Personalized Preference Optimization for Text-to-Image Generation using Large Language Models

Anonymous ACL submission

Abstract

Preference optimization is a crucial aspect of generative models, ensuring that the generated content aligns with users' preferences. While previous research has focused on optimizing for average preferences, text-to-image tasks require a personalized approach due to the diversity of individual preferences. In this study, we propose a two-stage framework for personalized preference optimization in text-to-image generation. The first stage, personalized image aesthetic assessment (PIAA), learns user preferences from a small amount of user image rating data. The second stage, prompt optimization, optimizes the text-to-image model's prompt to generate images that receive high scores from the learned preference model. We employ Large Language Models (LLMs) for the prompt optimization process. Through extensive experimentation with various configurations in the PIAA and prompt optimization stages, we demonstrate that our approach can generate novel images that align with individual user preferences, even with limited user data. Our research lays the foundation for future work on personalized content generation.

1 Introduction

As generative models continue to advance rapidly, the demand for personalized content creation has surged. Existing methods incorporate human preferences to create broadly appealing images but mainly cater to generalized aesthetics (Kirstain et al., 2023; Xu et al., 2023; Hao et al., 2023; Wu et al., 2023). The inherently emotional and artistic nature of text-to-image (T2I) generation highlights the need for personalization in creative outputs, as personal aesthetic preferences are increasingly important.

In recent developments, 'personalization' in text-to-image generation has often referred to generating images that reflect the user's input content (Ruiz et al., 2023; Shi et al., 2024). However, in this paper, we define personalization as understanding and generating images that match the user's unique preferences. To achieve this, we train a preference score network with real-world data to simulate user preferences, and employ a large language model (LLM) as an optimizer to refine prompts for the T2I model.

Overall, our novel T2I generation approach uses minimal user-provided image ratings to generate outputs that meets personal tastes. By focusing on personalized aesthetic modeling, we aim to bridge the gap between general image generation and personalized artistic expression, enhancing user satisfaction and advancing T2I generation by addressing diverse aesthetic inclinations. Experiments and ablation studies demonstrate that our framework effectively enhances user preference scores, showing the potential of LLMs in solving personalization optimization through advanced learning techniques.

2 Related Works

The advancement of text-to-image (T2I) generation models has emphasized the importance of preference optimization to enhance user satisfaction. Previous approaches such as ImageReward (Xu et al., 2023) and Pick-a-Pic (Kirstain et al., 2023) have incorporated human preferences to create appealing images, primarily focusing on generalized user preferences. For instance, Pick-a-Pic collected a large dataset of user preferences, enabling the training of the PickScore function, which predicts human preferences with remarkable accuracy. However, these methods often cater to a broad audience rather than individual preferences. Similarly, Pavlichenko and Ustalov, 2023 explored human-in-the-loop methods for optimizing prompts but did not focus on individual personalization.

LLMs as optimizers offer a promising approach to optimizing user-specific aesthetic preferences in T2I generation. Foundational concepts intro-

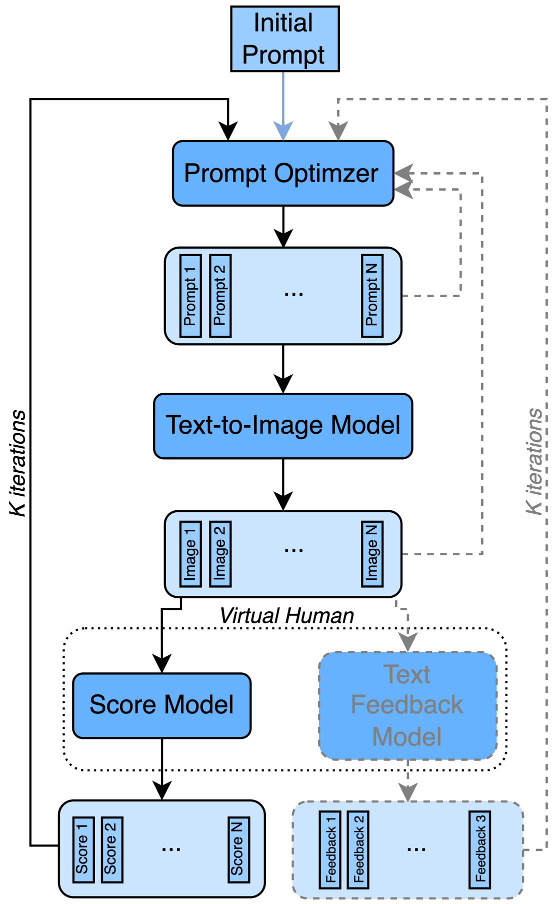

Figure 1: Overall framework. The gray dotted line indicates ablation choices.

duced by Brown et al., 2020 and Gao et al., 2020 highlighted the potential of pre-trained language models to adapt to new tasks with minimal data. Previous works (Pryzant et al., 2023; Yang et al., 2023; Yuksekgonul et al., 2024) demonstrate that an LLM can be used as a general optimizer. Use of LLMs in dynamically refining prompts to ensure personalized outputs. Additionally, works such as Yang et al., 2022 and Liu et al., 2023 provided insights into integrating detailed user feedback for better personalization. This integration of LLMs as optimizers and real-world user data for training a preference score network has shown effectiveness in enhancing personalized aesthetic modeling in T2I generation.

3 Method

We focus on personalized image generation, which we define as the optimization problem to generate an image that maximizes user's preference score. To solve this optimization problem, we first need a score function that provides a consistent and personalized score for images based on real-world hu-

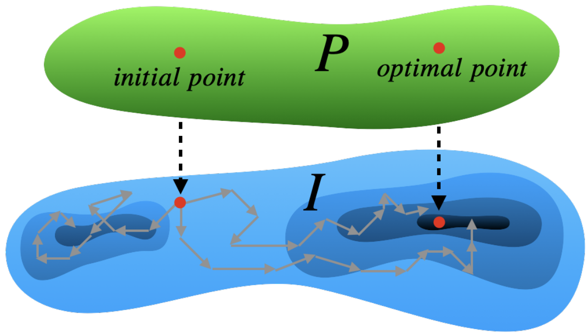

Figure 2: Relations between the prompt (P), image (I), and preference score are such that each prompt must pass through the intermediate image stage to determine the corresponding score.

man data. We simulate each user's score function by training a small score model using the Personalized Image Aesthetics Assessment (PIAA)(Ren et al., 2017; Yang et al., 2022) method and the Simulacra Aesthetic Captions (SAC) dataset (Pressman et al., 2022). Second, we need an optimizer that maximizes the user's score function. For this, we use an LLM as an optimizer. To connect the LLM's output text to the user's score function, which only accepts images as input, we use a text-to-image model. Each component of our method is explained in detail below.

3.1 PIAA

The Personalized Image Aesthetic Assessment (PIAA) (Ren et al., 2017; Yang et al., 2022) approach aims to model individual aesthetic preferences. PIAA databases typically contain images annotated by multiple subjects with various objective and subjective attributes, as well as desensitized subject information such as personality traits and experience levels. This comprehensive annotation allows for a nuanced understanding of aesthetic preferences. PIAA models often incorporate subject information as a prior, enhancing the prediction of personalized aesthetic preferences by leveraging detailed user feedback. The inclusion of such detailed information has been shown to significantly improve the performance of PIAA models compared to those without it.

3.2 Prompt Optimization Method

As illustrated in Figure 1, our framework aims to maximize the virtual user's score by iteratively modifying the prompts fed into the text-to-image model. This necessitates the prompt optimizer to determine the mapping function between prompts and the user's preference score. However, there are two significant challenges. First, as shown in Figure 2, the prompt and score are not directly connected, as there is an intermediate image space between them. Second, the personalized score space can be sparse due to the nature of personal preferences. These characteristics make it difficult for the prompt optimizer to discern the overall structure of the mapping function. To address this, we come up with various methods.

Multi-prompt Optimization To enhance the prompt optimizer's understanding of the preference score function, we utilize N prompts. Specifically, the initial prompt is fed into the prompt optimizer, which then generates N diverse prompts. In this initial process, the prompt optimizer aims to create diverse prompts to approximate the landscape of the score function. Subsequently, the prompt optimizer works to maximize the scores for all N prompts.

Comprehensive Feedback Integration To enhance the prompt optimizer's performance, we propose incorporating additional information beyond the previously generated prompt scores. First, we introduce a text feedback model, implemented by an LLM, which is conditioned on (image, score) pairs to capture user preferences and provide directional guidance for score improvement. Second, we include the image output from the text-to-image model to better understand the overall image structure. Finally, we utilize the prompt optimizer's history of generated prompts, paired with their corresponding scores and sorted for easy reference. The impact of these additional information sources is evaluated through ablation studies.

4 Experiments

If our framework can optimize personal preferences, it should also work for general preferences. To validate the effectiveness of our framework, we first conduct optimization for general preferences, where the virtual human's preferences are non-specific and general. Following this, we perform optimization for personal preferences.

4.1 Dataset

For the general preference optimization, we utilized prompts from the DiffusionDB dataset (Wang et al., 2022), which contains diverse text-to-image prompts and their corresponding AI-generated images, providing a robust foundation for training and evaluating generative models. For personal preference optimization, the Simulacra Aesthetic Captions (SAC) dataset (Pressman et al., 2022). The SAC dataset includes over 238,000 synthetic images generated from user-submitted prompts, rated on aesthetic value, which helps in refining and validating personalized aesthetic models by leveraging a large volume of user feedback. These datasets collectively enable comprehensive evaluation and optimization of both general and personal preferences in text-to-image generation models.

4.2 Evaluation Metric

For evaluating our personalized image generation framework, we employ several metrics that assess the effectiveness of both general and personal preference optimization. For personal preference optimization, we consider the maximum score (*max*) as the highest preference score assigned by the user to any generated image. We also calculate the improvement from the initial to the maximum score (*max - init*) and the number of iterations required to achieve the maximum score (*K**), reflecting the efficiency of the optimization process. Additionally, we measure the similarity between the highest-scoring generated image and the most similar previously high-rated image (*max sim*), given by $\max_{i \in I_{\text{high}}} \text{sim}(i, i^*)$. Lastly, we evaluate the score assigned to the image most similar to the highest-rated generated image among all previously evaluated images (*score of most sim*), defined as $\text{score}(i_{\text{MS}}), i_{\text{MS}} = \arg\max_{i \in I} \text{sim}(i, i^*)$.

For general preference optimization, we track the highest score achieved by any generated image (*max*), the improvement from the initial score to this maximum (*max - init*), and the number of iterations required to reach the maximum score (*K**). These metrics provide a comprehensive assessment of our framework's ability to enhance user satisfaction and align generated images with individual user preferences effectively.

4.3 Results

General Preference Optimization Table 2 shows the results for general preference optimization. Without any ablations, the maximum score achieved was 6.50, with an improvement of 1.78 from the initial score. It took an average of 5.33 iterations to reach this maximum score. Adding the score and prompt information (s+p) resulted in a lower maximum score of 5.93. Further adding feedback (s+p+f) and images (s+p+f+I) did not sub-

Table 1: Results of personal preference optimization. K^* represents the number of iterations to reach the maximum score. s, p, f, and I denote score, prompt, feedback, and images, respectively. Details on the evaluation metrics are provided in Section 4.2.

Ablations	*max*	*max - init*	K^*	*max sim*	*score of the most sim*
$s+p$	6.78	0.88	5.83	0.67	5.50
$s+p+f$	7.20	1.19	4.33	0.76	5.00
$s+p+f+I$	6.88	1.37	5.66	0.78	7.17

Table 2: General preference optimization results.

Ablations	max	max - init	K^*
none	6.50	1.78	5.33
$s+p$	5.93	1.21	5.62
$s+p+f$	5.88	1.16	5.51
$s+p+f+I$	5.92	1.20	5.66

Table 3: PIAA results

Method	SROCC	MSE	MAE
MAML	0.44	9.79	2.67
FineTune	0.45	11.14	2.78
kNN	0.12	9.01	2.53
MAML	0.42	9.56	2.63
FineTune	0.43	12.71	2.94
kNN	0.27	7.88	2.34

stantially change the results compared to the s+p ablation. These results suggest that for optimizing general preferences, the additional information beyond scores and prompts does not provide significant benefits.

PIAA Table 3 presents the results for the personalized image aesthetic assessment (PIAA) models. The MAML and FineTune approaches performed comparably, achieving Spearman rank order correlation coefficients (SROCC) of 0.44 and 0.45 respectively. The kNN method had a much lower SROCC of 0.12. In terms of mean squared error (MSE) and mean absolute error (MAE), the kNN approach had the lowest errors of 9.01 and 2.53 respectively. These results indicate that the MAML and FineTune methods are better at ranking images according to personal preferences, while the kNN method makes predictions with smaller absolute errors.

Personalized Preference Optimization The personalized preference optimization results are shown in Table 1. Using just the score and prompt information (s+p), a maximum score of 6.78 was reached, with an improvement of 0.88 over the initial score. Adding feedback (s+p+f) increased the maximum score to 7.20 and the improvement to 1.19, while reducing the number of iterations needed to 4.33. Further adding images (s+p+f+I) resulted in a maximum score of 6.88, an improvement of 1.37, and 5.66 iterations. The similarity metrics provide additional insights. The highest similarity of 0.78 between the optimized image and a previous highly rated image was achieved with all information included (s+p+f+I). However, the score of the most similar previous image was highest at 7.17 when using all information types. These results suggest that incorporating feedback and images helps the model generate novel highly rated images that still share similarities with the user's previous preferences.

5 Conclusion

We propose an optimization framework specifically focused on personal preference through in-context learning. To simulate real-world personal preferences, we utilized the PIAA method to train the personal model. For preference optimization, we introduced several novel techniques, including multi-prompt optimization and comprehensive feedback integration, to enhance the optimization process. Our results demonstrate that the proposed framework effectively optimizes personal preferences using in-context learning, which is a training-free approach, indicating its potential for real-world applications.

6 Limitations

Our experiments use a virtual human model to approximate real human responses. However, this approximation introduces discrepancies, making it uncertain whether the framework will perform effectively in real-world scenarios. Also, the score improvements achieved through optimization are not particularly substantial. There are two main reasons for this. First, the problem itself is inherently difficult to solve. Human preferences are complex, making it challenging to accurately model their structure. Second, the method's capability may not be sufficient for addressing this challenging problem. In-context learning may be less effective for such complex tasks compared to fine-tuning using gradient-based methods. Since the LLM is trained primarily on textual data, it may not fully capture preferences that are difficult to express in text.

References

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.

Tianyu Gao, Adam Fisch, and Danqi Chen. 2020. Making pre-trained language models better few-shot learners. *arXiv preprint arXiv:2012.15723*.

Yaru Hao, Zewen Chi, Li Dong, and Furu Wei. 2023. Optimizing prompts for text-to-image generation. *Preprint*, arXiv:2212.09611.

Yuval Kirstain, Adam Polyak, Uriel Singer, Shahbuland Matiana, Joe Penna, and Omer Levy. 2023. Pick-a-pic: An open dataset of user preferences for text-to-image generation. *Preprint*, arXiv:2305.01569.

Wenhao Liu, Xiaohua Wang, Muling Wu, Tianlong Li, Changze Lv, Zixuan Ling, Jianhao Zhu, Cenyuan Zhang, Xiaoqing Zheng, and Xuanjing Huang. 2023. Aligning large language models with human preferences through representation engineering. *arXiv preprint arXiv:2312.15997*.

Nikita Pavlichenko and Dmitry Ustalov. 2023. Best prompts for text-to-image models and how to find them. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 2067–2071.

John David Pressman, Katherine Crowson, and Simulacra Captions Contributors. 2022. Simulacra aesthetic captions. Technical Report Version 1.0, Stability AI. url https://github.com/JD-P/simulacra-aesthetic-captions .

Reid Pryzant, Dan Iter, Jerry Li, Yin Tat Lee, Chenguang Zhu, and Michael Zeng. 2023. Automatic prompt optimization with" gradient descent" and beam search. *arXiv preprint arXiv:2305.03495*.

Jian Ren, Xiaohui Shen, Zhe Lin, Radomir Mech, and David J Foran. 2017. Personalized image aesthetics. In *Proceedings of the IEEE international conference on computer vision*, pages 638–647.

Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. 2023. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22500–22510.

Jing Shi, Wei Xiong, Zhe Lin, and Hyun Joon Jung. 2024. Instantbooth: Personalized text-to-image generation without test-time finetuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8543–8552.

Zijie J. Wang, Evan Montoya, David Munechika, Haoyang Yang, Benjamin Hoover, and Duen Horng Chau. 2022. DiffusionDB: A large-scale prompt gallery dataset for text-to-image generative models. *arXiv:2210.14896 [cs]*.

Xiaoshi Wu, Keqiang Sun, Feng Zhu, Rui Zhao, and Hongsheng Li. 2023. Better aligning text-to-image models with human preference. *arXiv preprint arXiv:2303.14420*.

Jiazheng Xu, Xiao Liu, Yuchen Wu, Yuxuan Tong, Qinkai Li, Ming Ding, Jie Tang, and Yuxiao Dong. 2023. Imagereward: Learning and evaluating human preferences for text-to-image generation. *Preprint*, arXiv:2304.05977.

Chengrun Yang, Xuezhi Wang, Yifeng Lu, Hanxiao Liu, Quoc V Le, Denny Zhou, and Xinyun Chen. 2023. Large language models as optimizers. *arXiv preprint arXiv:2309.03409*.

Yuzhe Yang, Liwu Xu, Leida Li, Nan Qie, Yaqian Li, Peng Zhang, and Yandong Guo. 2022. Personalized image aesthetics assessment with rich attributes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19861–19869.

Mert Yuksekgonul, Federico Bianchi, Joseph Boen, Sheng Liu, Zhi Huang, Carlos Guestrin, and James Zou. 2024. Textgrad: Automatic" differentiation" via text. *arXiv preprint arXiv:2406.07496*.