

# Attractive or Faithful?

## Popularity-Reinforced Learning for Inspired Headline Generation

Yun-Zhu Song,<sup>1</sup> Hong-Han Shuai,<sup>1</sup> Sung-Lin Yeh,<sup>2</sup> Yi-Lun Wu,<sup>1</sup> Lun-Wei Ku,<sup>3</sup> Wen-Chih Peng<sup>1</sup>

<sup>1</sup>National Chiao Tung University, Taiwan, <sup>2</sup>National Tsing Hua University, Taiwan

<sup>3</sup>Academia Sinica, Taiwan

{yunzhusong.eed07g, hhshuai, w86763777.eed08g}@nctu.edu.tw,

ff936tw@gapp.nthu.edu.tw, lwku@iis.sinica.edu.tw, wcpeng@g2.nctu.edu.tw

### Abstract

With the rapid proliferation of online media sources and published news, headlines have become increasingly important for attracting readers to news articles, since users may be overwhelmed with the massive information. In this paper, we generate inspired headlines that preserve the nature of news articles and catch the eye of the reader simultaneously. The task of inspired headline generation can be viewed as a specific form of Headline Generation (HG) task, with the emphasis on creating an attractive headline from a given news article. To generate inspired headlines, we propose a novel framework called POPularity-Reinforced Learning for inspired Headline Generation (PORL-HG). PORL-HG exploits the extractive-abstractive architecture with 1) Popular Topic Attention (PTA) for guiding the extractor to select the attractive sentence from the article and 2) a popularity predictor for guiding the abstractor to rewrite the attractive sentence. Moreover, since the sentence selection of the extractor is not differentiable, techniques of reinforcement learning (RL) are utilized to bridge the gap with rewards obtained from a popularity score predictor. Through quantitative and qualitative experiments, we show that the proposed PORL-HG significantly outperforms the state-of-the-art headline generation models in terms of attractiveness evaluated by both human (71.03%) and the predictor (at least 27.60%), while the faithfulness of PORL-HG is also comparable to the state-of-the-art generation model.

### Introduction

*“A good basic selling idea, involvement and relevancy, of course, are as important as ever, but in the advertising din of today, unless you make yourself noticed and believed, you ain’t got nothing”*

— Leo Burnett (1891-1971)

Nowadays, users are overwhelmed with rapidly-increasing number of articles from not only news websites but also social media. Therefore, headlines have become more and more important for attracting readers to articles. Articles with eye-catching headlines often attract more attention and receive more views or shares, which is important to content providers since the number of views can be monetized with AD networks and bring revenue. To improve the view rate, some content farms use clickbait headlines to attract users,

e.g., “15 tweets that sum up married life perfectly. (number 13 is hilarious)”. Nevertheless, the clickbait approaches, though effective at the beginning, make users feel annoyed and eventually reluctant to read anything from these websites. It is important to generate attractive headlines while still being faithful to the content.

Headline generation can be regarded as a branch of the article summarization tasks and categorized into extractive-based methods and abstractive-based methods. Extractive-based methods generate headlines by selecting a sentence from the article (Higurashi et al. 2018). In contrast, abstractive methods generate the headline by understanding the article and summarizing the idea in one sentence (Takase et al. 2016; Hayashi and Yanagimoto 2018). These two kinds of approaches both generate faithful headlines that help readers understand the content at the first glance. However, most existing headline generation approaches do not take the attractiveness into consideration. Zhang et al. observe that interrogative headlines usually attract more clicks and thus formulate the headline generation task as Question Headline Generation (QHG) (Zhang et al. 2018). Nevertheless, the QHG approach is limited since 1) not every article is suitable for question headlines (e.g., obituaries) and 2) it looks annoying if every headline is in an interrogative form. To the best of our knowledge, this is the first work using the data-driven approach to generate both faithful and attractive headlines in a general form.

However, inspired headline generation introduces at least three new research challenges. First, there are currently only public datasets for the headline generation, and none of them contains information relating to attractiveness, e.g., views, comments, or shares. Second, even with datasets and the extractive-abstractive architecture, it is still challenging to incorporate the attractiveness and faithfulness of information for the extractive-abstractive architecture since i) attractiveness and faithfulness are evaluated based on the sentences rewritten by the abstractor but the gradient can not propagate through the non-differentiable operations of the extractor, and ii) the dependency between extractor and abstractor may lead to slow convergence, i.e., when the extractor is weak and selects a sentence without any popularity-related words, it is difficult for the abstractor to rewrite it into an attractive one. Third, generating headlines based on meaning faithfulness may sometimes conflict with attrac-

tiveness. It is challenging to strike a balance between faithfulness and attractiveness.

To tackle these challenges, we present in this paper a framework called POPularity-Reinforced Learning for inspired Headline Generation (PORK-HG), to generate attractive headlines while still preserving meaning faithfulness to the articles. Specifically, for the first challenge, we build two datasets, *CNNDM-DH* (CNN/Daily Mail-Document with Headline) and *DM-DHC* (Daily Mail-Document with Headline and Comment), based on the *CNN/Daily Mail* dataset (Hermann et al. 2015; Nallapati et al. 2016), which originally only contains documents with corresponding human written summaries. We further crawl the headlines and headlines with the number of comments for *CNNDM-DH* and *DM-DHC*, respectively.<sup>1</sup> Based on the datasets, we build a state-of-the-art popularity predictor (Lamprinidis, Hardt, and Hovy 2018) to provide the popularity information for unlabeled data.

Moreover, for the second challenge, we propose a new learning framework that exploits policy-based reinforcement learning to bridge the extractor and abstractor for propagating the attractiveness and faithfulness to the extractor. Moreover, to enhance the ability of the extractor for selecting the sentences containing popular information, we propose Popular Topic Attention (PTA) to incorporate the topic distributions of popular headlines as the auxiliary information. For the third challenge, we design a training pipeline to elegantly strike a balance between attractiveness and faithfulness to avoid generating clickbait-like headlines. Experimental results of the qualitative and quantitative analyses show that PORK-HG clearly attracts people’s attention and simultaneously preserves meaning faithfulness. The contributions of this paper are summarized as follows.

- Instead of the traditional headline generation, we propose the notion of inspired headline generation. To the best of our knowledge, this is the first work utilizing deep learning for generating headlines that are both attractive and faithful in **general form**. Moreover, the datasets will be released as a public download for future research.
- We introduce the PORK-HG, which adopts extractive-abstractive architecture. To incorporate the information of attractiveness and faithfulness for the extractor, we utilize the topic distributions of popular articles as auxiliary information and design a popularity-reinforced learning method with a training pipeline to strike the balance between attractiveness and faithfulness.
- The experimental results from both the user study and real datasets manifest that PORK-HG significantly outperforms the state-of-the-art headline generation models in terms of the attractiveness evaluated by both human (71.03%) and classifier (27.60%), while maintaining the relevance compared with the state-of-the-art headline generation models.

<sup>1</sup>Since CTR (click-through-rate) is only accessible for news platform owners, we use the comments as the popularity information to train our model, which has been proved to be highly related to CTR (Kuiken et al. 2017). The details of the datasets are discussed in Section Corpus.

## Related work

### Headline Generation

The headline generation task can be seen as a variant of the summarization task with one or two sentences, which is standardized in the DUC-2004 competitions (Over, Dang, and Harman 2007). Traditional summarization works are often statistical-based and mainly focus on extracting and compressing sentences (Knight and Marcu 2000; Cohn and Lapata 2008). Recently, with the large-scale corpora, many works have exploited neural networks (Rush, Chopra, and Weston 2015; Filippova et al. 2015; Cheng and Lapata 2016) to summarize articles via data-driven approaches. On the other hand, abstractive method generates summaries or headlines based on document comprehension, which can be considered as a machine translation task. Nallapati et al. propose several novel models to address critical problems, such as modeling keywords and capturing the sentence-to-word hierarchy structure (Nallapati et al. 2016). Since supervised learning often exhibits the “exposure bias” problem, reinforcement learning is also used for both of abstractive summarization (Paulus, Xiong, and Socher 2018; Chen and Bansal 2018) and extractive summarization (Narayan, Cohen, and Lapata 2018b). In addition, Narayan et al. generate extreme news summarization by creating an one-sentence summary answering the question “What is the article about?” (Narayan, Cohen, and Lapata 2018a). However, none of the existing approaches has considered generating attractive headlines with data-driven approaches. Zhang et al. formulate the attractive headline generation task as QHG (Question Headline Generation) (Zhang et al. 2018), according to the observation that interrogative sentences attract more clicks. However, in spite of the effectiveness achieved by a question form headline, it is still not suitable to generate every headline as a question.

### Popularity Prediction

Beside information faithfulness, the effectiveness of the headlines, i.e. click-through-rate, is also important as mentioned above. Kuiken et al. analyze the relationship between CTR and the textual/stylistic features of millions of headlines, which can provide insights for how to construct attractive headlines (Kuiken et al. 2017). Meanwhile, the online news on social media has also been analyzed by identifying the salient keyword combinations and recommending them to news editors based on the similarity between popular news headlines and keywords (Weng and Wu 2018; Szymanski, Orellana-Rodriguez, and Keane 2017). To predict the popularity of a headline, Lamprinidis et al. use an RNN with two auxiliary tasks, i.e., POS tagging and section prediction (Lamprinidis, Hardt, and Hovy 2018). Nevertheless, none of the existing approaches generate headlines in terms of popularity prediction.

Table 1: Dataset information

	Train	Val	Test
CNNDM-DH	281208	12727	10577
DM-DHC	138787	11862	10130

## Corpus

### Dataset and Headline Performance Analysis

To automatically generate a headline that is not only faithful but also eye-catching, the summarization dataset and the popularity statistical information are needed. Therefore, we build the datasets *CNN/DM-DH* (CNN/Daily Mail-Document with Headline) and *DM-DHC* (DailyMail-Document with Headline and Comment) based on *CNN/Daily Mail* dataset (Nallapati et al. 2016; Hermann et al. 2015), which contains online news articles paired with multi-sentence summaries without headlines. Hence, we access the original online news pages to crawl the headlines and popularity information for both CNN and the Daily Mail, and then remove the damaged data. Detail information is shown in Table 1.

Since only the comment counts and share counts are available for DM-DHC datasets, we first validate the idea of using them as the popularity scores for training. Following the previous research (Kuiken et al. 2017), which studies the relationship between “clickbait features” and CTR by extracting features from headlines to form 11 null hypotheses whose significance were examined using non-parametric Mann-Whitney  $U$  test<sup>2</sup>. We use the same null hypotheses to determine whether the significance of the CTR, comment counts, and share counts are similar. The result shows that the significance tests of the CTR and comment counts are almost the same (7/8), while the significance tests of the CTR and share counts exhibit more difference (6/8). Therefore, we use the comment counts as the ground truth of attractiveness for training and testing.<sup>3</sup>

### Popularity-Reinforced Learning for inspired Headline Generation

To strike a balance between faithfulness and attractiveness, we propose a novel framework called POPularity-Reinforced Learning for inspired Headline Generation (PORL-HG). The framework is shown in Figure 1. Specifically, PORL-HG first exploits a hybrid extractive-abstractive architecture to generate headline effectively, i.e., an extractor selects a candidate sentence from the article and an abstractor then rewrites the headline based on the candidate sentence. To provide the faithfulness and attractiveness information for the extractor, inspired by the pointer network (Vinyals, Fortunato, and Jaitly 2015), we propose Popular Topic Attention (PTA) by utilizing the topic distribution of the related and popular headlines. Moreover, to train the abstractor for writing eye-catching headlines, we build a headline popularity predictor by using CNN to extract features from headlines and an LSTM to predict the popularity score. The popularity score is integrated into the loss function of the abstractor to encourage the generation of attractive headlines and preserving the faithfulness simultaneously. However, this basic approach suffers from the is-

<sup>2</sup>Non-parametric Mann-Whitney  $U$  test is widely-used for the significance test of non-normal distributions.

<sup>3</sup>The details of the analysis and datasets are available at <https://github.com/yunzhusong/AAAI20-PORLHG>.




Figure 1: The framework of PORL-HG.

sue that the extractor is not trained with the popularity score to select attractive sentences. Therefore, the reinforcement learning (RL) techniques are used to train the extractor. In the following, we introduce each module in PORL-HG, and present the training pipeline afterward.

### Inspired Extractor

Given an article, the goal of the extractor is to choose a salient sentence from the article for the following rewriting task of the abstractor, which requires 1) the sentence representation and 2) the ground truth of the salient sentence for training. For the sentence representation, we first embed each word by word2vec and concatenate the word embeddings for each sentence in the article. Afterward, we exploit a convolutional neural network (Kim 2014) with different kernel sizes to capture the complete semantic meaning of each sentence, which is denoted as  $r_k$  for the  $k$ -th sentence in the article. However, the long-term relationship between sentences is not captured for generating the sentence embeddings. Therefore, a bidirectional LSTM is then applied to improve the embedding representation  $r_k$ , which is denoted as  $s_k$ . For the ground truth of the salient sentence, we use a proxy label by calculating *ROUGE-L* score of each sentence in the article (Chen and Bansal 2018; Nallapati, Zhai, and Zhou 2017) and marking the highest one as the proxy training label.

After deriving the sentence embedding and proxy label, one basic approach is to use existing extractive summarization models for selecting a sentence. However, since the proxy label only considers the faithfulness (*ROUGE-L*), the selected sentence may not contain any attractive keywords, which makes the rewriting task of the abstractor difficult. To solve this issue, we propose Popular Topic Attention (PTA) that uses attention with the popularity information as a pointer to select a sentence from the article. Specifically, we first exploit Latent Dirichlet Allocation (Blei, Ng, and Jordan 2003) to generate the topic distributions of articles and headlines. Let  $\theta_i^D$  and  $\theta_i^H$  denote the topic distributions of  $i$ -th article  $D_i$  and the corresponding headline  $H_i$ , respectively. Afterward, for input article  $D_i$ , we calculate the inner product of  $\theta_i^D$  and  $\theta_j^D$  ( $\forall j \neq i$ ), and retrieve the top- $m$  similar articles. Among the top- $m$  similar articles, the most popular article  $D_{j^*}$  is selected and the corresponding topic distribution of the headline, i.e.,  $\theta_{j^*}^H$ , is used as a reference. For each sentence  $s_k$ , we use  $\theta_k^S \otimes \theta_{j^*}^H$  as the popularity information, where  $\otimes$  denotes the element-wise multiplication. The operation preserves the topics appeared in both of  $\theta_k^S$  and  $\theta_{j^*}^H$ , allowing the model to find the topics that are both faithful and attractive.

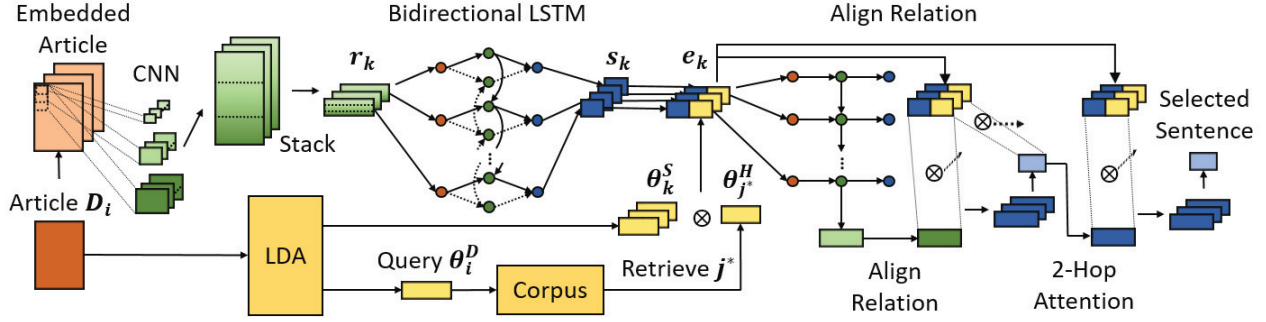


Figure 2: Extractor agent with PTA: a CNN is used as the sentence representation encoder, and sentence embedding  $r_k$  is further encoded by the bidirectional LSTM to acquire the context-aware sentence representation  $s_k$ . The popular topic distribution  $\theta_{j^*}^H$  related to the input article is multiplied with the topic distribution of sentences and concatenated with the  $s_k$ . Then, the last state of  $e_k$  is fed to the pointer for guiding the sentence selection.

To provide both faithfulness and attractiveness information for sentence selection, for each sentence  $s_k$ , PTA constructs  $e_k$  by concatenating  $s_k$  and  $\theta_k^S \otimes \theta_{j^*}^H$ , i.e.,

$$e_k = [s_k; \theta_k^S \otimes \theta_{j^*}^H],$$

where  $[\cdot; \cdot]$  denotes the concatenation. Then, the sentence information  $e_k$  will be fed into the sentence selector to make the selection, which is an LSTM equipped with 2-hop attention. The first attention step is to get the context vector  $c$  by applying the glimpse operation to every sentence  $e_k$  as follows:

$$\begin{aligned} u_k &= \nu_g^T \tanh(W_{g_1} e_k + W_{g_2} z) \\ \alpha_k &= \text{softmax}(u_k) \\ c &= \sum_k \alpha_k W e_k \end{aligned}$$

where  $z$  is the initial state of the LSTM and  $W$ ,  $W_{g_1}$ ,  $W_{g_2}$  and  $\nu_g$  are all trainable parameters, and  $\alpha_k$  is the attention coefficient of  $k$ -th sentence for deriving  $c$ . After this, the second attention step is to attend  $e_k$  again by the context vector  $c$ , which results  $o_k$ , i.e.,

$$o_k = \nu_p^T \tanh(W_{p_1} e_k + W_{p_2} c),$$

where  $W_{p_1}$ ,  $W_{p_2}$  and  $\nu_p$  are also trainable parameters for deriving the sentence selection probability. Let  $o$  denote the output vector composed of  $o_k$  for each sentence  $k$ . Finally, the extraction probability  $P(k)$  of extracting the  $k$ -th sentence can be obtained from  $o$  by using the softmax function as

$$P(k) = \text{softmax}(o). \quad (1)$$

Figure 2 shows the network architecture of the inspired extractor with the PTA mechanism. Specifically, this approach differs from the previous bidirectional LSTM that generates sentence representation, we add a new mechanism to encourage the model to select a sentence that is not only informative but also eye-catching.

## Inspired Abstractor

The abstractor compresses and paraphrases an extracted article sentence to a headline sentence, for which, we use the standard encoder-aligner-decoder with attention mechanism. To deal with the out-of-vocabulary (OOV) words, we apply a copy mechanism (See, Liu, and Manning 2017), which can directly copy words from the article. Additionally, to make the headline more eye-catching, the rewriting ability is also important. One basic approach of training the abstractor for generating attractive headlines is to construct a corpus containing multiple headlines with the corresponding popularity scores for the same article. However, deriving such datasets costs highly. In addition, the distribution of popularity scores (comment counts) follows a long-tailed distribution, which makes the prediction biased. Therefore, we transform the popularity score into a binary label, where 0 represents the popularity score which is smaller than the median and 1 otherwise. Next, we pre-train the binary classifier by using headlines, which are encoded with the same CNN used in the extractor. Then, we classify the results through an LSTM. The classification score will be returned to the abstractor as an auxiliary loss or reward.

## Training Pipeline

When training the POREL-HG, the gradient derived in the abstractor cannot be propagated back to the extractor. In order to perform an end-to-end training and to combine the popularity information, we apply the reinforcement learning with the standard policy gradient to connect the extractor, abstractor and auxiliary classifier. It is worth noting that training from a random initialization is difficult due to the dependent interplay between the extractor and abstractor, e.g., the abstractor cannot learn rewriting an attractive sentence when the extractor is not well-trained and thus selects meaningless sentences. Moreover, an abstractor without a good rewriting ability leads to a noisy estimation of the standard policy gradient, which deteriorates the training of the extractor. Hence, pre-training the abstractor, extractor and classifier before starting the reinforcement learning is necessary.

**Extractor Training** The task for the inspired extractor is to select an essential and eye-catching sentence from the article. Since most of the headline generation dataset does not include the extracted headline labels, we offer the proxy label similar to (Nallapati, Zhai, and Zhou 2017) for the extractive summarization task. The label is acquired by calculating *ROUGE-L* score for every sentence with respect to the ground-truth headline  $H_i$ . That is, the proxy target label  $y_i$  for the  $i$ -th article  $D_i$  is obtained by  $y_i = \operatorname{argmax}(ROUGE-L_{recall}(D_i, H_i))$ , and the loss function of the extractor is:

$$L_{ext} = -\frac{1}{N} \sum_{i=1}^N y_i \log(P(y_i)). \quad (2)$$

In addition to the traditional classification label for faithfulness, we propose a new pre-trained proxy label  $y'_i$  taking the popularity information into consideration. Specifically, since the topic distribution set  $\mathbf{e}$  represents the similarity between each sentence and the retrieved popular headline, the summation of topic values can be viewed as a popularity score  $\sum e_j \in \mathbb{R}^1$ . We normalize the summation by subtracting the mean and dividing the variance, then choosing the sentence with maximum value to be one of the extraction label  $y'_i$ , i.e.,  $y'_i = \operatorname{argmax}(\operatorname{normalize}(\sum e_1, \dots, \sum e_N))$ . The final loss of the extractor, denoted as  $L'_{ext}$ , is derived as follows:

$$L'_{ext} = L_{ext} - \frac{1}{N} \sum_{i=1}^N y'_i \log(P(y'_i)). \quad (3)$$

**Abstractor Training** The training data for the abstractor are pairs of extracted proxy headline  $h_{gen}$  (obtained from Eq.2) and the ground-truth headline. Specifically, the objective function  $L_{abs}$  has two main purposes: 1) to minimize the cross-entropy loss between the extracted proxy headline  $h_{gen}$  and ground-truth headline and 2) to increase the popularity score of  $h_{gen}$ . The objective function is derived as follows:

$$L_f(\theta_{abs}) = -\frac{1}{M} \sum_{m=1}^M \log P_{\theta_{abs}}(w_m | w_{1:m-1}),$$

$$L_a(h_{gen}) = -\operatorname{pop}(h_{gen}),$$

$$L_{abs} = L_f + L_a,$$

where  $w_m$  is the  $m$ -th token in the ground-truth headline and  $M$  is the headline length.

**Popularity Predictor Training** We train a binary classifier as an auxiliary model. There are two reasons for training the popularity predictor instead of the regression model. The first reason is that predicting the exact comment counts may result in overfitting to the outlier data, and the second reason is that, there is no information about other factors that affect the comment counts for a precise prediction, e.g., events. Therefore, we transform the popularity score into a binary label, where 0 represents the popularity score as being smaller than the median, otherwise the score is 1. A state-of-the-art popularity predictor (Lamprinidis, Hardt, and Hovy 2018) is trained to minimize the cross entropy. The final accuracy of popularity predictor is 65.46% on our test data.

**Reinforcement Learning** For the purpose of bridging the back propagation and introducing the classifier reward, we perform the RL training to optimize the whole model. We make the sentence extractor into an RL agent. For every extraction step, the agent observes the current state  $s = (D, \theta^D, \theta^S)$ , where  $D$  is the article,  $\theta^D$  is the topic distribution of article and  $\theta^S$  is the set of topic distributions for each sentence. After that, if the agent takes the action  $j$ , i.e.,

$$j \sim \pi(s) = P(j),$$

where  $P(j)$  is from Eq.1, it means that the agent selects  $j$ -th sentence from the article under the current policy  $\pi$ . The abstractor then rewrites the selected sentence and send it to the popularity predictor. Finally, the agent receives the reward  $r$  by adding (1) the ROUGE-L score between the target sentence and the rewritten sentence and (2) the score of the rewritten sentence from popularity predictor, i.e.,

$$r = ROUGE-L_{F_1}(abs(s_j), H) + \operatorname{pop}(abs(s_j)).$$

Moreover, due to the high variance of the vanilla policy gradient (Williams 1992), we add another mechanism, the Advantage Actor-Critic (A2C) (Mnih et al. 2016) to stabilize the training process.

It is worth noting that the role of the PTA is to guide the optimization of RL. Specifically, RL randomly selects sentences to explore the action space at the early training stage, which makes the training difficult. Without a good abstractor, RL can only receive little reward from the popularity predictor, which also makes the training of the extractor difficult. With the help of the PTA, the extractor can select a better sentence at the early training stage.

## Experimental Results

We conduct the qualitative and quantitative experiments with two real datasets to evaluate PORN-HG. For the qualitative evaluation, we provide the case study of generated headlines, and conduct a user study via asking users to evaluate the attractiveness, relevance, and grammaticality of the headlines generated by different approaches. For the quantitative evaluation, we compare different approaches in terms of attractiveness and faithfulness. To evaluate the attractiveness of generated headlines, we show the average score derived from the state-of-the-art popularity predictor (Lamprinidis, Hardt, and Hovy 2018). Moreover, we analyze the features of the popularity hypotheses mentioned in Sec. Corpus, provided along with the source code. To evaluate the faithfulness, we report the ROUGE scores for different approaches. In addition, we show the training reward curve of PORN-HG with and without the popularity information, and analyze the attention of the CNN features.

### Baselines

We implement the following baseline models and conduct an ablation task. Following the setting of (Chen and Bansal 2018), only an upper bound of the headline length is set (30 tokens) for the learning-based models.

- **IR<sub>BM25</sub>** is a bag-of-words retrieval function. It indexes every headline in the training corpus, and feeds documents

Table 2: Examples from the testing data showing the ground-truth headline and two generated headlines.

Ground-truth Headline	Chen et al.	PORL-HG
Come rain or hail! Surfers hit the southern California coast as freezing showers turn the beach white	Surf city’s surf city transformed into a white canvas ... but didn’t stop the surfers from hitting the shore	Beach! California beach transformed into white canvas after dumping of hail - but didn’t stop surfers from hitting the shore
From Rihanna to Madonna, new trend features designs of weapons, drugs and body parts	Madonna’s new black crocodile handbag causes stir for drug-related slogan	The worst offenders in the luggage? Designers create naked women and abandoned babies into their ranges
Netflix announce planet earth sequel our planet for 2019	Now, the sweeping documentary series ‘planet earth’ is getting a sequel	‘Planet earth’ is getting a sequel, says sweeping documentary

Table 3: Human evaluation results.

	PORL-HG	Chen et al.
Attractiveness	236 (63.10%)	138 (36.90%)
Relevance	77 (35.65%)	74 (34.26%)
Grammaticality	3.95	3.64

to search for the best matching headline as output according to the BM25 relevant score function.

- **Random** selects a sentence randomly from the article.
- **PREFIX** takes the first sentence as the headline.
- **Seq2Seq** employs a bidirectional LSTM as the sentence encoder, and use another bidirectional LSTM to obtain article level representation. A two-layer LSTM is then applied to decode the article representation. All models are equipped with the attention mechanism.
- **Chen et al. (Chen and Bansal 2018)** pre-train the extractor to minimize the cross-entropy loss, while the target is the proxy label. Then, they apply RL to train the extractor and use an abstractor to rewrite the sentence. Accordingly, the training target is the ground-truth headline.
- **See et al. (See, Liu, and Manning 2017)** uses the pointer network and coverage mechanism to generate headlines.

## Qualitative Results

To better understand what can be learned by our model, Table 2 shows the ground-truth and generated headlines from test data as a case study. PORL-HG can generate the headline with different forms, including interrogative sentence, exclamatory sentence or quoting the emphasis statement, depending on the suitability. Moreover, the headlines generated by PORL-HG sometimes express stronger sentiments (the first and second examples), which may make users feel stronger emotions for the headlines and lead to click. In contrast, without the information of attractiveness, (Chen and Bansal 2018) only headlines that summarize the articles are generated. Besides, PORL-HG sometimes uses the eye-catching words at the beginning of the headlines to draw attentions as shown in the first example.

To evaluate the performance of inspired headline generation, we conduct a user survey with 73 users, where 32 participated users have research experience in NLP/deep learning and the rest of the users are not familiar with NLP/deep

learning. For the human evaluation, we consider the following three modalities. 1) Attractiveness: given two headlines generated by PORL-HG and (Chen and Bansal 2018), we ask users to choose the one that he/she would click; 2) Relevance: given the human written summary that provided in the CNN/Daily Mail Dataset and the headlines generated by different approaches, users are asked to answer whether the headlines are related to the given summary and can select *more than one headline* as relevance; and 3) Grammaticality: given generated headlines, people are asked to rate the generated headlines from 1 to 5 (the higher score indicates a better result).

Table 3 shows the attractiveness, relevance and grammaticality of headlines generated from the state-of-the-art method and PORL-HG. The results manifest that 63.1% users think that the headlines generated by PORL-HG are more attractive, while only 36.9% users think that the headlines generated by (Chen and Bansal 2018) are more attractive. For the relevance, the scores of PORL-HG and (Chen and Bansal 2018) are close, indicating that PORL-HG generates headlines with higher attractiveness without affecting the faithfulness compared with the state-of-the-art method. Furthermore, the user survey shows that the grammatical quality of PORL-HG is slightly better. By our observation, the reason might be the readability of shorter words. The average token length of PORL-HG and (Chen and Bansal 2018) are 5.15 and 5.47 respectively.

## Quantitative Results

We statistically analyze the attractiveness and faithfulness of the headlines. First, the attractiveness is evaluated by the state-of-the-art popularity predictor (Lamprindis, Hardt, and Hovy 2018) and reports the percentage of headlines classified as attractive. Table 4 shows that PORL-HG achieves the best attractiveness. Second, we follow previous works (Nallapati, Zhai, and Zhou 2017; See, Liu, and Manning 2017; Zhang et al. 2018), and use ROUGE (Lin 2004) as a metric to evaluate the faithfulness. Table 5 is split into extraction and abstraction results for clear comparisons. For the extraction, the performances of  $IR_{BM25}$  and random models are poor, which suggests that memorizing the whole training corpus or randomly picking does not work. In contrast, the *PREFIX* model performs quite well, which is expected given that most news articles state the key points in the first sentence. We implement an extractor (denoted as *ext*) with a simple pointer (Vinyals, Fortu-

Table 4: Attractiveness evaluation results.

Models	Attractiveness
Ground truth	24.53
Seq2Seq	33.29
Chen et al.	34.53
See et al.	37.63
PORL-HG	<b>44.06</b>

Table 5: Faithfulness evaluation results, where CP denotes the copy rate. Note that the goal of PORL-HG is to “maintain” the faithfulness and “improve” the attractiveness instead of improving the faithfulness and attractiveness simultaneously.


Models	R-1	R-2	R-L	CP
Extraction results				
IR <sub>BM25</sub>	10.29	1.80	8.11	-
Random	13.74	2.92	11.04	-
PREFIX	32.36	13.66	26.67	-
ext	32.29	13.65	26.48	-
PORL-HG w/o abs	<b>32.51</b>	<b>13.79</b>	<b>26.82</b>	-
Abstraction results				
Seq2Seq	13.57	3.07	11.82	66.59
See et al.	27.80	12.06	23.30	97.00
Chen et al.	<b>34.84</b>	<b>15.91</b>	<b>30.26</b>	96.97
PORL-HG	34.23	15.35	29.48	95.30

nato, and Jaitly 2015), which performs slightly worse than *PREFIX*. The extraction result of PORL-HG shows that the way we incorporate the popularity information does not affect the ROUGE score. It is worth noting that the abstraction result of PORL-HG is slightly smaller than (Chen and Bansal 2018) in terms of ROUGE scores because the proposed PORL-HG adopts few more different words to make headlines attractive, which can be observed from the copy rates, i.e., 95.30% and 96.97% for PORL-HG and (Chen and Bansal 2018), respectively. Third, Table 6 shows the ablation task of our model evaluated by Meteor and Attractiveness. The Meteor metric is used for faithfulness comparison, which is calculated by the recall and the precision, and takes synonyms into consideration. The PORL-HG significantly outperforms the baseline by 27.60% for attractiveness while slightly improves the baseline in terms of Meteor. Moreover, the PTA contributes the most for attractiveness (from 40.76 to 44.06), suggesting that simultaneously considering topic distributions of sentences and popular headlines is effective.

To further validate the effect of the proposed PTA, Figure 3a illustrates the training reward curve of PORL-HG with and without PTA (red and blue curves, respectively). PORL-HG with PTA is more stable and achieves saturation more quickly. This is because RL without PTA randomly selects sentences to explore the action space at the early training stage, which makes the training of the attractive abstractor difficult. Without a good abstractor, RL can only receive a little reward from the popularity predictor, which also makes the training of the extractor difficult. With the help of the PTA, the extractor can select a better sentence to

Table 6: Ablation study evaluation results.

Models	Attractiveness	Meteor
Chen et al.	34.53	17.28
w/o pop, topic loss	40.76	17.25
w/o pop	40.96	17.53
PORL-HG	<b>44.06</b>	<b>18.07</b>



(a) The reward curve during (b) The attention to the feature training of PORL-HG. maps of the CNN.

Figure 3: Analysis of the proposed PTA.

gain the higher reward at the early training stage.

To analyze how the popularity predictor works, we visualize the convolution results for the generated headlines. Figure 3b shows the intensity of convolution feature map for different kernel sizes. Each column represents the feature map of a headline. From the heat map, we first observe that the features derived by  $conv1 \times 2$  tend to be more attended, meaning that the popularity predictor pays more attention on 2-word level. Moreover, since the headline 21 contains high attention values for  $conv1 \times 3$  (highlighted by the blue circle), we further investigate the headline, which is “*Glow-in-the-dark tampons shed light on water pollution*”. The popularity predictor attends on “Glow-in-the-dark tampons shed”, which shows the popularity predictor actually focuses on the important term (e.g., proper nouns). Moreover, we can also observe the spotting behavior of attention models for classifying the headline popularity, since the large values for  $conv1 \times 1$  and  $conv1 \times 2$  are only in one or two positions for a headline.

## Conclusion and Future Work

In this paper, we tackle the challenging task of generating an eye-catching headline with general form. To strike a balance between faithfulness and attractiveness, we propose the POPularity-Reinforced Learning for inspired Headline Generation incorporating the topic distributions of sentences and popular headlines. We also verify the effectiveness of each module in PORL-HG carefully by the quantitative and qualitative experiments, which demonstrate that our model outperforms the state-of-the-art baselines on the attractiveness while simultaneously maintaining the faithfulness. In the future, we plan to explore the possibility of designing a better RL algorithm purely on the abstractive system.

## Acknowledgement

This work was supported in part by the Ministry of Science and Technology of Taiwan under Grants MOST-

108-2221-E-009-088, MOST-108-2622-E-009-026-CC2, MOST-108-2634-F-009-006, MOST-108-2221-E-001-012-MY3, MOST-108-2218-E-009-050, and by the Higher Education Sprout Project by the Ministry of Education (MOE) in Taiwan through grant 108W267.

## References

- Blei, D. M.; Ng, A. Y.; and Jordan, M. I. 2003. Latent dirichlet allocation. *Journal of machine Learning research* 3(Jan):993–1022.
- Chen, Y.-C., and Bansal, M. 2018. Fast abstractive summarization with reinforce-selected sentence rewriting. In *Association for Computational Linguistics (ACL)*, 675–686.
- Cheng, J., and Lapata, M. 2016. Neural summarization by extracting sentences and words. In *Association for Computational Linguistics (ACL)*, 318–325. Association for Computational Linguistics.
- Cohn, T., and Lapata, M. 2008. Sentence compression beyond word deletion. In *International Conference on Computational Linguistics (COLING)*.
- Filippova, K.; Alfonseca, E.; Colmenares, C. A.; Kaiser, L.; and Vinyals, O. 2015. Sentence compression by deletion with lstms. In *Empirical Methods in Natural Language Processing (EMNLP)*, 360–368.
- Hayashi, Y., and Yanagimoto, H. 2018. Headline generation with recurrent neural network. In *New Trends in E-service and Smart Computing*. Springer. 81–96.
- Hermann, K. M.; Kocisky, T.; Grefenstette, E.; Espeholt, L.; Kay, W.; Suleyman, M.; and Blunsom, P. 2015. Teaching machines to read and comprehend. In *Advances in Neural Information Processing Systems (NIPS)*, 1693–1701.
- Higurashi, T.; Kobayashi, H.; Masuyama, T.; and Murao, K. 2018. Extractive headline generation based on learning to rank for community question answering. In *International Conference on Computational Linguistics (COLING)*, 1742–1753.
- Kim, Y. 2014. Convolutional neural networks for sentence classification. In *Empirical Methods in Natural Language Processing (EMNLP)*, 1746–1751. Doha, Qatar: Association for Computational Linguistics.
- Knight, K., and Marcu, D. 2000. Statistics-based summarization-step one: Sentence compression. *AAAI/IAAI* 2000:703–710.
- Kuiken, J.; Schuth, A.; Spitters, M.; and Marx, M. 2017. Effective headlines of newspaper articles in a digital environment. *Digital Journalism* 5(10):1300–1314.
- Lamprinidis, S.; Hardt, D.; and Hovy, D. 2018. Predicting news headline popularity with syntactic and semantic knowledge using multi-task learning. In *Empirical Methods in Natural Language Processing (EMNLP)*, 659–664.
- Lin, C.-Y. 2004. Rouge: A package for automatic evaluation of summaries. *Text Summarization Branches Out*.
- Mnih, V.; Badia, A. P.; Mirza, M.; Graves, A.; Lillicrap, T.; Harley, T.; Silver, D.; and Kavukcuoglu, K. 2016. Asynchronous methods for deep reinforcement learning. In *International conference on machine learning*, 1928–1937.
- Nallapati, R.; Zhou, B.; Gulcehre, C.; Xiang, B.; et al. 2016. Abstractive text summarization using sequence-to-sequence rnns and beyond. *SIGLL Conference on Computational Natural Language Learning (CoNLL)*.
- Nallapati, R.; Zhai, F.; and Zhou, B. 2017. Summarunner: A recurrent neural network based sequence model for extractive summarization of documents. In *AAAI Conference on Artificial Intelligence (AAAI)*.
- Narayan, S.; Cohen, S. B.; and Lapata, M. 2018a. Don’t give me the details, just the summary! topic-aware convolutional neural networks for extreme summarization. In *Empirical Methods in Natural Language Processing (EMNLP)*.
- Narayan, S.; Cohen, S. B.; and Lapata, M. 2018b. Ranking sentences for extractive summarization with reinforcement learning. In *North American Chapter of the Association for Computational Linguistics*.
- Over, P.; Dang, H.; and Harman, D. 2007. Duc in context. In *Information Processing & Management*.
- Paulus, R.; Xiong, C.; and Socher, R. 2018. A deep reinforced model for abstractive summarization. In *International Conference on Learning Representations (ICLR)*.
- Rush, A. M.; Chopra, S.; and Weston, J. 2015. A neural attention model for abstractive sentence summarization. In *Empirical Methods in Natural Language Processing (EMNLP)*.
- See, A.; Liu, P. J.; and Manning, C. D. 2017. Get to the point: Summarization with pointer-generator networks. In *Association for Computational Linguistics (ACL)*, 1073–1083.
- Szymanski, T.; Orellana-Rodriguez, C.; and Keane, M. T. 2017. Helping news editors write better headlines: A recommender to improve the keyword contents & shareability of news headlines. *arXiv preprint arXiv:1705.09656*.
- Takase, S.; Suzuki, J.; Okazaki, N.; Hirao, T.; and Nagata, M. 2016. Neural headline generation on abstract meaning representation. In *Empirical methods in Natural Language Processing (EMNLP)*, 1054–1059.
- Vinyals, O.; Fortunato, M.; and Jaitly, N. 2015. Pointer networks. In *Advances in Neural Information Processing Systems*, 2692–2700.
- Weng, S.-S., and Wu, J.-Y. 2018. Recommendation on keyword combination of news headlines. In *International Conference on Systems and Informatics (ICSAI)*, 1146–1151. IEEE.
- Williams, R. J. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning* 8(3-4):229–256.
- Zhang, R.; Guo, J.; Fan, Y.; Lan, Y.; Xu, J.; Cao, H.; and Cheng, X. 2018. Question headline generation for news articles. In *ACM International Conference on Information and Knowledge Management (CIKM)*, 617–626. ACM.