Subject:
Fw: Your ACL-IJCNLP 2021 Submission (Number 3)
From:
Jacob Parnell <Jacob.S.Parnell@student.uts.edu.au>
Date:
2/06/2021, 8:59 am
To:
Massimo Piccardi <Massimo.Piccardi@uts.edu.au>, Inigo Jauregi <inigo.jauregi@rozettatechnology.com>

Hi guys,

The paper was accepted!
Have a read of the reviewer comments below and letÂ me know what youÂ think.

Cheers,
Jacob

From: start@z.softconf.com <start@z.softconf.com> on behalf of SPNLP 2021 Workshop Organizers <w18_SPNLP21_acl2021@softconf.com>
Sent: Wednesday, 2 June 2021 3:44 AM
To: Jacob Parnell <Jacob.S.Parnell@student.uts.edu.au>
Subject: Your ACL-IJCNLP 2021 Submission (Number 3)
Â

Dear Jacob Parnell:

On behalf of the ACL-IJCNLP 2021 Program Committee, I am delighted to inform you that the following submission has been accepted to appear at the conference:

RewardsOfSum: Exploring Reinforcement Learning Rewards for Summarisation

The Program Committee worked very hard to thoroughly review all the submitted papers. Please repay their efforts, by following their suggestions when you revise your paper.

When you are finished, you can upload your final manuscript by June 7, 2021 at the following site:

https://www.softconf.com/acl2021/w18_SPNLP21/

You will be prompted to login to your START account. If you do not see your submission, you can access it with the following passcode:

3X-E9A3B3C2F6

Alternatively, you can click on the following URL, which will take you directly to a form to submit your final paper (after logging into your account):

https://www.softconf.com/acl2021/w18_SPNLP21/user/scmd.cgi?scmd=aLogin&passcode=3X-E9A3B3C2F6

The reviews and comments are attached below. Again, try to follow their advice when you revise your paper.

Congratulations on your fine work. If you have any additional questions, please feel free to get in touch.

Best Regards,
Organizers
ACL-IJCNLP 2021

==
ACL-IJCNLP 2021 Reviews for Submission #3
==

Title: RewardsOfSum: Exploring Reinforcement Learning Rewards for Summarisation
Authors: Jacob Parnell, Inigo Jauregi Unanue and Massimo Piccardi

==
REVIEWER #1
==

Reviewer's Scores

Appropriateness (1-5): 5
Clarity (1-5): 5
Originality / Innovativeness (1-5): 2
Soundness / Correctness (1-5): 4
Meaningful Comparison (1-5): 3
Thoroughness (1-5): 3
Impact of Ideas or Results (1-5): 3
Recommendation (1-5): 3
Reviewer Confidence (1-5): 4

Detailed Comments

This paper proposes RwB-Hinge and RISK based objective for summarization. The proposed objectives as used to fine-tune a negative log likelihood trained model to show improved performance.

Strong points:
The paper is well written and easy to understand. I like the wide range of experiments conducted to evaluate the method on different length summaries. Evaluation is clear with good details on validation and model parameters.

Weak points:

1. For majority of results, except Table 8, baseline is taken to be negative log likelihood. In my opinion at least Paulus et al. 2018 should be included, maybe others like Ranzato et al. (2015). This way we can compare clearly between non NLL methods.
2. The proposed RwB-Hinge method is a minor variation compared to the baseline. RISK method is similar to cost based classical risk loss.
3. I found it vague to understand what is the main aim to optimize for this objective. If the aim is to improve diversity then why not include that in the objective? There isn't much justification about how diversity connects to the performance and the connection with methods using RL to optimize Rouge.
4. Diversity can also be looked purely from sampling and decoding point of view, example: Holtzman et al. 2020, Li et al. 2016, Li and Jurafsky, 2016. Although these is not directly related but I think there is a close connection and there should be a small discussion on this aspect too.

Minor comment:
1. Equation 10 can be re-written to avoid multiple use of notation in the same work.

Questions for Authors

What happens if we train only on proposed function?

==

REVIEWER #2

==

Reviewer's Scores

Appropriateness (1-5): 5
Clarity (1-5): 4
Originality / Innovativeness (1-5): 4
Soundness / Correctness (1-5): 4
Meaningful Comparison (1-5): 4
Thoroughness (1-5): 4
Impact of Ideas or Results (1-5): 4
Recommendation (1-5): 4
Reviewer Confidence (1-5): 3

Detailed Comments

This paper explores different learning rewards in reinforcement learning for summarization, and proposes to use two reward functions: RwB-Hinge and RISK. RwB-Hinge reward is based on the reward difference between a sampled summary and the argmax prediction of the current model, but to avoid penalizing potentially good summaries, a hinge loss is added such that the model is updated only for the predictions that are considered to be good. RISK updates based on a set of candidate summaries, and the rewards are scaled relatively to those in the candidate set.

The experiment setting of this paper follows that of Zhang et al. (2020) with nine summarization datasets of different reference summary lengths. Experiment results show that the two proposed approaches lead to improved ROUGE scores across most of the datasets, though there is no clear winner approach that works well on all datasets. For the RwB-Hinge reward, there is also no clear conclusion regarding the use fo a hinge loss (Table 8). On the other hand, both approaches lead to improvements in terms on novel uni-, bi-, and tri-grams. This suggests that reinforcement learning based on multiple sampled sentences lead to more diverse outputs in the resulting summaries.

The proposed learning objective is a linear interpolation between cross entropy loss and the reinforcement learning objective, and the experiments indicate that the coefficient for the cross entropy loss is generally high across the board (0.7 or 0.9). This indicates the necessity of learning from gold reference summaries. On the other hand, the two proposed reward functions are solely based on the candidate sequences generated through sampling. It would be an interesting direction to consider a hybrid approach where the information/scoring of the gold reference summary is also incorporated into the RwB-Hinge or RISK.

--