

Generative Adversarial Network for Abstractive Text Summarization*

Linqing Liu,¹ Yao Lu,² Min Yang,¹ Qiang Qu,^{1,4} Jia Zhu,³ Hongyan Li⁴

¹Shenzhen Institutes of Advanced Technology, Chinese Academy of Sciences

²Alberta Machine Intelligence Institute ³School of Computer Science, South China Normal University

⁴MOE Key Laboratory of Machine Perception, Peking University

{likicode, 951uyao, min.yang1129}@gmail.com, qiang@siat.ac.cn

jzhu@m.scnu.edu.cn, lihy@cis.pku.edu.cn

Abstract

In this paper, we propose an adversarial process for abstractive text summarization, in which we simultaneously train a generative model G and a discriminative model D . In particular, we build the generator G as an agent of reinforcement learning, which takes the raw text as input and predicts the abstractive summarization. We also build a discriminator which attempts to distinguish the generated summary from the ground truth summary. Extensive experiments demonstrate that our model achieves competitive ROUGE scores with the state-of-the-art methods on CNN/Daily Mail dataset. Qualitatively, we show that our model is able to generate more abstractive, readable and diverse summaries¹.

Introduction

Abstractive text summarization is the task of generating a short and concise summary that captures the salient ideas of the source text. The generated summaries potentially contain new phrases and sentences that may not appear in the source text. In the past decades, a flurry of studies have been conducted on abstractive text summarization (Nallapati et al. 2016; See, Liu, and Manning 2017; Paulus, Xiong, and Socher 2017). Despite the remarkable progress of previous studies, abstractive summarization is still challenged by (i) Neural sequence-to-sequence models tend to generate trivial and generic summary, often involving high-frequency phrases; (ii) The generated summaries have limited grammaticality and readability; (iii) In most previous work the standard sequence-to-sequence models are trained to predict the next word in summary using the maximum-likelihood estimation (MLE) objective function. However, this strategy has two major shortcomings. First, the evaluation metric is different from the training loss. Second, the input of the decoder in each time step is often from the true summary during the training. Nevertheless, in the testing phase, the input of the next time step is the previous word

generated by the decoder. This exposure bias leads to error accumulation at test time.

To address the above challenge, in this paper, we propose an adversarial framework to jointly train a generative model G and a discriminative model D . Specifically, the generator G takes the original text as input and generate the summary. We use reinforcement learning (i.e., policy gradient) to optimize G for a highly rewarded summary. Thus, it effectively bypasses exposure bias and non-differentiable task metrics issues. We implement the discriminator D as a text classifier that learns to classify the generated summaries as machine or human generated. The generator G and the discriminator D are optimized with a minimax two-player game. The discriminator D tries to distinguish the ground truth summaries from the generated summaries by the generator G , while the training procedure of generator G is to maximize the probability of D making a mistake. Thus, this adversarial process can eventually adjust G to generate plausible and high-quality abstractive summaries.

Our model

Similar to the standard training strategy (Goodfellow et al. 2014), we simultaneously train two models in an adversarial manner: a generative model G and a discriminative model D . We first pre-train the generative model by generating summaries given the source text. Then we pre-train the discriminator by providing positive examples from the human-generated summaries and the negative examples produced from the pre-trained generator. After the pre-training, the generator and discriminator are trained alternatively.

Generative Model

The generator takes the source text $x = \{w_1, w_2, \dots, w_n\}$ as input and predicts the summary $\hat{y} = \{\hat{y}_1, \hat{y}_2, \dots, \hat{y}_m\}$. Here, the n is the length of the source text x and m is the length of the predicted summary. We use a bi-directional LSTM encoder to convert the input text x into a sequence of hidden states $h = \{h_1, \dots, h_n\}$. Following (See, Liu, and Manning 2017), on time step t , an attention-based LSTM decoder is then used to compute the hidden state s_t of the decoder and a context vector c_t . The reader can refer to the supplement of this paper (or (See, Liu, and Manning 2017)) for the implementation details. The parameters of the generator G are collectively represented by θ . The context vector c_t is concatenated with the decoder state s_t and fed through a fully

*The work was partially supported by CAS Pioneer Hundred Talents Program and the MOE Key Laboratory of Machine Perception at Peking University under grant number K-2017-02. Q. Qu is the corresponding author.

Copyright © 2018, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹Supplemental material: <http://likicode.com/textsum/>

connected layer and a softmax layer to produce the probability of predicting word from target vocabulary at each time step t :

$$P_{vocab}(\hat{y}_t) = \text{softmax}(V'(V[s_t, c_t] + b) + b')$$

where V' , V , b , b' are learnable parameters. Similar to the work of (See, Liu, and Manning 2017), we incorporate a switching pointer-generator network to use either word generator from fixed vocabulary or pointer copying rare or unseen from the input sequence. Finally, we can get the final probability $P(\hat{y}_t)$ of each token \hat{y}_t in the summary.

Discriminative Model

The discriminator is a binary classifier and aims at distinguishing the input sequence as originally generated by humans or synthesized by machines. We encode the input sequence with a CNN as it shows great effectiveness in text classification (Kim 2014). We use multiple filters with varying window sizes to obtain different features and then apply a max-over-time pooling operation over the features. These pooled features are passed to a fully connected softmax layer whose output is the probability of being original.

Updating model parameters

In the adversarial process, using the discriminator as a reward function can further improve the generator iteratively by dynamically updating the discriminator. Once we obtain more realistic and high-quality summaries generated by generator G , we re-train the discriminator as:

$$\min_{\phi} -\mathbf{E}_{Y \sim p_{data}}[\log D_{\phi}(Y)] - \mathbf{E}_{Y \sim G_{\theta}}[\log(1 - D_{\phi}(Y))]$$

When the discriminator D is obtained and fixed, we are ready to update the generator G . The loss function of our generator G consists two parts: the loss computed by policy gradient (denoted by J_{pg}) and the maximum-likelihood loss (denoted by J_{ml}). Formally, the objective function of G is $J = \beta J_{pg} + (1 - \beta) J_{ml}$, where β is the scaling factor to balance the magnitude difference between J_{pg} and J_{ml} . According to the policy gradient theorem (Sutton et al. 2000), we compute the gradient of J_{pg} w.r.t. the parameters θ :

$$\begin{aligned} \nabla_{\theta} J_{pg} &= \frac{1}{T} \sum_{t=1}^T \sum_{y_t} R_D^{G_{\theta}}((Y_{1:t-1}, X), y_t) \cdot \nabla_{\theta}(G_{\theta}(y_t|Y_{1:t-1}, X)) \\ &= \frac{1}{T} \sum_{t=1}^T \mathbf{E}_{y_t \in G_{\theta}} [R_D^{G_{\theta}}((Y_{1:t-1}, X), y_t) \nabla_{\theta} \log p(y_t|Y_{1:t-1}, X)] \end{aligned}$$

where $R_D^{G_{\theta}}((Y_{1:t-1}, X), y_t)$ is the action-value function, and we have $R_D^{G_{\theta}}((Y_{1:t-1}, X), y_t) = D_{\phi}(Y_{1:T})$, T is the length of the text. We update the parameters using stochastic gradient descent, $Y_{1:t}$ is the generated summary up to time step t , X is the source text to be condensed.

Experiments

Dataset - CNN/Daily Mail Corpus. The dataset (Nallapati et al. 2016) is widely used in abstractive summarization. It comprises news stories in CNN and Daily Mail websites paired with multi-sentence human generated abstractive summaries. It contains 287,226 training pairs, 13,368 validation pairs and 11,490 test pairs.

Methods	ROUGE-1	ROUGE-2	ROUGE-L	Human
ABS	35.46	13.30	32.65	2.07
PGC	39.53	17.28	36.38	3.81
DeepRL	39.87	15.82	36.90	3.04
Pretrain	38.82	16.81	35.71	3.70
Ours	39.92	17.65	36.71	4.01

Table 1: Quantitative evaluation results

Experimental Results. We compare our approach with three methods, including the abstractive model (ABS) (Nallapati et al. 2016), the pointer-generator coverage networks (PGC) (See, Liu, and Manning 2017), and the abstractive deep reinforced model (DeepRL) (Paulus, Xiong, and Socher 2017) (ML+RL version).

We firstly compare our model with the pre-trained generator. After adversarial training, ROUGE-1, ROUGE-2, ROUGE-L increase by 1.10, 0.84 and 1.00 absolute points respectively. In addition, our model exhibits competitive Rouge scores with the state-of-the-art methods. Specifically, our approach achieves the best ROUGE-1 and ROUGE-2 scores.

We also perform human evaluation to evaluate the readability and quality of summaries. We randomly select 50 test examples from the dataset. For each example, two human evaluators are asked to rank each summary generated by all 5 models based on their readability, where 1 indicates the lowest level of readability while 5 indicates the highest level. As we can observe from Table 1, our model contributes significantly to improving the readability of summaries.

To evaluate the proposed model qualitatively, we also report the generated summaries in supplementary files.

Conclusion

In this paper, we proposed an adversarial process for abstractive text summarization. Experimental results showed that our model could generate more abstractive, readable and diverse summaries.

References

- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. In *NIPS*, 2672–2680.
- Kim, Y. 2014. Convolutional neural networks for sentence classification. *arXiv preprint arXiv:1408.5882*.
- Nallapati, R.; Zhou, B.; Gulcehre, C.; Xiang, B.; et al. 2016. Abstractive text summarization using sequence-to-sequence rnns and beyond. *arXiv preprint arXiv:1602.06023*.
- Paulus, R.; Xiong, C.; and Socher, R. 2017. A deep reinforced model for abstractive summarization. *arXiv preprint arXiv:1705.04304*.
- See, A.; Liu, P. J.; and Manning, C. D. 2017. Get to the point: Summarization with pointer-generator networks. *arXiv preprint arXiv:1704.04368*.
- Sutton, R. S.; McAllester, D. A.; Singh, S. P.; and Mansour, Y. 2000. Policy gradient methods for reinforcement learning with function approximation. In *NIPS*, 1057–1063.