PLATO-Ad: A Unified Advertisement Text Generation Framework with Multi-Task Prompt Learning

Zeyang Lei*, **Chao Zhang***, **Xinchao Xu, Wenquan Wu, Zheng-Yu Niu,**
Hua Wu, Haifeng Wang, Yi Yang and Shuanglong Li
Baidu Inc., Beijing, China
{leizeyang, zhangchao38, xuxinchao, wuwenquan01}@baidu.com
{niuzhengyu, wu_hua, wanghaifeng, yangyi15, lishuanglong}@baidu.com

Abstract

Online advertisement text generation aims at generating attractive and persuasive text ads to appeal to users clicking ads or purchasing products. While pretraining-based models have achieved remarkable success in generating high-quality text ads, some challenges remain, such as ad generation in low-resource scenarios and training efficiency for multiple ad tasks. In this paper, we propose a novel unified text ad generation framework with multi-task prompt learning, called PLATO-Ad, to tackle these problems. Specifically, we design a three-phase transfer learning mechanism to tackle the low-resource ad generation problem. Furthermore, we present a novel multi-task prompt learning mechanism to efficiently utilize a single lightweight model to solve multiple ad generation tasks without loss of performance compared to training a separate model for each task. Finally, we conduct offline and online evaluations. Experiment results show that PLATO-Ad significantly outperforms the state-of-the-art on both offline and online metrics. PLATO-Ad has been deployed in a leading advertising platform with 3.5% CTR improvement on search ad descriptions and 10.4% CTR improvement on feed ad titles.

1 Introduction

In recent years, online advertising has been regarded as one of the most popular ways of internet monetization. A captivating and persuasive ad can greatly improve the probability of users clicking the ads or purchasing recommended products (Jansen and Resnick, 2005). Thus, advertisers usually spare no effort to improve the quality of displayed ads. Traditionally, some advertisers may manually design ads for high quality. However, this approach suffers from low efficiency and high labor costs. some other works design pre-defined templates (Fujita et al., 2010; Thomaidou et al., 2013)

Figure 1: Example ads in search ad and feed ad systems.

or use recurrent neural networks (RNN) (Hughes et al., 2019) to automatically create ads. However, these ads generated by templates are generally rigid and not unappealing enough to users. Over 15% ads generated by RNN are reported to be non-sense or bad (Hughes et al., 2019), which hinders the application of this method in real-world business scenarios with high-quality standards.

Recently, with the remarkable success of the pretraining plus fine-tuning paradigm, some works utilize pretraining-based natural language generation methods to generate content-rich, diverse and attractive ads based on large-scale corpora (Wang et al., 2021; Zhang et al., 2021a,b). Although these methods achieve significant success in generating high-quality ads, they still face great challenges when applied in real-world commercial scenarios. **Firstly**, the performance of these pretraining-based models relies heavily on large-scale high-quality training corpora, which may be yet difficult to be obtained in some low-resource [1] scenarios. **Secondly**, these work usually train a separate model for each task respectively when applied in multiple ad tasks, e.g., ad title and description generation (Wang et al., 2021; Zhang et al., 2021a), copywriting generation (Zhang et al., 2021b), selling point generation (Guo et al., 2021). This is inefficient and expensive for multiple ad tasks that require multiple copies of the model's parameters, with also ignoring the associations between multiple ad generation tasks (Lester et al., 2021).

*Equal contribution.

[1]Here, low resource denotes that relatively less training data can be available .

Proceedings of EMNLP 2022 Industry Track, pages 522–530
December 9–11, 2020. ©2022 Association for Computational Linguistics

To this end, we propose a novel unified text ad generation framework with multi-task prompt learning, called PLATO-Ad, to tackle the aforementioned problems. Concretely, the model architecture of PLATO-Ad is a Transformer-based pre-trained language model with 12 transformer blocks. **To effectively address the low-resource ad generation problem**, we propose a three-phase transfer learning mechanism to train PLATO-Ad. Specifically, we first pre-train PLATO-Ad on generic-domain text corpus to equip the model with the ability to generate fluent natural sentences. Then, in the second phase, we consecutively post-pretrain PLATO-Ad on the datasets of multiple resource-rich ad generation tasks where massive data is available and open-domain question answering (QA) datasets, which enables the model to learn to generate ad-domain and commonsense-enriched text. Finally, we use the prompting mechanism (Liu et al., 2021) to transfer the well-trained PLATO-Ad in the second phase to low-resource ad generation tasks. In this way, we can generate high-quality ads for those low-resource ad generation scenarios that lack a large amount of high-quality human-written data. Furthermore, **to improve training efficiency and reduce application costs for multiple ad tasks**, in the post-pretraining phase, we propose a novel multi-task prompt learning mechanism by introducing task prompts and multiple losses to better fuse multiple tasks into one single model without loss of performance.

The main contributions are summarized as follows:

- We present a unified ad generation framework with multi-task prompt learning, named PLATO-Ad, to effectively tackle the ad generation in low-resource scenarios and training efficiency for multiple resource-rich ad tasks. To our knowledge, this work is the first to study low-resource ad generation with a prompting mechanism in industrial scenarios.

- We propose a novel three-phase transfer learning mechanism to address ad generation in low-resource settings by transferring from generic domain text generation to resource-rich ad domain text generation, and finally to low-resource ad text generation.

- Furthermore, we devise a novel multi-task prompt learning mechanism by introducing task prompts and multiple training objectives to efficiently fuse multiple resource-rich ad tasks into a single lightweight model without loss of performance.

- The offline and online experiment results show PLATO-Ad significantly outperforms the state-of-the-art models and can generate high-quality ads in low-resource settings. In the A/B test, the advertisements generated by PLATO-Ad would bring about 3.5% CTR improvement on search ad descriptions and 10.4% CTR improvement on feed ad titles.

2 Related Work

Previous works on text ad generation rely on designing pre-defined templates to construct readable ad sentences (Bartz et al., 2008; Fujita et al., 2010; Thomaidou et al., 2013). However, these template-based ads are generally rigid and not diverse enough leading to unappealing to users. Afterward, Hughes et al. (2019) presents data-driven methods to learn to write text ads from existing examples. The model employs LSTMs and attention layers to encode product landing pages and decode text ads. Moreover, REINFORCE with baseline (Rennie et al., 2017) is leveraged to generate attractive ads that potentially have a larger click rate. This method achieves a certain degree of success in the automatic generation of text ads. However, over 15% text ads it generates are labeled as non-sense, broken, or bad, which fails to meet the high-quality standard for production (Hughes et al., 2019).

Recently, some work utilize pretraining-based natural language generation methods to generate content-rich, diverse and attractive ads based on large-scale corpus (Wang et al., 2021; Zhang et al., 2021a,b; Wei et al., 2022). However, these methods require large-scale high-quality training corpus and low training efficiency when applied to multiple ad generation scenarios. In this paper, we propose a unified text ad generation framework with multi-task prompt learning to tackle ad generation problems in low-resource settings and improve training efficiency for multiple resource-rich ad tasks.

3 Methodology

3.1 Problem Settings

Different ad generation tasks, such as ad description generation, ad title generation, selling point generation and tips generation (Li et al., 2019) for

Task	Input Text	Auxiliary Attributes	Output Text	Low-Resource	Training Phase
Dialog Gen.	Context	-	Response	✗	Pretrain
Ad Desc. Gen.	Ad Title	Product Landing Page (Text)	Ad Desc.	✗	Post-Pretrain
Ad Title Gen.	Product Entity	Product Landing Page	Ad Title	✗	Post-Pretrain
Sel. Point Gen.	Ad Title	Product Attributes	Product Selling Points	✗	Post-Pretrain
Comment Gen.	Product Desc.	Sentiment Polarity	Comment	✗	Post-Pretrain
QA	Question	Keyword	Answer	✗	Post-Pretrain
Commonsense-enriched Ad Desc. Gen.	Ad Title	Product Landing Page & Keyword	Commonsense-rich Ad Desc.	✓	Prompting
Tips Gen.	Ad Title	Focus Point & Sentiment Polarity	Tips	✓	Prompting

Table 1: Description of Different Ad Generation Tasks. Here, Gen. denotes generation.

recommendation, usually contain different inputs and outputs. To simplify description, we generalize elements of all ad generation tasks into the following fields: input text, auxiliary attributes and output text. Table 1 shows inputs and outputs of different ad generation tasks.

Formally, we refer to input text as $X = (x_1, x_2, ..., x_n)$, auxiliary attributes $A^j = (a^j_1, a^j_2, ..., a^j_k), j = 1, 2, ..l$ and output text as $Y = (y_1, y_2, ..., y_m)$. Here, A^j represents a word sequence of the j-th ad attribute, n, k, m denotes the sequence length of input text, the j-th ad attribute and output text respectively. l denotes the number of auxiliary attributes. $x_i, a^j_i, y_i \in \mathcal{V}$ denotes a word token. $\mathcal{V}$ denotes the vocabulary. The ad generation process can be defined as a sequence-to-sequence generation task as follow:

$$\begin{aligned} Y &\sim p(Y|X, A) \\ &= p(y_1, y_2, ..., y_m | x_1, x_2, ...x_n; A^1, A^2, ..., A^l). \end{aligned} \tag{1}$$

The goal of PLATO-Ad is to learn this function $p(Y|X, A)$.

3.2 Model Framework

3.2.1 Overview

PLATO-Ad is a transformer-based pretraining model with a three-phase transfer learning mechanism, as shown in Fig 2. Concretely, we first pretrain PLATO-Ad on a large-scale open-domain dialog corpus following the setting of PLATO-2 (Bao et al., 2021), a Chinese dialog-oriented pretraining model. The purpose of this step is to enable PLATO-Ad to generate fluent natural sentences and model correlation from input (context) to output(response). Then, in the second phase, we continue to pretrain PLTAO-Ad on the datasets of multiple resource-rich ad generation tasks with relatively massive training data available and open-domain QA datasets, which makes PLATO-Ad learn generating ad-domain and commonsense-enriched texts. This step enables that PLATO-Ad can transfer learning from generic domain text generation to ad domain text generation. Moreover, to apply PLATO-Ad to multiple ad generation scenarios more efficiently, we design a multi-task prompt learning mechanism by introducing task prompts and multiple losses to efficiently fuse multiple tasks into a single lightweight model without loss of performance. Finally, in the third phase, we use task prompts to transfer well-trained PLATO-Ad to low-resource text ad generation tasks. To simplify the description, we call the first phase as **Generic-Domain Pretraining**, the second phase as **Ad-Domain Post-Pretraining**, and the third phase as **Low-Resource Prompting**. We elaborate the details of all phases in the following.

3.2.2 Generic-Domain Pretraining

We first pretrain PLATO-Ad on generic-domain text following the settings of PLATO-2 (Bao et al., 2021). Specifically, the pretraining dataset is collected from public social medias, which contains 1.2B (context, response) samples. Unlike the original PLATO-2 with multi-step losses, we train PLATO-Ad only with negative log-likelihood (NLL) loss for fluent natural text generation.

3.2.3 Ad-Domain Post-Pretraining

Then, we post-pretrain PLATO-Ad on datasets of multiple resource-rich ad generation tasks and high-quality open-domain QA datasets to equip PLATO-Ad with the ability to generate ad-domain and commonsense-enriched text. In particular, the post-pretraining datasets contain four real-world resource-rich ad generation tasks, such as ad description generation, ad title generation, selling point generation and comment generation. The construction and data preprocessing of these datasets can be found in Section A.1.

Meanwhile, to more efficiently train lightweight PLATO-Ad for fusing multiple ad tasks, we design a multi-task prompt learning mechanism by introducing task prompts and multiple training ob-

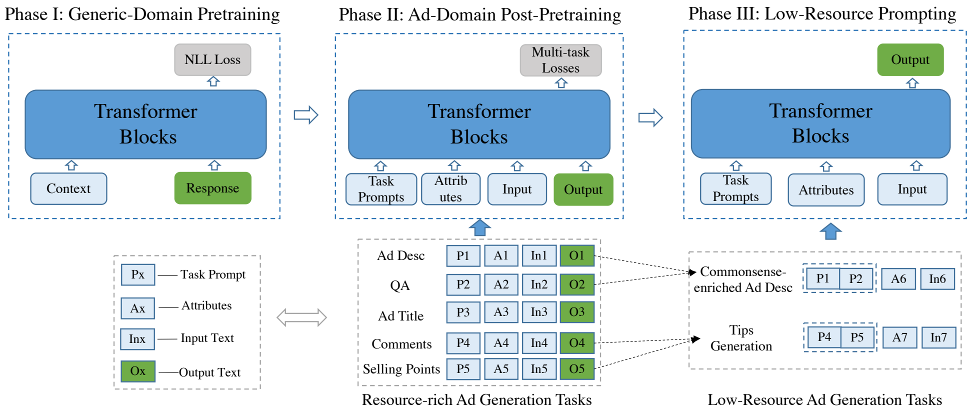

Figure 2: The architecture of our PLATO-Ad framework.

jectives. The model infrastructure is shown in Fig 2, which consists of input representation, transformer blocks and multi-task training objectives.

Input Representation. For each token of input text, auxiliary attributes and output text, its representation is the sum of the token embedding, the type embedding and the position embedding. Type embeddings are employed to differentiate inputs from different parts and position embeddings are set according to the token position in each sentence. Two special tokens [BOS] and [EOS] are inserted at the beginning and end of sentences for separation. Meanwhile, to effectively integrate multiple different tasks into one model, we add task prompts as the part of inputs following prompt tuning (Lester et al., 2021). Here, different tasks correspond to different task prompts and we randomly initialize the representation vectors of task prompts.

Multiple Training Objectives. To better train PLATO-Ad for meeting different ad tasks, we design multiple training losses to optimize the repetition problem, relevance of input and output, and controllability in these tasks.

a) Unlikelihood Loss for Repetition. To alleviate the repetition problem in text generation, especially long text generation for ad desc task, we adopt an unlikelihood loss following (Welleck et al., 2020). Formally, the unlikelihood loss L_U can be defined as follow.

$$L_U = -\sum_{y_t \in C^t} \log(1 - p_\theta(y_t|y < t, X, A)) \quad (2)$$

Where y_t denotes t-th output token, C^t refers to a set of negative candidate tokens (i.e. repetitive tokens), X and A denotes input text and attributes.

b) Relevance Loss. To improve relevance between input and output, we design a discriminative relevance loss L_R to estimate the relevance between them, which can be defined as follow.

$$\begin{aligned} L_R = &- \log p(l = 1|X, Y^+) \\ &- \log p(l = -1|X, Y^-) \end{aligned} \quad (3)$$

where $p(l = 1)$ denote positive samples and $p(l = -1)$ denote negative samples. Positive samples are from golden <X, Y> pairs and negative samples are obtained by random sampling <X, Y> pairs.

c) Keyword Loss for Controllability. Some ad generation tasks, for example, selling point generation and tip generation, expect the generated response to contain auxiliary attributes. Therefore, we design keyword loss L_C following (Kumar et al., 2021) to improve the token-level controllability of PLATO-Ad.

$$L_C = \min_{i=1}^{m}(-\log p(y_i = K|X, A)) \quad (4)$$

where K denotes the keyword controlled to generate. Finally, we train PLATO-Ad by summing multiple losses.

$$L = L_N + \lambda_u L_U + \lambda_r L_R + \lambda_c L_C \quad (5)$$

where L_N denotes negative log-likelihood (NLL) loss, $\lambda_u, \lambda_r, \lambda_c$ are model hyperparameters.

3.2.4 Low-Resource Prompting

After post-pretraining, we use prompting methods to apply well-trained PLATO-Ad to low-resource ad generation scenarios. Specifically, we select two low-resource ad generation, commonsense-enriched ad description generation (Zhang et al., 2021a) and tips generation (Li et al., 2019), both of which lack large-scale high-quality training data. Then we achieve task-level transfer by using the combination of task prompts from resource-rich ad tasks. In particular, the task prompt of commonsense-enriched ad description generation is a combination of task prompts of ad description generation and open-domain QA tasks, as shown in Fig 2. In this way, we will transfer common sense from open-domain QA to ad descriptions to generate commonsense-rich ad descriptions. Similarly, the task prompt of tips generation is a combination of task prompts of selling point generation and comment generation, which results in generating informative reviews for product selling points.

4 Experiments

4.1 Datasets

To our knowledge, there are no publicly available large-scale high-quality ad-domain datasets and we collect post-pretraining datasets from a leading advertising platform. Table 2 shows the statistics of these datasets. More details about dataset construction and data preprocessing can be found in Appendix A.1.

	Train	Dev	Test
Ad Desc. Gen.	10,800,000	50,000	10,000
Ad Title Gen.	4,288,167	50,000	10,000
Sel. Point Gen.	798,686	50,000	10,000
Comment Gen.	5,777,279	50,000	10,000
QA	17,859,294	50,000	10,000

Table 2: Dataset statistics.

4.2 Baselines

We select the following baselines to evaluate the effectiveness of our model.

CHASE (Zhang et al., 2021a): It is the online state-of-the-art model deployed on the commercial ad systems. We follow the same model settings.

PLATO-2-FT: It directly uses the PLATO-2 (Bao et al., 2021) model to finetune on multiple ad datasets. We use the released parameters [2].

PLATO-Ad: It is our proposed model in this paper with a multi-task prompting mechanism and three-phase transfer learning. These three models have the same magnitude of parameters (about 90M), which ensures a fair comparison.

[2] https://github.com/PaddlePaddle/Knover/tree/luge-dialogue

4.3 Evaluation Metrics

We use Perplexity (PPL) (Brown et al., 1992) and Pairwise-BLEU (Shen et al., 2019) to automatically measure the model quality and diversity of generation results. The more diverse the hypothesis set is, the lower the Pairwise-BLEU is. Meanwhile, we conduct a manual evaluation on 200 random samples from our test dataset. Three participants were recruited to measure the quality of the result generated by each baseline from three perspectives, including Readability (Read.), Relevance (Rele.), Information (Info.). Each perspective is measured by a 3-point Likert question where 0 is bad, 1 is neutral and 2 is good. The Overall (Over.) score is the average value of the above three scores. The detailed evaluation metrics can be found in appendix A.4.

4.4 Experimental Results on Resource-rich Ad Generation Tasks

We conduct a set of experiments to evaluate the effectiveness of PLATO-Ad on resource-rich ad generation tasks. As shown in Table 3, PLATO-Ad significantly outperforms the state-of-the-art CHASE and PLATO-2-FT in terms of all the metrics. It demonstrates that PLATO-Ad can generate more fluent (lower PPL) and more diverse (lower Pairwise-BLEU) ads in comparison with baselines. Meanwhile, PLATO-Ad removes multi-task losses causes more high PPL and Pairwise-BLEU scores, **indicating the effectiveness of multi-task losses**.

In addition, we conduct experiments to verify the efficiency of the multi-task prompt learning mechanism. From Table 4, we find that single lightweight PLATO-Ad obtains better performance than multiple separate models trained on different tasks. This demonstrates **the effectiveness of fusing multiple tasks into a single model with shared parameters via multi-task prompt learning mechanism**.

Models	PPL↓	Pairwise-BLEU↓
CHASE	8.30	47.90
PLATO-2-FT	3.29	41.93
PLATO-Ad	**2.21**	**38.46**
PLATO-Ad w/o multi losses	2.36	38.67

Table 3: Results on the test set of ad description generation. PLATO-Ad w/o multi losses denote that PLATO-Ad removes multi-task losses. Bold scores are the best.

Datasets	PLATO-Ad w/o Task Prompt	PLATO-Ad
Ad Desc Gen.	2.50	**2.21**
Ad Title Gen.	2.81	**2.65**
Sel. Point Gen.	7.31	**6.70**
Comment Gen.	**6.40**	6.42
QA	3.71	**3.31**

Table 4: PPL on test sets of multiple ad generation tasks. PLATO-Ad w/o Task Prompt represents that PLATO-Ad removes the multi-task prompting mechanism and separately trains on each corresponding ad generation task.

4.5 Transfer Learning on Low-Resource Ad Generation Tasks

We investigate the effectiveness of PLATO-Ad on two low-resource [3] ad generation tasks, commonsense-enriched ad description generation and tips generation. We compare the manual evaluation results of PLATO-2-FT (separately finetuning on these two datasets) and PLATO-Ad Prompting on 200 random samples of our test datasets. From Table 5, we can see that PLATO-Ad surpasses PLATO-2-FT on all manual metrics (especially Info.) for these two low-resource ad generation tasks, indicating that PLATO-Ad can generate more relevant and commonsense-enriched ad descriptions and tips. This also verifies **the effectiveness of three-phase transfer learning**.

Figure 3 shows the examples generated by PLATO-2-FT and PLATO-Ad. We can see that PLATO-Ad generates more commonsense-rich ad descriptions and more informative tips. *We analyze that this is because PLATO-Ad realizes the transfer learning of multiple tasks in the post-pretraining phase, such as transferring ad description and QA tasks to the commonsense-enrich ad description generation task and transferring comment generation and selling point tasks to the tip generation task.* Meanwhile, we also find that PLATO-Ad can generate less repetitive and more controllable ads in comparison with PLATO-2-FT.

4.6 Online A/B Test

We also conduct online A/B tests in a leading advertising platform to show the superiority of PLATO-Ad. We used 5% real-world web traffics in the advertising platform to conduct the A/B test. We had already gotten the permission from advertiser

[3] We manually evaluated only about 9.3% of ad descriptions dataset (93/1000) are commonsense-enriched ad descriptions and 8.6% of comment generation dataset (86/1000) are suitable for tips generation.

Tasks	Models	Read.	Rele.	Info.	Over.
Com.A.	PLATO-2-FT	1.70	1.25	0.68	1.21
	PLATO-Ad	**1.91**	**1.89**	**1.98**	**1.93**
Tip.	PLATO-2-FT	1.96	1.68	1.62	1.75
	PLATO-Ad	**1.96**	**1.86**	**1.93**	**1.92**

Table 5: Manual evaluation results on Commonsense-enriched Ad Description Generation (Com.A.) and Tip Generation (Tip.) tasks. The Kappa value is above 0.4, showing modest agreement among evaluators.

Commonsense-enriched Ad Desc	**Input:** 孩子注意力不集中怎么办？(What if children do not pay attention?) **Attributes:** 雨林教育从事3-14岁少年注意力培养和逻辑培训(Rainforest Education Corporation focuses on the attention and logic training of children aged 3-14.)
PLATO-2-FT	雨林教育,专注3-14岁少儿的注意力培养,逻辑思维训练,帮助孩子提升注意力, 逻辑思维能力训练。（Rainforest Education Corporation, focuses on the attention training and logical thinking ability training of children aged 3-14, and helps children improve their attention, and logical thinking ability training.）
PLATO-Ad	孩子注意力不集中,首先要培养孩子良好的生活习惯,让孩子养成早睡早起，多参加一些户外活动，培养孩子的专注力。雨林教育专业培训3-14岁孩子注意力和逻辑，值得信赖。(If the child is inattentive, first of all, it is necessary to cultivate the child's good living habits, let the child develop to go to bed early and get up early, participate in some outdoor activities, and cultivate the child's concentration. Rainforest Education is a trustworthy company that professionally trains the attention and logic for children aged 3-14.)

Tip Gen.	**Input:**男士纯棉休闲裤, 限时打折优惠(Men's cotton slacks with discount for a limited time) **Attributes:** Focus Points (价格#质地 Price # Material); Sentiment Polarity (正面Positive)
PLATO-2-FT	挺不错的，应该买的人挺多的。（It's pretty good, and there should be a lot of people buying it.）
PLATO-Ad	这款男士休闲裤，纯棉的，穿着很舒服。同时质量也不错, 男士穿上帅气又减龄, 价格也很优惠。(This pure cotton men's casual pant is very comfortable to wear.Meanwhile, the quality is also good, and men wear it to look handsome and age-reducing, and the price is also very favorable.)

Figure 3: Example ads generated by PLATO-Ad and PLATO-2-FT. PLATO-2-FT could generate repetitive ads (green color). Compare with PLATO-2-FT, PLATO-Ad can generate commonsense-rich ads (red color) and more controllable ads (blue color).

to use PLATO-Ad to generate advertisement descriptions/titles for this test. This online A/B test lasts for one week. On each day there were about 1 million page views (with ad shows) for the testing. We use Click-Through Rate (CTR) (Richardson et al., 2007) and Conversion Rate (CVR) (Lee et al., 2012) compared with CHASE to show the improvement of PLATO-Ad. Except for the displayed ad descriptions/titles, we keep other settings the same. Table 6 demonstrates that PLATO-Ad can bring more significant CTR and CVR improvement compared with the state-of-the-art CHASE. The details about deployed workflow can be found in Appendix A.5.

Tasks	△CTR	△CVR
Search Ad Description	+3.5%	+1.7%
Feed Ad Title	+10.4%	+4.1%
Feed Selling Point	+7.6%	+2.3%

Table 6: Online A/B Testings on different ad generation tasks. $\triangle\text{CTR/CVR} = \frac{\text{CTR/CVR of PLATO-Ad - CTR/CVR of CHASE}}{\text{CTR/CVR of CHASE}}$

4.7 Conclusion

In this paper, we propose a novel unified text ad generation framework with multi-task prompt learning, called PLATO-Ad, to tackle universal commercial ad generation tasks. Experiments show that PLATO-Ad can generate commonsense-rich and relevant ads in low-resource scenarios via **a three-phase transfer learning mechanism** and improve training efficiency for multiple resource-rich ad tasks by using **a multi-task prompt learning mechanism** to fuse multiple tasks into a single lightweight model without loss of performance. In the future, we will extend the idea of PLATO-Ad to more real-world text generation tasks.

Ethics Statement

We make sure that we have the copyright to use all datasets to train and deploy. Meanwhile, these datasets do not contain any user's private information. In manual evaluation, we ensure that all annotators were treated fairly. This includes but is not limited to, compensating them fairly, ensuring that they were able to give informed consent, and ensuring that they were voluntary participants who were aware of any risks of harm associated with their participation. During A/B testing and system deployment, all generated advertisements must be approved by the advertiser before using. Before online deployment, we conduct a post-processing procedure for all generated advertisements, including the basic correlation filtering (quality control) and business risk control system to strictly control the exposure risk of the displayed advertisements. Meanwhile, for badcases or harmful contents that are found or fed back from customers when displayed online, we also have an online blacklist procedure to filter them in real time.

References

Siqi Bao, Huang He, Fan Wang, Hua Wu, Haifeng Wang, Wenquan Wu, Zhen Guo, Zhibin Liu, and Xinchao Xu. 2021. Plato-2: Towards building an open-domain chatbot via curriculum learning. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 2513–2525.

Kevin Bartz, Cory Barr, and Adil Aijaz. 2008. Natural language generation for sponsored-search advertisements. In *Proceedings of the 9th ACM Conference on Electronic Commerce*, pages 1–9.

Peter F Brown, Stephen A Della Pietra, Vincent J Della Pietra, Jennifer C Lai, and Robert L Mercer. 1992. An estimate of an upper bound for the entropy of english. *Computational Linguistics*, 18(1):31–40.

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.

Atsushi Fujita, Katsuhiro Ikushima, Satoshi Sato, Ryo Kamite, Ko Ishiyama, and Osamu Tamachi. 2010. Automatic generation of listing ads by reusing promotional texts. In *Proceedings of the 12th International Conference on Electronic Commerce: Roadmap for the Future of Electronic Business*, pages 179–188.

Xiaojie Guo, Shugen Wang, Hanqing Zhao, Shiliang Diao, Jiajia Chen, Zhuoye Ding, Zhen He, Yun Xiao, Bo Long, Han Yu, et al. 2021. Intelligent online selling point extraction for e-commerce recommendation. *arXiv preprint arXiv:2112.10613*.

J Weston Hughes, Keng-hao Chang, and Ruofei Zhang. 2019. Generating better search engine text advertisements with deep reinforcement learning. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2269–2277.

Bernard J Jansen and Marc Resnick. 2005. Examining searcher perceptions of and interactions with sponsored results. In *Workshop on Sponsored Search Auctions*.

Shachi H. Kumar, Hsuan Su, Ramesh Manuvinakurike, Saurav Sahay, and Lama Nachman. 2021. Controllable response generation for assistive use-cases. *CoRR*, abs/2112.02246.

Kuang-chih Lee, Burkay Orten, Ali Dasdan, and Wentong Li. 2012. Estimating conversion rate in display advertising from past erformance data. In *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 768–776.

Brian Lester, Rami Al-Rfou, and Noah Constant. 2021. The power of scale for parameter-efficient prompt tuning. *arXiv preprint arXiv:2104.08691*.

Piji Li, Zihao Wang, Lidong Bing, and Wai Lam. 2019. Persona-aware tips generation? In *The World Wide Web Conference*, pages 1006–1016.

Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. 2021. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *arXiv preprint arXiv:2107.13586*.

Steven J Rennie, Etienne Marcheret, Youssef Mroueh, Jerret Ross, and Vaibhava Goel. 2017. Self-critical sequence training for image captioning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7008–7024.

Matthew Richardson, Ewa Dominowska, and Robert Ragno. 2007. Predicting clicks: estimating the click-through rate for new ads. In *Proceedings of the 16th international conference on World Wide Web*, pages 521–530.

Timo Schick and Hinrich Schütze. 2020. Few-shot text generation with pattern-exploiting training. *arXiv preprint arXiv:2012.11926*.

Tianxiao Shen, Myle Ott, Michael Auli, and Marc'Aurelio Ranzato. 2019. Mixture models for diverse machine translation: Tricks of the trade. In *International conference on machine learning*, pages 5719–5728. PMLR.

Stamatina Thomaidou, Ismini Lourentzou, Panagiotis Katsivelis-Perakis, and Michalis Vazirgiannis. 2013. Automated snippet generation for online advertising. In *Proceedings of the 22nd ACM international conference on Information & Knowledge Management*, pages 1841–1844.

Xiting Wang, Xinwei Gu, Jie Cao, Zihua Zhao, Yulan Yan, Bhuvan Middha, and Xing Xie. 2021. Reinforcing pretrained models for generating attractive text advertisements. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pages 3697–3707.

Penghui Wei, Xuanhua Yang, Shaoguo Liu, Liang Wang, and Bo Zheng. 2022. Creater: Ctr-driven advertising text generation with controlled pre-training and contrastive fine-tuning. *arXiv preprint arXiv:2205.08943*.

Sean Welleck, Ilia Kulikov, Stephen Roller, Emily Dinan, Kyunghyun Cho, and Jason Weston. 2020. Neural text generation with unlikelihood training. In *ICLR*.

Chao Zhang, Jingbo Zhou, Xiaoling Zang, Qing Xu, Liang Yin, Xiang He, Lin Liu, Haoyi Xiong, and Dejing Dou. 2021a. Chase: Commonsense-enriched advertising on search engine with explicit knowledge. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*, pages 4352–4361.

Xueying Zhang, Yanyan Zou, Hainan Zhang, Jing Zhou, Shiliang Diao, Jiajia Chen, Zhuoye Ding, Zhen He, Xueqi He, Yun Xiao, et al. 2021b. Automatic product copywriting for e-commerce. *arXiv preprint arXiv:2112.11915*.

A Appendix

A.1 Post-Pretraining Datasets

We collect all post-pretraining datasets from a Chinese leading advertising platform. Specifically,

- **Datasets for Ad Description Generation** We collect tens of millions of <search ad title, search ad description> pairs manually written by advertisers from the leading advertising platform as inputs and outputs of ad description task. For each pair, we match corresponding product landing page with it as auxiliary attributes.

- **Datasets for Ad Title Generation** For the ad title generation task, we use product entities as input text, corresponding product landing pages as auxiliary attributes and feed ad titles human-written by advertisers from a leading advertising platform as output text.

- **Datasets for Selling Point Generation** We construct selling point generation datasets by extracting snippets related to product attributes in the above ad descriptions and ad title dataset as product selling points via public universal information extraction tools [4].

- **Datasets for Comment Generation** We use hard prompt methods (Brown et al., 2020; Schick and Schütze, 2020) to obtain datasets for ad comment generation. Specifically, we use product descriptions and human-written prompt templates (e.g., what do you think about this product) as the context of the PLATO-Ad model in the pretraining phase to generate generic-domain comments (responses). We use publicly-available text sentiment analysis tools [5] to filter out negative text.

- **Datasets for Open-domain QA** We collect open-domain QA datasets from the Chinese community-based question-answering websites (just like Quora [6]) in the leading advertising platform. Here the question and answer of each QA item are treated as input text and output text respectively. Meanwhile, we use publicly available named entity recognition tools [7] to extract the entities of questions as auxiliary keywords (Attributes).

For all collected data, we will conduct a data-preprocess procedure, including filtering out low-quality ones via a set of heuristic rules and publicly-available tools (e.g., text error detection [8], sentence length constraint, repeat word constraint and harmful/abusive word vocabulary).

[4] https://github.com/PaddlePaddle/PaddleNLP/tree/develop/model_zoo/uie

[5] https://ai.baidu.com/tech/nlp_apply/sentiment_classify

[6] https://www.quora.com/

[7] https://ai.baidu.com/tech/nlp_basic/entity_analysis

[8] https://ai.baidu.com/tech/nlp_apply/text_corrector

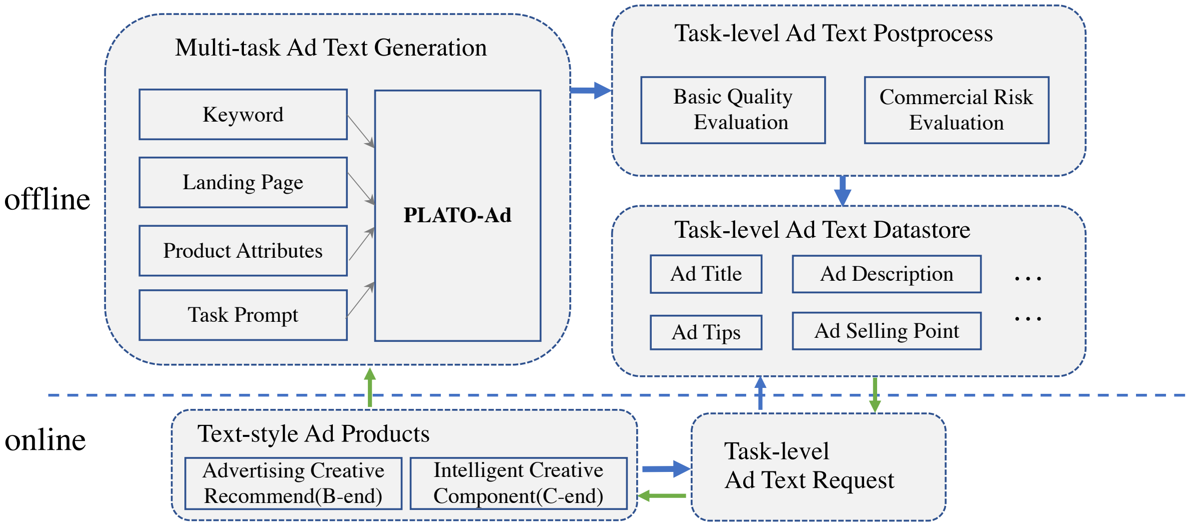

Figure 4: Deployment workflow of PLATO-Ad.

A.2 Model Settings

Our PLATO-Ad is a 93M parameter model with 12 transformer blocks and 12 attention heads, with the embedding dimension of 768. The model structure follows the setting of PLATO-2 (Bao et al., 2021). According to the number of post-pretraining tasks, we set five kinds of task prompts and the embedding size of task prompt is set to be 768. we randomly initialize the representation vectors of task prompts. We train PLATO-Ad in post-pretraining phase with the batch size of 65,536 on 8 A100 GPUs. During decoding, we adopt a topk-sampling decoding strategy with k=5. $\lambda_u, \lambda_r, \lambda_c$ are set to be 1.

A.3 Automatic Evaluation

Perplexity (Brown et al., 1992) is an evaluation metric to measure the model capacity for language modeling which is the normalized inverse probability of the dataset. BLEU is to use n-gram word matching to measure the similarity between golden truth and generated text. In this paper, we use Pairwise-BLEU (Shen et al., 2019) to measure the diversity of generation results. Specifically, Pairwise-BLEU measures similarity among the hypotheses (multiple generated candidate results). The more diverse the hypothesis set is, the lower the Pairwise-BLEU is.

A.4 Manual Evaluation

We conduct a manual evaluation on 200 random samples from our test dataset. Three participants were recruited to measure the quality of the result generated by each baseline from three perspectives. Each perspective is measured by a 3-point Likert question where 0 is bad, 1 is neutral and 2 is good.

- **Readability (Read.)**: measures how the generated ad text is smooth and grammatically corrects.
- **Relevance (Rele.)**: measure whether output text is relative with input text and whether the generated result is consistent with auxiliary attributes.
- **Information (Info.)**: measures how informative/knowledgable generated ad text is.

Overall (Over.): measures the overall quality of generated ad texts, which is calculated by the average of the above three scores.

A.5 Deployment Workflow

Figure 4 shows the deployment workflow of PLATO-Ad. We first use PLATO-Ad to generate ads offline. Then, before online deployment, we conduct a task-level ad text post-processing procedure for all generated advertisements, including the basic quality evaluation (quality control) and commercial risk control evaluation to strictly control the suitability of displayed advertisements. Finally we store the filtered ads for online retrieval. Overall, the suitability rate of ads generated by PLATO-Ad is around 91% to 96%, which means that PLATO-Ad can be deployed online in the industry.

Meanwhile, for badcases or harmful contents that are found or fed back from customers when displayed online, we also have an online blacklist procedure to filter out them in real time.