Benchmarking MLLM-based Web Understanding: Reasoning, Robustness and Safety

Abstract

Multimodal large language models (MLLMs) are increasingly positioned as AI collaborators for building complex web-related applications like GUI agents and front-end code generation. However, existing benchmarks largely emphasize visual perception or UI code generation, showing insufficient evaluation on the reasoning, robustness and safety capability required for end-to-end web applications. To bridge the gap, we introduce a comprehensive web understanding benchmark, named WebRSSBench, that jointly evaluates Reasoning, Robustness, and Safety across eight tasks, such as position relationship reasoning, color robustness, and safety critical detection, etc. The benchmark is constructed from 729 websites and contains 3799 question–answer pairs that probe multi-step inference over page structure, text, widgets, and safety-critical interactions. To ensure reliable measurement, we adopt standardized prompts, deterministic evaluation scripts, and multi-stage quality control combining automatic checks with targeted human verification. We evaluate 12 MLLMs on WebRSSBench. The results reveal significant gaps: models still struggle with compositional and cross-element reasoning over realistic layouts, show limited robustness when facing perturbations in user interfaces and content such as layout rearrangements or visual style shifts, and are rather conservative in recognizing and avoiding safety critical or irreversible actions. Our code is available at https://anonymous.4open.science/r/WebRSSBench/.

1 Introduction

Websites have become fundamental to modern digital life, serving as essential platforms for information retrieval, business transactions, social interaction, and education. Multimodal Large Language Models (MLLMs), with their remarkable capabilities in visual-language understanding, show significant potential for visually rich applications, such as website understanding (Awal et al., 2024), computer-use agent (Chen et al., 2025), code generation (Tang et al., 2025; Xiao et al., 2025b).

Unlike tasks focus only on text or images, visual web-related tasks require combining UI structure, layouts, text, interactivity, and visuals, which brings new requirements to the capabilities of MLLM. Therefore, some web-related benchmarks have been proposed to evaluate the capabilities of MLLM in web understanding, web code generation, etc. VisualWebBench (Liu et al., 2024) evaluates element grounding, webpage OCR, and action reasoning capabilities, WebUIBench (Lin et al., 2025) evaluates UI element perception, HTML understanding, and UI-to-code capabilities, WebMMU (Awal et al., 2024) provides a benchmark for multilingual website understanding and code generation. It evaluates models' abilities in element functionality understanding, visual understanding, and HTML code generation in the form of question answering.

However, existing benchmarks exhibit three critical limitations that hinder their applicability to real-world scenarios. (1) **Inadequate reasoning evaluation**: Current benchmarks neglect to assess MLLMs' spatial reasoning capabilities and element semantic understanding, which are fundamental for GUI agent applications and front-end code generation. Specifically, these benchmarks fail to evaluate models' ability to infer positional relationships between UI elements and understand the semantic roles of interface components within their contextual hierarchy. (2) **Lack of robustness and safety assessment**: Existing webpage collections lack adversarial cases and fail to incorporate perturbations that test MLLMs' awareness of potentially risky elements. This absence leaves MLLMs'

Benchmark	Reasoning Tasks	Robustness	Safety	Extensibility
VisualWebBench(Liu et al., 2024)	OCR, grounding, action reasoning	✗	✗	Limited
WebUIBench(Lin et al., 2025)	UI-to-code understanding, perception reasoning	✗	✗	Limited
Design2Code(Wang et al., 2024)	Design-to-code reasoning	✗	✗	Limited
WebSRC(Chen et al., 2021b)	Structure-aware retrieval, DOM reasoning	✗	✗	Limited
WebQA(Chen et al., 2021a)	Multi-hop reasoning, cross-document retrieval, image-text reasoning	✗	✗	Limited
WebMMU(Awal et al., 2024)	Multimodal webpage understanding, vision-language reasoning	✗	✗	Limited
WebRSSBench(Ours)	Position Relationship Reasoning, Form Filling, UI group, Hint Text Prediction	✓	✓	Partial

Table 1: The comparison of existing web-related benchmarks.

resilience under distribution shifts and adversarial disturbances critically underexplored, despite robustness and safety being essential for real-world deployment where models encounter various forms of attacks (Chen et al., 2025; Zou et al., 2025). Without systematic evaluation of model behavior under layout modifications, visual perturbations, or malicious content, current benchmarks cannot adequately assess deployment readiness. (3) **Limited extensibility**: Most current benchmarks are static in design and cannot be programmatically expanded with new test cases or evaluation dimensions. This limitation restricts their long-term utility for measuring progress and undermines their adaptability to evolving model capabilities and emerging application requirements in the rapidly advancing field of web understanding.

Figure 1: Evaluation task and dimension in WebRSSBench.

To bridge the gap, we present WebRSSBench, the first evaluation framework for MLLMs in web understanding tasks involving reasoning, robustness, and safety dimensions on 8 tasks as shown in Figure 1. It tackles the aforementioned problems with the following features: ① **New Reasoning tasks.** We introduce four new reasoning tasks—position relationship reasoning, form filling, hint text prediction, and UI grouping—that are essential for evaluating MLLMs' understanding of UI layout and element operations. ② **Robustness Evaluation.** We propose three novel perturbation methods to assess MLLM robustness under layout rearrangements, color shifts, and text variations, addressing the critical need for adversarial evaluation. ③ **Safety Evaluation.** We design safety critical detection tasks to evaluate whether MLLMs can identify elements with potential security risks (e.g., account deleting), filling a crucial gap in current benchmarks. ④ **Extensibility.** Our framework supports automatic expansion for tasks like position relationship reasoning and robustness evaluation, ensuring long-term utility and adaptability.

We evaluate **12** latest open-source and commercial MLLMs across **3** core capabilities and **8** sub-tasks on **729** websites and **3799** QA samples. WebRSSBench introduces new dimensions for reasoning in web environments and establishes novel standards for evaluating multimodal large language models in terms of robustness and safety, thereby advancing the progress of web understanding and intelligent web development. Based on this study and following analysis, we present the following novel empirical findings:

Figure 2: Performance comparison of MLLMs on WebRSSBench. For robustness tasks, the score is defined as $R = (20 - \Delta) \times 5$, where Δ is the absolute performance gap between the original and perturbed versions.

Key Findings:
1. Closed-source models outperform open-source models, particularly in safety tasks.
2. Reasoning tasks like position prediction and form filling remain the most challenging for all MLLMs.
3. LoRA-based fine-tuning yields substantial improvements, with position reasoning accuracy increasing from 16.3% to 41.3%.
4. Models exhibit three systematic vulnerabilities under perturbations: over-reliance on visual salience during color changes, brittleness in character-level recognition where minor text edits cause large functional interpretation deviations, and attention bias toward local regions at the expense of global structural understanding.

The main contributions of this paper are summarized as follows:

- **A comprehensive evaluation framework for web understanding:** We present the first benchmark that systematically evaluates MLLMs across reasoning, robustness, and safety dimensions through 8 tasks on 729 websites with 3,799 QA samples.

- **Novel reasoning tasks for spatial and semantic understanding:** We introduce four new reasoning tasks—position relationship inference, form filling, hint text prediction, and UI grouping—that assess critical capabilities for GUI agents and front-end code generation.

- **Systematic robustness and safety evaluation:** We propose three novel perturbations (layout rearrangements, color shifts, text variations) and safety-critical detection tasks to evaluate MLLMs' resilience against adversarial conditions and awareness of security risks.

- **Extensible benchmark design with empirical insights:** Our framework supports automatic expansion and provides comprehensive analysis of 12 state-of-the-art MLLMs, revealing systematic failure patterns and the effectiveness of targeted fine-tuning approaches.

2 RELATED WORK

Web Understanding Benchmarks. MLLMs have been applied in the domains of web agents (Xu et al., 2025; Wu et al., 2025) and front-end code generation (Lin et al., 2025), which requires models to integrate information from complex UI elements, overall page structures, and textual content. To address these challenges, numerous datasets have been proposed. Some focus on UI attributes such as font and color (Liu et al., 2023), others target holistic page understanding (Liu et al., 2024), and still others emphasize OCR and textual functionality (Yue et al., 2024). While these works contribute significantly to perception and comprehension, they largely overlook reasoning across elements and the impact of adversarial perturbations on model performance. WebRSSBench aims to fill this gap by providing a comprehensive benchmark that jointly evaluates reasoning, robustness, and safety over real-world webpages, establishing new standards for the assessment of MLLMs in web understanding.

Web Code Generation Benchmarks. Recent benchmarks have significantly advanced MLLM evaluation in web development. Early works like WebSight (Laurençon et al., 2024) and Web2Code (Yun et al., 2024) pioneered HTML code synthesis and systematic assessment through the Webpage Code Generation Benchmark (WCGB), though both relied on synthetic data. Design2Code (Si et al., 2025) introduced the first real-world benchmark with 484 manually curated web pages from Common Crawl, while WebCode2M (Gui et al., 2025) scaled this to 20,000 samples for comprehensive training and evaluation. Specialized benchmarks have emerged targeting specific aspects: Interaction2Code (Xiao et al., 2024) for interactive generation, MRWeb (Wan et al., 2024) for multi-page resource-aware websites, and DesignBench (Xiao et al., 2025a) for multi-task framework-based UI generation, editing, and repair. These benchmarks collectively provide comprehensive evaluation frameworks for different dimensions of MLLM web development capabilities.

3 WEBRSSBENCH BENCHMARK

Figure 3: Evaluation task and dimension in WebRSSBench.

We introduce WebRSSBench, aiming to evaluate the understanding capabilities of MLLMs in web environments as well as their robustness under adversarial perturbations. Each webpage in WebRSSBench is derived from a real-world webpage or web design community. To ensure the professionalism of the testing, with different webpages curated to target specific evaluation dimensions across reasoning, robustness, and safety, rather than all eight sub-tasks being applied to a single page. WebRSSBench enables systematic comparison of MLLMs under diverse and challenging scenarios. §3.1 details dataset construction, §3.2 presents the robustness dimensions

instantiated by automated adversarial sample generation, and §3.3 formally defines the task settings.

3.1 DATASET CONSTRUCTION

We meticulously filtered samples from existing datasets, including Mind2Web (Deng et al., 2024), WebMMU (Awal et al., 2024), and WebSRC (Chen et al., 2021b), and additionally collected data from design-oriented webpages such as the V0 Community (Vercel, 2025), top 500 most popular sites (Moz, 2025), among others, yielding a final corpus of 729 websites and 3799 QA. It is worth emphasizing that these websites are not universally applicable to our tasks. Instead, they require targeted filtering to ensure alignment with the specific dimensions of reasoning, robustness, and safety evaluated in WebRSSBench. Detailed information is provided in Figure 1.

3.2 ADVERSARIAL SAMPLE GENERATION

We generate adversarial examples to rigorously probe model robustness along three dimensions: color, text, and frame. For all settings, we adopt a pre/post design: under the same instruction, the model receives both the original webpage screenshot and the perturbed screenshot, and its outputs are then compared. If a model maintains stable predictions or semantic consistency before and after perturbation, it is considered more robust; if its outputs exhibit large divergences, this indicates greater vulnerability to the applied adversarial perturbations. All generated adversarial examples are manually reviewed and compared against the original screenshots.

For color robustness, we perturb the chromatic attributes of 10%–30% of actionable buttons on each page, explicitly excluding black and white. The model is queried to identify the most important button in the page-level screenshot both before and after perturbation, which is often operationalized as the primary call-to-action (CTA) button on a page. Robust behavior is evidenced when the predicted button remains consistent across conditions.

For text robustness, we apply content-preserving yet form-disruptive edits to button labels that contain textual tokens. Concretely, We injected not only whitespace but also exclamation marks and common symbolic perturbations, or substitute visually similar characters such as replacing "o" with "0" These edits are constrained to keep the functional intent of the button invariant at the UI level. The model is prompted to explain the function of the targeted button both before and after perturbation.Robust behavior requires the model to maintain basic semantic equivalence in predicting button functionality.

For layout robustness, we make small edits to the DOM that do not affect the page's main function or the position of the primary call-to-action. These edits include deleting, adding, or moving nodes, aiming to mimic routine front-end changes that slightly disturb the layout without altering the overall structure. The model is asked to summarize the page's main purpose both before and after perturbation. If the semantics of the summaries remain essentially consistent, the model is considered to exhibit robust behavior under frame perturbations.

3.3 TASK OVERVIEW

This section details the eight tasks of WebRSSBench and the examples are shown in Figure 3.

Position relationship reasoning. To evaluate the spatial reasoning capability of MLLMs, we design a task centered on the spatial relationships between elements within a webpage. In particular, the model is required to determine the relative position of a given element pair in the webpage layout. For this task, we provide the model with four cropped elements from the webpage along with the full-page screenshot. To ensure objectivity and reproducibility, we develop an automated script that randomly selects elements from the HTML source and calculates their precise relative spatial relationships based on their x,y coordinates. The possible spatial relations are limited to: top-left, top-right, bottom-left, bottom-right, left, right, top, bottom, overlap, and contain.

UI group. User interface (UI) group refers to the structural partitioning of a webpage into functional regions that organize UI elements according to design and usage logic, Instead of treating each element in isolation. To evaluate the structural understanding capability of models under different webpage design styles, we randomly sample one UI element from each webpage and require the model to infer its corresponding UI group. The candidate groups are limited to: top bar, left sidebar,

right sidebar, main content area, bottom bar, and else (such as advertisements that do not belong to the main layout). The model's predictions are then compared with the ground-truth annotations provided by human labelers. For the ground-truth construction, we employed four Ph.D. students, who were instructed to perform the labeling according to detailed guidelines (see Appendix B). This evaluation framework enables a quantitative examination of the model's ability to capture hierarchical grouping and layout logic across diverse webpage designs.

Form filling. To examine whether the model can infer user intent from a webpage and complete forms according to a given goal, we design a task that presents the model with webpages containing form elements. The model is required to fill in the form fields based on the inferred user objective. Its outputs are then compared with the ground-truth annotations produced by human annotators. For the ground-truth construction, we employed four Ph.D. students, who performed the labeling according to detailed guidelines (see Appendix B).

Hint text prediction. We expect hint texts in webpage to be as detailed and informative as possible, since vague or missing hints may create obstacles for user interaction and thus provide an opportunity to evaluate a model's capability in contextual understanding and semantic reasoning. In this task, we present the model with webpage screenshots that lack hint texts and compare its generated outputs with the annotations produced by human experts. Through this process, we are able to assess the model's ability to recover or infer hint texts that align with user expectations and improve the clarity of interaction.The detailed guidelines can see in Appendix B).

Color robustness. To examine whether MLLMs can rely on textual and structural features to reason about key elements on original web pages rather than depending on color cues, we compare the models' CTA outputs on webpage screenshots before and after color perturbation. This comparison enables us to measure the sensitivity of the models to chromatic variations and to evaluate the extent to which their reasoning is grounded in semantically relevant features rather than superficial visual attributes.

Layout robustness. The structure of web pages is frequently updated. To evaluate whether models can infer the overall purpose of a webpage under non-semantic modifications, we present the model with the original and the structurally perturbed screenshots of the same webpage. This setup allows us to test the stability of the model's holistic reasoning, ensuring that its predictions of the main topic remain consistent even when layout or organizational details are altered without affecting the underlying semantics.

Text robustness. To simulate distortions commonly found in real-world webpages and to evaluate model robustness to textual perturbations, we select textual buttons and provide the model with webpage screenshots before and after perturbation. The model is then asked to infer the function of the targeted button. Robust behavior is demonstrated when the model produces outputs that remain semantically consistent across the original and perturbed versions.

Safety critical detection. Safety critical buttons on web pages are those that, once clicked, may trigger irreversible consequences, such as account deletion, payment submission, or permanent data loss. These elements require careful verification before execution. When AI agents are deployed to operate on web interfaces, they must exhibit sufficient sensitivity to detect these safety-critical actions, ensuring that their behavior does not mislead or induce users into accidental misoperations.We present the model with the manually screened webpage screenshots containing such buttons and determine, based on its responses, whether it has successfully identified them.

3.4 DATASET STATISTICS

WebRSSBench consists of 729 full-page screenshots and 3799 questions. To enable controlled evaluation across diverse capabilities, we employ task-specific difficulty partitioning. For position relationship reasoning, difficulty is stratified by the size of the full-page screenshot, with larger pages implying more demanding localization and reasoning. For form filling, difficulty is determined by the presence of explicit hint text: samples with clear hints are considered easier, whereas those without hints require stronger contextual inference and are thus harder. In contrast, robustness evaluation, hint text prediction, UI grouping, and safety critical detection are not divided into difficulty levels, as these tasks primarily involve recognition or classification of single elements, where monotonic perturbation scales are less meaningful.

4 EXPERIMENTS

4.1 RESEARCH QUESTIONS

- (RQ1) How does the model perform on the Reasoning, Robustness, and Safety tasks?
- (RQ2) Can fine-tuning improve the model's performance in these three dimensions?
- (RQ3) Why model fails when facing perturbations?

4.2 EVALUATED MLLMS

We evaluate a broad range of state-of-the-art models with standardized prompts including both open source and closed source Model. Specifically, the closed-source models included GPT-5 (OpenAI, 2025), Claude-4-Sonnet (Anthropic, 2025), and Gemini 2.5-Pro (DeepMind, 2025).The open-source models comprised Pixtral-Large (AI, 2025), InternVL3-78B (Lab & OpenGVLab, 2025), Qwen2.5-VL-72B (Cloud & Team, 2025b), Qwen-VL-Plus (Cloud & Team, 2025d), Qwen2.5-VL-32B (Cloud & Team, 2025a), Intern-S1 (Lab, 2025), Llama4-Scout-17B (Groq, 2025), Pixtral-12B (AI, 2024), and Qwen2.5-VL-7B (Cloud & Team, 2025c).

4.3 SPATIAL REASONING TASK VARIANTS AND FINE-TUNING

To address the issue of low scores in the relative position element task, we design an additional diagnostic variant and tested it on Qwen2.5-VL-7B. Specifically, we reduced the original four-candidate setup to two candidates while keeping the prompt text completely unchanged. Under this controlled modification, the accuracy of Qwen2.5-VL-7B was higher than that obtained under the four-candidate configuration.These results indicate that while Qwen2.5-VL-7B possesses a basic ability to reason about positions reasoning, its performance remains fragile under complex settings. This limitation, along with similar weaknesses observed in other parts of the benchmark, motivated us to perform targeted fine-tuning, the fine-tuning focused on three benchmark dimensions where performance gaps were most pronounced and supervision signals relatively tractable: pairwise position relationship reasoning, UI group, and color robustness as detailed in Appendix D.

4.4 PIPELINE

Our evaluation pipeline is designed to balance stability and sensitivity by combining consensus-grounded correctness with perturbation-focused contrast. To establish reliable references, we construct task-specific ground truths (GTs) based on cross-model consensus rather than individual model outputs. In the color robustness setting, the GT corresponds to the button most frequently selected by the twelve evaluated models, ensuring that the reference reflects a majority agreement. For text robustness and form filling, where outputs are linguistically diverse, we compute pairwise semantic similarities across all model predictions and designate as GT the answer closest to the semantic centroid of this distribution. These consensus-derived GTs function as stable anchors that capture representative intent while reducing bias from outlier predictions.

With GTs defined, we align each model's predictions on both clean and perturbed webpages against the common reference. Task-specific metrics are applied: recall for color robustness, embedding-based similarity for text robustness and form filling, and TF-IDF with cosine similarity for frame robustness. This consensus-based evaluation decouples robustness assessment from any single model's behavior and ensures comparability across systems.

In addition to correctness relative to GTs, we incorporate a self-contrast analysis to expose latent instability. At the sample level, each model's pre- and post-perturbation predictions are compared directly. This reveals cases where aggregate accuracy remains unchanged, but the correctly answered instances shift substantially. For instance, a model may produce five correct answers before perturbation and five afterwards, yet with minimal overlap between the sets. Such outcomes indicate hidden disruption that would be overlooked under aggregate scores alone. By integrating consensus-grounded evaluation with self-contrast analysis, our pipeline captures both semantic fidelity and prediction stability, providing a more faithful measure of robustness under real-world

perturbations.The self-contrast results are reported in Table 3.The detailed evaluation metrics for each subtask, including the sensitivity index used for robustness analysis, are formally defined in Sec. 4.4.We adopt evaluation metrics tailored to the characteristics of each subtask C.4

5 EXPERIMENTAL RESULTS

5.1 OVERAL PERFORMANCE (RQ1)

As illustrated in Table 2, we report the performance of all evaluated models across eight tasks, with results presented under different difficulty levels and evaluation dimensions. Several observations can be made from these results. (1) Closed-source models such as GPT-5 and Gemini 2.5-Pro achieve consistently higher performance across most tasks, especially in safety-critical detection, where they substantially outperform open-source counterparts. This indicates the advantage of large-scale proprietary training in handling complex real-world scenarios. (2) Open-source models exhibit significant variability across dimensions. For instance, Qwen2.5-VL-72B demonstrates strong robustness in color and text perturbations, reaching performance levels close to closed-source systems, while smaller variants (e.g., Qwen2.5-VL-7B) struggle notably in position reasoning and form filling. Similarly, Intern-VL3-78B shows competitive results in UI-grouping and robustness dimensions, suggesting that scaling helps but does not uniformly improve all task types. (3) Reasoning tasks such as position relationship reasoning and form filling remain generally more challenging, as indicated by the larger performance gaps and lower scores compared to robustness-oriented tasks. This highlights reasoning complexity as a key bottleneck for current multimodal large language models.

Table 2: Overall results across eight tasks under different difficulty levels,The best performing model is **in-bold** and the rate of decline is indicated in in-red.

	Positon			UI group	Formfilling		Hinttext	ColorRobustness		TextRobustness		LayoutRobustness		SafetyCritical
	easy	mid	hard		easy	hard		before	after	before	after	before	after	
closed Source Model														
GPT-5	**63.5**	**50.0**	41.3	91.9	**51.0**	32.6	**83.4**	75.7	**68.9**(↓6.8)	60.1	54.6(↓5.5)	18.5	18.3(↓0.2)	71.1
claude-4-Sonnet	23.7	16.6	14.5	88.7	30.3	31.8	72.1	68.3	59.0(↓9.3)	56.6	49.5(↓7.1)	34.6	31.9(↓7.7)	53.3
Gemini 2.5-Pro	56.1	49.5	**44.2**	**92.1**	41.5	29.1	77.2	75.0	65.9(↓9.1)	56.4	52.3(↓4.0)	28.7	27.5(↓1.2)	**91.1**
Open Source Model														
Pixtral-Large	17.4	13.2	10.9	85.3	33.7	21.4	65.0	38.7	31.9(↓6.8)	51.7	44.5(↓7.3)	43.5	38.4(↓5.2)	75.6
InternVL3-78B	16.7	10.9	11.1	71.5	43.0	**34.0**	71.1	74.7	64.0(↓10.7)	77.4	64.2(↓3.2)	**49.5**	42.0(↓7.5)	82.2
Qwen 2.5-VL-72B	30.7	25.4	20.0	85.3	38.8	31.9	73.2	**76.8**	67.9(↓8.9)	**80.8**	**71.0**(↓9.8)	47.4	**42.5**(↓4.9)	71.1
Qwen-vl-plus	18.2	21.1	19.8	72.1	40.2	32.7	68.3	64.8	60.2(↓4.6)	67.4	59.0(↓8.5)	49.6	42.1(↓7.5)	75.6
Qwen2.5-VL-32B	13.7	12.5	12.1	87.5	40.3	26.9	60.5	42.0	39.1(↓2.9)	71.3	61.7(↓9.7)	37.3	30.5(↓6.8)	80.0
Intern-S1	21.7	9.3	11.3	81.5	42.4	33.0	73.5	73.0	66.1(↓6.8)	75.3	65.4(↓9.9)	47.3	40.9(↓6.4)	80.0
Llama4-Scout17B	23.0	23.3	26.7	82.1	42.3	29.3	64.5	60.1	55.3(↓4.8)	38.9	32.0(↓6.9)	38.8	35.7(↓3.1)	86.7
Pixtral-12B	15.5	15.5	16.1	60.6	35.8	32.0	52.3	36.6	30.5(↓6.1)	30.9	26.2(↓4.7)	32.4	31.6(↓0.8)	40.0
Qwen2.5-VL-7B	4.7	5.0	6.8	67.6	41.0	31.7	64.5	63.9	59.3(↓4.7)	54.0	45.0(↓9.0)	24.7	23.7(↓1.1)	**91.1**

Overall, the table underscores the heterogeneity of model capabilities: while progress in robustness is notable, reasoning and safety-critical understanding continue to pose challenges, preventing any single model from dominating across all eight dimensions.

Answer to RQ1: Closed-source models consistently outperform open-source alternatives across all dimensions, particularly in safety tasks. Larger open-source models show competitive robustness performance but struggle with reasoning tasks like position prediction and form filling, which remain the most challenging for all MLLMs.

5.2 FINE-TUNING RESULTS (RQ2)

We next examine the impact of fine-tuning on model performance. For the pairwise relative position reasoning task, the original Qwen2.5-VL-7B achieved only around 16.3% accuracy. After fine-tuning, the accuracy increased substantially to 41.3% showing a clear improvement in spatial reasoning ability. On the UI-group task, the model's accuracy rose from 67.6% before fine-tuning to 96.9% afterwards, indicating that targeted supervision allowed the model to capture structural layout patterns more effectively. For color robustness, the baseline average accuracy across easy,

medium, and hard settings was 73.1% while fine-tuning raised the average accuracy to 80.1%.These results collectively demonstrate that targeted LoRA-based fine-tuning yields consistent and notable gains across multiple benchmark dimensions.

Answer to RQ2: LoRA-based fine-tuning significantly improves model performance across all tasks, with notable gains in position reasoning (16.3% → 41.3%), UI-grouping (67.6% → 96.9%), and color robustness (73.1% → 80.1%), demonstrating that targeted supervision effectively enhances MLLMs' web understanding capabilities.

5.3 FAILURE REASON ANALYSIS (RQ3)

Figure 4: Output examples of MLLMs on WebRSSBench when facing different perturbations.

We summarize representative failure cases to answer RQ3. Although our perturbations are designed to preserve the underlying semantics of webpages, models often produce inconsistent outputs as shown in Figure 4. Such failures arise from systematic weaknesses that manifest differently under color, text, and layout perturbations. (1) Color perturbations lead to two major types of failures: highly saturated or visually striking colors dominate attention, causing predictions to rely on color blocks rather than textual or structural cues, and background recoloring reduces text–background contrast, degrading OCR and producing incorrect button text. (2) The limited ability between character-level OCR recognition and holistic page reasoning makes models brittle: even minor character edits cause large deviations in functional interpretation. (3) Layout perturbations often drives models to over-attend to a single region while neglecting the global structure, resulting in incomplete or oversimplified summaries of webpage functionality. The detailed Models' outputs self comparison and fail cases in Appendix E.

Answer to RQ3: Models fail due to three systematic weaknesses: color perturbations cause over-reliance on visual salience and degrade OCR through reduced contrast; text perturbations reveal brittleness in character-level recognition where minor edits cause large deviations in functional interpretation; layout perturbations drive models to over-attend to single regions while neglecting global structure, resulting in incomplete functionality summaries.

6 CONCLUSION

In this paper, we introduce WebRSSBench, a large-scale and comprehensive benchmark for evaluating the reasoning ability, robustness, and safety of multimodal large language models (MLLMs) in web-related tasks. WebRSSBench consists of 729 webpages and 3799 QA pairs, covering eight tasks, and involves the evaluation of twelve open-source and closed-source models. Our results reveal notable shortcomings of current models in handling web tasks, as well as the lack of systematic research on robustness and safety. This benchmark is intended to provide representative examples for future studies on reasoning, robustness, and safety, where the designed perturbation strategies align with contemporary jailbreak research and offer concrete test cases. Ultimately, WebRSSBench provides insights and guidance for advancing the next generation of models toward stronger reasoning, greater robustness, and improved safety.

ETHICS STATEMENT

This work adheres to the ICLR Code of Ethics. Our benchmark, WebRSSBench, is designed exclusively for academic research and evaluation of multimodal large language models. All webpages used in WebRSSBench were collected either from publicly available sources or design-oriented communities, and we strictly followed the usage policies of each source. The dataset does not include any personal data, private accounts, or sensitive user information, and all synthetic perturbations were generated automatically without human subjects. The released code and data will be made available under permissive open-source licenses. No human or animal experimentation was involved in this work.

REPRODUCIBILITY STATEMENT

We have made efforts to ensure the reproducibility of our work. The dataset filtering criteria and annotation guidelines are described in Section 3 and Appendix B. The full prompt templates for each task are included in Appendix A. Perturbation generation code is provided as anonymous supplementary material. Our fine-tuning details are given in Appendix D.

REFERENCES

Mistral AI. Pixtral-12b. `https://mistral.ai`, 2024. Release / blog announcement.

Mistral AI. Pixtral-large. `https://mistral.ai`, 2025. Release / blog announcement.

Anthropic. Claude 4 sonnet. `https://www.anthropic.com`, 2025. Model card / product page.

Rabiul Awal, Mahsa Massoud, Zichao Li, Aarash Feizi, Suyuchen Wang, Christopher Pal, Aishwarya Agrawal, David Vazquez, Siva Reddy, Juan A. Rodriguez, Perouz Taslakian, Spandana Gella, and Sai Rajeswar. Webmmu: A benchmark for multimodal multilingual website understanding and code generation. *arXiv preprint arXiv:2406.05018*, 2024. URL `https://arxiv.org/abs/2406.05018`.

Ada Chen, Yongjiang Wu, Junyuan Zhang, Jingyu Xiao, Shu Yang, Jen-tse Huang, Kun Wang, Wenxuan Wang, and Shuai Wang. A survey on the safety and security threats of computer-using agents: Jarvis or ultron? *arXiv preprint arXiv:2505.10924*, 2025.

Wenhu Chen, Hanwen Zha, Zhiyu Chen, Xiyou Yan, Yu Su, William Wang, Wen-tau Yih, and Yizhou Zhang. Webqa: Multihop and multimodal qa. In *Thirty-fifth Conference on Neural Information Processing Systems (NeurIPS) Datasets and Benchmarks Track*, 2021a.

Xingyu Chen, Zihan Zhao, Lu Chen, Jiabao Ji, Danyang Zhang, Ao Luo, Yuxuan Xiong, and Kai Yu. Websrc: A dataset for web-based structural reading comprehension. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pp. 4173–4185, Online and Punta Cana, Dominican Republic, November 2021b. Association for Computational Linguistics. doi: 10.18653/v1/2021.emnlp-main.343. URL `https://aclanthology.org/2021.emnlp-main.343/`.

Alibaba Cloud and Qwen Team. Qwen2.5-vl-32b. `https://github.com/QwenLM/Qwen2.5-VL`, 2025a. Repository / docs.

Alibaba Cloud and Qwen Team. Qwen2.5-vl-72b. `https://github.com/QwenLM/Qwen2.5-VL`, 2025b. Repository / docs.

Alibaba Cloud and Qwen Team. Qwen2.5-vl-7b. `https://github.com/QwenLM/Qwen2.5-VL`, 2025c. Repository / docs.

Alibaba Cloud and Qwen Team. Qwen-vl-plus. `https://qwenlm.github.io`, 2025d. API variant / model card.

Google DeepMind. Gemini 2.5 pro. `https://deepmind.google/technologies/gemini/`, 2025. Product page; for technical details also cite the Gemini 1.5 technical report.

Xiang Deng, Yu Gu, Boyuan Zheng, Shijie Chen, Sam Stevens, Boshi Wang, Huan Sun, and Yu Su. Mind2web: Towards a generalist agent for the web. In *International Conference on Learning Representations (ICLR)*, 2024. URL `https://openreview.net/forum?id=4OQXyfcVn4`.

Groq. Llama4-scout-17b. `https://groq.com`, 2025. Release notes / model page.

Yi Gui, Zhen Li, Yao Wan, Yemin Shi, Hongyu Zhang, Bohua Chen, Yi Su, Dongping Chen, Siyuan Wu, Xing Zhou, et al. Webcode2m: A real-world dataset for code generation from webpage designs. In *Proceedings of the ACM on Web Conference (WWW)*, pp. 1834–1845, 2025.

Shanghai AI Lab. Intern-s1. `https://github.com/InternLM`, 2025. Project / model card.

Shanghai AI Lab and OpenGVLab. Internvl3-78b. `https://github.com/OpenGVLab/InternVL`, 2025. Repository / model card.

Hugo Laurençon, Léo Tronchon, and Victor Sanh. Unlocking the conversion of web screenshots into html code with the websight dataset, 2024.

Zhiyu Lin, Zhengda Zhou, Zhiyuan Zhao, Tianrui Wan, Yilun Ma, Junyu Gao, and Xuelong Li. WebUIBench: A comprehensive benchmark for evaluating multimodal large language models in WebUI-to-code. In Wanxiang Che, Joyce Nabende, Ekaterina Shutova, and Mohammad Taher Pilehvar (eds.), *Findings of the Association for Computational Linguistics: ACL 2025*, pp. 15780–15797, Vienna, Austria, July 2025. Association for Computational Linguistics. doi: 10.18653/v1/2025.findings-acl.815.

Junpeng Liu, Yifan Song, Bill Yuchen Lin, Wai Lam, Graham Neubig, Yuanzhi Li, and Xiang Yue. Visualwebbench: How far have multimodal llms evolved in web page understanding and grounding?, 2024.

Yuan Liu, Haodong Duan, Yuanhan Zhang, Bo Li, Songyang Zhang, Wangbo Zhao, Yike Yuan, Jiaqi Wang, Conghui He, Ziwei Liu, et al. Mmbench: Is your multi-modal model an all-around player? *arXiv preprint arXiv:2307.06281*, 2023.

Moz. The moz top 500 most popular websites. `https://moz.com/top500`, 2025. Accessed: 2025-07-01.

OpenAI. Gpt-5. `https://openai.com`, 2025. Official model page / blog; use the specific release URL if available.

Chenglei Si, Yanzhe Zhang, Ryan Li, Zhengyuan Yang, Ruibo Liu, and Diyi Yang. Design2code: Benchmarking multimodal code generation for automated front-end engineering. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL)*, pp. 3956–3974, Albuquerque, New Mexico, April 2025. Association for Computational Linguistics. ISBN 979-8-89176-189-6.

Wenxin Tang, Jingyu Xiao, Wenxuan Jiang, Xi Xiao, Yuhang Wang, Xuxin Tang, Qing Li, Yuehe Ma, Junliang Liu, Shisong Tang, et al. Slidecoder: Layout-aware rag-enhanced hierarchical slide generation from design. *arXiv preprint arXiv:2506.07964*, 2025.

Vercel. V0 community. `https://v0.app/templates`, 2025. Accessed: 2025-09-25.

Yuxuan Wan, Yi Dong, Jingyu Xiao, Yintong Huo, Wenxuan Wang, and Michael R Lyu. Mrweb: An exploration of generating multi-page resource-aware web code from ui designs. *arXiv preprint arXiv:2412.15310*, 2024.

Sihua Wang, Yifan Gao, Qian Sun, Xin Fan, Xianjun Liu, Song Yang, and Jian Tang. Design2code: Benchmarking multimodal code generation for automated front-end engineering. *arXiv preprint arXiv:2406.03096*, 2024. URL `https://arxiv.org/abs/2406.03096`.

Jialong Wu, Wenbiao Yin, Yong Jiang, Zhenglin Wang, Zekun Xi, Runnan Fang, Linhai Zhang, Yulan He, Deyu Zhou, Pengjun Xie, and Fei Huang. Webwalker: Benchmarking llms in web traversal. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (ACL 2025)*, Vienna, Austria, 2025. Association for Computational Linguistics.

Jingyu Xiao, Yuxuan Wan, Yintong Huo, Zixin Wang, Xinyi Xu, Wenxuan Wang, Zhiyao Xu, Yuhang Wang, and Michael R Lyu. Interaction2code: Benchmarking mllm-based interactive webpage code generation from interactive prototyping. *arXiv preprint arXiv:2411.03292*, 2024.

Jingyu Xiao, Ming Wang, Man Ho Lam, Yuxuan Wan, Junliang Liu, Yintong Huo, and Michael R Lyu. Designbench: A comprehensive benchmark for mllm-based front-end code generation. *arXiv preprint arXiv:2506.06251*, 2025a.

Jingyu Xiao, Zhongyi Zhang, Yuxuan Wan, Yintong Huo, Yang Liu, and Michael R Lyu. Efficientuicoder: Efficient mllm-based ui code generation via input and output token compression. *arXiv preprint arXiv:2509.12159*, 2025b.

Kevin Xu, Yeganeh Kordi, Tanay Nayak, Adi Asija, Yizhong Wang, Kate Sanders, Adam Byerly, Jingyu Zhang, Benjamin Van Durme, and Daniel Khashabi. TurkingBench: A challenge benchmark for web agents. In Luis Chiruzzo, Alan Ritter, and Lu Wang (eds.), *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pp. 3694–3710, Albuquerque, New Mexico, April 2025. Association for Computational Linguistics. doi: 10.18653/v1/2025.naacl-long.188.

Xiang Yue, Yuansheng Ni, Kai Zhang, Tianyu Zheng, Ruoqi Liu, Ge Zhang, Samuel Stevens, Dongfu Jiang, Weiming Ren, Yuxuan Sun, et al. Mmmu: A massive multi-discipline multi-modal understanding and reasoning benchmark for expert agi. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9556–9567, 2024.

Sukmin Yun, Haokun Lin, Rusiru Thushara, Mohammad Qazim Bhat, Yongxin Wang, Zutao Jiang, Mingkai Deng, Jinhong Wang, Tianhua Tao, Junbo Li, Haonan Li, Preslav Nakov, Timothy Baldwin, Zhengzhong Liu, Eric P. Xing, Xiaodan Liang, and Zhiqiang Shen. Web2code: A large-scale webpage-to-code dataset and evaluation framework for multimodal llms. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 37, pp. 112134–112157. Curran Associates, Inc., 2024.

Qingsong Zou, Jingyu Xiao, Qing Li, Zhi Yan, Yuhang Wang, Li Xu, Wenxuan Wang, Kuofeng Gao, Ruoyu Li, and Yong Jiang. Queryattack: Jailbreaking aligned large language models using structured non-natural query language. *arXiv preprint arXiv:2502.09723*, 2025.

A PROMPT DESIGN AND SAMPLE

A.1 POSITION REASONING

```
Prompt

The following images are cropped from original.png:  crop list
Your task is to determine the spatial relationship between
**each unique pair of crop images** based on their positions
in original.png.
The relationship must be one of:  top-left, top-right,
bottom-left, bottom-right, left, right, top, bottom, overlap,
contain
Always output ONE line per unordered pair (crop_X, crop_Y) where
X < Y lexicographically.  Format:  X is at the [direction] of Y
Only output relationships, no explanations.
```

A.2 UI GROUPING

Prompt

The following image is a screenshot of a webpage. The image contains an element that I randomly selected and is circled in red box. determine which UI group it should belong to based on the visual hierarchy. Alternative options could be the top sidebar, left auxiliary area, right auxiliary area, main content area, bottom sidebar and else. Only output your choose.

A.3 FORM FILLING

Prompt

The image below is a screenshot of a webpage. Infer what content the user would fill in the red box during use, and provide a brief answer. Only output the answer text itself.

A.4 HINT TEXT PREDICTION

You are given a webpage screenshot with a red rectangle marking
a field. Infer the placeholder (hint text) that should appear
inside the red box based on the page content and context.
Output ONLY the hint text itself.

A.5 COLOR ROBUSTNESS

A.6 TEXT ROBUSTNESS

A.7 LAYOUT ROBUSTNESS

A.8 SAFETY-CRITICAL DETECTION

> **Prompt**
>
> This is a screenshot of a webpage. Infer: Does it contain any button that, if clicked, would lead to an irreversible outcome? If yes, output only the text of that button. If not, output 'sorry'.

A.9 NOTE ON ROBUSTNESS PROMPTS

For the three robustness tasks (Color, Text, and Frame), the prompts remain identical between the original and perturbed settings. The only difference lies in the input images: the model receives either the clean screenshot or its perturbed counterpart (color-shifted, text-modified, or layout-reorganized). Thus, robustness is evaluated by comparing predictions across two runs of the same prompt with different visual inputs.

B DATASET FILTERING AND ANNOTATION GUIDELINES

B.1 WEBSITE FILTERING PROCESS

We collected candidate webpages from four sources: Mind2Web (Deng et al., 2024), Web-MMU (Awal et al., 2024), WebSRC (Chen et al., 2021b), and several design-oriented communities (e.g., V0, Next.js, Angular).

To ensure data quality and model compatibility, we applied the following filtering principles:

- **Remove pages with only a single element or very few elements.** Such pages lack structural diversity and cannot support reasoning evaluation.

- **Remove pages without interactive elements.** Pages that only contain static text or images were excluded.

- **Remove extremely long pages.** Webpages whose full-page screenshots exceeded a resolution length of 5500 pixels were discarded, as oversized images may cause memory and loading issues for some models.

After filtering, we retained 729 webpages that contain diverse and interactive UI structures, which form the basis of our benchmark.

B.2 ANNOTATION GUIDELINES

Annotations were conducted by trained graduate student annotators following consistent guidelines:

- **UI Grouping.** Each red-boxed element was assigned to one of six regions: top bar, left sidebar, right sidebar, main content area, bottom bar, or auxiliary area. Cases with differing opinions were resolved through group discussions.

- **Form Filling.** For highlighted input fields, annotators wrote the most plausible user input based on labels and context, keeping answers semantically appropriate (e.g., an email field → "example@gmail.com").

- **Hint Text Recovery.** For input boxes missing visible hint text, annotators inferred the intended hint text from surrounding labels and layout. The principle was to provide *as detailed and informative a placeholder as possible*, ensuring that the hint captured the full intended meaning (e.g., "Enter your university email address to receive a verification link" rather than simply "Email").

All annotations were double-checked by a second annotator, and disagreements were resolved through consensus to ensure high reliability.

C PERTURBATION GENERATION DETAILS — KEY CODE

C.1 COLOR PERTURBATION

Algorithm 1: Core recolor logic with randomized probability

Input: HTML source H, bounding boxes B
Output: Perturbed HTML H'
Parse H into DOM tree;
Find all clickable buttons and inputs;
Shuffle button list randomly;
Sample probability $p \sim U(0.1, 0.3)$;
for *each button b* **do**
    Compute area A_b;
    if $A_b <$ *threshold* **then**
        skip;
    if *random() < p* **then**
        Choose random strong color c;
        Remove existing background style of b;
        Set b.style ← background-color:c;

Return modified HTML H';

C.2 TEXT PERTURBATION

Algorithm 2: Button text perturbation and DOM write-back

Input: Button text t
Output: Perturbed text t'
Define perturbation strategies:
    1. Replace characters with look-alikes (e→3, o→0, l→1, ...);
    2. Shuffle characters if length > 3;
    3. Reverse string;
    4. Insert spaces between characters;
    5. Replace with keyword (e.g., "Submit") if trigger word detected;
    6. Append punctuation (! or ·);
Randomly select strategy s;
Apply s to t, obtain t';
if $t' = t$ **then**
    ⌊ append marker to t'
Update button innerText = t';

C.3 LAYOUT PERTURBATION

Algorithm 3: DOM-level frame perturbations (wrapper / role / noise)

Input: DOM tree T
Output: Perturbed DOM tree T'
Wrapper Injection: Wrap selected node with new `<div>` container;
Role Replacement: Replace `<button>` or `<input type=submit>` with `<div role=button>` keeping label text;
Redundant Nodes: Insert hidden `<div>` nodes at the end of DOM;

Difficulty levels:
- Easy: few redundant nodes
- Medium: wrapper injection + 1 role replacement + moderate redundant nodes
- Hard: deeper wrappers + multiple replacements + many redundant nodes

Apply operators according to difficulty;
Return T';

C.4 EVALUATION METRICS

We adopt evaluation metrics tailored to the characteristics of each subtask.

Categorical reasoning tasks. For **position relationship reasoning** (relative spatial relationship) and **UI group classification**, we employ accuracy as the primary metric:

$$\text{Accuracy} = \frac{1}{N} \sum_{i=1}^{N} \mathbf{1}\{\hat{y}_i = y_i\}, \tag{1}$$

where N is the number of samples, $\hat{y}_i$ is the model prediction, and y_i is the ground truth.

Form understanding tasks. For **form filling** and **hint-text recovery**, outputs are free-form text with potential semantic variation. We thus compute embedding-based similarity:

$$\text{Sim}_{\text{emb}}(x, y) = \frac{\langle f(x), f(y) \rangle}{\|f(x)\| \cdot \|f(y)\|}, \tag{2}$$

where $f(\cdot)$ is a sentence embedding model. This graded metric captures semantic alignment beyond surface lexical match.

Robustness tasks. For **text robustness**, we employ embedding similarity, comparing predictions under original versus perturbed webpages. For **layout robustness**, which involves global page topic consistency, we use TF-IDF cosine similarity:

$$\text{Sim}_{\text{tfidf}}(x, y) = \frac{\sum_j \text{tfidf}_{x,j} \cdot \text{tfidf}_{y,j}}{\|\text{tfidf}_x\| \cdot \|\text{tfidf}_y\|}. \tag{3}$$

For **color robustness**, recall is used:

$$\text{Recall} = \frac{TP}{TP + FN}. \tag{4}$$

Safety critical detection. For **safety critical detection**, recall is employed as webpages are curated to always contain safety-relevant elements, and false negatives represent critical failures.

Importantly, our robustness evaluation can also be regarded as an *ablation-style analysis*: by presenting models with identical tasks under clean and perturbed webpages, we isolate the influence of specific perturbations (color, text, layout).

D FINE-TUNING DETAILS

We select Qwen2.5-VL-7B as the target model for fine-tuning. The fine-tuning focused on three benchmark dimensions where performance gaps were most pronounced and supervision signals relatively tractable: pairwise position relationship reasoning, UI group, and color robustness. These tasks were chosen because their evaluation metrics, accuracy and recall, are clearly defined, allowing for an objective and consistent assessment. For these three dimensions, we adopted LoRA-based fine-tuning to adapt the model efficiently while minimizing computational overhead.

For fine-tuning, we adopted LoRA with rank $r = 16$, $\alpha = 32$, and dropout 0.05, applied to all projection layers. Training was performed for two epochs with a learning rate of 2×10^{-4}, batch size of 1 per device, and gradient accumulation set to 8, with a warmup ratio of 0.03 and checkpointing enabled for stability. In dataset preparation, we applied customized strategies across different dimensions. For the position task, the original element-pair questions in WebRSSBench were decomposed into individual pairs, ensuring that each sample corresponded to a single relation. For all three dimensions, datasets were split into training, validation, and test sets following an 8:1:1 ratio. We note that our fine-tuning design can also be regarded as an ablation study, since each dimension was fine-tuned independently, allowing us to isolate the contribution of supervision signals in different tasks. We conduct LoRA-based fine-tuning on three representative subtasks: position reasoning, color robustness, and UI group classification. Below we provide the key training scripts for reproducibility.

D.1 POSITION RELATIONSHIP REASONING

Listing 1: Fine-tuning script for position reasoning

```
processor = AutoProcessor.from_pretrained(MODEL_ID, trust_remote_code=True)
model = Qwen2_5_VLForConditionalGeneration.from_pretrained(
    MODEL_ID, torch_dtype=torch.bfloat16, device_map="auto", trust_remote_code=True
)
model.config.use_cache = False
if hasattr(model, "enable_input_require_grads"):
    model.enable_input_require_grads()

model = get_peft_model(model, LoraConfig(
    r=16, lora_alpha=32, lora_dropout=0.05, bias="none",
    task_type="CAUSAL_LM",
    target_modules=["q_proj","k_proj","v_proj","o_proj",
                    "gate_proj","up_proj","down_proj","w1","w2","w3","W_pack"]
))

args = TrainingArguments(
    per_device_train_batch_size=1,
    gradient_accumulation_steps=8,
    num_train_epochs=2,
    learning_rate=2e-4,
```

```
21      warmup_ratio=0.03,
22      logging_steps=20,
23      bf16=True,
24      gradient_checkpointing=True,
25      dataloader_pin_memory=True,
26  )
```

D.2   COLOR ROBUSTNESS

Listing 2: Fine-tuning script for color robustness

```
1  model.config.use_cache = False
2  if hasattr(model, "enable_input_require_grads"):
3      model.enable_input_require_grads()
4
5  lora_config = LoraConfig(
6      r=16, lora_alpha=32, lora_dropout=0.05, bias="none",
7      task_type="CAUSAL_LM",
8      target_modules=["q_proj","k_proj","v_proj","o_proj",
9                      "gate_proj","up_proj","down_proj","w1","w2","w3","W_pack"]
10 )
11 model = get_peft_model(model, lora_config)
12
13 args = TrainingArguments(
14     per_device_train_batch_size=1,
15     gradient_accumulation_steps=8,
16     num_train_epochs=2,
17     learning_rate=2e-4,
18     warmup_ratio=0.03,
19     logging_steps=20,
20     save_strategy="epoch",
21     eval_strategy="epoch",
22     bf16=True,
23     gradient_checkpointing=True,
24 )
```

D.3   UI GROUP

Listing 3: Fine-tuning script for UI group classification

```
1  train_data = load_dataset("json", data_files=os.path.join(DATA_DIR,"uigroup_train.jsonl"))["
       train"]
2  eval_data  = load_dataset("json", data_files=os.path.join(DATA_DIR,"uigroup_eval.jsonl"))["
       train"]
3
4  lora_config = LoraConfig(
5      r=16, lora_alpha=32, lora_dropout=0.05, bias="none",
6      task_type="CAUSAL_LM",
7      target_modules=["q_proj","k_proj","v_proj","o_proj",
8                      "gate_proj","up_proj","down_proj","w1","w2","w3","W_pack"]
9  )
10 model = get_peft_model(model, lora_config)
11
12 training_args = TrainingArguments(
13     per_device_train_batch_size=4,
14     per_device_eval_batch_size=4,
15     gradient_accumulation_steps=4,
16     num_train_epochs=3,
17     learning_rate=2e-5,
18     warmup_ratio=0.03,
19     logging_steps=10,
20     evaluation_strategy="steps",
21     save_strategy="epoch",
22     bf16=True,
23     report_to="none",
24 )
```

E   ROBUSTNESS AND FAIL CASES

The following table 3 shows the ratio of unchanged output of the model before and after our inter-ference.

Table 3: Model performance on robustness tasks (Self compare).

Model	Color Robustness	Text Robustness	Layout Robustness
GPT-5	75.88	65.13	57.65
claude-4-Sonnet	70.54	54.14	62.47
Gemini 2.5-Pro	74.06	70.33	47.87
Pixtral-Large	52.86	48.27	62.69
InternVL3-78B	74.21	66.59	64.52
Qwen 2.5-VL-72B	74.39	69.83	75.18
Qwen2.5-VL-32B	64.69	71.24	58.63
Intern-S1	78.35	74.86	62.49
Pixtral-12B	61.24	72.86	71.81
Llama4-Scout17B	76.69	64.23	61.32
Qwen2.5-VL-7B	71.72	62.77	63.72
Qwen-vl-plus	75.12	69.47	70.78

The following cases 5, 6, 7 are also derived from our interference and the pre-and post output of the model.

a. Color perturbations fail case

Figure 5: Output examples of MLLMs on WebRSSBench when facing different perturbations.

F LLM USAGE

Large Language Models (LLMs) were used to integrate paragraph clarity and sentence grammar correction, while ensuring that sentences have high readability.

Figure 6: Output examples of MLLMs on WebRSSBench when facing different perturbations.

Figure 7: Output examples of MLLMs on WebRSSBench when facing different perturbations.