Controllable News Comment Generation based on Attribute Level Contrastive Learning

Hanyi Zou[1,2], Nan Xu[1,3], Qingchao Kong[1,2*], Wenji Mao[1,2]
[1]Institute of Automation, Chinese Academy of Sciences, Beijing, China
[2]University of Chinese Academy of Sciences, Beijing, China
[3]Beijing Wenge Technology Co.,Ltd, Beijing, China
{zouhanyi2021, xunan2015, qingchao.kong, wenji.mao}@ia.ac.cn

2023 IEEE International Conference on Intelligence and Security Informatics (ISI) | 979-8-3503-3773-0/23/$31.00 ©2023 IEEE | DOI: 10.1109/ISI58743.2023.10297146

***Abstract*—News comments provide a convenient way for people to express opinions and exchange ideas. Positive comments encourage a harmonious discussion atmosphere within news media communities. In contrast, offensive or insulting comments may result in cyberbullying and personal psychological trauma, which have particular practical impacts in security-related domains. The automatic generation of news comments with controllable attributes (e.g. sentiment) to assist users and news platform administrators is greatly needed. However, existing research for news comment generation has not addressed the controllable issue yet. On the other hand, existing methods for controllable text generation focus on token-level constraints, which are not applicable to controlling the sentence level attributes for news comment generation. To address this challenging issue, in this paper, we propose an attribute-level contrastive learning method for controllable news comment generation. To apply attribute level constraints on the generated text, our method considers the attributes of the generated comments and the pre-defined attributes as different views of the same attribute, and maximizes their similarity during the training process. We conduct experiments on two publicly available news comment datasets, and the experimental results show that our model achieves competitive performance in news comment generation and attribute controllability.**

***Index Terms*—controllable text generation, news comment generation, attribute level constraints, contrastive learning**

I. Introduction

With the development of Internet technology, online news platforms have become an essential way for the public to exchange information and express personal opinions. However, some users may also use the platforms to post insulting or offensive comments. These negative comments not only impact user experience but also pose potential challenges to cybersecurity [1], which is an important research issue in security-related domain applications. Therefore, the automatic generation of news comments with controllable attributes (e.g. sentiment) to assist users and news platform administrators is greatly needed. To address this challenging issue, we propose the new task of controllable news comment generation, aiming to generate comments that are relevant to the news context and comply with attribute level constraints, which has not been explored by previous studies.

Existing research on automatic news comment generation primarily focuses on how to understand important information in news articles [2] and generate diversity comments [3]. However, previous methods have the problem of generating generic and irrelevant comments, as news documents are usually too long for the models to learn important information that can guide comment generation. More importantly, existing methods for news comment generation have not considered the controllability issue, thus they are unable to control attribute level constraints, such as sentiment and topic of the generated comments.

On the other hand, existing controllable text generation methods are typically trained within the maximum likelihood estimation framework, where the objective is to minimize the negative log-likelihood between predicted tokens and ground truth. This is a token-level constraint instead of attribute level control. For controllable news comment generation, evaluation metrics such as ROUGE [4] compare the overall overlap between generated text and reference text, while the attribute-controlled accuracy measures the matching degree between predefined attributes and the attributes of the generated text. This mismatch brings about additional challenges to controllable news comment generation. To address the challenges mentioned above, this paper proposes an **Attri**bute level **C**ontrastive **L**earning (**AttriCL**) method for controllable news comment generation task. The objective of this task is to generate comments that adhere to pre-defined attributes, ensuring that the attributes of the generated comments align with the specified constraints. **AttriCL** aims to minimize the distance between representations of comments with the same attributes while maximizing the distances between representations of comments with different attributes. Thus we adopt the contrastive learning approach and consider the predefined attributes and the attributes of generated comments as contrastive samples. We construct an attribute discriminator to determine whether the attributes of the generated comments match the pre-defined attributes. Intuitively, this training approach encourages the model to generate comments that comply with the pre-defined attributes, thereby achieving a controllable generation of news comment.

The main contributions of this work are summarized as follows:

- We propose the new task of controllable news comment generation and develop an attribute level contrastive learning method AttriCL to address this challenging task.

* Corresponding Author

- Our method designs an attribute discriminator to impose the attribute level constraints on the generated text, and facilitates the fine-grained combination of controllable comment generation at content level.
- We extend a large-scale dataset for the controllable news comment generation task and the experimental results show the advantages of our method in terms of both text generation and attribute controllability.

II. Related Work

News comment generation builds on the combination of two tasks, news comment generation and controllable text generation. Below we review the related work from these two aspects, and also introduce the contrastive learning technique.

News comment generation is proposed by Qin et al. [5] and Zheng et al. [6]. Qin et al. [5] released a dataset containing millions of real comments and developed weighted-based automatic evaluation metrics for news comments generation. However, this work only utilized a seq2seq model as the baseline but did not propose a new model for news comment generation. Zheng et al. [6] proposed a gated attention neural network (GANN) to select news context self-adaptively, and applied generative adversarial nets to improve GANN. Yang et al. [7] simulated human habits when reading the news and designed a deep architecture called DeepCom for news comment generation. It consists of a reading network and a generation network. The reading network first extracts information segments from the news, and then the generation network uses these segments to generate comments. Experimental results have shown the effectiveness of this architecture. A similar architecture was also used by Li et al. [2]. They used graph convolutional networks as encoder and modeled the internal structure of news as a topic interaction graph. Recently, Wang et al. [3] proposed a variational generative clustering algorithm to mine topic information from reader comments. Bernoulli distribution estimating on news content was introduced to select saliency information. The above two important pieces of information were merged into the decoder to generate diverse and informative comments.

Controllable text generation focuses on generating sentences containing pre-specified attributes, such as sentiment and topic. Early controllable text generation methods are mainly based on recurrent neural network (RNN) [8] and variational auto-encoder (VAE) [9]. More recently, controllable text generation using Transformer-based models has become an extensive area of research, especially with the emergence of large pretrained language models like GPT-2. Keskar et al. [10] proposed 55 kinds of control codes to train a Transformer-based language model containing over one billion parameters. However, this model could only perform coarse-grained control according to control codes. Training large-scale language models with all parameters requires a significant amount of computational resources. Dathathri et al. [11] proposed a low-computational-cost approach to achieve controlled text generation, which developed a plug-and-play attribute discriminator to drive the hidden activations of language model and guide text generation. However, this method was slow during inference because it updates gradients at the token level. To solve this problem, Krause et al. [12] and Liu et al. [1] reweighted the probability distribution of the pre-trained language model's output so that guided text generation during decoding to improve inference speed.

Contrastive learning methods are mostly used in computer vision tasks, including object detection and image segmentation [13]. In natural language processing, contrastive learning methods are mainly applied in model pre-training or natural language understanding tasks. Take text summarization as an example domain, Iter et al. [14] applied contrastive learning to model pre-training by establishing sentence-level contrast to improve discourse-level representations. Xu et al. [15] employed the core idea of the contrastive learning method to consider the document, the gold summary and the model-generated summary as different views of the same latent representation, maximizing their similarities during training. Zhong et al. [16] defined extractive text summarization as a semantic text matching problem and solved it using contrastive learning methods. Liu and Liu [17] proposed a contrastive learning based text summarization framework, which calculates the contrastive loss between the generated summary and candidate summaries through a scoring model and optimizes the model parameters with contrastive loss.

However, the aforementioned methods have similar limitations in terms of controllable news comment generation tasks, which ignored the overall constraints on the generated comments. Inspired by Xu et al. [15], we propose AttriCL, a controllable news comment generation method based on attribute level contrastive learning. AttriCL enhances the similarities between pre-defined attributes and the attributes of the generated comments, thereby generating text that complies with attribute constraints.

III. Problem Definition

In text generation tasks, given a prompt text, the following text can be generated in an autoregressive manner. This process can be formalized as follows:

$$p(x_t, ..., x_l | x_1, ..., x_{t-1}) = \prod_{i=t}^{l} p(x_i | x_1, ..., x_{i-1}) \quad (1)$$

For the controllable news comment generation task, there exists a similar definition as in the controllable text generation task. Specifically, the target attributes [11] or control codes [10] can be used as conditions to control the sentiment or topic of text generation, i.e.,

$$p(x_t, ..., x_l | \textbf{att}, \{x_1, ..., x_{t-1}\}) = \prod_{i=t}^{l} p(x_i | \textbf{att}, \{x_1, ..., x_{-1}\}) \quad (2)$$

where **att** is the controllable attribute, $\{x_1, ..., x_{t-1}\}$ represents the news content. The output text $\{x_t, ..., x_l\}$ needs to comply with the constraints of the control attributes **att** and be contextually relevant to the news content $\{x_1, ..., x_{t-1}\}$.

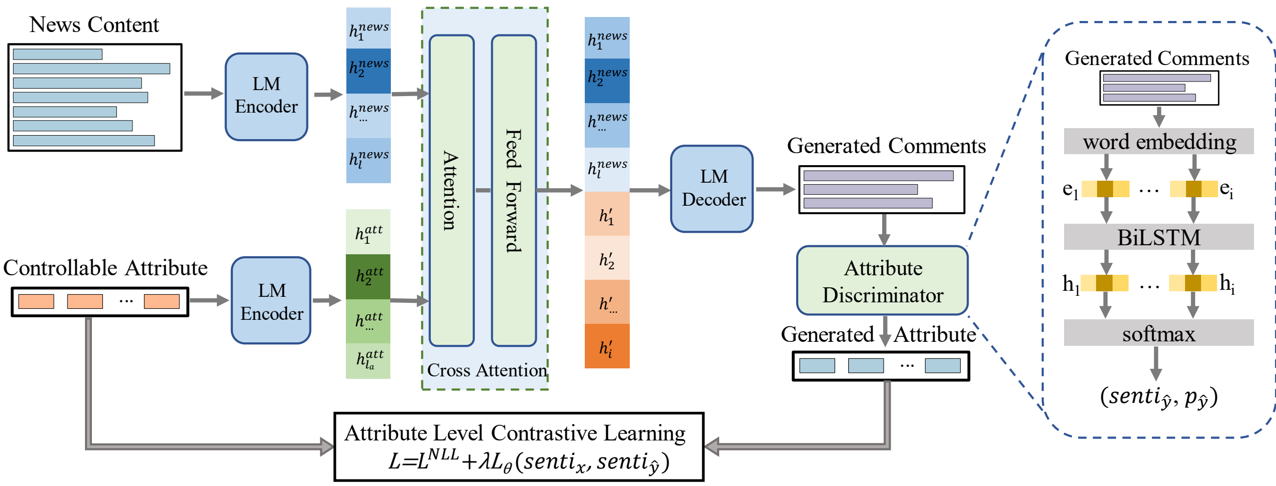

Fig. 1. Overall Architecture of Controllable News Comment Generation based on Attribute Level Contrastive Learning

IV. Proposed Method

In this section, we introduce our contrastive learning model **AttriCL** for controllable news comment generation. We first introduce the overall architecture of the proposed model. Then, we introduce how the contrastive learning method is applied to this model.

A. Model Design for Controllable News Comment Generation

Controllable news comment generation is essentially a text generation task, and the seq2seq framework [18] is widely used. In this work, we employ a pretrained Transformer-based language model (LM) GPT2 [19] as the foundation of our approach. The pretrained LM is stacks of Transformer blocks [20], each includes multi-head attention, position-wise feed-forward networks, layer normalization and other operations that have been proven to improve performance on natural language processing tasks. According to the definition of the seq2seq framework, we can divide the LM into two parts: the encoder (LM_{en}) and the decoder (LM_{de}).

$$o_t = LM(x_{:t-1}) = LM_{de}(LM_{en}(x_{:t-1})) \tag{3}$$

where o_t denotes the logits of next token x_t. $h_{:t}$ indicates the latent representation of input text.

To integrate control attributes and news content, we introduce a cross attention layer between the encoder and decoder. The specific model architecture is shown in Fig. 1.

We can obtain representations of the news content and controllable attribute through the encoder. This process can be formalized as follows:

$$h_{:l}^{news} = LM_{en}(news) = LM_{en}(x_{:l}) \tag{4}$$

$$h_{:l_a}^{att} = LM_{en}(attribute) = LM_{en}(att_{:l_a}) \tag{5}$$

where $h_{:l}^{news}$ and $h_{:l_a}^{att}$ denote the latent representation of news content $x_{:l}$ and controllable attribute $att_{:l_a}$ respectively. l and l_a are the length of news content and controllable attribute respectively. Then, we obtain an integrated representation of the news content and controllable attribute utilizing the cross attention layer:

$$h_i^{'} = CrossAtt(h_{:l}^{news}, h_{:l_a}^{att}) \tag{6}$$

By concatenating $h_i^{'}$ to $h_{:l}^{news}$ and passing them to the decoder, we can obtain the logits and probability distribution of the next token.

$$o_{i+1} = LM_{de}([h_{:l}^{news}; h_i^{'}]) \tag{7}$$

$$p(x_{i+1}|att, x_{:l}) = Softmax(o_{i+1}) \tag{8}$$

During the training process, the cross attention layer is optimized using contrastive loss.

B. Attribute Level Contrastive Learning

The contrastive loss is obtained from attribute similarity. In order to calculate the attribute similarity between the pre-defined attribute and the attribute of the generated comments, we utilize the sentiment analysis tool from paddleNLP[1] as the discriminator to extract the sentiment attributes of the generated comments. The discriminator was trained on Chinese corpus using BiLSTM, which is a competitive Chinese sentiment analysis model. For a generated comment $\hat{y}$, we can obtain its sentiment attribute $senti_{\hat{y}}$ and attribute probability $p_{\hat{y}}$ through the sentiment analysis tool.

$$senti_{\hat{y}}, p_{\hat{y}} = SentiTool(\hat{y}) \tag{9}$$

where *SentiTool* represents the sentiment analysis tool. The value of $p_{\hat{y}}$ ranges from 0.5 to 1, where a higher value indicates a stronger sentiment polarity. The similarity between input sentiment attribute $senti_x$ and the sentiment of generated comments $senti_{\hat{y}}$ is shown as follows:

$$sim(senti_x, senti_{\hat{y}}) = \begin{cases} -p_{\hat{y}} & \text{if } senti_x \neq senti_{\hat{y}} \\ p_{\hat{y}} & \text{if } senti_x = senti_{\hat{y}} \end{cases} \tag{10}$$

[1] https://github.com/PaddlePaddle/PaddleNLP

if the sentiment polarities of $senti_x$ and $senti_{\hat{y}}$ are different, the value of *sim(*$senti_x, senti_{\hat{y}}$*)* is negative sentiment attribute probability. Conversely, if their sentiment polarities are the same, the *sim(*$senti_x, senti_{\hat{y}}$*)* is assigned the value of the sentiment attribute probability $p_{\hat{y}}$. The larger the attribute probability, the greater the similarity.

To make $senti_x$ and $senti_{\hat{y}}$ closer, we minimize the contrastive loss function as follows:

$$\mathcal{L}_\theta(senti_x, senti_{\hat{y}}) = 1 - sim(senti_x, senti_{\hat{y}}) \tag{11}$$

where θ represents the parameters of the cross attention layer. Our objective is to minimize the distance between $senti_x$ and $senti_{\hat{y}}$ by optimizing the cross attention layer. However, if we only constrain the control attribute, it may result in a degradation of the generated text quality. To address this problem, we incorporate the negative log-likelihood loss ($\mathcal{L}^{NLL}$) into the overall loss. To enforce the similarities between the pre-defined sentiment attribute and the attribute of the generated comments, we employ the following loss function to train our model.

$$\mathcal{L} = \mathcal{L}^{NLL} + \lambda\mathcal{L}_\theta(senti_x, senti_{\hat{y}}) \tag{12}$$

where λ is the weight hyper-parameter of the sentiment similarity loss function. In the controllable news comment generation task, we consider ground truth comments to be closely related to the context of the news. To enhance the contextual relevance of generated comments, we aim to make the generated comments consistent with the ground truth comments during the training process. This can be achieved through the use of Negative Log Likelihood (NLL) loss. In addition to maintaining consistency with ground truth comments at the token level, generating comments also need to be consistent in terms of sentiment attributes. To achieve this objective, we use sentiment similarity loss to optimize the model parameters. The hyper-parameter λ is used to control the tradeoffs among contextual relevance and controllability.

Note that the pre-defined sentiment attribute $senti_x$ is fixed because it's an input controllable attribute of the comment generation model. Since the generated comment $\hat{y}$ can not be perfect, iteratively generating $\hat{y}$ makes the sentiment attribute of generated comment change toward pre-defined sentiment. This text generation model based on contrastive loss can move the generated text towards the direction that satisfies attribute constraints. Additionally, encouraging similarity between $senti_x$ and $senti_{\hat{y}}$ is not equivalent to optimizing NLL loss, as sentimental similarity loss operates at the sentiment attribute level while NLL loss operates at the token level.

V. Experiments

In this section, we first introduce the dataset used for this task. Next, we present the relevant baseline model. Subsequently, we describe the evaluation metrics employed for this task. Finally, we present and discuss the experimental results.

A. Datasets

This paper focuses on Chinese news comment generation and utilizes two datasets collected from Chinese most popular news websites.

Tencent News: This dataset was collected from Tencent News[2]. It was initially introduced in Qin et al. [5] and consists of news articles from the entertainment and sports domains. Each sample instance in the dataset includes a news title, news content, and a set of reader comments. The dataset comprises 198,112 news articles, with an average of 27 comments per article.

NetEase News: This dataset was crawled from NetEase News[3] and released by Zheng et al [6]. The dataset consists of a total of 82,787 news articles, with each article containing an average of 22.5 comments. The detailed statistics are shown in Table 1.

TABLE I
Statistics of Datasets

		train	dev	test
Tencent News	# News	191,502	5,000	1,610
	Avg. Cmts per News	27	27	27
NetEase News	# News	75,287	5,000	2,500
	Avg. Cmts per News	22.7	22.5	22.5

It should be noted that the comments in both of these datasets lack sentiment labels. To adapt to the task in this paper, we utilize sentiment analysis tools from paddleNLP to extract sentiment labels for news comments. After this preprocessing, these datasets include the following information: news title, news content, news comment and comment sentiment.

B. Baseline Models

In this section, we describe the baseline models used. For a fair comparison, we evaluate the performance of our model from two aspects: text generation quality and controllability. We select baseline models for both parts. The following models are used to evaluate the text generation quality of the proposed model.

Seq2seq [5]: This model employs a sequence-to-sequence framework, utilizing RNNs with attention mechanism as the encoder and decoder. The input can be either news title (**T**) or news title combined with news content (**TC**).

GANN [6]: This model adopts a framework similar to seq2seq and incorporates a gated attention layer between the encoder and decoder.

Self-attention [21]: This model also utilizes the seq2seq framework and incorporates multiple layers of LSTM in both the encoder and decoder. Additionally, it employs multi-head attention between the encoder and decoder.

CAVE [22]: This model introduces a novel framework based on conditional variational autoencoders (CVAE) to improve the diversity of neural dialogue models. In this paper, we use

[2]news.qq.com
[3]news.163.com

this model as a baseline for evaluating the diversity of our proposed model.

In addition to evaluating the quality of the generated text, it is also necessary to evaluate the controllable performance of the proposed model.

GPT2 Fine-tuning: This work selects the GPT2 language model pre-trained on a Chinese corpus as the foundational model. To verify the performance improvement of the proposed method, we fine-tune the GPT2 model with the sentiment-news-comment dataset and calculate the experimental results.

CoCon [23]: This model is a strong baseline for controllable text generation tasks, which integrates the target content into the generated text using content conditioner, thereby achieving fine-grained level controlled text generation.

C. Evaluation Metrics

Automatic metrics are used to evaluate the degree of overlap between the generated comments and the reference comments. In this paper, we employ ROUGE [4], CIDEr [24] and METEOR [25] evaluate the performance of different models using automatic metrics following Qin et al. [5].

ROUGE (Recall-Oriented Understudy for Gisting Evaluation): This metric is used to measure the overlap of n-grams between generated and reference summaries, used for text summarization and machine translation evaluation.

CIDEr (Consensus-based Image Description Evaluation): This metric is used to evaluate image captioning by measuring the similarity between generated and human reference descriptions, considering n-gram language statistics and consensus scores.

METEOR (Metric for Evaluation of Translation with Explicit ORdering): This metric is mainly used in machine translation tasks that consider both exact word matches and paraphrased variations, using global alignment optimization for accuracy.

We use a popular NLG evaluate tool NLG-eval[4] to calculate these metrics automatically. In addition to evaluating the quality of the generated comments, we also need to assess the control performance. Similar to Dathathri et al. [11], we use sentiment accuracy and perplexity to evaluate the model's controllability sentiment accuracy represents the ratio of the generated comments that align with the pre-defined sentiment attributes. The perplexity is used to evaluate the fluency of generated text. We utilize the language model GPT2 [26] pre-trained on Chinese corpus to conduct the evaluation.

D. Results and Discussions

For the controllable news comment generation task, we primarily evaluate the model's performance based on three aspects: contextual relevance, text diversity and attribute controllability. We use reference-based evaluation metrics such as ROUGE-L, CIDEr, and METEOR to assess contextual relevance. Since the reference comments used for comparison are collected from the news platform, these comments are manually written by individuals and are relevant to the news context. If the generated comments are related to the reference comments, it indicates that the generated comments are also contextually relevant to the news.

TABLE II
TEXT GENERATION QUALITY RESULTS

Datasets	Models	ROUGE-L	CIDEr	METEOR	Distinct-3	Distinct-4
Tencent	Seq2seq-T	0.261	0.015	0.076	0.088	0.079
	Seq2seq-TC	0.280	0.021	0.088	0.121	0.122
	GANN	0.267	0.017	0.081	0.087	0.081
	Self-attention	0.280	0.019	0.092	0.117	0.121
	CAVE	0.281	0.021	0.094	0.135	0.137
	TSGen	**0.289**	**0.024**	0.107	0.176	0.196
	OURS	0.237	0.007	**0.119**	**0.836**	**0.826**
NetEase	Seq2seq-T	0.263	0.025	0.105	0.149	0.169
	Seq2seq-TC	0.268	0.035	0.108	0.178	0.203
	GANN	0.258	0.022	0.105	0.129	0.146
	Self-attention	0.265	0.034	0.110	0.174	0.204
	CAVE	0.261	0.026	0.106	0.120	0.135
	TSGen	**0.269**	**0.034**	**0.111**	0.189	0.225
	OURS	0.217	0.013	0.1085	**0.843**	**0.830**

ROUGE-L, CIDEr, and METEOR are reference-based evaluation metrics, and these metrics are in contradiction with diversity metrics. The more diverse the generated text is, the lower the values of reference-based evaluation metrics will be, while diversity metrics will be higher. As shown in TABLE II, our model outperforms all other models in terms of the distinct metric, indicating that the comments generated by our model are more diverse. However, the performance on the ROUGE-L, CIDEr and METEOR metrics is not the strongest. This is because the reference-based metrics evaluate the overlap between the generated comments and the reference comments. Since generated comments with higher diversity, the overlap with the reference comments may decrease. This indicates that our model can generate diverse reviews while maintaining contextual relevance.

In addition to evaluating the quality of the generated comments using reference-based evaluation metrics, we also estimate the controllability of the model using attribute accuracy. We select two competitive models as baseline models: CoCon [23] and GPT2 [19]. The attribute accuracy results are as follows:

TABLE III
COMMENT ATTRIBUTE CONTROLLABILITY RESULTS

Datasets	Models	Accuracy ↑ better	Perplexity ↓ better
Tencent	GPT2 Fine-tuning	0.494	90.217
	CoCon	0.712	**71.743**
	OURS	**0.735**	73.177
NetEase	GPT2 Fine-tuning	0.501	41.875
	CoCon	0.576	75.179
	OURS	**0.600**	**40.948**

[4]https://github.com/Maluuba/nlg-eval

As shown in TABLE III, our model surpasses CoCon and GPT2 Fine-tuning in controllability and outperforms these two models in text fluency. This is because we use attribute contrastive loss and NLL loss as optimization goals during model training. This loss can constrain the model from the sentence level. Thereby improving the accuracy of attribute control. This is because we added attribute contrastive loss based on NLL loss during model training. These two losses can constrain the model from the sentence level, thereby improving the accuracy of attribute control.

VI. Conclusions

In the controllable news comment generation task, the attributes of the generated comments can be seen as a different view of pre-defined attributes. We propose a method called AttriCL, which leverages attribute level contrastive learning method for controllable news comment generation task. This method minimizes the distance between the pre-defined attributes and the generated attributes during the training process. Experimental results on two Chinese datasets (Tencent News and NetEase News) demonstrate that AttriCL achieves competitive performance in terms of evaluation metrics including contextual relevance, diversity, and controllability.

Acknowledgment

This work is supported in part by the Ministry of Science and Technology of China under Grant #2020AAA0108405, the National Natural Science Foundation of China under Grants #62206287 and #72293575.

References

[1] A. Liu, M. Sap, X. Lu, S. Swayamdipta, C. Bhagavatula, N. A. Smith, and Y. Choi, "Dexperts: Decoding-time controlled text generation with experts and anti-experts," in *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021*, pp. 6691–6706.

[2] W. Li, J. Xu, Y. He, S. Yan, Y. Wu, and X. Sun, "Coherent comments generation for chinese articles with a graph-to-sequence model," in *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019*, pp. 4843–4852.

[3] W. Wang, P. Li, and H. Zheng, "Generating diversified comments via reader-aware topic modeling and saliency detection," in *Proceedings of the 35th AAAI Conference on Artificial Intelligence, AAAI 2021*, pp. 13 988–13 996.

[4] C.-Y. Lin, "Rouge: A package for automatic evaluation of summaries," in *Text summarization branches out*, 2004, pp. 74–81.

[5] L. Qin, L. Liu, V. Bi, Y. Wang, X. Liu, Z. Hu, H. Zhao, and S. Shi, "Automatic article commenting: the task and dataset," in *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018*, pp. 151–156.

[6] H. Zheng, W. Wang, W. Chen, and A. K. Sangaiah, "Automatic generation of news comments based on gated attention neural networks," *IEEE Access*, vol. 6, pp. 702–710, 2018.

[7] Z. Yang, C. Xu, W. Wu, and Z. Li, "Read, attend and comment: A deep architecture for automatic news comment generation," in *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019*, pp. 5076–5088.

[8] A. Holtzman, J. Buys, M. Forbes, A. Bosselut, D. Golub, and Y. Choi, "Learning to write with cooperative discriminators," in *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018*, pp. 1638–1649.

[9] Z. Hu, Z. Yang, X. Liang, R. Salakhutdinov, and E. P. Xing, "Toward controlled generation of text," in *Proceedings of the 34th International Conference on Machine Learning, ICML 2017*, pp. 1587–1596.

[10] N. S. Keskar, B. McCann, L. R. Varshney, C. Xiong, and R. Socher, "Ctrl: A conditional transformer language model for controllable generation," *arXiv preprint arXiv:1909.05858*, 2019.

[11] S. Dathathri, A. Madotto, J. Lan, J. Hung, E. Frank, P. Molino, J. Yosinski, and R. Liu, "Plug and play language models: A simple approach to controlled text generation," in *8th International Conference on Learning Representations, ICLR 2020*.

[12] B. Krause, A. D. Gotmare, B. McCann, N. S. Keskar, S. R. Joty, R. Socher, and N. F. Rajani, "Gedi: Generative discriminator guided sequence generation," in *Findings of the Association for Computational Linguistics: EMNLP 2021*, pp. 4929–4952.

[13] K. He, H. Fan, Y. Wu, S. Xie, and R. B. Girshick, "Momentum contrast for unsupervised visual representation learning," in *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020*, pp. 9726–9735.

[14] D. Iter, K. Guu, L. Lansing, and D. Jurafsky, "Pretraining with contrastive sentence objectives improves discourse performance of language models," in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020*, pp. 4859–4870.

[15] S. Xu, X. Zhang, Y. Wu, and F. Wei, "Sequence level contrastive learning for text summarization," in *Proceedings of the 36th AAAI Conference on Artificial Intelligence, AAAI 2022*, pp. 11 556–11 565.

[16] M. Zhong, P. Liu, Y. Chen, D. Wang, X. Qiu, and X. Huang, "Extractive summarization as text matching," in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020*, pp. 6197–6208.

[17] Y. Liu and P. Liu, "SimCLS: A simple framework for contrastive learning of abstractive summarization," in *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing,ACL/IJCNLP 2021*, pp. 1065–1072.

[18] I. Sutskever, O. Vinyals, and Q. V. Le, "Sequence to sequence learning with neural networks," in *Annual Conference on Neural Information Processing Systems, NIPS 2014*, pp. 3104–3112.

[19] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, I. Sutskever *et al.*, "Language models are unsupervised multitask learners," *OpenAI blog*, vol. 1, no. 8, p. 9, 2019.

[20] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," in *Annual Conference on Neural Information Processing Systems, NIPS 2017*, pp. 5998–6008.

[21] M. X. Chen, O. Firat, A. Bapna, M. Johnson, W. Macherey, G. Foster, L. Jones, M. Schuster, N. Shazeer, N. Parmar, A. Vaswani, J. Uszkoreit, L. Kaiser, Z. Chen, Y. Wu, and M. Hughes, "The best of both worlds: Combining recent advances in neural machine translation," in *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018*, pp. 76–86.

[22] T. Zhao, R. Zhao, and M. Eskenazi, "Learning discourse-level diversity for neural dialog models using conditional variational autoencoders," in *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017*, pp. 654–664.

[23] A. Chan, Y. Ong, B. Pung, A. Zhang, and J. Fu, "Cocon: A self-supervised approach for controlled text generation," in *9th International Conference on Learning Representations, ICLR 2021*.

[24] R. Vedantam, C. Lawrence Zitnick, and D. Parikh, "Cider: Consensus-based image description evaluation," in *Proceedings of the IEEE conference on computer vision and pattern recognition, CVPR2015*, pp. 4566–4575.

[25] S. Banerjee and A. Lavie, "Meteor: An automatic metric for mt evaluation with improved correlation with human judgments," in *Proceedings of the acl workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*, 2005, pp. 65–72.

[26] Z. Zhao, H. Chen, J. Zhang, X. Zhao, T. Liu, W. Lu, X. Chen, H. Deng, Q. Ju, and X. Du, "UER: an open-source toolkit for pre-training models," in *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019*, pp. 241–246.