

Optimizing the Factual Correctness of a Summary: A Study of Summarizing Radiology Reports

Yuhao Zhang¹, Derek Merck², Emily Bao Tsai¹,
Christopher D. Manning¹, Curtis P. Langlotz¹

¹Stanford University ²University of Florida

{yuhaozhang, ebtsai, manning, langlotz}@stanford.edu
derek.merck@ufl.edu

Abstract

Neural abstractive summarization models are able to generate summaries which have high overlap with human references. However, existing models are not optimized for factual correctness, a critical metric in real-world applications. In this work, we develop a general framework where we evaluate the factual correctness of a generated summary by fact-checking it automatically against its reference using an information extraction module. We further propose a training strategy which optimizes a neural summarization model with a factual correctness reward via reinforcement learning. We apply the proposed method to the summarization of radiology reports, where factual correctness is a key requirement. On two separate datasets collected from hospitals, we show via both automatic and human evaluation that the proposed approach substantially improves the factual correctness and overall quality of outputs over a competitive neural summarization system, producing radiology summaries that approach the quality of human-authored ones.

1 Introduction

Neural abstractive summarization systems aim at generating sentences which compress a document while preserving the key facts in it (Nallapati et al., 2016b; See et al., 2017; Chen and Bansal, 2018). These systems are potentially useful in many real-world applications. For example, Zhang et al. (2018) have shown that customized neural abstractive summarization models are able to generate radiology summary statements with high quality by summarizing textual findings written by radiologists. This task has significant clinical value because of its potential to accelerate the radiology workflow, reduce repetitive human labor, and improve clinical communications (Kahn Jr et al., 2009).

Background: radiographic examination of the chest. clinical history: 80 years of age, male ...

Findings: frontal radiograph of the chest demonstrates repositioning of the right atrial lead possibly into the ivc. ... a right apical pneumothorax can be seen from the image. moderate right and small left pleural effusions continue. no pulmonary edema is observed. heart size is upper limits of normal.

Human Summary: pneumothorax is seen. bilateral pleural effusions continue.

Summary A (ROUGE-L = 0.77):

no pneumothorax is observed. bilateral pleural effusions continue.

Summary B (ROUGE-L = 0.44):

pneumothorax is observed on radiograph. bilateral pleural effusions continue to be seen.

Figure 1: A (truncated) radiology report and summaries with their ROUGE-L scores. Compared to the human summary, Summary A has high textual overlap (i.e., ROUGE-L) but makes a **factual error**; Summary B has a lower ROUGE-L score but is **factually correct**.

However, while existing abstractive summarization models are optimized to generate summaries that highly overlap with human references (Paulus et al., 2018), this does not guarantee factually correct summaries, as shown in Figure 1. Therefore, maintaining factual correctness of the generated summaries remains a critical yet unsolved problem. For example, Zhang et al. (2018) found that about 30% of the outputs from a radiology summarization model contain factual errors or inconsistencies. This has made such a system unusable in practice, as factual correctness is critically important in this domain to prevent medical errors.

Existing attempts at improving the factual correctness of abstractive summarization models have seen very limited success. For example, Cao et al. (2017) augmented the attention mechanism of neural models with factual triples extracted with open information extraction systems; Falke et al. (2019)

studied using natural language inference systems to rerank generated summaries based on their factual consistencies; Kryściński et al. (2019b) proposed to verify factual consistency of generated summaries with a weakly-supervised model. Despite these efforts, none of the existing work has focused explicitly on optimizing an abstractive summarization system with a correctness objective. As a result, even state-of-the-art systems trained with ample data still produce summaries with a substantial number of factual errors (Goodrich et al., 2019; Kryściński et al., 2019a).

In this work we aim to optimize the factual correctness of existing neural summarization systems, with a focus on summarizing radiology reports. This task has several key properties that make it ideal for studying factual correctness in summarization models. First, the clinical facts or observations present in radiology reports have less ambiguity compared to open-domain text, which allows objective comparison of facts. Second, radiology reports involve a relatively limited space of facts, which makes automatic measurement of factual correctness in the generated text approachable. Lastly, as factual correctness is a crucial metric in this domain, improving factual correctness will directly lead to an ability to use the system.

To this end, we design a framework where an external information extraction system is used to extract information in the generated summary and produce a factual accuracy score by comparing it against the human reference summary. We further develop a training strategy where we combine a factual correctness objective, a textual overlap objective and a language model objective, and jointly optimize them via reinforcement learning (RL).

On two datasets of radiology reports collected from different hospitals, we show that our training strategy substantially improves the factual correctness of the summaries generated by a competitive neural summarization system. Moreover, we observe for the first time that, even in the absence of a factual correctness objective, optimizing a textual overlap-based metric substantially improves the factual correctness of the resulting system compared to maximum likelihood training. We further show via human evaluation and analysis that our training strategy leads to summaries with higher overall quality and correctness and which are closer to the human-written ones.

Our main contributions are: (i) we propose a

general framework and a training strategy for improving the factual correctness of summarization models by optimizing a multi-part objective via RL; (ii) we apply the proposed strategy to radiology reports, and empirically show that it improves the factual correctness of the generated summaries; and (iii) we demonstrate via radiologist evaluation that our system is able to generate summaries with clinical validity close to human-written ones. To our knowledge, our work represents the first attempt at directly optimizing a neural summarization system with a factual correctness objective via RL.

2 Related Work

Neural Summarization Systems. Neural models for text summarization can be broadly divided into extractive approaches (Cheng and Lapata, 2016; Nallapati et al., 2016a) and abstractive approaches (Nallapati et al., 2016b; See et al., 2017). While existing models are often trained in an end-to-end manner by maximizing the likelihood of the reference summaries, RL has been shown useful in recent work (Chen and Bansal, 2018; Dong et al., 2018). Specifically, Paulus et al. (2018) found that directly optimizing an abstractive summarization model on the ROUGE metric via RL can improve the summary ROUGE scores. Our work extends the rewards used in existing work with a factual correctness reward to further improve the correctness of the generated summaries.

Factual Correctness in Summarization. Our work is closely related to recent work that studies factual correctness in summarization. Cao et al. (2017) proposed to improve summarization models by attending to fact triples extracted using open information extraction systems. Goodrich et al. (2019) compared different information extraction systems to evaluate the factual accuracy of generated text. Falke et al. (2019) explored using natural language inference systems to evaluate the correctness of generated summaries, and found models trained on existing datasets to be inadequate. Kryściński et al. (2019b) proposed to evaluate factual consistencies in the generated summaries using a weakly-supervised fact verification model. Despite these efforts, none of this work has shown success in directly optimizing a summarization system for factual correctness, and to our knowledge our work represents the first attempt in this direction. While our work is focused on improving neural summarization models, we note that the idea of

using information extraction systems to evaluate the fidelity of generated text has also been explored for data-to-text generation (Wiseman et al., 2017; Dhingra et al., 2019).

Summarization of Radiology Reports. Zhang et al. (2018) first studied the problem of automatic generation of radiology impressions by summarizing textual radiology findings, and showed that an augmented pointer-generator model achieves high overlap with human references. MacAvaney et al. (2019) extended this model with an ontology-aware pointer-generator and showed improved summarization quality. Li et al. (2019) and Liu et al. (2019) studied generating textual descriptions of radiology findings from medical images, and proposed RL-based approaches to tackle this problem. While Zhang et al. (2018) found that about 30% of the radiology summaries generated from neural models contain factual errors, improving factual correctness in radiology summarization remains unstudied.

3 Task & Baseline Pointer-Generator

We start by briefly introducing the task of summarizing radiology findings. Given a passage of radiology findings represented as a sequence of tokens $\mathbf{x} = \{x_1, x_2, \dots, x_N\}$, with N being the length of the findings, the task involves finding a sequence of tokens $\mathbf{y} = \{y_1, y_2, \dots, y_L\}$ that best summarizes the salient and clinically significant findings in \mathbf{x} . In routine radiology workflow, an output sequence \mathbf{y} is produced by the radiologist, which we treat as a reference summary sequence.¹

To model the summarization process, we use the background-augmented pointer-generator network (Zhang et al., 2018) as the backbone of our method. This abstractive summarization model extends a pointer-generator (See et al., 2017) with a separate background section encoder and is shown to be effective in summarizing radiology notes with multiple sections. We briefly describe this model and refer readers to the original papers for details.

At a high level, this model first encodes the input sequence \mathbf{x} into hidden states with a Bi-directional Long Short-Term Memory (Bi-LSTM) network, and then generates an output sequence \mathbf{y} with a separate LSTM decoder. To make the input information available at decoding time, an attention

¹While the name “impression” is often used in clinical settings, we use “summary” and “impression” interchangeably.

mechanism (Bahdanau et al., 2015) over the input hidden states is also added to the decoder.

The baseline pointer-generator model by Zhang et al. (2018) adds two augmentations to this attentional encoder-decoder model to make it suitable for summarizing radiology findings:

Copy Mechanism. To enable the model to copy words from the input, a copy mechanism (Vinyals et al., 2015; See et al., 2017) is added to calculate a generation probability at each step of decoding. This generation probability is then used to blend the original output vocabulary distribution and a copy distribution to generate the next word.

Background-guided Decoding. As shown in Figure 1, radiology reports often consist of a background section which documents the crucial study background information (e.g., purpose of the study, patient conditions), and a findings section which documents clinical observations. While words can be copied from the findings section to form the summary, Zhang et al. (2018) found it worked better to separately encode the background section, and inject the representation into the decoding process by concatenating it with the input.

4 Fact Checking in Summarization

Summarization models such as the one described in Section 3 are commonly trained with the teacher-forcing algorithm (Williams and Zipser, 1989) by maximizing the likelihood of the reference, human-written summaries. However, this training strategy results in a significant discrepancy between what the model sees during training and test time, often referred to as the *exposure bias* issue (Ranzato et al., 2016), leading to degenerate output at test time.

An alternative training strategy is to directly optimize standard metrics such as ROUGE scores (Lin, 2004) with RL and this was shown to improve summarization quality (Paulus et al., 2018). Nevertheless, this method still provides no guarantee that the generated summary is factually accurate and complete, since the ROUGE scores merely measure the superficial text overlap between two sequences and do not account for the factual alignment between them. To illustrate this, a reference sentence *pneumonia is seen* and a generated sentence *pneumonia is not seen* have substantial text overlap and thus the generated sentence would achieve a high ROUGE score, however the generated sentence conveys an entirely opposite fact. In this section

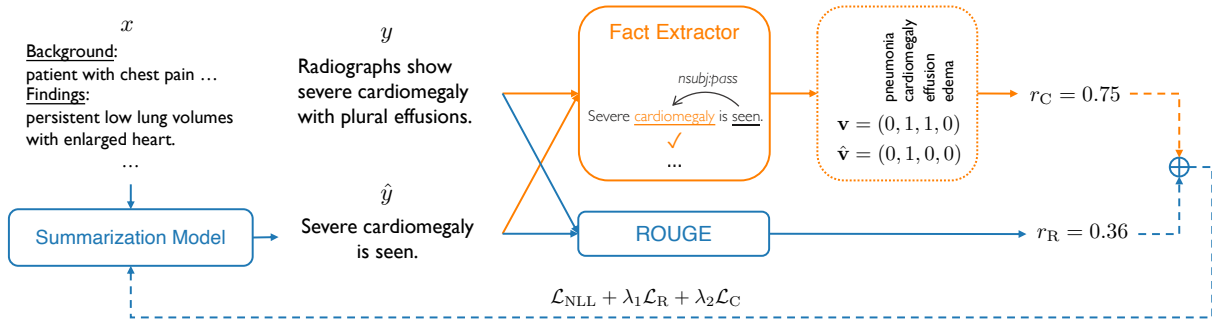


Figure 2: Our proposed training strategy. Compared to existing work which relies only on a ROUGE reward r_R , we add a factual correctness reward r_C which is enabled by a fact extractor. The summarization model is updated via RL, using a combination of the NLL loss, a ROUGE-based loss and a factual correctness-based loss. For simplicity we only show a subset of the clinical variables in the fact vectors \mathbf{v} and $\hat{\mathbf{v}}$.

we first introduce a method to verify the factual correctness of the generated summary against the reference summary, and then describe a training strategy to directly optimize a factual correctness objective to improve summary quality.

4.1 Evaluating Factual Correctness via Fact Extraction

A convenient way to explicitly measure the factual correctness of a generated summary against the reference is to first extract and represent the facts in a structured format. To this end, we define a *fact extractor* to be an information extraction (IE) module, denoted as f , which takes in a summary sequence y and returns a structured fact vector \mathbf{v} :

$$\mathbf{v} = f(y) = (v_1, \dots, v_m) \quad (1)$$

where v_i is a categorical variable that we want to measure via fact checking and m the total number of such variables. For example, in the case of summarizing radiology reports, v_i can be a binary variable that describes whether an event or a disease such as *pneumonia* is present or not in a radiology study.

Given a fact vector \mathbf{v} output by f from a reference summary and $\hat{\mathbf{v}}$ from a generated summary, we further define a *factual accuracy* score s to be the ratio of variables in $\hat{\mathbf{v}}$ which equal the corresponding variables in \mathbf{v} , namely:

$$s(\hat{\mathbf{v}}, \mathbf{v}) = \frac{\sum_{i=1}^m \mathbb{1}[v_i = \hat{v}_i]}{m} \quad (2)$$

where $s \in [0, 1]$. Note that this method requires a summary to be both precise and complete in order to achieve a high s score: missing out a positive variable or falsely claiming a negative variable will be equally penalized.

Our general definition of the fact extractor module f allows it to have different realizations for different domains. For our task of summarizing radiology findings, we make use of the open-source CheXpert radiology report labeler (Irvin et al., 2019).² At its core, the CheXpert labeler parses the input sentences into dependency structures and runs a series of surface and syntactic rules to extract the presence status of 14 clinical observations seen in chest radiology reports.³ It was evaluated to have over 95% overall F_1 when compared against oracle annotations from multiple radiologists on a large-scale radiology report dataset.

4.2 Improving Factual Correctness via Policy Learning

The fact extractor module introduced above not only enables us to measure the factual accuracy of a generated summary, but also provides us with an opportunity to directly optimize the factual accuracy as an objective. This can be achieved by viewing our summarization model as an agent, the actions of which are to generate a sequence of words to form the summary \hat{y} , conditioned on the input x .⁴ The agent then receives rewards $r(\hat{y})$ for its actions, where the rewards can be designed to measure the quality of the generated summary. Our goal is to learn an optimal policy $P_\theta(y|x)$ for the summarization model, parameterized by the network parameters θ , which achieves the highest expected reward under the training data.

Formally, we minimize loss \mathcal{L} , the negative ex-

²<https://github.com/stanfordmlgroup/chexpert-labeler>

³For this study we used a subset of these variables and discuss the reasons in Appendix A.

⁴For clarity, we drop the bold symbol and use x and y to represent the input and output sequences, respectively.

pectation of the reward $r(\hat{y})$ over the training data:

$$\mathcal{L}(\theta) = -\mathbb{E}_{\hat{y} \sim P_{\theta}(y|x)}[r(\hat{y})]. \quad (3)$$

The gradient can be calculated as (REINFORCE Williams, 1992):

$$\nabla_{\theta} \mathcal{L}(\theta) = -\mathbb{E}_{\hat{y} \sim P_{\theta}(y|x)}[\nabla_{\theta} \log P_{\theta}(\hat{y}|x) r(\hat{y})]. \quad (4)$$

In practice, we approximate this gradient over a training example with a single Monte Carlo sample and deduct a baseline reward to reduce the variance of the gradient estimation:

$$\nabla_{\theta} \mathcal{L}(\theta) \approx -\nabla_{\theta} \log P_{\theta}(\hat{y}_s|x)(r(\hat{y}_s) - \bar{r}), \quad (5)$$

where \hat{y}_s is a sampled sequence from the model and \bar{r} a baseline reward. Here we adopt the *self-critical training* strategy (Rennie et al., 2017), where we obtain the baseline reward \bar{r} by applying the same reward function r to a greedily decoded sequence \hat{y}_g , i.e., $\bar{r} = r(\hat{y}_g)$. We empirically find that using this self-critical baseline reward helps stabilize the training of our summarization model.

4.3 Reward Function

The learning strategy in Equation (5) provides us with the flexibility to optimize arbitrary reward functions. Here we decompose our reward function into two parts:

$$r = \lambda_1 r_R + \lambda_2 r_C, \quad (6)$$

where $r_R \in [0, 1]$ is a ROUGE reward, namely the ROUGE-L score (Lin, 2004) of the predicted sequence \hat{y} against the reference y ; $r_C \in [0, 1]$ is a correctness reward, namely the factual accuracy s of the predicted sequence against the reference sequence, as in Equation (2); $\lambda_1, \lambda_2 \in [0, 1]$ are scalar weights that control the balance between the two. To measure the similarity between the reference and the generation, we also experimented with more recent metrics that rely on neural representations of text, such as the BERTScore (Zhang et al., 2020). However, we found that these metrics, mostly trained on web and newswire data, generalize poorly to our domain of text.

Paulus et al. (2018) found that directly optimizing a reward function without the original negative log-likelihood (NLL) objective as used in teacher-forcing can hurt the readability of the generated summaries, and proposed to alleviate this problem by combining the NLL objective with the RL loss.

Split	Number of Examples	
	Stanford	RIH
Train	89,992 (68.8%)	84,194 (60.3%)
Dev	22,031 (16.8%)	25,966 (18.6%)
Test	18,827 (14.4%)	29,494 (21.1%)
Total	130,850	139,654

Table 1: Statistics of the Stanford and RIH datasets.

Here we adopt the same strategy, and our final loss during training is:

$$\mathcal{L} = \lambda_1 \mathcal{L}_R + \lambda_2 \mathcal{L}_C + \lambda_3 \mathcal{L}_{\text{NLL}}, \quad (7)$$

where $\lambda_3 \in [0, 1]$ is an additional scalar that controls the weight of the NLL loss.

Our overall training strategy is illustrated in Figure 2. Our final loss jointly optimizes three aspects of the summaries: \mathcal{L}_{NLL} serves as a conditional language model that optimizes the fluency and relevance of the generated summary, \mathcal{L}_R controls the brevity of the summary and encourages summaries which have high overlap with human references, and \mathcal{L}_C encourages summaries that are factually accurate when compared against human references.

5 Experiments

We collected two real-world radiology report datasets and describe our experiments using them as our main training and evaluation corpora.

5.1 Data Collection

We collected anonymized chest radiographic reports within a certain period of time from two collaborating hospitals: the Stanford University Hospital and the Rhode Island Hospital (RIH).⁵

For both datasets, we ran simple preprocessing following Zhang et al. (2018). To test the generalizability of the models, instead of using random stratification, we stratified each dataset over time into training, dev and test splits. We include statistics of both datasets in Table 1 and preprocessing and stratification details in Appendix B.

5.2 Models

As we use the augmented pointer-generator network described in Section 3 as the backbone of our method, we mainly compare against it as the

⁵Our retrospective study has been approved by the corresponding institutional review boards with waiver of consent.

System	Stanford				RIH			
	R-1	R-2	R-L	Factual F ₁	R-1	R-2	R-L	Factual F ₁
LexRank (Erkan and Radev, 2004)	26.8	16.3	23.6	—	20.6	10.7	18.3	—
BanditSum (Dong et al., 2018)	32.7	20.9	29.0	—	26.1	14.0	23.3	—
PG Baseline (Zhang et al., 2018)	48.3	38.8	46.6	55.9	54.1	44.7	52.2	69.3
PG + RL _R	52.0	41.1	49.5	63.2	58.0	47.2	55.7	73.3
PG + RL _C	50.7	39.7	48.0	65.9	55.2	45.4	52.9	75.4
PG + RL _{R+C}	52.0	41.0	49.3	64.5	57.0	46.6	54.7	74.8

Table 2: Main results on the two datasets. R-1, R-2, R-L represent the ROUGE scores. PG Baseline represents our baseline augmented pointer-generator; RL_R, RL_C and RL_{R+C} represent RL training with the ROUGE reward alone, with the factual correctness reward alone and with both. All the ROUGE scores have a 95% confidence interval of at most ± 0.6 . F₁ scores for extractive models were not evaluated for the reason discussed in Section 5.3.

baseline model (PG Baseline), and use the open implementation by Zhang et al. (2018).

For the proposed RL-based training, we compare three variants: training with only the ROUGE reward (RL_R), with only the factual correctness reward (RL_C), or with both (RL_{R+C}). All three variants have the NLL component in the training loss as in Equation (7). For all variants, we initialize the model with the best baseline model trained with standard teacher-forcing, and then finetune it on the training data with the corresponding RL loss, until it reaches the best validation score.

To understand the difficulty of the task and evaluate the necessity of using abstractive summarization models, we additionally evaluate two extractive summarization methods: (1) LexRank (Erkan and Radev, 2004), a widely-used non-neural extractive summarization algorithm; and (2) BanditSum (Dong et al., 2018), a state-of-the-art RL-based neural extractive summarization model. For both methods we use their open implementations. We include other model implementation and training details in Appendix C.

5.3 Evaluation

We use two sets of metrics to evaluate model performance at the corpus level. First, we use the standard ROUGE scores (Lin, 2004), and report the F₁ scores for ROUGE-1, ROUGE-2 and ROUGE-L, which compare the word-level unigram, bigram and longest common sequence overlap with the reference summary, respectively.

For factual correctness evaluation, we use a **Factual F₁** score. While the factual accuracy score s that we use in the reward function evaluates how factually accurate a specific summary is, compar-

ing it at the corpus level can be misleading, for the same reason that accuracy is a misleading measure in information retrieval (Manning et al., 2008). To understand this, imagine the case where a clinical variable v has rare presence in the corpus. A model which always generates a negative summary for it (i.e., $v = 0$; the disease is not present) can have high accuracy, but is useless in practice. Instead, for each variable, we obtain a model’s predictions over all test examples and calculate its F₁ score. We then macro-average the F₁ of all variables to obtain the overall factual F₁ score of the model.

Note that the CheXpert labeler that we use is specifically designed to run on radiology summaries, which usually have a different style of language compared to the radiology findings section of the reports (see further analysis in Section 7). As a result, we found the labeler to be less accurate when applied to the findings section. For this reason, we were not able to estimate the factual F₁ scores on the summaries generated by the two extractive summarization models.

6 Results

We first present our automatic evaluation results on the two collected datasets. We then present a human evaluation with board-certified radiologists where we compare the summaries generated by humans, the baseline and our proposed model.

6.1 Automatic Evaluation

Our main results on both datasets are shown in Table 2. We first notice that while the neural extractive model, BanditSum, outperforms the non-neural extractive method on ROUGE scores, our PG baseline model substantially outperforms both of them,

Variable	PG Baseline	RL _{R+C}	Δ
No Finding	77.3	81.5	+4.2*
Cardiomegaly	29.5	40.4	+10.9*
Airspace Opacity	64.6	74.9	+10.3*
Edema	58.4	70.9	+12.5*
Consolidation	46.3	53.2	+6.9*
Pneumonia	46.7	46.8	+0.2
Atelectasis	48.8	56.3	+7.5*
Pneumothorax	69.5	82.9	+13.4*
Pleural Effusion	62.0	73.4	+11.4*
Macro Avg.	55.9	64.5	+8.6*

Table 3: Test set factual F_1 scores for each variable on the Stanford dataset. * marks statistically significant improvements with $p < .01$ under a bootstrap test.

suggesting that on both datasets abstractive summarization is necessary to generate summaries comparable to human-written ones. We further show that this difference is likely due to the different styles of language (see Section 7): while radiologists tend to use more compressed language when writing the summaries, extractive methods produce more verbose summaries that fail to capture this difference.

On the Stanford dataset, training the pointer-generator model with ROUGE reward alone (RL_R) leads to improvements on all ROUGE scores, with a gain of 2.9 ROUGE-L scores. Training with the factual correctness reward alone (RL_C) leads to the best overall factual F_1 with a substantial gain of 10% absolute, however with consistent decline in the ROUGE scores compared to RL_R training. Combining the ROUGE and the factual correctness rewards (RL_{R+C}) achieves a balance between the two, leading to an overall improvement of 2.7 on ROUGE-L and 8.6% on factual F_1 compared to the baseline. This indicates that RL_{R+C} training leads to both higher overlap with references and improved factual correctness.

Most surprisingly, while ROUGE has been criticized for its poor correlation with human judgment of quality and insufficiency for evaluating correctness of the generated text (Chaganty et al., 2018), we find that optimizing ROUGE reward jointly with NLL leads to substantially more factually correct summaries than the baseline, shown by the notable gain of 7.3% factual F_1 from the RL_R training.

All of our findings are consistent on the RIH dataset, with RL_{R+C} achieving an overall improve-

Stanford Dataset
Background: radiographic examination of the chest ...
Findings: continuous rhythm monitoring device again seen projecting over the left heart. persistent low lung volumes with unchanged cardiomegaly. again seen is a diffuse reticular pattern with interstitial prominence demonstrated represent underlying emphysematous changes with superimposed increasing moderate pulmonary edema. small bilateral pleural effusions. persistent bibasilar opacities left greater than right which may represent infection versus atelectasis.
Human: increased moderate pulmonary edema with small bilateral pleural effusions . left greater than right basilar opacities which may represent infection versus atelectasis.
PG Baseline ($s = 0.33$): no significant interval change.
RL_{R+C} ($s = 1.00$): increasing moderate pulmonary edema . small bilateral pleural effusions . persistent bibasilar opacities left greater than right which may represent infection versus atelectasis.
RIH Dataset
Background: history: lobar pneumonia, unspecified organism ...
Findings: lines/tubes: none. lungs: right middle lobe airspace disease seen on prior radiographs from <date> and <date> is no longer evident . bilateral lungs appear clear. pleura: there is no pleural effusion or pneumothorax. heart and mediastinum: no cardiomegaly. thoracic aorta appears calcified and mildly tortuous. bones: ...
Human: no acute cardiopulmonary abnormality.
PG Baseline ($s = 0.75$): right middle lobe airspace disease could represent atelectasis, aspiration or pneumonia.
RL_{R+C} ($s = 1.00$): no acute cardiopulmonary abnormality.

Figure 3: Truncated examples from the test sets along with human, PG baseline and RL_{R+C} outputs. Factual accuracy scores (s) are also shown for the model outputs. For the Stanford example, [clinical observations](#) in the summaries are marked for clarity; for RIH, [a wrongly copied observation](#) is marked.

ment of 2.5 ROUGE-L and 5.5% factual F_1 scores.

Fine-grained Correctness. To understand how improvements in individual variables contribute to the overall improvement, we show the fine-grained factual F_1 scores for all variables on the Stanford dataset in Table 3 and include results on the RIH dataset in Appendix D. We find that on both datasets, improvements in RL_{R+C} can be observed on all variables tested. We further find that, as we change the initialization across different training runs, while the overall improvement on factual F_1 stays approximately unchanged, the distribution of the improvement on different variables can vary substantially. Developing a training strategy for fine-grained control over different variables is an interesting direction for future work.

Qualitative Results. In Figure 3 we present two example reports along with the human references, the PG baseline outputs and RL_{R+C} outputs. In the first example, while baseline output seems generic and does not include any meaningful observation, the summary from the RL_{R+C} model aligns well with the reference, and therefore achieves a higher

Metric	Win	Tie	Lose
Our Model vs. PG Baseline			
Fluency	7%	60%	33%
Factual Correctness	31%	55%	14%
Overall Quality	48%	24%	28%
Our Model vs. Human Reference			
Fluency	17%	54%	29%
Factual Correctness	23%	49%	28%
Overall Quality	44%	17%	39%

Table 4: Results of the radiologist evaluation. The top three rows present results when comparing our RL_{R+C} model output versus the baseline model output; the bottom three rows present results when comparing our model output versus the human-written summaries.

factual accuracy score. In the second example, the baseline model wrongly copied an observation from the findings although the actual context is *no longer evident*, while the RL_{R+C} model correctly recognizes this and produces a better summary.

6.2 Human Evaluation

To study whether the improvements in the factual correctness scores lead to improvement in summarization quality under expert judgment, we run a comparative human evaluation following previous work (Chen and Bansal, 2018; Dong et al., 2018; Zhang et al., 2018). We sampled 50 test examples from the Stanford dataset, and for each example we presented to two board-certified radiologists the full radiology findings along with blinded summaries from (1) the human reference, (2) the PG baseline and (3) our RL_{R+C} model. We shuffled the three summaries such that the correspondence cannot be guessed, and asked the radiologists to compare them based on the following three metrics: (1) **fluency**, (2) **factual correctness and completeness**, and (3) **overall quality**. For each metric we asked the radiologists to rank the three summaries, with ties allowed. After the evaluation, we converted each ranking into two binary comparisons: (1) our model versus the baseline model, and (2) our model versus human reference.

The results are shown in Table 4. Comparing our model against the baseline model, we find that: (1) in terms of fluency our model is less preferred, although a majority of the results (60%) are ties; (2) our model wins more on factual correctness and overall quality. Comparing our model against

System	Stanford pplx.	RIH pplx.
Human	6.7	5.5
LexRank	10.8	36.9
BanditSum	9.9	40.9
PG Baseline	4.8	3.8
PG + RL_{R+C}	6.5	4.8

Table 5: Perplexity scores as evaluated by the trained radiology impression LM on the test set human references and model predictions.

human references, we find that: (1) human wins more on fluency; (2) factual correctness results are close, with 72% of our model outputs being at least as good as human; (3) surprisingly, in terms of overall quality our model was slightly preferred by the radiologists compared to human references. Lastly, when comparing the baseline model against human references, we find that outputs from the baseline model are much less correct and lower-quality than human summaries.

7 Analysis & Discussion

Fluency and Style of Summaries. Our human evaluation results in Section 6.2 suggest that in terms of fluency our model output is less preferred than human reference and baseline output. To further understand the fluency and style of summaries from different models at a larger scale, we trained a neural language model (LM) for radiology summaries following previous work (Liu et al., 2018). Intuitively, radiology summaries which are more fluent and consistent with humans in style should be able to achieve a lower perplexity under this in-domain LM, and vice versa. To this end, we collected all human-written summaries from the training and dev split of both datasets, which in total gives us about 222,000 summaries. We then trained a strong Mixture of Softmaxes LM (Yang et al., 2018) on this corpus, and evaluated the perplexity of test set outputs for all models.

The results are shown in Table 5. We find that while extractive models can achieve non-trivial overlap with references, their perplexity scores tend to be much higher than humans. We conjecture that this is because radiologists are trained to write the summaries with more compressed language than when they are writing the findings, therefore sentences directly extracted from the findings tend to be more verbose than needed.

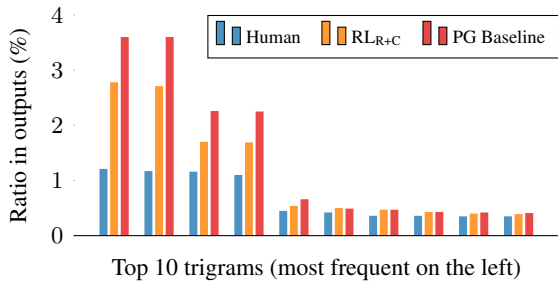


Figure 4: Distributions of the top 10 most frequent trigrams from model outputs on the Stanford test set.

We further observe that the baseline model achieves even lower perplexity than humans, and our proposed method leads to a perplexity score much closer to human references. We hypothesize that this is because models trained with teacher-forcing are prone to generic generations which are fluent and relevant but may not be factually correct. Training with the proposed rewards alleviates this issue, leading to summaries more consistent with humans in style. For example, we find that *no significant interval change* is a very frequent generation from the baseline, regardless of the actual input. This sentence occurs in 34% of the baseline outputs on the Stanford dev set, while the number for RL_{R+C} and human are only 24% and 17%. This hypothesis is further confirmed when we plot the distribution of the top 10 most frequent trigrams from different models in Figure 4: while the baseline heavily reuses the few most frequent trigrams, our model RL_{R+C} tends to have more diverse summaries which are closer to human references. The same trends are observed for 4-grams and 5-grams.

Limitations. While we showed the success of our proposed method on improving the factual correctness of a radiology summarization model, we also recognize several limitations of our work. First, our proposed training strategy crucially depends on the availability of an external IE module. While this IE module is relatively easy to implement for a domain with a limited space of facts, how to generalize this method to open-domain summarization remains unsolved. Second, our study was based on a rule-based IE system, and the use of a more robust statistical IE model can potentially improve the results. Third, we mainly focus on key factual errors which result in a flip of the binary outcome of an event (e.g., presence of disease), whereas factual errors in generated summaries can occur in other forms such as wrong adjectives or

coreference errors (Kryściński et al., 2019a). We leave the study of these problems to future work.

8 Conclusion

In this work we presented a general framework and a training strategy to improve the factual correctness of neural abstractive summarization models. We applied this approach to the summarization of radiology reports, and showed its success via both automatic and human evaluation on two separate datasets collected from hospitals.

Our general takeaways include: (1) in a domain with a limited space of facts such as radiology reports, a carefully implemented IE system can be used to improve the factual correctness of neural summarization models via RL; (2) even in the absence of a reliable IE system, optimizing the ROUGE metrics via RL can substantially improve the factual correctness of the generated summaries.

We hope that our work draws the community’s attention to the factual correctness issue of abstractive summarization models and inspires future work in this direction.

Acknowledgments

The authors would like to thank the anonymous reviewers, Peng Qi and Urvashi Khandelwal for their helpful comments, and Dr. Jonathan Movson for his help with obtaining the RIH data used in this study.

References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. *The 2015 International Conference on Learning Representations*.
- Ziqiang Cao, Furu Wei, Wenjie Li, and Sujian Li. 2017. Faithful to the original: Fact aware neural abstractive summarization. *The Thirty-First AAAI Conference on Artificial Intelligence (AAAI-2017)*.
- Arun Chaganty, Stephen Mussmann, and Percy Liang. 2018. The price of debiasing automatic metrics in natural language evaluation. *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (ACL 2018)*.
- Yen-Chun Chen and Mohit Bansal. 2018. Fast abstractive summarization with reinforce-selected sentence rewriting. *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (ACL 2018)*.

- Jianpeng Cheng and Mirella Lapata. 2016. Neural summarization by extracting sentences and words. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (ACL 2016)*.
- Bhuwan Dhingra, Manaal Faruqui, Ankur Parikh, Ming-Wei Chang, Dipanjan Das, and William Cohen. 2019. Handling divergent reference texts when evaluating table-to-text generation. *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (ACL 2019)*.
- Yue Dong, Yikang Shen, Eric Crawford, Herke van Hoof, and Jackie Chi Kit Cheung. 2018. BanditSum: Extractive summarization as a contextual bandit. *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP 2018)*.
- Günes Erkan and Dragomir R Radev. 2004. LexRank: Graph-based lexical centrality as salience in text summarization. *Journal of Artificial Intelligence Research*, 22:457–479.
- 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 (ACL 2019)*.
- Ben Goodrich, Vinay Rao, Peter J. Liu, and Mohammad Saleh. 2019. Assessing the factual accuracy of generated text. *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD 19)*.
- Jeremy Irvin, Pranav Rajpurkar, Michael Ko, Yifan Yu, Silvana Ciurea-Ilcus, Chris Chute, Henrik Marklund, Behzad Haghgoo, Robyn Ball, Katie Shpankaya, et al. 2019. CheXpert: A large chest radiograph dataset with uncertainty labels and expert comparison. *The Thirty-Third AAAI Conference on Artificial Intelligence (AAAI 2019)*.
- Charles E Kahn Jr, Curtis P Langlotz, Elizabeth S Burnside, John A Carrino, David S Channin, David M Hovsepian, and Daniel L Rubin. 2009. Toward best practices in radiology reporting. *Radiology*, 252(3):852–856.
- Diederik P Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *The 2015 International Conference for Learning Representations*.
- Wojciech Kryściński, Nitish Shirish Keskar, Bryan McCann, Caiming Xiong, and Richard Socher. 2019a. Neural text summarization: A critical evaluation. *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing (EMNLP 2019)*.
- Wojciech Kryściński, Bryan McCann, Caiming Xiong, and Richard Socher. 2019b. Evaluating the factual consistency of abstractive text summarization. *arXiv preprint arXiv:1910.12840*.
- Christy Y. Li, Xiaodan Liang, Zhiting Hu, and Eric P. Xing. 2019. Hybrid retrieval-generation reinforced agent for medical image report generation. *Advances in neural information processing systems*.
- Chin-Yew Lin. 2004. ROUGE: A package for automatic evaluation of summaries. *Text Summarization Branches Out: ACL Workshop*.
- Guanxiong Liu, Tzu-Ming Harry Hsu, Matthew McDermott, Willie Boag, Wei-Hung Weng, Peter Szolovits, and Marzyeh Ghassemi. 2019. Clinically accurate chest X-ray report generation. *arXiv preprint arXiv:1904.02633*.
- Peter J. Liu, Mohammad Saleh, Etienne Pot, Ben Goodrich, Ryan Sepassi, Lukasz Kaiser, and Noam Shazeer. 2018. Generating wikipedia by summarizing long sequences. *The 2018 International Conference for Learning Representations*.
- Sean MacAvaney, Sajad Sotudeh, Arman Cohan, Nazli Goharian, Ish Talati, and Ross W. Filice. 2019. Ontology-aware clinical abstractive summarization. *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR 19)*.
- Christopher D Manning, Prabhakar Raghavan, and Hinrich Schütze. 2008. *Introduction to information retrieval*. Cambridge University Press.
- Christopher D. Manning, Mihai Surdeanu, John Bauer, Jenny Finkel, Steven J. Bethard, and David McClosky. 2014. The Stanford CoreNLP natural language processing toolkit. In *Association for Computational Linguistics (ACL) System Demonstrations*.
- Ramesh Nallapati, Feifei Zhai, and Bowen Zhou. 2016a. SummaRuNNer: A recurrent neural network based sequence model for extractive summarization of documents. *The Thirty-First AAAI Conference on Artificial Intelligence (AAAI-2017)*.
- Ramesh Nallapati, Bowen Zhou, Cicero dos Santos, Caglar Gulcehre, and Bing Xiang. 2016b. Abstractive text summarization using sequence-to-sequence RNNs and beyond. *Proceedings of The 20th SIGNLL Conference on Computational Natural Language Learning*.
- Romain Paulus, Caiming Xiong, and Richard Socher. 2018. A deep reinforced model for abstractive summarization. In *International Conference on Learning Representations*.
- Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. GloVe: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*.

- Marc’Aurelio Ranzato, Sumit Chopra, Michael Auli, and Wojciech Zaremba. 2016. Sequence level training with recurrent neural networks. *The 2016 International Conference on Learning Representations*.
- Steven J Rennie, Etienne Marcheret, Youssef Mroueh, Jerret Ross, and Vaibhava Goel. 2017. Self-critical sequence training for image captioning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7008–7024.
- Abigail See, Peter J Liu, and Christopher D Manning. 2017. Get to the point: Summarization with pointer-generator networks. In *The 2017 Annual Meeting of the Association of Computational Linguistics (ACL 2017)*.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15(1):1929–1958.
- Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. 2015. Pointer networks. In *Advances in Neural Information Processing Systems*, pages 2692–2700.
- Ronald J Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, 8(3-4):229–256.
- Ronald J Williams and David Zipser. 1989. A learning algorithm for continually running fully recurrent neural networks. *Neural Computation*, 1(2):270–280.
- Sam Wiseman, Stuart Shieber, and Alexander Rush. 2017. Challenges in data-to-document generation. *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing (EMNLP 2017)*.
- Zhilin Yang, Zihang Dai, Ruslan Salakhutdinov, and William W. Cohen. 2018. Breaking the softmax bottleneck: A high-rank RNN language model. *The 2018 International Conference for Learning Representations*.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. 2020. BERTScore: Evaluating text generation with BERT. In *The 2020 International Conference for Learning Representations*.
- Yuhao Zhang, Daisy Yi Ding, Tianpei Qian, Christopher D. Manning, and Curtis P. Langlotz. 2018. Learning to summarize radiology findings. In *EMNLP 2018 Workshop on Health Text Mining and Information Analysis*.

Split	Time Coverage	
	Stanford	RIH
Train	2009/01 – 2014/04	2017/11 – 2018/06
Dev	2014/05 – 2014/08	2018/07 – 2018/09
Test	2014/09 – 2014/12	2018/10 – 2018/12

Table 6: Time coverage of different splits in the Stanford and RIH datasets.

A Clinical Variables Inclusion Criteria

While the CheXpert labeler that we use is able to extract status for 14 clinical variables, we found that several variables are very rarely represented in our corpora and therefore using all of them makes the calculation of the factual F_1 score very unstable. For example, we found that training the same model using different random initializations would result in highly varying F_1 scores for these variables. For this reason, for both datasets we removed from the factual F_1 calculation all variables which have less than 3% positive occurrences on the validation set. We further removed the variables “Pleural Other” and “Support Devices” due to their ambiguity. This process results in a total of 9 variables for the Stanford dataset and 8 for the RIH dataset.

Additionally, apart from the positive and negative status, the CheXpert labeler is also able to generate an *uncertain* status for a variable, capturing observations with uncertainty, such as in the sentence “*pneumonia is likely represented*”. While we can modify the factual accuracy score to take uncertainty into account, for simplicity in this work we do not make the distinction between a positive status and an uncertain status.

B Dataset Preprocessing and Stratification Details

We preprocessed both the Stanford and the RIH datasets following Zhang et al. (2018). All reports were first tokenized with Stanford CoreNLP (Manning et al., 2014). We then filtered the datasets by excluding reports where (1) no findings or impression (i.e., summary) section can be found; (2) multiple findings or impression sections can be found but cannot be aligned; or (3) the findings have fewer than 10 words or the impression has fewer than 2 words. Lastly, we replaced all date and time mentions with special tokens (e.g., $\langle DATE \rangle$).

For both datasets, we stratified them over time

Variable	PG Baseline	RL _{R+C}	Δ
No Finding	91.0	92.0	+1.0*
Cardiomegaly	21.1	33.8	+12.7*
Airspace Opacity	80.4	83.5	+3.1*
Edema	73.4	80.2	+6.8*
Pneumonia	63.5	69.2	+5.7*
Atelectasis	60.5	66.5	+6.0*
Pneumothorax	89.7	93.2	+3.5*
Pleural Effusion	74.3	79.9	+5.6*
Macro Avg.	69.3	74.8	+5.5*

Table 7: Test set performance for each variable on the RIH dataset. All numbers are F_1 scores. * marks statistically significant improvements with $p < .01$ under a bootstrap test.

into training, dev and test splits. We employed this stratification strategy to test whether our model generalizes to future data when trained on historical data. We show the time coverage of each split in Table 6.

C Model Implementation and Training Details

For the baseline background-augmented pointer-generator model, we use its open implementation.⁶ We use a 2-layer LSTM as the findings encoder, 1-layer LSTM as the background encoder, and a 1-layer LSTM as the decoder. For all LSTMs we use a hidden size of 200. For the embedding layer we use 100-dimensional GloVe vectors (Pennington et al., 2014) which we pretrained on about 4 million radiology reports. We apply dropout (Srivastava et al., 2014) with $p = 0.5$ to the embeddings. At decoding time, we use the standard beam search with a beam size of 5 and a maximum decoding length of 50.

For the training and finetuning of the models, we use the Adam optimizer (Kingma and Ba, 2015) with an initial learning rate of $1e^{-3}$. We use a batch size of 64 and clip the gradient with a norm of 5. During training we evaluate the model on the dev set every 500 steps and decay the learning rate by 0.5 whenever the validation score does not increase after 2500 steps. Since we want the model outputs to have both high overlap with the human references and high factual correctness, for training we always use the average of the dev ROUGE score

⁶<https://github.com/yuhaozhang/summarize-radiology-findings>

Stanford Dataset
<p>Background: radiographic examination of the chest: <date> <time>. clinical history: <age> years of age, with concern for pulmonary edema. procedure comments: 3 single views of the chest...</p> <p>Findings: in the first chest radiograph from <date> at <time> there is interval intubation. left arm-picc line remains in place. grossly unchanged persistent cardiomegaly, bilateral pleural effusion, and mild pulmonary edema. severe djd of the left gh joint is noted. in the second chest radiograph there is interval placement of a trialysis catheter in the left ij. no other significant changes are noted. in the third chest radiograph from <date> at <time> there is an increased left basilar opacity likely reflecting basilar consolidation, atelectasis or aspiration.</p> <p>Human: in the final chest radiograph there is increased left basilar opacity likely reflecting basilar consolidation, atelectasis or aspiration.</p> <p>PG Baseline: interval intubation with placement of a trialysis catheter in the left ij. grossly unchanged cardiomegaly, bilateral pleural effusion, and mild pulmonary edema.</p> <p>RL_{R+C}: interval placement of a trialysis catheter in the left ij. an increased left basilar opacity likely reflecting basilar consolidation, atelectasis or aspiration or aspiration.</p>
RIH Dataset
<p>Background: post op cardiac surgery - check lines and tubes. technique: single view of the chest obtained at <time> <date>...</p> <p>Findings: lines/tubes: right ij sheath with central venous catheter tip overlying the svc. on initial radiograph, endotracheal tube between the clavicular heads, and enteric tube with side port at the ge junction and tip below the diaphragm off the field-of-view; these are removed on subsequent film. mediastinal drains and left thoracostomy tube are unchanged. lungs: low lung volumes. retrocardiac airspace disease, slightly increased on most recent film. pleura: small left pleural effusion. no pneumothorax. heart and mediastinum: postsurgical widening of the cardiomeastinal silhouette. aortic arch calcification. bones: intact median sternotomy wires.</p> <p>Human: left basilar airspace disease and small left pleural effusion. lines and tubes positioned as above.</p> <p>PG Baseline: lines and tubes as above. retrocardiac airspace disease, which may represent atelectasis, aspiration, or pneumonia.</p> <p>RL_{R+C}: lines and tubes as described above. retrocardiac airspace disease, slightly increased on most recent film. small left pleural effusion.</p>

Figure 5: More examples from the test splits of both datasets along with human, PG baseline and RL_{R+C} summaries. In the first example, the baseline output successfully copied content from the context, but missed [important observations](#). In the second example, the baseline output included some [spurious facts](#) that were not mentioned, and again neglected some [important observations](#). In neither examples the RL_{R+C} outputs make perfect summaries, but they represent better summaries than the baseline outputs.

and the dev factual F_1 score as the stopping criteria. We tune the scalar weights in the loss function on the dev sets and use weights of $\lambda_1 = 0.97$, $\lambda_2 = 0.97$ and $\lambda_3 = 0.03$ for both datasets.

For the extractive LexRank and BanditSum models, we use their open implementations.⁷ For the BanditSum extractive summarization model, we use default values for all hyperparameters as in [Dong et al. \(2018\)](#). For both models we select the top 3 scored sentences to form the summary, which yields the highest ROUGE-L scores on the dev sets.

For ROUGE evaluation, we use the Python

⁷<https://github.com/miso-belica/sumy>;
<https://github.com/yuedongP/BanditSum>

ROUGE implementation released by Google Research.⁸ We empirically find it to provide very close results to the original Perl ROUGE implementation by [Lin \(2004\)](#).

D Fine-grained Correctness Results on the RIH Dataset

We show the fine-grained factual F_1 scores for all variables on the RIH dataset in Table 7.

E More Examples with Baseline and System Generations

In Figure 5 we present more examples from both datasets along with the generations from the baseline system and our approach.

⁸<https://github.com/google-research/google-research/tree/master/rouge>