

# AutoPP: Towards Automated Product Poster Generation and Optimization

Jiahao Fan\*, Yuxin Qin\*, Wei Feng<sup>†</sup>, Yanyin Chen, Yaoyu Li, Ao Ma, Yixiu Li, Li Zhuang, Haoyi Bian, Zheng Zhang, Jingjing Lv, Junjie Shen, Ching Law

JD.COM

{fanjiahao5,qinyuxin11,fengwei25,chenyanyin6,liyaoyu1,mao.8,liyixiu5,zhuangli10,bianhaoyi,zhangzheng11,lvjingjing1,shenjunjie,lawching}@jd.com

## Abstract

Product posters blend striking visuals with informative text to highlight the product and capture customer attention. However, crafting appealing posters and manually optimizing them based on online performance is laborious and resource-consuming. To address this, we introduce AutoPP, an automated pipeline for product poster generation and optimization that eliminates the need for human intervention. Specifically, the generator, relying solely on basic product information, first uses a unified design module to integrate the three key elements of a poster (background, text, and layout) into a cohesive output. Then, an element rendering module encodes these elements into condition tokens, efficiently and controllably generating the product poster. Based on the generated poster, the optimizer enhances its Click-Through Rate (CTR) by leveraging online feedback. It systematically replaces elements to gather fine-grained CTR comparisons and utilizes Isolated Direct Preference Optimization (IDPO) to attribute CTR gains to isolated elements. Our work is supported by AutoPP1M, the largest dataset specifically designed for product poster generation and optimization, which contains one million high-quality posters and feedback collected from over one million users. Experiments demonstrate that AutoPP achieves state-of-the-art results in both offline and online settings.

**Code&Dataset** — <https://github.com/JD-GenX/AutoPP>

## Introduction

Product posters play a crucial role in capturing user interest by artfully combining products, textual content, and visual backgrounds. However, creating posters manually offline and optimizing them through iterative analysis of online performance is both time-consuming and resource-intensive. This challenge has spurred growing interest in automated solutions to streamline both the generation and optimization processes.

With the advancement of text-to-image generation (Romach et al. 2022; Peebles and Xie 2023; Zhang, Rao, and Agrawala 2023), significant progress has been made in automated product poster generation (Gao et al. 2025; Wang

et al. 2024; Jia et al. 2023; Chen et al. 2025a). However, they still require additional manual input at certain stages as shown in Figure 1 (a) and (b). PAID (Chen et al. 2025a) employs a four-stage pipeline where three separate models sequentially generate prompts, layouts, and background images. Then, for text rendering, it uses manually defined rules to determine text attributes such as font and color, which inevitably restricts automation and often disrupts the visual harmony between text and background. PosterMaker (Gao et al. 2025) overcomes this issue by utilizing Stable Diffusion 3 (SD3) (Esser et al. 2024) with background and text ControlNets (Zhang, Rao, and Agrawala 2023), which automatically renders the text along with background generation. Nevertheless, it requires users to individually craft pleasing layouts and compelling selling points for each poster, making it inefficient for generating high-quality posters at scale. To address these issues, we propose an automated product poster generator as illustrated in Figure 1 (c), which requires only basic product information (a product image and candidate text) as input. Our model integrates a unified design module to jointly reason about three key poster elements (background, text, and layout). This comprehensive approach eliminates the fragmentation of task-specific models while ensuring design consistency. We further develop an element rendering module to jointly render the background and text according to the predicted layout. It encodes product images and text into condition tokens, then implements a decomposed attention mechanism with condition self-attention for intra-element modeling and image-condition cross-attention for generation guidance. Our approach reduces the computational overhead along with long token sequences while maintaining robust generation quality.

Despite the automation challenges in offline generation, existing methods also suffer from similar issues in online optimization as shown in Figure 1, often requiring manual adjustment of generation strategies based on empirical observations. Recently, CG4CTR (Yang et al. 2024) and CAIG (Chen et al. 2025b) have made initial progress by optimizing poster backgrounds for the online metric, like Click-Through-Rate (CTR). However, they represent only a partial solution, since CTR is also heavily influenced by textual content and layout design. A straightforward solution is to extend these methods to optimize background, text,

\*These authors contributed equally.

<sup>†</sup>Project Leader and Corresponding Author

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

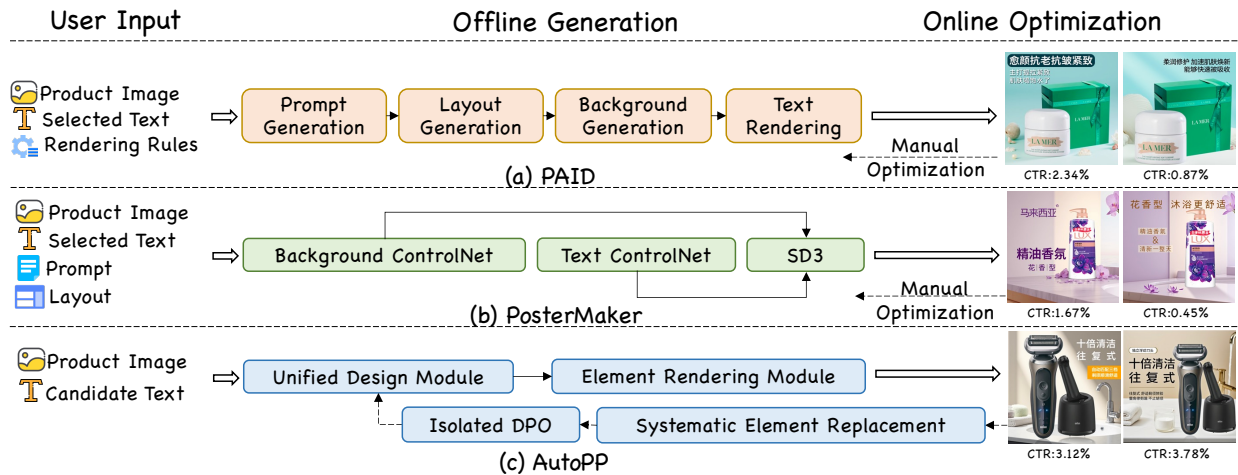


Figure 1: Comparison of different poster generation and optimization pipelines. AutoPP only relies on basic product information and can optimize the generated content automatically based on CTR feedback.

and layout jointly, but this coarse-grained approach cannot indicate which poster element actually drives the CTR improvement, leading to sub-optimal results and low learning efficiency. To address these limitations, we propose a product poster optimizer that automatically refines the content of posters through CTR feedback instead of manual modification. Together with the proposed product poster generator, these two components form *AutoPP*, an automated pipeline that generates and optimizes high-quality product posters for maximum CTR. We first collect CTR feedback on generated sample pairs, each constructed by systematically replacing one element while keeping others constant. Leveraging this fine-grained feedback, we propose Isolated Direct Preference Optimization (IDPO) that attributes CTR improvement to isolated elements. This decomposition enables precise attribution of performance improvements to specific elements, allowing targeted optimization. To support both generation and optimization tasks, we introduce *AutoPPIM*, the largest dataset of product posters, containing one million professionally designed posters and CTR feedback from over one million users. This comprehensive resource enables data-driven model training while providing the necessary performance signals for our optimization framework. To the best of our knowledge, we are the first to unify automated product poster generation and optimization within a single framework.

In summary, our main contributions are as follows:

- We propose an automated pipeline *AutoPP* for product poster generation and optimization. It contains a generator that produces high-quality posters with basic product information, eliminating the need for complex manual input.
- We introduce an automated optimizer to improve the CTR of generated posters without manual adjustments. With a systematic element replacement strategy and IDPO, we enable element-isolated CTR optimization for posters through controlled pairwise comparisons, over-

coming the attribution limitations of holistic approaches.

- We present *AutoPPIM*, the largest dataset of product posters, which facilitates new research in automated product poster generation and online performance optimization. Experiments show that our approach achieves state-of-the-art results in both offline and online settings.

## Related Work

### Product Poster Generation

The task of product poster generation aims to automatically generate visually appealing images that highlight both products and textual content. Existing approaches to product poster generation still struggle to achieve automated production. Recent methods (Lin et al. 2023; Jia et al. 2023; Wang et al. 2024; Jin et al. 2022; Liang et al. 2024a) rely on complicated processes like layout generation (Lu et al. 2025; Li et al. 2023a), attribute prediction, slogan creation, and rendering, and require manual tuning or domain-specific rules. While P&R (Li et al. 2023b) and PAID (Chen et al. 2025a) advance automation by learning to generate layouts directly from product metadata, they critically depend on hand-crafted rules for text rendering.

With the advent of diffusion-based text rendering methods (Ma et al. 2025; Tuo, Geng, and Bo 2024; Liu et al. 2024), poster text has become more coherent and diverse. TextDiffuser-2 (Chen et al. 2024) utilizes GlyphControl (Yang et al. 2023) and a language model within the diffusion model to render texts. While these methods have largely unified the poster generation process, they often fail to maintain product appearance consistency. Recently, PosterMaker (Gao et al. 2025) has been developed to incorporate product information, featuring background and text ControlNets (Zhang, Rao, and Agrawala 2023). This enables the simultaneous generation of backgrounds and text while preserving product authenticity. However, it still requires users to provide predefined layouts and compelling selling points. Our research aims to achieve fully automated prod-

uct poster generation, eliminating additional manual input and enhancing efficiency for large-scale creation and rapid deployment.

## Product Poster Optimization

As a subtask of image generation, product posters optimization primarily targets offline evaluation metrics (Wang et al. 2025a), particularly visual aesthetics. ImageReward (Xu et al. 2023), PickScore (Kirstain et al. 2023), and HPSv2 (Wu et al. 2023) leverage human preference data by training reward models on curated text-image pairs, effectively aligning model outputs with human aesthetic judgments. Du et al. (Du et al. 2024) introduce the multimodal Reliable Feedback Network (RFNet) to improve the availability of generated advertising images. Most recently, UnifiedRewardThink (Wang et al. 2025b) integrates chain-of-thought reasoning into reward models, enabling the internalization of reasoning processes and proposing the first unified reward model based on multimodal chains of thought.

Since these offline metrics often fail to align with actual user engagement metrics, recent approaches, like CG4CTR (Yang et al. 2024) and CAIG (Chen et al. 2025b), optimize for CTR to bridge this gap. CG4CTR develops a CTR-optimized reward model that jointly leverages multimodal features to accurately reflect user click preferences. CAIG advances the reward modeling architecture to a Multimodal Large Language Model (MLLM), enabling a deeper understanding of cross-modal preferences through semantic alignment. However, they only enhance CTR by optimizing poster backgrounds. In contrast, we investigate how multiple elements in poster generation affect CTR and propose IDPO: an optimization framework that decomposes CTR optimization into isolated poster elements.

## Dataset

The AutoPP1M dataset comprises one million product posters, divided into two subsets. The first subset, focused on poster generation, features posters with detailed annotations on visual elements. The second subset, dedicated to poster optimization, includes posters paired with high-quality preference labels.

## Collection and Annotation

For the product poster generation subset, we collect product posters from JD.COM, with all posters in a 1:1 aspect ratio and at least  $800 \times 800$  pixels in size. Starting from an initial candidate pool of approximately 10 million posters, we apply standard data cleaning procedures, including aesthetic filtering, blur detection, and watermark removal (Li et al. 2024; Gong et al. 2025), to eliminate low-quality posters, yielding roughly 2 million posters. Subsequently, we employ PaddleOCRv4 (Li et al. 2022) to remove posters containing excessive text or no textual content, resulting in a final refined set of one million posters. This final dataset covers a diverse range of everyday consumer products, spanning over 60 coarse-grained categories. To extract the elements from the poster, we first use Grounding SAM (Ren et al. 2024) and PaddleOCRv4 to obtain product masks, text content,

and text bounding boxes. Then, we generate the background prompt using Qwen2.5-VL (Bai et al. 2025), supplemented with product attribute information (e.g., product names and titles) to mitigate hallucination.

For the product poster optimization subset, we collect CTR feedback for product posters through a 10-day randomized display experiment within JD.COM. To ensure data reliability, each poster is viewed by at least 50 users, accumulating 1,118,140 users. Following the systematic element replacement strategy (detailed in Method), we construct 50,000 pairwise comparisons by matching posters of the same product. Additionally, we ensure a minimum relative CTR difference of 1% between paired variants to ensure distinguishable CTR differences.

## Characteristic

- **High Quality:** The generation subset comprises top-viewed product posters on the platform, reflecting well-designed posters in which merchants invested significant resources. For the optimization subset, the dataset benefits from extensive user engagement on JD.COM, with over one million users, ensuring a high level of diversity and reliability in user preference data.
- **Large Scale:** Existing product poster datasets, especially those focusing on Chinese product posters, are often limited in scale. With one million elegant product posters, AutoPP1M is an order of magnitude larger than existing datasets (Zhou et al. 2022; Lin et al. 2023; Hsu et al. 2023; Gao et al. 2025), enabling more robust training of modern generative models.
- **New Field:** Unlike existing human-feedback image datasets, which typically focus on pure scene images without text (Liang et al. 2024b; Xu et al. 2023), AutoPP1M uniquely addresses the significant impact of textual content on user attraction. By providing separate assessments of human preferences for background, text, and layout in posters, we enable more nuanced, element-level modeling of these preferences. This facilitates new research directions in fine-grained preference learning.

## Method

### Overview

In this work, we propose AutoPP, an automated pipeline for product poster generation and optimization. As illustrated in Figure 2, given the basic product information, the generator first generates three key poster elements through the unified design module. Then, the element rendering module synthesizes these elements into the product poster. The optimizer further systematically replaces the isolated elements in the generated poster. Finally, we employ isolated direct preference optimization to improve CTR performance.

### Automated Product Poster Generator

**Unified Design Module** Our unified design module generates coherent poster elements (background prompt, text, and layout) through a single integrated process as shown in Figure 2 (a). Unlike conventional approaches that rely on separate models for each element, often leading to inconsistent

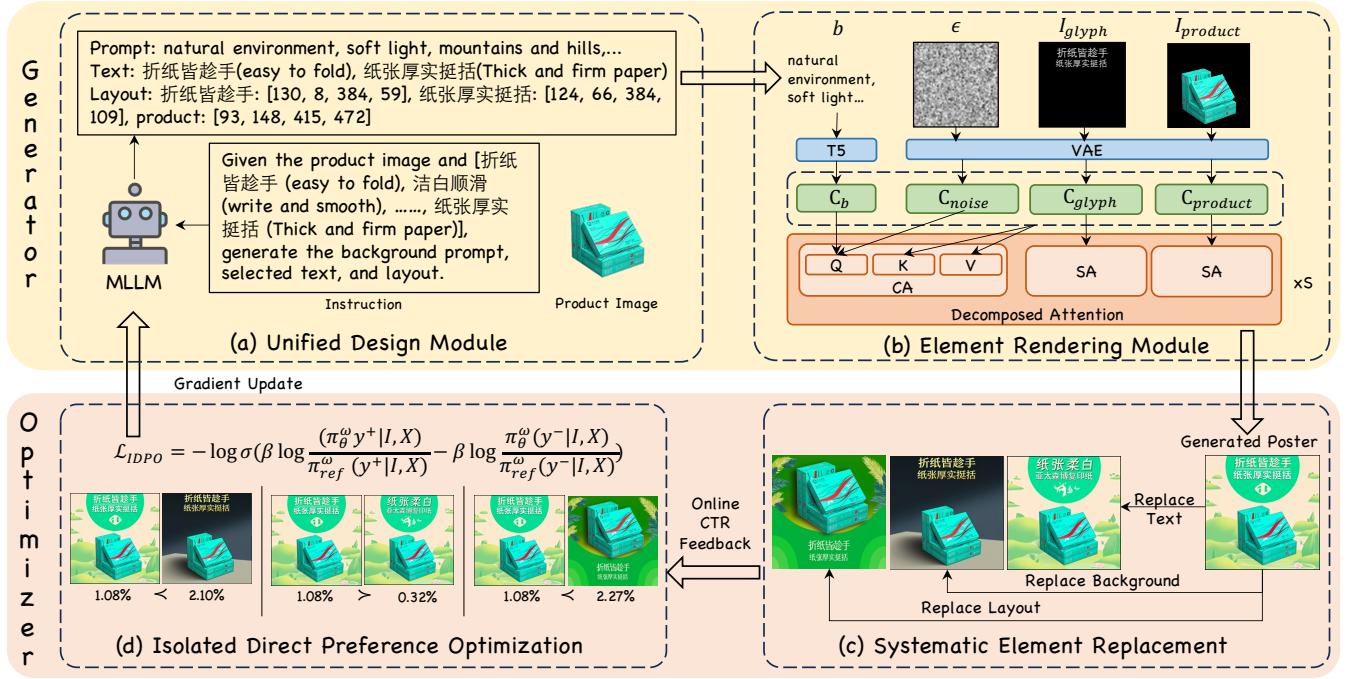


Figure 2: The framework of AutoPP. The generator synthesizes posters while the optimizer refines online CTR. For clarity, we only show the decomposed attention of MM-DiT blocks in the element rendering module. The English translations of the Chinese text are provided in parentheses.

designs, we employ a Multimodal Large Language Model (MLLM) to jointly produce all elements. Formally, given an original product image  $I_{product}$  and a set of candidate text  $T$ , we formulate a natural language instruction  $X_{instr}$  using a template function  $f_{instr}$ :

$$X_{instr} = f_{instr}(T). \quad (1)$$

A concrete instantiation of  $f_{instr}$  follows this pattern: *Given the product image and [candidate text], generate the background prompt, selected text, and layout.* [candidate text] is a placeholder.

Then, our MLLM models the joint distribution as an autoregressive sequence:

$$\pi(y|I_{product}, X_{instr}) = \prod_i p(y_i|I_{product}, X_{instr}, y_{<i}), \quad (2)$$

where  $y$  denotes the structured unified design output, containing three key elements: background prompt  $b$ , selected text  $T^*$ , and layout  $l$ , and  $y_i$  represents the  $i$ -th token. The module is trained to minimize the cross-entropy loss.

**Element Rendering Module** Building upon the outputs from the unified design module, our element rendering module generates the final poster through an efficient token-based architecture as shown in Figure 2 (b). We first prepare two specialized control signals: a product image  $I_{product}$  positioned on a black background following the predicted product layout, and a glyph image  $I_{glyph}$  with text rendered at the predicted text layout. These mask-form control images provide explicit spatial guidance through foreground-background separation.

Then, we employ a token-based mechanism for spatial robust glyph control. Note that it eliminates the need for pixel-level alignment between glyph and target images, significantly reducing annotation effort. Specifically, given the target image  $I_{target}$ , the glyph image  $I_{glyph}$ , the product image  $I_{product}$ , and the background prompt  $b$ , we first encode the visual part with VAE (Kingma, Welling et al. 2013), and the textual part with T5 (Ni et al. 2021):

$$\begin{aligned} C_x &= \text{VAE}(I_x) \in R^{N \times D}, \quad x \in \{\text{target, glyph, product}\} \\ C_b &= \text{T5}(b) \in R^{M \times D}, \end{aligned} \quad (3)$$

where  $C_x$  and  $C_b$  represent the encoded tokens.  $D, N, M$  denote the latent dimension, the flattened sequence length of image patches, and the tokenized text length, respectively. Since  $I_{target}$ ,  $I_{glyph}$ , and  $I_{product}$  share the same resolution, their sequence lengths are equal to  $N$ . Following the flow matching paradigm, we gradually perturb the target latent  $C_{target}$  through normalized time steps  $t \in [0, 1]$ :

$$C_{noise} = (1 - t)C_{target} + t\epsilon, \quad (4)$$

where  $\epsilon \sim \mathcal{N}(0, \mathbf{I})$  denotes Gaussian noise sampled from a standard normal distribution.

To effectively leverage the encoded token sequence described above, a straightforward approach is to concatenate all tokens and process them jointly through  $S$  MM-DiT blocks (BlackForestLabs 2024). However, this naive strategy leads to quadratic computational costs, making it impractical for long sequences. To alleviate this problem, we

modify the full attention in MM-DiT blocks into a decomposed attention (DA) mechanism, which contains two independent components: condition self-attention (SA) and image-condition cross-attention (CA). Specifically, glyph and product tokens are each processed with condition SA, allowing intra-element dependencies to be efficiently captured. To enable information exchange across image tokens and condition tokens, we perform image-condition CA, where the query is constructed by concatenating the prompt and noise tokens, while the key and value are formed by concatenating all types of tokens:

$$\begin{aligned} Q &= [Q_b; Q_{\text{noise}}], \\ K &= [K_b; K_{\text{noise}}; K_{\text{glyph}}; K_{\text{product}}], \\ V &= [V_b; V_{\text{noise}}; V_{\text{glyph}}; V_{\text{product}}], \end{aligned} \quad (5)$$

where  $Q$ ,  $K$ , and  $V$  denote query, key, and value in the cross-attention mechanism.

The training objective requires careful design to ensure precise text rendering—a critical requirement in poster generation. In addition to the core flow matching loss, we introduce an OCR perceptual loss that explicitly enforces legibility in text regions:

$$\mathcal{L}_{\text{render}} = \|v_{\text{pred}} - v_{\text{gt}}\|_2^2 + \lambda \|f_{\text{ocr}}(I_{\text{pred}}) - f_{\text{ocr}}(I_{\text{target}})\|_2^2, \quad (6)$$

where  $v_{\text{pred}}$  and  $v_{\text{gt}}$  are the predicted velocity field and the ground-truth velocity field at each denoising step, respectively.  $I_{\text{pred}}$  is the reconstructed image,  $f_{\text{ocr}}$  denotes intermediate features from the pretrained PaddleOCRv4 backbone, and  $\lambda = 0.1$  balances the loss terms.

## Automated Product Poster Optimizer

**Systematic Element Replacement** To assess the contribution of isolated elements to CTR improvement, we develop a strategy that systematically replaces elements in the poster, as shown in Figure 2 (c). Our approach begins with a generated poster  $P$  and its three key elements  $y$ : background prompt  $b$ , selected text  $T^*$ , where  $T^* \subseteq T$ , and layout  $l$ . To disentangle the impact of isolated elements, we create modified variants by perturbing individual elements while maintaining others fixed. Specifically, for background replacement, we provide the original product image and prompt  $b$  for GPT-4o (Hurst et al. 2024), and require it to generate a distinct prompt  $b'$  that is not a repetition of  $b$ , but is suitable for the product. Text content is replaced with length-matched alternative texts  $T^{*'}$  from the same candidate text set  $T$ . For layout replacement, we regenerate another layout  $l'$  by our unified design module. After replacing each element, we use the element rendering module to regenerate a product poster  $P'$  with the replaced element  $y'$ .

**Isolated Direct Preference Optimization** These posters then undergo a randomized display experiment on JD.COM to collect CTR feedback as shown in Figure 2 (d). For simplicity, we denote the product image  $I_{\text{product}}$  as  $I$ , the instruction  $X_{\text{instr}}$  as  $X$  in the following description. The experimental result establishes a preference relation  $P^+ \succ P^-$ , where  $P^+$  and  $P^-$  denote the higher-CTR and lower-CTR posters between  $(P, P')$ , respectively. Similar to

CAIG (Chen et al. 2025b), we focus solely on fine-tuning the unified design module to choose higher attractive elements  $y^+$  and reject less attractive elements  $y^-$ . The standard DPO (Rafailov et al. 2023) is a coarse-grained alignment method that directly optimizes the overall output of our unified design model  $\pi_\theta(y|I, X)$ . Given the policy model  $\pi_\theta$  and the reference model  $\pi_{\text{ref}}$ , the DPO objective is formulated as:

$$\mathcal{L}_{\text{DPO}} = -\log \sigma \left( \beta \log \frac{\pi_\theta(y^+|I, X)}{\pi_{\text{ref}}(y^+|I, X)} - \beta \log \frac{\pi_\theta(y^-|I, X)}{\pi_{\text{ref}}(y^-|I, X)} \right), \quad (7)$$

where  $\sigma$  denotes the sigmoid function and  $\beta$  is a temperature parameter controlling the deviation from the reference model.

We argue that background  $b$ , text  $T^*$ , and layout  $l$  contribute differently to user engagement. This motivates our IDPO, which incorporates CTR feedback through element-aware optimization:

$$\begin{aligned} \log \pi^w(y|I, X) &= \frac{\sum_i w_i \log p(y_i|I, X, y_{<i})}{\sum_i w_i}, \\ w_i &= \sum_{c \in b, T^*, l} \alpha_c \cdot I(y_i \in c), \end{aligned} \quad (8)$$

where  $w_i$  are element-specific weights and  $\alpha_c$  controls the importance of element  $c$ .  $I(\cdot)$  is an indicator function. For instance, emphasizing background modifications would set  $\alpha_b > 1$ ,  $\alpha_l = \alpha_{T^*} = 1$ . Additionally, the likelihood is normalized by  $\sum_i w_i$  to maintain scale consistency. We substitute  $\log \pi(y|I, X)$  in the DPO loss with  $\log \pi^w(y|I, X)$  to formulate the IDPO loss. This decoupled approach ensures CTR feedback accurately guides modifications to the most influential element, yielding more efficient and interpretable poster optimization compared to holistic approaches.

## Experiment

### Implementation Details

For the product poster generator, we initialize the unified design module from LLaVA (Liu et al. 2023) and the element rendering module from FLUX.1 dev (BlackForest-Labs 2024). The design module undergoes full-model fine-tuning for 10 epochs using a cosine learning rate scheduler (initialized as  $2e-6$ ), while the rendering module employs LoRA (Hu et al. 2022) adaptation with a learning rate of  $5e-6$  for 2 epochs, and is trained at the resolution of  $800 \times 800$ . The training phase takes approximately 1 and 2.5 days respectively. The product poster optimizer is trained using our proposed IDPO for 3 epochs, with hyperparameter  $\beta = 0.5$ ,  $\alpha = 1$  for the constant elements, and  $\alpha = 5$  for the replaced element, taking about 12 hours to complete. All experiments are conducted on a single node with 8 NVIDIA H100 GPUs.

### Product Poster Generation Performance

**Evaluation Metric** We evaluate our model from three perspectives on 500 product posters. For overall poster quality, we use FID (Heusel et al. 2017) and CLIP-T (Gao et al. 2025) to assess the visual quality of the generated posters. For layout rationality, we adopt Alignment (Ali) to measure how well the elements within the layout align with



Figure 3: Qualitative comparison with SOTA product poster generation methods. They often suffer from issues such as obstructed products, unclear text, and incorrect text, noted by colorful boxes.

Method	FID ↓	Clip-T ↑	Ali ↓	Ove ↓	MIoU ↑
P&R	104.05	27.21	0.014	0.024	0.203
PAID	83.55	28.92	0.013	0.041	0.215
GPT-4o	63.47	29.58	0.009	0.018	0.140
<b>AutoPP</b>	<b>60.71</b>	<b>29.75</b>	<b>0.007</b>	<b>0.011</b>	<b>0.256</b>

Table 1: Comparison of product poster generation methods.

Method	Sen. Acc ↑	NED ↓	FID ↓	Clip-T ↑
FLUX-Fill	33.14	49.07	51.53	28.51
FLUX-ControlNet	51.20	26.25	47.74	30.31
Glyph-byt5-v2	52.89	14.99	54.80	30.23
AnyText-v2	54.22	21.00	73.71	29.74
PosterMaker	57.87	21.93	49.76	30.43
<b>AutoPP</b>	<b>65.19</b>	<b>12.94</b>	<b>43.19</b>	<b>30.49</b>

Table 2: Comparison of visual text generation methods.

each other, Overlap (Ove) to calculate the area of overlap between any two elements, and Maximum IoU (MIoU) to evaluate the intersection over union between the generated layout and the ground-truth layout. We assess text rendering using sentence accuracy (Sen. Acc) and normalized edit distance (NED), employing PaddleOCRv4 to detect the text in the generated posters and compare the detected text with the ground-truth text.

**Comparison with SOTA** We compare AutoPP with both open-source methods (Li et al. 2023b; Chen et al. 2025a) and closed-source GPT-4o (Hurst et al. 2024), where open-source methods are trained on the same dataset as ours. See-dream 2.0 (Gong et al. 2025) is excluded from the comparison due to its potential to alter product appearances in ways

that could mislead consumers. For GPT-4o, we provide it with the same basic product information as ours. As shown in Table 1, our method outperforms all others in terms of visual quality and layout rationality. We also showcase several examples in Figure 3. P&R and PAID may cause unclear text and occlusion of products due to manual rendering rules. GPT-4o performs poorly in Chinese text generation and is prone to stroke errors in text. Some generated posters are shown in Figure 4, demonstrating that our method can generate diverse and exquisite product posters.

To further evaluate visual text generation performance, we use ground-truth layouts, background prompts, and text as inputs for all methods. For Glyph-byt5-v2 (Liu et al. 2024) and Anytext-v2 (Tuo, Geng, and Bo 2024), inpainting is applied to ensure that the product region remains unchanged. The results are summarized in Table 2. Our method achieves the best performance for both textual accuracy and visual quality metrics, proving the effectiveness of AutoPP in controllable visual text generation.

**Ablation Study** We conduct experiments to validate the efficient and robust generation capability of the element rendering module.

*Efficient Generation:* Thanks to the flexible token mechanism, our method does not add additional parameters when introducing more control conditions. In contrast, PosterMaker (Gao et al. 2025) and Flux-ControlNet (Zhang, Rao, and Agrawala 2023) increase the parameter count by 1.6B and 4.2B, respectively. We further compare DA and full attention’s cost. Results show that DA reduces GFLOPs of the MM-DiT block by 18% (800×800 resolution) and by 24% (1024×1024 resolution) with comparable performance (65.19 vs 66.37 Sen. Acc↑, 43.19 vs 44.28 FID↓).

*Robust Generation:* Our method directly encodes the glyph image as tokens, making it robust to the misalignment between the glyph image and the product poster. As shown in the first two rows of Table 2, both Flux-ControlNet and



Figure 4: Posters generated by AutoPP. Our method generates diverse layouts for various product categories.




Figure 5: Relative CTR improvement compared with SOTA.

Flux-Fill require pixel-level alignment, resulting in lower Sen. Acc and higher NED compared to our method.

### Product Poster Optimization Performance

**Evaluation Metric** We first use the Relative CTR to measure the CTR improvement compared with the non-optimized pretrained product poster generator. Then the Reward Accuracy quantifies the proportion of higher CTR posters that also receive higher model preference probabilities than their lower-CTR counterparts.

**Comparison with SOTA** To assess the CTR improvement capability of our product poster optimizer, we conduct an evaluation between our AutoPP framework and two state-of-the-art CTR optimization approaches: CG4CTR (Yang et al. 2024) and CAIG (Chen et al. 2025b). We randomly select 10,000 products and generate one poster per product using each of the three methods. The CTR performance is collected through a one-week online experiment on JD.COM. As shown in Figure 5, AutoPP achieves significant improvement in CTR performance, while comparison methods CG4CTR and CAIG show negative Relative CTR due to their failure to incorporate layout and text elements.

**Ablation Study** We conduct experiments to evaluate the effectiveness and impact of data volume for the optimizer.

*Effective Optimization:* Compared with the pretrained model, AutoPP achieves superior CTR (4.49%) through its automated online optimization capability, as shown in Figure 5. More importantly, while both AutoPP and DPO utilize the same training data, AutoPP’s isolated optimization strategy enables more effective learning of the complex relationships between poster elements, leading to better CTR performance (4.49% vs. 3.10%). Notably, even mere 0.5% CTR enhancements can yield significant financial returns, underscoring the practical value of our method.

*Data Volume Impact:* We first train the product poster optimizer across the range 10K, 30K, 50K samples and calculate Reward Accuracy on 1,000 test samples. Reward accuracy shows consistent improvement with increasing data volume (51.20% → 67.19% → 75.99%), confirming the framework’s ability to effectively leverage larger datasets without performance saturation.

### Conclusion

In this work, we introduce AutoPP, an automated pipeline for high-quality product poster generation and optimization. It integrates background, text, and layout design into a unified module, followed by an element rendering module, enabling generation with basic product information while overcoming the limitations of fragmented workflows. To further enhance online performance, we propose a systematic element replacement strategy coupled with IDPO, specifically designed for maximizing CTR. To advance research in the related field, we will release AutoPP1M—the largest dataset for poster generation and optimization to date. Experiments validate that AutoPP achieves SOTA results in both offline and online settings.

## References

- Bai, S.; Chen, K.; Liu, X.; Wang, J.; Ge, W.; Song, S.; Dang, K.; Wang, P.; Wang, S.; Tang, J.; et al. 2025. Qwen2. 5-v1 technical report. *arXiv preprint arXiv:2502.13923*.
- BlackForestLabs. 2024. FLUX. <https://github.com/black-forest-labs/flux>.
- Chen, H.; Zhou, M.; Jiang, J.; Chen, J.; Lu, Y.; Xiao, B.; Ge, T.; and Zheng, B. 2025a. PAID: A Framework of Product-Centric Advertising Image Design. *arXiv preprint arXiv:2501.14316*.
- Chen, J.; Huang, Y.; Lv, T.; Cui, L.; Chen, Q.; and Wei, F. 2024. Textdiffuser-2: Unleashing the power of language models for text rendering. In *European Conference on Computer Vision*, 386–402. Springer.
- Chen, X.; Feng, W.; Du, Z.; Wang, W.; Chen, Y.; Wang, H.; Liu, L.; Li, Y.; Zhao, J.; Li, Y.; et al. 2025b. CTR-Driven Advertising Image Generation with Multimodal Large Language Models. In *Proceedings of the ACM on Web Conference 2025*, 2262–2275.
- Du, Z.; Feng, W.; Wang, H.; Li, Y.; Wang, J.; Li, J.; Zhang, Z.; Lv, J.; Zhu, X.; Jin, J.; et al. 2024. Towards Reliable Advertising Image Generation Using Human Feedback. In *European Conference on Computer Vision*, 399–415. Springer.
- Esser, P.; Kulal, S.; Blattmann, A.; Entezari, R.; Müller, J.; Saini, H.; Levi, Y.; Lorenz, D.; Sauer, A.; Boesel, F.; et al. 2024. Scaling rectified flow transformers for high-resolution image synthesis. In *Forty-first international conference on machine learning*.
- Gao, Y.; Lin, Z.; Liu, C.; Zhou, M.; Ge, T.; Zheng, B.; and Xie, H. 2025. Postermaker: Towards high-quality product poster generation with accurate text rendering. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 8083–8093.
- Gong, L.; Hou, X.; Li, F.; Li, L.; Lian, X.; Liu, F.; Liu, L.; Liu, W.; Lu, W.; Shi, Y.; et al. 2025. Seedream 2.0: A native chinese-english bilingual image generation foundation model. *arXiv preprint arXiv:2503.07703*.
- Heusel, M.; Ramsauer, H.; Unterthiner, T.; Nessler, B.; and Hochreiter, S. 2017. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30.
- Hsu, H. Y.; He, X.; Peng, Y.; Kong, H.; and Zhang, Q. 2023. Posterlayout: A new benchmark and approach for content-aware visual-textual presentation layout. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6018–6026.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; Chen, W.; et al. 2022. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2): 3.
- Hurst, A.; Lerer, A.; Goucher, A. P.; Perelman, A.; Ramesh, A.; Clark, A.; Ostrow, A.; Welihinda, A.; Hayes, A.; Radford, A.; et al. 2024. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*.
- Jia, P.; Li, C.; Yuan, Y.; Liu, Z.; Shen, Y.; Chen, B.; Chen, X.; Zheng, Y.; Chen, D.; Li, J.; et al. 2023. COLE: A Hierarchical Generation Framework for Multi-Layered and Editable Graphic Design. *arXiv preprint arXiv:2311.16974*.
- Jin, C.; Xu, H.; Song, R.; and Lu, Z. 2022. Text2poster: Laying out stylized texts on retrieved images. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 4823–4827. IEEE.
- Kingma, D. P.; Welling, M.; et al. 2013. Auto-encoding variational bayes.
- Kirstain, Y.; Polyak, A.; Singer, U.; Matiana, S.; Penna, J.; and Levy, O. 2023. Pick-a-pic: An open dataset of user preferences for text-to-image generation. *Advances in Neural Information Processing Systems*, 36: 36652–36663.
- Li, C.; Liu, W.; Guo, R.; Yin, X.; Jiang, K.; Du, Y.; Du, Y.; Zhu, L.; Lai, B.; Hu, X.; et al. 2022. PP-OCRv3: More attempts for the improvement of ultra lightweight OCR system. *arXiv preprint arXiv:2206.03001*.
- Li, F.; Liu, A.; Feng, W.; Zhu, H.; Li, Y.; Zhang, Z.; Lv, J.; Zhu, X.; Shen, J.; Lin, Z.; et al. 2023a. Relation-aware diffusion model for controllable poster layout generation. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, 1249–1258.
- Li, Z.; Li, F.; Feng, W.; Zhu, H.; Li, Y.; Zhang, Z.; Lv, J.; Shen, J.; Lin, Z.; Shao, J.; et al. 2023b. Planning and Rendering: Towards Product Poster Generation with Diffusion Models. *arXiv preprint arXiv:2312.08822*.
- Li, Z.; Zhang, J.; Lin, Q.; Xiong, J.; Long, Y.; Deng, X.; Zhang, Y.; Liu, X.; Huang, M.; Xiao, Z.; et al. 2024. Hunyuan-dit: A powerful multi-resolution diffusion transformer with fine-grained chinese understanding. *arXiv preprint arXiv:2405.08748*.
- Liang, T.; Liu, J.; Huang, Y.; Jiang, S.; Shi, J.; Wang, C.; and Li, C. 2024a. TextCenGen: Attention-Guided Text-Centric Background Adaptation for Text-to-Image Generation. *arXiv preprint arXiv:2404.11824*.
- Liang, Y.; He, J.; Li, G.; Li, P.; Klimovskiy, A.; Carolan, N.; Sun, J.; Pont-Tuset, J.; Young, S.; Yang, F.; Ke, J.; Dvijotham, K. D.; Collins, K.; Luo, Y.; Li, Y.; Kohlhoff, K. J.; Ramachandran, D.; and Navalpakkam, V. 2024b. Rich Human Feedback for Text-to-Image Generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Lin, J.; Zhou, M.; Ma, Y.; Gao, Y.; Fei, C.; Chen, Y.; Yu, Z.; and Ge, T. 2023. Autoposter: A highly automatic and content-aware design system for advertising poster generation. In *Proceedings of the 31st ACM International Conference on Multimedia*, 1250–1260.
- Liu, H.; Li, C.; Wu, Q.; and Lee, Y. J. 2023. Visual instruction tuning. *Advances in neural information processing systems*, 36: 34892–34916.
- Liu, Z.; Liang, W.; Zhao, Y.; Chen, B.; Liang, L.; Wang, L.; Li, J.; and Yuan, Y. 2024. Glyph-byt5-v2: A strong aesthetic baseline for accurate multilingual visual text rendering. *arXiv preprint arXiv:2406.10208*.
- Lu, S.; Chen, Y.; Feng, W.; Fan, J.; Li, F.; Zhang, Z.; Lv, J.; Shen, J.; Law, C.; and Liang, J. 2025. Uni-layout: Integrating human feedback in unified layout generation and evaluation. In *Proceedings of the 33rd ACM International Conference on Multimedia*, 7709–7718.

- Ma, J.; Deng, Y.; Chen, C.; Du, N.; Lu, H.; and Yang, Z. 2025. Glyphdraw2: Automatic generation of complex glyph posters with diffusion models and large language models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 5955–5963.
- Ni, J.; Abrego, G. H.; Constant, N.; Ma, J.; Hall, K. B.; Cer, D.; and Yang, Y. 2021. Sentence-t5: Scalable sentence encoders from pre-trained text-to-text models. *arXiv preprint arXiv:2108.08877*.
- Peebles, W.; and Xie, S. 2023. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, 4195–4205.
- Rafailov, R.; Sharma, A.; Mitchell, E.; Manning, C. D.; Ermon, S.; and Finn, C. 2023. Direct preference optimization: Your language model is secretly a reward model. *Advances in neural information processing systems*, 36: 53728–53741.
- Ren, T.; Liu, S.; Zeng, A.; Lin, J.; Li, K.; Cao, H.; Chen, J.; Huang, X.; Chen, Y.; Yan, F.; Zeng, Z.; Zhang, H.; Li, F.; Yang, J.; Li, H.; Jiang, Q.; and Zhang, L. 2024. Grounded SAM: Assembling Open-World Models for Diverse Visual Tasks. *arXiv:2401.14159*.
- 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.
- Tuo, Y.; Geng, Y.; and Bo, L. 2024. AnyText2: Visual Text Generation and Editing With Customizable Attributes. *arXiv preprint arXiv:2411.15245*.
- Wang, H.; Feng, W.; Li, Y.; Zhang, Z.; Lv, J.; Shen, J.; Lin, Z.; and Shao, J. 2025a. Generate e-commerce product background by integrating category commonality and personalized style. In *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1–5. IEEE.
- Wang, S.; Ge, Y.; Chen, L.; Zhou, H.; Wang, Q.; Cheng, X.; and Yuan, L. 2024. Prompt2Poster: Automatically Artistic Chinese Poster Creation from Prompt Only. In *Proceedings of the 32nd ACM International Conference on Multimedia*, 10716–10724.
- Wang, Y.; Li, Z.; Zang, Y.; Wang, C.; Lu, Q.; Jin, C.; and Wang, J. 2025b. Unified multimodal chain-of-thought reward model through reinforcement fine-tuning. *arXiv preprint arXiv:2505.03318*.
- 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. 2023. ImageReward: learning and evaluating human preferences for text-to-image generation. In *Proceedings of the 37th International Conference on Neural Information Processing Systems*, 15903–15935.
- Yang, H.; Yuan, J.; Yang, S.; Xu, L.; Yuan, S.; and Zeng, Y. 2024. A New Creative Generation Pipeline for Click-Through Rate with Stable Diffusion Model. In *Companion Proceedings of the ACM Web Conference 2024*, 180–189.
- Yang, Y.; Gui, D.; Yuan, Y.; Liang, W.; Ding, H.; Hu, H.; and Chen, K. 2023. Glyphcontrol: glyph conditional control for visual text generation. *Advances in Neural Information Processing Systems*, 36: 44050–44066.
- Zhang, L.; Rao, A.; and Agrawala, M. 2023. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF international conference on computer vision*, 3836–3847.
- Zhou, M.; Xu, C.; Ma, Y.; Ge, T.; Jiang, Y.; and Xu, W. 2022. Composition-aware graphic layout GAN for visual-textual presentation designs. *arXiv preprint arXiv:2205.00303*.