

CLIFF: Contrastive Learning for Improving Faithfulness and Factuality in Abstractive Summarization

Shuyang Cao and Lu Wang

Computer Science and Engineering

University of Michigan

Ann Arbor, MI

{caoshuy, wangluxy}@umich.edu

Abstract

We study generating abstractive summaries that are faithful and factually consistent with the given articles. A novel contrastive learning formulation is presented, which leverages both reference summaries, as positive training data, and automatically generated erroneous summaries, as negative training data, to train summarization systems that are better at distinguishing between them. We further design four types of strategies for creating negative samples, to resemble errors made commonly by two state-of-the-art models, BART and PEGASUS, found in our new human annotations of summary errors. Experiments on XSum and CNN/Daily Mail show that our contrastive learning framework is robust across datasets and models. It consistently produces more factual summaries than strong comparisons with post error correction, entailment-based reranking, and unlikelihood training, according to QA-based factuality evaluation. Human judges echo the observation and find that our model summaries correct more errors.

1 Introduction

Large pre-trained Transformers have yielded remarkable performance on abstractive summarization (Liu and Lapata, 2019; Lewis et al., 2020; Zhang et al., 2020a) with impeccable fluency, yet their summaries often contain factually inconsistent content (Maynez et al., 2020; Zhang et al., 2020b; Goyal and Durrett, 2020), even for state-of-the-art models. Three types of remedies have been proposed: running a separately learned error correction component (Dong et al., 2020), removing noisy training samples (Nan et al., 2021; Goyal and Durrett, 2021), and modifying the Transformer architecture (Huang et al., 2020; Zhu et al., 2021). Yet they either rely on heuristically created data for error handling, falling short of generalization, or require learning a large number of new parameters, and summary informativeness is often sacrificed.

XSum Article: The Fermanagh MLA Phil Flanagan tweeted after Tom Elliott appeared on a BBC radio programme in May 2014. . . . “I wonder if he will reveal how many people he harassed and shot as a member of the UDR.” . . .

Contrastive learning (our method): A Sinn Féin MLA has been ordered to apologise and pay compensation to a former member of the **Ulster Defence Regiment** (UDR).

Cross-entropy: A Sinn Féin MLA has agreed to pay compensation to a former **Ulster Unionist Party** (UDR) MP after he tweeted that he had harassed and shot people as a member of the party.

Entailment reranking: A Sinn Féin MLA has agreed to pay compensation to a former **Ulster Unionist Party** (UDR) councillor for a tweet he sent about him.

Unlikelihood: An MLA has been ordered to apologise and pay compensation to a former **loyalist** MP for a remark he made about him while serving in the **Ministry of Defence**.

Figure 1: Sample article and system summaries by different methods. Our contrastive learning model trained on low confidence system outputs correctly generates the **full name**. Comparisons using cross-entropy loss, beam reranking by entailment scores (Kryscinski et al., 2020), and unlikelihood objective (Welleck et al., 2020) over negative samples all produce **unfaithful content**.

Our goal is to train abstractive summarization systems that generate both faithful and informative summaries in an end-to-end fashion. We observe that, while the commonly used maximum likelihood training optimizes over references, there is no guarantee for the model to distinguish references from incorrect generations (Holtzman et al., 2020; Welleck et al., 2020). Therefore, potential solutions reside in designing new learning objectives that can effectively inform preferences of factual summaries over incorrect ones.

Concretely, we hypothesize that including factually inconsistent summaries (i.e., *negative samples*) for training, in addition to references (i.e., *positive samples*), let models become better at differentiating these two types of summaries. Although using negative samples has been effective at text representation learning, e.g., word2vec (Mikolov et al.,

2013) and BERT (Devlin et al., 2019), there exist two major challenges for it to succeed in concrete language tasks. First, a *suitable training objective* is critical to avoid performance degradation (Saunshi et al., 2019). Second, it is nontrivial to *construct “natural” samples* that mimic the diverse errors made by state-of-the-art systems that vary in words and syntax (Goyal and Durrett, 2021).

To address both challenges, we first propose a new framework, **CLIFF**, that uses contrastive learning for improving faithfulness and factuality of the generated summaries.¹ Contrastive learning (CL) has obtained impressive results on many visual processing tasks, such as image classification (Khosla et al., 2020; Chen et al., 2020) and synthesis (Park et al., 2020; Zhang et al., 2021b). Intuitively, CL improves representation learning by compacting positive samples while contrasting them with negative samples. Here, we design a task-specific CL formulation that teaches a summarizer to expand the margin between factually consistent summaries and their incorrect peers.

Moreover, we design four types of strategies with different variants to construct negative samples by *editing reference summaries* via rewriting entity-/relation-anchored text, and using *system generated summaries* that may contain unfaithful errors. Importantly, these strategies are inspired by our new annotation study on errors made by state-of-the-art summarizers—models fine-tuned from BART (Lewis et al., 2020) and PEGASUS (Zhang et al., 2020a)—on two benchmarks: XSum (Narayan et al., 2018) and CNN/DailyMail (CNN/DM) (Hermann et al., 2015).

We fine-tune pre-trained large models with our contrastive learning objective on XSum and CNN/DM. Results based on QuestEval (Scialom et al., 2021), a QA-based factuality metric of high correlation with human judgments, show that our models trained with different types of negative samples uniformly outperform strong comparisons, including using a summarizer with post error correction and reranking beams based on entailment scores to the source. Moreover, compared with unlikelihood training method that penalizes the same negative samples (Welleck et al., 2020), our summaries also obtain consistently better QuestEval scores. Human evaluation further confirms that our models consistently reduce both extrinsic and

intrinsic errors over baseline across datasets.

2 Related Work

Factuality Improvement and Evaluation. Neural abstractive summaries often contain unfaithful content with regard to the source (Falke et al., 2019). To improve summary factuality, three major types of approaches are proposed. First, a separate correction model is learned to fix errors made by the summarizers (Zhao et al., 2020; Chen et al., 2021), including replacing entities absent from the source (Dong et al., 2020) or revising all possible errors (Cao et al., 2020). The second type targets at modifying the sequence-to-sequence architecture to incorporate relation triplets (Cao et al., 2018), knowledge graphs (Zhu et al., 2021), and topics (Aralikatte et al., 2021) to inform the summarizers of article facts. Yet additional engineering efforts and model retraining are often needed. Finally, discarding noisy samples from model training has also been investigated (Nan et al., 2021; Goyal and Durrett, 2021), however, it often leads to degraded summary informativeness. In comparison, our contrastive learning framework allows the model to be end-to-end trained and does not require model modification, thus providing a general solution for learning summarization systems.

Alongside improving factuality, we have also witnessed growing interests in automated factuality evaluation, since popular word-matching-based metrics, e.g., ROUGE, correlate poorly with human-rated factual consistency levels (Gabriel et al., 2021; Fabbri et al., 2021). Entailment-based scorers are designed at summary level (Kryscinski et al., 2020) and finer-grained dependency relation level (Goyal and Durrett, 2020). QA models are employed to measure content consistency by reading the articles to answer questions generated from the summaries (Wang et al., 2020; Durmus et al., 2020), or considering the summaries for addressing questions derived from the source (Scialom et al., 2019). Though not focusing on evaluation, our work highlights that models can produce a significant amount of world knowledge which should be evaluated differently instead of as extrinsic hallucination (Maynez et al., 2020). We also show that world knowledge can possibly be distinguished from errors via model behavior understanding.

Training with negative samples has been investigated in several classic NLP tasks, such as grammatical error detection (Foster and Andersen, 2009)

¹Our code and annotated data are available at https://shuyangcao.github.io/projects/cliff_summ.

and dialogue systems (Li et al., 2019). Notably, negative sampling plays a key role in word representation learning (Mikolov et al., 2013) and training large masked language models, such as BERT and ALBERT, to induce better contextual representations (Devlin et al., 2019; Lan et al., 2020). For text generation tasks, unlikelihood training is proposed to penalize the generation of negative tokens (e.g., repeated words) and sentences (e.g., contradictory responses in a dialogue system) (Welleck et al., 2020; Li et al., 2020; He and Glass, 2020). We use contrastive learning that drives enhanced representation learning to better distinguish between factual and incorrect summaries, which encourages more faithful summary generation.

Contrastive Learning (CL) for NLP. CL has been a popular method for representation learning, especially for vision understanding (Hjelm et al., 2019; Chen et al., 2020). Only recently has CL been used for training language models with self-supervision (Fang et al., 2020), learning sentence representations (Gao et al., 2021), and improving document clustering (Zhang et al., 2021a). With a supervised setup, Gunel et al. (2021) adopt the contrastive objective to fine-tune pre-trained models on benchmark language understanding datasets. Using a similar idea, Liu and Liu (2021) enlarge the distances among summaries of different quality as measured by ROUGE scores.

3 CLIFF: Contrastive Learning Framework for Summarization

We design a contrastive learning (CL)-based training objective that drives the summarization model to learn a preference of faithful summaries over summaries with factual errors. It is then used for fine-tuning BART (Lewis et al., 2020) and PEGASUS (Zhang et al., 2020a) for training summarization models. Formally, let an article x have a set of reference summaries P (henceforth *positive samples*) and another set of erroneous summaries N (*negative samples*). The contrastive learning objective is (Khosla et al., 2020; Gunel et al., 2021):

$$l_{cl}^x = -\frac{1}{\binom{|P|}{2}} \sum_{\substack{y_i, y_j \in P \\ y_j \neq y_i}} \log \frac{\exp(\text{sim}(\mathbf{h}_i, \mathbf{h}_j)/\tau)}{\sum_{\substack{y_k \in P \cup N \\ y_k \neq y_i}} \exp(\text{sim}(\mathbf{h}_i, \mathbf{h}_k)/\tau)} \quad (1)$$

where \mathbf{h}_i , \mathbf{h}_j , and \mathbf{h}_k are representations for summaries y_i , y_j , and y_k . $\text{sim}(\cdot, \cdot)$ calculates the cosine

similarity between summary representations. τ is a temperature and is set to 1.0.

Importantly, summaries in P and N are included in the *same batch* during training, so that the model acquires better representations to differentiate correct summaries from those with errors by comparing the two types of samples, thus maximizing the probabilities of the positive samples and minimizing the likelihoods of the corresponding negative samples. The **CL objective** on the full training set, denoted as \mathcal{L}_{CL} , is the sum of losses l_{cl}^x over all samples.

To effectively employ CL in summarization, we need to address two challenges: (1) how to automatically construct both positive and negative samples, which are critical for CL efficacy (Chen et al., 2020), and (2) how to represent the summaries (i.e., \mathbf{h}_*). Below we describe positive sample generation and options for \mathbf{h}_* , leaving the strategies for negative samples to § 5.

Positive Sample Construction (P). Summarization datasets often contain a single reference for each article. To create multiple positive samples, in our pilot study, we experiment with paraphrasing with synonym substitution (Ren et al., 2019), randomly replacing words based on the prediction of masked language models (Kobayashi, 2018), and back-translation (Mallinson et al., 2017). We find back-translation to be best at preserving meaning and offering language variation, and thus use NLPAug² to translate each reference to German and back to English. Together with the reference, the best translation is kept and added to P , if no new named entity is introduced.

Summary Representation (\mathbf{h}_*). We use the outputs of the decoder’s last layer, and investigate three options that average over *all tokens*, *named entity tokens*, and the *last token* of the decoded summary. Entities and other parsing results are obtained by spaCy (Honnibal et al., 2020). We further consider adding a multi-layer perceptron (MLP) with one hidden layer to calculate the final \mathbf{h}_* .

The final training objective combines the typical cross-entropy loss \mathcal{L}_{CE} and our contrastive learning objective: $\mathcal{L} = \mathcal{L}_{CE} + \lambda \mathcal{L}_{CL}$, where λ is a scalar and set to 1.0 for all experiments.

²<https://github.com/makcedward/nlpaug>

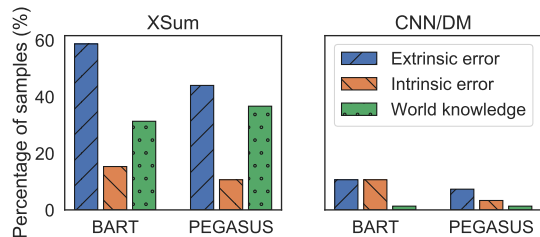


Figure 2: Percentage of samples with intrinsic and extrinsic error spans for models fine-tuned from BART and PEGASUS on XSum and CNN/DM.

4 Summary Error Annotation and Model Behavior Analysis

We first describe annotating unfaithfulness errors by state-of-the-arts, i.e., models fine-tuned from BART and PEGASUS on XSum and CNN/DM. We then probe into the model generation behavior that is indicative of errors, which guides the design of negative sample construction strategies.

600 ($150 \times 2 \times 2$) summaries are annotated to demonstrate how often do the models “hallucinate”, i.e., generating content not grounded by the source. To characterize errors, we annotate *text spans* in summaries with (i) **intrinsic errors** caused by misconstructing phrases or clauses from the source; and (ii) **extrinsic errors** which include words not in the source that are either unverifiable or cannot be verified by Wikipedia. Content not covered by the article but can be validated by Wikipedia is annotated as **world knowledge**, and the models’ behavior pattern when generating them differs from when they generate errors.

Two fluent English speakers with extensive experience in summary evaluation and error labeling are hired. For each sample, they are shown the article and two system summaries, and instructed to annotate text spans with the aforementioned errors and world knowledge. After labeling every 50 samples, the annotators discuss and resolve any disagreement. The Fleiss’s Kappas on XSum and CNN/DM are 0.35 and 0.45.

Error statistics are displayed in Fig. 2. Extrinsic errors dominate both datasets, especially on XSum. 58.7% of summaries by BART (and 44.0% by PEGASUS) contain at least one extrinsic error. Noticeably, PEGASUS is a newer model pre-trained with a larger amount of data, thus contains less errors than BART and other older models studied for error annotations by Maynez et al. (2020). This observation also highlights the usage of our anno-

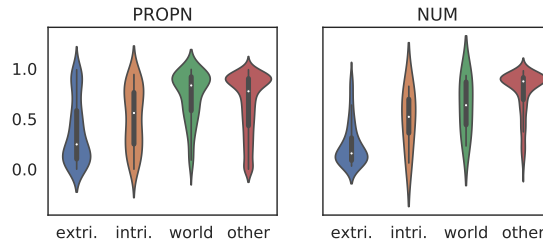


Figure 3: Probability distributions of generating the first tokens of proper nouns and numbers, grouped by extrinsic errors, intrinsic errors, world knowledge, and other correct tokens.

tations for future development and evaluation of summarization models.

Low confidence generation is indicative of extrinsic errors. Inspired by recent work that studies model prediction confidence (Liu et al., 2021), we examine generation probabilities for tokens of *different part-of-speech (POS) tags*. Fig. 3 shows salient results on the generation probabilities of the first token of a proper noun or a number (with additional analysis provided in Appendix A). As observed, model confidence tends to be lower for the first tokens of proper nouns and numbers if they are part of spans with *extrinsic errors*. Also note that world knowledge, which cannot be inferred from the source either, often has higher generation probability than extrinsic errors. Take this snippet generated by a fine-tuned BART as an example: “Manchester United captain Wayne Rooney’s testimonial game against Manchester City. . .”. “Manchester City” is an extrinsic error and “Wayne” is produced as world knowledge. The model assigns a low probability of 0.10 to the first token of “Manchester City” and a high probability of 0.92 to token “Wayne”. This implies that model confidence can be a useful indicator for negative sample collection.

5 Negative Sample Construction

Here we describe four strategies for constructing negative samples that *modify the references* (§ 5.1-5.3) or use *system generated summaries* (5.4).

5.1 Entity Swap

Entity swap imitates intrinsic errors, as over 55% of intrinsic errors in our annotations are found to contain named entities. We construct negative samples by swapping named entities in the references with other randomly selected entities of the same entity type in the source (SWAPENT). One sam-

REFERENCE: A “rare” short-eared owl found emaciated in Flintshire is now recuperating well, the RSPCA have said.
SWAPENT: Flintshire → Bettisfield ⇒ A “rare” short-eared owl found emaciated in Bettisfield is now recuperating well, the RSPCA have said.
MASKENT: A “rare” short-eared owl found emaciated in [MASK] is now recuperating well, the RSPCA have said. ⇒ A “rare” short-eared owl found emaciated in a field in South Yorkshire is now recuperating well, the RSPCA have said.
MASKREL: A “rare” short-eared owl found [MASK] in [MASK] is now recuperating well, the RSPCA have said. ⇒ A “rare” short-eared owl found dead in London is now recuperating well, the RSPCA have said.
REGENENT: A “rare” short-eared owl found emaciated in ____ ⇒ A “rare” short-eared owl found emaciated in Nottinghamshire is now at a wildlife centre to recover.
REGENREL: A “rare” short-eared owl found ____ ⇒ A “rare” short-eared owl found in the grounds of a former coal mine is being cared for by the RSPCA in Somerset.
SYSLOWCON: An injured golden owl found in a former coal mine in Lancashire is being cared for by the RSPCA.

Table 1: Negative sample construction strategies (§ 5). For summaries edited from the reference, their differences are **bolded**. Introduced errors are in **red**. Text before ____ is the prefix for regeneration.

ple is constructed for each entity in the reference. Though this idea has been studied by Kryscinski et al. (2020), they allow entities of different types to be used, e.g., a PERSON can be replaced by a LOCATION. Examples are displayed in Table 1.

SWAPENT has the advantage of not depending on any trained model. Yet it only introduces intrinsic errors and lacks the coverage for extrinsic errors, which is addressed by the following generation-based methods.

5.2 Mask-and-fill with BART

To simulate extrinsic errors, we leverage large unconditional language models’ capability of converting a sequence with masked tokens into a fluent and appropriate sequence. Specifically, we replace each named entity in a reference with a [MASK] token and encode it with BART (without any fine-tuning). BART then fills this partially masked summary with newly generated entities (**MASKENT**). BART is chosen since it can fill [MASK] with varying number of tokens. For each entity in the reference, we sample three summaries and only retain the ones containing at least one entity that is absent from both the source and the reference.

Up to now, the two introduced strategies both focus on incorrect named entities. To cover more diverse extrinsic and intrinsic errors (Goyal and Durrett, 2020), we extend MASKENT to contain

relations (**MASKREL**). We first obtain dependency relations using Stanza (Qi et al., 2020), with each relation denoted as $\langle \text{gov}, \text{rel}, \text{dep} \rangle$. To incorporate more context, we consider noun phrase spans enclosing the token of gov or dep if it is a content word and the noun phrase contains a named entity. Similar to MASKENT, three negative samples are generated by BART based on the input with both gov and dep spans masked in the reference. Only the samples that introduce any new dependency relation that is not contained in the source nor the reference are kept. Specifically, we consider a match of a dependency relation as the same form or synonyms of its gov and and dep is found in the source or the reference with the same relation.

Both MASKENT and MASKREL can create more extrinsic errors compared to other strategies introduced in this section, since negative samples are generated without being grounded on the source articles. However, their constructed negative samples may contained drifted topics that can be easily detected by a summarization model, resulting with less efficient training signals.

5.3 Source-conditioned Regeneration

To ground negative sample generation with the article, we further design a regeneration strategy based on conditional generation. For each named entity in the reference, we treat the text before it as a *prompt*. A summarizer, e.g., fine-tuned from BART or PEGASUS, first reads in the source using the encoder, then receives the prompt as the first part of the decoder output, and finally decodes the rest of the content based on nucleus sampling (Holtzman et al., 2020) with a cumulative probability threshold of 0.7. The prompt and the regenerated text comprise the final negative sample. This method is denoted as **REGENENT**.

We also extend entities to relations with expanded governor and dependent spans (**REGENREL**). Here, we consider a prompt as the text before the gov or dep span, whichever occurs first. For both REGENENT and REGENREL, three negative samples are generated for each prompt, and a sample is kept if it introduces any new entity (for REGENENT) or dependency relation (for REGENREL) with regard to the source and the reference.

Negative samples generated by both methods are more relevant to the article than the mask-and-fill

strategy, yet they may still miss certain types of errors and differ from real model outputs, since they are modified from the reference summaries.

5.4 System Generation

Motivated by the model confidence analysis in § 4, we explore using system generated summaries as negative samples. We first run fine-tuned BART or PEGASUS on the same training set to decode summaries. For each summary, we check the model confidence on the first token of each proper noun and number span. If the probability is below a threshold, we keep it as a negative sample (**SYSLOWCON**). The threshold is tuned by maximizing F1 based on our error annotations.

We consider all beams at the last decoding step as candidates. We use beam sizes of 6 and 4 for XSum and CNN/DM. Statistics of negative samples constructed by different strategies are in Appendix B.

6 Experiment Setup

Evaluation Metrics. QuestEval (Scialom et al., 2021) is used as the main metric to evaluate summaries’ factual consistency. Given an article and a summary, QuestEval first generates natural language questions for entities and nouns from both. A QA model then consumes the article to answer questions derived from the summary, producing a score. Another score is obtained from a QA model addressing article-based questions after reading the summary. The final QuestEval score is the harmonic mean of the two. We use the version with learned weights for questions, which has shown high correlation with human judged consistency and relevance.

We further use FactCC (Kryscinski et al., 2020), trained based on their negative sample construction method, to measure if the summary can be entailed by the source. We also report ROUGE-L (Lin, 2004). Both FactCC and ROUGE-L reasonably correlate with summary factuality as judged by human (Pagnoni et al., 2021).

Based on our error annotations, we report the correlations between each metric and the error rate—percentage of tokens being part of an error span, and the raw number of errors (Table 2). QuestEval correlates better on both aspects than other metrics.

Comparisons. In addition to the models fine-tuned with cross-entropy loss (**CRSENTROPY**), we consider reranking beams based on FactCC

Metric	XSum		CNN/DM	
	% of Err	# of Err	% of Err	# of Err
QuestEval	-0.43*	-0.25*	-0.33*	-0.29*
FactCC	-0.02	-0.15*	-0.13*	-0.12*
ROUGE-1	-0.16*	-0.02	-0.03	-0.06
ROUGE-2	-0.11*	-0.05	-0.02	-0.04
ROUGE-L	-0.13*	-0.03	-0.06	-0.08

Table 2: Pearson correlation between metrics and error rates and numbers of errors. *: p-value < 0.05.

score (also our metric) at the last decoding step (**ENTAILRANK**). We also include three common methods of improving factuality: (1) (**CORRECTION**) fine-tunes BART to fix summary errors as a separate step (Cao et al., 2020). (2) **SUBSETFT** fine-tunes large models based on training samples without any dependency relation error (Goyal and Durrett, 2021), with released checkpoint only available for XSum. (3) **FASUM** modifies Transformer by incorporating knowledge graphs for factual consistency (Zhu et al., 2021), with model outputs only on CNN/DM.

Moreover, we compare with **unlikelihood training** that penalizes the probabilities of all tokens in a negative sample (Li et al., 2020). Given a negative sample y' , the loss is defined as $-\sum_{t=1}^{|y'|} \log(1 - p(y'_t | y'_{1:t-1}, x))$, where $p(y'_t | y'_{1:t-1}, x)$ is the output probability at the t -th step. We combine the unlikelihood training objective with cross-entropy loss with equal weights for fine-tuning.

Lastly, we compare our negative sample strategies with negative samples constructed for training the FactCC scorer, denoted as **FCSAMPLE**. For CL only, we compare with using other samples in the same batch as negative samples (**BATCH**), a common practice for CL-based representation learning (Gao et al., 2021; Zhang et al., 2021a).

7 Results

7.1 Automatic Evaluation

We report results by models fine-tuned from BART and PEGASUS with different objectives and negative samples on XSum and CNN/DM in Tables 3 and 4. CLIFF models use a summary representation of averaging over all tokens with MLP projection, with other variants discussed in § 7.3. Unless explicitly stated, comparison models are fine-tuned from the same large model used by CLIFF.

First, comparing with other factuality improvement models (top of the tables), *almost all* CLIFF models trained with different negative samples uni-

Model	XSum			CNN/DM		
	QEval	FC	R-L	QEval	FC	R-L
<i>Comparisons without Negative Samples</i>						
CRSENTROPY	33.09	23.92	37.14	50.94	49.07	40.82
ENTAILRANK	32.95	38.45	36.55	51.02	49.84	40.89
CORRECTION	33.12	24.14	37.11	50.93	49.06	40.82
SUBSETFT	32.25	21.83	30.35	-	-	-
FASUM	-	-	-	50.73	50.58	37.18
<i>Comparisons with Unlikelihood Training</i>						
FCSAMPLE	32.93	24.46	33.96	50.60	35.09	41.22
SWAPENT	32.90	23.94	34.96	49.77	32.37	40.18
MASKENT	33.21	24.22	33.89	51.01	48.57	41.23
MASKREL	33.18	23.50	34.56	51.00	48.35	41.15
REGENENT	32.41	24.12	37.08	50.97	48.59	41.07
REGENREL	30.86	24.58	37.18	50.97	48.42	41.14
SYSLOWCON	32.01	26.30	32.04	50.82	48.66	40.81
<i>Our Method: CLIFF</i>						
BATCH	33.18	24.88	36.76	50.99	52.18	40.98
FCSAMPLE	33.15	24.50	36.72	51.02	49.62	41.06
SWAPENT	33.30	25.67*	35.60	51.05	50.96	40.89
MASKENT	33.32*	25.73*	36.02	50.98	49.04	41.06
MASKREL	33.35*	25.69*	35.86	51.03	49.89	41.04
REGENENT	33.15	24.64	36.32	51.04	49.91	41.11
REGENREL	33.22	25.39	36.21	50.96	49.48	41.11
SYSLOWCON	33.35*	25.47*	36.19	51.05	50.05	41.01

Table 3: Results of models fine-tuned from BART on XSum and CNN/DM. QEval: QuestEval; FC: FactCC; R-L: ROUGE-L. The best result per metric per dataset is **bolded**. For models of unlikelihood training and CLIFF that use the same negative samples, the better of the two is highlighted with **green**. *: our model is significantly better than CRSENTROPY (approximation randomization test, $p < 0.005$).

formly produce higher QuestEval scores across datasets with both large models, with the improvements more pronounced on XSum. Importantly, ROUGE scores for CLIFF models are comparable or better than baselines trained with cross-entropy, e.g., on CNN/DM as in Table 3. A similar trend is observed with the FactCC metric, especially when using PEGASUS as the base model (Table 4). Note that ENTAILRANK tends to yield significantly higher FactCC scores, though it obtains lower QuestEval scores than the cross-entropy baseline. Human inspection finds that ENTAILRANK can pick up beams with peculiar words of high FactCC scores, without improving factuality. Moreover, other comparisons based on post CORRECTION and model engineering (FASUM) only offer incremental gains. The sample selection-based method, SUBSETFT, sacrifices ROUGE scores significantly. Overall, CLIFF demonstrates stronger generalizability.

Second, CLIFF is more effective and robust than

Model	XSum			CNN/DM		
	QEval	FC	R-L	QEval	FC	R-L
<i>Comparisons without Negative Samples</i>						
CRSENTROPY	32.50	25.48	39.07	50.21	44.44	40.39
ENTAILRANK	32.42	41.90	38.47	50.15	61.04	40.67
CORRECTION	32.55	25.15	39.02	49.48	32.96	39.79
<i>Comparisons with Unlikelihood Training</i>						
FCSAMPLE	32.79	25.37	38.46	50.63	45.45	39.28
SWAPENT	32.88	24.76	37.91	50.43	43.02	38.96
MASKENT	33.04	26.30	37.51	51.11	52.19	39.34
MASKREL	33.08	24.38	38.05	51.14	52.93	39.31
REGENENT	32.89	24.46	38.47	51.11	52.90	39.23
REGENREL	32.91	24.80	38.46	51.07	53.68	39.45
SYSLOWCON	31.66	26.06	34.03	50.92	51.08	39.19
<i>Our Method: CLIFF</i>						
BATCH	32.64	24.96	38.42	51.03*	51.81*	39.38
FCSAMPLE	32.96*	25.28	38.58	51.00*	51.80*	39.37
SWAPENT	33.09*	25.09	38.58	51.16*	52.97*	38.95
MASKENT	33.09*	25.75	38.12	51.13*	53.60*	39.24
MASKREL	33.06*	25.28	38.37	51.17*	53.34*	39.36
REGENENT	33.09*	24.48	38.33	50.99*	52.18*	39.28
REGENREL	33.16*	24.82	38.30	51.16*	53.21*	39.25
SYSLOWCON	33.21*	25.18	38.18	50.85*	53.73*	39.30

Table 4: Results of models fine-tuned from PEGASUS on XSum and CNN/DM. We report results on 5,000 randomly selected samples on CNN/DM, due to long running time of QuestEval. For models of unlikelihood training and CLIFF that use the same negative samples, the better of the two is highlighted with **green**. *: our model is significantly better than CRSENTROPY (approximation randomization test, $p < 0.005$).

unlikelihood training with the same negative samples. According to Table 3, using 7 negative sample construction strategies on two datasets, CLIFF obtains higher QuestEval scores than unlikelihood training in 12 out of the 14 comparisons. Using PEGASUS, CLIFF also outperforms in 11 setups as listed in Table 4. Similar trends are found on FactCC and ROUGE-L. Another noteworthy piece is that CLIFF’s improvements over the cross-entropy baseline are more consistent, whereas unlikelihood training occasionally hurts factuality or ROUGE scores significantly. We believe the key advantage of CLIFF resides in its measure of representation similarities between positive and negative samples in the same batch, allowing models to better differentiate between correct and erroneous summaries.

Finally, among all variants, CLIFF trained with low confidence summaries as negative samples obtains the best QuestEval scores on the more abstract dataset. As seen in Table 3, using low confidence summaries also improves FactCC scores on both datasets, and enhances ROUGE-L on the more

Strategy	XSum			CNN/DM		
	QEval	FC	R-L	QEval	FC	R-L
SYSLOWCON	33.35	25.47	36.19	51.05	50.05	41.01
+ SWAPENT	33.40	25.50	35.50	51.32	53.95	40.57
+ MASKENT	33.21	25.47	35.91	51.16	51.90	40.66
+ MASKREL	33.39	25.20	35.70	51.24	52.48	40.80
+ REGENENT	33.31	25.07	35.94	51.21	51.86	40.91
+ REGENREL	33.38	24.97	36.03	51.13	50.85	40.97

Table 5: Results of fine-tuned BART with combinations of negative sample construction strategies.

extractive dataset CNN/DM. This indicates that *system generated summaries contribute more diverse errors made by existing models organically*, which are particularly suitable for our CL framework. As we use summaries generated by the same model for CLIFF training, one future direction is to use outputs by different models. For our mask-and-fill and source-conditioned regeneration strategies, we find that relation-anchored construction often beats their entity-anchored counterparts. This calls for efforts that steer the entity-driven methods to a more relation-focused direction.

Combining Strategies. We further show results by fine-tuning BARTs using samples based on combined negative sample construction strategies in Table 5. As can be seen, *combining SYSLOWCON and other strategies yields better QuestEval scores than models trained with negative samples by any single strategy*, except for MASKENT and REGENENT on XSum. This signifies the importance of covering diverse types of errors in negative samples.

7.2 Human Evaluation

Pairwise Comparison with Cross-entropy. We recruit the two human annotators for our summary error study, as well as another experienced annotator, to evaluate summary **informativeness** and **factual consistency**. For each article, the judges are shown summaries generated by the CRSENTROPY model and four other systems. They then rate each system summary against the CRSENTROPY summary. All four summaries generated by different factuality-improved models are shown in random order without system names shown, ensuring the fair comparison among them.

We randomly pick 100 articles from each dataset used in our error analysis study in § 4, and evaluate summaries generated by ENTAILRANK, unlikelihood training (ULL) with negative samples con-

Model	Inform.			Factual.		
	Win \uparrow	Tie	Lose \downarrow	Win \uparrow	Tie	Lose \downarrow
ENTAILRANK	3.3	84.3	12.3	23.7	71.7	4.7
ULL. MASKENT	6.3	80.7	13.0	26.3	62.0	11.7
CL. BATCH	5.3	80.0	14.7	21.7	68.0	10.3
CL. SYSLOWCON	8.7	78.3	13.0	31.3	61.7	7.0

(a) XSum

Model	Inform.			Factual.		
	Win \uparrow	Tie	Lose \downarrow	Win \uparrow	Tie	Lose \downarrow
ENTAILRANK	2.3	86.3	11.3	4.7	94.7	0.7
ULL. MASKENT	18.0	71.0	11.0	17.3	79.7	3.0
CL. BATCH	17.3	74.7	8.0	20.7	77.0	2.3
CL. SYSLOWCON	15.7	75.7	8.7	20.0	77.7	2.3

(b) CNN/DM

Table 6: Percentages of summaries that are better than, tied with, or worse than CRSENTROPY, in informativeness (Inform.) and factual consistency (Factual.) The Krippendorff’s α s are 0.33 and 0.62 for the two aspects on XSum, and 0.34 and 0.89 on CNN/DM. Our CL method using low confidence summaries is more frequently rated as better for informativeness and factuality on the more abstractive dataset XSum.

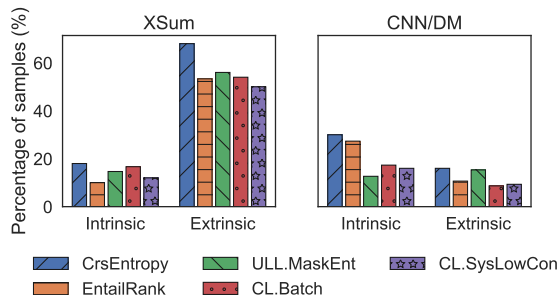


Figure 4: Portions of summaries with errors. CL models consistently reduce both types of errors.

structed by MASKENT, and CLIFF models trained with BATCH and SYSLOWCON negative samples. All are fine-tuned from BART. Detailed evaluation guidelines are in Appendix D.

Table 6 shows that on the more abstractive XSum data *CL trained with low confidence samples are more frequently rated as being more informative and more factual than CRSENTROPY summaries*. This echoes our automatic evaluations with QuestEval in § 7.1. On CNN/DM, all models trained with negative samples produce summaries with better informativeness and faithfulness. In contrast, ENTAILRANK summaries are less distinguishable from outputs by CRSENTROPY on both datasets, as more ties are found. We show sample outputs in Fig. 1, with additional examples in Appendix E.

Intrinsic vs. Extrinsic Errors. Next, the annotators are asked to label text spans with intrinsic and

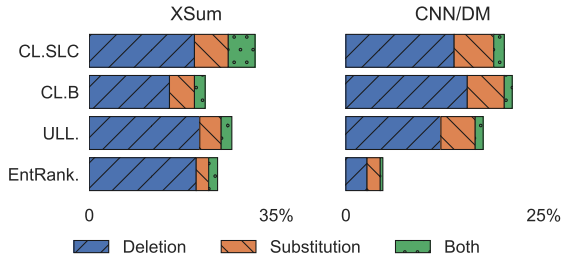


Figure 5: Summaries use different portions of error correction operations. Contrastive learning with SYS-LOWCON (CL.SLC) and BATCH (CL.B) substitute errors with correct content more often than unlikelihood training with MASKENT and ENTAILRANK.

extrinsic errors as done in § 4. Fig. 4 shows that *CL is more effective at reducing extrinsic errors than unlikelihood training can* on both datasets. We also observe slight decreases of world knowledge in the summaries (figure attached in Appendix D).

Error Correction Operations. Finally, with reference to CRSENTROPY summaries, human judges are instructed to label each system summary as whether it corrects any error by CRSENTROPY using **deletion** of the incorrect content, **substitution** with factual information, or **both**. As seen in Fig. 5, CL-based models restore factually consistent information, e.g., by replacing erroneous names and numbers with correct ones, more frequently than unlikelihood training or entailment reranking.

7.3 Variants of Summary Representation

Sample representation is critical for CL to be effective. Here we investigate summary representation variants as discussed in § 3. There are two major considerations: (1) Should we consider all tokens in a summary or only representative ones (e.g., entities or last token)? (2) Should additional transformation, i.e., an MLP, be used?

Experiments on XSum using three negative sample construction strategies demonstrate that *averaging the decoder outputs of all tokens and adding an MLP projection yield the best overall performance*, as shown in Table 7. The implications are at least two-fold. First, even for entity- or relation-triggered sample modifications, using more global context helps with CL training. Second, additional transformation can help avoid model degeneration. For instance, more nonsensical and repetitive content is produced by variants without MLP.

Rep.	MLP	SWAPENT		MASKREL		SYSLOWCON	
		QEval	FC	QEval	FC	QEval	FC
<i>BART</i>							
Last	✓	33.15	25.10	33.20	25.29	33.10	24.85
Last		+0.13	+0.02	-0.01	-0.32	-0.07	-0.10
Entity	✓	33.35	25.41	33.34	25.44	33.32	24.46
Entity		-0.13	-0.07	-0.14	-0.05	-0.29	+0.72
All	✓	33.30	25.67	33.35	25.69	33.35	25.47
All		-0.23	-0.80	-0.04	-0.48	-0.26	-0.40
<i>PEGASUS</i>							
Last	✓	33.07	25.45	32.99	25.09	33.18	24.94
Last		-0.07	-0.56	+0.01	-0.01	-0.02	-0.04
Entity	✓	33.03	25.43	33.05	24.77	33.20	24.59
Entity		+0.01	-0.34	-0.05	+0.64	-0.30	+0.05
All	✓	33.09	25.09	33.06	25.28	33.21	25.18
All		-0.11	+0.25	+0.03	-0.19	-0.02	-0.80

Table 7: Comparing different formulations of summary representation in CL. For models without MLP, we display score changes from their counterparts. Overall, using all tokens with MLP produces better summaries.

8 Conclusion

We present CLIFF, a contrastive learning-based framework to promote faithfulness and factuality of abstractive summaries. CLIFF uses both references and summaries that are factually inconsistent with the articles to train systems to be better at discriminating errors from factual and salient content. We further study strategies that automatically create erroneous summaries by editing from references or leveraging systems outputs, inspired by our new summary error analysis on state-of-the-art models. Both automatic evaluation and human ratings show that CLIFF achieves consistent improvements over competitive comparison methods, and is generalizable across datasets with systems fine-tuned from different large models.

Acknowledgements

This research is supported in part by Oracle for Research Cloud Credits, National Science Foundation through a CAREER award IIS-2046016, and the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via contract # FA8650-17-C-9116. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of ODNI, IARPA, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein. We thank three anonymous reviewers for their valuable suggestions.

References

- Rahul Aralikatte, Shashi Narayan, Joshua Maynez, Sascha Rothe, and Ryan McDonald. 2021. [Focus attention: Promoting faithfulness and diversity in summarization](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 6078–6095, Online. Association for Computational Linguistics.
- Meng Cao, Yue Dong, Jiapeng Wu, and Jackie Chi Kit Cheung. 2020. [Factual error correction for abstractive summarization models](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6251–6258, Online. Association for Computational Linguistics.
- Ziqiang Cao, Furu Wei, Wenjie Li, and Sujian Li. 2018. Faithful to the original: Fact aware neural abstractive summarization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.
- Sihao Chen, Fan Zhang, Kazoo Sone, and Dan Roth. 2021. [Improving faithfulness in abstractive summarization with contrast candidate generation and selection](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5935–5941, Online. Association for Computational Linguistics.
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. 2020. [A simple framework for contrastive learning of visual representations](#). In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 1597–1607. PMLR.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Yue Dong, Shuohang Wang, Zhe Gan, Yu Cheng, Jackie Chi Kit Cheung, and Jingjing Liu. 2020. [Multi-fact correction in abstractive text summarization](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9320–9331, Online. Association for Computational Linguistics.
- Esin Durmus, He He, and Mona Diab. 2020. [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*, pages 5055–5070, Online. Association for Computational Linguistics.
- Alexander R Fabbri, Wojciech Kryściński, Bryan McCann, Caiming Xiong, Richard Socher, and Dragomir Radev. 2021. Summeval: Re-evaluating summarization evaluation. *Transactions of the Association for Computational Linguistics*, 9:391–409.
- Tobias Falke, Leonardo F. R. Ribeiro, Prasetya Ajie Utama, Ido Dagan, and Iryna Gurevych. 2019. [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.
- Hongchao Fang, Sicheng Wang, Meng Zhou, Jiayuan Ding, and Pengtao Xie. 2020. [Cert: Contrastive self-supervised learning for language understanding](#).
- Jennifer Foster and Oistein Andersen. 2009. [GenERRate: Generating errors for use in grammatical error detection](#). In *Proceedings of the Fourth Workshop on Innovative Use of NLP for Building Educational Applications*, pages 82–90, Boulder, Colorado. Association for Computational Linguistics.
- Saadia Gabriel, Asli Celikyilmaz, Rahul Jha, Yejin Choi, and Jianfeng Gao. 2021. [GO FIGURE: A meta evaluation of factuality in summarization](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 478–487, Online. Association for Computational Linguistics.
- Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021. [Simcse: Simple contrastive learning of sentence embeddings](#).
- Tanya Goyal and Greg Durrett. 2020. [Evaluating factuality in generation with dependency-level entailment](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3592–3603, Online. Association for Computational Linguistics.
- Tanya Goyal and Greg Durrett. 2021. [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*, pages 1449–1462, Online. Association for Computational Linguistics.
- Beliz Gunel, Jingfei Du, Alexis Conneau, and Veselin Stoyanov. 2021. [Supervised contrastive learning for pre-trained language model fine-tuning](#). In *International Conference on Learning Representations*.
- Tianxing He and James Glass. 2020. [Negative training for neural dialogue response generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2044–2058, Online. Association for Computational Linguistics.
- Karl Moritz Hermann, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. [Teaching machines to read](#)

- and comprehend. In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems* 28, pages 1693–1701. Curran Associates, Inc.
- R Devon Hjelm, Alex Fedorov, Samuel Lavoie-Marchildon, Karan Grewal, Phil Bachman, Adam Trischler, and Yoshua Bengio. 2019. [Learning deep representations by mutual information estimation and maximization](#). In *International Conference on Learning Representations*.
- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2020. [The curious case of neural text de-generation](#). In *International Conference on Learning Representations*.
- Matthew Honnibal, Ines Montani, Sofie Van Landeghem, and Adriane Boyd. 2020. [spaCy: Industrial-strength Natural Language Processing in Python](#).
- Luyang Huang, Lingfei Wu, and Lu Wang. 2020. [Knowledge graph-augmented abstractive summarization with semantic-driven cloze reward](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5094–5107, Online. Association for Computational Linguistics.
- Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschiot, Ce Liu, and Dilip Krishnan. 2020. [Supervised Contrastive Learning](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 18661–18673. Curran Associates, Inc.
- Sosuke Kobayashi. 2018. [Contextual augmentation: Data augmentation by words with paradigmatic relations](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 452–457, New Orleans, Louisiana. Association for Computational Linguistics.
- Wojciech Kryscinski, Bryan McCann, Caiming Xiong, and Richard Socher. 2020. [Evaluating the factual consistency of abstractive text summarization](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9332–9346, Online. Association for Computational Linguistics.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2020. [Albert: A lite bert for self-supervised learning of language representations](#). In *International Conference on Learning Representations*.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. [BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- Jia Li, Chongyang Tao, Wei Wu, Yansong Feng, Dongyan Zhao, and Rui Yan. 2019. [Sampling matters! an empirical study of negative sampling strategies for learning of matching models in retrieval-based dialogue systems](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 1291–1296, Hong Kong, China. Association for Computational Linguistics.
- Margaret Li, Stephen Roller, Iliia Kulikov, Sean Welleck, Y-Lan Boureau, Kyunghyun Cho, and Jason Weston. 2020. [Don’t say that! making inconsistent dialogue unlikely with unlikelihood training](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4715–4728, Online. Association for Computational Linguistics.
- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Tianyu Liu, Yizhe Zhang, Chris Brockett, Yi Mao, Zhifang Sui, Weizhu Chen, and Bill Dolan. 2021. [A token-level reference-free hallucination detection benchmark for free-form text generation](#).
- Yang Liu and Mirella Lapata. 2019. [Text summarization with pretrained encoders](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3730–3740, Hong Kong, China. Association for Computational Linguistics.
- Yixin Liu and Pengfei Liu. 2021. [SimCLS: A simple framework for contrastive learning of abstractive summarization](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 1065–1072, Online. Association for Computational Linguistics.
- Jonathan Mallinson, Rico Sennrich, and Mirella Lapata. 2017. [Paraphrasing revisited with neural machine translation](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 881–893, Valencia, Spain. Association for Computational Linguistics.
- Joshua Maynez, Shashi Narayan, Bernd Bohnet, and Ryan McDonald. 2020. [On faithfulness and factuality in abstractive summarization](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1906–1919, Online. Association for Computational Linguistics.

- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. [Distributed representations of words and phrases and their compositionality](#).
- Feng Nan, Ramesh Nallapati, Zhiguo Wang, Cicero Nogueira dos Santos, Henghui Zhu, Dejjiao Zhang, Kathleen McKeown, and Bing Xiang. 2021. [Entity-level factual consistency of abstractive text summarization](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 2727–2733, Online. 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*, pages 1797–1807, Brussels, Belgium. Association for Computational Linguistics.
- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. [fairseq: A fast, extensible toolkit for sequence modeling](#). In *Proceedings of NAACL-HLT 2019: Demonstrations*.
- Artidoro Pagnoni, Vidhisha Balachandran, and Yulia Tsvetkov. 2021. [Understanding factuality in abstractive summarization with frank: A benchmark for factuality metrics](#).
- Taesung Park, Alexei A Efros, Richard Zhang, and Jun-Yan Zhu. 2020. [Contrastive learning for unpaired image-to-image translation](#). In *European Conference on Computer Vision*, pages 319–345. Springer.
- Peng Qi, Yuhao Zhang, Yuhui Zhang, Jason Bolton, and Christopher D. Manning. 2020. [Stanza: A python natural language processing toolkit for many human languages](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 101–108, Online. Association for Computational Linguistics.
- Shuhui Ren, Yihe Deng, Kun He, and Wanxiang Che. 2019. [Generating natural language adversarial examples through probability weighted word saliency](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1085–1097, Florence, Italy. Association for Computational Linguistics.
- Nikunj Saunshi, Orestis Plevrakis, Sanjeev Arora, Mikhail Khodak, and Hrishikesh Khandeparkar. 2019. [A theoretical analysis of contrastive unsupervised representation learning](#). In *International Conference on Machine Learning*, pages 5628–5637. PMLR.
- Thomas Scialom, Paul-Alexis Dray, Gallinari Patrick, Lamprier Sylvain, Piwowarski Benjamin, Staiano Jacopo, and Wang Alex. 2021. [Questeval: Summarization asks for fact-based evaluation](#). *arXiv preprint arXiv:2103.12693*.
- Thomas Scialom, Sylvain Lamprier, Benjamin Piwowarski, and Jacopo Staiano. 2019. [Answers unite! unsupervised metrics for reinforced summarization models](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3246–3256, Hong Kong, China. Association for Computational Linguistics.
- Alex Wang, Kyunghyun Cho, and Mike Lewis. 2020. [Asking and answering questions to evaluate the factual consistency of summaries](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5008–5020, Online. Association for Computational Linguistics.
- Sean Welleck, Ilya Kulikov, Stephen Roller, Emily Dinan, Kyunghyun Cho, and Jason Weston. 2020. [Neural text generation with unlikelihood training](#). In *International Conference on Learning Representations*.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Dejjiao Zhang, Feng Nan, Xiaokai Wei, Shangwen Li, Henghui Zhu, Kathleen McKeown, Ramesh Nallapati, Andrew Arnold, and Bing Xiang. 2021a. [Supporting clustering with contrastive learning](#).
- Han Zhang, Jing Yu Koh, Jason Baldridge, Honglak Lee, and Yinfei Yang. 2021b. [Cross-modal contrastive learning for text-to-image generation](#). *arXiv preprint arXiv:2101.04702*.
- Jingqing Zhang, Yao Zhao, Mohammad Saleh, and Peter Liu. 2020a. [Pegasus: Pre-training with extracted gap-sentences for abstractive summarization](#). In *International Conference on Machine Learning*, pages 11328–11339. PMLR.
- Yuhao Zhang, Derek Merck, Emily Tsai, Christopher D. Manning, and Curtis Langlotz. 2020b. [Optimizing the factual correctness of a summary: A study of summarizing radiology reports](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5108–5120, Online. Association for Computational Linguistics.

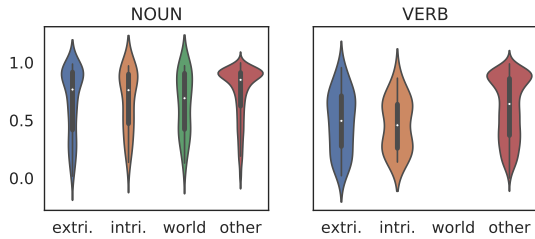


Figure 6: Probability distributions of generating the *first* tokens of nouns and verbs, grouped by extrinsic errors, intrinsic errors, world knowledge, and other correct tokens. No verb is annotated as world knowledge.

Zheng Zhao, Shay B. Cohen, and Bonnie Webber. 2020. [Reducing quantity hallucinations in abstractive summarization](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 2237–2249, Online. Association for Computational Linguistics.

Chenguang Zhu, William Hinthorn, Ruochen Xu, Qingkai Zeng, Michael Zeng, Xuedong Huang, and Meng Jiang. 2021. [Enhancing factual consistency of abstractive summarization](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 718–733, Online. Association for Computational Linguistics.

A Additional Analysis for Summary Error Annotation

We hire two fluent English speakers to annotate summary errors on XSum and CNN/DailyMail (CNN/DM). They annotate a common batch of 100 summaries generated by summarizers fine-tuned from BART and PEGASUS, with 50 articles in each batch. The two annotators are shown 50 HTML pages in a batch, each of which contains an article and two summaries generated by the two models. The detailed annotation guideline is given in Fig. 9.

For our analysis on token generation probabilities, we additionally show the distributions of the first token’s probability for nouns and verbs in Fig. 6. We also report the distributions of the non-first token’s probability for proper nouns, numbers, nouns, and verbs in Fig. 7. As can be seen, tokens within extrinsic and intrinsic errors have high generation probabilities when they are non-first tokens.

B Statistics for Datasets and Training Samples

Summarization Datasets. We follow the official data splits for the two datasets, with the number of samples in each split listed in Table 8.

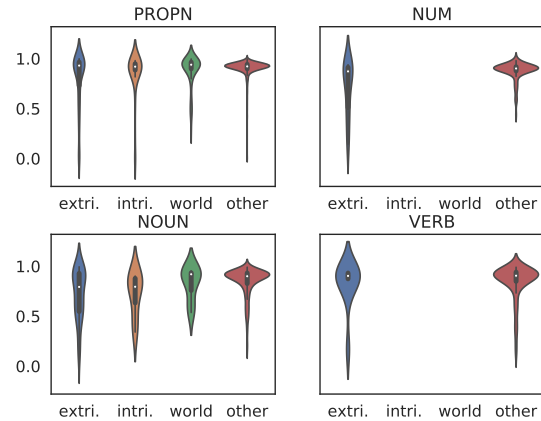


Figure 7: Probability distributions of generating the *non-first* tokens of proper nouns, numbers, nouns, and verbs, grouped by extrinsic errors, intrinsic errors, world knowledge, and other correct tokens. Non-first tokens do not exist for numbers and verbs, as they only contain single tokens.

Dataset	Train	Validation	Test
XSum	204,045	11,332	11,334
CNN/DM	287,227	13,368	11,490

Table 8: Numbers of samples in train/validation/test splits of XSum and CNN/DM.

Positive Samples. We observe unfaithful paraphrases by back-translation for some reference summaries, which are mainly due to the introduction of new entities and the rewriting of quoted text. Thus, we discard samples generated by back-translation that contain new entities and inconsistent quoted text. Finally, we obtain 182,114 and 91,468 positive samples by back-translation on XSum and CNN/DM.

Negative Samples. For consistency, we use the summarizer fine-tuned from BART in REGENENT, REGENREL (§ 5.3), and SYSLOWCON (§ 5.4) strategies. We tune a threshold to select negative samples from model generations in our SYSLOWCON strategy. The threshold is set to 0.21, with F1 scores of 73.99 and 40.49 on XSum and CNN/DM annotated samples generated by BART.

The numbers of negative samples constructed by each strategy for training on XSum and CNN/DM are shown in Table 9. SYSLOWCON constructs the least negative samples in total, while it achieves the best results as reported in our main paper (§ 7.1), indicating that its negative samples are more effective for training.

Strategy	XSum	CNN/DM
FCSAMPLE	936,164	1,291,710
SWAPENT	438,003	1,617,764
MASKENT	360,795	1,050,200
MASKREL	391,224	1,345,317
REGENENT	732,986	1,941,886
REGENREL	993,694	1,453,044
SYSLOWCON	401,112	502,768

Table 9: Numbers of negative samples constructed by different strategies on XSum and CNN/DM.

C Implementation Details

We use Fairseq (Ott et al., 2019) and Huggingface Transformers (Wolf et al., 2020) for our experiments with BART and PEGASUS. Our experiments are conducted on the RTX 8000 GPU with 48GB memory and the A100 GPU with 40GB memory.

Training Settings. For hyperparameters, we follow Lewis et al. (2020) for BART and Zhang et al. (2020a) for PEGASUS. During training, we randomly select 5 and 4 negative samples for each input article in XSum and CNN/DM. Mixed-precision training is not supported by the PEGASUS implementation and is utilized on BART only.

Decoding Settings. We use the beam search algorithm to decode summaries. For BART, we set the beam sizes as 6 and 4 on XSum and CNN/DM. A beam size of 8 is used for PEGASUS on both datasets.

Running Time and Model Sizes. The BART-based models take 6 and 13 hours for training on XSum and CNN/DM, and it takes 1.5 hour to decode on the two datasets. Meanwhile, training the PEGASUS-based models takes 8 and 25 hours for XSum and CNN/DM, and the decoding takes 1 hour.

As for model sizes, our BART-based models and PEGASUS-based models have 400M and 568M parameters.

D Human Evaluation

In § 7.2, we demonstrate the percentages of samples containing intrinsic errors and extrinsic errors for each model evaluated by human judges. Here, we report the percentages of samples containing world knowledge in Fig. 8. On XSum, all models produce less world knowledge compared to the model trained with cross-entropy loss, while gen-

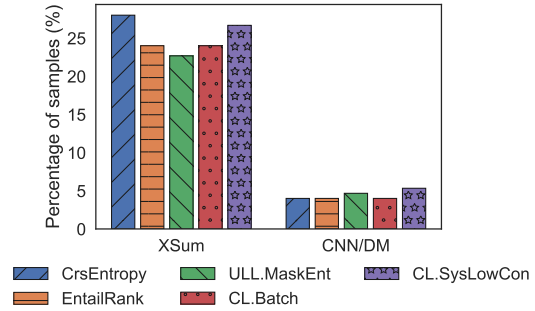


Figure 8: Percentages of samples containing world knowledge as labeled by human on the outputs of XSum and CNN/DM.

erating similar or greater amounts of samples with world knowledge on CNN/DM.

Our human evaluation guideline is shown in Fig. 10.

E Sample Outputs

We include more sample outputs in Fig. 11.

In this study, you will first read article-summary pairs and then identify three types of text spans in the summaries. These spans include content that is contradicted by or cannot be implied from the article. The description for each type is described below:

- **Intrinsic errors:** Text spans that misconstrue phrases or clauses from the article.
- **Extrinsic errors:** Text spans that include words that are not in the article and are not verifiable or cannot be verified by Wikipedia.
- **world knowledge:** Text spans that contain information that is not covered by the article but can be validated by Wikipedia.

When selecting spans, you should always make sure the spans are complete words.

In practice, you should follow these steps carefully: (1) read the article and summaries carefully; (2) figure out if there is content contradicted by or not presented in the article; (3) label the span as an intrinsic error if it misconstrues phrases or clauses from the article; (4) if the span does not belong to intrinsic errors, search within Wikipedia and determine whether the content in the span can be verified; (5) label it as world knowledge if it can be verified by Wikipedia, otherwise label it as an extrinsic error.

Example annotations 1

Article: Isis Academy in Oxford said it had rebranded as “Iffley Academy” to protect its “reputation, integrity and image”. **The name ‘Isis’ was originally chosen as the school is near to the section of the River Thames of the same name.** Formerly Iffley Mead School, it became Isis Academy in 2013. A statement issued by the school said it had changed name following “the unforeseen rise of ISIS (also known as ISIL and the Islamic State) and related global media coverage of the activities of the group”. “Our priority is to remove the detrimental impact which the name ‘Isis’ had on pupils, their families and our staff.” Last year a language school in the city removed Isis from its name for the same reason. The Isis is the name given to the part of the River Thames above Iffley Lock in Oxford. It is also the name of the goddess wife of the god Osiris in Egyptian beliefs.

Summary: A school that was named after the Islamic State (IS) militant group has changed its name.

Explanation: “was named after” is an intrinsic error contradicted by the article sentence in **bold**.

Example annotations 2

Article: Khalil Dale, 60, was abducted in **Quetta** in January 2012 and was found dead on a roadside a few months later. He had been beheaded. A note next to his body said he was killed because a ransom had not been paid. Mr Dale was born in York but lived in Dumfries. He spent 30 years working in countries including Somalia, Afghanistan and Iraq. An inquest into his death was held at Chesterfield Coroners Court because he is buried in Derbyshire. The court heard that the Muslim convert, who was formerly known as Kenneth, worked as a humanitarian assistance relief worker. Following his abduction, negotiations were undertaken by the International Committee of the Red Cross with the help of the UK government. His body was found on 29 April 2012. The inquest was told that he died as a result of decapitation. Senior coroner Dr Robert Hunter concluded that Mr Dale was unlawfully killed while providing international humanitarian assistance.

Summary: A British aid worker was unlawfully killed by **Islamist militants** in **Pakistan**, an inquest has heard.

Explanation: “*Islamist militant*” is an extrinsic error as it can not be found in or inferred from the article. The information is also not verifiable by Wikipedia. “*Pakistan*” is world knowledge as **Quetta** in the article is a city in Pakistan according to Wikipedia.

Figure 9: Guideline for our summary error annotation (§ 4).

In this study, you will evaluate 100 sets of summaries produced by four systems. For each set, its corresponding article and a baseline summary are shown before the four system summaries. The errors in the baseline summary are highlighted.

Please *first read the article and the baseline summary* and then *compare each system summary against the baseline summary* based on **informativeness** and **factual consistency**. In addition, please decide the **operations** made by the system to achieve better factual consistency.

For informativeness and factual consistency, you need to label whether the system summary is better or worse than the baseline summary. You can also label the system summary as tying with the baseline summary.

You need to consider two types of operations: **deletions** and **substitutions**. Please label the system summary as making deletions, substitutions, or both operations. Examples for the aspects and the operations are as follows.

Article: Alexys Brown, also known as Lexi, died at her home in Emmadale Close, Weymouth, on Thursday. An investigation is under way to discover how she became trapped. A post-mortem examination is due to be carried out this week. It was originally hoped the appeal would raise £2,000. Alison Record, who started the Just Giving appeal, said she was "heart broken" over the death. "Everybody by now has heard of the terrible tragedy the Brown family have suffered with the loss of their beautiful and beloved little girl Lexi," the appeal page reads. Many other comments have been posted on the appeal page. Steph Harris said: "Thinking of you all at this devastating time, fly high beautiful princess. Love Steph and family xxx" Lesley Andrews added: "No amount of money will take away the pain, but so much love comes with every penny. Take care. xx" Aster Group, the housing association responsible for managing the home, is assisting with the police investigation. The Health and Safety Executive (HSE) is also investigating. Dorset County Council said it had not installed the disabled lift at the property.

Baseline Summary: An appeal to raise 10,000 pounds for the family of a three-year-old girl who died after becoming trapped in a lift has raised more than 20,000 pounds.

Informativeness: Whether the summary captures salient content from the input article. Note that incorrect content should be considered as invalid.

Win. An appeal to raise money for the family of a three-year-old girl who died after getting stuck in a lift was originally hoped for raising £2,000. The target money of the appeal is a salient information.

Tie. An appeal to raise money for the family of a girl who died after getting stuck in a lift has raised more than £20,000. Compared to the baseline, missing incorrect information does not affect the informativeness.

Lose. An appeal to raise money for the family of a three-year-old girl has raised more than £20,000. This system summary does not mention the death of the girl, which is a salient content of the article.

Factual Consistency: Whether the summary is factually correctly based on the article and knowledge from Wikipedia.

Win. An appeal has been set up for the family of an eight-year-old girl who died after becoming trapped in a lift at her Dorset home. This system summary does not generate the incorrect numbers of money.

Tie. An appeal to raise 5,000 pounds for the family of a seven-year-old girl who died after becoming trapped in a lift has raised more than 20,000 pounds. This system summary makes similar errors to the baseline.

Lose. The family of an eight-year-old girl who died after becoming trapped in a lift at her Dorset home have set a fundraising target of 10,000 pounds. This system summary fabricates an event The family have set a fundraising target, which is more severe than errors of modifiers.

Deletion: The incorrect content in the baseline summary is deleted.

- An appeal for the family of a three-year-old girl who died after becoming trapped in a lift has raised more than 20,000 pounds. The error "10,000 pounds" is deleted.

Substitution: The incorrect content in the baseline summary is replaced with correct one.

- An appeal to raise 2,000 pounds for the family of a three-year-old girl who died after becoming trapped in a lift has raised more than 20,000 pounds. The error "10,000 pounds" is substituted with "2,000 pounds", which is the correct information.

Figure 10: Guideline for our human evaluation (§ 7.2).

Example 1

CNN/DM Article: At the grand old age of 75, Jack Nicklaus is still capable of hitting aces. The Golden Bear added another magic moment to his storied career at Augusta National in the Par-3 Contest. Stepping up to the tee on the 130-yard fourth, the greatest golfer of all time saw his shot sail beyond the flag before spinning back into the hole. Jack Nicklaus gave the crowd something to cheer with a hole in one on the fourth during the Par-3 Contest. Nicklaus holds up his ball to an adoring crowd as Gary Player (left) and Ben Crenshaw salute the great. Crenshaw and Nicklaus fist pump following his ace on the 130-yard hole at Augusta National. Nicklaus was playing alongside Gary Player and Ben Crenshaw as the first roar of Masters week filled the skies in a sunswept Georgia. 'I had an interview this morning, and I said "well, all I've got to do is go out and win the Par-3 and make a hole and one", and I make a hole in one,' Nicklaus said afterwards. . The ace was Nicklaus' first on the hallowed turf of Augusta, but the veteran saw his achievement eclipsed by that of Camilo Villegas on a day which saw five hole-in-ones, levelling the record from 2002. The Colombian followed Nicklaus with an ace on the fourth and also managed one on the eighth, but he couldn't quite land the title, losing out to world No 74 Kevin Streelman in a play-off. Nicklaus became the youngest person to wear a green jacket in 1963, and collected his sixth in 1986. He is one of five men to complete the career grand slam, an accolade which favourite Rory McIlroy can achieve if he wins his third major in succession.

Cross-entropy: Jack Nicklaus hit a hole in one on the fourth during the Par-3 Contest. The 75-year-old was playing alongside Gary Player and Ben Crenshaw. Camilo Villegas and Kevin Streelman also hit hole-in-ones on the day.

Entailment reranking: Jack Nicklaus hit a hole in one on the fourth during the Par-3 Contest. The 75-year-old was playing alongside Gary Player and Ben Crenshaw. Camilo Villegas and Kevin Streelman also hit hole-in-ones on day one.

Unlikelihood training: Jack Nicklaus hit a hole in one on the fourth during the Par-3 Contest. The 75-year-old was playing alongside Gary Player and Ben Crenshaw. Camilo Villegas followed Nicklaus with a hole-in-one on the same hole. Nicklaus is one of five men to complete the career grand slam.

Contrastive learning: Jack Nicklaus hit a hole in one on the fourth during the Par-3 Contest. The 75-year-old was playing alongside Gary Player and Ben Crenshaw. Camilo Villegas followed Nicklaus with a hole-in-one on the eighth.

Example 2

XSum Article: The British supermarket giant, the world's third largest retailer, has been looking to dispose of its loss-making US food chain for some time. Around 4,000 of Fresh & Easy's 5,000 employees are to transfer to the new business as part of the deal. Tesco said in a statement that this is an "orderly and efficient exit from the US market". Under the terms of the deal, Tesco will loan the new business around £80m. Fresh & Easy, which operates 200 stores across California, Nevada and Arizona, has lost money since opening in 2007. Tesco did not specify which stores are being transferred to Yucaipa, but it is understood the remaining 50 shops will be closed. It will cost Tesco £150m to dispose of the Fresh & Easy brand, but part of that is a loan which may be paid back. Also, if Yucaipa is successful and the business becomes profitable, Tesco has the option to buy a stake. Yucaipa is a Los Angeles based private equity company focused on supermarket and restaurant businesses and is run by billionaire investor Ron Burkle. In April of this year, Tesco reported its first fall in annual profits for 20 years.

Cross-entropy: Tesco has agreed to sell its US food chain Fresh & Easy to Yucaipa in a deal worth about 150m, the BBC understands.

Entailment reranking: Tesco is to sell its US food chain Fresh & Easy to Yucaipa for an undisclosed sum.

Unlikelihood training: Tesco has agreed to sell its US food business, Fresh & Easy, to a private equity firm called Yucaipa for an undisclosed sum.

Contrastive learning: Tesco has agreed to sell its US food chain, Fresh & Easy, to Yucaipa.

Example 3

XSum Article: Neil Aspin's promotion-chasing hosts have not lost in nine National League matches while Adam Lockwood's side are unbeaten in five. Guiseley went ahead on 15 minutes against the run of play when a throw-in found James Hurst who squared to Jake Lawlor to stroke into an empty net. Gateshead defender Liam Hogan superbly blocked Jordan Preston's effort and Guiseley keeper Jonny Maxted then saved well from Wesley York's shot just before the break. The hosts, who started the second half well, levelled on 62 minutes when a slip by half-time substitute Derek Asamoah let York curl sweetly into the top-right corner from the edge of the box. Match report supplied by the Press Association. Match ends, Gateshead 1, Guiseley 1. Second Half ends, Gateshead 1, Guiseley 1. Substitution, Guiseley. Michael Rankine replaces Jordan Preston. Substitution, Gateshead. Luke Hannant replaces Gus Mafuta. Gus Mafuta (Gateshead) is shown the yellow card. Substitution, Guiseley. Adam Boyes replaces Jake Cassidy. Goal! Gateshead 1, Guiseley 1. Wes York (Gateshead). Substitution, Guiseley. Derek Asamoah replaces Kevan Hurst. Second Half begins Gateshead 0, Guiseley 1. First Half ends, Gateshead 0, Guiseley 1. Simon Walton (Guiseley) is shown the yellow card. Goal! Gateshead 0, Guiseley 1. Jake Lawlor (Guiseley). First Half begins. Lineups are announced and players are warming up.

Cross-entropy: Gateshead and Guiseley maintained their unbeaten starts to the season with a 1-1 draw at the International Stadium.

Entailment reranking: Gateshead and Guiseley shared the spoils after a 1-1 draw at the International Stadium.

Unlikelihood training: Gateshead and Guiseley shared the spoils after a goalless draw in the National League.

Contrastive learning: Gateshead and Guiseley shared the spoils after a 1-1 draw at Gateshead.

Figure 11: Sample generated summaries by fine-tuned BART models. Intrinsic errors are highlighted in red and extrinsic errors are in blue. MASKENT and SYSLOWCON are used for negative sample construction with unlikelihood training and contrastive learning.