

ARTIST: TOWARDS DISENTANGLED TEXT PAINTER WITH DIFFUSION MODELS

Anonymous authors

Paper under double-blind review

ABSTRACT

Diffusion models have shown remarkable performance in generating a broad spectrum of visual content. However, their text rendering ability is still limited: they generate wrong characters or words that cannot blend well with the background image. To address this, we introduce a novel framework named ARTIST, which includes an additional textual diffusion model focusing on text structure learning. We first pre-train the textual diffusion model. Then we further fine-tune a visual model to learn how to inject textual structure information from the frozen textual model into the image. This disentangled architecture design and training strategy significantly enhance the text rendering ability of the diffusion models for text-rich image generation. Furthermore, we leverage pre-trained large-language models to infer the user’s intention leading to better generation quality. Empirical results on the MARIO-Eval benchmark underscore the effectiveness of the proposed method, showing an improvement of up to 15% in various metrics.

1 INTRODUCTION

The field of text-to-image generation has made remarkable progress, especially with the rise of diffusion models (Ho et al., 2020; Rombach et al., 2022; Ramesh et al., 2022; Saharia et al., 2022b; Yang et al., 2022; Zhou et al., 2023; Ruiz et al., 2023). These models not only have demonstrated prowess in generating high-fidelity images, but have also showcased versatile applications, such as image inpainting (Lugmayr et al., 2022; Esser et al., 2021; Jing et al., 2022), denoising (Xie et al., 2023; Yang et al., 2023), video generation (Han et al., 2022; Blattmann et al., 2023), and style transfer (Wang et al., 2023; Zhang et al., 2022).

However, along with these advancements, another realm of importance emerges. Specifically, the generation of images with coherent and readable text is increasingly critical. In daily life, images with embedded text are ubiquitous, from advertisements to road signs, posters, and book covers. Crafting these text-rich images manually requires professional expertise and a considerable investment of time. However, a significant shortcoming lies in the current state-of-the-art models; they often render text portions that are virtually unreadable, akin to gibberish, undermining the aesthetic and functional value of the generated images. Therefore, if AI, powered by diffusion models, can produce such images, it could revolutionize design workflows, inspire creativity, and alleviate designers’ workload.

To tackle the challenges posed by the quality of text generated by diffusion models, two primary pathways has been explored. Firstly, the traditional approach involves leveraging image editing tools to superimpose text onto images directly. However, this frequently introduces unnatural artifacts, especially when contending with intricate textures or varying lighting conditions in the image background. In contrast, recent research efforts aim to refine the diffusion models themselves for improved text quality. For example, recent innovations like Imagen (Saharia et al., 2022b), eDiff-I (Balaji et al., 2022), and DeepFloyd (Shonenkov et al., 2023) discovered that employing T5 series text encoders (Raffel et al., 2020) led to better results compared to using the CLIP text encoder (Radford et al., 2021b). Similarly, Liu et al. (2022) integrated character-aware text encoders to enhance the quality of text rendering. However, these advancements primarily revolve around optimizing text encoders and do not necessarily grant more control over the holistic generation process. On a parallel track, GlyphDraw (Ma et al., 2023) has enhanced model controllability by conditionally focusing on the positioning and architecture of Chinese characters. Still, its utility remains limited,

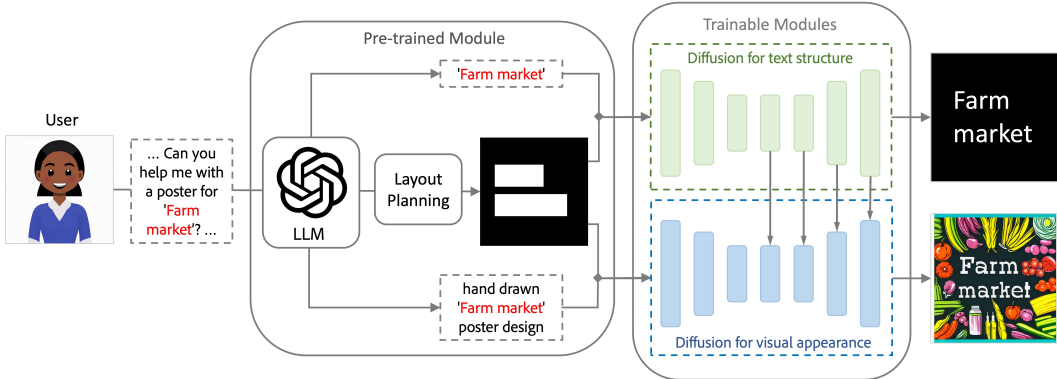


Figure 1: Illustration of the proposed ARTIST framework. A large-language model (LLM) is utilized to analyze the user’s intention. Two diffusion models will be trained to learn text structure and other visual appearance respectively. Given a user input, the LLM will output keywords, layout and text prompts, which will be fed into our trainable modules to generate target images.

as it cannot cater to scenarios that require multiple text bounding boxes, making it less suitable for prevalent text-image formats like posters and book covers. TextDiffuser (Chen et al., 2023) represents a recent advancement in the realm of enhancing the quality of text in generated images. While it undeniably marks a significant step forward, it is not without limitations. One of the prominent challenges is the model’s dependence on manual efforts when it comes to recognizing key terms from prompts, making the process less efficient than desired. Additionally, even though TextDiffuser exhibits improved capabilities, the accuracy rate of its Optical Character Recognition (OCR) evaluation on generated text still leaves room for optimization.

In this paper, our primary objective is to develop a more efficient end-to-end system for text-to-image generation. This system would eliminate the need for manual post-production adjustments, ensuring superior text quality within generated images. Upon meticulous investigation, we have identified the primary challenges we aim to address. Firstly, automation of accurately discerning which words from a provided text prompt should be incorporated into the image. In platforms such as TextDiffuser, this identification requires human involvement, typically by highlighting specific terms with quotation marks, diminishing the platform’s efficiency and automation. Second, there is the challenge of adeptly generating images that seamlessly integrate top-quality text, ensuring adherence to a predetermined layout, and emphasizing in the generated images.

To address the aforementioned challenges, we turned to the latest advancements in large language models (LLMs) (OpenAI, 2023; Touvron et al., 2023a;b). These models have showcased outstanding expertise in understanding and processing intricate linguistic patterns, making them perfectly suited for our needs. Inspired by their prowess, we formulated a strategy that employs large language models to surmount the initial hurdle of pinpointing the keywords. To tackle the subsequent challenge, we introduce a novel two-stage approach named ARTIST, detailed in Figure 1. Specifically, ARTIST’s dual stages are dedicated to mastering text structure and refining visual aesthetics in that order. The moniker “ARTIST” encapsulates the essence of *the Ability of Rendering Text can be Improved by diSentanglemenT*. By merging our proposed methodology with LLMs’ capabilities, we have achieved a marked improvement in the quality of text embedded within images.

Our contributions can be summarized as follows.

- We proposed an efficient end-to-end framework that improves the controllability and quality of text generated within images generated by diffusion models.
- Through combining LLMs and diffusion models, we leveraged the capabilities of large language models like GPT-4 (OpenAI, 2023) to devise a reasoning prompt for understanding complex, even open-domain, user instructions.
- We pioneered the first training strategy that separates learning text structure and visual appearance, boosting performance on existing text rendering benchmarks by up to 15% in terms of text ocr accuracy.

2 RELATED WORK

2.1 TEXT-TO-IMAGE GENERATION

Recent advances in text-to-image generation models can be categorized into two primary categories: autoregressive frameworks (Ramesh et al., 2021; Yu et al., 2022) and diffusion-based models (Ramesh et al., 2022; Saharia et al., 2022b), have been substantial. The latter, particularly diffusion-centric models, have gained significant momentum in recent times. The crux of these models lies in the conversion of textual prompts into latent representations, subsequently relying on diffusion mechanisms to formulate images. Several renowned models, such as Stable Diffusion (Rombach et al., 2022), DALL-E 2 (Ramesh et al., 2022) and Imagen (Saharia et al., 2022b) have set benchmarks in this domain. The challenge of controlled generation remains at the forefront. Contemporary image generation models, with an emphasis on diffusion, heavily prioritize text-based guidance, making it a promising vector for refining their controllability. The inherent complexities of gaining exact control through textual means are evident. For example, ControlNet (Zhang & Agrawala, 2023a) provides an architectural blueprint that adapts pretrained diffusion models to accommodate a variety of input conditions. Although offering a higher level of flexibility, obtaining such condition signals often necessitates manual intervention and might be restrictive for broader conceptual applications. Methods such as GLIGEN (Li et al., 2023) advocate for open-set image generation by employing grounding tokens to define the spatial parameters of objects. We agree with this approach, integrating elements of the grounding token methodology.

2.2 TEXT RENDERING & OPTICAL CHARACTER RECOGNITION

Despite the rapid development of diffusion-based Text-to-image Generation, existing methodologies still struggle with generating precise and consistent textual renderings. Various approaches, like Imagen (Saharia et al., 2022b), eDiff-I (Balaji et al., 2022), and DeepFolyd (Shonenkov et al., 2023), have leveraged the prowess of expansive language models (notably large T5 (Raffel et al., 2020)) to bolster their textual accuracy. The study in (Liu et al., 2022) highlights a limitation in which traditional text encoders overlook token length, prompting them to propose a character-sensitive alternative. Simultaneously, GlyphDraw (Ma et al., 2023) focuses on producing superior images integrated with Chinese texts, guided by textual positioning and glyph imagery. The techniques of Textdiffuser (Chen et al., 2023) harnesses the Transformer model to discern keyword layouts, promoting multiline text generation. Then it further employs character-based segmentation masks as a prior, offering flexibility in control to cater to user specifications. However, it still depends on manual efforts when it comes to recognizing key terms from prompts. From an evaluative standpoint, TextDiffuser introduced the benchmark called MARIO-Eval. This was accomplished by collecting and synthesizing insights from previous works, including the *DrawBench* from Imagen (Saharia et al., 2022b), the *DrawTextCreative* proposed by (Liu et al., 2022), among others.

OCR Optical Character Recognition (OCR) is a long-established academic endeavor (White & Rohrer, 1983; Cash & Hatamian, 1987). In the last decade, this field has seen remarkable progress, impacting a range of applications including recognition of car license plates (Omran & Jarallah, 2017), autonomous vehicle navigation (Schreiber et al., 2014; Wu et al., 2023), and its incorporation into foundational models such as GPT (Huang et al., 2023; Shen et al., 2023). In this paper, OCR serves as a pivotal evaluation metric used to critically assess the quality of the generated text and provide a comprehensive understanding of our model’s performance in realistic scenarios.

3 METHODOLOGY

As revealed in previous works (Ramesh et al., 2022; Saharia et al., 2022a; Chen et al., 2023), generating an image with text rendered on it is still challenging. We suspect that this happens because of two major reasons:

- It is challenging to simultaneously learn visual appearance and text structure with single model;
- Existing datasets are not able to cover all the words and their possible combinations, making it hard to learn text structure from these limited noisy data;

In this work, we propose to mitigate the aforementioned challenges by utilizing separate modules to learn text structure and visual appearance. Furthermore, these two modules are trained separately, making it possible to learn text structure with synthetic data constructed by ourselves, which also tackles the problem of limited data. To ensure disentanglement and prevent information leakage, our text module and visual module take different prompts as inputs. However, it can be inefficient and user-hostile if these prompts have to be manually designed by the user themselves at inference. Fortunately, because of the recent success of large-language models (LLMs) (Brown et al., 2020; OpenAI, 2023; Touvron et al., 2023a;b), we propose to utilize pre-trained LLMs to infer user’s intention, provide accurate prompts and layouts for both modules. With the help of LLM, user’s input can be either precise or vague, leading to better interactive experience.

Our proposed framework is termed ARTIST, because it illustrates that *the Ability of Rendering Text can be Improved by disentanglement*. Our proposed framework is illustrated in Figure 1, with details discussed below. As we show in Section 4, our proposed framework outperforms the previous state-of-the-art (SOTA) in terms of image fidelity, image-prompt alignment, and accuracy of generated texts.

3.1 LLM-BASED PROMPT UNDERSTANDING

TextDiffuser (Chen et al., 2023) propose to train a transformer model to extract s which are expected to be shown on generated image. However, there appears to be a limitation associated with this method: Their model is only capable of detecting words enclosed by quotation marks. This happens because most keywords from training dataset are enclosed by quotation marks inside the captions. A model trained on these samples will lack generalization ability or even be overfitting. Moreover, their model even fails to generate the desired image according to prompt “a poster of Batman with word Batman on it”, as it does not understand that the word “Batman” should be presented on the image.

On the contrary, we propose to provide pre-trained LLM with an open-domain prompt and relies on its capability to autonomously identify the essential s . Because of its vast open-domain knowledge, LLM is able to better understand the user intention and thus can generalize to complicated scenarios. Given vague prompts, LLM is able to discern and propose which words or text elements to incorporate, resulting in more coherent and aesthetically appealing suggestions.

3.2 LEARNING TEXT STRUCTURE

To learn text structures, a diffusion model denoted as text module is introduced. Specifically, this text module is trained to take bounding boxes as input and generates a black-white image with only text on it. Because we do not require it to learn any visual effects, the prompt for this module is the actual word that is desired to be rendered. The consequential benefit is that the dataset can be constructed by ourselves with the help of public Python libraries ¹. Specifically, we construct two large-scale datasets to train this module, as described below.

- Word-level dataset: our first dataset consists of around 10 millions of black-white images with only single word on it. Specifically, to construct each data sample, we first randomly select a word from the vocabulary of CLIP text encoder, then render this word with random font and size on a black image. Meanwhile, we can also obtain the ground-truth bounding box for each rendered word. This dataset will be further augmented during training by randomly moving both word and bounding box to a new position, resulting in infinite number of effective samples;
- Sentence-level dataset: although our word-level dataset contains massive structure information of English words, it contains no layout information about how should these words be placed and combined on the target image. To this end, we construct our second dataset which contains 50 millions of black-white images by utilizing MARIO-10M dataset (Chen et al., 2023). Specifically, we use the ground-truth text and layout information from MARIO-10M samples, and render the same text with randomly selected fonts on black images following the same layout.

¹Pillow 9.5.0 is used in our implementation.

After constructing both datasets, we train the text module in two stages, inside the latent space of a VAE (Kingma & Welling, 2013) following Rombach et al. (2022). In the first stage, the diffusion model is trained to take a single bounding box and a target word as inputs then generate an black-white image with the word rendered on it. Then, this module will be fine-tuned on our sentence-level dataset as the second stage so that it can take multiple bounding boxes and multiple s as inputs. Our s are denoted as $\mathcal{P} = \{p_1, p_2, \dots, p_n\}$, composed of the words to be generated in the image. Our input mask is an image contains the corresponding bounding boxes $\{m_1, m_2, \dots, m_n\}$ that indicate the position of each word. During training, the original text-only image and input mask image are first encoded into latent space features \mathbf{z} and \mathbf{m} . Then we sample a time step $t \sim \text{Uniform}(0, T)$ and a Gaussian noise ϵ to corrupt the original feature, yielding $\mathbf{z}_t = \sqrt{\bar{\alpha}_t}\mathbf{z} + \sqrt{1 - \bar{\alpha}_t}\epsilon$ where $\bar{\alpha}_T$ is the coefficient of the diffusion process introduced in (Ho et al., 2020). \mathbf{z}_t and \mathbf{m} are concatenated in the feature channel as input for diffusion model, which will be trained with the diffusion loss between the sampled ϵ and the predicated noise ϵ_θ :

$$\mathcal{L}_{\text{text}} = \mathbb{E} \left[\|\epsilon - \epsilon_\theta(\mathbf{z}_t, \mathbf{m}, \mathcal{P}, t)\|_2^2 \right]. \quad (1)$$

Some generated examples are shown in Figure 2, which illustrates that our text module can successfully generate a target image with desired texts on it in different styles.

3.3 LEARNING VISUAL APPEARANCE

After training the text module, we would like to utilize its learned knowledge to generate high-fidelity images containing text. To this end, we propose to inject intermediate features from our text module into our visual module, which is also a diffusion model. For each intermediate feature from the mid-block and up-block layers of text module, we propose to use a trainable convolutional layer to project the feature and add it element-wisely onto the corresponding intermediate output feature of the visual module. We have also tested different architectures, the comparison will be provided in later experiment section.

During the training of visual module, the text module will be frozen, and only the newly-introduced layers and visual module will be fine-tuned. Differently from the text module described in Section 3.2, whose training text only contains the target words to be rendered, the training data for our visual module have to contain visual descriptions of the image so that the model can successfully learn to generate visual contents based on the user’s input.

Let $\bar{\mathcal{P}}$ be the prompt which contains visual descriptions of the image, \mathcal{P} be the s as defined in Section 3.2. We feed $\bar{\mathcal{P}}$ and the input mask \mathbf{m} as mentioned in Section 3.2 into the pre-trained text module. The intermediate features from text module, denoted as $\{f_i(\mathbf{z}_t, \mathbf{m}, \mathcal{P}, t)\}_{i=1}^k$ will be injected into the visual module, where k stands for the number of intermediate features. The visual module will be trained with diffusion loss:

$$\mathcal{L}_{\text{visual}} = \mathbb{E} \left[\|\epsilon - \epsilon_\theta(\mathbf{x}_t, \{f_i(\mathbf{z}_t, \mathbf{m}, \mathcal{P}, t)\}_{i=1}^k, \mathbf{m}, \bar{\mathcal{P}}, t)\|_2^2 \right], \quad (2)$$

where $\mathbf{x}_t = \sqrt{\bar{\alpha}_t}\mathbf{x} + \sqrt{1 - \bar{\alpha}_t}\epsilon$ denotes the corrupted VAE feature of ground-truth image.

In our implementation, MARIO-10M dataset (Chen et al., 2023) is used to train this module following previous work for a fair comparison. $\bar{\mathcal{P}}, \mathcal{P}, \mathbf{m}$ have already been prepared in MARIO-10M. At inference time, LLM will be utilized to infer $\bar{\mathcal{P}}, \mathcal{P}, \mathbf{m}$ as mentioned in Section 3.1.

4 EXPERIMENTS

4.1 IMPLEMENTATION DETAILS

Our experiments are conducted on 8 Nvidia A100 GPUs, with Hugging Face Diffusers (von Platen et al., 2022). Both text and visual modules are initialized from the pre-trained Stable Diffusion

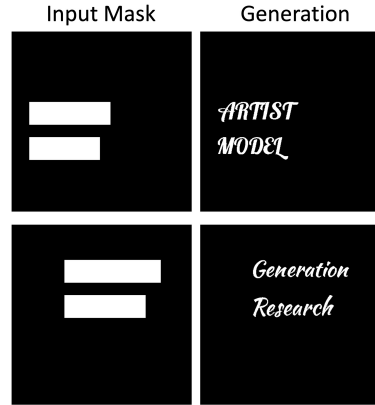


Figure 2: Generated examples from our text module, along with input masks.



Figure 3: Generated examples from the proposed ARTIST, along with the input prompts presented above the corresponding images.

(Rombach et al., 2021) checkpoint. The text module is first pre-trained on our word-level synthetic dataset for 400,000 steps, then further fine-tuned for 200,000 steps on our sentence-level synthetic dataset. With the frozen text module, our visual module is trained for 250,000 steps on MARIO-10M dataset (Chen et al., 2023). AdamW (Loshchilov & Hutter, 2017) optimizer is used in all training stages, with a learning rate of $1e-5$ and a weight decay of $1e-2$. The batch size is set to 128. Our code for training and constructing dataset will be publicly available upon acceptance.

Character-level segmentation masks are used in Chen et al. (2023), which has been proven to be very effective as it provides a strong signal to the diffusion model. Thus, we also utilize character-level segmentation masks in our implementation. But we also conduct experiments without the character-level segmentation mask, showing that the proposed framework can lead to improvements in both settings.

We compare our proposed framework with previous methods including Stable Diffusion (Rombach et al., 2021) (denoted as SD), ControlNet (Zhang & Agrawala, 2023b), DeepFloyd (Shonnikov et al., 2023) and TextDiffuser (Chen et al., 2023). We also fine-tune a Stable Diffusion on MARIO-10M dataset for a more comprehensive comparison. Two variants of our proposed framework are evaluated, which are denoted as ARTIST-TD and ARTIST respectively, indicating whether LLM is utilized. Specifically, ARTIST-TD directly uses the pre-trained transformer from TextDiffuser instead of LLM to obtain bounding boxes and keywords based on input prompts. Thus comparing ARTIST-TD with TextDiffuser can straightforwardly show the effectiveness of our training strategy, as they share exactly the same layout and keyword conditions.

Note that although there are two separate diffusion models in our framework, our computation requirement is still similar to the previous SOTA TextDiffuser. This is because TextDiffuser also utilizes an extra U-Net which is designed for character-aware loss as a regularization term.

4.2 MAIN RESULTS

MARIO-Eval benchmark To start with, we conduct experiments on MARIO-Eval benchmark proposed in (Chen et al., 2023), which contains 5,414 prompts in total. 4 images are generated for each prompt to compute the CLIP score (Radford et al., 2021a), Fréchet Inception Distance (FID) (Heusel et al., 2017) and OCR evaluations. Specifically, CLIP score is obtained by calculating the cosine similarity between generated images and prompts by using the features extracted with pre-trained ViT/B-32 CLIP model. OCR evaluation is performed with MaskSpotterv3 (Liao et al., 2020)

Table 1: Results on MARIO-Eval benchmark, our ARTIST outperforms previous methods.

Metrics	SD	Fine-tuned SD	ControlNet	DeepFloyd	TextDiffuser	ARTIST-TD	ARTIST
FID (↓)	51.295	28.761	51.485	34.902	38.758	36.579	38.43
CLIP Score (↑)	0.3015	0.3412	0.3424	0.3267	0.3436	0.3466	0.3482
OCR Accuracy (↑)	0.0178	0.0154	0.2705	0.0457	0.5712	0.6298	0.7373
OCR Precision (↑)	0.0192	0.1777	0.5391	0.1738	0.7795	0.8237	0.8681
OCR Recall (↑)	0.0260	0.2330	0.6438	0.2235	0.7498	0.7986	0.8677
OCR F-measure (↑)	0.0221	0.2016	0.5868	0.1955	0.7643	0.8110	0.8679



Figure 4: Comparison with TextDiffuser on MARIO-Eval benchmark. Layout generated by TextDiffuser is used as input conditions for both models for fair comparison.

following (Chen et al., 2023). The main results are presented in Table 1, from which we can see that our ARTIST outperforms the previous SOTA TextDiffuser in all metrics, even without the help of LLM. Some generated examples are provided in Figure 3, more results are provided in the Appendix because of the limited space.

For a more straightforward comparison, we also provide some generated examples in Figure 4. Although both TextDiffuser and ARTIST-TD are given the same layout and keyword conditions, we can see that ARTIST-TD generates images with better harmonization between painted text and background. ARTIST-TD is also able to generate many visual features that correspond to keywords such as “Green Lantern” and “transparent” in the prompt. This is because the learning of text structure and visual appearance is better disentangled, which leads to more efficient learning of both aspects compared to TextDiffuser.

ARTIST benchmark As the reader may notice from Figure 3 and 4, most keywords in MARIO-Eval prompts are enclosed by quotation marks. In other words, a model can easily obtain promising accuracy in keyword prediction even if it is over-fitting and simply extracts words enclosed by quotation marks. Thus MARIO-Eval benchmark is not a reasonable benchmark to justify model performance on open-domain instructions. To this end, we propose a new ARTIST benchmark, which contains 500 pairs of prompts and keywords. Details about constructing the benchmark and some examples are provided in the Appendix A. The keywords in ARTIST benchmark, which are expected to be shown in the resulting images, are designed to be contained in the prompts so that we can compare all the methods fairly. However, keywords are not always enclosed by quotation marks. Similar to MARIO-Eval benchmark, 4 images are generated for each prompt to calculate CLIP score and OCR evaluations. The results on ARTIST benchmark are presented in Table 2. From the results we can see that both ARTIST-TD and ARTIST outperform related methods, while ARTIST obtains a huge performance boost, indicating the benefits of using LLM.

Keywords Identification To quantitatively evaluate the contribution of LLM, we conduct an experiment of keywords identification. Different models are asked to extract keywords from prompts of ARTIST benchmark. The results are shown in Table 3, where we can conclude that LLM indeed leads to huge improvement in detecting keywords.

Table 2: Results on ARTIST benchmark, our proposed framework outperforms all previous methods.

Metrics	SD	Fine-tuned SD	ControlNet	DeepFloyd	TextDiffuser	ARTIST-TD	ARTIST
CLIP Score (↑)	0.3387	0.3440	0.3105	0.3419	0.3100	0.3225	0.3545
OCR Accuracy (↑)	0.0065	0.0300	0.0830	0.0545	0.1345	0.2185	0.6530
OCR Precision (↑)	0.1031	0.1931	0.1885	0.2225	0.2160	0.2850	0.8166
OCR Recall (↑)	0.1213	0.2352	0.2304	0.2669	0.1994	0.2779	0.8090
OCR F-measure (↑)	0.1114	0.2121	0.2074	0.2427	0.2073	0.2814	0.8128

Table 3: Large-language models improves keywords identification by a large margin.

Models	keywords Identification Evaluation			
	Accuracy (↑)	Precision (↑)	Recall (↑)	F-measure (↑)
TextDiffuser	0.6320	0.6412	0.6397	0.6404
GPT-3.5-Turbo	0.8300	0.9582	0.9496	0.9539
GPT-4	0.9380	0.9729	0.9893	0.9810




Figure 5: Generated examples in inpainting task, where the masked regions are indicated by red rectangles. Prompts used in these examples are “a book cover of Bed Times”, “a book cover for Kansas State”, “a poster for A Good Princess” and “a poster for The Man Who Said YES”.

Image inpainting Similar to TextDiffuser, the proposed ARTIST can also be trained to perform an image-inpainting task. We trained another ARTIST following the setting from Chen et al. (2023) so that the model is asked to perform image inpainting instead of image generation with a probability of 0.5 during training. The resulting ARTIST is able to perform both tasks at inference time. Some generated examples are provided in Figure 5, along with the comparison between ARTIST and TextDiffuser. Similarly to the whole-image generation task, our ARTIST leads to more accurate rendered texts and more harmonized images.

4.3 ABLATION STUDIES

Robustness Although character-level segmentation mask is used in our implementation following (Chen et al., 2023), the proposed ARTIST should lead to improvements no matter whether character-level segmentation mask is used or not. To verify this claim, we conduct an ablation study to show the robustness of ARTIST. Specifically, we trained an ARTIST model without character-level segmentation mask, which is to be compared with the fine-tuned Stable Diffusion model. For all the models, we apply 50 sampling steps and classifier-free guidance (Ho & Salimans, 2022) of $\omega \in \{2.5, 5, 7.5, 10, 15\}$ during inference. In general, a larger ω leads to better FID and CLIP score, while it may lead to a worse OCR evaluation. The comparison is provided in Figure 6, from which we can conclude that our ARTIST leads to improvement across different settings, because it leads to better alignment of the image and the text, image fidelity, and OCR accuracy compared to baselines.

Network architecture We also conduct an experiment in which a ControlNet-like architecture is used to introduce the learned features into the visual module. Specifically, intermediate features of down-blocks and mid-block from text module’s U-Net are projected into visual module. The




Figure 6: Ablation study of comparing proposed ARTIST with baseline models. The proposed framework is able to generate more accurate text and obtain better image-text similarity across different classifier-free guidance scales.




Figure 7: The generated text will follow the layout information in terms of mask positions and sizes. Prompts used here are “a neon light ‘ARTIST’ on the brick wall” and “word ‘ARTIST’ surrounded by hand-drawn flowers”.

Table 4: Intermediate features from text module’s U-Net decoder lead to better performance.

Features from	FID (\downarrow)	CLIP Score (\uparrow)	OCR Evaluation			
			Accuracy (\uparrow)	Precision (\uparrow)	Recall (\uparrow)	F-measure (\uparrow)
U-Net Encoder	37.454	0.3455	0.4935	0.7312	0.7293	0.7302
U-Net Decoder	36.579	0.3466	0.6298	0.8237	0.7986	0.8110

experiment is conducted on MARIO-Eval benchmark so that FID, CLIP, OCR results can all be reported. LLM is not used. The comparison is provided in Table 4, from which we can see that the intermediate features from text module’s U-Net decoder leads to better performance in this task.

Layout information We conduct ablation study to investigate the importance of layout information during inference. Some generated examples are presented in Figure 7, from which we can see that the generated text will always follow the provided position and size as expected.

5 CONCLUSIONS

In this paper, we proposed ARTIST, a novel framework that significantly enhances the text-rendering ability of diffusion models in text-rich image generation. Our proposed framework utilizes pre-trained large-language models to infer the user’s intention, provide accurate prompts, and improve the interactive experience. We also introduced a disentangled architecture design and training strategy which leads to better learning of both text structure and visual appearance. Our experimental results demonstrate that ARTIST outperforms the previous state-of-the-art in terms of image fidelity, image-prompt alignment, and accuracy of generated texts. In the future, we aim to further improve the disentangled learning, investigate the interpretability of the disentangled representations learned by ARTIST and their potential for downstream tasks. Overall, we believe that ARTIST represents a significant step forward in the field of text-rich image generation and has the potential to enable new applications and use cases in the future.

REFERENCES

- Yogesh Balaji, Seungjun Nah, Xun Huang, Arash Vahdat, Jiaming Song, Karsten Kreis, Miika Aittala, Timo Aila, Samuli Laine, Bryan Catanzaro, et al. ediffi: Text-to-image diffusion models with an ensemble of expert denoisers. *arXiv preprint arXiv:2211.01324*, 2022.
- Andreas Blattmann, Robin Rombach, Huan Ling, Tim Dockhorn, Seung Wook Kim, Sanja Fidler, and Karsten Kreis. Align your latents: High-resolution video synthesis with latent diffusion models. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- Glenn L Cash and Mehdi Hatamian. Optical character recognition by the method of moments. *Computer vision, graphics, and image processing*, 39(3):291–310, 1987.
- Jingye Chen, Yupan Huang, Tengchao Lv, Lei Cui, Qifeng Chen, and Furu Wei. Textdiffuser: Diffusion models as text painters. *arXiv preprint arXiv:2305.10855*, 2023.
- Patrick Esser, Robin Rombach, Andreas Blattmann, and Bjorn Ommer. Imagebart: Bidirectional context with multinomial diffusion for autoregressive image synthesis. *Advances in neural information processing systems*, 34:3518–3532, 2021.
- Xizewen Han, Huangjie Zheng, and Mingyuan Zhou. Card: Classification and regression diffusion models. *Advances in Neural Information Processing Systems*, 35:18100–18115, 2022.
- Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017.
- Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- Shaohan Huang, Li Dong, Wenhui Wang, Yaru Hao, Saksham Singhal, Shuming Ma, Tengchao Lv, Lei Cui, Owais Khan Mohammed, Qiang Liu, et al. Language is not all you need: Aligning perception with language models. *arXiv preprint arXiv:2302.14045*, 2023.
- Bowen Jing, Gabriele Corso, Renato Berlinghieri, and Tommi Jaakkola. Subspace diffusion generative models. In *European Conference on Computer Vision*, pp. 274–289. Springer, 2022.
- Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- Yuheng Li, Haotian Liu, Qingyang Wu, Fangzhou Mu, Jianwei Yang, Jianfeng Gao, Chunyuan Li, and Yong Jae Lee. Gligen: Open-set grounded text-to-image generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 22511–22521, 2023.
- Minghui Liao, Guan Pang, Jing Huang, Tal Hassner, and Xiang Bai. Mask textspotter v3: Segmentation proposal network for robust scene text spotting. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XI 16*, pp. 706–722. Springer, 2020.
- Rosanne Liu, Dan Garrette, Chitwan Saharia, William Chan, Adam Roberts, Sharan Narang, Irina Blok, RJ Mical, Mohammad Norouzi, and Noah Constant. Character-aware models improve visual text rendering. *arXiv preprint arXiv:2212.10562*, 2022.
- Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.

- Andreas Lugmayr, Martin Danelljan, Andres Romero, Fisher Yu, Radu Timofte, and Luc Van Gool. Repaint: Inpainting using denoising diffusion probabilistic models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11461–11471, 2022.
- Jian Ma, Mingjun Zhao, Chen Chen, Ruichen Wang, Di Niu, Haonan Lu, and Xiaodong Lin. Glyph-draw: Learning to draw chinese characters in image synthesis models coherently, 2023.
- Safaa S. Omran and Jumana A. Jarallah. Iraqi car license plate recognition using ocr. In *2017 Annual Conference on New Trends in Information & Communications Technology Applications (NTICT)*, pp. 298–303, 2017. doi: 10.1109/NTICT.2017.7976127.
- OpenAI. Gpt-4 technical report. *ArXiv*, abs/2303.08774, 2023. URL <https://api.semanticscholar.org/CorpusID:257532815>.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pp. 8748–8763. PMLR, 2021a.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pp. 8748–8763. PMLR, 2021b.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *The Journal of Machine Learning Research*, 21(1):5485–5551, 2020.
- Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *International Conference on Machine Learning*, pp. 8821–8831. PMLR, 2021.
- Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 1(2):3, 2022.
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models, 2021.
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 10684–10695, 2022.
- Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. 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*, pp. 22500–22510, 2023.
- Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in Neural Information Processing Systems*, 35:36479–36494, 2022a.
- Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in Neural Information Processing Systems*, 35:36479–36494, 2022b.
- Markus Schreiber, Fabian Poggenhans, and Christoph Stiller. Detecting symbols on road surface for mapping and localization using ocr. In *17th International IEEE Conference on Intelligent Transportation Systems (ITSC)*, pp. 597–602, 2014. doi: 10.1109/ITSC.2014.6957755.
- Yongliang Shen, Kaitao Song, Xu Tan, Dongsheng Li, Weiming Lu, and Yueting Zhuang. Hugginggpt: Solving ai tasks with chatgpt and its friends in huggingface. *arXiv preprint arXiv:2303.17580*, 2023.

- Alex Shonenkov, Misha Konstantinov, Daria Bakshandaeva, Christoph Schuhmann, Ksenia Ivanova, and Nadiia Klokova. Deepfloyd. <https://github.com/deep-floyd/if>, 2023.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023a.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023b.
- Patrick von Platen, Suraj Patil, Anton Lozhkov, Pedro Cuenca, Nathan Lambert, Kashif Rasul, Mishig Davaadorj, and Thomas Wolf. Diffusers: State-of-the-art diffusion models. <https://github.com/huggingface/diffusers>, 2022.
- Zhizhong Wang, Lei Zhao, and Wei Xing. Stylediffusion: Controllable disentangled style transfer via diffusion models. *arXiv preprint arXiv:2308.07863*, 2023.
- James M White and Gene D Rohrer. Image thresholding for optical character recognition and other applications requiring character image extraction. *IBM Journal of research and development*, 27(4):400–411, 1983.
- Zizhang Wu, Xinyuan Chen, Jizheng Wang, Xiaoquan Wang, Yuanzhu Gan, Muqing Fang, and Tianhao Xu. Ocr-rtps: an ocr-based real-time positioning system for the valet parking. *Applied Intelligence*, pp. 1–15, 2023.
- Yutong Xie, Minne Yuan, Bin Dong, and Quanzheng Li. Diffusion model for generative image denoising. *arXiv preprint arXiv:2302.02398*, 2023.
- Chenghan Yang, Lijing Liang, and Zhixun Su. Real-world denoising via diffusion model. *ArXiv*, abs/2305.04457, 2023. URL <https://api.semanticscholar.org/CorpusID:258557313>.
- Ling Yang, Zhilong Zhang, Yang Song, Shenda Hong, Runsheng Xu, Yue Zhao, Yingxia Shao, Wentao Zhang, Bin Cui, and Ming-Hsuan Yang. Diffusion models: A comprehensive survey of methods and applications. *arXiv preprint arXiv:2209.00796*, 2022.
- Jiahui Yu, Yuanzhong Xu, Jing Yu Koh, Thang Luong, Gunjan Baid, Zirui Wang, Vijay Vasudevan, Alexander Ku, Yinfei Yang, Burcu Karagol Ayan, et al. Scaling autoregressive models for content-rich text-to-image generation. *arXiv preprint arXiv:2206.10789*, 2(3):5, 2022.
- Lvmin Zhang and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. *arXiv preprint arXiv:2302.05543*, 2023a.
- Lvmin Zhang and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. *arXiv preprint arXiv:2302.05543*, 2023b.
- Yuxin Zhang, Nisha Huang, Fan Tang, Haibin Huang, Chongyang Ma, Weiming Dong, and Changsheng Xu. Inversion-based creativity transfer with diffusion models. *arXiv preprint arXiv:2211.13203*, 2022.
- Yufan Zhou, Bingchen Liu, Yizhe Zhu, Xiao Yang, Changyou Chen, and Jinhui Xu. Shifted diffusion for text-to-image generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10157–10166, 2023.

A ARTIST BENCHMARK

Here we provide some examples in our ARTIST benchmark:

- A vintage movie poster for Forrest Gump
- A modern movie poster for ‘Batman’
- A colorful book cover for “Iron Man”
- A minimalist movie poster for The Godfather
- An abstract movie poster for ‘Pulp Fiction’
- A gothic book cover for “Dracula”
- A romantic movie poster for The Notebook
- A futuristic movie poster for ‘Blade Runner’
- A watercolor book cover for “Pride and Prejudice”
- A playful movie poster for Finding Nemo

Our benchmark is constructed by prompting GPT-4 (OpenAI, 2023), the prompt we used is shown in Figure 8.

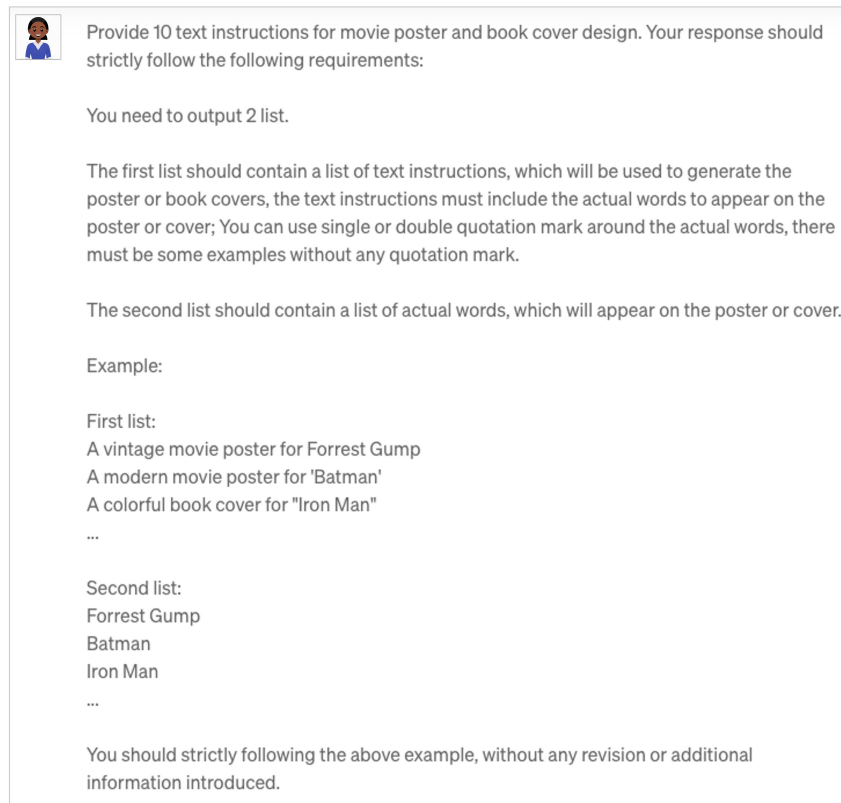


Figure 8: Prompt we used to construct our benchmark with GPT-4.