

Frontiers of Information Technology & Electronic Engineering
 www.jzus.zju.edu.cn; engineering.cae.cn; www.springerlink.com
 ISSN 2095-9184 (print); ISSN 2095-9230 (online)
 E-mail: jzus@zju.edu.cn



Enhancing low-resource cross-lingual summarization from noisy data with fine-grained reinforcement learning*

Yuxin HUANG^{†1,2}, Huailing GU^{1,2}, Zhengtao YU^{†1,2}, Yumeng GAO^{1,2}, Tong PAN^{1,2}, Jialong XU^{1,2}

¹Faculty of Information Engineering and Automation,
 Kunming University of Science and Technology, Kunming 650504, China

²Yunnan Key Laboratory of Artificial Intelligence,
 Kunming University of Science and Technology, Kunming 650504, China

[†]E-mail: huangyuxin2004@163.com; ztyu@hotmail.com

Received Apr. 27, 2023; Revision accepted Oct. 22, 2023; Crosschecked Nov. 3, 2023; Published online Dec. 27, 2023

Abstract: Cross-lingual summarization (CLS) is the task of generating a summary in a target language from a document in a source language. Recently, end-to-end CLS models have achieved impressive results using large-scale, high-quality datasets typically constructed by translating monolingual summary corpora into CLS corpora. However, due to the limited performance of low-resource language translation models, translation noise can seriously degrade the performance of these models. In this paper, we propose a fine-grained reinforcement learning approach to address low-resource CLS based on noisy data. We introduce the source language summary as a gold signal to alleviate the impact of the translated noisy target summary. Specifically, we design a reinforcement reward by calculating the word correlation and word missing degree between the source language summary and the generated target language summary, and combine it with cross-entropy loss to optimize the CLS model. To validate the performance of our proposed model, we construct Chinese-Vietnamese and Vietnamese-Chinese CLS datasets. Experimental results show that our proposed model outperforms the baselines in terms of both the ROUGE score and BERTScore.

Key words: Cross-lingual summarization; Low-resource language; Noisy data; Fine-grained reinforcement learning; Word correlation; Word missing degree

<https://doi.org/10.1631/FITEE.2300296>

CLC number: TP391

1 Introduction

Cross-lingual summarization (CLS) is the task of automatically generating a short target language summary based on the source language's long text, which can be regarded as a cross-lingual

text generation task. In recent years, data-driven sequence-to-sequence models have achieved considerable performance in cross-lingual generation tasks such as machine translation (Rippeth and Post, 2022), cross-lingual dialogue generation (Kim et al., 2021; Zhou et al., 2023), and video summarization (Li P et al., 2021; Javed and Ali Khan, 2022), and their performance is derived mainly from large-scale, high-quality training data. However, due to the paucity of data, the performance of cross-lingual text generation tasks for low-resource languages is unsatisfactory.

The main challenge encountered in CLS tasks is the construction of extensive and high-quality CLS datasets. Currently, a large number of monolingual

[‡] Corresponding author

* Project supported by the National Natural Science Foundation of China (Nos. U21B2027, 62266027, 61972186, 62241604), the Yunnan Provincial Major Science and Technology Special Plan Projects, China (Nos. 202302AD080003, 202103AA080015, and 202202AD080003), the General Projects of Basic Research in Yunnan Province, China (Nos. 202301AT070471 and 202301AT070393), and the Kunming University of Science and Technology "Double First-Class" Joint Project, China (No. 202201BE070001-021)

ORCID: Yuxin HUANG, <https://orcid.org/0000-0003-1277-6212>; Zhengtao YU, <https://orcid.org/0000-0002-4012-461X>

© Zhejiang University Press 2023

summarization datasets have been constructed for rich-resource languages such as Chinese and English. For instance, Hu et al. (2015) compiled a dataset named LCSTS, which includes Weibo posts as source texts and Weibo headlines as summaries. On the other hand, the Cable News Network/Daily Mail corpus primarily consists of news articles from the American Cable News Network and the Daily Mail (Hermann et al., 2015). However, acquiring and constructing CLS datasets through direct means remains a highly challenging task. The mainstream method is to use machine translation to translate the source language text or summary of a monolingual summary into the target language. For example, Zhu et al. (2019) used the round-trip translation (RTT) back-translation strategy to construct a CLS dataset. To ensure high-quality summaries, they applied several filtering criteria based on ROUGE (Lin, 2004) score, including length consistency, sentence fluency, and meaning preservation. This resulted in a high-quality parallel corpus suitable for CLS research. However, the construction process of CLS datasets based on translation heavily relies on the performance of machine translation models. However, for low-resource languages such as Chinese-Vietnamese, machine translation performance is unsatisfactory, resulting in a significant amount of noise during the data construction process. In particular, the model's ability to accurately generate summaries is greatly affected when the target language summary, which is used as the reference during training, contains translation errors.

According to statistics, there are about 50% of data with problems such as missing content words and improper word selection, as shown in Fig. 1. In Fig. 1a, the word “phản hồi (respond)” has been inaccurately translated as “处理 (handle),” while in Fig. 1b the content word “ví (wallet)” has been left untranslated. Employing such imprecise and incomplete pseudo-summaries as supervisory signals could

potentially misguide the model. Drawing from this insight, we propose to introduce the source language summary and align it with the generated target language summary to assess the adequacy of the generated summaries in terms of word omission and accuracy.

Based on the aforementioned analysis and leveraging the alignment information between the source language summary and the generated target language summary, we propose a fine-grained reinforcement learning based CLS approach to mitigate the errors caused by improper word selection and missing content words, which are prevalent in pseudo-target summaries and can lead to misguided model training. To address the issue of improper word selection during the decoding process, we design a reinforcement learning reward by the word correlations between the source language summary and the generated target summary. To tackle missing content words in the generated summary, we penalize the decoder based on the importance of the missing words relative to the source language summary.

We propose a fine-grained reinforcement learning reward that incorporates word correlation and missing degree between the source language summary and the generated target language summary. We combine this reward with the traditional cross-entropy loss to optimize the model, thus providing more effective guidance for generating the target language summary.

We conduct experiments on the Transformer framework using the Chinese-Vietnamese and Vietnamese-Chinese CLS datasets. The results show that our method achieves significant improvements compared with previous methods. Additionally, reinforcement learning rewards based on a combination of word correlation and missing degree can help generate a better summary. Our main contributions are as follows:

1. We model the relationship between the source




Fig. 1 Common translation errors: (a) improper choice of words; (b) missing content word

language summary and target language summary from a fine-grained perspective, alleviating error guidance caused by noisy data in CLS tasks.

2. The experimental results show that this method achieves significant improvements compared with previous methods on the Chinese-Vietnamese and Vietnamese-Chinese CLS datasets.

2 Related works

2.1 Cross-lingual summarization

CLS is the task of generating a summary in a target language from a document in a source language. Traditional CLS tasks usually adopt a technical framework of translating first and then summarizing (Leuski et al., 2003; Ouyang et al., 2019) or summarizing first and then translating (Lim et al., 2004; Orăsan and Chiorean, 2008). However, they are often affected by the error propagation of translation and summarization models, and the results are not satisfactory in low-resource languages. Neural network based CLS tasks (Jiang et al., 2022; Wang et al., 2022) are usually seen as similar to machine translation tasks, but the difference is that machine translation maintains the same amount of information in its input and output, whereas CLS requires compression and translation of information. There are typically two types of methods for low-resource language CLS tasks. The first type is summary methods based on zero-shot learning. Ayana et al. (2018) addressed the problem of the lack of source language to target language summary datasets by using a pre-trained machine translation model and a headline generation model as teacher networks to guide the learning of the cross-lingual headline generation model. This approach enables the model to have translation and summarization abilities and to generate cross-lingual summaries under zero-shot conditions. Nguyen and Luu (2022) employed a monolingual summarization model as the guiding teacher network to facilitate parameter learning in the CLS model. The second type of method is based on a multi-task joint learning approach that combines machine translation and summarization models to address the problem of sparse training data (Takase and Okazaki, 2020; Liang et al., 2022). Zhu et al. (2019) proposed an end-to-end CLS model based on the Transformer text generation framework. They jointly trained the CLS

and monolingual summarization and the CLS and machine translation tasks with parameter sharing at the encoding stage. During the training process, the two tasks were alternately trained to have the ability to generate cross-lingual summaries. Cao et al. (2020) used generative adversarial networks to align the contextual representations of two monolingual summarization models in the source and target languages, achieving bilingual alignment while performing monolingual summarization. Bai et al. (2021) argued that although joint learning of CLS and machine translation can enhance CLS performance by sharing encoder parameters, the decoders of the two tasks are independent of each other and cannot establish good alignment between CLS and machine translation tasks. Based on the analysis, most of the aforementioned works are based on machine translation to construct pseudo parallel CLS datasets from monolingual datasets, targeting languages with rich resources such as Chinese and English, where machine translation performs well and has fewer errors. However, for low-resource languages, the translation performance is unsatisfactory, and constructing CLS datasets through translation can introduce a large amount of data noise. Effective analysis and processing methods for CLS in low-resource and noisy data scenarios are still lacking.

2.2 Reinforcement learning

Reinforcement learning has been widely used in many tasks (Zhao J et al., 2022; Li HQ et al., 2023; Xiong et al., 2023), such as machine translation and text summarization, mainly through global decoding optimization to alleviate the exposure bias problem (Kumar et al., 2019; You et al., 2019). In the summarization task, Paulus et al. (2017) used the ROUGE value between the real summary and the generated summary as a reinforcement learning reward to reward or punish the model and to combine this reward with cross-entropy using linear interpolation as the training objective function, which partially alleviates the exposure bias problem. According to Böhm et al. (2019), the correlation between ROUGE and human evaluation is weak when it comes to summaries that have a diverse vocabulary, which suggests that using ROUGE as a reinforcement learning reward may not be reliable. They used the source text and the generated summary as the input to learn a reward function from human-scored summaries, achieving

better results than those using ROUGE as a reward. Yoon et al. (2021) calculated the semantic similarity between the generated summary and the reference summary based on a language model as a reinforcement learning reward, improving the reward acquisition method of word-level matching ROUGE. For CLS tasks, Dou et al. (2020) used the similarity between the source language summary and the generated target language summary as a reinforcement learning reward to constrain the model to generate better summaries. Inspired by this study, we believe that by better modeling the correlation between the source language summary and the generated summary, we can effectively use the noise-free source language summary to alleviate the noise problem caused by the translation.

3 Fine-grained reinforcement learning for low-resource cross-language summarization

To address the issue of noise in the supervision signal in low-resource cross-language summarization, we propose a fine-grained reinforcement learning method for cross-language summarization based on

the Transformer model (Vaswani et al., 2017). To improve the quality of generated summaries and mitigate the impact of noise in pseudo-target language summaries, we design reinforcement learning rewards based on the word correlation and missing degree between the source language summary and the generated target language summary. The reinforcement learning function is then combined with the maximum likelihood estimation function as the training objective to optimize the generated summary. The model structure is shown in Fig. 2.

3.1 Model

In traditional CLS models based on the Transformer architecture, given a training set $\{X^A, \hat{Y}^B\}$, where A represents the source language and B represents the target language, each document X^A is mapped to a high-dimensional vector to obtain an input document sequence $X^A = \{x_1, x_2, \dots, x_N\}$, which is then encoded by the encoder to obtain a vector representation $H = \{h_1, h_2, \dots, h_N\}$ of the document sequence (N is the length of sequence X^A). Finally, the decoder generates a summary sequence $Y^B = \{y_1, y_2, \dots, y_M\}$ based on the given H . During this process, the maximum likelihood estimate

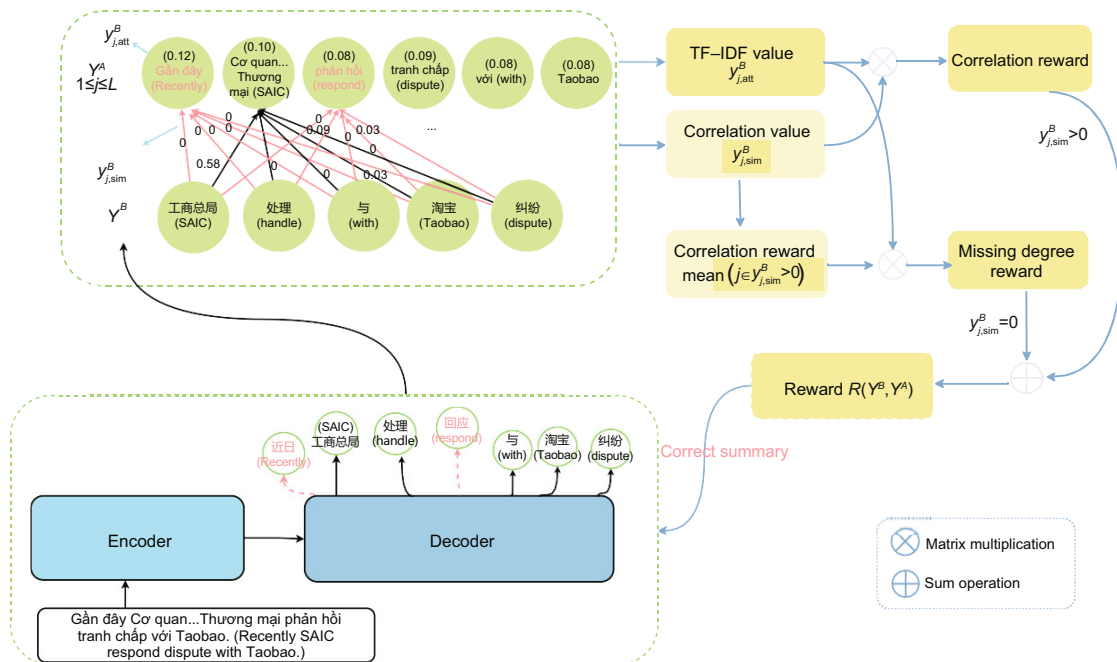


Fig. 2 Fine-grained reinforcement learning model structure to improve the quality of model-generated summary by computing word correlation and missingness between decoder-generated summary and source language summary (Words in red represent the updated words. References to color refer to the online version of this figure)

between the generated summary Y^B and the reference summary \tilde{Y}^B is used as the optimization objective, and the cross-entropy loss function is defined as follows:

$$L_{\text{cls}} = \sum_{t=1}^M \log P(y_t^B | y_{<t}^B, X^A; \theta), \quad (1)$$

where M is the length of summarization Y^B .

3.2 Reinforcement learning loss based on word correlation and word missing degree

Given the training set $\{X^A, \tilde{Y}^B\}$ obtained by translating a monolingual summary dataset, we investigate the noise types that account for a significant proportion of errors in \tilde{Y}^B data, namely, improper word selection and missing content words (detailed analysis is given in Section 4.1.2). Therefore, we introduce the source language summary Y^A as a reference and design a reinforcement learning reward by calculating the word correlation and word missing degree between the source language summary Y^A and the generated target language summary Y^B to weaken the error guidance caused by the pseudo-target language summary as the supervision signal.

In terms of word correlation, the quality of the generated target language summary $Y^B = \{y_1^B, y_2^B, \dots, y_M^B\}$ is directly evaluated by using the source language summary sequence $Y^A = \{y_1^A, y_2^A, \dots, y_L^A\}$. The fast-align (Dyer et al., 2013) tool is used to align the Chinese-Vietnamese bilingual parallel corpus in advance, obtaining the correlation table $Y_{\text{align}}^{A \rightarrow B}$ between each source language word and target language word. In the correlation table, the correlation values between each source language y_j^A and each target language y_t^B are recorded. Using the source language summary word y_j^A , it is matched with target language summary $Y^B = \{y_1^B, y_2^B, \dots, y_M^B\}$ in the correlation table $Y_{\text{align}}^{A \rightarrow B}$, denoted as $y_{j,\text{sim}}^B$, which represents the correlation value between each source language summary word y_j^A and the generated target language summary Y^B :

$$y_{j,\text{sim}}^B = \text{sim}(y_j^A, Y_{\text{align}}^{A \rightarrow B})_{\text{sum}}, \quad (2)$$

where $\text{sim}(\cdot)$ denotes obtaining the correlation values from the correlation table $Y_{\text{align}}^{A \rightarrow B}$ between y_j^A and Y^B , and sum represents the summation of the correlation values.

In terms of missing degree, the importance of each word in the source language summary $Y^A =$

$\{y_1^A, y_2^A, \dots, y_M^A\}$ is calculated, denoted as $y_{j,\text{idf}}^A$. As shown in Eq. (3), when the source language summary word y_j^A can match the words in the corresponding target language summary Y^B , the importance of word y_j^A is recorded as the importance of the matched word $y_{j,\text{att}}^B$ in the generated summary Y^B . When the source language summary word y_j^A cannot match the words in the corresponding target language summary Y^B , it is considered that the word is missing in the generated summary. In this case, the importance of word y_j^A is recorded as the word missing degree $y_{j,\text{att}}^B$ in the generated summary to avoid the impact of missing important words in the pseudo-target language on the model learning.

$$y_{j,\text{att}}^B = \begin{cases} y_{j,\text{idf}}^A, & y_{j,\text{sim}}^B > 0, \\ -y_{j,\text{idf}}^A, & y_{j,\text{sim}}^B = 0, \end{cases} \quad (3)$$

where $y_{j,\text{idf}}^A$ is calculated using term frequency-inverse document frequency (TF-IDF).

In the CLS model, we consider the model as an agent, with the context representation vector obtained at each decoding step t and the summary $y_{<t}^B$ generated in the previous time step $t - 1$ being perceived as the environment. During the process of summary generation, the agent needs to choose a word from the candidate word list as the summary word for the current time step t ; this selection process constitutes an action of the agent. Upon completing a summary generation, the model receives a reward $R(Y^B, Y^A)$. The reward function calculation process is shown in Algorithm 1. By assigning higher scores, the model is encouraged to generate better summaries. We use Eq. (4) to calculate the expected reward:

$$R(Y^B, Y^A) = \frac{1}{L} \sum_{j=1}^L \begin{cases} y_{j,\text{sim}}^B y_{j,\text{att}}^B, & y_{j,\text{sim}}^B > 0, \\ \frac{\sum_{j=1}^L y_{j,\text{sim}}^B}{\text{Count}(y_{j,\text{sim}}^B > 0)} y_{j,\text{att}}^B, & y_{j,\text{sim}}^B = 0, \end{cases} \quad (4)$$

where L represents the length of the source language summary and $\text{Count}(y_{j,\text{sim}}^B > 0)$ represents the number of times $y_{j,\text{sim}}^B$ being larger than 0.

The training objective of reinforcement learning is to maximize the expected reward, which is defined as the reinforcement learning loss:

$$L_{\text{rl}} = - \sum_{Y^B \in \mathcal{Y}} \log P(Y^B | X^A; \theta) R(Y^B, Y^A), \quad (5)$$

where \mathcal{Y} represents all the possibly generated candidate summaries. An exponential search space is constituted in the summary generation process.

Algorithm 1 Reward function design

```

1: Input:  $Y_{idf}, Y_{align}^{A \rightarrow B}$  /* Input TF-IDF values of source
   language summary words and correlation table */
2: score  $\leftarrow 0$  /* Total reward score */
3: sumwd  $\leftarrow 0$  /* Total word missingness penalty score */
4: Countcor  $\leftarrow 0$  /* Count the number of times  $y_{j,sim}^B$  being
   greater than 0 */
5: sumcor  $\leftarrow 0$  /* Total word correlation score */
6: for  $y_{j,idf}^A$  in  $Y_{idf}^A$  do
7:    $y_{j,sim}^B \leftarrow \text{sim}(y_{j,idf}^A, Y_{align}^{A \rightarrow B})_{sum}$ 
8:   if  $y_{j,sim}^B = 0$  then
9:     sumwd  $\leftarrow \text{sum}_{wd} - y_{j,idf}^A$ 
10:  else
11:    scorecor  $\leftarrow y_{j,sim}^B \cdot y_{j,idf}^A$ 
12:    score  $\leftarrow \text{score} + \text{score}_{cor}$ 
13:    sumcor  $\leftarrow \text{sum}_{cor} + y_{j,sim}^B$ 
14:    Countcor  $\leftarrow \text{Count}_{cor} + 1$ 
15:  end if
16: end for
17: if Countcor  $\neq 0$  then
18:   avgcor  $\leftarrow \frac{\text{sum}_{cor}}{\text{Count}_{cor}}$ 
19:   score  $\leftarrow \text{score} + \text{sum}_{wd} \cdot \text{avg}_{cor}$ 
20: end if
21: return score

```

In practice, a sequence Y^S is often sampled from the probability distribution function $P(Y^B|X^A; \theta)$ to optimize the expected reward; however, this can result in high variance. To address this issue, we adopt the same method as in previous research (Rennie et al., 2017; Kang et al., 2020) and introduce a baseline reward to reduce gradient variance. We use the self-critical policy gradient training algorithm in the training of the reinforcement learning objective, which involves two summary-generating strategies: one where Y^S is randomly sampled from the conditional probability function $P(Y^B|X^A; \theta)$ and the other where Y^G is generated by the greedy decoding. The training objective of a summary sentence in reinforcement learning is as follows:

$$L_{rl} = - (R(Y^S, Y^A) - R(Y^G, Y^A)) \cdot \sum_{t=1}^M \log P(Y^B|X^A; \theta). \quad (6)$$

Finally, the hybrid objective function for training the CLS model is a linear interpolation of the cross-entropy loss function and the reinforcement learning training objective function, expressed as

$$L_{mix} = \gamma L_{cls} + (1 - \gamma) L_{rl}, \quad (7)$$

where γ is the scale factor between the cross-entropy loss function and the reinforcement learning training objective function.

4 Experiments

4.1 Data analysis

4.1.1 Data construction

We constructed two types of CLS datasets: Chinese-Vietnamese and Vietnamese-Chinese. For Chinese-Vietnamese summarization, we used the first 200 000 data samples of LCSTS (Hu et al., 2015) for back-translation to obtain the Chinese-Vietnamese CLS dataset (Zh-Visum). For Vietnamese-Chinese summarization, we crawled Vietnamese monolingual datasets from various news websites, including Vietnam+, Vietnam News Agency, Vietnam Express, and Vietnam News Agency. The collected data were then cleaned and back-translated to obtain 115 798 samples of the Vietnamese-Chinese CLS dataset (Vi-Zhsum), where the translation was performed by YunLing translation (<http://yuntrans.vip>). We used ROUGE (Lin, 2004), BERTScore (Zhang et al., 2020), and MGF-Score (Lai et al., 2022) to filter the back-translated data. Taking Zh-Visum as an example, the specific workflow is shown in Fig. 3.

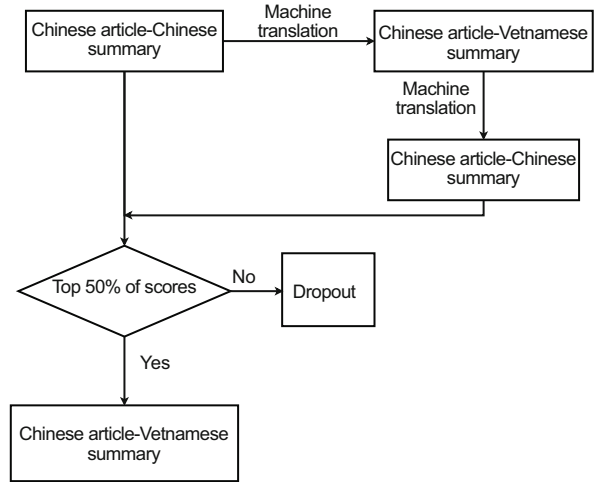


Fig. 3 Zh-Visum filtering flowchart

In Zh-Visum, the lowest 50% of the filtered data were removed, leaving 100 000 data samples; due to the relatively high quality of Vi-Zhsum back-translation, the lowest 30% of the data were filtered, leaving 81 000 data samples. Detailed information on back-translation scores before and after data filtering is shown in Table 1.

4.1.2 Noise analysis

From Table 1, it can be seen that the filtering of the back-translated data effectively improves the quality of Zh-Visum and Vi-Zhsum. However, further analysis of Zh-Visum and Vi-Zhsum revealed that only a small proportion of sentences were completely correct, and that high-quality and large-scale data were needed for training the model. We further analyzed the noise in the cross-lingual data constructed by machine translation according to the types of noise defined in the literature (Zhao H et al., 2013).

We randomly selected 100 pairs of source and target language summaries from the Chinese-Vietnamese and Vietnamese-Chinese CLS datasets and manually marked the noise types in the unfiltered and filtered data using three different filtering methods: ROUGE, BERTScore, and MGFScore. Table 2 shows the noise type statistics. Additionally, 50% of the Zh-Visum data and 30% of the Vi-Zhsum data have been filtered.

We can have the following conclusions from Table 2:

1. The proportion of error-free sentences in the constructed cross-lingual summary dataset is relatively small. Even though filtering by evaluation metrics can improve the data accuracy, it cannot avoid noisy data. Therefore, after improving the quality of the dataset, further research is needed on methods for cross-lingual summary generation under noise.

2. In both Vi-Zhsum and Zh-Visum, the types

of noise with the first and second highest proportions are inappropriate word selection and missing content words. In Zh-Visum, the data are obtained through back-translation from LCSTS, a short text summary dataset obtained from Weibo using headlines as summaries. The Chinese words used in these summaries are often concise and to the point, making it easy for machine translation to deviate from the correct understanding and ignore some content words. In Vi-Zhsum, errors in word order are also common as Vi-Zhsum is translated from longer texts, and machine translation tends to have weaker comprehension of the logical sequence between words in longer texts, making it prone to mistakes in word order. This type of noise has a weaker impact on the quality of the generated sentences.

In summary, in both Vi-Zhsum and Zh-Visum, the types of noise with the first and second highest proportions are improper word selection and missing content words. Therefore, it is necessary to use data filtering to improve the quality of pseudo-data and continue to weaken the noise.

To verify the effectiveness of the proposed model, we constructed a Chinese-Vietnamese CLS dataset Zh-Visum_{Filter} with 100 000 samples and a Vietnamese-Chinese CLS dataset Vi-Zhsum_{Filter} with 81 000 samples, using the method described in Section 3.1. The detailed data are shown in Table 3, where BERT represents using BERTScore to filter Zh-Visum data, MGF represents using MGFScore

Table 1 Back-translation score for Zh-Visum and Vi-Zhsum

Dataset	Zh-Visum				Vi-Zhsum			
	RG-1	RG-2	RG-L	BERTScore	RG-1	RG-2	RG-L	MGFScore
Filter-No	52.71	28.19	48.62	74.90	67.52	43.21	64.00	87.86
Filtered	65.70	41.10	61.34	81.18	73.55	50.89	70.86	90.36

ROUGE-1, ROUGE-2, and ROUGE-L are used to evaluate the quality of the generated summary, denoted as RG-1, RG-2, and RG-L, respectively. Filter-No: unfiltered Zh-Visum and Vi-Zhsum datasets; Filtered: filtered Zh-Visum and Vi-Zhsum datasets

Table 2 The proportion of data of different noise types in Zh-Visum and Vi-Zhsum

Noise type	Proportion for Zh-Visum (%)			Proportion for Vi-Zhsum (%)		
	Filter-No	Filter-RG	Filter-BERT	Filter-No	Filter-BERT	Filter-MGF
Improper choice of words	37.00	30.67	29.34	28.25	23.58	22.38
Missing content words	17.50	18.67	18.32	12.75	15.71	15.95
Wrong word order	6.33	5.66	4.67	17.42	15.24	16.19
Named entity error	12.67	13.00	14.67	7.75	8.57	8.57
Other	6.50	8.00	7.00	4.83	5.47	4.05
No error	20.00	24.00	26.00	29.00	31.43	32.86

Filter-No: unfiltered dataset. Filter-RG/BERT/MGF: the use of ROUGE/BERTScore/MGFScore for filtering

to filter Vi-Zhsum data, and RG represents using the ROUGE metric to filter the data. Regardless of the filtering and training method used, the test sets are the same for the same language.

4.2 Evaluation metrics

In this study, the quality of summaries generated by the CLS model was evaluated using two evaluation metrics. As with most summarization works, one is based on a statistical method called ROUGE (Lin, 2004), which calculates the co-occurrence degree of N -grams between the standard summary and the generated summary, and the formula is as follows:

$$\text{ROUGE-}N = \frac{\sum_{G \in \{\text{Ref}\}} \sum_{\text{gram}_N \in \{G\}} \text{Count}_{\text{match}}(\text{gram}_N)}{\sum_{G \in \{\text{Ref}\}} \sum_{\text{gram}_N \in \{G\}} \text{Count}(\text{gram}_N)}, \quad (8)$$

where G is the generated summary, Ref is the reference summary, gram_N is the N -gram phrase, $\text{Count}_{\text{match}}(\text{gram}_N)$ is the number of overlapping N -grams in the generated summary and the reference summary, and $\text{Count}(\text{gram}_N)$ is the number of N -grams in the reference summary. N is usually set to 1, 2, and L (L is the length of the longest common subsequence). In this study, ROUGE-1, ROUGE-2, and ROUGE- L are used to evaluate the quality of the generated summary, denoted as RG-1, RG-2, and RG- L , respectively.

The other evaluation method for measuring the quality of generated summaries in a CLS model is based on deep semantic matching, as proposed by Zhang et al. (2020). This method is called BERTScore and uses a pre-trained language model to calculate the semantic similarity between the generated and reference summaries. Nowadays, this method is widely used to evaluate the quality of gen-

erated summaries. In Chinese, the pre-trained model used for scoring is “bert-base-Chinese,” while in Vietnamese the pre-trained model used is “bert-base-multilingual-cased.” In the case of using BERTScore for evaluation, the “<unk>” tokens in the generated summaries are replaced with the “[unk]” tokens from the BERT vocabulary.

4.3 Experiment setup

The model code was implemented using the PyTorch framework. The Transformer encoder and decoder were both set to six layers, with eight attention heads and a hidden vector dimension of 512. The feedforward neural network was configured with a size of 1024. The model used a teacher-forcing strategy, with a label smoothing set to 0.1 and a dropout set to 0.1. The model was trained with the warmup phase of 3000 steps and accumulated a gradient every two steps. During decoding, a beam search strategy with a beam size of five was used. It is worth noting that similar to Wu et al. (2019) and Unanue et al. (2021), for models with reinforcement learning strategies, we used unfiltered Zh-Visum and Vi-Zhsum data for parameter initialization and then trained the model using filtered data.

4.4 Baselines

To verify the effectiveness of the proposed fine-grained reinforcement learning approach for Chinese-Vietnamese CLS, we trained and compared the following baseline models on the Zh-Visum_{Filter} and Vi-Zhsum_{Filter} datasets:

1. Sum-Tra: It is a traditional CLS method that generates a summary in the source language first and then translates it into the target language.
2. Tra-Sum: Similar to Sum-Tra, it is a two-step CLS method that translates the source language

Table 3 Experimental data details

Data type	Zh-Visum _{Filter}			Vi-Zhsum _{Filter}		
	Training	Valid	Testing	Training	Valid	Testing
Number of documents	94 000	3000	3000	77 000	2000	2000
AvgTextsWords-BERT/MGF	90.05	92.18	89.45	445.87	511.39	426.95
AvgSumsWords-BERT/MGF	14.59	13.13	17.33	40.60	41.99	40.13
AvgTextsWords-RG	89.92	92.17	89.45	446.93	516.23	426.95
AvgSumsWords-RG	14.59	13.06	17.33	41.22	42.07	40.13

AvgTextsWords-BERT/MGF: average article length using BERTScore to filter Zh-Visum data and using MGFScore to filter Vi-Zhsum data. AvgSumsWords-BERT/MGF: average summary length using BERTScore to filter Zh-Visum data and using MGFScore to filter Vi-Zhsum data. -RG: using ROUGE to filter Zh-Visum and Vi-Zhsum data

document into the target language and then generates a summary in the target language. In Sum-Tra and Tra-Sum, YunLing translation is used as the machine translation model, and an unsupervised extractive method called LexRank (Takase and Okazaki, 2020) is used as the summarization model.

3. NCLS (Zhu et al., 2019): This is an end-to-end neural network CLS model based on the Transformer framework. It incorporates two related tasks, monolingual summarization and machine translation, to further improve model performance.

4. MCLAS (Bai et al., 2021): This CLS method is based on the multi-task framework, which sequentially performs monolingual summarization and CLS, using BERT (mBERT) to initialize the Transformer encoder.

5. KDCLS (Nguyen and Luu, 2022): It is a novel knowledge distillation based framework for CLS, seeking to explicitly construct cross-lingual correlation by distilling the knowledge of the monolingual summarization teacher into the CLS student.

6. LR-ROUGE (Yoon et al., 2021): This method is similar to the proposed method, but it uses RG- L to calculate the expected reward.

7. XSIM (Dou et al., 2020): This method employs reinforcement learning to directly enhance a bilingual semantic similarity metric between the summaries generated in a target language and the gold summaries in a source language.

8. LR-MC: The proposed CLS model combines cross-entropy and reinforcement learning as the optimization objective. The expected reward is calculated based on the word missing degree and word correlation between the source language summary and the generated target language summary.

4.5 Analysis of experimental results

We designed experiments from different perspectives to verify the effectiveness of the Chinese-Vietnamese and Vietnamese-Chinese CLS method based on fine-grained reinforcement learning under noisy data.

First, the effects of fine-grained reinforcement learning proposed were compared with those of baseline models. Then, the improvements of word correlation reward and word missing penalty based on reinforcement learning designed under noise were explored, and the respective impacts of these parts on the model were analyzed. Second, the influence of

the proportion factor between the cross-entropy loss function and the reinforcement learning training objective function on model performance was investigated. Next, the neural network model was trained using data before and after noise filtering, and the performances of the model under different data were compared. Finally, a case study was conducted for the summaries generated by different models.

4.5.1 Experimental results

The results of comparison between the proposed model and the baselines are shown in Table 4. Here, γ represents the scale factor between cross-entropy loss and the expected reward; $\gamma = 1$ means that no reinforcement learning reward is added.

From Table 4, it can be seen that the proposed method performed the best on both the Chinese-Vietnamese and Vietnamese-Chinese cross-language summarization tasks (achieving the best performance with γ set to 0.6). The LR-MC model trained on noisy data and then fine-tuned on real data showed further improvement in model performance. LR-MC showed a larger improvement in RG-2 metric, which may be due to the higher quality and better coherence of pseudo-summary texts in real data. When compared with XSIM, the LR-MC model outperformed in both Chinese-Vietnamese and Vietnamese-Chinese cross-language summarization tasks. Additionally, in comparison to KDCLS, LR-MC achieved higher summary quality, particularly exhibiting significant improvements in the Chinese-Vietnamese cross-language summarization task. This outcome can be attributed to the significant linguistic disparities between these two languages, suggesting that the guidance provided by monolingual summarization or translation may not effectively support low-resource CLS.

Compared with directly using the cross-entropy loss function to optimize the model, adding the fine-grained expected reward proposed can effectively weaken the noise. Under the Zh-Visum data, RG-1 showed an improvement of 2.59%, RG-2 showed an improvement of 4.19%, RG- L showed an improvement of 3.50%, and BERTScore showed an improvement of 0.30%. Similarly, under the Vi-Zhsum data, RG-1 demonstrated an improvement of 2.78%, RG-2 showed an improvement of 1.97%, RG- L showed an improvement of 1.87%, and BERTScore showed an improvement of 0.34%. Using the reward expectation

calculated by the word correlation and word missing degree between the real source language summary and the generated target summary can further improve the model performance, compared with the method of using ROUGE as the reward expectation and the cross-entropy loss function in RG-*L*. This shows that the fine-grained reinforcement learning method proposed has good performance in both Chinese-Vietnamese and Vietnamese-Chinese CLS tasks, as well as in noisy data with short or long texts. It can weaken the impact of noise brought by pseudo-target language summaries to some extent.

4.5.2 Ablation experiment

To verify the effect of the reinforcement learning reward based on word correlation and word missing degree on the performance of the model, two single modules were used in the ablation experiment, and the results are shown in Table 5.

According to Table 5, both the word correlation and word missing degree between the source language summary and the generated target summary were helpful in improving the model performance. When calculating only the word missing degree between the source language summary and the generated target language summary as the expected reward (LR_{mis}), the performance decrease was more

significant, and when calculating only the word correlation as the expected reward (LR_{cor}), the performance decrease was relatively small. We believe that this is due to two reasons. First, when using only the word missing degree, the information obtained by the model is relatively limited. Second, the word missing degree is designed for the noise type of missing content words, and from the analysis of the noisy data, it can be seen that the proportion of missing content words is smaller than that of improper word selection.

4.5.3 γ parameter experiment

From Table 6, it can be seen that the model performed the best when γ was set to 0.6. As γ decreased, the proportion of reward expectation increased, and the model performance did not reach its optimal level. Based on the experimental results on the decoding of the test set data, it was found that the increase in the proportion of reward expectation resulted in a higher proportion of out-of-vocabulary words in the decoded summary, which was the main reason for the decrease in summary generation quality.

We believe that using reinforcement learning reward as the optimization objective function, the word-level reward based on the source language

Table 4 Results of comparison with baseline models

Model	Zh-Visum _{Filter}				Vi-Zhsum _{Filter}			
	RG-1	RG-2	RG- <i>L</i>	BERTScore	RG-1	RG-2	RG- <i>L</i>	BERTScore
Sum-Tra	17.78	8.13	16.54	63.29	19.65	6.76	15.70	60.59
Tra-Sum	20.92	10.24	18.97	65.36	22.61	8.24	16.89	61.61
NCLS (Zhu et al., 2019)	22.69	10.26	20.96	66.66	22.45	8.88	18.87	62.87
MCLAS (Bai et al., 2021)	23.19	10.88	21.35	66.81	22.71	9.30	19.36	62.82
KDCLS (Nguyen and Luu, 2022)	23.21	10.91	21.36	66.85	22.82	9.31	19.43	62.82
LR-ROUGE (Yoon et al., 2021)	23.16	10.83	21.36	66.79	22.63	9.30	19.31	62.81
XSIM (Dou et al., 2020)	22.89	10.53	20.96	66.76	22.76	9.25	19.23	62.71
LR-MC ($\gamma = 1$)	22.81	10.49	20.87	66.71	22.66	9.14	19.21	62.66
LR-MC ($\gamma = 0.6$)	23.40	10.93	21.60	66.91	23.29	9.32	19.57	62.87

Best results are in bold

Table 5 Ablation experiment results

Model	Zh-Visum _{Filter}				Vi-Zhsum _{Filter}			
	RG-1	RG-2	RG- <i>L</i>	BERTScore	RG-1	RG-2	RG- <i>L</i>	BERTScore
LR_{mis}	22.78	10.67	21.02	66.73	22.50	9.13	19.18	62.74
LR_{cor}	23.15	10.85	21.38	66.76	22.83	9.23	19.28	62.83
LR-MC ($\gamma=0.6$)	23.40	10.93	21.60	66.91	23.29	9.32	19.57	62.87

LR_{mis} : only the missing degree is calculated as the expectation reward. LR_{cor} : only the correlation is calculated as the expected reward. Best results are in bold

contains more word-level information, but does not include the logical relationships or sequence features between words in the target language. Compared with the Chinese-Vietnamese CLS of short texts, in the Vietnamese-Chinese CLS of long texts, the proportions of word order and logical relationships between words are larger. This is also the reason why the performance of the model decreases more quickly when increasing the proportion of expected reward in the Vietnamese-Chinese CLS. Therefore, even though the word-level reward based on the source language summary designed in this study has an encouraging effect in reducing noise, it is not recommended to use this reward alone to train the model. Using the expected reward in combination with cross-entropy loss can better learn the word order information between target language words while reducing noise.

4.5.4 Exploring the effect of noisy data on model performance

To fully investigate the impact of noisy data on neural network models, we conducted comparison experiments using the basic Transformer framework. The specific results are shown in Table 7.

Table 7 shows that neural network models were sensitive to noisy data, and that filtering out noisy

data was more conducive to model learning when using the same amount of data for training. In the Chinese-Vietnamese cross-language summarization dataset, the data noise was relatively high, and training the model with the top 100 000 high-quality data samples was still more conducive to generating readable summaries than training with unfiltered 200 000 data samples. In comparison, the Vi-Zhsum dataset had a smaller size but relatively high quality. Training the model by filtering out the top 30% of the data resulted in slightly worse performance in RG-2 and RG-*L* metrics than using all data for training, but the noise had a negative impact regardless of the dataset. Therefore, starting from noisy data, it is necessary to study cross-language summarization between Chinese and Vietnamese.

4.5.5 Case analysis

Table 8 presents examples of summaries generated by different summarization models, using the Vi-Zhsum task as an example. From Table 8, it can be seen that the method proposed generated summaries with the highest quality among several summarization models. The base model trained on unfiltered data (Transformer-all) generated less summary information. After further training the model using high-quality data, all models attempted to generate

Table 6 Experimental results at different γ values

γ	Zh-Visum _{Filter}				Vi-Zhsum _{Filter}			
	RG-1	RG-2	RG- <i>L</i>	BERTScore	RG-1	RG-2	RG- <i>L</i>	BERTScore
1.0	22.81	10.49	20.87	66.71	22.66	9.14	19.21	62.66
0.8	23.22	10.83	21.37	66.75	22.81	9.23	19.35	62.97
0.6	23.40	10.93	21.60	66.91	23.29	9.32	19.57	62.87
0.4	23.20	10.95	21.35	66.71	22.55	9.14	19.20	62.78
0.2	22.91	10.69	21.07	66.58	22.16	8.90	18.93	62.42

γ is the scale factor between the cross-entropy loss and the expected reward loss. Best results are in bold

Table 7 Experimental results under different noisy data

Model	Zh-Visum				Vi-Zhsum			
	RG-1	RG-2	RG- <i>L</i>	BERTScore	RG-1	RG-2	RG- <i>L</i>	BERTScore
All	21.77	9.98	20.09	66.37	22.26	9.00	19.03	62.56
Filter-BERT/MGF	22.69	10.26	20.96	66.66	22.46	8.88	18.87	62.87
Filter-ROUGE	22.64	10.19	20.89	66.56	22.37	8.79	18.87	62.63
Random	20.98	9.48	19.33	65.96	21.01	7.79	17.61	62.12

All: model training using 200 000 unfiltered Zh-Visum and 115 798 unfiltered Vi-Zhsum samples. Filter-BERT/MGF: model training with filtered Zh-Visum data using BERTScore and with Vi-Zhsum filtered data using MGFScore. Filter-ROUGE: model training with filtered Zh-Visum and Vi-Zhsum data using ROUGE. Random: model training with an equal amount of randomly selected data from unfiltered data. Best results are in bold

Table 8 Summaries generated by different models

Text: Có 2 trường hợp được cách ly ngay sau khi nhập cảnh tại An Giang. 76 ca ghi nhận trong nước tại TP. Trong đó, 71 ca được phát hiện trong khu cách ly hoặc khu đã được phong tỏa. Việt Nam có tổng cộng 11.289 ca ghi nhận trong nước và 1.689 ca nhập cảnh. Số lượng ca mắc mới tính từ ngày 27/4 đến nay là 9.719 ca, trong đó có 2.280 bệnh nhân đã được công bố khỏi bệnh...(Two cases were immediately isolated upon entering Anjiang, Ho Chi Minh City has recorded 76 cases, 71 of which have been isolated or blocked. Vietnam has accumulated 11 289 domestic cases and 1689 imported cases. From April 27 to now, there have been 9719 new cases, 2280 of which have recovered...)
Ref-Sum: 卫生部6月20日上午发布的COVID-19流行病学公报说还有76例COVID-19病例，胡志明市仍然是最多的有46。迄今为止越南共记录了12 978 名患者。(On the morning of June 20 th , the Ministry of Health released the COVID-19 epidemic bulletin, which stated that there are still 76 cases of COVID-19, with Ho Chi Minh City having the most cases at 46. So far, Vietnam has recorded 12 978 patients.)
Transformer-all: 卫生部#月#日下午发布的COVID-19流行病学公报说还有#<unk>例COVID-19病例，其中胡志明市继续保持最多为<unk>例。(On the afternoon of # #, the Ministry of Health released the COVID-19 epidemic bulletin, which stated that there are still <unk> cases of COVID-19, with Ho Chi Minh City still having the highest number of cases at <unk>.)
LR-ROUGE: 卫生部#月#日下午发布的COVID-19 流行病学公报说还有#<unk>例COVID-19病例，其中胡志明市仍然是最多的，是最多的有基础疾病。(On the afternoon of # #, the Ministry of Health released the COVID-19 epidemic bulletin, which stated that there are still <unk> cases of COVID-19, with Ho Chi Minh City still were the highest, has the highest number of patients with the basic disease.)
LR-MC ($\gamma=1.0$): 卫生部#月#日下午发布的COVID-19流行病学公报说还有#<unk>例COVID-19病例，其中胡志明市继续保持最多为<unk>例。白天宣布治愈<unk>名患者。(On the afternoon of # #, the Ministry of Health released the COVID-19 epidemic bulletin, which stated that there are still <unk> cases of COVID-19, with Ho Chi Minh City still having the highest number of cases at <unk>. Announced the cure of <unk> patients during the day.)
LR-MC ($\gamma=0.6$): 卫生部#月#日下午发布的COVID-19流行病学公报说还有#<unk>例COVID-19病例，其中胡志明市仍然是最多的有<unk>例。迄今为止越南共记录了<unk>名患者。(On the afternoon of # #, the Ministry of Health released the COVID-19 epidemic bulletin, which stated that there are still <unk> cases of COVID-19, with Ho Chi Minh City still having the highest number of cases at <unk>. So far, Vietnam has recorded <unk> patients.)

The corresponding English translations of the summaries directly generated by different models are in brackets

more informative summaries, but only the Vi-Zhsum fine-grained reinforcement learning summarization model proposed generated the key information: “So far, Vietnam has recorded <unk> patients.”

5 Conclusions

In this paper we analyze and study the noise problem in Chinese-Vietnamese cross-language summarization and propose a fine-grained reinforcement learning cross-language summarization method for the existence of two types of noise in pseudo-target summarization: improper word selection and lack of content words. Using the real source summary and the generated summary as a benchmark, the method calculates the expected reward based on the word correlation and missingness between the source and generated summaries to weaken noise. The traditional cross-entropy loss between the pseudo-target language summary and the generated summary is also retained to learn the word order relationship between the target languages. The combination of reinforcement learning loss and cross-entropy loss is

used as the optimization objective for model training, reducing the negative impact of noisy data when directly using pseudo-target language summaries to train the model and enhancing the quality of generated summaries. In addition, the experiments explore the impact of noisy data on neural network models, and the results show that high-quality data are more conducive to model training.

Contributors

Yuxin HUANG designed the research. Yumeng GAO processed the data. Huailing GU drafted the paper. Tong PAN and Zhengtao YU helped organize the paper. Huailing GU and Jialong XU revised and finalized the paper.

Compliance with ethics guidelines

Yuxin HUANG, Huailing GU, Zhengtao YU, Yumeng GAO, Tong PAN, and Jialong XU declare that they have no conflict of interest.

Data availability

Due to the nature of this research, participants of this study did not agree for their data to be shared publicly, so the supporting data are not available.

References

- Ayana, Shen SQ, Chen Y, et al., 2018. Zero-shot cross-lingual neural headline generation. *IEEE/ACM Trans Audio Speech Lang Process*, 26(12):2319-2327. <https://doi.org/10.1109/TASLP.2018.2842432>
- Bai Y, Gao Y, Huang HY, 2021. Cross-lingual abstractive summarization with limited parallel resources. Proc 59th Annual Meeting of the Association for Computational Linguistics and the 11th Int Joint Conf on Natural Language Processing, p.6910-6924. <https://doi.org/10.18653/v1/2021.acl-long.538>
- Böhm F, Gao Y, Meyer CM, et al., 2019. Better rewards yield better summaries: learning to summarise without references. Proc Conf on Empirical Methods in Natural Language Processing and the 9th Int Joint Conf on Natural Language Processing, p.3110-3120. <https://doi.org/10.18653/v1/D19-1307>
- Cao Y, Liu H, Wan XJ, 2020. Jointly learning to align and summarize for neural cross-lingual summarization. Proc 58th Annual Meeting of the Association for Computational Linguistics, p.6220-6231. <https://doi.org/10.18653/v1/2020.acl-main.554>
- Dou ZY, Kumar S, Tsvetkov Y, 2020. A deep reinforced model for zero-shot cross-lingual summarization with bilingual semantic similarity rewards. Proc 4th Workshop on Neural Generation and Translation, p.60-68. <https://doi.org/10.18653/v1/2020.ngt-1.7>
- Dyer C, Chahuneau V, Smith NA, 2013. A simple, fast, and effective reparameterization of IBM Model 2. Proc Conf of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, p.644-648.
- Hermann KM, Kočiský T, Grefenstette E, et al., 2015. Teaching machines to read and comprehend. Proc 28th Int Conf on Neural Information Processing Systems, p.1693-1701.
- Hu BT, Chen QC, Zhu FZ, 2015. LCSTS: a large scale Chinese short text summarization dataset. Proc Conf on Empirical Methods in Natural Language Processing, p.1967-1972. <https://doi.org/10.18653/v1/D15-1229>
- Javed A, Ali Khan A, 2022. Shot classification and replay detection for sports video summarization. *Front Inform Technol Electron Eng*, 23(5):790-800. <https://doi.org/10.1631/FITEE.2000414>
- Jiang SY, Tu DB, Chen XS, et al., 2022. ClueGraphSum: let key clues guide the cross-lingual abstractive summarization. <https://arxiv.org/abs/2203.02797>
- Kang XM, Zhao Y, Zhang JJ, et al., 2020. Dynamic context selection for document-level neural machine translation via reinforcement learning. Proc Conf on Empirical Methods in Natural Language Processing, p.2242-2254. <https://doi.org/10.18653/v1/2020.emnlp-main.175>
- Kim S, Jang JY, Jung M, et al., 2021. A model of cross-lingual knowledge-grounded response generation for open-domain dialogue systems. Findings of the Association for Computational Linguistics, p.352-365. <https://doi.org/10.18653/v1/2021.findings-emnlp.33>
- Kumar G, Foster G, Cherry C, et al., 2019. Reinforcement learning based curriculum optimization for neural machine translation. Proc Conf of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, p.2054-2061. <https://doi.org/10.18653/v1/N19-1208>
- Lai H, Gao YM, Huang YX, et al., 2022. Evaluation method of text generation based on multi-granularity feature. *J Chin Inform Process*, 36(3):45-53, 63 (in Chinese).
- Leuski A, Lin CY, Zhou L, et al., 2003. Cross-lingual C*ST*RD: English access to Hindi information. *ACM Trans Asian Lang Inform Process*, 2(3):245-269. <https://doi.org/10.1145/979872.979877>
- Li HQ, Huang J, Cao Z, et al., 2023. Stochastic pedestrian avoidance for autonomous vehicles using hybrid reinforcement learning. *Front Inform Technol Electron Eng*, 24(1):131-140. <https://doi.org/10.1631/FITEE.2200128>
- Li P, Tang C, Xu XH, 2021. Video summarization with a graph convolutional attention network. *Front Inform Technol Electron Eng*, 22(6):902-913. <https://doi.org/10.1631/FITEE.2000429>
- Liang YL, Meng FD, Zhou CL, et al., 2022. A variational hierarchical model for neural cross-lingual summarization. Proc 60th Annual Meeting of the Association for Computational Linguistics, p.2088-2099. <https://doi.org/10.18653/v1/2022.acl-long.148>
- Lim JM, Kang IS, Lee JH, 2004. Multi-document summarization using cross-language texts. Proc NTCIR-4.
- Lin CY, 2004. ROUGE: a package for automatic evaluation of summaries. Proc 4th Workshop on Annual Meeting of the Association for Computational Linguistics, p.74-81.
- Nguyen TT, Luu AT, 2022. Improving neural cross-lingual abstractive summarization via employing optimal transport distance for knowledge distillation. Proc 36th AAAI Conf on Artificial Intelligence, 36(10):11103-11111. <https://doi.org/10.1609/aaai.v36i10.21359>
- Orăsan C, Chiorean OA, 2008. Evaluation of a cross-lingual Romanian-English multi-document summariser. Proc Int Conf on Language Resources and Evaluation.
- Ouyang J, Song BY, McKeown K, 2019. A robust abstractive system for cross-lingual summarization. Proc Conf of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, p.2025-2031. <https://doi.org/10.18653/v1/N19-1204>
- Paulus R, Xiong CM, Socher R, 2017. A deep reinforced model for abstractive summarization. <https://arxiv.org/abs/1705.04304>
- Rennie SJ, Marcheret E, Mroueh Y, et al., 2017. Self-critical sequence training for image captioning. Proc IEEE Conf on Computer Vision and Pattern Recognition, p.1179-1195. <https://doi.org/10.1109/CVPR.2017.131>
- Rippeth E, Post M, 2022. Additive interventions yield robust multi-domain machine translation models. Proc 7th Conf on Machine Translation, p.220-232.
- Takase S, Okazaki N, 2020. Multi-task learning for cross-lingual abstractive summarization. <https://arxiv.org/abs/2010.07503>
- Unanue IJ, Parnell J, Piccardi M, 2021. BERTTune: fine-tuning neural machine translation with BERTScore. <https://arxiv.org/abs/2106.02208>
- Vaswani A, Shazeer N, Parmar N, et al., 2017. Attention is all you need. Proc 31st Int Conf on Neural Information Processing Systems, p.6000-6010.
- Wang JA, Meng FD, Lu ZY, et al., 2022. ClidSum: a benchmark dataset for cross-lingual dialogue summarization. Proc Conf on Empirical Methods in Natural Language

- Processing, p.7716-7729.
<https://doi.org/10.18653/v1/2022.emnlp-main.526>
- Wu LJ, Zhu JH, He D, et al., 2019. Machine translation with weakly paired documents. Proc Conf on Empirical Methods in Natural Language Processing and the 9th Int Joint Conf on Natural Language Processing, p.4375-4384. <https://doi.org/10.18653/v1/D19-1446>
- Xiong LL, Tang Y, Liu CS, et al., 2023. A home energy management approach using decoupling value and policy in reinforcement learning. *Front Inform Technol Electron Eng*, 24(9):1261-1272.
<https://doi.org/10.1631/FITEE.2200667>
- Yoon W, Yeo YS, Jeong M, et al., 2021. Learning by semantic similarity makes abstractive summarization better.
<https://arxiv.org/abs/2002.07767>
- You YJ, Jia WJ, Liu TY, et al., 2019. Improving abstractive document summarization with salient information modeling. Proc 57th Annual Meeting of the Association for Computational Linguistics, p.2132-2141.
<https://doi.org/10.18653/v1/P19-1205>
- Zhang TY, Kishore V, Wu F, et al., 2020. BERTScore: evaluating text generation with BERT.
<https://arxiv.org/abs/1904.09675>
- Zhao H, Xie J, Lv Y, et al., 2013. Common error analysis of machine translation output. The 9th China Workshop on Machine Translation.
- Zhao J, Zhao YP, Wang WX, et al., 2022. Coach-assisted multi-agent reinforcement learning framework for unexpected crashed agents. *Front Inform Technol Electron Eng*, 23(7):1032-1042.
<https://doi.org/10.1631/FITEE.2100594>
- Zhou J, Ke P, Qiu XP, et al., 2023. ChatGPT: potential, prospects, and limitations. *Front Inform Technol Electron Eng*, early access.
<https://doi.org/10.1631/FITEE.2300089>
- Zhu JN, Wang Q, Wang YN, et al., 2019. NCLS: neural cross-lingual summarization. Proc Conf on Empirical Methods in Natural Language Processing and the 9th Int Joint Conf on Natural Language Processing, p.3054-3064. <https://doi.org/10.18653/v1/D19-1302>