

# Converge to the Truth: Factual Error Correction via Iterative Constrained Editing

Jiangjie Chen<sup>1\*†</sup>, Rui Xu<sup>1\*</sup>, Wenxuan Zeng<sup>2</sup>, Changzhi Sun<sup>3‡</sup>, Lei Li<sup>4</sup>, Yanghua Xiao<sup>1,5‡</sup>

<sup>1</sup>Shanghai Key Laboratory of Data Science, School of Computer Science, Fudan University

<sup>2</sup>University of Electronic Science and Technology of China

<sup>3</sup>ByteDance AI Lab

<sup>4</sup>University of California, Santa Barbara

<sup>5</sup>Fudan-Aishu Cognitive Intelligence Joint Research Center

jjchen19@fudan.edu.cn, ruixu21@m.fudan.edu.cn, wenxuan.zeng@std.uestc.edu.cn,

sunchangzhi@bytedance.com, leili@cs.ucsb.edu, shawyh@fudan.edu.cn

## Abstract

Given a possibly false claim sentence, how can we automatically correct it with minimal editing? Existing methods either require a large number of pairs of false and corrected claims for supervised training or do not handle well errors spanning over multiple tokens within an utterance. In this paper, we propose VENCE, a novel method for factual error correction (FEC) with minimal edits. VENCE formulates the FEC problem as iterative sampling editing actions with respect to a target density function. We carefully design the target function with predicted truthfulness scores from an offline trained fact verification model. VENCE samples the most probable editing positions based on back-calculated gradients of the truthfulness score concerning input tokens and the editing actions using a distantly-supervised language model (T5). Experiments on a public dataset show that VENCE improves the well-adopted SARI metrics by 5.3 (or a relative improvement of 11.8%) over the previous best distantly-supervised methods.

## 1 Introduction

In the digital age, misinformation with textual *factual errors* is spreading widely and quickly with the help of social media. However, current text correction approaches mostly focus on *grammatical error correction* (GEC) (Yuan and Briscoe 2016; Chollampatt and Ng 2018; Omelianchuk et al. 2020; Yasunaga, Leskovec, and Liang 2021). The automatic correction of *factual errors* may help prevent them from fueling misunderstanding and misleading decision-making, attracting an increasing amount of research attention. Given a false claim and its evidence texts from trustworthy knowledge bases, *factual error correction* (FEC) aims to correct the non-factual text spans of the claim, so that it is better supported by evidence. For the example in Figure 1, an FEC system needs to correct that the movie that Will Smith starred in in 2006 is *The Pursuit for Happiness*, not *The Truman Show*.

\* Authors contributed equally.

† Work done during internship at ByteDance AI Lab.

‡ Corresponding authors.

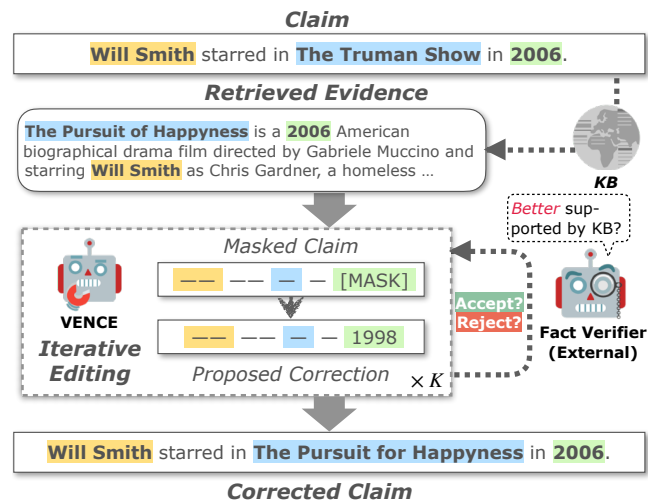


Figure 1: An overview of VENCE, an FEC method based on Metropolis-Hastings sampling. Given a claim and retrieved evidence texts, VENCE iteratively corrects the claim by claim masking and correction proposal. Based on whether a new proposal is better supported by the evidence, an external fact verifier contributes to the decision of proposal acceptance. After the sampling is converged, the most supported proposal is selected as the corrected claim.

The main factor limiting the application of FEC systems is the cost to obtain high-quality fine-grained annotations, making it relatively under-explored. To gain more training data, instead of annotating claims paired with their corrections, researchers (Shah, Schuster, and Barzilay 2020; Thorne and Vlachos 2021) start using distant supervision from fact verification (FV) datasets (Thorne et al. 2018). Specifically, for a claim SUPPORTED by evidence in the FV dataset, they first mask some tokens in the claim and train the FEC system to restore the original claim from the masked one.

Existing distantly supervised models usually follow a one-pass mask-then-correct paradigm (Shah, Schuster, and Barzilay 2020; Thorne and Vlachos 2021). The masker is designed to find problematic spans in the claim, where token-level

explanations of the FV model are usually exploited as masks. The corrector generates the corrected sentence given the masked claim. However, those models are limited due to two reasons: 1) The problem of semantic drift, *i.e.*, over-correcting the claim and making the intended meaning of the claim broken. 2) As a response to the lack of annotated data for FEC, the supervision from fact verification is under-utilized.

To address these problems, we propose to integrate the advantages of both *fact verification* and *iterative text editing* for FEC. *Our first key insight* is to correct errors in the manner of iterative editing. We break the correction process into unit-level (*e.g.*, tokens and entities), to overcome the over-erasure problem and explore and revise more correction choices. *Our second key insight* is to bridge FV with FEC in the iterative editing framework, where FV offers control and guidance to the correction in each editing iteration. Since resources for the FV task are significantly richer than that of FEC, an FEC system can then be able to iteratively edit small errors under the constraint of fact verification, guiding towards correctness during each iteration.

Motivated by these ideas, in this paper, we propose VENCE, a VERIFICATION-guided CONSTRAINED EDITING system for factual error correction. To realize our ideas, an FEC model should be able to utilize external constraints (FV) without direct supervision from FEC. Thus, we build our system based on the framework of constrained text generation with Metropolis-Hastings (MH) sampling (Metropolis et al. 1953), a typical Markov chain Monte Carlo method, which proved to be effective on various unsupervised constrained text generation tasks (Miao et al. 2019; Zhang et al. 2020a; Chen et al. 2022b; Mireshghallah, Goyal, and Berg-Kirkpatrick 2022). Compared with previous work, VENCE utilizes an external FV model for guidance in two ways: 1) the token gradients from the FV model guide the finding of edit tokens, and 2) the truthfulness score by the FV model contributes to a probabilistic energy-based model (Hinton 2002) that decides the acceptance of a correction proposal in each iteration. To maintain the fluency and minimum modification to the original claim (as opposed to over-erasure), we also consider a language modeling score and hamming distance in the energy model, respectively. To overcome the challenge of editing multi-token entities, we separate token space with entity space, and train a multi-tasking T5 model (Raffel et al. 2020) to propose from distributions of both spaces.

To summarize, this work contributes to solving factual error correction without direct supervision. To our knowledge, we are the first to adopt an iterative text editing method (VENCE) for this task, which alleviates the over-erasure problem in previous methods. Also, VENCE enjoys a more powerful error revision ability by effectively integrating external but coarse-grained verification signals during each editing iteration. Experimental results demonstrate the effectiveness of VENCE over previous methods, which achieves a new state-of-the-art for distantly supervised FEC task.<sup>1</sup>

<sup>1</sup>Resources of VENCE can be found at <https://github.com/jiangjiechen/VENCE>.

## 2 Related Work

**Text Factuality** The factuality of text has been widely studied by the community. During the automatic production of text, text generation systems are prone to hallucinate, harming the factuality of generated text (Raunak, Menezes, and Junczys-Dowmunt 2021; Xiao and Wang 2021). Methods have been proposed to alleviate hallucination (to the internal or external knowledge) in generation for text summarization (Cao et al. 2018; Cao, Dong, and Cheung 2022), dialogue systems (Honovich et al. 2021), text simplification (Devaraj et al. 2022), etc. In the post-verification stage, fact verification (or fact checking) aims to evaluate factuality of existing texts, which leverages external evidence to verify a given claim (Thorne et al. 2018; Augenstein et al. 2019; Wadden et al. 2020; Aly et al. 2021). It is worth noting that, many recent studies focus on interpretable fact verification (Stammach and Ash 2020; Kotonya and Toni 2020; Samarinas, Hsu, and Lee 2021; Wu et al. 2021), which is related to the objective of FEC - understanding and making the claims better supported by evidence. For example, ProoFVer (Krishna, Riedel, and Vlachos 2021) is trained to generate intermediate natural logic proofs for verification, where the proofs are annotated, while LOREN (Chen et al. 2022a) decomposes claim verification at phrase-level by logic and learns phrase veracity predictions without any intermediate annotations. Also aiming for the text factuality problem, we tackle it from the perspective of post-correction, which verifies and corrects nonfactual texts.

**Textual Error Correction** According to error types, research on textual error correction can be categorized as grammatical error correction (GEC) and factual error correction (FEC). GEC (Ng et al. 2014; Omelianchuk et al. 2020; Ishii, Tamura, and Ninomiya 2021; Qorib, Na, and Ng 2022) aims to correct grammatical errors without changing the semantic content of each sentence. Existing literature generally relies on human-labeled data, while Yasunaga, Leskovec, and Liang (2021) study GEC under an unsupervised paradigm, which leverages a critic model to assess the input and learns to repair bad examples without labeled data. Unlike GEC, FEC needs to modify the semantic content of sentences to achieve factual consistency against trustworthy knowledge sources, either textual corpus (Shah, Schuster, and Barzilay 2020; Thorne and Vlachos 2021) or structured tables (Iso, Qiao, and Li 2020). In this work, we study the problem of FEC based on textual knowledge bases.

**Text Editing** Many research efforts have been made to text editing (Kasner and Dušek 2020; Higashiyama et al. 2021; Agrawal and Carpuat 2022), especially for tasks that source and target text that has high overlap, such as sentence simplification (Inui et al. 2003; Mallinson et al. 2022), translation post-editing (Gu, Wang, and Zhao 2019; Lee et al. 2020), story rewriting (Chen et al. 2022b), text style transfer (Hu et al. 2022), etc. To impose constraints during editing, in recent years, a prevailing way is to adopt Markov chain Monte Carlo sampling for text generation. Miao et al. (2019) first utilize Metropolis-Hastings (MH) sampling for constrained text generation, which edits discrete tokens under the desired stationary distributions as constraints. As an extension,

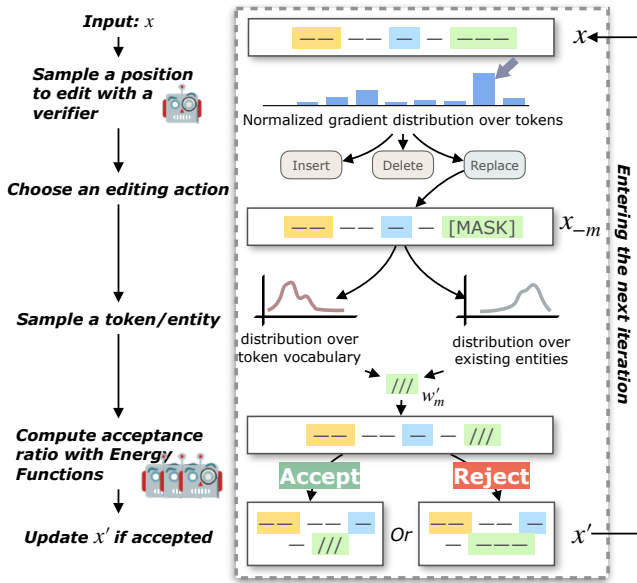


Figure 2: The editing procedure of VENCE of each iteration.

Zhang et al. (2020a) manage to satisfy combinatorial constraints for text generation, and Mireshghallah, Goyal, and Berg-Kirkpatrick (2022) propose a global score-based model that also leverages MH sampling for controllable text generation. We follow this line of work and extend with verification signals as constraints and guidance for FEC, and extends token-level editing to entity-level for semantic consistency.

### 3 Method

In this section, we present our FEC method VENCE. We will first sketch an overview in § 3.1. Then, we will detail the implementations in the sketch in § 3.2-3.4.

#### 3.1 Overview

The corrected sentence should be fluent, truthful and minimally edited, which are the constraints to be satisfied during correction. To this end, we adopt a text editing framework based on Metropolis-Hastings sampling (Metropolis et al. 1953; Miao et al. 2019), a MCMC method. The MH algorithm is able to ensure the edited text to be converged where the above constraints are satisfied, through the custom definition of the *stationary distribution* of the Markov chain.

Given an original claim  $x^0$  and retrieved evidence texts  $E$ , we iteratively execute steps of *mask-then-correct* over the claim. The correction proposal for the claim ( $x = \{w_1, \dots, w_{|x|}\}$ ) in current iteration has the following steps:

1. Sample an edit position  $m$  to be masked within  $x$ , i.e.,  $P_1(m | x)$ ,  $m = 1, 2, \dots, |x|$ . This yields a masked claim  $x_{-m} = \{w_1, \dots, w_{m-1}, [\text{MASK}], \dots, w_{|x|}\}$  (§ 3.3);
2. Sample an edit action  $a$  (*insert*, *delete* or *replace*), i.e.,  $P_2(a)$  (§ 3.4);
3. Sample a corrected token/entity according to the masked claim  $x_{-m}$  and the sampled edit action  $a$ , and get  $x'$ , i.e.,  $P_3(x' | x_{-m}, a)$  (§ 3.4).

Together, we have the *transition distribution*  $g(x' | x)$  in the Markov chain taking the action  $a$  to edit position  $m$ :

$$g(x' | x) = P_1(m | x)P_2(a)P_3(x' | x_{-m}, a). \quad (1)$$

Note that we omit evidence  $E$  in the condition for brevity.

In order to make sure the sampling to be converged on the *stationary distribution*  $\pi(x)$ , the acceptance of the proposed  $x'$  should be decided by the *acceptance ratio*  $A(x' | x)$ , and the accepted  $x'$  will enter the next iteration. We implement  $\pi(x)$  with the target objectives of FEC (factual, fluent and minimally modified), and use a probabilistic energy-based model  $\mathcal{E}(x)$ , where  $\pi(x) = \frac{e^{-\mathcal{E}(x)}}{\mathcal{Z}}$  (Mireshghallah, Goyal, and Berg-Kirkpatrick 2022), the Boltzmann distribution of the energy functions (§ 3.2). Thus,  $A(x' | x)$  is defined as in the MH algorithm:

$$\begin{aligned} A(x' | x) &= \min\left\{1, \frac{\pi(x')g(x | x')}{\pi(x)g(x' | x)}\right\} \\ &= \min\left\{1, \frac{e^{-\mathcal{E}(x')}}{e^{-\mathcal{E}(x)}}\frac{g(x | x')}{g(x' | x)}\right\}. \end{aligned} \quad (2)$$

Finally, the accepted texts are ranked based on  $\mathcal{E}(x)$ , and the highest one is chosen as the output. We visualize the process within each iteration in Figure 2. This process goes on until converged (20 iterations in practice). Keeping this workflow in mind, next, we will place the pieces together.

#### 3.2 Energy Functions for Target Objectives

First and foremost, we start by defining the target objectives in FEC. A corrected claim  $x$  should be fluent ( $\mathcal{E}_{\text{LM}}(x)$ ), truthful according to trustworthy evidence ( $\mathcal{E}_{\text{V}}(x)$ ) and edited with minimum modifications to preserve the intended meaning ( $\mathcal{E}_{\text{H}}(x)$ ), therefore we have an aggregated energy function  $\mathcal{E}(x)$  defined as:

$$\mathcal{E}(x) = \mathcal{E}_{\text{LM}}(x) + \mathcal{E}_{\text{V}}(x) + \mathcal{E}_{\text{H}}(x). \quad (3)$$

**Language Modeling Score** Constrained text generation with MH sampling should first guarantee text fluency, as shown by various previous work (Miao et al. 2019; Chen et al. 2022b; Mireshghallah, Goyal, and Berg-Kirkpatrick 2022). Following Goyal, Dyer, and Berg-Kirkpatrick (2022); Mireshghallah, Goyal, and Berg-Kirkpatrick (2022), we implement the energy function as

$$\mathcal{E}_{\text{LM}}(x) = - \sum_i \log P_{\text{MLM}}(w_i | x_{-i}) \quad (4)$$

using a masked language model (MLM) with a BERT-base (Devlin et al. 2019), which is the sum of the predicted logits over every masked token  $w_i$ . Note that this score can also be implemented with a causal language model such as GPT-2 (Radford et al. 2019).

**Truthfulness Score** An important point of view of our method is to incorporate verification signals deeply into each edit iteration. Therefore, the truthfulness score is indispensable for the energy function. To achieve this goal, we use an off-the-shelf discriminative verification model  $P_{\text{V}}$  that predicts whether the current claim  $x$  is supported by the evidence

$E$ , yielding the probability of the SUPPORTED label. Thus, the energy function is defined as:

$$\mathcal{E}_V(x) = -\log P_V(\text{SUPPORTED} \mid x, E). \quad (5)$$

In this way, powerful verifiers can be exploited to control the factuality of each correction proposal, since resources for verification datasets are generally much richer than that of the FEC task.

**Hamming Distance** One tricky pitfall for this task is that the corrections could be too often to keep the intended meaning of the original claim. For example, one can replace with a brand new correct text but completely different than the previous claim. Therefore, we also aim to constrain the number of edits of the proposed text, where we calculate a hamming distance between current claim  $x$  and the original claim  $x^0$ ,

$$\mathcal{E}_H(x) = \text{HAMMINGDISTANCE}(x, x^0). \quad (6)$$

### 3.3 Editing Position Proposal

In order to implement the transition distribution  $g(x' \mid x)$ , we start by designing  $P_1(m \mid x)$ , where only a handful of masks are applied during each iteration.

To find the position to edit, the MH algorithm generally randomly samples from the token sequence and edits a token at each time. However, this yields two challenges for the FEC task: 1) most tokens do not need to be changed: *how to accelerate the sampling process with only necessary edits?* and 2) factual errors usually exist in the wrong multi-token entities: *how to edit consecutive tokens in an entity in MH sampling?*

**Sampling from Gradient-based Distribution** To address the first challenge, we assume that the gradient by the verification model  $P_V$  of erroneous tokens should be larger than the correct ones, inspired by Sha (2020). Therefore, we reuse the  $\mathcal{E}_V$  from § 3.2 and substitute the random sampling with a normalized distribution over tokens of the token gradient calculated by back-propagating  $P_V$ . Then we have

$$P_1(i \mid x) \propto \|\nabla_{w_i} \mathcal{E}_V(x)\|_2 \quad (7)$$

where  $w_i$  is the word representation of the token  $w_i$  and  $\|\cdot\|_2$  denotes vector 2-norm operation.

**Multi-token Entity Masking** As for the second challenge, we believe tokens in an entity should be edited together during each proposal. Yet, sampling with one token at a time breaks the semantic meaning of the entity. To solve this problem, we choose to mask all the tokens in each named entity at a time. To do so, we first identify the named entities in the claim using an off-the-shelf NER model.<sup>2</sup> When any token within an entity is sampled, we select the whole entity for correction, *i.e.*, “Will Smith starred in [MASK] in 2006”, other than “Will Smith starred in The [MASK] show in 2006”. Thus, the probability of an entity  $e$  to be masked is the sum of the probability of its tokens, *i.e.*,  $\sum_{w_i \in e} P_1(i \mid x)$ . The burden of correcting the multi-token entity is now shifted to the edit proposal step (§ 3.4), which we will soon discuss.

<sup>2</sup><https://huggingface.co/Jean-Baptiste/camembert-ner>

### 3.4 Editing Proposal

After getting the masked sequence  $x_{-m}$  in each iteration, we will *randomly* sample an action from the uniform distribution  $P_2(a)$  of *insertion*, *deletion* and *replacement*, and propose a token/entity from  $P_3(x' \mid x_{-m}, a)$ .

However, with the aforementioned multi-token masking, challenges arise during sampling. The Markov chain must satisfy *detailed balance condition* in order for the sampling to converge on the stationary distribution  $\pi(x)$  with the target objectives (§ 3.2). To this end, the transition distribution  $g(x' \mid x)$  must be *reversible*. And yet, predicting an entity as a list of tokens given one mask is not reversible. For example, the state with a broken entity (such as “Will Smith starred in *The Pursuit for*”) will never be reached during a multi-token insertion, but it is necessary during a token-by-token deletion. Therefore, we separate entity space and token space with a generative proposal model, which we will next introduce. The correctness of our approach will be clarified in § 3.5.

**Generative Proposal Model: Token vs. Entity** Traditional proposal models (MLMs) have difficulty in dealing with multi-token prediction of a masked sentence. Instead of widely-used masked language models for token proposal, we fine-tune a sequence-to-sequence generative language model (T5, Raffel et al. 2020) to propose multi-token corrections. We separate the sampling spaces for entities and tokens, and define two generation tasks: token generation and entity generation, which are distinguished by adding different prefixes, a property enjoyed by T5.<sup>3</sup> During the training of T5, we use distant supervision by SUPPORTED claims. The model takes as input a masked sentence of  $a$  and its evidence, and outputs either a token or entity substituting that mask. The output choice depends on the types of correction action, which we will explain shortly.

With the proposal model, the transition distribution of corrected token/entity (denoted as  $w'_m$ , and  $x' = w'_m + x_{-m}$ ) is calculated accordingly. For the single token proposal, the probability distribution  $P_3^{\text{tok}}(w'_m \mid x_{-m}, a)$  naturally equals the vocabulary distribution of the token to be generated. For entity proposal, however, the probability distribution  $P_3^{\text{ent}}(w'_m \mid x_{-m}, a)$  should be computed over the possible entity space, which we approximate by *taking all the named entities* from both evidence  $E$  and the original claim  $x^0$ . Thus, the distribution can be acquired by normalizing the perplexity of each entity predicted by T5 in the candidate entity space.

**Replacement** For replacement ( $a = \text{rep}$ ), we define  $P_3(x' \mid x_{-m}, \text{rep})$  is either  $P_3^{\text{ent}}(x' \mid x_{-m}, \text{rep})$  or  $P_3^{\text{tok}}(x' \mid x_{-m}, \text{rep})$ , because T5 replaces the mask either with entity generation or token generation, depending on whether the mask is within an entity or not.

**Insertion** For *insertion* ( $a = \text{ins}$ ), we reload the notation  $x_{-m} = \{w_1, \dots, w_m, [\text{MASK}], w_{m+1}, \dots, w_{|x}]\}$ . In other words, we insert a [MASK] after the position  $m$  and keep the original tokens unchanged. Since we do not know whether an entity or a token should be generated, we take a weighted distribution of the two, *i.e.*,  $x$ , where  $\alpha$  is set as 0.5 in practice.

<sup>3</sup>Example prefixes: “*substituted word:*” and “*substituted entity:*”.

Label	# Train	# Valid	# Test
SUPPORTED	37,961	1,477	1,593
REFUTED	20,075	2,091	2,289

Table 1: Statistics of FECDATA (Thorne and Vlachos 2021), with data sample counts of each split and label.

**Deletion** For *deletion* ( $a = \text{del}$ ), the returned text is  $x' = \{w_1, \dots, w_{m-1}, w_{m+1}, \dots, w_{|x|}\}$ , where the token/entity  $w_m$  is deleted. So,  $P_3(x' | x_{-m}, \text{del}) = 1$ , and its reverse is actually *insertion*.

### 3.5 Convergence Analysis

We now give an analysis of the reversibility of the Markov chain, in response to the challenge mentioned at the beginning of § 3.4. In our implementation, the state transition within entities or tokens can naturally be reversible due to the separation of two spaces. To communicate two spaces, a delete-plus-insert action combination can be conducted. For example, when replacing an entity with a token, we can first delete the entity and then insert a token, thanks to the insertion action that balances two spaces. As a result, we can theoretically ensure the correctness and convergence of our implementation.

## 4 Experimental Setup

### 4.1 Evaluation

**Dataset** In this work, we evaluate our method using the dataset proposed by Thorne and Vlachos (2021), which is an evidence-based FEC task (referred to as FECDATA for brevity). FECDATA is collected from the intermediate annotations in FEVER (Thorne et al. 2018), a large and widely used fact verification dataset. A claim in FEVER is verified to be SUPPORTED or REFUTED by evidence, or labelled to be unverifiable, *i.e.*, NOTENOUGHINFO.

FECDATA takes the SUPPORTED and REFUTED samples from FEVER. Based on the REFUTED samples, FECDATA utilizes the mutated text spans that convert sentences from Wikipedia into REFUTED claims. Statistics of FECDATA are reported in Table 1. Note that, SUPPORTED instances are also included in the test set to test the models for keeping a correct claim as it is.

**Metrics** For automatic evaluation, we primarily use the SARI metric (Xu et al. 2016), which measures the goodness (F1) of words that are *added*, *deleted* or *kept* by the system from the source, compared with referenced ground truth. SARI-*final* is the average score of the three. We also use ROUGE score (RG-2) (Lin 2004) for evaluating the information recalled from the reference. Thorne and Vlachos (2021) show that, statistically, SARI scores and ROUGE score are *highly* correlated with human judgments in this task, especially SARI, according to Pearson’s  $r$ .

### 4.2 Baseline Methods

We adopt baselines of two settings: supervised and distantly supervised methods. The distinction between them lies in

whether direct supervision for FEC (the refuted claims and their corrections) is used for model training.

**Supervised Baselines** We train two models with REFUTED claims and their corrections from the training set of VENCE, covering paradigms of autoregressive generation and editing:

- T5 (Raffel et al. 2020), a strong pre-trained baseline that excels in sequence-to-sequence tasks. We fine-tune it by concatenating the original claim and evidence texts as input to generate a corrected claim.
- EdiT5 (Mallinson et al. 2022), a T5-based model that combines the best of non-autoregressive editing and autoregressive generation. The editing operations in EdiT5 include tagging, reordering, and insertion.

**Distantly Supervised Baselines** Since resources for FEC are limited, previous work mostly focuses on distantly supervised methods to address this problem. Specifically, two general types of strategies are used in previous work for FEC:

1. DS-1: evidence-based mask token/span prediction trained on SUPPORTED data samples (Shah, Schuster, and Barzilay 2020; Thorne and Vlachos 2021; Chen et al. 2022a), a denoising training objective similar to that used in BERT (Devlin et al. 2019), BART (Lewis et al. 2020) and T5; and
2. DS-2: distant supervision from external discriminative tasks, such as natural language inference (NLI) (Williams, Nangia, and Bowman 2018) and fact verification (Thorne et al. 2018) (which is also a type of NLI).

Applying these strategies (either or both), our distantly supervised baselines include:

- Masked Language Model (BERT, Devlin et al. 2019), applying DS-1 to greedily decode masked tokens.
- Dual encoder pointer network (2EncPtr for short) (Shah, Schuster, and Barzilay 2020), applying both DS-1 and DS-2 by querying a fact-checking model to predict the words to be masked, then conditionally replacing the masks with ground truth evidence.
- T5 Masker-Corrector (T5MC for short) (Thorne and Vlachos 2021), which is the previous SOTA and thus our primary baseline. T5MC follows a masker-corrector pipeline: SUPPORTED claims in FECDATA are used to apply DS-1 to train a corrector; during inference, masks are *heuristically* decided, where text spans that do not appear in evidence texts will be masked.
- T5 Masker-Corrector with verification signals (T5MC-V for short) (Thorne and Vlachos 2021), which is an extension to T5MC. T5MC-V follows (Shah, Schuster, and Barzilay 2020) and uses an external fact verification model trained on FEVER to decide the tokens to be masked. Thus, T5MC-V applies both DS-1 and DS-2.

What if these one-pass FEC baselines also have  $K$  candidates (same as the number iterations in VENCE) to be verified by FV? To make a fairer comparison, we upgrade T5MC and T5MC-V by taking  $K = 20$  generated sequences of them, and use the same FV model to rank for the best one, denoted as “Model+enumerate”.

Method	Verifier	SARI (%)				RG-2
		Keep	Delete	Add	Final	
<i>Fully Supervised</i>						
T5-base	-	79.6	90.2	59.2	76.4	72.7
EdiT5-base	-	<b>81.8</b>	<b>93.0</b>	<b>63.4</b>	<b>79.4</b>	<b>76.9</b>
<i>Distantly Supervised</i>						
MLM	-	56.1	52.9	7.8	38.9	42.7
2EncPtr	BERT <sub>b</sub>	34.5	48.1	1.7	28.1	34.8
T5MC	-	65.2	62.7	15.5	47.8	50.3
+enumerate	BERT <sub>b</sub>	66.2	<b>64.3</b>	17.1	49.2	51.2
T5MC-V	BERT <sub>b</sub>	61.1	54.3	19.4	44.9	42.0
+enumerate	BERT <sub>b</sub>	63.0	55.7	24.1	47.6	45.5
VENCE	BERT <sub>b</sub>	66.0	60.1	34.8	53.6	57.7
	RoBERTa <sub>l</sub>	<b>67.1</b>	61.9	<b>36.0</b>	<b>55.0</b>	<b>59.1</b>

Table 2: The automatic evaluation results of VENCE compared with baselines. Distantly supervised methods with verifiers apply the DS-2 strategy during training. Verifier<sub>b</sub> and Verifier<sub>l</sub> denote the base and large version of the pre-trained language model that the Verifier uses.

### 4.3 Implementation Details

**Evidence Retrieval** Since the retrieval of evidence is not the primary focus of this work, we adopt an existing approach and keep the retrieved evidence to be the same for all systems. Following previous work (Thorne and Vlachos 2021), the retrieval module consists of two main steps: 1) given a claim, the contained entities are predicted using GENRE (Cao et al. 2021), an autoregressive entity linking system, thus related pages are retrieved from Wikipedia; then 2) top- $k$  ( $k = 2$ ) passages are selected based on the similarity between the claim and these pages using DPR (Karpukhin et al. 2020), a dense retriever based on the BERT encoder.

**Training Details** We train the verifier (NLI) models and the proposal models based on the checkpoints hosted by HuggingFace (Wolf et al. 2020). Most of the training scripts and parameters are by default of the `transformers` library from HuggingFace. Specifically, 1) For the verifier (*i.e.*, NLI models), the BERT (base version) and RoBERTa (large version) checkpoints are fine-tuned on the FEVER dataset. We trained the NLI model for 5 epochs at a learning rate of  $1 * 10^{-5}$ . 2) For the proposal model, we fine-tune a T5 (base version) model with randomly masked SUPPORTED claims and the corresponding evidence in the FECDATA dataset. We train the model for 5 epochs at the learning rate of  $3 * 10^{-5}$ .

## 5 Results and Analysis

### 5.1 Overall Performance

We present the main results in the FECDATA dataset in Table 2. Generally, we have the following findings:

Our method outperforms previous distantly supervised baselines by large margins, including previous state-of-the-art (T5MC, Thorne and Vlachos 2021). VENCE achieves 53+ SARI final score and 57+ ROUGE-2 score. However, supervised methods still beat distantly supervised ones, envisioning a large room for improvement. Compared with

Dataset	Verifier	Acc (%)	SARI (%)			
			Keep	Delete	Add	Final
MultiNLI	BERT <sub>b</sub>	84.6	63.9	57.0	30.1	50.3
	RoBERTa <sub>l</sub>	<b>90.2</b>	<b>65.1</b>	<b>58.9</b>	<b>32.2</b>	<b>52.0</b>
FEVER	BERT <sub>b</sub>	71.7	66.0	60.1	34.8	53.6
	RoBERTa <sub>l</sub>	<b>72.9</b>	<b>67.1</b>	<b>61.9</b>	<b>36.0</b>	<b>55.0</b>

Table 3: Ablation results of the verification model used in VENCE w.r.t. model sizes and trained datasets. We report accuracy of verifier models on test set of MultiNLI and FEVER, respectively.

end-to-end text generation (T5), the EdiT5 wins due to the editing nature of the FEC task.

When adopting DS-2 (*i.e.*, with a verifier), we notice that the baselines do not fully exploit the verification signals for FEC. Especially, T5MC suffers from performance degradation when using a verifier to help select masks. In contrast, VENCE with a better verifier (RoBERTa-large, Liu et al. 2019) consistently achieves better results (+1.4 in SARI final score and +1.4 ROUGE-2 score). This indicates the success of our deep integration of verification guidance for FEC.

We notice that VENCE achieves more significant improvement on SARI-Add, indicating the added words of VENCE make much more sense than the baselines. This is because iterative correction offers a more delicate control over the generated text, whereas one-time correction cannot settle the number and conflict of proposed words for each mask.

When more corrections are enumerated and ranked, baselines results do improve 1.4 SARI points (T5MC) and 2.7 SARI points (T5MC-V), but not as good as VENCE. This shows the effectiveness of the iterative editing framework and deep integration of FV in VENCE.

### 5.2 Analysis

In response to the motivations of VENCE, we conduct detailed analyses to further understand VENCE and why it works. In particular, the analysis is designed to find out whether and how the verification model and the iterative editing process benefit FEC. For the rest of the analysis, we use a BERT-base verifier in VENCE, unless specified.

**How does verification affect correction?** To answer the first question, we explore different sets of verifiers in VENCE, each of them alters the gradient-based position finding and the energy model in the correction process. We conduct an ablation study of 1) the verification ability of the verifier models, instantiated with types and sizes of PLMs (BERT and RoBERTa); and 2) the datasets that the verifiers are trained upon, instantiated with MultiNLI (Williams, Nangia, and Bowman 2018) and FEVER (Thorne et al. 2018).

We draw two findings based on the results in Table 3: 1) Consistent with findings from Table 2, better verification performance contributes to better correction performance, which holds for both MultiNLI and FEVER verifier; 2) Even using an out-of-domain NLI model as a verifier, VENCE still achieves better performance than baselines, demonstrating




Figure 3: Correction results at each iteration of VENCE and its variants. Results on iteration 0 come from input claim.

Row	Verif.	Hamm.	LM	SARI-F $\uparrow$	HammS. $\downarrow$	BERTS. $\uparrow$
1	✓			41.2	3.42	0.90
2		✓		25.3	2.08	0.76
3			✓	37.1	<b>1.19</b>	0.82
4		✓	✓	39.8	2.10	0.87
5	✓		✓	49.7	2.36	0.92
6	✓	✓		45.0	2.30	0.89
7	✓	✓	✓	<b>53.6</b>	2.21	<b>0.93</b>

Table 4: Correction results of VENCE with different combinations of energy functions.

the effectiveness of our design. Our results in this part also call back to the fact that verifiers can be practically utilized to constrain language models for solving complex tasks, *e.g.*, math-word problems (Cobbe et al. 2021).

**Will more editing iterations help correction?** This question challenges the iterative editing design in VENCE. To answer this question, we evaluate the results VENCE produces during each iteration ranging from 0 to 30, where Iter#0 stands for the original claim. Moreover, we also test two variants of VENCE, *i.e.*, VENCE without verification model in the energy or without gradient-guided position finding. Results are presented in Figure 3. We find that 1) VENCE converges approximately at Iter#15; 2) when correcting without verification signals in the energy function  $\mathcal{E}(x)$ , VENCE suffers from a big performance drop ( $\sim 13$  points); 3) the results for random masking shows a big acceleration of convergence when masking is guided with gradient-based sampling, where the verifier also plays an important role.

**How do different combinations of energy functions affect correction?** One of the key modules in MH sampling and thus in VENCE is the stationary distribution  $\pi(x)$  implemented with the probabilistic energy-based model  $\mathcal{E}(x)$ . We enumerate the possible combinations of energy functions in VENCE and present the results in Table 4. SARI (final) measures correctness, HammingScore measures minimum modification, and BERTScore (Zhang et al. 2020b) measures fluency. We find that 1)  $\mathcal{E}_v$  for verified truthfulness contributes

Edit Choice	SARI (%)			
	Keep	Delete	Add	Final
Entity-only	33.6	30.0	1.7	21.8
Token-only	34.5	28.6	9.1	24.7
Both	<b>67.1</b>	<b>61.9</b>	<b>36.0</b>	<b>55.0</b>

Table 5: Correction results of VENCE, where the editing only happens on entities, tokens, or both.

Method	Evidence	SARI (%)			
		Keep	Delete	Add	Final
T5MC	<i>retrieved</i>	65.2	<b>62.7</b>	15.5	47.8
	<i>gold</i>	66.7	62.2	16.1	48.3
T5MC-V	<i>retrieved</i>	61.1	54.3	19.4	44.9
	<i>gold</i>	61.8	62.2	10.2	44.7
VENCE	<i>retrieved</i>	66.0	60.1	<b>34.8</b>	53.6
	<i>gold</i>	<b>67.5</b>	61.5	34.6	<b>54.5</b>

Table 6: Correction results of VENCE and baselines with retrieved (as in § 4.3) or ground truth evidence.

the most to the quality and correctness of the results, as seen in Row 1 & 2 & 3 and Row 4 & 7; 2) When only using  $\mathcal{E}_{LM}$ , the claim is not edited much, thus HammingScore achieves the lowest (Row 3), whereas poor SARI score reveals its falsity. Also, SARI drops much when removing  $\mathcal{E}_{LM}$  (Row 6 & 7), denoting the importance of fluency constraint during the correction. 3) Interestingly, as seen in Row 5 & 7,  $\mathcal{E}_H$  does not contribute to the Hamming distance score, but improves SARI for correctness (+3.9 points). These findings indicate that all energy functions are indispensable for VENCE, and VENCE is successfully constrained to find the best-suited modification for a corrected claim.

**What if we only edit tokens or entities?** As in § 3.3, 3.4, we propose new edits from both token and entity spaces in order to accommodate for consecutive tokens within entities. To show the effectiveness of this design, we conduct an ablation study where the correction happens *only* on tokens or entities. According to the results in Table 5, our design of generative proposal over two spaces greatly surpasses the token-only or entity-only counterparts, achieving a much higher score of over 30 points in SARI-Final.

**What if there is no error propagation from evidence retrieval?** Same to other retrieval-based tasks, the performance of correction is inevitably affected by the error propagation from the evidence retrieval module. Since our focus in this paper is not evidence retrieval, we here present the ceiling performance when FEC systems correct claims with ground truth evidence. Results in Table 6 show that correcting with gold evidence does improve the results in most cases, but marginally. We believe future upgrades of FEC methods should primarily concentrate on improving FEC itself.

**Human Evaluation** Finally, we manually evaluate the FEC performance of VENCE and baselines, even though

Method	Grammatical	Supported	Corrected
T5MC	92.0	81.3	55.3
T5MC-V	92.0	80.0	51.3
VENCE	<b>94.7</b>	<b>84.0</b>	<b>74.7</b>

Table 7: Human evaluation results for corrected claims.

# Iter.	Proposed Claims	Acc.
0	One True Thing is a German film.	-
1	<u>One True Thing is a film.</u>	✓
2	<u>One True Thing is film.</u>	✗
3	<u>One True Thing is a drama film.</u>	✓
4	One True Thing is drama film.	✗
5	One True Thing is American drama film.	✓
6	<u>One True Thing is an American drama film.</u>	✓
7	<u>One True Thing is an American film.</u>	✓
8	One True Thing is an American drama.	✗
9	One True Thing is an 1998 American film.	✗
10	One True Thing is an American.	✗
11	<b><u>One True Thing is an American drama film.</u></b>	✓
12	One True Thing is an American drama.	✗
13	One True Thing is an American film.	✗
14	One True Thing is a American drama film.	✗
15	One True Thing is an American film.	✗
-	One True Thing is an American film.	Gold

Table 8: The proposed claim of VENCE and its acceptance (Acc.) decisions for every iteration (# Iter.). Underlined texts are accepted claims. Bold text is the output corrected claim.

SARI and ROUGE scores are reported to be *highly* correlated with human judgment (Thorne and Vlachos 2021). We sample 100 cases and ask three students (undergraduate and graduate) to answer three Boolean questions for each instance 1) *is it grammatically correct?* 2) *is it supported by evidence?* 3) *are the original errors corrected?* We take the average of the results of three annotators. After a training session, in the annotation session, they reach an inter-annotator agreement at Fleiss’s  $\kappa = 0.56$ . The results are shown in Table 7. VENCE consistently outperforms baseline systems, especially in the third dimension (+20 points), *i.e.*, “are the errors corrected?”. Also, results by VENCE are better supported by the evidence and more grammatically accurate. These findings suggest that both our method and baselines can generate (or copy) reasonable results from evidence to the output, but VENCE better concentrates on the correction of given errors.

**Editing History** We present the history of the proposed corrected claim of each iteration and its acceptance decision in Table 8. The original claim gradually moves towards being supported by evidence after 15 rounds of iterative editing. Most accepted proposals happen in the early iterations because the correction of apparent errors is more readily accepted. Since the correctness of the claim has almost been reached in iteration 11, the gradient-based distribution is not as sharp as that in the previous iterations. Even if an erroneous position is carelessly sampled, the correction proposal

will not be accepted.

## 6 Conclusion

In this paper, we target the factual error correction problem and propose a novel approach (VENCE) for this task. VENCE is built on the idea of iterative editing and fact verification constrained correction, which further advances the new state-of-the-art of distantly supervised approach for FEC when high-quality training data is lacking. Further analysis also proves that the iterative editing converges well, and truthfulness signals from verification serves as an indispensable component for VENCE in FEC. However, our work is still limited by the external fact verification model that we used, which is not by design able to tell a claim that is *more* supported than another. To tackle this, future work could design better verification models that know *degrees* of factual errors. Another limitation is that FECDATA is still somehow too small and simple to reflect real-world scenarios of factual errors in texts, where we call for more comprehensive datasets for this task of increasing importance.

## Acknowledgements

We thank the anonymous reviewers for their valuable comments and suggestions. This work was supported by Science and Technology Commission of Shanghai Municipality Grant (No. 22511105902) and Shanghai Municipal Science and Technology Major Project (No.2021SHZDZX0103). Yanghua Xiao is also a member of Research Group of Computational and AI Communication at Institute for Global Communications and Integrated Media, Fudan University.

## References

- Agrawal, S.; and Carpuat, M. 2022. An Imitation Learning Curriculum for Text Editing with Non-Autoregressive Models. In *Proc. of ACL*, 7550–7563. Association for Computational Linguistics.
- Aly, R.; Guo, Z.; Schlichtkrull, M. S.; Thorne, J.; Vlachos, A.; Christodoulopoulos, C.; Cocarascu, O.; and Mittal, A. 2021. The Fact Extraction and VERification Over Unstructured and Structured information (FEVEROUS) Shared Task. In *Proceedings of the Fourth Workshop on Fact Extraction and VERification (FEVER)*, 1–13. Association for Computational Linguistics.
- Augenstein, I.; Lioma, C.; Wang, D.; Chaves Lima, L.; Hansen, C.; Hansen, C.; and Simonsen, J. G. 2019. MultiFC: A Real-World Multi-Domain Dataset for Evidence-Based Fact Checking of Claims. In *Proc. of EMNLP-IJCNLP*, 4685–4697. Association for Computational Linguistics.
- Cao, M.; Dong, Y.; and Cheung, J. 2022. Hallucinated but Factual! Inspecting the Factuality of Hallucinations in Abstractive Summarization. In *Proc. of ACL*, 3340–3354. Association for Computational Linguistics.
- Cao, N. D.; Izacard, G.; Riedel, S.; and Petroni, F. 2021. Autoregressive Entity Retrieval. In *ICLR*.
- Cao, Z.; Wei, F.; Li, W.; and Li, S. 2018. Faithful to the Original: Fact Aware Neural Abstractive Summarization. *Pro-*



- ceedings of the AAAI Conference on Artificial Intelligence, 32(1).
- Chen, J.; Bao, Q.; Sun, C.; Zhang, X.; Chen, J.; Zhou, H.; Xiao, Y.; and Li, L. 2022a. LOREN: Logic-Regularized Reasoning for Interpretable Fact Verification. *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(10): 10482–10491.
- Chen, J.; Gan, C.; Cheng, S.; Zhou, H.; Xiao, Y.; and Li, L. 2022b. Unsupervised Editing for Counterfactual Stories. *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(10): 10473–10481.
- Chollampatt, S.; and Ng, H. T. 2018. A Multilayer Convolutional Encoder-Decoder Neural Network for Grammatical Error Correction. *Proceedings of the AAAI Conference on Artificial Intelligence*, 32(1).
- Cobbe, K.; Kosaraju, V.; Bavarian, M.; Hilton, J.; Nakano, R.; Hesse, C.; and Schulman, J. 2021. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*.
- Devaraj, A.; Sheffield, W.; Wallace, B.; and Li, J. J. 2022. Evaluating Factuality in Text Simplification. In *Proc. of ACL*, 7331–7345. Association for Computational Linguistics.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proc. of NAACL-HLT, Volume 1 (Long and Short Papers)*, 4171–4186. Association for Computational Linguistics.
- Goyal, K.; Dyer, C.; and Berg-Kirkpatrick, T. 2022. Exposing the Implicit Energy Networks behind Masked Language Models via Metropolis–Hastings. In *ICLR*.
- Gu, J.; Wang, C.; and Zhao, J. 2019. Levenshtein Transformer. In Wallach, H.; Larochelle, H.; Beygelzimer, A.; d'Alché-Buc, F.; Fox, E.; and Garnett, R., eds., *NeurIPS*, volume 32. Curran Associates, Inc.
- Higashiyama, S.; Utiyama, M.; Watanabe, T.; and Sumita, E. 2021. A Text Editing Approach to Joint Japanese Word Segmentation, POS Tagging, and Lexical Normalization. In *Proceedings of the Seventh Workshop on Noisy User-generated Text (W-NUT 2021)*, 67–80. Association for Computational Linguistics.
- Hinton, G. E. 2002. Training products of experts by minimizing contrastive divergence. *Neural computation*, 14(8): 1771–1800.
- Honovich, O.; Choshen, L.; Aharoni, R.; Neeman, E.; Szpektor, I.; and Abend, O. 2021.  $Q^2$ : Evaluating Factual Consistency in Knowledge-Grounded Dialogues via Question Generation and Question Answering. In *Proc. of EMNLP*, 7856–7870. Association for Computational Linguistics.
- Hu, Z.; Lee, R. K.-W.; Aggarwal, C. C.; and Zhang, A. 2022. Text style transfer: A review and experimental evaluation. *ACM SIGKDD Explorations Newsletter*, 24(1): 14–45.
- Inui, K.; Fujita, A.; Takahashi, T.; Iida, R.; and Iwakura, T. 2003. Text Simplification for Reading Assistance: A Project Note. In *Proceedings of the Second International Workshop on Paraphrasing*, 9–16. Association for Computational Linguistics.
- Ishii, H.; Tamura, A.; and Ninomiya, T. 2021. Grammatical Error Correction via Supervised Attention in the Vicinity of Errors. In *Proc. of PACLIC*, 227–234. Association for Computational Linguistics.
- Iso, H.; Qiao, C.; and Li, H. 2020. Fact-based Text Editing. In *Proc. of ACL*. Association for Computational Linguistics.
- Karpukhin, V.; Oguz, B.; Min, S.; Lewis, P.; Wu, L.; Edunov, S.; Chen, D.; and Yih, W.-t. 2020. Dense Passage Retrieval for Open-Domain Question Answering. In *Proc. of EMNLP*, 6769–6781. Association for Computational Linguistics.
- Kasner, Z.; and Dušek, O. 2020. Data-to-Text Generation with Iterative Text Editing. In *Proceedings of the 13th International Conference on Natural Language Generation*, 60–67. Association for Computational Linguistics.
- Kotonya, N.; and Toni, F. 2020. Explainable Automated Fact-Checking: A Survey. In *Proc. of COLING*, 5430–5443. International Committee on Computational Linguistics.
- Krishna, A.; Riedel, S.; and Vlachos, A. 2021. ProoFVer: Natural Logic Theorem Proving for Fact Verification. *CoRR*, abs/2108.11357.
- Lee, W.; Shin, J.; Jung, B.; Lee, J.; and Lee, J.-H. 2020. Noising Scheme for Data Augmentation in Automatic Post-Editing. In *Proceedings of the Fifth Conference on Machine Translation*, 783–788. Association for Computational Linguistics.
- Lewis, M.; Liu, Y.; Goyal, N.; Ghazvininejad, M.; Mohamed, A.; Levy, O.; Stoyanov, V.; and Zettlemoyer, L. 2020. BART: Denoising Sequence-to-Sequence Pre-training for Natural Language Generation, Translation, and Comprehension. In *Proc. of ACL*, 7871–7880. Association for Computational Linguistics.
- Lin, C.-Y. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, 74–81.
- Liu, Y.; Ott, M.; Goyal, N.; Du, J.; Joshi, M.; Chen, D.; Levy, O.; Lewis, M.; Zettlemoyer, L.; and Stoyanov, V. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Mallinson, J.; Adamek, J.; Malmi, E.; and Severyn, A. 2022. EdiT5: Semi-Autoregressive Text-Editing with T5 Warm-Start. *arXiv preprint arXiv:2205.12209*.
- Metropolis, N.; Rosenbluth, A. W.; Rosenbluth, M. N.; Teller, A. H.; and Teller, E. 1953. Equation of state calculations by fast computing machines. *The journal of chemical physics*, 21(6): 1087–1092.
- Miao, N.; Zhou, H.; Mou, L.; Yan, R.; and Li, L. 2019. CGMH: Constrained Sentence Generation by Metropolis-Hastings Sampling. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01): 6834–6842.
- Mireshghallah, F.; Goyal, K.; and Berg-Kirkpatrick, T. 2022. Mix and Match: Learning-free Controllable Text Generation using Energy Language Models. In *Proc. of ACL*, 401–415. Association for Computational Linguistics.
- Ng, H. T.; Wu, S. M.; Briscoe, T.; Hadiwinoto, C.; Susanto, R. H.; and Bryant, C. 2014. The CoNLL-2014 Shared Task on Grammatical Error Correction. In *Proceedings of the Eighteenth Conference on Computational Natural Language*

- Learning: Shared Task*, 1–14. Association for Computational Linguistics.
- Omelianchuk, K.; Atrasevych, V.; Chernodub, A.; and Skurzhanyski, O. 2020. GECToR – Grammatical Error Correction: Tag, Not Rewrite. In *Proceedings of the Fifteenth Workshop on Innovative Use of NLP for Building Educational Applications*, 163–170. Association for Computational Linguistics.
- Qorib, M.; Na, S.-H.; and Ng, H. T. 2022. Frustratingly Easy System Combination for Grammatical Error Correction. In *Proc. of NAACL-HLT, 1964–1974*. Association for Computational Linguistics.
- Radford, A.; Wu, J.; Child, R.; Luan, D.; Amodei, D.; and Sutskever, I. 2019. Language models are unsupervised multi-task learners. *OpenAI Blog*, 1(8): 9.
- Raffel, C.; Shazeer, N.; Roberts, A.; Lee, K.; Narang, S.; Matena, M.; Zhou, Y.; Li, W.; Liu, P. J.; et al. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21(140): 1–67.
- Raunak, V.; Menezes, A.; and Junczys-Dowmunt, M. 2021. The Curious Case of Hallucinations in Neural Machine Translation. In *Proc. of NAACL-HLT*, 1172–1183. Association for Computational Linguistics.
- Samarinas, C.; Hsu, W.; and Lee, M. L. 2021. Improving Evidence Retrieval for Automated Explainable Fact-Checking. In *Proc. of NAACL-HLT: Demonstrations*, 84–91. Association for Computational Linguistics.
- Sha, L. 2020. Gradient-guided Unsupervised Lexically Constrained Text Generation. In *Proc. of EMNLP*, 8692–8703. Association for Computational Linguistics.
- Shah, D.; Schuster, T.; and Barzilay, R. 2020. Automatic Fact-Guided Sentence Modification. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(05): 8791–8798.
- Stammbach, D.; and Ash, E. 2020. e-fever: Explanations and summaries for automated fact checking. In *Proceedings of the 2020 Truth and Trust Online Conference (TTO 2020)*, 32. Hacks Hackers.
- Thorne, J.; and Vlachos, A. 2021. Evidence-based Factual Error Correction. In *Proc. of ACL-IJCNLP*, 3298–3309. Association for Computational Linguistics.
- Thorne, J.; Vlachos, A.; Christodoulopoulos, C.; and Mittal, A. 2018. FEVER: a Large-scale Dataset for Fact Extraction and VERification. In *Proc. of NAACL-HLT, Volume 1 (Long Papers)*, 809–819. Association for Computational Linguistics.
- Wadden, D.; Lin, S.; Lo, K.; Wang, L. L.; van Zuylen, M.; Cohan, A.; and Hajishirzi, H. 2020. Fact or Fiction: Verifying Scientific Claims. In *Proc. of EMNLP*, 7534–7550. Association for Computational Linguistics.
- Williams, A.; Nangia, N.; and Bowman, S. 2018. A Broad-Coverage Challenge Corpus for Sentence Understanding through Inference. In *Proc. of NAACL-HLT, Volume 1 (Long Papers)*, 1112–1122. Association for Computational Linguistics.
- Wolf, T.; Debut, L.; Sanh, V.; Chaumond, J.; Delangue, C.; Moi, A.; Cistac, P.; Rault, T.; Louf, R.; Funtowicz, M.; Davison, J.; Shleifer, S.; von Platen, P.; Ma, C.; Jernite, Y.; Plu, J.; Xu, C.; Le Scao, T.; Gugger, S.; Drame, M.; Lhoest, Q.; and Rush, A. 2020. Transformers: State-of-the-Art Natural Language Processing. In *Proc. of EMNLP: System Demonstrations*, 38–45. Association for Computational Linguistics.
- Wu, L.; Rao, Y.; Lan, Y.; Sun, L.; and Qi, Z. 2021. Unified Dual-view Cognitive Model for Interpretable Claim Verification. In *Proc. of ACL-IJCNLP*, 59–68. Association for Computational Linguistics.
- Xiao, Y.; and Wang, W. Y. 2021. On Hallucination and Predictive Uncertainty in Conditional Language Generation. In *Proc. of EACL: Main Volume*, 2734–2744. Association for Computational Linguistics.
- Xu, W.; Napoles, C.; Pavlick, E.; Chen, Q.; and Callison-Burch, C. 2016. Optimizing Statistical Machine Translation for Text Simplification. *TACL*, 4: 401–415.
- Yasunaga, M.; Leskovec, J.; and Liang, P. 2021. LM-Critic: Language Models for Unsupervised Grammatical Error Correction. In *Proc. of EMNLP*, 7752–7763. Association for Computational Linguistics.
- Yuan, Z.; and Briscoe, T. 2016. Grammatical error correction using neural machine translation. In *Proc. of NAACL-HLT*, 380–386. Association for Computational Linguistics.
- Zhang, M.; Jiang, N.; Li, L.; and Xue, Y. 2020a. Language Generation via Combinatorial Constraint Satisfaction: A Tree Search Enhanced Monte-Carlo Approach. In *Findings of EMNLP*, 1286–1298. Association for Computational Linguistics.
- Zhang, T.; Kishore, V.; Wu, F.; Weinberger, K. Q.; and Artzi, Y. 2020b. BERTScore: Evaluating Text Generation with BERT. In *ICLR*.