

FEQA: A Question Answering Evaluation Framework for Faithfulness Assessment in Abstractive Summarization

Esin Durmus*
Cornell
University
ed459@cornell.edu

He He
New York
University
hehe@cs.nyu.edu

Mona Diab
The George Washington
University
mtdiab@gwu.edu

Abstract

Neural abstractive summarization models are prone to generate content inconsistent with the source document, i.e. *unfaithful*. Existing automatic metrics do not capture such mistakes effectively. We tackle the problem of evaluating faithfulness of a generated summary given its source document. We first collected human annotations of faithfulness for outputs from numerous models on two datasets. We find that current models exhibit a trade-off between *abstractiveness* and *faithfulness*: outputs with less word overlap with the source document are more likely to be unfaithful. Next, we propose an automatic question answering (QA) based metric for faithfulness, FEQA,¹ which leverages recent advances in reading comprehension. Given question-answer pairs generated from the summary, a QA model extracts answers from the document; non-matched answers indicate unfaithful information in the summary. Among metrics based on word overlap, embedding similarity, and learned language understanding models, our QA-based metric has significantly higher correlation with human faithfulness scores, especially on highly abstractive summaries.

1 Introduction

Abstractive summarization models must aggregate salient content from the source document(s) and remain *faithful*, i.e. being factually consistent with information in the source documents. Neural abstractive models are effective at identifying salient content and producing fluent summaries (See et al., 2017; Chen and Bansal, 2018; Gehrmann et al., 2018). However, the generated summary may not always contain faithful information, which is vital for real-world applications.

*Most of the work is done while the authors were at Amazon Web Services AI.

¹Faithfulness Evaluation with Question Answering.

Source. The world’s oldest person has died a few weeks after celebrating her 117th birthday. **Born on March 5, 1898**, the great-grandmother had lived through two world wars, the invention of the television and the first successful powered aeroplane flight by the wright brothers...

Output sentence. The world ’s oldest person has **died on March 5, 1898**.

Table 1: An example of *unfaithful* output (highlighted in red); generated by Gehrmann et al. (2018).

Table 1 shows an example of unfaithful generation. Recent studies have shown that around 30% of generated summaries contain unfaithful information (Cao et al., 2018; Falke et al., 2019a; Kryściński et al., 2019), especially when the sentence combines content from multiple source sentences (Lebanoff et al., 2019).

In this paper, we address the problem of evaluating *faithfulness* of generated summaries given their source documents. Our key insight is that current models are limited by a trade-off between abstractiveness and faithfulness (Section 2). On a wide range of systems and two datasets with varying levels of abstractiveness (CNN/DM and XSum), we show that the number of unfaithful sentences (annotated by humans) increases as the summary becomes more abstractive (i.e. less overlap with the source document). Next, we investigate a diverse set of existing automatic evaluation metrics such as ROUGE, BERTScore (Zhang et al., 2019a), and learned entailment models. We find that their correlations with human scores of faithfulness drop significantly on highly abstractive summaries, where deeper text understanding beyond surface similarity is needed.

Recently, question answering (QA) based automatic metrics have been proposed for evaluating

content selection in summarization (Eyal et al., 2019; Scialom et al., 2019; Chen et al., 2018). Specifically, cloze-style QA is used to evaluate whether important information in the source is recovered from the summary. Inspired by prior work, we use *automatically generated* QA pairs to represent information in the summary and validate it against the source. Concretely, we generate a set of “groundtruth” QA pairs from the summary, using a learned model that converts a declarative sentence and an answer span to a question (Section 3). Then, off-the-shelf reading comprehension models are evaluated on this set by extracting answer spans from the *source documents*. High accuracy means that the summary and the source document tend to produce the same answers, thus they are factually consistent with respect to the questions. Compared to prior approaches using cloze tests, our question generation approach enables evaluation with a broader range of QA models and answer types (e.g. extractive and generative), thus maximally taking advantage of progress in QA.

Among automatic metrics based on n -gram overlap, word embeddings, and language understanding models (relation extraction and entailment), FEQA has significantly higher correlation with human scores of faithfulness and is the only metric that correlates with human scores on highly abstractive summaries from XSum.

2 The Abtractiveness-Faithfulness Tradeoff

While extractive summarizers are largely faithful (since they copy sentences from the source document), current abstractive models struggle to produce faithful summaries without copying. Similar to Lebanoff et al. (2019), we observe that factual errors occur more frequently as models generate more abstractive summary sentences, i.e. less overlap with the source document. In this section, we analyze generated summaries along two dimensions: abtractiveness and faithfulness. Specifically, we aim to answer the following questions: (1) How to quantify abtractiveness of a summary? (2) Is abtractiveness encouraged more by the data or the model? (3) How does being abstractive affect faithfulness?

2.1 Characterizing Abtractiveness of a Summary

Abstractive summarization involves rephrasing important content into brief statements, ranging from minor editing of a source sentence to condensing multiple sentences in new words. Given a source document and a summary, we want to measure the level of abtractiveness of the summary.

Prior work measures abtractiveness by overlapped text spans between the summary and the document (Grusky et al., 2018; Zhang et al., 2018), or indirectly by the effectiveness of extractive baselines such as LEAD-3 (Nallapati et al., 2016a). While metrics such as extractive fragment coverage and density (Grusky et al., 2018) provide a continuous measure of the level of abtractiveness, we define a more fine-grained categorization of abtractiveness by analyzing how each sentence in the summary is formed.

A more abstractive summary sentence aggregates content over a larger chunk of source text; consequently it must copy fewer words to maintain brevity. Therefore, we define the following abtractiveness types based on the amount of copying, e.g. copying a source sentence, one or more partial fragments from the source sentence, and individual words.

1. **Sentence extraction:** the summary sentence is exactly the same as one of the source sentences.
2. **Span extraction:** the summary sentence is a substring of one of the source sentences, e.g. “*the plane was coming back from the NCAA final*” is a span extracted from “*the plane was coming back from the NCAA final, according to spokesman John Twork*”.
3. **Word extraction:** the summary sentence is formed by a subset of the tokens in a source sentence, e.g. “*Capybara Joejoe has almost 60,000 followers*” is a result of deleting words in “*Capybara Joejoe who lives in Las Vegas has almost 60,000 followers on Instagram*”.
4. **Perfect fusion_k:** the summary sentence is constructed by piecing together the substrings from k ($k > 1$) source sentences in their original order, e.g. “*Capybara Joejoe has almost 60,000 followers*” is a perfect fusion of the sentences “*Capybara Joejoe lives*

in Las Vegas.” and “He has almost 60,000 followers on Instagram.”

To quantify the amount of abtractiveness of a set of summaries, we label each sentence with the first qualified type in the order above if it fits to one of these categories.

We then define the score of each type as the percentage of sentences labeled by that category. The types are ordered by increasing levels of abtractiveness. For example, a summary with higher fusion scores and lower extraction scores is considered more abtractive. In addition, we compute the percentage of **novel n -grams** that do not appear in the source document as another metric for abtractiveness.

2.2 Is abtractiveness from the model or the data?

Equipped with the metrics for abtractiveness above, we want to further understand how abtractive the generated summaries are, and whether the amount of abtractiveness is a result of the training data or the model. Therefore, we compute abtractiveness scores for both the reference summaries and summaries generated from a diverse set of models on two datasets.

Datasets. We use the CNN/DailyMail (Hermann et al., 2015; Nallapati et al., 2016b) (CNN/DM) and the XSum (Narayan et al., 2018) datasets, which are both used for single-document news summarization tasks. CNN/DM consists of articles from the CNN and Daily Mail websites, where the summaries comprise highlights in bullet points. XSum consists of BBC articles, where the summaries comprise a single-sentence summary that is written as the opening introductory sentence for the article. XSum was released in particular to promote research on highly abtractive summarization systems. Appendix A provides statistics on CNN/DM and XSum datasets: they contain around 288k and 204k training examples, respectively; CNN/DM includes longer documents and summaries on average.

Models. Most neural abtractive summarization models are based on sequence-to-sequence models. They differ in how summarization-specific operations such as copying/extraction are instantiated. We consider 5 prominent models and sum-

Systems	Extractor	Encoder	Decoder
PGC	–	LSTM	LSTM+copy
FASTRL	sentences	LSTM	LSTM+copy
BOTTOMUP	words	LSTM	LSTM+copy
TCONV	–	CNN+topic	CNN
BERTSUM	–	BERT-based	Transformer

Table 2: Comparison of summarization systems in terms of model architecture.

marize their characteristics in Table 2.² Details of each model can be found in Appendix B. **PGC** (See et al., 2017) uses the copy mechanism during decoding to allow extraction. **FASTRL** (Chen and Bansal, 2018) and **BOTTOMUP** (Gehrmann et al., 2018) decouple extraction and abtractive generation by learning to select sentences and words respectively in the first step; this model has been shown to generate more abtractive summaries compared to PGC. **TCONV** (Narayan et al., 2018) is initially designed for XSum, thus it does not include any explicit copying/extraction components and focuses on long text representation using convolutional neural networks. **BERTSUM** (Liu and Lapata, 2019) consists of a BERT-based encoder and a 6-layer Transformer decoder. It incorporates extraction implicitly by first fine-tuning the encoder on the extractive summarization task.³

Results. Our goal is to understand the level of abtractiveness of summaries generated by different models, and the influence on abtractiveness from the training data. Therefore, we analyzed summaries generated by the above models on CNN/DM and XSum. We computed the metrics described in Section 2.1 for both the generated summaries and the reference summaries on the test sets. The results are shown in Table 3.

First, *CNN/DM is more extractive than XSum*. Extraction scores of the reference summaries in CNN/DM shows that almost half of the sentences are formed by deleting words in one of the source sentences. This shows that sentence *compression* (Knight and Marcu, 2002) is the main technique used for this dataset. In contrast, none of the summary sentences in XSum are formed by copying from a single source sentence. They are generated mostly by paraphrasing the input content, indicated by the large fraction of novel n -grams.

²We use state-of-the-art models proposed for each dataset at the time of writing.

³We use the BERTSUMEXTABS variation.

Dataset	Model	Extraction			Perfect fusion		Novel n -grams		
		Sentence	Span	Word	$k = 2$	$k \geq 2$	$n = 1$	$n = 2$	$n = 3$
CNN/DM	Ref	1.39	2.14	9.27	12.92	14.87	12.40	51.03	71.22
	PGC	35.45	34.18	15.45	10.90	1.61	0.62	3.33	7.42
	FASTRL	8.94	40.06	39.64	4.22	0.84	0.82	10.89	20.74
	BOTTOMUP	7.65	17.98	36.75	21.86	6.77	0.86	11.44	22.40
	BERTSUM	—	13.73	53.40	16.18	4.39	5.23	14.55	23.09
XSum	Ref	—	—	—	0.87	0.77	39.20	84.98	96.05
	PGC	—	—	—	0.41	3.47	30.08	74.27	91.27
	TCONV	—	—	—	0.35	2.31	34.07	80.62	95.12
	BERTSUM	—	—	—	0.33	3.15	28.93	75.85	91.41

Table 3: Abtractiveness measures of the models on CNN/DM and XSum datasets. The numbers for Extraction and Perfect fusion indicate % of sentences generated with these strategies. Numbers for novel n -grams indicate % of n -grams that are present in the output sentence but is not present in the source.

Second, *training data has a larger influence on the abtractiveness of model outputs*. Similar to Zhang et al. (2018), we find that models trained on CNN/DM are near-extractive. However, the same models trained on XSum are significantly more abtractive. In fact, none of the models produced any sentence that copies words/phrases from a single source sentence, which is consistent with characteristics of the reference summaries in XSum. The content is more often rephrased in novel words/phrases. However, on both datasets, current models struggle to achieve the same level of abtractiveness as the reference summaries, indicating that additional inductive bias is needed to condense multiple sentences by rephrasing.

Third, *different models have different ways of doing extraction*. When trained on CNN/DM, PGC generates the majority of sentences by copying complete source sentences, whereas FASTRL, BOTTOMUP and BERTSUM do simple compression by deletion more often. In addition, BOTTOMUP does more fusion compared to PGC, FASTRL and BERTSUM.

2.3 Annotating Summary Faithfulness⁴

To understand faithfulness of current systems and its relation to abtractiveness, we crowd-sourced human annotations on the output of each model-dataset pair described in Section 2.2. Since a near-extractive sentence is very likely to be grammatical and faithful, we focus on more abtractive cases by excluding output sentences that are either an exact copy or a substring of one of the source sentences.

A key challenge to reliable human annotation is that the inter-annotator agreement on faithfulness is relatively low (Lebanoff et al., 2019). Our pi-

⁴We make our data and code available for reproducibility at: <https://github.com/esdurmus/summary-faithfulness>.

lot study shows that workers often do not agree on incoherent sentences, e.g. whether “*Chelsea beat Chelsea 5 – 3 in the Premier League on Saturday.*” is faithful or not. To standardize the annotation process, we design hierarchical questions to distinguish among failed generation that render a sentence meaningless, low-level grammatical errors that hardly affect semantic understanding, and faithfulness errors that convey incorrect (yet meaningful) information.

Figure 1 shows the decision tree of our human annotation steps. We first evaluate the grammaticality of generated sentences (independent from the source document). We show annotators a summary sentence and ask them to choose whether the given sentence is *meaningful* or *nonsensical* to determine if the given sentence is structurally and semantically sound. If the annotator can make sense of the sentence, we then ask whether it is *grammatical* or *has minor grammaticality problems* which a person can easily correct.

Next, for sentences labeled as meaningful in the first step, we ask workers whether they are faithful to the provided source document. In case the worker labels a sentence as unfaithful, we conduct a simple error analysis by asking them to indicate if the sentence contains information that is absent from or conflicting with the source document, which corresponds to *hallucination* and *contradiction* errors, respectively. More details about the annotation schema and guidelines are included in the Appendix C. Next, we describe our human evaluation results.

2.3.1 Human Annotation Results

For each dataset-model pair described in Section 2.2, we randomly sampled 1000 sentence-source pairs eliminating output sentences that are either an exact copy or substring of a source sen-




Figure 1: The decision diagram of our human annotation process. Decision nodes are rectangular and outcome nodes are circular. We show the annotation path of two summary sentences, S1 (green arrows) and S2 (red arrows). S2 is annotated as *nonsensical* thus is not considered for faithfulness. S1 is annotated as unfaithful due to hallucinated content.

Dataset	Model	Grammaticality			Faithfulness		
		Score	Agreement	Abstractiveness	Score	Agreement	Abstractiveness
CNN/DM	PGC	93.34	94.04	10.05	70.05	77.28	13.35
	FASTRL	83.06	88.05	44.46	68.27	77.45	49.74
	BOTTOMUP	85.83	89.19	29.62	64.17	76.04	42.36
	BERTSUM	97.53	97.65	29.44	95.03	95.14	39.16
XSum	PGC	65.85	81.03	91.10	40.33	71.63	97.06
	TCONV	70.85	85.03	94.94	38.96	69.90	98.81
	BERTSUM	90.44	91.80	91.50	60.54	70.00	97.60

Table 4: Grammaticality and faithfulness results of human annotations. **Score** is computed by taking the percentage of annotators that selected “meaningful” and “faithful” for grammaticality and faithfulness annotation tasks, respectively, and then averaging these values across all the examples for the given annotation task. **Agreement** is computed by taking the percentage of the workers that annotate the majority class for the given example. **Abstractiveness** is measured by the percentage of novel trigrams in a given sentence.

tence. We collected grammaticality annotations for these sentences from 5 annotators. We consider a sentence meaningful if at least 4 out of 5 annotators label it as meaningful in the first stage. We sampled 200 meaningful sentences randomly to collect annotations for faithfulness. Table 4 shows the results of the grammaticality and faithfulness human evaluations.

Grammaticality. Overall, outputs from all models are scored high on grammaticality with high inter-annotator agreement. However, on more *abstractive* summaries (i.e. when trained on XSum), the grammaticality scores drop significantly. One exception is BERTSUM, which maintains good performance on XSum and achieves the highest grammaticality score on both datasets.⁵

Faithfulness. Near-extractive summaries generated from models trained on CNN/DM have significantly higher faithfulness scores than highly

abstractive summaries from models trained on XSum. We find that PGC and TCONV has faithfulness errors in more than half of the sentences they generate when trained on XSum. Although BERTSUM generates fewer unfaithful sentences, it still suffers from performance drop on XSum. Interestingly, human agreement on faithfulness is also lower for abstractive summaries from XSum. This suggests that faithfulness errors are harder to catch for humans as well in more abstractive settings. We further observe conflicting information is more common among models trained on CNN/DM while hallucination is more common among models trained on XSum. Table 5 shows examples of meaningful but unfaithful sentences.

3 FEQA: Faithfulness Evaluation with Question Answering

Our analysis above shows that the number of unfaithful sentences increases significantly as more abstractive summaries are generated. Thus the key challenge to faithfulness evaluation is to verify highly abstractive sentences against the source document, where surface similarity match-

⁵Majority of the sentences (> 70%) identified as “meaningful” are annotated as “perfectly grammatical” for each model-dataset pair.

Source	Output Sentence	Domain	Category
...However, Winger Ross Wallace (knee) and right-back Steven Reid (calf) could return for the Barclays premier league contest...	Dean Marney and Steven Reid could return for the Barclays Premier League match.	CNN/DM	IC
...Odom also played for the US in the 2004 Athens Olympics, winning the bronze medal. His condition is unknown but well-wishers tweeted their support following the news...	NBA basketball player Odom has been found dead in a helicopter crash in the US state of Nevada.	XSum	H

Table 5: Examples of *meaningful* but *unfaithful* sentences. Category corresponds to the faithfulness *error* type for the output sentence. **IC**: Incorrect Concatenation, **H**: Hallucination. More examples are provided in Table 11.

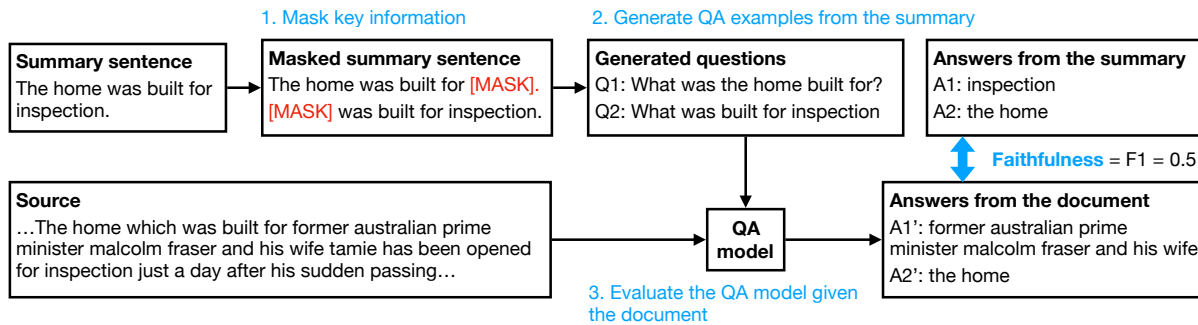


Figure 2: Overview of FEQA. Given a summary sentence and its corresponding source document, we first mask important text spans (e.g. noun phrases, entities) in the summary. Then, we consider each span as the “gold” answer and generate its corresponding question using a learned model. Lastly, a QA model finds answers to these questions in the documents; its performance (e.g. F1 score) against the “gold” answers from the summary is taken as the faithfulness score.

ing would fail. If we have a good semantic representation of the sentence abstracting away its surface form (e.g. a list of facts about *who* did *what* to *whom*), we can simply compare the sentence representation to the document representation (e.g. check whether the fact list from the summary is a subset of the list from the document). Ideally, the representation should be domain-general and interpretable for easy error analysis.

Motivated by the fast progress in reading comprehension (Chen, 2018; Gao et al., 2018) we propose to use QA pairs as a generic meaning representation of sentences for faithfulness evaluation. Given a summary sentence, we produce a list of questions asking about key information in the sentence and their corresponding answers. To verify this information against the source, we use a QA model to predict answers from the document. The questions and the QA model thus extract comparable information from two pieces of text. More matched answers from the document implies a more faithful summary since the information addressing these questions are consistent between the summary and the source document. Figure 2 shows the workflow of FEQA.

Question generation. Prior work (Eyal et al., 2019; Scialom et al., 2019) uses cloze tests as questions by masking entities. To go beyond cloze-style QA and leverage more recent extractive (Rajpurkar et al., 2016) or even generative (Alec et al., 2019) QA models, we generate natural language questions from the summary sentence automatically. Specifically, we mask important text spans in a sentence, including noun phrases extracted by a constituency parser (Kitaev and Klein, 2018) and named entities extracted by the Stanford CoreNLP NER model (Finkel et al., 2005; Manning et al., 2014). We consider each span as the gold answer and generate its corresponding question by fine-tuning a pretrained BART language model (Lewis et al., 2019). To train the question generator, we adapt the QA2D dataset Demszky et al. (2018). The input is a declarative sentence with masked answers and the output is a question. A training example might look like:

Input: Sally was born in <m> 1958 </m>
Output: When was Sally born ?

Since the transformation from declarative sen-

tences to questions is almost rule-based without much paraphrasing, we expect the model to generalize to various domains.

Answer verification. Given the QA pairs generated from a summary sentence, we run off-the-shelf QA models to get answers to these questions from the source document. We then measure the average F1 score against the “gold” answers from the summary, which is our faithfulness score for the given sentence. This step does not have any constraint on the QA model. We experiment with the pretrained BERT-base model (Devlin et al., 2019) fine-tuned on SQuAD-1.1 (Rajpurkar et al., 2016) and SQuAD-2.0 (Rajpurkar et al., 2018). Note that in the case of SQuAD-2.0, the model may be able to hypothesize that a question is unanswerable. This case is equivalent to getting an answer incorrect (i.e. unfaithful).

4 Experiments

We aim to understand to what extent the proposed QA-based metric and existing metrics capture faithfulness of a summary. Given pairs of documents and summary sentences *without reference summaries*, we measure correlations between human-annotated faithfulness scores (Section 2.3) and scores computed using each metric described below.

4.1 Automated Metrics for Faithfulness

Word overlap-based metrics. A straightforward metric for faithfulness is the word overlap between the summary sentence and the document. We compute ROUGE (R), BLEU (B),⁶ between the output sentence and each of the source sentences (i.e. taking the source sentence as the reference). We then take the average scores and maximum score across all the source sentences. Since according to our analysis taking the average score consistently has higher correlation, we report only the correlation for the average.

Embedding-based metrics. Word embeddings extend word overlap-based metrics beyond exact match. Recently, BERTScore (Zhang et al., 2019b) was proposed to compute the similarity between two sentences using contextual word embeddings from BERT. It has higher correlation

⁶We report only BLUE-4 since it performed the best for CNN/DM and no variation of BLEU has significant correlation with faithfulness for XSum.

with human judgements on image captioning and machine translation than word overlap based metrics. We compute BERTScore (BERTSc) between each source sentence and the summary sentence.⁷ To get the final score, we experiment with both the average and the maximum scores computed from each source sentence and the summary sentence. We report results using the maximum score since it has better performance.

Model-based metrics. In addition to QA, recent work has used relation extraction and textual entailment models for faithfulness evaluation (Falke et al., 2019a; Goodrich et al., 2019). For the relation extraction metric (RE), we compute the precision for the relation triplets extracted from the summary sentence and the source document using an off-the-shelf model (Angeli et al., 2015) from Stanford Open IE. For the textual entailment metric (ENT), we measure whether the summary sentence is entailed by the source using the pretrained ESIM model (Chen et al., 2017) from AllenNLP (Gardner et al., 2018).

4.2 Results

Metric Comparison. We first compute scores for each metric on document and output sentence pairs on both CNN/DM and XSum datasets (748 and 286 pairs respectively). We then compute Pearson and Spearman correlation coefficients between scores given by each metric and human-annotated scores. Table 7 includes correlation coefficients for the examples from CNN/DM and XSum, respectively. We observe that for both CNN/DM and XSum, the score of QA-based evaluation has a higher correlation with faithfulness than other metrics. Although word-overlap based metrics are correlated with the faithfulness in more extractive settings (i.e. for CNN/DM), these metrics have no correlation with faithfulness in more abstractive settings (i.e. for XSum). We further notice that all the metrics have significantly lower correlation with human scores for XSum, suggesting that evaluating faithfulness is more difficult in highly abstractive settings; deeper understanding of the source and the summary sentence is necessary here.

Consistent with the findings of Falke et al. (2019b), the entailment metric does not have a significant correlation with faithfulness in most cases. These models fail to distinguish entailed (faithful)

⁷https://github.com/Tiiiger/bert_score.

Source Sentence	Output Sentence	Metric	Score
Health Inspectorate Wales said Wrexham Maelor Hospital staff were under “considerable pressure” for long periods as ambulances waited outside.	A hospital ward in Wrexham has been rated “inadequate” by inspectors after inspectors found patients at risk of harm.	Entailment	72.83%
The Black Poplar is one of the rarest native trees in the UK, with only 2,500 thought to be left.	Northern Ireland’s first trees are among those recognised in the Welsh Architecture Trust’s list of the year’s best trees.	BertScore	83.06%

Table 6: Unfaithful examples missed by Entailment and BertScore. **Score:** Output score of the metrics; higher score indicates stronger entailment and similarity respectively.

Metric	CNN/DM		XSum	
	P	S	P	S
Word overlap-based				
R-1	12.02**	15.86**	-2.57	0.07
R-2	13.25**	15.99**	-5.78	-8.47
R-L	12.58**	16.49**	-6.37	-9.68
B-4	12.09**	11.68**	-6.76	-10.02
Embedding-based				
BERTSc	11.07*	10.70*	10.06	10.69
Model-based				
RE	8.58*	5.52	1.62	2.32
ENT	2.80	3.65	-5.62	-3.85
FEQA	32.01**	28.23**	26.31**	21.34**

Table 7: Pearson (P) and Spearman (S) correlation between human-annotated *faithfulness scores* and the metric scores. *,** indicates p -values $< 0.05, < 0.001$, respectively. FEQA has the highest correlation with human scores for both CNN/DM and XSum.

and non-entailed (unfaithful) summary sentences when both overlap largely with the source document, because models trained on current entailment datasets may rely on simple heuristics such as lexical overlap (McCoy et al., 2019). Similarly, BERTScore tends to give higher scores when there are overlapping concepts between the sentences even though the content is not the same. See Table 6 for examples.

Content selection and faithfulness. Current evaluation metrics for summarization produce a single measure of the overall quality of the summary. Typically, the output summary is compared against the reference summary in terms of n -gram overlap. These metrics mainly evaluate *content selection*, i.e. whether the content of the output is similar to the content of the reference. In contrast, to evaluate faithfulness, we compare the output summary against the source document. One natural question that follows is whether high content matching sufficient for faithfulness. We compute the correlation coefficients between human-annotated faithfulness scores and ROUGE scores computed from the reference and the output sentence. As shown in Table 8, while there is a weak

Metric	CNN/DM		XSum	
	P	S	P	S
ROUGE-1	15.31**	14.92**	5.44	5.79
ROUGE-2	15.10**	16.39**	8.25	6.79
ROUGE-L	13.33**	13.35**	4.61	3.97

Table 8: Pearson (P) and Spearman (S) correlation between human-annotated faithfulness scores and ROUGE scores of content selection (computed between the reference and the output sentence). High content selection scores (typical ROUGE score for summarization) do not necessarily imply faithfulness of the summary.

correlation between ROUGE scores of content selection and faithfulness on CNN/DM, the correlation is significantly lower than ROUGE scores of faithfulness (i.e. computed between the *source* and the output sentence). For XSum, there is no significant correlation between the content selection metrics and faithfulness. We provide unfaithful examples with high content selection scores in Appendix D.3. This suggests that content selection and faithfulness should be measured separately as opposed to using a unified score.

Analysis and limitations of QA-based evaluation. Table 9 shows examples for a faithful and an unfaithful output sentence and the corresponding QA pairs. Note that the QA system is able to capture common errors such as conflicting information in the output sentence. To measure the reliability of FEQA, we further perform a manual error analysis using 100 randomly sampled QA pairs. We observe that around 94% of generated questions are mostly grammatical and correct given the mask. For 78% of the questions, the QA system has the correct behaviour: it answers the question correctly if the sentence is faithful to the article, otherwise it produces “unanswerable” or an incorrect answer. Majority of the errors of the QA system are because it either didn’t detect unanswerable questions or produces “unanswerable” when there exists an answer (14%). More-

Source	Output Sentence	Question	OA	SA
...However, Winger Ross Wallace (knee) and right-back Steven Reid (calf) could return for the Barclays premier league contest...	Dean Marney and Steven Reid could return for the Barclays Premier League match.	Who and Steven Reid could return for the premier league match?	Dean Marney	Ross Wallace
...Miss Bruck, 22, from maybe has not been seen since the early hours of October 26, 2014. She has not been seen for six months...	Miss Bruck, 22, from maybe has not been seen for six months .	How long has Miss Bruck, 22 from not been seen for?	six months	six months

Table 9: Examples detection results from FEQA. **OA**:Output Answer, **SA**:Source Answer. The output sentence in the first example is *unfaithful*, whereas the one for the second example is *faithful*. Bold text indicates the span that was masked to generate the question.

over, when the article is long, QA system tends to make more mistakes. Especially for more abstractive settings, F1-score penalizes the correct answers when the answer from the article does not exactly match with the gold answer (i.e. “Donald Trump” vs. “the President of the United States Donald Trump”) (16%).

5 Related Work

Problems in current neural generation models. Since the beginning of neural text generation, problems with repetition and generic responses have received lots of attention (Sordoni et al., 2015; Li et al., 2016; Holtzman et al., 2019). Recently, more work has focused on semantic errors in model outputs, such as adequacy in machine translation (Tu et al., 2017), faithfulness in summarization (Cao et al., 2018), and consistency in dialogue (Li et al., 2019). Our analysis on the abstractiveness-faithfulness tradeoff reveals additional limitation of current models, and suggests that we need new inductive bias on how to summarize beyond copying.

QA as a proxy. Question answering is a broad format that subsumes many tasks (Gardner et al., 2019). To the best of our knowledge, Mani et al. (1999) first use QA as an extrinsic evaluation for summarization: A good summary should answer key questions a reader might have about an article. Later, QA is incorporated in human evaluation where one person writes questions and another person answers them based on the summary (Clarke and Lapata, 2010; Liu and Lapata, 2019). The closest to our work are recent efforts in automating this protocol, including rule-based approaches (Chen et al., 2018) and cloze-test QA (Eyal et al., 2019; Scialom et al., 2019). Our work is the first to apply automated question generation. While we focus on faithfulness, our QA-based metric is applicable to semantic comparison

between any two pieces of text.

Automated evaluation for NLG. Automated NLG evaluation is challenging as it often requires deep understanding of the text. Although metrics based on word overlap with the reference text are commonly used, it is widely known that they do not correlate well with human judgments (Novikova et al., 2017; Liu et al., 2016). Recently, more work has focused on model-based evaluation using discriminators (Lowe et al., 2017; Hashimoto et al., 2019), entailment models (Falke et al., 2019a), information extraction (Wiseman et al., 2017; Goodrich et al., 2019), and question answering (Chen et al., 2018; Eyal et al., 2019).

6 Conclusion

We investigate the faithfulness problem in neural abstractive summarization and propose a QA-based metric for evaluating summary faithfulness. We show that current models suffer from an inherent trade-off between abstractiveness and faithfulness. They are good at copying important source content, but tend to concatenate unrelated spans and hallucinate details when generating more abstractive sentences. A new inductive bias or additional supervision is needed for learning reliable models. While our QA-based metric correlates better with human judgment and is useful for model development, it is limited by the quality of the QA model. The final evaluation should still rely on human annotation or human-in-the-loop methods (Chaganty et al., 2018).

Acknowledgement

We would like to thank Faisal Ladhak, the Lex and Comprehend groups at Amazon Web Services AI, and the anonymous reviewers for their feedback on this work.

References

- R. Alec, W. Jeff, C. Rewon, L. David, A. Dario, and S. Ilya. 2019. Language models are unsupervised multitask learners.
- Gabor Angeli, Melvin Jose Johnson Premkumar, and Christopher D. Manning. 2015. [Leveraging linguistic structure for open domain information extraction](#). 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 344–354, Beijing, China. Association for Computational Linguistics.
- Z. Cao, F. Wei, W. Li, and S. Li. 2018. Faithful to the original: Fact aware neural abstractive summarization. In *Association for the Advancement of Artificial Intelligence (AAAI)*.
- A. Chaganty, S. Mussmann, and P. Liang. 2018. The price of debiasing automatic metrics in natural language evaluation. In *Association for Computational Linguistics (ACL)*.
- Danqi Chen. 2018. *Neural Reading Comprehension and Beyond*. Ph.D. thesis, Stanford University.
- P. Chen, F. Wu, T. Wang, and W. Ding. 2018. A semantic QA-based approach for text summarization evaluation. In *Association for the Advancement of Artificial Intelligence (AAAI)*.
- Qian Chen, Xiaodan Zhu, Zhen-Hua Ling, Si Wei, Hui Jiang, and Diana Inkpen. 2017. [Enhanced LSTM for natural language inference](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1657–1668, Vancouver, Canada. Association for Computational Linguistics.
- Yen-Chun Chen and Mohit Bansal. 2018. Fast abstractive summarization with reinforce-selected sentence rewriting. In *ACL*.
- J. Clarke and M. Lapata. 2010. Discourse constraints for document compression. *Computational Linguistics*, 36.
- Dorottya Demszky, Kelvin Guu, and Percy Liang. 2018. Transforming question answering datasets into natural language inference datasets. *ArXiv*, abs/1809.02922.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*.
- M. Eyal, T. Baumel, and M. Elhadad. 2019. Question answering as an automatic evaluation metric for news article summarization. In *Human Language Technology and North American Association for Computational Linguistics (HLT/NAACL)*.
- T. Falke, L. F. R. Ribeiro, P. A. Utama, I. Dagan, and I. Gurevych. 2019a. Ranking generated summaries by correctness: An interesting but challenging application for natural language inference. In *Association for Computational Linguistics (ACL)*.
- Tobias Falke, Leonardo F. R. Ribeiro, Prasetya Ajie Utama, Ido Dagan, and Iryna Gurevych. 2019b. [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*, pages 2214–2220, Florence, Italy. Association for Computational Linguistics.
- Jenny Rose Finkel, Trond Grenager, and Christopher Manning. 2005. [Incorporating non-local information into information extraction systems by gibbs sampling](#). In *Proceedings of the 43rd Annual Meeting on Association for Computational Linguistics, ACL '05*, pages 363–370, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Jianfeng Gao, Michel Galley, and Lihong Li. 2018. [Neural approaches to conversational AI](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics: Tutorial Abstracts*, pages 2–7, Melbourne, Australia. Association for Computational Linguistics.
- M. Gardner, J. Berant, H. Hajishirzi, A. Talmor, and S. Min. 2019. Question answering is a format; when is it useful? *arXiv preprint arXiv:1909.11291*.
- 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.
- Sebastian Gehrmann, Yuntian Deng, and Alexander Rush. 2018. [Bottom-up abstractive summarization](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4098–4109, Brussels, Belgium. Association for Computational Linguistics.
- B. Goodrich, V. Rao, P. J. Liu, and M. Saleh. 2019. Assessing the factual accuracy of generated text. In *International Conference on Knowledge Discovery and Data Mining (KDD)*.
- M. Grusky, M. Naaman, , and Y. Artzi. 2018. Newsroom: A dataset of 1.3 million summaries with diverse extractive strategies. In *North American Association for Computational Linguistics (NAACL)*.
- T. Hashimoto, H. Zhang, and P. Liang. 2019. Unifying human and statistical evaluation for natural language generation. In *North American Association for Computational Linguistics (NAACL)*.

- Karl Moritz Hermann, Tomáš Kočiský, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. [Teaching machines to read and comprehend](#). In *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 1*, NIPS'15, pages 1693–1701, Cambridge, MA, USA. MIT Press.
- A. Holtzman, J. Buys, M. Forbes, and Y. Choi. 2019. The curious case of neural text degeneration. *arXiv preprint arXiv:1904.09751*.
- Nikita Kitaev and Dan Klein. 2018. [Constituency parsing with a self-attentive encoder](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2676–2686, Melbourne, Australia. Association for Computational Linguistics.
- K. Knight and D. Marcu. 2002. Summarization beyond sentence extraction: A probabilistic approach to sentence compression. *Artificial Intelligence*, 139:91–107.
- W. Kryściński, N. S. Keskar, B. McCann, C. Xiong, and R. Socher. 2019. Neural text summarization: A critical evaluation. In *Empirical Methods in Natural Language Processing (EMNLP)*.
- L. Lebanoff, J. Muchovej, F. Dernoncourt, D. S. Kim, S. Kim, W. Chang, and F. Liu. 2019. Analyzing sentence fusion in abstractive summarization. *arXiv preprint arXiv:1910.00203*.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. 2019. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *ArXiv*, abs/1910.13461.
- J. Li, M. Galley, C. Brockett, J. Gao, and W. B. Dolan. 2016. A diversity-promoting objective function for neural conversation models. In *Human Language Technology and North American Association for Computational Linguistics (HLT/NAACL)*, pages 110–119.
- M. Li, S. Roller, I. Kulikov, S. Welleck, Y. Boureau, K. Cho, and J. Weston. 2019. Don't say that! making inconsistent dialogue unlikely with unlikelihood training. *arXiv preprint arXiv:1911.03860*.
- C. Liu, R. Lowe, I. V. Serban, M. Noseworthy, L. Charlin, and J. Pineau. 2016. How NOT to evaluate your dialogue system: An empirical study of unsupervised evaluation metrics for dialogue response generation. In *Empirical Methods in Natural Language Processing (EMNLP)*.
- Y. Liu and M. Lapata. 2019. Text summarization with pretrained encoders. In *Empirical Methods in Natural Language Processing (EMNLP)*.
- R. Lowe, M. Noseworthy, I. V. Serban, N. Angelard-Gontier, Y. Bengio, and J. Pineau. 2017. Towards an automatic turing test: Learning to evaluate dialogue responses. In *Association for Computational Linguistics (ACL)*.
- I. Mani, G. Klein, L. Hirschman, T. Firmin, D. House, and B. Sundheim. 1999. The TIPSTER SUMMAC text summarization evaluation. In *European Association for Computational Linguistics (EACL)*.
- Christopher Manning, Mihai Surdeanu, John Bauer, Jenny Finkel, Steven Bethard, and David McClosky. 2014. [The Stanford CoreNLP natural language processing toolkit](#). In *Proceedings of 52nd Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 55–60, Baltimore, Maryland. Association for Computational Linguistics.
- R. T. McCoy, E. Pavlick, and T. Linzen. 2019. Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference. *arXiv preprint arXiv:1902.01007*.
- Ramesh Nallapati, Feifei Zhai, and Bowen Zhou. 2016a. Summarunner: A recurrent neural network based sequence model for extractive summarization of documents. In *AAAI*.
- Ramesh Nallapati, Bowen Zhou, Cicero dos Santos, Çağlar Gülçehre, and Bing Xiang. 2016b. [Abstractive text summarization using sequence-to-sequence RNNs and beyond](#). In *Proceedings of The 20th SIGNLL Conference on Computational Natural Language Learning*, pages 280–290, Berlin, Germany. Association for Computational Linguistics.
- Shashi Narayan, Shay B. Cohen, and Mirella Lapata. 2018. Don't give me the details, just the summary! Topic-aware convolutional neural networks for extreme summarization. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, Brussels, Belgium.
- J. Novikova, O. Dušek, A. C. Curry, and V. Rieser. 2017. Why we need new evaluation metrics for NLG. In *Empirical Methods in Natural Language Processing (EMNLP)*.
- Pranav Rajpurkar, Robin Jia, and Percy Liang. 2018. [Know what you don't know: Unanswerable questions for SQuAD](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 784–789, Melbourne, Australia. Association for Computational Linguistics.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. [SQuAD: 100,000+ questions for machine comprehension of text](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392, Austin, Texas. Association for Computational Linguistics.

- T. Scialom, S. Lamprier, B. Piwowarski, and J. Staiano. 2019. Answers unite! unsupervised metrics for reinforced summarization models. In *Empirical Methods in Natural Language Processing (EMNLP)*.
- Abigail See, Peter J. Liu, and Christopher D. Manning. 2017. Get to the point: Summarization with pointer-generator networks. In *ACL*.
- A. Sordoni, M. Galley, M. Auli, C. Brockett, Y. Ji, M. Mitchell, J. Nie, J. Gao, and B. Dolan. 2015. A neural network approach to context-sensitive generation of conversational responses. In *North American Association for Computational Linguistics (NAACL)*.
- Z. Tu, Y. Liu, L. Shang, X. Liu, and H. Li. 2017. Neural machine translation with reconstruction. In *Association for the Advancement of Artificial Intelligence (AAAI)*.
- Zhaopeng Tu, Zhengdong Lu, Yang Liu, Xiaohua Liu, and Hang Li. 2016. [Modeling coverage for neural machine translation](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 76–85, Berlin, Germany. Association for Computational Linguistics.
- Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. 2015. [Pointer networks](#). In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems 28*, pages 2692–2700. Curran Associates, Inc.
- S. Wiseman, S. M. Shieber, and A. M. Rush. 2017. Challenges in data-to-document generation. In *Empirical Methods in Natural Language Processing (EMNLP)*.
- F. Zhang, J. Yao, and R. Yan. 2018. On the abstractiveness of neural document summarization. In *Empirical Methods in Natural Language Processing (EMNLP)*.
- T. Zhang, V. Kishore, F. Wu, K. Q. Weinberger, and Y. Artzi. 2019a. BERTSCORE: Evaluating text generation with BERT. *arXiv preprint arXiv:1904.09675*.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2019b. Bertscore: Evaluating text generation with bert. *ArXiv*, abs/1904.09675.

A Summarization Datasets

All of our experiments are run on the CNN/DM and XSum datasets. We show basic statistics of the two datasets in Table 10.

	CNN/DM	XSum
# Training Documents	287,227	204,045
# Validation Documents	13,368	11,332
# Test Documents	11,490	11,334
Document: avg # of tokens	781.00	431.07
Document: avg # of sents.	40.00	33.00
Summary: avg # tokens	56.00	23.26
Summary: avg # of sents.	3.75	1.00

Table 10: Statistics of CNN/DM and XSum datasets.

B Summarization Models

The characteristics of each model used in our experiments are detailed below.

Pointer Generator Model with Coverage (PGC) (See et al., 2017) uses the copy mechanism (Vinyals et al., 2015) to allow copying words from the source. The adapted *coverage* mechanism (Tu et al., 2016) is incorporated to alleviate repetition by keeping track of source words that have been summarized. This copy mechanism is widely adopted by subsequent models.

Fast Abstractive Summarization with Reinforce (FASTRL) (Chen and Bansal, 2018) first uses an extractor agent to select salient sentences from the document, then condenses the extracted sentences using the Pointer-Generator summarizer.

Bottom-up Summarization Model (BOTTOMUP) (Gehrmann et al., 2018) first selects words from the source document that are likely to appear in the summary, then generates using the Pointer-Generator model, where the copying mechanism is constrained to the previously selected words. It improves upon PGC by explicitly learning the selector to avoid copying long text spans.

Topic-aware Convolutional Sequence-to-Sequence model (TCONVS2S) (Narayan et al., 2018) is a convolutional neural network-based model conditioned on the topics of the article. It is shown to be effective in capturing long-range dependencies in the documents.

BERT-based model (BERTSUM) (Liu and Lapata, 2019) is a two-stage fine-tuning approach

where the BERT-based encoder is first fine-tuned on the extractive summarization task and then on the abstractive summarization task with the decoder (denoted as BERTSUMEXTABS in the original paper).

C Details of Human Annotations

C.1 Grammaticality Annotation Guidelines

For grammaticality annotation, we present only the output sentence to the workers. We collect annotations from 5 workers for both of the tasks. For this task, given the output sentence, we provide workers the following guidelines:

1. First select whether the given sentence is “Nonsensical” or “Makes sense”.
2. If the given text is not a complete sentence, mark it as “Nonsensical”.
3. If you can understand the meaning of the sentence, despite grammaticality errors, and you are able to make sense of it, select “Makes sense”.
4. If you did not select “Nonsensical”, evaluate whether the sentence is “Grammatical” or “Has Minor Grammaticality Issues”.

C.2 Faithfulness Annotation Guidelines

We present workers both the source and the output sentence and provide the following guidelines:

1. Read the sentence and the source fully.
2. If the information conveyed by the sentence is not expressed in the source, select “unfaithful”.
3. Avoid using general knowledge, and check if the sentence is consistent with the source.
4. If you select “unfaithful”, for the second part, select whether the information expressed by the sentence is *not contained* in the source or *conflicting* with the source.

D Additional Analysis

D.1 Examples for *nonsensical* sentences

- Sandals, £34, office.co.uk, luluguinness.com. (generated by PGC for CNN/DM)

Source	Output Sentence	Category
...Although her due date has not officially been confirmed, the duchess of Cambridge told wellwishers at a charity event last month: I am due mid-April, to the end of April...	The duchess of Cambridge told wellwishers at a charity event last month: “The duke’s intention is to be at the commemorations” .	IC
...Carragher spoke to a local TV starton during his time in Girona. Carragher posted a picture on his Instagram account of the opening ceremony...	Carragher posted a picture on his son play in the famous youth tournament.	IC
A body was found by a member of the public on private land near Leighton, about 10 miles (16.09km) away from the centre of Shrewsbury, on Monday. Mr Bebbington’s family has been informed, West Mercia Police confirmed.	The death of a man whose body was found in a river in Cumbria has been identified as murder.	H
The incident happened near Dr. Gray’s hospital shortly after 10:00. The man was taken to the hospital with what police said were serious but not life-threatening injuries. The a96 was closed in the area for several hours, but it has since reopened.	A man has been taken to hospital after he was hit by a lorry in Dumfries.	H

Table 11: Examples of *meaningful* but *unfaithful* sentences. Category corresponds to the category of unfaithfulness error for the output sentence. **IC**: Incorrect Concatenation, **H**: Hallucination.

Reference	Output Sentence
... University of Nebraska researcher has revealed why stress is bad for you. Limited periods of stress are good, as they release cortisol...	University of Nebraska researcher has revealed why stress is bad for you, stimulating your body to produce an important hormone called cortisol.
...Indian air force and Nepalese army medical team launch rescue mission to bring injured people to hospitals in Kathmandu. Forshani Tamang’s family carried her for four hours to reach help after she was wounded when their home was destroyed...	Indian air crew and Nepalese army medical team were killed in Nepal’s Sindhupalchok quake.

Table 12: Examples of unfaithful sentence with high content overlap (computed by ROUGE-L) with the reference.

- He says easter triduum is a progression , although the word itself – triduum. (generated by FASTRL for CNN/DM)
- Chelsea beat Chelsea 5 – 3 in the Premier League on Saturday. (generated by FASTRL for CNN/DM)
- 12 years a slave actress Lupita Woodley and oily vegetables. (generated by BOTTOMUP for CNN/DM)
- A judge in Japan has ordered a judge to order a woman who has absconded from Japan to Japan. (generated by PGC for XSum)
- Stoke City moved up to third in the Premier League with victory over Stoke City at Stoke. (generated by TCONV for XSum)
- Johnny Depp’s management group is suing his management group over his “lavish lifestyle”. (generated by BERTSUM for XSum)

D.2 Examples for *meaningful but unfaithful sentences*

Table 11 includes examples that are annotated as *meaningful but unfaithful*. First three examples are picked from the models trained on CNN/DM, and last three are from the models trained on XSum. We observe that majority of sentences with faithfulness errors for CNN/DM dataset are generated by incorrect concatenation (IC). The models fuse two sentences from the source and generate a new sentence that is not consistent with the context of the source. Within this category, however, the models make a wide-range of mistakes such as copying the wrong entity, date, and quote.

For XSum, the faithfulness mistakes are mostly hallucinations. Models tend to hallucinate information (e.g. entities, events, date) that is not present in the source.

D.3 Examples for sentences with high content overlap with reference that are *unfaithful*

Although current summarization models are evaluated with respect to the content overlap between the reference and the output, these metrics do not necessarily provide any guarantees for the faithfulness of the output. Table 12 includes examples with similar content overlap scores as the faithful

examples but are unfaithful. We can see that although the output sentences include similar words and refer to similar topics, they include hallucinations and inaccurate information.

D.4 Limitations of the datasets

Since CNN/DM and XSum datasets are automatically crawled, we find that there is noise in the data. For example, source documents can include phrases such as “click here for the latest news”. We further observe that *reference* can carry information that is not in the source document since some of these one sentence highlights are written using additional world knowledge. Table 13 shows an example where the reference is unfaithful since it includes information that is not in the source (i.e. the fact that Ms. Wood’s first name is Leanne and she is Plaid Cymru leader.).

Source	Reference
<p>Ms Wood blamed the Conservatives in particular for claiming the SNP posed a threat to the future of the UK. She claimed "progressive" parties like hers were offering a "collaborative" alternative to "combative" politics. "This election presents an opportunity for harmonious co-existence between our nations," she said. Ms Wood's comments followed Conservative claims that Labour dependence on support from the SNP to form a government after the election on 7 May would threaten the break-up of the UK. Campaigning in south Wales on Monday, she said: "The parties advocating progressive, inclusive non-partisan cooperation in this election are not those who claim to cherish the political union above all others, but the national parties of Wales and Scotland. Along with the Greens in England, our parties have provided people across these islands with a collaborative alternative to the traditional combative Westminster politics.". Ms Wood added that she had received "hundreds" of supportive messages from people in England following the televised debates.</p>	<p>Plaid Cymru leader Leanne Wood has accused rival parties of "dangerous and divisive rhetoric" in a "desperate" attempt to win votes.</p>

Table 13: Example where reference includes information that **is not** in the source.