_t) is given to the model as a single sentence, where the edits are between curly brackets. The human-written claim is provided for reference. The rightmost column contains the prediction of our ALBERT-xlarge $f_{\text{verdict}}(c, s_t)$ model (trained on VITAMINC) when using the generated claim.

we also explore how GPT-3 handles this task (we provide two demonstrations in the prompt). Compared to the BART model trained on VITAMINC, GPT-3 appears to make more factually inconsistent or unsupported generations (see Appendix C for more details). Encouragingly, our f_{verdict} classifier is still able to pick up on this—as demonstrated by the predictions in the rightmost column of Table 7. For example, classifying the report about 20 deaths as NEI since it is not part of the source. Once again, this serves to qualitatively demonstrate the effectiveness of leveraging VITAMINC.

6 Conclusion

We presented VITAMINC, a large-scale dataset for training and evaluating fact verification models using contrastive contexts. Our novel method of leveraging factual revisions to Wikipedia enabled us to create challenging examples in which a claim is paired with contexts that are lexically similar, yet factually opposing. Our results illustrated that training on VITAMINC improves classifier sensitivity to subtle changes in evidence, and increases their robustness to adversarial examples.

Furthermore, we formulated several new, important tasks for fact verification that VITAMINC allows us to test. We showed how the dataset’s unique “before and after” structure lends itself to training classifiers to flag factual revisions. In addition, for factual revisions, the edits reveal which words in the evidence are the most critical—which helps supervise word-level rationale models for better interpretability. Finally, we demonstrated that VITAMINC can help with factually consistent text generation. We hope that this work and the range of tasks it presents will motivate and support the fact verification field in developing reliable models that can adapt to dynamically changing evidence.

Acknowledgements

We thank the TransPerfect team, Darsh J Shah and Enrico Santus for helpful discussions, as well as the members of the MIT NLP group and Andreas Vlachos for valuable feedback. This work is supported in part by the Facebook Online Safety Benchmark Award. TS is supported in part by DSO grant DSOCL18002. AF is supported in part by a NSF Graduate Research Fellowship.

References

- Pepa Atanasova, Jakob Grue Simonsen, Christina Lioma, and Isabelle Augenstein. 2020. [Generating fact checking explanations](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7352–7364, Online. Association for Computational Linguistics.
- 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 *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 4685–4697, Hong Kong, China. Association for Computational Linguistics.
- Omer Benjakob and Stephen Harrison. 2019. [From anarchy to wikiality, glaring bias to good cop: Press coverage of wikipedia’s first two decades](#). *Wikipedia @ 20*. <https://wikipedia20.pubpub.org/pub/tw83imoy>.
- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. [Enriching word vectors with subword information](#). *Transactions of the Association for Computational Linguistics*, 5:135–146.
- 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*, pages 632–642, Lisbon, Portugal. Association for Computational Linguistics.
- Ronan Le Bras, Swabha Swayamdipta, Chandra Bhagavatula, Rowan Zellers, Matthew Peters, Ashish Sabharwal, and Yejin Choi. 2020. [Adversarial filters of dataset biases](#).
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#).
- Oana-Maria Camburu, Tim Rocktäschel, Thomas Lukasiewicz, and Phil Blunsom. 2018. [e-snli: Natural language inference with natural language explanations](#). In *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc.
- Wenhu Chen, Hongmin Wang, Jianshu Chen, Yunkai Zhang, Hong Wang, Shiyang Li, Xiyu Zhou, and William Yang Wang. 2020. [Tabfact: A large-scale dataset for table-based fact verification](#). In *International Conference on Learning Representations*.
- 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.
- Jay DeYoung, Sarthak Jain, Nazneen Fatema Rajani, Eric Lehman, Caiming Xiong, Richard Socher, and Byron C. Wallace. 2020. [ERASER: A benchmark to evaluate rationalized NLP models](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4443–4458, Online. Association for Computational Linguistics.
- Angela Fan, Aleksandra Piktus, Fabio Petroni, Guillaume Wenzek, Marzieh Saeidi, Andreas Vlachos, Antoine Bordes, and Sebastian Riedel. 2020. [Generating fact checking briefs](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7147–7161, Online. Association for Computational Linguistics.
- Manaa Faruqui, Ellie Pavlick, Ian Tenney, and Dipanjan Das. 2018. [WikiAtomicEdits: A multilingual corpus of Wikipedia edits for modeling language and discourse](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 305–315, Brussels, Belgium. Association for Computational Linguistics.
- Adam Fisch, Tal Schuster, Tommi S. Jaakkola, and Regina Barzilay. 2021. [Efficient conformal prediction via cascaded inference with expanded admission](#). In *International Conference on Learning Representations*.
- JL Fleiss. 1971. [Measuring nominal scale agreement among many raters](#). *Psychological bulletin*, 76(5):378–382.
- Matt Gardner, Yoav Artzi, Victoria Basmov, Jonathan Berant, Ben Bogin, Sihao Chen, Pradeep Dasigi, Dheeru Dua, Yanai Elazar, Ananth Gottumukkala, Nitish Gupta, Hannaneh Hajishirzi, Gabriel Ilharco, Daniel Khashabi, Kevin Lin, Jiangming Liu, Nelson F. Liu, Phoebe Mulcaire, Qiang Ning, Sameer Singh, Noah A. Smith, Sanjay Subramanian, Reut Tsarfaty, Eric Wallace, Ally Zhang, and Ben Zhou. 2020. [Evaluating models’ local decision boundaries via contrast sets](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1307–1323, Online. Association for Computational Linguistics.
- Suchin Gururangan, Swabha Swayamdipta, Omer Levy, Roy Schwartz, Samuel Bowman, and Noah A. Smith. 2018. [Annotation artifacts in natural language inference data](#). In *Proceedings of the 2018*

- Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 107–112, New Orleans, Louisiana. Association for Computational Linguistics.
- Andreas Hanselowski, Hao Zhang, Zile Li, Daniil Sorokin, Benjamin Schiller, Claudia Schulz, and Iryna Gurevych. 2018. **UKP-athene: Multi-sentence textual entailment for claim verification**. In *Proceedings of the First Workshop on Fact Extraction and VERification (FEVER)*, pages 103–108, Brussels, Belgium. Association for Computational Linguistics.
- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2020. **The curious case of neural text degeneration**. In *International Conference on Learning Representations*.
- Alon Jacovi and Yoav Goldberg. 2020. **Towards faithfully interpretable NLP systems: How should we define and evaluate faithfulness?** In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4198–4205, Online. Association for Computational Linguistics.
- Eric Jang, Shixiang Gu, and Ben Poole. 2017. **Categorical reparameterization with gumbel-softmax**. In *International Conference on Learning Representations*.
- Yichen Jiang, Shikha Bordia, Zheng Zhong, Charles Dognin, Maneesh Singh, and Mohit Bansal. 2020. **HoVer: A dataset for many-hop fact extraction and claim verification**. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3441–3460, Online. Association for Computational Linguistics.
- Saehan Jo, Immanuel Trummer, Weicheng Yu, Xuezhi Wang, Cong Yu, Daniel Liu, and Niyati Mehta. 2019. **Aggchecker: A fact-checking system for text summaries of relational data sets**. *Proc. VLDB Endow.*, 12(12):1938–1941.
- Rabeeh Karimi Mahabadi, Yonatan Belinkov, and James Henderson. 2020. **End-to-end bias mitigation by modelling biases in corpora**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8706–8716, Online. Association for Computational Linguistics.
- Wojciech Kryscinski, Bryan McCann, Caiming Xiong, and Richard Socher. 2020. **Evaluating the factual consistency of abstractive text summarization**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9332–9346, Online. Association for Computational Linguistics.
- Srijan Kumar, Robert West, and Jure Leskovec. 2016. **Disinformation on the web: Impact, characteristics, and detection of wikipedia hoaxes**. In *WWW*, pages 591–602.
- Ouyu Lan, Xiao Huang, Bill Yuchen Lin, He Jiang, Liyuan Liu, and Xiang Ren. 2020. **Learning to contextually aggregate multi-source supervision for sequence labeling**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2134–2146, Online. Association for Computational Linguistics.
- Tao Lei, Regina Barzilay, and Tommi Jaakkola. 2016. **Rationalizing neural predictions**. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 107–117, Austin, Texas. Association for Computational Linguistics.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020a. **BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- Patrick Lewis, Ethan Perez, Aleksandara Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020b. **Retrieval-augmented generation for knowledge-intensive nlp tasks**.
- Tom McCoy, Ellie Pavlick, and Tal Linzen. 2019. **Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference**. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3428–3448, Florence, Italy. Association for Computational Linguistics.
- Feng Nie, Jin-Ge Yao, Jinpeng Wang, Rong Pan, and Chin-Yew Lin. 2019a. **A simple recipe towards reducing hallucination in neural surface realisation**. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2673–2679, Florence, Italy. Association for Computational Linguistics.
- Yixin Nie, Haonan Chen, and Mohit Bansal. 2019b. **Combining fact extraction and verification with neural semantic matching networks**. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01):6859–6866.
- Yixin Nie, Adina Williams, Emily Dinan, Mohit Bansal, Jason Weston, and Douwe Kiela. 2020. **Adversarial NLI: A new benchmark for natural language understanding**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4885–4901, Online. Association for Computational Linguistics.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. **Bleu: a method for automatic evaluation of machine translation**. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia,

- Pennsylvania, USA. Association for Computational Linguistics.
- Ankur Parikh, Xuezhi Wang, Sebastian Gehrmann, Manaal Faruqui, Bhuwan Dhingra, Diyi Yang, and Dipanjan Das. 2020. [ToTTo: A controlled table-to-text generation dataset](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1173–1186, Online. Association for Computational Linguistics.
- Mithun Paul Panenghat, Sandeep Sunthal, Faiz Rafique, Rebecca Sharp, and Mihai Surdeanu. 2020. [Towards the necessity for debiasing natural language inference datasets](#). In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 6883–6888, Marseille, France. European Language Resources Association.
- Fabio Petroni, Aleksandra Piktus, Angela Fan, Patrick Lewis, Majid Yazdani, Nicola De Cao, James Thorne, Yacine Jernite, Vassilis Plachouras, Tim Rocktäschel, and Sebastian Riedel. 2020. [Kilt: a benchmark for knowledge intensive language tasks](#).
- Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander Miller. 2019. [Language models as knowledge bases?](#) 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 2463–2473, Hong Kong, China. Association for Computational Linguistics.
- Adam Poliak, Jason Naradowsky, Aparajita Haldar, Rachel Rudinger, and Benjamin Van Durme. 2018. [Hypothesis only baselines in natural language inference](#). In *Proceedings of the Seventh Joint Conference on Lexical and Computational Semantics*, pages 180–191, New Orleans, Louisiana. Association for Computational Linguistics.
- Kashyap Papat, Subhabrata Mukherjee, Andrew Yates, and Gerhard Weikum. 2018. [DeClarE: Debunking fake news and false claims using evidence-aware deep learning](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 22–32, Brussels, Belgium. Association for Computational Linguistics.
- Beatrice Portelli, Jason Zhao, Tal Schuster, Giuseppe Serra, and Enrico Santus. 2020. [Distilling the evidence to augment fact verification models](#). In *Proceedings of the Third Workshop on Fact Extraction and VERification (FEVER)*, pages 47–51, Online. Association for Computational Linguistics.
- Adithya Pratapa, Sai Muralidhar Jayanthi, and Kavya Nerella. 2020. [Constrained Fact Verification for FEVER](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7826–7832, Online. Association for Computational Linguistics.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, W. Li, and Peter J. Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21:140:1–140:67.
- Marco Tulio Ribeiro, Tongshuang Wu, Carlos Guestrin, and Sameer Singh. 2020. [Beyond accuracy: Behavioral testing of NLP models with CheckList](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4902–4912, Online. Association for Computational Linguistics.
- Adam Roberts, Colin Raffel, and Noam Shazeer. 2020. [How much knowledge can you pack into the parameters of a language model?](#) In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5418–5426, Online. Association for Computational Linguistics.
- Kevin Roitero, Michael Soprano, Beatrice Portelli, Damiano Spina, Vincenzo Della Mea, Giuseppe Serra, Stefano Mizzaro, and Gianluca Demartini. 2020. [The covid-19 infodemic: Can the crowd judge recent misinformation objectively?](#) CIKM '20, page 1305–1314, New York, NY, USA. Association for Computing Machinery.
- Alexis Ross, Ana Marasović, and Matthew E. Peters. 2020. [Explaining nlp models via minimal contrastive editing \(mice\)](#).
- Aalok Sathe, Salar Ather, Tuan Manh Le, Nathan Perry, and Joonsuk Park. 2020. [Automated fact-checking of claims from Wikipedia](#). In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 6874–6882, Marseille, France. European Language Resources Association.
- Tal Schuster, Roei Schuster, Darsh J. Shah, and Regina Barzilay. 2020. [The limitations of stylometry for detecting machine-generated fake news](#). *Computational Linguistics*, 46(2):499–510.
- Tal Schuster, Darsh Shah, Yun Jie Serene Yeo, Daniel Roberto Filizzola Ortiz, Enrico Santus, and Regina Barzilay. 2019. [Towards debiasing fact verification models](#). 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 3419–3425, Hong Kong, China. Association for Computational Linguistics.
- Thibault Sellam, Dipanjan Das, and Ankur Parikh. 2020. [BLEURT: Learning robust metrics for text generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7881–7892, Online. Association for Computational Linguistics.
- Darsh J Shah, Tal Schuster, and Regina Barzilay. 2020. [Automatic fact-guided sentence modification](#). In *Association for the Advancement of Artificial Intelligence (AAAI)*.

- James Thorne and Andreas Vlachos. 2018. [Automated fact checking: Task formulations, methods and future directions](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 3346–3359, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- James Thorne and Andreas Vlachos. 2020. [Avoiding catastrophic forgetting in mitigating model biases in sentence-pair classification with elastic weight consolidation](#).
- 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, New Orleans, Louisiana. Association for Computational Linguistics.
- 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 2944–2953, Hong Kong, China. Association for Computational Linguistics.
- James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. 2019b. [Generating token-level explanations for natural language inference](#). 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 963–969, Minneapolis, Minnesota. Association for Computational Linguistics.
- James Thorne, Andreas Vlachos, Oana Cocarascu, Christos Christodoulopoulos, and Arpit Mittal. 2019c. [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.
- Ran Tian, Shashi Narayan, Thibault Sellam, and Ankur P. Parikh. 2020. [Sticking to the facts: Confident decoding for faithful data-to-text generation](#).
- Masatoshi Tsuchiya. 2018. [Performance impact caused by hidden bias of training data for recognizing textual entailment](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Prasetya Ajie Utama, Nafise Sadat Moosavi, and Iryna Gurevych. 2020a. [Mind the trade-off: Debiasing NLU models without degrading the in-distribution performance](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8717–8729, Online. Association for Computational Linguistics.
- Prasetya Ajie Utama, Nafise Sadat Moosavi, and Iryna Gurevych. 2020b. [Towards debiasing NLU models from unknown biases](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7597–7610, Online. 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, Hong Kong, China. Association for Computational Linguistics.
- William Yang Wang. 2017. [“liar, liar pants on fire”: A new benchmark dataset for fake news detection](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 422–426, Vancouver, Canada. Association for Computational Linguistics.
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. [A broad-coverage challenge corpus for sentence understanding through inference](#). 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 1112–1122, New Orleans, Louisiana. Association for Computational Linguistics.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi 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 M. Rush. 2019. [Huggingface’s transformers: State-of-the-art natural language processing](#). *ArXiv*, abs/1910.03771.
- Wei Xu, Courtney Napoles, Ellie Pavlick, Quanze Chen, and Chris Callison-Burch. 2016. [Optimizing statistical machine translation for text simplification](#). *Transactions of the Association for Computational Linguistics*, 4:401–415.
- Diyi Yang, Aaron Halfaker, Robert Kraut, and Eduard Hovy. 2017. [Identifying semantic edit intentions from revisions in Wikipedia](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2000–2010, Copenhagen, Denmark. Association for Computational Linguistics.
- Takuma Yoneda, Jeff Mitchell, Johannes Welbl, Pontus Stenetorp, and Sebastian Riedel. 2018. [UCL machine reading group: Four factor framework for fact](#)

finding (HexaF). In *Proceedings of the First Workshop on Fact Extraction and VERification (FEVER)*, pages 97–102, Brussels, Belgium. Association for Computational Linguistics.

Yuan Zhang, Jason Baldrige, and Luheng He. 2019. PAWS: Paraphrase adversaries from word scrambling. 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 1298–1308, Minneapolis, Minnesota. Association for Computational Linguistics.

Chunting Zhou, Jiatao Gu, Mona Diab, Paco Guzman, Luke Zettlemoyer, and Marjan Ghazvininejad. 2020. Detecting hallucinated content in conditional neural sequence generation.

A VITAMINC: Complementary details

We provide additional details about the VITAMINC dataset.

A.1 Claim Statistics

Topic Distribution. Figure A.1 shows the distribution of claims in the VITAMINC dataset by the topic of the Wikipedia article they are based on. The information was collected from DBpedia,¹² retrieving the parent class of the pages. Labels for about 25% of the articles were missing, and left blank in the diagram.

The “synthetic” part of VITAMINC, which is based on the claims of the FEVER dataset, contains many claims about specific human entities. About 15% of the claims in VITAMINC real are about COVID-19.

Category Distribution. We sample 100 examples from the “real” and “synthetic” subsets of VITAMINC and manually categorize their claims. The results are presented in Table A.1. Due to the creation methodology of VITAMINC real, its claims mostly describe frequently updating facts, or facts that tend to be corrected. We find about half of these claims to describe changes in numerical values (e.g., number of COVID-19 cases, earnings or ratings of movies, number of awards etc.). In contrast, VITAMINC synthetic mostly covers general facts about specific entities, (e.g., place of birth, date of birth, occupation, etc.). This is a result of the synthetic claims being based on the FEVER dataset, where annotators were asked to come up with claims on popular Wikipedia pages. Combined, the VITAMINC dataset holds a diverse set of claims about various topics.

¹²<http://dbpedia.org/ontology/>




Figure A.1: Distribution of claims in the VITAMINC dataset by the topic of the originated Wikipedia article.

A.2 Inter-annotator Agreement

We ask three additional annotators to independently annotate a random set of two thousand claim-evidence pairs, evenly distributed between the development and test splits of the real and synthetic sets. The Fleiss κ score (Fleiss, 1971) between the four annotations is 0.7065, which means substantial agreement. Similar agreement scores of 0.6841 and 0.7 were reported for fact verification (Thorne et al., 2018) and NLI datasets (Bowman et al., 2015), respectively.

A.3 Claim-only Classification

Annotation artifacts are common in crowd-sourced sentence-pair inference datasets such as fact verification and NLI. Models can leverage these idiosyncrasies to achieve unexpectedly high performance when given only one sentence of the pair. For example, Schuster et al. (2019) showed that a claim-only classifier can obtain 61.7% accuracy. The VITAMINC dataset avoids this bias by pairing each claim with two contrastive contexts.

Claim category	real	synthetic	Example
Quantitative	48%	9%	The COVID-19 pathogen may last less than 10 days on some surfaces.
Calendrical	9%	15%	Italy surpassed the 10,000 coronavirus-related deaths on a Saturday.
Entity	23%	58%	Mary of Teck was queen-consort.
Event	14%	14%	In the last EFL Cup, Manchester defeated Chelsea.
Other	6%	4%	Most genes need further research to better understand the function of their RNA products.

Table A.1: Estimated distribution of claims in the VITAMINC, based on manual annotations of 100 randomly sampled claims from the development split of the real and synthetic subsets. An example claim from each category is provided for reference.




Figure A.2: Probability density function of claim-evidence overlap for different labels in the dataset. The overlap is computed as the ratio of mutual bigrams in the two sentences.

All claims in the VITAMINC-synthetic are paired with one refuting and one supporting evidence, making it impossible for a claim-only to perform better than random. Each claim in the VITAMINC-real is paired with one refuting or neutral evidence, in addition to a supporting one. To evaluate whether models can utilize lexical cues in claims, we train a claim-only classifier on VITAMINC-real and find it to achieve 50% accuracy—the same as always predicting SUP.

A.4 Claim-evidence Word Overlap

Naturally, when pairing claims to evidence sentences, the overlapping words will be higher on average for claims with their supporting evidence. In VITAMINC dataset, we want to minimize this bias in order to create challenging examples that require sentence-pair inference and cannot be solved by simple word matching techniques. Therefore, we asked annotators, when possible, to avoid copying exact phrases from the evidence to the claim (see §3.2).

Figure A.2 shows the probability density function of bigram overlaps between the claim and evidence for each relation. Similar to FEVER, the overlap ratio of supporting pairs in the VITAMINC dataset is only slightly higher than the one of refuting pairs. Also, the overlap ratio of the NEI pairs of the VITAMINC real dataset is on average higher than FEVER.

B Experimental Setting

We implement all our models with the HuggingFace Transformers library (Wolf et al., 2019). When comparing across training datasets of different sizes, we train the model for the same amount of update steps, upsampling the smaller datasets. We pick the checkpoint with the highest accuracy on the development set of the training task and report performance on the test set. More details are available at <https://github.com/TalSchuster/VitaminC>

C GPT-3 Evaluation

The GPT-3 model has recently demonstrated impressive results in zero-shot and few-shot generation and classification tasks (Brown et al., 2020). This 175B parameters language model was trained on billions of words from online sources, including the English Wikipedia. As result, it can be applied on many tasks without any further fine-tuning—instead, one need only provide a task-specific prefix (i.e., “prompt”) with a few examples that direct the language model towards the desired output format. For example, GPT-3 achieves better than random results on ANLI with only a single example in the prompt, and over 40% accuracy with 50 examples (Brown et al., 2020).

We used OpenAI’s beta API to query GPT-3. Due to our limited quota, we could not perform

extensive experiments. Instead, we performed a qualitative evaluation using several examples from VITAMINC test set for the claim extraction (factually consistent generation) and the fact verification tasks. Therefore, these results should be viewed as exploratory only.

GPT-3 for Claim Extraction. We examine a two-shot setting for the claim extraction task. The model is asked to convert a revision into a short claim that expresses the fact that is true after the edit. To guide the model for this task, we provide a prompt with two random examples from the VITAMINC training set (see Figure C.1). One of the main concerns regarding large language models is the limited control it allows for ensuring that the facts in the generated output align with the source (Schuster et al., 2020). The generation tasks of VITAMINC provide a useful test-bed for evaluating the factual consistency with the input. Importantly, our VITAMINC-trained fact verification classifiers (f_{verdict}) allow strong automatic evaluation for the factual agreement of the generation with the source.

We use GPT-3 to extract claims for four revisions with a sampling temperature value (\mathcal{T}) set to either 0 or 0.7. The zero value is recommended for maximizing the factual consistency as the model follows its most certain predictions. Using low temperature, however, can result in less fluent generations (Holtzman et al., 2020). Therefore, high values of \mathcal{T} are also commonly used.

The results are reported in Tables 7 and E.3. With only two guiding examples, GPT-3 is able to follow the desired format and create a short claim. Yet, some of its generations follow s_{t-1} instead of s_t or add new, unsupported facts. f_{verdict} provides an indication for the factual correctness of the output. For example, it correctly classifies the output of the $\mathcal{T} = 0.7$ setting for the top example in Table 7 as “Not Enough Information” since GPT-3 reported about 20 deaths even though the input doesn’t mention death numbers at all.

We expect GPT-3 to improve with longer prompts or fine-tuning and leave this to future research due to our limited quota.

GPT-3 for Fact Verification. We also experiment with using GPT-3 few-shot classification capabilities for the fact verification task. We follow the ANLI few-shot format of Brown et al. (2020) and compose prompts with 6 examples (2 from each class) with random examples from VITAMINC

training set. We use only numerical examples to evaluate numerical claims (Figure C.3), and mixed examples for other claims (Figure C.2). We set $\mathcal{T} = 0$ as recommended for classification.

Table C.1 summarizes the results. Even with only six examples, GPT-3 seems to perform significantly better than random. Yet, its verdict is wrong in several cases that can be easily classified by humans. For example, we find it to refrain from predicting a True/False verdict even when the evidence is clear. We observe this both for a date-based (line 3.2 in Table C.1), numerical (lines 4.1-4.2), and entity-focused claims (line 5.2).

To experiment with the sensitivity of the model to the provided context, we manually modified some of the examples to provide even stronger evidence. For example, while GPT-3’s prediction for line 5.2 is acceptable as actually, Turner Broadcasting System merged with WarnerMedia in 1996, changing the evidence to another disconnected entity (The Walt Disney Company) did not change the prediction (line 5.3) as expected. Even when explicitly stating that there is no other owner GPT-3 didn’t modify its verdict (line 5.4). Similarly, when evaluating the claim about the population of Beaverton being less than 90K, GPT-3 ignores the supporting evidence and outputs a false verdict (lines 1.4-1.5). Changing the claim to state “approximately 86K” instead of “less than 90,000” modified the prediction to “Neither” (line 1.6). Only repeating the exact same number as the evidence led to a true verdict (line 1.7).

D Complementary Experiments

We report fact verification results with a fine-tuned BERT-base (Devlin et al., 2019) model in Table D.1. We find ALBERT-base to outperform BERT-base on most of the evaluated datasets. ALBERT-xlarge performed better than the two base models in all datasets except for Triggers. The Triggers dataset is very small (186 examples) and contains some unnaturally looking claims, which could explain the high variance across models.

E Example Outputs

We provide examples of predicted word-level rationales in Table E.1 and of outputs for the two generation tasks in Tables E.2 and E.3.

```

Life Is Peachy: Life Is Peachy is the{*| first -> second |}studio album by the
American nu metal band Korn, released on October 15, 1996 through both Immortal
Records and Epic Records.
Claim: Life Is Peachy is Korn's second studio album.
###
2020 coronavirus pandemic in Kerala: As of 14 March 2020, there are{*| 19 -> 22
|}confirmed cases of the virus and more than 4000 people are under surveillance in
Kerala.
Claim: As of 14 March, there have been more than 20 confirmed COVID-19 cases in
Kerala.
###
<Visualized edit>
Calim: <prediction>

```

Figure C.1: The prompt used for GPT-3 few-shot claim extraction.

```

Manchester is a major city and metropolitan borough in Greater Manchester,
England, with a population of 545,500 as of 2017 (5th most populous English
district).
Question: Manchester had a population of more than 540,000 in 2017 and was the
5th most populous English district. True, False, or Neither? True
###
As of March 2018, the apps have achieved more than 8 billion downloads.
Question: Talking Tom and Friends apps have less than 8 billion downloads. True,
False, or Neither? False
###
He won the Premier League in 2018.
John Stones won both the Premier League and EFL Cup in 2018. True, False, or
Neither? Neither
###
Neck Deep are a emo band.
Question: Neck deep is an emo band. True, False, or Neither? True
###
Critics generally gave The Final Frontier mixed to poor reviews.
Question: The film Star Trek V: The Final Frontier got negative reviews only.
True, False, or Neither? False
###
The series was favorably compared to the HBO series The Jinx and the podcast
Serial.
Question: The follow-up of the series Making a Murderer, was released in 2018.
True, False, or Neither? Neither
###
<Examined evidence>
Question: <Examined claim>. True, False, or Neither? <prediction>

```

Figure C.2: The prompt used for GPT-3 few-shot fact verification predictions on **non-numerical** claims (examples 2 and 4 in Table C.1. We follow the few-shot setting of [Brown et al. \(2020\)](#) for ANLI.

#	Claim	Evidence	GPT-3	Ours	Gold
1.1	Less than 90,000 people live in Beaverton , Oregon	its population is estimated to be 91,757 , almost 14% more than the 2000 census figure of 76,129	False	False	False
1.2	More than 90K people live in Beaverton	its population is estimated to be 91,757 , almost 14% more than the 2000 census figure of 76,129	True	True	True
1.3	More than 90K people live in Beaverton	its population is estimated to be 86,205 , almost 14% more than the 2000 census figure of 76,129	Neither	False	False
1.4	Less than 90,000 people live in Beaverton, Oregon	its population is estimated to be 86,205 , almost 14% more than the 2000 census figure of 76,129	False	True	True
1.5	Less than 90,000 people live in Beaverton, Oregon	Beaverton’s population is estimated to be 86,205	False	True	True
1.6	Approximately 86k people live in Beaverto, Oregon	Beaverton’s population is estimated to be 86,205	Neither	True	True
1.7	Approximately 86,205 people live in Beaverton, Oregon	Beaverton’s population is estimated to be 86,205	True	True	True
2.1	Diego Corrales’ father was Puerto Rican and his mother Dominican	Corrales was born to a African American father and a Mexican mother	False	False	False
2.2	Diego Corrales’ father was Puerto Rican and his mother Dominican	Corrales was born to a Puerto Rican father and a Dominican mother	True	True	True
3.1	COVID-19 outbreak was identified before December	The outbreak was first identified in Wuhan, Hubei, China in December 2019 and recognized as a pandemic	False	False	False
3.2	COVID-19 outbreak was identified before December	The outbreak was first identified in Wuhan, Hubei, China in 17 November 2019 and recognized as a pandemic	Neither	True	True
4.1	There have been more than 400 confirmed coronavirus cases in Germany	There have been 444 confirmed cases and 16 recoveries of coronavirus in Germany	Neither	True	True
4.2	There have been more than 400 confirmed coronavirus cases in Germany	There have been less than 349 confirmed cases and 16 recoveries of coronavirus in Germany	Neither	False	False
5.1	Cartoon Network is owned by Turner Broadcasting System	Cartoon Network is an American pay television channel owned by Turner Broadcasting System , a subsidiary of AT&T’s WarnerMedia	True	True	True
5.2	Cartoon Network is owned by Turner Broadcasting System	Cartoon Network is an American pay television channel owned by Warner Bros. Entertainment , a subsidiary of AT&T’s WarnerMedia	Neither	False	False
5.3	Cartoon Network is owned by Turner Broadcasting System	Cartoon Network is an American pay television channel owned by The Walt Disney Company	Neither	False	False
5.4	Cartoon Network is owned by Turner Broadcasting System	The Walt Disney Company is the only owner of Cartoon Network	Neither	False	False

Table C.1: GPT-3 fact verification predictions on examples from the VITAMINC test dataset (examples 1.5-1.7 and 5.3-5.4 were manually modified to examine the model’s behavior). We follow the few-shot setting of [Brown et al. \(2020\)](#) for ANLI (see Figures C.2 and C.3). The bold spans are for presentation and are not part of the input. Our VITAMINC-trained ALBERT classifiers predicted correctly on all these examples (though they weren’t picked this way). The GPT-3 few-shot succeeds on some examples and even expresses sensitivity to evidence in lines 2.1-2.2. In several cases, however, GPT-3 abstains from a True/False verdict, even when provided with strong evidence (see “Neither” predictions). Line 1.4 shows an example where GPT-3’s verdict is opposite of the provided evidence. Only when rephrasing the claim to exactly overlap with the evidence, it predicts an agreement.

Manchester is a major city and metropolitan borough in Greater Manchester, England, with a population of 545,500 as of 2017 (5th most populous English district).

Question: Manchester had a population of more than 540,000 in 2017 and was the 5th most populous English district. True, False, or Neither? True

###

As of March 2018, the apps have achieved more than 8 billion downloads.

Question: Talking Tom and Friends apps have less than 8 billion downloads. True, False, or Neither? False

###

As of January 2015, JFC had a total of more than 3,000 stores worldwide, with system-wide retail sales totaling 82.1 billion pesos for the fiscal year 2011.

Question: Jollibee had a total of more than 20,000 stores worldwide after January 2016. True, False, or Neither? Neither

###

As of March 2018, the apps have achieved more than 8 billion downloads.

Question: Talking Tom and Friends apps have over 8 billion downloads. True, False, or Neither? True

###

Bet365 has more than 35 million customers globally.

Question: Bet365 has less than 30 million customers worldwide. True, False, or Neither? False

###

The series was favorably compared to the HBO series The Jinx and the podcast Serial.

Question: The follow-up of the series Making a Murderer, was released in 2018. True, False, or Neither? Neither

###

<Examined evidence>

Question: <Examined claim>. True, False, or Neither? <prediction>

Figure C.3: The prompt used for GPT-3 few-shot fact verification predictions on **numerical** claims (examples 1 and 3 in Table C.1).

Model	Train dataset	VitC real	VitC syn	FEVER	MNLI	Adversarial	Symmetric	Triggers	ANLI	Contrast
BERT-base	FEVER	60.55	71.35	87.16	61.90	52.09	73.60	69.89	34.53	54.05
	MNLI	46.31	69.01	70.06	83.80	50.13	73.88	65.05	26.88	51.92
	FEVER + MNLI	56.24	81.80	95.59	85.06	63.05	85.11	37.63	29.63	60.63
	VitC	85.80	90.63	74.21	66.66	76.24	90.17	63.98	33.19	72.49
	VitC + MNLI	84.47	91.00	74.88	83.70	63.05	84.55	66.13	31.00	84.88
	VitC + FEVER	84.72	89.16	87.55	69.28	64.75	90.73	72.58	34.06	84.01

Table D.1: Fact verification Complementary results for Table 4 with a BERT-base model.

Claim	the youtube channel chuchu tv is placed 42nd and has more than 25 million subscribers .
Evidence	chuchu tv is the 43rd most subscribed youtube channel in the world , with over 20 million subscribers .
Claim	the ramus has sold less than 4.5 million albums worldwide .
Evidence	the rasmus has sold 5 million albums worldwide , 310,000 copies in their native finland alone .
Claim	darren randolph is spanish .
Evidence	humes dated irish footballer darren randolph in 2005 .
Claim	astravyets is near vilnius .
Evidence	his father may have migrated to the us in the 1860s from astravyets near vilnius .
Claim	the pace controlling stamina meter is a new feature in the game series .
Evidence	new to the series is a three-tier stamina meter which controls the pace of a match .
Claim	the movie will be released on 25 november 2015 .
Evidence	[...] are producing the film which columbia pictures will release on november 25 , 2015 .

Table E.1: Example masks produced by the word-level rationale model for identifying **anchoring words** in the evidence that are responsible for the classifiers verdict regarding the claim. Masking these words leads the classifier to predict NEI instead of what would have been SUP or REF.

Claim	Stephen Bruner was born on October 19 , 1984 .
s_{t-1}	Stephen Bruner (born October 18 , 1984) , better known by his stage name Thundercat , is an American multi-genre bass player , producer and singer fr om Los Angeles , California .
Output	Stephen Bruner (born October 19, 1984) , better known by his stage name Thundercat, is an American multi-genre bass player, producer and singer from Los Angeles, California.
s_t	Stephen Bruner (born October 19 , 1984) , better known by his stage name Thundercat , is an American multi-genre bass player , producer and singer from Los Angeles , California .
Claim	The House that Jack Built has an approval rate of over 45 % and over 13 critics .
s_{t-1}	On Rotten Tomatoes , the film has an approval rating of 42 % , based on 12 critics , and an average rating of 5.9/10
Output	On Rotten Tomatoes, the film has an approval rating of 47 % , based on 14 critics, and an average rating of 5.9/10.
s_t	On Rotten Tomatoes , the film has an approval rating of 47 % , based on 15 critics , and an average rating of 6.3/10.
Claim	Cartoon Network is owned by Warner Bros. Entertainment .
s_{t-1}	Cartoon Network is an American pay television channel owned by Turner Broadcasting System , a subsidiary of AT & T 's WarnerMedia .
Output	Cartoon Network is an American pay television channel owned by Warner Bros. Entertainment, a subsidiary of AT & T's WarnerMedia.
s_t	Cartoon Network is an American pay television channel owned by Warner Bros. Entertainment , a subsidiary of AT & T 's WarnerMedia .
Claim	The New York Times reviewed the film Mad Max : Fury Road .
s_{t-1}	Lindsay Bahr of The Associated Press wrote , “ Miller has reminded us that blockbusters have the potential to not only be art , but radically visionary – even the fourth in a series .
Output	Lindsay Bahr of The New York Times wrote, “ Miller has reminded us that blockbusters have the potential to not only be art, but radically visionary – even the fourth in a series.
s_t	The New York Times wrote , “ Miller has reminded us that blockbusters have the potential to not only be art , but radically visionary – even the fourth in a series .

Table E.2: Example outputs of the BART-base used for generating factually consistent revisions given the old version s_{t-1} and the updated claim we wish to support. The “ground-truth” s_t is provided for reference.

(s_{t-1}, s_t)	WWE 2K15: As of August 2015 , WWE 2K15 has shipped over{* six -> 7.5 }million units.	
BART (VitC)	WWE 2K15 sold more than 7 million units.	$f_{\text{verdict}}(c, s_t) = \text{SUP}$
GPT-3 $\mathcal{T}=0$	As of August 2015, WWE 2K15 has shipped over seven million units	SUP
GPT-3 $\mathcal{T}=0.7$	As of August 2015, WWE 2K15 has shipped over seven million units	SUP
Reference	By August 2015 , WWE 2K15 shipped over 7.5 million units .	SUP
(s_{t-1}, s_t)	Pat Jennings: He has played for League of Ireland clubs UCDDerry and is now at {* Shamrock Rovers -> Dublin }.	
BART (VitC)	Pat Jennings is currently playing for Dublin club UCDDerry.	SUP
GPT-3 $\mathcal{T}=0$	He has played for League of Ireland clubs UCD and is now at Shamrock Rovers	REF
GPT-3 $\mathcal{T}=0.7$	He played for Shamrock Rovers and is now at Dublin	REF
Reference	Pat Jennings currently plays for the Dublin club .	SUP

Table E.3: Additional examples for Table 7. Example outputs for extracting claims that express the factual change in a Wikipedia revision. The BART-base model is trained on VITAMINC data and GPT-3 is applied in a 2-shot setting with a temperature of 0 or 0.7. The revision (s_{t-1}, s_t) is given to the model as a single sentence visualization where the edits are between curly brackets, preceded by the article’s title. The human-written claim is provided for reference. The prediction of our ALBERT-xlarge VITAMINC-trained model $f_{\text{verdict}}(c, s_t)$ on the generated claim against s_t is also reported in the rightmost column.