Quality of syntactic implication of RL-based sentence summarization

Hoa T Le, Christophe Cerisara, Claire Gardent

► To cite this version:

Hoa T Le, Christophe Cerisara, Claire Gardent. Quality of syntactic implication of RL-based sentence summarization. AAAI Workshop on Engineering Dependable and Secure Machine Learning Systems 2020, Feb 2020, New York, United States. hal-02883327

HAL Id: hal-02883327

<https://hal.science/hal-02883327>

Submitted on 29 Jun 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Quality of syntactic implication of RL-based sentence summarization

Hoa T. LE¹, Christophe Cerisara², Claire Gardent²

¹ LORIA, UMR 7503, Nancy, France; ² CNRS-LORIA, France

Abstract

Work on summarization has explored both reinforcement learning (RL) optimization using ROUGE as a reward and syntax-aware models, such as models whose input is enriched with part-of-speech (POS)-tags and dependency information. However, it is not clear what is the respective impact of these approaches beyond the standard ROUGE evaluation metric. Especially, RL-based for summarization is becoming more and more popular. In this paper, we provide a detailed comparison of these two approaches and of their combination along several dimensions that relate to the perceived quality of the generated summaries: number of repeated words, distribution of part-of-speech tags, impact of sentence length, relevance and grammaticality. Using the standard Gigaword sentence summarization task, we compare an RL self-critical sequence training (SCST) method with syntax-aware models that leverage POS tags and Dependency information. We show that on all qualitative evaluations, the combined model gives the best results, but also that only training with RL and without any syntactic information already gives nearly as good results as syntax-aware models with less parameters and faster training convergence.

Introduction

Early neural approaches to text generation tasks such as machine translation, summarization and image captioning mostly relied on sequence-to-sequence models (Sutskever, Vinyals, and Le 2014) where the model was trained using cross-entropy and features were learned automatically. More recent work however, shows that using reinforcement learning or explicitly enriching the input with additional linguistic features helps to improve performance.

Reinforcement learning was proposed to address two shortcomings of cross entropy training. First, there is a discrepancy between how the model is trained (conditioned on the ground truth) and used at test time (using argmax or beam search), namely the *exposure bias* problem. Second, the evaluation metrics (for ex. ROUGE, METEOR, BLEU, etc.) differ from the objective that is maximized with the standard cross-entropy on each token; this is known as the *loss mismatch* problem. Typically, RL is used to

optimize task-specific objectives such as ROUGE for text summarization systems (Paulus, Xiong, and Socher 2018; Narayan, Cohen, and Lapata 2018; Celikyilmaz et al. 2018; Pasunuru and Bansal 2018) and SARI (Xu et al. 2016) for sentence simplification models (Zhang and Lapata 2017).

Similarly, while neural networks allow for features to be learned automatically, explicitly enriching the input with linguistic features was repeatedly shown to improve performance. For instance, (Sennrich and Haddow 2016; Li et al. 2017) show that adding morphological features, part-of-speech (POS) tags, syntactic dependency and or parse trees as input features improves the performance of neural machine translation (NMT) systems; and (Nallapati et al. 2016) that integrating linguistic features such as POS tags and named-entities helps to improve summarization.

In this paper, we explore the relative impact of these two approaches on sentence summarization. More precisely, we assess and compare the quality of the summaries generated by syntax-aware, RL-trained and combined models with regard to several qualitative aspects that strongly impact the perceived quality of the generated texts: number of repetitions, sentence length, distribution of part-of-speech tags, relevance and grammaticality.

Using the standard Gigaword benchmark corpus, we compare and combine an RL self-critical sequence training (SCST) method with syntax-aware models that leverage POS tags and/or dependency information. We show that both enhancements, syntactic information and RL training, benefit to a sequence-to-sequence summarization model with attention and copy-pointer mechanism. While the combined model gives the best quality of summaries, we also show that the reinforcement learning approach alone may be preferred when computational complexity is an issue, as it gives nearly as good results as the syntax-aware model but with less parameters and faster training convergence.

Related Work

We briefly discuss previous work on syntax-aware and RL-based models for text-to-text generation focusing on summarization and NMT and we position our work with respect to these approaches.

Syntax models: Explicit modeling of syntax has frequently

been used in text generation applications in particular, for NMT and summarization. Thus (Sennrich and Haddow 2016) enrich the input to NMT with dependency labels, POS tags, subword tags and lemmas so that each input token is represented by the concatenation of word and features embeddings. Similarly, (Nallapati et al. 2016) enrich the encoder side of a neural summarization model with POS tag, NER tag, TF-IDF features. The intuition is that words will be better disambiguated by taking syntactic context into account. Speculating that full parse trees can be more beneficial for NMT than shallow syntactic information, (Li et al. 2017) enrich the input with a linearization of its parse tree and compares three ways of integrating parse tree and word information (parallel, hierarchical and mixed). Other works have focused on integrating syntax in the decoder or through multi-task learning. Thus, (Le et al. 2017) defines machine translation as a sequence-to-dependency task in which the decoder generates both words and a linearized dependency tree while (Kiperwasser and Ballesteros 2018) propose a scheduled multi-task learning approach where the main task is translation but the model is alternatively trained on POS tag, Dependency Tree and translation sequences.

Our model for syntax-aware summarization is similar to (Li et al. 2017) in that we use a hierarchical model to integrate syntax in the encoder. We depart from it in that (i) we enrich the input with POS tag and/or dependency information rather than constituency parse trees; (ii) we apply our model to summarization rather than translation.

RL sequence models: Various RL models have been proposed for sequence-to-sequence models. (Ranzato et al. 2016) introduce an adaptation of REINFORCE (Williams 1992) to sequence model and a curriculum learning strategy to alternate between ground truth and the sample from the RL model. This vanilla REINFORCE is known to have high variance. Thus, a learned baseline is equipped to mitigate this issue. (Bahdanau et al. 2016) propose another reinforcement learning model, namely actor-critic, to have lower variance of model estimations, offsetting by a little bias. The impact of the bias, as well as the goodness of the model, relies particularly on the careful design of the *critic*. In practice, to ensure convergence, intricate techniques must be used including an additional target network Q' , delayed actor, a critic value penalizing term and reward shaping. In contrast, (Rennie et al. 2017) introduce a very simple and effective way to construct a better baseline for REINFORCE, namely self-critical sequence training (SCST) method. Instead of looking for and constructing a baseline or using a real critic as above, SCST uses its own prediction normally used at inference time to construct the sequence and uses this to normalize the reward. (Kryscinski et al. 2018) adapt this training method to improve the abstraction of text summarization via a combination of cross-entropy, policy learning, pretrained language model and novel phrase reward. Similarly, we use SCST to train a summarization model. However, our model uses ROUGE as a reward and we focus on comparing models trained using different learning strategies (RL vs Cross Entropy) and informed by different sources (with and without syntax).

Models

We train and compare models that differ along two main dimensions: training (cross-entropy vs. RL) and syntax (with and without syntactic information).

Baseline

The baseline is a sequence-to-sequence model consisting of a bidirectional LSTM encoder and a decoder equipped with an attention and a copy pointer-generator mechanism.

Encoder. The source sentence is encoded using two recurrent neural networks (denoted *bi-RNN*) (Chung et al. 2014): one reads the whole sequence of words $x = (x_1, \dots, x_m)$ from left to right and the other from right to left. This results in a forward and backward sequence of hidden states $(\vec{h}_1, \dots, \vec{h}_m)$ and $(\overleftarrow{h}_1, \dots, \overleftarrow{h}_m)$ respectively. The representation of each source word x_j is the concatenation of the hidden states $h_j = [\vec{h}_j, \overleftarrow{h}_j]$.

Decoder. An RNN is used to predict the target summary $y = (y_1, \dots, y_n)$. At each decoder timestep, a multi-layer perceptron (MLP) takes as input the recurrent hidden state s_i , the previously predicted word y_{i-1} and a source-side context vector c_i to predict the target word y_i . c_i is the weighted sum of the source-side vectors $(h_1, \dots, h_m)$. The weights in this sum are obtained with an attention model (Bahdanau, Cho, and Bengio 2014), which computes the similarity between the target vector s_{i-1} and every source vector h_j .

Copy-Pointer. The attention encoder-decoder tends to ignore the presence of rare words in the source, which might be important especially for the summarization task. The Copy-Pointer (See, Liu, and Manning 2017) enables to either copy source words via a pointer or generate words from a fixed vocabulary. A soft switch p_{gen} is learned to choose between *generating* a word from the vocabulary by sampling from the output distribution of the decoder, or *copying* a word from the input sequence by sampling from the attention distribution.

Integrating Syntax

Conventional neural summarization models only rely on the sequence of raw words and ignore syntax information. We include syntactic information in our summarization model using the hierarchical-RNN topology introduced by (Li et al. 2017) and comparing three sources of information: POS tags (Postag), dependency relations (Deptag) and their combination (Pos+Deptag). Figure 1 shows a graphical depiction of the Pos+tag model. In essence, each source of information (sequence of tokens, of POS tags, of dependency relations) is encoded using a bidirectional LSTM and each input token is then represented by the concatenation of the hidden-states produced by each information source considered. For instance, in the Postag model, the POS tag bi-LSTM takes as input a sequence of POS tags and outputs a sequence of hidden states $(hp_j = [\vec{hp}_j, \overleftarrow{hp}_j])_{1 \leq j \leq m}$ similarly to the word bi-RNN. Each hp_j is then concatenated with the input word embeddings ew_j and passed on to the word bi-RNN.

Figure 1: Pos+Deptag Model

For the Deptag model, the input sequence to the Deptag bi-LSTM includes, for each input tokens, the name of the dependency relation that relates this token to its syntactic head (e.g., *nsubj* for the token “Mark” in the sentence shown at the top of Figure 1). The Deptag bi-LSTM then output a sequence of hidden states $(hd_j = [\overrightarrow{hd_j}; \overleftarrow{hd_j}])_{1 \leq j \leq m}$ which are concatenated with the corresponding word embeddings ew_j and passed to the word bi-RNN.

Finally, for the Pos+Deptag model, both the POS tag and the syntactic hidden states are concatenated with the words embeddings to give the final input vector $[\overrightarrow{ew_j}; \overleftarrow{ew_j}; \overrightarrow{hp_j}; \overleftarrow{hp_j}; \overrightarrow{hd_j}; \overleftarrow{hd_j}]$ which is passed on to the upper-level word bi-RNN.

RL Learning

Summarization as an RL problem Neural summarization models are traditionally trained using the cross entropy loss. (Ranzato et al. 2016) propose to directly optimize Natural Language Processing (NLP) metrics by casting sequence generation as a Reinforcement Learning problem. As most NLP metrics (BLEU, ROUGE, METEOR,...) are non-differentiable, RL is appropriate to reach this objective. The parameters θ of the neural network define a *natural policy* p_θ , which can be used to predict the next word. At each iteration, the decoder RNN updates its internal state (hidden states, attention weights, copy-pointer weights...). After generating the whole sequence, a reward $r(\cdot)$ is computed, for instance the ROUGE score. This reward is evaluated by comparing the generated sequence and the gold sequence. The RL training objective is to minimize the negative expected reward:

$$L^{RL}(\theta) = -\mathbb{E}_{w^s \sim p_\theta} [r(w^s)] \quad (1)$$

where $w^s = (w_1^s, \dots, w_T^s)$ and w_t^s is the word sampled from the decoder at time step t . Following (Williams 1992), the gradient $\nabla_\theta L^{RL}(\theta)$ can be computed as follows:

$$\nabla_\theta L^{RL}(\theta) = -\mathbb{E}_{w^s \sim p_\theta} [r(w^s) \nabla_\theta \log p_\theta(w^s)] \quad (2)$$

In practice, the vanilla REINFORCE yields a very high variance during training. In order to help the model stabilize and converge to good local optima, vanilla REINFORCE is extended to compute the reward *relative* to a *baseline* b :

$$\nabla_\theta L^{RL}(\theta) = -\mathbb{E}_{w^s \sim p_\theta} [(r(w^s) - b) \nabla_\theta \log p_\theta(w^s)] \quad (3)$$

This baseline can be an arbitrary function (function of θ or t), as long as it does not depend on w^s (Sutton and Barto 1998).

Self-critical sequence training There are various ways to reduce RL variance and choose a proper baseline: for instance, using a second decoder (Ranzato et al. 2016) or building a *critic network* and optimizing with a *value network* instead of real reward (Bahdanau et al. 2016). In the following, we have chosen the self-critical sequence training (SCST) technique (Rennie et al. 2017), which has been shown to be very simple and effective. The main idea of SCST is to use, as baseline in the vanilla REINFORCE algorithm, the reward obtained with the inference algorithm used at test time. Equation 3 then becomes:

$$\nabla_\theta L^{RL}(\theta) = -\mathbb{E}_{w^s \sim p_\theta} [(r(w^s) - r(\hat{w})) \nabla_\theta \log p_\theta(w^s)] \quad (4)$$

where $r(\hat{w})$ is the reward obtained by the current model with the inference algorithm used at test time. As demonstrated by (Zaremba and Sutskever 2015), we can rewrite this gradient formula as:

$$\frac{\partial L^{RL}(\theta)}{\partial s_t} = (r(w^s) - r(\hat{w})) (p_\theta(w_t | w_{t-1}^s, h_t) - 1_{w_t^s}) \quad (5)$$

where s_t is the input to the final softmax function in the decoder. The term on the right side resembles logistic regression, except that the ground truth w_t is replaced by sampling w_t^s . In logistic regression, the gradient is the difference between the prediction and the actual 1-of-N representation of the target word:

$$\frac{\partial L^{XENT}(\theta)}{\partial s_t} = p_\theta(w_t | w_{t-1}, h_t) - 1_{w_t} \quad (6)$$

We see that samples that return a higher reward than $r(\hat{w})$ will be encouraged while samples that result in a lower reward will be discouraged. Therefore, SCST intuitively tackles well the exposure bias problem as it forces to improve the performance of the model with the inference algorithm used at test time. In order to speed up sequence evaluation at training time, we use greedy decoding with $\hat{w}_t = \arg \max_{w_t} p(w_t | h_t)$.¹

¹Two variants of this training method exist: TD-SCST and the “True” SCST, but both variants do not lead to significant additional

Training objective and Reward The number of words in the vocabulary may be quite large in text generation, which leads to a large state space that may be challenging for reinforcement learning to explore. To reduce this effect, we follow (Kryscinski et al. 2018) and adopt a final loss that is a linear combination of the cross-entropy loss and the policy learning loss:

$$L = (1 - \alpha)L^{XENT} + \alpha L^{RL} \quad (7)$$

α is a hyper-parameter that is tuned on the development set.

We use ROUGE- F_1 as the reward for the reinforce agent as the generation should be as concise as the gold target sentence.

Experiments

Data We evaluate our approach on the Gigaword corpus (Rush, Chopra, and Weston 2015), a corpus of 3.8M sentence-headline pairs and where the average input sentence length is 31.4 words (in the training corpus) and the average summary length is 8.3 words. The test set consists of 1951 sentence/summary pairs. As (Rush, Chopra, and Weston 2015), we use 2000 sample pairs (among 189K pairs) as development set.

Automatic Evaluation Metric We adopt ROUGE (Lin 2004) for automatic evaluation. It measures the quality of the summary by computing overlapping lexical units between the candidate and gold summaries. We report ROUGE-1 (unigram), ROUGE-2 (bi-gram) and ROUGE-L (longest common sequence) F1 scores. ROUGE-1 and ROUGE-2 mainly represent informativeness while ROUGE-L is supposed to be related to readability (Cao et al. 2018b).

Implementation Our models implementations are based on the Fast-Abs-RL (Chen and Bansal 2018) code ². Although this code is not optimized to give the best possible performances, it is flexible enough to allow for the integration of syntactic features.

The hyperparameter α in Eq 7 needs careful tuning. Figure 2 illustrates a problematic case when α continuously increases until it reaches $\alpha = 1$ at iteration 10^5 , where the RL models forget the previously learned patterns and degenerate. A good balance between exploration and supervised learning is thus necessary to prevent such catastrophic forgetting. We have found on the development set that the Reinforcement Learning weight α may increase linearly by (step/ 10^5) with the number of training iterations, until it reaches a maximum of 0.82 for the RL-s2s model and of 0.4 for the RL-s2s-syntax model.

For all models, we use the Adam optimizer (Kingma and Ba 2014) with a learning rate of 0.001 (tuned on the dev set). The word vocabulary size is 30k, number of part-of-speech tags 40 and number of dependency tags 244. We have

gain on image captioning (Rennie et al. 2017). So we didn’t explore these two variants for summarization as greedy decoding already obtains quite good result. We leave this for future work.

²https://github.com/ChenRocks/fast_abs_rl

chosen the default size (from the codebase) of 128 for word embeddings, and arbitrarily 30 dimensions both for the part-of-speech and dependency embeddings. Similarly, we have chosen the default values of 256 hidden states for every bi-directional RNN, 32 samples for the batch size, a gradient clipping of 2 and early-stopping on the development set. Our adapted code is given as supplementary material and will be published with an open-source licence.

Figure 2: Catastrophic forgetting of the RL decoder on the Gigaword dev set

Results and Analysis

ROUGE. Table 1 shows the performances measured in ROUGE score. State-of-the-art summarization system come from (Cao et al. 2018a), which appears to be the best system on the Gigaword corpus reported in <http://nlpprogress.com/english/summarization.html>, as of May 2019.

Both Syntactic and RL models outperform the baseline.

Syntax-aware models outperform the baseline by +1.84 (Postag), +1.86 (Deptag) and +2.01 (Dep+Postag) ROUGE-2 points.

While RL without syntax slightly under-performs syntactic models, it still achieves an improvement of +1.35 rouge-2 over the baseline. In other words, directly optimizing the ROUGE metric helps improve performance almost as much as integrating syntactic information. The combination of reinforcement learning with syntax information keeps increasing the score. However, the resulting improvement is smaller than when adding syntax without RL. We speculate that because the search space with syntax has a larger number of dimensions than without syntax, it may also be more difficult to explore with RL.

Parameters. The baseline model has 6.617M parameters. This is increased by roughly 300K paramters for the Postag and the Deptag model and correspondingly by roughly 600K paramters for the Pos+Deptag model. In comparison, RL optimization does not involve any additional parameters. However, it requires two more decoder passes for the sampling and greedy predictions.

Models	#Params	Time for 1 epoch	R-1	R-2	R-L
Re3Sum (Cao et al. 2018a)	-	-	37.04	19.03	34.46
Our Baseline s2s	6.617M	13h18m	27.57	10.29	26.02
Postag s2s	+312k	+54m	30.52	12.13	28.8
Deptag s2s	+318k	+1h13m	30.25	12.15	28.73
Pos+Deptag s2s	+630k	+1h34m	30.8	12.3	29.22
RL s2s	+0	-1h35m	29.94	11.64	28.54
RL postag s2s	-	-	30.82	12.19	29.12
RL deptag s2s	-	-	30.58	12.08	29.01
RL pos+deptag s2s	-	-	30.76	12.31	29.11

Table 1: Performance comparisons between models

Figure 3: Evolution on test set during training

Speed. Syntax-aware models are slightly longer to train than the baseline. Running on a single GPU GeForce GTX 1080, the baseline model requires 13h18m per epoch with 114k updates while the training time of syntax-aware models increases by about 6% (Postag s2s). Also, it takes one week to get the pre-processing tag labels of these syntactic features for the whole 3.8M training samples of Giga-word corpus on 16 cores cpu machine Dell Precision Tower 7810. Surprisingly, adding the RL loss (which requires re-evaluating the ROUGE score for every generated text at every timestep) reduces training time by 12%. We speculate that the RL loss may act as a regularizer by smoothing the search space and help gradient descent to converge faster.

Learning Curve. Figure 3 shows the evolution of Rouge-2 on the test set over 1 epoch. We can observe that syntactic models obtain a better performance than the baseline after the first 6k iterations. Sequence models with RL also quickly reach the same performance than syntactic models, even though the RL loss is only weakly taken into account at the start of training. As learning continues, the gap between the top models (with syntax and/or RL) and the baseline stays constant. The increased speed of training with RL, especially at the beginning of training, constitutes an important advantage in many experimental conditions.

Figure 4: Repetition comparisons by length (lower is better)

Repetitions. Repetition is a common problem for sequence-to-sequence models (Tu et al. 2016; Sankaran et al. 2016). To assess the degree of repetitions in the generated summaries, we use (Le et al. 2017)’s repetition rate metric which is defined as follows:

$$rep_rat = \sum_{i=1}^{T(y)} \frac{1 + r(\tilde{y}_i)}{1 + r(Y)} \quad (8)$$

where $\tilde{y}_i$ and Y_i are the i^{th} generated sentence and i^{th} gold abstract target sentence respectively, and r is the number of repeated words: $r(X) = len(X) - len(set(X))$. $len(X)$ is the length of sentence X and $len(set(X))$ is the number of words that are not repeated in sentence X . Figure 4 compares the repetition rate of several models; the horizontal axis is the length of sentences, and the vertical axis is the repetition rate. The proposed RL-model combined with syntactic information performs the best on long sentences, with less repeated words than all other models. Indeed, short sentences are less likely to contain repetitions, but it is interesting to observe that RL-training enriched with syntax improves the quality of long sentences on this aspect.

Analysis by Postags. To further investigate the linguistic structure of the generated output, we compute for each POS tag class T , the proportion of POS tags of type T relative to the number of generated words (on the test set). We group POS tags into 9 classes: cardinal numbers (CD), determiners (DT), nouns and proper nouns (NN), verbs (VV), adjectives

Models	Content words				Function words					MSE to gold to gold
	NN	VV	JJ	RB	CD	DT	TO	IN	SYM	
Gold target	49	12.5	12.9	1.6	1.3	1.5	2.5	10.6	4	0
Our baseline s2s	43.4	13.8	10.8	1.4	1.3	1.6	3.5	8.9	11.3	10.52
Postag s2s	50.4	14.5	12.1	1.3	1.3	1.1	3.9	8.1	2.1	2.07
Deptag s2s	49.8	14.5	12	1.7	1.3	1.1	3.7	8.9	1.9	1.59
Pos+Deptag s2s	51.4	14.7	12.6	1.6	1.1	1	3.7	7.9	1.8	2.72
RL s2s	50	14.1	11.9	1.3	1.6	1.2	4.1	9.1	1.6	1.71
RL pos+deptag s2s	49.9	14.3	12.4	1.3	1.5	1.2	3.6	9	2.2	1.28

Table 2: Proportion of generated postags

Source	the us space shuttle atlantis separated from the orbiting russian mir space station early saturday , after three days of test runs for life in a future space facility , nasa announced .
Abstract	atlantis mir part ways after three-day space collaboration by emmanuel unk
Baseline s2s	<i>atlantis atlantis atlantis separated from mir space station</i>
Postag s2s	us space shuttle atlantis separated from mir space station
Deptag s2s	atlantis separated from russian space station
RL s2s	us shuttle atlantis separated from mir
RL pos+deptag s2s	us shuttle atlantis separated from russian space station
Source	swedish telecommunications giant ericsson has reached a basic agreement to sell its relay production to japanese electronics company UNK corp , ericsson said thursday .
Abstract	ericsson sells relay production to unk 's unk corp
Baseline s2s	<i>ericsson to sell its its to sell its</i>
Postag s2s	ericsson to sell relay production to unk
Deptag s2s	ericsson reaches basic agreement to sell relay production
RL s2s	ericsson reaches basic agreement to sell relay production
RL pos+deptag s2s	ericsson sells relay production to unk corp
Source	the shooting down of the largest transport plane in the sri lankan air force has wrecked supply lines and slowed a major government offensive against the tamil rebel citadel of jaffna , analysts said .
Abstract	downing of plane slows sri lanka 's army onslaught on jaffna by amal jayasinghe
Baseline s2s	<i>sri lankan air force has</i>
Postag s2s	sri lankan air force plane shooting down
Deptag s2s	sri lanka 's largest transport plane shooting kills supply lines
RL s2s	sri lankan air force wrecked supply lines
RL pos+deptag s2s	sri lankan air force shooting down

Table 3: Generated examples of different models.

Figure 5: Performance comparisons by length

(JJ), adverbs (RB), to (TO), prepositions and subordinating conjunctions (IN) and symbols (SYM).

We evaluate whether the generated summary has a simi-

Models	Rel.	Grammar.
Baseline s2s	2.13 (+/-0.14)	2.47 (+/-0.18)
Postag s2s	3.26 (+/-0.18)	4.19 (+/-0.16)
Deptag s2s	3.17 (+/-0.19)	4.2 (+/-0.17)
RL s2s	3.23 (+/-0.19)	4.26 (+/-0.16)
RL pos+deptag	3.45 (+/-0.18)	4.49 (+/-0.13)

Table 4: Human Evaluations

lar or different POS tags distribution than the ground truth by computing for each model the mean square error (MSE) between every generated and gold POS tag class. These errors are shown in Table 2.

On average and for all POS tag classes, both syntax-aware and RL models are much closer (about 5 times) to the gold than the baseline. In a similar way as with repetitions, the best summarization model in terms of POS tag classes is the combined RL and syntax model.

Effects on Long Sentences. We group sentences of similar lengths together and compute the Rouge score. Figure 5 reports the Rouge-2 scores for various lengths of generated texts, with a 95% *t-distribution* confidence interval. It shows that the RL and syntax models perform globally better than the baseline as sentences get longer. For long sentences (more than 10 words), this effect is more pronounced, the syntax(+RL) models outperform significantly the RL and baseline models.

Human Evaluation. In order to evaluate the quality of the summaries produced by the models, we asked 3 annotators to assess the relevance and grammaticality quality of the summaries. Each criterion is rated with a score from 0 (bad) to 5 (excellent). Annotators are instructed to evaluate 50 samples randomly selected from the test set. The model information is anonymous to the annotators. The evaluation results with a 95% *t-distribution* confidence interval is shown in Table 4. We see that RL performs on par with postag, deptag on relevance and grammaticality criteria and they all outperform baseline. This is consistent with the results on POS tag classes above which indicate that these models generate more content words and less function words than the baseline. Once again, RL with pos+deptag obtains the best result.

Qualitative Analysis. Table 3 shows some sample outputs of several models. Example 1 presents a typical repetition problem (the word “atlantis”) often found in the baseline. Both syntax and RL models manage to avoid repetitions. Example 2 shows that RL (without any syntactic information) can search and find surprisingly the same structure as the syntax-aware model. In the last example, the baseline fails as it stops accidentally after a modal verb while syntax and RL models can successfully generate well-formed sentences with subject-verb-object. However, semantically, RL and RL with pos+dep tag (like the baseline model) fail to capture the true meaning of the gold summary (“transport plane” instead of “air force” should be the real subject in this case). Deptag s2s seems the best in terms of summarizing syntactic and semantic content on these examples.

Conclusion

We have studied in details in this work the quality of syntactic implication of the summaries that are generated by both syntactically-enriched summarization models and reinforcement-learning trained models, beyond the traditional ROUGE-based metric classically used to evaluate summarization tasks. We have thus focused on the quality of the generated summaries, in terms of the number of repeated words, which is a common issue with summarization systems, but also in terms of the distribution of various types of words (through their POS-tags) as compared to the gold. Because these aspects strongly depend on sentence length, we have also studied the impact of sentence length. Finally, we have manually evaluated the quality of the generated sentences in terms of relevance and grammaticality. Our results

suggest that enriching summarization models with both syntactic information and RL training improves the quality of generation in all of these aspects. Furthermore, when computational complexity is a concern, we have shown that RL-only models may be a good choice because they provide nearly as good results as syntactic-aware models but with less parameters and faster convergence time. We plan to extend this work by further applying similar qualitative evaluations to other types of summarization models and text generation tasks.

Acknowledgments

This work has been funded by Lorraine Universit dExcellence; experiments realized in this work have been done partly on Grid5000 Inria/Loria Nancy. We would like to thank all consortiums for giving us access to their resources.

References

- Bahdanau, D.; Brakel, P.; Xu, K.; Goyal, A.; Lowe, R.; Pineau, J.; Courville, A. C.; and Bengio, Y. 2016. An actor-critic algorithm for sequence prediction. *CoRR* abs/1607.07086.
- Bahdanau, D.; Cho, K.; and Bengio, Y. 2014. Neural machine translation by jointly learning to align and translate. *arXiv e-prints* abs/1409.0473.
- Cao, Z.; Li, W.; Li, S.; and Wei, F. 2018a. Retrieve, rerank and rewrite: Soft template based neural summarization. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 152–161. Melbourne, Australia: Association for Computational Linguistics.
- Cao, Z.; Wei, F.; Li, W.; and Li, S. 2018b. Faithful to the original: Fact aware neural abstractive summarization. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, 4784–4791.
- Celikyilmaz, A.; Bosselut, A.; He, X.; and Choi, Y. 2018. Deep communicating agents for abstractive summarization. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 1662–1675.
- Chen, Y.-C., and Bansal, M. 2018. Fast abstractive summarization with reinforce-selected sentence rewriting. In *Proceedings of ACL*.
- Chung, J.; Gulcehre, C.; Cho, K.; and Bengio, Y. 2014. Empirical evaluation of gated recurrent neural networks on sequence modeling. In *NIPS 2014 Workshop on Deep Learning, December 2014*.
- Kingma, D. P., and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Kiperwasser, E., and Ballesteros, M. 2018. Scheduled multi-task learning: From syntax to translation. *Transactions of the Association for Computational Linguistics* 6:225–240.

- Kryscinski, W.; Paulus, R.; Xiong, C.; and Socher, R. 2018. Improving abstraction in text summarization. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, 1808–1817.
- Le, A. N.; Martinez, A.; Yoshimoto, A.; and Matsumoto, Y. 2017. Improving sequence to sequence neural machine translation by utilizing syntactic dependency information. In *IJCNLP*.
- Li, J.; Xiong, D.; Tu, Z.; Zhu, M.; Zhang, M.; and Zhou, G. 2017. Modeling source syntax for neural machine translation. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 688–697. Vancouver, Canada: Association for Computational Linguistics.
- Lin, C.-Y. 2004. Rouge: A package for automatic evaluation of summaries. In Marie-Francine Moens, S. S., ed., *Text Summarization Branches Out: Proceedings of the ACL-04 Workshop*, 74–81. Barcelona, Spain: Association for Computational Linguistics.
- Nallapati, R.; Zhou, B.; dos Santos, C.; Gulcehre, C.; and Xiang, B. 2016. Abstractive text summarization using sequence-to-sequence RNNs and beyond. In *Proceedings of The 20th SIGNLL Conference on Computational Natural Language Learning*, 280–290. Berlin, Germany: Association for Computational Linguistics.
- Narayan, S.; Cohen, S. B.; and Lapata, M. 2018. Ranking sentences for extractive summarization with reinforcement learning. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, 1747–1759. Association for Computational Linguistics.
- Pasunuru, R., and Bansal, M. 2018. Multi-reward reinforced summarization with saliency and entailment. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 646–653.
- Paulus, R.; Xiong, C.; and Socher, R. 2018. A deep reinforced model for abstractive summarization. In *Proceedings of the 6th International Conference on Learning Representations*.
- Ranzato, M.; Chopra, S.; Auli, M.; and Zaremba, W. 2016. Sequence level training with recurrent neural networks. In *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*.
- Rennie, S. J.; Marcheret, E.; Mroueh, Y.; Ross, J.; and Goel, V. 2017. Self-critical sequence training for image captioning. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* 1179–1195.
- Rush, A. M.; Chopra, S.; and Weston, J. 2015. A neural attention model for abstractive sentence summarization. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, 379–389. Association for Computational Linguistics.
- Sankaran, B.; Mi, H.; Al-Onaizan, Y.; and Ittycheriah, A. 2016. Temporal attention model for neural machine translation. *CoRR* abs/1608.02927.
- See, A.; Liu, P. J.; and Manning, C. D. 2017. Get to the point: Summarization with pointer-generator networks. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 1073–1083. Vancouver, Canada: Association for Computational Linguistics.
- Sennrich, R., and Haddow, B. 2016. Linguistic input features improve neural machine translation. In *Proceedings of the First Conference on Machine Translation*, 83–91. Berlin, Germany: Association for Computational Linguistics.
- Sutskever, I.; Vinyals, O.; and Le, Q. V. 2014. Sequence to sequence learning with neural networks. In Ghahramani, Z.; Welling, M.; Cortes, C.; Lawrence, N. D.; and Weinberger, K. Q., eds., *Advances in Neural Information Processing Systems 27*. Curran Associates, Inc. 3104–3112.
- Sutton, R. S., and Barto, A. G. 1998. *Reinforcement Learning: An Introduction*. MIT Press.
- Tu, Z.; Lu, Z.; Liu, Y.; Liu, X.; and Li, H. 2016. Modeling coverage for neural machine translation. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 76–85. Berlin, Germany: Association for Computational Linguistics.
- Williams, R. J. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. In *Machine Learning*, 229–256.
- Xu, W.; Napoles, C.; Pavlick, E.; Chen, Q.; and Callison-Burch, C. 2016. Optimizing statistical machine translation for text simplification. *Transactions of the Association for Computational Linguistics* 4:401–415.
- Zaremba, W., and Sutskever, I. 2015. Reinforcement learning neural turing machines. *CoRR* abs/1505.00521.
- Zhang, X., and Lapata, M. 2017. Sentence simplification with deep reinforcement learning. In *Proceedings of EMNLP*.