

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 POSTERCRAFT: RETHINKING HIGH-QUALITY AESTHETIC POSTER GENERATION IN A UNIFIED FRAMEWORK

Anonymous authors

Paper under double-blind review

ABSTRACT

Generating aesthetic posters is more challenging than simple design images: it requires not only precise text rendering but also the seamless integration of abstract artistic content, striking layouts, and overall stylistic harmony. To address this, we propose PosterCraft, a unified framework that abandons prior modular pipelines and rigid, predefined layouts, allowing the model to freely explore coherent, visually compelling compositions. PosterCraft employs a carefully designed, cascaded workflow to optimize the generation of high-aesthetic posters: (i) large-scale text-rendering optimization on our newly introduced Text-Render-2M dataset; (ii) region-aware supervised fine-tuning on HQ-Poster-100K; (iii) aesthetic-text reinforcement learning via best-of-n preference optimization; and (iv) joint vision–language feedback refinement. Each stage is supported by a fully automated data-construction pipeline tailored to its specific needs, enabling robust training without complex architectural modifications. Evaluated on multiple experiments, PosterCraft significantly outperforms open-source baselines in rendering accuracy, layout coherence, and overall visual appeal—approaching the quality of SOTA commercial systems.



Figure 1: **Aesthetic posters generated by PosterCraft** demonstrate that backgrounds, layouts, and typographic designs are produced directly from textual input without modular designs, highlighting its unified framework for visually consistent and appealing designs.

1 INTRODUCTION

Despite recent advances in automated visual design, aesthetic poster generation remains a formidable challenge and is still relatively underexplored. Existing generative approaches primarily focus on foundational tasks such as text rendering Chen et al. (2023b); Liu et al. (2024); Tuo et al. (2023); Chen et al. (2023a) or the creation of specific product-oriented posters Wang et al. (2025); Gao et al. (2025); Podell et al. (2023), offering limited capacity to produce high-quality, aesthetically compelling outputs. These methods often fall short of addressing the multifaceted demands of aesthetic poster design, which requires not only *(i)* accurate and stylistically coherent text, but also *(ii)* the creation of abstract and visually appealing artistic content and *(iii)* the striking layouts and holistic stylistic consistency. Therefore, aesthetic poster generation demands a more comprehensive synthesis of content, form, and communicative intent.

Recent approaches to aesthetic poster generation Chen et al. (2025a); Yang et al. (2024b); Seol et al. (2024); Tang et al. (2023) have primarily followed a modular design paradigm. Typically, a fine-tuned vision-language model (VLM) acts as a layout planner, suggesting text content and positioning. The suggestions are then overlaid onto a separately generated background, or used as hard constraints for the generative model to follow. However, this design strategy presents several limitations. *(i) Lack of aesthetic consistency*: it undermines the visual and stylistic coherence essential for aesthetic poster creation. *(ii) Limited visual quality*: it constrains the upper bound of visual quality due to the decoupled design process and heavy reliance on the VLM’s accuracy and robustness. In contrast, existing end-to-end design-centered generation approaches Chen et al. (2023b); Inoue et al. (2024); Wang et al. (2025); Gao et al. (2025) remain limited to relatively simple tasks, such as greeting cards or product compositions, which lack the visual and structural complexity of high-quality aesthetic posters. Additionally, while powerful foundation models <https://github.com/black-forest/labs/flux> (2024); Esser et al. (2024); AI (2024) have demonstrated impressive capabilities in generating complex natural images, they still fall short of meeting all the specific requirements of aesthetic posters. (e.g. precise text rendering, abstract artistic content, and holistic stylistic coherence). For this, we classify it as *(iii) Simplified use cases*. More importantly, the absence of large-scale, versatile datasets tailored specifically for aesthetic poster generation has further constrained the development of fully generative solutions—*(iv) Absence of targeted datasets*.

To move beyond the limitations of modular and simply scoped designs, we leverage foundation models to explore unified generation for aesthetic posters, aiming for visually coherent and compelling results. We argue that incremental, component-level improvements are insufficient for major aesthetic gains. Instead, we propose a unified framework, *PosterCraft*, which includes a comprehensive workflow to systematically perform four critical stages: *(i)* scalable text rendering optimization, *(ii)* high-quality poster fine-tuning, *(iii)* aesthetic-text reinforcement learning, and *(iv)* vision-language feedback refinement. We construct specialized datasets for each stage through automated pipelines, enabling robust training and future research in aesthetic poster generation. This framework empowers the trained model to generate high-quality posters end-to-end. Experiments show our approach outperforms existing baselines and is competitive with several closed-source models.

Overall, our contributions can be summarized as follows:

- **A unified framework for aesthetic poster generation:** We revisit aesthetic poster generation through an end-to-end approach for high-quality, visually coherent posters, surpassing prior modular pipelines and methods focused on simpler or product-centric designs.
- **A cascade workflow for high-quality poster optimization:** We propose a unified training pipeline with four stages: (i) scalable text rendering optimization, (ii) high-quality poster fine-tuning, (iii) aesthetic-text reinforcement learning, and (iv) vision-language feedback refinement. Each stage targets a key challenge in aesthetic poster generation, enabling the model to produce artistically compelling results at inference time.
- **Stage-specific, fully automated dataset construction:** We construct specialized datasets for each workflow stage using automated collection and filtering, tailored to the unique demands of aesthetic poster generation. These datasets overcome the limitations of resources and support more robust, transferable training.
- **Superior performance over existing baselines:** Extensive experiments show that our method significantly outperforms open-source baselines in terms of aesthetic quality and layout structure, and achieves competitive performance compared to commercial systems.

2 RELATED WORKS

Image Generation for Design Images. Recently, design-centric image generation has transitioned from early GAN- and VAE-based generators to more powerful diffusion-based frameworks Chen et al. (2023b); Tuo et al. (2023); Wang et al. (2025); Gao et al. (2025); Inoue et al. (2023); Zheng et al. (2023); Chen et al. (2023a). Notably, LayoutDiffusion Zheng et al. (2023) reformulated layout generation as a discrete token denoising process, achieving significant gains. Advanced pipelines like TextDiffuser Chen et al. (2023b;a) adopted a two-stage approach: a transformer-based layout planner followed by a diffusion model conditioned on OCR-derived masks to generate coherent text images from prompts. DesignDiffusion Wang et al. (2025) improved text-centric design by enhancing prompts with character-level embeddings and applying a localization loss to boost text rendering accuracy. Despite recent progress, many methods imposed rigid pre-layout constraints on

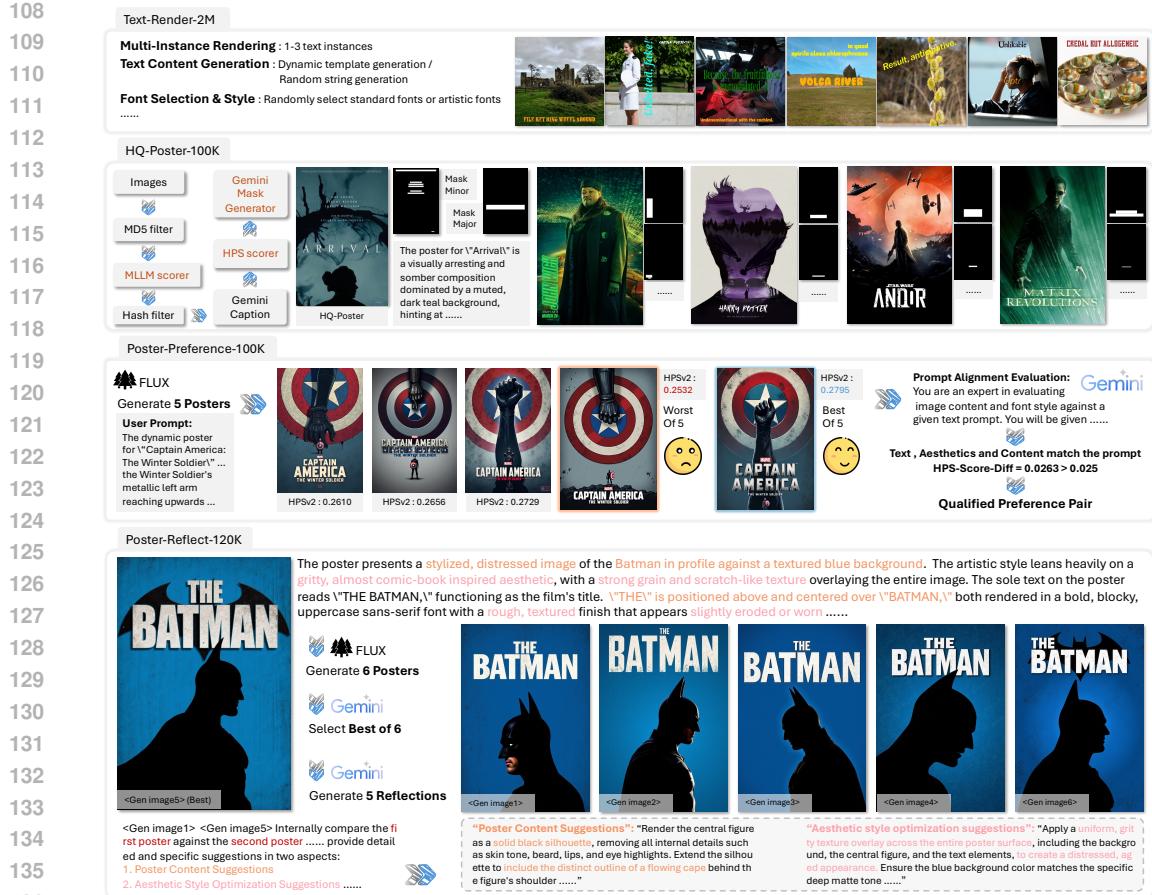


Figure 2: **Four datasets of PosterCraft** across four stages: (i) Text-Render-2M for text rendering optimization, (ii) HQ-Poster-100K, comprising over 100K high-quality posters with masks and captions, (iii) Poster-Preference-100K, yielding 6K high-quality preference pairs from 100K generated samples, (iv) Poster-Reflect-120K, constructing 64K feedback pairs from 120K generated posters.

text and layout, often undermining overall aesthetic coherence. Furthermore, their focus on specific domains such as product ads or greeting cards also limited the complexity and creativity of the tasks they address. In contrast, our approach adopts a unified framework that integrates text rendering, artistic content, and layout design within a single inference process.

VLM for Image Generation. Vision–language models served dual roles as both high-level planners Lin et al. (2023); Luo et al. (2024); Feng et al. (2023); Yang et al. (2024b;a); Chen et al. (2025a); Tang et al. (2023) and fully end-to-end generators Zhou et al. (2024a); Xie et al. (2024); Team (2024a); Ma et al. (2024); Zhou et al. (2024a); Wu et al. (2024). Fine-tuned adaptations PosterLlama Seol et al. (2024) and PosterLLava Yang et al. (2024b) employed HTML- or JSON-formatted tokens to produce content-aware graphic designs. POSTA Chen et al. (2025a) further leveraged a fine-tuned VLM as an “aesthetic designer”, applying modular design atop existing high-quality images. Unified transformer architectures like TransFusion Zhou et al. (2024a), and JanusFlow Ma et al. (2024) advanced this paradigm by generating image and text tokens in one architecture, enabling one-shot synthesis of complete visual compositions. Complementing these approaches, feedback-driven pipelines Li et al. (2024; 2025) incorporated VLM-based critics to assess object accuracy during training. Nevertheless, these VLM-driven methods continued to struggle with maintaining cohesive aesthetic consistency, and faced an upper bound in layout complexity imposed by the models’ architectural capacity and the scope of their limited training data.

3 UNIFIED WORKFLOW AND SPECIFIC DATASET

In this work, we rethink aesthetic poster generation: a high-capacity model optimized through a unified workflow can directly produce high-quality, fully rendered posters without modular design.

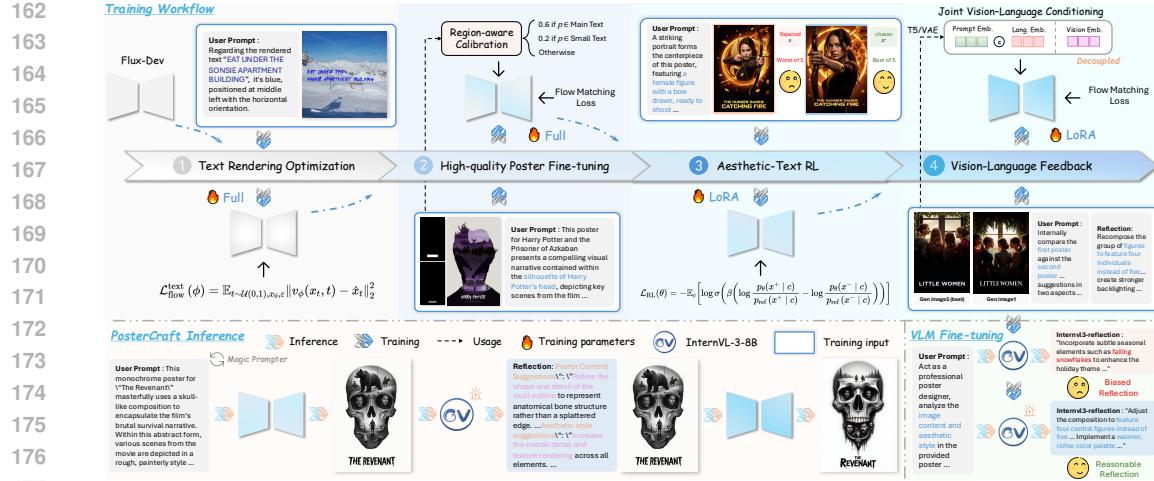


Figure 3: **The PosterCraft pipeline** has four stages: (1) Text Rendering Optimization to improve text accuracy and fidelity; (2) High-Quality Poster Fine-Tuning with region-aware calibration to poster styling across text and non-text regions; (3) Aesthetic-Text Reinforcement Learning to instill detailed aesthetic and content preferences; and (4) Joint Vision-Language Feedback, integrating multimodal reflections for refined outputs. At inference, the fine-tuned model generates high-quality aesthetic posters end-to-end from a single prompt, with an optional VLM-driven critique loop.

This allows holistic integration of text, visuals, and layout, while leveraging vision-language feedback during inference for greater coherence and aesthetic appeal. Unlike prior methods that rely on layout embeddings Zheng et al. (2023); Gao et al. (2025); Chen et al. (2023b) or external VLM-based designers Chen et al. (2025a); Yang et al. (2024b); Seol et al. (2024) (i.e. inherently restrict a model’s expressive freedom), we unlock the potential of a standard diffusion backbone via workflow optimization rather than intricate architectural modifications. Our paradigm remains compatible with existing techniques and provides a flexible foundation for future advances. Fig.3 shows our unified optimization workflow, and Fig.2 the dataset pipeline supporting each stage.

3.1 SCALABLE TEXT RENDERING OPTIMIZATION

In the first stage of our workflow, we target the challenge of accurate text rendering, a persistent bottleneck in poster generation. Progress is hindered by two factors: (i) the scarcity of large-scale, high-quality datasets with perfectly rendered text, and (ii) most available text datasets feature plain or low-quality backgrounds, which easily make the model lose the ability to represent common backgrounds. To overcome these issues, we construct Text-Render-2M via an automated pipeline, producing 2 million samples with diverse text (varying in content, size, count, placement, and rotation) rendered crisply onto high-quality backgrounds. Each text instance is paired with precise captions merged seamlessly with existing image captions. This dataset ensures both 100% text rendering accuracy and rich background diversity, enhancing fidelity and robustness in real-world scenarios. Fig.2 illustrates Text-Render-2M, with construction details and examples provided in the supplementary. We then fine-tune foundation models on paired Text-Render-2M using the flow matching loss Esser et al. (2024) to enhance text rendering:

$$\mathcal{L}_{\text{flow}}^{\text{text}}(\phi) = \mathbb{E}_{t \sim \mathcal{U}(0,1), x_0, \varepsilon} \| v_{\phi}(x_t, t) - \dot{x}_t \|_2^2, \quad (1)$$

where $x_t = \alpha_t x_0 + \sigma_t \varepsilon$ follows the forward noising schedule, \dot{x}_t is its time derivative, and v_{ϕ} predicts the velocity field. Discretized over timesteps, this loss encourages the model to match the true data flow and yields markedly improved text rendering. As shown in Fig.4, the model augmented with our text-rendering optimization achieves significant gains in both rendering accuracy and text alignment across poster typography, plain-text scenes, and long-text posters.

3.2 HIGH-QUALITY POSTER FINE-TUNING

HQ-Poster-100K. To build a high-quality dataset for supervised fine-tuning, we introduce HQ-Poster-100K, a carefully filtered poster dataset. The pipeline first removes exact duplicates via MD5 hashing. To exclude posters with large credit/billing blocks, we employ an MLLM-based scoring

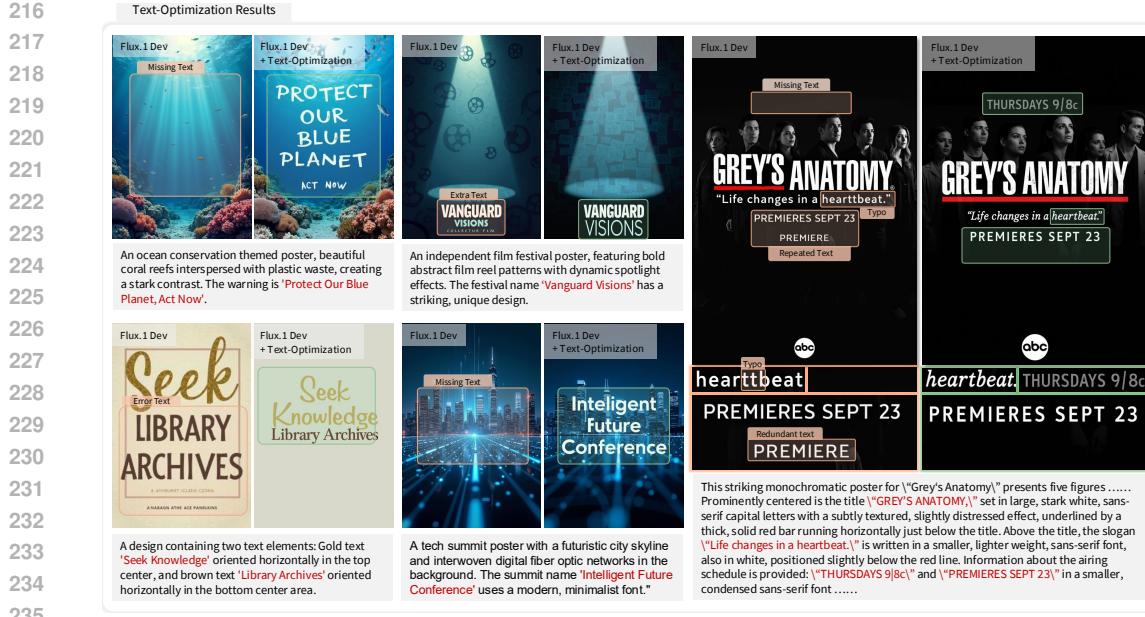


Figure 4: **Comparison of text rendering on poster typography, plain-text scenes, and long-form text posters.** Each pair shows the Flux.1 dev baseline (left), exhibiting missing, repeated, or error text, alongside the optimized output (right) after our scalable text rendering optimization, demonstrating marked gains in text fidelity, alignment, and accuracy.

system that asks a binary question about their presence, outputs logits, and applies Softmax to obtain probabilistic scores. Scores closer to 1 indicate better alignment with our criteria:

$$prob_x = \frac{e^{l_x}}{\sum_{x \in L} e^{l_x}}, \quad score = \sum_{x \in L} prob_x \cdot w_x \quad (2)$$

Where set L contains all the option letters, l_x represents the logit for option x . For two-choice questions (A, B), we set $w_A = 0$ and $w_B = 1$. In our pipeline, logits are computed using VLM (InternVL2.5-8B-MPO InternVL Team (2024)) with a 0.98 threshold to control filtering stringency. We then apply perceptual hashing to remove visually similar posters. The remaining posters are captioned by Gemini2.5-Flash and scored with HPS, filtering out those below 0.25 to ensure aesthetic alignment with human preferences.

To support Region-aware Calibration, HQ-Poster-100K provides precise text region masks for each poster. Since traditional OCR struggles with artistic typography, we use Gemini2.5-Flash to extract text region coordinates. Masks are then classified as major or minor based on their relative size, denoting large and small text areas respectively, as shown in Figure 2.

Region-aware Calibration. In poster design, harmony between text and background is crucial. Since our first stage has already improved the text rendering capability of the model, this phase shifts the focus to overall poster style. Therefore, we propose Region-aware Calibration to achieve it. Specifically, essential text carries the core message and is assigned moderate weight to ensure clarity and integration with the background; by contrast, small text—occupying minimal space and prone to rendering errors—is downweighted to prevent distracting collapse. Non-text regions, which define the visual style of the poster, receive full emphasis to guarantee a smooth transition from high-quality imagery to a unified aesthetic layout. This balanced weighting strategy allows the fine-tuned model to preserve text accuracy while strengthening the artistic integrity of the poster.

We implement this via an empirical per-pixel weight map $w(p)$:

$$w(p) = \begin{cases} 0.6 & \text{if } p \in \text{LargeTextMask}, \\ 0.2 & \text{if } p \in \text{SmallTextMask}, \\ 1.0 & \text{otherwise,} \end{cases} \quad (3)$$

270 and define our weighted loss as:
 271

$$\mathcal{L}_{\text{flow}}^{\text{poster}} = \mathbb{E}_{t \sim \mathcal{U}(0,1), x_0, \varepsilon} \left\| (v_\phi(x_t, t) - \dot{x}_t) \odot w \right\|_2^2, \quad (4)$$

272 where " \odot " denotes point-wise multiplication by the weight map w . Different from previous scalable
 273 text rendering optimization, here we multiply the scaling factor w , which encourages the model to
 274 learn both crisp text information and a cohesive aesthetic style.
 275

276 3.3 AESTHETIC-TEXT REINFORCEMENT LEARNING

277 **Poster-Preference-100K.** To enhance poster aesthetics and text rendering, we construct the Poster-
 278 Preference-100K dataset. Using about 20K prompts, we generate 5 images per prompt with the
 279 model after Region-aware calibration, producing 100K posters as the basis for preference pairs.
 280 With HPSv2 Wu et al. (2023), we score each group of 5 and select the highest- and lowest-scoring
 281 posters as preferred and rejected samples. Since HPSv2 evaluates only content and aesthetics, we
 282 use Gemini2.5-Flash to verify text accuracy and style consistency in the preferred posters, filtering
 283 out inconsistent pairs. This yields 6K preference pairs meeting two criteria: (i) the HPSv2 score gap
 284 ≥ 0.025 , and (ii) complete text accuracy in preferred posters.
 285

286 **Aesthetic-Text Preference Optimization.** While earlier stages ensure pixel-level text fidelity and
 287 calibrated styles, they miss higher-order trade-offs that make a poster genuinely compelling: In par-
 288 ticular, (i) detailed preferences, such as subtle layout balance, color harmony, and typographic co-
 289 hesion, which require global evaluation beyond per-pixel accuracy; (ii) even after achieving crisp text
 290 rendering, further corrective tuning is necessary to alleviate residual errors and seamlessly integrate
 291 text with the holistic aesthetic. To address these gaps, we frame poster generation as a reinforcement
 292 learning problem in this stage: the model must not only denoise accurately but also preferentially
 293 generate outputs that satisfy holistic aesthetic criteria.
 294

295 Concretely, for each prompt, we sample n poster variants $\{x^{(i)}\}_{i=1}^n$ under the current diffusion
 296 policy and collapse them into a single “winning” and “losing” pair via best-of- n selection under the
 297 combined aesthetic-text reward $R(x)$:
 298

$$x^+ = \arg \max_i R(x^{(i)}), \quad x^- = \arg \min_i R(x^{(i)}). \quad (5)$$

301 We then optimize the Direct Preference Optimization (DPO) Rafailov et al. (2023) objective:

$$\mathcal{L}_{\text{RL}}(\theta) = -\mathbb{E}_c \left[\log \sigma \left(\beta \left(\log \frac{p_\theta(x^+ | c)}{p_{\text{ref}}(x^+ | c)} - \log \frac{p_\theta(x^- | c)}{p_{\text{ref}}(x^- | c)} \right) \right) \right], \quad (6)$$

302 where p_{ref} denotes the fixed reference distribution, and p_θ is learned diffusion policy parameterized
 303 by θ . Because the marginal $p_\theta(x_0 | c)$ is intractable, we employ the ELBO over the full diffusion
 304 chain to evaluate these log-ratio rewards, following prior work Wallace et al. (2024); Wang et al.
 305 (2025). In this way, best-of- n Aesthetic-Text Preference Optimization directly injects a unified
 306 preference signal into the diffusion training process.
 307

308 3.4 VISION-LANGUAGE FEEDBACK REFINEMENT.

309 **Poster-Reflect-120K.** To address potential deficiencies in content and aesthetic quality of initially
 310 posters, we apply reflection optimization to improve accuracy and quality. We build the Poster-
 311 Reflect-120K dataset by generating six posters per prompt with our preference-learned model, to-
 312 taling 120K images. Gemini2.5-Flash then selects the optimal poster from each set as the feedback
 313 target, required to meet three conditions: accurate prompt alignment, superior aesthetics, and correct
 314 text rendering. This process yields 5 reflection-pairs from each set of 6 generated images.
 315

316 During the feedback collection phase, we gather suggestions in two key areas: **Poster Content Sug-
 317 gestions** and **Aesthetic Style Optimization Suggestions**, with Gemini2.5-Flash analyzing both the
 318 target poster and the poster requiring optimization to provide comprehensive feedback. Additionally,
 319 to optimize our prompting strategy for both feedback and VLM fine-tuning, we implement specific
 320 guidelines: the model is instructed to perform internal comparisons without explicitly referencing
 321 the second reference poster, and feedback is structured as concrete editing instructions.
 322

324 Table 1: **A text quality comparison** with SOTA poster generation models, demonstrates that Poster-
 325 Craft achieves superior performance across recall, F-score, and accuracy, while only slightly below
 326 the recent closed-source Gemini2.0-Flash-Gen. We highlight the **best** and **second** metrics. Open
 327 and **Close** denote open-source and closed-source.

328 Method	329 Text Recall \uparrow	330 Text F-score \uparrow	331 Text Accuracy \uparrow
330 OpenCOLE Inoue et al. (2024) (Open)	0.082	0.076	0.061
330 Playground-v2.5 Yang et al. (2024b) (Open)	0.157	0.146	0.132
331 PosterMaker Gao et al. (2025) (Open)	0.522	0.488	0.467
332 BizGen Peng et al. (2025) (Open)	0.689	0.661	0.641
333 SD3.5 AI (2024) (Open)	0.565	0.542	0.497
333 Flux1.dev https://github.com/black-forest-labs/flux (2024) (Open)	0.723	0.707	0.667
334 Ideogram-v2 v2. https://ideogram.ai/launch (2024) (Close)	0.711	0.685	0.680
335 BAGEL Deng et al. (2025) (Open)	0.543	0.536	0.463
336 Gemini2.0-Flash-Gen Team et al. (2023) (Close)	0.798	0.786	0.746
337 PosterCraft (ours)	0.787	0.774	0.735

338 **Reflect VLM fine-tuning.** To obtain optimization feedback during inference, we construct VQA
 339 samples by embedding the original caption in the prompt alongside the generated poster requiring
 340 optimization, and using Gemini2.5-Flash-generated feedback as supervision. Specifically, this in-
 341 put configuration maintains consistency between training and inference phases, excluding reference
 342 target posters in both cases and using the original prompt as the baseline for suggestions. Addition-
 343 ally, when generating feedback with Gemini2.5-Flash, we deliberately utilize only target posters as
 344 references, omitting original captions to preserve model creativity.

345 **Joint Vision-Language Conditioning.** For the poster construction, iterative critique—combining
 346 visual inspection with targeted verbal feedback—is essential for refining both aesthetic content and
 347 background harmony. Inspired by this, we introduce a joint vision–language feedback loop for mul-
 348 timodal corrections in a unified workflow. For each generated–ground-truth pair, Gemini produces
 349 two textual reflections, f_c (Poster Content Suggestions) and f_s (Aesthetic Style Suggestions). Rather
 350 than appending these strings to the original prompt—which would exceed the encoder’s length and
 351 degrade performance—we jointly encode them via a text encoder E_t , yielding $e_{c,s} = E_t(f_c, f_s)$,
 352 and then concatenate this with the original prompt embedding e_p (with positional encodings to pre-
 353 serve order and semantics). Additionally, drawing on OmniControl Tan et al. (2024), we inject
 354 the image-level feedback signal v_{img} (encoded by VAE) directly into the conditioning branch. The
 355 resulting multimodal context is:

$$355 \quad c = [e_p; e_{c,s}; v_{\text{img}}] \quad (7)$$

356 where c serves as the conditioning input. Finally, we fine-tune the model under the conditional flow
 357 matching objective:

$$360 \quad \mathcal{L}_{\text{flow}}^{VL}(\theta) = \mathbb{E}_{t,x_0,\varepsilon} \|v_{\theta}(x_t, t \mid c) - \dot{x}_t\|_2^2, \quad (8)$$

361 which enables the model to iteratively refine its outputs in response to structured textual reflections
 362 and semantically enriched visual feedback.

364 3.5 INFERENCE

366 During inference, PosterCraft generates a complete poster from a single user prompt. The prompt
 367 is first processed by a MLLM (e.g., Qwen3-based Yang et al. (2025) Magic Promter), which en-
 368 riches the input with detailed aesthetic cues to guide layout and content generation. PosterCraft then
 369 produces the poster in a unified manner, without any additional inputs or modular designs. Addi-
 370 tionally, an VLM-based critique loop can further enhance the result by providing structured multi-modal
 371 feedback. This enables iterative refinement to improve aesthetic quality and semantic alignment.

372 4 EXPERIMENTS

373 4.1 IMPLEMENTATION

375 For PosterCraft, we initialize from the Flux-dev <https://github.com/black-forest-labs/flux> (2024)
 376 backbone and train in mixed precision. Text rendering optimization involves 300K full-parameter
 377 iterations on Text-Render-2M with Adafactor Shazeer & Stern (2018) at lr = 2e-6. Stage 2 fine-tunes
 6,000 steps on HQ-Poster-100K using Adafactor (lr = 1e-5) with per-pixel flow-matching weights

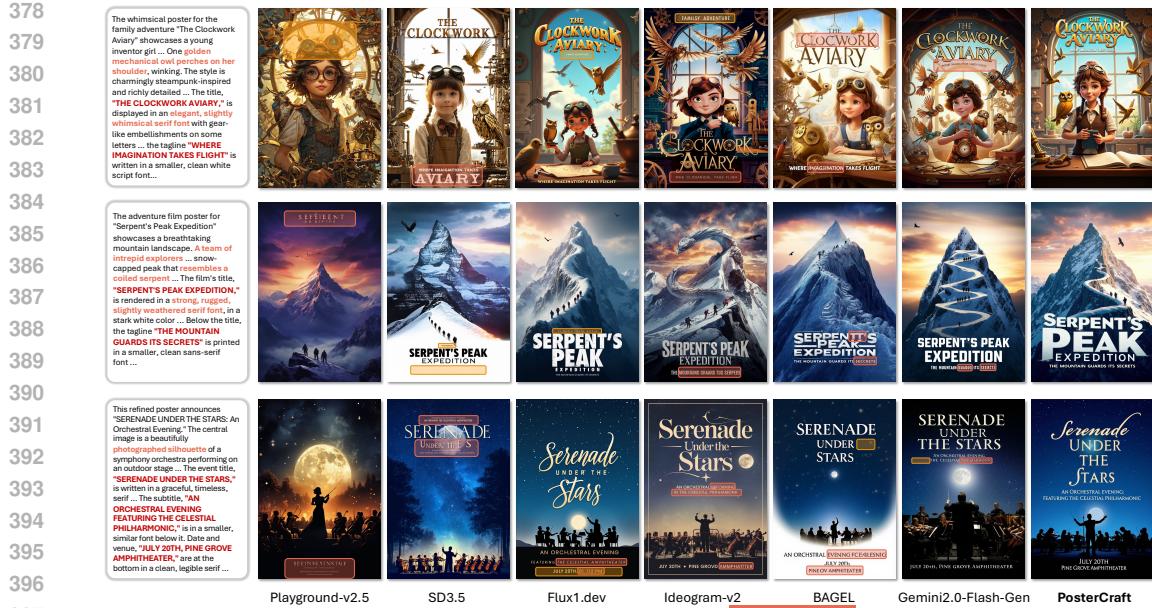


Figure 6: **Visual comparison** of different model outputs. Red boxes highlight misspelled or distorted text, while yellow boxes indicate redundant or missing text elements. Within the prompts, orange text denotes content and style requirements, and red text indicates textual elements. From the visual comparison, it is evident that our method achieves superior aesthetic appeal compared to existing state-of-the-art approaches. In terms of text presentation, our rendered fonts blend more naturally with the scene content, and text rendering errors are nearly eliminated.

0.6, 0.2, 1.0. For reinforcement learning, we sample $n=5$ candidates per prompt and optimize via AdamW Loshchilov & Hutter (2017) ($lr = 1e-4$) for 1500 steps, tuning only LoRA (rank 64). In vision-language feedback refinement, dual-language reflections f_c and f_s are encoded by T5 Raffel et al. (2020), followed by LoRA (rank 128) fine-tuning under conditional flow-matching loss for 6000 steps using AdamW ($lr = 1e-4$). We employ Internvl3-8B InternVL Team (2025) for feedback generation, fine-tuning 2 epochs and using temperature 0 at inference. More details about the dataset can be found in the supplementary material.

4.2 QUANTITATIVE RESULTS AND COMPARISONS



Figure 5: **User study comparisons** between PosterCraft and both SOTA open-source and closed-source models. It consistently outperforms all open-source baselines and several proprietary systems cross multiple dimensions, with performance marginally below that of the leading closed-source model, Gemini2.0-Flash-Gen.

commercial solutions (Ideogram-v2). The table also includes poster-specific generators BizGen and

We conduct a quantitative comparison of PosterCraft against seven leading models, both open-source and commercial. Using Gemini2.0-Flash-Gen Team et al. (2023), we randomly generate 100 aesthetic poster prompts—balanced across short, medium, and long lengths—and sample three outputs per model, yielding 300 test images. Posters are generated with OpenCOLE, Playground-v2.5, SD3.5, Flux1.dev, Ideogram-v2, BAGEL, and Gemini2.0-Flash-Gen. We then apply the OCR engine of the SOTA VLM to each image and report three precision-oriented metrics—text recall, text F-score, and text accuracy—averaged over all 300 samples.

As shown in Tab. 1, PosterCraft achieves substantially higher recall and F1 scores, capturing a more complete character set, while its accuracy surpasses leading open-source methods (Flux1.dev, SD3.5) and competitive commercial solutions (Ideogram-v2). The table also includes poster-specific generators BizGen and

432 PosterMaker, for which we provide the same five representative layouts (and no product image for
 433 PosterMaker); despite their more structured inputs, PosterCraft still attains better recall, F-score,
 434 and accuracy. PosterCraft’s performance closely approaches that of some mature SOTA commer-
 435 cial systems such as Gemini2.0-Flash-Gen, demonstrating that its text-rendering not only advances
 436 academic benchmarks but also offers near production-level competitiveness. Besides, inspired by
 437 previous works Chen et al. (2024a); Esser et al. (2024), we conduct the user study to assess model
 438 quality. Twenty experienced poster designers evaluate outputs across multiple dimensions, as shown
 439 in Fig. 5. Experiments show that our method significantly enhances aesthetics and text rendering
 440 over the base model Flux1.dev, and outperforms all state-of-the-art open-source and several closed-
 441 source models across all evaluation metrics. Additionally, PosterCraft performs only slightly below
 442 Gemini2.0-Flash-Gen, validating the effectiveness of our unified workflow in fully unlocking the
 443 generative potential of the powerful baseline model. More detailed evaluation strategies and experi-
 444 ments on other benchmarks can be found in the supplementary material.
 445

446 4.3 QUALITATIVE RESULTS AND COMPARISONS

447 To further evaluate the superiority of vision, we conduct a comprehensive visual comparison be-
 448 tween PosterCraft and six other models—four open-source and two commercial systems, as shown
 449 in Fig.6. Playground-v2.5 completely fails in text rendering, while SD3.5 renders titles only
 450 partially, with both showing poor aesthetics and weak prompt adherence (e.g., animated styles).
 451 Flux1.dev and BAGEL improve title rendering but still contain textual errors and fall short in aes-
 452 thetics and prompt alignment, missing details such as “One golden mechanical owl perches on her
 453 shoulder.” Among proprietary models, Ideogram-v2 and Gemini2.0-Flash-Gen show only minor
 454 text errors in small text and achieve superior aesthetics, but both have prompt adherence issues:
 455 Ideogram-v2 misrepresents mountain formations, Gemini lacks photorealism, and neither generates
 silhouette effects in the third image. More comparisons appear in the supplementary material.

456 5 ABLATION STUDY

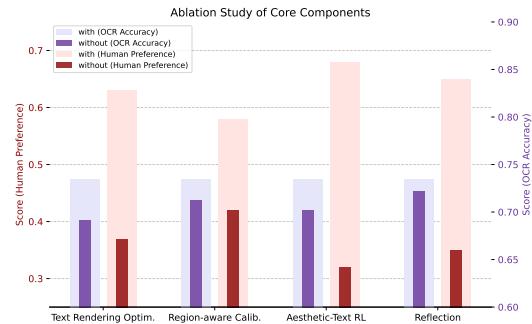
457 In this section, we validate our workflow by independently assessing each stage, keeping all param-
 458 eters and conditions identical to the preceding experiments.

459 460 **Text Rendering Optimization is critical for 461 both accuracy and perception.**

462 Text Rendering Optimization ensures clarity and fidelity. Its
 463 removal leads to drops in both OCR accuracy
 464 and human preference, confirming that large-
 465 scale, high-quality text data significantly im-
 466 proves text rendering. The diverse and realistic
 467 backgrounds also preserve visual quality, while
 468 models without this optimization often fail to
 469 maintain legibility and accuracy, as shown in
 470 Fig.4 and Fig. 7. **Region-aware Calibration**
 471 **improves poster consistency.** Fig. 7 demon-
 472 strates that it helps the model adapt to spatial
 473 context and balance text-background. Without
 474 it, all regions are treated equally, weakening
 475 stylistic coherence in visually complex posters
 476 while suffering text bias. **Aesthetic-Text Re-**
 477 **inforcement Learning and Reflection refine**
 478 **outputs with vision-language feedback.** As
 479 shown in Fig. 7, Aesthetic-Text RL directly en-
 480 hances the harmony between aesthetic appeal and
 481 textual accuracy, while Reflection further im-
 482 proves overall quality through iterative vision-language
 483 guidance. We provide more detailed ablation studies in the supplementary material.

484 6 CONCLUSION

485 PosterCraft presents a unified, cascaded workflow that integrates scalable text-rendering optimiza-
 486 tion, poster fine-tuning, RL-driven aesthetic enhancement, and joint vision–language feedback,
 487 demonstrating powerful foundation model can directly produce posters with both striking visuals



488 **Figure 7: Abl. experiments on the core compo-
 489 nents of our workflow:** Removing any of these
 490 four improvements results in sustained declines
 491 in OCR accuracy (purple) and human preference
 492 (peach), showing our design and optimizations are
 493 effective and our motivation valid.

494 As shown in Fig. 7, Aesthetic-Text RL directly en-
 495 hances the harmony between aesthetic appeal and
 496 textual accuracy, while Reflection further im-
 497 proves overall quality through iterative vision-language
 498 guidance. We provide more detailed ablation studies in the supplementary material.

486 and precise text. Our fully automated dataset pipelines support scalable, task-specific training, and
 487 we achieve substantial gains over open-source baselines, approaching leading commercial quality.
 488

489 7 ETHICS STATEMENT

490 This work does not involve human subjects, personally identifiable information, or sensitive pri-
 491 vate data. All datasets used in this study were either automatically constructed through synthetic
 492 pipelines (e.g., Text-Render-2M, Poster-Preference-100K, Poster-Reflect-120K) or collected from
 493 publicly available poster resources (HQ-Poster-100K) with automated filtering to ensure quality and
 494 compliance.
 495

496 The HQ-Poster-100K dataset contains materials protected by third-party copyrights. Their inclu-
 497 sion follows the principle of fair use under copyright law and is intended solely for non-commercial
 498 scientific research. We do not claim copyright ownership of these materials, and users are responsi-
 499 ble for ensuring that any use complies with applicable laws and regulations. The publishers of this
 500 project assume no responsibility for copyright disputes arising from use of the dataset.
 501

502 Our method focuses on aesthetic poster generation and does not target manipulative, deceptive, or
 503 harmful applications. We are committed to releasing our datasets and models under appropriate
 504 open-source licenses to support transparency, reproducibility, and fair access for the research com-
 505 munity.

506 8 REPRODUCIBILITY STATEMENT

507 We have made substantial efforts to ensure the reproducibility of our work. The main paper and sup-
 508 plementary materials provide detailed descriptions of the proposed framework, dataset construction
 509 pipelines, training configurations, and evaluation protocols. Hyperparameters and ablation studies
 510 are explicitly documented to facilitate independent verification. With the growing availability of
 511 alternative models, many of our data filtering designs based on closed-source models can be eas-
 512 ily reproduced using state-of-the-art open-source vision–language models. In addition, we include
 513 the data processing scripts, evaluation code, and all prompts used in testing in the supplementary
 514 materials to further support transparency and reproducibility. All datasets and models used in our
 515 experiments are described with sufficient detail to ensure full reproducibility and to enable future
 516 research extensions.
 517

519 520 REFERENCES

521 Stability AI. Stable diffusion 3.5 large. [https://huggingface.co/stabilityai/](https://huggingface.co/stabilityai/stable-diffusion-3.5-large)
 522 stable-diffusion-3.5-large, 2024.

523 Qi Cai, Jingwen Chen, Yang Chen, Yehao Li, Fuchen Long, Yingwei Pan, Zhaofan Qiu, Yiheng
 524 Zhang, Fengbin Gao, Peihan Xu, et al. Hidream-i1: A high-efficient image generative foundation
 525 model with sparse diffusion transformer. *arXiv preprint arXiv:2505.22705*, 2025.

526 Haoyu Chen, Xiaojie Xu, Wenbo Li, Jingjing Ren, Tian Ye, Songhua Liu, Ying-Cong Chen, Lei
 527 Zhu, and Xinchao Wang. Posta: A go-to framework for customized artistic poster generation.
 528 *arXiv preprint arXiv:2503.14908*, 2025a.

529 Jingye Chen, Yupan Huang, Tengchao Lv, Lei Cui, Qifeng Chen, and Furu Wei. Textdiffuser:
 530 Diffusion models as text painters. *Advances in Neural Information Processing Systems*, 36:9353–
 531 9387, 2023a.

532 Jingye Chen, Yupan Huang, Tengchao Lv, Lei Cui, Qifeng Chen, and Furu Wei. Textdiffuser-2:
 533 Unleashing the power of language models for text rendering. *arXiv preprint arXiv:2311.16465*,
 534 2023b.

535 Juhai Chen, Zhiyang Xu, Xichen Pan, Yushi Hu, Can Qin, Tom Goldstein, Lifu Huang, Tianyi
 536 Zhou, Saining Xie, Silvio Savarese, et al. Blip3-o: A family of fully open unified multimodal
 537 models-architecture, training and dataset. *arXiv preprint arXiv:2505.09568*, 2025b.

540 Junsong Chen, Chongjian Ge, Enze Xie, Yue Wu, Lewei Yao, Xiaozhe Ren, Zhongdao Wang, Ping
 541 Luo, Huchuan Lu, and Zhenguo Li. Pixart- σ : Weak-to-strong training of diffusion transformer for
 542 4k text-to-image generation. In *European Conference on Computer Vision*, pp. 74–91. Springer,
 543 2024a.

544 Xiaokang Chen, Zhiyu Wu, Xingchao Liu, Zizheng Pan, Wen Liu, Zhenda Xie, Xingkai Yu, and
 545 Chong Ruan. Janus-pro: Unified multimodal understanding and generation with data and model
 546 scaling. *arXiv preprint arXiv:2501.17811*, 2025c.

548 Zhennan Chen, Yajie Li, Haofan Wang, Zhibo Chen, Zhengkai Jiang, Jun Li, Qian Wang, Jian Yang,
 549 and Ying Tai. Region-aware text-to-image generation via hard binding and soft refinement. *arXiv*
 550 *preprint arXiv:2411.06558*, 2024b.

551 Chaorui Deng, Deyao Zhu, Kunchang Li, Chenhui Gou, Feng Li, Zeyu Wang, Shu Zhong, Wei-
 552 hao Yu, Xiaonan Nie, Ziang Song, Guang Shi, and Haoqi Fan. Emerging properties in unified
 553 multimodal pretraining, 2025. URL <https://arxiv.org/abs/2505.14683>.

555 Nikai Du, Zhennan Chen, Shan Gao, Zhizhou Chen, Xi Chen, Zhengkai Jiang, Jian Yang, and Ying
 556 Tai. Textcrafter: Accurately rendering multiple texts in complex visual scenes. *arXiv preprint*
 557 *arXiv:2503.23461*, 2025.

558 Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam
 559 Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers for
 560 high-resolution image synthesis. In *Forty-first International Conference on Machine Learning*,
 561 2024.

562 Weixi Feng, Wanrong Zhu, Tsu-jui Fu, Varun Jampani, Arjun Akula, Xuehai He, Sugato Basu,
 563 Xin Eric Wang, and William Yang Wang. Layoutgpt: Compositional visual planning and gen-
 564 eration with large language models. *Advances in Neural Information Processing Systems*, 36:
 565 18225–18250, 2023.

566 Yifan Gao, Zihang Lin, Chuanbin Liu, Min Zhou, Tiezheng Ge, Bo Zheng, and Hongtao Xie.
 567 Postermaker: Towards high-quality product poster generation with accurate text rendering. *arXiv*
 568 *preprint arXiv:2504.06632*, 2025.

569 Zigang Geng, Yibing Wang, Yeyao Ma, Chen Li, Yongming Rao, Shuyang Gu, Zhao Zhong, Qinglin
 570 Lu, Han Hu, Xiaosong Zhang, et al. X-omni: Reinforcement learning makes discrete autoregres-
 571 sive image generative models great again. *arXiv preprint arXiv:2507.22058*, 2025.

572 Black Forest Labs. FLUX. <https://github.com/black-forest-labs/flux>, 2024.

573 Naoto Inoue, Kotaro Kikuchi, Edgar Simo-Serra, Mayu Otani, and Kota Yamaguchi. Layoutdm:
 574 Discrete diffusion model for controllable layout generation. In *Proceedings of the IEEE/CVF*
 575 *Conference on Computer Vision and Pattern Recognition*, pp. 10167–10176, 2023.

576 Naoto Inoue, Kento Masui, Wataru Shimoda, and Kota Yamaguchi. Opencole: Towards reproducible
 577 automatic graphic design generation. *arXiv preprint arXiv:2406.08232*, 2024.

578 InternVL Team. Internvl2.5: Expanding performance boundaries of open-source multimodal mod-
 579 els with model, data, and test-time scaling. <https://internvl.github.io/blog/2024-12-05-InternVL-2.5>, 2024.

580 InternVL Team. Internvl3: Advancing open-source multimodal models with native multimodal
 581 pretraining. <https://internvl.github.io/blog/2025-04-11-InternVL-3.0>,
 582 2025.

583 Shuangqi Li, Hieu Le, Jingyi Xu, and Mathieu Salzmann. Enhancing compositional text-to-image
 584 generation with reliable random seeds. *arXiv preprint arXiv:2411.18810*, 2024.

585 Shufan Li, Konstantinos Kallidromitis, Akash Gokul, Arsh Koneru, Yusuke Kato, Kazuki Kozuka,
 586 and Aditya Grover. Reflect-dit: Inference-time scaling for text-to-image diffusion transformers
 587 via in-context reflection. *arXiv preprint arXiv:2503.12271*, 2025.

594 Jiawei Lin, Jiaqi Guo, Shizhao Sun, Zijiang Yang, Jian-Guang Lou, and Dongmei Zhang. Layout-
 595 prompter: awaken the design ability of large language models. *Advances in Neural Information
 596 Processing Systems*, 36:43852–43879, 2023.

597 Zeyu Liu, Weicong Liang, Yiming Zhao, Bohan Chen, Ji Li, and Yuhui Yuan. Glyph-byt5-
 598 v2: A strong aesthetic baseline for accurate multilingual visual text rendering. *arXiv preprint
 599 arXiv:2406.10208*, 2024.

600 Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint
 601 arXiv:1711.05101*, 2017.

602 Chuwei Luo, Yufan Shen, Zhaoqing Zhu, Qi Zheng, Zhi Yu, and Cong Yao. Layoutllm: Layout
 603 instruction tuning with large language models for document understanding. In *Proceedings of the
 604 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 15630–15640, 2024.

605 Yiyang Ma, Xingchao Liu, Xiaokang Chen, Wen Liu, Chengyue Wu, Zhiyu Wu, Zizheng Pan,
 606 Zhenda Xie, Haowei Zhang, Liang Zhao, et al. Janusflow: Harmonizing autoregression and recti-
 607 fied flow for unified multimodal understanding and generation. *arXiv preprint arXiv:2411.07975*,
 608 2024.

609 Yuhang Ma, Yunhao Shui, Xiaoshi Wu, Keqiang Sun, and Hongsheng Li. Hpsv3: Towards wide-
 610 spectrum human preference score. *arXiv preprint arXiv:2508.03789*, 2025.

611 Yuyang Peng, Shishi Xiao, Keming Wu, Qisheng Liao, Bohan Chen, Kevin Lin, Danqing Huang,
 612 Ji Li, and Yuhui Yuan. Bizgen: Advancing article-level visual text rendering for infographics
 613 generation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp.
 614 23615–23624, 2025.

615 Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe
 616 Penna, and Robin Rombach. Sdxl: Improving latent diffusion models for high-resolution image
 617 synthesis. *arXiv preprint arXiv:2307.01952*, 2023.

618 Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea
 619 Finn. Direct preference optimization: Your language model is secretly a reward model. *Advances
 620 in Neural Information Processing Systems*, 36:53728–53741, 2023.

621 Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi
 622 Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text
 623 transformer. *Journal of machine learning research*, 21(140):1–67, 2020.

624 Jaejung Seol, Seojun Kim, and Jaejun Yoo. Posterllama: Bridging design ability of langauge model
 625 to contents-aware layout generation. *arXiv preprint arXiv:2404.00995*, 2024.

626 Noam Shazeer and Mitchell Stern. Adafactor: Adaptive learning rates with sublinear memory cost.
 627 In *International Conference on Machine Learning*, pp. 4596–4604. PMLR, 2018.

628 Zhenxiong Tan, Songhua Liu, Xingyi Yang, Qiaochu Xue, and Xinchao Wang. Ominicontrol: Min-
 629 imal and universal control for diffusion transformer. *arXiv preprint arXiv:2411.15098*, 2024.

630 Zecheng Tang, Chenfei Wu, Juntao Li, and Nan Duan. Layoutnuwa: Revealing the hidden layout
 631 expertise of large language models. *arXiv preprint arXiv:2309.09506*, 2023.

632 Chameleon Team. Chameleon: Mixed-modal early-fusion foundation models. *arXiv preprint
 633 arXiv:2405.09818*, 2024a.

634 Gemini Team, Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut,
 635 Johan Schalkwyk, Andrew M Dai, Anja Hauth, Katie Millican, et al. Gemini: a family of highly
 636 capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023.

637 Kolors Team. Kolors: Effective training of diffusion model for photorealistic text-to-image synthe-
 638 sis. *arXiv preprint*, 2024b.

639 Yuxiang Tuo, Wangmeng Xiang, Jun-Yan He, Yifeng Geng, and Xuansong Xie. Anytext: Multilin-
 640 gual visual text generation and editing. *arXiv preprint arXiv:2311.03054*, 2023.

648 Ideogram AI. Ideogram v2. <https://ideogram.ai/launch>, 2024.
 649

650 Bram Wallace, Meihua Dang, Rafael Rafailov, Linqi Zhou, Aaron Lou, Senthil Purushwalkam,
 651 Stefano Ermon, Caiming Xiong, Shafiq Joty, and Nikhil Naik. Diffusion model alignment using
 652 direct preference optimization. In *Proceedings of the IEEE/CVF Conference on Computer Vision*
 653 and *Pattern Recognition*, pp. 8228–8238, 2024.

654 Zhendong Wang, Jianmin Bao, Shuyang Gu, Dong Chen, Wengang Zhou, and Houqiang Li. De-
 655 signdiffusion: High-quality text-to-design image generation with diffusion models. *arXiv preprint*
 656 *arXiv:2503.01645*, 2025.

657 Chengyue Wu, Xiaokang Chen, Zhiyu Wu, Yiyang Ma, Xingchao Liu, Zizheng Pan, Wen Liu,
 658 Zhenda Xie, Xingkai Yu, Chong Ruan, et al. Janus: Decoupling visual encoding for unified
 659 multimodal understanding and generation. *arXiv preprint arXiv:2410.13848*, 2024.

660 Chenyuan Wu, Pengfei Zheng, Ruiran Yan, Shitao Xiao, Xin Luo, Yueze Wang, Wanli Li, Xiyan
 661 Jiang, Yexin Liu, Junjie Zhou, et al. Omnipgen2: Exploration to advanced multimodal generation.
 662 *arXiv preprint arXiv:2506.18871*, 2025.

663 Xiaoshi Wu, Yiming Hao, Keqiang Sun, Yixiong Chen, Feng Zhu, Rui Zhao, and Hongsheng Li.
 664 Human preference score v2: A solid benchmark for evaluating human preferences of text-to-
 665 image synthesis. *arXiv preprint arXiv:2306.09341*, 2023.

666 Jinheng Xie, Weijia Mao, Zechen Bai, David Junhao Zhang, Weihao Wang, Kevin Qinghong Lin,
 667 Yuchao Gu, Zhijie Chen, Zhenheng Yang, and Mike Zheng Shou. Show-o: One single transformer
 668 to unify multimodal understanding and generation. *arXiv preprint arXiv:2408.12528*, 2024.

669 An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu,
 670 Chang Gao, Chengan Huang, Chenxu Lv, et al. Qwen3 technical report. *arXiv preprint*
 671 *arXiv:2505.09388*, 2025.

672 Ling Yang, Zhaochen Yu, Chenlin Meng, Minkai Xu, Stefano Ermon, and Bin Cui. Mastering text-
 673 to-image diffusion: Recaptioning, planning, and generating with multimodal llms. In *Forty-first*
 674 *International Conference on Machine Learning*, 2024a.

675 Tao Yang, Yingmin Luo, Zhongang Qi, Yang Wu, Ying Shan, and Chang Wen Chen. Posterllava:
 676 Constructing a unified multi-modal layout generator with llm. *arXiv preprint arXiv:2406.02884*,
 677 2024b.

678 Guangcong Zheng, Xianpan Zhou, Xuewei Li, Zhongang Qi, Ying Shan, and Xi Li. Layoutdiffusion:
 679 Controllable diffusion model for layout-to-image generation. In *Proceedings of the IEEE/CVF*
 680 *Conference on Computer Vision and Pattern Recognition*, pp. 22490–22499, 2023.

681 Chunting Zhou, Lili Yu, Arun Babu, Kushal Tirumala, Michihiro Yasunaga, Leonid Shamis, Jacob
 682 Kahn, Xuezhe Ma, Luke Zettlemoyer, and Omer Levy. Transfusion: Predict the next token and
 683 diffuse images with one multi-modal model. *arXiv preprint arXiv:2408.11039*, 2024a.

684 Dewei Zhou, Ji Xie, Zongxin Yang, and Yi Yang. 3dis: Depth-driven decoupled instance synthesis
 685 for text-to-image generation. *arXiv preprint arXiv:2410.12669*, 2024b.

686

687

688

689

690

691

692 **A APPENDIX**

693

694 This is supplementary material for **PosterCraft: Rethinking High-Quality Aesthetic Poster Generation in a Unified Framework**.

695

696

697 We present the following materials in this supplementary material:

698

699

- Sec.A.1 Detailed dataset construction on Text-Render-2M.

700

701

- Sec.A.2 More information about the automatic processing pipeline for HQ-Poster-100K.

- 702 • Sec.A.3 More examples and explanations of the Poster-Preference-100K.
- 703
- 704 • Sec.A.4 Detailed illustration of the Poster-Reflect-120K.
- 705
- 706 • Sec.A.5 Gemini for OCR calculation and preference evaluator.
- 707
- 709 • Sec.A.6 Discussion about reflection and instruction-based editing.
- 710
- 711 • Sec.A.7 An extra evaluation strategy comparing PosterCraft with other methods.
- 712
- 713 • Sec.A.8 More performance comparisons of PosterCraft and other methods on different
- 714 task benchmarks.
- 715
- 716 • Sec.A.9 More key ablation experiments and effectiveness demonstrations of our frame-
- 717 work.
- 718
- 719 • Sec.A.10 Additional visual results and comparisons generated by our PosterCraft.
- 720
- 721 • Sec.A.11 Limitations.
- 722
- 724 • Sec.A.12 Future work.
- 725
- 726 • Sec.A.13 Usage of LLM.
- 727
- 728 • Sec.A.14 Generalizability beyond Flux-dev: Experiments on SD3.5.
- 729
- 730 • Sec.A.15 Why a unified PosterCraft framework works.
- 731
- 732 • Sec.A.16 Dataset availability and statistical analysis of our main datasets.
- 733
- 734
- 735

A.1 TEXT-RENDER-2M CONSTRUCTION PIPELINE

736 In this section, we provide a detailed overview of the automated construction process behind Text-
 737 Render-2M, a large-scale synthetic dataset designed to improve text rendering quality in the baseline
 738 model. This dataset plays a critical role in the text rendering optimization stage of our workflow. By
 739 overlaying diverse textual elements onto high-resolution background images, it allows the model to
 740 learn accurate text generation while preserving its ability to represent rich visual content.

741 Multi-Instance Text Rendering Each image contains a variable number of independently placed text
 742 instances, typically ranging from one to three. This multi-instance setup introduces compositional
 743 complexity and better reflects natural layouts found in aesthetic posters.

744 Text Content Generation Text content is synthesized using a hybrid strategy:

- 745 • A majority of the samples are generated using template-based grammars, with phrases
 746 constructed from predefined structures filled with curated vocabulary lists.
- 747 • A smaller portion uses random alphanumeric strings to simulate noisy or unstructured tex-
 748 tual inputs.
- 749
- 750
- 751

752 The generator supports rich variations in punctuation, casing , and structure (e.g., both single-word
 753 and short-phrase constructions), ensuring linguistic diversity.

754 Font Selection and Style Variability Fonts are randomly drawn from a categorized library contain-
 755 ing both standard and artistic typefaces. When multiple styles are available, a roughly even split



Figure 8: More samples are shown, which are high-quality paired samples in Text-Render-2M.

is enforced between classic and stylized fonts. The system filters fonts that do not support lower-case letters to avoid invalid renderings. This selection mechanism ensures visual variability while maintaining text legibility.

Layout and Placement Strategy Text is positioned using a 3×3 grid-based partitioning scheme (e.g., top-left, center, bottom-right). Before final placement, bounding box collision checks are performed. If an overlap is detected, the system will retry placement with different positions or reduced font sizes, often within 3–5 attempts per instance. This strategy enables dense yet legible arrangements while minimizing visual clutter.

Orientation and Alignment Orientation options include:

- Horizontal,
- Vertically rotated (rotated 90 degrees),
- Vertically stacked (one character per line).

In the horizontal mode, a small proportion of texts receive a random rotation, and a subset of longer text fragments are automatically wrapped across multiple lines. Alignment is randomly selected among left, center, or right justification.

Prompt Generation Each image is paired with a structured natural language prompt. Prompts include:

- The text content,
- Position (e.g., “bottom right”),
- Orientation (e.g., “vertically stacked”),
- Color category, and
- Optionally (included in 50% of cases), the font style.

When multiple texts are rendered, their prompts are numbered and concatenated. If no text is successfully rendered, a fallback prompt indicating the absence of text is generated. We provide a number of samples to view in Fig.8.

A.2 AUTOMATIC PROCESSING PIPELINE FOR HQ-POSTER-100K

For collecting high-quality poster images, we design an automated image filtering and annotation pipeline for HQ-Poster-100K utilizing three MLLMs:

(1) MLLM Scorer: We employ Internvl2.5-8B-MPO with a binary classification task to filter out posters containing extensive Credit Blocks, Billing Blocks, or “4K ultrahd” cover texts. The complete prompt is shown in Fig.17.

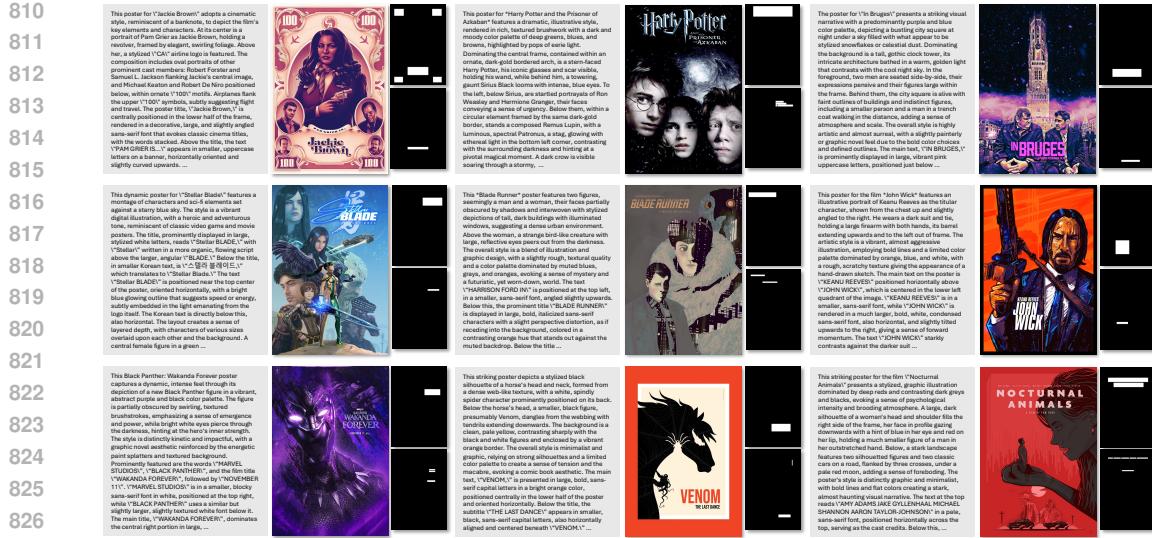


Figure 9: More pictures are shown, which are high-quality paired samples with masks in HQ-Poster-100K.

(2) Gemini Caption Generation: We utilize Gemini2.5-Flash-Preview-04-17 for automated poster caption annotation, systematically describing the poster content, visual style, and textual elements. The complete prompt structure is illustrated in Fig.18

(3) Gemini Mask Generation: The final step in poster data collection involves generating text region masks. We employ Gemini2.5-Flash-Preview-04-17’s robust OCR capabilities to extract text masks from posters, which are subsequently used for Region-aware Calibration. The prompt is shown in Fig.19

In addition, we provide more masks, images, prompts triplets in Fig.9 to demonstrate the advantages of our dataset.

A.3 EXPLANATIONS OF THE POSTER-PREFERENCE-100K

In constructing the Poster-Preference-100K dataset, we first utilize HPSv2 to filter out poster pairs that meet aesthetic requirements and exhibit sufficient diversity. Subsequently, we employ Gemini2.5-Flash-Preview-04-17 to evaluate the alignment between the Best of 5 posters and their corresponding prompts. The complete prompt is shown in Fig.20 Our criteria require that text in the Best of 5 posters must be completely accurate while maximizing alignment with both the content and aesthetic style requirements specified in the prompt. We implement a binary classification system where the model assigns 0 to unqualified samples and 1 to samples meeting all requirements. Fig.10 presents additional preference pairs, demonstrating that our pipeline built upon HPSv2 and Gemini effectively constructs preference data.

A.4 ILLUSTRATION OF THE POSTER-REFLECT-120K

In constructing the Poster-Reflect-120K dataset, we first employ a preference-optimized model to generate 6 posters per prompt, totaling 120K posters. Subsequently, our feedback collection pipeline utilizes Gemini2.5-Flash-Preview-04-17 in two phases: first to select the Best of 6 posters as optimization targets, and second to gather direct feedback.

(1) Best of 6 Selection: We sequentially input six images to Gemini, which selects the optimal image based on predetermined priorities, returning the index number of the best image. The complete prompt is shown in Fig.21.

(2) Feedback Collection: We input two images sequentially - the first being the image requiring improvement, and the second being the optimization target selected in phase (1). Feedback is collected across two dimensions: “Poster Content Suggestions” and “Aesthetic Style Optimization Sugges-

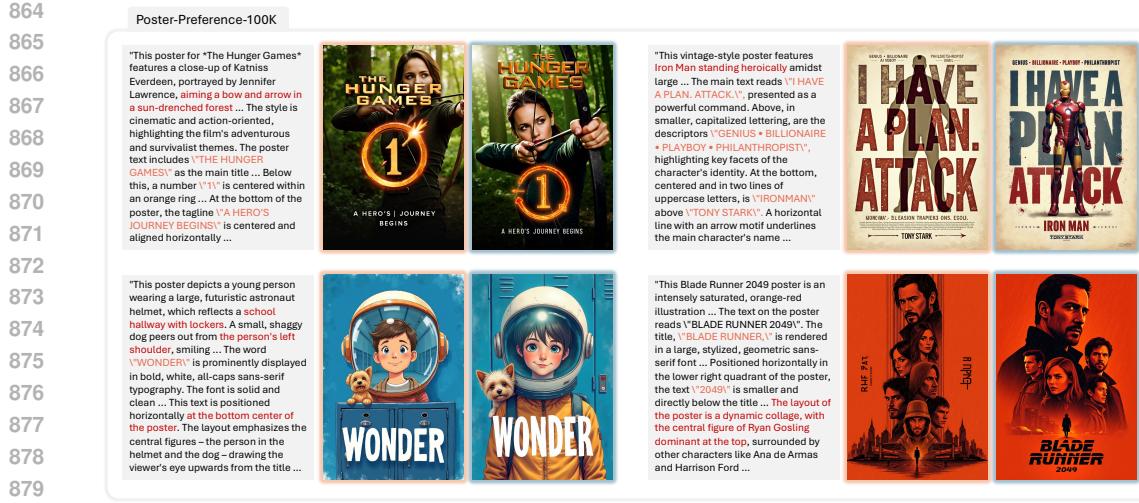


Figure 10: **Additional Preference Pairs** in Poster-Preference-100K. The images on the left are Rejected Samples, while those on the right are Preferred Samples. Orange text indicates textual content, and red text corresponds to content, style, or layout requirements.

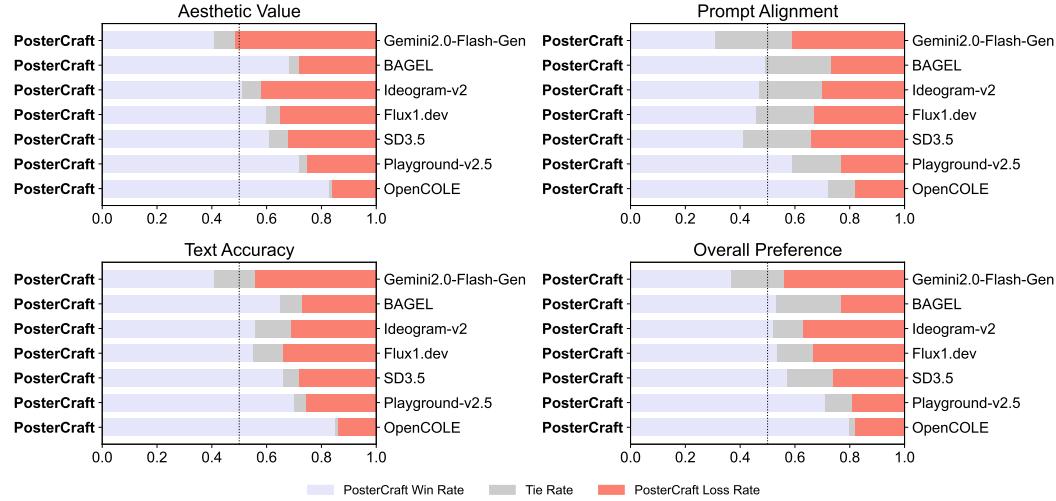


Figure 11: **Gemini serves as an authoritative evaluator for human preference comparisons** across our method and other baselines. PosterCraft outperforms most state-of-the-art generative models in aesthetic coherence, prompt alignment, text rendering, and overall preference. It achieves performance nearly on par with the leading commercial model Gemini2.0-Flash-Gen in text rendering, while showing only a slight gap in other aspects.

tions”. Each image pair receives 10 pieces of feedback through 5 Gemini requests. The complete prompt is shown in Fig.22.

A.5 GEMINI FOR OCR CALCULATION AND PREFERENCE EVALUATOR

During the experimental phase, we employ Gemini2.5-Flash-Preview-05-20 as both an OCR metric calculator and a multi-dimensional preference evaluator. The robust perceptual and comprehension capabilities of Gemini2.5-Flash establish a solid foundation for our experimental accuracy.

(1) OCR Evaluator: Despite the significant challenges that artistic fonts in posters pose to traditional OCR algorithms, Gemini2.5-Flash-Preview-05-20, as a state-of-the-art MLLM model, accurately extracts textual information from images. We further utilize Gemini2.5-Flash’s reasoning and math-

918
 919 **Table 2: Comparison on the our proposed dataset Geng et al. (2025) using HPSv3 Ma et al.**
 920 **(2025).** We highlight the **best** and **second** performance. The upward arrow (\uparrow) indicates that
 921 higher values are better.

Method	HPSv3 \uparrow
OpenCOLE Inoue et al. (2024) (Open)	8.2280
PosterMaker Gao et al. (2025) (Open)	8.5825
Playground-v2.5 Yang et al. (2024b) (Open)	8.7831
BizGen Peng et al. (2025) (Open)	9.4879
SD3.5 AI (2024) (Open)	10.3104
Flux1.dev https://github.com/black-forest-labs/flux (2024) (Open)	10.4066
Ideogram-v2 v2. https://ideogram.ai/launch (2024) (Close)	10.0911
HiDream-II-Full Cai et al. (2025) (Open)	10.4533
BAGEL Deng et al. (2025) (Open)	8.9116
Gemini2.0-Flash-Gen Team et al. (2023) (Close)	10.5251
PosterCraft (ours)	10.7402

933
 934 **ematical** capabilities to directly compute and output OCR metrics. The specific prompt is shown in
 935 Fig.23.

936 (2) Multi-Dimensional Preference Evaluator: Gemini2.5-Flash demonstrates superior assessment
 937 capabilities with its reasoning abilities. We present two images side by side and instruct Gemini2.5-
 938 Flash to select from "L", "R", or "None", representing left image superior, right image superior, or
 939 indeterminate (either both excellent or both inadequate). The specific prompt is shown in Fig.24 and
 940 Fig.25.

941
 942 **A.6 DISCUSSION ON THE DIFFERENCES BETWEEN REFLECTION AND INSTRUCTION-BASED**
 943 **EDITING**

945 While our refinement stage may appear similar to instruction-based editing, the underlying mecha-
 946 nisms and objectives differ substantially. We summarize the key distinctions as follows:

- 948 • **Source of instruction.** In instruction-based editing, the model typically receives explicit,
 949 human-written commands describing object-level changes (e.g., "remove the chair" or
 950 "change background color"). By contrast, our reflection mechanism relies on an offline aes-
 951 thetic model that generates structured feedback from imperfect images. This enables multi-
 952 dimensional improvement, covering aesthetics, typography, and layout balance, rather than
 953 focusing solely on localized object edits or simple manipulations.
- 954 • **Role of the original prompt.** Instruction-based editing often discards or overrides the
 955 initial text-to-image (T2I) prompt. In our design, the original T2I prompt remains central,
 956 while the reflection-derived instruction serves only as auxiliary guidance. This ensures that
 957 semantic alignment with the initial user intent is preserved, while aesthetics and overall
 958 visual quality are enhanced.
- 959 • **Objective difference.** Instruction-based editing aims for literal compliance with user-
 960 specified edits. In contrast, reflection targets holistic enhancement—refining text render-
 961 ing, composition, and aesthetic appeal. In this sense, reflection is closer to an iterative
 962 self-improvement loop rather than one-shot command execution.

963
 964 **A.7 MORE EVALUATION STRATEGY**

966 Inspired by PixArt- Σ Chen et al. (2024a), we further employ Gemini2.0-Flash as an automated eval-
 967 uator to provide strict preference comparisons across models, as illustrated in Fig. 11. Leveraging
 968 carefully designed evaluation prompts, Gemini consistently confirms that PosterCraft achieves su-
 969 perior aesthetic quality, faithful prompt alignment, and accurate text rendering compared to both open-
 970 source and several closed-source baselines. Although slightly behind the commercial Gemini2.0-
 971 Flash-Gen in text rendering, PosterCraft remains competitive overall, demonstrating the robustness
 of our unified workflow.

972
973
974
975
976
977
978
979
980
981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025
Table 3: **Quantitative results on CVTG-2K Du et al. (2025) dataset.** We highlight the **best** and **second** performance for each metric. The upward arrow (\uparrow) indicates that higher values are better.

Model	Word Accuracy \uparrow	NED \uparrow	CLIPScore \uparrow	VQAScore \uparrow
FLUX.1-dev https://github.com/black-forest-labs/flux (2024)	0.4965	0.6879	0.7401	0.8886
AnyText Tuo et al. (2023)	0.1804	0.4675	0.7432	0.6935
TextDiffuser-2 Chen et al. (2023b)	0.2326	0.4353	0.6765	0.5627
RAG-Diffusion Chen et al. (2024b)	0.2648	0.4498	0.6688	0.6397
3DIS Zhou et al. (2024b)	0.3813	0.6505	0.7767	0.8684
TextCrafter Du et al. (2025)	0.7370	0.8679	0.7868	0.9140
PosterCraft (Ours)	0.5814	0.7581	0.7493	0.9007
PosterCraft (Ours) with TextCrafter	0.7544	0.8863	0.7952	0.9241

983
984
985
986
987
988
989
990
991
992
993
994
995
996
Table 4: **Comparison on the LongText-Bench (English) dataset Geng et al. (2025).** We highlight the **best** and **second** performance. The upward arrow (\uparrow) indicates that higher values are better.

Method	LongText-Bench-Eng \uparrow
Janus-Pro Chen et al. (2025c)	0.019
BLIP3-o Chen et al. (2025b)	0.021
Kolors 2.0 Team (2024b)	0.258
PosterMaker (without image input) Gao et al. (2025)	0.315
BAGEL Deng et al. (2025)	0.373
BizGen (random layouts input) Peng et al. (2025)	0.478
HiDream-II-Full Cai et al. (2025)	0.543
OmniGen2 Wu et al. (2025)	0.561
FLUX.1-dev https://github.com/black-forest-labs/flux (2024)	0.607
X-Omni Geng et al. (2025)	0.900
PosterCraft (Ours)	0.631

To further strengthen the aesthetic and prompt alignment preference evaluation, we additionally benchmarked PosterCraft against existing methods using HPSv3 on the prompts from our proposed poster generation benchmark (Tab. 2). In this comparison we also include poster-oriented systems such as **BizGen** and **PosterMaker**, adapting them to our text-only setting by supplying five representative poster layouts and, for **PosterMaker**, disabling the product-image input. PosterCraft achieves the highest HPSv3 score (10.7402), surpassing both open-source and commercial baselines, with Gemini2.0-Flash-Gen ranking second (10.5251). That PosterCraft outperforms these task-specific poster generators while using only a single text prompt further emphasizes the strength of our unified backbone design. These results highlight the superiority of our unified workflow in capturing human-preferred aesthetics beyond text accuracy and layout coherence. Taken together, our evaluation covers a comprehensive spectrum—from user studies, to win-rate analysis with state-of-the-art VLM evaluators, to objective metrics of poster text accuracy and aesthetic quality—demonstrating that PosterCraft delivers outstanding performance in aesthetic poster generation across multiple perspectives.

A.8 MORE BENCHMARK COMPARISONS

Our primary task is poster generation, and in the main manuscript we have already demonstrated that PosterCraft can match closed-source models in that domain. Here we additionally report results on more general text-rendering benchmarks to show that our training workflow also brings consistent gains beyond aesthetic posters. On the partitioned text dataset CVTG-2K Du et al. (2025), PosterCraft achieves Word Accuracy = 0.5814, NED = 0.7581, CLIPScore = 0.7493, and VQAScore = 0.9007. Although we do not outperform methods specifically designed for text rendering (for instance, TextCrafter, which uses attention- and partition-aware strategies), PosterCraft surpasses almost all open-source baselines (such as FLUX.1-dev, AnyText, TextDiffuser-2, RAG-Diffusion, and 3DIS). Moreover, when we apply the training-free TextCrafter optimizer on top of PosterCraft, the combined system (PosterCraft + TextCrafter) further improves all four CVTG-2K metrics and clearly surpasses TextCrafter itself, as shown in Table 3. This confirms that our four-stage training pipeline has already unlocked strong text-rendering capacity in the backbone and remains fully compatible with specialized test-time plug-ins when desired. Notably, we also significantly outperform our base model FLUX, which demonstrates that our proposed training workflow yields substan-



Figure 12: **Aesthetic-text reinforcement learning (top row) and vision-language feedback (bottom row) qualitative comparisons.** orange boxes denote the text biases. Different color fonts represent key feedback information from VLM. The top examples demonstrate that our reinforcement learning stage effectively improves overall aesthetic quality and partially corrects text rendering errors after supervised fine-tuning. The bottom examples illustrate the impact of vision-language reflection, where feedback from a VLM is integrated into the generation loop. This results in noticeable enhancements to both visual aesthetics and semantic coherence in the final poster outputs.

tive improvements even on tasks different from poster generation. This strongly indicates that our workflow possesses robust generalization and resilience.

To further validate PosterCraft’s capability in longer text generation in image contexts, we evaluate it on the LongText-Bench (English) dataset. PosterCraft attains a score of 0.631, ranking second among public models (only X-Omni is ahead with a score of 0.900). Although we do not surpass X-Omni, we outperform many prior methods (such as Janus-Pro, BLIP3-o, Kolors 2.0, BAGEL, HiDream-I1-Full, OmniGen2, FLUX.1-dev, etc.), including the very strong HiDream-I1-Full. For poster-specific baselines *BizGen* and *PosterMaker* (needing layout), we randomly sample one of five predefined layouts per prompt and omit the product image for *PosterMaker*—so that they match the testing form of LongText-Bench. Given that X-Omni is built on a much larger backbone and substantially more text-centric data, this gap is expected and mainly reflects differences in base-model capacity and supervision rather than a limitation of PosterCraft. Yet the fact that we outperform most open models, and significantly surpass our base model, highlights the generalizability and effectiveness of our method. These additional benchmark results serve as supportive evidence, reinforcing the argument that PosterCraft is not limited to poster generation, but also has promise across broader visual text generation tasks.

A.9 MORE ABLATION EXPERIMENTS AND EFFECTIVENESS DEMONSTRATIONS

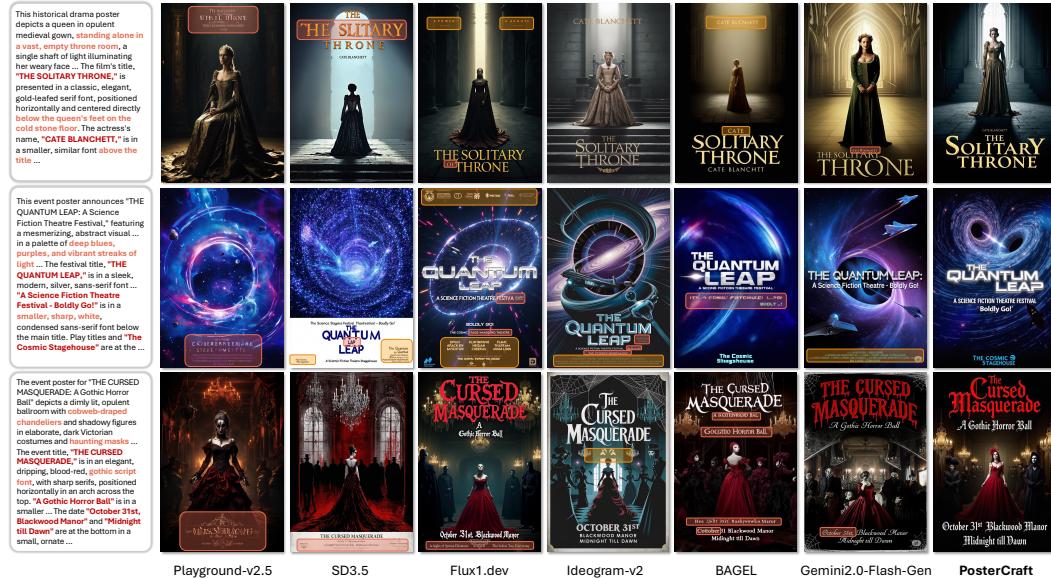
A.9.1 AESTHETIC-TEXT REINFORCEMENT LEARNING AND VISION-LANGUAGE REFLECTION

Aesthetic-Text Reinforcement Learning. The reinforcement learning stage significantly boosts visual appeal and human preference. By optimizing higher-order aesthetic cues (e.g., layout balance, color harmony, typographic cohesion), it enables the model to generate posters that are more visually compelling and coherent. The top-row results in Fig. 12 clearly show these improvements: reinforcement learning corrects text rendering errors left after supervised fine-tuning and enhances holistic preference, confirming its effectiveness.

Vision-Language Reflection. The reflection stage integrates structured feedback from a vision-language model (VLM) into the generation loop. This mechanism allows the system to iteratively refine stylistic integration, semantic alignment, and overall aesthetic appeal. The bottom-row ex-

1080	Stage-4 (Reflection)	Feedback generator	Feedback Quality (BERTScore-F1 \uparrow)	Win rate \uparrow
1081	No (w/o Stage-4)	—	N/A	0.33
1082	Yes	InternVL3 (original)	0.8641 ± 0.0099	0.30
1083	Yes	InternVL3 (fine-tuned)	0.8957 ± 0.0095	0.37

1085 Table 5: Ablation on the VLM used to produce Stage-4 feedback. Fine-tuning the VLM with high-
1086 quality, domain-specific feedback pairs improves both the quality of feedback text and downstream
1087 human preference.



1108 Figure 13: **Visual comparison** of different model outputs. Red boxes highlight misspelled or dis-
1109 torted text, while yellow boxes indicate redundant or missing text elements. Within the prompts,
1110 orange text denotes content and style requirements, and red text indicates textual elements. It in-
1111 dicates that our method significantly outperforms existing SOTA approaches in generating high-
1112 quality posters under long-prompt conditions, with notably improved prompt alignment. In terms of
1113 text rendering, our model produces fonts that align closely with the visual context of the scene, with
1114 minimal rendering errors.

1115 ames in Fig. 12 highlight that reflection leads to more consistent layouts and improved text-
1116 background harmony, particularly in visually complex posters.

1117 These ablations and Fig. 7 in the manuscript demonstrate that both stages—reinforcement learning
1118 and vision-language reflection—are indispensable for unlocking the full potential of foundation dif-
1119 fusion models in end-to-end aesthetic poster generation.

1123 A.9.2 THE NECESSITY OF FEEDBACK MODEL FINE-TUNING

1124 Fine-tuning the VLM used for Stage-4 feedback substantially improves both the feedback text and
1125 downstream generation quality. Compared with the original VLM, our fine-tuned version yields
1126 higher feedback quality (BERTScore-F1 0.8957 ± 0.0095 vs. 0.8641 ± 0.0099 , $+0.0316$ absolute,
1127 $\sim 3.66\%$ relative). A higher BERTScore-F1 indicates better semantic alignment between our gen-
1128 erated feedback and reference (high-quality) feedback: it means that the fine-tuned VLM’s feedback
1129 is more semantically faithful to the gold feedback, and it increases the human win rate of gener-
1130 ated posters from 0.30 to 0.37 ($+0.07$ absolute, $\sim 23.33\%$ relative). Notably, using the original
1131 VLM hurts Stage-4 compared with removing reflection entirely (0.30 vs. 0.33), indicating that non-
1132 domain or ill-formatted feedback can misguide the reflection loop. Enforcing a structured schema
1133 and training the VLM on high-quality, domain-specific feedback pairs is therefore essential for sta-
1134 ble, effective reflection.



Figure 14: **Visual comparison** of different model outputs. Red boxes highlight misspelled or distorted text, while yellow boxes indicate redundant or missing text elements. Within the prompts, orange text denotes content and style requirements, and red text indicates textual elements. Compared to other methods, our approach produces cleaner layouts, better theme alignment, and more accurate text rendering under short prompts.

A.10 ADDITIONAL VISUAL COMPARISON AND EXAMPLES

In the long-prompt setting of Fig.13, PosterCraft demonstrates clear advantages in both prompt alignment and aesthetic quality. Our model consistently integrates complex scene descriptions with textual elements, producing posters that closely follow the narrative and stylistic cues in the prompts. Competing models often struggle with either missing or distorted text (e.g., Playground-v2.5, SD3.5, Flux1.dev, Gemini2.0-Flash-Gen), or fail to faithfully capture scene details (e.g., BAGEL, Ideogram-v2). By contrast, PosterCraft generates outputs with coherent layouts, accurate text rendering, and strong adherence to detailed style requirements, achieving more natural visual-text alignment even under demanding long-prompt conditions.

In the short-prompt setting of Fig.14, PosterCraft maintains a strong balance between visual appeal and accurate text rendering. It consistently integrates title and scene elements, for example, seamlessly embedding 'PixelPlay Retro Game Console' or 'International Animation Festival' in stylized compositions. Competing models often display legibility issues (e.g., SD3.5, Flux1.dev), omit key poster elements (e.g., BAGEL, Ideogram-v2), or suffer from text-scene disconnection (e.g., Playground-v2.5). While Gemini2.0-Flash-Gen performs well in text rendering, it still suffers from aesthetic limitations, often producing visually monotonous outputs with missing or underdeveloped design elements. Our results stand out with their vibrant layout, theme adherence, and natural visual-text coherence, even under minimal input conditions.

We provide additional visual results generated by our PosterCraft framework to further demonstrate its capability in producing high-quality, aesthetically consistent posters. As shown in Fig.15 and Fig.16, our model is able to seamlessly integrate text and imagery without requiring any external layout templates or modular refinement. The generated posters exhibit strong visual coherence—text elements are not only stylistically aligned with the visual content but are sometimes cleverly embedded into the composition, enhancing the overall design fluency. From an aesthetic perspective, the model captures genre-specific styles across diverse themes, such as cinematic sci-fi, educational charts, cultural festivals, and commercial advertisements. It achieves a fine balance between visual richness and layout readability, effectively modeling principles such as symmetry, emphasis, and hierarchy. These results underscore PosterCraft’s potential as a powerful end-to-end tool for automatic poster generation with minimal input while maintaining high visual and semantic fidelity.



Figure 15: **Examples generated by our PosterCraft** demonstrating high diversity and aesthetic quality across themes including education, entertainment, and science fiction. All generation results showcase genre-specific fidelity, text rendering, and layout aesthetic.

A.11 LIMITATIONS

In this work, we propose PosterCraft to explore how unified workflow design can unlock the aesthetic design potential of powerful foundation models. Our results demonstrate that with carefully crafted design strategies, the model’s capabilities can be significantly enhanced—making it competitive with leading proprietary commercial systems. This validates the soundness of our motivation. However, our approach is not without limitations. Specifically, our model is fundamentally built upon the current flux1.dev baseline. As such, if the pre-trained flux model has never encountered certain types of samples or contains significant flaws, our method may not be able to fully correct these shortcomings. That said, our workflow is highly unified and readily transferable to stronger baselines, ensuring full compatibility with other models in the community.

A.12 FUTURE WORK

In future work, we plan to enhance our unified workflow in several directions. First, we will explore integrating PosterCraft with stronger backbone models and broader pretraining data, so that our four-stage pipeline can further close the gap to the most advanced commercial systems while remaining architecture-agnostic. Second, we aim to scale up training with larger and more diverse datasets, and to adapt our text-aesthetic preference design to product-poster scenarios where both product images and explicit layout hints are provided, making PosterCraft complementary to frameworks such as PosterMaker. Third, we plan to extend our unified pipeline beyond aesthetic posters to other text-heavy domains (e.g., infographics and dense scene text) by redesigning the data and text-specific rewards while still avoiding heavy modular architectures. Finally, we will push towards multilingual (Chinese/English and more) poster generation, where cross-language typography and



Figure 16: **Examples generated by our PosterCraft** demonstrating high diversity and aesthetic quality across themes including movies, product, and virtual reality. All generation results showcase genre-specific fidelity, text rendering, and layout aesthetic.

layout bring additional challenges but also offer an important real-world testbed for the scalability and cross-cultural applicability of our framework.

A.13 USAGE OF LLM

In preparing this manuscript, we employed large language models exclusively for editorial purposes, such as polishing grammar, enhancing readability, and shortening overly complex sentences. At no stage were ideas, methodologies, experiments, results, figures, or references generated by an LLM, nor was unverifiable content introduced. All scientific contributions—including the design of the study, implementation of methods, execution of experiments, and validation of findings—were entirely carried out and verified by the authors, who retain full responsibility for the final content.

Prompt A.1 (MLLM Scorer Prompt)

Does this poster contain a large Billing Block or Credit Block at the bottom or "4K ultrahd" text at the top?

Based on your judgment, use the closest option to answer, and only return the label:

A. Yes. There is a large Billing Block or Credit Block at the bottom or "4K ultrahd" text at the top.

B. No. There is no Billing Block or Credit Block at the bottom and no "4K ultrahd" text at the top.

Figure 17: **Prompt** for MLLM Scorer in HQ-Poster-100K.

1296
 1297
 1298
 1299
 1300
 1301
 1302
 1303
 1304
 1305
 1306
 1307
 1308
 1309
 1310
 1311
 1312
 1313

1314
 1315
 1316
 1317
 1318
 1319
 1320
 1321
 1322
 1323
 1324
 1325
 1326
 1327
 1328
 1329
 1330

1331
 1332
 1333
 1334
 1335
 1336
 1337
 1338
 1339
 1340
 1341
 1342
 1343
 1344
 1345
 1346
 1347
 1348
 1349

Prompt A.2 (Gemini Caption Generation)

Please write a structured and detailed caption in a single paragraph for this poster, covering the following five aspects in order:

Poster Content—Describe what is visually depicted.

Poster Style—Describe the visual or artistic tone, such as cinematic, surreal, minimalist, or other distinct aesthetics.

Poster Text—Provide the exact words shown in the image (title, subtitle, slogan, etc.) and their overall communicative intent.

Text Style and Position—Describe the typography in detail, including font style, size, texture, and how it visually blends or contrasts with the background (e.g., carved into a surface, embedded in light, wrapped by natural objects, etc.); also specify where each piece of text is positioned and its orientation angle in the frame.

Layout—Describe how the all elements are arranged to guide the viewer’s focus.

Be specific, descriptive, and cohesive. Keep the response between 200 and 300 words, written as a single paragraph. Avoid listing or enumeration. Do not mention any design tools or generation methods. Write as if for a professional design catalog, highlighting how visual and typographic design choices form a unified and compelling narrative.

Figure 18: **Prompt** for Gemini Caption Generation in HQ-Poster-100K.

1350
1351
1352
1353
1354

Prompt A.3 (Gemini Mask Generation)

1355 Detect all text regions in the image. For each text region, provide its bounding box in
1356 box_2d format [ymin, xmin, ymax, xmax]. The coordinates for each bounding
1357 box must be a list of four integers [ymin, xmin, ymax, xmax], normalized to the
1358 range [0, 1000]. Ensure the box completely covers the text area.
1359

MANDATORY GUIDELINES:

- The box_2d coordinates [ymin, xmin, ymax, xmax] should be integers normalized to 0-1000.
- If no text is found in the image, the "text_regions" list in the JSON output should be empty.

STRICT CONSTRAINTS:

- Adhere strictly to the JSON output format specified below.
- Do not include any explanations, apologies, or conversational text outside of the JSON structure.
- Ensure the provided normalized coordinates are accurate.

RESPONSE FORMAT:

- Respond with a single JSON object. Do NOT use markdown (e.g., ````json ...`).
- The JSON object must have a single key "text_regions".
- The value of "text_regions" must be a list of bounding_boxes.
- Each bounding_box must be a list of four integer coordinates [ymin, xmin, ymax, xmax], normalized to [0, 1000].
- Example of the required JSON structure for "text_regions" containing two bounding boxes:

```
[  
  [ymin1, xmin1, ymax1, xmax1],  
  [ymin2, xmin2, ymax2, xmax2]  
]
```

- The complete JSON object should look like this:

```
{  
  "text_regions": [  
    // List of bounding_boxes as shown above  
    // e.g., [[20, 10, 50, 100], [70, 150, 100, 250]]  
  ]  
}
```

- If no text is found, the output should be: `{"text_regions": []}`.
- Provide ONLY this JSON object.

1396 Now, based SOLELY on your comprehensive image analysis, provide ONLY the JSON
1397 object detailing all detected text regions and their normalized box_2d coordinates [ymin,
1398 xmin, ymax, xmax] as specified.
1399

1400
1401
1402
1403

Figure 19: **Prompt** for Gemini Mask Generation in HQ-Poster-100K.

1404
1405
1406
1407
1408
1409
1410
1411
1412
1413

Prompt A.4 (Prompt Alignment Evaluation)

1414
1415
1416
1417
1418
1419
1420
1421
1422
1423
1424
1425
1426
1427
1428
1429
1430
1431
1432
1433
1434
1435
1436
1437
1438
1439
1440
1441
1442
1443
1444
1445
1446
1447
1448
1449
1450
1451
1452
1453
1454
1455
1456
1457

You are an expert in evaluating image content and font style against a given text prompt. You will be given an image and an original text prompt that was intended to generate an image similar to the one provided. Your task is to assess whether the image is substantially consistent with the original text prompt based on the criteria below.

Original Text Prompt: "{original_prompt_text}"

Evaluation Criteria (Prioritized):

1. Text Accuracy:

- Thoroughly analyze all text visible in the image. Check for any inaccuracies such as typos, missing characters/words, or extra characters/words when compared to the "Original Text Prompt". This is the MOST CRITICAL factor. If ANY such error is found, the decision MUST be "0".

2. Text Style and Positioning:

- If text is present, does its style (font, color, decoration) and positioning (layout, orientation) in the image reasonably align with what is described or implied in the "Original Text Prompt"?

3. Overall Content, Artistic Style, and Visual Appeal:

- Does the overall image content (subjects, scene, objects) and artistic style align well with the "Original Text Prompt"?
- Is the image generally clear, well-composed, and visually appealing in the context of the prompt?

Output Format: Based on your assessment, output ONLY a JSON object in the following format: `{{"final_decision": "1"}}` if the image is substantially consistent with the original prompt across the prioritized criteria (especially if no text errors are found when text is intended) and should be kept. `{ {"final_decision": "0"}}` if there are ANY discrepancies in Text Accuracy (typos, missing/extraneous characters/words), or significant issues in other critical criteria, or overall poor alignment, meaning the image should be discarded.

Strict constraints:

- Only output the JSON object.
- Do NOT include any additional text, explanation, or markdown.
- Use exactly "0" or "1" as the value for "final_decision".

Figure 20: **Prompt** for Prompt Alignment Evaluation in Poster-Preference-100K

1458
1459
1460
1461
1462
1463

Prompt A.5 (Best-of-6 Selection)

1464
1465 You are a professional Poster Designer. Your task is to evaluate six generated posters based
1466 on a design brief ("Original Prompt") and select the single best poster, or indicate if none
1467 are suitable.

Evaluation Process:

1468
1469 1. **Textual Accuracy (Paramount Importance):**
1470 • First, assess all posters for textual accuracy against the "Original Prompt".
1471 Text (if any is specified or implied by the brief) **MUST** be **perfectly accurate**:
1472 – No typographical errors.
1473 – No missing or extra characters/words.
1474 • **A poster with any textual flaw cannot be chosen as the best IF an alterna-**
1475 **tive poster with perfect text exists.**

1476 2. **Content Alignment and Aesthetic Value:**

1477 • This criterion is used to select among posters that have passed the textual
1478 accuracy check.
1479 • The chosen poster should:
1480 – Provide content as close as possible to the "Original Prompt".
1481 – Demonstrate the highest possible aesthetic value (considering composi-
1482 tion, color palette, typography, imagery, and overall visual impact).
1483

Selection Logic:

1484
1485 • **Ideal Case:** If one or more posters have **perfect textual accuracy**, select from
1486 THIS group the single poster that best meets Criterion 2 (Content Alignment and
1487 Aesthetic Value).
1488 • **Special Case (All Posters Have Textual Flaws):** If ALL six posters have some
1489 textual inaccuracies, then no poster meets the primary standard for "best." In this
1490 situation, you **MUST** output "none".
1491 • **Fallback Case (This should ideally not be reached if "Special Case" is handled**
1492 **correctly):** If the logic leads here unexpectedly after "Special Case" consideration,
1493 and no poster has perfect text, but a selection is still forced, choose the poster
1494 that, despite its textual flaws, is superior when evaluated **SOLELY** on Criterion 2
1495 (Content Alignment and Aesthetic Value across all six flawed images). However,
1496 prioritize outputting "none" if all have text flaws.

1497 Original Prompt (Design Brief): ``original_prompt``

1498 Select the image ("1", "2", "3", "4", "5", "6", or "none") that best meets these requirements.

1499 Respond **ONLY** with a JSON object in ONE of these exact formats: `>{"best_image": "1"}` OR `>{"best_image": "none"}`

1500 Strict constraints:

1501
1502 • Only output the JSON object.
1503 • Do NOT include any additional text or markdown.
1504 • Use exactly "1", "2", "3", "4", "5", "6", or "none" to refer to your selection.
1505

1506
1507
1508
1509
1510
1511

Figure 21: **Prompt** for Best-of-6 Selection in Poster-Reflect-120K.

1512
1513
1514
1515
1516
1517
1518
1519
1520
1521
1522
1523
1524
1525
1526
1527

Prompt A.6 (Feedback Collection)

1528
1529
1530
1531
1532
1533

Internally compare the first poster against the second poster, focusing strictly on visual content layout and overall aesthetic style. Based on this internal comparison, provide detailed and specific suggestions in two aspects: 1. Poster Content Suggestions 2. Aesthetic Style Optimization Suggestions. Act as a professional poster designer. Deliver highly detailed, specific, and actionable feedback in the form of standardized image editing instructions.
MANDATORY GUIDELINES:

1534
1535
1536
1537
1538

- The second poster must be fully followed as the standard. Identify and correct all visual layout and style discrepancies based on this reference.
- Focus exclusively on content and visual/aesthetic design. Completely ignore any issues related to text, typography, wording, spelling, rendering, or text styling.

1539
1540

STRICT CONSTRAINTS:

1541
1542
1543
1544

- NEVER mention the second poster, reference, or target.
- NEVER use comparative phrases such as "similar to the second poster" or "make it like the second poster".
- ONLY describe the editing instructions for Poster 1, framed as standalone improvement tasks.

1545
1546
1547
1548
1549

RESPONSE FORMAT: Response should be formatted as clearly structured json schema: {`Poster Content Suggestions`: str, `Aesthetic style optimization suggestions`: str} Return ONLY the JSON object itself, without any introductory text or markdown formatting.

1550
1551
1552
1553
1554
1555
1556
1557
1558
1559
1560
1561
1562
1563
1564
1565

Figure 22: **Prompt** for Feedback Collection in Poster-Reflect-120K.

1566
1567
1568
1569
1570

Prompt A.7 (OCR Evaluation)

1571 You are an OCR evaluation assistant. Follow these steps exactly on the attached image:

1. Ground-Truth Extraction (from the design prompt only):
 - Do NOT read text from the image for GT.
 - Parse ONLY the following design prompt and extract ALL text strings that should appear on the poster (titles, subtitles, dates, slogans, venue, etc.), preserving spaces and punctuation exactly: `original_prompt_text`
 - Order them in spatial sequence (top→bottom, left→right) and concatenate into **raw GT text**.
2. OCR Extraction (from the attached image):
 - Run OCR on the provided image and extract ALL rendered text exactly as it appears.
 - Preserve visual reading order (top-left→bottom-right). This is your **raw OCR text**.
3. Text Normalization (apply to BOTH raw GT and raw OCR before comparison):
 - Convert all letters to lowercase.
 - Remove ALL punctuation characters: . , ; : ! ? ' " - () [] { } . . . `
 - Collapse any sequence of whitespace/newlines into a single space.
 - Trim leading and trailing spaces.
4. Character-Level Alignment & Error Counting:
 - Align the **normalized** GT text and OCR text **character by character**.
 - Count four categories:
 - **Insertion (I)**: extra character in OCR not in GT ("more").
 - **Deletion (D)**: GT character missing in OCR ("less").
 - **Substitution (S)**: OCR character differs from GT character ("render error").
 - **Correct match (C)**: identical characters.
5. Metrics Calculation:
 - Let N = total normalized GT characters = $C + D + S$.
 - Let P = total normalized OCR characters = $C + I + S$.
 - Let T = total compared characters = $C + I + D + S$.
 - **Character Accuracy** = C / T .
 - **Text Precision** = $C / (C + I + S)$.
 - **Text Recall** = $C / (C + D + S)$.
 - **Text F-score** = $2 * \text{Precision} * \text{Recall} / (\text{Precision} + \text{Recall})$.
6. Final JSON Output (strictly this format, no extra keys or commentary):


```
{
    "GT_text": "<normalized GT text>",
    "OCR_text": "<normalized OCR text>",
    "total_GT_chars": N,
    "correct_chars": C,
    "insertions": I,
    "deletions": D,
    "substitutions": S,
    "accuracy": "XX.XX%",
    "precision": "YY.YY%",
    "recall": "ZZ.ZZ%",
    "f_score": "WW.WW%"
  }
```

1616
1617
1618
1619

Figure 23: **Prompt** for OCR Evaluation.

1620
1621
1622
1623
1624

1625 **Prompt A.8 (Preference Evaluation)**

1626 Your task is to evaluate a single input image containing two sub-images side-by-side (Left:
1627 Right: R), both generated from the "Original Prompt". Compare them on Aesthetic Value,
1628 Prompt Alignment, Text Accuracy, and Overall Preference.

1629 **General Evaluation Protocol:** For each of the four categories:

1630 1. Provide a brief textual analysis justifying your choice.
1631 2. Make a definitive choice: "L" (Left is superior), "R" (Right is superior), or "none".

1633 **When to Choose "none":** You **must select "none"** for a category if:

1634 a) L and R are tied or indistinguishable in quality for that category.
1635 b) The category is not applicable.
1636 c) After careful review, you cannot definitively determine a superior side.
1637 d) **Crucially: L and R exhibit clear, offsetting strengths and weaknesses *within that specific category*. If L excels in one aspect of the category while R excels in another, and these trade-offs make declaring an overall winner for that category difficult or misleading, choose "none". Do not attempt to weigh these distinct, offsetting pros and cons to force a preference.**

1643 Your careful judgment is vital.

1644 Original Prompt: ``original_prompt``

1645 Please provide your evaluation in the JSON format specified below.

1646 1. **Aesthetic Value:**

- 1647 • Evaluate visual appeal: harmony and consistency of background style, text
1648 style (if present), thematic consistency between background/text, overall con-
1649 tent/text layout, and how the artistic style (background, content, text) aligns
1650 with the "Original Prompt".
- 1651 • `aesthetic_value_explanation`: Your brief analysis.
- 1652 • `aesthetic_value`: Choose "L", "R", or "none" (if L/R are equally
1653 pleasing/coherent, a choice is impossible, or they exhibit offsetting aesthetic
1654 strengths/weaknesses as per the "When to Choose 'none'" protocol).
- 1655 • Respond with: `{ "aesthetic_value": "L/R/none", "aesthetic_value_explanation": "Your analysis..." }`

1657 2. **Prompt Alignment (excluding text elements and artistic style):**

- 1659 • Evaluate how well non-textual elements (subjects, objects, scene) in L and R
1660 match the "Original Prompt".
- 1661 • `prompt_alignment_explanation`: Your brief analysis.
- 1662 • `prompt_alignment`: Choose "L", "R", or "none" (if L/R align equally
1663 well/poorly, it's too close to call, or they exhibit offsetting strengths in align-
1664 ment as per the "When to Choose 'none'" protocol).
- 1665 • Respond with: `{ "prompt_alignment": "L/R/none", "prompt_alignment_explanation": "Your
1666 analysis..." }`

1668 Figure 24: **Prompt** for Preference Evaluation.

1669
1670
1671
1672
1673

1674

1675

1676

1677

1678

1679

1680

1681

1682

1683

1684

1685

1686

1687

1688

1689

1690

1691

1692

1693

1694

1695

1696

1697

1698

1699

1700

1701

1702

1703

1704

1705

1706

1707

1708

1709

1710

1711

1712

1713

1714

1715

1716

1717

1718

1719

1720

1721

1722

1723

1724

1725

1726

1727

Prompt A.9 (Preference Evaluation)**1. Text Accuracy (if applicable):**

- Evaluate text in L and R based *only* on textual content specified/implied in the "Original Prompt".
- Focus *only* on:
 - **Accuracy:** All prompt-specified words/characters present, no typos/mis-spellings/alterations?
 - **Recall:** All intended textual elements from prompt included? Any missing words/phrases?
- **Ignore text style, font, visual appeal, legibility (unless it prevents determining accuracy/recall), and placement.**
- `text_accuracy_explanation`: Your brief analysis.
- `text_accuracy`:
 - First, determine if "none" is appropriate (as per the general "When to Choose 'none'" protocol, especially if L is better on Accuracy but R on Recall, or vice-versa; or if performance is identical/text N/A).
 - If "none" is not chosen, select "L" if L is demonstrably superior overall in combined text accuracy and recall, or "R" if R is.
- Respond with: `{ "text_accuracy": "L/R/none", "text_accuracy_explanation": "Your analysis..." }`

2. Overall Preference:

- Considering all above aspects (aesthetics, alignment, text accuracy) and any other factors relevant to the "Original Prompt".
- `overall_preference_explanation`: Your brief analysis.
- `overall_preference`: Choose "L", "R", or "none" (if L/R are equally preferred, a choice is impossible, or they present compelling but different and offsetting strengths across categories making neither holistically superior, as per the "When to Choose 'none'" protocol).
- Respond with: `{ "overall_preference": "L/R/none", "overall_preference_explanation": "Your analysis..." }`

Respond ONLY with a single JSON object in the following format: `{ "aesthetic_value": "your_choice_for_aesthetic", "aesthetic_value_explanation": "Your brief analysis for aesthetics...", "prompt_alignment": "your_choice_for_alignment", "prompt_alignment_explanation": "Your brief analysis for prompt alignment...", "text_accuracy": "your_choice_for_text", "text_accuracy_explanation": "Your brief analysis for text accuracy...", "overall_preference": "your_choice_for_overall", "overall_preference_explanation": "Your brief analysis for overall preference..." }` Replace placeholders with your choices ("L", "R", "none") and analyses.

Strict constraints:

- Only output the JSON object.
- No additional text or markdown.
- Each choice value (e.g., "aesthetic_value") must be "L", "R", or "none".
- Explanation fields must contain your textual analysis.

Figure 25: Prompt for Preference Evaluation.

1728

1729

1730

1731

Prompt A.10 (Final Retouching Instructions)

1732

Character

1733

You are a professional image retouching artist tasked with finalizing a single retouching approach based on the user's preferences and previous proposals. Your expertise ensures that the final approach integrates key aspects from different suggestions or follows a single selected approach in full.

1737

Background

1738

The user has reviewed previous retouching approaches and provided feedback or specific instructions for a final retouching plan that aligns with their creative goals.

1739

Ambition

1740

Your goal is to either choose one of the previously proposed approaches that best matches the user's vision or create a new, cohesive retouching approach by combining elements from different suggestions. Ensure that the final approach fully respects the user's instructions and creative intent.

1741

User Instruction

1742

User says: "*{user_instruction}*"

1743

Task

1744

1. Review the provided retouching approaches and the user's feedback or instructions.
2. Decide whether to:

1745

- Select a single approach that fits the user's description.
- Create a new approach that integrates relevant aspects from different suggestions.

1746

3. Final Approach:

1747

- Describe the adjustments to **Light** (exposure, contrast, highlights, shadows, blacks, and whites) and **Color** (temperature, tint, vibrance, and saturation).
- For each adjustment, specify **which objects or areas of the image** are most affected and describe the specific details (e.g., "the intricate carvings on the roof are highlighted by a gentle increase in exposure").
- Explain the expected **visual effect on these objects**, such as "the water reflections appear richer and more defined" or "the sky becomes softer and more inviting."
- For each **individual HSL adjustment** (Red, Orange, Yellow, Green, Cyan, Blue, Purple, Magenta), explain why it is necessary and describe the expected visual change for specific objects (e.g., "the red tones in the window frames become more vivid to emphasize their ornate design").
- Organize the description as a step-by-step plan, indicating the sequence of adjustments.

1748

1749

Guidelines for Description:

1750

1751

1752

1753

1754

1755

1756

1757

1758

1759

1760

1761

1762

1763

1764

1765

1766

1767

1768

1769

1770

1771

1772

1773

1774

1775

1776

1777

1778

1779

1780

1781

Figure 26: **Prompt** for Final Retouching Instructions.

1782
1783
1784
1785
1786 Table 6: Effect of applying the full PosterCraft pipeline to SD3.5. Text F-score and HPSv3 are
1787 computed on our aesthetic text-to-poster prompts. The Gemini win rate vs. SD3.5 reflects overall
1788 preference between SD3.5 with our training pipeline and the vanilla SD3.5 baseline.
1789

Model variant	Text F-score \uparrow	HPSv3 \uparrow	Win rate vs. SD3.5 \uparrow
SD3.5 (baseline)	0.479	10.122	50.0
SD3.5 + full PosterCraft (4 stages)	0.554	10.330	58.1

1790 1791 A.14 GENERALIZABILITY BEYOND FLUX-DEV: EXPERIMENTS ON SD3.5 1792

1793 To verify that our workflow is not tied to Flux-dev, we further apply the same four-stage post-training
1794 pipeline (text rendering \rightarrow poster SFT with region-aware calibration \rightarrow aesthetic-text preference RL
1795 \rightarrow VLM reflection) to Stable Diffusion 3.5-Medium (SD3.5). We keep the SD3.5 architecture and
1796 sampler unchanged and only reuse our datasets and method designs.

1797 Figure 27 presents a visual comparison on several poster prompts: the left column shows raw SD3.5
1798 outputs, and the right column shows results after our PosterCraft pipeline. The optimized posters
1799 exhibit sharper and more consistent text, clearer layout structure, and improved global aesthetics,
1800 while preserving the semantics of the original prompts.

1801 Quantitatively, Table 6 reports text F-score on our text-to-poster prompts, the HPSv3 aesthetic score,
1802 and the Gemini win rate when comparing SD3.5 with and without our post-training. Our pipeline
1803 brings consistent gains on all three metrics, mirroring the trends observed on Flux-dev. This confirms
1804 that the proposed framework is architecture-agnostic and can serve as a general post-training recipe
1805 for modern diffusion backbones, rather than a Flux-specific trick.
1806

1807 A.15 WHY A UNIFIED POSTERCRAFT FRAMEWORK WORKS 1808

1809 Our goal is not to propose a new backbone, but to show that a single diffusion model can learn text,
1810 layout, and visual style for posters through a carefully staged, unified workflow. Traditional poster
1811 systems typically decompose these factors into separate networks or ControlNet-style branches (e.g.,
1812 independent text, layout, and style modules), which increases inference cost and makes global opti-
1813 mization difficult. PosterCraft instead keeps one backbone and upgrades it via four stages that are
1814 all defined on simple image-text pairs.

1815 This unified view is feasible because each stage is poster-specific yet complementary: (i) a text-
1816 rendering stage first calibrates stroke-level accuracy and small-font legibility on synthetic Text-
1817 Render-2M without changing the architecture; (ii) a region-aware SFT stage on HQ-Poster-100K
1818 rebalances losses over major text, minor text, and non-text regions so that the model learns realistic
1819 poster layouts and typography while preserving the strengthened text ability; (iii) an aesthetic-text
1820 preference RL stage optimizes a single preference signal that jointly measures aesthetic quality and
1821 text correctness, instead of focusing on only one of them; and (iv) a customized VLM-reflection
1822 stage feeds critique embeddings back into the prompt space, further refining layout and style with-
1823 out adding heavy controllable signals or liking a general editing models.
1824

1825 Because all four stages share the same backbone and operate on unified supervision, PosterCraft
1826 learns to trade off text, layout, and style within one model rather than across loosely coupled mod-
1827 ules. Empirically, we observe consistent gains over the base FLUX and SD3.5 backbones on both
1828 poster benchmarks and generic text-rendering datasets, suggesting that the unified framework is
1829 not tied to a particular architecture and can serve as a general post-training recipe for high-quality,
1830 text-aware image generation.

1831 A.16 DATASET AVAILABILITY AND STATISTICAL ANALYSIS 1832

1833 **Public availability.** As clarified in the main paper, all datasets in this work are automatically
1834 constructed from publicly available imagery rather than private or sensitive data. Text-Render-2M
1835 is built by synthetically rendering diverse texts onto high-quality real images drawn from publicly
1836 available image datasets (natural scenes, products, artworks, historical photos, etc.). HQ-Poster-
1837 100K is collected from publicly available poster resources with automatic filtering and text-mask

1836 generation. We will release: (i) the complete Text-Render-2M dataset (rendered image + prompt),
 1837 (ii) the full HQ-Poster-100K poster images with prompts and text masks, and (iii) the PosterCraft
 1838 model checkpoints trained on these data, under a non-commercial research license. HQ-Poster-
 1839 100K contains third-party copyrighted materials and is released strictly for non-commercial research
 1840 under a fair-use principle; we do not claim ownership of these materials, and downstream users are
 1841 responsible for complying with applicable copyright regulations.

1842 A.16.1 STATISTICS OF TEXT-RENDER-2M

1843 Text-Render-2M is a synthetic data of rendered texts on real photographic backgrounds. We compute
 1844 corpus-level and scene-level statistics over all rendered samples.

1845 Table 7: Statistics of Text-Render-2M (rendered texts + background images).

1849 Aspect	1850 Category	1851 Portion
1852 Text instances per image	1 instance	43%
	2 instances	36%
	3 instances	21%
1853 Text content type	Template-based grammatical phrases	71%
	Random alphanumeric strings	29%
1854 Text length (words per instance)	1–5 words / very short strings	23%
	Short phrases (\approx 5–10 words)	47%
	Sentence-level text (\approx 10–20 words)	22%
	Long sentences or short paragraphs ($>$ 20 words)	8%
1855 Orientation / layout	Horizontal	61%
	Vertically rotated (90°)	17%
	Vertically stacked (one character per line)	22%
1856 Background scene type (underlying images)	Macro nature & ecology (insects, plants, fungi)	21%
	Everyday outdoors & landscapes	18%
	Urban scenes & architecture	16%
	People & portraits	13%
	Products & still-life	14%
	Artistic paintings & illustrations	10%
	Historical / archival / documentary photos	8%

1860 These statistics show that Text-Render-2M covers multi-instance layouts, several orientation modes,
 1861 and a substantial share of sentence- and paragraph-level text, while the photographic backgrounds
 1862 span diverse natural, urban, artistic, and historical scenes rather than being limited to simple natural
 1863 images.

1864 A.16.2 STATISTICS OF HQ-POSTER-100K

1865 Unlike the synthetic Text-Render-2M data, HQ-Poster-100K consists of real-world aesthetic posters.
 1866 Most posters revolve around a main title, an optional tagline, and occasionally compact auxiliary
 1867 information.

1868 Taken together, these statistics highlight the complementary roles of the two datasets in our training
 1869 pipeline. Text-Render-2M provides controlled diversity in text length, orientation, and cross-domain
 1870 photographic backgrounds, which is crucial for robust text rendering. HQ-Poster-100K, in turn, sup-
 1871 plies fully designed aesthetic posters where layout structure, visual style, and text–image integration
 1872 are jointly optimized by human designers—exactly the style and layout distribution that the later
 1873 stages of PosterCraft are trained to model.

1874
 1875
 1876
 1877
 1878
 1879
 1880
 1881
 1882
 1883
 1884
 1885
 1886
 1887
 1888
 1889

Table 8: Text statistics of HQ-Poster-100K.

Aspect (Text)	Category	Portion	Notes
Text instances per poster	1–2 blocks	≈48%	Mostly only the main title.
	3–4 blocks	≈36%	Title plus one or two short taglines or secondary blocks.
	>4–5 blocks	≈16%	Additional taglines or other information.
Text content type	Titles / named entities	≈46%	Titles, character names, locations, etc.
	Taglines / short phrases	≈34%	Promotional slogans, mood-setting phrases.
	Others	≈20%	Cast lists, logos, copyright lines, release information.
Text length (per block)	1–10 words (short)	≈72%	Short titles and concise taglines.
	10–20 words (medium)	≈21%	Longer slogans or brief descriptions.
	>20 words (long)	≈7%	Mainly dense billing blocks at the bottom.

Table 9: Visual and aesthetic statistics of HQ-Poster-100K.

Aspect (Visual)	Category	Portion	Key descriptors
Visual style / art direction	Cinematic / realistic	≈63%	Dramatic lighting, clear depth of field, photographic look.
	Illustrative / graphic	≈24%	Hand-drawn art, comic/graphic-novel styles, painterly or pop-art treatments.
	Minimalist / symbolic	≈13%	Strong negative space, silhouettes, emblematic icons.
Subject matter / composition	Character portrait	≈43%	One or two close-up faces, expressive emotions.
	Ensemble / group composition	≈26%	Multiple characters in montage or tiered layouts.
	Action or scene-driven	≈19%	Dynamic poses, chase scenes, large landscapes or cityscapes.
	Object- / concept-centric	≈12%	Key prop or symbolic object (helmet, device, broken glass, etc.).



Figure 27: **Visual comparison of poster generation performance.** The left column displays the raw outputs from the original **SD3.5-Medium**. The right column presents the results optimized via our **PosterCraft** pipeline. The text boxes on the right contain the complete input prompts used for generation. For a fair comparison, the random seed is fixed at 0 for all inference steps.

1998

1999

2000

2001

2002

2003

2004

2005

2006

2007

2008

2009

2010

2011

2012

2013

2014

2015

2016

2017

2018

2019

2020

2021

2022

2023

2024

2025

2026

2027

2028

2029

2030

2031

2032

2033

2034

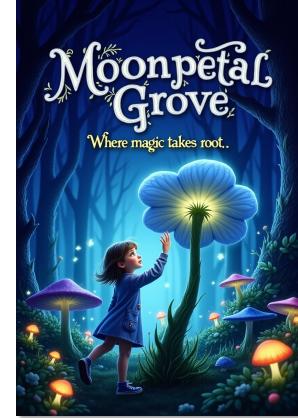
"The whimsical poster for the fantasy film 'Moonpetal Grove' depicts a young girl with luminous blue eyes reaching out to touch a giant, glowing moonpetal flower in an enchanted forest at twilight. The forest is filled with oversized, fantastical flora, faintly glowing mushrooms, and sparkling fireflies, with silhouetted, ancient trees in the deep blue background. The artistic style is richly illustrative and magical, using a palette of deep blues, purples, and bioluminescent greens and yellows to create a dreamlike atmosphere. The film's title, 'Moonpetal Grove,' is rendered in an ornate, flowing, silver script font, with delicate, leafy tendrils entwining the capital letters. The title is arched gracefully across the top of the poster, positioned horizontally. Beneath it, the tagline 'Where magic takes root' is written in a smaller, whimsical gold serif font, also horizontal. The layout guides the viewer's eye from the enchanting title down to the central interaction between the girl and the magical flower. The glowing elements and intricate details invite exploration, promising a journey into a world of wonder and enchantment, perfectly captured by the fantastical typography."



Postermaker



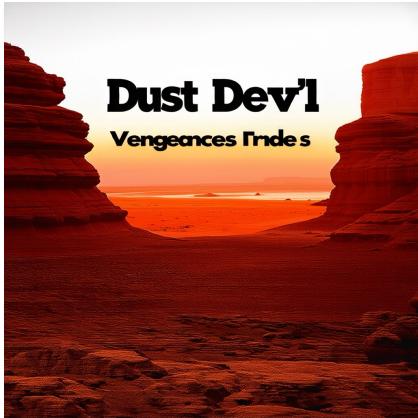
bizgen



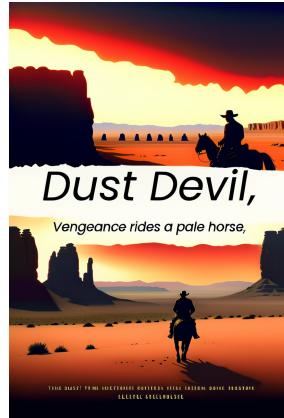
ours

Figure 28: **Qualitative comparison against layout based methods (Part 1).** We present generation results from Postermaker (left), Bizgen (middle), and our **PosterCraft** (right). The corresponding creative briefs are displayed in the red boxes above each example. As illustrated, PosterCraft demonstrates superior performance in strict text rendering, layout arrangement, and aesthetic alignment with the prompt, whereas baseline models frequently exhibit text hallucinations or fail to capture the specified atmosphere.

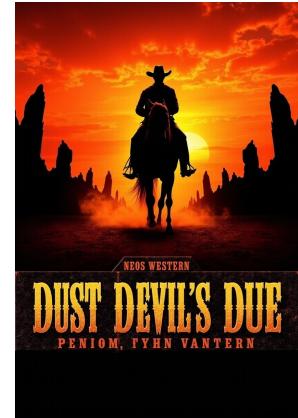
"This poster for the neo-western 'Dust Devil's Due' features a lone, silhouetted cowboy on horseback, riding into a fiery, dust-choked sunset over a desolate desert landscape. Jagged rock formations frame the horizon. The artistic style is stark and cinematic, using a dramatic color palette of deep oranges, reds, and shadowy blacks to create a sense of rugged isolation and impending conflict. The film's title, 'DUST DEVIL'S DUE,' is rendered in a bold, heavily distressed, and spurred serif font, reminiscent of classic Western typography. The letters are a weathered, sandy yellow and appear to be branded or burned into a dark, leather-like strip that runs horizontally across the bottom of the poster. The title is slightly arched. The tagline, 'Vengeance rides a pale horse,' is etched in a smaller, sharper, dark red font above the title. The layout is epic and atmospheric, with the vast desert and dramatic sky dwarfing the lone rider, emphasizing his solitary journey. The rugged, impactful title treatment at the bottom grounds the image and immediately establishes the film's genre and tone."



Postermaker



bizgen



ours

Figure 29: **Qualitative comparison against layout based methods (Part 2).**

2052
2053
2054
20552056 "A post-apocalyptic survival film poster features a lone survivor, heavily cloaked, trudging through a snow-covered, ruined city, a loyal dog by their side.
2057 The sky is a perpetual, toxic yellow-grey. The style is desolate and gritty, emphasizing hardship and resilience, with a muted, cold color palette. The title,
2058 'THE LONG WINTER,' is rendered in a stark, blocky, ice-blue sans-serif font with a cracked texture, positioned horizontally at the very bottom, almost
2059 buried in the snow. The actor's name, 'ANYA TAYLOR-JOY,' is subtly placed in a small, white, weathered font at the top. The layout conveys a sense of
2060 vast emptiness and the arduous journey of the protagonist, with the icy typography reflecting the harsh environment."

Postermaker

bizgen

ours

Figure 30: Qualitative comparison against layout based methods (Part 3).

2073
20742075
2076
2077
2078
2079
2080
2081
20822083
2084
2085
2086

"This epic historical adventure poster shows a fleet of Viking longships sailing through a stormy, dark sea towards a rugged, mist-shrouded coastline. The style is dramatic and painterly, with a sense of rugged power and impending conflict, using a palette of dark blues, greys, and flashes of lightning. The title, 'NORTHWIND SAGA,' is rendered in a heavy, runic-inspired, metallic silver font with a weathered texture, positioned horizontally at the bottom, appearing forged and ancient. The names of the lead actors, 'ALEXANDER SKARSGÅRD' and 'NICOLE KIDMAN,' are in a smaller, classic serif font above the title. The typography strengthens the film's Norse mythological and historical themes."



Postermaker

bizgen

ours

Figure 31: Qualitative comparison against layout based methods (Part 4).

2100
21012102
2103
2104
2105

2106

2107

2108

2109

2110 "A romantic drama poster features a silhouette of a couple embracing against a vibrant, abstract watercolor background of swirling pinks, purples, and blues, representing their emotional world. The style is artistic and emotive, focusing on the feeling of love rather than realistic depiction. The film's title, 'COLORS OF THE HEART,' is written in an elegant, flowing, white script font, positioned diagonally across the center of the watercolor swirl, blending with the artwork. The tagline, 'Love paints its own picture,' is in a smaller, delicate serif font below. The typography is soft and romantic, perfectly complementing the abstract, emotional visuals."

2111

2112



2113

2114

2115

2116

2117

2118

2119

2120

2121

2122

2123

2124

2125

2126

Postermaker

bizgen

ours

Figure 32: Qualitative comparison against layout based methods (Part 5).

2127

2128

2129

2130

2131

2132

2133

2134

2135

2136

2137 "This advertisement for 'Oasis Infusions' botanical gin features a sleek, clear glass bottle of gin, with delicate illustrations of juniper berries, citrus peel, and elderflower visible on the label, surrounded by fresh botanicals and ice cubes in a crystal glass. The background is a softly lit, deep teal. The style is sophisticated, refreshing, and artisanal, emphasizing natural ingredients and premium quality. The brand name, 'Oasis Infusions,' is in a refined, contemporary serif font in a subtle silver, positioned horizontally at the top. The tagline 'A Sip of Serenity' is in a smaller, elegant script font in white, centered below the bottle and glass. 'Crafted with Wild Botanicals' is discreetly printed on the bottle label. The layout is elegant and visually appealing, with the gin bottle and cocktail as the clear focus, and the sophisticated typography reinforcing the brand's premium, nature-inspired identity."



2138

2139

2140

2141

2142

2143

2144

2145

2146

2147

2148

2149

2150

2151

2152

2153

2154

2155

2156

2157

2158

2159

Postermaker

bizgen

ours

Figure 33: Qualitative comparison against layout based methods (Part 6).

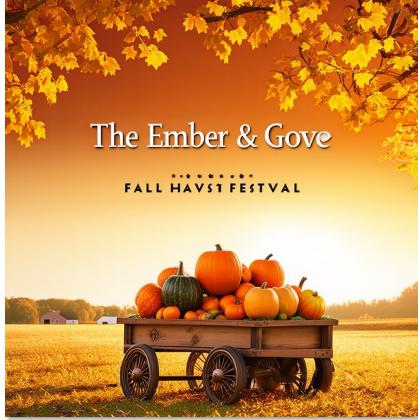
2160

2161

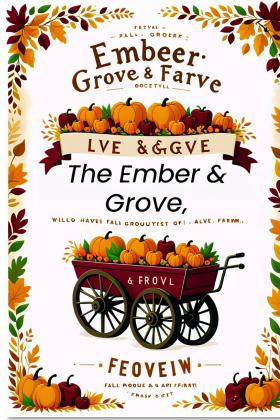
2162

2163

2164 "This poster for 'The Ember & Grove' Fall Harvest Festival depicts a charming, rustic scene: a wooden cart overflowing with pumpkins, apples, and
 2165 gourds, set against a backdrop of golden autumn trees and a distant farmhouse. Warm, dappled sunlight filters through the leaves. The style is warm,
 2166 inviting, and traditional, celebrating the bounty of autumn. The festival title, 'Ember & Grove', is in a friendly, hand-lettered, slightly whimsical serif font in
 2167 a rich burgundy, arcing above the cart. 'Fall Harvest Festival' is in a smaller, rustic sans-serif in dark orange, beneath the title. 'Live Music | Artisan Crafts |
 2168 Fresh Produce | Oct 15-16 | Willow Creek Farm' is at the bottom in a clean, brown sans-serif. The layout is picturesque and abundant, with the harvest
 2169 cart as the centerpiece, and the charming typography enhancing the festival's warm, community-focused atmosphere."



Postermaker



bizgen



ours

Figure 34: Qualitative comparison against layout based methods (Part 7).

2171

2172

2173

2174

2175

2176

2177

2178

2179

2180

2181

2182

2183

2184

2185

2186

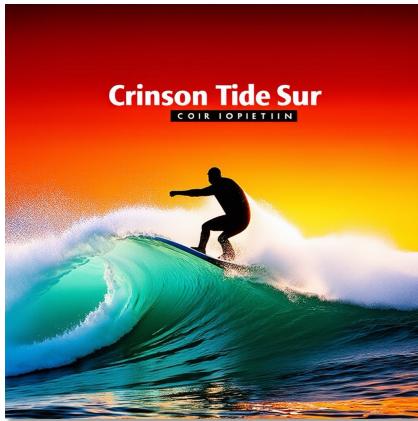
2187

2188

2189

2190

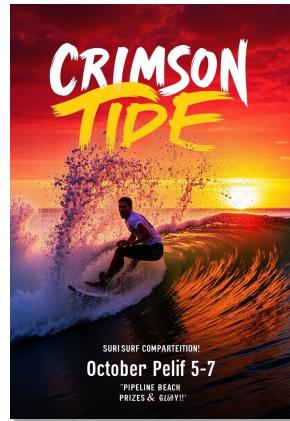
2191 "This poster for the 'Crimson Tide Surf Competition' showcases a dramatic, action-packed photograph of a surfer riding a massive, curling wave,
 2192 silhouetted against a fiery red and orange sunset. Water spray is frozen in mid-air. The style is exhilarating, powerful, and adventurous, capturing the raw
 2193 energy of surfing. The competition title, 'CRIMSON TIDE', is in a bold, dynamic, sans-serif font with a slight wave-like distortion, rendered in a stark white
 2194 that cuts through the vibrant sunset colors, positioned horizontally across the top third of the poster. 'Annual Surf Competition' is in a smaller, sharp
 2195 sans-serif in yellow, beneath the title. 'Pipeline Beach | October 5-7 | Prizes & Glory!' is at the bottom in a clean white sans-serif. The layout is impactful,
 2196 with the surfer and wave dominating the visual field, conveying the intensity of the sport, while the energetic typography enhances the event's thrilling
 2197 and competitive nature."



Postermaker



bizgen



ours

Figure 35: Qualitative comparison against layout based methods (Part 8).

2209

2210

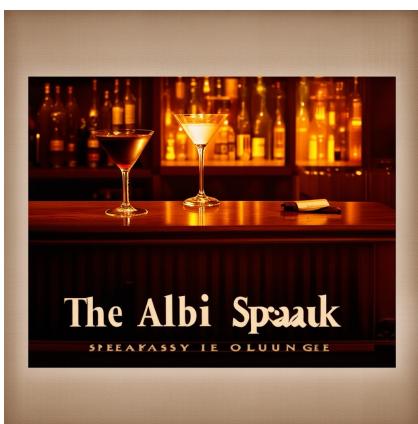
2211

2212

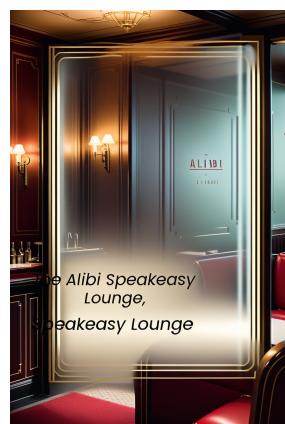
2213

2214
2215
2216
2217
2218
2219
2220
2221
2222
2223
2224
2225
2226
2227
2228
2229
2230

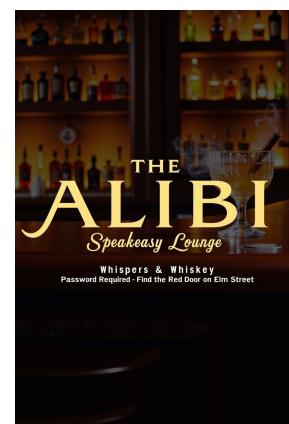
2231 "A poster for 'The Alibi Speakeasy Lounge' features a dimly lit, vintage interior shot: a polished mahogany bar, amber glow from behind bottles, a classic
2232 cocktail in a crystal glass, and a hint of cigar smoke. The style is clandestine, sophisticated, and nostalgic, evoking the Prohibition era. The lounge name,
2233 'THE ALIBI', is in an elegant, art-deco-inspired serif font, in a muted gold with a subtle shadow, positioned horizontally in the center, appearing as if
2234 etched into a frosted glass panel. 'Speakeasy Lounge' is in a smaller, classic sans-serif in cream, beneath the name. 'Whispers & Whiskey | Password
2235 Required | Find the Red Door on Elm Street' is at the bottom in a discreet, white, typewriter-style font. The layout is atmospheric and intriguing, with the
2236 dimly lit bar setting a mysterious tone, and the vintage typography perfectly capturing the speakeasy's exclusive and secretive charm."



2237 Postermaker



2238 bizgen



2239 ours

2240 Figure 36: **Qualitative comparison against layout based methods (Part 9).**

2241
2242
2243
2244
2245
2246
2247
2248
2249
2250
2251
2252
2253
2254
2255
2256
2257
2258
2259
2260
2261
2262
2263
2264
2265
2266
2267