

Fact-based Text Editing

Hayate Iso^{†*} Chao Qiao[‡] Hang Li[‡]

[†]Nara Institute of Science and Technology [‡]ByteDance AI Lab

hyate.iso@gmail.com, {qiaochao, lihang.lh}@bytedance.com

Abstract

We propose a novel text editing task, referred to as *fact-based text editing*, in which the goal is to revise a given document to better describe the facts in a knowledge base (e.g., several triples). The task is important in practice because reflecting the truth is a common requirement in text editing. First, we propose a method for automatically generating a dataset for research on fact-based text editing, where each instance consists of a draft text, a revised text, and several facts represented in triples. We apply the method into two public table-to-text datasets, obtaining two new datasets consisting of 233k and 37k instances, respectively. Next, we propose a new neural network architecture for fact-based text editing, called FACTEDITOR, which edits a draft text by referring to given facts using a buffer, a stream, and a memory. A straightforward approach to address the problem would be to employ an encoder-decoder model. Our experimental results on the two datasets show that FACTEDITOR outperforms the encoder-decoder approach in terms of fidelity and fluency. The results also show that FACTEDITOR conducts inference faster than the encoder-decoder approach.

1 Introduction

Automatic editing of text by computer is an important application, which can help human writers to write better documents in terms of accuracy, fluency, etc. The task is easier and more practical than the automatic generation of texts from scratch and is attracting attention recently (Yang et al., 2017; Yin et al., 2019). In this paper, we consider a new and specific setting of it, referred to as *fact-based text editing*, in which a draft text and several facts (represented in triples) are given, and the system

* The work was done when Hayate Iso was a research intern at ByteDance AI Lab.

Set of triples

{(Baymax, creator, Duncan_Rouleau),
(Duncan_Rouleau, nationality, American),
(Baymax, creator, Steven_T_Seagle),
(Steven_T_Seagle, nationality, American),
(Baymax, series, Big_Hero_6),
(Big_Hero_6, starring, Scott_Adsit)}

Draft text

Baymax was created by Duncan_Rouleau, a winner of Eagle_Award. Baymax is a character in Big_Hero_6 .

Revised text

Baymax was created by American creators Duncan_Rouleau and Steven_T_Seagle . Baymax is a character in Big_Hero_6 which stars Scott_Adsit .

Table 1: Example of fact-based text editing. Facts are represented in triples. The **facts in green** appear in both draft text and triples. The **facts in orange** are present in the draft text, but absent from the triples. The **facts in blue** do not appear in the draft text, but in the triples. The task of *fact-based text editing* is to edit the draft text on the basis of the triples, by deleting **unsupported facts** and inserting **missing facts** while retaining **supported facts**.

aims to revise the text by adding missing facts and deleting unsupported facts. Table 1 gives an example of the task.

As far as we know, no previous work did address the problem. In a text-to-text generation, given a text, the system automatically creates another text, where the new text can be a text in another language (machine translation), a summary of the original text (summarization), or a text in better form (text editing). In a table-to-text generation, given a table containing facts in triples, the system automatically composes a text, which describes the facts. The former is a text-to-text problem, and the latter a table-to-text problem. In comparison, fact-based text editing can be viewed as a ‘text&table-to-text’ problem.

First, we devise a method for automatically creating a dataset for fact-based text editing. Recently, several table-to-text datasets have been created and released, consisting of pairs of facts and corresponding descriptions. We leverage such kind of data in our method. We first retrieve facts and their descriptions. Next, we take the descriptions as revised texts and automatically generate draft texts based on the facts using several rules. We build two datasets for fact-based text editing on the basis of WEBNLG (Gardent et al., 2017) and ROTOWIRE, consisting of 233k and 37k instances respectively (Wiseman et al., 2017)¹.

Second, we propose a model for fact-based text editing called FACTEDITOR. One could employ an encoder-decoder model, such as an encoder-decoder model, to perform the task. The encoder-decoder model implicitly represents the actions for transforming the draft text into a revised text. In contrast, FACTEDITOR explicitly represents the actions for text editing, including *Keep*, *Drop*, and *Gen*, which means retention, deletion, and generation of word respectively. The model utilizes a buffer for storing the draft text, a stream to store the revised text, and a memory for storing the facts. It also employs a neural network to control the entire editing process. FACTEDITOR has a lower time complexity than the encoder-decoder model, and thus it can edit a text more efficiently.

Experimental results show that FACTEDITOR outperforms the baseline model of using encoder-decoder for text editing in terms of fidelity and fluency, and also show that FACTEDITOR can perform text editing faster than the encoder-decoder model.

2 Related Work

2.1 Text Editing

Text editing has been studied in different settings such as automatic post-editing (Knight and Chan-der, 1994; Simard et al., 2007; Yang et al., 2017), paraphrasing (Dolan and Brockett, 2005), sentence simplification (Inui et al., 2003; Wubben et al., 2012), grammar error correction (Ng et al., 2014), and text style transfer (Shen et al., 2017; Hu et al., 2017).

The rise of encoder-decoder models (Cho et al., 2014; Sutskever et al., 2014) as well as the attention (Bahdanau et al., 2015; Vaswani et al., 2017)

¹The datasets are publicly available at <https://github.com/isomap/factedit>

and copy mechanisms (Gu et al., 2016; Gulcehre et al., 2016) has dramatically changed the landscape, and now one can perform the task relatively easily with an encoder-decoder model such as Transformer provided that a sufficient amount of data is available. For example, Li et al. (2018) introduce a deep reinforcement learning framework for paraphrasing, consisting of a generator and an evaluator. Yin et al. (2019) formalize the problem of text edit as learning and utilization of edit representations and propose an encoder-decoder model for the task. Zhao et al. (2018) integrate paraphrasing rules with the Transformer model for text simplification. Zhao et al. (2019) proposes a method for English grammar correction using a Transformer and copy mechanism.

Another approach to text editing is to view the problem as sequential tagging instead of encoder-decoder. In this way, the efficiency of learning and prediction can be significantly enhanced. Vu and Haffari (2018) and Dong et al. (2019) conduct automatic post-editing and text simplification on the basis of edit operations and employ Neural Programmer-Interpreter (Reed and De Freitas, 2016) to predict the sequence of edits given a sequence of words, where the edits include *KEEP*, *DROP*, and *ADD*. Malmi et al. (2019) propose a sequential tagging model that assigns a tag (*KEEP* or *DELETE*) to each word in the input sequence and also decides whether to add a phrase before the word. Our proposed approach is also based on sequential tagging of actions. It is designed for fact-based text editing, not text-to-text generation, however.

2.2 Table-to-Text Generation

Table-to-text generation is the task which aims to generate a text from structured data (Reiter and Dale, 2000; Gatt and Krahmer, 2018), for example, a text from an infobox about a term in biology in wikipedia (Lebret et al., 2016) and a description of restaurant from a structured representation (Novikova et al., 2017). Encoder-decoder models can also be employed in table-to-text generation with structured data as input and generated text as output, for example, as in (Lebret et al., 2016). Puduppully et al. (2019) and Iso et al. (2019) propose utilizing an entity tracking module for document-level table-to-text generation.

One issue with table-to-text is that the style of generated texts can be diverse (Iso et al., 2019). Re-

y'	AGENT-1	performed	as	PATIENT-3	on	BRIDGE-1	mission	that	was	operated	by	PATIENT-2	.
\hat{x}'	AGENT-1	served	as	PATIENT-3	was a crew member of the	BRIDGE-1	mission	.					
x'	AGENT-1	performed	as	PATIENT-3	on	BRIDGE-1	mission	.					

(a) Example for insertion. The revised template y' and the reference template \hat{x}' share subsequences. The set of triple templates $\mathcal{T} \setminus \hat{\mathcal{T}}$ is $\{(BRIDGE-1, operator, PATIENT-2)\}$. Our method removes “that was operated by PATIENT-2” from the revised template y' to create the draft template x' .

y'		AGENT-1				was	created	by	BRIDGE-1	and	PATIENT-2	.				
\hat{x}'	The character of	AGENT-1	,	whose	full	name	is	PATIENT-1	,	was	created	by	BRIDGE-1	and	PATIENT-2	.
x'		AGENT-1	,	whose	full	name	is	PATIENT-1	,	was	created	by	BRIDGE-1	and	PATIENT-2	.

(b) Example for deletion. The revised template y' and the reference template \hat{x}' share subsequences. The set of triple templates $\hat{\mathcal{T}} \setminus \mathcal{T}$ is $\{(AGENT-1, fullName, PATIENT-1)\}$. Our method copies “whose full name is PATIENT-1” from the reference template \hat{x}' to create the draft template x' .

Table 2: Examples for insertion and deletion, where **words in green** are matched, **words in gray** are not matched, **words in blue are copied**, and **words in orange** are removed. Best viewed in color.

searchers have developed methods to deal with the problem using other texts as templates (Hashimoto et al., 2018; Guu et al., 2018; Peng et al., 2019). The difference between the approach and fact-based text editing is that the former is about table-to-text generation based on other texts, while the latter is about text-to-text generation based on structured data.

3 Data Creation

In this section, we describe our method of data creation for fact-based text editing. The method automatically constructs a dataset from an existing table-to-text dataset.

3.1 Data Sources

There are two benchmark datasets of table-to-text, WEBNLG (Gardent et al., 2017)² and ROTOWIRE (Wiseman et al., 2017)³. We create two datasets on the basis of them, referred to as WEBEDIT and ROTOEDIT respectively. In the datasets, each instance consists of a table (structured data) and an associated text (unstructured data) describing almost the same content.⁴

For each instance, we take the table as triples of facts and the associated text as a revised text, and we automatically create a draft text. The set of triples is represented as $\mathcal{T} = \{t\}$. Each triple t consists of subject, predicate, and object, denoted

²The data is available at <https://github.com/ThiagoCF05/webnlg>. We utilize version 1.5.

³We utilize the ROTOWIRE-MODIFIED data provided by Iso et al. (2019) available at <https://github.com/aistairc/rotowire-modified>. The authors also provide an information extractor for processing the data.

⁴In ROTOWIRE, we discard redundant box-scores and unrelated sentences using the information extractor and heuristic rules.

as $t = (subj, pred, obj)$. For simplicity, we refer to the nouns or noun phrases of subject and object simply as entities. The revised text is a sequence of words denoted as y . The draft text is a sequence of words denoted as x .

Given the set of triples \mathcal{T} and the revised text y , we aim to create a draft text x , such that x is not in accordance with \mathcal{T} , in contrast to y , and therefore text editing from x to y is needed.

3.2 Procedure

Our method first creates templates for all the sets of triples and revised texts and then constructs a draft text for each set of triples and revised text based on their related templates.

Creation of templates

For each instance, our method first delexicalizes the entity words in the set of triples \mathcal{T} and the revised text y to obtain a set of triple templates \mathcal{T}' and a revised template y' . For example, given $\mathcal{T} = \{(Baymax, voice, Scott_Adsit)\}$ and $y = \text{“Scott_Adsit does the voice for Baymax”}$, it produces the set of triple templates $\mathcal{T}' = \{(AGENT-1, voice, PATIENT-1)\}$ and the revised template $y' = \text{“AGENT-1 does the voice for PATIENT-1”}$. Our method then collects all the sets of triple templates \mathcal{T}' and revised templates y' and retains them in a key-value store with y' being a key and \mathcal{T}' being a value.

Creation of draft text

Next, our method constructs a draft text x using a set of triple templates \mathcal{T}' and a revised template y' . For simplicity, it only considers the use of either *insertion* or *deletion* in the text editing, and one can easily make an extension of it to a more complex

setting. Note that the process of data creation is reverse to that of text editing.

Given a pair of \mathcal{T}' and \mathbf{y}' , our method retrieves another pair denoted as $\hat{\mathcal{T}}'$ and $\hat{\mathbf{x}}'$, such that \mathbf{y}' and $\hat{\mathbf{x}}'$ have the longest common subsequences. We refer to $\hat{\mathbf{x}}'$ as a reference template. There are two possibilities; $\hat{\mathcal{T}}'$ is a subset or a superset of \mathcal{T}' . (We ignore the case in which they are identical.) Our method then manages to change \mathbf{y}' to a draft template denoted as \mathbf{x}' on the basis of the relation between \mathcal{T}' and $\hat{\mathcal{T}}'$. If $\hat{\mathcal{T}}' \subsetneq \mathcal{T}'$, then the draft template \mathbf{x}' created is for insertion, and if $\hat{\mathcal{T}}' \supsetneq \mathcal{T}'$, then the draft template \mathbf{x}' created is for deletion.

For *insertion*, the revised template \mathbf{y}' and the reference template $\hat{\mathbf{x}}'$ share subsequences, and the set of triples $\mathcal{T} \setminus \hat{\mathcal{T}}$ appear in \mathbf{y}' but not in $\hat{\mathbf{x}}'$. Our method keeps the shared subsequences in \mathbf{y}' , removes the subsequences in \mathbf{y}' about $\mathcal{T} \setminus \hat{\mathcal{T}}$, and copies the rest of words in \mathbf{y}' , to create the draft template \mathbf{x}' . Table 2a gives an example. The shared subsequences “AGENT-1 performed as PATIENT-3 on BRIDGE-1 mission” are kept. The set of triple templates $\mathcal{T} \setminus \hat{\mathcal{T}}$ is $\{(\text{BRIDGE-1}, \text{operator}, \text{PATIENT-2})\}$. The subsequence “that was operated by PATIENT-2” is removed. Note that the subsequence “served” is not copied because it is not shared by \mathbf{y}' and $\hat{\mathbf{x}}'$.

For *deletion*, the revised template \mathbf{y}' and the reference template $\hat{\mathbf{x}}'$ share subsequences. The set of triples $\hat{\mathcal{T}} \setminus \mathcal{T}$ appear in $\hat{\mathbf{x}}'$ but not in \mathbf{y}' . Our method retains the shared subsequences in \mathbf{y}' , copies the subsequences in $\hat{\mathbf{x}}'$ about $\hat{\mathcal{T}} \setminus \mathcal{T}$, and copies the rest of words in \mathbf{y}' , to create the draft template \mathbf{x}' . Table 2b gives an example. The subsequences “AGENT-1 was created by BRIDGE-1 and PATIENT-2” are retained. The set of triple templates $\hat{\mathcal{T}} \setminus \mathcal{T}$ is $\{(\text{AGENT-1}, \text{full-Name}, \text{PATIENT-1})\}$. The subsequence “whose full name is PATIENT-1” is copied. Note that the subsequence “the character of” is not copied because it is not shared by \mathbf{y}' and $\hat{\mathbf{x}}'$.

After getting the draft template \mathbf{x}' , our method lexicalizes it to obtain a draft text \mathbf{x} , where the lexicons (entity words) are collected from the corresponding revised text \mathbf{y} .

We obtain two datasets with our method, referred to as WEBEDIT and ROTOEDIT, respectively. Table 3 gives the statistics of the datasets.

In the WEBEDIT data, sometimes entities only appear in the *subj*’s of triples. In such cases, we also make them appear in the *obj*’s. To do so, we

	WEBEDIT			ROTOEDIT		
	TRAIN	VALID	TEST	TRAIN	VALID	TEST
$\#\mathcal{D}$	181k	23k	29k	27k	5.3k	4.9k
$\#\mathcal{W}_d$	4.1M	495k	624k	4.7M	904k	839k
$\#\mathcal{W}_r$	4.2M	525k	649k	5.6M	1.1M	1.0M
$\#\mathcal{S}$	403k	49k	62k	209k	40k	36k

Table 3: Statistics of WEBEDIT and ROTOEDIT, where $\#\mathcal{D}$ is the number of instances, $\#\mathcal{W}_d$ and $\#\mathcal{W}_r$ are the total numbers of words in the draft texts and the revised texts, respectively, and $\#\mathcal{S}$ is total the number of sentences.

introduce an additional triple (`ROOT`, `IsOf`, *subj*) for each *subj*, where `ROOT` is a dummy entity.

4 FACTEDITOR

In this section, we describe our proposed model for fact-based text editing referred to as FACTEDITOR.

4.1 Model Architecture

FACTEDITOR transforms a draft text into a revised text based on given triples. The model consists of three components, a buffer for storing the draft text and its representations, a stream for storing the revised text and its representations, and a memory for storing the triples and their representations, as shown in Figure 1.

FACTEDITOR scans the text in the buffer, copies the parts of text from the buffer into the stream if they are described in the triples in the memory, deletes the parts of the text if they are not mentioned in the triples, and inserts new parts of next into the stream which is only presented in the triples.

The architecture of FACTEDITOR is inspired by those in sentence parsing Dyer et al. (2015); Watanabe and Sumita (2015). The actual processing of FACTEDITOR is to generate a sequence of words into the stream from the given sequence of words in the buffer and the set of triples in the memory. A neural network is employed to control the entire editing process.

4.2 Neural Network

Initialization

FACTEDITOR first initializes the representations of content in the buffer, stream, and memory.

There is a feed-forward network associated with the memory, utilized to create the embeddings of triples. Let M denote the number of triples. The

embedding of triple $t_j, j = 1, \dots, M$ is calculated as

$$t_j = \tanh(\mathbf{W}_t \cdot [e_j^{\text{subj}}; e_j^{\text{pred}}; e_j^{\text{obj}}] + \mathbf{b}_t),$$

where \mathbf{W}_t and \mathbf{b}_t denote parameters, $e_j^{\text{subj}}, e_j^{\text{pred}}, e_j^{\text{obj}}$ denote the embeddings of subject, predicate, and object of triple t_j , and $[\cdot]$ denotes vector concatenation.

There is a bi-directional LSTM associated with the buffer, utilized to create the embeddings of words of draft text. The embeddings are obtained as $\mathbf{b} = \text{BiLSTM}(\mathbf{x})$, where $\mathbf{x} = (x_1, \dots, x_N)$ is the list of embeddings of words and $\mathbf{b} = (\mathbf{b}_1, \dots, \mathbf{b}_N)$ is the list of representations of words, where N denotes the number of words.

There is an LSTM associated with the stream for representing the hidden states of the stream. The first hidden state is initialized as

$$\mathbf{s}_1 = \tanh\left(\mathbf{W}_s \cdot \left[\frac{\sum_{i=1}^N \mathbf{b}_i}{N}; \frac{\sum_{j=1}^M \mathbf{t}_j}{M}\right] + \mathbf{b}_s\right)$$

where \mathbf{W}_s and \mathbf{b}_s denotes parameters.

Action prediction

FACTEDITOR predicts an action at each time t using the LSTM. There are three types of action, namely Keep, Drop, and Gen. First, it composes a context vector $\tilde{\mathbf{t}}_t$ of triples at time t using attention

$$\tilde{\mathbf{t}}_t = \sum_{j=1}^M \alpha_{t,j} \mathbf{t}_j$$

where $\alpha_{t,j}$ is a weight calculated as

$$\alpha_{t,j} \propto \exp\left(\mathbf{v}_\alpha^\top \cdot \tanh(\mathbf{W}_\alpha \cdot [\mathbf{s}_t; \mathbf{b}_t; \mathbf{t}_j])\right)$$


where \mathbf{v}_α and \mathbf{W}_α are parameters. Then, it creates the hidden state \mathbf{z}_t for action prediction at time t

$$\mathbf{z}_t = \tanh(\mathbf{W}_z \cdot [\mathbf{s}_t; \mathbf{b}_t; \tilde{\mathbf{t}}_t] + \mathbf{b}_z)$$

where \mathbf{W}_z and \mathbf{b}_z denote parameters. Next, it calculates the probability of action a_t

$$P(a_t | \mathbf{z}_t) = \text{softmax}(\mathbf{W}_a \cdot \mathbf{z}_t)$$


where \mathbf{W}_a denotes parameters, and chooses the action having the largest probability.



(a) The **Keep** action, where the top embedding of the buffer \mathbf{b}_t is popped and the concatenated vector $[\tilde{\mathbf{t}}_t; \mathbf{b}_t]$ is pushed into the stream LSTM.



(b) The **Drop** action, where the top embedding of the buffer \mathbf{b}_t is popped and the state in the stream is reused at the next time step $t + 1$.



(c) The **Gen** action, where the concatenated vector $[\tilde{\mathbf{t}}_t; \mathbf{W}_p \mathbf{y}_t]$ is pushed into the stream, and the top embedding of the buffer is reused at the next time step $t + 1$.

Figure 1: Actions of FACTEDITOR.

Action execution

FACTEDITOR takes action based on the prediction result at time t .

For **Keep** at time t , FACTEDITOR pops the top embedding \mathbf{b}_t in the buffer, and feeds the combination of the top embedding \mathbf{b}_t and the context vector of triples $\tilde{\mathbf{t}}_t$ into the stream, as shown in Fig. 1a. The state of stream is updated with the LSTM as $\mathbf{s}_{t+1} = \text{LSTM}([\tilde{\mathbf{t}}_t; \mathbf{b}_t], \mathbf{s}_t)$. FACTEDITOR also copies the top word in the buffer into the stream.

For **Drop** at time t , FACTEDITOR pops the top embedding in the buffer and proceeds to the next state, as shown in Fig. 1b. The state of stream is updated as $\mathbf{s}_{t+1} = \mathbf{s}_t$. Note that no word is inputted into the stream.

For **Gen** at time t , FACTEDITOR does not pop the top embedding in the buffer. It feeds the

Draft text \mathbf{x}	Bakewell_pudding is Dessert that can be served Warm or cold .
Revised text \mathbf{y}	Bakewell_pudding is Dessert that originates from Derbyshire.Dales .
Action sequence \mathbf{a}	Keep Keep Keep Keep Gen(originate)s Gen(from) Gen(Derbyshire.Dales) Drop Drop Drop Drop Keep

Table 4: An example of action sequence derived from a draft text and revised text.

combination of the context vector of triples $\tilde{\mathbf{t}}_t$ and the linearly projected embedding of word w into the stream, as shown in Fig. 1c. The state of stream is updated with the LSTM as $\mathbf{s}_{t+1} = \text{LSTM}([\tilde{\mathbf{t}}_t; \mathbf{W}_p \mathbf{y}_t], \mathbf{s}_t)$, where \mathbf{y}_t is the embedding of the generated word y_t and \mathbf{W}_p denotes parameters. In addition, FACTEDITOR copies the generated word y_t into the stream.

FACTEDITOR continues the actions until the buffer becomes empty.

Word generation

FACTEDITOR generates a word y_t at time t , when the action is Gen,

$$P_{\text{gen}}(y_t | \mathbf{z}_t) = \text{softmax}(\mathbf{W}_y \cdot \mathbf{z}_t)$$

where \mathbf{W}_y is parameters.

To avoid generation of OOV words, FACTEDITOR exploits the copy mechanism. It calculates the probability of copying the object of triple t_j

$$P_{\text{copy}}(o_j | \mathbf{z}_t) \propto \exp(\mathbf{v}_c^\top \cdot \tanh(\mathbf{W}_c \cdot [\mathbf{z}_t; \mathbf{t}_j]))$$

where \mathbf{v}_c and \mathbf{W}_c denote parameters, and o_j is the object of triple t_j . It also calculates the probability of gating

$$p_{\text{gate}} = \text{sigmoid}(\mathbf{w}_g^\top \cdot \mathbf{z}_t + b_g)$$

where \mathbf{w}_g and b_g are parameters. Finally, it calculates the probability of generating a word w_t through either generation or copying,

$$P(y_t | \mathbf{z}_t) = p_{\text{gate}} P_{\text{gen}}(y_t | \mathbf{z}_t) + (1 - p_{\text{gate}}) \sum_{j=1: o_j=y_t}^M P_{\text{copy}}(o_j | \mathbf{z}_t),$$

where it is assumed that the triples in the memory have the same subject and thus only objects need to be copied.

4.3 Model Learning

The conditional probability of sequence of actions $\mathbf{a} = (a_1, a_2, \dots, a_T)$ given the set of triples \mathcal{T} and

the sequence of input words \mathbf{x} can be written as

$$P(\mathbf{a} | \mathcal{T}, \mathbf{x}) = \prod_{t=1}^T P(a_t | \mathbf{z}_t)$$

where $P(a_t | \mathbf{z}_t)$ is the conditional probability of action a_t given state \mathbf{z}_t at time t and T denotes the number of actions.

The conditional probability of sequence of generated words $\mathbf{y} = (y_1, y_2, \dots, y_T)$ given the sequence of actions \mathbf{a} can be written as

$$P(\mathbf{y} | \mathbf{a}) = \prod_{t=1}^T P(y_t | a_t)$$

where $P(y_t | a_t)$ is the conditional probability of generated word y_t given action a_t at time t , which is calculated as

$$P(y_t | a_t) = \begin{cases} P(y_t | \mathbf{z}_t) & \text{if } a_t = \text{Gen} \\ 1 & \text{otherwise} \end{cases}$$

Note that not all positions have a generated word. In such a case, y_t is simply a null word.

The learning of the model is carried out via supervised learning. The objective of learning is to minimize the negative log-likelihood of $P(\mathbf{a} | \mathcal{T}, \mathbf{x})$ and $P(\mathbf{y} | \mathbf{a})$

$$\mathcal{L}(\boldsymbol{\theta}) = - \sum_{t=1}^T \{\log P(a_t | \mathbf{z}_t) + \log P(y_t | a_t)\}$$

where $\boldsymbol{\theta}$ denotes the parameters.

A training instance consists of a pair of draft text and revised text, as well as a set of triples, denoted as \mathbf{x} , \mathbf{y} , and \mathcal{T} respectively. For each instance, our method derives a sequence of actions denoted as \mathbf{a} , in a similar way as that in (Dong et al., 2019). It first finds the longest common subsequence between \mathbf{x} and \mathbf{y} , and then selects an action of Keep, Drop, or Gen at each position, according to how \mathbf{y} is obtained from \mathbf{x} and \mathcal{T} (cf., Tab. 4). Action Gen is preferred over action Drop when both are valid.




Figure 2: Model architectures of the baselines. All models employ attention and copy mechanism.

4.4 Time Complexity

The time complexity of inference in FACTEDITOR is $\mathcal{O}(NM)$, where N is the number of words in the buffer, and M is the number of triples. Scanning of data in the buffer is of complexity $\mathcal{O}(N)$. The generation of action needs the execution of attention, which is of complexity $\mathcal{O}(M)$. Usually, N is much larger than M .

4.5 Baseline

We consider a baseline method using the encoder-decoder architecture, which takes the set of triples and the draft text as input and generates a revised text. We refer to the method as ENCDECEDITOR. The encoder of ENCDECEDITOR is the same as that of FACTEDITOR. The decoder is the standard attention and copy model, which creates and utilizes a context vector and predicts the next word at each time.

The time complexity of inference in ENCDECEDITOR is $\mathcal{O}(N^2 + NM)$ (cf., Britz et al. (2017)). Note that in fact-based text editing, usually N is very large. That means that ENCDECEDITOR is less efficient than FACTEDITOR.

5 Experiment

We conduct experiments to make comparison between FACTEDITOR and the baselines using the two datasets WEBEDIT and ROTOEDIT.

5.1 Experiment Setup

The main baseline is the encoder-decoder model ENCDECEDITOR, as explained above. We further consider three baselines, No-Editing, Table-to-Text, and Text-to-Text. In No-Editing, the draft text is directly used. In Table-to-Text, a revised text is generated from the triples using encoder-decoder. In Text-to-Text, a revised text is created from the draft text using the encoder-decoder model. Figure 2 gives illustrations of the baselines.

We evaluate the results of revised texts by the models from the viewpoint of fluency and fidelity.

We utilize ExactMatch (EM), BLEU (Papineni et al., 2002) and SARI (Xu et al., 2016) scores⁵ as evaluation metrics for fluency. We also utilize precision, recall, and F1 score as evaluation metrics for fidelity. For WEBEDIT, we extract the entities from the generated text and the reference text and then calculate the precision, recall, and F1 scores. For ROTOEDIT, we use the information extraction tool provided by Wiseman et al. (2017) for calculation of the scores.

For the embeddings of subject and object for both datasets and the embedding of the predicate for ROTOEDIT, we simply use the embedding lookup table. For the embedding of the predicate for WEBEDIT, we first tokenize the predicate, lookup the embeddings of lower-cased words from the table, and use averaged embedding to deal with the OOV problem (Moryossef et al., 2019).

We tune the hyperparameters based on the BLEU score on a development set. For WEBEDIT, we set the sizes of embeddings, buffers, and triples to 300, and set the size of the stream to 600. For ROTOEDIT, we set the size of embeddings to 100 and set the sizes of buffers, triples, and stream to 200. The initial learning rate is $2e-3$, and AMS-Grad is used for automatically adjusting the learning rate (Reddi et al., 2018). Our implementation makes use of AllenNLP (Gardner et al., 2018).

5.2 Experimental Results

Quantitative evaluation

We present the performances of our proposed model FACTEDITOR and the baselines on fact-based text editing in Table 5. One can draw several conclusions from the results.

First, our proposed model, FACTEDITOR, achieves significantly better performances than the main baseline, ENCDECEDITOR, in terms of almost all measures. In particular, FACTEDITOR

⁵We use a modified version of SARI where β equals 1.0, available at https://github.com/tensorflow/tensor2tensor/blob/master/tensor2tensor/utils/sari_hook.py

Model	FLUENCY						FIDELITY		
	BLEU	SARI	KEEP	ADD	DELETE	EM	P%	R%	F1%
Baselines									
No-Editing	66.67	31.51	78.62	3.91	12.02	0.	84.49	76.34	80.21
Table-to-Text	33.75	43.83	51.44	27.86	52.19	5.78	98.23	83.72	90.40
Text-to-Text	63.61	58.73	82.62	25.77	67.80	6.22	81.93	77.16	79.48
Fact-based text editing									
ENCDECEDITOR	71.03	69.59	89.49	43.82	75.48	20.96	98.06	87.56	92.51
FACTEDITOR	75.68	72.20	91.84	47.69	77.07	24.80	96.88	89.74	93.17

(a) WEBEDIT

Model	FLUENCY						FIDELITY		
	BLEU	SARI	KEEP	ADD	DELETE	EM	P%	R%	F1%
Baselines									
No-Editing	74.95	39.59	95.72	0.05	23.01	0.	92.92	65.02	76.51
Table-to-Text	24.87	23.30	39.12	14.78	16.00	0.	48.01	24.28	32.33
Text-to-Text	78.07	60.25	97.29	13.04	70.43	0.02	63.62	41.08	49.92
Fact-based text editing									
ENCDECEDITOR	83.36	71.46	97.69	44.02	72.69	2.49	78.80	52.21	62.81
FACTEDITOR	84.43	74.72	98.41	41.50	84.24	2.65	78.84	52.30	63.39

(b) ROTOEDIT

Table 5: Performances of FACTEDITOR and baselines on two datasets in terms of Fluency and Fidelity. EM stands for exact match.

obtains significant gains in DELETE scores on both WEBEDIT and ROTOEDIT.

Second, the fact-based text editing models (FACTEDITOR and ENCDECEDITOR) significantly improve upon the other models in terms of fluency scores, and achieve similar performances in terms of fidelity scores.

Third, compared to No-Editing, Table-to-Text has higher fidelity scores, but lower fluency scores. Text-to-Text has almost the same fluency scores, but lower fidelity scores on ROTOEDIT.

Qualitative evaluation

We also manually evaluate 50 randomly sampled revised texts for WEBEDIT. We check whether the revised texts given by FACTEDITOR and ENCDECEDITOR include all the facts. We categorize the factual errors made by the two models. Table 6 shows the results. One can see that FACTEDITOR covers more facts than ENCDECEDITOR and has less factual errors than ENCDECEDITOR.

FACTEDITOR has a larger number of correct editing (CQT) than ENCDECEDITOR for fact-based text editing. In contrast, ENCDECEDITOR often includes a larger number of unnecessary rephrasings (UPARA) than FACTEDITOR.

	Covered facts		Factual errors			
	CQT	UPARA	RPT	MS	USUP	DREL
ENCDECEDITOR	14	7	16	21	3	12
FACTEDITOR	24	4	9	19	1	3

Table 6: Evaluation results on 50 randomly sampled revised texts in WEBEDIT in terms of numbers of correct editing (CQT), unnecessary paraphrasing (UPARA), repetition (RPT), missing facts (MS), unsupported facts (USUP) and different relations (DREL)

There are four types of factual errors: fact repetitions (RPT), fact missings (MS), fact unsupported (USUP), and relation difference (DREL). Both FACTEDITOR and ENCDECEDITOR often fail to insert missing facts (MS), but rarely insert unsupported facts (USUP). ENCDECEDITOR often generates the same facts multiple times (RPT) or facts in different relations (DREL). In contrast, FACTEDITOR can seldomly make such errors.

Table 7 shows an example of results given by ENCDECEDITOR and FACTEDITOR. The revised texts of both ENCDECEDITOR and FACTEDITOR appear to be fluent, but that of FACTEDITOR has higher fidelity than that of ENCDECEDITOR. ENCDECEDITOR cannot effectively eliminate the

Set of triples	{(Ardmore_Airport, runwayLength, 1411.0), (Ardmore_Airport, 3rd_runway_SurfaceType, Poaceae), (Ardmore_Airport, operatingOrganisation, Civil_Aviation_Authority_of_New_Zealand), (Ardmore_Airport, elevationAboveTheSeaLevel, 34.0), (Ardmore_Airport, runwayName, 03R/21L)}
Draft text	Ardmore_Airport , ICAO Location Identifier UTAA . Ardmore_Airport 3rd runway is made of Poaceae and Ardmore_Airport . 03R/21L is 1411.0 m long and Ardmore_Airport is 34.0 above sea level .
Revised text	Ardmore_Airport is operated by Civil_Aviation_Authority_of_New_Zealand . Ardmore_Airport 3rd runway is made of Poaceae and Ardmore_Airport name is 03R/21L . 03R/21L is 1411.0 m long and Ardmore_Airport is 34.0 above sea level .
ENCDECEDITOR	Ardmore_Airport , ICAO Location Identifier UTAA , is operated by Civil_Aviation_Authority_of_New_Zealand . Ardmore_Airport 3rd runway is made of Poaceae and Ardmore_Airport . 03R/21L is 1411.0 m long and Ardmore_Airport is <u>34.0 m long</u> .
FACTEDITOR	Ardmore_Airport is operated by Civil_Aviation_Authority_of_New_Zealand . Ardmore_Airport 3rd runway is made of Poaceae and Ardmore_Airport . 03R/21L is 1411.0 m long and Ardmore_Airport is 34.0 above sea level .

Table 7: Example of generated revised texts given by ENCDECEDITOR and FACTEDITOR on WEBEDIT. **Entities in green** appear in both the set of triples and the draft text. **Entities in orange** only appear in the draft text. **Entities in blue** should appear in the revised text but do not appear in the draft text.

	WEBEDIT	ROTOEDIT
Table-to-Text	4,083	1,834
Text-to-Text	2,751	581
ENCDECEDITOR	2,487	505
FACTEDITOR	3,295	<u>1,412</u>

Table 8: Runtime analysis (# of words/second). Table-to-Text always shows the fastest performance (**Bold-faced**). FACTEDITOR shows the second fastest runtime performance (Underlined).

description about an unsupported fact (in orange) appearing in the draft text. In contrast, FACTEDITOR can deal with the problem well. In addition, ENCDECEDITOR conducts an unnecessary substitution in the draft text (underlined). FACTEDITOR tends to avoid such unnecessary editing.

Runtime analysis

We conduct runtime analysis on FACTEDITOR and the baselines in terms of number of processed words per second, on both WEBEDIT and ROTOEDIT. Table 8 gives the results when the batch size is 128 for all methods. Table-to-Text is the fastest, followed by FACTEDITOR. FACTEDITOR is always faster than ENCDECEDITOR, apparently because it has a lower time complexity, as explained in Section 4. The texts in WEBEDIT are relatively short, and thus FACTEDITOR and ENCDECEDITOR have similar runtime speeds. In contrast, the texts in ROTOEDIT are relatively long, and thus FACTEDITOR executes approximately two times faster than ENCDECEDITOR.

6 Conclusion

In this paper, we have defined a new task referred to as *fact-based text editing* and made two contributions to research on the problem. First, we have proposed a data construction method for fact-based text editing and created two datasets. Second, we have proposed a model for fact-based text editing, named FACTEDITOR, which performs the task by generating a sequence of actions. Experimental results show that the proposed model FACTEDITOR performs better and faster than the baselines, including an encoder-decoder model.

Acknowledgments

We would like to thank the reviewers for their insightful comments.

References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. *Neural Machine Translation by Jointly Learning to Align and Translate*. In *International Conference on Learning Representations*.
- Denny Britz, Melody Guan, and Minh-Thang Luong. 2017. *Efficient Attention using a Fixed-Size Memory Representation*. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 392–400, Copenhagen, Denmark. Association for Computational Linguistics.
- Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. *Learning Phrase Representations using RNN Encoder-Decoder for Statistical Machine Translation*. In

- Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1724–1734, Doha, Qatar. Association for Computational Linguistics.
- William B. Dolan and Chris Brockett. 2005. [Automatically constructing a corpus of sentential paraphrases](#). In *Proceedings of the Third International Workshop on Paraphrasing (IWP2005)*.
- Yue Dong, Zichao Li, Mehdi Rezagholizadeh, and Jackie Chi Kit Cheung. 2019. [EditNTS: An Neural Programmer-Interpreter Model for Sentence Simplification through Explicit Editing](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3393–3402, Florence, Italy. Association for Computational Linguistics.
- Chris Dyer, Miguel Ballesteros, Wang Ling, Austin Matthews, and Noah A. Smith. 2015. [Transition-Based Dependency Parsing with Stack Long Short-Term Memory](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 334–343, Beijing, China. Association for Computational Linguistics.
- Claire Gardent, Anastasia Shimorina, Shashi Narayan, and Laura Perez-Beltrachini. 2017. [Creating training corpora for NLG micro-planners](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 179–188, Vancouver, Canada. Association for Computational Linguistics.
- Matt Gardner, Joel Grus, Mark Neumann, Oyvind Tafjord, Pradeep Dasigi, Nelson F. Liu, Matthew Peters, Michael Schmitz, and Luke Zettlemoyer. 2018. [AllenNLP: A Deep Semantic Natural Language Processing Platform](#). In *Proceedings of Workshop for NLP Open Source Software (NLP-OSS)*, pages 1–6, Melbourne, Australia. Association for Computational Linguistics.
- Albert Gatt and Emiel Kraemer. 2018. [Survey of the State of the Art in Natural Language Generation: Core tasks, applications and evaluation](#). *Journal of Artificial Intelligence Research (JAIR)*, 61:65–170.
- Jiatao Gu, Zhengdong Lu, Hang Li, and Victor O.K. Li. 2016. [Incorporating copying mechanism in sequence-to-sequence learning](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1631–1640, Berlin, Germany. Association for Computational Linguistics.
- Caglar Gulcehre, Sungjin Ahn, Ramesh Nallapati, Bowen Zhou, and Yoshua Bengio. 2016. [Pointing the unknown words](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 140–149, Berlin, Germany. Association for Computational Linguistics.
- Kelvin Guu, Tatsunori B Hashimoto, Yonatan Oren, and Percy Liang. 2018. [Generating Sentences by Editing Prototypes](#). *Transactions of the Association for Computational Linguistics*, 6:437–450.
- Tatsunori B Hashimoto, Kelvin Guu, Yonatan Oren, and Percy S Liang. 2018. [A Retrieve-and-Edit Framework for Predicting Structured Outputs](#). In *Advances in Neural Information Processing Systems*, pages 10052–10062. Curran Associates, Inc.
- Zhiting Hu, Zichao Yang, Xiaodan Liang, Ruslan Salakhutdinov, and Eric P. Xing. 2017. [Toward Controlled Generation of Text](#). In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 1587–1596, International Convention Centre, Sydney, Australia. PMLR.
- Kentaro Inui, Atsushi Fujita, Tetsuro Takahashi, Ryu Iida, and Tomoya Iwakura. 2003. [Text simplification for reading assistance: A project note](#). In *Proceedings of the Second International Workshop on Paraphrasing*, pages 9–16, Sapporo, Japan. Association for Computational Linguistics.
- Hayate Iso, Yui Uehara, Tatsuya Ishigaki, Hiroshi Noji, Eiji Aramaki, Ichiro Kobayashi, Yusuke Miyao, Naoaki Okazaki, and Hiroya Takamura. 2019. [Learning to Select, Track, and Generate for Data-to-Text](#). In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 2102–2113, Florence, Italy.
- Kevin Knight and Ishwar Chander. 1994. [Automated Postediting of Documents](#). In *Proceedings of the AAAI Conference on Artificial Intelligence.*, volume 94, pages 779–784.
- Rémi Lebret, David Grangier, and Michael Auli. 2016. [Neural Text Generation from Structured Data with Application to the Biography Domain](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1203–1213, Austin, Texas. Association for Computational Linguistics.
- Zichao Li, Xin Jiang, Lifeng Shang, and Hang Li. 2018. [Paraphrase Generation with Deep Reinforcement Learning](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3865–3878, Brussels, Belgium. Association for Computational Linguistics.
- Eric Malmi, Sebastian Krause, Sascha Rothe, Daniil Mirylenka, and Aliaksei Severyn. 2019. [Encode, Tag, Realize: High-Precision Text Editing](#). 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 5053–5064, Hong Kong, China. Association for Computational Linguistics.

- Amit Moryossef, Yoav Goldberg, and Ido Dagan. 2019. [Step-by-Step: Separating Planning from Realization in Neural Data-to-Text Generation](#). 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 2267–2277, Minneapolis, Minnesota. Association for Computational Linguistics.
- Hwee Tou Ng, Siew Mei Wu, Ted Briscoe, Christian Hadiwinoto, Raymond Hendy Susanto, and Christopher Bryant. 2014. [The CoNLL-2014 Shared Task on Grammatical Error Correction](#). In *Proceedings of the Eighteenth Conference on Computational Natural Language Learning: Shared Task*, pages 1–14, Baltimore, Maryland. Association for Computational Linguistics.
- Jekaterina Novikova, Ondřej Dušek, and Verena Rieser. 2017. [The E2E Dataset: New Challenges For End-to-End Generation](#). In *Proceedings of the 18th Annual SIGdial Meeting on Discourse and Dialogue*, pages 201–206, Saarbrücken, Germany. 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.
- Hao Peng, Ankur Parikh, Manaal Faruqui, Bhuwan Dhingra, and Dipanjan Das. 2019. [Text Generation with Exemplar-based Adaptive Decoding](#). 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 2555–2565, Minneapolis, Minnesota. Association for Computational Linguistics.
- Ratish Puduppully, Li Dong, and Mirella Lapata. 2019. ["Data-to-text Generation with Entity Modeling"](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2023–2035, Florence, Italy. Association for Computational Linguistics.
- Sashank J Reddi, Satyen Kale, and Sanjiv Kumar. 2018. [On the convergence of adam and beyond](#). In *International Conference on Learning Representations*.
- Scott Reed and Nando De Freitas. 2016. [Neural Programmer-Interpreters](#). In *International Conference on Learning Representations*.
- Ehud Reiter and Robert Dale. 2000. *Building Natural Language Generation Systems*. Studies in Natural Language Processing. Cambridge University Press.
- Tianxiao Shen, Tao Lei, Regina Barzilay, and Tommi Jaakkola. 2017. [Style Transfer from Non-Parallel Text by Cross-Alignment](#). In *Advances in Neural Information Processing Systems 30*, pages 6830–6841. Curran Associates, Inc.
- Michel Simard, Cyril Goutte, and Pierre Isabelle. 2007. [Statistical Phrase-Based Post-Editing](#). In *Human Language Technologies 2007: The Conference of the North American Chapter of the Association for Computational Linguistics; Proceedings of the Main Conference*, pages 508–515, Rochester, New York. Association for Computational Linguistics.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. [Sequence to Sequence Learning with Neural Networks](#). In *Advances in Neural Information Processing Systems 27*, pages 3104–3112. Curran Associates, Inc.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is All you Need](#). In *Advances in Neural Information Processing Systems 30*, pages 5998–6008. Curran Associates, Inc.
- Thuy-Trang Vu and Gholamreza Haffari. 2018. [Automatic Post-Editing of Machine Translation: A Neural Programmer-Interpreter Approach](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3048–3053, Brussels, Belgium. Association for Computational Linguistics.
- Taro Watanabe and Eiichiro Sumita. 2015. [Transition-based neural constituent parsing](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1169–1179, Beijing, China. Association for Computational Linguistics.
- Sam Wiseman, Stuart Shieber, and Alexander Rush. 2017. ["Challenges in Data-to-Document Generation"](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2253–2263, Copenhagen, Denmark. Association for Computational Linguistics.
- Sander Wubben, Antal van den Bosch, and Emiel Krahmer. 2012. [Sentence simplification by monolingual machine translation](#). In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1015–1024, Jeju Island, Korea. Association for Computational Linguistics.
- 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 Edward 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.

Pengcheng Yin, Graham Neubig, Miltiadis Allamanis, Marc Brockschmidt, and Alexander L Gaunt. 2019. [Learning to Represent Edits](#). In *International Conference on Learning Representations*.

Sanqiang Zhao, Rui Meng, Daqing He, Andi Saptono, and Bambang Parmanto. 2018. [Integrating Transformer and Paraphrase Rules for Sentence Simplification](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3164–3173, Brussels, Belgium. Association for Computational Linguistics.

Wei Zhao, Liang Wang, Kewei Shen, Ruoyu Jia, and Jingming Liu. 2019. ["Improving Grammatical Error Correction via Pre-Training a Copy-Augmented Architecture with Unlabeled Data"](#). 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 156–165, Minneapolis, Minnesota. Association for Computational Linguistics.