Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Just ClozE! A Novel Framework for Evaluating the Factual Consistency Faster in Abstractive Summarization

Yiyang Li
Beijing University of Posts and Telecommunications
Lei Li
leili@bupt.edu.cn
Beijing University of Posts and Telecommunications
Dingxin Hu
Beijing University of Posts and Telecommunications
Yuze Li
Beijing University of Posts and Telecommunications
Yanquan Zhou
Beijing University of Posts and Telecommunications
Marina Litvak
Sami Shamoon College of Engineering
Natalia Vanetik
Sami Shamoon College of Engineering

Research Article

Keywords: Factual Consistency, Cloze-based Metric, Abstractive Summarization

Posted Date: December 5th, 2023

DOI: https://doi.org/10.21203/rs.3.rs-3695932/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Just ClozE! A Novel Framework for Evaluating the Factual Consistency Faster in Abstractive Summarization

Yiyang Li[1†], Lei Li[1*†], Dingxin Hu[1], Yuze Li[1], Yanquan Zhou[1*], Marina Litvak[2], Natalia Vanetik[2]

[1*]School of Artificial Intelligence, Beijing University of Posts and Telecommunications, Beijing, 100876, P. R. China.
[2]Department of Software Engineering, Shamoon College of Engineering, Beer Sheva, 84100,Israel.

*Corresponding author(s). E-mail(s): leili@bupt.edu.cn; zhouyanquan@bupt.edu.cn;
Contributing authors: kenlee@bupt.edu.cn; hudingxin@bupt.edu.cn; lyzbupt@bupt.edu.cn; marinal@ac.sce.ac.il; natalyav@sce.ac.il;
†These authors contributed equally to this work.

Abstract

The issue of factual consistency in abstractive summarization has received extensive attention in recent years, and the evaluation of factual consistency between summary and document has become an important and urgent task. Most of the current metrics are adopted from the question answering (QA) or natural language inference (NLI) task. However, the application of QA-based metrics is extremely time-consuming in practice while NLI-based metrics are lack of interpretability. In this paper, we propose a cloze-based evaluation framework called ClozE and show its great potential. It inherits strong interpretability from QA, while maintaining the speed of NLI-level reasoning. We demonstrate that ClozE can reduce the evaluation time by nearly 96% relative to QA-based metrics while retaining their interpretability and performance through experiments on six human-annotated datasets and a meta-evaluation benchmark GO FIGURE[1]. Finally, we discuss three important facets of ClozE in practice, which further shows better overall performance of ClozE compared to other metrics. The codes and models are released at https://github.com/Mr-KenLee/ClozE.

Keywords: Factual Consistency, Cloze-based Metric, Abstractive Summarization

1 Introduction

The factual consistency problem in abstractive summarization is a major challenge that limits the application of the summaries generated by state-of-the-art models in practice. A number of works[2][3][4] have shown that even the best abstractive summarization models still suffer from many factual inconsistencies. Thus, determining the factual consistency between a summary and its document is critical to addressing this challenge.

<table>
<tr><th colspan="3">Document & Summary</th></tr>
<tr><td colspan="3">D: England coach peter moores talks to the news media during a press conference at the adelaide oval on sunday.
S: Peter moores talks to the adelaide oval on sunday.</td></tr>
<tr><th>NLI-based</th><th>QA-based</th><th>Cloze-based</th></tr>
<tr><td rowspan="3">Entailment Matrix
$\begin{pmatrix}\text{Entail.}\\ \text{Neutral}\\ \text{Contra.}\end{pmatrix} = \begin{pmatrix}0.2\\ 0.1\\ 0.7\end{pmatrix}$</td><td>Extract Factual Factor
the adelaide oval</td><td>Extract Factual Factor
the adelaide oval</td></tr>
<tr><td>Question Generation (QG)
Who does Peter talk to?</td><td>Masking Summary
Peter moores talks to [M] on sunday.</td></tr>
<tr><td>Question Answering (QA)
the news media</td><td>Filling in the blank
the news media</td></tr>
<tr><td>Entailment Score
0.2</td><td>Similarity Score
0.4</td><td>Similarity Score
0.4</td></tr>
</table>

Table 1: A simple example of three categories of metrics. NLI-based metrics: The factual consistency score is computed with the probability of the summary being entailed by the document. QA-based metrics: (1) Extracting factual factors from the summary as answers. (2) Generating associated questions through the QG module. (3) Answering the questions based on the document through the QA module. (4) Computing the score based on the similarity between the extracted factual factors and generated answers. Cloze-based metrics (ClozE): The overall process are similar to QA-based metrics, with the difference that the QG module is replaced with a model-free strategy. The factual factors in the summary are masked at first and then filled by a cloze model.

In recent years, researchers have proposed some abstractive summarization evaluation metrics for factual consistency, which can be classified mainly into two categories: NLI-based metrics[5][6][7] and QA-based metrics[8–10]. We provide a simple comparison and description for them in Table 1.

As shown in Table 1, due to the one-stage computation with NLI modeling, NLI-based metrics have faster evaluation speed but poorer interpretability. They cannot locate the factual inconsistency, which is also an essential task in analyzing factual consistency. In contrast, QA-based metrics have strong interpretability based on the explicit definition of facts. However, their multi-stage computational process results in large computational overheads[11]. Furthermore, QA-based metrics are calculated in several independent modules, which results in accumulated errors. For example,

the QA module fails to generate answers to roughly 14% of the questions that can be answered[9], and questions generated by the QG module need to be carefully filtered.

Faced with the pros and cons of the two categories of metrics, we focus on the cloze task[12] and propose a cloze-based metric ClozE (**Cloz**e-based Factual Consistency **E**valuation), which is a fast and lightweight factual consistency evaluation framework that leverages the best of both NLI-based and QA-based metrics while avoiding their limitations. To the best of our knowledge, ClozE is the first cloze-based metric for evaluating factual consistency. Specifically, our contributions are the following:

(1) In order to maintain the interpretability, we leverage the extracting-comparing-facts framework in QA-based metrics. But instead of generating questions and answers with QG and QA modules, we simply use a cloze module, where the questions are defined as text spans whose factual factors are replaced with `[MASK]` and the cloze results are answers. This one-stage modeling can compute faster than QA-based metrics while being more interpretable than NLI-based ones. Meanwhile, it can maximize the use of knowledge from masked language modeling (MLM) based pre-trained language models (PLMs) due to the similarity between the upstream and downstream tasks.

(2) Experiments on two benchmarks show that ClozE can achieve a comparable correlation with human judgments as QA-based and NLI-based metrics. Compared to QA-based metrics, ClozE reduces the time consumption by 96% on average and obtains higher accuracy in generating answers, 30% above QA-based metrics. Meanwhile, ClozE is not only more interpretable but also more accurate for more abstract summaries than NLI-based metrics.

(3) We further analyze three important facets of ClozE in practice, including the balance of speed and performance, optimization of metric and location of factual errors. We demonstrate that ClozE performs better than QA-based and NLI-based metrics on these three aspects combined, fully showing the usefulness and great potential of our proposed framework.

2 Related Work

2.1 Non-factual metrics

N-gram overlap-based metrics were generally used for summarization before the concept of factual consistency was proposed, such as ROUGE[13], BLEU[14] and METEOR[15]. Recent work has applied pre-trained models to evaluate summaries. BERTScore[16], BLANC[17] and COMET[18] give scores by calculating the similarity between a candidate summary and reference summaries with Bert[19]. [20] proposes BLEURT, which pre-trains a fully learned metric on large amounts of synthetic data and fine-tunes it on human ratings.

2.2 NLI-based metrics

[5],[21] and [6] trained an NLI model to score candidate summaries for factual consistency and obtained a high correlation with human judgments.[7] divided a document

and a summary into multiple blocks and applied the NLI model to calculate the blocks for an entailment matrix to resolve the problem of the granularity mismatch.

2.3 QA-based metrics

SummaQA[22], QAGS[8], FEQA[9] and QuestEval[10] are all QA-based metrics which used a similar process as Table 1 to obtain the final score by comparing the extracted summary with the factual factors of the document. [23] proposed QUALS, which greatly improves the speed of Q&A by allowing the model to generate questions and answers simultaneously. [24] have experimented and optimized each module of the QA-based methods to achieve better performance than other QA-based metrics.

2.4 Factual Consistency Benchmark

Assessing the validity of factual consistency metrics is also a challenging task. In recent years, more work has given relevant fact consistency annotated datasets as well as meta-evaluation methods. [8][25] separately proposed a crowdsourced annotated dataset on the factual consistency, as a way to evaluate the correlation of factual consistency metrics to human judgments. [7] fused previous works and standardized their datasets. Meanwhile, the Pearson correlation coefficient is commonly used to evaluate the correlation between the above metrics and human judgments. In contrast, [1][26] proposed the meta-evaluation of factual consistency with more refined criteria for the performance evaluation of factual consistency metrics.

3 Methods

We now introduce our ClozE for factual consistency evaluation. Firstly, the overall framework of ClozE is shown in Table 1 and Figure 1. In the next subsections, we introduce the modules in our ClozE framework according to a practical evaluation process. In order to show the great potential and application value of our framework, the modules in our framework are very simple and basic.

3.1 ClozE

The modeling approach in QA-based metrics is naturally interpretable so that we adopted its framework[1] . In order to alleviate the problems existing in QA-based metrics which are mainly caused by the multi-module collaboration, we consider merging the QG and QA parts so that only one model will be computed under the whole framework.

Following the motivation, we introduce cloze task, an elegant but lightweight modeling method. We transform a question into a summary text with k factual factors masked off and the answer is generated by the cloze model. In this case, since questions can be obtained directly, we can omit the QG module and substitute the QA module with a cloze model. We present a formula similar to the one proposed by [8].

Given the document X, the summary Y and the masked summary (cloze question) $Y^{'}$, each of which consists of a series of tokens, we can obtain $P(A|X, Y^{'})$ and $P(A|Y)$,

[1] Extract factual factors and compare them.

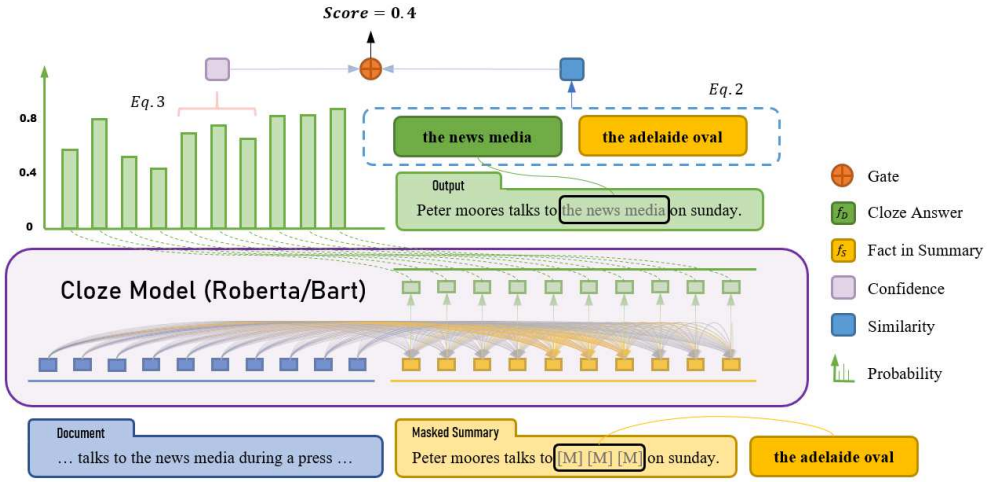


Fig. 1: We show the overall framework of ClozE. We first extract all factual factors f_s according to the summary and replace f_s with `[MASK]` identifiers at the token-level in the summary (`[M]` part of the Summary in the figure). By inputting the masked summary into the cloze model together with the document, we can later obtain the cloze answers f_D corresponding to the document at the position previously replaced with `[MASK]`. Finally, we use the equations given in Eq.2 and Eq.3 to calculate the final ClozE score.

which respectively represent the cloze answers generated by the cloze model and the factual factors in the summary Y. Therefore, the factual consistency score can be computed by the following formula.

$$E_{A\sim P(A|Y)}[D(P(A|X,Y^{'}),P(A|Y))] \tag{1}$$

where $D(*,*^{'})$ denotes a function that measures the distribution of both $*$ and $*^{'}$.

Due to the similarity to MLM, the cloze model in ClozE can be trained by prompt learning[27] and maximize the use of knowledge in PLMs. Meanwhile, it is also a self-supervised task and does not require annotated datasets, which prevents from severe data bias[28, 29] caused by data augmentation similar to FactCC[6]. Besides, it still maintains the same interpretability and is capable of factual error correction in some level.

3.2 Factual Factors

In the work of [26], factual errors shaped as named entities and noun phrases account for the majority of all factual errors, and both fit well in line with the task of cloze in our framework. As a result, we define named entities and noun phrases as factual factors in our framework.

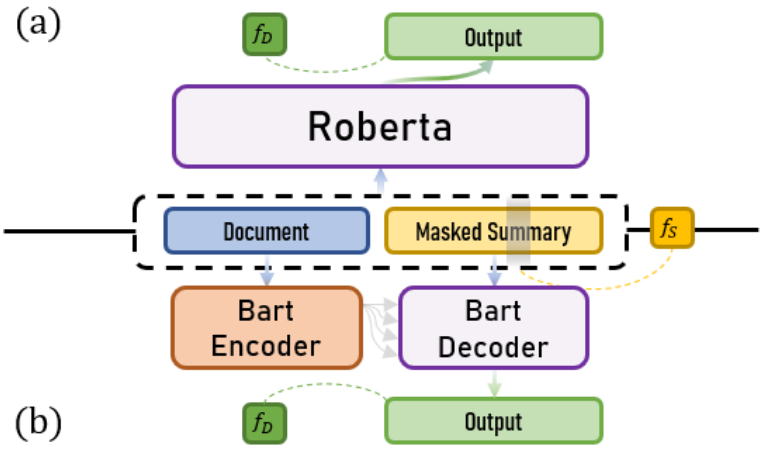


Fig. 2: Two kinds of cloze models are used for the ClozE. The first one is Roberta (a), where we concatenate the document and the masked summary as the input of the model. The second one is Bart (b), where we take the document as the Encoder input and the masked summary as the Decoder input.

3.3 Cloze Model

Cloze models occupy the key position in our work. For fairness considerations of most previous metrics (both QA-based and NLI-based metrics), we conduct experiments using two types of basic pre-trained models, Roberta [30] and Bart [31], as shown in Figure 2.

The approach using Roberta is a general form of our framework. The pre-training task of Roberta is a better match to the cloze task, where we concatenate the document and the masked summary together by feeding them into the model, expecting the model output to be a complete text.

We also use the Bart model because of its training and prediction runtimes. In our task, we actually do not require the model to regenerate the part of the document, so this part can be combined with the masked summary by cross-attention mechanism[2] instead of self-attention mechanism [32]. It is important to state that we leverage MLM task in the decoder instead of the original autoregressive task in Bart. This setting is intended to facilitate comparison with the Roberta-based ClozE model.

3.4 The ClozE Score

In line with the approach adopted in the QA-based metrics in recent years, we use token-level F1 as the quality for the comparison of factual factors. Because F1 focuses on whether a token appears in both two sequences, it is not sensitive to errors such as misalignment and missing tokens. This characteristic can help to mitigate the errors caused by the performance of the cloze model to some extent.

[2]It only computes the attention scores between the decoder hidden states and the encoder outputs in the decoding stage.

For a summary with N factual factors, we extract all of the factual factors and split them into $\lceil N/k \rceil$ groups, where k is the number of factual factors for a single masking. Therefore, a summary will need to be calculated $\lceil N/k \rceil$ times by the cloze model, and the final ClozE score is computed using the following equations:

$$Score = \frac{\sum_{i=1}^{\lceil N/k \rceil} \sum_j F1(G_{ij}, S_{ij})}{N} \tag{2}$$

where $F1(*, *^{'})$ denotes the function which computes the $F1$ score between the two symbols in the tuple. G_{ij} and S_{ij} respectively denote the factual factor in the summary and generated with the cloze model.

3.5 Granularity

Most of the factual consistency metrics calculate the factual consistency score with the sentence-level units. Out of preciseness, we take the summary level into account. We concatenate one document with one masked sentence instead of a full masked summary. This means that a masked summary is divided into multiple sentences as the inputs for cloze model. Afterwards, the multiple cloze answers are collected and the evaluation scores are calculated uniformly using Eq.2. We also conduct experiments with both levels.

3.6 Confidence

We take into account the confidence mentioned in the work of [22] and [33]. We propose a convenient and effective way to calculate the confidence by Eq.3, which is the average prediction probability of the tokens contained by the masked factual factors.

$$Confidence = \frac{1}{|T_f|} \sum_{t \in T_f} p(t) \tag{3}$$

where t denotes a token in T_f while T_f denotes all the tokens contained in factual factor f.

In addition, we introduce two thresholds α and β. When the confidence is less than α and the current ClozE score is less than β, the ClozE score will be set to 0. This trick is designed to make ClozE more robust by ignoring some of the score increases due to incorrectly filled blanks.

4 Experiments

4.1 Baseline Metrics

Since there has been a lot of previous work[7, 8, 10, 25] that has demonstrated the inability of model-free metrics to effectively measure factual consistency, we only consider model-based metrics as our baselines.

Non-Factual Baselines

We use BERTScore and BLANC-help to verify the sensitivity of the proposed method for factual consistency. We do not perform speed tests on the above metrics because they do not address the factual consistency issue.

Factual Baselines

We use QAGS, FEQA, SummaQA, QUALS, QuestEval and QAFactEval[24] as QA-based baselines and FactCC [6], ANLI[21] and SUMMAC[7] as NLI-based baselines. We verify that our method performs comparable with these factual consistency metrics while avoiding their limitations through experiments.

We test most of the baselines above with the code and models they provided for fairness. For ANLI, we used the results from [24] because of the lack of its code and model. Besides, due to some problems with the official QAGS code, we opted for the factsumm[3] implementation for QAGS. In addition, we do not use any acceleration strategies such as multi-processing and half-precision inference for all baselines.

4.2 Factual Evaluation Benchmark

We use two benchmarks to evaluate the baselines and our framework, including Pearson Benchmark and GO FIGURE Benchmark. We believe that more meta-evaluation will give a more objective evaluation for these metrics. In this section, we briefly describe these two benchmarks while providing a more detailed description in Appendix A.

4.2.1 Pearson Benchmark

Following [24], we use six human-annotated datasets to evaluate factual consistency metrics with Pearson correlation coefficient. The summaries used for evaluation are generated by a variety of different models on CNN/DM and XSum[34], respectively. These datasets are XSumFaith (XSF)[3], SummEval (SE)[25], FRANK[26] and QAGS[8]. We partition FRANK and QAGS into FRANK-CNN/DM (FCD), FRANK-XSum (FXS), QAGS-CNN/DM (QCD) and QAGS-XSum (QXS), resulting in a total of six test datasets.

4.2.2 GO FIGURE Benchmark

In addition, as suggested in [17], human scores are error-prone and being correlated with human scores should not be considered as a sufficient and unique validation, but rather as an independent confirmation of the reasonableness of the metrics. Therefore, we use GO FIGURE [1], which has developed a systematic criterion for factual consistency metrics, as another benchmark to evaluate baselines and our work from a different aspect.

[3] https://github.com/Huffon/factsumm

4.3 Experimental Details

4.3.1 Extractor

We use SpaCy [4] to extract named entities and noun phrases. We first used `en_core_web_sm`, the model adopted in previous QA-based metrics work, to extract factual factors. After that, we also adopted the `en_core_web_trf` model with higher accuracy to extract factual factors. `en_core_web_sm` and `en_core_web_trf` are both models used in SpaCy for text analysis, the former is designed for cpu only and the latter is based on Roberta's extraction model.

By using SpaCy, we obtained too many noun phrases, and some of them have overlapping parts with named entities. Therefore, we take named entities as the main part and remove those noun phrases containing overlapping parts, to get a set of factual factors that do not interfere with each other.

4.3.2 Cloze Model

We use transformers library [35] to build and employ Roberta-base and Bart-base pre-trained models. We denote the ClozE based on Roberta-base model and Bart-base model as ClozE-R and ClozE-B respectively.

4.3.3 Cloze Task Dataset

Since the input of our cloze model is a document and its summary, it is a good choice to use the summarization dataset to train the cloze model. And we believe that the golden summaries in the summarization dataset have high factual consistency. Due to this, we adopt the summarization dataset CNN/DM[36] as the dataset for our cloze task. We use their training set as the training set for our cloze task and their test set to evaluate the performance of our cloze model.

4.3.4 Training Settings

The input lengths of the models are limited to 512. For every training sample, we first select a random sentence in the summary and randomly mask a number of factual factors in that sentence for cloze task. Cross entropy loss with Adam [37] optimizer is used, with a batch size of 96, a learning rate of 1e-4 and a warming up step of 8000. We trained and tested our models on two NVIDIA GeForce RTX 3090 GPUs for 2-3 days with the parameters $k = 1$ (§3.4) and $\alpha = \beta = 0.5$ (§3.6).

4.4 Results and Analysis

For objectivity and comprehensiveness, we evaluate ClozE and baselines in multiple aspects of performance, time consumption, interpretability and convenience. We emphasize that all these aspects are equally important and our main purpose is to propose a new factual consistency evaluation framework with a strong potential rather than pursuing all aspects of state of the arts (SOTA).

[4] https://spacy.io/api

Type	Metric	XSF	SE	FCD	FXS	QCD	QXS	Overall	Sec./Sum	Inc.
NF	BERTScore	20.00	24.19	41.69	19.26	57.19	1.83	27.36	/	1
	BLANC-help	10.94	17.08	23.68	12.22	18.50	4.01	14.41		1
NLI	FactCC	4.78	42.11	51.87	8.15	57.63	34.16	33.12	**0.06**	1
	ANLI*	16.00	43.00	53.00	18.00	65.00	**39.00**	39.00	/	1
	SUMMAC	6.93	**49.67**	**56.19**	**32.24**	57.30	**44.01**	41.06	0.19	2
QA	SummaQA	14.82	14.69	16.14	6.19	28.28	9.90	15.00	3.39	1
	FEQA	29.44	43.71	52.50	14.96	55.81	31.72	39.28	3.00	2
	QAGS(factsumm impl.)	11.90	20.32	26.06	10.64	22.65	22.22	18.31	2.20	2
	QUALS	**36.81**	27.32	49.21	20.32	56.97	19.47	35.01	1.01	1
	QuestEval	**46.38**	37.75	52.66	**30.05**	53.39	27.81	**44.05**	4.37	2
	QAFactEval(lerc)	**33.93**	**61.06**	**67.60**	**38.42**	**71.66**	**49.03**	**54.53**	1.96	3
ClozE-R	**Base**	16.45	44.29	49.70	21.33	**65.76**	27.79	39.50	0.13	1
	`+sentence-level`	16.45	47.17	**57.27**	21.27	65.20	27.79	41.47	0.11	1
	`+en_core_web_trf`	18.66	47.09	55.38	23.10	**66.12**	32.36	42.07	0.23	1
	`+confidence`	19.22	**48.01**	55.12	23.54	65.59	29.37	**42.30**	0.23	1
ClozE-B	**Base**	20.05	38.17	45.21	19.31	63.39	20.56	37.23	**0.09**	1
	`+sentence-level`	20.08	35.33	42.12	19.27	61.07	20.56	25.57	**0.08**	1
	`+en_core_web_trf`	20.44	36.82	43.26	20.40	61.68	19.42	36.52	0.19	1
	`+confidence`	20.44	36.47	42.12	20.89	61.01	21.99	36.18	0.19	1

Table 2: Pearson correlation coefficients of each metric relative to human judgments, the time spent on the evaluation (**Sec./Sum**) and the inconvenience of metrics (**Inc.**). The coefficients under different signatures separately denote the **best**, **second** and **third** ones. **Sec./Sum** means the number of seconds it takes to evaluate a summary. The table also shows the results of the ablation study in the Section ClozE-R and ClozE-B, where **Base** denotes `en_core_web_sm+summary-level` setting.

4.4.1 Pearson Benchmark

We show the Pearson correlation coefficients between human-annotated judgements and various metrics in Table 2. To facilitate the overall analysis, we follow [7] and compute the average as the overall performance on the benchmark. For all the datasets, the correlation scores are sufficiently plausible with $p < 0.01$.

Our framework ranks 3rd out of 11 baselines by overall scores on six datasets. And for multiple benchmarks, our framework has a high ranking compared to the previous methods. Though we do not obtain the best performance, the results also show that our framework has a strong potential and a comparable performance with other framework metrics.

We also analyze the gap between our framework and the better ones. Compared with QuestEval and QUALS, our framework performs worse on XSF while better on the other datasets. It indicates that we are required to pay more attention to the special data distribution on XSF. And compared with QAFactEval, the scores of our framework are lower than theirs. We believe that the better performances of QAFactEval are due to their optimization for each module in QA-based metric, while our ClozE only uses the most basic PLMs. We will carry out further optimization work afterward.

And under different cloze models, ClozE-B does not perform as well as ClozE-R. The reason is that the Bart decoder uses a causal mask mechanism, which makes the model unable to perceive the following information of the summary text and leads to the inability to fill in some blanks correctly.

	FactCC	SUMMAC	SummaQA	FEQA	QUALS	QuestEval	QAFactEval(f1)	ClozE-R	ClozE-B
Upper Bound (%)	44.04	27.14	10.87	37.56	-75.55	34.63	46.63	70.10	48.90
Level 1 (%)	36.44/38.71	16.07/11.28	9.58/10.63	33.35/36.64	-86.26/-82.43	31.92/32.31	37.27/37.70	64.80/69.28	45.58/48.44
Level 2 (%)	32.30/33.80	10.70/0.30	8.53/10.38	28.46/36.13	-99.83/-87.86	30.41/30.67	31.66/32.07	61.00/69.05	43.14/47.97
Level 3 (%)	29.58/28.89	2.99/-9.99	7.54/10.26	25.12/35.63	-107.22/-92.61	28.45/29.21	27.34/25.79	57.16/68.40	40.38/47.60
Lower Bound (%)	22.29	-18.87	0.12	0.58	-373.01	0.68	0.01	18.77	7.38
Correlation	-0.99/-1.00	-0.99/-1.00	-1.00/-0.98	-0.99/-1.00	-0.98/-1.00	-1.00/-1.00	-0.99/-0.99	-1.00/-0.96	-1.00/-1.00
p-value	0.08/<0.01**	0.07/0.01*	0.01*/0.13	0.07/<0.01**	0.11/0.02*	0.04*/0.02*	0.05*/0.02*	<0.01**/0.17	0.02*/0.04*

Table 3: Results of various metrics under the GO FIGURE benchmark on CNN/DM. Upper and Lower Bound denote the scores given by the metrics under golden summaries and random texts respectively. Level * denotes the scores given by the metrics under error level *. Correlation denotes the correlation between the metrics and error level. In each cell (*/*), entity factual errors are on the left and non-entity factual errors are on the right. Results marked red and green respectively denote bad and good cases that deserve attention.

	FactCC	SUMMAC	SummaQA	FEQA	QUALS	QuestEval	QAFactEval(f1)	ClozE-R	ClozE-B
Upper Bound (%)	22.40	-33.35	4.48	27.87	-141.69	19.95	14.70	65.00	44.13
Level 1 (%)	15.59/13.96	-35.70/-43.58	3.80/4.31	23.20/26.94	-153.31/-150.59	18.14/17.47	9.79/9.48	57.94/64.43	45.58/48.44
Level 2 (%)	11.97/7.86	-37.13/-49.84	3.40/4.22	20.05/26.55	-163.81/-155.65	16.78/16.01	7.15/7.15	51.63/63.97	43.14/47.97
Level 3 (%)	11.34/5.67	-38.96/-54.28	3.13/4.14	15.81/26.06	-176.09/-159.92	14.84/14.44	5.11/5.11	47.55/63.09	40.38/47.60
Lower Bound (%)	23.68	-43.16	0.26	1.18	-330.03	0.23	0.00	35.95	10.70
Correlation	-0.92/-0.96	-1.00/-1.00	-0.99/-1.00	-1.00/-0.99	-0.99/-0.99	-0.99/-1.00	-0.99/-0.99	-0.99/-0.98	-1.00/-1.00
p-value	0.25/0.17	0.04*/0.06	0.07/0.03*	0.05*/0.04*	0.03*/0.03*	0.06/0.01*	0.05*/0.02*	0.07/0.11	0.02*/0.04*

Table 4: Results of various metrics under the GO FIGURE benchmark on XSum. The symbolic meaning is the same as that in Table 3.

4.4.2 GO FIGURE Benchmark

We present the results for each baseline under GO FIGURE benchmark in Table 3 and Table 4. We analyze the GO FIGURE results in terms of three aspects, including factual sensitivity, lower bound and upper bound.

Factual Sensitivity

Factual sensitivity is mainly evaluated by Correlation and p-value. The results show that all baselines and ClozE are negatively correlated with the factual error level and the p-values are small enough to be credible on both CNN/DM and XSum datasets. It indicates that ClozE is sensitive to both entity and non-entity factual errors, which is consistent with the existing factual consistency evaluation metrics.

Lower Bound

The "Lower Bound" in GO FIGURE represents the random text, which is not relevant to the document. We find only FactCC gains a higher score in "Lower Bound" than that in "Upper Bound" on XSum, which is a bad case for factual consistency metrics. It shows that FactCC is not robust on more abstractive text.

Upper Bound

"Upper Bound" in GO FIGURE represents the summary text without factual errors, which is the golden summary actually. This means that for extracting-comparing-facts metrics like QA-based metrics and ClozE, the scores they give to "Upper Bound"

is the accuracy of the factual factors they generate. The results show that most of QA-based metrics generate low-accurate answers while ClozE-R can generate more accurate ones (**+30**% on average). In addition, the scores given by baselines on XSum are much lower than those on CNN/DM while ClozE maintains a better performance, which shows that our framework is better and more robust while evaluating abstract texts.

4.4.3 Time Consumption

We use QAGS-CNN/DM to measure the time consumption on the same device. As shown in Table 2, our method has an extremely significant improvement over the QA-based metrics and comparable with NLI-based metrics, which reduces the time consumption by nearly 96% on average compared to QA-based metrics. This great improvement is mainly due to : (1) we use a smaller model compared to the models in most of QA-based metrics, and (2) most importantly, we discard the most time-consuming QG module, which greatly reduces the runtimes for evaluation.

We also observe that ClozE-R takes nearly twice the evaluation time of ClozE-B, whose reason is mentioned in §3.3.

4.4.4 Interpretability

We report the interpretability of ClozE by combining the computational procedure of the three metrics in Table 1 and Figure 1. Obviously, our ClozE framework has the same procedure as QA-based metrics, i.e., extracting factual factors for comparison. This fact indicates that our ClozE has the same interpretability as QA-based metrics, which is stronger than that of NLI-based metrics.

4.4.5 Convenience

We believe that the convenience of the metric is as important as the aspects above. To simplify the convenience measurement, we evaluate the inconvenience of the metrics instead of the convenience. We give the scores directly through the number of models involved in the metrics. As shown in Inc. section in Table 2, most of the QA-metrics have two or more models, while others, including our ClozE, require only one model.

4.5 Ablation Study

In Table 2, we show the main ablation experiment results, and in Table B1 in Appendix B, we present the full ablation experimental results.

Granularity

In terms of granularity, there is not much difference between summary-level and sentence-level. We believe that the main reason for this is that the candidate summary sentences in the dataset are not very related to each other. Therefore, neither the single-sentence input nor the whole candidate summary input affects the accuracy of the sentence-level cloze.

Factual Factors Extractor

We used two factual factors extractors, `en_core_web_sm` (`SM`) and `en_core_web_trf` (`TRF`) respectively. `SM` has less time spent and, relatively, it extracts a much lower number and accuracy of factual factors than `TRF`. Due to this, we can observe that the time spent on the results with `SM` is much less than that with `TRF`, but the correlation is also basically lower than that with `TRF`.

Confidence

We also conducted experiments on the Confidence proposed in §3.6. Because it is represented simply by the generation probabilities of factual factors, it has no impact on the computational speed of ClozE. Despite its simple design, it is also a certain effect enhancement for ClozE. It can help ClozE discard some uncertain cloze results to improve the accuracy of evaluating factual consistency.

4.6 Case Study

In Table 5, we present three examples, each of which contains document, summary, masked summary, correct answers and cloze results with similarity scores. Obviously, there are only two cases for the cloze results, which are filling in correct factual factors and incorrect ones.

On most of the examples, ClozE fills in the same factual factors when the masked ones are correct, which is a common phenomenon throughout ClozE and one of the reasons for its high performance. Meanwhile, we also note that ClozE replaces some incorrect factual factors with correct ones, such as `"The president"` → `"brady"` on the first example and `"cameron"` → `"miliband"` on the third example. This is a delightful surprise, indicating that ClozE may be able to correct the wrong factual factors.

As shown on the second and third examples, ClozE sometimes fills in incorrect answers when the masked factual factors are incorrect either. We believe that the major reason is due to the limitations of autoencoding pre-trained models, which are weak on non-fixed-length generating tasks [38].

However, this issue does not lead to poor performance in factual consistency evaluation. As shown in Eq.2, the ClozE score is calculated by comparing the factual factors in the summary with those filled in by the cloze model. This means that even if a factual factor generated by the cloze model is wrong, we can still obtain a precise ClozE score, which is verified on the second and third examples.

5 Discussion

In this section, we further discuss three important facets of ClozE, including the balance of speed and performance, the optimization of metric and location of factual errors.

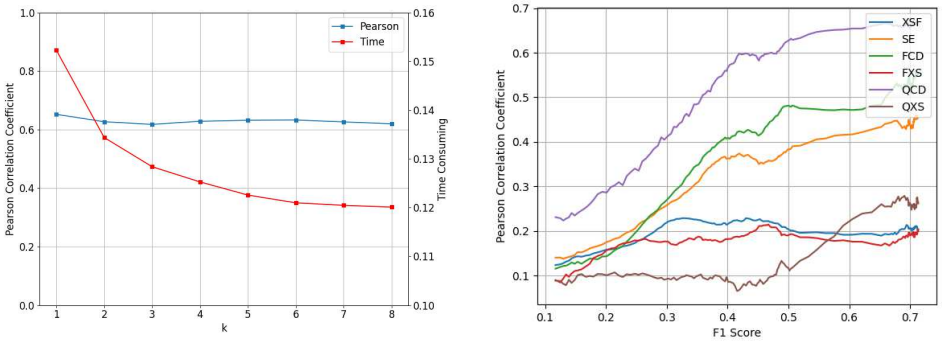
Balance of Speed and Performance

The number of questions that can be answered at once is critical to the speed of evaluation. The low speed of QA-based metrics is due to the limitation of answering

Example #1	
Document	[. . .] The patriots beat seattle seahawks 28 - 24 to win the nfl super bowl in glendale, arizona with brady named as mvp for the third time in his fourth super bowl victory. Although brady wasn't able to visit the white house with [. . .] He joins the likes of joe montana and terry bradshaw on four rings, and is now the all-time leader with 12 super bowl touchdown passes and completions on 37. [. . .]
Summary	The new england patriots beat seattle seahawks 28-24 to win super bowl xlix in glendale, arizona with brady named as mvp for the third time in his fourth super bowl victory. [...] The president was a senator with 12 super bowl touchdown passes and completions on 37.
Masked Summary	The new england patriots beat [MASK] 28-24 to win super bowl xlix in glendale, [MASK] with brady named as mvp for the third time in his fourth super bowl victory. [MASK] was a senator with 12 super bowl touchdown passes and completions on 37.
Correct Answers	seattle seahawks (1.0), arizona (1.0), brady (0.0)
Cloze Results	seattle seahawks (1.0), arizona (1.0), brady (0.0)
Example #2	
Document	[. . .] Nominations are open for cnn heroes 2015. A decade - long civil war had just ended in the country, and doyne witnessed its effects firsthand. She met women and children who were suffering, struggling to survive. ' 'it changed me," said doyne, now 28. [. . .] Today, kopila - - which means " flower bud" in nepali - - is home to about 50 children, from infants to teenagers. [. . .]
Summary	Nominations are open for cnn heroes 2015. Doyne, nepal, met women and children in nepal. The school is home to about 50 children.
Masked Summary	Nominations are open for [MASK] heroes 2015. Doyne, nepal, met women and children in [MASK]. [MASK] is home to about 50 children.
Correct Answers	cnn (1.0), [None] (0.0), kopila (0.0)
Cloze Results	cnn (1.0), maggie do (0.0), doyne (0.0)
Example #3	
Document	[. . .] Cameron splayed his legs, he smiled on cue and when he swivelled his huge eyes direct to camera to repeatedly use the phrase 'if i am prime minister, it was [. . .] Miliband's coaching team will be patting themselves on the back and declaring a win following the debate. [. . .]
Summary	Nicola sturgeon is a prime minister and miliband's coaching team. [...] David cameron is a coaching team for cameron.
Masked Summary	[MASK] is a prime minister and miliband's coaching team. [...] David cameron is a coaching team for [MASK].
Correct Answers	cameron (0.0), miliband (0.0)
Cloze Results	m cameronons (0.0), miliband (0.0)

Table 5: Examples generated by ClozE on the QAGS-CNN/DM dataset. Where green words indicate correct factual factors, red ones indicate incorrect factual factors, and [None] indicates that it cannot be answered. The (*) after each factual factor is the similarity score obtained by Eq.2.

one question at once. And for NLI-based metrics, it can be considered as answering infinite questions at once, which improves the speed of evaluation greatly. Following this characteristic of NLI-based metrics, our ClozE can mask k factual factors at once (§3.4), which is equivalent to simultaneously generate and answer k questions. We present the performance and speed of ClozE with different k in Figure 3(a) and more results on different datasets in Figure B1 in Appendix B. We use summary-level granularity for a larger range of k and greedily mask the k consecutive factual factors together one by one. We observe that the time (in seconds) taken by ClozE decreases linearly as the number k increases. The performance of the model is the highest when $k = 1$ and it decreases slowly when $k > 1$, reflected by the Pearson



(a) (b)

Fig. 3: (a) The trend of ClozE evaluation speed and correlation with human judgments given different k settings on QAGS-CNN/DM. (b) The trend of cloze accuracy (F1 Score) and correlation with human judgments on different datasets.

correlation coefficient. The reason is that there may be causal relationships between factual factors, and when we replace them with `[MASK]` at the same time, the cloze model will have a degradation of performance due to the lack of information. In some cases, we can sacrifice a little performance in this way in exchange for faster evaluation.

Optimization of Metric

One direction of optimization for model-based metrics is to reduce the cumulative errors between modules. The QG and QA modules of QA-based metrics lead to errors accumulating between modules and complicates the study of how to improve the QA-based metrics. In contrast, NLI-metrics alleviate this issue due to the use of only one model. Following this insight, we perform further experiments and analyze on ClozE to explore the correlation between cloze performance and metric performance. Because the performance of cloze depends on the structure of the model as well as the degree of training, we test both Roberta-base and Bart-base model and their checkpoints at different training steps (Figure 3(b)) and observe a clear upward trend on most of the datasets. Although there is a slight decrease on XSF and FXS, the decrease is not significant and acceptable. Meanwhile, §4.6 shows that ClozE does only care whether the generated answers are correct when the factual factors in the summary are correct, which means our method is equivalent to the autoencoding pre-trained models with MLM. Overall, we can assume that the performance of our ClozE can be improved when we have a better autoencoding pre-trained model as the cloze model, which brings great convenience in further improving factual consistency metrics.

Location of Factual Errors

In practice, we usually not only need the metric to give us a score, but also want it to help us analyze where the factual inconsistency errors occur, so as to help us

improve our work. As mentioned above, NLI-metrics are lack of interpretability, so it cannot locate where the errors occur. In contrast, our ClozE can naturally locate the incorrect factual factors when it computes scores, which is similar to QA-based metrics. As shown in Table 5, the factual factors in the summary that differ from the answers generated by the cloze model will be considered as the factual errors.

6 Conclusion

In this paper, we propose ClozE – a cloze-based factual consistency evaluation framework, which leverages the best of both NLI-based and QA-based metrics while avoiding their limitations. It carries the strong interpretability and fast evaluation speed while develops more advantages and the comparisons at all aspects in §4.4 and §5 also show its great potential. We believe that a faster and effective factual consistency metric can greatly promote the study of factual consisitency in abstractive summarization and hope our findings in this paper will provide insights to future work in all the factual consistency related research. Moreover, we will continue to further research for evaluating more abstractive summarization, higher overall performance and faster speed.

Declarations

This work was supported by the National Natural Science Foundation of China (Grant No. 62176024) and Engineering Research Center of Information Networks, Ministry of Education.

Appendix A Details of Benchmarks

We will show the details of the benchmarks we use in this section.

XSumFaith

[3] focused on XSum because the golden summaries of this dataset are more abstract and can bring out more of the factual consistency issues of the system summaries. They first asked the annotators to label various types of hallucinations in the system summary and then calculated the factual consistency score of the system summary based on the number of these hallucinations as a reference.

SummEval

[25] randomly selected 100 articles from the CNN/DailyMail test set and generated numerous system summaries for crowd-sourced and expert annotation by means of an existing neural summary model. Also, the annotation was done in the same way as [39], and the model output was evaluated according to the four aspects of consistency, coherence, fluency, and relevance. We used the score of consistency as the result of the factual consistency assessment of the systematic summary of the document.

FRANK

[26] included model summaries from CNN/DM and XSum datasets. They annotated each sentence of a summary with numerous error types based on their proposed criteria. In order to obtain the human evaluation score, we consider all error types uniformly and compute the score based on the percentage of errors.

QAGS

[8] also took into account the system summaries generated on both the CNNDM and XSum datasets. They collected 3 annotations for each summary and gave a score for a single sentence by voting. Then all sentences were averaged to get the final human evaluation score.

GO FIGURE

[1] differed from previous work on evaluating the factuality of summaries in that they tried to give a criterion to rate the performance of these aforementioned metrics for evaluating factual consistency. They gave five objective conditions: (1) the evaluation score in the generated summary should be greater than that in the completely irrelevant document and less than or equal to that in the golden summary; (2) the evaluation score is negatively correlated with the number of factual errors; (3) the performance is not affected by different factual error types; (4) the performance is not affected by different domain data; (5) the evaluation score is highly correlated with the human judgments.

Appendix B More Experiments

Type	Metric	XSF	SE	FCD	FXS	QCD	QXS	Sec./Sum
ClozE-R	**summary-level**							
	+en_core_web_sm	16.45	44.29	49.70	21.33	65.76	27.79	0.13
	+confidence	17.19	45.60	49.39	21.93	**66.30**	24.97	0.13
	summary-level							
	+en_core_web_trf	18.67	43.81	49.31	23.15	65.24	**32.36**	0.23
	+confidence	19.20	45.05	48.71	**23.63**	65.90	29.37	0.23
	sentence-level							
	+en_core_web_sm	16.45	47.17	**57.27**	21.27	65.20	27.79	0.11
	+confidence	17.09	47.97	56.82	21.89	64.93	24.97	0.11
	sentence-level							
	+en_core_web_trf	18.66	47.09	55.38	23.10	66.12	**32.36**	0.21
	+confidence	19.22	**48.01**	55.12	23.54	65.59	29.37	0.21
ClozE-B	**summary-level**							
	+en_core_web_sm	20.05	38.17	45.21	19.31	63.39	20.56	0.09
	+confidence	19.63	37.52	43.84	20.12	64.39	21.62	0.09
	summary-level							
	+en_core_web_trf	**20.47**	40.08	46.74	20.38	63.72	19.60	0.19
	+confidence	20.46	39.52	45.47	20.86	65.13	22.16	0.19
	sentence-level							
	+en_core_web_sm	20.08	35.33	42.12	19.27	61.07	20.56	**0.08**
	+confidence	19.64	34.77	40.48	20.09	60.69	21.62	**0.08**
	sentence-level							
	+en_core_web_trf	20.44	36.82	43.26	20.40	61.68	19.42	0.18
	+confidence	20.44	36.47	42.12	20.89	61.01	21.99	0.18

Table B1: Full ablation experiments on six datasets. We also set the parameters $k = 1$ and $\alpha = \beta = 0.5$.

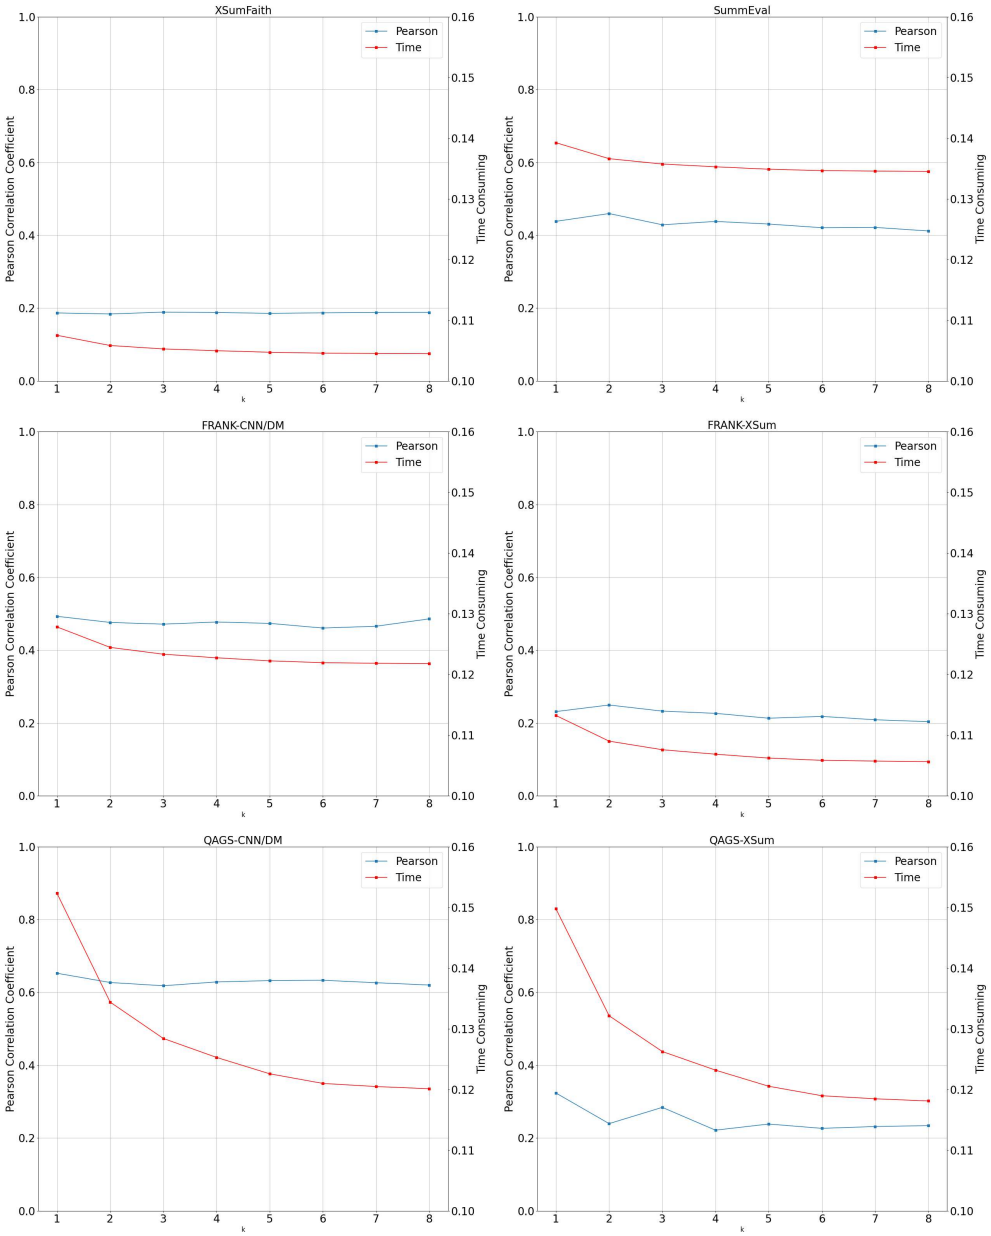


Fig. B1: The effect of different k settings on ClozE performance on six datasets.

References

[1] Gabriel, S., Celikyilmaz, A., Jha, R., Choi, Y., Gao, J.: Go figure: A meta evaluation of factuality in summarization. In: Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021, pp. 478–487 (2021)

[2] Cao, Z., Wei, F., Li, W., Li, S.: Faithful to the original: Fact aware neural abstractive summarization. In: Thirty-second AAAI Conference on Artificial Intelligence (2018)

[3] Maynez, J., Narayan, S., Bohnet, B., McDonald, R.: On faithfulness and factuality in abstractive summarization. In: Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, pp. 1906–1919 (2020)

[4] Deutsch, D., Roth, D.: Understanding the extent to which content quality metrics measure the information quality of summaries. In: Proceedings of the 25th Conference on Computational Natural Language Learning, pp. 300–309 (2021)

[5] Falke, T., Ribeiro, L.F., Utama, P.A., Dagan, I., Gurevych, I.: Ranking generated summaries by correctness: An interesting but challenging application for natural language inference. In: Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, pp. 2214–2220 (2019)

[6] Kryściński, W., McCann, B., Xiong, C., Socher, R.: Evaluating the factual consistency of abstractive text summarization. In: Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP), pp. 9332–9346 (2020)

[7] Laban, P., Schnabel, T., Bennett, P.N., Hearst, M.A.: Summac: Re-visiting nli-based models for inconsistency detection in summarization. Transactions of the Association for Computational Linguistics **10**, 163–177 (2022)

[8] Wang, A., Cho, K., Lewis, M.: Asking and answering questions to evaluate the factual consistency of summaries. In: Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, pp. 5008–5020 (2020)

[9] Durmus, E., He, H., Diab, M.: Feqa: A question answering evaluation framework for faithfulness assessment in abstractive summarization. In: Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, pp. 5055–5070 (2020)

[10] Scialom, T., Dray, P.-A., Gallinari, P., Lamprier, S., Piwowarski, B., Staiano, J., Wang, A.: Questeval: Summarization asks for fact-based evaluation. In: Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, pp. 6594–6604 (2021). Association for Computational Linguistics

[11] Koto, F., Baldwin, T., Lau, J.H.: Ffci: A framework for interpretable automatic

evaluation of summarization. Journal of Artificial Intelligence Research **73**, 1553–1607 (2022)

[12] Taylor, W.L.: "cloze procedure": A new tool for measuring readability. Journalism quarterly **30**(4), 415–433 (1953)

[13] Lin, C.-Y.: Rouge: A package for automatic evaluation of summaries. In: Text Summarization Branches Out, pp. 74–81 (2004)

[14] Papineni, K., Roukos, S., Ward, T., Zhu, W.-J.: Bleu: a method for automatic evaluation of machine translation. In: Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics, pp. 311–318 (2002)

[15] Banerjee, S., Lavie, A.: Meteor: An automatic metric for mt evaluation with improved correlation with human judgments. In: Proceedings of the Acl Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation And/or Summarization, pp. 65–72 (2005)

[16] Zhang, T., Kishore, V., Wu, F., Weinberger, K.Q., Artzi, Y.: Bertscore: Evaluating text generation with bert. In: International Conference on Learning Representations (2019)

[17] Vasilyev, O., Dharnidharka, V., Bohannon, J.: Fill in the blanc: Human-free quality estimation of document summaries. In: Proceedings of the First Workshop on Evaluation and Comparison of NLP Systems, pp. 11–20 (2020)

[18] Rei, R., Stewart, C., Farinha, A.C., Lavie, A.: Comet: A neural framework for mt evaluation. In: Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP), pp. 2685–2702 (2020)

[19] Kenton, J.D.M.-W.C., Toutanova, L.K.: Bert: Pre-training of deep bidirectional transformers for language understanding. In: Proceedings of NAACL-HLT, pp. 4171–4186 (2019)

[20] Sellam, T., Das, D., Parikh, A.: Bleurt: Learning robust metrics for text generation. In: Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, pp. 7881–7892 (2020)

[21] Barrantes, M., Herudek, B., Wang, R.: Adversarial nli for factual correctness in text summarisation models. arXiv preprint arXiv:2005.11739 (2020)

[22] Scialom, T., Lamprier, S., Piwowarski, B., Staiano, J.: Answers unite! unsupervised metrics for reinforced summarization models. In: 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), pp. 3237–3247 (2019). Association for Computational Linguistics

[23] Nan, F., Santos, C.N.d., Zhu, H., Ng, P., McKeown, K., Nallapati, R., Zhang, D., Wang, Z., Arnold, A.O., Xiang, B.: Improving factual consistency of abstractive summarization via question answering. arXiv preprint arXiv:2105.04623 (2021)

[24] Fabbri, A.R., Wu, C.-S., Liu, W., Xiong, C.: Qafacteval: Improved qa-based factual consistency evaluation for summarization. arXiv preprint arXiv:2112.08542 (2021)

[25] Fabbri, A.R., Kryściński, W., McCann, B., Xiong, C., Socher, R., Radev, D.: Summeval: Re-evaluating summarization evaluation. Transactions of the Association for Computational Linguistics **9**, 391–409 (2021)

[26] Pagnoni, A., Balachandran, V., Tsvetkov, Y.: Understanding factuality in abstractive summarization with frank: A benchmark for factuality metrics. In: Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, pp. 4812–4829 (2021)

[27] Liu, P., Yuan, W., Fu, J., Jiang, Z., Hayashi, H., Neubig, G.: Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. arXiv preprint arXiv:2107.13586 (2021)

[28] Cao, M., Dong, Y., Wu, J., Cheung, J.C.K.: Factual error correction for abstractive summarization models. In: Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP), pp. 6251–6258 (2020)

[29] Goyal, T., Durrett, G.: Annotating and modeling fine-grained factuality in summarization. In: Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, pp. 1449–1462 (2021)

[30] Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., Stoyanov, V.: Roberta: A robustly optimized bert pretraining approach. arXiv preprint arXiv:1907.11692 (2019)

[31] Lewis, M., Liu, Y., Goyal, N., Ghazvininejad, M., Mohamed, A., Levy, O., Stoyanov, V., Zettlemoyer, L.: Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In: Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, pp. 7871–7880 (2020)

[32] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. Advances in neural information processing systems **30** (2017)

[33] Qasemi, E., Kezar, L., Pujara, J., Szekely, P.: Evaluating machine common sense via cloze testing. arXiv preprint arXiv:2201.07902 (2022)

[34] Narayan, S., Cohen, S., Lapata, M.: Don't give me the details, just the summary! topic-aware convolutional neural networks for extreme summarization. In: 2018 Conference on Empirical Methods in Natural Language Processing, pp. 1797–1807 (2018). Association for Computational Linguistics

[35] Wolf, T., Debut, L., Sanh, V., Chaumond, J., Delangue, C., Moi, A., Cistac, P., Rault, T., Louf, R., Funtowicz, M., et al.: Huggingface's transformers: State-of-the-art natural language processing. arXiv preprint arXiv:1910.03771 (2019)

[36] Hermann, K.M., Kocisky, T., Grefenstette, E., Espeholt, L., Kay, W., Suleyman, M., Blunsom, P.: Teaching machines to read and comprehend. Advances in neural information processing systems **28** (2015)

[37] Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)

[38] Du, Z., Qian, Y., Liu, X., Ding, M., Qiu, J., Yang, Z., Tang, J.: Glm: General language model pretraining with autoregressive blank infilling. In: Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pp. 320–335 (2022)

[39] Kryściński, W., Keskar, N.S., McCann, B., Xiong, C., Socher, R.: Neural text summarization: A critical evaluation. 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), pp. 540–551 (2019)