CAN VISION-LANGUAGE MODELS ASSESS GRAPHIC DESIGN AESTHETICS? A BENCHMARK, EVALUATION, AND DATASET PERSPECTIVE

Anonymous authorsPaper under double-blind review

ABSTRACT

Assessing the aesthetic quality of graphic design is central to visual communication, yet remains underexplored in vision—language models (VLMs). We investigate whether VLMs can evaluate design aesthetics in ways comparable to humans. Prior work faces three key limitations: benchmarks restricted to narrow principles and coarse evaluation protocols, a lack of systematic VLM comparisons, and limited training data for model improvement. In this work, we introduce AesEval-Bench, a comprehensive benchmark spanning four dimensions, twelve indicators, and three fully quantifiable tasks: aesthetic judgment, region selection, and precise localization. Then, we systematically evaluate proprietary, open-source, and reasoning-augmented VLMs, revealing clear performance gaps against the nuanced demands of aesthetic assessment. Moreover, we construct a training dataset to fine-tune VLMs for this domain, leveraging human-guided VLM labeling to produce task labels at scale and indicator-grounded reasoning to tie abstract indicators to concrete design regions. Together, our work establishes the first systematic framework for aesthetic quality assessment in graphic design.

1 Introduction

The rapid development of *vision-language models (VLMs)* (Yang et al., 2025; Li et al., 2024a; Hurst et al., 2024) has opened new opportunities for understanding and acting upon multimodal information. While they have already achieved remarkable progress in traditional vision tasks such as image captioning (Lin et al., 2024b; Luo et al., 2024; Li et al., 2025b) and visual question answering (An et al., 2024; Zhang et al., 2024; Lin et al., 2025), VLMs are increasingly expected to contribute to new applications. Among these, *graphic design*—which integrates textual and visual elements to convey information across advertising, branding, and digital media—represents a promising direction given its broad societal and practical impact. Crucially, the success of graphic design critically depends on its *aesthetic quality*, shaped by principles such as balance, contrast, and hierarchy.

In this work, we aim to explore a central question: can VLMs understand and evaluate aesthetic quality of graphic design in a manner comparable to humans? This question is of significant importance for at least three reasons. First, it can assist human designers by identifying where a design falls short and explaining why, thereby enabling more effective improvement. Second, for generative AI systems, it provides the basis for automatic feedback loops that can guide iterative refinement without extensive human intervention. Third, it suggests an opportunity to extend VLMs beyond factual recognition toward aesthetic evaluation.

Despite its importance, research along this direction remains limited (see Table 1). First, benchmarks are inadequate. Those (Zhou et al., 2024; Huang et al., 2024) developed for natural photos often ignore design-specific factors such as typography, while early ones for graphic design typically cover only a narrow subset of design principles (Lin et al., 2023; 2024a; Jiang & Chen, 2025). Moreover, existing evaluation protocols are limited. Scoring-based methods fail to indicate where poor aesthetics occur (Haraguchi et al., 2024), while description-based ones provide qualitative feedback that is difficult to quantify (Jung et al., 2025). Second, comparisons between VLMs are missing. There has been no systematic evaluation across different VLMs, whether open-source or closed-source, in

056

057

060

061

062

063

064

065

066

067

068

069

071

072

073

074

075

076

077

079

081

082 083

084

085

087

880

089

090

091

092

094

095

096

097

098

099

100

101

102

103

104

105

106

107

the context of design aesthetics. Third, *training datasets are lacking*, leaving the question of how to further improve VLM performance in this domain underexplored.

As a first step toward addressing these limitations, we introduce AesEval-Bench, a new benchmark for evaluating the aesthetic quality of graphic designs (see Figure 1). Drawing on prior literature (Li & Chen, 2009; Wangwiwattana & Meeyen, 2024), we identify four critical dimensions—typography, layout, color, and graphics—that together comprehensively capture design aesthetics. These dimensions reflect the major factors that humans consistently emphasize when assessing visual appeal. To provide finer granularity, we further define twelve *indicators* that specify concrete aspects within each dimension, such as hierarchy and legibility under typography. For each indicators, we then design three challenging tasks: 1) aesthetic judgment, which asks models to decide whether a design is aesthetically pleasing (yes/no), providing a straightforward measure of overall perception; 2) region selection, which requires models to choose from candidate regions where unpleasing elements appear, testing their ability to pinpoint problematic areas beyond a global judgment; 3) precise localization, which challenges models to predict the exact bounding box (bbox) coordinates of unpleasing areas, offering the most detailed diagnosis and reflecting a deeper understanding of aesthetics. Unlike prior benchmarks (), AesEval-Bench not only covers a broad range of aesthetic factors through its dimensions and indicators, but also defines evaluation tasks that are fully quantifiable via choice or bbox prediction formats, enabling systematic and reproducible assessment of design aesthetics.

With AesEval-Bench, we systematically evaluate VLMs on their ability to assess the aesthetic quality of graphic designs (see Table 1). For each design, models perform three tasks across twelve indicators. We use human annotation as ground truth, and measure performance by accuracy (for aesthetic judgment and region selection) and bbox IoU (for precise localization). Overall, our results reveal clear gaps between current state-of-the-art VLMs and the nuanced demands of aesthetic quality assessment. Specifically, proprietary VLMs (e.g., GPT series (Achiam et al., 2023)) outperform open-source ones (e.g., Qwen-VL (Wang et al., 2024), Intern-VL (Zhu et al., 2025), LLaVA (Liu et al., 2023)). Among open-source models, larger variants (32B, 72B) generally achieve better performance than smaller ones (3B, 7B). Surprisingly, reasoning-augmented VLMs (e.g., GPT-o1 (Jaech et al., 2024), GPT-o3, Gemini-2.5-Pro (Comanici et al., 2025)) offer no clear advantage over their non-reasoning counterparts. Together, these findings expose the limitations of existing VLMs and underscore the need for domain-specific training tailored to aesthetic quality assessment.

Building on these findings, we move beyond evaluation and turn to training VLMs for aesthetic quality assessment. To this end, we construct a training dataset consisting of three components: the task (what the model is asked to do), the task label (the expected answer), and the reasoning path (the explanation leading to the answer) (see Figure 2). We treat the reasoning path as essential, since generic reasoning has shown little benefit. However, constructing such data poses two key challenges: producing task labels at scale is costly, and generating reasoning paths that genuinely improve performance requires new approaches. To address these, we introduce two solutions. First, human-guided VLM labeling, where a small set of human annotations serve as in-context examples to instruct powerful VLMs in producing task labels. This approach maintains alignment with human understanding while reducing manual annotation costs. Second, indicator-grounded reasoning, where abstract indicators (e.g., hierarchy or layering) are explicitly tied to concrete regions in the design. Each reasoning path consists of bounding-box coordinates linked to the indicator and textual explanations for their relevance, providing fine-grained and interpretable supervision. We fine-tune VLMs with the task as input and both the reasoning path and the task label as supervision, and evaluate on AesEval-Bench. The results show consistent performance gains across all tasks and indicators (e.g., 7.97%, 7.70% and 17.17%), demonstrating that human-guided VLM labeling yields reliable labels and indicator-grounded reasoning supplies effective supervision.

To sum up, our contributions are as follows:

- We introduce AesEval-Bench, a comprehensive benchmark for assessing aesthetic quality
 of graphic designs spanning four dimensions, twelve indicators and three quantifiable tasks.
- We systematically evaluate proprietary, open-source, and reasoning-augmented VLMs, revealing clear performance gaps in aesthetic quality assessment.
- We construct a training dataset to fine-tune VLMs for this domain. Our approach introduces human-guided VLM labeling to produce task labels at scale and indicator-grounded reasoning to tie abstract indicators to concrete design regions. Experiments on AesEvalBench show that this dataset consistently improves performance across all tasks.

Table 1: A comparison of AesEval-Bench with existing benchmarks for both image aesthetics and design aesthetics. We highlight key differences in their scale, task formats, source data, covered design dimensions, and the inclusion of reasoning paths.

Benchmark	#Data	Task	Source	Source Type		Din	nension		Training Set	Reasoning Path	Open-source		
Denemmark	"Dutu	Format	Source	Source Type	Font	Layout	Graphics	Color	Truming Sec	reasoning rum			
				Image Aesthetic	s Bench	mark							
AesBench (Huang et al., 2024)	~10k	Free-form	Photographic Image	Image-only	×	×	✓	✓	×	×	✓		
UNIAA-Bench (Zhou et al., 2024)	${\sim}6k$	Free-form	Photographic Image	Image-only	×	✓	×	✓	×	×	✓		
FineArtBench (Jiang & Chen, 2025)	-	Choice Free-form	Art Work + Photographic Image	Image-only	×	×	×	✓	×	✓	×		
	Design Aesthetics Benchmark												
DesignBench (Lin et al., 2023)	-	Choice Free-form	Graphic Design	Image+Json	✓	✓	×	✓	×	×	✓		
DesignProbe (Lin et al., 2024a)	$\sim\!1.6k$	Choice	Graphic Design	Image+Json	✓	✓	×	✓	×	×	×		
GPT-Eval Bench (Haraguchi et al., 2024)	${\sim}2k$	Scoring	Graphic Design	Image+Json	×	✓	×	×	×	×	×		
UI-Bench (Jung et al., 2025)	${\sim}3k$	Choice Description	UI Design	Image-only	✓	✓	×	✓	×	×	×		
UICrit (Jung et al., 2025)	${\sim}3k$	Free-form Bbox regression	UI Design	Image-only	✓	✓	×	✓	×	×	✓		
AesEval-Bench (Ours)	${\sim}4.5k$	Choice Bbox Regression	Graphic Design	Image+Json	✓	✓	✓	✓	✓	✓	✓		

2 RELATED WORKS

Aesthetic Quality Assessment. Aesthetic quality assessment (Deng et al., 2017) aims to automatically evaluate visual appeal, serving as a computational proxy for human judgment. Within this area, two major lines of research have emerged (see Table 1). *Image aesthetics assessment* (Huang et al., 2024; Zhou et al., 2024; Jiang & Chen, 2025) focuses on photographic images, where quality is determined by factors such as color harmony, lighting, and subject placement. *Design aesthetics assessment* (Lin et al., 2023; 2024a; Haraguchi et al., 2024; Jung et al., 2025) targets graphic designs such as posters, advertisements, or user interfaces, which depend on design-related factors including typography, hierarchy, and alignment. Our work falls within design aesthetics assessment.

Despite its importance, design aesthetics assessment remains underexplored. Existing benchmarks capture only a narrow subset of design dimensions. For instance, (Lin et al., 2024a) omits graphics-related factors, while (Haraguchi et al., 2024) ignores both fonts and graphics. Furthermore, their task formulations lack rigor. Some adopt free-form question answering, which is difficult to quantify (Lin et al., 2023), while others provide only holistic scores without identifying problematic regions, limiting interpretability and actionability (Jung et al., 2025). Our work introduces a benchmark that comprehensively covers design-related aesthetic factors across font, layout, graphics and color, defining well-structured and quantifiable tasks using choice and bbox prediction formats.

Vision-Language Models (VLMs). Vision-Language Models (VLMs) (Wang et al., 2024; Comanici et al., 2025; Li et al., 2024a) have achieved remarkable performance on tasks such as image captioning (Luo et al., 2025; Li et al., 2024b; Zhang et al., 2025b) and visual question answering (An et al., 2024; Lin et al., 2024b; Luo et al., 2024). Yet, their ability to assess the aesthetic quality of graphic designs remains largely unexplored. Prior work typically evaluates only one or two VLMs (e.g., (Haraguchi et al., 2024) studies GPT). We provide a systematic comparison across a broad set of popular VLMs, including proprietary, open-source, and reasoning-augmented models.

Recently, increasing attention has been devoted to the reasoning capabilities of VLMs (Li et al., 2025a; Zhang et al., 2025a). For example, (Sarch et al., 2025) employs tree-based search to improve reasoning chains, (Li et al., 2025a) visualizes reasoning trajectories for transparency, and (Sun et al., 2024; Shao et al., 2024; Cao et al., 2025; Wu et al., 2025) explores grounded visual reasoning by jointly generating bounding boxes and textual explanations. In our work, we observe that generic reasoning in current VLMs provides limited benefit for assessing design aesthetics. To address this, we construct a training dataset with reasoning paths that explicitly link abstract design indicators to concrete regions of the design. Unlike grounded visual reasoning, which localizes semantically salient entities (e.g., a "dog" or "chair"), our regions are indicator-centric, capturing higher-level concepts such as hierarchy, alignment, and spacing that directly embody design principles.

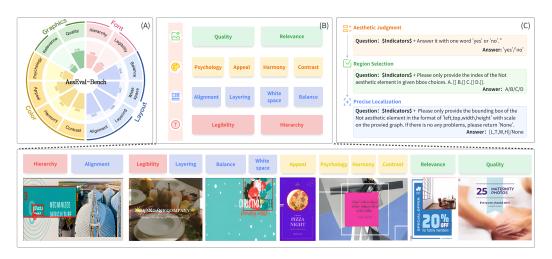


Figure 1: Overview of AesEval-Bench. (A) The four dimensions and twelve indicators considered in the benchmark. Numbers inside the circles indicate how many designs are labeled as flawed for each indicator. (B) Example designs illustrating the indicators, with regions exhibiting aesthetic issues highlighted by red boxes. Detailed textual explanations of all indicators are provided in the Appendix. (C) The three tasks, along with example questions and their expected answers.

3 BENCHMARK CONSTRUCTION

3.1 Overview

AesEval-Bench formulates design aesthetics assessment as a question—answering task. The input contains a *task* description and a *design image*, optionally accompanied by metadata such as layout, font, or color information in JSON format. The output is the *answer* corresponding to the task.

To capture different aspects of design aesthetics assessment, we introduce three task types (Figure 1(C)): 1) *aesthetic judgment* asks whether a design is aesthetically pleasing (yes/no), providing a measure of overall perception. 2) *region selection* requires choosing from candidate regions where aesthetic issues appear, testing the ability to localize problematic areas beyond a global judgment. 3) *precise localization* requires predicting the exact bounding box coordinates of problematic regions, offering a fine-grained diagnosis. Each task is accompanied by the explanation of an indicator—the key factor humans consistently emphasize when evaluating visual appeal (e.g., hierarchy, layering, contrast). We consider twelve indicators (Figure 1(B)), grouped into four dimensions (Figure 1(A)).

For design images, we sample 1500 designs from the test split of Crello dataset (Yamaguchi, 2021), which contains professional graphic designs with both the design image and its metadata. The expected answers differ across tasks. For aesthetic judgment, the answer is yes or no. For region selection, it is the index of one region among four candidates. For precise localization, it is the bounding box coordinates of the identified region or None if the design has no aesthetic issues.

Overall, AesEval-Bench comprises 4500 base question—answer pairs (three tasks across 1500 designs), each further instantiated across twelve indicators to enable fine-grained evaluation.

3.2 CURATION PIPELINE

Establishing Dimensions and Indicators. Aesthetic quality in graphic design is inherently multidimensional. To define a rigorous benchmark, we first conducted a comprehensive literature review of classical and contemporary design principles (McCormack & Lomas, 2020; Lou et al., 2022; Lu et al., 2020). We then consulted professional designers to refine this taxonomy, ensuring alignment with both theoretical foundations and practical expertise. This process yielded four core dimensions—layout, font, graphics, and color (Figure 1(A))—each further specified by con-

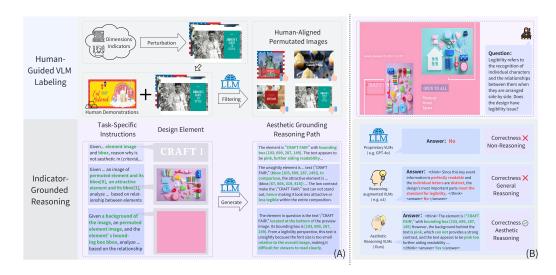


Figure 2: (A) Illustration of two key steps in training data construction. Human-guided VLM labeling enables scalable determination of whether designs exhibit aesthetic issues. Indicator-grounded reasoning generates reasoning paths that explicitly link abstract indicators to concrete design regions (represented as bbox coordinates). (B) Example highlighting the difference between non-reasoning models, generic reasoning models, and our indicator-grounded reasoning model.

crete indicators consistently emphasized in human aesthetic judgment. In total, we distilled twelve indicators that together capture the essential factors of design aesthetics (Figure 1(B)).

Constructing Potentially Flawed Designs. As introduced in Section 3.1, the design images in AesEval-Bench are sourced from the Crello dataset, which contains professional-quality graphic designs. To effectively evaluate design aesthetics, the benchmark must include not only well-designed but also less appealing examples. We therefore repurpose Crello by introducing controlled perturbations, such as repositioning elements, altering font choices, or adjusting colors. These perturbations may either degrade the visual quality or leave it largely intact. For example, slightly enlarging a heading might preserve hierarchy, whereas shifting it left could disrupt balance. Each base design undergoes one to three random perturbations, generating a spectrum of variations that range from aesthetically unchanged to noticeably flawed, while still appearing realistic. Since Crello Yamaguchi (2021) provides element-level metadata in JSON format along with separate design layers, these perturbations can be applied directly at the JSON level and rendered into new design images by recombining the modified metadata with the corresponding layers.

Human-in-the-Loop Aesthetic Review. We engage human annotators to verify whether the perturbed designs truly exhibit aesthetic issues. Before annotation, all annotators receive a tutorial that includes examples of both well-designed and flawed cases, along with detailed explanations of the underlying reasons. During the review, each annotator is shown a design image together with a description of the focal indicator and asked to determine whether the design contains the corresponding flaw (yes or no). For each design, we derive the final label by applying majority voting across multiple annotators to ensure consensus.

Generating Question-Answer Pairs. With metadata in JSON format, records of applied perturbations and human annotations, we can systematically construct answers corresponding to each task. For aesthetic judgment, the rule is straightforward: if human annotators label a design as good, the ground-truth answer is no (i.e., no aesthetic issue); otherwise, it is yes. For region selection, if a design is labeled as good, the four answer choices consist of three randomly sampled bboxes from the metadata and a None option, with the ground-truth answer being None. If the design is labeled as flawed, the four choices include the bbox of the perturbed element, two randomly sampled bboxes from the metadata, and None, with the ground-truth answer set to the bbox of the perturbed element. For precise localization, if a design is labeled as good, the ground-truth answer is None; otherwise, it corresponds to the exact bbox of the perturbed element.

3.3 EVALUATION PROTOCOLS

For aesthetic judgment and region selection, both formulated as choice problems, we adopt accuracy as the metric, measuring the exact match between model predictions and the ground truth. For precise localization, the task combines two components: a choice problem (predicting None when no aesthetic issue exists) and a bounding box regression problem (predicting the exact bbox when an issue is present). Accordingly, we use accuracy for cases where the ground truth is None, and intersection over union (IoU)—which quantifies the overlap between the predicted and ground-truth bboxes—for cases where a bbox is required.

4 TRAINING DATA CONSTRUCTION

Evaluation on popular VLMs reveals clear gap between the capabilities of current state-of-thearts VLMs and the nuanced requirements of aesthetic quality assessment. Moreover, reasoningaugmented VLMs show no clear performance gains (see Section 5.1).

To this end, we construct a training dataset, named **AesEval-Train**, to fine-tune VLMs for this domain. First, we adopt the same procedure as benchmark construction to *construct potentially flawed designs* (see Section 3.2). Next, since relying solely on human annotation to determine whether perturbed designs exhibit aesthetic issues is neither scalable nor cost-effective for training at large scale, we introduce *human-guided VLM labeling*. Then, we follow the benchmark construction to *generate question–answer pairs* (see Section 3.2). Finally, we introduce *indicator-grounded reasoning* to generate domain-specific reasoning paths aimed at improving task performance. In the following, we describe in detail the two steps that differ from benchmark construction.

Human-Guided VLM Labeling. We leverage a small set of human annotations as demonstrations, together with the bbox coordinates of perturbed regions, as input to strong VLMs. The model is instructed to generate a binary label indicating whether the perturbed design exhibits an aesthetic issue (see Figure 2(A)). By incorporating human annotations, we preserve alignment with human judgment while substantially reducing manual annotation costs. Moreover, providing the perturbation region as a prior, which is unavailable in real-world scenarios, simplifies the labeling process and improves reliability. With these two sources of guidance, while the generated labels may not be perfectly accurate, they yield a training set of sufficient quality to enhance fine-tuning performance.

Indicator-Grounded Reasoning. As illustrated in Figure 2(B), generic reasoning often explains or analyzes a given indicator and task without grounding the discussion in relevant regions of the design. To address this limitation, we propose explicitly linking abstract indicators to concrete regions within the design. Specifically, we include both the bounding box (bbox) coordinates of relevant regions and textual explanations of their relevance to the indicator in the reasoning path.

To obtain such reasoning paths, we instruct powerful VLMs (e.g., GPT in our experiments) by providing them with the bbox coordinates of the target regions and the corresponding design layers (Figure 2(A)). The model is required to output the provided coordinates alongside an explanation of how the region relates to the indicator, thereby ensuring that the reasoning path consistently contains the desired information. We further adopt task-specific strategies to determine the regions of interest. For aesthetic judgment, we directly use the bbox of the perturbed regions. For region selection, we include both the perturbed and non-perturbed regions to strengthen the model's ability to discriminate among candidate regions. For precise localization, we not only highlight the bbox of perturbed regions but also emphasize their relationship to the overall design, enabling the model to better localize problematic regions within the global design context.

5 EXPERIMENT

5.1 BENCHMARKING VLMs on AesEval-Bench

Setups. We conduct a comprehensive evaluation of 10 VLMs spanning diverse model families and parameter scales. For non-reasoning models, we consider open-source representatives such as LLaVA (Liu et al., 2023), Qwen2.5-VL (Bai et al., 2025), and Intern-VL3 (Zhu et al., 2025), as well as closed-source GPT models (Jaech et al., 2024). For reasoning-augmented models, we evaluate

Table 2: Evaluation on aesthetic judgment task. Overall acc is the average value of all indicators.

Model	Overall Acc		I	ayout			Co	olor		Font		Graphics	
Model	O Terum Tree	balance	layering	whitespace	alignment	harmony	contrast	appeal	psycholoy	legibility	hierarchy	quality	relevance
					Non-re	asoning Mo	dels						
LLaVA-7B	0.5343	0.5488	0.5066	0.5135	0.5704	0.6061	0.6161	0.1185	0.0642	0.6815	0.6131	0.6509	0.5572
LLaVA-13B	0.5636	0.5982	0.4539	0.6451	0.6235	0.3887	0.1885	0.2091	0.2327	0.7136	0.7736	0.6862	0.6209
Qwen-VL-3B	0.5643	0.5862	0.5925	0.5925	0.5652	0.5731	0.5157	0.4989	0.6146	0.5720	0.5336	0.5583	0.5409
Qwen-VL-7B	0.6190	0.8314	0.4150	0.7918	0.8255	0.8065	0.9178	0.3651	0.0997	0.9472	0.2742	0.8225	0.3812
Qwen-VL-32B	0.6258	0.6066	0.4993	0.6166	0.6245	0.5613	0.5080	0.4252	0.1166	0.7327	0.6922	0.7012	0.8305
Qwen-VL-72B	0.6524	0.6361	0.6334	0.6540	0.6140	0.6330	0.7213	0.2029	0.1756	0.8348	0.7261	0.7038	0.8243
Intern-VL3-8B	0.6331	0.4043	0.5878	0.7003	0.7177	0.2913	0.2755	0.2508	0.4606	0.8912	0.6377	0.7997	0.8917
Intern-VL3-14B	0.6883	0.6982	0.3402	0.7139	0.7177	0.7282	0.8170	0.4852	0.1488	0.7885	0.7880	0.7949	0.7305
GPT-40	0.7031	0.7423	0.6624	0.6432	0.4523	0.7964	0.8025	0.8072	0.3127	0.7342	0.7179	0.7692	0.7995
GPT-5	0.7252	0.7962	0.7416	0.6859	0.6094	0.6537	0.7959	0.3584	0.4821	0.7056	0.7003	0.8607	0.8135
					Reasoning	-augmented	Models						
GPT-o1	0.6705	0.7225	0.7372	0.5363	0.3239	0.6990	0.5280	0.7268	0.6591	0.6890	0.7107	0.7359	0.7871
GPT-o3	0.7105	0.7225	0.7372	0.6363	0.3739	0.7490	0.6780	0.7768	0.6091	0.7390	0.7107	0.7859	0.7371
Gemini-2.5-Pro	0.6368	0.7403	0.6972	0.5984	0.5137	0.7381	0.6543	0.6824	0.5652	0.5936	0.7045	0.5265	0.6851
				Expert	Models for I	mage Aestl	netics Asse	ssment					
AesExpert-7B	0.4056	0.4863	0.3980	0.3155	0.4474	0.2855	0.3985	0.3091	0.2156	0.2508	0.4911	0.6004	0.4742
UNIAA-LLaVA	0.2900	0.1065	0.0792	0.1143	0.0879	0.1405	0.1988	0.1713	0.1745	0.1565	0.3938	0.1082	0.1551

Table 3: Evaluation on region selection task. Overall acc is the average value of all indicators.

Model	Overall Acc		I	ayout			Co	olor		Font		Graphics							
Wiodei	Overall Acc	balance	layering	whitespace	alignment	harmony	contrast	appeal	psycholoy	legibility	hierarchy	quality	relevance						
					Non-re	asoning Mo	odels			40 0.6551 0.4107 0.5466 0 47 0.6021 0.6894 0.6220 0 20 0.3349 0.3838 0.3060 0 83 0.5373 0.5778 0.6157 0 57 0.7445 0.7651 0.7077 0 18 0.6459 0.6033 0.5907 0 38 0.7148 0.6774 0.6013 0 38 0.7148 0.6774 0.6013 0 58 0.7959 0.7486 0.6874 0 80 0.6066 0.6838 0.5682 0 80 0.6066 0.6838 0.5682 0 60 0.6481 0.7639 0.6496 0									
LLaVA-7B	0.5500	0.6077	0.4023	0.5524	0.5925	0.6709	0.5950	0.6166	0.7140	0.6551	0.4107	0.5466	0.4218						
LLaVA-13B	0.6065	0.5714	0.5604	0.5503	0.5809	0.6062	0.5057	0.6210	0.6747	0.6021	0.6894	0.6220	0.5636						
Qwen-VL-3B	0.3377	0.3264	0.3181	0.3791	0.3207	0.4311	0.3617	0.3181	0.3302	0.3349	0.3838	0.3060	0.2839						
Qwen-VL-7B	0.5295	0.4527	0.4769	0.5147	0.4842	0.5221	0.4832	0.4811	0.5783	0.5373	0.5778	0.6157	0.5089						
Qwen-VL-32B	0.5911	0.5255	0.3119	0.4334	0.4574	0.5000	0.6787	0.6089	0.5457	0.7555	0.5600	0.6972	0.6855						
Qwen-VL-72B	0.6626	0.4822	0.5556	0.4264	0.5065	0.5694	0.6942	0.6212	0.7657	0.7445	0.7651	0.7077	0.7732						
Intern-VL3-8B	0.5799	0.5118	0.4824	0.5482	0.5239	0.5403	0.5444	0.5681	0.6218	0.6459	0.6033	0.5907	0.6291						
Intern-VL3-14B	0.6378	0.5776	0.5108	0.5849	0.5649	0.6186	0.6901	0.6323	0.6938	0.7148	0.6774	0.6013	0.6722						
GPT-40	0.6745	0.4332	0.4512	0.4625	0.5629	0.7024	0.8209	0.6784	0.7753	0.6251	0.8698	0.6062	0.6483						
GPT-5	0.6989	0.6405	0.5850	0.6551	0.6317	0.7150	0.7135	0.7431	0.6768	0.7959	0.7486	0.6874	0.6991						
					Reasoning	-augmented	Models												
GPT-o1	0.6347	0.6065	0.5492	0.6143	0.5680	0.5924	0.6069	0.6626	0.7480	0.6066	0.6838	0.5682	0.5051						
GPT-o3	0.6581	0.6378	0.6158	0.5220	0.3548	0.8167	0.5381	0.7008	0.7876	0.6481	0.7639	0.6496	0.7361						
Gemini-2.5-Pro	0.6100	0.6264	0.6435	0.5550	0.2446	0.5504	0.5992	0.6132	0.5504	0.5618	0.5993	0.5150	0.6059						
				Expert	Models for I	mage Aestl	netics Asse	ssment											
AesExpert-7b	0.2883	0.2954	0.2280	0.3174	0.2631	0.3426	0.2646	0.3176	0.3176	0.3176	0.2588	0.2765	0.2602						
UNIAA-LLaVA	0.2418	0.1510	0.2806	0.1443	0.3966	0.1591	0.1407	0.1651	0.2370	0.2668	0.3687	0.2752	0.1859						

GPT-o1, GPT-o3, and Gemini-2.5-Pro (Comanici et al., 2025). In addition, we include expert models specifically designed for image aesthetic assessment, namely AesExpert (Huang et al., 2024) and UNIAA-LLAVA (Zhou et al., 2024). All models are evaluated under the same input setting, which consists of a question (see Figure 1), a design image, and metadata in JSON format.

Results. The performance of VLMs are evaluated following the protocols introduced in Section 3.3. Specifically, in addition to reporting scores for each individual indicator, we also provide an overall score computed as the average across all indicators.

Aesthetic Judgment. Table 2 presents the results. First, among non-reasoning models, GPT-5 achieves the highest performance, with an overall accuracy of 0.7252. This suggests that even the strongest VLMs still struggle with design aesthetics assessment. Second, reasoning-augmented models do not outperform their non-reasoning counterparts (e.g., GPT-01 Jaech et al. (2024) and GPT-03 vs. GPT-40 Hurst et al. (2024) and GPT-5), indicating that generic reasoning provides little benefit in this domain. Third, expert models designed for image aesthetics assessment perform worse overall, highlighting a substantial gap between design aesthetics and image aesthetics. Finally, model performance varies across indicators. For instance, the Qwen-VL series tends to perform better on legibility but worse on psychology compared to other VLMs.

Region Selection. Table 3 reports the results. First, VLM performance on this task is generally worse than on aesthetic judgment, likely because it requires not only assessing whether a design is pleasing but also identifying where flaws occur. Second, consistent with aesthetic judgment, GPT-5 achieves the best performance, while reasoning-augmented models show no clear advantage. Finally, across model families, larger models (e.g., 32B, 72B) typically outperform smaller ones (e.g., 3B, 7B).

<u>Precise Localization.</u> As described in Section 3.3, this task consists of two components, each evaluated separately: a choice problem, where the model predicts None if no aesthetic issue exists

Table 4: Evaluation on precise localization task for the choice component where the model should predict None if no aesthetic issues are present. Overall score is the average accuracy of all indicators.

Model	Overall Score		I	ayout		Color				Font		Graphics	
1120461	o veran score	balance	layering	whitespace	alignment	harmony	contrast	appeal	psycholoy	legibility	hierarchy	quality	relevance
					Non-rea	asoning Mo	dels						
LLaVA-13B	0.4455	0.4328	0.5935	0.5504	0.4519	0.3161	0.3528	0.2703	0.2279	0.7106	0.2258	0.5416	0.4378
Qwen-VL-72B	0.5192	0.4821	0.5556	0.4263	0.5064	0.4093	0.5342	0.5212	0.5056	0.4245	0.5051	0.5077	0.5132
GPT-4o	0.5680	0.6378	0.2024	0.5915	0.1640	0.8587	0.6333	0.7161	0.6674	0.4555	0.5125	0.7723	0.5579
GPT-5	0.6090	0.6314	0.6150	0.6918	0.6255	0.6065	0.6178	0.5651	0.5997	0.5472	0.6142	0.6225	0.5812
	Reasoning-augmented Models												
GPT-o1	0.5295	0.4527	0.4769	0.5147	0.4842	0.5221	0.4832	0.4811	0.5783	0.5373	0.5778	0.6157	0.5089
GPT-o3	0.5800	0.6077	0.4023	0.5524	0.6725	0.7709	0.6750	0.6166	0.7140	0.7551	0.4107	0.5466	0.4218
Gemini-2.5-Pro	0.6047	0.6065	0.5492	0.5143	0.5680	0.5924	0.6069	0.6626	0.7480	0.5066	0.6838	0.4082	0.5051
				Expert 1	Models for I	mage Aesth	etics Asses	sment					
AesExpert-7b	0.3377	0.3264	0.3181	0.3791	0.3207	0.4311	0.3617	0.3181	0.3302	0.3349	0.3838	0.3060	0.2839

Table 5: Evaluation on precise localization task for the bbox prediction component where the model should output coordinates of the aesthetic issues. Overall score is the average IoU of all indicators.

Model	Overall Score		I	ayout	Color				Font		Graphics		
Model	Overall Score	balance	layering	whitespace	alignment	harmony	contrast	appeal	psycholoy	legibility	hierarchy	quality	relevance
					Non-rea	asoning Mo	dels						
LLaVA-13B	0.0559	0.0646	0.0073	0.0295	0.0165	0.1017	0.0785	0.0420	0.0296	0.0181	0.0095	0.1503	0.1153
Owen-VL-72B	0.0935	0.0268	0.1011	0.1178	0.0907	0.1131	0.1216	0.0030	0.1008	0.1149	0.1094	0.1129	0.0238
GPT-40	0.1712	0.1737	0.2889	0.1314	0.2467	0.0798	0.1268	0.0579	0.2101	0.2059	0.0501	0.2622	0.1185
GPT-5	0.1993	0.1609	0.1289	0.1820	0.1356	0.1572	0.1268	0.1272	0.3091	0.1488	0.1578	0.2410	0.2081
	GP1-5 0.1995 0.1609 0.1289 0.1820 0.1536 0.1572 0.1268 0.1272 0.3091 0.1488 0.1578 0.2410 0.20 Reasoning-augmented Models												
O1	0.1286	0.0772	0.0632	0.1277	0.1091	0.1584	0.0069	0.1411	0.1101	0.1305	0.1306	0.1777	0.1482
O3	0.1418	0.0510	0.1806	0.0443	0.2966	0.0591	0.0407	0.0651	0.1370	0.1668	0.2687	0.1752	0.0859
Gemini-2.5-Pro	0.0977	0.0567	0.1669	0.1101	0.0759	0.0809	0.0536	0.0539	0.2306	0.1012	0.0952	0.0994	0.1063
				Expert 1	Models for I	mage Aesth	etics Asses	ssment					
AesExpert-7b	0.0327	0.0458	0.0221	0.0111	0.1102	0.0879	0.0081	0.0067	0.0136	0.0019	0.0366	0.0619	0.0085

(Table 4), and a bbox prediction problem, where the model outputs the exact bbox of the aesthetic issue (Table 5). We exclude some VLMs (e.g., Intern-VL series and small Qwen-VL models) because they failed to produce meaningful bbox predictions. For the choice problem, VLMs achieve reasonable performance, with the best model reaching an overall score of 0.6090. For the bbox prediction problem, even the best-performing model, GPT-5, scores below 0.20, highlighting the substantial difficulty of precisely localizing aesthetic issues.

Discussions on Input Components. When benchmarking VLMs, the input consists of three components: (1) the question, which includes a detailed explanation of the target indicator; (2) the design image; and (3) metadata in JSON format, containing layout, color, and font information. We analyze the contribution of each component to model performance using GPT-40 as a representative example. Figure 3 presents the results, where Full Model denotes the setting that uses all three components; Without Images removes the design image; Without Explanation omits the detailed indicator description; and Without Metainfo excludes the metadata. Our findings reveal three key insights. First, across all tasks, the design image is indispensable—its removal results in the largest performance drop. Second, indicator explanations have limited influence for more intuitive indicators (e.g., balance), but they play a crucial role for subjective indicators (e.g., relevance or psychology), where clearer definitions are necessary. Finally, metadata has the greatest effect on precise localization, where its absence causes a larger decline in performance compared to aesthetic judgment or region selection. We hypothesize that this is because metadata provides explicit layout information, which aids bbox prediction in localization tasks.

5.2 FINE-TUNING VLMs WITH AESEVAL-TRAIN

Setups. We construct the training set following the pipeline described in Section 4, resulting in 30k question–answer pairs. In our experiments, we use Qwen2.5-VL-7B-Instruct (Bai et al., 2025) as a representative model and adopt full-parameter finetuning on the constructed dataset. The learning rate is set to 1e-6, with a cosine scheduler and a 3% warmup ratio. For computational efficiency, training is performed with bfloat16 mixed precision and FlashAttention-2 (Dao et al., 2022). The vision encoder is kept frozen, while the language model parameters are tuned.

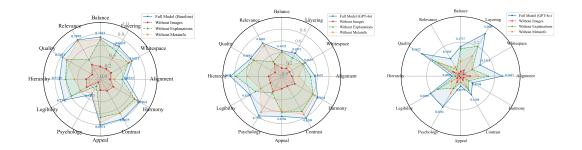


Figure 3: Results for model variants using different input components.

Table 6: Results and ablation study of fine-tuning VLMs using our constructed training set.

Model Setting	Overall Score	Layout					Co	olor		Font		Graphics	
odei Setting	O reruii Score	balance	layering	whitespace	alignment	harmony	contrast	appeal	psycholoy	legibility	hierarchy	quality	relevance
				Ac	sthetic Judg	ment (Accı	iracy)						
Qwen-VL-7B (Base)	0.6190	0.8314	0.4150	0.7918	0.8255	0.8065	0.9178	0.3651	0.0997	0.9472	0.2742	0.8225	0.3812
+ AesEval-Train	0.6987	0.7123	0.6789	0.7215	0.6868	0.7031	0.6654	0.7329	0.6577	0.7436	0.6482	0.7096	0.7244
- Reasoning Path	0.6576	0.6511	0.6589	0.6413	0.6687	0.6309	0.6791	0.6207	0.6893	0.6115	0.6985	0.6508	0.6904
- Positive Samples	0.2072	0.2101	0.1999	0.2202	0.1893	0.2058	0.2147	0.1955	0.2246	0.1855	0.2296	0.2004	0.2108
	Region Selection (Accuracy)												
Owen-VL-7B (Base)	0.5295	0.4527	0.4769	0.5147	0.4842	0.5221	0.4832	0.4811	0.5783	0.5373	0.5778	0.6157	0.5089
+ AesEval-Train	0.6065	0.5827	0.6108	0.5963	0.6289	0.5714	0.6236	0.6541	0.6412	0.6389	0.6487	0.5899	0.5915
- Reasoning Path	0.5795	0.5732	0.5279	0.5697	0.5322	0.5741	0.5322	0.5341	0.6283	0.5953	0.6318	0.6667	0.5885
- Positive Samples	0.5327	0.5089	0.5411	0.5257	0.5543	0.5013	0.5587	0.5291	0.5709	0.5212	0.5788	0.5146	0.4878
					Precise Loc	alization (Io	U)						
Owen-VL-7B (Base)	0.0514	0.0067	0.1669	0.0101	0.0259	0.0109	0.0036	0.0039	0.2306	0.0012	0.0452	0.0994	0.0063
+ AesEval-Train	0.2231	0.2518	0.1982	0.3103	0.0857	0.2204	0.2846	0.1552	0.3901	0.0607	0.2313	0.1152	0.2745
- Reasoning Path	0.0782	0.1523	0.0211	0.1987	0.0095	0.0750	0.1204	0.0348	0.0812	0.0159	0.0555	0.0601	0.1139
- Positive Samples	0.0641	0.1452	0.0231	0.0888	0.0079	0.1711	0.0456	0.0613	0.0199	0.0924	0.0337	0.0578	0.0224

Main Results. Table 6 presents the results, where *Qwen-VL-7B* (*Base*) denotes the base model without finetuning, and +*AesEval-Train* refers to the model finetuned on our constructed training set. First, across all three tasks, finetuning with AesEval-Train yields substantial performance improvements. Moreover, on aesthetic judgment, the finetuned model surpasses even the largest Qwen-VL variant (72B parameters), and on precise localization, it outperforms GPT-5 despite the latter having far more parameters. These results demonstrate that our proposed pipeline effectively constructs training data that significantly enhances model performance.

Ablation Studies. We investigate the impact of different data recipes on model performance. Table 6 reports the results, where *-Reasoning Path* denotes training with plain question–answer pairs without the proposed indicator-grounded reasoning, and *-Positive Samples* denotes training only on flawed designs. We observe that *-Reasoning Path* still improves performance across all three tasks, suggesting that incorporating domain-specific knowledge of design aesthetics is beneficial. However, its performance remains notably lower than that of the full variant with reasoning paths (+AesEval-Train), underscoring the effectiveness of indicator-grounded reasoning. In addition, *-Positive Samples* performs worse than both +AesEval-Train and *-Reasoning Path*, highlighting the importance of maintaining label balance in the training set.

6 Conclusion

In this work, we introduce AesEval-Bench for design aesthetics assessment, which spans four dimensions, twelve indicators and three quantifiable tasks. Based on it, we systematically evaluate proprietary, open-source and reasoning-augmented VLMs, revealing clear gaps in design aesthetics assessment. Furthermore, we construct a training dataset to fine-tune VLMs for this domain. Experiments show that this dataset significantly improves model performance across all tasks.

Limitations. First, as Crello serves as the source dataset, the benchmark does not cover all types of graphic design, such as infographics or mobile UIs. Second, the considered indicators are not fully disentangled, leaving room for a more refined and systematic taxonomy in future work. Third, highly subjective aspects of design, such as creativity, are not included. Finally, leveraging reinforcement learning to further enhance reasoning capabilities is left for future exploration.

REFERENCES

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- Ruichuan An, Sihan Yang, Ming Lu, Renrui Zhang, Kai Zeng, Yulin Luo, Jiajun Cao, Hao Liang, Ying Chen, Qi She, et al. Mc-llava: Multi-concept personalized vision-language model. *arXiv* preprint arXiv:2411.11706, 2024.
- Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibo Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, et al. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*, 2025.
- Meng Cao, Haoze Zhao, Can Zhang, Xiaojun Chang, Ian Reid, and Xiaodan Liang. Groundr1: Incentivizing grounded visual reasoning via reinforcement learning. *arXiv* preprint *arXiv*:2505.20272, 2025.
- Gheorghe Comanici, Eric Bieber, Mike Schaekermann, Ice Pasupat, Noveen Sachdeva, Inderjit Dhillon, Marcel Blistein, Ori Ram, Dan Zhang, Evan Rosen, et al. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv* preprint arXiv:2507.06261, 2025.
- Tri Dao, Dan Fu, Stefano Ermon, Atri Rudra, and Christopher Ré. Flashattention: Fast and memory-efficient exact attention with io-awareness. *Advances in neural information processing systems*, 35:16344–16359, 2022.
- Yubin Deng, Chen Change Loy, and Xiaoou Tang. Image aesthetic assessment: An experimental survey. *IEEE Signal Processing Magazine*, 34(4):80–106, 2017.
- Daichi Haraguchi, Naoto Inoue, Wataru Shimoda, Hayato Mitani, Seiichi Uchida, and Kota Yamaguchi. Can gpts evaluate graphic design based on design principles? In *SIGGRAPH Asia 2024 Technical Communications*, pp. 1–4, 2024.
- Yipo Huang, Quan Yuan, Xiangfei Sheng, Zhichao Yang, Haoning Wu, Pengfei Chen, Yuzhe Yang, Leida Li, and Weisi Lin. Aesbench: An expert benchmark for multimodal large language models on image aesthetics perception. *arXiv preprint arXiv:2401.08276*, 2024.
- Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, et al. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*, 2024.
- Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden Low, Alec Helyar, Aleksander Madry, Alex Beutel, Alex Carney, et al. Openai o1 system card. *arXiv* preprint arXiv:2412.16720, 2024.
- Ruixiang Jiang and Changwen Chen. Multimodal llms can reason about aesthetics in zero-shot. *arXiv preprint arXiv:2501.09012*, 2025.
- Sam Jung, Agustin Garcinuno, and Spencer Mateega. Ui-bench: A benchmark for evaluating design capabilities of ai text-to-app tools. *arXiv preprint arXiv:2508.20410*, 2025.
- Bo Li, Yuanhan Zhang, Dong Guo, Renrui Zhang, Feng Li, Hao Zhang, Kaichen Zhang, Peiyuan Zhang, Yanwei Li, Ziwei Liu, et al. Llava-onevision: Easy visual task transfer. *arXiv preprint arXiv:2408.03326*, 2024a.
- Chengzu Li, Wenshan Wu, Huanyu Zhang, Yan Xia, Shaoguang Mao, Li Dong, Ivan Vulić, and
 Furu Wei. Imagine while reasoning in space: Multimodal visualization-of-thought. *arXiv preprint arXiv:2501.07542*, 2025a.
 - Congcong Li and Tsuhan Chen. Aesthetic visual quality assessment of paintings. *IEEE Journal of selected topics in Signal Processing*, 3(2):236–252, 2009.

- Jian Li, Weiheng Lu, Hao Fei, Meng Luo, Ming Dai, Min Xia, Yizhang Jin, Zhenye Gan, Ding Qi, Chaoyou Fu, et al. A survey on benchmarks of multimodal large language models. *arXiv preprint arXiv:2408.08632*, 2024b.
 - Junxian Li, Di Zhang, Xunzhi Wang, Zeying Hao, Jingdi Lei, Qian Tan, Cai Zhou, Wei Liu, Yaotian Yang, Xinrui Xiong, et al. Chemvlm: Exploring the power of multimodal large language models in chemistry area. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 415–423, 2025b.
 - Jieru Lin, Danqing Huang, Tiejun Zhao, Dechen Zhan, and Chin-Yew Lin. Designprobe: A graphic design benchmark for multimodal large language models. *arXiv preprint arXiv:2404.14801*, 2024a.
 - Kevin Lin, Zhengyuan Yang, Linjie Li, Jianfeng Wang, and Lijuan Wang. Designbench: Exploring and benchmarking dall-e 3 for imagining visual design. *arXiv preprint arXiv:2310.15144*, 2023.
 - Weifeng Lin, Xinyu Wei, Ruichuan An, Peng Gao, Bocheng Zou, Yulin Luo, Siyuan Huang, Shanghang Zhang, and Hongsheng Li. Draw-and-understand: Leveraging visual prompts to enable mllms to comprehend what you want. *arXiv preprint arXiv:2403.20271*, 2024b.
 - Weifeng Lin, Xinyu Wei, Ruichuan An, Tianhe Ren, Tingwei Chen, Renrui Zhang, Ziyu Guo, Wentao Zhang, Lei Zhang, and Hongsheng Li. Perceive anything: Recognize, explain, caption, and segment anything in images and videos. *arXiv preprint arXiv:2506.05302*, 2025.
 - Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *Advances in neural information processing systems*, 36:34892–34916, 2023.
 - Yun Lou, Weiyue Gao, Pei Chen, Xuanhui Liu, Changyuan Yang, and Lingyun Sun. Designeva: A design-supported tool with multi-faceted perceptual evaluation. In *International Conference on Human-Computer Interaction*, pp. 508–519. Springer, 2022.
 - Yue Lu, Chao Guo, Yilun Lin, Fan Zhuo, and FY Wang. Computational aesthetics of fine art paintings: The state of the art and outlook. *Acta Automatica Sinica*, 46(11):2239–2259, 2020.
 - Meng Luo, Shengqiong Wu, Liqiang Jing, Tianjie Ju, Li Zheng, Jinxiang Lai, Tianlong Wu, Xinya Du, Jian Li, Siyuan Yan, Jiebo Luo, William Yang Wang, Hao Fei, Mong-Li Lee, and Wynne Hsu. Dr.v: A hierarchical perception-temporal-cognition framework to diagnose video hallucination by fine-grained spatial-temporal grounding, 2025. URL https://arxiv.org/abs/2509.11866.
 - Yulin Luo, Ruichuan An, Bocheng Zou, Yiming Tang, Jiaming Liu, and Shanghang Zhang. Llm as dataset analyst: Subpopulation structure discovery with large language model. In *European Conference on Computer Vision*, pp. 235–252. Springer, 2024.
 - Jon McCormack and Andy Lomas. Understanding aesthetic evaluation using deep learning. In *International conference on computational intelligence in music*, sound, art and design (part of EvoStar), pp. 118–133. Springer, 2020.
 - Gabriel Sarch, Snigdha Saha, Naitik Khandelwal, Ayush Jain, Michael J Tarr, Aviral Kumar, and Katerina Fragkiadaki. Grounded reinforcement learning for visual reasoning. *arXiv* preprint arXiv:2505.23678, 2025.
 - Hao Shao, Shengju Qian, Han Xiao, Guanglu Song, Zhuofan Zong, Letian Wang, Yu Liu, and Hongsheng Li. Visual cot: Advancing multi-modal language models with a comprehensive dataset and benchmark for chain-of-thought reasoning. *Advances in Neural Information Processing Systems*, 37:8612–8642, 2024.
 - Guangyan Sun, Mingyu Jin, Zhenting Wang, Cheng-Long Wang, Siqi Ma, Qifan Wang, Tong Geng, Ying Nian Wu, Yongfeng Zhang, and Dongfang Liu. Visual agents as fast and slow thinkers. *arXiv* preprint arXiv:2408.08862, 2024.
 - Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, et al. Qwen2-vl: Enhancing vision-language model's perception of the world at any resolution. *arXiv preprint arXiv:2409.12191*, 2024.

- Chatchai Wangwiwattana and Areerat Meeyen. Design through ai eyes: Automating aesthetic assessments for learning graphic design concepts. In *International Conference on Multi-disciplinary Trends in Artificial Intelligence*, pp. 336–347. Springer, 2024.
- Qiong Wu, Xiangcong Yang, Yiyi Zhou, Chenxin Fang, Baiyang Song, Xiaoshuai Sun, and Rongrong Ji. Grounded chain-of-thought for multimodal large language models. *arXiv preprint arXiv:2503.12799*, 2025.
- Kota Yamaguchi. Canvasvae: Learning to generate vector graphic documents. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 5481–5489, 2021.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, et al. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*, 2025.
- Huanyu Zhang, Chengzu Li, Wenshan Wu, Shaoguang Mao, Yifan Zhang, Haochen Tian, Ivan Vulić, Zhang Zhang, Liang Wang, Tieniu Tan, et al. Scaling and beyond: Advancing spatial reasoning in mllms requires new recipes. *arXiv preprint arXiv:2504.15037*, 2025a.
- Qintong Zhang, Bin Wang, Victor Shea-Jay Huang, Junyuan Zhang, Zhengren Wang, Hao Liang, Conghui He, and Wentao Zhang. Document parsing unveiled: Techniques, challenges, and prospects for structured information extraction. *arXiv* preprint arXiv:2410.21169, 2024.
- YiFan Zhang, Huanyu Zhang, Haochen Tian, Chaoyou Fu, Shuangqing Zhang, Junfei Wu, Feng Li, Kun Wang, Qingsong Wen, Zhang Zhang, et al. Mme-realworld: Could your multimodal llm challenge high-resolution real-world scenarios that are difficult for humans? In *The Thirteenth International Conference on Learning Representations*, 2025b.
- Zhaokun Zhou, Qiulin Wang, Bin Lin, Yiwei Su, Rui Chen, Xin Tao, Amin Zheng, Li Yuan, Pengfei Wan, and Di Zhang. Uniaa: A unified multi-modal image aesthetic assessment baseline and benchmark. *arXiv preprint arXiv:2404.09619*, 2024.
- Jinguo Zhu, Weiyun Wang, Zhe Chen, Zhaoyang Liu, Shenglong Ye, Lixin Gu, Hao Tian, Yuchen Duan, Weijie Su, Jie Shao, et al. Internvl3: Exploring advanced training and test-time recipes for open-source multimodal models. arXiv preprint arXiv:2504.10479, 2025.

Dimension-Indicators Prompts:

Graphic-Quality: Resolution is a measure of image detail. Appropriate resolution ensures the best picture quality and readability. Does element have quality issue?

Graphic-Relevance: Relevance refers to the direct connection between a graphic element and the meaning it conveys. Does element have relevance issue?

Color-Harmony: Color Harmony refers to the overall coordination, pleasure and beauty of the entire color when there are two or more colors in an image.

Color-Contrast: Color contrast refers to the contrasts, oppositions, and differences existing among various colors.

Color-Appreal: Color appeal refers to the fact that the selection and combination of colors can attract the attention of the audience.

Color-Psychology: Color psychology refers to the idea that color can trigger subjective psychological experiences and influence emotions, feelings, and behaviors.

Layout-Balance: Balance is the distribution of visual weight in design. It can be symmetrical (with equal weights on both sides) or asymmetrical (with unequal weights but still achieving visual balance).

Layout-Layering: Use size, color, contrast, and other visual cues to establish a hierarchical structure of design elements to guide the audience's eyes.

Layout-Whitespace: White space refers to the blank area around the elements in a design. Effective utilization of white space can create balance, visual hierarchy, and clarity.

Layout-Alignment: Alignment refers to the arrangement of design elements relative to each other or a specific axis or grid. Proper alignment creates a sense of order and organization, making the design easier to understand and navigate.

Font-Hierarchy: The presentation of the font has a hierarchical structure, so users can scan the text to obtain key information.

Font-Legibility: Legibility refers to the recognition of individual characters and the relationships between them when they are arranged side by side.

Table 7: Showcase of descriptions of each indicator.

Task Reasoning Instructions:

Aesthetic Judgment : Given the preview image and the element image and element bounding box bbox, please reason why the element is not aesthetic in the aspect of criteria. Your reason should only contain the analysis about criteria. Please think and reasoning. Once the element is mentioned, you should add its bounding box after the element(only bounding box).

Region Selection: Given a preview image, an image of an unsightly element and its bounding box bbox[0], and an image of an attractive element and its bounding box bbox[1], analyze why the first element is unsightly in terms of criteria based on the relationship between the elements. Your reasoning should only include an analysis of criteria. Please think and reason. Once an element is mentioned, you should add its bounding box after the element (bounding box only).

Precise Localization: Given a preview image, an element image, and the element's bounding box bbox, analyze why the element is unsightly when viewed from the perspective of criteria, based on their relationship. Your reasoning should only include analysis of criteria. Think and reason. Once you mention the element, you should follow it with its bounding box (only the bounding box).

Table 8: Showcase of task reasoning instructions.

A STATEMENT OF LLM USAGE

Large language models (LLMs) were consulted for technical guidance during implementation and debugging; following the collaborative drafting of the manuscript, we further employed LLMs to refine the prose and enhance the overall exposition.

B PROMPTS AND INSTRUCTIONS

First, Dimension-Indicator Prompts establish a clear set of evaluation criteria. These are organized into four core dimensions: Graphics, Color, Layout, and Font, each containing specific indicators like Font-Legibility. As shown in Table 7, every indicator is defined and paired with a guiding question to standardize the analysis.

Second, Task Reasoning Instructions (Table 8) provide operational guidance for creating the reasoning paths. They direct the analysis to focus on an element's intrinsic flaws, its relationship with other elements, or its immediate context, while critically mandating the inclusion of the element's bounding box (bbox) to ground the reasoning in precise spatial evidence.

C BENCHMARK SHOWCASE

In this section, we provide visual examples to better illustrate the evaluation criteria for the various aesthetic dimensions within the AesEval-Benchmark. As shown in Table 9, Table 13, Table 16, Table 14, Table 12, Table 11 and Table 15, each showcase presents a side-by-side comparison of designs that exemplify positive and negative attributes for a specific criterion. These examples serve to clarify the standards used for judging aspects such as layout alignment, color harmony, graphic quality, and font legibility, offering a tangible guide to our benchmark's methodology.

Benchmark Showcase

> Layout-Alignment, Graphic-Quality





Explanation: Clear background, center-aligned text.

The background is blurry and the words in the middle are not aligned.

Table 9: Examples of Layout-Alignment and Graphic-Quality in AesEval-Benchmark.

Benchmark Showcase

⊳ Font-Legbility





Explanation: The yellow font in the lower left corner is clearly visible.

The yellow text in the lower left corner becomes blurred and invisible.

Table 10: Examples of Font-Legbility in AesEval-Benchmark.

Benchmark Showcase

⊳ Graphic-Relevance





Explanation: The background is a woman wearing a The background is a beautiful landscape photo, which mask, which is relevant. does not fit the theme.

Table 11: Examples of Graphic-Relevance in AesEval-Benchmark.

Benchmark Showcase

▷ Layout-Whitespace, Layout-Layering





Explanation: The text is in the blank space, with no blank space or overlap.

The text has been moved to the animal's feet, resulting in white space below and stacked elements.

Table 12: Examples of Layout-Whitespace and Layout-Layering in AesEval-Benchmark.

Benchmark Showcase

▷ Layout-Balance





Explanation: The overall layout of the picture is bal- The background of the picture is moved upwards, and anced and coordinated. the balance of the whole picture is broken.

Table 13: Examples of Layout-Balance in AesEval-Benchmark.

Benchmark Showcase

⊳ Layout-Hierarchy





Explanation: All the fonts in the picture are consis- The font in the picture is disturbed and looks like it is tent, which gives a sense of hierarchy.

The font in the picture is disturbed and looks like it is not on the same level as the previous font. There is no sense of hierarchy.

Table 14: Examples of Layout-Hierarchy in AesEval-Benchmark.

Benchmark Showcase

▷ Graphic-Quality, Color-Contrast





Explanation: Clear background, blue and white colors have good contrast.

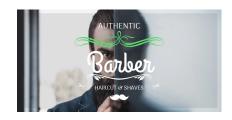
The background is blurred and the color changes from white to blue, with no contrast to the background.

Table 15: Examples of Graphic-Quality and Color-Contrast in AesEval-Benchmark.

Benchmark Showcase

▷ Color-Harmony, Color-Appealing, Color-Psychology





Explanation: The whole picture has harmonious and beautiful colors.

The middle element turns green, which makes the whole picture look disharmonious and unappealing. Green is strange and cause bad psychological effects.

Table 16: Examples of Color-Harmony, Color-Appealing and Color-Psychology in AesEval-Benchmark.