

000 001 REINFORCEMENT LEARNING FOR SYMBOLIC GRAPH- 002 ICS CODE GENERATION WITH VISUAL FEEDBACK 003 004

005 **Anonymous authors**

006 Paper under double-blind review

007 008 ABSTRACT 009

011 Symbolic graphics code generation, particularly text-to-SVG generation, plays
012 a critical role in numerous practical applications, including web design, digital
013 publishing, and user interface prototyping. However, current open large language
014 models face significant challenges in handling these visually intricate and struc-
015 turally precise tasks, often exhibiting a considerable performance gap compared to
016 leading proprietary models. In this paper, we present a novel approach aimed at sub-
017 stantially improving the capabilities in text-to-SVG tasks. Our main contributions
018 are threefold: First, we propose a reinforcement learning framework that leverages
019 vision-language models (VLMs) as visual reward model, providing comprehensive
020 visual feedback that guides LLMs towards generating more accurate and visually
021 coherent SVG outputs. Second, we investigate inference-time scaling methods
022 through extended long Chain-of-Thought (CoT) reasoning combined with large-
023 scale RL, revealing that such methods inherently counteract reward hacking by
024 refining prompt engineering and making task objectives more explicit and concrete.
025 Third, we introduce a new, high-quality benchmark alongside a rigorously curated
026 training dataset dedicated to text-to-SVG generation, addressing the notable ab-
027 sence of specialized benchmarks and datasets in this domain. Experiments on open
028 model, i.e., Qwen3 demonstrate that our approach achieves results comparable to
029 state-of-the-art proprietary and larger models, including Claude-4.0-Sonnet. This
030 work substantially narrows the performance gap and provides both methods and
031 resources to advance symbolic code generation research.

032 1 INTRODUCTION

033 Large language models (LLMs), have made remarkable progress across a wide range of do-
034 mains (Hurst et al., 2024; Jaech et al., 2024; OpenAI, 2025c; Team et al., 2023), including question
035 answer (Ouyang et al., 2022; Yang et al., 2025a; Team et al., 2024), code generation (Hui et al., 2024),
036 and complex problem-solving (Guo et al., 2025; Wang et al., 2024b; Yang et al., 2024; Team, 2025).
037 Nevertheless, generating symbolic graphics code, particularly from natural language to Scalable
038 Vector Graphics (SVG), remains a persistent challenge (Nishina & Matsui, 2024). Unlike conven-
039 tional code generation, text-to-SVG tasks require not only syntactic correctness but also adherence
040 to structural precision and visual semantics (Cai et al., 2023). Bridging the gap between textual
041 descriptions and visually faithful SVG would significantly advance a variety of practical applications,
042 including digital publishing, web design, educational illustration, and user interface prototyping.

043 Although proprietary models have achieved strong performance in symbolic graphics tasks (Yang
044 et al., 2025b), open-source models still fall noticeably behind. One major factor underlying to this
045 gap is the lack of visual feedback during pre-training and post-training. Many current methods
046 focus primarily on textual correctness, overlooking the visual quality of the rendered outputs. Yet
047 in symbolic graphics generation, success is ultimately determined by how well the generated image
048 reflects the intended meaning of the input. Without explicit visual feedback, models often produce
049 outputs that are syntactically correct but visually misaligned with the user’s instructions.

050 To address this limitation, we propose a reinforcement learning (RL) framework that incorporates
051 visual feedback into the post-training of LLMs. Central to our approach is the use of frozen vision-
052 language models (VLMs) (Bai et al., 2025; OpenAI, 2023; Lu et al., 2024; Wu et al., 2024b) as
053 perceptual judges, which evaluate the alignment between the rendered SVG output and the reference

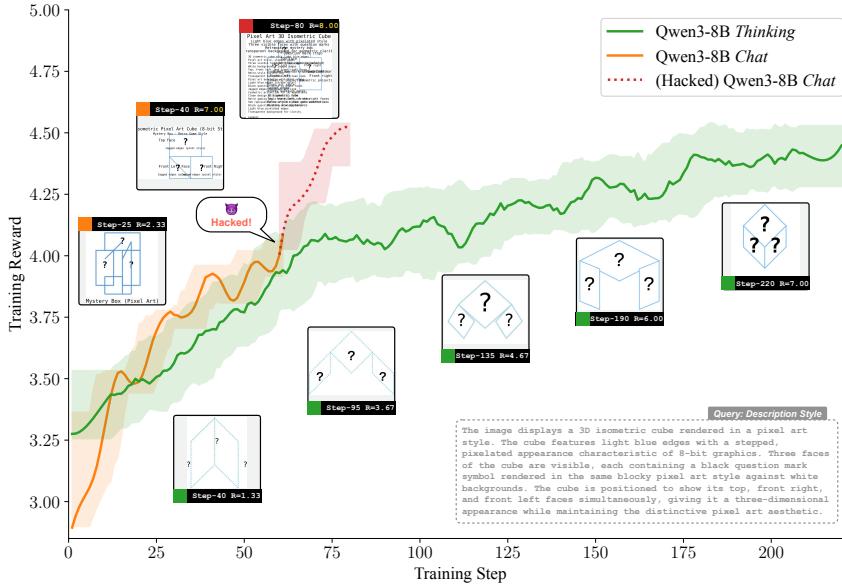


Figure 1: Training curves for Qwen3-8B under different settings. The *Chat* model (orange) initially improves but exhibits reward hacking (red dotted) by inserting descriptive text instead of graphical elements. In contrast, the *Thinking* model (green), enabled by inference-time reasoning, maintains consistent gains without reward hacking, producing visually faithful SVG code.

image. Rather than relying on token-level or syntactical correctness, the reward signal is grounded in perceptual fidelity across multiple visual dimensions (Baumli et al., 2023). These include object presence and accuracy, spatial arrangement, and overall stylistic quality (Rodriguez et al., 2023). The resulting reward functions are not only sensitive to visual details but are also fine-grained, enabling scalable optimization. Compared to conventional methods, our approach introduces a stronger inductive bias toward generating SVG code that is both semantically faithful and visually coherent.

However, the introduction of perceptual rewards also brings new challenges. As shown in Figure 1, reward hacking is observed (Weng, 2024), where the model exploits weaknesses in the reward signal, such as embedding descriptive text into the image in place of rendering the graphical elements. This behavior is largely driven by biases in vision-language models pretraining stage, many of which are pretrained with grounding data that assigns comparable importance to textual overlays and visual features (Wang et al., 2024a; Bai et al., 2025). To mitigate such failure modes, we investigate inference-time scaling as a complementary strategy (Team et al., 2025; Guo et al., 2025; Jaech et al., 2024; OpenAI, 2025c; Yang et al., 2025a). Specifically, we employ reasoning model and prompt it to engage in long chain-of-thought (CoT) decomposition prior to code generation (Yang et al., 2025a; Guo et al., 2025). This “thinking mode” encourages the model to reflect on the instruction and internalize task constraints before producing output. We find that such reflective generation significantly reduces reward hacking and leads to improved output quality across both visual fidelity and instruction compliance.

In parallel with these technical contributions, we also address a data limitation in this domain. Existing benchmarks are either absent or insufficiently tailored to capture the nuances of symbolic graphics generation. To fill this gap, we construct a high-quality benchmark and a curated training dataset using a scalable pipeline. Our process ensures both structural diversity and semantic alignment between code, rendered images, and textual instructions, supporting reliable evaluation and training.

In summary, we present a unified framework for symbolic graphics code generation that integrates perceptual feedback and inference-time scaling. Applied to open-source models such as Qwen3 (Yang et al., 2025a), we substantially narrows the performance gap with state-of-the-art proprietary models. Beyond demonstrating empirical gains, this work contributes scalable methodologies and resources that advance the broader goal of aligning language models with visually grounded semantics.

108 **2 TASK FORMULATION**
 109

110 Let \mathcal{Q} be the space of natural language query (instructions), and let \mathcal{O} denote the set of syntactically
 111 valid SVG code. The text-to-SVG task is to learn a mapping $\pi_\theta : \mathcal{Q} \rightarrow \mathcal{O}$, parameterized by θ , that
 112 generates an SVG program $o \in \mathcal{O}$ given an input instruction $q \in \mathcal{Q}$. Each SVG code o is rendered
 113 into an image $I \in \mathcal{I}$ via a deterministic function $\text{render} : \mathcal{O} \rightarrow \mathcal{I}$. The goal is to learn a policy π_θ
 114 that produces code o whose rendered image $I = \text{render}(o)$ faithfully reflects the visual semantics
 115 described by q . Due to the inherent ambiguity of the task, multiple distinct codes may yield visually
 116 identical or semantically equivalent images. As a result, evaluation is conducted primarily in the
 117 visual domain, rather than based on token-level similarity between codes. This task is motivated by
 118 real-world applications such as web design, UI/UX, and digital publishing, where SVGs provide a
 119 compact and resolution-independent representation of icons, charts, and other vector graphics.
 120

121 **3 METHOD**
 122

123 **3.1 PRELIMINARY**
 124

125 **Group Relative Policy Optimization (GRPO)** Reinforcement learning excels at code and mathematical
 126 reasoning tasks (Guo et al., 2025; Jaech et al., 2024; Yang et al., 2025a), where programmatic
 127 oracles (e.g., unit tests, symbolic solvers) provide explicit correctness signals. Text-to-SVG genera-
 128 tion lacks such ground truth: many syntactically different codes render the same image, and textual
 129 metrics correlate poorly with visual fidelity. Unlike actor-critic methods such as PPO (Schulman
 130 et al., 2017) that learn a separate critic, GRPO (Shao et al., 2024) eliminates the need for a critic by
 131 comparing a group of completions directly. For a query q , the behavior policy $\pi_{\theta_{\text{old}}}$ samples a group
 132 of G completions $\{o_i\}_{i=1}^G$. Each completion o_i receives a task-specific reward $R_i = R(q, o_i)$. Its
 133 group-relative advantage is then computed via z-score normalization:
 134

$$\hat{A}_i = \frac{R_i - \text{mean}(\{R_j\}_{j=1}^G)}{\text{std}(\{R_j\}_{j=1}^G)}. \quad (1)$$

135 Policy improvement is performed at the token level. Let $o_{i,t}$ denote the t -th token of code o_i . GRPO
 136 maximizes a clipped surrogate objective with a KL regularization term:
 137

$$\begin{aligned} J_{\text{GRPO}}(\theta) = & \mathbb{E}_{(q,a) \sim \mathcal{D}} \left[\mathbb{E}_{\{o_i\}_{i=1}^G \sim \pi_{\theta_{\text{old}}}(\cdot | q)} \left[\frac{1}{G} \sum_{i=1}^G \frac{1}{|o_i|} \sum_{t=1}^{|o_i|} \min \left(r_{i,t}(\theta) \hat{A}_i, \text{clip}(r_{i,t}(\theta), 1 - \varepsilon, 1 + \varepsilon) \hat{A}_i \right) \right] \right. \\ & \left. - \beta D_{\text{KL}}(\pi_\theta \parallel \pi_{\text{ref}}) \right], \end{aligned} \quad (2)$$

138 where the importance ratio $r_{i,t}(\theta)$ is defined as:
 139

$$r_{i,t}(\theta) = \frac{\pi_\theta(o_{i,t} \mid q, o_{i,<t})}{\pi_{\theta_{\text{old}}}(o_{i,t} \mid q, o_{i,<t})}. \quad (3)$$

140 Here, ε is the clipping threshold, β controls the strength of the KL regularization toward a fixed
 141 reference policy π_{ref} , and expectations are taken first over tokens, then averaged over the G sampled
 142 codes. GRPO thus optimizes the policy toward completions that outperform others.
 143

144 **3.2 REINFORCEMENT LEARNING WITH VISUAL FEEDBACK**
 145

146 Reinforcement learning excels in code and math reasoning because deterministic oracles supply
 147 explicit correctness signals (Guo et al., 2025). However, text-to-SVG generation lacks such ground-
 148 truth supervision. A visually accurate SVG may have multiple structurally distinct implementations,
 149 and textual similarity metrics often fail to reflect perceived visual fidelity. We overcome this gap with
 150 a frozen state-of-the-art vision-language model (VLM). At each training step, the model generates
 151 SVG code, we render it, pair the image with the prompt, and query VLM for a score. This dense
 152 scalar reward proxies human preference, pushing the model toward syntactically valid, semantically
 153 faithful, and visually coherent graphics.
 154

Prompt for Visual-Langauge Model

Please help me evaluate SVG images against specified instructions and a reference image through three major assessment areas. Each area is scored independently, with scores summed for a final rating.

- **First Image:** the generated image to be evaluated
- **Second Image:** the reference “ground-truth” image

Here's the specified instructions for SVG code writing:

{THE_SVG_INSTRUCTION}

1. Object and Text Accuracy (0-3 points)

Criteria: Object Presence, Object Completeness, Shape Accuracy, Text Accuracy, Typography

- **0 points:** Significant deviation from requirements, critical objects missing or severely distorted
- **1 points:** [Base GOOD score] All required objects present and identifiable, though may have minor flaws in shape or execution
- **2 points:** [Perfect shapes] All objects match specified shapes, with correct proportions and proper sizing relative to each other
- **3 points:** [Outstanding accuracy] Perfect shaping with precise edges, and perfect text implementation matching reference exactly

2. Positioning and Stroke Precision (0-4 points)

Criteria: Relative Positioning, Size and Proportion, Stroke Accuracy, Clean Layout, Viewbox Utilization

- **0 points:** Completely incorrect layout or missing stroke elements
- **1 points:** Significant layout issues, problematic stroke implementations or poor positioning with major overlapping issues
- **2 points:** [Base GOOD score] Objects positioned correctly with proper spacing and appropriate stroke weights
- **3 points:** [Excellent positioning] Perfect layout matching reference image with precise spacing and optimal viewbox utilization
- **4 points:** [Masterful execution] Perfect positioning that match reference image with exceptional accuracy down to the pixel level

3. Color and Overall Quality (0-3 points)

Criteria: Color Matching, Opacity/Transparency, Rendering Quality, Detail Precision, Overall Impression

- **0 points:** Incorrect colors, severe rendering failures or major quality problems
- **1 points:** [Base GOOD score] Colors match specifications, rendering is clean with no artifacts
- **2 points:** [Perfect coloring] Exact color matching to reference image, with appropriate use of opacity/transparency if specified
- **3 points:** [Outstanding quality] Perfect color implementation that perfectly matches or exceeds reference image

Evaluation Guideline

- Reference the second image (the reference “ground-truth” image) when assessing
- Always compare to the reference when assigning higher scores than the [Base GOOD score]
- If an element does not match the reference image, it must be noted and reflected in the score

Please write your evaluation in the following format:

```
```xml
<comparison_summary>
...brief overall comparison between the generated image and the reference image...
</comparison_summary>
<object_text_accuracy> <review>...</review> <score>...a integer...</score>
</object_text_accuracy>
<positioning_stroke> <review>...</review> <score>...</score> </positioning_stroke>
<color_overall> <review>...</review> <score>...</score> </color_overall>
<final_score>...a integer...</final_score>
```

```

Figure 2: Prompt used for the VLM judge, specifying evaluation criteria and XML output format.

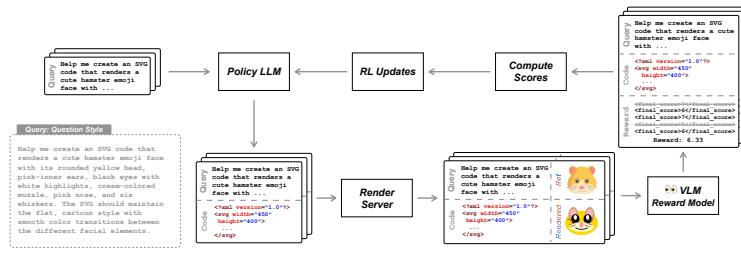


Figure 3: Overview of reinforcement learning with visual feedback for text-to-SVG. The policy model generates SVG based on a textual query, which is then rendered into an image and evaluated by a vision-language model (VLM). The reward provides visual feedback to guide policy updates.

VLM-as-a-Judge The evaluation prompt used for reinforcement learning is designed based on three key principles: **(1) Comprehensive coverage.** We decompose “visual quality” into three independent dimensions critical to SVG graphics: (i) *object and text correctness*, (ii) *geometric alignment and stroke accuracy*, and (iii) *color fidelity and overall polish*. Each dimension is supported by a checklist, ensuring that every visible flaw maps to a concrete criterion. **(2) Fine-grained scoring.** Each dimension is rated on a short integer scale with human-interpretable levels (e.g., adequate, perfect, excellent). The overall score ranges from 0 to 10, with each increment representing a meaningful visual improvement. This dense feedback provides smooth reward gradients and supports stable policy updates. **(3) Machine-readable output.** The judge returns its decision in a fixed XML format that contains both the written assessments and the numeric scores. The consistent output structure allows direct extraction of rewards during training, eliminating the need for custom parsing and ensuring deterministic behavior across runs.

216 **Mitigating Reward Hacking via Inference-Time Scaling**

217 **Instruction:**
218 Help me create an SVG code that renders a simple house with an orange trapezoid roof, white walls, a set of four steps on the left side
219 with purple accents, a green base, and some dashed curved lines on the right side to represent wind or rain.
220 **Response:**
<think>
221 Okay, I need to create an SVG of a simple house based on the instruction. Let's start by breaking down the components. First, the house has **an orange trapezoid roof**. A trapezoid has two parallel sides, so **I'll need to define a polygon or use a path** ... So the walls are white. Let's plan the structure.
222 Let's say the house has a square base, **like a rectangle for the walls** ... So for the roof: trapezoid. **Let's make it large enough** ...
</think>
223 <xml
224 <svg viewBox="0 0 300 300" xmlns="http://www.w3.org/2000/svg">
225 <!-- Orange trapezoid roof -->
226 <polygon points="80,100 220,100 250,150 50,150" fill="#FF8C00"/>
227 <!-- White walls -->
228 <rect x="50" y="150" width="200" height="150" fill="#FFFFFF"/>
229 <!-- Green base -->
230 <rect x="50" y="150" width="200" height="150" fill="#008000"/>
231 <!-- Purple steps -->
232 <rect x="50" y="150" width="200" height="150" fill="#800080"/>
233 <!-- Dashed curved lines -->
234 <rect x="50" y="150" width="200" height="150" fill="#000000"/>
235 </svg>

229 Figure 5: An example of inference-time scaling using Qwen3 with a `<think>` directive, where the
230 model reflects step by step before generating SVG.
231

232 These designs yield a reward function that is *comprehensive*, *modifiable*, and *reproducible*, which are
233 essential properties for reliable reinforcement learning in text-to-SVG generation.
234

235 **Reward Score** Each training sample consists of a textual instruction q and its reference SVG code
236 o^* . The policy π_θ generates a candidate o . Both are rendered into images:

$$I^* = \text{render}(o^*), \quad I = \text{render}(o). \quad (4)$$

237 The tuple (q, I, I^*) , together with the prompt (Figure 2), is passed to a frozen vision-language model,
238 which returns an integer score $r \in [0, 10]$, labelled as `<final_score></final_score>` in the
239 XML output. To reduce the effect of randomness in VLM outputs, we query the VLM five times
240 ($\mathcal{R} = \{r_1, \dots, r_5\}$), drop the highest and lowest scores, and average the rest to define the reward:
241

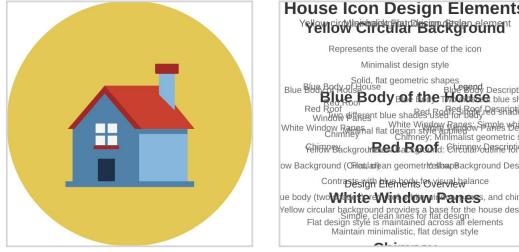
$$R_{\text{vlm}}(q, o) = \text{mean}(\mathcal{R} \setminus \{\min \mathcal{R}, \max \mathcal{R}\}). \quad (5)$$

245 3.3 MITIGATING REWARD HACKING VIA INFERENCE-TIME SCALING

246 **Reward Hacking** While reinforcement learning based on visual feedback can improve text-
247 to-SVG performance in the early stages of training, we observe that it often leads to reward
248 hacking in later stages (Weng, 2024; Di Langosco et al., 2022), which results in degradation
249 of output quality (Pan et al., 2022; Skalse et al., 2022). A prominent failure mode occurs when
250 the model generates SVGs that include textual
251 annotations describing visual elements, such as
252 “Red Roof” or “Blue Body of the House”, in-
253 stead of rendering the corresponding graphics.
254 These outputs frequently receive high reward
255 scores from the vision-language model, even
256 though they fail to capture the intended visual
257 semantics.

258 This behavior stems from biases in the reward signal introduced by vision-language models, many of
259 which are pretrained with grounding-based supervision (Wang et al., 2024a; Bai et al., 2025). In such
260 settings, both textual overlays and object-level features contribute similarly during alignment. As a
261 result, rendering relevant words within the SVG becomes a shortcut for maximizing reward, even
262 when the output fails to reflect the intended graphical content.

263 **Inference-time Scaling** Accordingly, we choose a thinking-enabled model as the RL policy back-
264 bone, using inference-time reasoning to avoid reward hacking and promote visually grounded, seman-
265 tically coherent outputs. Figure 5 shows our implementation with Qwen3, whose explicit `<think>`



266 Figure 4: A case of reward hacking where the
267 model embeds descriptive text like “Red Roof” in-
268 stead of rendering the intended visual concept.
269

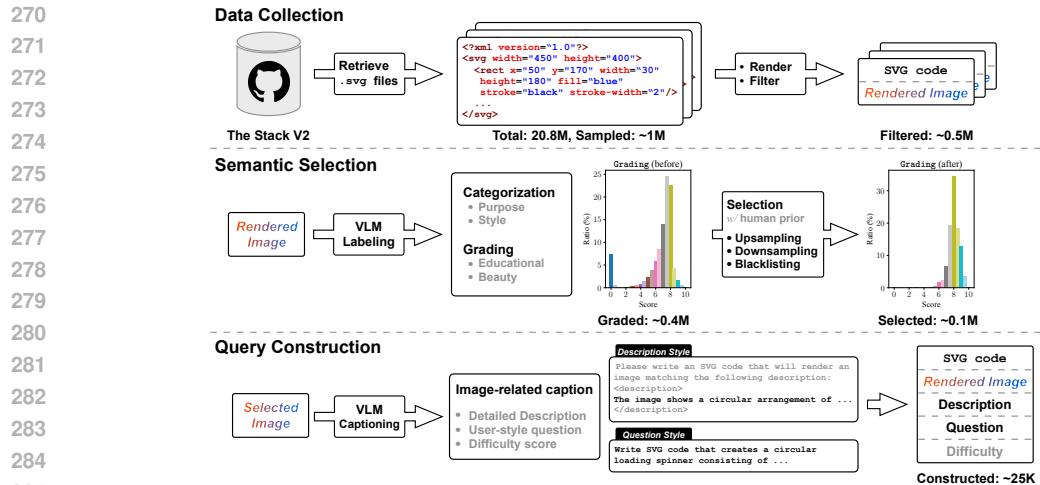


Figure 6: Data pipeline for constructing text-to-SVG training examples.

directive triggers a chain-of-thought (CoT, Wei et al. (2022)) before SVG synthesis. Reflection lets the model parse the instruction, decompose the visual specification, and embed implicit structure and task-specific constraints. For example, it reasons about each SVG component before coding, reducing reliance on reward shortcuts and producing outputs that remain both visually grounded and semantically faithful.

3.4 DATA SYNTHESIS PIPELINE

Manual text-to-SVG annotation is prohibitively expensive, demanding design expertise, coding skills, and pixel-level alignment. We instead build a synthetic pipeline that programmatically converts raw SVGs into large, diverse, high-quality instruction–code pairs with minimal human effort (Figure 6). The pipeline has three stages: data collection, semantic selection, and query construction.

Data Collection. We begin by retrieving all SVG files from The Stack V2 (Lozhkov et al., 2024), a permissively licensed code dataset. We randomly sample around 1 million SVG files out of 20.8 million. Files that fail to render, produce blank or corrupt outputs, or have invalid dimensions or extreme aspect ratios are filtered out, resulting in approximately 0.5 million valid code-image pairs.

Semantic Selection. We then apply a frozen vision-language model (VLM) to each rendered image. The model performs two core tasks: (i) semantic categorization, which assigns coarse labels for *purpose* (e.g., UI, chart, icon) and *style* (e.g., flat, line, modern); and (ii) scalar grading, which evaluates each image’s educational clarity and visual appeal on a [0–10] scale. From approximately 400K graded samples, we apply a human-informed selection strategy based on labels’ combinations. This process yields a curated set of roughly 100K examples with improved quality and diversity.

Query Construction. For each selected image, we generate textual instructions using VLM-based captioning. We synthesize two complementary prompt styles: (i) a **description-style** instruction providing a detailed natural language account of the image, and (ii) a **question-style** instruction emulating typical user queries (e.g., “Create an SVG of...”). A difficulty score is assigned to each sample based on structural and visual complexity. The final dataset contains around 25K instances, each consisting of the SVG code, rendered image, associated textual instructions, and difficulty score, forming tuples of the form \langle Instruction, SVG Code, Image, Difficulty \rangle .

By synthesizing data grounded in both symbolic structure and visual semantics, our pipeline provides scalable, controllable supervision for training LLMs via reinforcement learning in text-to-SVG tasks.

324
 325 Table 1: Comparison of closed-source and open-source on the text-to-SVG benchmark. Qwen3-8B
 326 with proposed RL method achieves competitive performance.

| 327
328
329
330
331
332
333
334
335 | 336
337
338
339
340
341
342
343
344
345
346
347
348
349
350 | 351
352
353
354
355
356
357
358
359
360
361
362
363
364
365
366
367
368
369
370
371
372
373
374
375
376
377 | 351
352
353
354
355
356
357
358
359
360
361
362
363
364
365
366
367
368
369
370
371
372
373
374
375
376
377 | 351
352
353
354
355
356
357
358
359
360
361
362
363
364
365
366
367
368
369
370
371
372
373
374
375
376
377 | | | | 351
352
353
354
355
356
357
358
359
360
361
362
363
364
365
366
367
368
369
370
371
372
373
374
375
376
377 | | | |
|---|---|---|---|---|---|---|---|---|-------------|--|--|
| | | | | 351
352
353
354
355
356
357
358
359
360
361
362
363
364
365
366
367
368
369
370
371
372
373
374
375
376
377 | 351
352
353
354
355
356
357
358
359
360
361
362
363
364
365
366
367
368
369
370
371
372
373
374
375
376
377 | 351
352
353
354
355
356
357
358
359
360
361
362
363
364
365
366
367
368
369
370
371
372
373
374
375
376
377 | 351
352
353
354
355
356
357
358
359
360
361
362
363
364
365
366
367
368
369
370
371
372
373
374
375
376
377 | 351
352
353
354
355
356
357
358
359
360
361
362
363
364
365
366
367
368
369
370
371
372
373
374
375
376
377 | | | |
| <i>Closed-Source LLMs</i> | | | | | | | | | | | |
| Claude-4-Sonnet | ✓ | 5.87 | 5.61 | 4.95 | 5.69 | - | - | - | - | | |
| Claude-4-Sonnet | ✓ | 5.78 | 5.88 | 4.64 | 5.73 | - | - | - | - | | |
| Claude-3.7-Sonnet | ✓ | 6.08 | 5.80 | 4.21 | 5.81 | 5.67 | 5.23 | 4.08 | 5.36 | | |
| Claude-3.7-Sonnet | ✓ | 5.89 | 5.71 | 3.59 | 5.64 | 5.76 | 5.48 | 4.28 | 5.53 | | |
| Claude-3.5-Sonnet | | 5.86 | 5.26 | 3.46 | 5.41 | 5.56 | 4.87 | 2.97 | 5.07 | | |
| GPT-4.1 | | 5.86 | 5.42 | 5.15 | 5.62 | - | - | - | - | | |
| GPT-4.5-Preview | | 5.90 | 5.77 | 4.13 | 5.71 | 5.63 | 5.23 | 4.31 | 5.35 | | |
| ChatGPT-4o | | 5.73 | 5.70 | 4.56 | 5.62 | 5.59 | 5.25 | 3.79 | 5.31 | | |
| GPT-4o-mini | | 5.02 | 4.79 | 3.87 | 4.83 | 4.73 | 4.01 | 3.46 | 4.33 | | |
| <i>Open-Source LLMs</i> | | | | | | | | | | | |
| R1-Distill-Llama-70B | ✓ | 4.47 | 3.78 | 2.20 | 4.00 | 4.37 | 3.74 | 2.28 | 3.94 | | |
| R1-Distill-Qwen-32B | ✓ | 4.35 | 3.85 | 2.23 | 3.97 | 4.33 | 3.44 | 2.02 | 3.77 | | |
| R1-Distill-Qwen-14B | ✓ | 3.99 | 3.17 | 1.65 | 3.46 | 3.84 | 3.11 | 1.76 | 3.37 | | |
| R1-Distill-Qwen-7B | ✓ | 1.71 | 1.26 | 0.51 | 1.43 | 1.60 | 1.19 | 0.54 | 1.34 | | |
| Llama-4-Maverick | | 5.13 | 4.72 | 3.37 | 4.82 | 4.73 | 4.25 | 2.70 | 4.37 | | |
| Llama-4-Scout | | 4.37 | 3.99 | 2.91 | 4.10 | 4.46 | 3.79 | 2.28 | 4.01 | | |
| Llama-3.1-70B | | 4.70 | 4.09 | 2.57 | 4.28 | 4.48 | 3.88 | 2.03 | 4.03 | | |
| Llama-3.1-8B | | 3.31 | 2.64 | 1.58 | 2.89 | 3.20 | 2.62 | 1.09 | 2.79 | | |
| Qwen2.5-Coder-32B | | 4.82 | 4.43 | 2.78 | 4.49 | 4.66 | 4.07 | 2.44 | 4.24 | | |
| Qwen2.5-Coder-14B | | 4.43 | 3.73 | 2.38 | 3.97 | 4.33 | 3.63 | 2.03 | 3.85 | | |
| Qwen2.5-Coder-7B | | 3.98 | 3.24 | 1.75 | 3.50 | 3.68 | 3.10 | 1.61 | 3.27 | | |
| Qwen3-235B-A22B | | 5.40 | 5.06 | 3.55 | 5.11 | 5.28 | 4.64 | 2.94 | 4.83 | | |
| Qwen3-235B-A22B | ✓ | 5.28 | 5.18 | 3.52 | 5.10 | 5.18 | 4.70 | 3.34 | 4.83 | | |
| Qwen3-32B | | 5.13 | 4.69 | 2.67 | 4.75 | 4.91 | 4.36 | 2.66 | 4.50 | | |
| Qwen3-32B | ✓ | 5.03 | 4.93 | 3.40 | 4.86 | 5.04 | 4.63 | 3.06 | 4.71 | | |
| Qwen3-30B-A3B | | 4.95 | 4.66 | 2.53 | 4.63 | 4.75 | 4.15 | 2.46 | 4.32 | | |
| Qwen3-30B-A3B | ✓ | 4.98 | 4.83 | 3.48 | 4.80 | 4.80 | 4.44 | 2.97 | 4.50 | | |
| Qwen3-14B | | 4.96 | 4.49 | 2.84 | 4.60 | 4.85 | 4.17 | 2.17 | 4.35 | | |
| Qwen3-14B | ✓ | 4.93 | 4.75 | 3.35 | 4.73 | 4.90 | 4.37 | 2.58 | 4.49 | | |
| Qwen3-8B | | 4.63 | 4.09 | 2.79 | 4.26 | 4.36 | 3.66 | 2.11 | 3.89 | | |
| *Qwen3-8B | ✓ | 4.78 | 4.51 | 3.12 | 4.54 | 4.80 | 4.15 | 2.34 | 4.33 | | |
| *Qwen3-8B w/ RL | ✓ | 5.79 | 5.48 | 4.15 | 5.53 | 5.58 | 5.24 | 3.72 | 5.29 | | |

4 EXPERIMENTS

4.1 EVALUATION

Benchmark While emerging datasets like SVG-Bench¹ provide initial resources, the field still lacks a standardized evaluation suite. To address this, we curate a high-quality test set of 300 text-to-SVG instances disjoint from training data, following the filtering pipeline in Section 3.4. After iterative multi-dimensional validation, we retain **164** verified samples. To support detailed performance analysis, we categorize the test set into three difficulty levels: *Easy* (82), *Medium* (69), and *Hard* (13). The difficulty annotations are inherited from the scalar grading stage of our data pipeline, where a vision-language model assigns a complexity score to each image. These automatic scores are further verified. There’s two prompt styles for each instance to reflect different forms of user intent: (1) a natural *Description*-style prompt and (2) a *Question*-style prompt.

Compared Models We conducted evaluations using a diverse set of models on our text-to-SVG benchmark. For proprietary closed-source models, we selected the largest variants of the Claude 3.5, 3.7 and 4 Sonnet series (Anthropic, 2023; 2025a;b), with the 3.7 and 4 generation supporting extended reasoning capabilities (referred to as “thinking” models). Additionally, we included OpenAI’s GPT-4.1, GPT-4.5-Preview, GPT-4o and GPT-4o-mini (OpenAI, 2025a;b; Hurst et al., 2024) in our assessment. For open-source models, we first evaluated several variants distilled from DeepSeek-R1 (Guo et al., 2025) across different sizes. We also considered Llama 3.1 and the latest Llama 4 series (Grattafiori et al., 2024; AI@Meta, 2024; 2025). Within the Qwen family, we selected the Qwen 2.5 Coder series (Hui et al., 2024) and the Qwen3 series (Yang et al., 2025a) featuring Hybrid Thinking Modes, where the “Hybrid” refers to the ability to control the model’s depth of reasoning via prompt-based instructions.

¹Available at: <https://github.com/johnbean393/SVGBench>

| 378 | Instruction | 379 | Reference | 380 | Qwen3-8B w/ RL | 381 | Qwen3 8B | 382 | Claude-3.7 Sonnet | 383 | Qwen3-235B-A22B |
|-----|--|-----|-----------|-----|----------------|-----|----------|-----|-------------------|-----|-----------------|
| 384 | Write SVG code that creates this letter H logo with two blue eye-like circles, the two blue circular elements have centered black dots functioning as pupils, and the entire design is contained within a perfect black circular border. | | | | | | | | | | |
| 385 | Write SVG code that creates this mountain landscape scene with four evergreen trees in the foreground and layered teal mountains with snow caps in the background. | | | | | | | | | | |
| 386 | Generate SVG code for a necktie symbol with light blue outlines, semi-transparent overlapping triangles at the top, and a long pointed rectangle at the bottom. Use pink and light blue with transparency to achieve a soft gradient and layered look. | | | | | | | | | | |

Figure 8: Case study comparing text-to-SVG outputs across models.

4.2 TRAINING DETAILS

We conduct RL experiments on Qwen3-8B, which supports “Hybrid Thinking Modes” (Yang et al., 2025a), allowing us to investigate the impact of enabling the *Thinking Mode* within the same model. To provide the perceptual reward signal essential for our RL, we employ Qwen2.5-VL 72B (Bai et al., 2025) as the frozen vision-language model judge. For training runs without the *Thinking Mode* (i.e., standard instruction model), we set the maximum generation length to 8,192 tokens. In contrast, for experiments with *Thinking Mode* enabled, the maximum generation length was extended to 16,384 tokens. Each RL training step involved inference over a batch of 256 queries, with 8 rollouts per sample. We set the GRPO mini-batch size to 32 and set clipping parameter to 0.2 (Schulman et al., 2017). The learning rate was fixed at 1.5e-6. All other training details followed the standard GRPO algorithm (Shao et al., 2024). All experiments were conducted on 16 NVIDIA A800 80GB GPUs, with a total runtime of approximately 300 hours.

4.3 RESULTS

Main Results Table 1 shows that closed-source models outperform open-source models on the text-to-SVG benchmark. Within the open-source Qwen family, performance improves with model size, with Qwen3-32B yielding higher scores than 14B and 8B. Moreover, enabling the thinking mode in Qwen models, particularly in smaller ones, brings consistent gains. Notably, incorporating our proposed reinforcement learning method with visual feedback into Qwen3-8B leads to a substantial improvement. The RL-enhanced model achieves an overall score of 5.29, significantly outperforming its supervised counterpart (4.33). This result demonstrates that proposed method can effectively bridge the performance gap, enabling open-source models to approach the capabilities of stronger proprietary models in text-to-SVG tasks. A qualitative comparison in Figure 8 further illustrates the improvement, where RL-trained outputs show greater semantic fidelity and visual completeness.

Findings To better understand the behavior of our reinforcement learning method, we monitor two key metrics throughout training: *response length* and *entropy loss*. The results for Qwen3-8B under both *chat* and *thinking* modes are shown in Figure 7. We observe that when trained with proposed method in thinking mode, the model exhibits a steady increase in response length. This suggests that the policy is actively encouraged to perform inference-time scaling during training. At the same time, the entropy loss also increases consistently, indicating a growing degree of exploration. This enables the model to sample potentially higher-reward SVG candidates, further improving learning dynamics.

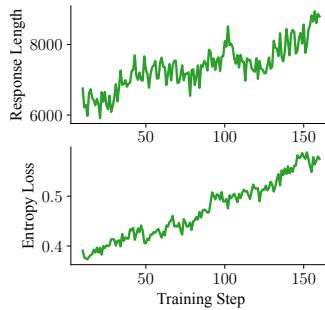


Figure 7: Trends of response length and entropy loss during training with Qwen3-8B in thinking mode.

432 4.4 ABLATION STUDY
433

434 As discussed earlier, Qwen3-8B exhibits reward hacking
435 during training. We investigate whether this can be
436 mitigated by **(a)** scaling model size and **(b)** using a code-
437 specialized model. Thus, we conduct training on Qwen2.5-
438 Coder-32B-Instruct, which satisfies both conditions. Fig-
439 ure 9 shows its training reward curve: reward hacking still
440 occurs, though its onset is delayed compared to Qwen3-
441 8B, suggesting that scaling and specialization slow down
442 hacking but do not eliminate it. In our experiments, only
443 enabling the thinking mode consistently prevents reward
444 hacking, highlighting the critical role of inference-time
445 scaling in preserving reward integrity.

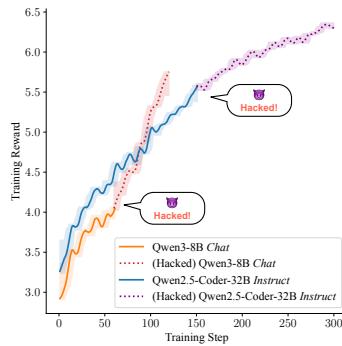
446 5 RELATED WORK
447

448 Recent efforts have explored enhancing LLMs’ capabilities in SVG understanding and generation.
449 LLM4SVG (Xing et al., 2024) proposed a modular architecture combining semantic tagging and
450 vector encoders, supported by a 580K-sample dataset. Chat2SVG (Wu et al., 2024a) combined
451 LLMs with diffusion models to generate visually expressive SVGs from text. StarVector (Rodriguez
452 et al., 2023) fused vision encoders and CodeLLMs to generate SVGs. Follow-up study (Rodriguez
453 et al., 2025) have also applied reinforcement learning with rendering feedback, but typically rely
454 on VLM-based image reconstruction and holistic visual similarity for rewards. Several works have
455 also investigated SVG as an evaluation or reasoning medium. SVGEditBench (Nishina & Matsui,
456 2024) introduced a benchmark for assessing LLMs on structured SVG editing tasks. Cai et al. (2023)
457 demonstrated that LLMs can perform vision-language reasoning via SVG-based representations.
458 However, a common limitation of these reward mechanisms (e.g., FID (Theis et al., 2015), CLIP
459 Score (Radford et al., 2021), FID-CLIP (Wu et al., 2023) or image reconstruction) is their struggle
460 with fine-grained fidelity, as they are ill-suited for a task where a single prompt can have multiple
461 valid visual outputs. In contrast, our method introduces a generative reward model that produces a
462 semantic checklist, enabling fine-grained feedback through RL.

463 6 DISCUSSION
464

465 **Limitation and future directions** This work has several limitations. **First**, due to resource
466 constraints, we only applied our method to Qwen-8B, as larger models require significantly more
467 compute. **Second**, regarding reward hacking, our work explores a different path from directly
468 engineering the reward function. While methods like using perceptual metrics or masking textual
469 artifacts are valid strategies, we focused on improving the model’s intrinsic reasoning process. Our
470 findings suggest that the “thinking mode” acts as an effective regularizer, compelling the model to
471 follow instructions more faithfully rather than exploiting reward signals. **Third**, our evaluation relied
472 on a single VLM as the judge. To validate that our improvements are robust and not an artifact of
473 overfitting to a specific judge’s biases, we plan to re-evaluate key model outputs using a distinct,
474 powerful VLM (e.g., (InternVL-3 and Qwen3-VL) and report the agreement. **Fourth**, the model
475 still struggles with prompts involving complex spatial reasoning, such as 3D shapes, likely due to
476 insufficient visual feedback. Building on these points, our future work will focus on applying visual
477 feedback–driven reinforcement learning to broader symbolic code domains, such as front-end web
478 generation, while employing a more rigorous multi-judge evaluation framework.

479 **Conclusion** In this paper, we introduce a reinforcement learning framework for symbolic graphics
480 code generation, with a focus on text-to-SVG tasks. By leveraging vision-language models as visual
481 reward models, we align model outputs with perceptual semantics. To address reward hacking,
482 we introduce inference-time scaling with thinking-enabled policies, encouraging the generation of
483 visually grounded and semantically faithful code. Furthermore, we construct a high-quality training
484 dataset through a VLM-guided selection and captioning pipeline. These components significantly
485 narrow the performance gap between open-source and proprietary models, and establish a foundation
for scalable, perceptually aligned code generation.



446 Figure 9: Training reward curves for
447 Qwen3-8B and Qwen2.5-Coder-32B-
448 Instruct.

486 REPRODUCIBILITY STATEMENT
487

488 We have taken several steps to support the reproducibility of our work. The construction of our
489 datasets, the design of evaluation prompts, and the overall experimental setup are detailed in Section 3,
490 along with Figures 2, 3, and 6. Section 4 specifies the models used and reports the corresponding
491 evaluation scores. To further facilitate replication, we provide a zipped supplementary material,
492 which include: (1) all prompts used in our experiments, (2) the full evaluation datasets, (3) end-to-end
493 evaluation code described in the paper, and (4) a step-by-step README that guides users to conduct
494 evaluations, which can also be integrated into any standard RL framework.

496 REFERENCES
497

498 AI@Meta. Llama 3.1 model card. 2024. URL https://www.llama.com/docs/model-cards-and-prompt-formats/llama3_1/.

500 AI@Meta. Llama 4 model card. 2025. URL <https://www.llama.com/docs/model-cards-and-prompt-formats/llama4/>.

502 Anthropic. The claude 3 model family: Opus, sonnet, haiku. Technical report, Anthropic, 2023.

504 Anthropic. Claude 3.7 sonnet system card. Technical report, Anthropic, 2025a.

506 Anthropic. System card: Claude opus 4 & claude sonnet 4. Technical report, Anthropic, 2025b.

507 Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan,
508 Ellen Jiang, Carrie Cai, Michael Terry, Quoc Le, et al. Program synthesis with large language
509 models. *arXiv preprint arXiv:2108.07732*, 2021.

510 Shuai Bai, Kefin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibo Song, Kai Dang, Peng Wang,
511 Shijie Wang, Jun Tang, et al. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*,
512 2025.

514 Kate Baumli, Satinder Baveja, Feryal Behbahani, Harris Chan, Gheorghe Comanici, Sebastian
515 Flennherag, Maxime Gazeau, Kristian Holsheimer, Dan Horgan, Michael Laskin, et al. Vision-
516 language models as a source of rewards. *arXiv preprint arXiv:2312.09187*, 2023.

517 Mu Cai, Zeyi Huang, Yuheng Li, Utkarsh Ojha, Haohan Wang, and Yong Jae Lee. Leveraging
518 large language models for scalable vector graphics-driven image understanding. *arXiv preprint
519 arXiv:2306.06094*, 2023.

521 Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde De Oliveira Pinto, Jared
522 Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. Evaluating large
523 language models trained on code. *arXiv preprint arXiv:2107.03374*, 2021.

524 Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser,
525 Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. Training verifiers to solve
526 math word problems. *arXiv preprint arXiv:2110.14168*, 2021.

528 Lauro Langosco Di Langosco, Jack Koch, Lee D Sharkey, Jacob Pfau, and David Krueger. Goal
529 misgeneralization in deep reinforcement learning. In *International Conference on Machine
530 Learning*, pp. 12004–12019. PMLR, 2022.

531 Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad
532 Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. The llama 3 herd of
533 models. *arXiv preprint arXiv:2407.21783*, 2024.

534 Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu,
535 Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms
536 via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025.

538 Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and
539 Jacob Steinhardt. Measuring massive multitask language understanding. *arXiv preprint
arXiv:2009.03300*, 2020.

540 Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song,
 541 and Jacob Steinhardt. Measuring mathematical problem solving with the math dataset. *arXiv*
 542 *preprint arXiv:2103.03874*, 2021.

543 Binyuan Hui, Jian Yang, Zeyu Cui, Jiaxi Yang, Dayiheng Liu, Lei Zhang, Tianyu Liu, Jiajun Zhang,
 544 Bowen Yu, Keming Lu, et al. Qwen2. 5-coder technical report. *arXiv preprint arXiv:2409.12186*,
 545 2024.

546 Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Os-
 547 trow, Akila Welihinda, Alan Hayes, Alec Radford, et al. Gpt-4o system card. *arXiv preprint*
 548 *arXiv:2410.21276*, 2024.

549 Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden Low, Alec
 550 Helyar, Aleksander Madry, Alex Beutel, Alex Carney, et al. Openai o1 system card. *arXiv preprint*
 551 *arXiv:2412.16720*, 2024.

552 Naman Jain, King Han, Alex Gu, Wen-Ding Li, Fanjia Yan, Tianjun Zhang, Sida Wang, Armando
 553 Solar-Lezama, Koushik Sen, and Ion Stoica. Livecodebench: Holistic and contamination free
 554 evaluation of large language models for code. *arXiv preprint arXiv:2403.07974*, 2024.

555 Jiawei Liu, Chunqiu Steven Xia, Yuyao Wang, and Lingming Zhang. Is your code generated by
 556 chatgpt really correct? rigorous evaluation of large language models for code generation. *Advances*
 557 *in Neural Information Processing Systems*, 36:21558–21572, 2023.

558 Anton Lozhkov, Raymond Li, Loubna Ben Allal, Federico Cassano, Joel Lamy-Poirier, Nouamane
 559 Tazi, Ao Tang, Dmytro Pykhtar, Jiawei Liu, Yuxiang Wei, et al. Starcoder 2 and the stack v2: The
 560 next generation. *arXiv preprint arXiv:2402.19173*, 2024.

561 Haoyu Lu, Wen Liu, Bo Zhang, Bingxuan Wang, Kai Dong, Bo Liu, Jingxiang Sun, Tongzheng Ren,
 562 Zhusuo Li, Hao Yang, et al. Deepseek-vl: towards real-world vision-language understanding.
 563 *arXiv preprint arXiv:2403.05525*, 2024.

564 Kunato Nishina and Yusuke Matsui. Sveditbench: A benchmark dataset for quantitative assessment
 565 of llm’s svg editing capabilities. In *Proceedings of the IEEE/CVF Conference on Computer Vision*
 566 *and Pattern Recognition*, pp. 8142–8147, 2024.

567 OpenAI. Gpt-4v(ision) system card. Technical report, OpenAI, 2023.

568 OpenAI. Introducing GPT-4.1 in the API. Technical report, OpenAI, 2025a.

569 OpenAI. Openai gpt-4.5 system card. Technical report, OpenAI, 2025b.

570 OpenAI. Openai o3 and o4-mini system card. Technical report, OpenAI, 2025c.

571 Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong
 572 Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow
 573 instructions with human feedback. *Advances in neural information processing systems*, 35:27730–
 574 27744, 2022.

575 Alexander Pan, Kush Bhatia, and Jacob Steinhardt. The effects of reward misspecification: Mapping
 576 and mitigating misaligned models. *arXiv preprint arXiv:2201.03544*, 2022.

577 Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal,
 578 Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual
 579 models from natural language supervision. In *International conference on machine learning*, pp.
 580 8748–8763. PMLR, 2021.

581 Juan A Rodriguez, Shubham Agarwal, Issam H Laradji, Pau Rodriguez, David Vazquez, Christopher
 582 Pal, and Marco Pedersoli. Starvector: Generating scalable vector graphics code from images.
 583 *arXiv preprint arXiv:2312.11556*, 2023.

584 Juan A Rodriguez, Haotian Zhang, Abhay Puri, Aarash Feizi, Rishav Pramanik, Pascal Wichmann,
 585 Arnab Mondal, Mohammad Reza Samsami, Rabiul Awal, Perouz Taslakian, et al. Rendering-aware
 586 reinforcement learning for vector graphics generation. *arXiv preprint arXiv:2505.20793*, 2025.

594 John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy
 595 optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.

596

597 Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang,
 598 Mingchuan Zhang, YK Li, Y Wu, et al. Deepseekmath: Pushing the limits of mathematical
 599 reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024.

600 Joar Skalse, Nikolaus Howe, Dmitrii Krasheninnikov, and David Krueger. Defining and characterizing
 601 reward gaming. *Advances in Neural Information Processing Systems*, 35:9460–9471, 2022.

602 Gemini Team. Gemini 2.5: Our most intelligent ai model. Technical report, Google, 2025.

603

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

607 Gemini Team, Petko Georgiev, Ving Ian Lei, Ryan Burnell, Libin Bai, Anmol Gulati, Garrett
 608 Tanzer, Damien Vincent, Zhufeng Pan, Shibo Wang, et al. Gemini 1.5: Unlocking multimodal
 609 understanding across millions of tokens of context. *arXiv preprint arXiv:2403.05530*, 2024.

610

611 Kimi Team, Angang Du, Bofei Gao, Bowei Xing, Changjiu Jiang, Cheng Chen, Cheng Li, Chenjun
 612 Xiao, Chenzhuang Du, Chonghua Liao, et al. Kimi k1.5: Scaling reinforcement learning with
 613 llms. *arXiv preprint arXiv:2501.12599*, 2025.

614 Lucas Theis, Aäron van den Oord, and Matthias Bethge. A note on the evaluation of generative
 615 models. *arXiv preprint arXiv:1511.01844*, 2015.

616 Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu,
 617 Jialin Wang, Wenbin Ge, et al. Qwen2-vl: Enhancing vision-language model’s perception of the
 618 world at any resolution. *arXiv preprint arXiv:2409.12191*, 2024a.

619

620 Xingyao Wang, Boxuan Li, Yufan Song, Frank F Xu, Xiangru Tang, Mingchen Zhuge, Jiayi Pan,
 621 Yueqi Song, Bowen Li, Jaskirat Singh, et al. Openhands: An open platform for ai software develop-
 622 ers as generalist agents. In *The Thirteenth International Conference on Learning Representations*,
 623 2024b.

624 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny
 625 Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in
 626 neural information processing systems*, 35:24824–24837, 2022.

627 Lilian Weng. Reward hacking in reinforcement learning. *lilianweng.github.io*, Nov 2024. URL
 628 <https://lilianweng.github.io/posts/2024-11-28-reward-hacking/>.

629

630 Ronghuan Wu, Wanchao Su, Kede Ma, and Jing Liao. Iconshop: Text-guided vector icon synthesis
 631 with autoregressive transformers. *ACM Transactions on Graphics (TOG)*, 42(6):1–14, 2023.

632 Ronghuan Wu, Wanchao Su, and Jing Liao. Chat2svg: Vector graphics generation with large language
 633 models and image diffusion models. *arXiv preprint arXiv:2411.16602*, 2024a.

634

635 Zhiyu Wu, Xiaokang Chen, Zizheng Pan, Xingchao Liu, Wen Liu, Damai Dai, Huazuo Gao, Yiyang
 636 Ma, Chengyue Wu, Bingxuan Wang, et al. Deepseek-vl2: Mixture-of-experts vision-language
 637 models for advanced multimodal understanding. *arXiv preprint arXiv:2412.10302*, 2024b.

638 Ximing Xing, Juncheng Hu, Guotao Liang, Jing Zhang, Dong Xu, and Qian Yu. Empowering llms to
 639 understand and generate complex vector graphics. *arXiv preprint arXiv:2412.11102*, 2024.

640 An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang
 641 Gao, Chengan Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu,
 642 Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin
 643 Yang, Jiaxi Yang, Jing Zhou, Jingren Zhou, Junyang Lin, Kai Dang, Keqin Bao, Kexin Yang,
 644 Le Yu, Lianghao Deng, Mei Li, Mingfeng Xue, Mingze Li, Pei Zhang, Peng Wang, Qin Zhu, Rui
 645 Men, Ruize Gao, Shixuan Liu, Shuang Luo, Tianhao Li, Tianyi Tang, Wenbiao Yin, Xingzhang
 646 Ren, Xinyu Wang, Xinyu Zhang, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yinger
 647 Zhang, Yu Wan, Yuqiong Liu, Zekun Wang, Zeyu Cui, Zhenru Zhang, Zhipeng Zhou, and Zihan
 Qiu. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*, 2025a.

648 John Yang, Carlos Jimenez, Alexander Wettig, Kilian Lieret, Shunyu Yao, Karthik Narasimhan,
 649 and Ofir Press. Swe-agent: Agent-computer interfaces enable automated software engineering.
 650 *Advances in Neural Information Processing Systems*, 37:50528–50652, 2024.

651
 652 Yiyi Yang, Wei Cheng, Sijin Chen, Xianfang Zeng, Jiaxu Zhang, Liao Wang, Gang Yu, Xingjun
 653 Ma, and Yu-Gang Jiang. Omnisvg: A unified scalable vector graphics generation model. *arXiv*
 654 *preprint arXiv:2504.06263*, 2025b.

655 656 A EVALUATION OF GENERAL CAPABILITIES AFTER RL

657 A.1 GENERAL CAPABILITIES

658 A central challenge in model development is the specialization-generalization dilemma: task-specific
 659 fine-tuning, while effective for the target domain, often risks degrading a model’s broader, general-
 660 purpose capabilities. This phenomenon is a significant concern, particularly when using reinforcement
 661 learning (RL) on highly structured, synthetic data like text-to-SVG. To investigate this issue, we
 662 conducted a comprehensive evaluation of our final model, Qwen3-8B *w/ RL*, against its original base
 663 model, Qwen3-8B.

664 Our evaluation spans a suite of established text-only benchmarks designed to measure founda-
 665 tional abilities, including general knowledge (MMLU, Hendrycks et al. 2020), mathematical
 666 reasoning (GSM8K, Cobbe et al. 2021 and MATHHendrycks et al. 2021), and code generation
 667 (HumanEval, Chen et al. 2021 and MBPP, Austin et al. 2021) with their enhanced version Human-
 668 Eval+ and MBPP+ in EvalPlus (Liu et al., 2023). To ensure a fair and thorough comparison, we
 669 evaluated both models under two distinct decoding strategies: with and without “thinking mode”.
 670 The aggregated results are presented in Table 2. We can draw the following conclusions:

671
 672 Table 2: Performance comparison on general text-based benchmarks. The results demonstrate that RL
 673 fine-tuning for SVG generation does not degrade, and in some cases enhances, general capabilities.
 674 Higher is better.

| 675 Model | Thinking | MMLU | HumanEval | HumanEval+ | MBPP | MBPP+ | GSM8K | MATH |
|---------------------------|----------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|
| 676 Qwen3-8B | ✓ | 89.5 | 89.6 | 81.7 | 85.0 | 69.7 | 94.8 | 48.5 |
| 677 Qwen3-8B <i>w/ RL</i> | ✓ | 88.6 | 93.9 | 86.6 | 86.2 | 72.7 | 93.6 | 46.7 |
| 678 Qwen3-8B | | 83.8 | 82.9 | 80.5 | 69.2 | 59.4 | 88.2 | 25.0 |
| 679 Qwen3-8B <i>w/ RL</i> | | 84.5 | 82.3 | 79.9 | 71.9 | 60.9 | 91.5 | 25.0 |

680
 681 **Robust Preservation of Foundational Reasoning and Knowledge** Across benchmarks, the
 682 performance of Qwen3-8B *w/ RL* remains remarkably stable. On challenging tasks like MMLU,
 683 GSM8K, and MATH, the scores are either marginally higher or exhibit only negligible decreases
 684 (typically within a 1–2 point margin). This stability demonstrates that the specialized training for SVG
 685 generation did not impair the model’s fundamental language understanding and complex reasoning
 686 abilities.

687
 688 **Clear Evidence of Positive Transfer to Code Generation** More strikingly, the model after RL
 689 shows consistent and significant improvements across multiple code generation benchmarks, including
 690 a +4.3 point gain on HumanEval and a +4.9 point gain on HumanEval+ in the “thinking mode”. This
 691 suggests a powerful positive transfer effect. We hypothesize that the process of learning to generate
 692 structured SVG code acts as a beneficial regularizer. This training helps the model learn to follow
 693 instructions precisely, think logically, and stick to syntax constraints.

694 A.2 TARGETED VALIDATION OF CODING SKILLS

695 To further probe this observed enhancement in coding capabilities, we performed a more targeted
 696 evaluation on **LiveCodeBench** (Jain et al. 2024; versions 5 and 6), a dynamic benchmark for
 697 real-world coding challenges. We report **pass@1** and **pass@8** scores in the non-thinking mode.

702
 703 Table 3: LiveCodeBench results (non-thinking mode), reported as “pass@1 / pass@8”. The model
 704 after RL maintains or improves performance, corroborating the positive transfer effect on coding
 705 tasks.

| 706 Model | 707 LiveCodeBench (v5) | 708 LiveCodeBench (v6) |
|---------------------------|-------------------------------|-------------------------------|
| 709 Qwen3-8B | 710 24.2 / 34.1 | 711 25.9 / 34.3 |
| 712 Qwen3-8B w/ <i>RL</i> | 713 24.0 / 35.3 | 714 26.8 / 33.5 |

715 As shown in Table 3, the Qwen3-8B *w/ RL* model maintains highly competitive performance. Notably,
 716 it improves the pass@8 rate on v5 and the pass@1 rate on v6, suggesting that the RL training may
 717 encourage the model to generate a broader set of correct candidates or improve its first-attempt
 718 accuracy, depending on the contest’s nature.

719 In summary, the collective evidence from this multi-faceted evaluation is compelling. Our specialized
 720 RL training pipeline for text-to-SVG does not lead to a trade-off. Instead, it enhances the target
 721 capability *without sacrificing*—and in the important domain of code generation, even *meaningfully*
 722 *improving*—the model’s general-purpose performance. This suggests our approach yields a model
 723 that is both a master of its specialized domain and an even more capable generalist.

724 B LLM USAGE STATEMENT

725 In this work, we used a large language model as a general-purpose writing assistance tool. Specifically,
 726 the LLM was used to help refine the clarity, grammar, and overall fluency of the text during the
 727 drafting and revision stages. The model provided suggestions for rephrasing sentences, improving
 728 paragraph structure, and ensuring consistent academic tone. However, all ideas, research design,
 729 analysis and interpretation of results were made by the authors.