

DiffExp: Efficient Exploration in Reward Fine-tuning for Text-to-Image Diffusion Models

Daewon Chae^{1*}, June Suk Choi^{2*}, Jinkyu Kim^{1†}, Kimin Lee^{2†}

¹Korea University, South Korea

²KAIST, South Korea

cdw098@korea.ac.kr, w_choi@kaist.ac.kr, jinkyukim@korea.ac.kr, kiminlee@kaist.ac.kr

Abstract

Fine-tuning text-to-image diffusion models to maximize rewards has proven effective for enhancing model performance. However, reward fine-tuning methods often suffer from slow convergence due to online sample generation. Therefore, obtaining diverse samples with strong reward signals is crucial for improving sample efficiency and overall performance. In this work, we introduce DiffExp, a simple yet effective exploration strategy for reward fine-tuning of text-to-image models. Our approach employs two key strategies: (a) dynamically adjusting the scale of classifier-free guidance to enhance sample diversity, and (b) randomly weighting phrases of the text prompt to exploit high-quality reward signals. We demonstrate that these strategies significantly enhance exploration during online sample generation, improving the sample efficiency of recent reward fine-tuning methods, such as DDPO and AlignProp.

Introduction

Reward fine-tuning (Black et al. 2024; Fan et al. 2024; Clark et al. 2024; Prabhudesai et al. 2023; Lee et al. 2023) has recently emerged as a powerful method for improving text-to-image diffusion models (Podell et al. 2023; Rombach et al. 2022). Unlike the conventional optimization strategy of likelihood maximization, this framework focuses on maximizing reward scores that measure the quality of model outputs, such as image-text alignment (Kirstain et al. 2024; Lee et al. 2023; Wu et al. 2023; Xu et al. 2024) and image fidelity (Schuhmann et al. 2022). Several methods including policy gradient (Black et al. 2024; Fan et al. 2024) and direct backpropagation (Clark et al. 2024; Prabhudesai et al. 2023) have been studied for reward maximization. These methods have shown promising results in improving image-text alignment (Black et al. 2024; Fan et al. 2024), reducing undesired biases (Fan et al. 2024), and removing artifacts from generated images (Wu et al. 2023).

Since this reward-based fine-tuning involves online sample generation, the reward optimization depends on what samples are produced during the generation process. In other words, if samples with good reward signals are not obtained during the generation process, the model will con-

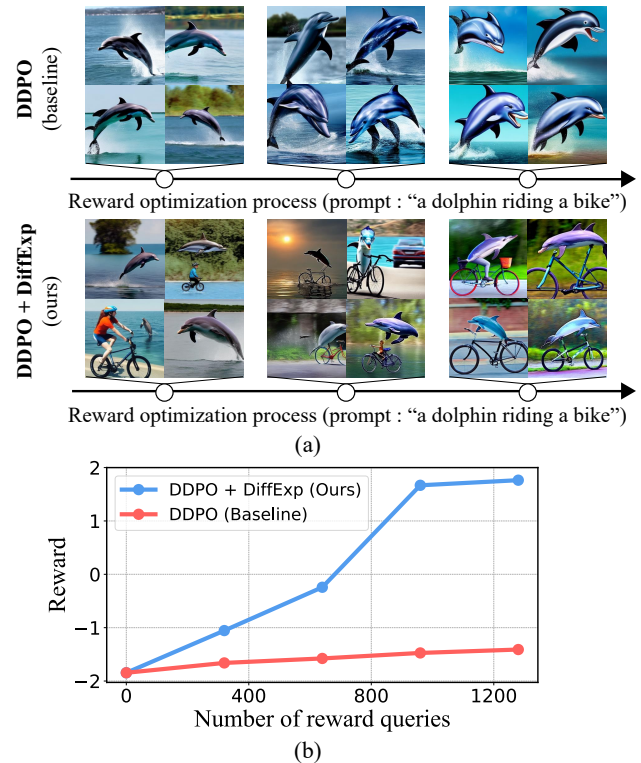


Figure 1: (a) Generated images during the reward optimization process with our proposed method called DiffExp given the prompt “a dolphin riding a bike,” which is often challenging for existing reward fine-tuning approaches, including our baseline DDPO (Black et al. 2024). (b) We also provide corresponding reward curves against the number of reward queries, where our method indeed improves its sample efficiency during reward optimization process, capturing good reward signals for reward fine-tuning.

verge slowly due to the lack of these signals. To demonstrate this, we conduct reward-based fine-tuning using the single prompt “a dolphin riding a bike”, which is known to be challenging for reward optimization (Black et al. 2024). In this experiment, we fine-tune Stable Diffusion 1.5 (Rombach et al. 2022) using the policy gradient method (i.e.,

*Equal contribution. †Equal advising.

DDPO (Black et al. 2024)) to maximize image-text alignment based on ImageReward (Xu et al. 2024) scores. As shown in Figure 1, the current method fails to find a good reward signal (i.e., images containing a bicycle) during the optimization process, resulting in trivial solutions that change the style of the image without improving image-text alignment. This highlights the need for enhanced exploration in the sample generation process to better capture good reward signals.

In this work, we introduce DiffExp (DiffusionExplore), an exploration method designed to efficiently fine-tune text-to-image diffusion models using rewards. Our method relies on two main strategies: dynamically controlling the scale of classifier-free guidance (CFG) (Ho and Salimans 2022) and randomly weighting certain phrases in the text prompt. CFG is a sampling technique that balances sample diversity and fidelity. Unlike conventional approaches that fix the CFG scale during the denoising process, we propose initially setting this scale to an extremely low value and increasing it in the later step. This dynamic scheduling improves the diversity of the online samples while maintaining their fidelity. To further enhance diverse sample generation, we perform additional exploration by randomly assigning weights to certain phrases within the text prompt. This approach creates different images where specific elements of the text prompt are emphasized randomly, thereby promoting the emergence of valuable reward signals. We find that these two key strategies significantly boost the effectiveness of reward-based fine-tuning (see blue curve in Figure 1).

To verify the effectiveness of our exploration method, DiffExp, we integrate it with two popular reward fine-tuning techniques: policy gradient method (Black et al. 2024) and direct reward backpropagation method (Prabhudesai et al. 2023). Specifically, we fine-tune the Stable Diffusion model (Rombach et al. 2022) using our method to optimize various reward functions, including Aesthetic scorer (Schuhmann et al. 2022) and PickScore (Kirstain et al. 2024). Our experiments demonstrate that DiffExp improves the sample-efficiency of both the policy gradient and reward backpropagation methods by encouraging more efficient exploration, which in turn increases the diversity of online samples. Finally, we show that our method also significantly improves the quality of high-resolution text-to-image models like SDXL (Podell et al. 2023). We summarize our contributions as follows:

- We propose DiffExp: an exploration method that improves diversity of online sample generation for efficient reward fine-tuning of text-to-image diffusion models.
- We show that our method enhances sample-efficiency of reward optimization in various fine-tuning methods, such as policy gradient and direct backpropagation.
- We demonstrate that our method is also effective when applied to the more recent SDXL model, as well as on the more challenging prompt set, DrawBench.

Related Work

Fine-tuning diffusion models using reward functions. Optimizing text-to-image diffusion models on rewards has

been proven effective in supervising the final output of a diffusion model by a given reward. They are especially useful when training objectives are difficult to define given a set of images, such as human preference. Early studies (Lee et al. 2023) use supervised approaches to fine-tune the pre-trained models on images, which are weighted according to the reward function. As they are not trained online on examples, recent works (Black et al. 2024; Fan et al. 2024) further explore optimizing rewards using policy gradient algorithms by formulating the denoising process as a multi-step decision-making task. Though these RL approaches can flexibly train models on non-differentiable rewards, they might lose useful signals as analytic gradients of many reward functions are, in practice, available. AlignProp (Prabhudesai et al. 2023) and DRaFT (Clark et al. 2024) explored optimizing diffusion models with differentiable rewards of human preference (Kirstain et al. 2024; Wu et al. 2023), effectively improving the image generation quality. In this work, we explore another strategy by focusing on an exploration technique toward a more efficient reward optimization process.

Efficient reward fine-tuning with exploration. In reinforcement learning (RL), where agents learn to make decisions by interacting with the environment, exploration is a crucial aspect. Exploration directly impacts the discovery of optimal policies, as it allows agents to gather diverse experiences. Intrinsic Curiosity Module (ICM) (Pathak et al. 2017) motivates exploration by rewarding actions that lead to unseen states, promoting discovery of unexplored areas of the environment. Soft Actor-Critic (SAC) (Haarnoja et al. 2018) enhances exploration by employing entropy regularization, encouraging a wider range of actions. These works promote exploration by driving the agent to gather new experiences. There are works that investigate efficient reward fine-tuning using generative diffusion models as well. A concurrent work (Uehara et al. 2024a) proposed to endorse exploration by promoting diversity, formulating the fine-tuning of diffusion models as entropy-regularized control against pre-trained diffusion model. Another work (Uehara et al. 2024b) facilitates exploration by integrating an uncertainty model and KL regularization into the diffusion model tuning process. Inspired by the success of exploration strategy in RL literature and the reward fine-tuning of diffusion models, we propose an effective exploration method for reward fine-tuning of text-to-image diffusion models.

Preliminaries

Text-to-image diffusion models. Diffusion models (Ho, Jain, and Abbeel 2020; Song et al. 2021) are a class of generative models that represent the data distribution $p(x)$ by iteratively transforming Gaussian noise into data samples through a denoising process. These models can be extended to model the conditional distribution $p(x|c)$ by incorporating an additional condition c . In this work, we consider text-to-image diffusion models that utilize textual conditions to guide image generation. The models consist of parameterized denoising function ϵ_θ and are trained by predicting added noise ϵ to the image x at timestep t . Formally, given

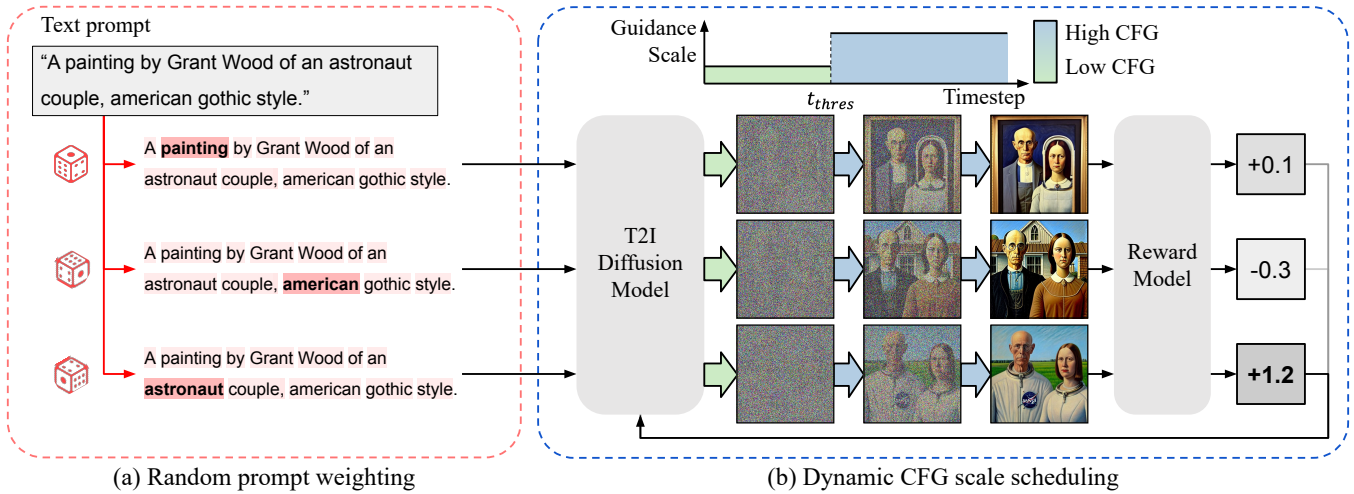


Figure 2: An overview of our proposed method called DiffExp, which consists of two main steps: (a) random prompt weighting and (b) dynamic scheduling of the CFG (classifier-free guidance) scale. In (a), word embeddings of the given prompt are randomly and differently weighted, which are then consumed by the image generation process, increasing the diversity of generated images. Further, in (b), the CFG scale of the denoising process is dynamically scheduled to control models to generate high-quality and diverse images, which is often challenging with a constantly set CFG scale.

dataset D comprising of image-text pairs (x, c) , the training objective is as follows:

$$\mathcal{L}_{DM} := E_{x,c,\epsilon,t} \left[\|\epsilon - \epsilon_\theta(x_t, t, c)\|_2^2 \right], \quad (1)$$

where ϵ is a Gaussian noise $\sim \mathcal{N}(0, I)$, t is a timestep sampled from uniform distribution $\mathcal{U}(0, T)$, and x_t is the noised image by diffusion process at a timestep t . At inference phase, sampling begins with drawing $x_T \sim \mathcal{N}(0, I)$, and denoising process is iteratively conducted to sample a denoised image using the estimated noise by ϵ_θ . Specifically, the estimated noise is utilized to obtain the mean of transition distribution $p(x_{t-1}|x_t, c)$ in the denoising process, as follows:

$$\mu_\theta(x_t, t, c) = \frac{1}{\sqrt{\alpha_t}} \left(x_t - \frac{\beta_t}{\sqrt{1 - \alpha_t}} \epsilon_\theta(x_t, t, c) \right), \quad (2)$$

$$p_\theta(x_{t-1} | x_t, c) = \mathcal{N}(\mu_\theta(x_t, t, c), \Sigma_t), \quad (3)$$

where α_t, β_t are pre-defined constants used for timestep dependent denoising, Σ_t represents covariance matrix of denoising transition, and $\bar{\alpha}_t := \prod_{s=1}^t \alpha_s$.

Classifier-free guidance (CFG). Ho and Salimans (2022) proposed *classifier-free guidance* (CFG), which serves as the basis of recent text-to-images diffusion models. In CFG, the denoising network predicts the noise based on the linear combination of conditional and unconditional noise estimates as follows:

$$\tilde{\epsilon}_\theta(x_t, t, c) = \epsilon_\theta(x_t, t) + w * (\epsilon_\theta(x_t, t, c) - \epsilon_\theta(x_t, t)), \quad (4)$$

where w is the guidance scale of CFG. By incorporating the unconditional noise estimate into the prediction, we can achieve a trade-off between sample quality and diversity. Specifically, increasing the guidance weight w enhances the fidelity of the samples but reduces their variety.

Method

In this section, we introduce DiffExp (DiffusionExplore): an exploration method for efficient fine-tuning of text-to-image diffusion models using rewards. We first describe the problem setup regarding reward optimization of text-to-image diffusion models. Then, we present our approach, which promotes exploration in sample generation by scheduling CFG scale and randomly weighting a phrase of prompt.

Problem Formulation

We consider the problem of fine-tuning a text-to-image model using a reward model. The goal is to maximize the expected reward $r(x_0, c)$ for image x_0 generated by the model p_θ using text prompt c . Formally, it can be formulated as follows:

$$\max_{\theta} E_{p(c)} E_{p_\theta(x_0|c)} [r(x_0, c)], \quad (5)$$

where $p(c)$ is the training prompt distribution and $p_\theta(\cdot|c)$ is the sample distribution for the generated image x_0 . To achieve this goal, we consider online optimization methods that continuously generate image samples while optimizing models with rewards. Specifically, we utilize two types of reward fine-tuning methods: *the policy gradient method* (Fan et al. 2024; Black et al. 2024) and *direct reward backpropagation method* (Clark et al. 2024; Prabhudesai et al. 2023). In the policy gradient method, the denoising process is defined as a multi-step Markov Decision Process, and the transition distribution $p_\theta(x_{t-1} | x_t, c)$ (described in Equation 3) is treated as a policy. Based on this framework, the gradient of the objective in Equation 1 is computed as follows to update the model:

$$E_{p(c), p_\theta(x_0|c)} \left[\sum_{t=1}^T \nabla_{\theta} \log p_\theta(x_{t-1} | x_t, c) r(x_0, c) \right]. \quad (6)$$

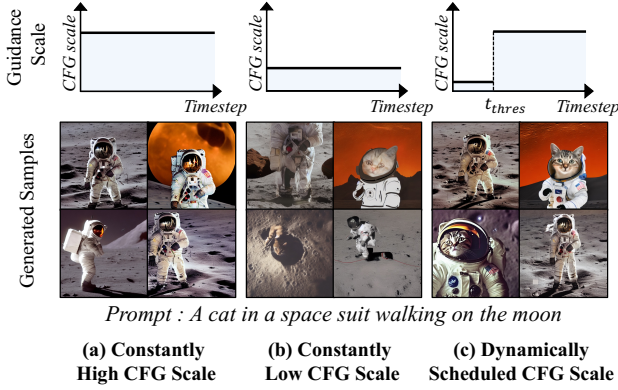


Figure 3: Comparison of the generated images with different CFG scale scheduling strategies: (a) constantly high CFG scale, (b) constantly low CFG scale, and (c) dynamically scheduled CFG scale. We observe that a model often shows increased sample diversity with a low CFG scale but suffers from degraded image quality, which is generally the opposite with a high CFG scale. Instead, dynamically scheduling the CFG scale balances sample diversity and image quality.

As an alternative approach, the direct reward backpropagation method assumes that the reward model is differentiable and directly backpropagates the gradient of the reward function through the denoising process to optimize the objective in Equation 5.

Dynamic Scheduling of CFG Scale

To effectively discover good reward signals, it is crucial to generate diverse image samples through exploration. To achieve this, we propose dynamically scheduling the CFG scale of the denoising process during online optimization. As detailed in the preliminaries, the CFG scale controls the trade-off between sample quality and diversity: a high CFG scale yields high-quality but low-diversity samples, whereas a low CFG scale promotes diversity at the cost of reduced quality (see Figure 3). To optimize this balance, we start with a low CFG scale during the early stages of the denoising process and switch to a high CFG scale after the t_{thres} denoising step, as follows:

$$\tilde{\epsilon}_\theta(x_t, t, c) = \epsilon_\theta(x_t, t) + w(t) \cdot (\epsilon_\theta(x_t, t, c) - \epsilon_\theta(x_t, t)),$$

where

$$w(t) = \begin{cases} w_l & \text{if } t > t_{thres}, \\ w_h & \text{if } t \leq t_{thres}, \end{cases} \quad (7)$$

and w_l, w_h ($w_l < w_h$) are the CFG scale values. Note that in the denoising process, we get Gaussian noise at $t = T$ while the final image is at $t = 0$. We set w_l to an extremely low value and w_h to an ordinary CFG value (i.e., 5.0 or 7.5). This dynamic scheduling adaptively balances between image quality and diversity, allowing for generating diverse image samples without sacrificing overall sample quality.

Random Prompt Weighting

To further promote diverse sample generation, we propose an additional exploration method that alters text prompts.

Specifically, we increase the weight of a random word (i.e. token) in the text prompt embedding. This adjustment emphasizes the selected word in the generated image, leading to variations where different elements are highlighted. Consider the text prompt "A dolphin riding a bike" for example, which typically results in images featuring only a dolphin, with the bike often omitted. By increasing the weight of the word "bike", there is a high chance that text-to-image diffusion models generate the image containing a bike. This leads to more complete representations of the original prompt, thereby achieving higher prompt alignment rewards and enhancing image diversity.

Formally, given a text embedding c with N_{words} words, we choose a random word with index $0 \leq i < N_{words}$. Then, we increase the weight of the chosen word as follows:

$$c[i] \leftarrow c_{null} + w_{prompt} * (c[i] - c_{null}), \quad (8)$$

where c_{null} is the text embedding that corresponds to the empty text (""), and w_{prompt} is the prompt weight. We find that sampling w_{prompt} randomly from $\mathcal{U}(1, 1.2)$ every time is generally successful.

Experiments

We design our experiments to investigate the following questions:

1. Can our exploration method enhance the sample-efficiency of reward fine-tuning methods?
2. Can our exploration method improve the quality of generated samples?
3. Can our exploration method exhibit generalization ability for unseen prompts that are not used during fine-tuning?

Baselines. As discussed in the Problem Formulation section, we evaluate the effectiveness of our exploration method using two types of state-of-the-art reward fine-tuning methods: policy gradient method (Fan et al. 2024; Black et al. 2024) and direct reward backpropagation method (Clark et al. 2024; Prabhudesai et al. 2023). We combine our approach with each of these methods and compare the performance with the original methods. Specifically, we use (1) DDPO (Black et al. 2024) for the policy gradient method and (2) AlignProp (Prabhudesai et al. 2023) for direct reward backpropagation method.

Reward functions. To evaluate our method across different types of rewards, we conduct experiments using two distinct reward functions. First, we utilize an Aesthetic Score (Schuhmann et al. 2022), which is trained to predict the aesthetic quality of images. Following the baseline (Black et al. 2024; Prabhudesai et al. 2023), we use 45 animal names as training prompts for the aesthetic quality task. Second, in order to improve image-text alignment, we employ PickScore (Kirstain et al. 2024), an open-source reward model trained on a large-scale human feedback dataset. Based on the baseline (Black et al. 2024), we use a total of 135 prompts for the image-text alignment task, combining 45 different animal names with 3 different activities (e.g., "a monkey washing the dishes"). We provide the entire set of prompts used for training in the supplementary materials.




Figure 4: Reward curves for training prompts. At each checkpoint, we generate 10 images per seen prompt and use the average of their reward scores. Our sampling method is employed only during fine-tuning, not for plotting this curve.

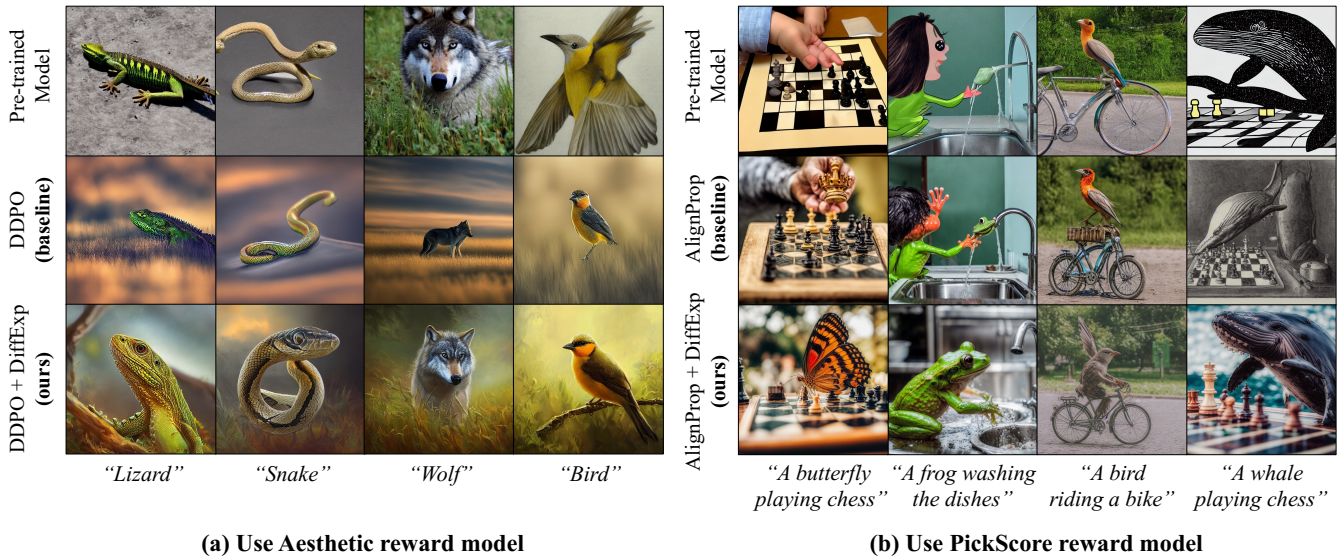


Figure 5: Generated samples from baselines and ours with (a) Aesthetic and (b) PickScore reward models. Notably, ours generates images with comparably high aesthetic quality (see (a)) and produces images with better image-text alignment given the prompts (see (b)). Note that images in the same column are generated with the same random seed.

Implementation details. In our experiment, we use Stable Diffusion v1.5 (Rombach et al. 2022) as the pre-trained text-to-image diffusion model. Following the baselines, we employ Low-Rank Adaptation (LoRA) (Hu et al. 2022) rather than fine-tuning the weights of pre-trained text-to-image model. As for scheduling exploration, we apply our exploration method only up to the three-fourths of the entire fine-tuning. We provide more implementation details in the supplemental material.

Experimental Results

Sample efficiency in reward fine-tuning. We begin by conducting experiments and comparing our method with DDPO and AlignProp across two reward functions: Aesthetic and PickScore. Figure 4 presents the results, highlighting the achieved reward scores and sample efficiency. Notably, our method outperforms the baselines in terms of reward scores, while demonstrating better sample efficiency. Specifically, our approach required approximately

20% fewer samples to achieve the same reward score as DDPO and AlignProp.

Qualitative comparison of image quality. We provide visual examples generated by each method in Figure 5. With the Aesthetic reward function, we observe that our method generates images with relatively high aesthetic quality, which aligns with the intent of reward function (see Figure 5 (a)). Notably, with the PickScore reward function, we observe that our method plays an important role in generating images with high image-text alignment (see Figure 5 (b)). For example, given a text prompt “A butterfly playing chess,” all baselines result in generating a human (not a butterfly) playing chess (see 1st column). We provide more diverse generated samples in the supplemental materials.

Quantitative comparison of image quality. To further quantitatively evaluate the sample quality, we use an “unseen” reward function called ImageReward (Xu et al. 2024) (which is trained on a large-scale human preference dataset)

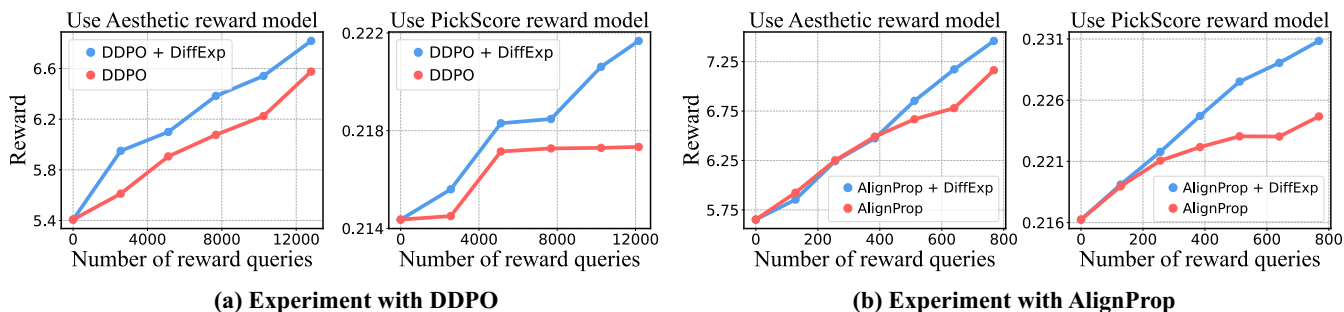


Figure 6: Reward curve for unseen prompts. At each checkpoint, we generate 10 images per unseen prompt and use the average of their reward scores. Our sampling method is employed only during fine-tuning, not for plotting this curve.

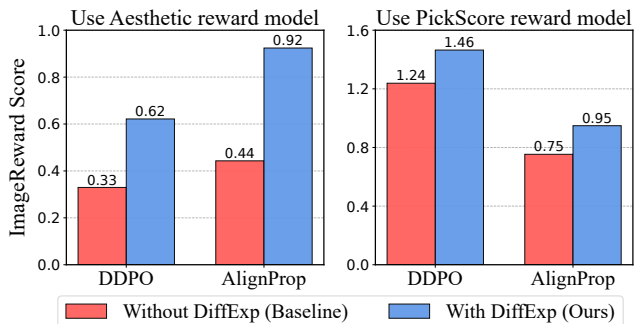


Figure 7: Comparison of an unseen ImageReward (Xu et al. 2024) score after reward fine-tuning with and without our proposed DiffExp on our baselines, such as DDPO and AlignProp. The proposed method DiffExp achieves higher reward compared to all other baselines. Note that models are optimized with either Aesthetic or PickScore as a reward function.

that was not used during reward fine-tuning (i.e., not in the optimization objectives). In Figure 7, we observe that ours clearly achieves higher ImageReward scores compared to baselines, such as DDPO and AlignProp fine-tuned with Aesthetic and PickScore reward functions. This demonstrates that our DiffExp indeed improves the overall image quality rather than merely increasing the reward through reward hacking (Gao, Schulman, and Hilton 2023; Kim et al. 2024).

Generalization to unseen prompts. We also investigate the model’s capacity for generalization to new prompts, which were not used during reward fine-tuning. Following the standard evaluation protocol (Black et al. 2024; Prabhudesai et al. 2023), we consider a novel test set of animal names (for models fine-tuned with Aesthetic reward model) and activities (for models fine-tuned with PickScore reward model) that were not encountered during the training phase. A more detailed explanation of the experimental setup is provided in the supplemental material. As shown in Figure 6, which shows reward curves to compare ours with baselines, a model with DiffExp consistently outper-

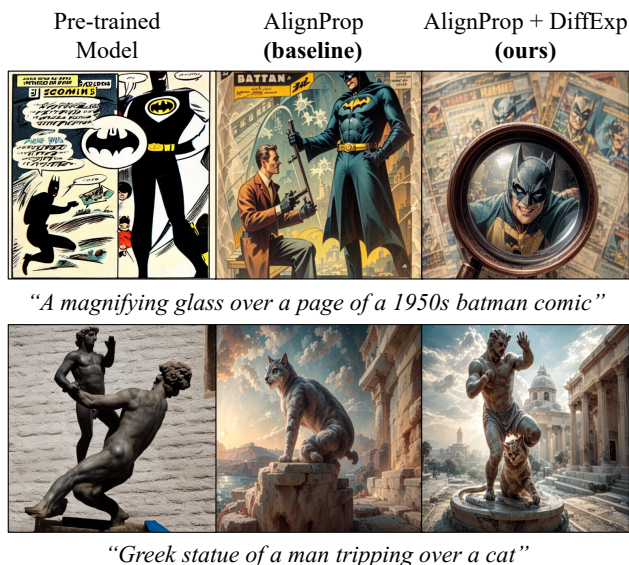


Figure 8: Generated image samples with DrawBench prompts, which are known to be challenging for the current text-to-image models. Our method successfully generates images with high sample quality and image-text alignment, capturing contexts of the given prompts such as “a magnifying glass” and “a man tripping over a cat”.

forms the baseline models in terms of rewards and sample efficiency, confirming its ability to generalize well to unseen prompts.

Experiment with challenging prompts. Our proposed DiffExp provides an effective way to encourage the generation of diverse samples during reward fine-tuning. Thus, we argue that DiffExp is advantageous for learning more complicated prompts that are often challenging with conventional approaches, such as AlignProp. To evaluate this, we collect 58 challenging prompts from DrawBench (Saharia et al. 2022), which is well-known to be challenging for the current text-to-image models. We use AlignProp as our baseline model and fine-tune with an ensemble of HPS (Wu et al. 2023), Aesthetic, and PickScore reward mod-

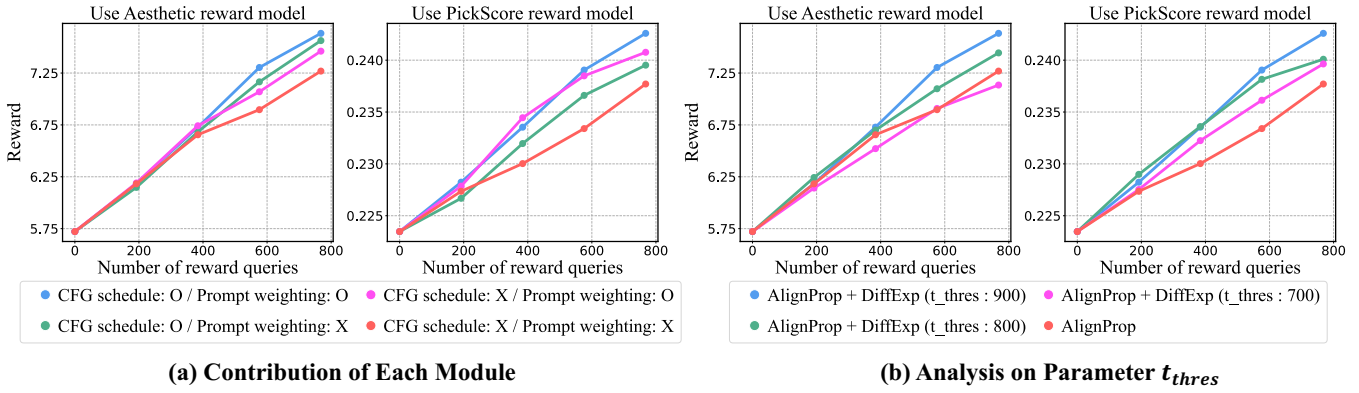


Figure 9: The results of the ablation study represent (a) the effect of each main module and (b) the impact of different hyper-parameter values of t_{thres} .

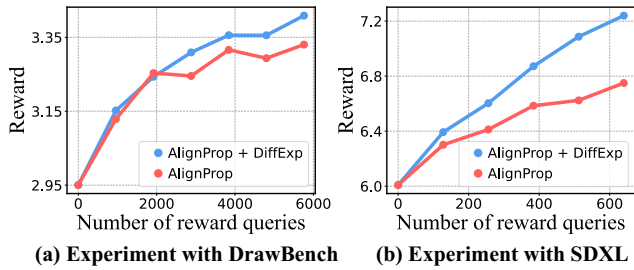


Figure 10: Reward curves for our experiments with (a) DrawBench and (b) Stable Diffusion XL (SDXL). The former evaluates the ability to generate images with challenging prompts, while the latter evaluates the adaptability to the current SOTA model.

els, which is reported as effective in previous studies (Clark et al. 2024; Deng et al. 2024), especially for complicated prompts. As shown in Figure 8, the pre-trained model and AlignProp struggle to generate images aligned well with the given prompts. For example, such models are often overfitted to certain phrases (e.g., “Batman”), ignoring other contexts (e.g., “magnifying glass”). In contrast, our proposed DiffExp successfully generates samples with high image quality and high text fidelity (see the last column), confirming that our exploration method effectively improves reward fine-tuning with complicated (and challenging) prompts. We provide more diverse examples in supplemental material. We further provide reward curves for this experiment in Figure 10 (a), which is consistent with our previous analysis, demonstrating the effectiveness of our approach even to the challenging scenarios.

Experiment with Stable Diffusion XL (SDXL). In order to verify the usefulness of DiffExp with state-of-the-art models, we fine-tune SDXL (Podell et al. 2023) to maximize Aesthetic Score on animal prompts using DiffExp. As shown in Figure 10, using AlignProp with DiffExp achieves a significantly higher maximum reward value as well as better sample efficiency compared to AlignProp without Diff-

Exp. This trend is in accordance with our experiments using Stable Diffusion v1.5, which are presented in Figure 4. We include the visual examples in the supplementary material.

Ablation studies. We conduct an ablation study to see the effect of each proposed module: (i) random prompt weighting and (ii) dynamic scheduling of the CFG scale. In Figure 9 (a), we provide reward curves with variants of our models with and without those modules. We observe that all our variants generally perform better than our baseline, AlignProp, in terms of rewards and sample efficiency, while using those modules together provides the best results. This trend is consistent with different reward models, such as Aesthetic and PickScore. Further, we experiment with different values of hyper-parameter, t_{thres} , which determines how long the CFG scale should be maintained to a low value. In Figure 9 (b), we provide reward curves for variants of our models with $t_{thres} = \{900, 800, 700\}$. Note that in the denoising process, we have Gaussian noise at $t = 1000$ while the final image is at $t = 0$. We observe that higher t_{thres} generally provides better results in terms of rewards, with the best performance with $t_{thres} = 900$, and lowering it degrades the image quality (consistent with our analysis in Figure 3).

Conclusion

In this work, we propose DiffExp, an exploration method for promoting diverse reward signals during the reward optimization of text-to-image models. We demonstrate that adjusting the CFG scale of denoising process and randomly weighting certain phrase within prompts can serve as effective exploration strategy by improving the diversity of on-line sample generation. We conduct extensive experiments to verify that DiffExp improves both sample efficiency and generated image quality, and demonstrate this across various reward fine-tuning methods such as DDPO or AlignProp. Furthermore, we conduct analysis using more advanced prompt sets such as DrawBench, and apply our method to the more advanced diffusion models such as SDXL, both of which result in significant performance improvements.

Acknowledgements

This work was supported by Artificial intelligence industrial convergence cluster development project funded by the Ministry of Science and ICT(MSIT, Korea)&Gwangju Metropolitan City (10%) and supported by Institute of Information & communications Technology Planning & Evaluation(IITP) under the Leading Generative AI Human Resources Development(IITP-2025-RS-2024-00397085, 20%), Adaptive Personality for Intelligent Agents (RS-2022-II220043, 20%), ICT Creative Consilience Program (IITP-2025-2020-0-01819, 10%), Global AI Frontier Lab (RS-2024-00509279, 20%), Artificial Intelligence Graduate School Program(KAIST) (RS-2019-II190075, 10%), and Development and Study of AI Technologies to Inexpensively Conform to Evolving Policy on Ethics (RS-2022-II220184, 2022-0-00184, 10%) grant funded by the Korea government(MSIT). This work was also supported by the Hyundai Motor Chung Mong-Koo Foundation.

References

- Black, K.; Janner, M.; Du, Y.; Kostrikov, I.; and Levine, S. 2024. Training Diffusion Models with Reinforcement Learning. In *The Twelfth International Conference on Learning Representations*.
- Clark, K.; Vicol, P.; Swersky, K.; and Fleet, D. J. 2024. Directly Fine-Tuning Diffusion Models on Differentiable Rewards. In *The Twelfth International Conference on Learning Representations*.
- Deng, F.; Wang, Q.; Wei, W.; Grundmann, M.; and Hou, T. 2024. PRDP: Proximal Reward Difference Prediction for Large-Scale Reward Finetuning of Diffusion Models. *arXiv preprint arXiv:2402.08714*.
- Fan, Y.; Watkins, O.; Du, Y.; Liu, H.; Ryu, M.; Boutilier, C.; Abbeel, P.; Ghavamzadeh, M.; Lee, K.; and Lee, K. 2024. Reinforcement learning for fine-tuning text-to-image diffusion models. *Advances in Neural Information Processing Systems*, 36.
- Gao, L.; Schulman, J.; and Hilton, J. 2023. Scaling laws for reward model overoptimization. In *International Conference on Machine Learning*, 10835–10866. PMLR.
- Haarnoja, T.; Zhou, A.; Abbeel, P.; and Levine, S. 2018. Soft Actor-Critic: Off-Policy Maximum Entropy Deep Reinforcement Learning with a Stochastic Actor.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33: 6840–6851.
- Ho, J.; and Salimans, T. 2022. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*.
- Hu, E. J.; yelong shen; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; and Chen, W. 2022. LoRA: Low-Rank Adaptation of Large Language Models. In *International Conference on Learning Representations*.
- Kim, K.; Jeong, J.; An, M.; Ghavamzadeh, M.; Dvijotham, K. D.; Shin, J.; and Lee, K. 2024. Confidence-aware Reward Optimization for Fine-tuning Text-to-Image Models. In *The Twelfth International Conference on Learning Representations*.
- Kirstain, Y.; Polyak, A.; Singer, U.; Matiana, S.; Penna, J.; and Levy, O. 2024. Pick-a-pic: An open dataset of user preferences for text-to-image generation. *Advances in Neural Information Processing Systems*, 36.
- Lee, K.; Liu, H.; Ryu, M.; Watkins, O.; Du, Y.; Boutilier, C.; Abbeel, P.; Ghavamzadeh, M.; and Gu, S. S. 2023. Aligning text-to-image models using human feedback. *arXiv preprint arXiv:2302.12192*.
- Pathak, D.; Agrawal, P.; Efros, A. A.; and Darrell, T. 2017. Curiosity-driven Exploration by Self-supervised Prediction.
- Podell, D.; English, Z.; Lacey, K.; Blattmann, A.; Dockhorn, T.; Müller, J.; Penna, J.; and Rombach, R. 2023. Sdxl: Improving latent diffusion models for high-resolution image synthesis. *arXiv preprint arXiv:2307.01952*.
- Prabhudesai, M.; Goyal, A.; Pathak, D.; and Fragkiadaki, K. 2023. Aligning text-to-image diffusion models with reward backpropagation. *arXiv preprint arXiv:2310.03739*.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10684–10695.
- Saharia, C.; Chan, W.; Saxena, S.; Li, L.; Whang, J.; Denton, E. L.; Ghasemipour, K.; Gontijo Lopes, R.; Karagol Ayan, B.; Salimans, T.; et al. 2022. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in Neural Information Processing Systems*, 35: 36479–36494.
- Schuhmann, C.; Beaumont, R.; Vencu, R.; Gordon, C.; Wightman, R.; Cherti, M.; Coombes, T.; Katta, A.; Mullis, C.; Wortsman, M.; et al. 2022. Laion-5b: An open large-scale dataset for training next generation image-text models. *Advances in Neural Information Processing Systems*, 35: 25278–25294.
- Song, Y.; Sohl-Dickstein, J.; Kingma, D. P.; Kumar, A.; Ermon, S.; and Poole, B. 2021. Score-Based Generative Modeling through Stochastic Differential Equations. In *International Conference on Learning Representations*.
- Uehara, M.; Zhao, Y.; Black, K.; Hajiramezanali, E.; Scalia, G.; Diamant, N. L.; Tseng, A. M.; Biancalani, T.; and Levine, S. 2024a. Fine-Tuning of Continuous-Time Diffusion Models as Entropy-Regularized Control. *arXiv:2402.15194*.
- Uehara, M.; Zhao, Y.; Black, K.; Hajiramezanali, E.; Scalia, G.; Diamant, N. L.; Tseng, A. M.; Levine, S.; and Biancalani, T. 2024b. Feedback Efficient Online Fine-Tuning of Diffusion Models. *arXiv:2402.16359*.
- Wu, X.; Hao, Y.; Sun, K.; Chen, Y.; Zhu, F.; Zhao, R.; and Li, H. 2023. Human preference score v2: A solid benchmark for evaluating human preferences of text-to-image synthesis. *arXiv preprint arXiv:2306.09341*.
- Xu, J.; Liu, X.; Wu, Y.; Tong, Y.; Li, Q.; Ding, M.; Tang, J.; and Dong, Y. 2024. Imagereward: Learning and evaluating human preferences for text-to-image generation. *Advances in Neural Information Processing Systems*, 36.