# WEBGEN-R1: INCENTIVIZING LLMS TO GENERATE FUNCTIONAL AND AESTHETIC WEBSITES WITH REINFORCEMENT LEARNING

**Anonymous authors**Paper under double-blind review

000

001

002

004

006

008 009 010

011 012

013

014

015

016

017

018

019

021

024

025

026

027

028

029

031

033

034

039

040

041

042

043

044

046

047

048

049

050

051

052

# **ABSTRACT**

Large Language Models (LLMs) have demonstrated strong capabilities in functional-level code generation, yet their performance remains limited in projectlevel scenarios such as generating large-scale multi-page websites. Such tasks require coherent multi-file structures, handling of intricate cross-page dependencies, and visually appealing designs. Prior works address only partial aspects of this challenge. For instance, WebDev Arena<sup>1</sup> focuses exclusively on single-page static sites, while agent-based frameworks decompose tasks into subtasks coordinated through multi-turn execution, often relying on proprietary models and suffering from fragile integration, particularly in visual coherence and stylistic consistency. In this work, we introduce WebGen-R1, pushing toward a more ambitious and practically relevant goal of training a small-scale LLM via reinforcement learning (RL) to generate the entire multi-page websites in an end-to-end manner. A key obstacle lies in reward design. Unlike functional code generation where correctness can be verified by passing automated test suites, web aesthetics covering layout harmony, typographic consistency, and stylistic alignment are inherently subjective, and functional verification often requires dynamic execution across pages where rule-based reward function tend to be brittle. To address these limitations, we design a vision-language-model-based reward model that jointly optimizes functional correctness and aesthetic quality, enabling the model to produce websites that are both visually coherent and faithful to the intended task specification. Extensive experiments across real-world benchmarks demonstrate that WebGen-R1 consistently outperforms, or is comparable to, strong proprietary and open-source baselines in a multi-dimensional evaluation protocol. To facilitate future research in end-to-end multi-page website generation, we release our code and data at https://anonymous.4open.science/r/WebGen-R1.

# 1 Introduction

Recent advances in Large Language Models (LLMs) have markedly expanded their capabilities in automated code generation (Jaech et al., 2024; Hui et al., 2024; Jiang et al., 2025; Guo et al., 2025; Team et al., 2025; Yang et al., 2025a; Zheng et al., 2025), achieving human competitive performance on established functional-level benchmarks such as the HumanEval (Chen et al., 2021) dataset or even complex International Olympiad in Informatics (IOI) programming tasks (Li et al., 2022). These achievements demonstrate impressive syntactic and semantic reasoning over standalone problems. However, moving from constrained, function-level snippets to project-level code generation that meets real-world software engineering demands remains a frontier challenge (Jimenez et al., 2023; Bi et al., 2024; Zan et al., 2025; Badertdinov et al., 2025). Among the various categories of such tasks, end-to-end website generation, encompassing multi-page routing, dynamic functionality, modern user interface (UI) design, and responsive layouts, presents especially rich and consequential opportunities for LLM research. Website generation represents a particularly challenging instantiation of project-level code generation (Wan et al., 2024; Xiao et al., 2024; Lu et al., 2025; Zhang et al., 2025). Unlike single-function problems, real-world websites demand consistent architectural patterns, multi-file codebases with intricate dependencies, long-range contextual coherence,

<sup>&</sup>lt;sup>1</sup>https://blog.lmarena.ai/blog/2025/webdev-arena

and conformity to design principles that balance functionality with visual appeal. This necessitates reasoning not only over software engineering constraints but also over aesthetic and user-experience considerations, which have traditionally been difficult to formalize and evaluate in automated code generation.

Despite promising early steps, current approaches to LLM-driven website generation exhibit notable limitations. One line of work has chosen to simplify the generation task drastically. For example, WebDev Arena (LMArena, 2025) is constrained to generating single-page static sites. While this makes evaluation tractable, it abstracts away essential complexities such as dynamic routing, state management, authentication flows, and cross-page navigation. Another line of research adopts multi-agent orchestration frameworks, in which different specialized LLMs handle discrete subtasks (UI layout, backend logic, testing), and their outputs are subsequently integrated (Hong et al., 2023; He et al., 2024a;b; Lu et al., 2025). However, such modularity introduces brittle inter-agent dependency chains, where small inconsistencies in contracts, file names, or interface definitions can cascade into non-functional builds. Moreover, both paradigms rarely incorporate formal optimization of aesthetic quality or human-aligned design sensibilities, an omission that leads to websites which, while functional, often fail to meet the expectations of end-users in visual polish. Further discussion of related work is provided in Appendix D.

In this work, we push toward a more ambitious and practically relevant goal of *training a small-scale LLM to generate an entire multi-page, functional, and visually aesthetic website project from scratch, in an end-to-end manner, without external decomposition into subtasks.* This departure invites several formidable challenges: ① Global structural reasoning over project-level architectures, including framework-specific conventions (e.g., Next.js routing, Vue plugin registration) and modular directory organization. ② Maintaining multi-file consistency for cross-referencing components, dependencies, and dynamic import paths across the codebase. ③ Ensuring cohesive visual design and modern aesthetics, beyond merely placing elements on a page, by rendering layouts that exhibit balance, alignment, accessible color contrasts, and brand coherence. ④ Capturing complex interactive behaviors such as animations, drag-and-drop, and responsive state updates within the generated code. ⑤ Overcoming limited long-context reasoning, as project code often exceeds the context window available in current LLMs, creating difficulties in tracking dependencies over hundreds of lines across multiple files.

To address these challenges, we propose WebGen-R1, a novel framework that integrates reinforcement learning (RL) directly into the end-to-end website generation process. Nevertheless, a central obstacle in bringing RL into such open-ended generative tasks lies in the design of a reliable reward signal (Guo et al., 2025; Zeng et al., 2025; Wen et al., 2025; Yang et al., 2025b; Mroueh, 2025). Different from tasks like complex mathematics and competitive programming, where correctness is objectively verifiable by exactly matching unambiguous ground-truth answers or passing automated test suites, it faces a serious challenges with website generation: (1) Not all desired qualities are reducible to scalar outcomes, as visual appeal, design cohesion, and user experience have no trivial Boolean test. (2) Comprehensive functional verification for complex, multi-page websites often requires running the site and inspecting behaviors in varied scenarios, making static rule definition brittle. (3) The scarcity of annotated, fully verifiable website outputs limits scalability for this domain.

To overcome these limitations, we replace handcrafted rule-based scorers with a reward model that incorporates both task specification cognition and visual rendering perception. Concretely, after the model generates the complete web project code and directory structure, we execute a standardized front-end development pipeline, which includes parsing and verifying the scaffolded file organization, installing dependencies, building the project, launching a local development or production server, and rendering the pages in a browser. The rendered page screenshots, together with the original task specification, are then fed into a state-of-the-art vision-language model (VLM) for joint functional and aesthetic evaluation, producing a graded scalar reward. This pipeline preserves verifiability for functional aspects while enabling nuanced assessment of open-ended design quality, making large-scale RL training both stable and meaningful.

We conduct extensive experiments on real-world benchmarks to evaluate the effectiveness of WebGen-R1. When employing Qwen2.5-Coder-Instruct-7B as our base model, WebGen-R1 exhibits substantial gains in multi-page website generation, achieving an improvement in functional quality metrics from 1.59% to 29.21%, a 44.32% increase in aesthetic scoring, and a drastic increase

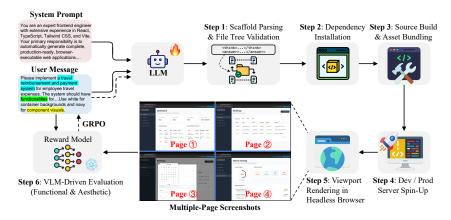


Figure 1: The architecture of our proposed WebGen-R1 for end-to-end multi-page website generation using a single LLM. A user provides a natural-language website design request, which the LLM generates the entire project source code. The generated scaffold is then processed through a standardized front-end build pipeline, comprising parsing and validation of the file organization, dependency installation, project build, server spin-up (in development or production mode), and rendering of the resulting webpages in a headless browser environment to obtain multiple page screenshots. These rendered page images, along with the original user request, are subsequently fed into a state-of-the-art VLM for joint evaluation of functional correctness and visual design quality, producing a scalar reward signal in the discrete range 0-5. This automated evaluation framework ensures both objective verification of executable functionality and nuanced assessment of open-ended aesthetic attributes. The reward signal is then used to fine-tune the LLM via the GRPO objective, enabling stable and semantically meaningful reinforcement learning at scale.

in valid render ratio from 30.56% to 95.89%. Importantly, despite having significantly fewer parameters, WebGen-R1 attains aesthetic quality and valid render ratios that surpass those of much larger and more powerful baselines such as Gemini-2.5-Pro and DeepSeek-R1, while maintaining comparable performance in functional evaluation. These results provide strong empirical evidence that our WebGen-R1 with reinforcement learning can effectively enhance both the functional reliability and visual appeal of generated websites. In summary, our contributions are as follows:

- To our best knowledge, we are the first to introduce reinforcement learning, namely WebGen-R1, for end-to-end generating the entire multi-page websites without relying on task decomposition or proprietary model orchestration.
- We design a vision-language-model-based reward model that jointly measures functional
  correctness and aesthetic quality, capturing layout harmony, typographic consistency, and
  stylistic alignment. This resolves the brittleness of rule-based rewards in project-level code
  generation and aligns optimization signals with human preferences.
- We establish a multi-dimensional evaluation protocol covering functional correctness, visual coherence, deployability, and human-perceived quality, and demonstrate through extensive experiments on real-world benchmarks that our WebGen-R1 outperforms or matches advanced proprietary and open-source baselines, even in challenging scenarios requiring fine-grained inter-page coordination.
- We release our codebase, datasets, RL training framework, and trained model checkpoints to enable reproducibility and to foster future research on end-to-end, LLM-driven multi-page website generation.

#### 2 METHODOLOGY

#### 2.1 END-TO-END OPEN-ENDED WEBSITE CODE GENERATION

Given a task specification  $x \in \mathcal{X}$ , the policy  $\pi_{\theta}$  must produce an entire project  $y \in \mathcal{Y}$  containing a coherent multi-file structure, framework-specific directories, and interdependent code modules. Unlike template-based generation, the space  $\mathcal{Y}$  is open-ended  $\mathcal{Y} = \bigcup_{K=1}^{\infty} \mathcal{Y}_K$ ,  $\mathcal{Y}_K = \{y : y \in \mathcal{Y}_K : y \in \mathcal{Y}_K$ 

163

164

165

166

167

168

169

170 171

172

173

174

175

176

177

178

179

181

182

183

184 185

186 187

188

189

190

191

192

193

195 196

197

199 200

201

202

203

204

205

206

207

208

209

210

211

212 213

214

215

y contains K valid files in a consistent project graph. We model generation as an autoregressive process  $\pi_{\theta}(y \mid x) = \prod_{t=1}^{T} \pi_{\theta}(y_t \mid y_{< t}, x)$ , where the token sequence  $y_{1:T}$  is subsequently parsed into a file set  $\mathcal{F}(y) = \{f_k\}_{k=1}^K$  and a directory graph  $\mathcal{S}(y)$ . We denote the resulting website project instance as  $Web(y) = \{\mathcal{F}(y), \mathcal{S}(y)\}$ . By emitting all code in one globally conditioned sequence, the policy is forced to maintain cross-file variable references, routing conventions, and naming coherence. These properties are frequently broken in iterative, file-by-file generation, and their absence leads to non-functional builds. This holistic treatment is particularly advantageous for web projects, where front-end routing, shared styles, and state management are tightly coupled.

#### 2.2 AUTOMATED FRONT-END BUILD AND RENDERING PIPELINE

A generated website codebase carries no operational meaning until it is executed within an actual front-end development pipeline. We model this execution as a deterministic environment transformation:

$$o = \mathcal{E}(y) = \mathcal{R}(\mathcal{L}(\mathcal{B}(\mathcal{I}(\mathcal{W}eb(y))))), \tag{1}$$

where  $\mathcal{I}$  installs dependencies,  $\mathcal{B}$  builds the compiled bundle,  $\mathcal{L}$  launches a server, and  $\mathcal{R}$  renders pages through a headless browser. As a result, the output o contains:

$$o = \left( \left\{ I_p \right\}_{p=1}^P, \Gamma \right), \tag{2}$$

with  $\{I_p\}_{p=1}^P$  denoting rendered screenshots for P routes and  $\Gamma$  collecting build and runtime logs. This execution step enforces realizability constraints so that only codebases that install, build, and run successfully yield complete visual evidence, directly grounding learning signals in executable behavior. In website generation, this is critical as projects that pass linters can still fail at runtime due to subtle integration errors, broken imports, or misconfigured frameworks.

#### 2.3 REWARD MODEL FOR FUNCTIONAL AND AESTHETIC EVALUATION

While functional correctness covering aspects such as hyperlink validity, navigation integrity, and component responsiveness can be partially verified via deterministic, rule-based checks, purely algorithmic verification fails to capture subjective, perceptual qualities of web design. Aesthetic factors, including layout harmony, typographic consistency, and stylistic coherence with the task specification, require perceptual evaluation signals that are inherently non-deterministic. To jointly capture both dimensions, we propose a multimodal reward evaluator  $\phi_{\psi}$  built upon a state-of-the-art vision-language model (VLM). Formally, given an input tuple  $z = (x, \{I_p\}_{p=1}^P, \Gamma)$ , where x denotes the natural language specification,  $\{I_p\}_{p=1}^P$  are rendered page snapshots from multiple responsive breakpoints, and  $\Gamma$  represents runtime execution logs, the evaluator jointly attends to textual, visual, and behavioral modalities to produce a scalar score:

$$s_{\langle \text{func,vis} \rangle} = g_{\phi_{\psi}}[s_{\text{func}}, s_{\text{vis}}] = \phi_{\psi}^{\text{VLM}}(z),$$
 (3)

 $s_{\langle \text{func}, \text{vis} \rangle} = g_{\phi_{\psi}}[s_{\text{func}}, \ s_{\text{vis}}] = \phi_{\psi}^{\text{VLM}}(z), \tag{3}$  where  $s_{\text{func}}$  measures functional soundness and  $s_{\text{vis}}$  evaluates visual design quality. The unified score  $s_{(\text{func,vis})} \in [0,5]$  thus reflects both executable fidelity and human-perceived aesthetics in a single quantitative metric.

Beyond these two primary criteria, we incorporate additional signals to enforce structural and reasoning quality in generated code. Inspired by (Guo et al., 2025), we assess code format correctness  $s_{\text{(code)}}$ , where the generated website code is parsed into a file set  $\mathcal{F}(y) = \{f_k\}_{k=1}^K$  and a corresponding directory graph S(y). This check ensures that the output can be executed without modification, providing a direct, verifiable reward signal for structural validity. Moreover, to incentivize long-horizon and explicit reasoning in project organization, such as planning directory hierarchies, configuring frameworks appropriately, and maintaining coherent shared state, we define a reasoning format reward  $s_{(cot)}$ . Here, the model is required to externalize its reasoning process between dedicated <think>...</think>tags, enabling downstream parsing and reward estimation on reasoning quality. Finally, these components are fused into a multi-dimensional scalar reward:

$$R(y) = s_{\langle \text{func,vis} \rangle} + \gamma \, s_{\langle \text{code} \rangle} + \lambda \, s_{\langle \text{cot} \rangle}, \tag{4}$$

with weight coefficients  $\gamma, \lambda \in [0,1]$  controlling the contribution of code format and reasoning quality to the overall signal. Thus, in the context of LLM-based website generation, this integrative scoring mitigates the common pitfall of over-optimizing for mechanically verifiable checks at the expense of human-perceived functionality and visual appeal, thereby aligning model outputs more closely with end-user expectations.

#### 2.4 Group Relative Policy Optimization for Website Generation

We formulate our reinforcement learning objective as maximizing the expected task-specific reward under the learned policy  $\pi_{\theta} \max_{\theta} \mathbb{E}_{x \sim \mathcal{D}, \ y \sim \pi_{\theta}(\cdot|x)}[R(y)]$ , where x denotes a website generation prompt, y is a candidate structured website output, and  $R(\cdot)$  measures a composite reward capturing both functional correctness and visual quality (see Eq. 4). Unlike standard PPO(Schulman et al., 2017), GRPO (Shao et al., 2024; Guo et al., 2025; Yu et al., 2025) removes the need for an explicit value function by normalizing rewards within a group of sampled responses for the same prompt, enabling more stable optimization and mitigating inter-prompt variance in difficulty. Specifically, for each website specification x, we sample a group of G candidate responses  $\{y_i\}_{i=1}^G$  from the behavior policy  $\pi_{\theta_{\text{old}}}(\cdot \mid x)$ , where each sequence is tokenized as  $y_i = (y_{i,1}, y_{i,2}, \dots, y_{i,|y_i|})$ . We obtain a scalar reward  $R_i$  for each  $y_i$  by executing the generated website and evaluating functional validation and visual fidelity score. The group-relative normalized advantage for token position t in response i is then defined as:

$$\hat{A}_{i,t} \triangleq \frac{R_i - \text{mean}(\{R_j\}_{j=1}^G)}{\text{std}(\{R_j\}_{j=1}^G)},\tag{5}$$

Following the PPO-style clipped surrogate objective with an explicit KL penalty to constrain divergence from a frozen reference policy  $\pi_{ref}$ , the GRPO objective is:

$$\mathcal{J}_{GRPO}(\theta) = \mathbb{E}_{x \sim \mathcal{D}, \{y_i\}_{i=1}^G \sim \pi_{\theta_{\text{old}}}(\cdot | x)} \left[ \frac{1}{G} \sum_{i=1}^G \frac{1}{|y_i|} \sum_{t=1}^{|y_i|} \right] \\
\left( \min \left[ r_{i,t}(\theta) \, \hat{A}_{i,t}, \operatorname{clip}(r_{i,t}(\theta), 1 - \varepsilon, 1 + \varepsilon) \, \hat{A}_{i,t} \right] - \beta \, D_{\text{KL}} \left( \pi_{\theta}(\cdot \mid x) \parallel \pi_{\text{ref}}(\cdot \mid x) \right) \right) \right], \\
r_{i,t}(\theta) \triangleq \frac{\pi_{\theta} \left( y_{i,t} \mid x, y_{i, < t} \right)}{\pi_{\theta_{\text{old}}} \left( y_{i,t} \mid x, y_{i, < t} \right)}, \tag{7}$$

where  $\varepsilon > 0$  is the clipping parameter, and  $\beta$  controls the strength of the KL regularization. The detailed procedure of the algorithm is presented in Algorithm 1 (see Appendix A).

# 3 EXPERIMENTS

#### 3.1 EXPERIMENTAL SETUP

Datasets and Benchmarks. We leverage the WebGen-Instruct (Lu et al., 2025) as our training corpus. This dataset comprises 6,667 end-to-end website generation tasks spanning a broad spectrum of real-world web application domains. We employ the WebGen-Bench (Lu et al., 2025), which contains 101 carefully curated website generation tasks. These tasks range from minimal single-page designs to complex corporate websites with rich interactivity and data-driven dashboards. Each benchmark instance is paired with a comprehensive, repeatedly validated test suite, ensuring reliable measurement of both the functional behavior and stylistic conformance of generated websites. The natural-language task descriptions explicitly specify functional requirements and visual design expectations, enabling precise evaluation. The detailed statistics of the dataset and the benchmark are summarized in Table 2 (Appendix B). Figure 2 illustrates the distributions of prompt and response lengths, where each prompt is obtained by concatenating the system prompt with the web task query, for several state-of-the-art commercial LLMs on the end-to-end website generation task.

**Baselines.** We benchmark WebGen-R1 against a broad spectrum of state-of-the-art LLMs, encompassing eight proprietary models and seven open-source models. We exclude WebGen-LM (Lu et al., 2025) from comparison as it is specifically fine-tuned on Bolt.diy <sup>2</sup> website-generation trajectories gathered via DeepSeek-V3 (Liu et al., 2024), making it bound to a particular agent-based framework and not directly applicable in our evaluation setting.

**Metrics.** To thoroughly evaluate LLMs for end-to-end website generation, we use several quantitative metrics, including (1) Functional Success Rate (FSR): the percentage of generated websites

<sup>&</sup>lt;sup>2</sup>https://github.com/stackblitz-labs/bolt.diy

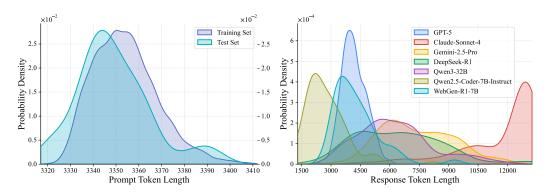


Figure 2: Token length distributions of prompts and generated responses for several state-of-the-art LLMs on the end-to-end multi-page website generation task across WebGen-Instruct (training set) and WebGen-Bench (test set). The response length distributions are aggregated over test set.

Table 1: Performance comparison of WebGen-R1 and various state-of-the-art LLMs from multiple institutions on the WebGen-Bench benchmark, evaluated by FSR, AAS, and VRR metrics. We show the score improvement  $(\pm)$  of our model relative to its base. Bold values indicate the best results.

Institution	Model	FSR(%)	AAS	VRR(%)
	GPT-5	46.53	3.34	90.43
	GPT-4.1	43.91	6.53 3.34 3.91 3.78 2.86 3.55 7.29 3.31 11.60 3.31 6.13 3.86 7.72 3.90 6.31 3.89 0.25 3.67 1.59 2.73 2.54 3.14 3.72 2.57 8.69 3.39 9.30 2.60	82.09
OpenAI	03	42.86	3.55	81.08
-	o4-mini	27.29	3.31	56.52
	GPT-4o 21.60		3.31	85.71
A) A41	Claude-Sonnet-4	46.13	3.86	86.05
A\ Anthropic	Claude-3.7-Sonnet		3.90	84.00
<b>G</b> Google	Gemini-2.5-Pro	36.31	3.89	83.33
♥ DeepSeek	DeepSeek-R1	30.25	3.67	42.86
	Qwen2.5-Coder-7B-Instruct	1.59	2.73	30.56
_	Qwen2.5-72B-Instruct	2.54	3.14	8.86
Alibaba	Qwen3-8B	3.72	2.57	12.50
	Qwen3-32B	42.86   27.29   21.60   46.13   <b>57.72</b>   36.31   30.25   1.59   2.54   3.72   18.69   9.30	3.39	59.42
	Qwen3-30B-A3B-Thinking-2507	9.30	2.60	23.81
	Qwen3-Coder-30B-A3B-Instruct	6.06	2.90	32.81
Ours	WebGen-R1-7B	29.21+27.62	<b>3.94</b> +44.32	<b>95.89</b> +65.33

that pass predefined interactive checks such as button clicks and form submissions; (2) Aesthetic Alignment Score (AAS): the average reward model score, measuring how well function and design match human aesthetic preferences; (3) Valid Render Ratio (VRR): the percentage of generated websites that render without major errors; (4) Lint & Dependency Pass Rate (LDPR): the fraction of projects that pass static code analysis (like ESLint) and resolve dependencies automatically, indicating readiness for deployment. These metrics form a multi-dimensional evaluation protocol that jointly accounts for execution correctness, visual attractiveness, deployability, and human-perceived quality, providing a holistic assessment of LLM-driven website generation.

Implementation Details. We conduct all experiments on a cluster equipped with  $64 \times \text{NVIDIA}$  H100 GPUs (80 GB memory each), using the open-source Open-R1 framework <sup>3</sup> (Hugging Face, 2025). We fine-tune the Qwen2.5-Coder-Instruct-7B model for 400 optimization steps using the GRPO objective. Following (Lu et al., 2025), the VLM reward model leverages a state-of-the-art GPT-4o evaluator to score both the functional correctness and visual aesthetics of generated websites. The reward function is calibrated to align closely with human preferences through prompt engineering and normalization, ensuring stable optimization. More details, including the hyperparameter settings, are provided in Appendix C.

<sup>&</sup>lt;sup>3</sup>https://github.com/huggingface/open-r1

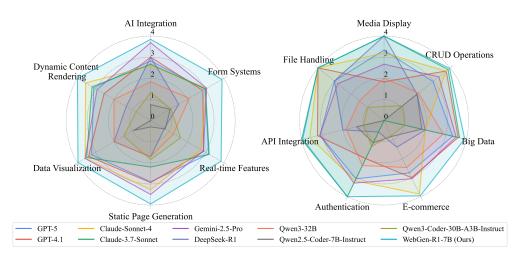


Figure 3: Performance comparison of WebGen-R1 and various state-of-the-art LLMs on 13 heterogeneous multi-scenario front-end development tasks encompasses diverse requirements, providing a rigorous testbed for both functionality and visual fidelity, sourced from the WebGen-Bench.

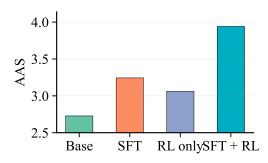
#### 3.2 Main Results

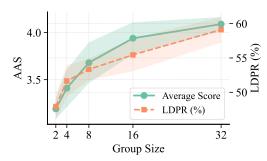
We compare WebGen-R1 with state-of-the-art LLMs on the WebGen-Bench using FSR, AAS, and VRR to examine whether reinforcement learning can optimize functionality, design, and render reliability. As shown in Table 1, Claude-3.7-Sonnet achieves the highest FSR (57.72%), showing strong execution, while WebGen-R1 gets the highest AAS score (3.94), outperforming all models, including Claude. WebGen-R1 also attains the best VRR (95.89%), clearly higher than all baselines, which we attribute to our reward design focusing on code format correctness. This balance of visual quality and reliable rendering is achieved even though WebGen-R1 is a 7B model, proving the effectiveness of our RL training. Smaller Qwen3 models generally have lower FSR, while some larger variants like Qwen3-32B (AAS 3.39) and Qwen2.5-72B-Instruct (AAS 3.14) reach relatively high AAS. *This pattern suggests that generating functionally correct website is substantially more challenging than achieving visual appeal, as functional requirements in front-end development often involve complex logic and interactions beyond visual layout.* Overall, our RL optimization enables the model to produce visually attractive websites while still maintaining functional correctness, demonstrating that targeted reward shaping can balance these multiple objectives in website generation.

# 3.3 IN-DEPTH ANALYSIS AND INSIGHTS.

Multi-Scenario Web Environments. We evaluate WebGen-R1 and a variety of state-of-the-art LLMs on multi-scenario front-end development tasks, including AI Integration, Form Systems, Real-time Features, Static Page Generation, Data Visualization, Dynamic Content Rendering, Media Display, CRUD Operations, Big Data, E-commerce, Authentication, API Integration, and File Handling, sourced from WebGen-Bench (Lu et al., 2025). As shown in Figure 3, our WebGen-R1 achieves superior performance across all 13 categories on AAS metric, which indicates a consistent improvements in both functional correctness and UI/UX quality. Such uniform improvements across diverse front-end scenarios demonstrate that our WebGen-R1 is task-agnostic, robust to domain shifts, and effective in harmonizing execution-level correctness with aesthetic alignment, highlighting the practical applicability to real-world, mixed-requirement web development.

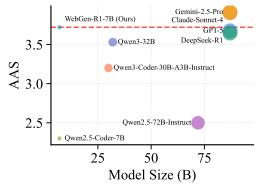
**RL Fine-Tuning.** We compare RL and supervised fine-tuning (SFT) on the web generation task. Using 600 GPT-4.1-generated examples for SFT, we observe an 18.68% performance boost over the baseline model (Qwen2.5-Coder-7B-Instruct), significantly improving functional and coherent webpage generation, as shown in Figure 4. RL-only (R1-Zero) gives a 12.09% improvement, showing the value of vision–language model rewards for optimizing both function and appearance. SFT brings higher gains than RL alone, but combining SFT initialization with RL achieves even better results, with performance 21.60% and 28.76% higher than SFT and RL-only respectively. These findings indicate that SFT enables the model to acquire a robust structural and semantic prior for





mark under the AAS metric.

Figure 4: Performance comparison of SFT, RL- Figure 5: Impact of group size G in GRPO on only, and SFT+RL on the WebGen-Bench bench- WebGen-Bench performance measured by AAS and LDPR.



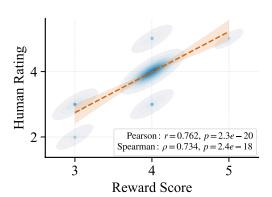


Figure 6: Performance of WebGen-R1 on the WebDev Arena benchmark, which differs in do- Figure 7: Alignment between reward model evalmain and prompt distribution from the WebGen- uations and human ratings on WebGen-Bench Instruct and WebGen-Bench.

websites with strong correlations.

webpage generation, while RL expands the model's exploration capabilities and reward-driven optimization, allowing it to discover and produce higher-quality, more aesthetically pleasing websites.

**Group Size in GRPO.** We investigate the effect of the group size parameter G in GRPO, with  $G \in \{2, 4, 8, 16, 32\}$ , while keeping all other hyperparameters fixed. For each setting, we track the evolution of both the mean and standard deviation of the reward per training step, and assess the final policy on the WebGen-Bench using the AAS and LDPR metrics. As shown in Figure 5, larger group sizes consistently yield superior performance across both evaluation metrics. We hypothesize that this improvement arises from the enhanced exploration capacity provided by larger groups, which increases the diversity of candidate trajectories and improves the likelihood of discovering high-quality website designs.

**Generalization and Robustness.** We evaluate the ability of WebGen-R1 to generalize its reasoning processes and visual design sensibility to settings where both the domain and prompt distributions differ substantially from those seen during RL, which was conducted on the WebGen-Instruct. To this end, we adopt the WebDev Arena benchmark, which features instruction distributions and task categories that are not covered in our training set. Detailed dataset statistics are provided in Table 2 in the Appendix. The quantitative results are presented in Figure 6. WebGen-R1 consistently outperforms a range of state-of-the-art proprietary and open-source baselines (e.g., DeepSeek-R1, GPT-5, and Qwen3-32B) across AAS metric. This performance suggests that WebGen-R1 has learned architecture-level and style-level abstractions that remain effective in previously unseen web domains. These findings further indicate that WebGen-R1 maintains strong applicability and robustness in practical deployment scenarios where specification distributions may vary over time.

Human Alignment Study. We examine whether our vision-language-model-based reward accurately reflects human preferences for functionality and aesthetics. Since RL relies entirely on this reward, any misalignment could lead to outputs that do not meet user expectations. To evaluate this, we compare our reward model to human judgments by having three experienced front-end devel-

**[WebGen-Bench] Instruction:** Please implement a Q&A website that answers user-submitted questions. The website should have functionalities for submitting questions, answering questions, and viewing answers. Users should be able to submit questions, view answers, and rate the answers. Set overall background to beige, then style components with saddle brown.



GPT-5 Gemini-2.5-Pro WebGen-R1-7B (Ours)

**[WebDev Arena] Instruction:** Make a website that fetches data (joke) from an external API and displays it on the screen using react for use example.



Figure 8: Case study of our WebGen-R1-7B against two leading baselines on in-distribution (WebGen-Bench) and out-of-distribution (WebDev Arena) tasks, showing superior layout organization, visual coherence, and functional fidelity.

opers rate 101 websites from WebGen-Bench on functionality and visual appeal. We aggregate the human scores and compare them to the reward model's outputs. As shown in Figure 7, the results show strong correlations (Pearson r=0.762, Spearman  $\rho=0.734$ ), indicating that the model's ratings closely match human evaluations. This demonstrates that our vision–language model can reliably assess both functional fidelity and aesthetic appeal in generated websites.

#### 3.4 CASE STUDIES

To qualitatively evaluate WebGen-R1's ability to improve website functionality and visual quality, we conduct case studies using two user instructions from in-distribution (WebGen-Bench) and out-of-distribution (WebDev Arena) sources. Figure 8 shows that, compared to strong baseline models, WebGen-R1 generates websites with organized layouts, coherent and attractive designs, and responsive behaviors that match the instructions. These findings support our quantitative results, showing that WebGen-R1 achieves both functional correctness and aesthetics preferred by humans. We attribute these strengths to end-to-end code generation, multimodal reward modeling, and GRPO optimization. Overall, the results show that LLMs trained with perceptually grounded RL can effectively balance engineering requirements and design principles in web development.

# 4 CONCLUSION

In this work, we introduce WebGen-R1, a reinforcement learning framework that enables small-scale LLMs to generate the entire multi-page websites in an end-to-end manner while meeting both functional and aesthetic requirements. Extensive experiments across two real-world benchmarks demonstrate that our WebGen-R1 consistently improves functional robustness, visual coherence, and deployability, surpassing or matching to advanced proprietary and open-source baselines. Notably, our work demonstrates that RL with perceptually grounded rewards can substantially advance structured, multi-modal generation tasks beyond the functional level. We believe these insights open new directions for end-to-end LLM training in full-stack application development, and we release all resources to catalyze future research in this emerging area.

# 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.
- Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan, Ellen Jiang, Carrie Cai, Michael Terry, Quoc Le, et al. Program synthesis with large language models. *arXiv preprint arXiv:2108.07732*, 2021.
- Ibragim Badertdinov, Alexander Golubev, Maksim Nekrashevich, Anton Shevtsov, Simon Karasik, Andrei Andriushchenko, Maria Trofimova, Daria Litvintseva, and Boris Yangel. Swe-rebench: An automated pipeline for task collection and decontaminated evaluation of software engineering agents. *arXiv preprint arXiv:2505.20411*, 2025.
- Yuntao Bai, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones, Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, et al. Constitutional ai: Harmlessness from ai feedback. *arXiv preprint arXiv:2212.08073*, 2022.
- Zhangqian Bi, Yao Wan, Zheng Wang, Hongyu Zhang, Batu Guan, Fangxin Lu, Zili Zhang, Yulei Sui, Hai Jin, and Xuanhua Shi. Iterative refinement of project-level code context for precise code generation with compiler feedback. In *Findings of the Association for Computational Linguistics ACL 2024*, pp. 2336–2353, 2024.
- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*, 2021.
- Xinyun Chen, Maxwell Lin, Nathanael Schärli, and Denny Zhou. Teaching large language models to self-debug. *arXiv preprint arXiv:2304.05128*, 2023.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. Scaling instruction-finetuned language models. *Journal of Machine Learning Research*, 25(70):1–53, 2024.
- Shihan Dou, Yan Liu, Haoxiang Jia, Limao Xiong, Enyu Zhou, Wei Shen, Junjie Shan, Caishuang Huang, Xiao Wang, Xiaoran Fan, et al. Stepcoder: Improve code generation with reinforcement learning from compiler feedback. *arXiv preprint arXiv:2402.01391*, 2024.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025.
- Hongliang He, Wenlin Yao, Kaixin Ma, Wenhao Yu, Yong Dai, Hongming Zhang, Zhenzhong Lan, and Dong Yu. Webvoyager: Building an end-to-end web agent with large multimodal models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics* (Volume 1: Long Papers), pp. 6864–6890, 2024a.
- Hongliang He, Wenlin Yao, Kaixin Ma, Wenhao Yu, Hongming Zhang, Tianqing Fang, Zhenzhong Lan, and Dong Yu. Openwebvoyager: Building multimodal web agents via iterative real-world exploration, feedback and optimization. *arXiv preprint arXiv:2410.19609*, 2024b.
- Sirui Hong, Xiawu Zheng, Jonathan Chen, Yuheng Cheng, Jinlin Wang, Ceyao Zhang, Zili Wang, Steven Ka Shing Yau, Zijuan Lin, Liyang Zhou, et al. Metagpt: Meta programming for multiagent collaborative framework. *arXiv preprint arXiv:2308.00352*, 2023.
- Hugging Face. Open r1: A fully open reproduction of deepseek-r1, January 2025. URL https://github.com/huggingface/open-r1.
- Binyuan Hui, Jian Yang, Zeyu Cui, Jiaxi Yang, Dayiheng Liu, Lei Zhang, Tianyu Liu, Jiajun Zhang, Bowen Yu, Keming Lu, et al. Qwen2. 5-coder technical report. *arXiv preprint arXiv:2409.12186*, 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.
  - Juyong Jiang, Fan Wang, Jiasi Shen, Sungju Kim, and Sunghun Kim. A survey on large language models for code generation. *ACM Trans. Softw. Eng. Methodol.*, July 2025. ISSN 1049-331X. doi: 10.1145/3747588. URL https://doi.org/10.1145/3747588. Just Accepted.
  - Carlos E Jimenez, John Yang, Alexander Wettig, Shunyu Yao, Kexin Pei, Ofir Press, and Karthik Narasimhan. Swe-bench: Can language models resolve real-world github issues? *arXiv preprint arXiv:2310.06770*, 2023.
  - Hung Le, Yue Wang, Akhilesh Deepak Gotmare, Silvio Savarese, and Steven Chu Hong Hoi. Coderl: Mastering code generation through pretrained models and deep reinforcement learning. *Advances in Neural Information Processing Systems*, 35:21314–21328, 2022.
  - Harrison Lee, Samrat Phatale, Hassan Mansoor, Kellie Lu, Thomas Mesnard, Colton Bishop, Victor Carbune, and Abhinav Rastogi. Rlaif: Scaling reinforcement learning from human feedback with ai feedback. *arXiv preprint arXiv:2309.00267*, 2023.
  - Jierui Li, Hung Le, Yingbo Zhou, Caiming Xiong, Silvio Savarese, and Doyen Sahoo. Codetree: Agent-guided tree search for code generation with large language models. *arXiv preprint arXiv:2411.04329*, 2024.
  - Yujia Li, David Choi, Junyoung Chung, Nate Kushman, Julian Schrittwieser, Rémi Leblond, Tom Eccles, James Keeling, Felix Gimeno, Agustin Dal Lago, et al. Competition-level code generation with alphacode. *Science*, 378(6624):1092–1097, 2022.
  - Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, et al. Deepseek-v3 technical report. *arXiv preprint arXiv:2412.19437*, 2024.
  - LMArena. Webdev arena: Ai battle to build the best website. https://web.lmarena.ai/, 2025.
  - Zimu Lu, Yunqiao Yang, Houxing Ren, Haotian Hou, Han Xiao, Ke Wang, Weikang Shi, Aojun Zhou, Mingjie Zhan, and Hongsheng Li. Webgen-bench: Evaluating Ilms on generating interactive and functional websites from scratch. *arXiv preprint arXiv:2505.03733*, 2025.
  - Michael Luo, Naman Jain, Jaskirat Singh, Sijun Tan, Ameen Patel, Qingyang Wu, Alpay Ariyak, Colin Cai, Shang Zhu Tarun Venkat, Ben Athiwaratkun, Manan Roongta, Ce Zhang, Li Erran Li, Raluca Ada Popa, Koushik Sen, and Ion Stoica. Deepswe: Training a state-of-the-art coding agent from scratch by scaling rl. https://pretty-radio-b75.notion.site/DeepS WE-Training-a-Fully-Open-sourced-State-of-the-Art-Coding-Agent-by-Scaling-RL-22281902c1468193aabbe9a8c59bbe33, 2025a. Notion Blog.
  - Michael Luo, Sijun Tan, Roy Huang, Ameen Patel, Alpay Ariyak, Qingyang Wu, Xiaoxiang Shi, Rachel Xin, Colin Cai, Maurice Weber, Ce Zhang, Li Erran Li, Raluca Ada Popa, and Ion Stoica. Deepcoder: A fully open-source 14b coder at o3-mini level. https://pretty-radio-b75.notion.site/DeepCoder-A-Fully-Open-Source-14B-Coder-at-O3-mini-Level-1cf81902c14680b3bee5eb349a512a51, 2025b. Notion Blog.
  - Youssef Mroueh. Reinforcement learning with verifiable rewards: Grpo's effective loss, dynamics, and success amplification. *arXiv* preprint arXiv:2503.06639, 2025.
  - Theo X Olausson, Jeevana Priya Inala, Chenglong Wang, Jianfeng Gao, and Armando Solar-Lezama. Is self-repair a silver bullet for code generation? In *The Twelfth International Conference on Learning Representations*, 2023.
  - Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35: 27730–27744, 2022.

- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
  - Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Yang Wu, et al. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024.
  - Bo Shen, Jiaxin Zhang, Taihong Chen, Daoguang Zan, Bing Geng, An Fu, Muhan Zeng, Ailun Yu, Jichuan Ji, Jingyang Zhao, et al. Pangu-coder2: Boosting large language models for code with ranking feedback. *arXiv preprint arXiv:2307.14936*, 2023.
  - Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. Reflexion: Language agents with verbal reinforcement learning. *Advances in Neural Information Processing Systems*, 36, 2024.
  - Parshin Shojaee, Aneesh Jain, Sindhu Tipirneni, and Chandan K Reddy. Execution-based code generation using deep reinforcement learning. *Transactions on Machine Learning Research*, 2023.
  - Disha Shrivastava, Hugo Larochelle, and Daniel Tarlow. Repository-level prompt generation for large language models of code. In *International Conference on Machine Learning*, pp. 31693–31715. PMLR, 2023.
  - Kimi Team, Angang Du, Bofei Gao, Bowei Xing, Changjiu Jiang, Cheng Chen, Cheng Li, Chenjun Xiao, Chenzhuang Du, Chonghua Liao, et al. Kimi k1. 5: Scaling reinforcement learning with llms. *arXiv preprint arXiv:2501.12599*, 2025.
  - 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.
  - Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A Smith, Daniel Khashabi, and Hannaneh Hajishirzi. Self-instruct: Aligning language models with self-generated instructions. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 13484–13508, 2023.
  - Xumeng Wen, Zihan Liu, Shun Zheng, Zhijian Xu, Shengyu Ye, Zhirong Wu, Xiao Liang, Yang Wang, Junjie Li, Ziming Miao, et al. Reinforcement learning with verifiable rewards implicitly incentivizes correct reasoning in base llms. *arXiv preprint arXiv:2506.14245*, 2025.
  - Jingyu Xiao, Yuxuan Wan, Yintong Huo, Zhiyao Xu, and Michael R Lyu. Interaction2code: How far are we from automatic interactive webpage generation? *arXiv e-prints*, pp. arXiv–2411, 2024.
  - 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*, 2025a.
  - Zhicheng Yang, Zhijiang Guo, Yinya Huang, Yongxin Wang, Dongchun Xie, Yiwei Wang, Xiaodan Liang, and Jing Tang. Depth-breadth synergy in rlvr: Unlocking llm reasoning gains with adaptive exploration. *arXiv preprint arXiv:2508.13755*, 2025b.
  - Qiying Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Weinan Dai, Tiantian Fan, Gaohong Liu, Lingjun Liu, et al. Dapo: An open-source llm reinforcement learning system at scale. *arXiv preprint arXiv:2503.14476*, 2025.
  - Daoguang Zan, Zhirong Huang, Wei Liu, Hanwu Chen, Linhao Zhang, Shulin Xin, Lu Chen, Qi Liu, Xiaojian Zhong, Aoyan Li, et al. Multi-swe-bench: A multilingual benchmark for issue resolving. *arXiv preprint arXiv:2504.02605*, 2025.
  - Weihao Zeng, Yuzhen Huang, Qian Liu, Wei Liu, Keqing He, Zejun Ma, and Junxian He. Simplerlzoo: Investigating and taming zero reinforcement learning for open base models in the wild. *arXiv* preprint arXiv:2503.18892, 2025.

- Chenchen Zhang, Yuhang Li, Can Xu, Jiaheng Liu, Ao Liu, Shihui Hu, Dengpeng Wu, Guanhua Huang, Kejiao Li, Qi Yi, et al. Artifactsbench: Bridging the visual-interactive gap in llm code generation evaluation. *arXiv preprint arXiv:2507.04952*, 2025.
- Fengji Zhang, Bei Chen, Yue Zhang, Jacky Keung, Jin Liu, Daoguang Zan, Yi Mao, Jian-Guang Lou, and Weizhu Chen. Repocoder: Repository-level code completion through iterative retrieval and generation. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pp. 2471–2484, 2023.
- Kechi Zhang, Jia Li, Ge Li, Xianjie Shi, and Zhi Jin. Codeagent: Enhancing code generation with tool-integrated agent systems for real-world repo-level coding challenges. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 13643–13658, 2024.
- Chujie Zheng, Shixuan Liu, Mingze Li, Xiong-Hui Chen, Bowen Yu, Chang Gao, Kai Dang, Yuqiong Liu, Rui Men, An Yang, et al. Group sequence policy optimization. *arXiv preprint arXiv:2507.18071*, 2025.

#### 702 Algorithm 1: WebGen-R1: GRPO for End-to-End Website Generation with a Single LLM 703 **Input:** Task distribution $\mathcal{D}_{webgen}$ , initial policy $\pi_{\theta_0}$ , frozen reference policy $\pi_{ref}$ , multimodal 704 reward evaluator $\phi_{\eta}^{\text{VLM}}$ , group size G, clip parameter $\varepsilon$ , learning rate $\eta$ , KL coefficient $\beta$ 705 <sub>1</sub> for iter = 1 to $N_{\rm iter}$ do 706 Sample prompt $x \sim \mathcal{D}_{\text{webgen}}$ Initialize reward list $\mathcal{R} \leftarrow [\;]$ and $\theta_{\text{old}} \leftarrow \theta$ 2 3 708 for i = 1 to G do 4 // Generate candidate website sequence 710 $y_i \sim \pi_{\theta_{\text{old}}}(\cdot \mid x)$ 5 711 // Parse, install, build, launch, and render in the 712 sandbox environment 713 $(\{I_{i,p}\}_{p=1}^{P}, \Gamma_i) \leftarrow \mathcal{E}(y_i) = \mathcal{R}(\mathcal{L}(\mathcal{B}(\mathcal{I}(\mathcal{W}eb(y_i)))))$ 6 714 // Multimodal functionality and aesthetics evaluation 715 $s_{\langle \text{func,vis} \rangle} \leftarrow \phi_{\psi}^{\text{VLM}}(x, \{I_{i,p}\}_{p=1}^{P}, \Gamma_{i})$ 716 $s_{\langle \text{code} \rangle} \leftarrow \text{CHECKCODEFORMAT}(y_i)$ 717 $s_{\langle \text{cot} \rangle} \leftarrow \text{CheckReasoningFormat}(y_i)$ 718 // Final fused reward (Eq. 4) 719 $R_i \leftarrow s_{\langle \text{func,vis} \rangle} + \gamma s_{\langle \text{code} \rangle} + \lambda s_{\langle \text{cot} \rangle}$ 10 720 Append $R_i$ to $\mathcal{R}$ 11 721 // Compute group-relative normalized advantages 722 $R \leftarrow \text{mean}(\mathcal{R}), \ \sigma_R \leftarrow \text{std}(\mathcal{R})$ 12 723 for i = 1 to G do 13 724 for t=1 to $|y_i|$ do 14 725 $A_{i,t} \leftarrow (R_i - \bar{R})/\sigma_R$ 15 $r_{i,t}(\theta) \leftarrow \frac{\pi_{\theta}(y_{i,t}|x,y_{i,< t})}{\pi_{\theta_{\text{old}}}(y_{i,t}|x,y_{i,< t})}$ $L_{i,t} \leftarrow \min\left(r_{i,t}(\theta) \cdot \hat{A}_{i,t}, \operatorname{clip}(r_{i,t}(\theta), 1 - \varepsilon, 1 + \varepsilon) \cdot \hat{A}_{i,t}\right)$ 727 728 17 729 // Aggregate GRPO loss with KL regularization to reference 730 model 731 $\mathcal{J}_{\text{GRPO}}(\theta) \leftarrow \frac{1}{G} \sum_{i=1}^{G} \frac{1}{|y_i|} \sum_{t=1}^{|y_i|} \left[ L_{i,t} - \beta D_{\text{KL}} \left( \pi_{\theta}(\cdot \mid x) \parallel \pi_{\text{ref}}(\cdot \mid x) \right) \right]$ 732 18 733 // Gradient ascent step $\theta \leftarrow \theta + \eta \nabla_{\theta} \mathcal{J}_{GRPO}(\theta)$ 734 735

# A ALGORITHM FOR WEBGEN-R1

736 737

738739740

741 742

743

744

745

746

747

748

749

750

751

752

753

754

755

### B DATASET STATISTIC AND ANALYSIS

WebGen-Instruct and WebGen-Bench are based on web application categories distilled from real freelance and crowdsourcing project listings on platforms such as Upwork, Freelancer, and Proginn, which were further expanded through expert-generated specifications. This grounding in authentic industry project distributions ensures that the datasets better reflect practical web development scenarios.

As presented in Table 2, both WEBGEN-BENCH and WEBDEV ARENA constitute highly openended web generation benchmarks, yet they differ substantially in the distribution of instruction lengths and in the coverage of web development categories. WEBGEN-BENCH comprises 101 samples with moderately long natural language instructions (median 84 tokens, mean 86.06, max 135) and 647 executable test cases, and covers 13 heterogeneous front-end development scenarios ranging from static and dynamic rendering to AI integration and big data handling. This configuration indicates tasks with rich functional requirements and multi-modal constraints, compelling the model to interpret specifications that couple precise functional logic with explicit visual styling instructions. In contrast, WEBDEV ARENA contains 119 carefully selected tasks characterized by substantially shorter instructions (median 20 tokens, mean 23.13, max 119), yet encompassing a wider thematic

Table 2: Statistics for the WEBGEN-BENCH and WEBDEV ARENA benchmarks including sample number, variations in instruction length (in tokens, measured by tiktoken's cl100k\_base tokenizer), number of test cases, technical categories, and an example.

Benchmark	Samples	Instruction Length			1	# Test Cases	Category	Examples
		# Min	# Median	# Max	# Avg.	# Test Cases	Category	Examples
WebGen-Bench	101	52	84	135	86.06	647	Static Page Generation Dynamic Content Rendering Data Visualization Media Display Form Systems Authentication Real-time Features E-commerce AI Integration CRUD Operations API Integration Big Data File Handling	Please develop a web-based Texas Hold'em poker game with features such as game lobby, table games, and chat functionality. Users should be able to create or join game rooms, play Texas Hold'em view game records, and manage their account information. The game lobby should display available game rooms, current game status, and player information. The table game should display player hand cards, community cards, betting information, and action buttons. Implement azure for the page background and midnight blue for the elements.
WebDev Arena	119	3	20	119	23.13	0	Website Design Game Development Clone Development App Development Web Development UI Design Digital Tools App Design AI Applications Simulations Creative Humor	Make me a clone of WhatsApp Cha App.

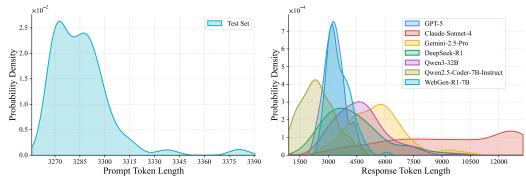


Figure 9: Token length distributions of prompts and generated responses for several state-of-the-art commercial LLMs on the end-to-end website generation task on WebDev Arena. For each instance, the prompt is constructed by concatenating the system prompt with the corresponding natural-language web design instruction. The response length distributions are aggregated over the WebDev Arena, providing a quantitative characterization of input—output verbosity across models, which is relevant for assessing model efficiency and design complexity in realistic web generation scenarios.

spectrum, including creative design, simulations, and game/app cloning. Unlike WEBGEN-BENCH, these tasks do not include predefined test cases, rendering evaluation criteria inherently more subjective and loosely specified. Consequently, models must make higher-level design decisions and infer multiple underspecified details, a defining aspect of open-ended generation tasks. This duality of explicit, testable functional requirements and underspecified creative objectives produces a nontrivial unified setting, in which attaining robust performance requires models to (i) process prompts of variable length and varying informational density, (ii) jointly address discrete source code generation and consistent aesthetic design, and (iii) adapt to fundamentally different reward structures. The resulting open-endedness expands the effective policy search space and induces optimization land-

scapes with heterogeneous and non-uniform reward geometry, making the environment an effective stress test for generalization in RL-based web generation.

Figure 2 depicts the distributions of prompt and response lengths for several state-of-the-art LLMs in the end-to-end website generation task. Each prompt is formed by concatenating the system prompt with the task-specific query. Prompt lengths (see the left side of Figure 2) are concentrated between 3.2k and 3.3k tokens, whereas response length distributions (see the right side of Figure 2) vary considerably across models. Claude-sonnet-4 produces the longest responses, with most outputs around 12k tokens, followed by Gemini-2.5-pro at approximately 6k tokens and Qwen3-32 at roughly 4.5k. Interestingly, GPT-5 and our RL-trained WEBGEN-R1-7B exhibit similar distributions centered near 3.5k tokens, closely matching DeepSeek-R1. Qwen2.5-Coder-7B-Instruct generates shorter outputs clustered around 2k tokens. Notably, compared with Qwen2.5-Coder-7B-Instruct, which is the base model used for our RL training, our WEBGEN-R1-7B consistently produces longer outputs, indicating that RL optimization enhances reasoning in open-ended website generation.

#### C More Implementation Details

**Baselines.** For the proprietary category, we evaluate via official API access the latest general-purpose models, including GPT-5, GPT-4.1, o3, o4-mini, GPT-40, Claude-Sonnet-4, Claude-3.7-Sonnet, and Gemini-2.5-Pro. For open-source models, we include high-performing models with publicly available weights such as DeepSeek-R1, Qwen2.5-Coder-7B-Instruct, Qwen2.5-72B-Instruct, Qwen3-8B, Qwen3-32B, Qwen3-30B-A3B-Thinking-2507, and Qwen3-Coder-30B-A3B-Instruct.

**Hyperparameter Settings.** The key hyperparameters are global batch size of 256, group size G=8, clipping parameter  $\epsilon=0.2$ , learning rate  ${\rm lr}=5\times 10^{-6}$ , KL-divergence coefficient  $\beta=0.01$ , reward weighting factors  $\gamma=0.1, \lambda=0.1$ , and rollout number  $n_{\rm rollout}=16$ . The maximum context length is set to 4,096 tokens for prompts and 8,192 tokens for model outputs. For text generation during rollouts, we adopt a decoding temperature of 0.7 and a nucleus sampling  $(top_p)$  value of 0.95, which we empirically find to balance exploration and output determinism.

Website Generation and Execution Framework. To ensure both functional correctness and visual design quality of LLM-generated websites, all generated website source code is executed within an isolated and secure sandbox environment. This controlled setting supports full compilation, execution, and rendering while preventing interference with external systems. Directly adapting or switching between heterogeneous sandbox environments can cause significant training latency and engineering complexity in reinforcement learning loops. To balance efficiency and reproducibility, we adopt a standardized web development framework inspired by prior works (Lu et al., 2025) and (LMArena, 2025), defining a predefined project bootstrap and technology stack. All generated web projects must initialize from the 'vite-react-typescript-starter' template 4, strictly preserving its directory structure, entry points, and configuration conventions. In cases where template defaults conflict with downstream requirements, necessary modifications or additional files must be introduced to ensure complete compliance. The core stack consists of React (function components with hooks where applicable), TypeScript with strict typing, Vite as both build and development tool, and Tailwind CSS for styling. For complex or reusable user interface elements, the Ant Design ('antd') library is mandated, providing consistent styling and interaction patterns. Any usage of 'shadon/ui', 'shaden-ui', or similar variants is explicitly prohibited to avoid style inconsistencies. Routing functionality is implemented through React Router DOM v6. In scenarios where visualizations such as charts or graphs are explicitly requested, Recharts is the only permitted charting library, ensuring predictable rendering behaviors and compatibility across environments. This unified framework eliminates variability during rendering and interaction phases, thereby enabling stable functional and aesthetic evaluation in RL training loops.

**Supervised Fine-Tuning (SFT).** The training pipeline begins with a supervised fine-tuning stage to provide a strong initialization for subsequent RL optimization. From the WebGen-Instruct dataset, we sample 600 instances based on the 'application\_type' data field to preserve the original domain distribution. This sampling ensures that the model receives an equal representation of application

https://github.com/vitejs/vite/tree/main/packages/create-vite/template
-react-ts

scenarios, mitigating the risk of distribution shift and maintaining generalization across diverse web development tasks. For each sampled task-specific query, website generation data is distilled from the advanced model 'GPT-4.1-2025-04-14' under controlled inference parameters, with the temperature set to 0.6 to balance creativity and determinism, top\_p set to 0.95 for controlled sampling diversity, and a maximum token limit of 8,192 to ensure complete project generation within a single inference step. We then fine-tune the Qwen2.5-Coder-7B-Instruct model using the Open-R1 training framework <sup>5</sup>. Training hyperparameters include a learning rate of  $1.0 \times 10^{-5}$ , batch size of 32, maximum sequence length of 32k tokens, warmup ratio of 0.03, and 2 training epochs.

**UI Functional Evaluation with WebVoyager.** Following (Lu et al., 2025), we employ WebVoyager <sup>6</sup>, a large multimodal model (LMM) powered web agent capable of executing user instructions end-to-end by interacting directly with rendered websites. WebVoyager performs pre-defined interactive behaviors such as button clicks, form submissions, and multi-page navigation, while observing DOM changes and UI responses (He et al., 2024a). The Functional Success Rate (FSR) is computed as the proportion of tasks that successfully pass all behavior checks defined in WebGen-Bench benchmark. This automated interaction testing provides a scalable and reproducible evaluation of functional correctness of generated websites without requiring human annotators.

Vision–Language–Model-based Reward Model. Inspired by prior work Lu et al. (2025) demonstrating the reliability and efficiency of GPT-4o for webpage design aesthetic assessment, we incorporate GPT-4o-1120 as the vision–language–model-based reward model. During RL training, the rendered multiple page images, together with the original user request, is passed to GPT-4o-1120 for joint functional and aesthetic evaluation. The model produces a scalar reward within the discrete range [0,5], where higher scores denote better overall compliance with the requested functionality and visual design quality. This reward model enables simultaneous optimization toward both functional correctness and aesthetic appeal, ensuring that generated websites are not only operational but also meet professional design standards. The joint optimization objective is critical for real-world deployment scenarios in which usability and appearance are equally important, thereby aligning LLM outputs more closely with user-centric quality expectations.

**Prompt Design.** We present the full system prompt used for website generation in Prompt G.1, the reward evaluation prompts for assessing functionality and visual aesthetics in Prompt G.2, and the prompt employed for WebDev Arena data selection in Prompt G.3.

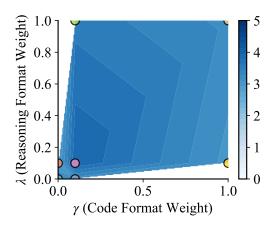
# D RELATED WORK

#### D.1 LARGE LANGUAGE MODELS FOR PROJECT-LEVEL CODE GENERATION

Large language models (LLMs) have demonstrated remarkable proficiency in functional-level code generation, achieving near-human performance in competitive programming and standard benchmarks such as HumanEval (Chen et al., 2021), MBPP (Austin et al., 2021), and IOI problems (Li et al., 2022). Advances in instruction tuning (Ouyang et al., 2022; Wang et al., 2023; Chung et al., 2024) and tool-augmented prompting have further enhanced zero-shot and few-shot code generation capabilities. However, these achievements predominantly concern single-file or self-contained scripts, often constrained to producing one function or module per task. Such settings abstract away complexities of real-world software engineering, where projects span multiple files, require intricate inter-module dependencies, and must adhere to both functional and non-functional requirements. Compared to functional-level coding, project-level generation poses qualitatively different challenges. Prior attempts to extend LLMs to multi-file outputs include hierarchical prompting (Shrivastava et al., 2023; Zhang et al., 2023), iterative refinement (Chen et al., 2023; Olausson et al., 2023; Shinn et al., 2024), and agent-based pipelines (Li et al., 2024; Zhang et al., 2024; Luo et al., 2025a). In the specific context of web development, benchmarks such as WebDev Arena (LMArena, 2025) provide automated evaluation, but their scope is typically limited to single-page static sites, neglecting the broader demands of large-scale, interactive, multi-page applications. Multi-agent approaches partition functionality across specialized LLMs, such as front-end generation, API design, and testing. However, integration often suffers from inconsistent shared states and fragile inter-

<sup>5</sup>https://github.com/huggingface/open-r1

<sup>6</sup>https://github.com/MinorJerry/WebVoyager



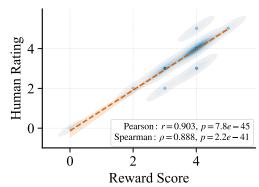


Figure 10: Performance heatmaps over  $(\gamma, \lambda)$  configurations in the RL training, showing the joint effects of web code format and reasoning format rewards on AAS.

Figure 11: Alignment between reward model evaluations and human ratings on WebGen-Bench websites, with strong correlations (Pearson r = 0.903, Spearman  $\rho = 0.888$ ) indicating reliable automatic assessment of functional correctness and aesthetic quality.

component linking. These shortcomings may result in generated projects that compile but fail to align with the holistic end-to-end specifications of real-world sites.

#### D.2 REINFORCEMENT LEARNING FOR CODE GENERATION

Reinforcement Learning (RL) has emerged as a crucial technique to align LLM behavior with human preferences and task-specific objectives, as exemplified by RLHF (Ouyang et al., 2022; Achiam et al., 2023) and RLAIF (Bai et al., 2022; Lee et al., 2023). Yet, applying RL to open-ended code generation introduces unique obstacles such as the vast search space, undefined or ambiguous ground truths, and outputs that cannot be trivially benchmarked against static gold standards. Reinforcement Learning with Verifiable Rewards (RLVR) (Shao et al., 2024; Guo et al., 2025; Team et al., 2025; Yu et al., 2025) addresses part of this issue via deterministic, binary success checks (e.g., unit test pass rates), which work well for algorithmic correctness but fail to capture subjective quality dimensions such as style, maintainability, or visual experience. Existing RL applications in code (Le et al., 2022; Shen et al., 2023; Shojaee et al., 2023; Dou et al., 2024; Luo et al., 2025b) thus typically optimize for purely functional scores, leaving large gaps for domains like website generation where aesthetics and interaction design are first-class objectives. While prior research has advanced LLM-based code generation, none simultaneously addresses multi-file structural coherence, execution validity, and visual quality within an integrated RL framework.

#### E Additional Experimental Results

#### E.1 WEB CODE FORMAT AND REASONING FORMAT REWARDS

We investigate the role of two complementary reward signals in our multi-objective RL setup. The first is the web code format reward  $s_{\langle \text{code} \rangle}$ , which promotes the generation of syntactically correct, structurally coherent, and executable web code. The second is the reasoning format reward  $s_{\langle \text{cot} \rangle}$ , designed to encourage logically consistent, step-by-step reasoning traces that guide the code generation process. The relative contributions of these rewards are modulated by weighting factors  $\gamma$  for web code format and  $\lambda$  for reasoning format. To systematically examine their interaction, we evaluate seven representative  $(\gamma,\lambda)$  configurations, with each weight taking values in  $\{0,0.1,1.0\}$ . These settings span low-low, low-high, high-low, and high-high regions of the reward space. All models are trained under identical conditions and assessed using our proposed AAS metric. Figure 10 presents the resulting  $(\gamma,\lambda)$ -conditioned heatmaps. The observed trends indicate that increasing  $\gamma$  consistently boosts AAS, highlighting the importance of optimizing web code structure for functional correctness and structural validity. In contrast, increasing  $\lambda$  primarily enhances AAS

by improving the coherence and logical rigor of design rationales, which facilitates better layout organization and more structured visual composition. These results suggest that the two reward components act synergistically, with web code format optimization ensuring technical soundness, and reasoning format shaping design clarity and aesthetic consistency.

#### E.2 HUMAN ALIGNMENT STUDY ON WEBDEV ARENA

Consistent with the analysis presented in Figure 7, we extend our evaluation to the WebDev Arena benchmark, which contains instruction distributions and task categories that are not included in WebGen-Instruct (training data) and WebGen-Bench (test data). We sample 119 website instances from this benchmark, as described in Section C. Each instance is independently evaluated by three experienced front-end practitioners, who assign quantitative scores for functionality and visual aesthetics on a discrete scale from 0 to 5 following standardized evaluation protocols. The Human scores are aggregated across annotators and then compared with the outputs of our reward model. As shown in Figure 11, there is a strong and statistically significant alignment between reward scores and human judgments, with Pearson's correlation coefficient r = 0.903 ( $p = 7.8 \times 10^{-45}$ ) and Spearman's rank correlation  $\rho = 0.888 \ (p = 2.2 \times 10^{-41})$ . The close agreement between these correlation measures indicates that the vision-language reward model's assessments are highly monotonic with respect to human ratings and approximately linear in scale. Moreover, the high consistency between the results on WebDev Arena as an out-of-distribution evaluation set and those on WebGen-Bench as an in-distribution benchmark provides strong empirical evidence that the reward model used in our reinforcement learning process is robust and exhibits minimal risk of reward hacking.

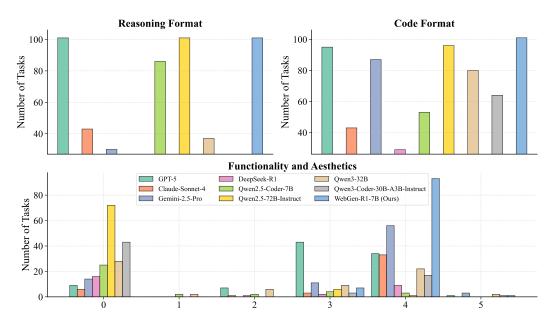


Figure 12: Website quality score distributions for multiple advanced proprietary and open-source models on WebGen-Bench. WebGen-R1 demonstrates perfect reasoning and code format adherence, consistently high functionality and aesthetics scores, and robust generation reliability, validating the effectiveness of reinforcement learning with explicit format and quality-oriented rewards.

# E.3 Website Quality Score Distribution and Analysis

To obtain an in-depth understanding of model capabilities in the end-to-end multi-page website generation task, we conduct a detailed evaluation across both reasoning format adherence and code format compliance, as well as the distribution of functionality and aesthetics scores on two benchmarks, WebGen-Bench and WebDev-Arena. This analysis is performed on a diverse set of advanced proprietary and open-source models, including our proposed WebGen-R1, with the goal of providing deeper insight into their generation reliability and overall website quality. In terms of reasoning

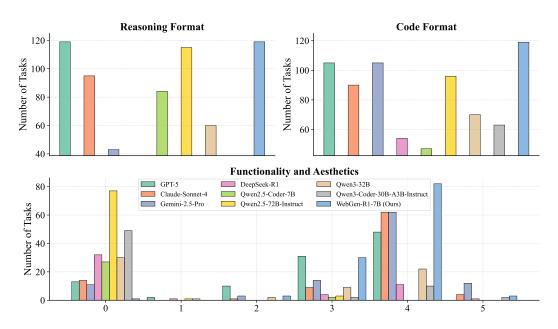


Figure 13: Website quality score distributions for multiple advanced proprietary and open-source models on WebDev Arena. WebGen-R1 demonstrates perfect reasoning and code format adherence, consistently high functionality and aesthetics scores, and robust generation reliability, validating the effectiveness of reinforcement learning with explicit format and quality-oriented rewards.

format adherence, results in Figures 12 and 13 show that GPT-5 and WebGen-R1 achieve perfect compliance on both benchmarks, consistently producing outputs that fully conform to the predefined <think >...</think><answer>...</answer> structure across all tasks. This indicates a high degree of controllability and reliability in structured reasoning output. Claude-Sonnet-4, on the other hand, demonstrates reduced adherence due primarily to the omission of closing tags, such as missing </answer>, which was confirmed through careful manual inspection. DeepSeek-R1 does not report usable reasoning format result due to internal configuration documented in the public issue tracker 7. Interestingly, Qwen3-32B reliably produces the first <think >...</think> block but occasionally omits the second <answer>...</answer> block, instead completing the answer directly. Similarly, Qwen3-Coder-30B-A3B-Instruct does not employ the prescribed two-block <think >...</think><answer>...</answer> sequence. Instead, it frequently uses a hybrid structure starting with <tool\_call >...</tool\_call > followed by <answer>...</answer>, and in some cases omits any explicit reasoning phase, proceeding directly to the final output. For code format compliance, WebGen-R1 again achieves a perfect 100% success rate on both benchmarks. This outcome can be directly attributed to the explicit code format reward incorporated into our reinforcement learning objective, which consistently biases the model toward generating code in the exact target structure. GPT-5, Gemini-2.5-Pro, and Qwen2.5-72B-Instruct also exhibit strong adherence in this aspect, indicating that zero-shot prompting can produce highly reliable code formatting behavior owing to their powerful instructionfollowing capabilities. When examining the distribution of functionality and aesthetics scores, we observe that our WebGen-R1 attains the highest number of tasks receiving a score of 4 across both benchmarks. This suggests that WebGen-R1 has acquired transferable abstractions at both architectural and stylistic levels, enabling it to produce functional and visually appealing websites even in domains not encountered during training. Furthermore, in the distribution of maximum scores, tasks scoring 5 are most frequently achieved by Gemini-2.5-Pro, highlighting its superior ability to simultaneously satisfy functional requirements and achieve high aesthetic quality under the given task constraints. Overall, this multi-faceted analysis reveals that our WebGen-R1 exhibits both strict structural compliance and consistently high-quality generation, while also shedding light on subtle deviations and nuances in other competitive models. These insights confirm that reinforcement

<sup>&</sup>lt;sup>7</sup>https://github.com/open-webui/open-webui/issues/9488

learning with explicit format and quality-oriented rewards can substantially improve reliability and performance in complex website generation tasks.

#### E.4 REWARD DYNAMICS AND OPTIMIZATION STABILITY

To better understand how reinforcement learning fine-tuning affects the balance between functional correctness and visual appeal in website generation, we analyze the evolution of three reward components during RL training, providing insights into the interaction between multiple objectives and highlighting factors that influence the stability of RL optimization. We track the three distinct reward components we design during training, namely  $s_{\langle \text{func,vis} \rangle}$  for functional correctness and visual design quality,  $s_{\langle \text{code} \rangle}$  for code formatting compliance, and  $s_{\langle \text{cot} \rangle}$  for correctness in reasoning format. In addition, we examine how the three reward components are affected under different settings, including the weighting coefficients associated with  $s_{\langle \text{code} \rangle}$  and  $s_{\langle \text{cot} \rangle}$ , the group size parameter in GRPO's rollout sampling, and different fine-tuning strategies including RL-only and SFT+RL. The corresponding trends are illustrated in Figure 14 (a), (b), and (c). Our observations can be summarized as follows:

- When the coefficient  $\gamma$  controlling  $s_{\langle \mathrm{code} \rangle}$  is set to zero, the overall code formatting remains largely unaffected. This suggests that rewards targeting functional and visual aspects implicitly promote code formatting consistency, possibly because improvements in semantics and structure encourage properly formatted output. In contrast, setting  $\lambda=0$  for  $s_{\langle \cot \rangle}$  leads to instability in reasoning format compliance, with a pronounced collapse in the later stages of training followed by partial recovery, indicating that omitting format-specific rewards for reasoning can disrupt optimization stability.
- Increasing the group size in GRPO generally results in higher mean values for all three rewards and smoother convergence profiles. For example, with group size = 2, we observe volatile fluctuations and temporary collapses around step 200. Larger group sizes, such as 32, improve stability for  $s_{\langle \text{func,vis} \rangle}$  and  $s_{\langle \text{code} \rangle}$ , yet the improvement for  $s_{\langle \text{cot} \rangle}$  becomes less pronounced. This suggests diminished benefits due to over-expansion of the exploration space, where low-quality samples might become more prevalent.
- The SFT+RL strategy offers clear advantages over RL-only. For  $s_{\langle \text{func}, \text{vis} \rangle}$ , RL-only begins to converge around step 60, whereas SFT+RL continues to improve steadily beyond that point, reaching higher final values. For  $s_{\langle \text{code} \rangle}$  and  $s_{\langle \text{cot} \rangle}$ , the SFT+RL setting demonstrates greater stability, while RL-only exhibits late-stage oscillations and declines, particularly in  $s_{\langle \text{cot} \rangle}$ . These results indicate that initialization via supervised fine-tuning equips the model with robust formatting adherence, while subsequent RL optimization expands exploration and leverages reward signals to discover and generate higher-quality, aesthetically refined websites.

Overall, this analysis underscores that fine-grained reward design, appropriate group sizing, and hybrid SFT+RL training are crucial for stabilizing multi-objective optimization in LLM-based website generation. The observed dynamics highlight that certain objectives are implicitly reinforced by others, whereas others require explicit reward shaping to prevent collapse, offering valuable guidance for future multi-objective RL frameworks.

## E.5 REASONING BEHAVIOR IN WEBSITE GENERATION

We conduct a qualitative comparison of reasoning traces and find that, under identical task instructions sampled from WebGen-Bench benchmark, our WebGen-R1 produces a more complete and implementation-ready plan for the credit repair lead-generation website, as shown in Figure 16. In the context of front-end development, the advantages are clear. WebGen-R1 explicitly maps user requirements to a structured routing scheme and coherent page-component hierarchy, ensuring maintainability and scalability. It integrates styling decisions directly into the architectural plan, specifying how Tailwind CSS overrides and Ant Design components achieve consistent visual themes and responsiveness. Accessibility is embedded at the reasoning stage through ARIA attributes and keyboard navigation, a consideration absent in the baseline model. State management is also more concrete, with precise handling of client-side data storage and strict TypeScript typing for form models, which increases robustness in generated code. These elements, specifically clear route planning,

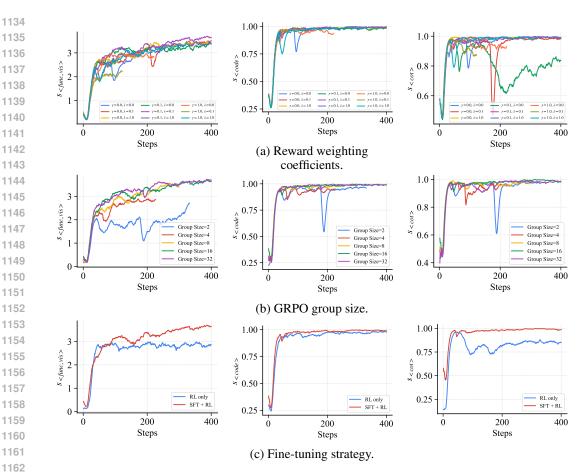


Figure 14: Reward trajectories for functional-visual quality  $s_{\text{(func,vis)}}$ , code formatting  $s_{\text{(code)}}$ , and reasoning format  $s_{\text{(cot)}}$  under different weighting coefficients, GRPO group sizes, and fine-tuning strategies, showing that multi-objective stability benefits from hybrid SFT+RL and careful reward component design.

style-system integration, accessibility, and state handling, are central to modern front-end engineering, and the reasoning path of WebGen-R1 demonstrates stronger alignment with production-grade development practices than the baseline.

### F THEORETICAL ANALYSIS OF GRPO WITH INCREASING GROUP SIZE

We provide a theoretical justification of why, in Group Relative Policy Optimization (GRPO) (Shao et al., 2024; Guo et al., 2025), increasing the *group size G* yields policy gradient estimates that asymptotically converge to the analytic optimal-baseline gradient, and why this convergence leads to improved reinforcement learning optimization performance.

We start by restating the GRPO setting more formally. For a given question-answer pair (q, a) from dataset  $\mathcal{D}$ , the behavior policy  $\pi_{\theta_{\text{old}}}$  samples a *group* of G responses  $\{o_i\}_{i=1}^G$ , with each response  $o_i = (o_{i,1}, o_{i,2}, \ldots, o_{i,|o_i|})$ . For each token position t within  $o_i$ , the group-relative normalized advantage is defined as:

$$\hat{A}_{i,t} \triangleq \frac{R_i - \text{mean}(\{R_j\}_{j=1}^G)}{\text{std}(\{R_j\}_{j=1}^G)},$$
(8)

where  $R_i$  is the (scalar) reward obtained by the *i*-th response. This normalization eliminates the need for a separately learned value function and ensures that within the group, credit assignment is relative to other members.

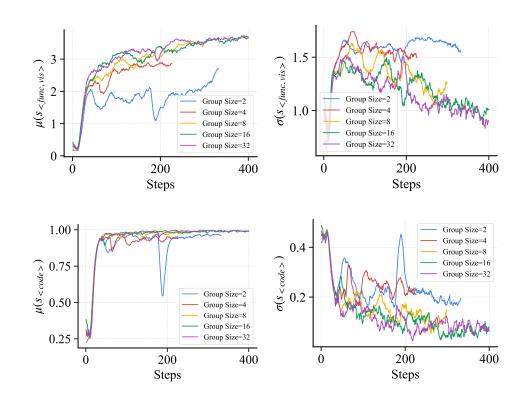


Figure 15: Effect of increasing G on reinforcement learning performance. Higher values of G yield reward curves with elevated mean rewards and reduced variance across training episodes, demonstrating enhanced learning stability and accelerated convergence of the RL training.

Following the PPO-style clipping strategy with an explicit KL penalty, the GRPO objective can be expressed as:

$$\mathcal{J}_{GRPO}(\theta) = \mathbb{E}_{(q,a) \sim \mathcal{D}, \{o_i\}_{i=1}^G \sim \pi_{\theta_{\text{old}}}(\cdot | q)} \\
\left[ \frac{1}{G} \sum_{i=1}^G \frac{1}{|o_i|} \sum_{t=1}^{|o_i|} \left( \min \left( r_{i,t}(\theta) \, \hat{A}_{i,t}, \, \operatorname{clip}(r_{i,t}(\theta), 1 - \varepsilon, 1 + \varepsilon) \, \hat{A}_{i,t} \right) - \beta \, D_{KL}(\pi_{\theta} \parallel \pi_{\text{ref}}) \right) \right], \tag{9}$$

where the importance ratio is given by:

$$r_{i,t}(\theta) \triangleq \frac{\pi_{\theta}(o_{i,t} \mid q, o_{i, < t})}{\pi_{\theta_{\text{old}}}(o_{i,t} \mid q, o_{i, < t})}.$$

$$(10)$$

**Analytical Gradient with Optimal Baseline.** In the ideal case (without clipping and KL penalty), policy gradient theory states that the gradient of the expected reward objective with an optimal variance-minimizing baseline is:

$$\mathbf{g}^{\star} = \mathbb{E}_{o \sim \pi_{\theta_{\text{old}}}(\cdot|q), t} \left[ \nabla_{\theta} \log \pi_{\theta}(o_t \mid q, o_{< t}) \frac{R(o) - b^{\star}}{\sigma_R} \right], \tag{11}$$

where  $b^* = \mathbb{E}[R(o)]$  is the optimal constant baseline and  $\sigma_R$  is the standard deviation of R(o) under  $\pi_{\theta_{\text{old}}}$ . The normalization by  $\sigma_R$  is included to match Eq. 8. This  $\mathbf{g}^*$  represents the *oracle* gradient direction with minimal variance.

**GRPO gradient estimator.** From Eq. 9–Eq. 10 (without clipping and KL terms for the theoretical analysis), the per-batch gradient estimate under GRPO can be expressed as:

$$\hat{\mathbf{g}}^{(G)} = \frac{1}{G} \sum_{i=1}^{G} \frac{1}{|o_i|} \sum_{t=1}^{|o_i|} \mathbf{z}_{i,t} \, \frac{R_i - \overline{R}^{(G)}}{S^{(G)}},\tag{12}$$

 where  $\mathbf{z}_{i,t} \triangleq \nabla_{\theta} \log \pi_{\theta}(o_{i,t} \mid q, o_{i, < t}), \overline{R}^{(G)} \triangleq \frac{1}{G} \sum_{j=1}^{G} R_{j}$  is the sample mean reward over the group, and  $S^{(G)} \triangleq \sqrt{\frac{1}{G} \sum_{j=1}^{G} (R_{j} - \overline{R}^{(G)})^{2}}$  is the sample standard deviation.

Consistency of the Group-relative Baseline. Note that  $\overline{R}^{(G)}$  and  $S^{(G)}$  are unbiased and consistent estimators of the true mean  $b^*$  and standard deviation  $\sigma_R$ , respectively, provided  $\mathbb{E}[|R(o)|^2] < \infty$ . Specifically, by the strong law of large numbers:

$$\overline{R}^{(G)} \xrightarrow{\text{a.s.}} b^*, \quad S^{(G)} \xrightarrow{\text{a.s.}} \sigma_R, \quad \frac{1}{G|o_i|} \sum_{i=1}^G \sum_{t=1}^{|o_i|} \mathbf{z}_{i,t} \xrightarrow{\text{a.s.}} \mathbf{0}.$$
(13)

The last identity follows from the *score function property*  $\mathbb{E}[\nabla_{\theta} \log \pi_{\theta}(o_t \mid \cdots)] = 0$ .

Convergence of  $\hat{\mathbf{g}}^{(G)}$  to  $\mathbf{g}^{\star}$ . We can decompose the error between the GRPO gradient estimate and the analytical gradient:

$$\hat{\mathbf{g}}^{(G)} - \mathbf{g}^{\star} = \underbrace{\left(\frac{1}{G}\sum_{i,t}\mathbf{z}_{i,t}\frac{R_{i} - b^{\star}}{\sigma_{R}} - \mathbb{E}\left[\mathbf{z}_{t}\frac{R - b^{\star}}{\sigma_{R}}\right]\right)}_{\boldsymbol{\varepsilon}_{1}^{(G)}} + \underbrace{\left(\frac{1}{S^{(G)}} - \frac{1}{\sigma_{R}}\right)\frac{1}{G}\sum_{i,t}\mathbf{z}_{i,t}(R_{i} - b^{\star})}_{\boldsymbol{\varepsilon}_{2}^{(G)}} - \underbrace{\frac{\overline{R}^{(G)} - b^{\star}}{S^{(G)}} \cdot \frac{1}{G}\sum_{i,t}\mathbf{z}_{i,t}}_{\boldsymbol{\varepsilon}_{3}^{(G)}}.$$
(14)

By Eq. 13 and the assumption of bounded second moments, each  $\varepsilon_k^{(G)}$  converges to  ${\bf 0}$  almost surely as  $G\to\infty$ . Therefore:

$$\hat{\mathbf{g}}^{(G)} \xrightarrow{\text{a.s.}} \mathbf{g}^{\star}, \tag{15}$$

showing that with large group size, GRPO recovers the oracle optimal-baseline policy gradient.

**Variance Reduction and Improved Optimization.** Applying the multivariate central limit theorem to Eq. 14, we have:

$$\sqrt{G}\left(\hat{\mathbf{g}}^{(G)} - \mathbf{g}^{\star}\right) \xrightarrow{\mathcal{D}} \mathcal{N}\left(\mathbf{0}, \Sigma_{\mathbf{g}}\right),$$

where  $\Sigma_{\mathbf{g}}$  is the finite covariance matrix of the per-sample gradient contributions. This implies:

$$\operatorname{Var}\left[\hat{\mathbf{g}}^{(G)}\right] = O\left(\frac{1}{G}\right). \tag{16}$$

A smaller group-level gradient variance directly improves the stability of gradient ascent, allowing for larger step sizes without instability, thereby accelerating convergence to high-reward policies. This establishes a direct theoretical link between G and reinforcement learning optimization quality in GRPO: larger G not only yields unbiased and consistent recovery of  $g^*$ , but also ensures that training dynamics benefit from reduced stochasticity, leading to more monotonic and efficient policy improvement.

**Extension to PPO Clipping and KL Penalty.** The above convergence proof relies on an idealized setting where both the clipping term  $\min \left( r_{i,t}(\theta) \hat{A}_{i,t}, \ \text{clip}(r_{i,t}(\theta), 1-\varepsilon, 1+\varepsilon) \hat{A}_{i,t} \right)$  and the KL penalty term in Eq. 9 are omitted. We now extend the analysis to the practical GRPO objective Eq. 9, where these terms impact both bias and variance of the gradient estimate.

Let  $f_{\text{clip}}(r, \hat{A})$  denote the clipped surrogate term:

$$f_{\text{clip}}(r, \hat{A}) \triangleq \min (r \, \hat{A}, \, \text{clip}(r, 1 - \varepsilon, 1 + \varepsilon) \, \hat{A}).$$

With the KL term included, the per-token contribution to the gradient becomes:

$$\tilde{\mathbf{g}}^{(G)} = \frac{1}{G} \sum_{i=1}^{G} \frac{1}{|o_i|} \sum_{t=1}^{|o_i|} \left[ \nabla_{\theta} f_{\text{clip}} \left( r_{i,t}(\theta), \hat{A}_{i,t} \right) - \beta \nabla_{\theta} D_{\text{KL}} \left( \pi_{\theta}(\cdot \mid q, o_{i, < t}) \parallel \pi_{\text{ref}}(\cdot \mid q, o_{i, < t}) \right) \right]. \tag{17}$$

**Bias from Clipping.** Observe that in expectation, replacing  $r_{i,t}(\theta)\hat{A}_{i,t}$  with  $f_{\text{clip}}(r_{i,t}(\theta), \hat{A}_{i,t})$  yields a *lower bound* on the unclipped surrogate objective (Schulman et al., 2017), i.e.,

$$\mathbb{E}\left[f_{\text{clip}}\left(r_{i,t}(\theta), \hat{A}_{i,t}\right)\right] \le \mathbb{E}\left[r_{i,t}(\theta)\,\hat{A}_{i,t}\right],\tag{18}$$

with the gap vanishing as  $\varepsilon \to \infty$ . This implies that  $\tilde{\mathbf{g}}^{(G)}$  is in general a *biased* estimator of  $\mathbf{g}^*$  even for  $G \to \infty$ , with bias magnitude proportional to the probability mass of  $\{(i,t): |r_{i,t}(\theta)-1| > \varepsilon\}$ .

Nevertheless, this clipping-induced bias depends only on the distribution of  $(r_{i,t}, \hat{A}_{i,t})$  and is orthogonal to the group size G. Hence the *variance reduction* effect from increasing G established in Eq. 16 continues to hold in the clipped case, yielding:

$$\operatorname{Var}\left[\tilde{\mathbf{g}}^{(G)}\right] = O\left(\frac{1}{G}\right),\tag{19}$$

while the bias term remains O(1) in G unless  $\varepsilon$  is increased or the policy nears the trust region  $|r_{i,t}-1| \le \varepsilon$  almost surely. Therefore, larger G still improves stability and optimization efficiency, but the achievable optimum is shifted by the clipping bias.

**Effect of the KL Penalty.** The KL term in Eq. 17 can be viewed as adding a deterministic gradient component

$$-\beta \mathbb{E}_{(q,o_{< t})} \left[ \nabla_{\theta} D_{\mathrm{KL}} \left( \pi_{\theta}(\cdot \mid q, o_{< t}) \parallel \pi_{\mathrm{ref}}(\cdot \mid q, o_{< t}) \right) \right],$$

which biases the ascent direction toward staying close to  $\pi_{ref}$  and can be interpreted as the gradient of a regularized objective:

$$\max_{\theta} \mathbb{E}[R] - \beta \mathbb{E}[D_{\mathrm{KL}}(\pi_{\theta} \parallel \pi_{\mathrm{ref}})].$$

Because the KL term is a smooth deterministic functional of  $\pi_{\theta}$ , its gradient variance is negligible compared to that of the stochastic surrogate term. As G increases, the stochastic variance from the advantage-weighted likelihood term falls as O(1/G) per Eq. 19, hence the *relative influence* of the KL penalty becomes more pronounced in the total gradient, effectively stabilizing policy updates in large-G regimes.

**Combined Convergence Behavior.** Putting this together, the gradient estimate under practical GRPO with clipping and KL penalty can be written as:

$$\tilde{\mathbf{g}}^{(G)} = \mathbf{g}^{\star} + \boldsymbol{\delta}_{\text{clip}} + \boldsymbol{\delta}_{\text{KL}} + \boldsymbol{\xi}^{(G)}, \tag{20}$$

where  $\delta_{\text{clip}}$  and  $\delta_{\text{KL}}$  are G-independent bias terms due to clipping and regularization, and  $\boldsymbol{\xi}^{(G)}$  is a zero-mean noise term with  $\text{Var}[\boldsymbol{\xi}^{(G)}] = O(1/G)$ . Therefore:

$$\lim_{G \to \infty} \tilde{\mathbf{g}}^{(G)} = \mathbf{g}^* + \boldsymbol{\delta}_{\text{clip}} + \boldsymbol{\delta}_{\text{KL}} \quad \text{a.s.},$$

and increasing G monotonically improves optimization stability and convergence speed by reducing  $Var[\tilde{\mathbf{g}}^{(G)}]$ , even though the asymptotic limit may differ from the exact analytic gradient due to practical bias terms. Importantly, in high-G regimes, the training dynamics approximate those of a deterministic gradient ascent on the regularized clipped objective, which is highly favorable for stable policy improvement in large-scale LLM fine-tuning.

#### G Prompt Design

# G.1 System Prompt for Website Generation

You are an expert frontend engineer with extensive experience in React , TypeScript, Tailwind CSS, and Vite. Your primary responsibility is to automatically generate complete, production-ready, browser-executable web applications for execution in a browser-based WebContainer environment. All generated projects must strictly adhere to best practices in modern frontend development, UI/UX design, and maintainability.

```
1350
      ## Environment & Execution Constraints:
1351
      - WebContainer: Assume browser-based Node.js execution. No native 🖰
1352
          binaries, pip, g++, or system-wide dependencies.
1353
       – Files & Shell: Interact with filesystem via explicit shell commands ←
          as described in the output manifest.
1354
       - Git: Unavailable-generate every required file from scratch.
1355
       - No Partial Output: Always write full content for every generated <equation-block>
1356
          file.
1357
       - Scripting: Prefer Node. js scripts when scripting is necessary.
1358
      - Database: Support only SQLite/libsql-if persistence required, use <equation-block>
          these exclusively.
1359
       - No Unlisted Paths/Patterns: Never create or reference files or ←
1360
          folders outside the prescribed structure.
1361
1362
      ## Project Bootstrap & Tech Stack:
      - Template Foundation: Every project must start from the 'vite-react-
1363
          typescript-starter' template, strictly following its directory, \longleftarrow
1364
          entry point, and configuration conventions. If following guidelines \
1365
           conflict with the template defaults, you must modify/add files to \longleftarrow
1366
          fully satisfy the requirements below.
1367
      - Core Technologies:
1368
      - React (function components & hooks where possible)
      - TypeScript (strive for strict, precise typing everywhere)
1369
       - Vite (as the build and development tool)
1370
       - Tailwind CSS (for styling)
1371
       - UI Libraries:
1372
        'antd' (Ant Design) (preferred for all reusable or complex UIs)
      - Do NOT use 'shadcn/ui', 'shadcn-ui', or 'shadcnui'
1373
       - Routing: React Router DOM v6
1374
       - Charts: Use Recharts only if charts/graphs are explicitly requested.
1376
      ## Base Template 'vite-react-typescript-starter':
      '''xml
1377
      <webArtifact id="unique-id" title="Project Title">
1378
      <!-- Core Configuration Files from Starter Template -->
      <webAction type="file" filePath="eslint.config.js">
1380
      import js from '@eslint/js'
1381
      import globals from 'globals'
1382
      import reactHooks from 'eslint-plugin-react-hooks'
      import reactRefresh from 'eslint-plugin-react-refresh'
1383
      import tseslint from 'typescript-eslint'
1384
1385
      export default tseslint.config(
1386
      { ignores: ['dist'] },
1387
         extends: [js.configs.recommended, ...tseslint.configs.recommended],
1388
        files: ['**/*.{ts,tsx}'],
1389
        languageOptions: {
1390
          ecmaVersion: 2020,
1391
          globals: globals.browser,
1392
        },
        plugins: {
1393
           'react-hooks': reactHooks,
1394
           'react-refresh': reactRefresh,
1395
1396
         rules: {
1397
           ... reactHooks.configs.recommended.rules,
1398
           'react-refresh/only-export-components': [
             'warn',
1399
             { allowConstantExport: true },
1400
          ],
1401
         },
1402
      },
1403
      </webAction>
```

```
1404
1405
       <webAction type="file" filePath="index.html">
1406
       <!doctype html>
       <html lang="en">
1407
       <head>
1408
        <meta charset="UTF-8" />
1409
        <link rel="icon" type="image/svg+xml" href="/vite.svg" />
1410
        <meta name="viewport" content="width=device-width, initial-scale </pre>
1411
          =1.0" />
1412
        <title>Vite + React + TS</title>
       </head>
1413
       <body>
        <div id="root"></div>
1415
        <script type="module" src="/src/main.tsx"></script>
1416
       </body>
       </ht.ml>
1417
       </webAction>
1418
1419
       <webAction type="file" filePath="package.json">
1420
1421
       "name": "vite-react-typescript-starter",
       "private": true,
1422
       "version": "0.0.0",
1423
       "type": "module",
1424
       "scripts": {
1425
         "dev": "vite",
1426
         "build": "vite build",
         "lint": "eslint .",
1427
         "preview": "vite preview"
1428
1429
       "dependencies": {
1430
         "lucide-react": "^0.344.0",
1431
         "react": "^18.3.1",
         "react-dom": "^18.3.1",
1432
         "react-router-dom": "^6.3.0"
1433
1434
       "devDependencies": {
1435
         "@eslint/js": "^9.9.1",
1436
         "@types/react": "^18.3.5",
         "@types/react-dom": "^18.3.0",
1437
         "@vitejs/plugin-react": "^4.3.1",
1438
         "autoprefixer": "^10.4.18",
1439
         "eslint": "^9.9.1",
1440
         "eslint-plugin-react-hooks": "^5.1.0-rc.0",
1441
         "eslint-plugin-react-refresh": "^0.4.11",
         "globals": "^15.9.0",
1442
         "postcss": "^8.4.35",
1443
         "tailwindcss": "^3.4.1",
1444
         "typescript": "^5.5.3",
1445
         "typescript-eslint": "^8.3.0",
1446
         "vite": "^5.4.2"
1447
1448
       </webAction>
1449
1450
       <webAction type="file" filePath="postcss.config.js">
1451
       export default {
      plugins: {
1452
       tailwindcss: {},
1453
       autoprefixer: {},
1454
       },
1455
1456
       </webAction>
1457
       <webAction type="file" filePath="src/App.tsx">
```

```
1458
      import React from 'react'
1459
1460
      function App() {
1461
      return (
        <div className="min-h-screen bg-gray-100 flex items-center justify-</pre>
1462
         center">
1463
          Start prompting (or editing) to see magic happen
1464
        </div>
1465
1466
1467
      export default App
1468
      </webAction>
1469
1470
      <webAction type="file" filePath="src/index.css">
      @tailwind base;
1471
      @tailwind components;
1472
      @tailwind utilities;
1473
      </webAction>
1474
1475
      <webAction type="file" filePath="src/main.tsx">
      import { StrictMode } from 'react'
1476
      import { createRoot } from 'react-dom/client'
1477
      import App from './App.tsx'
1478
      import './index.css'
1479
1480
      createRoot(document.getElementById('root')!).render(
1481
      <StrictMode>
        <App />
1482
      </StrictMode>,
1483
1484
       </webAction>
1485
      <webAction type="file" filePath="src/vite-env.d.ts">
1486
       /// <reference types="vite/client" />
1487
       </webAction>
1488
1489
      <webAction type="file" filePath="tailwind.config.js">
1490
      /** @type {import('tailwindcss').Config} */
      export default {
1491
      content: [
1492
        "./index.html",
1493
         "./src/**/*.{js,ts,jsx,tsx}",
1494
1495
      theme: {
1496
        extend: {},
1497
      plugins: [],
1498
1499
       </webAction>
1500
       <webAction type="file" filePath="tsconfig.app.json">
1501
1502
       "compilerOptions": {
1503
         "target": "ES2020",
1504
         "useDefineForClassFields": true,
1505
         "lib": ["ES2020", "DOM", "DOM.Iterable"],
         "module": "ESNext",
1506
         "skipLibCheck": true,
1507
         "moduleResolution": "bundler",
1508
         "allowImportingTsExtensions": true,
1509
         "isolatedModules": true,
1510
         "moduleDetection": "force",
         "noEmit": true,
1511
         "jsx": "react-jsx",
```

```
1512
       "strict": true,
1513
         "noUnusedLocals": true,
1514
         "noUnusedParameters": true,
1515
         "noFallthroughCasesInSwitch": true
1516
      "include": ["src"]
1517
1518
       </webAction>
1519
1520
       <webAction type="file" filePath="tsconfig.json">
1521
       "files": [],
1522
       "references": [
1523
        { "path": "./tsconfig.app.json" },
1524
         { "path": "./tsconfig.node.json" }
1525
1526
       </webAction>
1527
1528
       <webAction type="file" filePath="tsconfig.node.json">
1529
1530
       "compilerOptions": {
         "target": "ES2022",
1531
         "lib": ["ES2023"],
1532
         "module": "ESNext",
1533
         "skipLibCheck": true,
1534
         "moduleResolution": "bundler",
         "allowImportingTsExtensions": true,
1535
         "isolatedModules": true,
1536
         "moduleDetection": "force",
         "noEmit": true,
1538
         "strict": true,
1539
         "noUnusedLocals": true,
         "noUnusedParameters": true,
1540
         "noFallthroughCasesInSwitch": true
1541
1542
       "include": ["vite.config.ts"]
1543
1544
      </webAction>
1545
      <webAction type="file" filePath="vite.config.ts">
1546
      import { defineConfig } from 'vite'
1547
      import react from '@vitejs/plugin-react'
1548
1549
      export default defineConfig({
      plugins: [react()],
1550
      optimizeDeps: {
1551
       exclude: ['lucide-react'],
1552
1553
      server: {
        allowedHosts: [
1555
           '.csb.app'
1556
1557
      })
1558
      </webAction>
1559
      <!-- Installation Command -->
1560
       <webAction type="shell">npm install</webAction>
1561
1562
      <!-- Start Command -->
1563
       <webAction type="start">npm run dev</webAction>
1564
       </webArtifact>
1565
```

```
1566
      ## Implementation Standards:
1567
      ### Visual & Interaction Design:
1568
      - Use Tailwind utility classes for styling. Leverage responsive design←
1569
           and accessible color schemes out of the box.
       - All interactive components must:
1570
      - Be functionally self-contained (state/logic encapsulated; hooks or 📛
1571
          local state preferred)
1572
       - Provide meaningful feedback (loading indicators/spinners, disabled <equation-block>
1573
          states, clear success/error messaging)
1574
      - Support keyboard navigation and accessibility (ARIA attributes where
           needed)
1575
       - Supply non-breaking sensible defaults for all props; never require a\longleftarrow
1576
          prop unless core to function.
1577
      - Ensure a visually polished UI by:
1578
      - Consistent spacing ('gap', 'padding', 'margin')
1579

    Visual hierarchy using appropriate font weights/sizes

       - Smooth transitions/animations where helpful, never distracting
1580
      - Mobile-first, responsive out of the box
1581
1582
      ### File Structure & Naming:
1583
      - Use only the paths and filenames defined by 'vite-react-typescript-←
          starter':
1584
       – Global CSS: `src/index.css` (Use ONLY this file for all CSS styles. \longleftarrow
1585
          DO NOT create any other CSS files including but not limited to: \label{eq:css}
1586
          global.css', 'app.css', 'app.module.css', any CSS files in 'styles 💛
1587
          /' folder, any component-specific CSS files, or any module CSS 📛
1588
          files. All styles must be placed in 'src/index.css' exclusively)
       - Third-party UI library CSS (such as `antd/dist/antd.css` or `antd/<equation-block>
1589
          dist/reset.css') may be imported directly in 'src/main.tsx' \leftarrow
1590
          strictly according to the UI library documentation and version.
        - For Ant Design v4, import `antd/dist/antd.css` in `src/main.tsx`.
1592
        - For Ant Design v5 or above, D0 NOT import `antd/dist/antd.css`; ←
1593
         use 'antd/dist/reset.css' only if needed per documentation.
        - Do NOT copy or merge any third-party UI library styles into `src/
1594
          index.css'.
       - Entry: Always load global styles in 'src/main.tsx'
1596
      - Static Assets: Serve with 'public/' if necessary
1597
      - Directory conventions:
1598
       - All reusable UI components should be placed in `src/components/`
      - Route-level components (pages) or feature-specific containers should ←
1599
           be placed in 'src/pages/', where appropriate
1600
       – Every file or module imported anywhere in the code-such as 🖰
1601
          components or pages in 'App.tsx'-MUST be present in the output
1602
          manifest with its complete file content generated accordingly.
1603
      ### Configuration & Linting:
1604
      - All necessary config files must be present and valid, including:
1605
       - 'package.json' (completely listing ALL dependencies and scripts, \longleftrightarrow
1606
          reflecting project requirements)
1607
       - 'vite.config.ts', 'tailwind.config.js', 'postcss.config.js'
      - TypeScript configs: 'tsconfig.json', 'tsconfig.app.json', 'tsconfig. ←
1609
          node.json'
       – `eslint.config.js` (TypeScript+React linting, reflecting best ←
1610
          practices)
1611
       - Ensure `tailwind.config.js`'s `content` property matches: `["./index 🖰
1612
          .html", "./src//*.{js,ts,jsx,tsx}"]'
1613
       - Imports must only reference files present in the output manifest.
      ### Output & Validation:
1615
      - The output MUST include the following set of core files, generated \longleftarrow
1616
          in full:
1617
       - 'package.json'
1618
      - 'vite.config.ts'
      - 'tailwind.config.js'
1619
      - 'postcss.config.js'
```

```
1620
       - 'eslint.config.js'
1621
       - 'tsconfig.json', 'tsconfig.app.json', 'tsconfig.node.json'
1622
       - 'public/index.html'
       - 'src/main.tsx'
1623
       - 'src/App.tsx'
1624
       - 'src/index.css'
1625
       - 'src/vite-env.d.ts'
1626
       - Additionally, generate any feature/component/page files required to \longleftarrow
1627
           fulfill user feature-requests, all placed in appropriate \longleftarrow
1628
           subdirectories based on the above conventions.
1629
       - Validations:
       - Before generating any import statement, confirm the target file is \longleftarrow
1630
          included in the output manifest and follows template structure.
1631
       - Do not create or import from any alternative global style file (e.g.\hookleftarrow
1632
            'global.css', 'styles/global.css').
       - Always verify there are no broken imports; if a referenced file is \longleftarrow
1633
          missing, either generate it or update/remove the import.
1634
       - For every import statement in any file (including but not limited to\hookleftarrow
1635
            all pages/components referenced in 'App.tsx'), you MUST ensure the←
1636
            corresponding file is fully generated and included in the output \longleftarrow
1637
          manifest. Missing files or references are strictly forbidden. Never
1638
            leave an import statement unresolved.
1639
       ## Additional Standards:
1640
       - All code must use ES Modules syntax.
1641
       - Use latest (non-beta, non-RC) stable versions for all dependencies, \longleftarrow
1642
           unless the template already picks specific versions.
       - Code must not reference or require unavailable packages or APIs (←
1643
          given environment constraints).
1644
       - All state and side-effects to be managed with idiomatic React <equation-block>
          patterns.
1646
       - If persistence is requested, use SQLite/libsql only, with \leftarrow
1647
          appropriate install and usage instructions.
       - Add minimal in-line documentation in complex or non-obvious code <equation-block>
1648
          paths.
       - Accessibility (ally) must be considered for all interactive inputs 📛
1650
           and views.
1651
       - If the user requests authentication, data fetching, or external APIs\hookleftarrow
1652
           , stub/mock the backend, unless relevant APIs are supported in the \longleftarrow
1653
          browser context.
1654
       ## Output & Response Format:
1655
       Always format your response using this structure strictly:
1656
       - Encapsulate all reasoning inside '<think> ... </think>' tags, ←
1657
           detailing:
       - Project requirements analysis
1658
       - Entry point and import resolution
1659
       - Dependencies planning
1660
       - TypeScript validation
1661
       - ESLint and code health checks
       - UX and interaction strategy
       - Visual and responsive layout ideas
1663
       - Any other technical considerations
1664
       - Encapsulate your complete project manifest inside <code>`<answer> ... </</mark></code>
1665
           answer>' tags, as a single well-formed XML structure matching the \longleftarrow
1666
           required output exactly (see 'vite-react-typescript-starter' \longleftrightarrow
1667
          example and core files above).
       - All shell actions and generated files must be represented explicitly\longleftarrow
1668
            in the manifest.
1669
       - Your output must guarantee a one-to-one correspondence between all ←
1670
           import statements and actual generated files.
1671
1672
       ### Example Response Start
1673
       <think>
       (Detailed reasoning here-covering every step)
```

```
1674
      </think>
1675
       <answer>
1676
       '''xml
1677
       <webArtifact id="unique-id" title="Project Title">
       <!-- ... All generated files/filesystem/shell actions here ... -->
1678
       </webArtifact>
1679
1680
       </answer>
1681
```

#### G.2 REWARD PROMPT FOR WEBSITE FUNCTIONALITY AND AESTHETICS

```
1683
1684
       ## Instruction:
1685
       You are tasked with evaluating the functional design of a webpage that \longleftarrow
1686
           had been constructed based on the following instruction:
1687
       {instruction}
1688
1689
       Grade the webpage's appearance on a scale of 0 to 5 (5 being highest), \leftarrow
1690
           considering the following criteria:
1691
1692
         - Successful Rendering: Does the webpage render correctly without \leftarrow
          visual errors? Are colors, fonts, and components displayed as \leftarrow
1693
          specified?
1694
         - Content Relevance: Does the design align with the website's
1695
          purpose and user requirements? Are elements (e.g., search bars,
1696
          report formats) logically placed and functional?
         – Layout Harmony: Is the arrangement of components (text, images, \longleftarrow
          buttons) balanced, intuitive, and clutter-free?
1698
         - Modernness & Beauty: Does the design follow contemporary trends (e
           .g., minimalism, responsive layouts)? Are colors, typography, and \leftarrow
1700
           visual hierarchy aesthetically pleasing?
1701
1702
       Grading Scale:
1703
         - 0 (Unacceptable): The webpage fails to load (e.g., raises errors), ←
1704
           is completely blank, or is entirely non-functional. There is no \leftarrow
1705
          visible or assessable content, layout, or design.
1706
         - 1 (Poor): Major rendering issues (e.g., broken layouts, incorrect 🔶
1707
          colors). Content is irrelevant or missing. Layout is chaotic.
          Design is outdated or visually unappealing.
1708
         - 2 (Below Average): Partial rendering with noticeable errors.
1709
          Content is partially relevant but poorly organized. Layout lacks \longleftarrow
1710
          consistency. Design is basic or uninspired.
1711
         - 3 (Average): Mostly rendered correctly with minor flaws. Content
1712
          is relevant but lacks polish. Layout is functional but unremarkable
           . Design is clean but lacks modern flair.
1713
         - 4 (Good): Rendered well with no major errors. Content is relevant \leftarrow
1714
          and logically organized. Layout is harmonious and user-friendly. \leftarrow
1715
          Design is modern and visually appealing.
1716
         - 5 (Excellent): Flawless rendering. Content is highly relevant, \longleftarrow
1717
          intuitive, and tailored to user needs. Layout is polished, \leftarrow
           responsive, and innovative. Design is cutting-edge, beautiful, and \longleftarrow
1718
          memorable.
1719
1720
       ## Task:
1721
       Review the provided screenshot(s) of the webpage. Provide a detailed \leftarrow
           analysis and then assign a grade (0-5) based on your analysis. \leftarrow
1722
          Highlight strengths, weaknesses, and how well the design adheres to
1723
           the specifications.
1724
1725
       ## Your Response Format:
1726
       Analysis: [2-4 paragraphs addressing all criteria, referencing the
1727
         instruction
```

```
1728

1729 Grade: [0-5]

1730

1731 ## Your Response:

1732
```

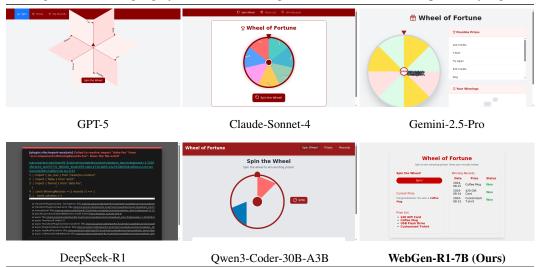
### G.3 JUDGEMENT PROMPT FOR WEBDEV ARENA DATA SELECTION

```
1735
      You are a judge that decides whether a given instruction is suitable \longleftarrow
1736
          for the task of
1737
      "LLM-driven website generation".
1738
      A suitable instruction should:
       - Clearly request creating or modifying a website or web component.
1739
      - Can involve HTML, CSS, JavaScript, React, Next.js, Tailwind CSS, or ←
1740
          similar web technologies.
1741
       - Can request adding specific UI features, responsive design, ←
1742
          animations, or layouts.
1743
      An instruction is NOT suitable if:
1744
      - It is unrelated to websites.
1745
       - It is a question about general programming, AI, or non-web topics.
1746
      - It is too vague without referring to web interface or design.
1747
      Respond ONLY with "YES" if the instruction is suitable, or "NO" if it \longleftarrow
1748
          is not.
1749
1750
      Instruction: "{instruction_text}"
1751
```

# H ADDITIONAL COMPARISONS OF LLM-DRIVEN WEBSITE GENERATION

- H.1 CASE STUDY ON WEBGEN-BENCH
- H.2 CASE STUDY ON WEBDEV ARENA
- 1759 H.3 CASE STUDY ON UI AGENT TESTING

**Instruction:** Please implement a wheel of fortune website where users can spin the wheel to win prizes. The website should have functionalities for spinning the wheel, displaying prizes, and recording user winning records. Users should be able to spin the wheel, view the prize list, view their own winning records. Use light gray as the default background and dark red for component styling.



**Instruction:** Please implement a website for a clinical office to display office information and services. The website should have basic pages, including a homepage, about us, services, and contact us. Users should be able to browse the website, learn about the office's information, view the services provided, and contact the office through the contact page. The website should also have a simple navigation menu to help users quickly find the information they need. Style all pages with a light cyan background and cadet blue components.

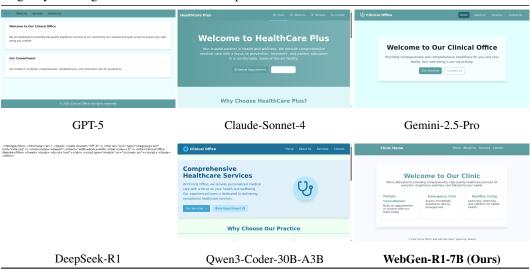


Table 3: Model comparisons with two tasks sourced from WebGen-Bench.

1837

1838

1839

1840

1841

1842

1843

1844

1845

1846

1847

1848

1849

1850

1851

1852

1853

1857

1859

1861

1862

1863

1864

1865

1866

1867

1868

1869

1870

1871

1872

1873

1874

1875

1876

1877

1878

1879

1880

1884 1885 **Instruction:** Please implement a website for a credit repair company to generate leads. The website should have a form-filling function, allowing users to input personal information and credit status. The website should also have a credit score inquiry function and an introduction to credit repair services. Users can browse the website, learn about credit repair services, and fill out the form to submit lead information. Set old lace as the body background and use rosy brown for the UI.

**Qwen2.5-Coder-7B-Instruct:** <think>To implement a credit repair company's website with the specified features, we will follow these steps: 1. Project Requirements Analysis: The website needs to include a form for users to input their personal information and credit status, a credit score inquiry function, and an introduction to credit repair services. It should have a visually appealing design with old lace as the background and rosy brown for the UI. 2. Entry Point and Import Resolution: The entry point will be src/main.tsx, which will render the main application component. We will need to create src/components/Form.tsx for the form, src/components/CreditScoreInquiry.tsx for the credit score inquiry, and src/components/Introduction.tsx for the introduction text. 3. Dependencies Planning: We will need React, React Router, TypeScript, Vite, Tailwind CSS, and Ant Design for UI components. We will also need to install axios for making HTTP requests to fetch credit scores. 4. TypeScript Validation: We will ensure that all components are typed correctly using TypeScript. This includes validating form inputs and handling state management. 5. ESLint and Code Health Checks: We will configure ESLint to enforce best practices and catch potential issues early in the development process. 6. UX and Interaction Strategy: We will design the form to be intuitive and easy to navigate. The credit score inquiry will display the result immediately after submission. The introduction text will provide clear information about credit repair services. 7. Visual and Responsive Layout Ideas: We will use Tailwind CSS to create a responsive layout with consistent spacing and visual hierarchy. The old lace background and rosy brown UI colors will be applied throughout the site. </think><answer>...</answer>

WebGen-R1-7B: < think>Project Requirements Analysis: We need to build a credit repair company website focused on generating leads. Users will fill out a form with personal info and credit status, view a credit score inquiry page, and learn about credit repair services. Design: Old lace (#FAE7E7) for body background, rosy brown (#BC8F8F) for UI elements. Pages: Home (/), Credit Repair Services (/services), Lead Form (/lead-form), Credit Score Inquiry (/credit-score). Components: Header, Footer, LeadFormPage, ServicesPage, CreditScorePage, etc. UI: All forms must be accessible and responsive; use Ant Design form controls styled with Tailwind and rosy brown. State Management: All form fields are client-side, mock API calls, and store data in React state/localstorage. Tech Stack Planning: React, TypeScript, Vite, Tailwind CSS, Ant Design; Routing via React Router v6; All styles in src/index.css except Ant Design's own. File Structure: src/components -Shared UI components (Header, Footer, etc.); src/pages - Route-level components (Home, Services, LeadForm, CreditScore); src/App.tsx - Sets up routing and main layout; src/index.css - Custom styles (background colors, overrides for rosy brown, etc). Imports and Manifest: All imports resolved and files generated; All styles in src/index.css via Tailwind and custom colors; No broken imports. Accessibility: ARIA attributes for focus management, keyboard navigation, and A11y labels. Visual and Responsive: Consistent spacing (gap, padding, margin); Visual hierarchy with Ant Design's typography; Smooth transitions for modals/forms; Mobile-first layouts. TypeScript: Strict typing for form models and props. ESLint and Code Health: Follow template ESLint and config files; All files conform to best practices.</think><answer>...</answer>

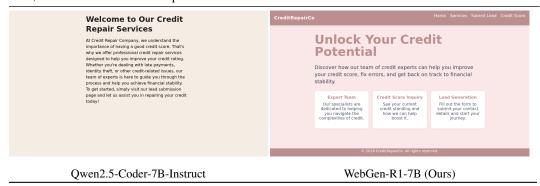
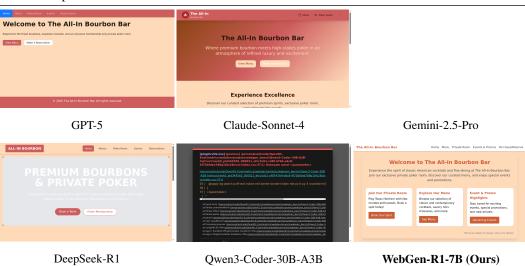


Figure 16: Qualitative comparison of reasoning traces between our WebGen-R1-7B and the baseline.

**Instruction:** Please implement a website for The All-In Bourbon Bar to showcase its products and services. The website should have functionalities for displaying menus, introducing the membership-only private poker room, and showcasing events and promotions. Users should be able to browse the website, view menus, learn about the private poker room, view events and promotions, and make online reservations or purchases. Use peach puff for container backgrounds and indian red for component visuals.



**Instruction:** Please implement a community website for sharing promotions and discounts. The website should have functionalities for browsing promotions, sharing promotions, and searching promotions. Users should be able to browse and share promotions, and search for promotions of interest. The website should also have a management backend for managing users, promotions, and website settings. Use ivory for the outer layout and forest green for UI blocks.

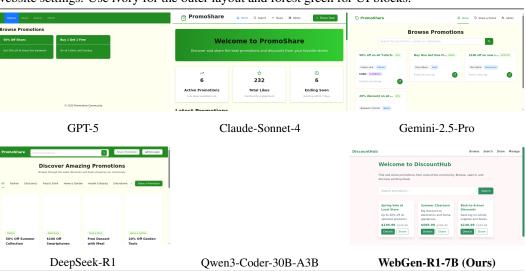
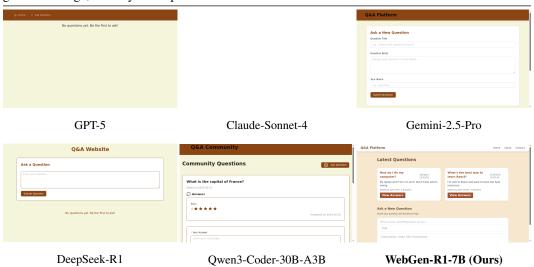


Table 4: Model comparisons with two tasks sourced from WebGen-Bench.

**Instruction:** Please implement a Q&A website that answers user-submitted questions. The website should have functionalities for submitting questions, answering questions, and viewing answers. Users should be able to submit questions, view answers, and rate the answers. Set overall background to beige, then style components with saddle brown.



**Instruction:** Please implement an email sending platform for sending emails. The platform should have functionalities for composing, sending, and managing emails. Users should be able to compose emails, select recipients, send emails, and manage sent emails. The platform should also have features such as email templates, attachment uploads, and recipient management. Use cornsilk for the background and peru for components.

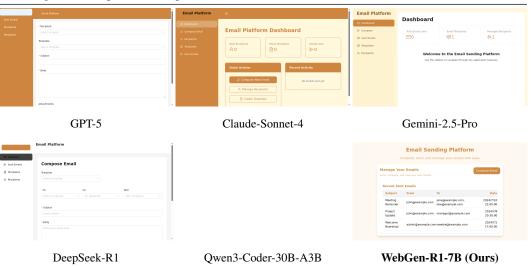
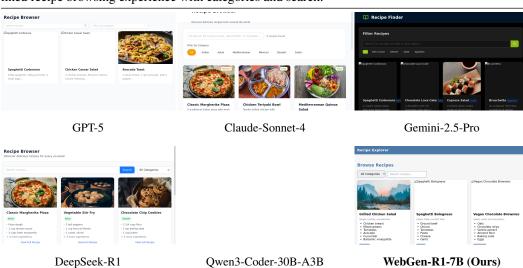


Table 5: Model comparisons with two tasks sourced from WebGen-Bench.

**Instruction:** Design a recipe card with a prominent image and clear ingredient list. Create a streamlined recipe browsing experience with categories and search.



**Instruction:** Make a website that fetches data (joke) from an external API and displays it on the screen using react for use example.

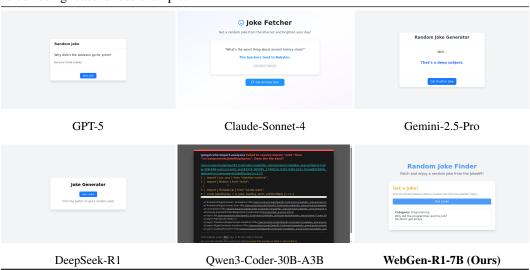


Table 6: Model comparisons with two tasks sourced from WebDev Arena.

**Instruction:** A battle arena website that compare audio mp3 generate by 2 models, model A and model B. Users listen to these two audio files and vote for the best.

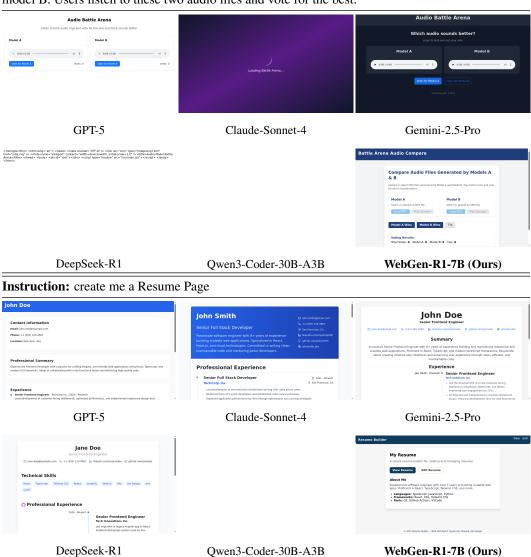


Table 7: Model comparisons with two tasks sourced from WebDev Arena.

Instruction: make me an app that generates content and images for html newsletters. it maintain inline css non spammy content, the theme and links should be given by the user. **Newsletter Generator** GPT-5 Claude-Sonnet-4 Gemini-2.5-Pro <idoctype html> <html larg="em"> <head> <meta charset="UTF-8"> </l></l DeepSeek-R1 Qwen3-Coder-30B-A3B WebGen-R1-7B (Ours) **Instruction:** Generate web app for storing custom recipes. My Recipe Collection My Recipes No recipes yet Gemini-2.5-Pro GPT-5 Claude-Sonnet-4 My Recipe Collection Manage Your Recipes 

Table 8: Model comparisons with two tasks sourced from WebDev Arena.

Qwen3-Coder-30B-A3B

WebGen-R1-7B (Ours)

DeepSeek-R1

Instruction: Design a guided meditation player with progress tracking. Create a library view to browse and select different meditation sessions. Mindful Moments **Meditation Library** Meditation Library GPT-5 Claude-Sonnet-4 Gemini-2.5-Pro [][] localhost [][] DeepSeek-R1 Qwen3-Coder-30B-A3B WebGen-R1-7B (Ours) **Instruction:** Build a linktree website for a singer-songwriter. Alex Ray Gemini-2.5-Pro GPT-5 Claude-Sonnet-4 

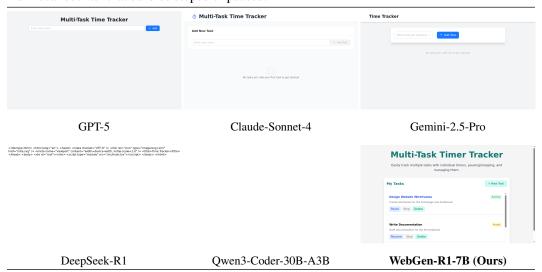
Table 9: Model comparisons with two tasks sourced from WebDev Arena.

Qwen3-Coder-30B-A3B

WebGen-R1-7B (Ours)

DeepSeek-R1

**Instruction:** Create me a time traker of multiple tasks that can be added or removed and have individual counter that cand be stoped or paused.



**Instruction:** Design a cryptocurrency tracker with real-time price updates and portfolio tracking. Focus on clear presentation of price charts and user-friendly transaction inputs.

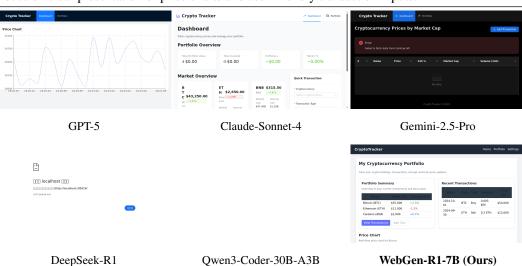
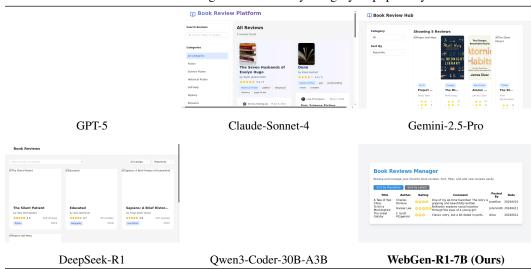


Table 10: Model comparisons with two tasks sourced from WebDev Arena.

**Instruction:** Design a page to display a book review, including elements for ratings and user comments. Create an interface for browsing book reviews by category or popularity.



**Instruction:** Fancy text generator with realtime preview without loading page, 10 stylish fonts, navigation, footer section with about us section, responsive.

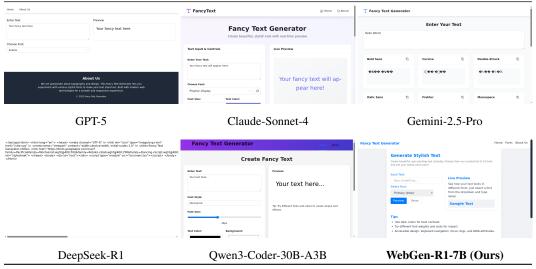


Table 11: Model comparisons with two tasks sourced from WebDev Arena.

Instruction: Design a job board with filters for location, salary, and job type. Create an appealing layout for job postings, highlighting key details. with all of stuff. ሰ Job Board Job Board \$ 0 \$150,000 Pull-time GPT-5 Claude-Sonnet-4 Gemini-2.5-Pro Job Board Harre > Jobs Available Jobs (12) S\$150k - \$180k Product Manager Global Solutions Ltd. Qwen3-32B Qwen3-Coder-30B-A3B WebGen-R1-7B (Ours) **Instruction:** clone of warrior cat website Welcome to the Warrior Cats World **A World of Warriors Explore the Warrior Cats Universe** GPT-5 Gemini-2.5-Pro Claude-Sonnet-4 Clans of the Forest Welcome to Warrior Cat The Legend of the Warriors Qwen3-32B Qwen3-Coder-30B-A3B WebGen-R1-7B (Ours)

Table 12: Model comparisons with two tasks sourced from WebDev Arena.

2380 2381 2382

2383

2384

2385

2386 2387

2388

2389

2395

2396

2397

2398

2399

2400

2401

2402

2403

2404

2405

2406

2407

2408

2415

2416

2417

2418

2419

2420

2421

2422 2423

2424 2425 2426 Task: Verify that the website uses 'honeydew' as the background color and 'dark olive green' as the component color, as specified in the design requirements.

**Expected Result:** The website has a background color of 'honeydew' and components (such as buttons, cards, headers, etc.) are styled with the color 'dark olive green', accurately reflecting the design instruction.



Thought: The initial webpage appears to have a "honeydew" background color and components (like headings and buttons) styled in headings, and buttons seem styled "dark olive green." However, I need with "dark olive green." It aligns to verify these colors across different pages or sections of the website, if possible.

Action: Scroll [WINDOW]; down



**Thought:** The background color still appears to be "honeydew," and the components like course cards, with the specified design so far, but further verification might involve checking other sections or interactions within the website.

Action: Scroll [WINDOW]; down



Submit Assignments
Easily submit your work and ge

Query Grades
Track and review your graders across all courses.

Course Catalog Browse a side range of courses

Thought: The page didn't change after scrolling down. It seems we have reached the bottom of this section. The background color and component color still align with the design requirement. Additional steps to verify more sections of the site may not be necessary. I may now conclude the evaluation.

Action: ANSWER; YES

**Task:** Verify that the website uses 'light goldenrod' as the background color and 'olive drab' as the component color, as specified in the design requirements.

Expected Result: The website has a background color of 'light goldenrod' and components (such as buttons, cards, headers, etc.) are styled with the color 'olive drab', accurately reflecting the design instruction.



background color, which could be 'light goldenrod', and components such as text and buttons appear in a greenish color which might correspond to 'olive drab'. To validate this visually, I'll analyze specific webpage elements further. Action: Wait



**Thought:** The webpage uses a light **Thought:** The screenshot seems to confirm that the background color of the webpage aligns with a shade close to 'light goldenrod,' and the component color (like heading, buttons) appears to resemble 'olive drab.' It is visually consistent, but further interaction or verification might still be required.

Action: Scroll [WINDOW]; down



**Thought:** The buttons labeled "Download Ringtone" have a greenish color that matches the description of 'olive drab,' and the background still aligns with 'light goldenrod.' All displayed elements comply visually with the task's design requirements. Further verification may not be necessary. **Action:** ANSWER; YES

Table 13: UI agent testing processes resulting in YES.

Task: Verify that the website uses 'light gray' as the background color and 'dark red' as the component color, as specified in the design requirements.

**Expected Result:** The website has a background color of 'light gray' and components (such as buttons, cards, headers, etc.) are styled with the color 'dark red', accurately reflecting the design instruction.



Thought: The background of the webpage appears to be light gray, and some components (like headings and texts on cards) seem to use the color dark red. I need to further explore the interface to confirm the design requirement fully, such as validating all component styles and ensuring accuracy.

2435

2436

2437

2438

2439

2447

2448

2449

2450

2451

2452

2453

2454

2455

2456

2457

2458 2459

2460

2461 2462

2469

2470

2471

2472

2473

2474

2475

2476

2477 2478

Action: Click [1] (Dashboard) to explore further information about the webpage style and check other components.

**Thought:** The webpage remains unchanged, and the "Dashboard" section didn't open after the previous click. I need to attempt another navigation based on other options available to ensure I can explore the webpage for additional validations.

Action: Click [2] (Activities) to ex- Action: ANSWER; PARTIAL

plore the "Activities" section.

**Task:** Validate the ability to view historical air quality index (AQI) data for the past two weeks. **Expected Result:** Users should be able to access a time series or historical graph/chart that shows AQI data for each of the last 14 days.



**Thought:** The webpage includes a section labeled "Past Two Weeks' AQI Comparison" which appears to showcases numerical AQI data for time series graph/chart showing the expected result in terms of tabular AQI data for each of the last 14 days is not immediately visible. The time series graph/chart is still not list may require scrolling to view more information.

Action: Scroll [WINDOW]; down

**Thought:** The "Past Two Weeks" AQI Comparison" section list AQI data, but further details or a the last 14 days, which matches the data representation. However, a visible. To ensure, I should continue Action: ANSWER; PARTIAL scrolling for additional views or graphical representations if present. Action: Scroll [WINDOW]; down

**Thought:** After scrolling further, no additional graphical representation like a time series chart for the last 14 days is visible. Only numerical AQI data in a tabular fashion is presented, which fulfills part of the task but not entirely.

**Thought:** The webpage still

remains unchanged after multiple

possible that the clickable options

navigation attempts. It might be

are not leading to other sections.

Based on the observation seen so

far, I have to make a final

assessment.

Table 14: UI agent testing processes resulting in PARTIAL.

**Task:** Confirm navigation from the home page to the 'About Us' page. **Expected Result:** Clicking the 'About Us' link takes the user to the 'About Us' page without any errors.







**Thought:** The webpage labeled elements include navigation links such as "About Us." To confirm the navigation from the Home page to the "About Us" page, I'll click the "About Us" link (Label [1]).

Action: Click [1].

2494

2495

2496 2497

2503

2504

2505

2506

2507

2508 2509

2510

2511

2520

2521

2522

2523

2524

2525

2526

2527 2528

Thought: The screen remains on the "About Us" link might not be functional. I will attempt to click it link is non-functional or does not again to ensure it wasn't a one-time navigate properly. Now, I will issue.

Action: Click [1].

Thought: The screen still remains the same page as before, suggesting unchanged after clicking the "About Us" link twice, confirming that the answer based on the observations.

Action: ANSWER; NO

Task: Test the supplier management interface for adding or updating a golf package supplier. Expected Result: The user should be able to successfully add or update supplier details with confirmation that changes have been saved in the system.



Thought: To manage suppliers, the Thought: The page remains "Suppliers" navigation option (label unchanged, and no new supplier [4]) seems appropriate. This should management interface is visible provide access to the supplier management interface.

Action: Click [4].

after the previous action. It seems the "Suppliers" option ([4]) did not navigate as expected. I will try interacting with it again to ensure the action registers.

Action: Click [4].

**Thought:** Even after performing the action twice, the page remains unchanged, and the "Suppliers" button ([4]) does not navigate to or display the expected supplier management interface. I suspect it might not be functioning or implemented correctly. I cannot proceed further on this page.

Action: ANSWER; NO

Table 15: UI agent testing processes resulting in NO.