

# Learning with Rejection for Abstractive Text Summarization

Meng Cao    Yue Dong    Jingyi He    Jackie Chi Kit Cheung

Mila / McGill University, Montréal, QC, Canada

{meng.cao@mail, yue.dong2@mail, jingyi.he@mail, jcheung@cs}.mcgill.ca

## Abstract

State-of-the-art abstractive summarization systems frequently hallucinate content that is not supported by the source document, mainly due to noise in the training dataset. Existing methods opt to drop the noisy samples or tokens from the training set entirely, reducing the effective training set size and creating an artificial propensity to copy words from the source. In this work, we propose a training objective for abstractive summarization based on *rejection learning*, in which the model learns whether or not to reject potentially noisy tokens. We further propose a regularized decoding objective that penalizes non-factual candidate summaries during inference by using the rejection probability learned during training. We show that our method considerably improves the factuality of generated summaries in automatic and human evaluations when compared to five baseline models, and that it does so while increasing the abstractiveness of the generated summaries.<sup>1</sup>

## 1 Introduction

Abstractive summarization has progressed significantly in recent years, as evidenced by automatic evaluation scores like ROUGE (Lin, 2004). However, existing abstractive summarization systems are prone to produce non-factual summaries, which contain information that is not supported by the source (Kryscinski et al., 2019; Maynez et al., 2020; Durmus et al., 2020; Zhao et al., 2020; Kryscinski et al., 2020; Narayan et al., 2021). According to recent studies, one cause of this phenomenon is noise in the training set, such as unsupported facts in the references (Kang and Hashimoto, 2020; Goyal and Durrett, 2021; Raunak et al., 2021). Specifically, the widely-used negative log likelihood objective aims to match the underlying data distribution. If summarization models are trained on noisy datasets

containing erroneous references, they will tend to imitate and even amplify undesired generations (Kang and Hashimoto, 2020).

Different approaches have been proposed to reduce non-factual hallucinations. Falke et al. (2019); Zhao et al. (2020) employ re-ranking based approaches to up-rank candidate summaries based on their factuality scores from out-of-the-box factual consistency evaluation systems. Kang and Hashimoto (2020); Goyal and Durrett (2021); Goyal et al. (2022) use loss truncation-based method that drops high loss sentences/tokens during training. Cao et al. (2022) propose reinforcement learning (RL)-based methods that penalize hallucinated tokens with negative rewards.

One instance of the reference summary containing unsupported information is shown in Table 1. We can see that the reference summary contains details on the precise amount of government funding (£12.8m), yet this amount is absent from the source document. This kind of noise in the reference summary is highly deceptive and could cause the summarization model to hallucinate. Furthermore, the noise can “break” the beam search process since the phrases associated with the noise (e.g., *has been awarded [...]*) are decoded with high probability during inference. Therefore, we think it is crucial for the summarization model to learn what information is unknown and reject contexts in which hallucination is likely to happen.

In this work, we propose a training objective that enables the summarization model to recognize and reject unsupported target text spans. Inspired by rejection learning for image classification (Grandvalet et al., 2008; Cortes et al., 2016; Thulasidasan et al., 2019), we introduce a new rejection class at training time that offers the model an option to reject confusing tokens rather than assigning high probability to all reference tokens. Besides, our method does not require labeling or data pre-processing. The learning is based on the training

<sup>1</sup><https://github.com/mcao516/rej-summ>

<p><b>Source:</b>  (...) Chris Cox, the university’s director of development, said the research centre would translate discoveries made in the laboratory into new treatments. He said it would house 150 additional researchers who will be developing new ideas and treatments. Research will focus on radiation therapy, lung cancer, women’s cancers, melanoma and haematological oncology. The centre is the result of a partnership between The University of Manchester, The Christie NHS Foundation Trust and Cancer Research UK. The Christie’s chief executive Caroline Shaw said the funding would “help facilitate groundbreaking research right here in Manchester”. (...)</p>
<p><b>Reference:</b>  Plans for a new cancer research centre in Manchester have received a £12.8m funding boost from the government.</p>
<p><b>BART (beam size=6):</b>  The University of Manchester has been awarded £3.5m to help fund a new cancer research centre.  The University of Manchester has been awarded £1.3m to help fund a new cancer research centre.  The University of Manchester has been awarded £1.2m to help fund a new cancer research centre.  The University of Manchester has been awarded £1.5m to help fund a new cancer research centre.  The University of Manchester has been awarded £5m to help fund a new cancer research centre.  The University of Manchester has been awarded £4.5m to help fund a new cancer research centre.</p>
<p><b>BART trained with rejection (greedy decoding):</b>  The University of Manchester has been awarded £&lt;REJ&gt;m to help fund its new cancer research centre.</p>
<p><b>BART trained with rejection (regularized decoding):</b>  Details of a new cancer research centre at the University of Manchester have been revealed.</p>

Table 1: Example of a BBC article from XSUM test set. In this example, the source document does not have information on the specific value of the government funding. However, BART model (Lewis et al., 2020) trained using the MLE objective hallucinates different funding values in its generation (highlighted). In the bottom, we show that the BART model trained using the rejection objective (see Section 3.1) work successfully reject the hallucinated token during inference. After applying the regularized decoding (Section 3.2), the model generates a factual summary without hallucinated information.

dynamics of abstractive summarization systems where clean samples are learned quickly and noisy samples have relatively high loss (Goyal et al., 2022; Kang and Hashimoto, 2020). We demonstrate that the rejection learning objective allows the summarization model to learn features that are indicative of unsupported information. As shown in the example in Table 1, the model trained with the rejection objective correctly rejects after the phrase “has been awarded” since such information is absent from the source. According to analysis on human-annotated factuality datasets, our model can accurately reject 77.6% of non-factual hallucinations generated by state-of-the-art abstractive summarization model.

Furthermore, we also provide a regularized decoding objective for summarization models trained with the rejection training objective. We add a regularization term to the original decoding objective to penalize candidates using the rejection probability learned during training. We demonstrate that the regularized decoding objective, combined with rejection training, can vastly improve the factuality of summarization models using both automatic and human evaluation. Compared with the MLE baseline, our method improves the sen-

tence factuality rate by 25.4% (34.6% → 43.4%, measured by DAE (Goyal and Durrett, 2021)) and 41.7% (22.8% → 32.3%, measured by FACTCC (Kryscinski et al., 2020)). A recent study has shown that part of the improvement of faithfulness brought by prior methods stems from an increased level of extractiveness of the model outputs (Ladhak et al., 2022). We demonstrate that our method improves model’s factuality while simultaneously increasing the abstractiveness of the generated summaries.

Our main contributions in this work are: (i) We introduce a rejection learning objective for training abstractive summarization models. We demonstrate how the additional training objective enables the summarization model to precisely reject hallucinated tokens during inference (Section 5.1). (ii) We propose a novel factuality-aware regularized decoding objective based on the rejection probability learned during training. Our method significantly improves the factuality of generated summaries in automatic and human evaluations when compared to five baseline models. (iii) We demonstrate that, in comparison to the baselines, the summaries produced by our method also exhibit a greater level of abstractiveness with more novel  $n$ -grams and less copying.

## 2 Related Work

With the advent of transformer-based summarization models (Lewis et al., 2020; Zhang et al., 2020; Qi et al., 2020; Liu and Liu, 2021; Liu et al., 2022), current system-generated summaries are often fluent and salient. However, detailed analysis of these systems’ outputs reveals that the majority of the generation contains hallucinations (Kryscinski et al., 2020; Pagnoni et al., 2021). Moreover, annotations from Maynez et al. (2020) suggest that the majority of these hallucinations are extrinsic; namely, they contain information that is not directly inferable from the source.

Many prior studies consider extrinsic hallucinations generated by models as undesirable, and propose different methods to reduce these hallucinations (Falke et al., 2019; Dong et al., 2020; Zhu et al., 2021; Cao et al., 2020; Pagnoni et al., 2021; Nan et al., 2021). On the other hand, some prior work (Maynez et al., 2020; Cao et al., 2022; Dong et al., 2022) suggest many of these hallucinations are actually factual and can be useful to provide additional background information.

The most similar work to ours is the approach of using loss truncation to reduce noise in the training stage. Kang and Hashimoto (2020) and Goyal et al. (2022) propose sentence and token-level loss truncation for hallucination reduction. These strategies are based on the assumption that noisy examples are difficult to learn, resulting in high training loss. Simply removing these difficult-to-learn high log loss samples might reduce model hallucination. However, dropping too many samples could make the training process inefficient. It also makes the summarization model more extractive. Dropping high loss token might not be effective either since the model does not learn the right token to generate during decoding.

We instead use the so-called *rejection* (or abstention) approach to combat noise in training. The loss abstention approach was proposed in Thulasidasan et al. (2019) for classification tasks, where a classifier is trained to abstain on mislabeled images during training. Inspired by this idea, we introduce a rejection loss for generation tasks (Section 3.1). The advantage of loss abstention over loss truncation is its ability to learn text features that are likely to lead to hallucination rather than simply ignore them. Such features can be used during inference to avoid hallucinations.

## 3 Method

### 3.1 Learning with Rejection

**Problem Formulation** We consider abstractive summarization as a conditional language generation task. Given a source document  $\mathbf{x} = (x_1, \dots, x_L)$  with  $L$  tokens, the task is to learn a probabilistic model  $\hat{p}_\theta(\mathbf{y}|\mathbf{x})$  that generates a summary  $\mathbf{y} = (y_1, \dots, y_{|\mathbf{y}|})$ , where  $y_i$  comes from a pre-defined vocabulary  $\mathcal{V}$ .  $\theta$  denotes the parameters of the model.  $\hat{p}_\theta$  is modeled by decomposing the probability of a sequence into conditional probabilities of each token given the previous context:  $\hat{p}_\theta(\mathbf{y}) = \prod_{t=1}^{|\mathbf{y}|} \hat{p}_\theta(y_t | \mathbf{y}_{<t}, \mathbf{x})$ . Usually,  $\hat{p}_\theta$  is trained using the maximum likelihood estimation (MLE) objective, which aims to maximize the likelihood of the human-written reference summaries:

$$\theta^* = \arg \max_{\theta} \sum_{\mathbf{x}, \mathbf{y}^* \sim \mathcal{D}} \log \hat{p}_\theta(\mathbf{y}^* | \mathbf{x}; \theta), \quad (1)$$

where  $\mathcal{D}$  denotes the training set and  $\mathbf{y}^*$  is the reference summary. For a given document  $\mathbf{x}$  and reference summary  $\mathbf{y}^*$ , the MLE objective is equivalent to minimizing the sum of negative log-likelihoods of the ground-truth tokens:

$$\begin{aligned} \mathcal{L}(\theta) &= -\log \hat{p}_\theta(\mathbf{y}^* | \mathbf{x}; \theta) \\ &= -\sum_{t=1}^{|\mathbf{y}^*|} \log \hat{p}_\theta(y_t^* | \mathbf{x}, \mathbf{y}_{<t}^*; \theta), \end{aligned} \quad (2)$$

where  $y_t^* \in \mathcal{V}$  is the ground-truth token from the reference summary. Equation 2 is also equivalent to the cross-entropy loss objective where the true distribution is a one-hot distribution.

For the sake of notational brevity, we define  $p_t^i = \hat{p}_\theta(y_t^i | \mathbf{y}_{<t}^*, \mathbf{x})$ , ( $0 \leq i \leq |\mathcal{V}| - 1$ ) as the conditional probability of the  $i$ -th vocabulary token at the  $t$ -th time-step given the ground-truth prefix.

**The Rejection Loss** The MLE objective seeks to maximize the likelihood of all training samples. Knowing that the underlying training data is noisy, we want the model to automatically recognize unsupported text spans in references and reject them during training (Thulasidasan et al., 2019). To do this, we can add an additional rejection class  $r$  out of the original  $|\mathcal{V}|$  classes (i.e., tokens). Let  $p_t^r = \hat{p}_\theta(r | \mathbf{x}, \mathbf{y}_{<t}^*)$  denote the probability of rejection at the  $t$ -th step, the modified loss function is

$$\mathcal{L}(\theta) = - \sum_{t=1}^{|\mathbf{y}|} \left[ (1 - p_t^r) \log \frac{\hat{p}_\theta(y_t^* | \mathbf{x}, \mathbf{y}_{<t}^*; \theta)}{1 - p_t^r} - \alpha \cdot \mathcal{R}_T(p_t^r) \right] \quad (3)$$

where  $\mathcal{R}_T(p) = \log(1 - p)$  is the regularization function for rejection. In the first term, increasing either  $\hat{p}_\theta(y^* | \mathbf{x}, \mathbf{y}_{<t}^*; \theta)$  or  $p_t^r$  will minimize the loss function. During training, when the ground-truth token contains misannotations or unsupported information, term  $\log \hat{p}_\theta(y_t^* | \mathbf{x}, \mathbf{y}_{<t}^*; \theta)$  will be abnormally high (Kang and Hashimoto, 2020). In this situation, minimizing the log loss results in a suboptimal model that hallucinates. By introducing the rejection class, the model has the alternative of increasing the probability of rejection in order to minimize the total loss. In an extreme case where  $p_t^r = 1$ , the first loss term will be zero. To prevent the model from deteriorating into the extreme case, we add a second regularization term that penalizes large rejection probability. During training, we set  $p_t^r = 0$  to fully recover the original cross-entropy loss (Equation 2) for the first  $m$  steps. This allows the model to quickly absorb easy patterns.

### 3.2 Decoding with Rejection

Let’s consider how to do decoding for a model trained with the rejection loss objective above. At inference time, the rejection probability reflects the degree of uncertainty the model has about its prediction. When the rejection class is assigned a high probability at time step  $t$ , the generated content is likely to be hallucinated and non-factual at this step (see Section 5.1). Therefore, we can make use of the rejection probability to penalize non-factual candidate summaries. We consider the following regularized decoding objective:

$$\mathbf{y}^* = \arg \max_{\mathbf{y} \in \mathcal{Y}} \left( \log \hat{p}_\theta(\mathbf{y} | \mathbf{x}; \theta) - \lambda \cdot \mathcal{R}_D(\mathbf{y}) \right) \quad (4)$$

where

$$\mathcal{R}_D(\mathbf{y}) = \frac{1}{|\mathbf{y}|} \sum_{t=1}^{|\mathbf{y}|} \log^k \frac{1}{1 - p_t^r} \quad (5)$$

and  $\mathcal{Y}$  is the hypothesis space.  $k$  is the exponent for the log that can amplify the effect of the regularization term when  $p_t^r$  is large. We set  $k = 1$  in the following experiments. We also consider a max regularizer that finds the token with the highest rejection probability in the generated summary:

$$\mathcal{R}_{D,\max}(\mathbf{y}) = \max_{t=1}^{|\mathbf{y}|} \log^k \frac{1}{1 - p_t^r} \quad (6)$$

During inference, the rejection token is only used for regularization and will not be generated. At each decoding step  $t$ , we only select tokens from the original vocabulary  $\mathcal{V}$ . The probability of each token  $p_t^i$  is divided by  $1 - p_t^r$  to make sure the model outputs a proper probability distribution over vocabulary words.

## 4 Experiments

### 4.1 Experimental Settings

**Datasets** We evaluated our method on the XSUM (Narayan et al., 2018) dataset. XSUM contains 226,711 online British Broadcasting Corporation (BBC) articles covering domains like news, politics, sports, business and more. Each article is paired with a single sentence summary written by the BBC journalists. The dataset is split into three subsets: training (204,045, 90%), validation (11,332, 5%), and test (11,334, 5%) sets. We choose XSUM because it is more abstractive than other summarization datasets. It also relatively noisy due to the way it is created (Maynez et al., 2020).

For fine-grained entity-level factuality analysis (see Section 5.1), we use the XENT dataset created by Cao et al. (2022). XENT contains 800 article and summary pairs where the articles are sampled from XSUM test set and the summaries are generated using a fine-tuned BART model (Lewis et al., 2020). In XENT, each entity in the summary is annotated with three labels: non-hallucinated, factual hallucination, and non-factual hallucination. Factual hallucinations refer to an entity that is verifiable by world knowledge but not directly inferable from the source text. We evaluate our model on the test set of XENT which contains 240 summaries and 835 entities.

**Baselines** We compare our method with five baselines: (1) State-of-the-art abstractive summarization model BART (Lewis et al., 2020). (2) The loss truncation method proposed by Kang and Hashimoto (2020). (3) The token-level loss truncation method proposed by Goyal et al. (2022). (4) GOLD (generation by off-policy learning from demonstrations), an off-policy RL algorithm proposed by Pang and He (2021). (5) A factuality-aware off-policy RL algorithm for text summarization proposed by (Cao et al., 2022).

BART is a frequently used transformer-based model for summarization. It employs the encoder-decoder architecture, which is essentially a denoising autoencoder that reconstructs the original text with injected noise, such as token deletion, replacement, masking and sentence permutation. The loss truncation method proposed by Kang and Hashimoto (2020) adaptively discards high log loss examples as a way to reduce hallucination in generation. Similarly, the token-level loss truncation in Goyal et al. (2022) removes high-loss and low-loss tokens that are challenging or too easy to learn. GOLD (Pang and He, 2021) is proposed to alleviate the discrepancies between training and evaluation objectives in MLE (e.g, over-generalization, exposure bias). Pang and He (2021) shows that GOLD outperforms BART on text summarization in terms of ROUGE scores and human evaluation. Cao et al. (2022) augmented GOLD with factuality rewards generated using their entity-level factuality classifier.

**Implementation Details** In the following experiments, we use pre-trained BART-LARGE as the backbone for our model. The size of hidden layers is set to 1024 and the max sequence length is set to 1024 for both the encoder and decoder modules. We fine-tuned our model on the XSUM training set for 20,000 steps with a warm-up period of 500 steps. We use the Adam optimizer (Kingma and Ba, 2014) with  $\epsilon = 1e-8$  and an initial learning rate  $3e-5$ . Both dropout and attention dropout are applied in each layer with a dropout rate set to 0.1. The  $\alpha$  in Equation 3 is set to 1.0 unless otherwise specified. We use the sum regularizer defined in Equation 5 for our main results. All experiments are conducted on 4 Tesla V100 GPUs with 32GB of memory.

**Automated Evaluation Metrics** For sentence-level faithfulness evaluation, we use FACTCC (Kryscinski et al., 2020) and DAE (Goyal and Durrett, 2021). FACTCC is a sentence-level faithfulness classifier trained on synthetic dataset that is created based on analysis of errors made by state-of-the-art summarization models. DAE is trained on fine-grained, human annotations of factual errors made by abstractive summarization models. Given a document and summary pair  $(x, y)$ , both methods predict a label  $y \in \{0, 1\}$  where 0 indicates the summary being non-faithful and 1 being faithful. The percentage of sentences that are predicted

as faithful by the classifier is used as the sentence factuality rate.

For entity-level factuality evaluation, we use ENTFA proposed by Cao et al. (2022). ENTFA is a  $k$ -NN model that uses prior and posterior probabilities as features to predict the factuality and hallucination status of the entities in a given summary. We use the percentage of entities in the output summaries that are predicted as *non-factual hallucinations* as the entity hallucination rate in Table 2. For all three evaluation methods, we use the evaluation scripts and parameters provided in the papers. To measure the impact of each method on summary abstractiveness, we calculate the percentage of novel  $n$ -grams and *coverage*, the percentage of words in the summary that are part of an extractive fragment of the source document (Grusky et al., 2018)

## 4.2 Factuality Evaluation

In this section, we conduct automatic and human evaluation on summary factuality.

### 4.2.1 Automatic Evaluation Results

The findings of the factuality evaluation on the XSUM test set are presented in Table 2. We evaluate the effect of rejection loss without regularized decoding by setting  $\lambda = 0$ . We make the following observations:

(1) **In the evaluation of sentence-level summary factuality, our proposed method – the rejection loss objective combined with regularized decoding – outperforms all baseline methods.** Compared with MLE training, we observe a 41.7% improvement (22.8%  $\rightarrow$  32.3%) on FACTCC score and 25.4% improvement (34.6%  $\rightarrow$  43.4%) on DAE score (when  $\lambda = 5.0$ ). Our method also outperforms MLE and the loss-truncation baseline in terms of entity-level factuality evaluation, but it underperforms the RL+FR baseline. We speculate that this is because the RL method is specifically optimized toward the entity-level factuality rewards generated using ENTFA.

(2) **We observe (almost) consistent improvement of factuality scores as we increase the value of  $\lambda$ .** This is in line with what we expected, as candidate summaries with uncertainty will be penalized more severely with a larger  $\lambda$ . An intriguing finding is that summary factuality is substantially increased when the rejection loss objective is used alone without regularized decoding. Compared to the MLE baseline, the summariza-

Model	Sentence Factuality Rate		Entity Hal. Rate	Novel $n$ -gram		ROUGE-1
	FactCC $\uparrow$	DAE $\uparrow$	ENTFA $\downarrow$	Unigram $\uparrow$	Bigram $\uparrow$	
Sentence-level Loss Truncation						
MLE <sup>‡</sup>	22.8%	34.6%	17.2%	27.9%	74.5%	45.1
+SLT <sup>†</sup> ( $c = 0.3$ )	$\uparrow$ 4.8% 23.9%	$\uparrow$ 5.2% 36.4%	$\downarrow$ 6.4% 16.1%	$\downarrow$ 3.9% 26.8%	$\downarrow$ 1.5% 73.4%	$\downarrow$ 2.2% 44.1
+SLT <sup>†</sup> ( $c = 0.5$ )	$\uparrow$ 7.5% 24.5%	$\uparrow$ 7.7% 37.5%	$\downarrow$ 4.7% 16.4%	$\downarrow$ 4.3% 26.7%	$\downarrow$ 1.9% 73.1%	$\downarrow$ 3.1% 43.7
+SLT <sup>†</sup> ( $c = 0.7$ )	$\uparrow$ 11.8% 25.5%	$\uparrow$ 12.1% 38.8%	$\downarrow$ 7.6% 15.9%	$\downarrow$ 4.7% 26.6%	$\downarrow$ 1.7% 73.2%	$\downarrow$ 5.3% 42.7
Token-level Loss Truncation						
+TLT <sup>†</sup> ( $c = 0.3$ )	$\uparrow$ 11.4% 25.4%	$\uparrow$ 0.3% 34.7%	$\downarrow$ 3.5% 16.6%	$\uparrow$ 0.4% 28.0%	$\uparrow$ 1.6% 75.7%	$\downarrow$ 4.0% 43.3
+TLT <sup>‡</sup> ( $c = 0.5$ )	$\uparrow$ 3.0% 23.5%	$\uparrow$ 5.8% 36.6%	$\uparrow$ 3.5% 17.8%	$\downarrow$ 9.3% 25.3%	$\downarrow$ 1.6% 73.3%	$\downarrow$ 6.4% 42.2
+TLT <sup>†</sup> ( $c = 0.7$ )	$\uparrow$ 19.3% 27.2%	$\uparrow$ 20.5% 41.7%	$\uparrow$ 15.7% 19.9%	$\downarrow$ 14.0% 24.0%	$\downarrow$ 4.3% 71.3%	$\downarrow$ 21.3% 35.5
RL with Factuality Rewards						
RL*	22.1%	33.3%	17.8%	28.1%	74.7%	45.8
+FR* ( $r = -2.0$ )	$\uparrow$ 4.5% 23.1%	$\uparrow$ 12.0% 37.3%	$\downarrow$ 44.4% 9.9%	$\downarrow$ 1.4% 27.7%	$\uparrow$ 0.3% 74.9%	$\downarrow$ 2.6% 44.6
+FR* ( $r = -4.0$ )	$\uparrow$ 8.6% 24.0%	$\uparrow$ 22.5% 40.8%	$\downarrow$ 57.9% 7.5%	$\downarrow$ 4.3% 26.9%	$\downarrow$ 0.8% 74.1%	$\downarrow$ 6.1% 43.0
Rejection Learning & Regularized Decoding (Ours)						
REJL ( $\lambda = 0.0$ )	26.1%	39.1%	13.5%	28.1%	75.6%	45.1
+RD ( $\lambda = 1.0$ )	$\uparrow$ 7.3% 28.0%	$\uparrow$ 5.9% 41.4%	$\downarrow$ 10.4% 12.1%	$\downarrow$ 1.1% 27.8%	$\uparrow$ 0.3% 75.8%	$\downarrow$ 1.6% 44.4
+RD ( $\lambda = 2.0$ )	$\uparrow$ 12.3% 29.3%	$\uparrow$ 9.0% <b>42.6%</b>	$\downarrow$ 14.8% 11.5%	$\uparrow$ 0.0% 28.1%	$\uparrow$ 0.8% 76.2%	$\downarrow$ 3.8% 43.4
+RD ( $\lambda = 3.0$ )	$\uparrow$ 16.5% <b>30.4%</b>	$\uparrow$ 8.2% 42.3%	$\downarrow$ 11.9% 11.9%	$\uparrow$ 1.1% <b>28.4%</b>	$\uparrow$ 1.1% <b>76.4%</b>	$\downarrow$ 5.8% 42.5
+RD ( $\lambda = 5.0$ )	$\uparrow$ 23.8% <b>32.3%</b>	$\uparrow$ 11.0% <b>43.4%</b>	$\downarrow$ 20.0% 10.8%	$\uparrow$ 3.2% <b>29.0%</b>	$\uparrow$ 1.7% <b>76.9%</b>	$\downarrow$ 8.2% 41.4

Table 2: Comparison between our proposed method and five baseline models. Numbers that are highlighted **blue** represents the percentage of improvement and **orange** denotes the percentage of performance drop. All models are trained on the XSUM training set and evaluated on 3K samples from the XSUM test set. For MLE, we use the BART (Lewis et al., 2020) model. SLT and TLT stands for sentence- and token-level loss truncation respectively.  $c$  is the percentage of data to be dropped. In this experiment, all summaries are generated using beam search with beam size equals to 6. \*: results reported in the original papers. ‡: we run evaluation ourselves using checkpoints the authors provided. †: model checkpoints are not available and we train the models ourselves using the author’s code. Training hyper-parameters are based on the original paper.

tion model trained with the rejection loss objective (without regularized decoding) achieves the same ROUGE-1 score 45.1 but significantly higher factuality scores (FACTCC: 22.8%  $\rightarrow$  26.1%; DAE: 34.6%  $\rightarrow$  39.1%; ENTFA: 15.5%  $\rightarrow$  13.5%) and has more novel  $n$ -grams. These findings imply that when the underlying dataset is noisy, the rejection loss objective performs generally better than the native cross-entropy objective for text summarization. It is unclear whether this generalizes to other text generation tasks, such as knowledge-grounded dialogue generation, where training datasets frequently contain noise (Dziri et al., 2022). This will be left for future work.

(3) **The resulting summaries become more abstractive when  $\lambda$  increases.** This is supported by the observation from Table 2, where the generated summaries contain more novel  $n$ -grams as  $\lambda$  increases. Thus, our method can simultaneously increase the abstractiveness and factuality of the generated summaries. In contrast, the proportion of novel  $n$ -grams falls if we try to make the generated

summaries more factual in loss truncation and RL methods. The drop of novel  $n$ -grams suggests that factuality improvement brought by the baselines may come from increased level of extractiveness. This is consistent with the findings of Ladhak et al. (2022).

We also plot the Faithfulness-Abstractiveness Tradeoff curve following the recent work by Ladhak et al. (2022), as seen in Figure 1. We adopt *coverage* defined in Grusky et al. (2018) to measure the level of extractiveness: a summary with higher coverage is considered more extractive. Faithfulness is measured using DAE score. In Figure 1, the abstractiveness and factuality of our model are positively correlated. The loss truncation approach, in contrast, exhibits a negative correlation. This result further confirms our point.

(4) **We observe a decline in ROUGE scores as the summary factuality increases.** According to our calculation, the average length of the generated summaries for different  $\lambda$  remains constant (around eighteen words). Hence, it is unlikely that sum-

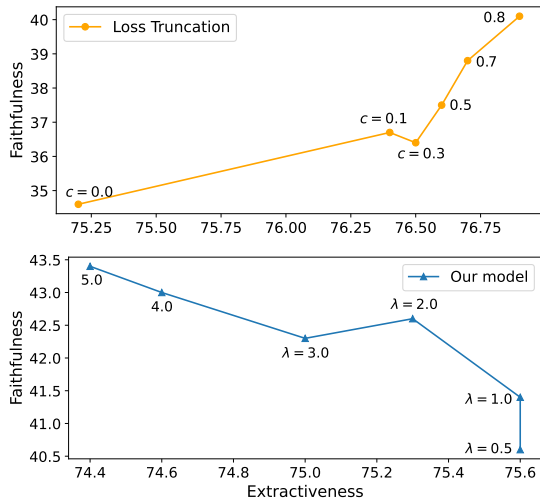


Figure 1: Faithfulness-Abstractiveness trade-off curves for our model and the loss truncation baseline (Kang and Hashimoto, 2020). Extractiveness is measured using *coverage* defined in (Grusky et al., 2018). Faithfulness is measured using the DAE score (Goyal and Durrett, 2021).

	vs. MLE	vs. Loss Trunc.
Win	39.3%	32.0%
Lose	23.3%	26.0%
Equally Factual	17.3%	22.7%
Equally Non-factual	20.0%	19.3%

Table 3: Human comparisons on 150 randomly selected articles from XSUM test set. Win: percentage of outputs from our model that are preferred by the annotators over the baseline’s outputs, given the same source document.

mary length played a role in the declining ROUGE score. Instead, we speculate that the decline in the ROUGE scores is caused by the increase in abstractiveness. For instance, the regularized decoded summary in Table 1 is highly abstract. It makes no mention about the funding, as the value of the fund is not specified in the source. In contrast, the summary generated by BART matches the reference more closely in terms of  $n$ -gram overlap even though it contains non-factual hallucinations. In light of this example, it appears that the MLE objective may encourage the summarization model to hallucinate in order to boost ROUGE scores.

#### 4.2.2 Human Evaluation

In addition to automatic evaluation, we also conducted a human evaluation, comparing our model ( $\lambda=2.0$ ) to the MLE and the token-level loss truncation  $c = 0.5$  baseline. We showed participants

The University of Manchester has been awarded ____		
<b>MLE:</b>	£3.5m to help fund a new research centre.	
<b>Ours:</b>	£3.5m (...)	% same ent.
	£1.5m (...)	% different hal.
	£12.8m (...)	% factual
	a large amount of money to (...)	% remove ent.
	<REJ>	% rejection token

Table 4: An example of how we evaluate our rejection model on XENT (Cao et al., 2022). In this example, £3.5m is generated by BART and labeled as non-factual hallucination by human annotator. We provide the same context to our model and analyze the distribution of five different types of outputs from our model. We did this analysis for both factual and non-factual entities in XENT.

the output of our model and a comparison model, and asked them to make a pairwise preference judgment between the two summaries based on their faithfulness.

Specifically, we randomly sampled 150 articles from the XSUM test set for summarization. To make the evaluation process more effective, we filtered out samples where the generated summary pair has more than 70% unigram overlap. We presented the two summaries in random order and ask participants to carefully consider the faithfulness of each summary. We then asked them to select the summary that they prefer, or to say they are equally factual/non-factual. A more detailed annotation guide can be found in Appendix A.

Table 3 shows the human evaluation results. From the table, we can find that summaries generated by our model are preferred by the human annotators over two baselines.<sup>2</sup>

## 5 Analysis

### 5.1 The Rejection Probability of Hallucinated Tokens

In this section, we analyze how well the rejection class calibrates with hallucinated entities. We use the test set of XENT dataset, which contains 240 machine generated summaries with entity-level hallucination annotations. Each entity is annotated by human experts for its hallucination and factuality status.

For each annotated entity in XENT, we give our model the same context as input and analyze the model’s output. For non-factual entities, we expect

<sup>2</sup>Discussion of inter-annotator agreement is included in Appendix B.

	Factual Entity		Non-factual
	Non-hal.	Factual Hal.	
Same Ent.	81.8%	64.4%	17.2%
Different Hal.	0.0%	2.5%	0.7%
Factual Ent.	0.2%	0.0%	0.0%
Remove Ent.	5.1%	7.6%	4.5%
Rejection	13.0%	25.4%	77.6%
Total	583	118	134

Table 5: Distribution of five different types of outputs for different entities. Non-hal. stands for non-hallucinated factual entity. Factual Hal. are hallucination that cannot be directly entailed in their generated context from the source document but can be based on world knowledge.

the model will generate the rejection token. Table 4 shows the evaluation setup. In this example, “*The University of Manchester has been awarded £3.5m to help fund a new cancer research centre.*” is one annotated sample from the XENT test set. The entity “£3.5m” is labeled as non-factual hallucination by human annotators. We use the left context (i.e., the first row in Table 4) as well as the source document as input to our model. There are five possible outputs from our model: (1) the same entity (2) a different entity that is hallucinated and non-factual (3) a different entity this is factual (4) removing the entity and replace it with an abstract description (e.g., “an 90-year-old woman” → “an elderly woman”) (5) a rejection token. The results are summarized in Table 5.

As shown in Table 5, when the entity is factual, our model will not change the generation in most cases (78.9%). When the entity is non-factual, however, 77.6% of the time our model will generate the rejection token and 4.5% of the time the hallucination will be removed. These results indicate that the rejection model works as we expect. We also notice that the model will reject a factual entity about 15% of the time. This number can be reduced by increasing the regularization weight  $\alpha$  in Equation 3. For example, by increasing  $\alpha$  from 1.0 to 2.0, the rejection rate for factual entities drops from 15% to 6.8%. However, the rejection rate for non-factual entities also drops to 68.7%.

## 5.2 The Effect of Beam Size on Factuality

Beam search is widely used to decode sentences from neural text generation models. Previous work has shown that increasing the beam size (i.e., the number of beam candidates held during searching)

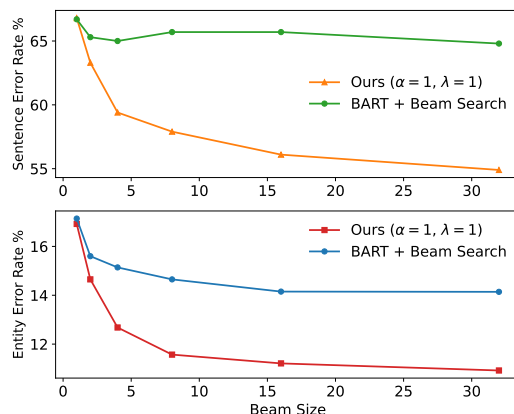


Figure 2: Beam size and factuality. We calculate the sentence error rate using DAE (Goyal and Durrett, 2021) and non-factual entity hallucination rate using ENTFA (Cao et al., 2022).

does not necessarily improve the model’s performance (Stahlberg and Byrne, 2019). However, little work has been done to investigate the effects of beam size on summary factuality. In this section, we evaluate the factuality of summaries generated using different beam size.

We compare the factuality of our model (rejection loss + regularized decoding) and the pre-trained BART model when decoded with different beam sizes. We select the beam sizes by powers of 2 from 1 to 32. In Figure 2, we observe that the factuality of summaries generated by BART slightly improves as the beam size increases. Our model, on the other hand, is able to achieve much better factuality score with a larger beam size. It is worth noting that our method has almost the same error rate as the baseline when doing greedy decoding (i.e., beam size = 1). Then the error rate decreases rapidly as the beam size increases. These results demonstrate that our method has a better inductive bias to identify non-factual beam candidates and remove them.

## 6 Conclusion

In this work, we propose a rejection learning based training objective for abstractive summarization models that is robust to the presence of noise in the dataset. The proposed training objective allows the model to learn features that are associated with the noise and reject them during training. We additionally propose a novel regularized decoding objective that effectively penalizes non-factual can-



didate summaries during inference. We show that our method vastly improves the factuality of generated summaries over five baseline models using both automatic and human evaluation. Further analysis reveals that our model enhances the factuality of the summaries while simultaneously increasing the abtractiveness of the generated summaries.

## 7 Limitations

Our assumption in this work is that noisy references contain information that is absent from the source, which creates an impossible learning task for the summarization model. As a result, noisy samples have high loss during training, which allows rejection learning to identify and reject them. However, it is not clear whether there also exist “difficult” but clean samples that also have high loss because they are simply difficult for the model to learn. For example, some articles in the training set contain numerical data, and arithmetic reasoning may be required to generate a good summary. If such samples exist, it is possible that our method would reject them instead of learning. Whether this happens remains to be evaluated.

In addition, the proposed method is only tested on the summarization task. Although our method can be applied to other NLG tasks such as machine translation and knowledge-grounded dialogue generation, it remains to be seen whether a similar advantage in factuality can be obtained.

## References

- Meng Cao, Yue Dong, and Jackie Cheung. 2022. [Hallucinated but factual! inspecting the factuality of hallucinations in abstractive summarization](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3340–3354, Dublin, Ireland. 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.
- Corinna Cortes, Giulia DeSalvo, and Mehryar Mohri. 2016. [Learning with rejection](#). In *Algorithmic Learning Theory - 27th International Conference, ALT 2016, Proceedings*, Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics), pages 67–82. Springer Verlag. Funding Information: This work was partly funded by NSF IIS-1117591, CCF-1535987, and DGE-1342536. Publisher Copyright: © Springer International Publishing Switzerland 2016.; 27th International Conference on Algorithmic Learning Theory, ALT 2016 ; Conference date: 19-10-2016 Through 21-10-2016.
- 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.
- Yue Dong, John Wieting, and Pat Verga. 2022. Faithful to the document or to the world? mitigating hallucinations via entity-linked knowledge in abstractive summarization. *arXiv preprint arXiv:2204.13761*.
- 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.
- Nouha Dziri, Sivan Milton, Mo Yu, Osmar Zaiane, and Siva Reddy. 2022. [On the origin of hallucinations in conversational models: Is it the datasets or the models?](#) *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*.
- 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.
- 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.
- Tanya Goyal, Jiacheng Xu, Junyi Jessy Li, and Greg Durrett. 2022. [Training dynamics for text summarization models](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 2061–2073, Dublin, Ireland. Association for Computational Linguistics.
- Yves Grandvalet, Alain Rakotomamonjy, Joseph Keshet, and Stéphane Canu. 2008. [Support vector machines with a reject option](#). In *Advances in Neural Information Processing Systems*, volume 21. Curran Associates, Inc.

- Max Grusky, Mor Naaman, and Yoav Artzi. 2018. [Newsroom: A dataset of 1.3 million summaries with diverse extractive strategies](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 708–719, New Orleans, Louisiana. Association for Computational Linguistics.
- Daniel Kang and Tatsunori B. Hashimoto. 2020. [Improved natural language generation via loss truncation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 718–731, Online. Association for Computational Linguistics.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Wojciech Kryscinski, Nitish Shirish Keskar, Bryan McCann, Caiming Xiong, and Richard Socher. 2019. [Neural text summarization: A critical evaluation](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 540–551, Hong Kong, China. 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.
- Faisal Ladhak, Esin Durmus, He He, Claire Cardie, and Kathleen McKeown. 2022. [Faithful or extractive? on mitigating the faithfulness-abstractiveness trade-off in abstractive summarization](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1410–1421, Dublin, Ireland. Association for Computational Linguistics.
- 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.
- 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.
- 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.
- Yixin Liu, Pengfei Liu, Dragomir Radev, and Graham Neubig. 2022. [BRIO: Bringing order to abstractive summarization](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2890–2903, Dublin, Ireland. 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.
- Feng Nan, Ramesh Nallapati, Zhiguo Wang, Cicero Nogueira dos Santos, Henghui Zhu, Dejiao 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*, Brussels, Belgium.
- Shashi Narayan, Yao Zhao, Joshua Maynez, Gonçalo Simões, Vitaly Nikolaev, and Ryan McDonald. 2021. [Planning with learned entity prompts for abstractive summarization](#). *Transactions of the Association for Computational Linguistics*, 9:1475–1492.
- Artidoro Pagnoni, Vidhisha Balachandran, and Yulia Tsvetkov. 2021. [Understanding factuality in abstractive summarization with FRANK: A benchmark for factuality metrics](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4812–4829, Online. Association for Computational Linguistics.
- Richard Yuanzhe Pang and He He. 2021. [Text generation by learning from demonstrations](#). In *International Conference on Learning Representations*.
- Weizhen Qi, Yu Yan, Yeyun Gong, Dayiheng Liu, Nan Duan, Jiusheng Chen, Ruofei Zhang, and Ming Zhou. 2020. [ProphetNet: Predicting future n-gram for sequence-to-SequencePre-training](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 2401–2410, Online. Association for Computational Linguistics.

- Vikas Raunak, Arul Menezes, and Marcin Junczys-Dowmunt. 2021. [The curious case of hallucinations in neural machine translation](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1172–1183, Online. Association for Computational Linguistics.
- Felix Stahlberg and Bill Byrne. 2019. [On NMT search errors and model errors: Cat got your tongue?](#) 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 3356–3362, Hong Kong, China. Association for Computational Linguistics.
- Sunil Thulasidasan, Tanmoy Bhattacharya, Jeff Bilmes, Gopinath Chennupati, and Jamal Mohd-Yusof. 2019. Combating label noise in deep learning using abstention. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 6234–6243, Long Beach, California, USA. PMLR.
- Jingqing Zhang, Yao Zhao, Mohammad Saleh, and Peter Liu. 2020. Pegasus: Pre-training with extracted gap-sentences for abstractive summarization. In *International Conference on Machine Learning*, pages 11328–11339. PMLR.
- 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 Annotation Guideline

We conduct human evaluation based on randomized pair comparison. With the following guidelines, we present annotators with 1) the source 2) and two systems’ outputs in random order:

1. Read the source completely.
2. Based on the source, compare the generations of system 1 with system 2 in terms of level of “factuality” with respect to the source document.
3. Try to consider the following questions when doing factuality evaluation:
  - Does the summary accurately capture the important information in the original document?
  - Is the salient information in each summary, such as location, date, quantity, etc., consistent with the original document?
  - Does the generated summary contain information that cannot be directly inferred from the source document? (avoid using external knowledge unless it is common sense)
  - Does the summary unfoundedly expand the meaning of the source document?
  - Are the relationship between characters, the sequence of events, and the causal relationship in the summary consistent with the original text?
4. Select the summary that you think is more faithful with respect to the source document (i.e., contain less unfounded facts, more accurate). Or they are equally factual/non-factual.

## B Inter-Annotator Agreement

We create a common set of 50 samples for inter-annotator agreement evaluation. We ask three annotators to annotate the common set. We report Fleiss’s Kappa ( $\kappa$ ) to assess the reliability of agreement between annotators. On the four-category annotation, we obtain a fair agreement with  $\kappa = 0.309$  ( $0.20 \leq \kappa \leq 0.40$ ). We also report the percentage  $\mu$  of annotators that annotate the majority class for the given example (Durmus et al., 2020). We get  $\mu = 0.633$  on our four-category annotation.

	Factuality		ROUGE-1	Novel Unigram
	FACTCC $\uparrow$	DAE $\uparrow$		
$\lambda = 1.0$	26.8%	40.0%	44.7	27.9%
$\lambda = 3.0$	28.3%	41.0%	44.4	27.9%
$\lambda = 5.0$	28.5%	41.2%	43.9	28.1%
$\lambda = 10.0$	30.0%	41.7%	42.7	28.6%

Table 6: Evaluation results on XSUM test set using the max regularizer proposed in Section 3.2 (Equation 6). We set  $\alpha = 1.0$ .

We speculate that the reason for the relatively low Fleiss’s Kappa score is because of the introduction of the *equally factual/non-factual* class which makes the annotation task hard. For example, when the two summaries both contain errors, some annotators prefer to select the one with less errors (as we specified in the guideline) but some annotators tend to mark both summaries as equally non-factual. When all three annotators only annotate win or loss, we get a substantial agreement with  $\kappa = 0.615$  ( $0.61 \leq \kappa \leq 0.80$ ) and  $\mu = 0.912$ .

## C Max Regularizer

Table 6 shows the evaluation results when we use the max regularizer defined in Section 3.2. Compared with the sum regularizer (Equation 5), we find that the max regularizer performs slightly worse in terms of automatic factuality evaluation scores. Thus we use the sum regularizer in our main results.

## D Example Appendix

---

**Source:**

St Michael's Hospital is joining other UK foetal medicine centres to set up a twin-to-twin transfusion syndrome (TTTS) registry to share information. Consultant Mark Denbow said they were "constantly learning" about TTTS. He said the "rare and often devastating condition" occurs in about 10-15% of identical twin pregnancies. St Michael's Hospital is also one of a few UK centres offering laser ablation surgery, where doctors can operate on the babies while they are in the womb. Jo and Finbarr O'Halloran, from Backwell, said they were "devastated" when they discovered their twins had the syndrome during a hospital scan. Mrs O'Halloran was given laser ablation surgery at about 22 weeks into the pregnancy in 2013. A few weeks later, daughters Eve and Amy were born by emergency caesarean section. Jo O'Halloran said: "I am literally thankful every single day that they survived and that they're here." Keith Reed, from the Twins and Multiple Birth Association (Tamba), said they were "still only skimming the surface" in terms of data collection. He said: "In order to build the best possible picture of TTTS cases in the UK, and help see which treatments offer the best possible outcomes, we need more hospitals with foetal medicine departments to sign up." TTTS affects twins who share a placenta, where the babies are not sharing blood equally. One baby gets too much blood and the other baby does not get enough. If left untreated 90% of these babies will die and even with treatment there is only up to 70% chance of both babies surviving.

---

**Reference:**

A study to improve the survival rate of unborn twins, with a potentially life threatening syndrome, is under way in Bristol.

---

**Baseline BART:**

A Swansea hospital is to set up a registry of twin-to-twin transfusion syndrome.

---

**Baseline Sentence Loss Truncation ( $c = 0.5$ ):**

A Worcestershire hospital has become the first in the UK to set up a registry for twin transfusion syndrome.

---

**Ours:**

A new registry is being set up to help save the lives of babies born with a rare blood disorder.

---

---

**Source:**

The 30-year-old is to remain there following the completion of a psychiatric report. Belfast Magistrates' Court was told she would not be able to leave without the hospital's permission. The woman cannot be named amid claims that identifying her would increase the risk of her taking her own life. A press challenge to the temporary reporting restrictions is due to be heard next month. The woman was arrested by detectives investigating the child's death following an incident in Belfast in March. She was charged with murder and then held under the Mental Health Act. Her barrister revealed on Wednesday that a medical report had now been prepared. Based on its contents he sought a termination of the current arrangements for keeping his client at the facility. "The application is that she be remanded on bail, subject to the condition that she continues to reside (there)," he said. A doctor who assessed the accused confirmed that the health trust consented to the proposal, provided the accused was there as a detained person. The judge was informed that under those arrangements the woman would not be able to leave. Granting the application, she listed the criminal proceedings for a further update in eight weeks time.

---

**Reference:**

A woman accused of murdering her baby son has been remanded on bail at a mental health facility.

---

**Baseline BART:**

A woman charged with murdering her two-year-old daughter has been remanded in custody at the Royal Victoria Hospital in Belfast.

---

**Baseline Sentence Loss Truncation ( $c = 0.5$ ):**

A woman charged with murdering her two-year-old daughter has been remanded in custody at Belfast's Royal Victoria Hospital.

---

**Ours:**

A woman charged with murder has been remanded in custody after a judge granted her bail on the condition that she remains in hospital.

---

Table 7: Comparison of summaries generated by the standard BART model, BART trained with sentence loss truncation and BART trained using our method. The factual errors made by the models are highlighted orange.