

TEXT-TO-IMAGE DIFFUSION MODELS CANNOT COUNT, AND PROMPT REFINEMENT CANNOT HELP

Anonymous authors

Paper under double-blind review

ABSTRACT

Generative modeling is widely regarded as one of the most essential problems in today’s AI community, with text-to-image generation having gained unprecedented real-world impacts. Among various approaches, diffusion models have achieved remarkable success and have become the de facto solution for text-to-image generation. However, despite their impressive performance, these models exhibit fundamental limitations in adhering to numerical constraints in user instructions, frequently generating images with an incorrect number of objects. While several prior works have mentioned this issue, a comprehensive and rigorous evaluation of this limitation remains lacking. To address this gap, we introduce **T2ICountBench**, a novel benchmark designed to rigorously evaluate the counting ability of state-of-the-art text-to-image diffusion models. Our benchmark encompasses a diverse set of generative models, including both open-source and private systems. It explicitly isolates counting performance from other capabilities, provides structured difficulty levels, and incorporates human evaluations to ensure high reliability. Extensive evaluations with T2ICountBench reveal that all state-of-the-art diffusion models fail to generate the correct number of objects, with accuracy dropping significantly as the number of objects increases. Additionally, an exploratory study on prompt refinement demonstrates that such simple interventions generally do not improve counting accuracy. Our findings highlight the inherent challenges in numerical understanding within diffusion models and point to promising directions for future improvements.

1 INTRODUCTION

Generative modelling is widely regarded as one of the most essential problems in today’s AI community, encompassing tasks such as natural language generation (Brown et al., 2020; Achiam et al., 2023; Liu et al., 2024), image synthesis (Donahue & Simonyan, 2019; Dhariwal & Nichol, 2021; Yang et al., 2023), video generation (Tulyakov et al., 2018; Ho et al., 2022; Singer et al., 2023), and speech synthesis (Oord et al., 2018; Radford et al., 2023; Tan et al., 2024). Among various generative approaches, Diffusion Models (DMs) have demonstrated remarkable success across multiple domains, particularly in text-to-image and text-to-video generation (Ruiz et al., 2023; Wu et al., 2023a; Yang et al., 2024c). Notable models like Diffusion Transformers (DiTs) (Peebles & Xie, 2023) and Video LDM (Blattmann et al., 2023) have been shown to produce high-resolution and realistic images and videos, forming the foundation of advanced generative AI tools, including OpenAI Sora (OpenAI, 2024) and Kling (Kuaishou, 2024).

Despite these advancements, diffusion-based models exhibit fundamental limitations in adhering to numerical constraints in user instructions. Prior empirical studies have shown that text-to-image diffusion models often struggle with basic object counting tasks (Saharia et al., 2022; Huang et al., 2023; Petsiuk et al., 2022). Specifically, when given prompts specifying an exact number of objects (e.g., “generate an image with 7 apples on a wooden table”), the generated content frequently fails to match the requested quantity. These limitations become even more pronounced in complex scenarios, such as “generate an image with 7 apples on a table, separated by 3 oranges.” Such failures raise concerns about the reliability of such generative models and highlight their inherent difficulty in following precise numerical constraints.

054 However, existing empirical studies on the counting ability of text-to-image models suffer from key
055 limitations. Many benchmark studies evaluate only a small number of possibly outdated generative
056 models (Saharia et al., 2022; Petsiuk et al., 2022), with most models dating back to 2022–2023.
057 Additionally, some benchmarks are too general and fail to disentangle counting ability from other
058 factors such as adherence to style and shape constraints (Huang et al., 2023; Peng et al., 2024;
059 Wu et al., 2024). These shortcomings necessitate the need for a comprehensive, up-to-date, and
060 specialized benchmark dedicated to evaluating the counting ability of text-to-image models.

061 To address this gap, we introduce **T2ICountBench**, a novel benchmark designed to rigorously as-
062 sess the counting ability of state-of-the-art text-to-image models in 2025. Our benchmark cov-
063 ers a diverse set of generative models, including both open-source and private image generation
064 systems (Podell et al., 2024; Baldrige et al., 2024; Yang et al., 2024b). Unlike prior works,
065 T2ICountBench explicitly isolates counting performance from other capabilities and provides struc-
066 tured difficulty levels, spanning object counts from 1 to 15. Additionally, our benchmark incorpo-
067 rates human evaluations to ensure high reliability and robustness.

068 With the proposed T2ICountBench, we conduct a comprehensive evaluation to determine whether
069 diffusion-based text-to-image models can accurately generate objects under numerical constraints.
070 Our results show that most existing models exhibit significant failures in simple counting tasks,
071 frequently generating the wrong number of objects. To highlight the non-trivial nature of this limi-
072 tation, we also explore whether simple prompt refinements—decomposing a difficult counting task
073 (e.g., generating 15 objects) into smaller subtasks—can improve performance. Our contributions
074 are summarized as follows:

- 075 • We present a comprehensive and rigorous benchmark, T2ICountBench, for evaluating the
076 counting ability of text-to-image diffusion models. This benchmark effectively exposes the
077 inherent limitations of these models in generating the exact number of objects.
- 078 • We conduct extensive ablation studies on various factors influencing counting performance,
079 including the number of objects, scene type, and style. Our findings indicate that as the
080 number of objects increases from 1 to 15, model accuracy significantly drops, reaching
081 around 10% for higher counts. We also find that complex background scenes will further
082 adversely affect counting ability.
- 083 • We performed an exploratory study to investigate whether simple prompt refinements could
084 alleviate counting limitations. Our results indicate that such refinements generally do not
085 improve counting performance, highlighting the inherent challenge of text-to-image diffu-
086 sion models in counting.

088 **Roadmap.** In Section 3, we introduce our new benchmark to evaluate the counting capability of text-
089 to-image diffusion models. In Section 4, we show the main findings from our counting benchmark.
090 In Section 5, we discuss the possibility of improving text-to-image diffusion models with prompt
091 refinement. In Section 6, we show the conclusion of this paper.

093 2 RELATED WORKS

095 **Benchmarks on Text-to-Image Generation.** The rapid advancement and real-world impact of
096 text-to-image models have driven the development of evaluation benchmarks, particularly follow-
097 ing the emergence of diffusion models. Early benchmarks (Ramesh et al., 2022; Cho et al., 2023;
098 Hu et al., 2023) primarily relied on captions sourced from well-established datasets such as MS
099 COCO, focusing on generating simple objects and scenes that could be automatically evaluated us-
100 ing pre-trained vision models. For instance, DALL-Eval (Cho et al., 2023) employs a 3D renderer to
101 generate synthetic scenes for training text-to-image models, subsequently assessing them with ob-
102 ject detection models. It also incorporates fairness considerations by evaluating social biases such as
103 gender and skin tone. [GenEval \(Ghosh et al., 2023\) as an object-focused automatic evaluation frame-
104 work that uses object detection and related vision models to assess fine-grained compositional
105 and text-to-image alignment.](#) Addressing DALL-Eval’s limited scope, TIFA (Hu et al., 2023) expands
106 evaluation criteria by leveraging a pretrained visual question-answering (VQA) model, enabling as-
107 sessments beyond synthetic captions and 3D-rendered scenes to include more diverse conditions
such as geolocation and weather variations.

108 More recent benchmarks have shifted toward evaluating advanced capabilities of text-to-image mod-
109 els. HPDv2 (Wu et al., 2023b) and Gecko (Wiles et al., 2024) incorporate human preference-based
110 ranking to assess alignment with aesthetic preferences. Another key research direction focuses on
111 compositional text-to-image generation, which involves associating arbitrary attributes with objects
112 beyond predefined datasets like COCO and reasoning about complex object relationships. Repre-
113 sentative benchmarks in this area include T2I-CompBench (Huang et al., 2023), ConceptMix (Wu
114 et al., 2024), and GenAI-Bench (Li et al., 2024a). Additionally, Commonsense-T2I (Fu et al., 2024)
115 and PhyBench (Meng et al., 2024) further extend these evaluations by incorporating real-world com-
116 monsense reasoning, such as physical constraints. Despite the progress in benchmarking various as-
117 pects of text-to-image models, ranging from basic object recognition to complex compositional and
118 commonsense reasoning, the fundamental ability of these models to accurately count objects still
119 requires a rigorous evaluation. This paper aims to address this gap through a rigorous evaluation of
120 the counting capability of state-of-the-art text-to-image models.

121
122 **Diffusion Models for Text-to-Image Generation.** As a fundamental paradigm shift in generative
123 AI, diffusion models have substantially enhanced the quality and resolution of generated images,
124 surpassing earlier approaches such as Variational Autoencoders (VAEs) (Kingma & Welling, 2014;
125 Razavi et al., 2019) and Generative Adversarial Networks (GANs) (Goodfellow et al., 2014; Xu
126 et al., 2018). Recent diffusion-based backbone models (Ho et al., 2020; Song et al., 2021b;a; Lip-
127 man et al., 2023) have achieved impressive results in high-fidelity image synthesis without control
128 conditions. However, the challenge of precisely controlling image content via language prompts has
129 motivated the development of more controllable text-to-image generation methods (Rombach et al.,
130 2022; Ramesh et al., 2022).

131 Text-to-image diffusion models can be broadly classified into two categories: pixel space mod-
132 els (Nichol et al., 2022; Saharia et al., 2022; Chen et al., 2023) and latent space models (Rombach
133 et al., 2022; Samuel et al., 2023; Podell et al., 2024). Pixel space models directly perturb image
134 pixels with noise and iteratively denoise them. For example, GLIDE (Nichol et al., 2022) adapts
135 class-conditioned diffusion models by replacing class labels with text tokens and employs both clas-
136 sifier guidance and classifier-free guidance to align images with text. Imagen (Saharia et al., 2022)
137 similarly leverages classifier-free guidance but utilizes a pretrained large language model for text
138 encoding to enhance image fidelity and text alignment. Re-Imagen (Chen et al., 2023) further aug-
139 ments this approach by incorporating Retrieval-Augmented Generation (RAG) to improve image
140 quality by grounding from multi-modal knowledge bases. In contrast, DALL-E 2 (Ramesh et al.,
141 2022) uses a diffusion decoder that inverts a CLIP image encoder, effectively bridging text embed-
142 dings and image generation in a semantically rich manner.

143 Owing to the substantial computational demands of pixel space models for high-resolution syn-
144 thesis, latent space models have emerged as a more efficient alternative. These models perform
145 the diffusion process in a compressed latent space derived from pretrained autoencoders such as
146 VQ-VAE (Van Den Oord et al., 2017), which reduces computational load while maintaining image
147 quality. A well-known example is Stable Diffusion (Podell et al., 2024), which builds on the lat-
148 ent diffusion framework to generate high-resolution images efficiently. Additionally, NAO (Samuel
149 et al., 2023) investigates the structure of the latent space to further enhance performance, especially
150 in long-tail and few-shot scenarios. Despite these advances, a rigorous evaluation of these models’
151 ability to accurately count objects in generated images remains largely unexplored, motivating the
152 empirical studies in this paper. Our findings in this paper may also inspire future directions for
153 enhancing current text-to-image and text-to-video diffusion models, particularly regarding control-
154 lability (Wang et al., 2024c;a; Cheng et al., 2025; Cao et al., 2025a) and expressiveness (Cao et al.,
155 2025c; Chen et al., 2025; Gong et al., 2025; Cao et al., 2025b), thereby providing novel insights into
156 the synthesis process and benchmark performance.

157 3 THE T2I COUNTBENCH

158
159
160 In this section, we first introduce the baseline models used in our benchmark in Section 3.1, fol-
161 lowed by the prompts designed to evaluate the counting ability of text-to-image diffusion models in
Section 3.2. We then describe our evaluation protocol in Section 3.3.

Table 1: Basic information of the Evaluated Text-to-Image Diffusion Models.

Model Name	Organization	Year	# Params	Open
Recraft V3 (AI, 2024a)	Recraft AI	2024	N/A	No
Imagen-3 (Baldrige et al., 2024)	Google	2024	N/A	No
Grok 3 (xAI, 2025)	xAI	2025	N/A	No
Gemini 2.0 Flash (Google, 2025)	Google	2025	N/A	No
FLUX 1.1 (Labs, 2024)	Black Forest	2024	N/A	No
Firefly 3 (Adobe, 2024)	Adobe	2024	N/A	No
Dall-E 3 (Betker et al., 2023)	OpenAI	2024	N/A	No
SD 3.5 Large Turbo (AI, 2024b)	Stability AI	2024	8.1B	Yes
Doubao (Team, 2025)	Bytedance	2023	N/A	No
Qwen2.5-Max (Yang et al., 2024a)	Alibaba	2025	N/A	No
WanX2.1 (Cloud, 2025)	Alibaba	2025	14B	Yes
Kling (Kuaishou, 2024)	Kwai	2024	N/A	No
Star-3 Alpha (LiblibAI, 2024)	LiblibAI	2024	N/A	No
Hunyuan (Li et al., 2024b)	Tencent	2024	1.5B	Yes
GLM-4 (GLM et al., 2024)	ZhipuAI	2024	9B	Yes

3.1 BASELINE MODELS

A rigorous evaluation of the counting ability of text-to-image diffusion models requires a diverse and up-to-date selection of models. However, existing benchmarks often fall short in this issue. For instance, a human evaluation benchmark that includes counting tasks (Petsiuk et al., 2022) considers only Stable Diffusion (Rombach et al., 2022) and DALL·E 2 (Ramesh et al., 2022), both released in 2022, covering a limited subset of available models. Similarly, several recent benchmarks (Li et al., 2024a; Meng et al., 2024; Fu et al., 2024) evaluate at most ten text-to-image diffusion models, failing to provide a comprehensive assessment of counting capabilities across the latest systems.

To address these limitations, our benchmark includes 15 state-of-the-art text-to-image diffusion models, encompassing both open-source and privately owned commercial models. This selection ensures broad coverage of models widely used in generative AI research and applications, most of which have been introduced after 2024. By incorporating a more extensive set of models, we provide a trustworthy and representative evaluation of counting performance. Basic information on the selected models is presented in Table 1, and further implementation details on baseline model evaluation (e.g., model type, length-to-width ratio) are presented in Appendix B.

3.2 GENERATION PROMPTS

The design of generation prompts is the key to effectively evaluating text-to-image models. Although counting is a fundamental capability of diffusion models, many existing benchmarks (e.g., ConceptMix (Wu et al., 2024), Commonsense-T2I (Fu et al., 2024), and PhyBench (Meng et al., 2024)) do not include object quantity in their prompts. Moreover, previous studies on evaluating the counting ability of diffusion models have offered only preliminary explorations without a comprehensive, multi-level evaluation (Saharia et al., 2022; Li et al., 2024a). For instance, while GenAI-Bench (Li et al., 2024a) provides a broad evaluation of text-to-image generation, only 339 of its prompts address counting. These prompts are also combined with a wide range of additional conditions, limited to numbers below 10, and often generate fewer than 3 objects.

In contrast, our approach uses a simple yet effective prompt design that directly tests the counting ability while minimizing irrelevant factors. Our prompt template used in most experiments is:

Prompt Template 1: Generate <number> <object> in/on <scene> in <style>.

Here, <number> denotes an integer between 1 and 15 in Arabic numeral form, providing a more comprehensive range than those used in previous benchmarks. The <object> field covers 6 common categories: fruit, human, animal, abstract shape, furniture, and plant. In addition, we vary the scene and style by including 3 different types for each to assess the models’ performance under different conditions. Overall, our benchmark evaluates 525 prompts for each baseline model. These prompts cover all 15 numbers, 7 object categories, and combinations of 3 scenes and 3 styles. For example:

Example Prompt 1.1: Generate 13 chairs on a wooden floor in a watercolor style.

3.3 EVALUATION PROTOCOLS

To ensure a rigorous and thorough evaluation, we adopt a full human evaluation process. Five graduate students with expertise in AI and visual perception assess each generated image. An image is marked as “correct” if it contains exactly the number of objects specified in the prompt; otherwise, it is labeled as “incorrect”. [To ensure a fair comparison, we have each model generate four images per prompt, and we consider the task successful if at least one of the four images is correct.](#) This comprehensive human evaluation offers more reliable results than previous approaches that rely on object detection (Cho et al., 2023) or visual question answering models (Hu et al., 2023), both of which may introduce biases.

[Our primary evaluation metric is counting accuracy, which considers only whether the generated images contain the correct number of objects.](#) Each unique combination of object, scene, and style is treated as a distinct task, and overall accuracy is computed from correct outputs across all 15 numbers and relevant prompts. [This design allows us to more directly and intuitively compare the counting capabilities of different text-to-image models.](#)

4 EXPERIMENTS

In this section, we present our experimental results using the proposed T2ICountBench. Section 4.1 reports the overall counting performance of all baseline models, while Section 4.2 investigates the impact of various factors on the counting ability of text-to-image diffusion models. [Finally, Section 4.3 presents our analysis of variance across human annotators.](#)

4.1 OVERALL COUNTING RESULTS

To evaluate the fundamental counting ability of diffusion models, we employ the general prompt described as Prompt Template 1 in Section 3.2. Specifically, the four key elements in the prompt template are instantiated as follows:

- `<number>`: 1, 2, 3, ..., 15;
- `<object>`: 'fruit', 'human', 'animal', 'shape', 'furniture', 'plant';
- `<scene>`: 'home', 'nature', 'city';
- `<style>`: 'plain', 'watercolor', 'cartoon'.

For each model, we generate outputs for all possible combinations of these properties and record the number of cases in which the generated image contains the correct quantity of objects. All counting results are evaluated through a full human evaluation process as described in Section 3.3. We then categorize the results by object class and present them in Table 2.

The overall results lead to several observations. First, when considering both per-category and overall average accuracy, all state-of-the-art text-to-image diffusion models struggle to generate objects in the correct quantities. No model achieves an average accuracy above 50%, and for each category, accuracy does not exceed 60%. Additionally, the variance across different object categories is minimal, indicating that models consistently perform poorly across all categories. These findings highlight a significant gap in the counting ability of diffusion models.

Furthermore, a comparison among models reveals a large disparity in performance. For instance, the strongest models, such as Imagen-3 (with an average accuracy of 43%) and Gemini 2.0 Flash (with an average accuracy of 39%), significantly outperform models like Recraft V3 and SD 3.5, which achieve average accuracies of 25% and 26%, respectively. This represents nearly a 150% improvement in accuracy between the best and worst performing models.

Observation 4.1. *Overall, state-of-the-art models exhibit a significant gap in accurately counting objects, and the performance difference between the strongest and weakest models is notable.*

Table 2: **Counting Accuracy Across Different Object Categories.** The models are sorted in ascending order based on their average accuracy across six object categories.

Model	Fruit	Human	Animal	Shape	Furniture	Plant	Avg. Acc.
Recraft V3	0.23	0.36	0.27	0.23	0.24	0.20	0.25
SD 3.5	0.14	0.33	0.35	0.27	0.27	0.23	0.26
Grok 3	0.21	0.55	0.23	0.35	0.23	0.17	0.29
DALL-E 3	0.29	0.31	0.43	0.23	0.31	0.23	0.30
GLM-4	0.27	0.33	0.44	0.25	0.37	0.24	0.32
Qwen2.5-Max	0.35	0.29	0.41	0.33	0.33	0.19	0.32
Firefly 3	0.27	0.41	0.51	0.29	0.33	0.20	0.34
FLUX 1.1	0.31	0.43	0.40	0.27	0.36	0.31	0.35
Kling	0.30	0.51	0.45	0.19	0.33	0.35	0.35
Doubao	0.35	0.43	0.4	0.35	0.39	0.33	0.37
Hunyuan	0.33	0.36	0.49	0.37	0.36	0.32	0.37
Star-3 Alpha	0.38	0.39	0.44	0.39	0.39	0.27	0.37
WanX2.1	0.37	0.52	0.47	0.32	0.36	0.20	0.37
Gemini 2.0 Flash	0.28	0.48	0.48	0.51	0.29	0.32	0.39
Imagen-3	0.31	0.53	0.51	0.44	0.43	0.33	0.43

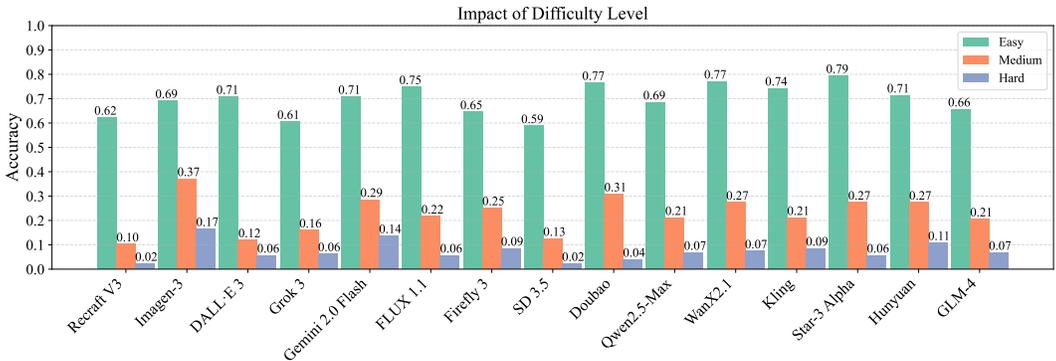


Figure 1: **Impact of Difficulty Levels.** This figure presents the comparison of the accuracy of various models across three difficulty levels (Easy, Medium, Hard). The horizontal axis lists the models, while the vertical axis represents accuracy. Each bar in the figure represents the accuracy for a specific model under the corresponding prompt difficulty level.

4.2 ABLATION STUDY

In this section, we present several ablation studies to examine the impact of different factors on the counting ability of text-to-image diffusion models, including difficulty levels, scenes, and styles.

Impact of Difficulty Levels. To assess the models’ ability to handle counting tasks of varying difficulty, we leverage our wide range of counting numbers (1–15). Specifically, we use the same prompt template and evaluation process as described in Section 4.1 to compute the overall accuracy of each model across different difficulty levels. We define three levels: (i) *Easy*: counting tasks with numbers 1, 2, 3, 4, 5; (ii) *Medium*: counting tasks with numbers 6, 7, 8, 9, 10; (iii) *Hard*: counting tasks with numbers 11, 12, 13, 14, 15.

Figure 1 clearly shows a significant gap in counting accuracy across the three difficulty levels. For all 15 models, the *Easy* level yields accuracies between approximately 60% and 80%, while the *Medium* level drops to between 10% and 30%. The most striking results are observed at the *Hard* level, where almost all models achieve accuracies below 10%. Only Imagen-3, Gemini 2.0 Flash, and Hunyuan exceed 10%, with some models such as Recraft V3 and SD 3.5 achieving accuracies as low as 2%. This indicates that these models nearly fail the counting task at higher difficulty levels. We thus make the following observation:

Observation 4.2. *As the counting task becomes more difficult (i.e., as the number of objects increases from 1 to 15), the models’ accuracies drop drastically. For tasks involving 11–15 objects, nearly all models exhibit accuracies below 10%.*



354 **Figure 2: Qualitative Study of Different Difficulty Levels.** A high-resolution version of this image
355 is available in Figure 8 in Appendix D.

356
357
358 To further support our observations, we present qualitative results on the impact of difficulty levels
359 in Figure 2. We observe that at higher difficulty levels, all models make significant mistakes, with
360 the generated object counts deviating markedly from the target. In contrast, images generated at
361 lower difficulty levels sometimes succeed (e.g., Imagen-3 achieved perfect counts in both easy and
362 medium prompts). Additionally, higher difficulty levels tend to result in reduced image detail and
363 fidelity; for example, in the “15 cats” prompt on GLM-4, the generated image depicts a cat with
364 two tails. This indicates that increased task difficulty not only exacerbates counting errors but also
365 adversely affects overall image quality, which resonates with our quantitative results in this study.
366 Due to space limitations, we moved the statement on the impact of the scene and the impact of style
367 to the Appendix C.

368 4.3 HUMAN ANNOTATOR VARIANCE ANALYSIS

369
370 For each prompt and model, four images were generated and independently evaluated by five anno-
371 tators. An annotator considered the model’s result correct if at least one of the four images had a
372 correct count; otherwise, it was marked as incorrect. To assess consistency among annotators, we
373 computed Fleiss’ Kappa using Eq. (1). The resulting value of 0.58 indicates moderate inter-annotator
374 agreement.

375
376
377

$$\kappa = \frac{P - P_e}{1 - P_e}, \quad (1)$$

where $P = \frac{1}{N} \sum_{i=1}^N \frac{n_{i0}(n_{i0}-1)+n_{i1}(n_{i1}-1)}{n(n-1)}$, $P_e = \left(\frac{1}{Nn} \sum_{i=1}^N n_{i0}\right)^2 + \left(\frac{1}{Nn} \sum_{i=1}^N n_{i1}\right)^2$, N represents the number of evaluated image groups, n is the number of five annotators. n_{i0} indicates that the i th sample is marked as an incorrect count, and n_{i1} indicates that the i th sample is marked as a correct count.

5 PROMPT REFINEMENT

In this section, we address the counting limitations of text-to-image diffusion models through prompt refinement. Specifically, we first introduce our proposed prompt refinements in Section 5.1, followed by experimental results in Section 5.2. Finally, in Section 5.3 we discuss several open questions and conjectures regarding the counting ability of text-to-image models.

5.1 THE PROPOSED PROMPTS

Due to the poor performance observed when directly generating a large number of objects (as shown in Section 4), we adopt a simple work-around by refining the prompts, which verifies whether such counting limitations can be solved by straightforward improvements. Our exploratory study takes a task-decomposition approach by breaking the generation task into smaller subtasks, which mirrors how humans draw many objects on a single canvas. We consider four types of prompt refinement mechanisms: Multiplicative Decomposition, Additive Decomposition, Grid Prior, and Position Guidance.

Multiplicative Decomposition. For example, when drawing 15 apples on a table, a human might consider drawing 5 apples in a row and repeating this process 3 times. In this prompt refinement, we decompose the task by instructing the model to generate a large number of objects as smaller groups. Specifically, let the number of objects to be generated be N , and let a be a factor of N smaller than $\sqrt{N/2}$, and b be a factor larger than $\sqrt{N/2}$, satisfying $N = ab$. When N is prime, its only factors are 1 and N , so we set $a = 1$ and $b = N$. Our prompt can be shown as follows:

Prompt Template 2: Generate a times b <object> in/on <scene> in <style>.

For example, considering the task of generating 12 watermelons on a wooden table in a cartoon style, the refined prompt would be:

Example Prompt 2.1: Generate 3 times 4 watermelons on a wooden table in cartoon style.

Besides the basic prompt refinement introduced above, we also explore three mechanisms, namely **Additive Decomposition**, **Grid Prior**, and **Position Guidance**. Due to space limitations, their detailed descriptions and examples are provided in Appendix C.

5.2 PROMPT REFINEMENT RESULTS

Building on the prompt refinement approaches introduced in the previous subsection, we systematically evaluate whether these refinements can mitigate the counting limitations of text-to-image diffusion models. In this study, we consider the prompt templates in Prompts 2–5 from Section 5.1 and fill in the properties as follows:

<number>: 1, 2, 3, ..., 15; <object>: 'fruit', 'human', 'animal', 'shape', 'furniture', 'plant'.

In order to focus on the simplest generation scenarios and eliminate the impact of extraneous factors, we fix the <scene> and <style> to 'Home' and 'Plain', respectively. Specifically, we compute the average accuracy for each model across all six object types under each prompt refinement strategy, and our results are presented in Table 3.

The results in Table 3 reveal that all four types of prompt refinement lead to worse performance compared with the original prompt. The performance drop is particularly pronounced for multiplicative decomposition, additive decomposition, and position guidance, where the average accuracy across 15 models decreases by more than 40% relative to the original accuracy (dropping from 42% to 26%, 23%, and 20%, respectively). In some cases, such as with Firefly 3, the reduction is as steep as

75% (from 42% to 10% under multiplicative decomposition). Among the refinement strategies, grid prior shows the most promise, as its performance drop is relatively marginal compared with other methods.

Observation 5.1. *For most models and in most cases, prompt refinement degrades the counting performance of text-to-image diffusion models, with particularly severe drops observed for multiplicative decomposition, additive decomposition, and position guidance.*

Table 3: **Prompt Refinement Results.** Each entry represents the average accuracy across all object categories for a specific prompt refinement method.

Model	Original	Multiplicative	Additive	Grid	Position	Avg. Acc.
Recraft V3	0.37	0.29	0.26	0.26	0.15	0.26
Imagen-3	0.58	0.33	0.29	0.49	0.26	0.39
DALL-E 3	0.36	0.38	0.20	0.33	0.16	0.29
Grok 3	0.34	0.26	0.35	0.26	0.22	0.29
Gemini 2.0 Flash	0.46	0.39	0.30	0.36	0.33	0.37
FLUX 1.1	0.38	0.16	0.24	0.32	0.13	0.25
Firefly 3	0.42	0.10	0.16	0.39	0.17	0.25
SD 3.5	0.29	0.18	0.15	0.30	0.13	0.21
Doubao	0.48	0.20	0.12	0.40	0.08	0.26
Qwen2.5-Max	0.41	0.27	0.27	0.35	0.19	0.30
WanX2.1	0.50	0.35	0.19	0.34	0.30	0.34
Kling	0.40	0.17	0.10	0.33	0.07	0.22
Star-3 Alpha	0.35	0.25	0.15	0.29	0.22	0.25
Hunyuan	0.45	0.28	0.26	0.21	0.26	0.29
GLM-4	0.46	0.36	0.35	0.39	0.28	0.37
Avg. Acc.	0.42	0.26	0.23	0.34	0.20	0.29

5.3 DISCUSSION

We discuss possible reasons for the observed counting failures in text-to-image diffusion models. One key reason for poor counting performance is that several early text-to-image models (e.g., Stable Diffusion (Rombach et al., 2022), SDXL (Podell et al., 2024), unCLIP (Ramesh et al., 2022)) use CLIP (Radford et al., 2021) as their text encoder. Previous studies have demonstrated that CLIP has inherent counting issues (Paiss et al., 2023; Jiang et al., 2023; Zhang et al., 2024). This limitation also contributes to the failure of prompt refinement, as CLIP is not designed to process complex, instruction-based prompts. In contrast, models such as Imagen (Saharia et al., 2022) and DALL-E 3 (Betker et al., 2023) employ large language models like T5 (Raffel et al., 2020) for prompt processing, which offer improved language understanding. Nonetheless, their counting failures may stem from insufficient alignment with human preferences, preventing strict adherence to detailed instructions.

To improve the counting capability in existing text-to-image diffusion models, there are several open directions, including CLIP counting ability improvement, automatic prompt refinement, and human preference alignment. Due to the space limitation, we defer the more details of the potential directions are presented in Appendix A.

6 CONCLUSION

In this paper, we introduced T2ICountBench, a comprehensive benchmark to rigorously evaluate the counting ability of text-to-image diffusion models. Our extensive evaluations reveal that even state-of-the-art models struggle to adhere to numerical constraints, with accuracy dropping sharply as the number of objects increases and under complex scene conditions. We also show that simple prompt refinements generally fail to improve counting performance, underscoring inherent challenges in numerical understanding within these models. These findings motivate further research to address these limitations and enhance the reliability of diffusion-based generative systems.

486 ETHIC STATEMENT
487

488 This paper does not involve human subjects, personally identifiable data, or sensitive applications.
489 We do not foresee direct ethical risks. We follow the ICLR Code of Ethics and affirm that all aspects
490 of this research comply with the principles of fairness, transparency, and integrity.
491

492 REPRODUCIBILITY STATEMENT
493

494 We ensure the reproducibility of our empirical findings. For all experiments, we describe the sources
495 of the LLM models, datasets, evaluation metrics, and experiment setup in the main text. All prompt
496 templates used are also provided to support the reproducibility of our results.
497

498 REFERENCES
499

- 500 Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Ale-
501 man, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical
502 report. *arXiv preprint arXiv:2303.08774*, 2023.
503
- 504 Adobe. Adobe introduces firefly image 3 foundation model to take creative exploration and ideation
505 to new heights, 2024. URL [https://news.adobe.com/news/news-details/2024/
506 adobe-introduces-firefly-image-3-foundation-model-to-take-creat
507 ive-exploration-and-ideation-to-new-heights](https://news.adobe.com/news/news-details/2024/adobe-introduces-firefly-image-3-foundation-model-to-take-creative-exploration-and-ideation-to-new-heights).
- 508 Recraft AI. Recraft introduces a revolutionary ai model that thinks in design language, 2024a. URL
509 [https://www.recraft.ai/blog/recraft-introduces-a-revolutionary-a
510 i-model-that-thinks-in-design-language](https://www.recraft.ai/blog/recraft-introduces-a-revolutionary-ai-model-that-thinks-in-design-language).
- 511 Stability AI. Introducing stable diffusion 3.5, 2024b. URL [https://stability.ai/news/
512 introducing-stable-diffusion-3-5](https://stability.ai/news/introducing-stable-diffusion-3-5).
513
- 514 Jason Baldridge, Jakob Bauer, Mukul Bhutani, Nicole Brichtova, Andrew Bunner, Kelvin Chan,
515 Yichang Chen, Sander Dieleman, Yuqing Du, Zach Eaton-Rosen, et al. Imagen 3. *arXiv preprint
516 arXiv:2408.07009*, 2024.
- 517 James Betker, Gabriel Goh, Li Jing, Tim Brooks, Jianfeng Wang, Linjie Li, Long Ouyang, Juntang
518 Zhuang, Joyce Lee, Yufei Guo, et al. Improving image generation with better captions. *Computer
519 Science*. <https://cdn.openai.com/papers/dall-e-3.pdf>, 2(3):8, 2023.
520
- 521 Andreas Blattmann, Robin Rombach, Huan Ling, Tim Dockhorn, Seung Wook Kim, Sanja Fidler,
522 and Karsten Kreis. Align your latents: High-resolution video synthesis with latent diffusion mod-
523 els. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*,
524 pp. 22563–22575, 2023.
- 525 Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal,
526 Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are
527 few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
528
- 529 Yang Cao, Bo Chen, Xiaoyu Li, Yingyu Liang, Zhizhou Sha, Zhenmei Shi, Zhao Song, and Mingda
530 Wan. Force matching with relativistic constraints: A physics-inspired approach to stable and
531 efficient generative modeling. *arXiv preprint arXiv:2502.08150*, 2025a.
- 532 Yang Cao, Zhao Song, and Chiwun Yang. Video latent flow matching: Optimal polynomial projec-
533 tions for video interpolation and extrapolation. *arXiv preprint arXiv:2502.00500*, 2025b.
- 534 Yuefan Cao, Chengyue Gong, Xiaoyu Li, Yingyu Liang, Zhizhou Sha, Zhenmei Shi, and Zhao Song.
535 Richspace: Enriching text-to-video prompt space via text embedding interpolation. *arXiv preprint
536 arXiv:2501.09982*, 2025c.
537
- 538 Bo Chen, Chengyue Gong, Xiaoyu Li, Yingyu Liang, Zhizhou Sha, Zhenmei Shi, Zhao Song,
539 and Mingda Wan. High-order matching for one-step shortcut diffusion models. *arXiv preprint
arXiv:2502.00688*, 2025.

- 540 Wenhu Chen, Hexiang Hu, Chitwan Saharia, and William W. Cohen. Re-imagen: Retrieval-
541 augmented text-to-image generator. In *The Eleventh International Conference on Learning Rep-*
542 *resentations*, 2023. URL <https://openreview.net/forum?id=XSEBx0iSjFQ>.
543
- 544 Dabing Cheng, Haosen Zhan, Xingchen Zhao, Guisheng Liu, Zemin Li, Jinghui Xie, Zhao Song,
545 Weiguo Feng, and Bingyue Peng. Text-to-edit: Controllable end-to-end video ad creation via
546 multimodal llms. *arXiv preprint arXiv:2501.05884*, 2025.
- 547 Jaemin Cho, Abhay Zala, and Mohit Bansal. Dall-eval: Probing the reasoning skills and social biases
548 of text-to-image generation models. In *Proceedings of the IEEE/CVF international conference*
549 *on computer vision*, pp. 3043–3054, 2023.
550
- 551 Alibaba Cloud. Alibaba cloud unveiled wanx 2.1: Redefining ai-driven video generation, 2025.
552 URL [https://www.alibabacloud.com/blog/alibaba-cloud-unveiled-wan](https://www.alibabacloud.com/blog/alibaba-cloud-unveiled-wanx-2-1-redefining-ai-driven-video-generation_601930)
553 [x-2-1-redefining-ai-driven-video-generation_601930](https://www.alibabacloud.com/blog/alibaba-cloud-unveiled-wanx-2-1-redefining-ai-driven-video-generation_601930).
554
- 555 Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances*
556 *in neural information processing systems*, 34:8780–8794, 2021.
- 557 Jeff Donahue and Karen Simonyan. Large scale adversarial representation learning. *Advances in*
558 *neural information processing systems*, 32, 2019.
559
- 560 Ying Fan, Olivia Watkins, Yuqing Du, Hao Liu, Moonkyung Ryu, Craig Boutilier, Pieter Abbeel,
561 Mohammad Ghavamzadeh, Kangwook Lee, and Kimin Lee. Dpok: Reinforcement learning for
562 fine-tuning text-to-image diffusion models. *Advances in Neural Information Processing Systems*,
563 36:79858–79885, 2023.
- 564 Xingyu Fu, Muyu He, Yujie Lu, William Yang Wang, and Dan Roth. Commonsense-t2i challenge:
565 Can text-to-image generation models understand commonsense? In *First Conference on Lan-*
566 *guage Modeling*, 2024. URL <https://openreview.net/forum?id=MI52iXSSNy>.
567
- 568 Dhruva Ghosh, Hannaneh Hajishirzi, and Ludwig Schmidt. Geneval: an object-focused framework
569 for evaluating text-to-image alignment. In *Proceedings of the 37th International Conference on*
570 *Neural Information Processing Systems*, pp. 52132–52152, 2023.
- 571 Team GLM, Aohan Zeng, Bin Xu, Bowen Wang, Chenhui Zhang, Da Yin, Dan Zhang, Diego Rojas,
572 Guanyu Feng, Hanlin Zhao, et al. Chatglm: A family of large language models from glm-130b to
573 glm-4 all tools. *arXiv preprint arXiv:2406.12793*, 2024.
574
- 575 Chengyue Gong, Yekun Ke, Xiaoyu Li, Yingyu Liang, Zhizhou Sha, Zhenmei Shi, and Zhao
576 Song. On computational limits of flowar models: Expressivity and efficiency. *arXiv preprint*
577 *arXiv:2502.16490*, 2025.
- 578 Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair,
579 Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information*
580 *processing systems*, 27, 2014.
581
- 582 Google. Gemini 2.0 is now available to everyone, 2025. URL [https://blog.google/tech](https://blog.google/technology/google-deepmind/gemini-model-updates-february-2025/)
583 [nology/google-deepmind/gemini-model-updates-february-2025/](https://blog.google/technology/google-deepmind/gemini-model-updates-february-2025/).
584
- 585 Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in*
586 *neural information processing systems*, 33:6840–6851, 2020.
- 587 Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J
588 Fleet. Video diffusion models. *Advances in Neural Information Processing Systems*, 35:8633–
589 8646, 2022.
590
- 591 Yushi Hu, Benlin Liu, Jungo Kasai, Yizhong Wang, Mari Ostendorf, Ranjay Krishna, and Noah A
592 Smith. Tifa: Accurate and interpretable text-to-image faithfulness evaluation with question an-
593 swering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp.
20406–20417, 2023.

- 594 Kaiyi Huang, Kaiyue Sun, Enze Xie, Zhenguo Li, and Xihui Liu. T2i-compbench: A compre-
595 hensive benchmark for open-world compositional text-to-image generation. *Advances in Neural*
596 *Information Processing Systems*, 36:78723–78747, 2023.
- 597 Ruixiang Jiang, Lingbo Liu, and Changwen Chen. Clip-count: Towards text-guided zero-shot object
598 counting. In *Proceedings of the 31st ACM International Conference on Multimedia*, pp. 4535–
599 4545, 2023.
- 600 Diederik P Kingma and Max Welling. Auto-encoding variational bayes. In *International Conference*
601 *on Learning Representations (ICLR)*, 2014.
- 602 Kuaishou. Kling ai: Next-generation ai creative studio. <https://klingai.com/>, 2024.
- 603 Black Forest Labs. Announcing the flux pro finetuning api, 2024. URL <https://blackforestlabs.ai/announcements/>.
- 604 Baiqi Li, Zhiqiu Lin, Deepak Pathak, Jiayao Li, Yixin Fei, Kewen Wu, Xide Xia, Pengchuan Zhang,
605 Graham Neubig, and Deva Ramanan. Evaluating and improving compositional text-to-visual
606 generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recog-*
607 *niton*, pp. 5290–5301, 2024a.
- 608 Zhimin Li, Jianwei Zhang, Qin Lin, Jiangfeng Xiong, Yanxin Long, Xincheng Deng, Yingfang Zhang,
609 Xingchao Liu, Minbin Huang, Zedong Xiao, et al. Hunyuan-dit: A powerful multi-resolution
610 diffusion transformer with fine-grained chinese understanding. *arXiv preprint arXiv:2405.08748*,
611 2024b.
- 612 LiblibAI. Star3 alpha, 2024. URL [https://www.liblib.art/modelinfo/8b86d942a](https://www.liblib.art/modelinfo/8b86d942adc241228900c7c5f0a47ad9)
613 [dc241228900c7c5f0a47ad9](https://www.liblib.art/modelinfo/8b86d942adc241228900c7c5f0a47ad9).
- 614 Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow
615 matching for generative modeling. In *The Eleventh International Conference on Learning Repre-*
616 *sentations*, 2023. URL <https://openreview.net/forum?id=PqvMRDCJT9t>.
- 617 Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao,
618 Chengqi Deng, Chenyu Zhang, Chong Ruan, et al. Deepseek-v3 technical report. *arXiv preprint*
619 *arXiv:2412.19437*, 2024.
- 620 Fanqing Meng, Wenqi Shao, Lixin Luo, Yahong Wang, Yiran Chen, Quanfeng Lu, Yue Yang, Tian-
621 shuo Yang, Kaipeng Zhang, Yu Qiao, et al. Phybench: A physical commonsense benchmark for
622 evaluating text-to-image models. *arXiv preprint arXiv:2406.11802*, 2024.
- 623 Wenyi Mo, Tianyu Zhang, Yalong Bai, Bing Su, Ji-Rong Wen, and Qing Yang. Dynamic prompt
624 optimizing for text-to-image generation. In *Proceedings of the IEEE/CVF Conference on Computer*
625 *Vision and Pattern Recognition*, pp. 26627–26636, 2024.
- 626 Alexander Quinn Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob
627 McGrew, Ilya Sutskever, and Mark Chen. Glide: Towards photorealistic image generation and
628 editing with text-guided diffusion models. In *International Conference on Machine Learning*, pp.
629 16784–16804. PMLR, 2022.
- 630 Aaron Oord, Yazhe Li, Igor Babuschkin, Karen Simonyan, Oriol Vinyals, Koray Kavukcuoglu,
631 George Driessche, Edward Lockhart, Luis Cobo, Florian Stimberg, et al. Parallel wavenet: Fast
632 high-fidelity speech synthesis. In *International conference on machine learning*, pp. 3918–3926.
633 PMLR, 2018.
- 634 OpenAI. Video generation models as world simulators. Technical report, OpenAI, February 2024.
635 URL [https://openai.com/index/video-generation-models-as-world-s](https://openai.com/index/video-generation-models-as-world-simulators)
636 [imulators](https://openai.com/index/video-generation-models-as-world-simulators).
- 637 Roni Paiss, Ariel Ephrat, Omer Tov, Shiran Zada, Inbar Mosseri, Michal Irani, and Tali Dekel.
638 Teaching clip to count to ten. In *Proceedings of the IEEE/CVF International Conference on*
639 *Computer Vision*, pp. 3170–3180, 2023.

- 648 William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of*
649 *the IEEE/CVF International Conference on Computer Vision*, pp. 4195–4205, 2023.
- 650
- 651 Yuang Peng, Yuxin Cui, Haomiao Tang, Zekun Qi, Runpei Dong, Jing Bai, Chunrui Han, Zheng Ge,
652 Xiangyu Zhang, and Shu-Tao Xia. Dreambench++: A human-aligned benchmark for personalized
653 image generation. *arXiv preprint arXiv:2406.16855*, 2024.
- 654 Vitali Petsiuk, Alexander E Siemenn, Saisamrit Surbehera, Zad Chin, Keith Tyser, Gregory Hunter,
655 Arvind Raghavan, Yann Hicke, Bryan A Plummer, Ori Kerret, et al. Human evaluation of text-to-
656 image models on a multi-task benchmark. *arXiv preprint arXiv:2211.12112*, 2022.
- 657
- 658 Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe
659 Penna, and Robin Rombach. SDXL: Improving latent diffusion models for high-resolution image
660 synthesis. In *The Twelfth International Conference on Learning Representations*, 2024. URL
661 <https://openreview.net/forum?id=di52zR8xgf>.
- 662 Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal,
663 Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual
664 models from natural language supervision. In *International conference on machine learning*, pp.
665 8748–8763. PmLR, 2021.
- 666
- 667 Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever.
668 Robust speech recognition via large-scale weak supervision. In *International conference on ma-*
669 *chine learning*, pp. 28492–28518. PMLR, 2023.
- 670 Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi
671 Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text
672 transformer. *Journal of machine learning research*, 21(140):1–67, 2020.
- 673
- 674 Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-
675 conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 1(2):3, 2022.
- 676
- 677 Ali Razavi, Aaron Van den Oord, and Oriol Vinyals. Generating diverse high-fidelity images with
678 vq-vae-2. *Advances in neural information processing systems*, 32, 2019.
- 679 Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-
680 resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF confer-*
681 *ence on computer vision and pattern recognition*, pp. 10684–10695, 2022.
- 682
- 683 Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman.
684 Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In *Pro-*
685 *ceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 22500–
686 22510, 2023.
- 687
- 688 Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar
689 Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic
690 text-to-image diffusion models with deep language understanding. *Advances in neural informa-*
691 *tion processing systems*, 35:36479–36494, 2022.
- 692
- 693 Dvir Samuel, Rami Ben-Ari, Nir Darshan, Haggai Maron, and Gal Chechik. Norm-guided latent
694 space exploration for text-to-image generation. *Advances in Neural Information Processing Sys-*
695 *tems*, 36:57863–57875, 2023.
- 696
- 697 Uriel Singer, Adam Polyak, Thomas Hayes, Xi Yin, Jie An, Songyang Zhang, Qiyuan Hu, Harry
698 Yang, Oron Ashual, Oran Gafni, Devi Parikh, Sonal Gupta, and Yaniv Taigman. Make-a-video:
699 Text-to-video generation without text-video data. In *The Eleventh International Conference on*
700 *Learning Representations*, 2023. URL [https://openreview.net/forum?id=nJfy1D](https://openreview.net/forum?id=nJfy1Dvgz1q)
701 [vgz1q](https://openreview.net/forum?id=nJfy1Dvgz1q).
- 702
- 703 Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *Interna-*
704 *tional Conference on Learning Representations*, 2021a. URL [https://openreview.net](https://openreview.net/forum?id=StlgiaRCHLP)
705 [/forum?id=StlgiaRCHLP](https://openreview.net/forum?id=StlgiaRCHLP).

- 702 Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben
703 Poole. Score-based generative modeling through stochastic differential equations. In *International
704 Conference on Learning Representations*, 2021b. URL [https://openreview.net
705 /forum?id=PxtIG12RRHS](https://openreview.net/forum?id=PxtIG12RRHS).
- 706 Xu Tan, Jiawei Chen, Haohe Liu, Jian Cong, Chen Zhang, Yanqing Liu, Xi Wang, Yichong Leng,
707 Yuanhao Yi, Lei He, et al. Naturalspeech: End-to-end text-to-speech synthesis with human-level
708 quality. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(6):4234–4245,
709 2024.
- 710 Doubao Team. Doubao1.5 pro, 2025. URL [https://team.doubao.com/en/special/d
711 oubao_1_5_pro](https://team.doubao.com/en/special/doubao_1_5_pro).
- 712 Sergey Tulyakov, Ming-Yu Liu, Xiaodong Yang, and Jan Kautz. Mocogan: Decomposing motion
713 and content for video generation. In *Proceedings of the IEEE conference on computer vision and
714 pattern recognition*, pp. 1526–1535, 2018.
- 715 Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *Advances in
716 neural information processing systems*, 30, 2017.
- 717 Yilin Wang, Haiyang Xu, Xiang Zhang, Zeyuan Chen, Zhizhou Sha, Zirui Wang, and Zhuowen Tu.
718 Omnicontrolnet: Dual-stage integration for conditional image generation. In *Proceedings of the
719 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 7436–7448, 2024a.
- 720 Zhijie Wang, Yuheng Huang, Da Song, Lei Ma, and Tianyi Zhang. Promptcharm: Text-to-image
721 generation through multi-modal prompting and refinement. In *Proceedings of the 2024 CHI
722 Conference on Human Factors in Computing Systems*, pp. 1–21, 2024b.
- 723 Zirui Wang, Zhizhou Sha, Zheng Ding, Yilin Wang, and Zhuowen Tu. Tokencompose: Text-to-
724 image diffusion with token-level supervision. In *Proceedings of the IEEE/CVF Conference on
725 Computer Vision and Pattern Recognition*, pp. 8553–8564, 2024c.
- 726 Olivia Wiles, Chuhan Zhang, Isabela Albuquerque, Ivana Kajić, Su Wang, Emanuele Bugliarello,
727 Yasumasa Onoe, Chris Knutsen, Cyrus Rashtchian, Jordi Pont-Tuset, et al. Revisiting text-
728 to-image evaluation with gecko: On metrics, prompts, and human ratings. *arXiv preprint
729 arXiv:2404.16820*, 2024.
- 730 Jay Zhangjie Wu, Yixiao Ge, Xintao Wang, Stan Weixian Lei, Yuchao Gu, Yufei Shi, Wynne Hsu,
731 Ying Shan, Xiaohu Qie, and Mike Zheng Shou. Tune-a-video: One-shot tuning of image diffusion
732 models for text-to-video generation. In *Proceedings of the IEEE/CVF International Conference
733 on Computer Vision*, pp. 7623–7633, 2023a.
- 734 Xiaoshi Wu, Yiming Hao, Keqiang Sun, Yixiong Chen, Feng Zhu, Rui Zhao, and Hongsheng Li.
735 Human preference score v2: A solid benchmark for evaluating human preferences of text-to-
736 image synthesis. *arXiv preprint arXiv:2306.09341*, 2023b.
- 737 Xindi Wu, Dingli Yu, Yangsibo Huang, Olga Russakovsky, and Sanjeev Arora. Conceptmix: A
738 compositional image generation benchmark with controllable difficulty. In *NeurIPS 2024 Work-
739 shop on Compositional Learning: Perspectives, Methods, and Paths Forward*, 2024. URL
740 <https://openreview.net/forum?id=2UWnlxVNPz>.
- 741 xAI. Grok 3 beta, 2025. URL <https://x.ai/blog/grok-3>.
- 742 Tao Xu, Pengchuan Zhang, Qiuyuan Huang, Han Zhang, Zhe Gan, Xiaolei Huang, and Xiaodong
743 He. AttnGAN: Fine-grained text to image generation with attentional generative adversarial net-
744 works. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp.
745 1316–1324, 2018.
- 746 An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li,
747 Dayiheng Liu, Fei Huang, Haoran Wei, et al. Qwen2.5 technical report. *arXiv preprint
748 arXiv:2412.15115*, 2024a.
- 749 Chenglin Yang, Celong Liu, Xueqing Deng, Dongwon Kim, Xing Mei, Xiaohui Shen, and Liang-
750 Chieh Chen. 1.58-bit flux. *arXiv preprint arXiv:2412.18653*, 2024b.

756 Ling Yang, Jingwei Liu, Shenda Hong, Zhilong Zhang, Zhilin Huang, Zheming Cai, Wentao Zhang,
757 and Bin Cui. Improving diffusion-based image synthesis with context prediction. *Advances in*
758 *Neural Information Processing Systems*, 36:37636–37656, 2023.

759
760 Zhuoyi Yang, Jiayan Teng, Wendi Zheng, Ming Ding, Shiyu Huang, Jiazheng Xu, Yuanming Yang,
761 Wenyi Hong, Xiaohan Zhang, Guanyu Feng, et al. Cogvideox: Text-to-video diffusion models
762 with an expert transformer. *arXiv preprint arXiv:2408.06072*, 2024c.

763 Zeliang Zhang, Zhuo Liu, Mingqian Feng, and Chenliang Xu. Can clip count stars? an empirical
764 study on quantity bias in clip. In *Findings of the Association for Computational Linguistics:*
765 *EMNLP 2024*, pp. 1081–1086, 2024.

766
767
768
769
770
771
772
773
774
775
776
777
778
779
780
781
782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809

Appendix

In Section A, we discuss some future directions. In Section B, we present the details of the evaluated generative models. In Section C, we show the details of two prompt refinement mechanisms and some additional quantitative experiments. In Section D, we present more qualitative studies. In Section E, we discuss the potential risks of this paper. In Section F, we show a full list of the results for every single experiment. In Section G, we present all the generated images in this benchmark study.

A FUTURE WORKS

To enhance counting ability in text-to-image models, one direction is to improve CLIP-based models by incorporating recent advances that address CLIP’s counting shortcomings (Paiss et al., 2023; Jiang et al., 2023). Another promising approach is automatic prompt refinement (Mo et al., 2024; Wang et al., 2024b), which translates complex human instructions into simpler forms that diffusion models can more reliably interpret. For models leveraging large language models, reinforcement learning techniques may further align generated images with human preferences and improve the processing of task-decomposition prompts (Fan et al., 2023).

B IMPLEMENTATION DETAILS

In this section, we provide the implementation details for generating images using the baseline models outlined in Table 1. Specifically, the details for all 15 models are listed as follows:

- Model 1: **Recraft V3** (AI, 2024a). Recraft V3 is a close-sourced text-to-image model from Recraft AI company, released in 2024. We use default mode of this model for experiment.
- Model 2: **Imagen-3** (Baldrige et al., 2024). Imagen-3 is a close-sourced text-to-image model from Google company, released in 2024. We use best quality mode of this model for experiment. Since the default setting for landscape ratio is 16:9, we change it to 1:1 to ensure fair comparison.
- Model 3: **Grok 3** (xAI, 2025). Grok 3 is a close-sourced multi-modal model from xAI company, released in 2025. We use default mode of this model for experiment.
- Model 4: **Gemini 2.0 Flash** (Google, 2025). Gemini 2.0 Flash is a close-sourced multi-modal model from Google company, released in 2025. We use default mode of this model for experiment.
- Model 5: **FLUX 1.1** (Labs, 2024). FLUX 1.1 is a close-sourced text-to-image model from Black Forest Labs company, released in 2024. We use default mode of this model for experiment. Since the default setting for landscape ratio is 4:3, we change it to 1:1 to ensure fair comparison.
- Model 6: **Firefly 3** (Adobe, 2024). Firefly 3 is a close-sourced multi-modal model from Adobe company, released in 2024. We use fast mode of this model for experiment.
- Model 7: **Dall-E 3** (Betker et al., 2023). Dall-E 3 is a close-sourced text-to-image model from OpenAI company, released in 2024. We use default mode of this model for experiment.
- Model 8: **Stable Diffusion 3.5 Large Turbo** (AI, 2024b). Stable Diffusion 3.5 Large Turbo is an open-sourced text-to-image model from Stability AI company, released in 2024. We use default mode of this model for experiment.
- Model 9: **Doubao** (Team, 2025). Doubao is a close-sourced multi-modal model from Bytedance company, released in 2023. We use default mode of this model for experiment.
- Model 10: **Qwen2.5-Max** (Yang et al., 2024a). Qwen2.5-Max is a close-sourced multi-modal model from Alibaba Cloud company, released in 2025. We use default mode of this model for experiment.

- 864 • Model 11: **WanX2.1** (Cloud, 2025). WanX2.1 is a close-sourced multi-modal model from
865 Alibaba Cloud company, released in 2025. We use default mode of this model for experi-
866 ment.
- 867 • Model 12: **Kling** (Kuaishou, 2024). Kling is a close-sourced multi-modal model from
868 Kwai company, released in 2024. We use default mode of this model for experiment.
- 869 • Model 13: **Star-3 Alpha** (LiblibAI, 2024). Star-3 Alpha is a close-sourced text-to-image
870 model from LiblibAI company, released in 2024. We use default mode of this model for
871 experiment. Since the default setting for landscape ratio is 4:3, we change it to 1:1 to ensure
872 fair comparison.
- 873 • Model 14: **Hunyuan** (Li et al., 2024b). Hunyuan is an open-sourced multi-modal model
874 from Tencent company, released in 2024. We use default mode of this model for experi-
875 ment.
- 876 • Model 15: **GLM-4** (GLM et al., 2024). GLM-4 is an open-sourced multi-modal model
877 from ZhipuAI company, released in 2024. We use default mode of this model for experi-
878 ment.
879

880 C ADDITIONAL EXPERIMENTS

881
882
883 **Additive Decomposition.** Another human-inspired approach to generating a large number of ob-
884 jects is to first create a subset of objects and then generate the remainder. Unlike the multiplicative
885 decomposition, this method breaks the task into two smaller parts that are subsequently combined.
886 Specifically, let the number of objects to be generated be N , where $N \geq 2$, and let $\lfloor x \rfloor$ denote the
887 floor function, which returns the largest integer less than or equal to x . We define our prompt as
888 follows:

889 **Prompt Template 3:** Generate $\lfloor N/2 \rfloor$ plus $N - \lfloor N/2 \rfloor$ <object> in/on <scene> in <style>.

890
891 An example for such prompt refinement on generating 11 objects would be:

892
893 **Example Prompt 3.1:** Generate 5 plus 6 triangles on a painting on a wall.

894
895 **Grid Prior.** An extension of the multiplicative decomposition is to provide an explicit spatial ar-
896 rangement for the objects. Without such guidance, the model might be uncertain about where to
897 place the generated objects. Therefore, we use a grid layout to structure the output. This process
898 resembles a chain-of-thought strategy by breaking down the task into simpler, sequential steps, in
899 which the first step determines the positions, and the second step puts the objects. Specifically, let
900 the number of objects to be generated be N , and let a be its largest factor smaller than $N/2$, and b
901 be the smallest factor larger than $N/2$, so that $N = ab$. Our prompt can be shown as follows:

902
903 **Prompt Template 4:** Generate <number> <object> in/on <scene> in <style>, with a a
row b column grid.

904
905 Extending the example of 12 watermelons from the multiplicative decomposition, we have the fol-
906 lowing instance:

907
908 **Example Prompt 4.1:** Generate 12 watermelons on a wooden table in cartoon style, with a 3
row 4 column grid.

909
910 **Position Guidance.** A further extension of the additive decomposition approach is to provide ex-
911 plicit positional guidance. In this method, the two groups of objects are placed in designated areas
912 on the canvas, which reduces the cognitive load on the model and provides clearer instructions. In
913 our template, the first group is positioned on the left and the second group on the right. We designed
914 the prompt carefully to ensure that the positional instructions integrate seamlessly with the scene
915 and style constraints:

916
917 **Prompt Template 5:** Generate $\lfloor N/2 \rfloor$ <object> on the left, $N - \lfloor N/2 \rfloor$ <object> on the
right, in/on <scene> in <style>.

Extending the triangles example from the additive decomposition, an example for position guidance would be:

Example Prompt 5.1: Generate 5 triangles on the left, 6 triangles on the right, on a painting on a wall.

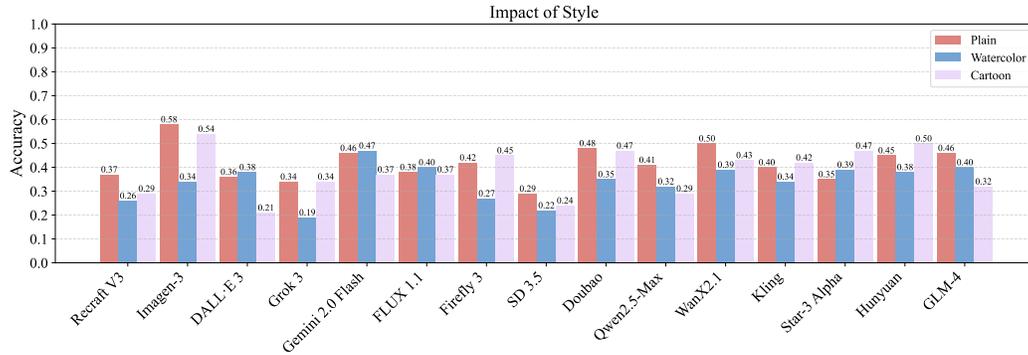


Figure 3: **Impact of Style.** This figure presents the comparison of the accuracy of various models across three styles (Plain, Watercolor, Cartoon). The horizontal axis lists the models, while the vertical axis represents accuracy. Each bar in the figure represents the accuracy for a specific model under corresponding prompt style setting.

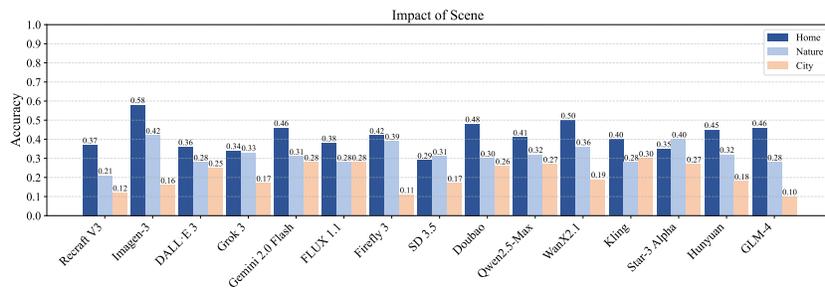


Figure 4: **Impact of Scene.** This figure presents the comparison of the accuracy of various models across three scenes (Home, Nature, City). The horizontal axis lists the models, while the vertical axis represents accuracy. Each bar in the figure represents the accuracy for a specific model under the corresponding prompt scene setting.

Impact of Scene. In this study, we investigate how the scene in which objects are presented affects counting performance. The intuition is that complex environments, such as cityscapes with multiple irrelevant elements, may pose a greater challenge compared to simpler settings like a simple wooden table in a home environment. We use the general prompt described in Prompt Template 1 with the following settings:

<number>: 1,2,3,...,15; <object>: 'fruit', 'human', 'animal', 'shape', 'furniture', 'plant'; <scene>: 'home', 'nature', 'city'.

- <number>: 1,2,3,...,15;
- <object>: 'fruit', 'human', 'animal', 'shape', 'furniture', 'plant';
- <scene>: 'home', 'nature', 'city';

The <style> keyword is fixed to 'plain' to exclude the effect of styles and focus on the effect of scene modifications. All generation results are evaluated by human annotators, and the results for each scene are summarized in Figure 4 and Figure 5.

The experimental results reveal a significant variance in counting accuracy across different scenes. When averaging the results of all 15 models, the home scene achieves an average accuracy of

972
973
974
975
976
977
978
979
980
981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025

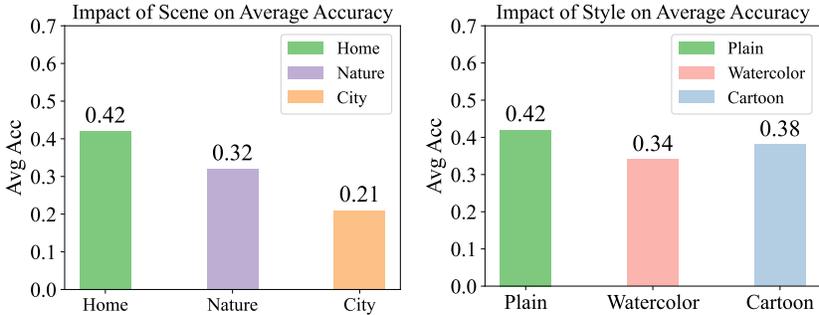


Figure 5: **Impact of Scene and Style on Average Accuracy.** **Left:** This figure presents a comparison of the average accuracy across three scenes (Home, Nature, City) for 15 models. Each bar represents the average accuracy of the 15 models under the corresponding prompt scene setting. **Right:** This figure presents a comparison of the average accuracy across three styles (Plain, Watercolor, Cartoon) for 15 models. Each bar represents the average accuracy of the 15 models under the corresponding prompt style setting.

42%, whereas the `city` scene falls to 21%—a reduction of nearly 50%. This indicates that the compositional complexity of a scene strongly influences a model’s counting ability. Moreover, the variation in accuracy for individual models across scenes can be even more pronounced than the average difference across all models. For example, Imagen-3 achieves an accuracy of 58% in the `home` scene but only 16% in the `city` scene, while GLM-4 scores 46% in `home` compared to just 10% in `city`. This leads us to the following observation:

Observation C.1. *The models’ counting ability is significantly affected by the scene. Complex scenes such as `city` and `nature` lead to a drop in counting performance.*

Another interesting finding is that a model performing well in one scene does not necessarily excel in other scenes. For instance, in the `home` scene, FLUX 1.1 ranks among the worst in counting accuracy; however, in the `city` scene, its accuracy rises to 28%, making it the second best in that category. Similarly, the best model in the `home` and `nature` scenes, Imagen-3, shows relatively poor performance in the `city` scene compared to other models. This variability suggests a notable instability in the counting ability of text-to-image models across different scenes, indicating a potential direction for future research. We summarize this observation as follows:

Observation C.2. *Models that perform well in one scene may not maintain high performance in other scenes, highlighting an instability in counting ability under varying scene conditions.*

Impact of Style. In this study, we examine the effect of image style on the counting ability of text-to-image diffusion models. Unlike the scene, which can introduce many irrelevant objects, style is an important property while imposing less generation burden on the generative models. We use the previously used prompt described in Prompt Template 1 and follow the same human evaluation protocols as in our other experiments. Specifically, we use the following property composition to fill in the prompt template:

- `<number>`: 1, 2, 3, ..., 15;
- `<object>`: 'fruit', 'human', 'animal', 'shape', 'furniture', 'plant';
- `<style>`: 'plain', 'watercolor', 'cartoon'.

To exclude the effect of scenes and focus on the style categories, the `<scene>` keyword is fixed to 'home'. By aggregating the accuracy results into three style categories, we present the findings in Figure 3 and Figure 5.

The results indicate that style has a less significant impact on the counting performance compared to the scene. For example, the average accuracy across all 15 models for the styles `plain`, `watercolor`, and `cartoon` are 42%, 34%, and 38%, respectively, which are on a similar scale. Furthermore, for specific models such as FLUX 1.1 and SD 3.5, the variance in accuracy across different style classes is minimal. Thus, we summarize the following observation:

Observation C.3. *Style categories have a less significant impact on models’ counting abilities.*

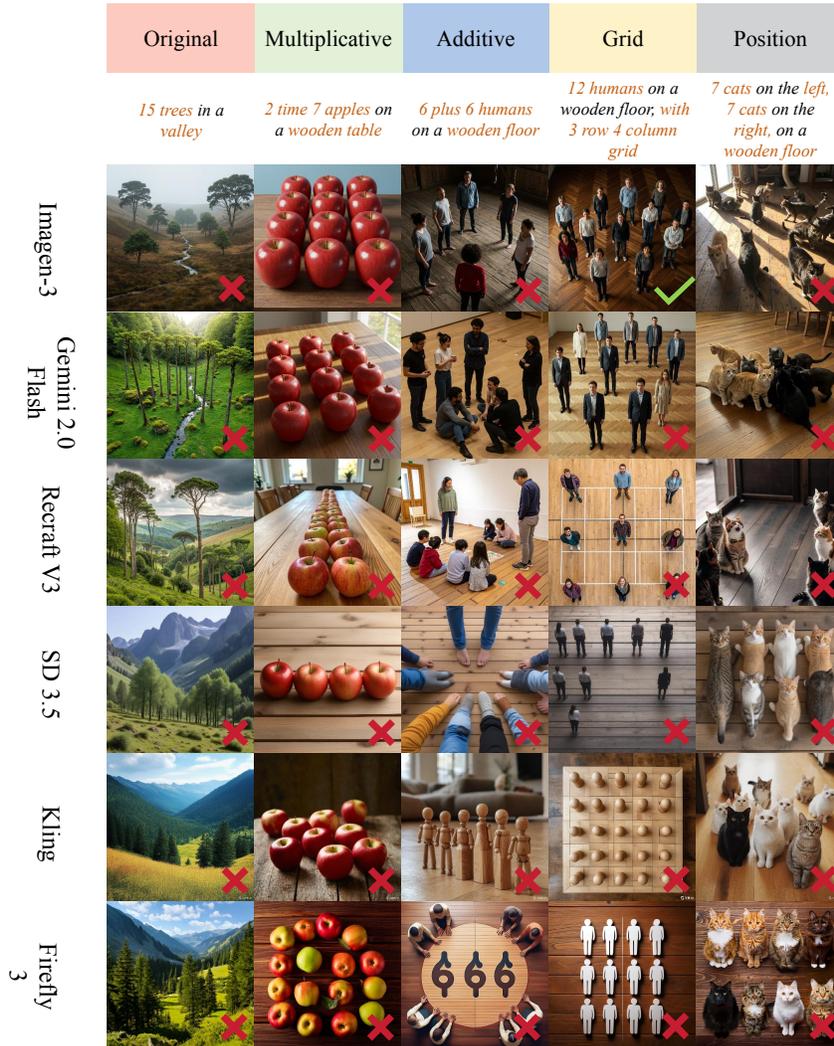


Figure 6: **Qualitative Study of Prompt Refinement Results.** This figure presents the qualitative study of the prompt refinement results in Section 5.1. A high-resolution version of this image is available in Figure 11 in Appendix D.

To support our observations on prompt refinement, we present a qualitative study in Figure 6. The figure shows that for the multiplicative refinement prompt “2 times 7 apples,” almost all models fail to adhere to both numbers (2 and 7), completely disregarding the instruction; only Recraft V3 manages to generate two columns, while still failing to produce the correct number of rows. For the additive prompt, models appear to misinterpret “6 + 6” as simply 6. Furthermore, with more complex prompts such as grid prior or position guidance, nearly all models struggle with the subtasks—they fail to correctly interpret directions (e.g., left and right) or generate a grid with the correct number of rows and columns, and in some cases, do not generate a grid at all. These diverse failure cases further reinforce our quantitative findings that prompt refinement does not overcome the counting limitations of text-to-image diffusion models.

D QUALITATIVE STUDY

In this section, we introduce the qualitative study based on our experiments across all models.

Qualitative Study on Main Results. We present a qualitative study on the main results in Figure 7. The images generated by most models exhibit satisfactory fidelity and aesthetics, with minimal distortions or incorrect spatial relationships (e.g., misplaced eyes or noses on human and cat faces). One notable negative example is the fruit result of the “5 watermelons” prompt on SD 3.5 in Figure 7, where the watermelons are irregularly cut and arranged messily. This issue may be attributable to the relatively small number of parameters of the model. Despite the overall high fidelity, many models still encounter counting errors, as demonstrated by the “5 flowers” prompt on Gemini 2.0 Flash in the plant result. This observation suggests that fidelity and counting accuracy are not necessarily correlated—a model that produces high-fidelity images may still fail to count objects correctly. Interestingly, some models misinterpret the word “earthy.” For instance, in the shape results, the “3 triangles” prompt on Doubao produced an output in which an Earth-like model is depicted on land with triangles superimposed on its surface. Although the counting outcome is correct, this example indicates that these models may benefit from additional human preference alignment to better follow user instructions.

Qualitative Study on the Impact of Different Difficulty Levels. We present a qualitative study on the impact of different difficulty levels in Figure 8. Our observations indicate that as the difficulty level increases, the quality of the generated images deteriorates. For example, at the medium difficulty level, the “9 humans” prompt on Imagen-3 demonstrates that the model can generate the correct number of objects; however, at higher difficulty levels, the “15 cats” prompt on Imagen-3 reveals that the model tends to produce unsatisfactory results, underscoring the inherent limitations of diffusion-based text-to-image models. Furthermore, we observe that higher difficulty levels are associated with diminished image detail and fidelity. For instance, in the “15 cats” prompt on GLM-4, the generated image features a cat with two tails. This not only indicates counting difficulties but also suggests that increased task difficulty adversely affects other aspects of the model’s performance in certain cases.

Qualitative Study on the Impact of Scene. We present a qualitative study on the impact of scene in Figure 9. We observe that the scene context can adversely influence the models’ counting ability. For instance, when prompted with “8 trees,” the models often generate more trees than specified, frequently relegating many trees to the background. This behavior may stem from a conflict between the models’ large-scale prior knowledge (e.g., that many trees typically line streets) and the instruction to produce only a limited number of trees. Additionally, scene context can impact image fidelity; in the “8 trees” example with Dall·E 3, trees are placed in the middle of the road, which contradicts common sense.

Qualitative Study on the Impact of Style. We present a qualitative study on the impact of scene in Figure 10. The results indicate that altering the style does not overcome the inherent counting limitations of text-to-image diffusion models. Specifically, for the “8 flowers” prompt in a cartoon style, all models fail to produce the correct number of flowers. Moreover, in the “10 humans” example rendered in cartoon style, the image generated by GLM-4 exhibits noticeable facial distortions. These findings suggest that while style variations can modify visual aesthetics, they may also affect the overall fidelity of the generated images.

Qualitative Study on Prompt Refinement Results. To support our observations on prompt refinement, we present a qualitative study in Figure 11. This figure shows that for the multiplicative refinement prompt “2 times 7 apples,” almost all models fail to adhere to both numbers (2 and 7), completely disregarding the instruction; only Recraft V3 manages to generate 2 columns, while still failing to generate the correct number of objects. For the additive prompt, models appear to misinterpret “6 + 6” as simply 6. Furthermore, with more complex prompts such as grid prior or position guidance, nearly all models struggle with the subtasks—they fail to correctly interpret directions (e.g., left and right) or generate a grid with the correct number of rows and columns, and in some cases, do not generate a grid at all. These diverse failure cases further reinforce our quantitative findings that prompt refinement does not overcome the counting limitations of text-to-image diffusion models.

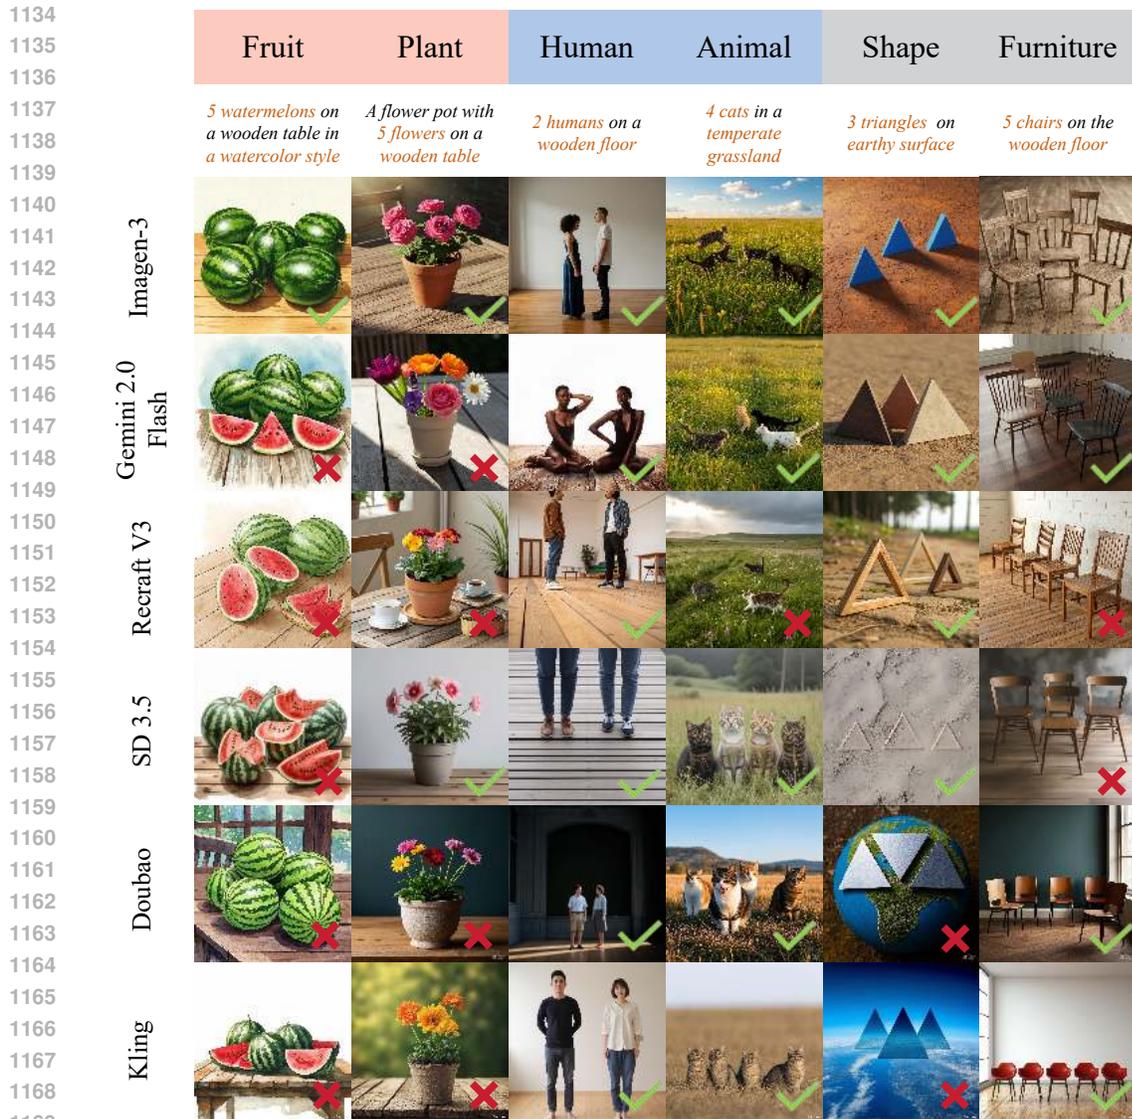
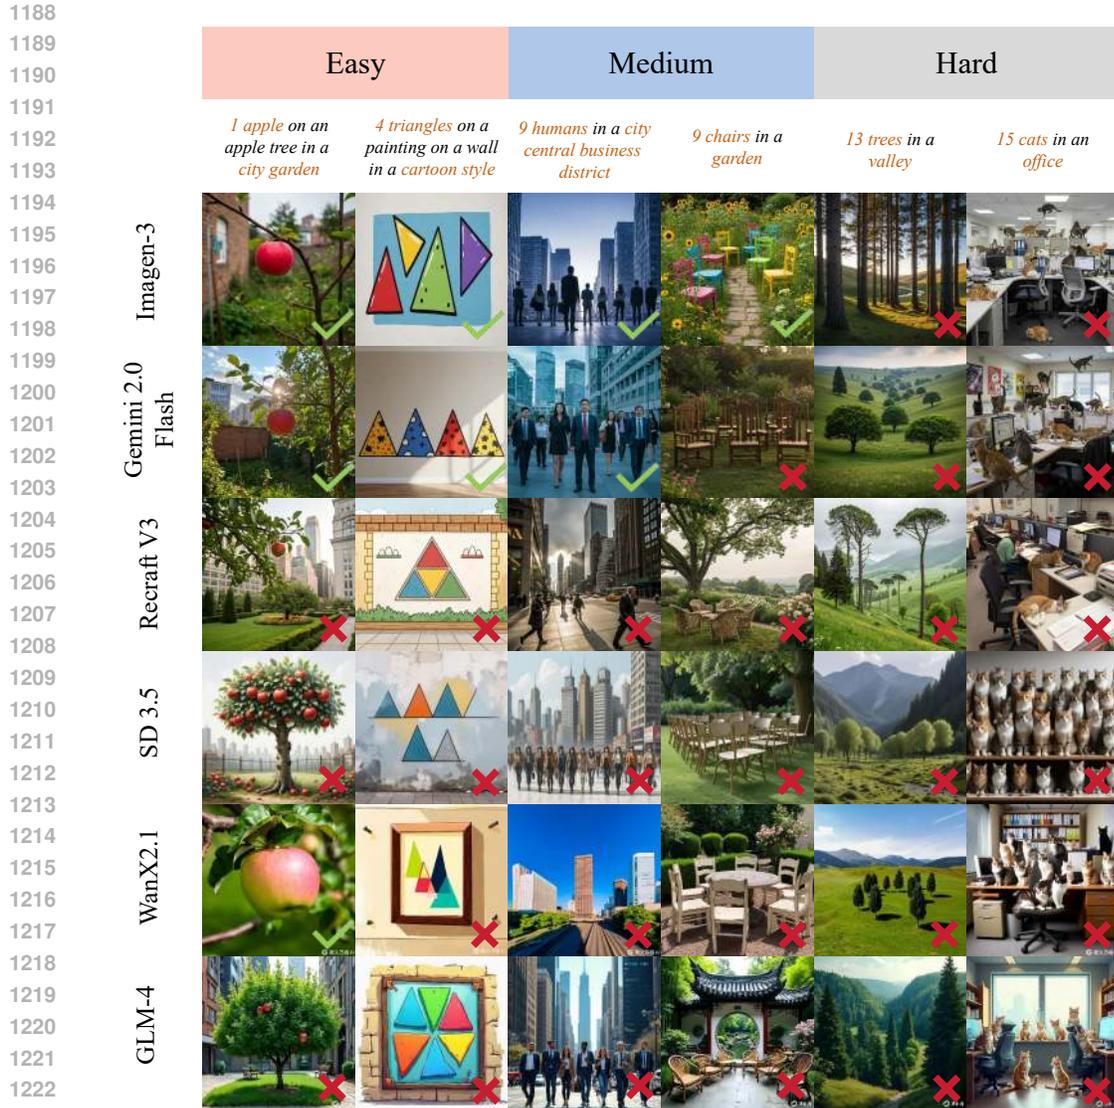


Figure 7: **Qualitative Study on Main Results.** This figure presents the qualitative study of the main results in Section 4.1. We selected the two best models (top two rows) with the highest average accuracy in Table 2, the two worst models (middle two rows) with the lowest average accuracy in Table 2, and two additional models (bottom two rows) that exhibit distinct behaviors. Correct images are marked with a tick, whereas erroneous images are indicated with a cross.

E POTENTIAL RISKS

One potential risk of our work is that the suggested directions for improving counting abilities in diffusion-based text-to-image models may lead to more realistic image generation, which could be misused to mislead the public. However, we believe that existing safeguard mechanisms for diffusion models remain effective for mitigating such risks. Moreover, our work focuses solely on benchmarking and does not involve releasing any new large pretrained models. Therefore, we do not foresee any negative societal impact resulting from this study.



1224
1225
1226
1227
1228
1229
1230

Figure 8: **Qualitative Study on Different Difficulty Levels.** This figure presents the qualitative study on the different difficulty levels in Section 4.2. We selected the two best models (top two rows in this figure) with the highest average accuracy in Table 2, the two worst models (middle two rows in this figure) with the lowest average accuracy in Table 2, and two additional models (bottom two rows in this figure) that exhibit distinct behaviors. Correct images are marked with a tick, whereas erroneous images are indicated with a cross.

1231 F DETAILED RESULTS

1232
1233
1234
1235
1236
1237
1238
1239
1240
1241

In this section, we present the detailed results for each prompt used in our experiments across all models in tables 4–39. The models’ capability in generating exactly 1-15 objects is thoroughly demonstrated here, instead of focusing on difficulty levels.

1242
1243
1244
1245
1246
1247
1248
1249
1250
1251
1252
1253
1254
1255
1256
1257
1258
1259
1260
1261
1262
1263
1264
1265
1266
1267
1268
1269
1270
1271
1272
1273
1274
1275
1276
1277
1278
1279
1280
1281
1282
1283
1284
1285
1286
1287
1288
1289
1290
1291
1292
1293
1294
1295

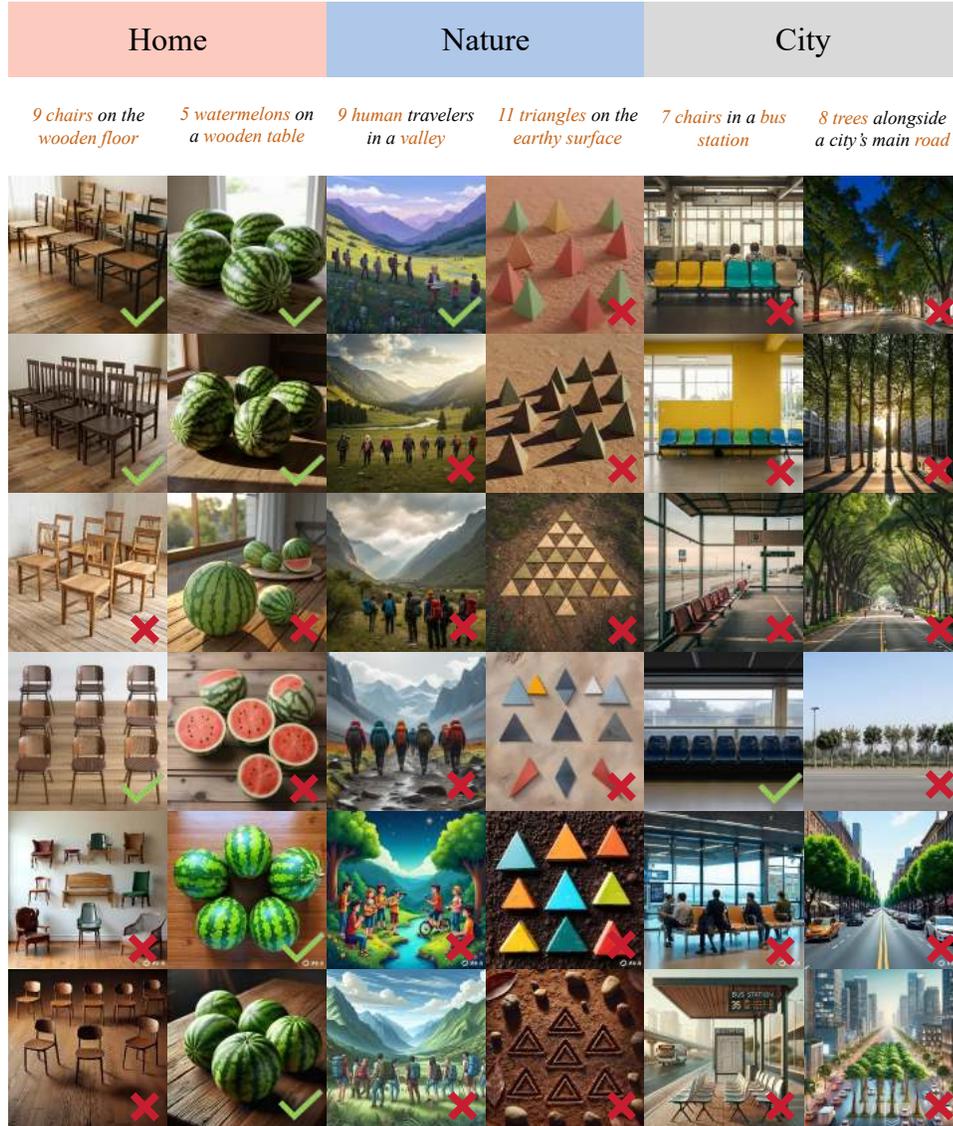


Figure 9: **Qualitative Study on the Impact of Scene.** This figure presents the qualitative study of the impact of scene in Section 4.2. We selected the two best models (top two rows in this figure) with the highest average accuracy in Table 2, the two worst models (middle two rows) with the lowest average accuracy in Table 2, and two additional models (bottom two rows) that exhibit distinct behaviors. Correct images are marked with a tick, whereas erroneous images are indicated with a cross.

1296
1297
1298
1299
1300
1301
1302
1303
1304
1305
1306
1307
1308
1309
1310
1311
1312
1313
1314
1315
1316
1317
1318
1319
1320
1321
1322
1323
1324
1325
1326
1327
1328
1329
1330
1331
1332
1333
1334
1335
1336
1337
1338
1339
1340
1341
1342
1343
1344
1345
1346
1347
1348
1349



Figure 10: **Qualitative Study on the Impact of Style.** This figure presents the qualitative study on the impact of style in Section 4.2. We selected the two best models (top two rows in this figure) with the highest average accuracy in Table 2, the two worst models (middle two rows in this figure) with the lowest average accuracy in Table 2, and two additional models (bottom two rows in this figure) that exhibit distinct behaviors. Correct images are marked with a tick, whereas erroneous images are indicated with a cross.

1350
1351
1352
1353
1354
1355
1356
1357
1358
1359
1360
1361
1362
1363
1364
1365
1366
1367
1368
1369
1370
1371
1372
1373
1374
1375
1376
1377
1378
1379
1380
1381
1382
1383
1384
1385
1386
1387
1388
1389
1390
1391
1392
1393
1394
1395
1396
1397
1398
1399
1400
1401
1402
1403

	Original	Multiplicative	Additive	Grid	Position
	<i>15 trees in a valley</i>	<i>2 time 7 apples on a wooden table</i>	<i>6 plus 6 humans on a wooden floor</i>	<i>12 humans on a wooden floor, with 3 row 4 column grid</i>	<i>7 cats on the left, 7 cats on the right, on a wooden floor</i>
Imagen-3	 ❌	 ❌	 ❌	 ✅	 ❌
Gemini 2.0 Flash	 ❌	 ❌	 ❌	 ❌	 ❌
Recraft V3	 ❌	 ❌	 ❌	 ❌	 ❌
SD 3.5	 ❌	 ❌	 ❌	 ❌	 ❌
Kling	 ❌	 ❌	 ❌	 ❌	 ❌
Firefly 3	 ❌	 ❌	 ❌	 ❌	 ❌

Figure 11: **Qualitative Study on Prompt Refinement Results.** This figure presents the qualitative study of the prompt refinement results in Section 5.1. We selected the two best models (top two rows in this figure) with the highest average accuracy in Table 2, the two worst models (middle two rows in this figure) with the lowest average accuracy in Table 2, and two additional models (bottom two rows in this figure) that exhibit distinct behaviors. Correct images are marked with a tick, whereas erroneous images are indicated with a cross.

1404

1405

1406

1407

1408

1409

1410

1411

1412

1413

1414

1415

1416

1417

1418

1419

1420

1421

1422

1423

1424

1425

1426

1427

Objects	Model	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	Correct	Wrong
Apples	Recraft V3	1	0	1	1	1	0	0	1	1	0	0	0	1	0	0	7	8
	Imagen-3	1	1	1	1	1	1	0	1	0	1	1	1	0	0	0	10	5
	Dall-E 3	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	4	11
	Grok 3	1	1	1	1	0	0	0	0	0	0	0	1	0	0	0	5	10
	Gemini 2.0 Flash	1	1	1	1	1	1	1	0	0	0	0	1	0	0	0	8	7
	FLUX 1.1	1	1	1	1	1	0	0	0	1	0	0	1	0	1	0	7	8
	Firefly 3	1	1	1	1	0	0	0	0	1	0	0	1	0	0	0	6	9
	SD 3.5	1	1	0	1	1	1	0	1	1	0	0	0	0	0	0	7	8
	Doubao	1	1	1	1	1	0	1	0	1	0	0	0	0	0	0	7	8
	Qwen2.5-Max	1	1	1	1	1	1	0	0	1	0	0	0	0	0	0	7	8
	WanX2.1	1	1	1	1	1	1	1	0	1	0	0	0	0	0	0	8	7
	Kling	1	1	1	1	0	0	1	0	0	0	0	1	0	0	0	6	9
	Star-3 Alpha	1	1	1	1	1	0	1	1	1	1	0	0	0	0	0	9	6
	Hunyuan	1	0	1	1	1	1	1	0	0	0	0	1	0	0	0	8	7
	GLM-4	1	1	1	1	1	0	0	1	1	0	0	0	0	0	0	7	8
Watermelons	Recraft V3	1	1	1	1	0	1	0	0	0	0	0	0	0	0	0	5	10
	Imagen-3	1	1	1	1	1	1	1	1	0	0	1	1	1	1	0	12	3
	Dall-E 3	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	5	10
	Grok 3	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	4	11
	Gemini 2.0 Flash	1	1	1	1	1	1	0	0	1	0	0	1	0	0	0	8	7
	FLUX 1.1	1	1	1	1	0	0	0	0	0	0	0	1	0	1	0	6	9
	Firefly 3	1	1	1	1	0	0	0	1	1	0	0	1	0	0	0	7	8
	SD 3.5	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	15
	Doubao	1	1	1	1	1	1	1	0	1	0	0	0	0	0	0	8	7
	Qwen2.5-Max	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	6	9
	WanX2.1	1	1	1	1	1	0	0	1	0	0	0	1	0	0	0	7	8
	Kling	1	1	1	1	1	0	1	0	0	0	0	0	0	0	0	6	9
	Star-3 Alpha	0	1	0	0	0	0	1	0	0	0	0	1	0	0	0	3	12
	Hunyuan	1	1	0	1	1	1	0	0	1	0	1	0	1	0	0	7	8
	GLM-4	1	1	1	1	1	0	1	0	1	0	0	0	0	0	0	7	8

Table 4: Counting Apples and Watermelons in Home Scene Results.

1428

1429

1430

1431

1432

1433

1434

1435

1436

1437

1438

1439

1440

1441

1442

1443

1444

1445

1446

1447

1448

1449

1450

1451

1452

1453

1454

1455

1456

1457

Objects	Model	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	Correct	Wrong
Humans	Recraft V3	1	1	1	1	1	1	1	1	0	1	0	0	0	0	0	9	6
	Imagen-3	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0	11	4
	Dall-E 3	1	1	1	1	1	0	0	1	0	1	0	0	0	0	0	7	8
	Grok 3	1	1	1	1	0	1	1	1	0	1	1	1	1	0	0	11	4
	Gemini 2.0 Flash	1	1	1	1	1	1	1	1	0	0	0	0	0	0	0	8	7
	FLUX 1.1	1	1	1	1	1	1	1	0	0	0	0	0	0	0	0	7	8
	Firefly 3	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	6	9
	SD 3.5	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	5	10
	Doubao	1	1	1	1	1	1	1	0	1	1	1	1	0	0	0	11	4
	Qwen2.5-Max	1	1	1	0	1	0	1	1	1	0	1	0	0	0	0	8	7
	WanX2.1	1	1	1	1	1	1	1	1	1	1	0	1	0	0	0	11	4
	Kling	1	1	1	1	1	1	0	1	0	0	1	0	0	1	0	9	6
	Star-3 Alpha	1	1	1	0	1	0	0	0	0	0	0	0	1	0	0	5	10
	Hunyuan	1	1	1	1	1	0	0	0	1	0	1	0	0	0	0	7	8
	GLM-4	1	1	1	1	1	0	1	0	0	0	0	0	0	0	0	6	9
Cats	Recraft V3	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	5	10
	Imagen-3	1	1	1	1	1	0	1	1	1	0	0	0	1	0	0	9	6
	Dall-E 3	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	5	10
	Grok 3	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	4	11
	Gemini 2.0 Flash	1	1	1	1	0	1	0	1	0	0	0	1	0	0	0	7	8
	FLUX 1.1	1	1	1	1	1	1	1	0	0	0	0	0	0	0	0	7	8
	Firefly 3	1	1	1	1	1	1	1	0	0	0	0	0	1	0	0	8	7
	SD 3.5	1	1	1	1	1	0	1	0	0	0	0	0	0	0	0	6	9
	Doubao	1	1	1	1	1	0	0	0	0	1	0	0	0	0	0	6	9
	Qwen2.5-Max	1	1	1	1	0	1	0	0	0	1	1	0	0	0	0	6	9
	WanX2.1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	9	6
	Kling	1	1	1	1	1	0	0	0	1	0	0	0	0	0	1	8	7
	Star-3 Alpha	1	1	1	0	1	1	0	1	0	0	0	0	0	0	0	6	8
	Hunyuan	1	1	1	1	1	1	1	0	1	1	1	1	1	1	0	13	2
	GLM-4	1	1	1	1	1	0	1	0	0	0	0	0	0	0	0	6	9

Table 5: Counting Human and Animals in Home Scene Results.

1458

1459

1460

1461

1462

Objects	Model	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	Correct	Wrong
Triangles	Recraft V3	1	1	1	1	0	0	0	0	0	1	1	0	0	0	0	6	9
	Imagen-3	1	1	1	0	1	0	1	0	1	1	0	1	0	0	0	8	7
	Dall-E 3	1	1	1	1	1	0	0	1	0	0	0	1	0	0	0	7	8
	Grok 3	1	1	0	0	0	1	0	1	1	0	0	0	0	0	0	5	10
	Gemini 2.0 Flash	1	1	1	1	1	1	0	1	1	1	0	0	0	1	0	10	5
	FLUX 1.1	1	1	1	1	1	0	1	0	0	0	0	1	0	0	0	7	8
	Firefly 3	1	1	1	1	1	0	0	0	1	0	0	0	0	0	0	6	9
	SD 3.5	1	1	1	1	0	0	0	1	1	0	0	0	0	0	0	6	9
	Doubao	1	1	1	1	1	0	0	1	0	0	0	0	0	0	0	6	9
	Qwen2.5-Max	1	1	0	0	1	1	0	0	1	0	0	0	0	0	0	4	11
	WanX2.1	1	1	1	1	1	0	0	1	0	0	0	0	0	0	0	6	9
	Kling	1	1	1	1	1	0	0	0	1	0	0	0	0	0	0	6	9
	Star-3 Alpha	1	1	1	0	0	0	0	0	1	1	0	0	0	0	0	5	10
	Hunyuan	1	1	1	1	1	0	0	0	1	0	0	0	0	0	0	6	9
	GLM-4	1	1	1	1	1	1	0	0	0	0	0	1	0	0	1	8	7
Chairs	Recraft V3	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	2	13
	Imagen-3	1	1	1	0	0	1	1	1	0	0	0	0	0	0	0	6	9
	Dall-E 3	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	4	11
	Grok 3	1	1	1	0	1	0	0	0	0	0	0	0	1	0	0	5	10
	Gemini 2.0 Flash	1	1	0	0	1	0	0	0	0	0	0	0	0	0	0	3	12
	FLUX 1.1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	2	13
	Firefly 3	1	1	1	1	0	1	0	1	1	0	0	0	0	1	0	8	7
	SD 3.5	0	1	0	1	0	0	0	0	0	0	0	0	0	0	0	2	13
	Doubao	1	1	1	0	1	1	1	0	0	0	0	0	0	0	0	6	9
	Qwen 2.5-Max	1	0	1	1	1	0	0	0	1	0	0	0	0	1	1	8	7
	WanX2.1	1	1	0	1	1	0	0	0	0	0	0	0	0	1	0	4	11
	Kling	1	1	0	1	0	0	0	0	0	0	0	0	0	0	0	3	12
	Star-3 Alpha	1	1	0	0	1	0	0	0	1	0	1	0	0	0	0	5	10
	Hunyuan	1	1	1	0	0	0	0	0	1	0	0	0	0	0	0	4	11
	GLM-4	1	1	1	1	1	1	0	0	1	0	0	0	0	0	0	7	8

1485

Table 6: Counting Triangles and Chairs in Home Scene Results.

1486

1487

1488

1489

1490

1491

1492

1493

1494

1495

Objects	Model	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	Correct	Wrong
Flowers	Recraft V3	1	1	1	1	0	0	0	0	0	0	0	0	0	1	0	5	10
	Imagen-3	1	1	0	0	1	0	0	0	0	1	0	0	0	0	1	5	10
	Dall-E 3	1	1	1	1	1	0	1	0	0	0	0	0	0	0	0	6	9
	Grok 3	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	2	13
	Gemini 2.0 Flash	1	1	1	0	0	0	0	1	0	0	0	0	0	0	0	4	11
	FLUX 1.1	1	1	0	1	0	0	0	0	0	0	0	0	0	0	0	4	11
	Firefly 3	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	3	12
	SD 3.5	1	1	1	0	1	0	0	0	0	0	0	0	0	0	0	4	11
	Doubao	1	1	1	1	0	1	0	1	0	0	0	0	0	0	0	6	9
	Qwen2.5-Max	1	0	1	0	0	0	0	0	0	0	0	0	0	0	0	2	13
	WanX2.1	1	1	1	1	1	1	0	0	0	0	0	1	0	0	0	7	8
	Kling	1	1	1	0	0	1	0	0	0	0	0	0	0	0	0	4	11
	Star-3 Alpha	1	1	0	1	0	1	0	0	0	0	0	0	0	0	0	4	11
	Hunyuan	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	3	12
	GLM-4	1	1	1	0	0	1	1	1	1	0	0	0	0	0	0	7	8

1507

Table 7: Counting Flowers in Home Scene Results.

1508

1509

1510

1511

Objects	Model	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	Correct	Wrong	
Apples	Recraft V3	1	0	0	1	0	0	0	0	0	0	0	0	0	0	0	2	13	
	Imagen-3	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	4	11	
	Dall-E 3	1	1	1	1	0	0	0	0	1	0	0	0	0	0	0	5	10	
	Grok 3	1	1	0	1	0	0	0	0	0	0	0	0	0	0	0	3	12	
	Gemini 2.0 Flash	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	3	12	
	FLUX 1.1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	2	13	
	Firefly 3	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	14	
	SD 3.5	0	0	0	1	0	0	0	0	0	0	0	0	0	0	0	1	14	
	Doubao	1	1	1	0	0	1	1	0	0	0	0	0	0	0	0	5	10	
	Qwen2.5-Max	1	1	1	1	1	0	0	0	0	1	0	0	0	0	0	6	9	
	WanX2.1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	4	11	
	Kling	0	0	0	0	0	0	1	0	1	0	0	0	0	0	0	2	13	
	Star-3 Alpha	1	1	1	1	0	1	1	0	0	0	0	0	0	0	0	6	9	
	Hunyuan	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	15	
	GLM-4	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	1	14	
	Watermelons	Recraft V3	1	0	1	0	0	0	0	0	0	0	0	0	0	0	0	2	13
		Imagen-3	1	0	1	0	1	1	0	0	0	0	0	0	0	0	0	4	11
Dall-E 3		1	1	1	1	0	0	0	0	0	1	0	0	0	0	0	5	10	
Grok 3		1	1	1	1	0	1	0	0	1	0	0	0	0	0	0	6	9	
Gemini 2.0 Flash		1	1	1	0	0	1	0	0	0	0	0	0	0	0	0	4	11	
FLUX 1.1		1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	3	12	
Firefly 3		1	1	1	1	1	1	0	1	0	1	0	0	0	0	0	8	7	
SD 3.5		1	1	1	0	1	0	0	0	0	0	0	0	0	0	0	4	15	
Doubao		1	1	1	0	0	1	0	1	0	0	0	0	0	0	0	5	10	
Qwen2.5-Max		1	1	1	1	1	0	0	0	0	0	1	0	0	0	0	6	9	
WanX2.1		1	1	1	1	1	0	0	1	0	0	0	0	0	0	0	7	8	
Kling		1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	5	10	
Star-3 Alpha		1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	6	9	
Hunyuan		1	1	1	1	1	1	0	0	1	0	0	1	0	0	0	7	8	
GLM-4		1	1	0	1	1	0	0	0	0	0	0	0	0	0	0	4	11	

Table 8: Counting Apples and Watermelons in Nature Scene Results.

1534

1535

1536

1537

Objects	Model	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	Correct	Wrong	
Humans	Recraft V3	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	4	11	
	Imagen-3	1	1	1	1	1	1	1	1	1	0	0	0	1	0	0	10	5	
	Dall-E 3	1	1	1	0	1	0	0	0	0	0	0	1	0	0	0	5	10	
	Grok 3	1	1	1	1	1	1	1	1	1	0	1	1	0	0	0	11	4	
	Gemini 2.0 Flash	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	4	11	
	FLUX 1.1	1	1	1	1	0	1	1	0	0	0	0	0	0	0	0	6	9	
	Firefly 3	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	5	10	
	SD 3.5	1	1	1	1	1	1	0	1	0	0	0	0	0	0	0	7	8	
	Doubao	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	5	10	
	Qwen2.5-Max	1	0	0	1	0	0	0	0	0	0	0	0	0	0	0	2	13	
	WanX2.1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	5	10	
	Kling	1	1	0	1	1	0	1	0	1	0	0	0	0	0	0	4	11	
	Star-3 Alpha	1	1	1	1	1	0	0	1	0	1	0	0	0	0	0	7	8	
	Hunyuan	1	1	1	1	0	0	1	1	0	0	0	0	0	0	0	6	9	
	GLM-4	1	1	0	0	1	1	0	0	0	0	0	0	0	0	0	4	11	
	Cats	Recraft V3	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	3	12
		Imagen-3	1	1	1	1	0	0	0	0	0	1	0	0	0	0	0	5	10
Dall-E 3		1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	5	10	
Grok 3		1	1	1	0	0	0	0	0	1	0	0	0	0	0	0	4	11	
Gemini 2.0 Flash		1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	5	10	
FLUX 1.1		1	1	1	0	0	1	1	0	0	0	0	0	0	0	0	6	9	
Firefly 3		1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	9	6	
SD 3.5		1	1	1	1	0	0	0	0	1	0	0	0	0	0	0	5	10	
Doubao		1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	6	9	
Qwen2.5-Max		1	1	1	1	0	0	1	1	0	0	0	0	0	0	0	6	9	
WanX2.1		1	1	1	1	1	1	0	0	1	0	0	0	0	0	0	8	7	
Kling		1	1	1	1	1	0	1	0	0	0	0	0	0	0	0	6	9	
Star-3 Alpha		1	1	1	1	1	0	1	0	0	0	0	0	0	1	0	7	8	
Hunyuan		1	1	1	0	1	1	0	0	0	0	0	0	0	0	0	5	10	
GLM-4		1	1	1	1	1	1	0	0	1	0	0	0	0	0	0	7	8	

Table 9: Counting Humans and Cats in Nature Scene Results.

1560

1561

1562

1563

1564

1565

1566

1567

1568

1569

1570

Objects	Model	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	Correct	Wrong
Chairs	Recraft V3	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	3	12
	Imagen-3	1	1	1	1	1	0	0	0	1	1	0	0	0	1	0	8	7
	Dall-E 3	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	4	11
	Grok 3	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	3	12
	Gemini 2.0 Flash	1	1	1	0	0	0	0	0	0	0	0	1	0	0	0	4	11
	FLUX 1.1	1	1	1	1	0	0	0	1	1	0	0	0	0	0	0	6	9
	Firefly 3	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	6	9
	SD 3.5	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	5	10
	Doubao	1	1	1	1	0	1	0	0	0	1	0	0	0	0	0	6	9
	Qwen2.5-Max	1	1	1	0	1	0	0	0	0	0	0	0	0	0	0	4	11
	WanX2.1	1	1	1	1	0	1	0	0	0	1	0	0	0	0	0	6	9
	Kling	1	1	1	1	0	0	0	1	0	0	0	1	0	1	0	8	7
	Star-3 Alpha	1	1	0	0	0	1	0	0	1	0	0	0	0	0	0	4	11
	Hunyuan	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	5	10
	GLM-4	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	5	10
Triangles	Recraft V3	1	1	1	0	1	1	0	0	1	0	0	0	0	0	0	6	9
	Imagen-3	1	1	1	1	1	0	0	0	1	1	0	0	0	0	1	6	9
	Dall-E 3	1	0	1	0	0	0	0	0	0	0	0	0	0	0	0	2	13
	Grok 3	1	1	1	0	0	1	0	0	0	0	0	0	0	0	0	4	11
	Gemini 2.0 Flash	1	1	1	0	1	0	1	0	1	0	0	0	0	1	0	7	8
	FLUX 1.1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	2	13
	Firefly 3	1	1	0	1	0	0	0	0	1	0	0	1	1	0	0	6	9
	SD 3.5	1	0	1	1	1	0	1	0	0	0	0	0	0	1	0	6	9
	Doubao	0	0	0	0	0	0	0	0	1	0	0	0	0	0	0	1	14
	Qwen2.5-Max	1	1	1	0	0	1	1	0	0	0	0	0	0	0	0	5	10
	WanX2.1	1	1	1	1	0	0	0	0	1	0	0	0	0	0	0	5	10
	Kling	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	15
	Star-3 Alpha	1	1	1	0	0	0	1	0	1	0	0	1	1	0	0	7	8
	Hunyuan	1	1	1	0	0	0	0	0	0	1	0	1	0	0	0	5	10
	GLM-4	1	0	0	1	0	0	0	0	1	0	0	1	0	0	0	4	11

1593

Table 10: Counting Triangles and Chairs in Nature Scene Results.

1594

1595

1596

1597

1598

1599

1600

1601

1602

1603

Objects	Model	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	Correct	Wrong
Trees	Recraft V3	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	2	13
	Imagen-3	1	1	1	0	1	1	0	0	0	0	0	0	0	0	0	5	10
	Dall-E 3	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	3	12
	Grok 3	1	1	0	1	0	1	0	0	0	0	0	0	0	0	0	4	11
	Gemini 2.0 Flash	1	1	0	1	0	1	1	0	0	1	0	0	0	0	0	6	9
	FLUX 1.1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	5	10
	Firefly 3	1	1	1	1	0	1	0	1	0	0	0	0	0	0	0	6	9
	SD 3.5	1	1	1	1	0	1	0	0	0	0	0	0	0	0	0	5	10
	Doubao	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	4	11
	Qwen2.5-Max	1	1	1	0	1	1	0	0	0	0	0	0	0	0	0	5	10
	WanX2.1	1	1	0	0	0	0	0	0	0	0	0	1	0	0	0	3	12
	Kling	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	3	12
	Star-3 Alpha	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	5	10
	Hunyuan	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	5	10
	GLM-4	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	4	11

1615

1616

1617

1618

1619

Table 11: Counting Trees in Nature Scene Results.

1620

1621

1622

1623

1624

1625

1626

1627

1628

1629

1630

1631

1632

1633

1634

1635

1636

1637

1638

1639

1640

1641

1642

1643

Objects	Model	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	Correct	Wrong
Apples	Recraft V3	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	2	13
	Imagen-3	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	14
	Dall-E 3	1	1	1	1	0	1	0	0	0	0	0	0	0	0	0	5	10
	Grok 3	0	0	0	0	1	0	0	0	0	0	0	0	0	0	0	1	14
	Gemini 2.0 Flash	1	1	0	1	0	0	1	0	0	0	0	0	0	0	1	5	10
	FLUX 1.1	1	1	1	0	0	1	0	0	0	0	0	1	0	0	0	5	10
	Firefly 3	1	1	0	0	0	0	0	0	0	1	0	0	0	0	0	3	12
	SD 3.5	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	15
	Doubao	1	1	1	1	0	1	0	0	0	0	0	0	0	0	0	5	10
	Qwen2.5-Max	1	1	1	0	0	0	0	0	0	0	0	0	0	0	1	4	11
	WanX2.1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	4	11
	Kling	0	0	1	0	0	1	0	1	0	0	0	0	0	0	0	3	12
	Star-3 Alpha	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	6	9
	Hunyuan	0	1	1	1	0	1	0	0	0	0	0	0	0	0	0	4	11
	GLM-4	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	15
Watermelons	Recraft V3	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	2	13
	Imagen-3	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	14
	Dall-E 3	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	6	9
	Grok 3	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	14
	Gemini 2.0 Flash	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	2	13
	FLUX 1.1	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	6	9
	Firefly 3	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	14
	SD 3.5	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	15
	Doubao	0	1	1	1	1	1	0	0	0	0	0	0	0	0	0	5	10
	Qwen2.5-Max	1	1	0	0	0	1	1	1	0	0	0	0	0	0	0	5	10
	WanX2.1	0	0	0	0	1	0	1	0	0	0	0	0	0	0	0	2	13
	Kling	1	1	1	0	0	0	0	1	0	0	0	0	0	0	0	4	11
	Star-3 Alpha	1	1	1	0	1	0	0	0	0	0	0	0	0	0	0	4	11
	Hunyuan	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	1	14
	GLM-4	1	1	0	0	0	1	0	0	0	0	0	0	0	0	0	3	12

Table 12: Counting Apples and Watermelons in City Scene Results.

1644

1645

1646

1647

1648

1649

1650

1651

1652

1653

1654

1655

1656

1657

1658

1659

1660

1661

1662

1663

1664

1665

1666

1667

1668

1669

1670

1671

1672

1673

Objects	Model	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	Correct	Wrong
Humans	Recraft V3	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	4	11
	Imagen-3	1	0	0	0	0	0	0	0	1	0	0	0	0	0	0	2	13
	Dall-E 3	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	15
	Grok 3	1	0	1	0	1	1	1	0	0	1	0	0	0	0	0	6	9
	Gemini 2.0 Flash	1	0	0	0	1	1	0	0	1	0	0	0	0	0	0	4	11
	FLUX 1.1	0	1	0	1	0	0	0	0	0	0	0	0	0	0	0	2	13
	Firefly 3	1	1	1	0	0	0	1	0	0	0	0	0	0	0	0	4	11
	SD 3.5	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	15
	Doubao	0	0	0	1	0	0	0	0	0	0	0	0	0	0	0	1	14
	Qwen2.5-Max	1	1	0	0	0	0	1	0	0	0	0	0	0	0	0	3	12
	WanX2.1	1	1	0	1	0	0	0	0	0	0	0	0	0	0	0	3	12
	Kling	1	1	1	1	1	0	0	0	1	0	0	0	0	0	0	6	9
	Star-3 Alpha	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	3	12
	Hunyuan	0	0	1	0	1	1	0	0	0	0	0	0	0	0	0	3	12
	GLM-4	0	0	0	0	0	0	0	0	0	0	1	0	0	0	0	1	14
Cats	Recraft V3	1	1	1	0	1	0	0	0	0	0	0	0	0	0	0	4	11
	Imagen-3	1	0	1	1	0	1	0	0	0	0	0	1	1	0	0	6	9
	Dall-E 3	1	1	1	1	0	0	1	0	1	0	0	1	0	0	0	7	8
	Grok 3	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	3	12
	Gemini 2.0 Flash	1	1	1	0	0	1	1	1	0	1	1	0	0	0	0	8	7
	FLUX 1.1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	4	11
	Firefly 3	0	1	1	0	1	0	0	0	0	0	0	0	0	0	0	3	12
	SD 3.5	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	5	10
	Doubao	1	1	1	0	0	0	1	0	0	0	0	0	0	0	0	4	11
	Qwen2.5-Max	1	1	0	1	1	0	1	0	0	0	0	0	0	0	0	5	10
	WanX2.1	1	1	0	0	0	0	1	1	0	0	0	0	0	0	0	4	11
	Kling	1	1	1	1	1	0	0	0	0	0	1	0	0	0	0	6	9
	Star-3 Alpha	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	5	10
	Hunyuan	0	1	0	1	1	0	1	0	0	0	0	0	0	0	0	4	11
	GLM-4	1	1	1	1	0	0	0	0	0	0	1	0	0	0	0	5	10

Table 13: Counting Humans and Cats in City Scene Results.

1674
1675
1676
1677

Objects	Model	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	Correct	Wrong
Triangles	Recraft V3	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	15
	Imagen-3	1	0	1	0	0	1	0	0	0	0	0	0	0	0	0	3	12
	Dall-E 3	1	0	1	0	0	0	1	0	0	0	0	0	0	0	0	3	12
	Grok 3	1	1	0	0	1	0	1	0	0	0	0	0	0	0	0	4	11
	Gemini 2.0 Flash	1	1	1	0	0	0	0	0	0	0	0	1	1	0	0	5	10
	FLUX 1.1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	3	12
	Firefly 3	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	1	14
	SD 3.5	0	1	1	1	0	0	0	0	0	0	0	0	0	0	0	3	12
	Doubao	0	0	0	0	0	1	0	1	1	0	0	0	0	0	0	3	12
	Qwen2.5-Max	1	0	1	1	0	1	0	1	0	0	0	0	0	0	0	5	10
	WanX2.1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	3	12
	Kling	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	14
	Star-3 Alpha	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	3	12
	Hunyuan	0	1	0	1	0	1	0	0	0	0	0	0	0	0	0	3	12
	GLM-4	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	15
Chairs	Recraft V3	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	14
	Imagen-3	1	0	0	1	0	1	0	0	1	0	0	0	0	0	0	4	11
	Dall-E 3	1	0	1	0	0	0	0	0	0	0	0	0	1	0	0	3	12
	Grok 3	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	14
	Gemini 2.0 Flash	1	0	1	0	0	0	0	0	0	0	0	0	0	0	0	2	13
	FLUX 1.1	0	1	1	1	0	1	0	0	1	0	0	0	0	0	0	5	10
	Firefly 3	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	15
	SD 3.5	0	1	1	1	0	0	1	0	1	0	1	0	0	0	0	6	9
	Doubao	0	1	1	1	1	0	1	0	0	0	0	0	0	0	0	5	10
	Qwen2.5-Max	0	1	1	1	1	0	0	0	0	0	0	0	0	0	0	4	11
	WanX2.1	0	1	1	1	1	0	0	0	0	0	0	0	0	0	0	4	11
	Kling	0	1	1	1	1	0	1	0	0	0	0	0	0	0	0	5	10
	Star-3 Alpha	1	1	1	1	1	0	1	0	0	0	0	0	0	0	0	6	9
	Hunyuan	0	1	0	0	1	1	0	0	0	0	0	0	0	0	0	3	12
	GLM-4	0	0	0	0	0	1	0	0	0	0	0	0	0	0	0	1	15

Table 14: Counting Triangles and Chairs in City Scene Results.

1701
1702
1703
1704
1705
1706
1707
1708
1709
1710

Objects	Model	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	Correct	Wrong
Trees	Recraft V3	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	15
	Imagen-3	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	15
	Dall-E 3	1	0	1	0	0	0	0	0	0	0	0	0	0	0	0	2	13
	Grok 3	1	0	0	0	1	0	0	0	0	0	0	0	0	0	0	2	13
	Gemini 2.0 Flash	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	3	12
	FLUX 1.1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	4	11
	Firefly 3	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	15
	SD 3.5	0	0	1	1	1	1	0	0	0	0	0	0	0	0	0	4	11
	Doubao	0	1	1	1	0	0	0	0	1	0	0	0	0	0	0	4	11
	Qwen2.5-Max	0	1	0	0	1	0	0	0	0	0	0	0	0	0	0	2	13
	WanX2.1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	15
	Kling	1	1	0	1	1	1	1	0	0	0	0	0	0	0	0	6	9
	Star-3 Alpha	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	1	14
	Hunyuan	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	1	14
	GLM-4	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	15

Table 15: Counting Trees in City Scene Results.

1723
1724
1725
1726
1727

1728
1729
1730
1731
1732
1733
1734
1735
1736
1737
1738
1739
1740
1741
1742
1743
1744
1745
1746
1747
1748
1749
1750
1751

Objects	Model	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	Correct	Wrong
Apples	Recraft V3	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	4	11
	Imagen-3	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	9	6
	Dall-E 3	1	1	1	1	0	1	1	0	0	0	0	0	0	0	0	6	9
	Grok 3	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	2	13
	Gemini 2.0 Flash	1	1	1	1	0	1	1	0	0	0	0	0	0	0	0	6	9
	FLUX 1.1	0	1	1	1	1	1	0	0	1	0	0	0	0	0	0	6	9
	Firefly 3	0	1	1	1	0	0	0	0	1	1	0	1	0	0	0	6	9
	SD 3.5	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	3	12
	Doubao	1	1	1	1	0	0	1	0	0	0	0	0	0	0	0	5	10
	Qwen2.5-Max	1	1	1	1	0	0	1	0	0	0	0	0	0	0	1	6	9
	WanX2.1	1	1	1	1	0	0	1	0	0	1	0	0	0	0	0	6	9
	Kling	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	3	12
	Star-3 Alpha	1	1	1	1	0	0	0	0	0	1	0	0	0	0	0	5	10
	Hunyuan	1	1	1	1	0	1	0	0	1	1	0	0	0	0	0	7	8
	GLM-4	1	1	1	1	1	0	0	1	1	0	0	0	0	0	0	7	8
	Watermelons	Recraft V3	1	0	1	0	0	0	0	0	0	0	0	0	0	0	0	2
Imagen-3		0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	15
Dall-E 3		1	0	0	1	0	0	0	0	0	0	0	1	0	0	0	3	12
Grok 3		1	0	1	0	0	0	0	0	0	0	0	0	0	0	0	2	13
Gemini 2.0 Flash		1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	14
FLUX 1.1		1	0	1	1	0	0	0	0	1	0	0	0	0	0	0	4	11
Firefly 3		0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	15
SD 3.5		1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	14
Doubao		1	0	0	0	1	0	0	0	0	0	0	0	0	0	0	2	13
Qwen2.5-Max		1	1	1	0	0	1	0	0	0	0	0	1	0	0	0	5	10
WanX2.1		1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	2	13
Kling		1	1	1	0	0	0	1	0	0	0	0	0	0	0	0	4	12
Star-3 Alpha		0	1	1	1	0	0	0	0	0	0	0	0	0	0	0	3	12
Hunyuan		1	1	0	1	0	1	0	0	0	0	0	0	0	0	0	4	11
GLM-4		0	0	0	1	0	0	0	0	0	0	0	1	0	0	0	2	13

Table 16: Counting Apples and Watermelons With Cartoon style Results.

1752
1753
1754
1755
1756
1757
1758
1759
1760
1761
1762
1763
1764
1765
1766
1767
1768
1769
1770
1771
1772
1773
1774
1775
1776
1777
1778
1779
1780
1781

Objects	Model	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	Correct	Wrong
Humans	Recraft V3	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	4	11
	Imagen-3	1	1	1	1	1	0	1	0	0	1	0	0	0	0	0	7	8
	Dall-E 3	1	1	1	1	1	0	1	0	0	0	1	0	0	0	0	7	8
	Grok 3	1	1	1	0	1	0	0	1	0	0	0	0	0	0	0	5	10
	Gemini 2.0 Flash	1	1	1	1	1	1	1	1	1	0	1	1	0	0	0	11	4
	FLUX 1.1	1	1	1	1	0	1	1	1	1	0	0	1	0	0	0	9	6
	Firefly 3	1	1	1	1	1	1	0	0	0	0	1	0	0	1	0	8	7
	SD 3.5	1	1	1	1	0	0	0	1	0	0	0	0	0	1	0	6	9
	Doubao	1	1	1	1	1	0	1	1	0	1	0	0	0	0	0	8	7
	Qwen2.5-Max	1	1	1	0	1	1	0	1	0	0	0	0	0	0	0	6	9
	WanX2.1	1	1	1	1	1	1	1	1	1	0	1	1	0	1	0	12	3
	Kling	1	1	1	1	0	0	0	1	1	1	1	1	0	0	0	9	6
	Star-3 Alpha	1	1	1	1	0	0	0	1	0	0	0	0	0	0	1	6	9
	Hunyuan	0	1	1	1	1	0	0	1	0	0	0	0	0	0	0	5	10
	GLM-4	1	1	1	1	1	0	0	1	0	0	0	0	1	0	0	7	8
	Cats	Recraft V3	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	5
Imagen-3		1	0	1	0	0	0	1	0	0	0	0	0	0	1	1	5	10
Dall-E 3		1	1	1	1	1	1	0	0	1	0	0	1	0	0	0	8	7
Grok 3		1	1	0	0	1	0	0	0	0	0	0	0	0	0	0	3	12
Gemini 2.0 Flash		1	0	1	1	1	0	1	1	0	0	1	0	0	0	0	7	8
FLUX 1.1		0	1	1	1	1	1	1	0	0	1	0	1	0	0	0	8	7
Firefly 3		0	1	1	1	0	1	0	1	1	0	0	0	0	0	0	6	9
SD 3.5		0	1	1	1	0	1	0	0	0	0	0	0	0	0	0	4	11
Doubao		1	1	1	1	0	0	0	0	0	1	0	1	0	0	0	6	9
Qwen2.5-Max		1	1	1	1	0	1	0	1	0	0	0	0	1	1	0	8	7
WanX2.1		1	1	1	1	1	0	1	0	0	0	0	0	0	0	0	6	9
Kling		1	1	1	1	1	0	0	0	0	0	0	1	0	0	0	6	9
Star-3 Alpha		1	1	1	0	1	1	1	1	1	1	0	0	0	0	0	9	6
Hunyuan		1	1	1	1	1	0	0	0	0	1	1	0	0	0	0	7	8
GLM-4		1	1	1	1	1	1	1	0	1	0	0	1	0	0	0	9	6

Table 17: Counting Humans and Cats With Cartoon style Results.

Objects	Model	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	Correct	Wrong
Triangles	Recraft V3	0	0	0	1	0	1	1	0	0	0	0	0	0	0	0	3	12
	Imagen-3	1	1	1	1	0	0	1	0	0	1	0	0	0	0	0	6	9
	Dall-E 3	0	1	1	1	0	0	0	0	0	0	0	0	0	0	0	3	12
	Grok 3	1	1	0	1	0	0	0	0	0	0	0	0	0	0	0	3	12
	Gemini 2.0 Flash	1	1	1	1	1	1	1	0	1	0	0	1	0	1	0	10	5
	FLUX 1.1	1	1	1	0	0	0	1	0	0	0	0	0	0	0	0	4	11
	Firefly 3	0	1	1	0	0	0	0	0	1	0	0	0	0	0	0	3	12
	SD 3.5	1	1	1	0	0	0	0	0	1	0	0	0	0	0	0	4	11
	Doubao	1	1	1	1	0	0	1	0	1	0	0	1	0	0	0	7	8
	Qwen2.5-Max	1	1	0	1	0	1	0	0	0	0	0	0	0	0	0	4	11
	WanX2.1	1	1	1	0	0	1	0	0	0	0	0	0	0	0	0	4	11
	Kling	1	1	1	0	0	1	0	0	0	0	0	0	0	0	0	4	11
	Star-3 Alpha	1	1	1	0	1	0	1	0	1	0	0	0	0	0	1	7	8
	Hunyuan	1	1	1	0	1	0	0	0	0	0	0	0	0	0	0	4	11
	GLM-4	1	1	0	0	0	0	0	0	1	0	0	0	0	0	0	3	12
Chairs	Recraft V3	1	1	1	1	1	1	1	0	0	0	0	0	0	0	0	7	8
	Imagen-3	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	3	12
	Dall-E 3	1	1	1	1	1	0	1	0	0	0	0	1	0	0	0	7	8
	Grok 3	1	1	0	1	0	0	0	0	0	0	0	0	0	0	0	3	12
	Gemini 2.0 Flash	1	1	1	0	1	0	0	0	0	1	0	1	1	0	0	7	8
	FLUX 1.1	1	1	1	1	1	1	0	1	0	0	0	0	0	0	0	7	8
	Firefly 3	0	1	1	0	1	0	1	0	1	0	0	0	0	0	0	5	10
	SD 3.5	1	1	1	0	0	1	0	0	0	0	0	0	0	0	0	4	11
	Doubao	1	1	1	0	1	1	0	0	0	0	0	0	0	0	0	5	10
	Qwen2.5-Max	1	1	1	0	0	0	0	1	0	0	0	0	0	0	0	4	11
	WanX2.1	1	1	1	1	1	0	0	0	1	0	0	0	0	1	0	7	8
	Kling	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	5	10
	Star-3 Alpha	1	1	1	1	0	0	0	1	1	0	0	0	0	0	0	6	9
	Hunyuan	1	1	1	1	1	0	0	1	1	0	0	0	0	0	0	7	8
	GLM-4	1	1	1	1	1	1	0	1	1	1	0	0	0	0	1	10	5

Table 18: Counting Triangles and Chairs With Cartoon style Results.

1805
1806
1807
1808
1809
1810
1811
1812
1813
1814
1815
1816
1817
1818
1819
1820
1821
1822
1823
1824
1825
1826
1827
1828
1829
1830
1831
1832
1833
1834
1835

Objects	Model	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	Correct	Wrong
1836	Flowers	Recraft V3	1	0	1	0	0	0	0	0	0	0	0	0	0	0	2	13
1837		Imagen-3	1	1	1	0	1	1	0	0	0	0	0	0	0	1	6	9
1838		Dall-E 3	1	1	1	1	0	1	0	0	0	1	0	0	0	0	6	9
1839		Grok 3	1	0	0	0	0	0	1	0	0	0	0	0	0	0	2	13
1840		Gemini 2.0 Flash	1	1	1	1	0	0	1	1	0	0	1	0	0	0	7	8
1841		FLUX 1.1	1	1	1	0	0	0	1	0	0	0	0	0	0	0	4	11
1842		Firefly 3	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	15
1843		SD 3.5	0	0	1	0	0	0	0	0	0	0	0	0	0	0	1	14
1844		Doubao	1	1	1	0	0	1	0	0	0	0	0	0	0	0	4	11
1845		Qwen2.5-Max	0	0	0	0	0	0	1	0	0	0	0	0	0	0	1	14
1846		WanX2.1	1	1	1	0	0	1	0	0	0	0	0	0	0	0	4	11
1847		Kling	0	0	1	0	1	1	0	0	0	0	1	1	0	0	5	10
1848		Star-3 Alpha	1	1	1	0	1	0	0	1	0	0	0	0	0	0	5	10
1849		Hunyuan	1	1	1	0	0	0	0	0	1	0	0	0	1	1	6	9
1850		GLM-4	1	1	1	0	0	0	0	0	0	0	0	0	0	1	4	11

Table 19: Counting Flowers With Cartoon Style Results.

Objects	Model	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	Correct	Wrong	
1851	Apples	Recraft V3	1	1	1	0	1	1	0	1	0	0	0	0	0	0	6	9	
1852		Imagen-3	1	1	0	0	0	0	0	0	1	0	0	0	0	0	1	4	11
1853		Dall-E 3	1	1	0	0	0	0	1	0	0	0	0	0	0	0	0	3	12
1854		Grok 3	1	1	0	1	0	0	0	0	0	0	0	0	0	0	0	3	12
1855		Gemini 2.0 Flash	1	1	1	0	0	1	0	0	1	0	0	0	0	0	0	5	10
1856		FLUX 1.1	1	1	1	0	1	0	1	0	0	0	0	0	0	0	0	5	10
1857		Firefly 3	1	1	1	1	1	1	0	0	1	0	0	0	0	0	1	8	7
1858		SD 3.5	1	1	0	0	1	0	0	0	0	0	0	0	0	0	0	3	12
1859		Doubao	1	1	1	1	0	0	1	0	0	1	0	0	0	0	0	6	9
1860		Qwen2.5-Max	1	1	0	0	1	0	0	0	0	1	0	0	0	0	0	4	11
1861		WanX2.1	1	1	1	1	1	1	1	1	1	0	0	1	0	0	0	10	5
1862		Kling	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	6	9
1863		Star-3 Alpha	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	6	9
1864		Hunyuan	1	1	1	1	0	0	1	1	1	0	0	1	0	0	0	8	7
1865		GLM-4	1	1	1	1	1	0	1	0	0	0	0	0	0	0	1	7	8
1866	Watermelons	Recraft V3	1	0	0	0	1	0	0	0	0	0	0	0	0	0	2	13	
1867		Imagen-3	0	0	0	1	1	0	0	0	0	0	0	0	0	0	0	2	13
1868		Dall-E 3	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	14
1869		Grok 3	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	4	11
1870		Gemini 2.0 Flash	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	15
1871		FLUX 1.1	1	0	0	0	0	0	0	1	0	0	0	0	0	0	0	2	13
1872		Firefly 3	0	0	0	0	0	0	0	0	0	0	0	1	0	0	0	1	14
1873		SD 3.5	0	1	0	0	0	0	0	0	0	0	0	1	0	0	0	2	13
1874		Doubao	0	1	1	1	1	1	0	0	0	0	0	0	0	0	0	5	10
1875		Qwen2.5-Max	1	0	0	1	0	0	0	0	0	0	0	0	0	0	1	3	12
1876	WanX2.1	1	1	1	1	0	0	1	0	1	0	0	0	0	0	0	6	9	
1877	Kling	1	1	1	1	0	1	0	0	1	0	0	0	0	0	0	6	9	
1878	Star-3 Alpha	1	1	1	1	1	1	1	1	0	0	1	0	0	0	0	9	6	
1879	Hunyuan	0	1	0	0	1	0	0	1	0	0	0	0	0	0	0	3	12	
1880	GLM-4	0	0	0	1	0	0	1	0	0	0	0	0	0	0	0	2	13	

Table 20: Counting Apples and Watermelons With Watercolor style Results.

1890
1891
1892
1893
1894
1895
1896
1897
1898
1899
1900
1901
1902
1903
1904
1905
1906
1907
1908
1909
1910
1911
1912
1913
1914
1915
1916
1917
1918
1919
1920
1921
1922
1923
1924
1925
1926
1927
1928
1929
1930
1931
1932
1933
1934
1935
1936
1937
1938
1939
1940
1941
1942
1943

Objects	Model	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	Correct	Wrong
Humans	Recraft V3	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	6	9
	Imagen-3	1	1	1	1	1	1	1	1	1	0	0	0	0	0	1	10	5
	Dall-E 3	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	4	11
	Grok 3	1	1	1	1	1	1	0	1	0	0	1	0	0	0	0	8	7
	Gemini 2.0 Flash	1	1	1	1	1	0	1	1	1	0	0	1	0	0	0	9	6
	FLUX 1.1	1	1	1	1	1	0	0	1	0	0	1	1	0	0	0	8	7
	Firefly 3	1	1	1	1	1	0	0	0	1	1	1	0	0	0	0	8	7
	SD 3.5	1	1	1	1	1	1	0	1	0	0	0	0	0	0	0	7	8
	Doubao	1	1	1	1	1	1	1	0	0	0	0	0	0	0	0	7	8
	Qwen2.5-Max	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	3	12
	WanX2.1	1	1	1	1	1	1	1	0	0	0	0	0	1	0	0	8	7
	Kling	1	1	1	1	1	1	1	0	0	0	1	0	0	0	0	8	7
	Star-3 Alpha	1	1	1	1	1	1	0	1	1	0	0	0	0	0	0	8	7
	Hunyuan	1	1	1	1	0	0	1	1	0	0	0	0	0	0	0	6	9
	GLM-4	1	1	1	1	1	1	1	0	0	0	0	0	0	0	0	7	8
	Cats	Recraft V3	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	3
Imagen-3		1	1	1	1	1	1	1	1	1	1	0	1	0	1	1	13	2
Dall-E 3		1	1	1	0	1	1	1	0	0	1	0	0	0	0	0	7	8
Grok 3		1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	3	12
Gemini 2.0 Flash		1	1	1	1	1	1	0	0	1	1	0	0	0	0	1	9	6
FLUX 1.1		1	1	1	1	1	0	1	0	0	0	0	0	0	0	0	6	9
Firefly 3		1	1	1	1	1	1	1	1	1	0	0	1	1	0	1	12	3
SD 3.5		1	1	1	1	0	1	0	0	1	0	0	0	0	0	0	6	9
Doubao		1	1	1	1	1	1	1	0	1	0	0	0	0	0	0	8	7
Qwen2.5-Max		1	1	1	0	1	0	0	0	0	1	0	0	0	0	0	5	10
WanX2.1		1	1	1	1	1	1	0	0	1	0	0	0	0	1	0	8	7
Kling		1	1	1	1	1	1	1	0	1	0	0	0	0	0	0	8	7
Star-3 Alpha		1	1	1	1	0	1	1	0	0	0	0	0	0	0	0	6	9
Hunyuan		1	1	1	1	1	1	1	0	1	0	0	0	0	0	0	8	7
GLM-4		1	1	1	1	1	0	0	0	0	0	0	1	0	0	0	6	9

Table 21: Counting Humans and Cats With Watercolor style Results.

Objects	Model	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	Correct	Wrong	
1944	Triangles	Recraft V3	1	0	0	1	0	0	0	0	0	0	0	0	0	0	2	13	
1945		Imagen-3	1	1	1	1	1	0	0	1	1	0	0	0	0	0	1	8	7
1946		Dall-E 3	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	2	13
1947		Grok 3	1	1	0	1	1	1	1	0	1	1	1	0	0	0	1	10	5
1948		Gemini 2.0 Flash	1	1	1	1	0	1	0	0	0	0	0	1	0	0	0	6	9
1949		FLUX 1.1	1	1	1	0	1	0	0	0	0	0	0	0	0	0	0	4	11
1950		Firefly 3	1	1	1	1	0	1	0	0	1	0	0	0	0	0	0	6	9
1951		SD 3.5	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	1	14
1952		Doubao	1	1	1	1	1	1	0	0	1	0	0	1	0	0	1	9	6
1953		Qwen2.5-Max	1	1	1	1	1	0	0	0	0	0	0	0	0	1	0	6	9
1954		WanX2.1	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	6	9
1955		Kling	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	3	12
1956		Star-3 Alpha	1	1	1	1	1	0	0	1	0	0	0	0	0	0	1	7	8
1957		Hunyuan	1	1	0	1	1	0	0	1	1	1	1	1	0	0	1	10	5
1958		GLM-4	1	1	0	0	0	1	0	0	1	0	0	0	0	0	0	4	11
1959		Chairs	Recraft V3	1	1	1	1	1	0	0	0	0	0	0	0	0	0	5	10
1960	Imagen-3		1	1	1	1	1	1	1	1	1	0	0	0	1	1	0	11	4
1961	Dall-E 3		1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	5	10
1962	Grok 3		1	1	1	0	1	0	0	0	0	0	0	0	1	0	0	5	10
1963	Gemini 2.0 Flash		1	1	1	1	0	0	0	0	0	0	1	0	0	1	0	6	9
1964	FLUX 1.1		1	1	1	1	1	1	0	1	0	0	0	0	0	0	0	7	8
1965	Firefly 3		1	1	1	1	1	0	0	0	1	0	0	0	0	0	0	6	9
1966	SD 3.5		1	1	0	1	0	0	0	0	0	0	0	0	0	0	0	3	12
1967	Doubao		1	1	1	1	0	1	1	0	0	1	0	0	0	0	0	7	8
1968	Qwen2.5-Max		1	1	1	0	1	1	0	0	0	0	0	0	0	0	0	5	10
1969	WanX2.1		1	1	1	1	0	1	0	0	0	0	0	0	0	0	1	6	9
1970	Kling		1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	5	10
1971	Star-3 Alpha		1	1	1	1	1	1	1	1	0	0	0	0	0	0	0	8	7
1972	Hunyuan		1	1	1	1	1	0	1	1	0	0	0	0	0	0	1	9	6
1973	GLM-4		1	1	1	1	0	1	0	0	0	0	0	0	0	0	0	5	10

Table 22: Counting Triangles and Chairs With Watercolor style Results.

1967
1968
1969
1970
1971
1972
1973
1974
1975
1976
1977
1978
1979
1980
1981
1982
1983
1984
1985
1986
1987
1988
1989
1990
1991
1992
1993
1994
1995
1996
1997

1998	Objects	Model	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	Correct	Wrong
1999	Flowers	Recraft V3	1	1	1	1	0	1	0	0	0	0	0	0	1	0	0	6	9
2000		Imagen-3	1	1	0	1	1	0	1	0	0	1	1	1	0	0	1	9	6
2001		Dall-E 3	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	15
2001		Grok 3	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	3	12
2002		Gemini 2.0 Flash	1	1	1	0	0	1	0	0	0	0	0	0	0	0	0	4	11
2002		FLUX 1.1	1	1	1	1	1	0	1	1	0	0	0	0	0	0	0	7	8
2003		Firefly 3	1	1	1	1	1	0	0	1	0	0	0	0	0	0	0	6	9
2004		SD 3.5	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	3	12
2005		Doubao	1	1	1	1	1	0	0	0	0	1	1	0	0	0	0	7	8
2005		Qwen2.5-Max	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	4	11
2006		WanX2.1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	14
2007		Kling	1	1	1	1	0	1	0	0	1	0	0	0	1	1	0	8	7
2007		Star-3 Alpha	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	5	10
2008		Hunyuan	0	1	1	1	1	1	1	0	1	0	0	1	1	0	0	9	6
2009		GLM-4	0	1	1	0	1	0	0	0	0	0	0	0	0	0	0	3	12

Table 23: Counting Flowers With Watercolor Style Results.

2010
2011
2012
2013
2014
2015
2016
2017
2018
2019
2020
2021
2022
2023
2024
2025
2026
2027
2028
2029
2030
2031
2032
2033
2034
2035
2036
2037
2038
2039
2040
2041
2042
2043
2044
2045
2046
2047
2048
2049
2050
2051

Objects	Model	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	Correct	Wrong	
2052	Apples	Recraft V3	1	1	1	1	0	0	1	0	1	0	0	0	0	0	6	9	
2053		Imagen-3	1	1	1	1	1	0	0	0	0	0	0	0	0	0	5	10	
2054		Dall-E 3	1	1	1	0	1	0	0	0	0	0	0	1	0	0	5	10	
2055		Grok 3	1	1	1	0	0	0	0	0	0	0	0	0	0	0	3	12	
2056		Gemini 2.0 Flash	1	1	1	1	0	0	1	0	0	0	0	0	0	0	5	10	
2057		FLUX 1.1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	2	13	
2058		Firefly 3	1	0	0	1	0	0	0	0	0	0	0	0	0	0	2	13	
2059		SD 3.5	1	1	0	0	0	0	0	0	0	0	0	0	0	0	2	13	
2060		Doubao	1	1	1	0	0	0	1	0	0	0	0	0	0	0	4	11	
2061		Qwen2.5-Max	1	1	1	0	0	0	0	0	0	0	0	0	0	0	3	12	
2062		WanX2.1	1	1	1	0	0	0	1	0	1	0	0	0	0	0	5	10	
2063		Kling	1	0	1	0	1	0	1	0	0	0	0	0	0	0	4	11	
2064		Star-3 Alpha	1	1	1	0	1	0	0	0	0	0	1	0	0	0	5	10	
2065		Hunyuan	1	1	1	0	0	0	0	0	0	0	0	0	0	0	3	12	
2066		GLM-4	1	1	1	1	1	1	1	0	0	0	0	0	0	0	7	8	
2067		Watermelons	Recraft V3	1	1	1	1	0	0	1	0	1	0	0	0	0	0	6	9
2068			Imagen-3	1	1	1	1	1	0	1	0	0	0	0	0	0	0	6	9
2069			Dall-E 3	1	1	1	1	1	0	1	0	0	1	1	1	1	0	10	5
2070			Grok 3	1	1	1	0	0	0	0	0	0	0	0	0	0	0	3	12
2071			Gemini 2.0 Flash	1	1	1	0	0	0	1	0	0	0	0	0	0	0	4	11
2072	FLUX 1.1		0	0	1	0	0	0	0	0	0	0	0	0	0	0	1	14	
2073	Firefly 3		0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	15	
2074	SD 3.5		0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	15	
2075	Doubao		0	1	1	0	0	0	0	0	0	0	0	0	0	0	2	13	
2076	Qwen2.5-Max		1	1	1	1	0	0	0	0	0	0	0	0	0	0	4	11	
2077	WanX2.1		1	0	0	0	1	0	0	0	0	0	0	0	1	0	3	12	
2078	Kling		0	0	1	0	0	0	0	0	0	0	0	0	0	0	1	14	
2079	Star-3 Alpha		1	1	0	0	0	0	1	0	0	0	1	0	0	0	4	11	
2080	Hunyuan		1	1	1	0	1	0	0	0	1	0	0	0	0	0	5	10	
2081	GLM-4	0	1	1	1	1	0	0	1	0	0	0	0	0	0	5	10		

Table 24: Counting Apples and Watermelons with Multiplicative Decomposition Results.

Objects	Model	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	Correct	Wrong	
2077	Humans	Recraft V3	1	1	1	0	1	0	0	0	0	0	0	0	0	0	4	11	
2078		Imagen-3	1	1	1	0	1	0	1	0	0	0	1	0	0	0	6	9	
2079		Dall-E 3	1	1	1	0	1	0	1	0	0	0	0	0	0	0	5	10	
2080		Grok 3	1	0	0	1	0	1	1	0	0	0	0	0	0	1	5	10	
2081		Gemini 2.0 Flash	1	1	1	1	1	0	1	0	0	0	1	0	0	0	7	8	
2082		FLUX 1.1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	2	13	
2083		Firefly 3	1	0	0	0	0	0	0	0	0	0	0	0	0	0	1	14	
2084		SD 3.5	1	1	1	0	0	0	0	0	0	0	0	0	0	0	3	12	
2085		Doubao	0	1	1	0	1	0	0	0	0	0	1	0	0	0	4	11	
2086		Qwen2.5-Max	1	1	0	1	0	0	0	1	0	0	1	0	0	0	5	10	
2087		WanX2.1	1	1	1	1	1	0	0	1	0	0	1	0	0	0	7	8	
2088		Kling	0	0	0	0	1	0	0	0	0	0	0	0	0	0	1	14	
2089		Star-3 Alpha	0	1	1	0	1	0	0	0	0	0	0	0	0	0	3	12	
2090		Hunyuan	1	1	1	0	0	0	1	0	0	0	0	0	0	0	4	11	
2091		GLM-4	1	1	1	1	1	0	0	0	0	0	0	0	0	0	5	10	
2092		Cats	Recraft V3	1	1	1	1	0	1	0	0	0	0	0	0	0	5	10	
2093			Imagen-3	1	1	0	0	1	1	0	1	0	1	0	0	0	0	6	9
2094			Dall-E 3	1	1	1	1	1	1	0	0	0	0	0	0	0	0	7	8
2095			Grok 3	1	1	1	0	0	0	0	0	0	0	0	0	0	0	3	12
2096			Gemini 2.0 Flash	1	1	1	0	0	1	1	1	0	0	0	0	0	0	6	9
2097	FLUX 1.1		1	1	1	0	1	0	0	0	0	0	0	0	0	0	4	11	
2098	Firefly 3		1	1	1	0	0	0	0	0	0	0	0	0	0	0	3	12	
2099	SD 3.5		1	1	1	0	1	0	0	0	0	0	0	0	0	0	4	11	
2100	Doubao		1	1	1	0	0	0	0	0	0	0	0	0	0	0	3	12	
2101	Qwen2.5-Max		1	1	1	0	0	0	0	0	0	0	1	0	0	0	4	11	
2102	WanX2.1		1	1	1	1	1	0	1	0	0	0	0	0	0	0	6	9	
2103	Kling		1	1	1	0	1	0	0	0	0	0	0	0	0	0	4	11	
2104	Star-3 Alpha		1	1	1	0	1	0	1	0	0	0	0	1	0	0	6	9	
2105	Hunyuan		1	1	1	0	0	0	1	0	0	0	0	1	0	0	5	10	
2106	GLM-4	1	1	1	0	0	0	1	1	0	0	0	1	0	0	6	9		

Table 25: Counting Humans and Cats with Multiplicative Decomposition Results.

Objects	Model	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	Correct	Wrong	
2106	Triangles	Recraft V3	0	0	1	0	0	0	0	0	0	0	0	0	0	0	1	14	
2107		Imagen-3	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	14
2108		Dall-E 3	1	1	0	0	0	1	0	0	0	0	0	0	0	0	0	3	12
2109		Grok 3	1	1	1	1	1	0	1	0	0	0	0	0	0	0	0	5	10
2110		Gemini 2.0 Flash	1	1	0	1	1	0	1	1	0	0	0	0	0	1	0	7	8
2111		FLUX 1.1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	2	13
2112		Firefly 3	1	0	1	0	0	0	0	0	0	0	0	0	0	0	0	2	13
2112		SD 3.5	1	0	1	0	0	0	0	0	0	0	0	0	0	0	0	2	13
2113		Doubao	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	1	14
2113		Qwen2.5-Max	1	1	0	0	0	0	1	0	0	0	0	0	0	0	0	3	12
2114		WanX2.1	1	0	1	0	1	0	0	0	0	0	0	0	0	0	0	3	12
2115		Kling	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	2	13
2116		Star-3 Alpha	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	1	14
2116		Hunyuan	0	1	1	0	0	0	0	0	0	0	0	0	0	0	0	2	12
2117		GLM-4	1	0	0	0	0	0	0	0	0	0	1	0	0	0	0	2	13
2118		Chairs	Recraft V3	1	1	1	1	1	0	0	0	0	0	0	0	0	0	5	10
2119	Imagen-3		1	1	1	1	1	0	0	0	0	0	0	1	0	0	0	6	9
2119	Dall-E 3		1	1	1	1	0	0	0	0	0	1	1	0	0	0	0	6	9
2120	Grok 3		1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	4	11
2121	Gemini 2.0 Flash		1	1	1	1	0	0	0	0	0	1	0	0	1	0	0	6	9
2122	FLUX 1.1		1	1	1	0	1	0	0	0	0	0	0	0	0	0	0	4	11
2122	Firefly 3		1	0	1	0	0	0	0	0	0	0	0	0	0	0	0	2	13
2123	SD 3.5		1	1	1	0	1	0	0	0	0	0	0	0	0	0	0	4	11
2124	Doubao		1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	3	12
2124	Qwen2.5-Max		1	1	1	1	1	1	1	0	0	0	0	0	0	0	0	7	8
2125	WanX2.1		1	1	1	1	1	1	1	0	0	0	0	0	0	0	0	7	8
2126	Kling		1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	3	12
2126	Star-3 Alpha		1	1	1	0	1	0	0	0	0	0	0	0	0	0	0	4	11
2127	Hunyuan		1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	5	10
2128	GLM-4		1	1	1	1	1	0	0	1	0	0	0	0	0	0	1	7	8

Table 26: Counting Triangles and Chairs with Multiplicative Decomposition Results.

Objects	Model	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	Correct	Wrong	
2132	Flowers	Recraft V3	1	1	1	0	0	0	0	0	0	0	0	0	0	0	3	12	
2133		Imagen-3	1	1	1	0	1	0	0	0	0	0	0	1	0	0	0	5	10
2134		Dall-E 3	1	1	1	0	0	0	1	0	0	0	0	0	0	0	0	4	11
2135		Grok 3	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	4	11
2136		Gemini 2.0 Flash	1	0	0	0	1	0	1	0	1	1	0	0	1	0	0	6	9
2136		FLUX 1.1	1	0	0	0	0	0	1	0	0	0	0	0	0	0	0	2	13
2137		Firefly 3	0	0	0	0	0	1	0	0	0	0	0	0	0	0	0	1	14
2138		SD 3.5	1	1	1	0	0	0	0	0	0	0	0	0	1	0	0	4	11
2138		Doubao	1	1	1	0	0	0	0	0	0	1	0	0	0	0	0	4	11
2139		Qwen2.5-Max	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	2	13
2140		WanX2.1	1	1	1	1	0	0	0	0	1	0	1	0	0	0	0	6	9
2141		Kling	1	0	1	1	0	0	0	0	0	0	0	0	0	0	0	3	12
2141		Star-3 Alpha	1	1	0	1	0	0	0	0	0	0	0	0	0	0	0	3	12
2142		Hunyuan	1	0	1	1	0	0	1	0	0	0	0	1	0	0	0	5	10
2143		GLM-4	1	1	0	0	1	1	0	1	0	1	0	0	0	0	0	6	9

Table 27: Counting Flowers with Multiplicative Decomposition Results.

2160

2161

2162

2163

2164

2165

2166

2167

2168

2169

2170

2171

2172

2173

2174

2175

2176

2177

2178

2179

2180

2181

2182

2183

Objects	Model	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	Correct	Wrong	
Apples	Recraft V3	1	1	1	1	0	1	0	0	0	0	0	0	0	0	0	5	10	
	Imagen-3	1	1	1	1	0	1	0	0	0	0	0	0	0	0	0	5	10	
	Dall-E 3	0	0	0	1	0	0	0	0	0	0	0	0	0	0	0	1	14	
	Grok 3	1	1	1	0	0	0	0	1	0	1	0	0	1	1	0	7	8	
	Gemini 2.0 Flash	1	0	1	0	0	0	0	0	0	0	1	0	0	0	0	3	12	
	FLUX 1.1	1	1	1	1	1	0	1	0	1	0	0	0	0	0	0	7	8	
	Firefly 3	1	0	0	0	1	0	0	0	0	0	0	0	0	0	0	2	13	
	SD 3.5	1	1	0	0	1	0	0	0	0	0	0	0	0	0	0	3	12	
	Doubao	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	2	13	
	Qwen2.5-Max	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	1	4	11
	WanX2.1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	3	12	
	Kling	1	0	1	0	0	0	0	0	0	0	0	0	0	0	0	2	13	
	Star-3 Alpha	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	2	13	
	Hunyuan	1	1	1	1	1	0	0	0	1	0	0	1	0	0	0	7	8	
	GLM-4	1	1	1	0	1	0	0	0	0	0	0	0	0	1	0	5	10	
Watermelons	Recraft V3	1	0	0	0	0	0	0	0	0	1	0	0	0	0	0	2	13	
	Imagen-3	1	0	1	1	1	0	0	0	1	0	0	0	0	0	0	5	10	
	Dall-E 3	1	1	0	0	0	1	0	0	0	0	0	0	0	0	0	3	12	
	Grok 3	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	4	11	
	Gemini 2.0 Flash	1	1	1	1	1	0	0	0	0	0	0	0	0	1	0	6	9	
	FLUX 1.1	1	0	1	0	1	1	0	0	0	0	0	0	0	0	0	4	11	
	Firefly 3	1	0	1	0	0	0	0	0	0	0	0	0	0	0	0	2	13	
	SD 3.5	1	0	0	1	0	1	1	0	0	0	0	0	0	0	0	4	11	
	Doubao	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	2	13	
	Qwen2.5-Max	1	1	0	0	0	0	0	1	0	0	0	0	0	0	0	3	12	
	WanX2.1	1	0	1	0	1	0	0	0	0	0	0	0	0	0	0	3	12	
	Kling	1	1	1	0	1	0	0	0	0	0	0	0	0	0	0	4	11	
	Star-3 Alpha	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	14	
	Hunyuan	1	1	0	0	0	0	1	0	0	0	0	0	0	0	0	3	12	
	GLM-4	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	5	10	

Table 28: Counting Apples and Watermelons With Additive Decomposition Results.

2184

2185

2186

2187

2188

2189

2190

2191

2192

2193

2194

2195

2196

2197

2198

2199

2200

2201

2202

2203

2204

2205

2206

2207

2208

2209

2210

2211

2212

2213

Objects	Model	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	Correct	Wrong
Humans	Recraft V3	1	1	1	1	1	0	1	0	0	0	0	0	0	0	0	6	9
	Imagen-3	1	1	1	1	1	1	1	0	1	0	0	0	0	0	0	8	7
	Dall-E 3	1	1	1	0	0	0	0	0	0	0	1	0	0	0	0	4	11
	Grok 3	1	1	1	1	1	1	0	1	1	0	0	0	1	1	0	10	5
	Gemini 2.0 Flash	1	1	1	1	0	0	0	0	1	0	1	0	0	0	0	6	9
	FLUX 1.1	1	1	0	1	0	0	0	0	0	0	0	0	0	0	0	3	12
	Firefly 3	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	3	12
	SD 3.5	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	14
	Doubao	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	14
	Qwen2.5-Max	1	1	0	0	0	0	0	0	0	0	1	0	0	0	0	3	12
	WanX2.1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	3	12
	Kling	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	14
	Star-3 Alpha	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	1	14
	Hunyuan	1	0	0	0	0	0	1	0	0	0	0	0	0	0	0	2	13
	GLM-4	1	1	1	1	1	0	1	0	0	1	0	0	0	0	0	7	8
Cats	Recraft V3	1	1	0	1	0	0	0	0	0	0	0	0	0	0	0	3	12
	Imagen-3	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	2	13
	Dall-E 3	1	1	1	1	0	0	0	0	0	1	0	0	0	0	0	5	10
	Grok 3	1	1	1	0	1	1	0	0	0	0	0	0	0	0	0	5	10
	Gemini 2.0 Flash	1	1	0	0	1	0	0	0	0	0	0	0	0	0	0	3	12
	FLUX 1.1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	2	13
	Firefly 3	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	4	11
	SD 3.5	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	3	12
	Doubao	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	2	13
	Qwen2.5-Max	1	1	1	1	1	1	0	0	0	0	0	0	1	0	0	7	8
	WanX2.1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	2	13
	Kling	1	0	1	0	0	0	0	0	0	0	0	0	0	0	0	2	13
	Star-3 Alpha	1	1	1	0	0	0	1	0	1	0	0	0	0	0	0	5	10
	Hunyuan	1	0	1	0	1	1	0	0	0	0	0	0	0	0	0	4	11
	GLM-4	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	4	11

Table 29: Counting Humans and Cats With Additive Decomposition Results.

2214
2215
2216
2217

Objects	Model	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	Correct	Wrong
Triangles	Recraft V3	1	0	1	0	0	0	0	0	0	0	0	0	0	0	0	2	13
	Imagen-3	1	0	0	0	0	0	0	1	0	0	0	0	0	0	0	2	13
	Dall-E 3	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	14
	Grok 3	1	1	1	1	0	0	0	0	0	1	0	0	1	0	0	6	9
	Gemini 2.0 Flash	1	0	1	0	1	0	0	0	0	0	0	0	0	1	0	4	11
	FLUX 1.1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	2	13
	Firefly 3	1	0	0	0	0	1	0	0	0	0	0	0	0	0	0	2	13
	SD 3.5	0	0	0	1	0	0	0	0	0	0	0	0	0	0	0	1	14
	Doubao	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	2	13
	Qwen2.5-Max	1	1	0	1	0	0	0	0	1	0	0	0	0	0	0	4	11
	WanX2.1	1	0	1	0	1	0	0	0	0	0	0	0	0	0	0	3	12
	Kling	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	15
	Star-3 Alpha	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	14
	Hunyuan	1	0	0	1	0	1	0	0	0	0	0	0	0	0	0	3	12
	GLM-4	1	1	1	1	0	1	1	0	0	0	0	0	0	0	0	6	9
Chairs	Recraft V3	1	1	1	1	0	1	0	0	1	0	0	0	0	0	6	9	
	Imagen-3	1	1	0	1	1	0	0	0	0	0	0	0	0	0	4	11	
	Dall-E 3	1	1	0	0	0	0	0	0	1	0	0	1	0	0	4	11	
	Grok 3	1	1	0	0	0	0	0	0	0	0	0	1	0	0	3	12	
	Gemini 2.0 Flash	1	1	1	1	1	1	0	0	0	0	0	0	0	0	6	9	
	FLUX 1.1	1	1	0	0	0	0	0	0	1	0	0	0	0	0	3	12	
	Firefly 3	1	1	1	0	0	0	0	0	0	0	0	0	0	0	3	12	
	SD 3.5	1	1	0	0	0	0	0	0	0	0	0	0	0	0	2	13	
	Doubao	1	1	0	0	0	0	0	0	0	0	0	0	0	0	2	13	
	Qwen2.5-Max	1	1	0	0	0	1	0	0	1	0	1	0	0	0	5	10	
	WanX2.1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	2	13	
	Kling	1	0	0	0	0	0	0	0	0	0	0	0	0	0	1	14	
	Star-3 Alpha	1	1	0	0	0	0	0	0	0	0	0	0	0	0	2	13	
	Hunyuan	1	1	1	1	0	0	1	0	0	0	1	0	0	0	5	10	
	GLM-4	1	1	1	1	0	1	1	1	0	0	0	0	0	1	8	7	

Table 30: Counting Triangles and Chairs With Additive Decomposition Results.

2241
2242
2243
2244
2245
2246
2247
2248
2249
2250

Objects	Model	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	Correct	Wrong
Flowers	Recraft V3	1	1	0	1	0	0	0	0	0	0	0	0	0	0	3	12	
	Imagen-3	1	0	0	0	0	0	0	0	0	1	0	1	0	1	0	4	11
	Dall-E 3	1	1	1	0	0	0	0	0	0	0	0	0	0	0	3	12	
	Grok 3	1	1	0	0	0	0	0	0	0	0	0	0	0	0	2	13	
	Gemini 2.0 Flash	1	1	1	0	0	1	0	0	0	0	0	0	0	0	4	11	
	FLUX 1.1	1	0	1	0	0	0	0	0	0	1	0	1	0	0	4	11	
	Firefly 3	1	0	0	0	0	0	0	0	0	0	0	0	0	0	1	14	
	SD 3.5	1	0	1	0	0	0	0	0	0	0	0	0	0	0	2	13	
	Doubao	1	1	0	0	0	0	0	0	0	0	0	0	0	0	2	13	
	Qwen2.5-Max	1	1	0	0	0	0	0	0	0	0	0	0	0	0	2	13	
	WanX2.1	1	1	0	0	0	0	1	0	0	0	0	0	0	1	0	4	11
	Kling	0	0	1	0	0	0	0	0	0	0	0	0	0	0	1	14	
	Star-3 Alpha	1	1	1	0	0	0	1	0	0	0	0	0	0	0	4	11	
	Hunyuan	0	0	0	1	0	0	0	0	1	0	0	0	0	0	2	13	
	GLM-4	0	1	1	0	0	0	0	0	0	0	0	0	0	0	2	13	

Table 31: Counting Flowers With Additive Decomposition Results.

2263
2264
2265
2266
2267

2268

2269

Objects	Model	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	Correct	Wrong
Apples	Recraft V3	1	0	0	1	0	0	0	0	1	0	0	0	0	0	0	3	12
	Imagen-3	1	0	1	1	1	1	0	0	1	1	0	0	0	0	0	7	8
	Dall-E 3	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	5	10
	Grok 3	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	3	12
	Gemini 2.0 Flash	1	1	0	1	0	1	0	0	1	0	0	1	0	0	0	6	9
	FLUX 1.1	1	1	1	1	0	0	0	0	1	0	0	0	0	0	0	5	10
	Firefly 3	1	1	1	1	1	0	0	0	1	0	0	1	0	0	1	8	7
	SD 3.5	1	1	0	0	0	1	0	0	0	0	0	0	0	0	0	3	12
	Doubao	1	1	1	1	1	0	0	0	1	0	0	0	0	0	0	6	9
	Qwen2.5-Max	1	1	1	1	1	1	1	1	1	0	0	1	0	0	0	9	6
	WanX2.1	1	0	1	1	0	0	0	0	1	0	0	0	0	0	1	5	10
	Kling	1	1	1	1	0	0	0	0	1	0	0	0	1	0	0	6	9
	Star-3 Alpha	1	1	0	1	0	0	0	0	1	0	0	0	0	0	0	4	11
	Hunyuan	1	0	0	1	0	0	0	0	1	0	0	0	0	0	0	3	12
	GLM-4	1	0	1	1	0	0	0	0	1	0	0	1	0	0	0	5	10
	Watermelons	Recraft V3	0	1	0	1	0	0	0	0	1	0	0	0	0	0	0	3
Imagen-3		1	1	1	1	0	1	0	0	1	0	0	1	0	0	0	7	8
Dall-E 3		1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	5	10
Grok 3		1	1	0	1	0	0	0	0	0	0	0	0	0	0	1	4	11
Gemini 2.0 Flash		1	1	0	1	1	0	0	0	1	0	0	1	0	0	0	6	9
FLUX 1.1		1	1	1	1	0	1	0	1	1	0	0	0	0	0	0	7	8
Firefly 3		1	0	1	0	1	0	0	0	1	0	0	0	0	0	0	4	11
SD 3.5		1	1	0	0	0	0	0	0	1	0	0	0	0	0	0	3	11
Doubao		0	1	1	1	1	1	0	0	1	0	0	0	0	1	0	7	8
Qwen2.5-Max		1	0	1	1	0	0	0	0	0	0	0	0	0	0	0	3	12
WanX2.1		1	0	1	1	0	1	0	0	1	0	0	0	0	0	0	5	10
Kling		1	0	1	0	1	0	0	1	1	0	0	0	0	0	0	5	10
Star-3 Alpha		1	1	1	1	0	0	0	0	1	0	0	1	0	0	0	6	9
Hunyuan		1	0	0	0	0	0	0	0	0	0	0	1	0	0	0	2	13
GLM-4		1	1	1	1	0	0	0	0	1	0	0	0	0	0	0	5	10

2292

Table 32: Counting Apples and Watermelons with Grid Prior Results.

2293

2294

2295

2296

Objects	Model	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	Correct	Wrong
Humans	Recraft V3	1	1	1	1	0	0	0	1	0	0	0	0	0	1	6	9	
	Imagen-3	1	1	1	1	1	1	1	0	1	0	0	1	0	1	0	10	5
	Dall-E 3	1	1	1	1	1	0	1	0	0	0	0	0	0	0	0	6	9
	Grok 3	1	1	0	1	0	0	0	1	1	0	0	0	1	0	0	6	6
	Gemini 2.0 Flash	1	1	1	1	0	1	1	0	1	0	0	0	0	1	0	8	9
	FLUX 1.1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	4	11
	Firefly 3	1	1	1	1	1	0	0	0	0	0	0	1	0	0	0	6	9
	SD 3.5	1	1	1	1	0	0	1	1	0	0	0	0	0	0	0	6	9
	Doubao	1	1	1	1	0	0	1	1	0	0	0	0	0	0	0	6	9
	Qwen2.5-Max	1	0	1	0	0	1	0	1	0	1	0	0	0	0	0	5	10
	WanX2.1	1	1	1	1	1	0	1	0	1	0	0	0	0	0	0	7	8
	Kling	0	1	1	1	0	1	0	1	1	0	0	0	1	0	0	7	8
	Star-3 Alpha	1	0	1	0	0	0	1	0	0	0	0	0	0	0	0	3	12
	Hunyuan	1	1	0	1	0	0	0	1	0	0	0	0	0	0	0	4	11
	GLM-4	1	1	1	1	1	0	0	0	1	1	0	0	0	0	0	7	8
	Cats	Recraft V3	1	1	1	1	1	0	0	0	0	0	0	0	0	0	5	10
Imagen-3		1	1	1	1	1	0	0	0	1	1	0	1	0	0	1	9	6
Dall-E 3		1	1	1	1	0	0	0	0	1	0	0	0	0	0	0	5	10
Grok 3		1	1	0	1	1	0	0	1	1	0	0	0	0	0	0	6	9
Gemini 2.0 Flash		1	0	1	1	0	0	0	0	1	0	0	1	0	0	0	5	10
FLUX 1.1		1	1	0	1	1	1	0	1	0	0	0	0	0	0	0	6	9
Firefly 3		1	1	1	1	1	1	0	1	1	0	0	0	0	0	0	8	7
SD 3.5		1	1	1	1	0	0	0	1	1	0	0	0	0	0	0	6	9
Doubao		1	1	1	1	0	1	0	0	1	0	0	0	0	0	0	6	9
Qwen2.5-Max		1	1	1	1	0	0	0	0	1	0	0	0	0	0	0	5	10
WanX2.1		1	0	0	1	1	0	0	0	0	0	0	0	0	1	0	4	11
Kling		1	1	1	1	0	1	0	0	1	0	0	0	0	0	0	6	9
Star-3 Alpha		1	1	1	1	0	1	0	0	1	0	0	0	0	0	0	6	9
Hunyuan		1	1	0	1	0	1	0	0	1	0	0	0	0	0	0	5	10
GLM-4		1	1	1	1	0	0	0	0	1	0	0	0	0	0	0	5	10

2320

Table 33: Counting Humans and Cats with Grid Prior Results.

2321

2322

2323

2324

2325

2326

2327

2328

2329

2330

2331

2332

2333

2334

2335

2336

2337

2338

2339

2340

2341

2342

2343

2344

2345

2346

2347

2348

Objects	Model	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	Correct	Wrong
Triangles	Recraft V3	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	15
	Imagen-3	1	0	0	1	0	1	0	0	1	0	0	0	0	0	0	4	11
	Dall-E 3	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	4	11
	Grok 3	1	0	0	0	0	1	0	0	0	0	0	0	0	0	0	2	13
	Gemini 2.0 Flash	1	1	1	0	0	1	0	0	0	0	0	0	0	0	0	4	11
	FLUX 1.1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	14
	Firefly 3	1	0	1	1	0	1	0	1	1	0	0	0	0	0	0	6	9
	SD 3.5	0	0	0	1	0	1	0	0	0	0	0	0	0	0	0	2	13
	Doubao	1	1	1	1	0	1	0	0	1	0	0	0	0	0	0	6	9
	Qwen2.5-Max	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	6	9
	WanX2.1	1	1	0	1	0	0	0	0	0	0	0	0	0	0	0	3	12
	Kling	1	0	1	0	0	0	0	0	1	0	0	0	0	0	0	3	12
	Star-3 Alpha	1	1	1	0	0	0	0	0	1	0	0	0	0	0	0	4	11
	Hunyuan	1	0	0	1	0	0	0	0	0	0	0	0	0	0	0	2	13
	GLM-4	1	1	0	0	0	0	0	1	1	0	1	1	0	0	0	6	9
Chairs	Recraft V3	1	1	1	1	0	1	0	1	1	0	0	0	0	1	0	8	7
	Imagen-3	1	1	1	1	1	1	0	1	0	0	0	0	0	0	1	8	7
	Dall-E 3	1	1	1	1	1	1	1	0	0	0	0	0	0	0	0	7	8
	Grok 3	1	1	0	1	0	0	0	0	0	0	0	0	0	0	0	3	12
	Gemini 2.0 Flash	1	1	0	1	0	1	0	0	1	1	0	0	0	0	0	6	9
	FLUX 1.1	1	1	1	0	0	1	0	1	1	1	0	1	0	0	0	8	7
	Firefly 3	1	1	1	1	1	1	0	1	1	0	0	0	0	0	0	8	7
	SD 3.5	1	1	1	0	1	0	0	1	0	0	0	1	0	0	0	6	9
	Doubao	1	1	1	1	1	1	1	0	0	0	0	0	0	0	0	7	8
	Qwen2.5-Max	1	1	1	0	0	1	1	1	0	0	0	0	0	0	0	6	9
	WanX2.1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	4	11
	Kling	1	1	1	1	1	0	0	0	1	0	0	0	0	0	0	6	9
	Star-3 Alpha	1	1	0	0	0	1	0	0	1	0	0	0	0	0	0	4	11
	Hunyuan	1	1	0	1	1	0	0	0	1	0	0	0	0	0	0	5	10
	GLM-4	1	1	1	1	1	1	0	0	1	0	0	0	0	0	0	7	8

Table 34: Counting Triangles and Chairs with Grid Prior Results.

2349

2350

2351

2352

2353

2354

2355

2356

2357

2358

2359

2360

2361

2362

2363

2364

2365

2366

2367

2368

2369

2370

Objects	Model	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	Correct	Wrong
Flowers	Recraft V3	0	0	0	1	0	0	0	0	1	0	0	0	0	0	0	2	15
	Imagen-3	1	0	1	0	0	1	0	1	1	0	1	0	0	0	0	6	9
	Dall-E 3	1	0	1	0	1	0	0	0	0	0	0	0	0	0	0	3	12
	Grok 3	1	0	1	0	1	0	0	0	0	0	0	0	0	0	0	3	12
	Gemini 2.0 Flash	1	0	0	1	0	0	0	0	1	0	0	0	0	0	0	3	12
	FLUX 1.1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	3	12
	Firefly 3	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	14
	SD 3.5	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	6	9
	Doubao	1	1	1	0	0	0	0	0	0	0	0	0	0	0	1	4	11
	Qwen2.5-Max	1	1	0	1	0	0	0	0	0	0	0	0	0	0	0	3	12
	WanX2.1	1	1	0	1	1	1	0	1	1	0	0	0	1	0	0	8	7
	Kling	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	2	13
	Star-3 Alpha	1	0	1	0	0	0	1	0	0	0	0	0	0	0	0	3	12
	Hunyuan	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	1	14
	GLM-4	1	1	1	0	1	1	0	0	0	0	0	0	1	0	0	6	9

Table 35: Counting Flowers with Grid Prior Results.

2371

2372

2373

2374

2375

2376

2377

Objects	Model	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	Correct	Wrong
2378	Recraft V3	0	1	0	1	0	0	0	0	0	0	0	0	0	0	0	2	13
2379	Imagen-3	1	1	1	1	0	1	0	0	0	1	0	0	0	0	0	6	9
2380	Dall-E 3	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	2	13
2381	Grok 3	0	1	0	0	0	0	1	0	0	1	0	0	0	0	0	3	12
2382	Gemini 2.0 Flash	1	1	1	1	1	1	1	0	1	0	0	0	0	0	1	9	6
2383	FLUX 1.1	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	1	14
2384	Firefly 3	0	1	1	1	1	0	0	0	0	0	0	0	0	0	0	4	11
2385	SD 3.5	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	1	14
2386	Doubao	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	1	14
2387	Qwen2.5-Max	0	1	1	0	1	1	0	0	0	0	0	1	1	0	0	6	9
2388	WanX2.1	0	1	1	1	0	0	1	1	1	1	0	0	0	0	0	7	8
2389	Kling	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	1	14
2390	Star-3 Alpha	1	1	1	0	0	0	0	1	1	0	0	0	0	0	0	5	10
2391	Hunyuan	0	1	1	1	1	0	0	0	0	0	0	0	0	0	0	4	11
2392	GLM-4	0	1	1	1	0	0	0	1	0	0	0	1	1	0	0	6	9
2393	Recraft V3	0	1	0	0	0	0	0	0	1	0	0	0	0	0	0	2	13
2394	Imagen-3	0	1	1	0	0	0	1	0	0	0	0	0	0	0	0	3	12
2395	Dall-E 3	0	0	0	1	0	0	0	0	0	0	0	0	1	0	0	2	13
2396	Grok 3	0	0	1	0	0	0	1	1	0	0	0	0	0	0	0	3	12
2397	Gemini 2.0 Flash	1	1	1	0	1	0	0	0	0	0	0	0	0	0	0	4	11
2398	FLUX 1.1	0	1	0	1	1	0	0	0	0	0	0	0	0	0	0	3	12
2399	Firefly 3	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	1	14
2400	SD 3.5	0	0	1	1	1	0	0	0	0	0	0	0	0	0	0	3	12
2401	Doubao	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	1	12
2402	Qwen2.5-Max	0	0	0	0	0	0	0	0	1	0	0	0	0	0	0	1	14
2403	WanX2.1	1	0	0	0	1	0	0	0	0	0	0	0	0	0	0	2	13
2404	Kling	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	1	14
2405	Star-3 Alpha	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	1	14
2406	Hunyuan	0	1	1	0	1	0	0	0	0	0	0	0	0	1	0	4	11
2407	GLM-4	0	1	1	0	0	0	0	0	0	0	0	0	0	1	1	4	11

Table 36: Counting Apples and Watermelons with Position Guidance Results.

2400

2401

2402

2403

2404

Objects	Model	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	Correct	Wrong
2405	Recraft V3	1	1	1	0	1	0	0	0	0	0	0	0	0	0	0	4	11
2406	Imagen-3	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	2	13
2407	Dall-E 3	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	2	13
2408	Grok 3	0	1	1	0	0	0	0	0	0	0	0	0	0	0	0	2	12
2409	Gemini 2.0 Flash	1	1	1	1	1	1	1	0	0	0	0	0	0	0	0	7	8
2410	FLUX 1.1	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	1	14
2411	Firefly 3	1	1	1	0	0	0	0	0	0	0	0	0	0	1	0	4	11
2412	SD 3.5	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	1	14
2413	Doubao	0	0	1	1	0	0	0	0	0	0	0	0	0	0	0	2	13
2414	Qwen2.5-Max	0	1	0	1	0	0	0	0	0	0	0	0	0	0	0	2	13
2415	WanX2.1	1	1	1	1	0	0	0	0	0	0	1	0	0	0	0	5	10
2416	Kling	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	1	14
2417	Star-3 Alpha	1	1	1	1	1	0	0	0	1	0	0	0	0	0	0	6	9
2418	Hunyuan	1	0	1	0	1	0	1	1	0	1	0	0	0	0	0	6	9
2419	GLM-4	1	1	1	0	0	0	0	0	0	0	0	1	0	0	0	4	11
2420	Recraft V3	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	1	14
2421	Imagen-3	1	1	0	1	1	0	0	0	0	0	0	0	0	0	0	4	11
2422	Dall-E 3	0	1	1	0	0	0	0	0	0	0	0	0	0	0	0	2	13
2423	Grok 3	1	1	1	0	1	1	1	0	0	0	0	0	0	0	0	6	9
2424	Gemini 2.0 Flash	1	1	1	0	1	1	1	0	0	1	0	0	0	0	0	7	8
2425	FLUX 1.1	0	1	0	0	0	0	0	1	1	0	0	0	0	0	0	3	12
2426	Firefly 3	0	1	1	0	0	0	0	0	0	0	0	0	0	0	0	2	13
2427	SD 3.5	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	1	14
2428	Doubao	0	1	0	0	1	0	0	0	0	0	0	0	0	0	0	2	13
2429	Qwen2.5-Max	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	2	13
2430	WanX2.1	1	1	0	1	0	1	0	0	1	0	0	0	0	0	1	6	9
2431	Kling	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	1	14
2432	Star-3 Alpha	1	1	0	0	0	0	1	0	0	0	0	0	0	1	0	4	11
2433	Hunyuan	0	1	0	1	1	0	0	0	0	0	0	1	1	0	0	5	10
2434	GLM-4	0	1	1	0	1	0	0	0	0	0	0	0	0	0	0	3	12

Table 37: Counting Humans and Cats with Position Guidance Results.

2428

2429

2430

2431

2432

2433

Objects	Model	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	Correct	Wrong
2434	Recraft V3	0	1	0	0	0	0	1	1	0	1	0	0	0	0	0	4	11
2435	Imagen-3	0	1	1	1	0	0	0	0	0	0	0	0	0	0	0	3	12
2436	Dall-E 3	1	1	1	0	1	0	0	0	0	0	0	0	0	0	0	4	11
2437	Grok 3	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	3	12
2438	Gemini 2.0 Flash	0	1	1	0	0	1	0	0	0	0	0	0	0	0	0	3	12
2439	FLUX 1.1	0	0	0	0	0	0	0	0	0	0	1	0	0	0	1	2	13
2440	Firefly 3	0	0	0	0	0	1	0	0	0	0	0	0	0	0	0	1	14
2441	SD 3.5	0	0	0	1	0	0	0	0	0	0	0	0	0	0	0	1	14
2442	Doubao	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	15
2443	Qwen2.5-Max	0	0	0	0	1	0	0	1	0	0	0	0	0	0	0	2	13
2444	WanX2.1	0	1	1	1	0	0	0	0	0	0	0	0	0	0	0	3	12
2445	Kling	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	15
2446	Star-3 Alpha	0	0	0	0	1	1	1	0	0	0	0	0	0	0	0	3	12
2447	Hunyuan	0	0	0	0	0	0	0	0	0	0	1	0	0	0	0	1	14
2448	GLM-4	0	0	0	1	0	0	0	0	0	1	1	1	0	1	0	5	10
2449	Recraft V3	0	1	0	0	0	0	0	1	0	0	0	0	0	0	0	2	13
2450	Imagen-3	0	0	0	1	0	0	0	0	0	1	0	0	0	0	0	2	13
2451	Dall-E 3	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	1	14
2452	Grok 3	1	1	1	0	0	0	1	0	0	0	0	0	0	0	0	4	11
2453	Gemini 2.0 Flash	0	1	1	1	0	0	0	1	0	0	0	0	0	0	0	4	11
2454	FLUX 1.1	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	1	14
2455	Firefly 3	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	3	12
2456	SD 3.5	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	1	14
2457	Doubao	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	1	14
2458	Qwen2.5-Max	1	1	0	1	0	0	0	1	0	0	0	0	0	1	0	5	10
2459	WanX2.1	1	1	0	1	0	1	0	1	0	0	0	0	0	0	1	6	9
2460	Kling	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	14
2461	Star-3 Alpha	1	1	0	1	0	0	0	0	0	0	0	0	0	0	0	3	12
2462	Hunyuan	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	6	9
2463	GLM-4	0	1	1	1	0	0	0	0	0	0	0	0	0	0	0	3	12

Table 38: Counting Triangles and Chairs with Position Guidance Results.

2457

2458

2459

2460

2461

2462

2463

2464

2465

2466

2467

2468

2469

2470

2471

2472

2473

2474

2475

2476

2477

2478

Objects	Model	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	Correct	Wrong
2467	Recraft V3	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	14
2468	Imagen-3	1	1	1	0	1	0	0	0	0	0	1	0	0	1	1	7	8
2469	Dall-E 3	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	4	11
2470	Grok 3	0	1	0	0	0	1	0	0	0	0	0	0	0	0	0	2	13
2471	Gemini 2.0 Flash	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	14
2472	FLUX 1.1	0	0	1	0	0	1	0	1	0	0	0	0	0	0	0	3	12
2473	Firefly 3	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	3	12
2474	SD 3.5	0	0	1	1	1	0	0	0	0	0	0	1	1	1	0	6	9
2475	Doubao	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	1	14
2476	Qwen2.5-Max	0	0	0	1	0	1	0	0	0	0	0	0	0	0	0	2	13
2477	WanX2.1	1	0	1	1	0	0	0	0	0	0	0	0	0	0	0	3	12
2478	Kling	1	0	0	1	0	0	0	0	0	0	0	0	0	0	0	2	13
2479	Star-3 Alpha	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	1	14
2480	Hunyuan	0	0	0	0	1	0	0	0	0	0	0	0	0	0	0	1	14
2481	GLM-4	1	1	0	0	0	0	1	0	1	0	0	0	0	0	0	4	11

Table 39: Counting Flowers with Position Guidance Results.

2479

2480

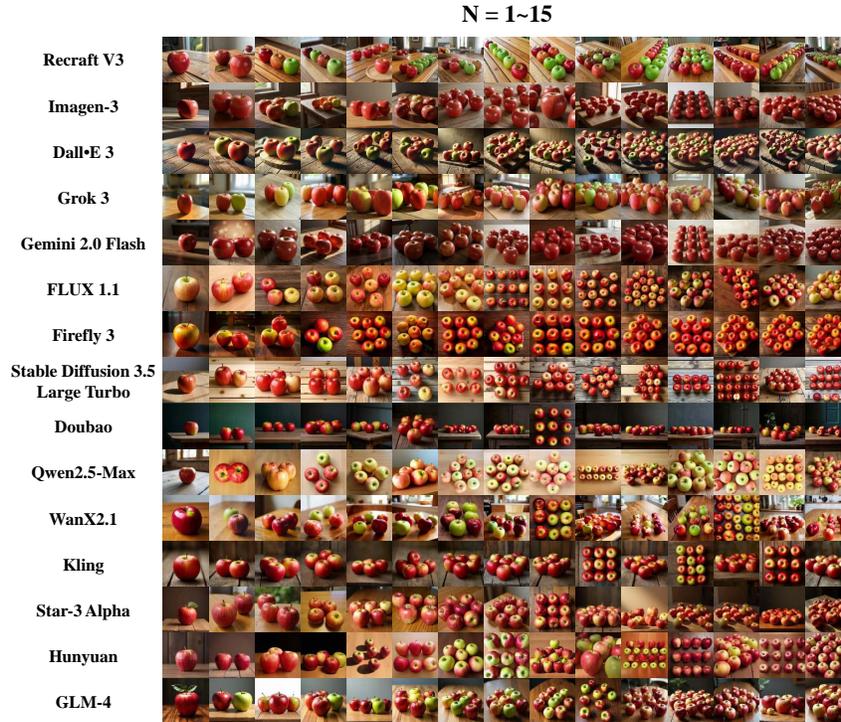
2481

2482

2483

2484 G IMAGE EXAMPLES
2485

2486 In this section, we present all the generated images used in our benchmarking results and prompt
2487 refinement experiments in Figures 12–74.
2488



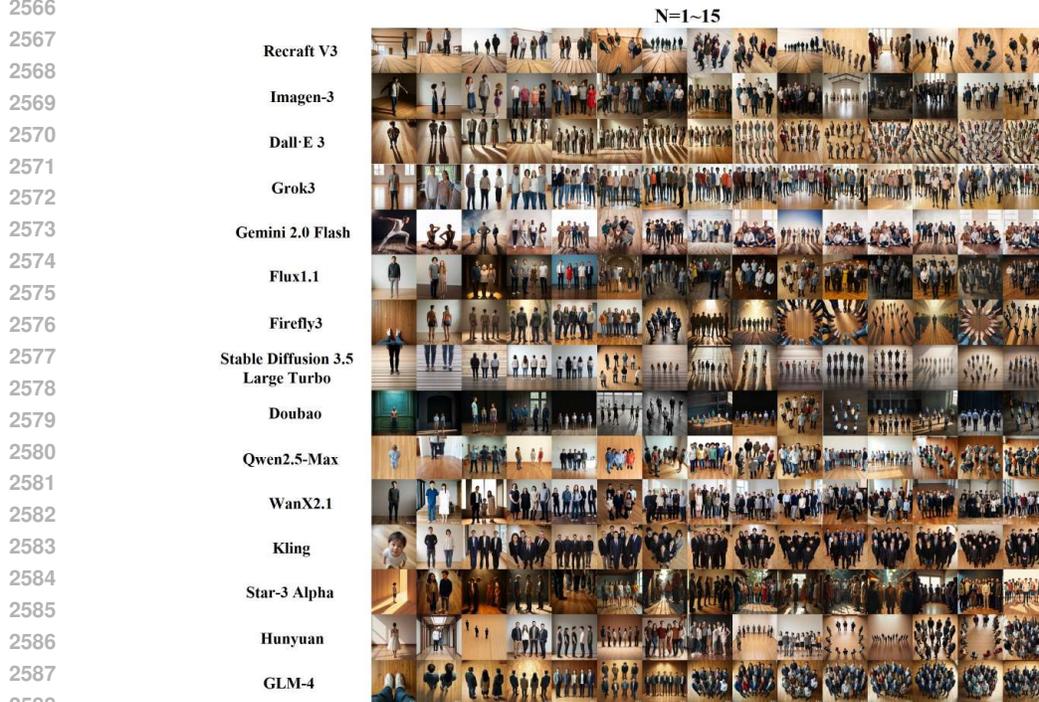
2506 Figure 12: **Counting Apples in Home Scene Results on 15 Models.** This figure presents the
2507 generation results of counting apples. We use the Prompt:“ N apples on a wooden table.”, where
2508 $N \in [1, 15]$ denotes the number of objects expected to be generated.
2509
2510
2511

2512
2513
2514
2515
2516
2517
2518
2519
2520
2521
2522
2523
2524
2525
2526
2527
2528
2529
2530
2531
2532
2533
2534
2535
2536
2537



2562
2563
2564
2565

Figure 13: **Counting Watermelons in Home Scene Results on 15 Models.** This figure presents the generation results of counting watermelons. We use the Prompt: “ N watermelons on a wooden table.”, where $N \in [1, 15]$ denotes the number of objects expected to be generated.



2589
2590
2591

Figure 14: **Counting Humans in Home Scene Results on 15 Models.** This figure presents the generation results of counting Humans. We use the Prompt: “ N Humans on the wooden floor.”, where $N \in [1, 15]$ denotes the number of objects expected to be generated.

2592
2593
2594
2595
2596
2597
2598
2599
2600
2601
2602
2603
2604
2605
2606
2607
2608
2609
2610
2611
2612
2613
2614
2615

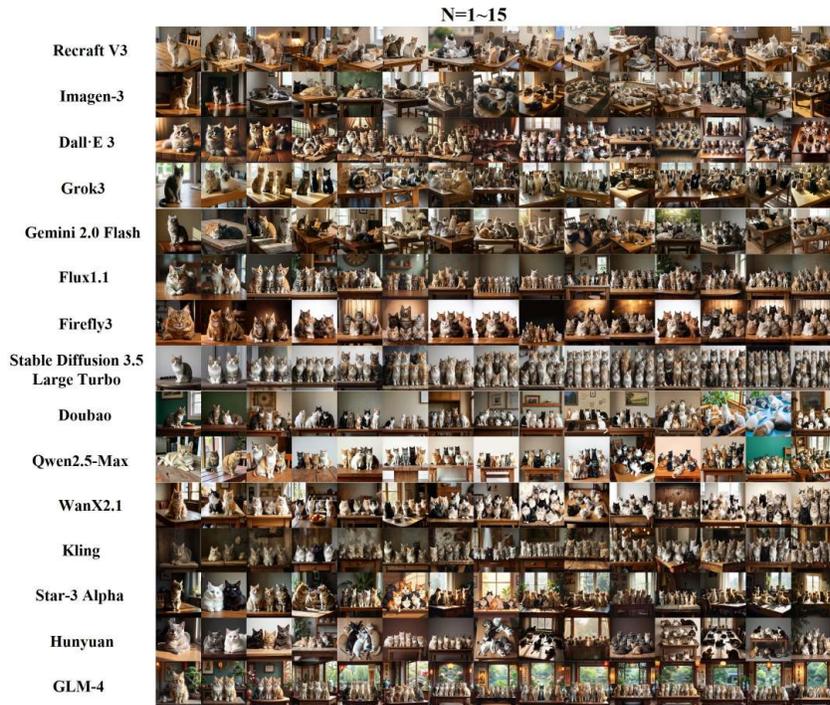


Figure 15: **Counting Cats in Home Scene Results on 15 Models.** This figure presents the generation results of counting cats. We use the Prompt:“ N cats on a wooden table.”, where $N \in [1, 15]$ denotes the number of objects expected to be generated.

2616
2617
2618
2619
2620
2621
2622
2623
2624
2625
2626
2627
2628
2629
2630
2631
2632
2633
2634
2635
2636
2637
2638
2639
2640
2641
2642

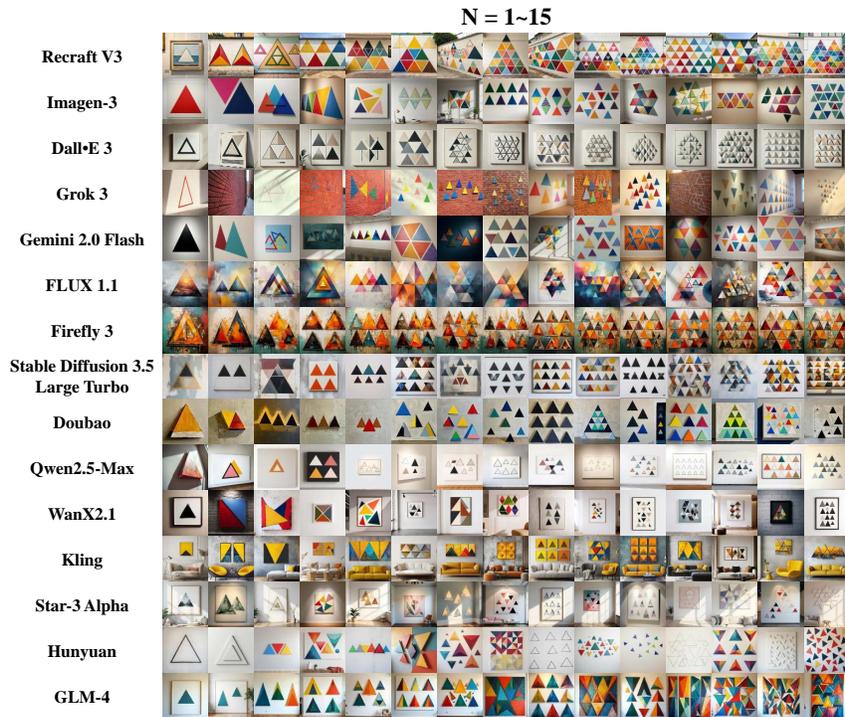
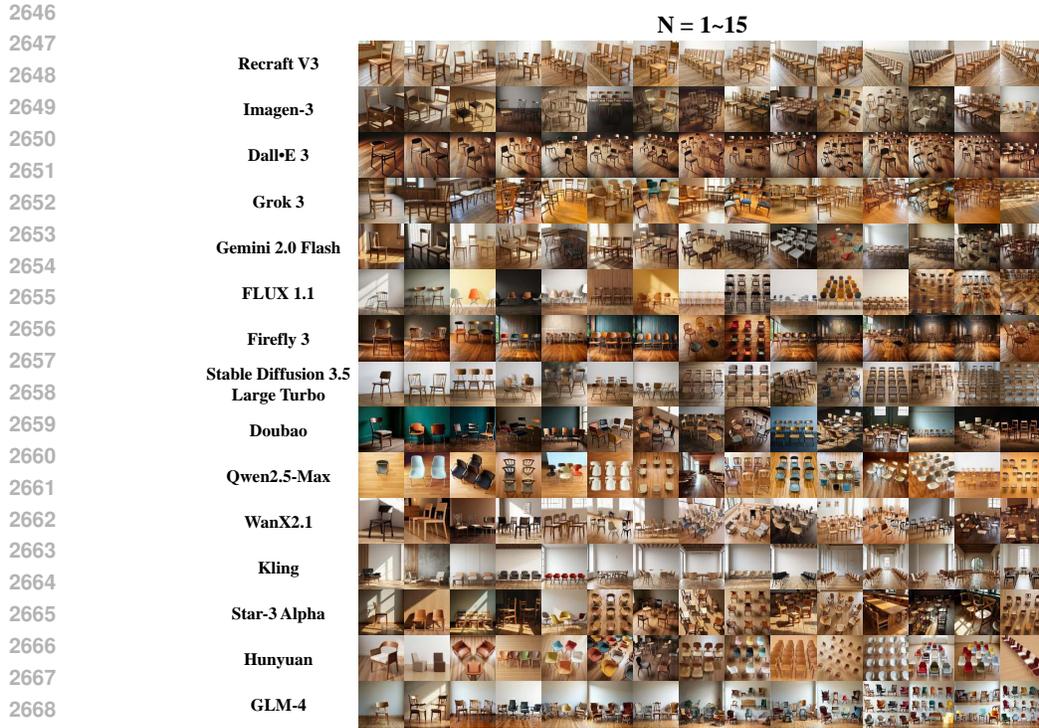
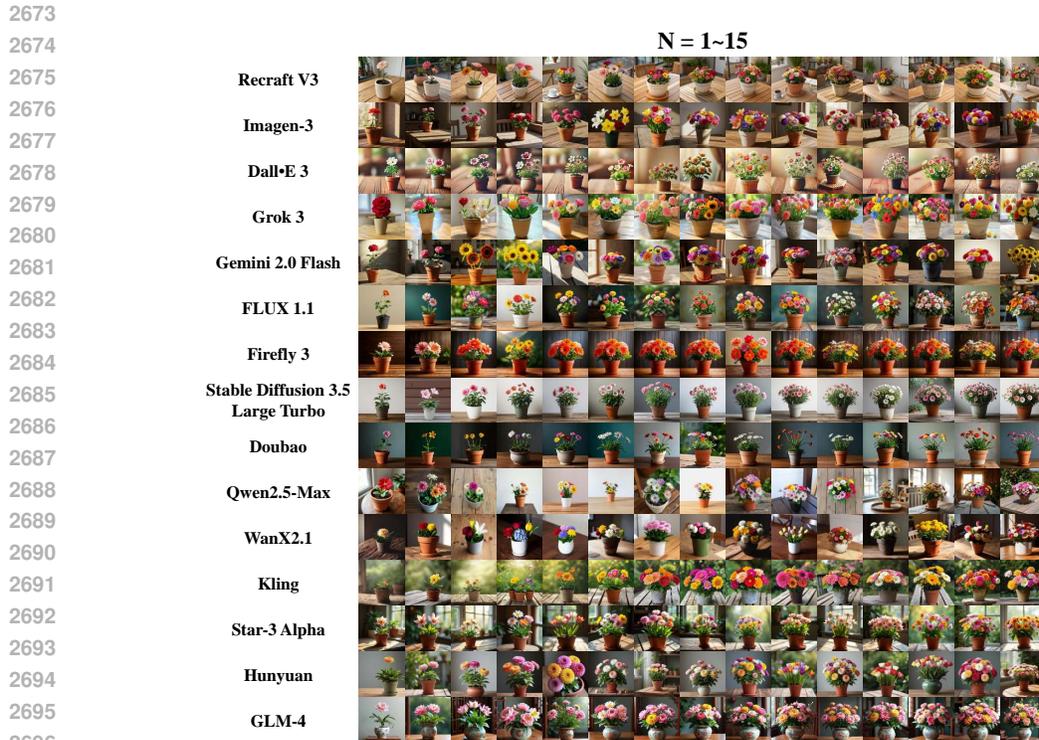


Figure 16: **Counting Triangles in Home Scene Results on 15 Models.** This figure presents the generation results of counting triangles. We use the Prompt:“ N triangles on a painting on a wall.”, where $N \in [1, 15]$ denotes the number of objects expected to be generated.



2670 **Figure 17: Counting Chairs in Home Scene Results on 15 Models.** This figure presents the
2671 generation results of counting chairs. We use the Prompt:“ N chairs on the wooden floor.”, where
2672 $N \in [1, 15]$ denotes the number of objects expected to be generated.



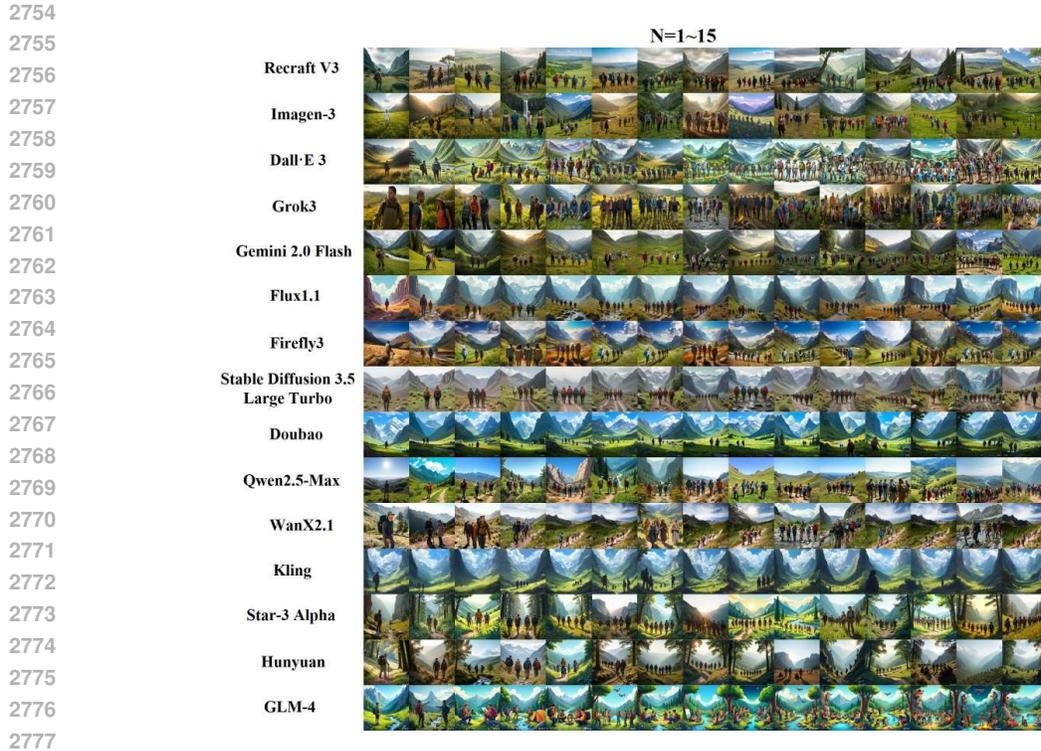
2697 **Figure 18: Counting Flowers in Home Scene Results on 15 Models.** This figure presents the
2698 generation results of counting flowers. We use the Prompt:“A flowerpot with N flowers on a wooden
2699 table.”, where $N \in [1, 15]$ denotes the number of objects expected to be generated.



2724 **Figure 19: Counting Apples in Nature Scene on 15 Models.** This figure presents the generation
2725 results of counting apples in a nature scene. We use the Prompt:“ N apples on an apple tree in a
2726 grassland”, where $N \in [1, 15]$ denotes the number of objects expected to be generated.



2751 **Figure 20: Counting Watermelons in Nature Scene on 15 Models.** This figure presents the gen-
2752 eration results of counting watermelons in a nature scene. We use the Prompt:“ N watermelon in a
2753 farmland”, where $N \in [1, 15]$ denotes the number of objects expected to be generated.



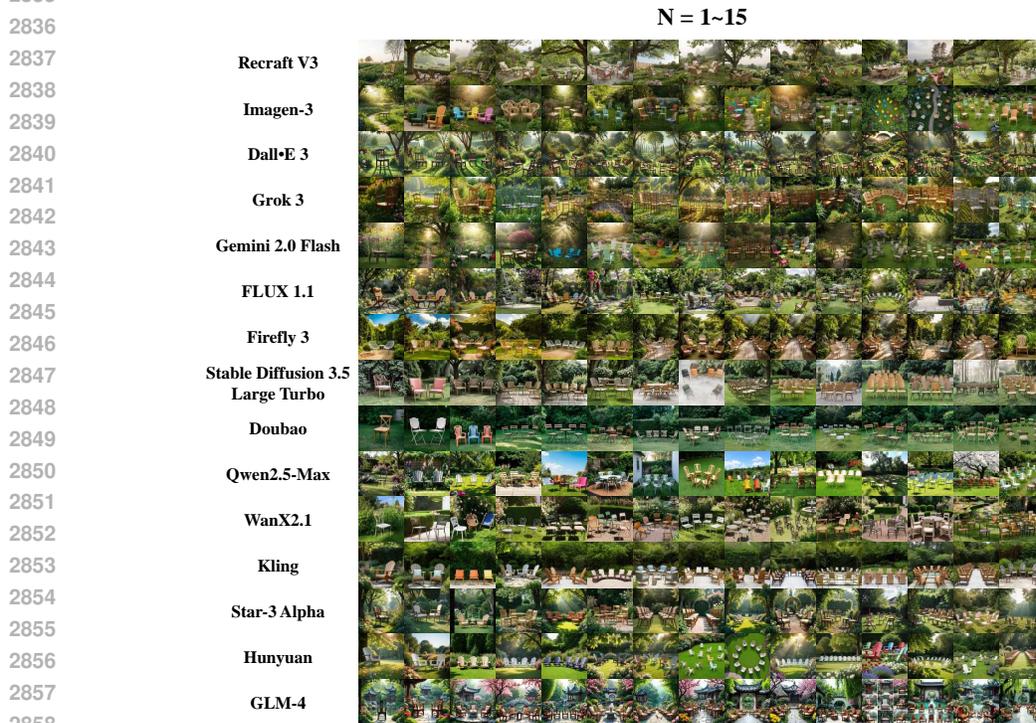
2778 **Figure 21: Counting Humans in Nature Scene on 15 Models.** This figure presents the generation
2779 results of counting humans in a nature scene. We use the Prompt:“ N human travelers in a valley.”,
2780 where $N \in [1, 15]$ denotes the number of objects expected to be generated.



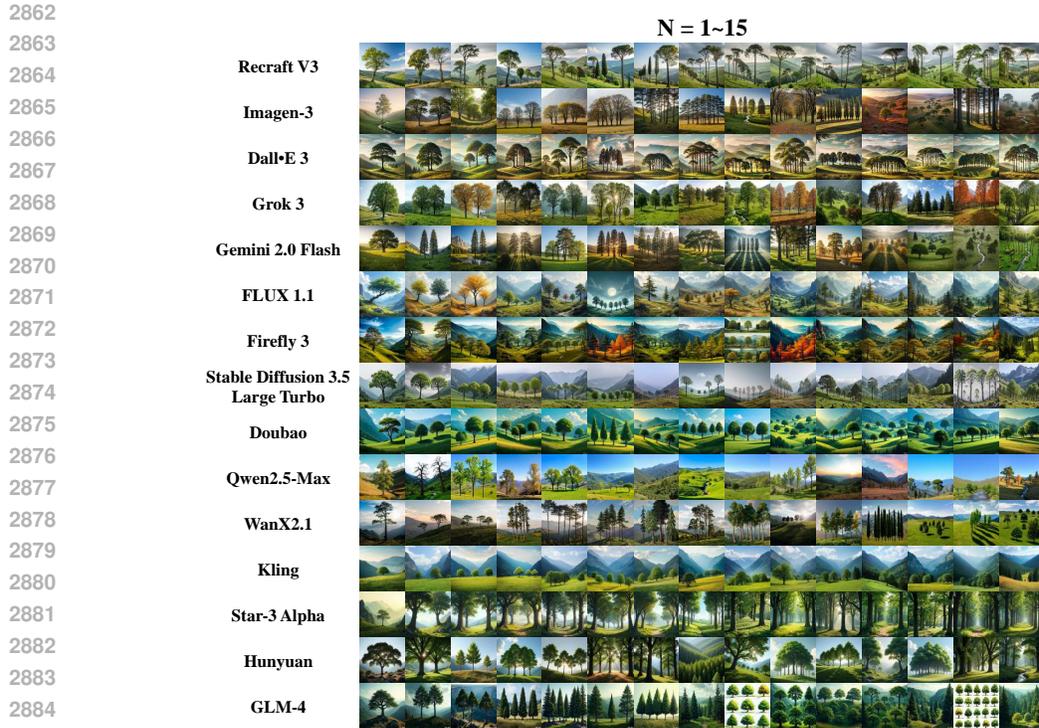
2805 **Figure 22: Counting Cats in Nature Scene on 15 Models.** This figure presents the generation
2806 results of counting cats in a nature scene. We use the Prompt:“ N cats in a temperate grassland.”,
2807 where $N \in [1, 15]$ denotes the number of objects expected to be generated.



2832 **Figure 23: Counting Triangles in Nature Scene on 15 Models.** This figure presents the generation
2833 results of counting triangles in a nature scene. We use the Prompt:“ N triangles on an earthy sur-
2834 face.”, where $N \in [1, 15]$ denotes the number of objects expected to be generated.



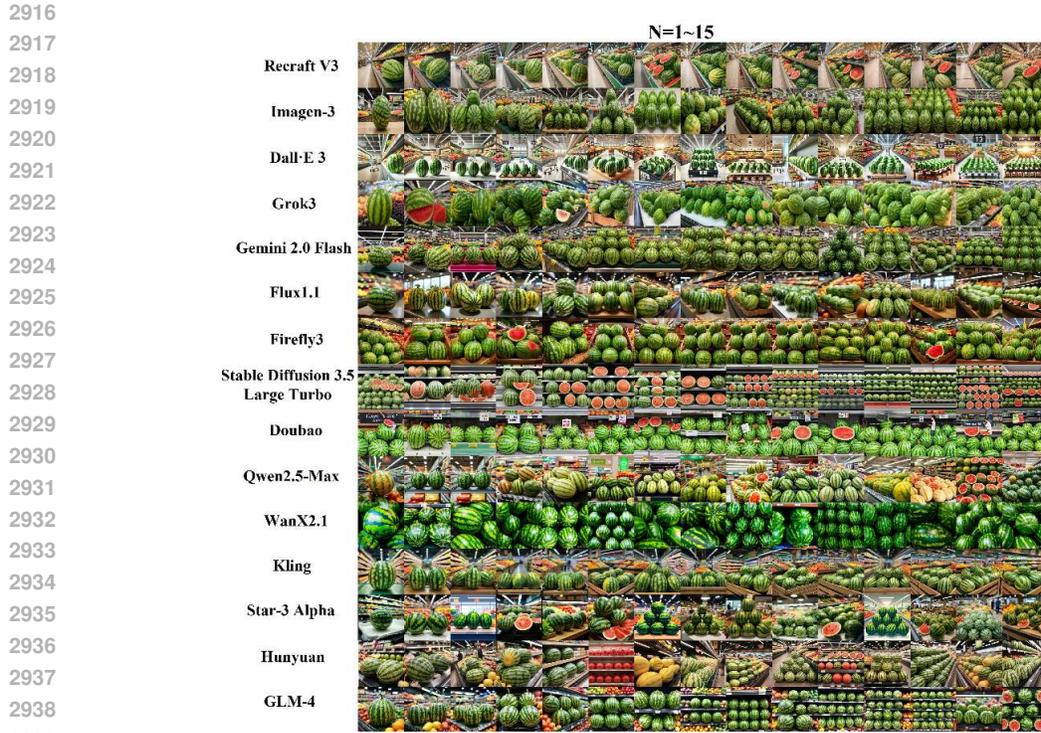
2859 **Figure 24: Counting Chairs in Nature Scene on 15 Models.** This figure presents the generation
2860 results of counting chairs in a nature scene. We use the Prompt:“ N chairs in a garden.”, where
2861 $N \in [1, 15]$ denotes the number of objects expected to be generated.



2886 **Figure 25: Counting Trees in Nature Scene on 15 Models.** This figure presents the generation
2887 results of counting trees in a nature scene. We use the Prompt:“ N trees in a valley.”, where $N \in$
2888 $[1, 15]$ denotes the number of objects expected to be generated.



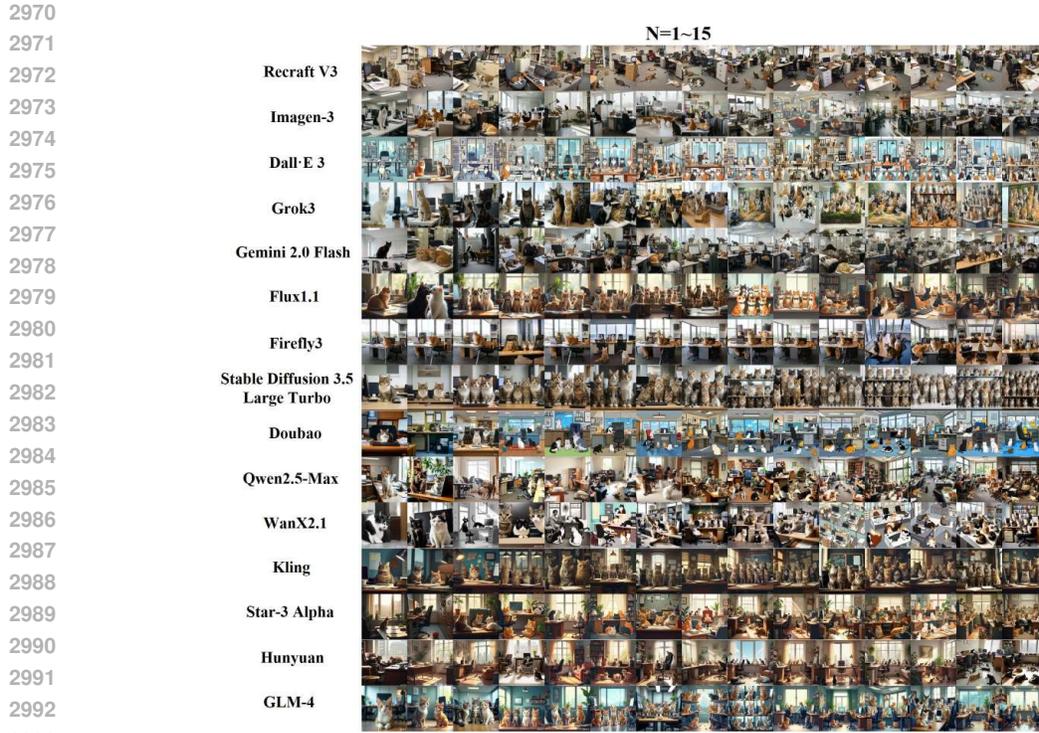
2913 **Figure 26: Counting Apples in City Scene on 15 Models.** This figure presents the generation
2914 results of counting apples on apple trees. We use the Prompt:“ N apples on an apple tree in a city
2915 garden. ”, where $N \in [1, 15]$ denotes the number of objects expected to be generated.



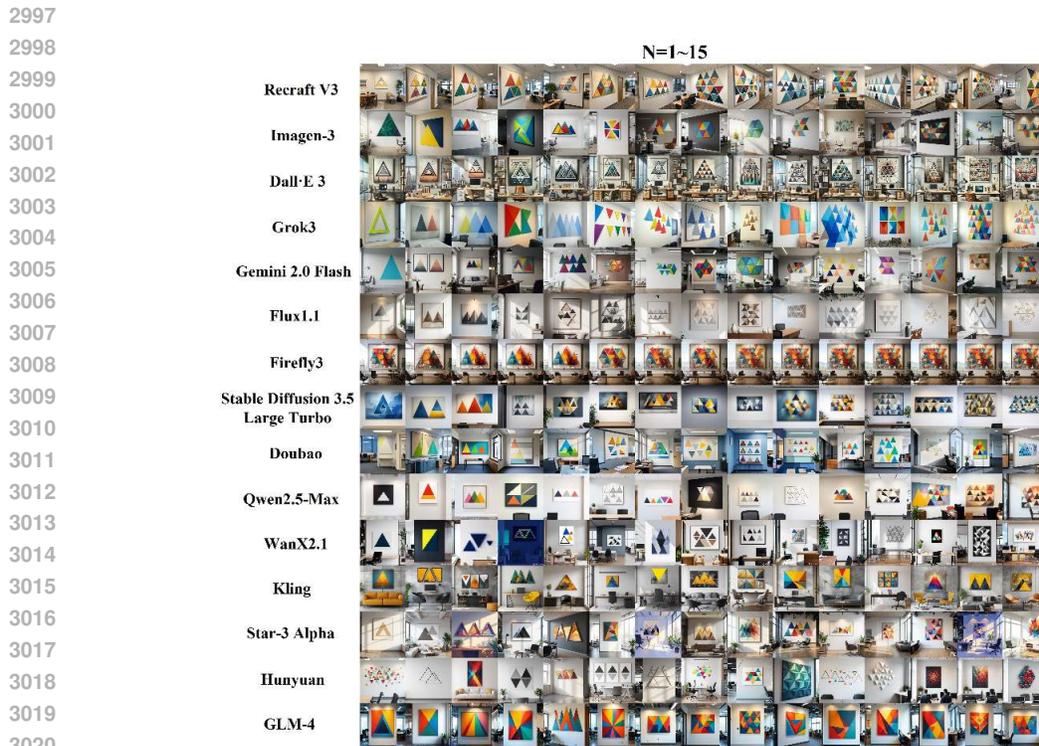
2940 **Figure 27: Counting Watermelons in City Scene on 15 Models.** This figure presents the generation
2941 results of counting watermelons. We use the Prompt:“ N watermelons in a supermarket.”, where
2942 $N \in [1, 15]$ denotes the number of objects expected to be generated.



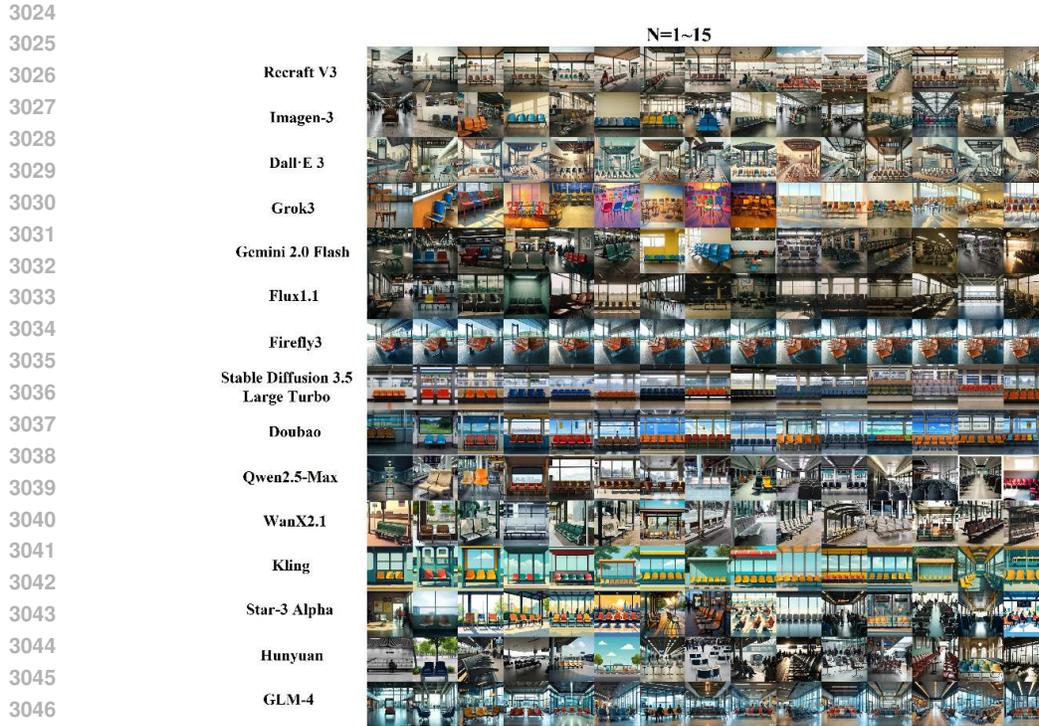
2967 **Figure 28: Counting Humans in City Scene on 15 Models.** This figure presents the generation
2968 results of counting humans. We use the Prompt:“ N humans in a city central business district.”,
2969 where $N \in [1, 15]$ denotes the number of objects expected to be generated.



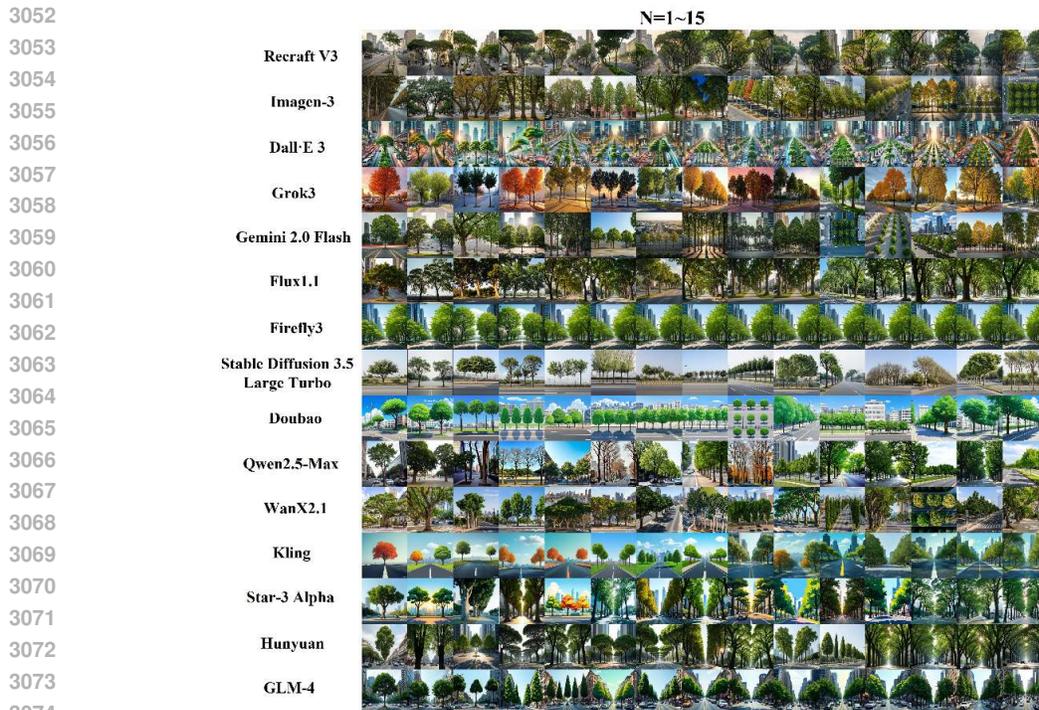
2994 **Figure 29: Counting Cats in City Scene on 15 Models.** This figure presents the generation results
2995 of counting cats. We use the Prompt:“ N cats in an office.”, where $N \in [1, 15]$ denotes the number
2996 of objects expected to be generated.



3021 **Figure 30: Counting Triangles in City Scene on 15 Models.** This figure presents the generation
3022 results of counting triangles. We use the Prompt:“ N triangles in a painting on the wall of an office.”,
3023 where $N \in [1, 15]$ denotes the number of objects expected to be generated.



3048 **Figure 31: Counting Chairs in City Scene on 15 Models.** This figure presents the generation results of counting chairs. We use the Prompt:“ N chairs in a bus station.”, where $N \in [1, 15]$
3049 denotes the number of objects expected to be generated.
3050
3051



3075 **Figure 32: Counting Trees in City Scene on 15 Models.** This figure presents the generation results
3076 of counting trees. We use the Prompt:“ N trees alongside a city’s main road.”, where $N \in [1, 15]$
3077 denotes the number of objects expected to be generated.



Figure 33: **Counting Apples With Cartoon Style on 15 Models.** This figure presents the generation results of counting apples on apple trees. We use the Prompt: “ N apple on a wooden table in a cartoon style.”, where $N \in [1, 15]$ denotes the number of objects expected to be generated.



Figure 34: **Counting Watermelons With Cartoon Style on 15 Models.** This figure presents the generation results of counting watermelons. We use the Prompt: “ N watermelons on a wooden table in a cartoon style.”, where $N \in [1, 15]$ denotes the number of objects expected to be generated.

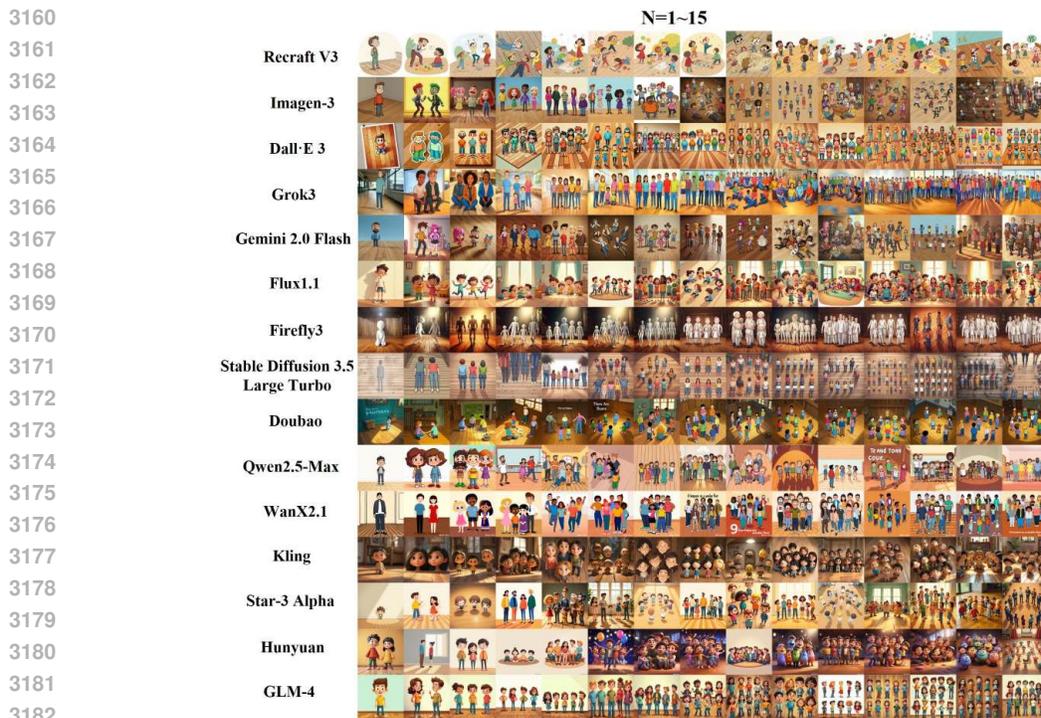


Figure 35: **Counting Humans With Cartoon Style on 15 Models.** This figure presents the generation results of counting humans. We use the Prompt: “ N humans on a wooden floor in a cartoon style.”, where $N \in [1, 15]$ denotes the number of objects expected to be generated.

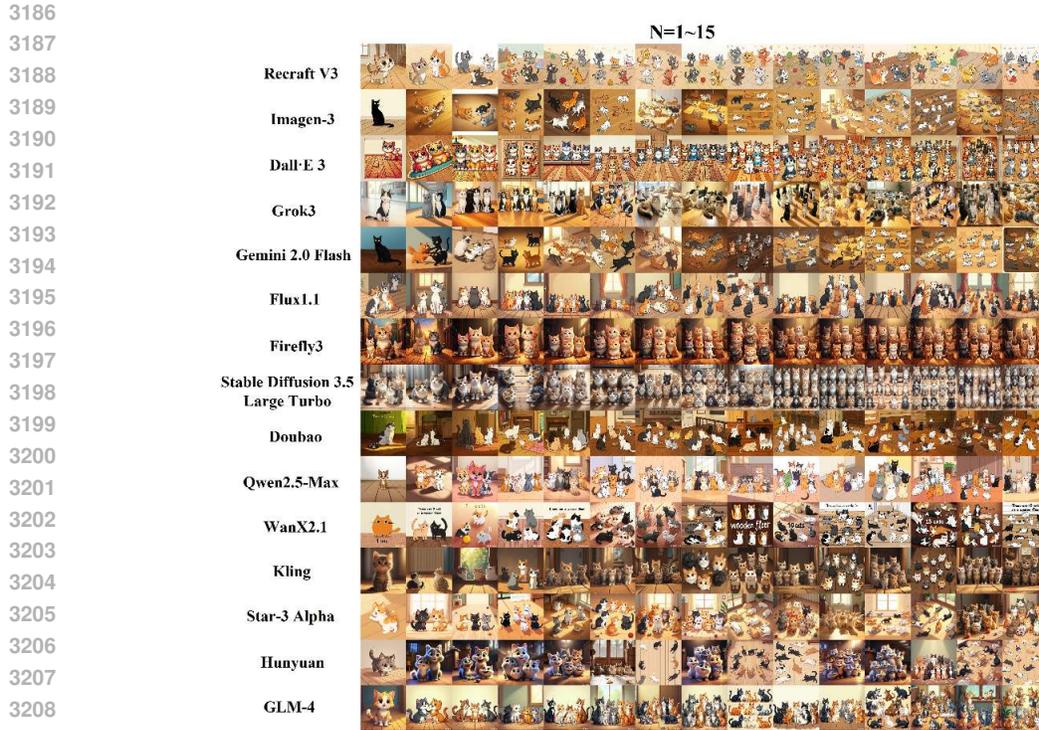
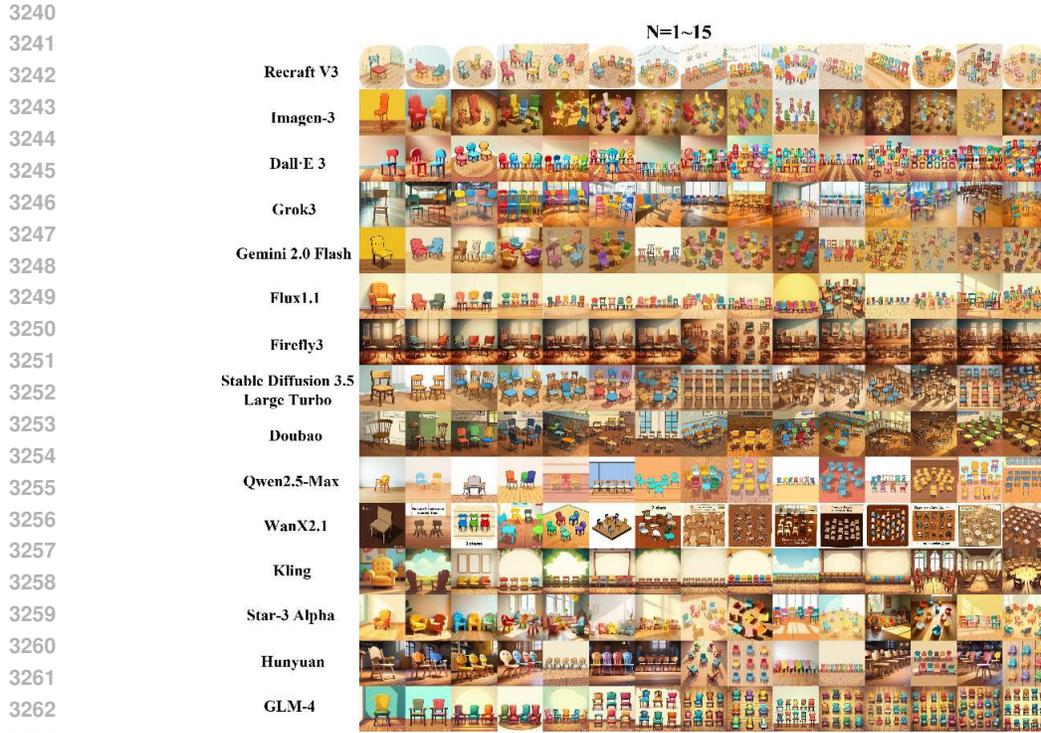


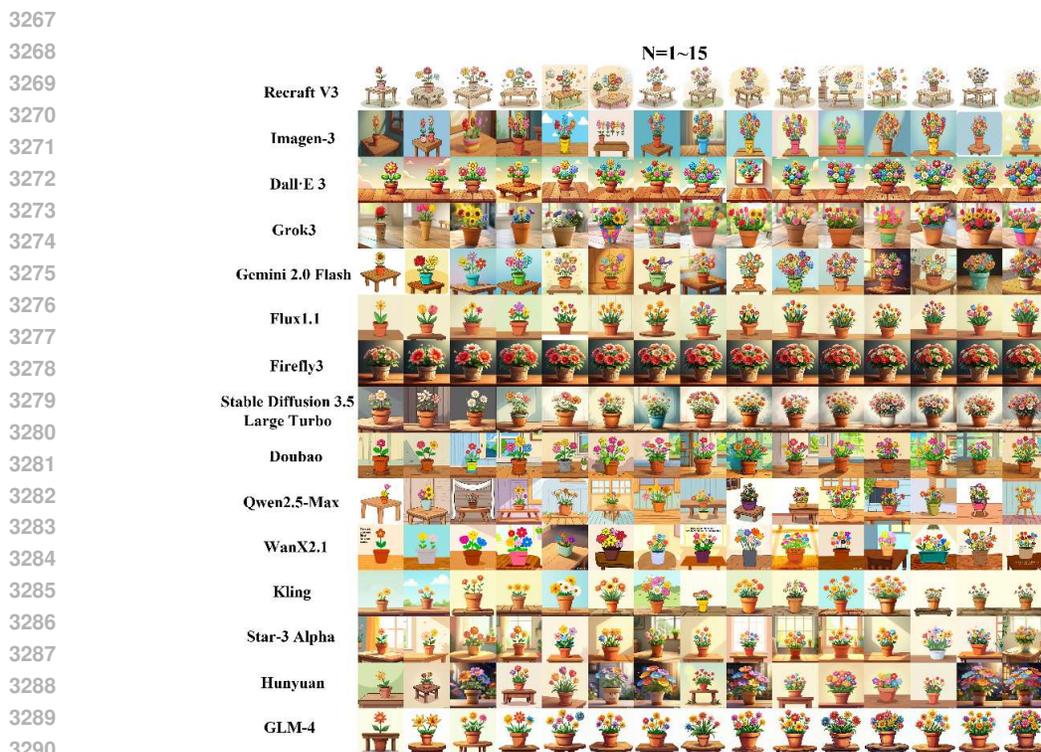
Figure 36: **Counting Cats With Cartoon Style on 15 Models.** This figure presents the generation results of counting cats. We use the Prompt: “ N cats on a wooden floor in a cartoon style.”, where $N \in [1, 15]$ denotes the number of objects expected to be generated.



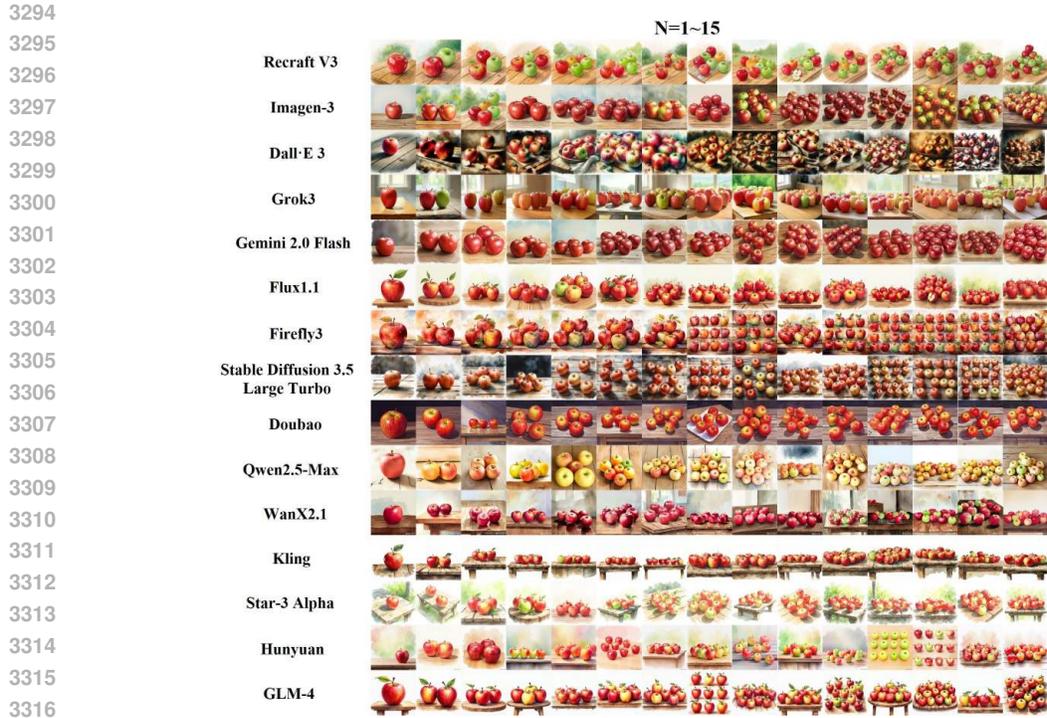
Figure 37: **Counting Triangles With Cartoon Style on 15 Models.** This figure presents the generation results of counting triangles. We use the Prompt: “ N triangles on a painting on a wall in a cartoon style.”, where $N \in [1, 15]$ denotes the number of objects expected to be generated.



3264 **Figure 38: Counting Chairs With Cartoon Style on 15 Models.** This figure presents the generation
3265 results of counting chairs. We use the Prompt:“ N chairs on a wooden floor in a cartoon style.”, where
3266 $N \in [1, 15]$ denotes the number of objects expected to be generated.



3291 **Figure 39: Counting Flowers With Cartoon Style on 15 Models.** This figure presents the generation
3292 results of counting trees. We use the Prompt:“a flower pot with N flowers on a wooden table in
3293 a cartoon style.”, where $N \in [1, 15]$ denotes the number of objects expected to be generated.



3317
3318
3319
3320

Figure 40: **Counting Apples with Watercolor style on 15 Models.** This figure presents the generation results of counting apples with watercolor style. We use the Prompt:“ N apples on a table in a watercolor style.”, where $N \in [1, 15]$ denotes the number of objects expected to be generated.



3344
3345
3346
3347

Figure 41: **Counting Watermelons with Watercolor style on 15 Models.** This figure presents the generation results of counting watermelons with watercolor style. We use the Prompt:“ N watermelons on a table in a watercolor style.”, where $N \in [1, 15]$ denotes the number of objects expected to be generated.

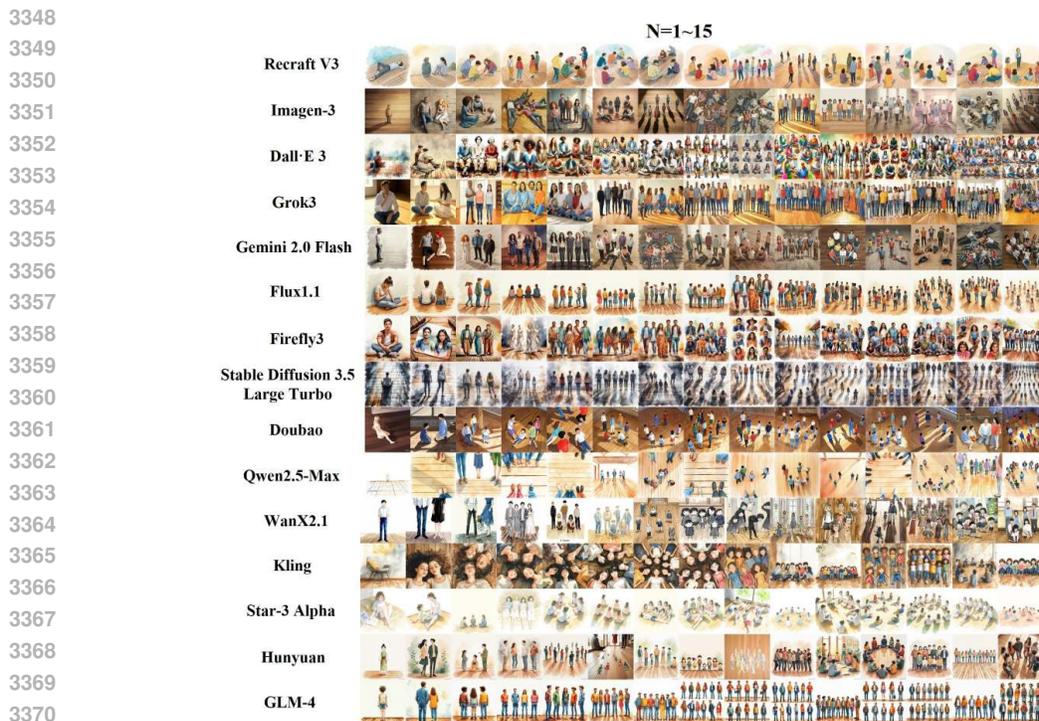


Figure 42: **Counting Humans with Watercolor style on 15 Models.** This figure presents the generation results of counting humans with watercolor style. We use the Prompt:“ N humans on a wooden in a watercolor style.”, where $N \in [1, 15]$ denotes the number of objects expected to be generated.

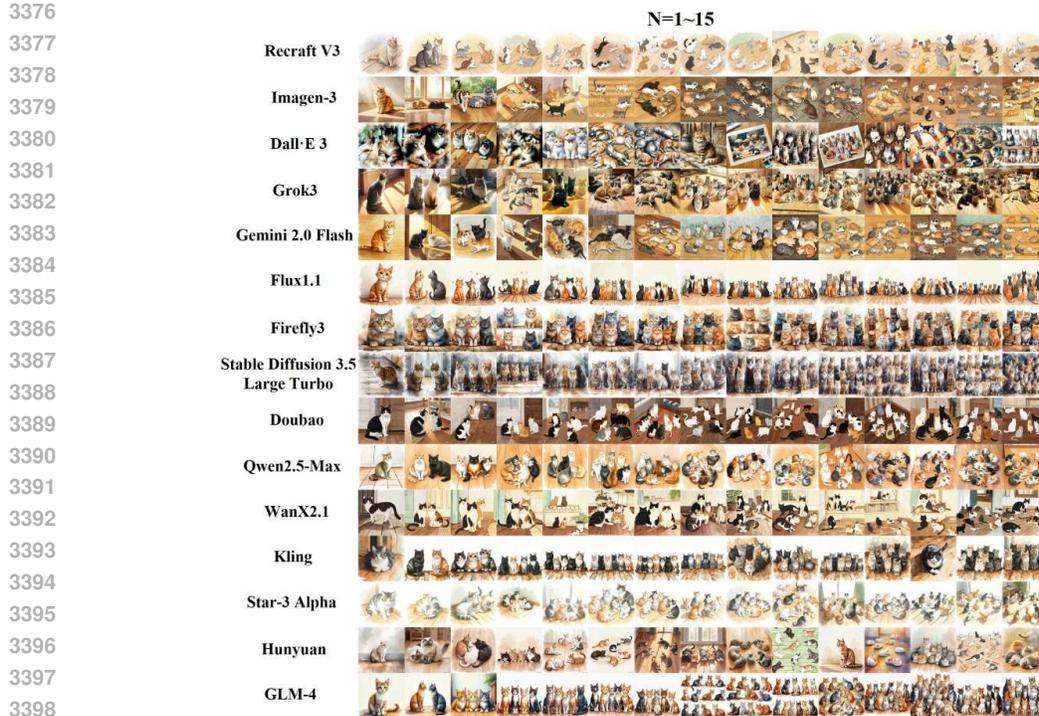
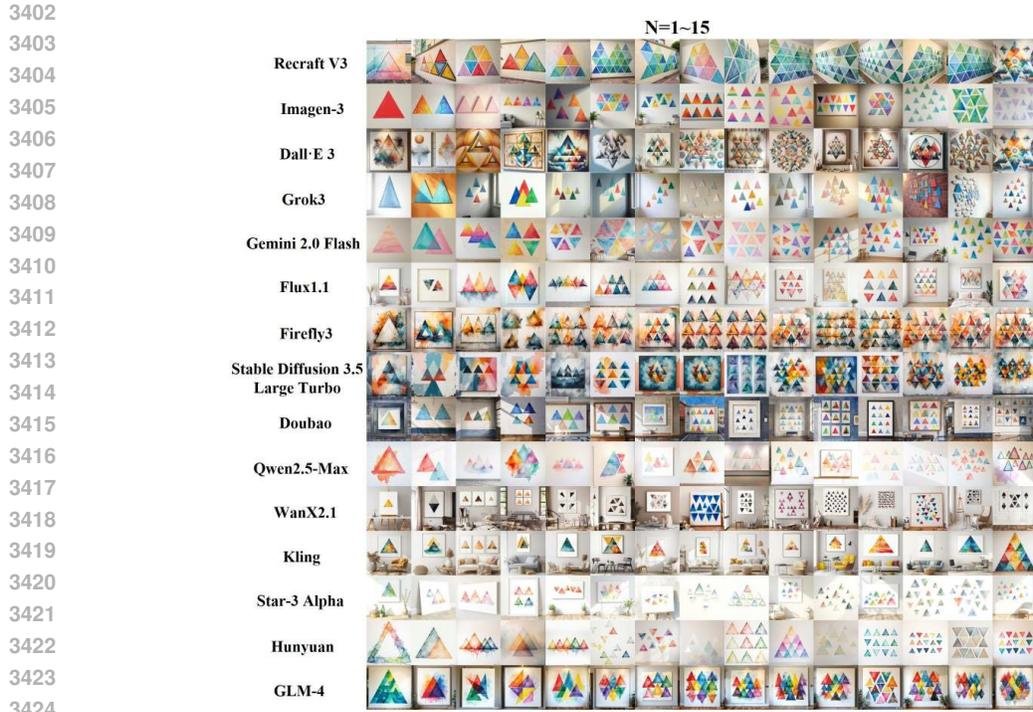
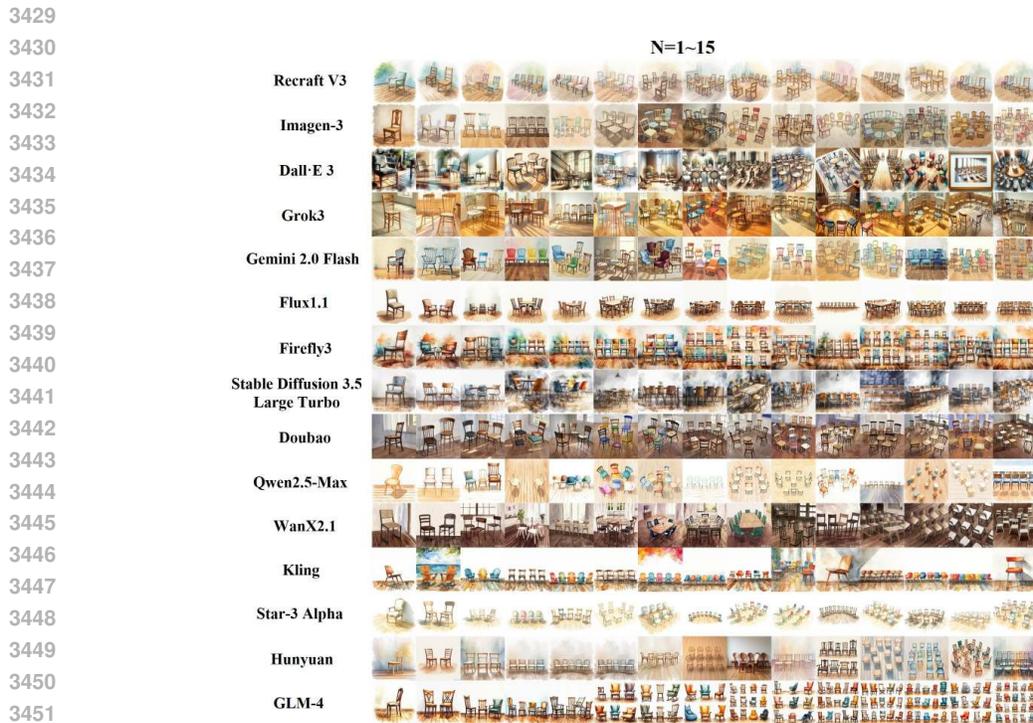


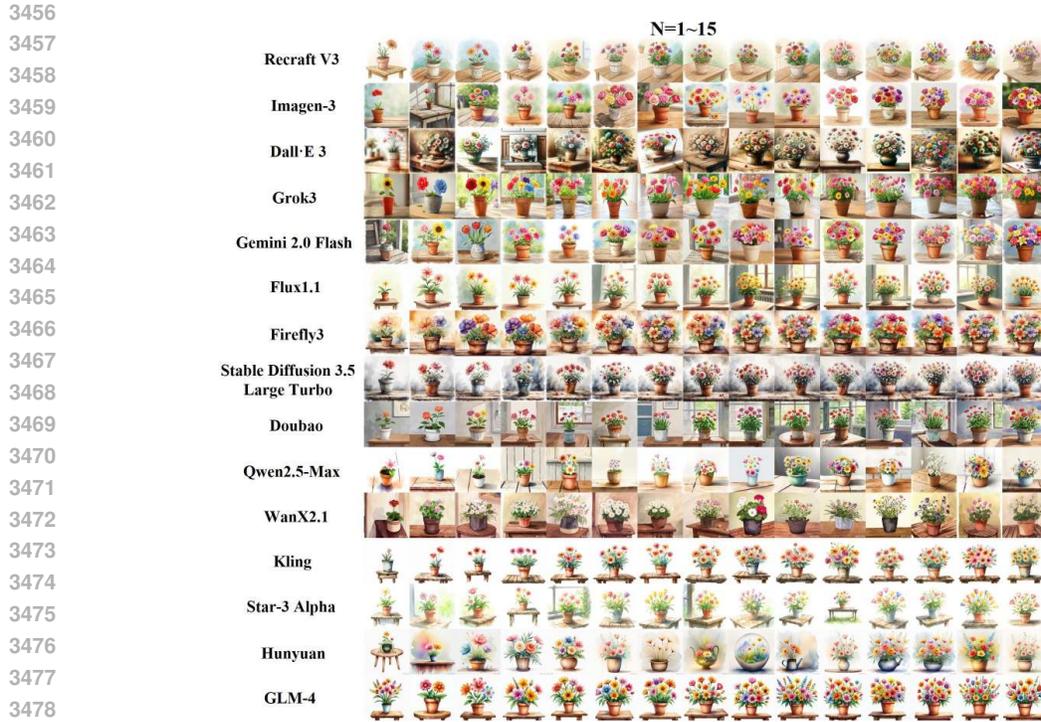
Figure 43: **Counting Cats with Watercolor style on 15 Models.** This figure presents the generation results of counting cats with watercolor style. We use the Prompt:“ N cats on a wooden in a watercolor style.”, where $N \in [1, 15]$ denotes the number of objects expected to be generated.



3425 **Figure 44: Counting Triangles with Watercolor style on 15 Models.** This figure presents the
3426 generation results of counting triangles with watercolor style. We use the Prompt:“ N triangles on a
3427 painting on a wall in a watercolor style.”, where $N \in [1, 15]$ denotes the number of objects expected
3428 to be generated.



3452 **Figure 45: Counting Chairs with Watercolor style on 15 Models.** This figure presents the gen-
3453 eration results of counting chairs with watercolor style. We use the Prompt:“ N chairs on a wooden
3454 floor in a watercolor style.”, where $N \in [1, 15]$ denotes the number of objects expected to be gener-
3455 ated.



3480

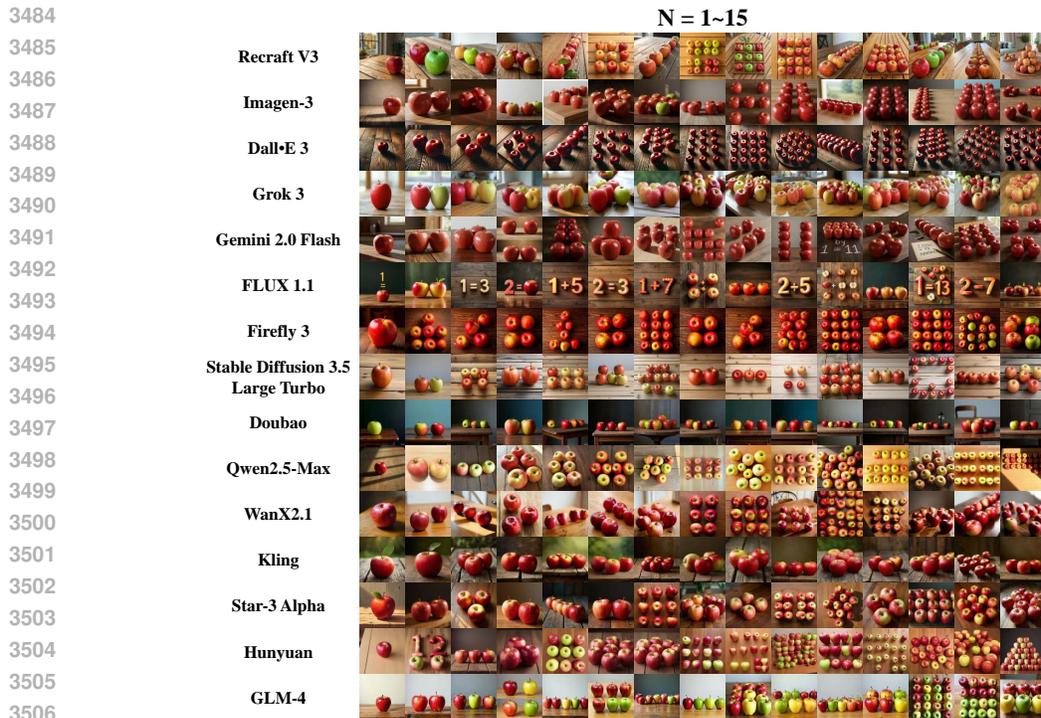
3481

3482

3483

3484

Figure 46: **Counting Flowers with Watercolor style on 15 Models.** This figure presents the generation results of counting flowers with watercolor style. We use the Prompt:“ a flower pot with N flowers on a wooden table in a watercolor style.”, where $N \in [1, 15]$ denotes the number of objects expected to be generated.

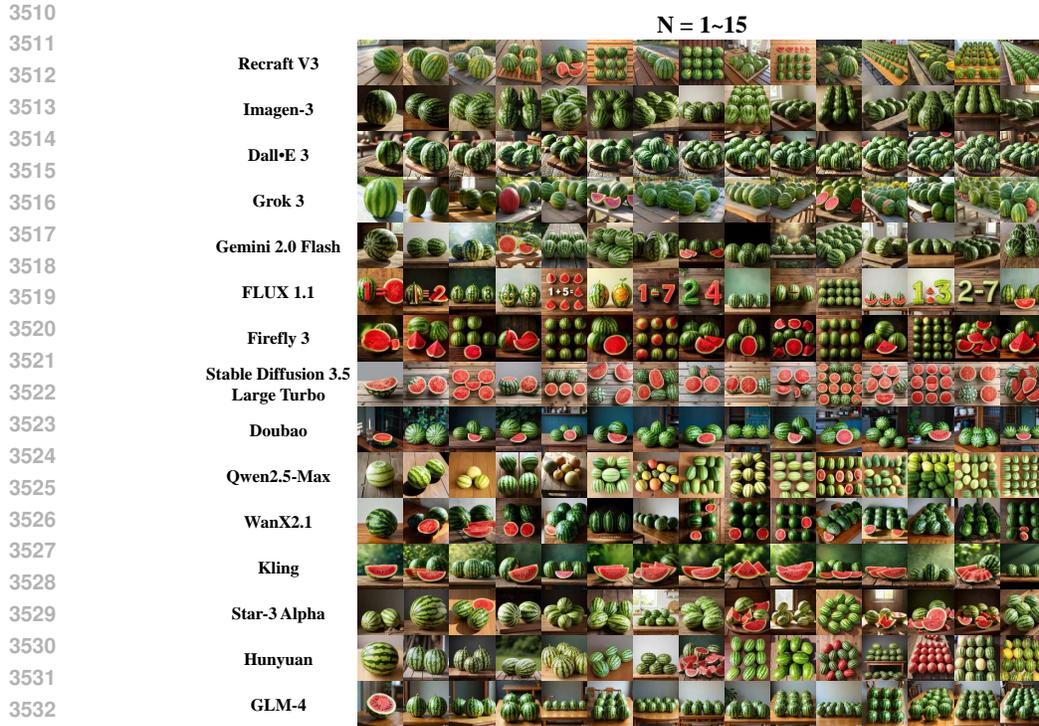


3507

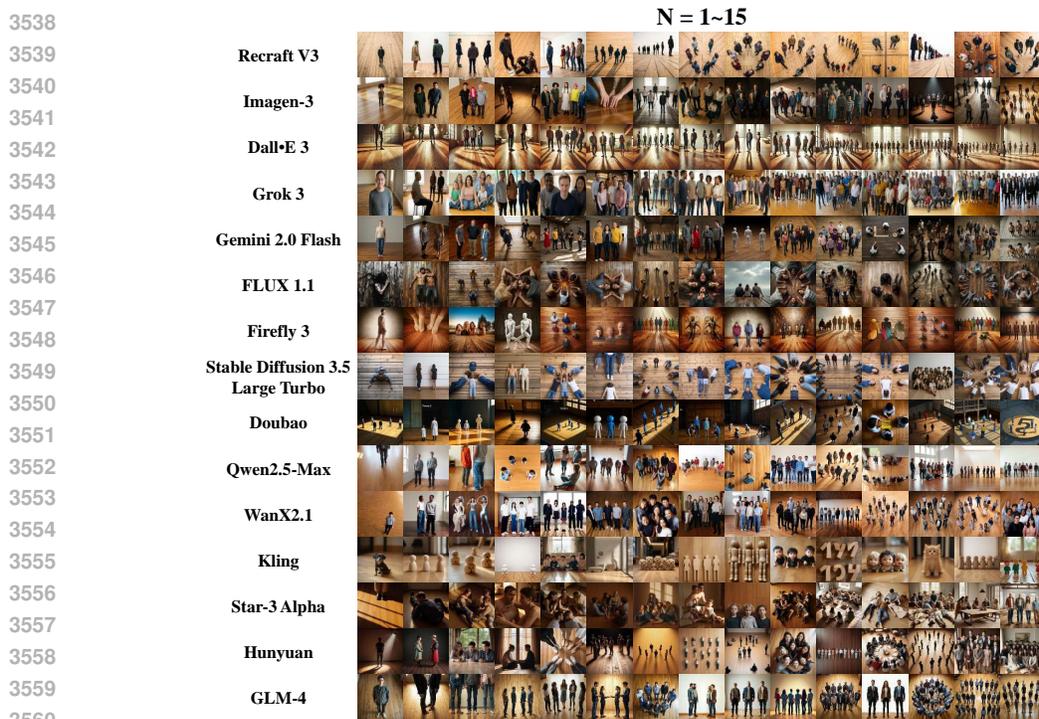
3508

3509

Figure 47: **Counting Apples with Multiplicative Decomposition on 15 Models.** This figure presents the generation results of counting apples with multiplicative decomposition. We use r times c to replace the N in the Prompt:“ N apples on a wooden table.”



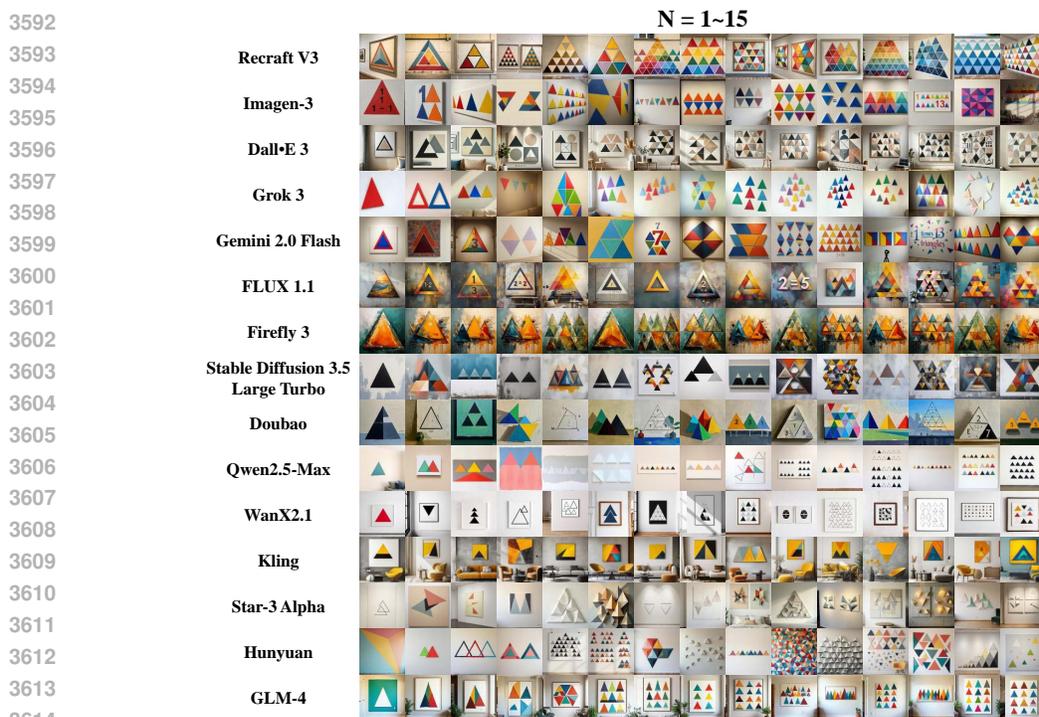
3534 **Figure 48: Counting watermelons with Multiplicative Decomposition on 15 Models.** This figure
 3535 presents the generation results of counting watermelons with multiplicative decomposition. We use
 3536 r times c to replace the N in the Prompt: “ N watermelons on a wooden table.”



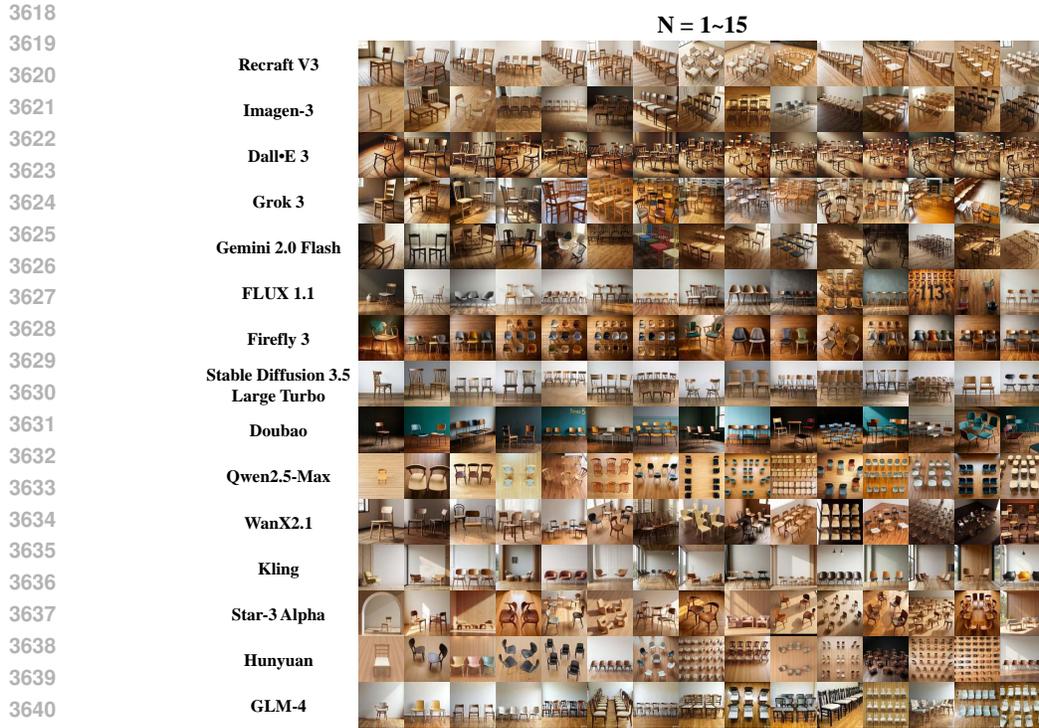
3561 **Figure 49: Counting Trees with Multiplicative Decomposition on 15 Models.** This figure presents
 3562 the generation results of counting humans with multiplicative decomposition. We use
 3563 r times c to replace the N in the Prompt: “ N humans on a wooden floor.”



3588 **Figure 50: Counting Cats with Multiplicative Decomposition on 15 Models.** This figure presents
3589 the generation results of counting cats with multiplicative decomposition. We use r times c to replace
3590 the N in the Prompt: “ N cats on a wooden floor.”



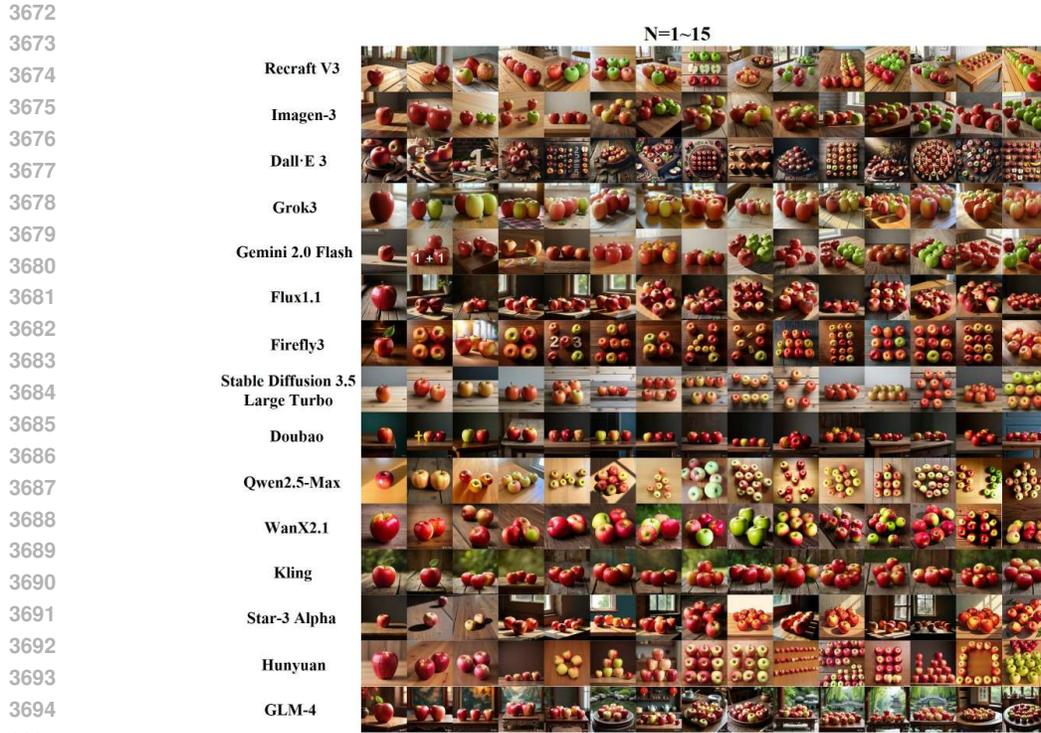
3615 **Figure 51: Counting Triangles with Multiplicative Decomposition on 15 Models.** This figure
3616 presents the generation results of counting triangles with multiplicative decomposition. We use r
3617 times c to replace the N in the Prompt: “ N triangles on a wooden floor.”



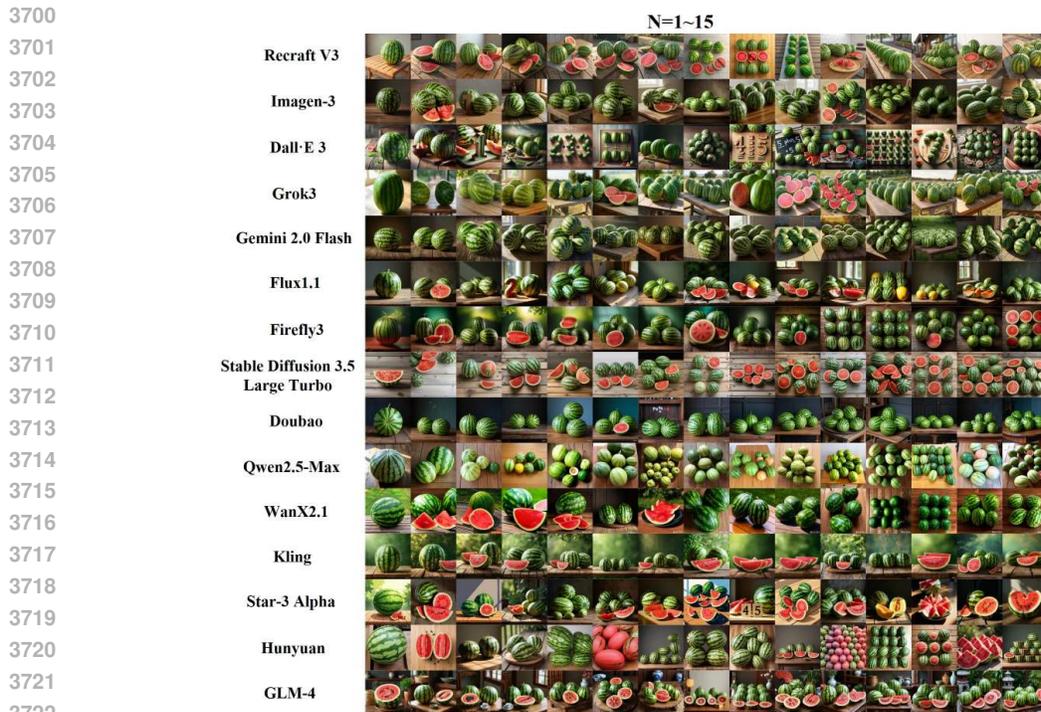
3642 **Figure 52: Counting Chairs with Multiplicative Decomposition on 15 Models.** This figure
3643 presents the generation results of counting chairs with multiplicative decomposition. We use r
3644 times c to replace the N in the Prompt: “ N chairs on a wooden floor.”



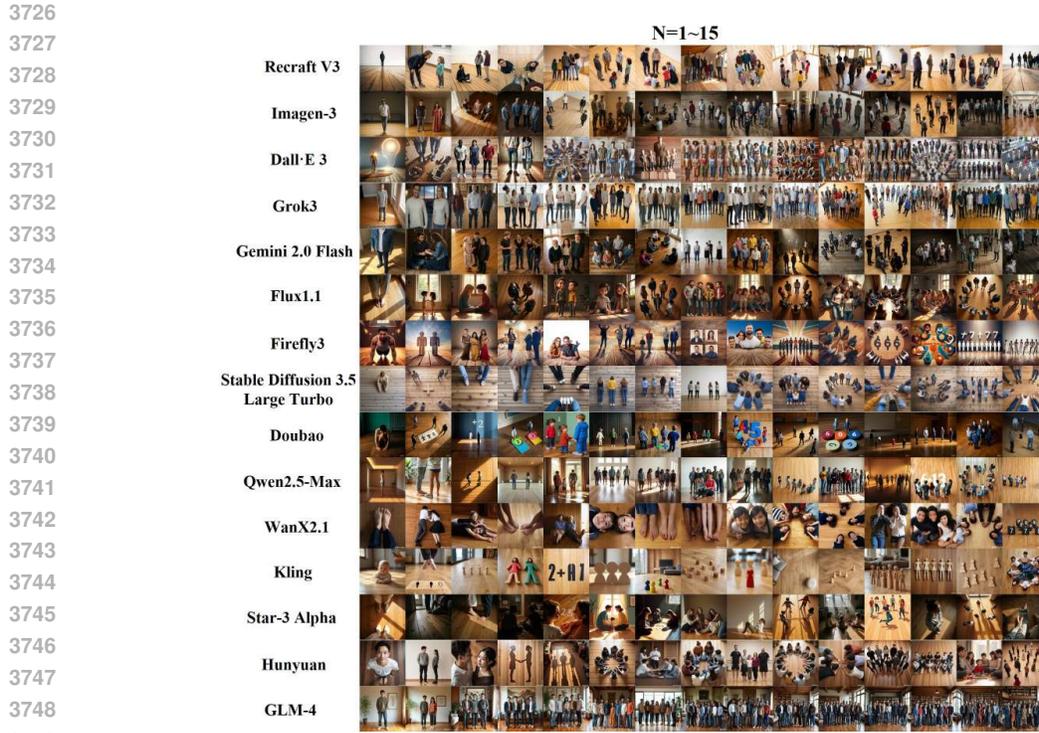
3669 **Figure 53: Counting Flowers with Multiplicative Decomposition on 15 Models.** This figure
3670 presents the generation results of counting flowers with multiplicative decomposition. We use r
3671 times c to replace the N in the Prompt: “A flower pot with N flowers on a wooden table.”



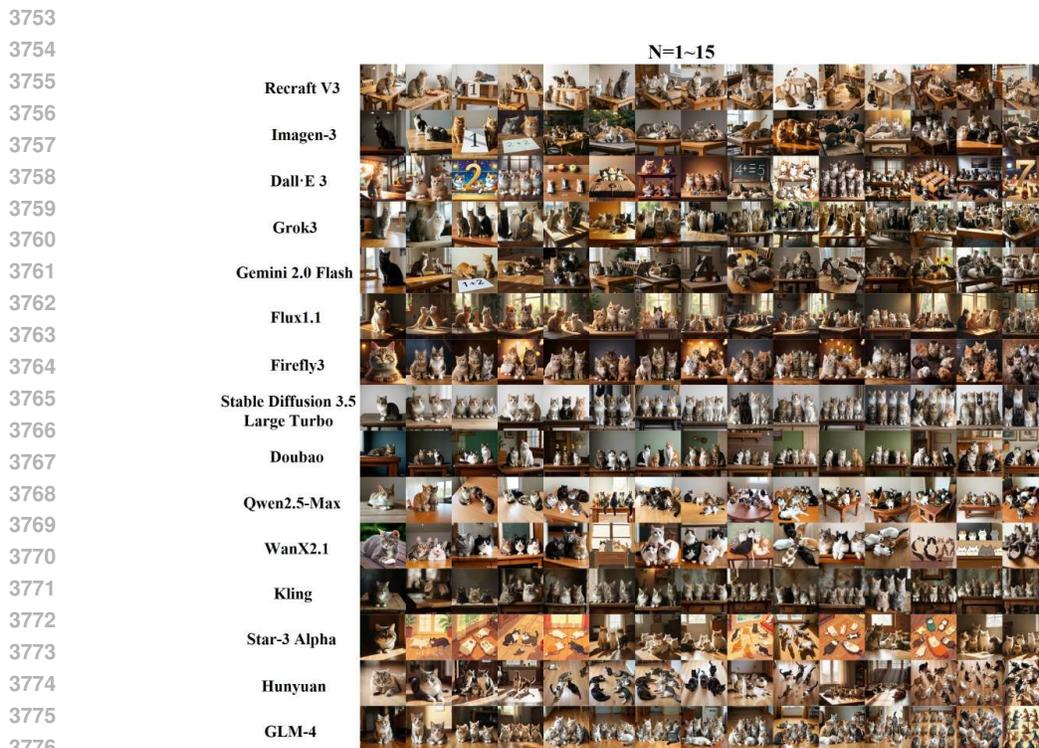
3696 **Figure 54: Counting Apples with Additive Decomposition on 15 Models.** This figure presents the
3697 generation results of counting apples with additive decomposition. We use $\lfloor N/2 \rfloor$ plus $N - \lfloor N/2 \rfloor$
3698 to replace the N in the Prompt: “ N apples on a wooden table.”
3699



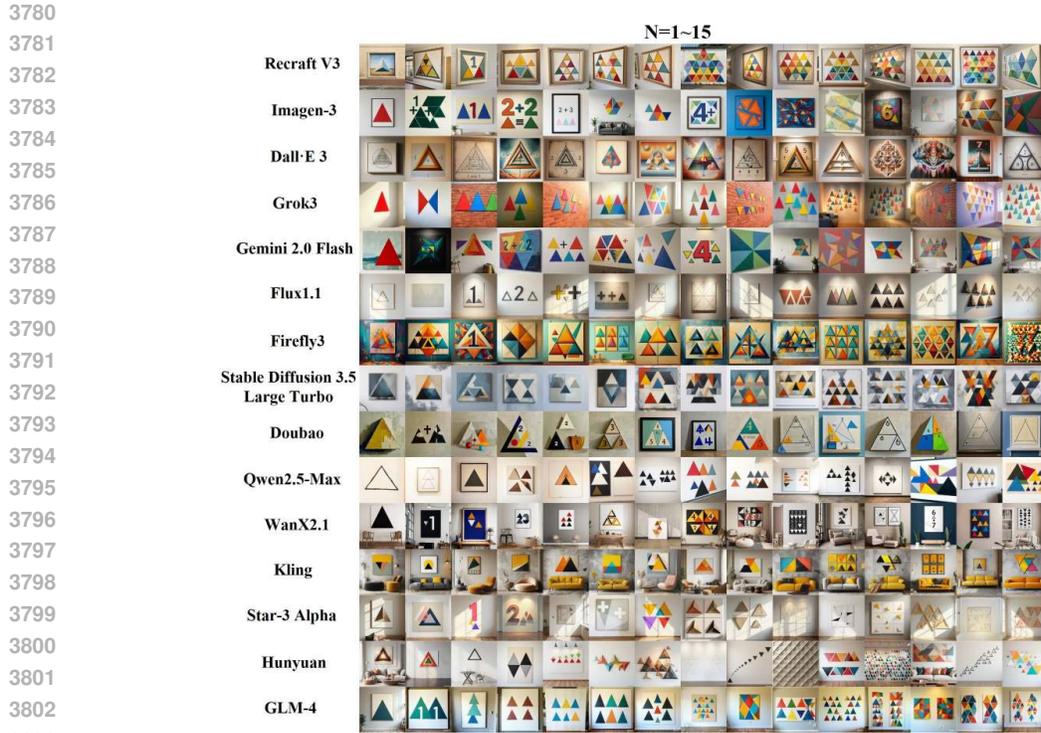
3723 **Figure 55: Counting Watermelons with Additive Decomposition on 15 Models.** This figure
3724 presents the generation results of counting watermelons with additive decomposition. We use $\lfloor N/2 \rfloor$
3725 plus $N - \lfloor N/2 \rfloor$ to replace the N in the Prompt: “ N watermelons on a wooden table.”



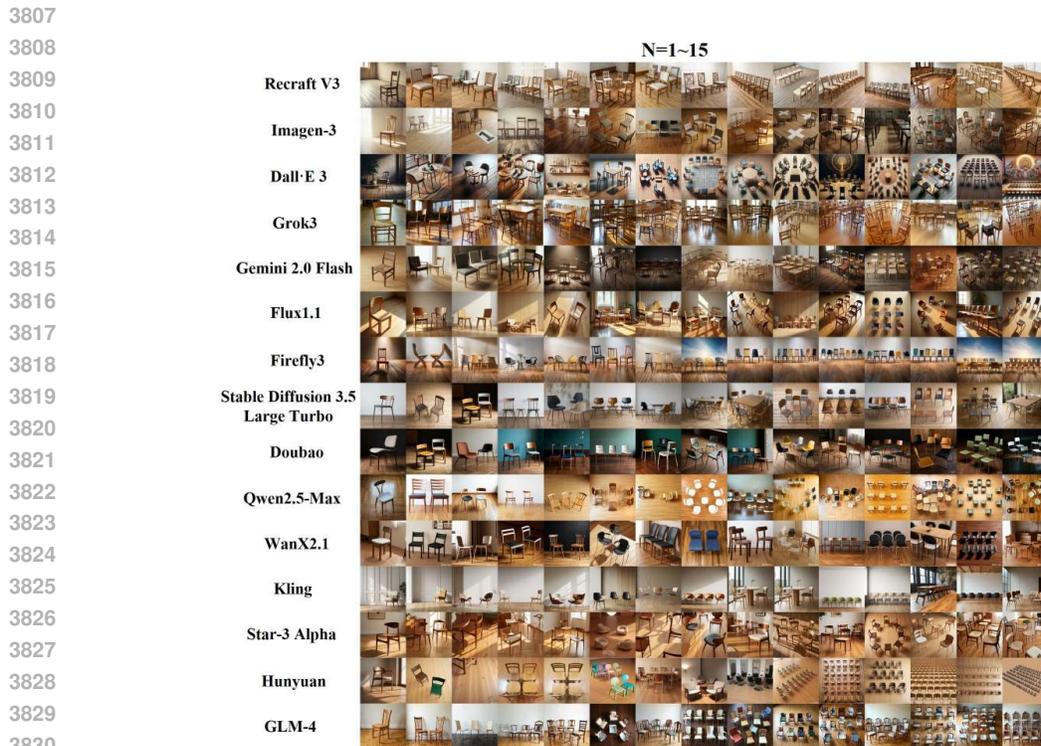
3750 **Figure 56: Counting Humans with Additive Decomposition on 15 Models.** This figure presents
3751 the generation results of counting humans with additive decomposition. We use $\lfloor N/2 \rfloor$ plus $N -$
3752 $\lfloor N/2 \rfloor$ to replace the N in the Prompt: “ N humans on a wooden floor.”



3777 **Figure 57: Counting Cats with Additive Decomposition on 15 Models.** This figure presents the
3778 generation results of counting cats with additive decomposition. We use $\lfloor N/2 \rfloor$ plus $N - \lfloor N/2 \rfloor$ to
3779 replace the N in the Prompt: “ N cats on a wooden table.”



3804 **Figure 58: Counting Triangles with Additive Decomposition on 15 Models.** This figure presents
3805 the generation results of counting triangles with additive decomposition. We use $\lfloor N/2 \rfloor$ plus $N -$
3806 $\lfloor N/2 \rfloor$ to replace the N in the Prompt: “ N triangle on a painting on a wall.”



3831 **Figure 59: Counting Chairs with Additive Decomposition on 15 Models.** This figure presents the
3832 generation results of counting chairs with additive decomposition. We use $\lfloor N/2 \rfloor$ plus $N - \lfloor N/2 \rfloor$
3833 to replace the N in the Prompt: “ N chairs on a wooden floor.”

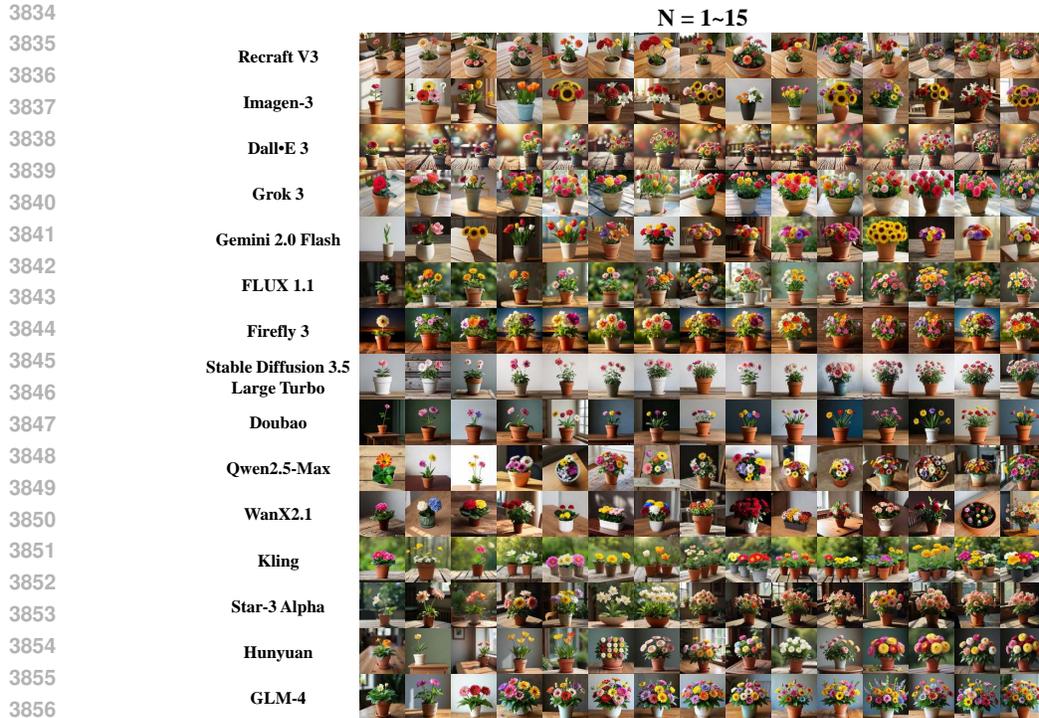


Figure 60: **Counting Flowers with Additive Decomposition on 15 Models.** This figure presents the generation results of counting flowers with additive decomposition. We use $\lfloor N/2 \rfloor$ plus $N - \lfloor N/2 \rfloor$ to replace the N in the Prompt: “A flower pot with N flowers on a wooden table.”

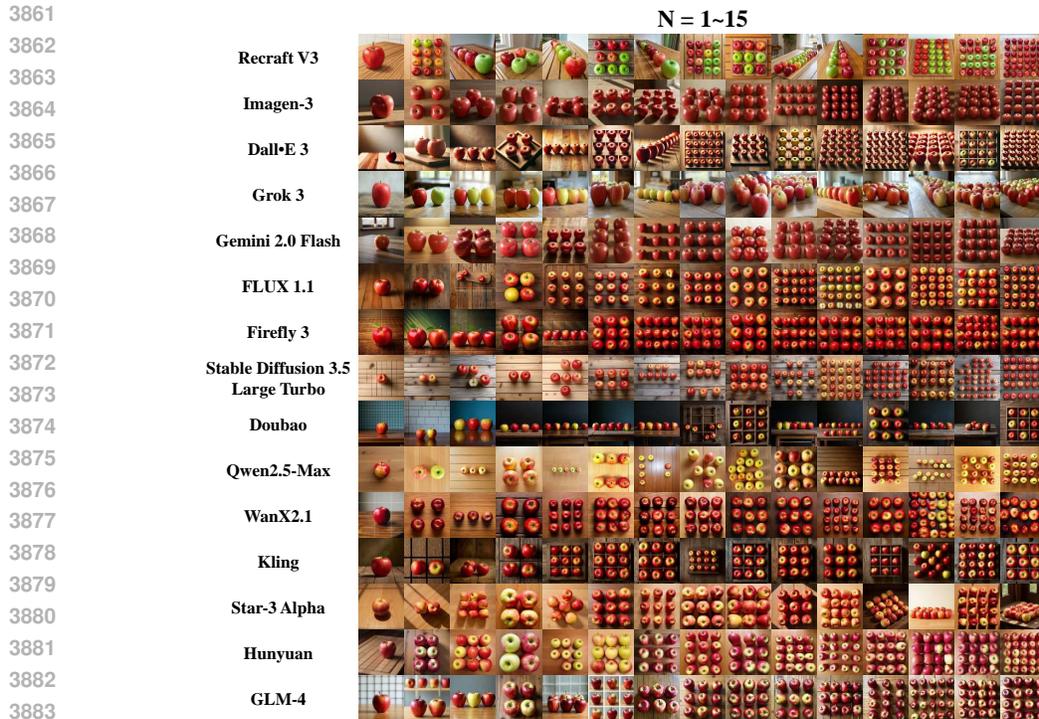
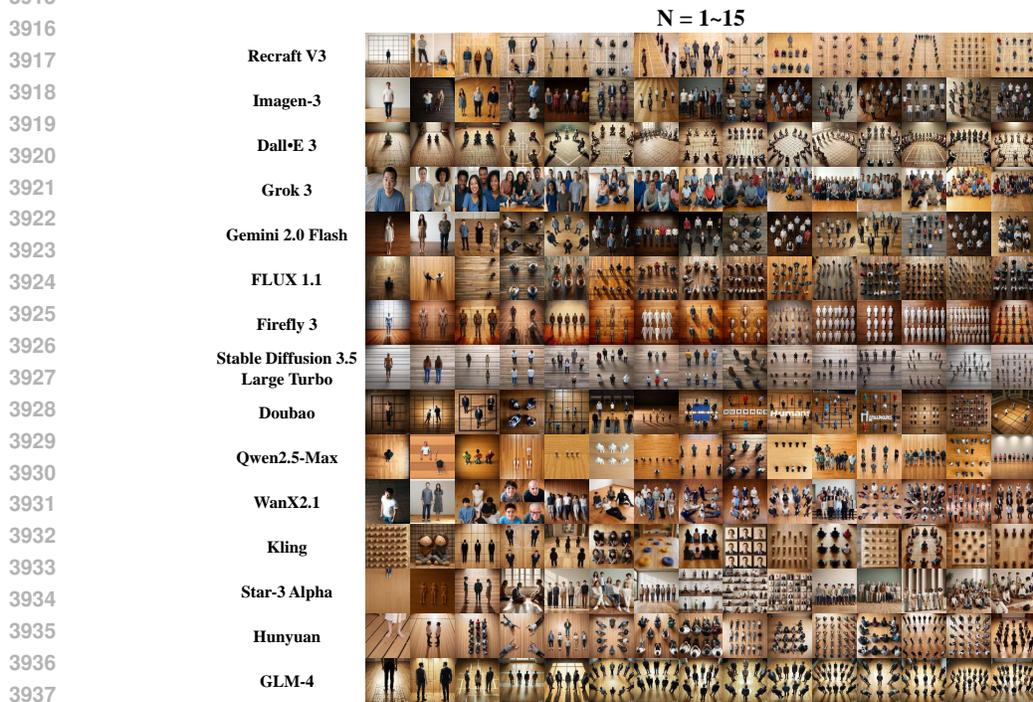


Figure 61: **Counting Apples with Grid Prior on 15 Models.** This figure presents the generation results of counting apples with grid prior. We use the prompt “ N apples on a wooden table, with r row c column grid”, where $N \in [1, 15]$ denotes the number of objects expected to be generated and $N = r \times c$.



3911 **Figure 62: Counting Watermelons with Grid Prior on 15 Models.** This figure presents the generation results of counting watermelons with grid prior. We use the prompt “ N watermelons on a wooden table, with r row c column grid”, where $N \in [1, 15]$ denotes the number of objects expected to be generated and $N = r \times c$.



3938 **Figure 63: Counting Humans with Grid Prior on 15 Models.** This figure presents the generation results of counting humans with grid prior. We use the prompt “ N humans on a wooden floor, with r row c column grid”, where $N \in [1, 15]$ denotes the number of objects expected to be generated and $N = r \times c$.

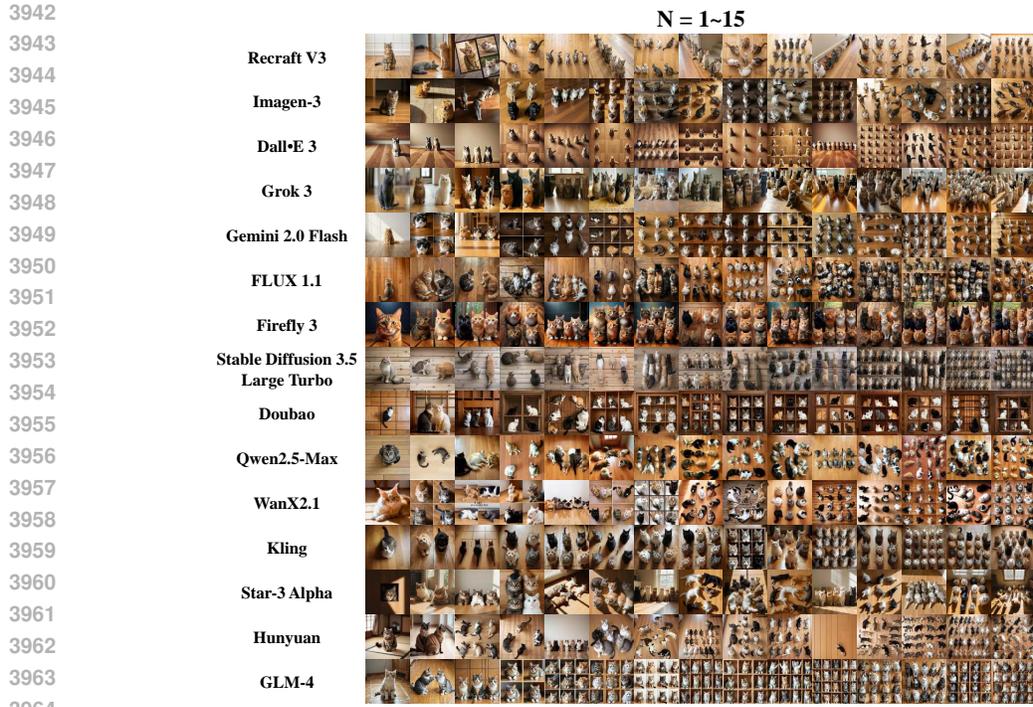


Figure 64: **Counting Cats with Grid Prior on 15 Models.** This figure presents the generation results of counting cats with grid prior. We use the prompt “ N cats on a wooden table, with r row c column grid”, where $N \in [1, 15]$ denotes the number of objects expected to be generated and $N = r \times c$.

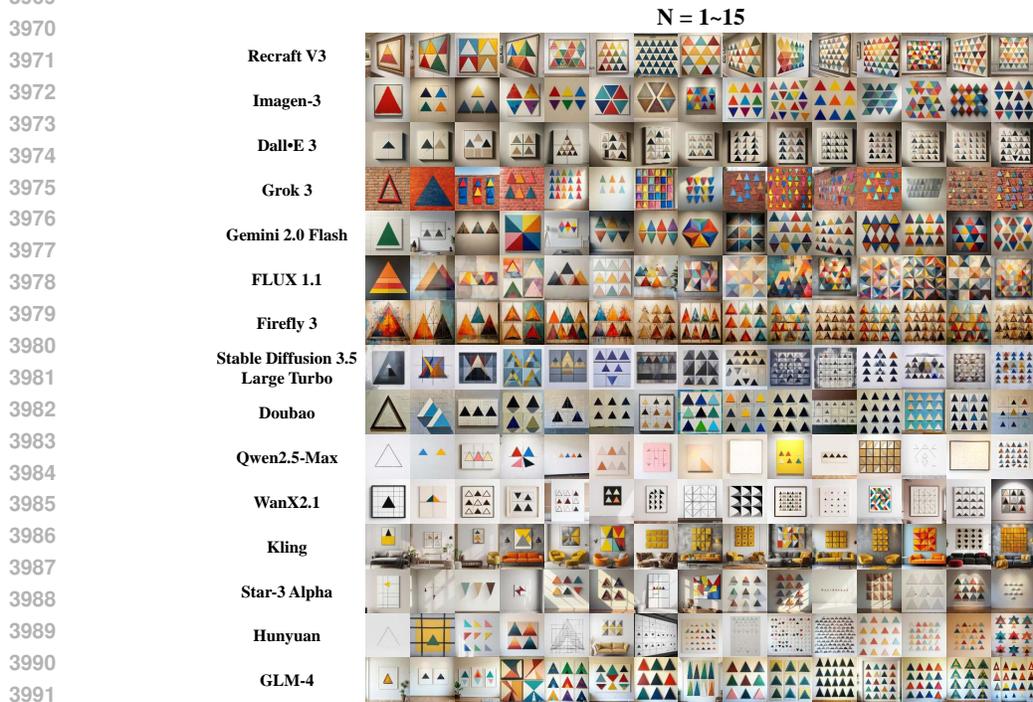
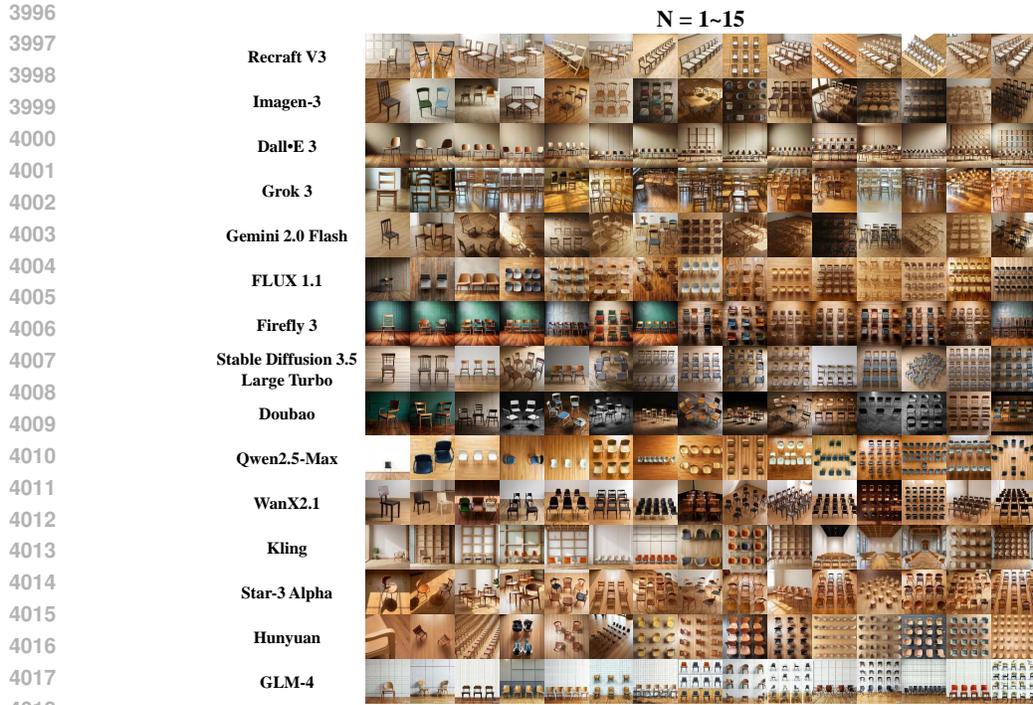
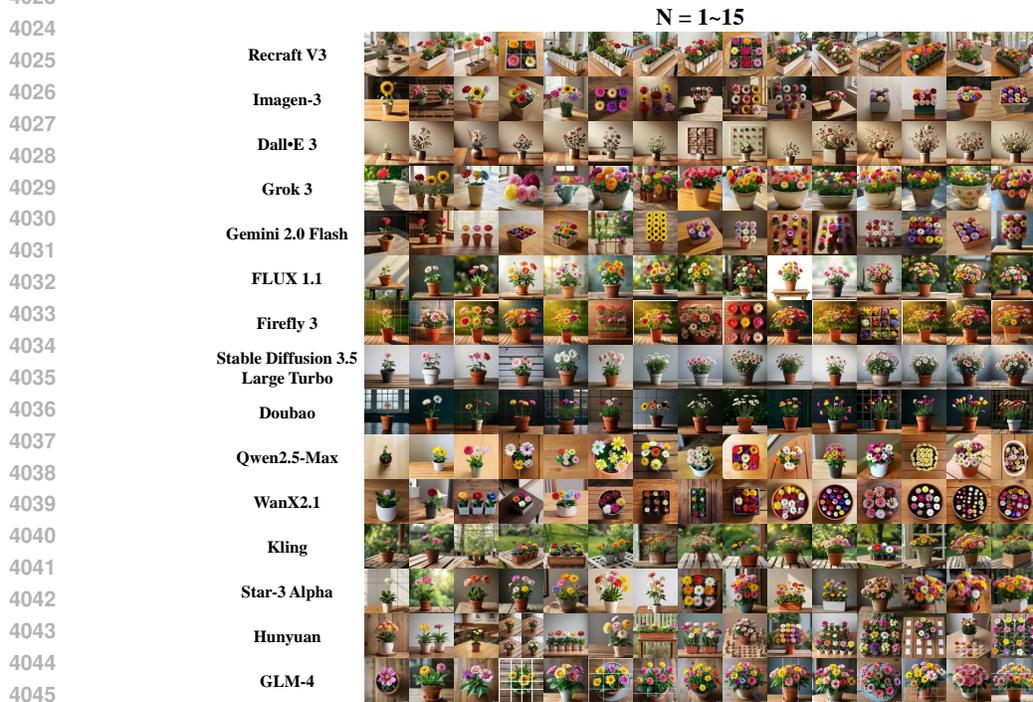


Figure 65: **Counting Triangles with Grid Prior on 15 Models.** This figure presents the generation results of counting triangles with grid prior. We use the prompt “ N triangles on a painting on a wall, with r row c column grid”, where $N \in [1, 15]$ denotes the number of objects expected to be generated and $N = r \times c$.



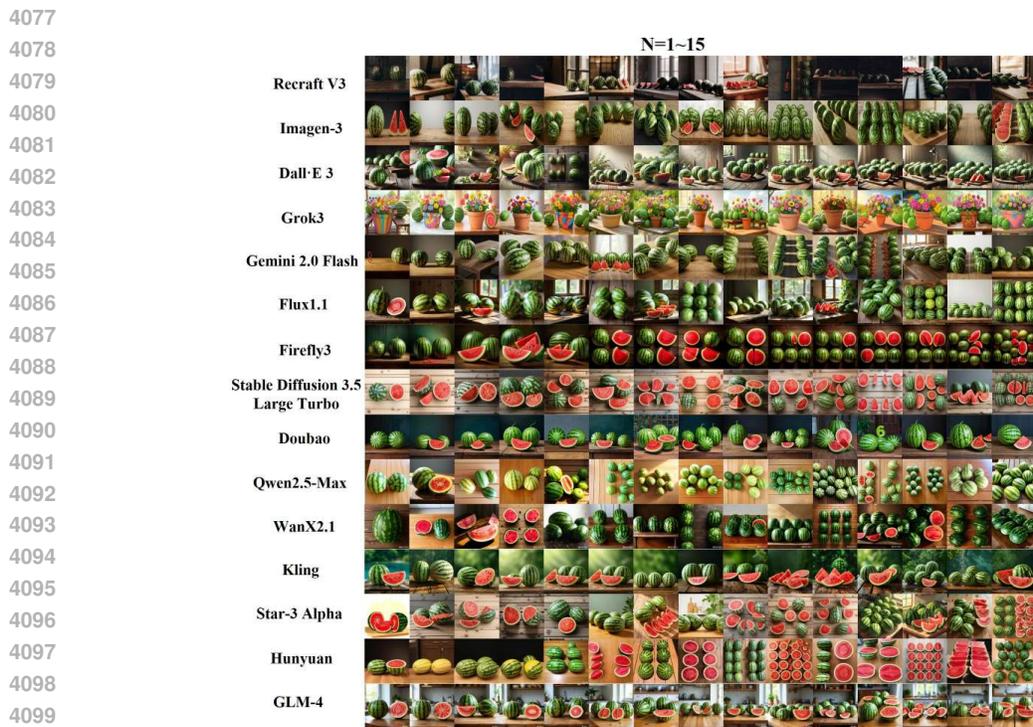
4019 **Figure 66: Counting Chairs with Grid Prior on 15 Models.** This figure presents the generation
4020 results of counting chairs with grid prior. We use the prompt “ N chairs on a wooden floor, with r
4021 row c column grid”, where $N \in [1, 15]$ denotes the number of objects expected to be generated and
4022 $N = r \times c$.



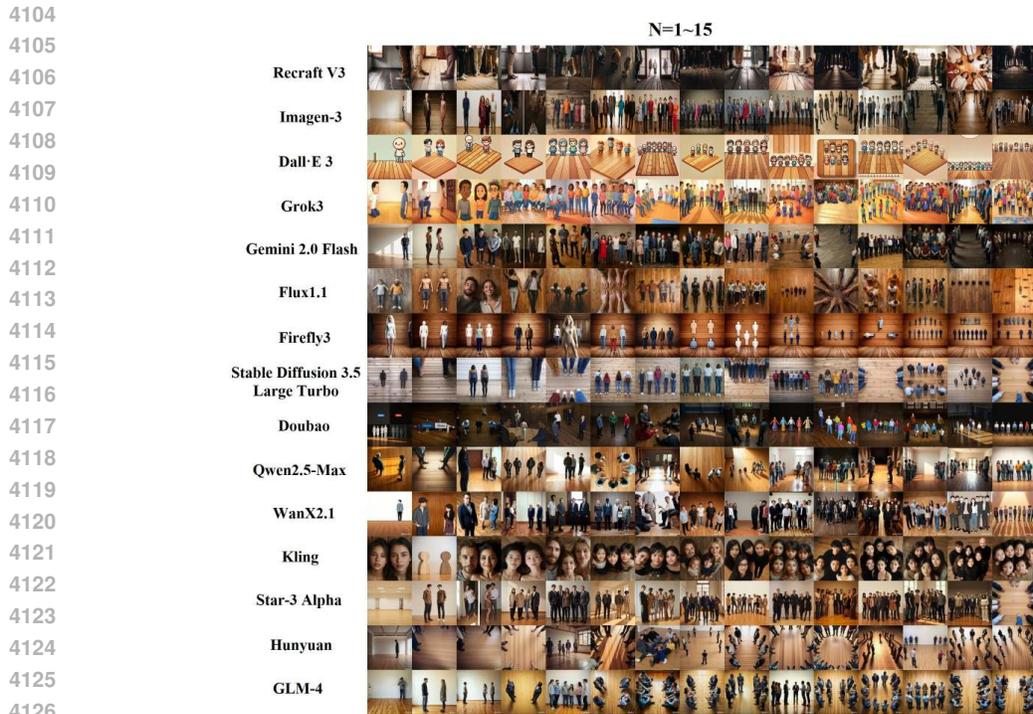
4046 **Figure 67: Counting Flowers with Grid Prior on 15 Models.** This figure presents the generation
4047 results of counting flowers with grid prior. We use the prompt “A flower pot with N flowers on a
4048 wooden table, with r row c column grid”, where $N \in [1, 15]$ denotes the number of objects expected
4049 to be generated and $N = r \times c$.



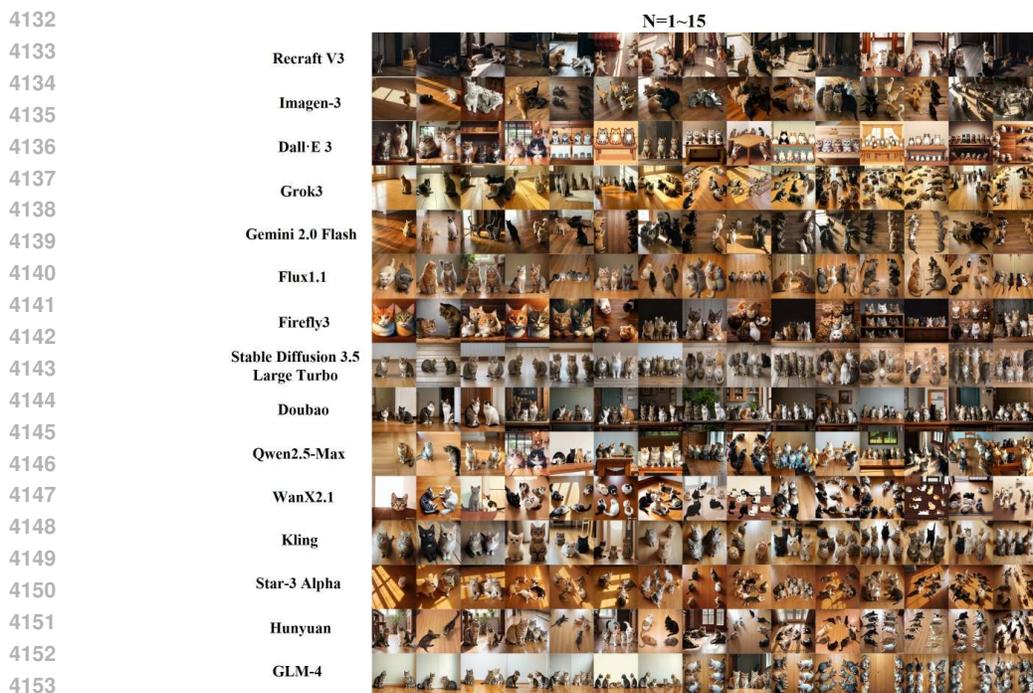
4073 **Figure 68: Counting Apples with Position Guidance on 15 Models.** This figure presents the
 4074 generation results of counting apples with position guidance. We use the prompt “[$\lfloor N/2 \rfloor$ apples on
 4075 the left, $N - \lfloor N/2 \rfloor$ apples on the right, on a wooden table”, where $N \in [1, 15]$ denotes the number
 4076 of objects expected to be generated.



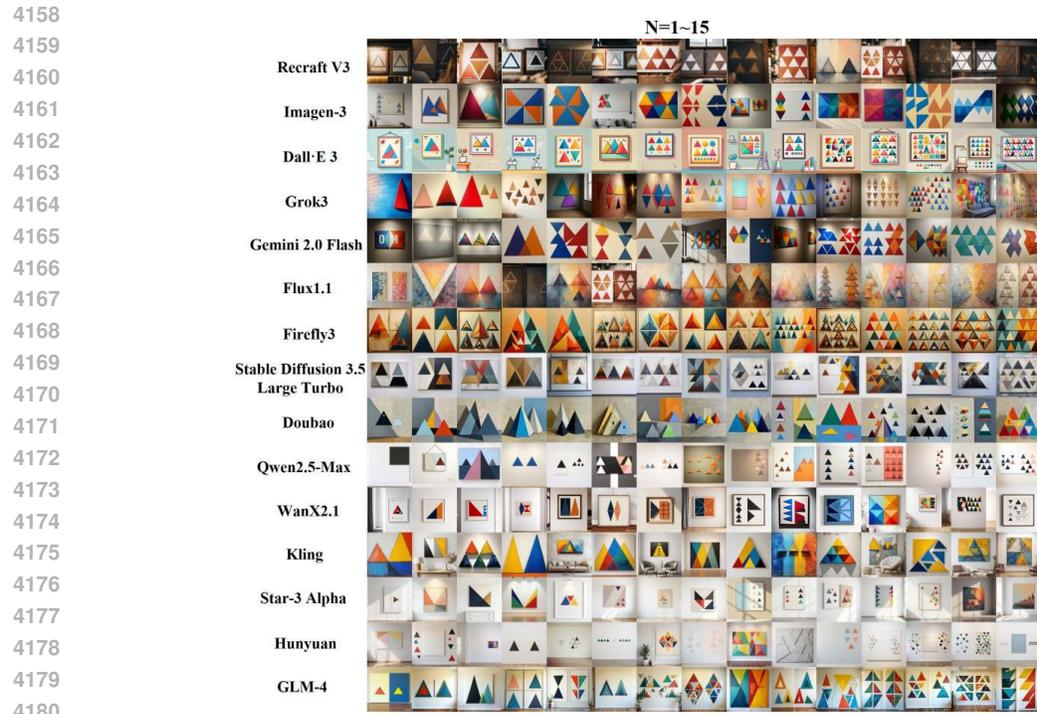
4100 **Figure 69: Counting Watermelons with Position Guidance on 15 Models.** This figure presents
 4101 the generation results of counting watermelons with position guidance. We use the prompt “[$\lfloor N/2 \rfloor$
 4102 watermelons on the left, $N - \lfloor N/2 \rfloor$ watermelons on the right, on a wooden table”, where $N \in$
 4103 $[1, 15]$ denotes the number of objects expected to be generated.



4127 **Figure 70: Counting Humans with Position Guidance on 15 Models.** This figure presents the
4128 generation results of counting humans with position guidance. We use the prompt “[$\lfloor N/2 \rfloor$ humans
4129 on the left, $N - \lfloor N/2 \rfloor$ humans on the right, on a wooden floor”, where $N \in [1, 15]$ denotes the
4130 number of objects expected to be generated.



4154 **Figure 71: Counting Cats with Position Guidance on 15 Models.** This figure presents the gen-
4155 eration results of counting cats with position guidance. We use the prompt “[$\lfloor N/2 \rfloor$ humans on the
4156 left, $N - \lfloor N/2 \rfloor$ humans on the right, on a wooden floor”, where $N \in [1, 15]$ denotes the
4157 number of objects expected to be generated.



4181 **Figure 72: Counting Triangles with Position Guidance on 15 Models.** This figure presents the
4182 generation results of counting triangles with position guidance. We use the prompt “ $\lfloor N/2 \rfloor$ triangles
4183 on the left, $N - \lfloor N/2 \rfloor$ triangles on the right, on a painting on a wall”, where $N \in [1, 15]$ denotes
4184 the number of objects expected to be generated.



4209 **Figure 73: Counting Chairs with Position Guidance on 15 Models.** This figure presents the
4210 generation results of counting chairs with position guidance. We use the prompt “ $\lfloor N/2 \rfloor$ chairs on
4211 the left, $N - \lfloor N/2 \rfloor$ apples on the right, on a wooden floor”, where $N \in [1, 15]$ denotes the number
of objects expected to be generated.

4212
4213
4214
4215
4216
4217
4218
4219
4220
4221
4222
4223
4224
4225
4226
4227
4228
4229
4230
4231
4232
4233
4234
4235
4236
4237
4238
4239
4240
4241
4242
4243
4244
4245
4246
4247
4248
4249
4250
4251
4252
4253
4254
4255
4256
4257
4258
4259
4260
4261
4262
4263
4264
4265



Figure 74: **Counting Flowers with Position Guidance on 15 Models.** This figure presents the generation results of counting flowers with position guidance. We use the prompt “[$N/2$] flowers on the left, $N - \lfloor N/2 \rfloor$ flowers on the right, on a wooden table”, where $N \in [1, 15]$ denotes the number of objects expected to be generated.

4266 LLM USAGE DISCLOSURE

4267
4268
4269
4270
4271
4272
4273
4274
4275
4276
4277
4278
4279
4280
4281
4282
4283
4284
4285
4286
4287
4288
4289
4290
4291
4292
4293
4294
4295
4296
4297
4298
4299
4300
4301
4302
4303
4304
4305
4306
4307
4308
4309
4310
4311
4312
4313
4314
4315
4316
4317
4318
4319

LLMs were used only to polish language, such as grammar and wording. These models did not contribute to idea creation or writing, and the authors take full responsibility for this paper's content.