

UniCorn: Towards Self-Improving Unified Multimodal Models through Self-Generated Supervision

Zhen Fang^{1*†} Ruiyan Han^{2*} Xinyu Sun^{2*} Yuchen Ma² Ziheng Wang² Yu Zeng^{1†}
Zehui Chen¹ Lin Chen¹ Wenxuan Huang^{3,4} Weijie Xu⁵ Yi Cao⁶ Feng Zhao^{1‡}

¹MoE Key Lab of BIPC, USTC ²FDU

³ECNU ⁴CUHK ⁵NJU ⁶SUDA

fazii@mail.ustc.edu.cn fzha0956@ustc.edu.cn

Abstract

While Unified Multimodal Models (UMMs) have achieved remarkable success in cross-modal comprehension, a significant gap persists in their ability to leverage such internal knowledge for high-quality generation. We formalize this discrepancy as *Conduction Aphasia*, a phenomenon where models accurately interpret multimodal inputs but struggle to translate that understanding into faithful and controllable synthesis. To address this, we propose *UniCorn*, a simple yet elegant self-improvement framework that **eliminates the need for external data or teacher supervision**. By partitioning a single UMM into three collaborative roles: Proposer, Solver, and Judge, *UniCorn* generates high-quality interactions via self-play and employs cognitive pattern reconstruction to distill latent understanding into explicit generative signals. To validate the restoration of multimodal coherence, we introduce *UniCycle*, a cycle-consistency benchmark based on a *Text* \rightarrow *Image* \rightarrow *Text* reconstruction loop. Extensive experiments demonstrate that *UniCorn* achieves comprehensive and substantial improvements over the base model across six general image generation benchmarks. Notably, it achieves **state-of-the-art (SOTA)** performance on T1IF(73.8), DPG(86.8), CompBench(88.5), and *UniCycle*(46.5), while further delivering substantial gains of **+5.0** on WISE and **+6.5** on OneIG. These results highlight that our method significantly enhances T2I generation while maintaining robust comprehension, demonstrating the scalability of fully self-supervised refinement for unified multimodal intelligence. The code is available at <https://github.com/Hungryyan1/UniCorn>.

1 Introduction

The realization of Artificial General Intelligence (AGI) requires a tight synergy between comprehension and generation, wherein comprehension enables the internalization of knowledge and generation allows its coherent and expressive externalization. By integrating multiple modalities into a shared representational space, Unified Multimodal Models (UMMs) (Deng et al., 2025a; Chen et al., 2025b; Xie et al., 2025c) naturally couple comprehension and generation as two complementary phases of a unified cognitive process, supporting both knowledge grounding and coherent reasoning.

Despite these advances, a fundamental disparity remains between comprehension and generation in current UMMs. This mismatch, which we formalize as **Conduction Aphasia**, arises when a model demonstrates strong domain understanding yet fails to translate that knowledge into high-quality generative outputs. As shown in Fig. 1, a representative case appears in image generation: although the model can accurately recognize what an image depicts and reliably assess its visual quality, it often cannot act on this knowledge during generation. This disconnect motivates a central research question: *how can a model’s robust understanding guide and strengthen its generative behavior?*

Driven by this simple yet fundamental question, we propose *UniCorn*, a post-training framework that enables self-improvement through a unified cycle of proposal, execution, and evaluation. Requiring no external data or teacher-model supervision, *UniCorn* allows UMMs to autonomously narrow the comprehension–generation gap by acting as their own instructor within a single parameter space. Motivated by the observation that a single UMM can exhibit distinct capabilities for proposing, executing, and evaluating, we treat the model as a modular system in which comprehension can explicitly guide generation. This design turns the

* Equal contribution, ordered alphabetically.

† Project lead.

‡ Corresponding author.




Figure 1: **Motivation of UniCORN.** UMMs often exhibit an understanding-generation gap: they can accurately understand and critique errors in an image, yet fail to generate the same scene correctly. This conduction aphasia motivates our framework to leverage the model’s superior internal understanding to strengthen and refine its generative capabilities through self-contained feedback.

model’s latent interpretive capability into an internal training signal, enabling autonomous generative improvement without external supervision.

Specifically, **UniCORN** operates through a self multi-agent framework that functionalizes the UMM into three distinct internal roles. The process begins with the model acting as a **Proposer** to propose diverse and expansive prompts, followed by its transition into a **Solver** to synthesize corresponding image candidates. Finally, it assumes the role of a **Judge** to provide evaluative rewards based on its superior comprehension.

By simulating structured collaboration within a single parameter space, this design yields rich interaction data that are refined through data reconstruction. Concretely, we convert raw multi-agent outputs into structured training signals, including descriptive captions, evaluative judgments, and reflective feedback, thereby distilling latent understanding into explicit supervision for effective self-improvement.

To determine whether internal collaboration produces general multimodal intelligence instead of narrow task fitting, we introduce **UniCycle**, a cycle-consistency benchmark that probes cognitive alignment via informational integrity. Existing evaluations often separate comprehension and generation, which can lead to piecemeal measurements and biased conclusions. In contrast, **UniCycle** frames evaluation as a Text \rightarrow Image \rightarrow Text reconstruction process. It compares the model’s original intent with its reconstructed description, using the resulting semantic gap as a holistic, training-free indicator of conceptual coherence, while reducing the bias that arises when capabilities are tested in isolation.

Across extensive experiments, we find that our model achieves reliable self-improvement without



Figure 2: **Visualization results of UniCORN.**

heuristic reward engineering, curriculum design, or external supervision. Compared with prior self-improvement approaches (Jin et al., 2025a) and methods that depend on external guidance, our approach learns from internally generated training signals, generalizes well, and remains stable under out-of-distribution (OOD) conditions. These results support the effectiveness of a fully self-contained learning paradigm.

- We identify the **Conduction Aphasia** phenomenon in UMMs, where strong understanding fails to translate into accurate generation, and propose **UniCORN**, which repurposes internal comprehension as self-supervision through Proposer, Solver, and Judge roles with data reconstruction.
- To assess whether multimodal understanding and generation remain conceptually consistent across modality transitions, we introduce **UniCycle**, a training-free evaluation protocol that measures multimodal coherence through a Text \rightarrow Image \rightarrow Text cycle.
- Experimental results demonstrate that our method consistently outperforms prior approaches, achieving **SOTA** performance on TIF (73.8), DPG (86.8), CompBench (88.5), and **UniCycle**(46.5), together with substantial improvements of **+4.0** on Geneval, **+5.0** on WISE, and **+6.5** on OneIG.

2 Related Work

Unified Multimodal Models UMMs aim to unify cross-modal understanding and generation,

yet strong understanding often fails to yield equally strong native generation. Existing designs fall into two paradigms: *pure autoregressive* models that jointly predict text and visual tokens over interleaved sequences (Chen et al., 2025e; Cui et al., 2025; Tong et al., 2025) and *hybrid* models that combine autoregressive language modeling with diffusion-based image synthesis, either within a unified backbone (Xie et al., 2024; Zhao et al., 2024) or via modular routing and sparse experts (Shi et al., 2024; Liang et al., 2024b; Deng et al., 2025b), with related guidance schemes such as Diffusion Forcing (Chen et al., 2024a). Beyond architecture, self-improvement methods convert self-generated signals into training objectives (Yu et al., 2025; Zhou et al., 2024b; Wang et al., 2025c); for UMMs, SRUM derives internal rewards from understanding (Jin et al., 2025a), and UniRL jointly optimizes understanding and generation (Mao et al., 2025). However, most pipelines depend on auxiliary components or task-specific feedback, limiting scalability and generalization.

Multi-Agent and Self-Improvement Learning

Multi-agent systems decompose reasoning through role specialization and interaction, enabling solution diversity and cross-verification, but often incur high coordination cost and brittle verification (Chen et al., 2024d; Liang et al., 2024a; Cemri et al., 2025). In parallel, LLM self-improvement converts self-generated tasks and evaluations into training signals, supporting zero-data learning via self-play and self-rewarding mechanisms (Silver et al., 2017; Huang et al., 2025a; Zhao et al., 2025a; Yuan et al., 2024). Unified Multimodal Models (UMMs) naturally unify understanding and generation within a single parameter space, making them particularly well-suited for lightweight role instantiation and fully model-driven self-improvement without external supervision.

3 Method

In this section, we begin by presenting the motivation through an analysis of the mismatch between generation and understanding capabilities in UMMs. Building on these observations, we introduce *UniCorn*, a simple yet elegant post-training framework that enables self-improvement without any external annotated data or teacher models.




Figure 3: **Results of BAGEL (Deng et al., 2025a) and GPT-4o (Hurst et al., 2024) on four understanding benchmarks.** For Omni-RewardBench (Jin et al., 2025b) and MMRB2 (Hu et al., 2025), we evaluate the T2I task. Performances are normalized with GPT-4 (Achiam et al., 2023) results for better visualization.

3.1 Motivation

Just as a child who associates the word “apple” with the fruit can spontaneously name it upon seeing it, cognitive symmetry (Blanco, 2018) enables a bidirectional mapping between internal concepts and external expressions. This alignment is reminiscent of escaping Plato’s Cave: true intelligence must move beyond observing surface data to mastering the reciprocal relationship between an appearance and its underlying source.

However, current UMMs suffer from a functional deficit akin to **Conduction Aphasia**: while the model exhibits profound comprehension, its generative performance remains fractured, failing to produce the very content it can inherently understand. Bridging this gap is critical; without aligning these dual processes, a model remains a “passive observer,” capable of grounding symbols but incapable of utilizing them. Mastering the synergy between understanding and generation is thus not merely a functional upgrade but the essential step toward achieving the cognitive integrity required for AGI.

On the one hand, as illustrated in Fig. 3, current UMMs demonstrate formidable perception and comprehension capabilities. Specifically, when serving as a reward model for Text-to-Image (T2I) generation, the UMM exhibits a sophisticated grasp of cross-modal semantics. This suggests that the model has already internalized a robust ‘world model’ and possesses the necessary latent knowledge to discern high-quality visual-textual alignments.

On the other hand, the model’s generative capability remains markedly constrained, primarily due to its failure to bridge the gap between internal recognition and active synthesis. This functional dissociation means that the UMM’s own sophisticated understanding remains a ‘silent pas-

senger’ during the generative process, unable to inform or correct its outputs. Building on this observation, our key insight is that **the UMM’s formidable comprehension can be repurposed as an autonomous supervisory signal to steer its generative behavior**. By transforming latent interpretive depth into explicit guidance, we promote a tighter coupling between these two processes, ultimately restoring the cognitive symmetry essential for a truly integrated multimodal intelligence.

3.2 Problem Definition

We study UMMs that process interleaved image-text inputs and outputs. A UMM is formulated as a policy π_θ that maps a multimodal input sequence

$$X = (x_1, \dots, x_N), \quad x_n \in T \cup I, \quad (1)$$

to an interleaved multimodal output sequence $Y = \pi_\theta(X)$. This unified input-output formulation supports both Image-to-Text (I2T) understanding and Text-to-Image (T2I) generation. We operationalize understanding as I2T and generation as T2I, and leverage the model’s stronger I2T understanding to supervise and refine its weaker T2I generation.

3.3 UniCorn

UniCorn operates via two core stages: Self Multi-Agent Sampling and Cognitive Pattern Reconstruction (CPR). First, the UMM concurrently assumes three roles: Proposer, Solver, and Judge (§ 3.3.1), to simulate a collaborative loop. Then, the CPR stage reconstructs these raw interactions into three training patterns: caption, judgement, and reflection (§ 3.3.2), which are combined with high-quality self-sampled T2I generation data for post-training. Critically, the entire process is **fully self-contained, requiring no external teacher models or human-annotated data**.

3.3.1 Stage 1: Self Multi-Agent Sampling

LLMs are naturally suited for self-play in multi-task settings (Radford et al., 2019). For UMMs, interleaved multimodal inputs and functional diversity allow prompting, generation, and judgement to coexist within a shared model, enabling role-conditioned behaviors under different prompts. We leverage this property to functionalize a single UMM into collaborative roles, bridging the comprehension–generation gap through internal synergy.

Proposer $\pi_\theta(T | T)$ The proposer is designed to generate a diverse set of challenging prompts for the unified multimodal model, which are subsequently used to produce training images. To this end, inspired by LAION-5B (Schuhmann et al., 2022) and COYO-700M (Byeon et al., 2022), we partition all T2I task prompts into ten categories and designed fine-grained generation rules for each category. Next, we prompt UMM to generate an initial batch of prompts and act as the judge to select the best candidate for subsequent iterations. Leveraging the strong in-context learning (ICL) capabilities of LLMs (Dong et al., 2024), the initial example serves as a few-shot demonstration to guide the generation of subsequent prompts. To further enhance diversity, we introduce a dynamic seeding mechanism. After generating a predefined number of prompts, several examples are sampled from the prompt library for evaluation and then used to construct new demonstrations that guide the next round of prompt generation. Compared with prior approaches that either directly rely on training set (Jin et al., 2025a) or employ external models for prompt construction (Mao et al., 2025), our method requires no external data and generates more diverse prompts, thereby improving generalization.

Solver $\pi_\theta(I | T)$ The solver is responsible for producing a diverse set of outputs in response to the prompts generated by the proposer. Therefore, we encourage the UMM to generate images under random seeds and different hyperparameters. Following DeepSeek-R1 (Guo et al., 2025a), we perform 8 rollouts per prompt to strike a favorable trade-off between sample quality, diversity, and computational efficiency.

Judge $\pi_\theta(T | T, I)$ The judge is responsible for assigning scores to the images generated by the solver in response to prompts proposed by the proposer, which are then used for rejection sampling during training.

Previous work has relied on heuristic reward functions based on keywords (Mao et al., 2025) or on powerful external models to provide dense reward maps (Jin et al., 2025a). Such reward judges depend heavily on parameter tuning and the performance of external models, which varies across tasks, thereby severely limiting the generalization of self-improvement. As illustrated in Fig. 3, UMMs exhibit strong reward modeling capabilities. Thus, we formulate reward evaluation for

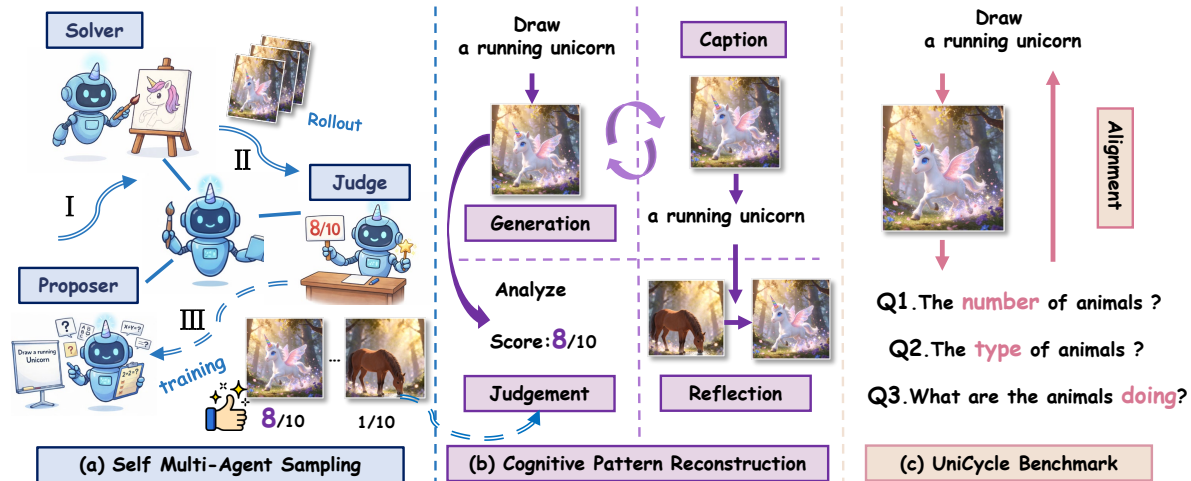


Figure 4: **Overview of the UniCORN Framework.** (a) Illustrates the self-multi-agent collaboration for high-quality data sampling. (b) Details the Cognitive Pattern Reconstruction process, which reorganizes data to facilitate robust and efficient learning. (c) Presents the UniCycle benchmark evaluation, verifying whether the model can accurately reconstruct key textual information from its own generated content.

all T2I tasks using discrete scores ranging from 0 to 10, following a widely adopted LLM-as-a-judge paradigm (Radford et al., 2019; Kim et al., 2023). To further enhance judgement quality, we transfer generation reward models (Liu et al., 2025c), which have demonstrated strong potential in LLMs, to T2I evaluation. Specifically, we design task-specific rubrics for each category and encourage the model to explicitly articulate its reasoning before producing the final score.

3.3.2 Stage 2: Cognitive Pattern Reconstruction

Through self multi-agent rejection sampling using the Proposer–Solver–Judge pipeline, we obtain a batch of high-quality prompt–image pairs. While these pairs reflect a mapping from abstract conceptual spaces to high-dimensional visual manifolds, directly optimizing this cross-domain alignment remains stochastic and inefficient, often leading to mode collapse (Chen et al., 2025a; Zhenyu et al., 2024). To move beyond this "black-box" optimization, we draw inspiration from metacognitive theory (Dunlosky and Metcalfe, 2008), which identifies monitoring, evaluation, and regulation as the pillars of robust learning. Based on this insight, we propose a tripartite data architecture that reclaims and structures the overlooked trajectories from the self-play cycle. By replaying these latent interactions as explicit caption, judgement, and reflection patterns, we respectively ground abstract concepts in visual features, provide evaluative signals, and encode self-correction processes. This

design transforms the previously discarded internal "inner monologue" into a structured supervisory signal, fostering cognitive symmetry without external intervention.

CAPTION To establish robust **semantic grounding**, this pattern ensures the model internalizes the conceptual essence of its own creations by optimizing the inverse mapping $\pi_{\theta}(T | I^*)$. By treating the highest-scoring image I^* as the input and its originating prompt T as the ground truth, the model learns to **anchor** abstract concepts within the specific visual manifolds it is capable of synthesizing, thereby reinforcing the bidirectional cognitive symmetry between internal concepts and external manifestations.

JUDGEMENT This pattern focuses on **evaluative calibration** to refine the model’s internal value system. We train the model to predict the evaluative signal J for any generated pair, formulated as $\pi_{\theta}(J | T, I)$. By leveraging the task-specific rubrics and reasoning traces provided by the Judge, the model develops an acute perception of the latent gap between its current output and the ideal objective, providing a critical diagnostic signal for stabilizing the generative process.

REFLECTION Inspired by Reflexion (Shinn et al., 2023), this pattern introduces **iterative regulation** to enhance the model’s capacity for self-evolution. Leveraging the Solver’s multiple rollouts $\{I_1, \dots, I_n\}$, we utilize the rewards assigned by the Judge to identify pairs of contrasting quality, specifically selecting a high-reward "winning" image I^* and a lower-reward "losing" image I_{lose} .

from the same prompt. We then construct reflection trajectories formulated as $\pi_\theta(I^* | T, I_{lose}, J)$, which explicitly encode the transition from sub-optimal states to superior ones. By learning to transform the lower-quality manifestation I_{lose} into its optimized counterpart I^* , the model internalizes a mechanism for self-correcting generative errors, effectively mitigating mode collapse without the need for external supervision.

These three data types are combined with high-quality self-sampled T2I generation data to fine-tune the UMM. Note that the whole reconstruction procedure is rule-based and does not introduce any complexity. Detailed generation pipeline and examples can be found in Appendix B.

3.4 UniCycle

To assess whether internal collaboration yields genuine multimodal intelligence rather than task-specific performance gains, we introduce **UniCycle**, a cycle-consistency benchmark that measures information preservation under a **Text** \rightarrow **Image** \rightarrow **Text** loop. Given an instruction, **UniCycle** evaluates whether a unified multimodal model can recover instruction-critical semantics from its *own* generated image through subsequent visual understanding.

Based on TIIF (Wei et al., 2025), we generate QA pairs to probe instruction-implied attributes grounded in the generated image, extending the original TIIF benchmark from the T2I setting to the Text-to-Image-to-Text (T2I2T) setting. After annotation, we obtain 1,401 TIIF-style instances that cover more than ten task categories and span multiple question formats, including multiple-choice, binary (yes/no), and open-ended questions.

For evaluation, given a prompt T , the model first generates an image and then answers each question q_k independently conditioned on the generated image. An external judge model assesses whether each predicted answer \hat{y}_k is consistent with the initial prompt T and the reference answer a_k , and produces a score for each question.

We define a unified metric to quantify this T2I2T consistency.

Let $\mathcal{Q}(T)$ denote the set of questions associated with a prompt T . We define

$$\begin{aligned} \text{Soft}(T) &= \frac{1}{|\mathcal{Q}(T)|} \sum_{k \in \mathcal{Q}(T)} s_k, \\ \text{Hard}(T) &= \mathbb{1}[\forall k \in \mathcal{Q}(T), s_k = 1]. \end{aligned} \quad (2)$$

where s_k denotes the judge score for question q_k , defined as a binary indicator for non-text questions and as the proportion of correctly recovered Keywords to enable a more fine-grained and continuous metric for text-type questions.

The final Soft and Hard scores are obtained by averaging over all prompts. Additional details on data construction and evaluation prompt templates are provided in Appendix E.

4 Experiments

In this section, we first introduce the experiment setup, and conduct extensive experiments to demonstrate the effectiveness of our method.

4.1 Experiment Setup

Implementation We adopt BAGEL (Deng et al., 2025a) as the base model for our main experiments. The Proposer generates 5,000 prompts, then the Solver rolls out 8 times for each prompt. Training is conducted for 600 steps on 8 NVIDIA H800 GPUs for about 7 hours with a constant learning rate of 1×10^{-5} . Additional details are provided in the Appendix B.

Baselines and Benchmarks To validate our method, we compare it against three categories of approaches. First, we consider baseline models, including T2I frameworks: SD3 Medium (Esser et al., 2024), FLUX.1-dev (Labs, 2024) and unified multimodal models: Janus-Pro (Chen et al., 2025d), Show-o2 (Xie et al., 2025c), BLIP3-o (Chen et al., 2025b), UniGen (Tian et al., 2025), TwiG (Guo et al., 2025c) and T2I-R1 (Jiang et al., 2025a).

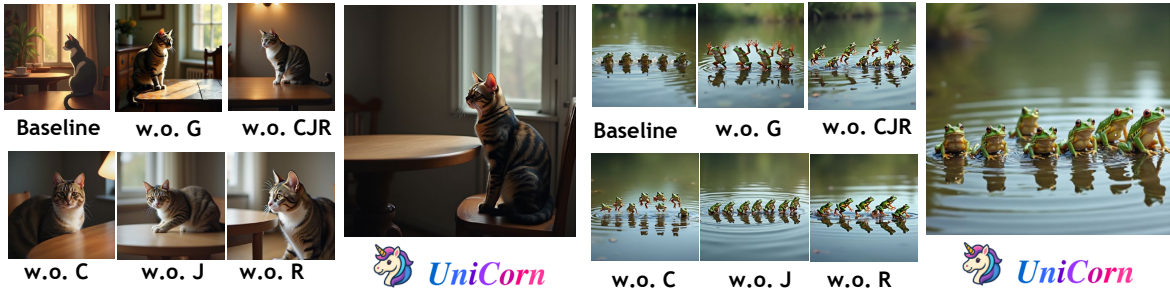
Regarding evaluation, we focus on TIIF (Wei et al., 2025), WISE (Niu et al., 2025), OneIG-EN (Chang et al., 2025), CompBench (Kil et al., 2024), DPG (Hu et al., 2024), and Geneval (Ghosh et al., 2023) to assess generation performance. To evaluate understanding, we further report results on standard benchmarks including MME (Fu et al., 2023), MMB (Liu et al., 2024b), MMMU-val (Yue et al., 2024), MMVP (Tong et al., 2024), and MM-Star (Chen et al., 2024c).

4.2 Main Results

As shown in Table 1, **UniCorn** achieves highly competitive performance across five T2I benchmarks. Our method significantly enhances fine-grained instruction following on TIIF, particularly improving robustness to short prompts (+3.7 points). On the comprehensive OneIG benchmark,

Model	TIIF \uparrow		WISE \uparrow			OneIG-EN \uparrow			CompBench \uparrow			DPG \uparrow	Geneval \uparrow
	Short	Long	Physics	Chemistry	Overall	Text	Alignment	Overall	Numeracy	3d Spatial	Overall	Score	Score
<i>Generation Only Models</i>													
SD3 Medium	64.8	64.8	47.0	29.0	42.0	40.7	80.6	42.8	72.8	77.8	84.3	84.1	74
FLUX.1 dev	66.2	66.7	51.0	35.0	50.0	<u>52.3</u>	78.6	43.4	75.3	76.4	83.1	83.8	<u>82</u>
<i>Unified Multimodal Models</i>													
Janus-Pro	65.4	61.1	42.0	26.0	35.0	0.1	55.3	26.7	56.4	76.2	74.0	84.3	80.0
show-o2	62.8	63.9	<u>63.0</u>	49.0	61.0	0.2	<u>81.7</u>	30.8	69.7	88.6	82.8	<u>86.1</u>	76.0
BLIP3-o	58.8	58.7	<u>63.0</u>	37.0	52.0	1.3	71.1	30.7	71.7	81.7	84.7	80.7	84.0
OmniGen2	70.2	70.3	52.0	34.0	48.0	68.0	80.4	47.5	72.0	82.2	<u>85.8</u>	83.6	80.0
TwIG \dagger	-	-	-	-	-	-	-	-	61.9	38.9	-	-	-
T2I-R1	67.6	68.3	55.0	30.0	54.0	7.3	80.4	27.7	<u>83.3</u>	79.4	81.9	-	77.0
BAGEL	<u>71.0</u>	<u>71.8</u>	57.0	43.0	50.0	24.4	76.9	36.1	70.4	78.0	82.2	84.0	78.0
UniCorn(Ours)	74.7	72.9	67.0	<u>47.0</u>	<u>55.0</u>	46.8	84.1	<u>42.6</u>	83.5	84.1	88.5	86.8	<u>82.0</u>
Δ (vs. BAGEL)	+3.7	+1.1	+10.0	+4.0	+5.0	+22.4	+7.2	+6.5	+13.1	+6.1	+6.3	+2.8	+4.0

Table 1: **Evaluation results on TIIF, WISE, OneIG-EN, CompBench, DPG, and Geneval benchmarks.** Arrows (\uparrow) denote that higher is better. **Bold** indicates the best performance across all models, and the second best is underlined. The WISE score is normalized to a 0–100 scale for visualization. Detailed comparison is listed in Appendix F.1.



Prompt: A cat is positioned to the **right** of a table.

Prompt: A photo of **seven** frogs are on the lake.

Figure 5: **Qualitative comparison between UniCorn, BAGEL and UniCorn’s adifferent data settings.** Our method jointly balances visual aesthetics, prompt fidelity, and realism in generation.

UniCorn yields a 6.5-point overall improvement, with a remarkable 22.4-point gain in the Text sub-task, indicating superior internalization of underlying knowledge. Furthermore, **UniCorn** achieves a 5 point gain on the knowledge-intensive WISE benchmark and a 6.3 point boost on CompBench. Notably, the substantial improvements in Numeracy (+13.1) and 3D Spatial (+6.1) tasks demonstrate the effective transfer of structured understanding into faithful synthesis, with **UniCorn** even surpassing **GPT-4o** on DPG benchmark (86.8 vs 86.2). These results consistently demonstrate that our self-play framework enables UMMs to bridge the gap between multimodal understanding and controllable generation, achieving robust performance that rivals state-of-the-art closed-source models.

4.3 Ablation Study

This section conducts ablation studies on data pattern, model architecture, and dataset size to further analyze our method.

4.3.1 Data Pattern

This section deconstructs multimodal data patterns to demonstrate how Cognitive Pattern Reconstruction bridges the gap between understanding and generation within a unified framework.

Table 2 reveals a hierarchical synergy between data patterns: while relying solely on generation (w.o. CJR) maintains basic instruction following (TIIF-S: 72.3), it triggers a catastrophic collapse of the latent space, evidenced by the sharp drop in MME-P from 1685.0 to 311.0. This proves that unconstrained generative training without semantic grounding leads to mode collapse. Conversely, incorporating Cognitive Pattern Reconstruction patterns (C, J, R) stabilizes the model; Judgment and Reflection provide evaluative signals that boost complex generative quality (TIIF-R: 78.4), while Captioning preserves the multimodal foundation and spatial reasoning capabilities. Finally, although removing generation (w.o. G) maintains comprehension metrics like MME-P (1669.0), it stalls generative growth, resulting in lower TIIF scores (73.4). Qualitative comparisons are shown in Fig. 5.

Setting	TIIF-S	TIIF-R	MME-P	MME-C	MMB	MMU	MMVP	MMStar
Base	71.0	70.7	1685.0	696.0	84.6	52.8	69.3	65.0
Ours	74.7	78.4	1660.0	677.0	84.1	53.8	<u>70.0</u>	65.0
w.o. CJR	72.3	74.0	311.0	92.0	24.3	23.0	7.10	21.0
w.o. R	73.8	75.9	1632.0	655.0	84.2	<u>53.3</u>	71.3	65.0
w.o. J	74.2	74.8	1542.0	478.0	82.6	51.9	65.3	61.0
w.o. C	<u>74.5</u>	<u>76.4</u>	1653.0	701.0	<u>84.3</u>	50.9	68.0	64.0
w.o. G	73.4	72.3	<u>1669.0</u>	685.0	84.2	53.0	<u>70.0</u>	64.0

Table 2: **Ablation study on data composition.** Each variant is trained independently by removing exactly one data type from the full GCJR setting while keeping all other components fixed. S and R denote Short Score and Real Score, respectively.

Ultimately, these results confirm a reciprocal reinforcement: generative trajectories reconstructed as interpretive signals refine semantic boundaries, which in turn guide higher-fidelity synthesis, allowing **UniCorn** to significantly improve generation while preserving its core multimodal intelligence.

4.3.2 Model Architecture

We first demonstrate the effectiveness of our method on BAGEL, where understanding and generation components are decoupled. To evaluate its generalization to tightly coupled architectures, we conduct a base model ablation on the purely autoregressive Janus-Pro-7b (Chen et al., 2025d). Table 4 shows that our method improves Janus by +2.7 on TIIF, +7.0 on WISE, and +4.7 on OneIG-EN, with the most pronounced gain on the knowledge-intensive WISE benchmark. This suggests that the proposed approach enhances knowledge expression by guiding generation through improved understanding, a mechanism that generalizes across different model architectures.

4.3.3 Scaling Law for UniCorn

Scaling laws guide architectural design and optimization (Kaplan et al., 2020; Chen et al., 2024e), but prior methods scale poorly due to reliance on external models or fixed prompts. In contrast, **UniCorn** achieves scalable self-improvement through unbounded self-sampling and efficient Cognitive Pattern Reconstruction. To explicitly assess scalability, we conduct scale-up experiments by varying the amount of self-generated data across {1k, 5k, 8k, 10k, 20k}.

As shown in Fig. 6, with only 1K training samples, our method already surpasses RecA (Xie et al., 2025b) on TIIF. As the data scale increases, the model’s generative performance continues to improve, with more pronounced gains on long-prompt generation; notably, with just **5k** samples, it outperforms IRG (Huang, 2025) trained on **30k**

Model	Hard score
Bagel	36.6
Show-o2	36.1
Janus-Pro	9.9
UniCorn*	40.0
UniCorn	46.5

Table 3: **Hard score results on UniCycle.** Soft score results are reported in Appendix 17.

Model	TIIF	WISE	OneIG-EN
Janus Pro	63.2	35.0	26.7
+UniCorn	65.9+2.7	42.0+7.0	31.4+4.7
UniCorn	73.8	55.0	42.6
UniCorn*	74.4+0.6	54.0-1.0	44.9+2.3

Table 4: **Ablation studies on the base model (top) and the self-play framework (bottom).** Unicorn* denotes the BAGEL model trained on data constructed using Qwen3-VL-235B-A22B-Instruct.




Figure 6: **Data scaling result on TIIF.** The score consistently improves when the dataset size scales up. Notably, **UniCorn** surpasses many powerful models only using 5k training data.

GPT-4o distilled data as well as the strong closed-source model DALL-E 3 (Betker et al., 2023a). These results reveal a favorable scaling regime in which self-generated data alone suffices to drive continual and efficient improvements in generative capability.

4.4 Analysis

We design a series of experiments for in-depth analysis to address the following two questions.

Q1: Is self-play necessary?

For self-play assessment, we use Qwen3-VL-235B-A22B-Instruct (Yang et al., 2025a) for data construction (UniCorn*). As shown in Table 4, employing stronger proposers/judges yields diminishing returns, where high costs and training time outweigh performance gains. This likely stems from the difficulty of fitting high-entropy teacher distributions, increasing latency without proportional

information gain. We then compare **UniCorn** with four unified models on **UniCycle** (Judge: Qwen3-235B-A22B (Yang et al., 2025a)). **UniCycle** requires both generation and understanding, reducing task bias and evaluating the model’s self-reflection, thus signaling comprehensive multimodal intelligence.

As shown in Table 3, **UniCorn** achieves the highest Hard score (46.5), outperforming its base BAGEL by nearly 10 points and others by over 3 points. UniCorn* lags by 6.5 points, suggesting strong external supervision yields disproportionate costs and insufficient unified coordination. This demonstrates that self-play enhances unified capabilities by distilling understanding into generation without degradation, achieving SOTA on **UniCycle**. In contrast, Janus-Pro significantly underperforms comparable-scale models on **UniCycle**, revealing a gap between its generation and self-understanding.

Q2: Why UniCorn works?

UniCorn addresses the asymmetry in Unified Multimodal Models where strong understanding capabilities remain inactive during generation. We identify three critical limitations: (1) the model lacks a holistic perception of the content it is about to generate, (2) it does not actively assess the quality of its own outputs during generation, and (3) it lacks the ability to reflect on and revise suboptimal generations. UniCorn resolves this by enabling understanding to supervise generation through a unified self-improvement loop involving captioning, evaluation, and reflection, restoring alignment for more faithful and controllable results.

Theoretically, we justify this approach using Mutual Information and Bayes’ theorem, demonstrating that our task decomposition effectively minimizes Negative Log Likelihood (NLL). This guarantees that the auxiliary understanding tasks mathematically optimize the final unified objective. Detailed derivation is presented in Appendix D.

5 Conclusion

In this paper, we propose *UniCorn*, a self-supervised post-training framework that unifies multimodal comprehension and generation within a single model via multi-agent self-play and Cognitive Pattern Reconstruction, distilling internal latent knowledge into high-quality generative signals without external supervision. Extensive experiments, including the UniCycle cycle-consistency benchmark, demonstrate significant improvements

in T2I generation while preserving multimodal intelligence, highlighting self-contained feedback loops as a scalable path for unified multimodal models.

Limitations

Despite achieving robust performance in both T2I generation and multimodal understanding, UniCorn possesses certain limitations. First, the current self-improvement framework operates in a single-turn manner and primarily enhances generative capabilities, with no significant gains observed in understanding metrics. In future work, we intend to explore multi-turn iterative self-play to foster the co-evolution of both capabilities. Second, the self-play mechanism requires the UMM to handle prompt generation, rollout, and judgment, which inevitably introduces additional computational costs. We plan to investigate more efficient methodologies to streamline this process in subsequent research.

Ethical Statement

The development of UniCorn adheres to ethical standards for AI research. We utilize publicly available open-source models as our foundation and conduct all experiments using standard public benchmarks. Our self-improvement framework aims to enhance generative quality through internal feedback, thereby reducing the need for massive external data collection. While we implement internal filters during the self-play process to improve output alignment, we acknowledge that multimodal models may still reflect biases present in their pre-training data. We are committed to transparency and encourage the responsible use of our framework in downstream applications.

Acknowledgments

This work was supported by the Anhui Provincial Natural Science Foundation under Grant 2108085UD12. We acknowledge the support of GPU cluster built by MCC Lab of Information Science and Technology Institution, USTC. The AI-driven experiments, simulations and model training were performed on the robotic AI-Scientist platform of Chinese Academy of Sciences.

References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, and 1 others. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Ruichuan An, Sihan Yang, Ziyu Guo, Wei Dai, Zijun Shen, Haodong Li, Renrui Zhang, Xinyu Wei, Guopeng Li, Wenshan Wu, and 1 others. 2026. Genius: Generative fluid intelligence evaluation suite. *arXiv preprint arXiv:2602.11144*.
- Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, and 1 others. 2025. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*.
- James Betker, Gabriel Goh, Li Jing, Tim Brooks, Jianfeng Wang, Linjie Li, Long Ouyang, Juntang Zhuang, Joyce Lee, Yufei Guo, and 1 others. 2023a. Improving image generation with better captions. *Computer Science*. <https://cdn.openai.com/papers/dall-e-3.pdf>, 2(3):8.
- James Betker, Gabriel Goh, Li Jing, Tim Brooks, Jianfeng Wang, Linjie Li, Long Ouyang, Juntang Zhuang, Joyce Lee, Yufei Guo, and 1 others. 2023b. Improving image generation with better captions. *OpenAI blog*.
- Ignacio Matte Blanco. 2018. *The unconscious as infinite sets: An essay in bi-logic*. Routledge.
- Minwoo Byeon, Beomhee Park, Haecheon Kim, Sungjun Lee, Woonhyuk Baek, and Saehoon Kim. 2022. Coyo-700m: Image-text pair dataset.
- Huanqia Cai, Sihan Cao, Ruoyi Du, Peng Gao, Steven Hoi, Shijie Huang, Zhaohui Hou, Dengyang Jiang, Xin Jin, Liangchen Li, and 1 others. 2025. Z-image: An efficient image generation foundation model with single-stream diffusion transformer. *arXiv preprint arXiv:2511.22699*.
- Mert Cemri, Melissa Z Pan, Shuyi Yang, Lakshya A Agrawal, Bhavya Chopra, Rishabh Tiwari, Kurt Keutzer, Aditya Parameswaran, Dan Klein, Kanann Ramchandran, and 1 others. 2025. Why do multi-agent llm systems fail? *arXiv preprint arXiv:2503.13657*.
- Jingjing Chang, Yixiao Fang, Peng Xing, Shuhan Wu, Wei Cheng, Rui Wang, Xianfang Zeng, Gang Yu, and Hai-Bao Chen. 2025. **Oneig-bench: Omni-dimensional nuanced evaluation for image generation**. *Preprint*, arXiv:2506.07977.
- Boyuan Chen, Diego Martí Monsó, Yilun Du, Max Simchowitz, Russ Tedrake, and Vincent Sitzmann. 2024a. Diffusion forcing: Next-token prediction meets full-sequence diffusion. In *NeurIPS*.
- Chieh-Yun Chen, Min Shi, Gong Zhang, and Humphrey Shi. 2025a. T2i-copilot: A training-free multi-agent text-to-image system for enhanced prompt interpretation and interactive generation. In *International Conference on Computer Vision*, pages 19396–19405.
- Jiuhai Chen, Zhiyang Xu, Xichen Pan, Yushi Hu, Can Qin, Tom Goldstein, Lifu Huang, Tianyi Zhou, Saining Xie, Silvio Savarese, and 1 others. 2025b. Blip3-o: A family of fully open unified multimodal models-architecture, training and dataset. *arXiv preprint arXiv:2505.09568*.
- Jiuhai Chen, Zhiyang Xu, Xichen Pan, Yushi Hu, Can Qin, Tom Goldstein, Lifu Huang, Tianyi Zhou, Saining Xie, Silvio Savarese, and 1 others. 2025c. Blip3-o: A family of fully open unified multimodal models-architecture, training and dataset. *arXiv preprint arXiv:2505.09568*.
- Junsong Chen, Chongjian Ge, Enze Xie, Yue Wu, Lewei Yao, Xiaozhe Ren, Zhongdao Wang, Ping Luo, Huchuan Lu, and Zhenguo Li. 2024b. Pixart-sigma: Weak-to-strong training of diffusion transformer for 4k text-to-image generation. In *European Conference on Computer Vision*.
- Lin Chen, Jinsong Li, Xiaoyi Dong, Pan Zhang, Yuhang Zang, Zehui Chen, Haodong Duan, Jiaqi Wang, Yu Qiao, Dahua Lin, and 1 others. 2024c. Are we on the right way for evaluating large vision-language models? *Advances in Neural Information Processing Systems*, 37:27056–27087.
- Weize Chen, Yusheng Su, Jingwei Zuo, Cheng Yang, Chenfei Yuan, Chi-Min Chan, Heyang Yu, Yaxi Lu, Yi-Hsin Hung, Chen Qian, and 1 others. 2024d. Agentverse: Facilitating multi-agent collaboration and exploring emergent behaviors. In *International Conference on Learning Representations*.
- Xiaokang Chen, Chengyue Wu, Zhiyu Wu, Yiyang Ma, Xingchao Liu, Zizheng Pan, Wen Liu, Zhenda Xie, Xingkai Yu, Chong Ruan, and Ping Luo. 2025d. Janus-pro: Unified multimodal understanding and generation with data and model scaling. *arXiv preprint arXiv:2501.17811*.
- Xiaokang Chen, Zhiyu Wu, Xingchao Liu, Zizheng Pan, Wen Liu, Zhenda Xie, Xingkai Yu, and Chong Ruan. 2025e. Janus-pro: Unified multimodal understanding and generation with data and model scaling. *arXiv preprint arXiv:2501.17811*.
- Yixing Chen, Yiding Wang, Siqi Zhu, Haofei Yu, Tao Feng, Muhan Zhang, Mostofa Patwary, and Jiaxuan You. 2025f. Multi-agent evolve: Llm self-improve through co-evolution. *arXiv preprint arXiv:2510.23595*.
- Zehui Chen, Kuikun Liu, Qiuchen Wang, Wenwei Zhang, Jiangning Liu, Dahua Lin, Kai Chen, and Feng Zhao. 2024e. Agent-flan: Designing data and methods of effective agent tuning for large language models. *Annual Meeting of the Association for Computational Linguistics*.

- Zihao Cheng, Zeming Liu, Yingyu Shan, Xinyi Wang, Xiangrong Zhu, Yunpu Ma, Hongru Wang, Yuhang Guo, Wei Lin, and Yunhong Wang. 2026. [Mem²evolve: Towards self-evolving agents via co-evolutionary capability expansion and experience distillation](#). *Preprint*, arXiv:2604.10923.
- Eugene Choi, Arash Ahmadian, Matthieu Geist, Olivier Pietquin, and Mohammad Gheshlaghi Azar. 2024. Self-improving robust preference optimization. *arXiv preprint arXiv:2406.01660*.
- Yufeng Cui, Honghao Chen, Haoge Deng, Xu Huang, Xinghang Li, Jirong Liu, Yang Liu, Zhuoyan Luo, Jinsheng Wang, Wenxuan Wang, and 1 others. 2025. Emu3. 5: Native multimodal models are world learners. *arXiv preprint arXiv:2510.26583*.
- Chaorui Deng, Deyao Zhu, Kunchang Li, Chenhui Gou, Feng Li, Zeyu Wang, Shu Zhong, Weihao Yu, Xiaonan Nie, Ziang Song, Guang Shi, and Haoqi Fan. 2025a. Emerging properties in unified multimodal pretraining. *arXiv preprint arXiv:2505.14683*.
- Chaorui Deng, Deyao Zhu, Kunchang Li, Chenhui Gou, Feng Li, Zeyu Wang, Shu Zhong, Weihao Yu, Xiaonan Nie, Ziang Song, and 1 others. 2025b. Emerging properties in unified multimodal pretraining. *arXiv preprint arXiv:2505.14683*.
- Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Jingyuan Ma, Rui Li, Heming Xia, Jingjing Xu, Zhiyong Wu, Baobao Chang, and 1 others. 2024. A survey on in-context learning. In *Proceedings of the 2024 conference on empirical methods in natural language processing*, pages 1107–1128.
- John Dunlosky and Janet Metcalfe. 2008. *Metacognition*. Sage Publications.
- Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, and 1 others. 2024. Scaling rectified flow transformers for high-resolution image synthesis. In *ICML*.
- Zhen Fang, Zhuoyang Liu, Jiaming Liu, Hao Chen, Yu Zeng, Shiting Huang, Zehui Chen, Lin Chen, Shanghang Zhang, and Feng Zhao. 2025. Dualvla: Building a generalizable embodied agent via partial decoupling of reasoning and action. *arXiv preprint arXiv:2511.22134*.
- Yuanning Feng, Sinan Wang, Zhengxiang Cheng, Yao Wan, and Dongping Chen. 2025. Are we on the right way to assessing llm-as-a-judge? *arXiv preprint arXiv:2512.16041*.
- Chaoyou Fu, Peixian Chen, Yunhang Shen, Yulei Qin, Mengdan Zhang, Xu Lin, Jinrui Yang, Xiawu Zheng, Ke Li, Xing Sun, and 1 others. 2023. Mme: A comprehensive evaluation benchmark for multimodal large language models. *arXiv preprint arXiv:2306.13394*.
- Yuying Ge, Sijie Zhao, Jinguo Zhu, Yixiao Ge, Kun Yi, Lin Song, Chen Li, Xiaohan Ding, and Ying Shan. 2024. Seed-x: Multimodal models with unified multi-granularity comprehension and generation. *arXiv preprint arXiv:2404.14396*.
- Dhruba Ghosh, Hannaneh Hajishirzi, and Ludwig Schmidt. 2023. Geneval: An object-focused framework for evaluating text-to-image alignment. In *NeurIPS*.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, and 1 others. 2025a. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, and 1 others. 2025b. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- Ziyu Guo, Renrui Zhang, Hongyu Li, Manyuan Zhang, Xinyan Chen, Sifan Wang, Yan Feng, Peng Pei, and Pheng-Ann Heng. 2025c. Thinking-while-generating: Interleaving textual reasoning throughout visual generation. *arXiv preprint arXiv:2511.16671*.
- Xiwei Hu, Rui Wang, Yixiao Fang, Bin Fu, Pei Cheng, and Gang Yu. 2024. Ella: Equip diffusion models with llm for enhanced semantic alignment. *arXiv preprint arXiv:2403.05135*.
- Yushi Hu, Reyhane Askari-Hemmat, Melissa Hall, Emily Dinan, Luke Zettlemoyer, and Marjan Ghazvininejad. 2025. Multimodal rewardbench 2: Evaluating omni reward models for interleaved text and image. *arXiv preprint arXiv:2512.16899*.
- Chengsong Huang, Wenhao Yu, Xiaoyang Wang, Hongming Zhang, Zongxia Li, Ruosen Li, Jiabin Huang, Haitao Mi, and Dong Yu. 2025a. R-zero: Self-evolving reasoning llm from zero data. *arXiv preprint arXiv:2508.05004*.
- Shiting Huang, Zhen Fang, Zehui Chen, Siyu Yuan, Junjie Ye, Yu Zeng, Lin Chen, Qi Mao, and Feng Zhao. 2025b. Critictool: Evaluating self-critique capabilities of large language models in tool-calling error scenarios. *arXiv preprint arXiv:2506.13977*.
- Shiting Huang, Zecheng Li, Yu Zeng, Qingnan Ren, Zhen Fang, Qisheng Su, Kou Shi, Lin Chen, Zehui Chen, and Feng Zhao. 2026a. Internalizing meta-experience into memory for guided reinforcement learning in large language models. *arXiv preprint arXiv:2602.10224*.
- Wenxuan Huang. 2025. [Interleaving reasoning: Next-generation reasoning systems for agi](#). GitHub repository. Accessed 2025-08-19.

- Wenxuan Huang, Shuang Chen, Zheyong Xie, Shaosheng Cao, Shixiang Tang, Yufan Shen, Qingyu Yin, Wenbo Hu, Xiaoman Wang, Yuntian Tang, and 1 others. 2025c. Interleaving reasoning for better text-to-image generation. *arXiv preprint arXiv:2509.06945*.
- Wenxuan Huang, Yu Zeng, Qiuchen Wang, Zhen Fang, Shaosheng Cao, Zheng Chu, Qingyu Yin, Shuang Chen, Zhenfei Yin, Lin Chen, and 1 others. 2026b. Vision-deepresearch: Incentivizing deepresearch capability in multimodal large language models. *arXiv preprint arXiv:2601.22060*.
- Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, and 1 others. 2024. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*.
- Dongzhi Jiang, Ziyu Guo, Renrui Zhang, Zhuofan Zong, Hao Li, Le Zhuo, Shilin Yan, Pheng-Ann Heng, and Hongsheng Li. 2025a. T2i-r1: Reinforcing image generation with collaborative semantic-level and token-level cot. *arXiv preprint arXiv:2505.00703*.
- Dongzhi Jiang, Renrui Zhang, Haodong Li, Zhuofan Zong, Ziyu Guo, Jun He, Claire Guo, Junyan Ye, Rongyao Fang, Weijia Li, and 1 others. 2025b. Draco: Draft as cot for text-to-image preview and rare concept generation. *arXiv preprint arXiv:2512.05112*.
- Weiyang Jin, Yuwei Niu, Jiaqi Liao, Chengqi Duan, Aoxue Li, Shenghua Gao, and Xihui Liu. 2025a. Srum: Fine-grained self-rewarding for unified multimodal models. *arXiv preprint arXiv:2510.12784*.
- Zhuoran Jin, Hongbang Yuan, Kejian Zhu, Jiachun Li, Pengfei Cao, Yubo Chen, Kang Liu, and Jun Zhao. 2025b. Omni-reward: Towards generalist omni-modal reward modeling with free-form preferences. *arXiv preprint arXiv:2510.23451*.
- Jared Kaplan, Sam McCandlish, Tom Henighan, Tom Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. 2020. Scaling laws for neural language models. In *ICML*.
- Jihyung Kil, Zheda Mai, Justin Lee, Zihe Wang, Kerrie Cheng, Lemeng Wang, Ye Liu, Arpita Chowdhury, and Wei-Lun Chao. 2024. Compbench: A comparative reasoning benchmark for multimodal llms.
- Seungone Kim, Jamin Shin, Yejin Cho, Joel Jang, Shayne Longpre, Hwaran Lee, Sangdoon Yun, Seongjin Shin, Sungdong Kim, James Thorne, and 1 others. 2023. Prometheus: Inducing fine-grained evaluation capability in language models. In *The Twelfth International Conference on Learning Representations*.
- Black Forest Labs. 2024. [Flux](#).
- Daiqing Li, Aleks Kamko, Ehsan Akhgari, Ali Sabet, Linmiao Xu, and Suhail Doshi. 2024a. Playground v2. 5: Three insights towards enhancing aesthetic quality in text-to-image generation. *arXiv preprint arXiv:2402.17245*.
- Dawei Li, Bohan Jiang, Liangjie Huang, Alimohammad Beigi, Chengshuai Zhao, Zhen Tan, Amrita Bhat-tacharjee, Yuxuan Jiang, Canyu Chen, Tianhao Wu, and 1 others. 2025a. From generation to judgment: Opportunities and challenges of llm-as-a-judge. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 2757–2791.
- Haitao Li, Qian Dong, Junjie Chen, Huixue Su, Yujia Zhou, Qingyao Ai, Ziyi Ye, and Yiquan Liu. 2024b. Llms-as-judges: a comprehensive survey on llm-based evaluation methods. *arXiv preprint arXiv:2412.05579*.
- Haodong Li, Chunmei Qing, Huanyu Zhang, Dongzhi Jiang, Yihang Zou, Hongbo Peng, Dingming Li, Yuhong Dai, ZePeng Lin, Juanxi Tian, and 1 others. 2026a. Coco: Code as cot for text-to-image preview and rare concept generation. *arXiv preprint arXiv:2603.08652*.
- Haodong Li, Jingwei Wu, Quan Sun, Guopeng Li, Juanxi Tian, Huanyu Zhang, Yanlin Lai, Ruichuan An, Hongbo Peng, Yuhong Dai, and 1 others. 2026b. Gebench: Benchmarking image generation models as gui environments. *arXiv preprint arXiv:2602.09007*.
- Mengdi Li, Jiaye Lin, Xufeng Zhao, Wenhao Lu, Peilin Zhao, Stefan Wermtter, and Di Wang. 2025b. Curriculum-rlaif: Curriculum alignment with reinforcement learning from ai feedback. *arXiv preprint arXiv:2505.20075*.
- Songze Li, Xiaoke Guo, Tianqi Liu, Biao Yi, Zhaoyan Gong, Zhiqiang Liu, Huajun Chen, and Wen Zhang. 2026c. [What’s missing in screen-to-action? towards a ui-in-the-loop paradigm for multimodal gui reasoning](#). *Preprint*, arXiv:2604.06995.
- Xiaoyuan Li, Keqin Bao, Yubo Ma, Moxin Li, Wenjie Wang, Rui Men, Yichang Zhang, Fuli Feng, Dayiheng Liu, and Junyang Lin. 2025c. Mtr-bench: A comprehensive benchmark for multi-turn reasoning evaluation. *arXiv preprint arXiv:2505.17123*.
- Zhimin Li, Jianwei Zhang, Qin Lin, Jiangfeng Xiong, Yanxin Long, Xincheng Deng, Yingfang Zhang, Xingchao Liu, Minbin Huang, Zedong Xiao, Dayou Chen, Jiajun He, Jiahao Li, Wenyue Li, Chen Zhang, Rongwei Quan, Jianxiang Lu, Jiabin Huang, Xiaoyan Yuan, and 26 others. 2024c. Hunyuan-dit: A powerful multi-resolution diffusion transformer with fine-grained chinese understanding.
- Tian Liang, Zhiwei He, Wenxiang Jiao, Xing Wang, Yan Wang, Rui Wang, Yujiu Yang, Shuming Shi, and Zhaopeng Tu. 2024a. Encouraging divergent thinking in large language models through multi-agent debate. In *Proceedings of the 2024 conference on*

- empirical methods in natural language processing*, pages 17889–17904.
- Weixin Liang, Lili Yu, Liang Luo, Srinivasan Iyer, Ning Dong, Chunting Zhou, Gargi Ghosh, Mike Lewis, Wen-tau Yih, Luke Zettlemoyer, and 1 others. 2024b. Mixture-of-transformers: A sparse and scalable architecture for multi-modal foundation models. *arXiv preprint arXiv:2411.04996*.
- Honglin Lin, Zheng Liu, Yun Zhu, Chonghan Qin, Juekai Lin, Xiaoran Shang, Conghui He, Wentao Zhang, and Lijun Wu. 2026a. Mmfineason: Closing the multimodal reasoning gap via open data-centric methods. *arXiv preprint arXiv:2601.21821*.
- Honglin Lin, Chonghan Qin, Zheng Liu, Qizhi Pei, Yu Li, Zhanping Zhong, Xin Gao, Yanfeng Wang, Conghui He, and Lijun Wu. 2026b. Scientific image synthesis: Benchmarking, methodologies, and downstream utility. *arXiv preprint arXiv:2601.17027*.
- Jiaye Lin, Yifu Guo, Yuzhen Han, Sen Hu, Ziyi Ni, Licheng Wang, Mingguang Chen, Hongzhang Liu, Ronghao Chen, Yangfan He, and 1 others. 2025. Se-agent: Self-evolution trajectory optimization in multi-step reasoning with llm-based agents. *arXiv preprint arXiv:2508.02085*.
- Hao Liu, Wilson Yan, Matei Zaharia, and Pieter Abbeel. 2024a. World model on million-length video and language with ringattention. *arXiv preprint arXiv:2402.08268*.
- Yuan Liu, Haodong Duan, Yuanhan Zhang, Bo Li, Songyang Zhang, Wangbo Zhao, Yike Yuan, Jiaqi Wang, Conghui He, Ziwei Liu, and 1 others. 2024b. Mmbench: Is your multi-modal model an all-around player? In *European Conference on Computer Vision*.
- Zheng Liu, Hao Liang, Xijie Huang, Wentao Xiong, Qinhan Yu, Linzhuang Sun, Chong Chen, Conghui He, Bin Cui, and Wentao Zhang. 2024c. Synthlm: High-efficiency and high-quality synthetic data for vision language models. *arXiv preprint arXiv:2407.20756*, 3.
- Zheng Liu, Honglin Lin, Chonghan Qin, Xiaoyang Wang, Xin Gao, Yu Li, Mengzhang Cai, Yun Zhu, Zhanping Zhong, Qizhi Pei, and 1 others. 2026. Chartverse: Scaling chart reasoning via reliable programmatic synthesis from scratch. *arXiv preprint arXiv:2601.13606*.
- Zheng Liu, Mengjie Liu, Jingzhou Chen, Jingwei Xu, Bin Cui, Conghui He, and Wentao Zhang. 2025a. Fusion: Fully integration of vision-language representations for deep cross-modal understanding. *arXiv preprint arXiv:2504.09925*.
- Zheng Liu, Mengjie Liu, Siwei Wen, Mengzhang Cai, Bin Cui, Conghui He, and Wentao Zhang. 2025b. From uniform to heterogeneous: Tailoring policy optimization to every token’s nature. *arXiv preprint arXiv:2509.16591*.
- Zijun Liu, Peiyi Wang, Runxin Xu, Shirong Ma, Chong Ruan, Peng Li, Yang Liu, and Yu Wu. 2025c. Inference-time scaling for generalist reward modeling. *arXiv preprint arXiv:2504.02495*.
- Hongxu Ma, Kai Tian, Tao Zhang, Xuefeng Zhang, Han Zhou, Chunjie Chen, Han Li, Jihong Guan, and Shuigeng Zhou. 2024a. Generative regression based watch time prediction for short-video recommendation. *arXiv preprint arXiv:2412.20211*.
- Hongxu Ma, Han Zhou, Kai Tian, Xuefeng Zhang, Chunjie Chen, Han Li, Jihong Guan, and Shuigeng Zhou. 2026a. Gor: A unified and extensible generative framework for ordinal regression. In *The Fourteenth International Conference on Learning Representations*.
- Kexin Ma, Ruochun Jin, Wang Haotian, Wang Xi, Huan Chen, Yuhua Tang, and Qian Wang. 2024b. Context-driven index trimming: A data quality perspective to enhancing precision of ralms. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 4886–4901.
- Kexin Ma, Bojun Li, Yuhua Tang, Liting Sun, and Ruochun Jin. 2026b. Cast: Character-and-scene episodic memory for agents. *Preprint, arXiv:2602.06051*.
- Xueqi Ma, Shuo Yang, Yanbei Jiang, Shu Liu, Zhenzhen Liu, Jiayang Ao, Xingjun Ma, Sarah Monazam Erfani, and James Bailey. 2026c. Attention in space: Functional roles of vlm heads for spatial reasoning. *arXiv preprint arXiv:2603.20662*.
- Weijia Mao, Zhenheng Yang, and Mike Zheng Shou. 2025. Unirl: Self-improving unified multimodal models via supervised and reinforcement learning. *arXiv preprint arXiv:2505.23380*.
- Midjourney. 2025. [midjourney v7](#).
- Sumeet Ramesh Motwani, Chandler Smith, Rocktim Jyoti Das, Rafael Rafailov, Ivan Laptev, Philip HS Torr, Fabio Pizzati, Ronald Clark, and Christian Schroeder de Witt. 2024. Malt: Improving reasoning with multi-agent llm training. *arXiv preprint arXiv:2412.01928*.
- Yuwei Niu, Munan Ning, Mengren Zheng, Bin Lin, Peng Jin, Jiaqi Liao, Kunpeng Ning, Bin Zhu, and Li Yuan. 2025. Wise: A world knowledge-informed semantic evaluation for text-to-image generation. *arXiv preprint arXiv:2503.07265*.
- OpenAI. 2025. [Introducing 4o image generation](#).
- Xichen Pan, Satya Narayan Shukla, Aashu Singh, Zhuokai Zhao, Shlok Kumar Mishra, Jialiang Wang, Zhiyang Xu, Jiuhai Chen, Kunpeng Li, Felix Juefei-Xu, Ji Hou, and Saining Xie. 2025. Transfer between modalities with metaqueries. *arXiv preprint arXiv:2504.06256*.

- Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. 2024. Sdxl: Improving latent diffusion models for high-resolution image synthesis. In *International Conference on Learning Representations*.
- Yukun Qi, Yiming Zhao, Yu Zeng, Xikun Bao, Wenxuan Huang, Lin Chen, Zehui Chen, Jie Zhao, Zhongang Qi, and Feng Zhao. 2025. Vcr-bench: A comprehensive evaluation framework for video chain-of-thought reasoning. *arXiv preprint arXiv:2504.07956*.
- Qi Qin, Le Zhuo, Yi Xin, Ruoyi Du, Zhen Li, Bin Fu, Yiting Lu, Jiakang Yuan, Xinyue Li, Dongyang Liu, and 1 others. 2025. Lumina-image 2.0: A unified and efficient image generative framework. *arXiv preprint arXiv:2503.21758*.
- Liao Qu, Huichao Zhang, Yiheng Liu, Xu Wang, Yi Jiang, Yiming Gao, Hu Ye, Daniel K Du, Zehuan Yuan, and Xinglong Wu. 2024. Tokenflow: Unified image tokenizer for multimodal understanding and generation. *arXiv preprint arXiv:2412.03069*.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, and 1 others. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Vyas Raina, Adian Liusie, and Mark Gales. 2024. Is llm-as-a-judge robust? investigating universal adversarial attacks on zero-shot llm assessment. *arXiv preprint arXiv:2402.14016*.
- Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. 2022. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*.
- Qingnan Ren, Shiting Huang, Zhen Fang, Zehui Chen, Lin Chen, Lijun Li, and Feng Zhao. 2026. Adora: Training reasoning models with dynamic advantage estimation on reinforcement learning. *arXiv preprint arXiv:2602.10019*.
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-resolution image synthesis with latent diffusion models. In *CVPR*, pages 10684–10695.
- Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, and 1 others. 2022. Laion-5b: An open large-scale dataset for training next generation image-text models. *Advances in neural information processing systems*, 35:25278–25294.
- Weijia Shi, Xiaochuang Han, Chunting Zhou, Weixin Liang, Xi Victoria Lin, Luke Zettlemoyer, and Lili Yu. 2024. Lmfusion: Adapting pretrained language models for multimodal generation. *arXiv preprint arXiv:2412.15188*.
- Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. 2023. Reflexion: Language agents with verbal reinforcement learning. *Advances in Neural Information Processing Systems*, 36:8634–8652.
- David Silver, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, Laurent Sifre, Dhharshan Kumaran, Thore Graepel, and 1 others. 2017. Mastering chess and shogi by self-play with a general reinforcement learning algorithm. *arXiv preprint arXiv:1712.01815*.
- Chameleon Team. 2024. Chameleon: Mixed-modal early-fusion foundation models. *arXiv preprint arXiv:2405.09818*.
- Aman Singh Thakur, Kartik Choudhary, Venkat Srinik Ramayapally, Sankaran Vaidyanathan, and Dieuwke Hupkes. 2025. Judging the judges: Evaluating alignment and vulnerabilities in llms-as-judges. In *Proceedings of the Fourth Workshop on Generation, Evaluation and Metrics (GEM²)*, pages 404–430.
- Rui Tian, Mingfei Gao, Mingze Xu, Jiaming Hu, Jiasen Lu, Zuxuan Wu, Yinfei Yang, and Afshin Dehghan. 2025. Unigen: Enhanced training & test-time strategies for unified multimodal understanding and generation. *arXiv preprint arXiv:2505.14682*.
- Shengbang Tong, David Fan, Jiachen Li, Yunyang Xiong, Xinlei Chen, Koustuv Sinha, Michael Rabbat, Yann LeCun, Saining Xie, and Zhuang Liu. 2025. Metamorph: Multimodal understanding and generation via instruction tuning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 17001–17012.
- Shengbang Tong, Zhuang Liu, Yuexiang Zhai, Yi Ma, Yann LeCun, and Saining Xie. 2024. Eyes wide shut? exploring the visual shortcomings of multimodal llms. In *CVPR*, pages 9568–9578.
- Chunwei Wang, Guansong Lu, Junwei Yang, Runhui Huang, Jianhua Han, Lu Hou, Wei Zhang, and Hang Xu. 2024a. Illume: Illuminating your llms to see, draw, and self-enhance. *arXiv preprint arXiv:2412.06673*.
- Eileen Wang, Hiba Arnaout, Dhita Pratama, Shuo Yang, Danyang Liu, Jie Yang, Josiah Poon, Jeff Pan, and Caren Han. 2026a. Mmcomet: A large-scale multimodal commonsense knowledge graph for contextual reasoning. *arXiv preprint arXiv:2603.01055*.
- Junxi Wang, Te Sun, Jiayi Zhu, Junxian Li, Haowen Xu, Zichen Wen, Xuming Hu, Zhiyu Li, and Linfeng Zhang. 2026b. Streammeco: Long-term agent memory compression for efficient streaming video understanding. *Preprint*, arXiv:2604.09000.
- Junxi Wang, Yaxiong Wang, Lechao Cheng, and Zhun Zhong. 2025a. Fakesv-vlm: Taming vlm for detecting fake short-video news via progressive mixture-of-experts adapter. *Preprint*, arXiv:2508.19639.

- Qiuchen Wang, Ruixue Ding, Yu Zeng, Zehui Chen, Lin Chen, Shihang Wang, Pengjun Xie, Fei Huang, and Feng Zhao. 2025b. Vrag-rl: Empower vision-perception-based rag for visually rich information understanding via iterative reasoning with reinforcement learning. *arXiv preprint arXiv:2505.22019*.
- Xinlong Wang, Xiaosong Zhang, Zhengxiong Luo, Quan Sun, Yufeng Cui, Jinsheng Wang, Fan Zhang, Yueze Wang, Zhen Li, Qiyang Yu, and 1 others. 2024b. Emu3: Next-token prediction is all you need. *arXiv preprint arXiv:2409.18869*.
- Yibin Wang, Yuhang Zang, Hao Li, Cheng Jin, and Jiaqi Wang. 2025c. Unified reward model for multimodal understanding and generation. *arXiv preprint arXiv:2503.05236*.
- Zi-Han Wang, Lam Nguyen, Zhengyang Zhao, Mengyue Yang, Chengwei Qin, Yujiu Yang, and Linyi Yang. 2026c. Creativebench: Benchmarking and enhancing machine creativity via self-evolving challenges. *arXiv preprint arXiv:2603.11863*.
- Xinyu Wei, Jinrui Zhang, Zeqing Wang, Hongyang Wei, Zhen Guo, and Lei Zhang. 2025. [Tiif-bench: How does your t2i model follow your instructions?](#) *Preprint*, arXiv:2506.02161.
- Chenfei Wu, Jiahao Li, Jingren Zhou, Junyang Lin, Kaiyuan Gao, Kun Yan, Sheng-ming Yin, Shuai Bai, Xiao Xu, Yilei Chen, and 1 others. 2025a. Qwen-image technical report. *arXiv preprint arXiv:2508.02324*.
- Chengyue Wu, Xiaokang Chen, Zhiyu Wu, Yiyang Ma, Xingchao Liu, Zizheng Pan, Wen Liu, Zhenda Xie, Xingkai Yu, Chong Ruan, and 1 others. 2025b. Janus: Decoupling visual encoding for unified multimodal understanding and generation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 12966–12977.
- Shitao Xiao, Yueze Wang, Junjie Zhou, Huaying Yuan, Xingrun Xing, Ruiran Yan, Chaofan Li, Shuting Wang, Tiejun Huang, and Zheng Liu. 2025a. Omnigen: Unified image generation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 13294–13304.
- Xi Xiao, Chenrui Ma, Yunbei Zhang, Chen Liu, Zhuxuanzi Wang, Yanshu Li, Lin Zhao, Guosheng Hu, Tianyang Wang, and Hao Xu. 2026. Not all directions matter: Toward structured and task-aware low-rank adaptation. *arXiv preprint arXiv:2603.14228*.
- Xi Xiao, Yunbei Zhang, Xingjian Li, Tianyang Wang, Xiao Wang, Yuxiang Wei, Jihun Hamm, and Min Xu. 2025b. Visual instance-aware prompt tuning. In *33rd ACM International Conference on Multimedia, MM 2025*, pages 2880–2889. Association for Computing Machinery, Inc.
- Xi Xiao, Yunbei Zhang, Lin Zhao, Yiyang Liu, Xiaoying Liao, Zhenda Mai, Xingjian Li, Xiao Wang, Hao Xu, Jihun Hamm, and 1 others. 2025c. Prompt-based adaptation in large-scale vision models: A survey. *arXiv preprint arXiv:2510.13219*.
- Enze Xie, Junsong Chen, Yuyang Zhao, Jincheng Yu, Ligeng Zhu, Chengyue Wu, Yujun Lin, Zhekai Zhang, Muyang Li, Junyu Chen, and 1 others. 2025a. Sana 1.5: Efficient scaling of training-time and inference-time compute in linear diffusion transformer. *arXiv preprint arXiv:2501.18427*.
- Ji Xie, Trevor Darrell, Luke Zettlemoyer, and XuDong Wang. 2025b. Reconstruction alignment improves unified multimodal models. *arXiv preprint arXiv:2509.07295*.
- Jinheng Xie, Weijia Mao, Zechen Bai, David Junhao Zhang, Weihao Wang, Kevin Qinghong Lin, Yuchao Gu, Zhijie Chen, Zhenheng Yang, and Mike Zheng Shou. 2024. Show-o: One single transformer to unify multimodal understanding and generation. *arXiv preprint arXiv:2408.12528*.
- Jinheng Xie, Zhenheng Yang, and Mike Zheng Shou. 2025c. Show-o2: Improved native unified multimodal models. *arXiv preprint arXiv:2506.15564*.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, and 1 others. 2025a. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.
- Shuo Yang, Caren Han, Yihao Ding, Shuhe Wang, and Eduard Hovy. 2026. Tooltree: Efficient llm tool planning via dual-feedback monte carlo tree search and bidirectional pruning. In *The Fourteenth International Conference on Learning Representations*, volume 3.
- Shuo Yang, Caren Han, Siwen Luo, and Eduard Hovy. 2025b. Magic-vqa: Multimodal and grounded inference with commonsense knowledge for visual question answering. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 16967–16986.
- Wenhao Yu, Zhenwen Liang, Chengsong Huang, Kishan Panaganti, Tianqing Fang, Haitao Mi, and Dong Yu. 2025. Guided self-evolving llms with minimal human supervision. *arXiv preprint arXiv:2512.02472*.
- Weizhe Yuan, Richard Yuanzhe Pang, Kyunghyun Cho, Xian Li, Sainbayar Sukhbaatar, Jing Xu, and Jason E Weston. 2024. Self-rewarding language models. In *Forty-first International Conference on Machine Learning*.
- Xiang Yue, Yuansheng Ni, Kai Zhang, Tianyu Zheng, Ruoqi Liu, Ge Zhang, Samuel Stevens, Dongfu Jiang, Weiming Ren, Yuxuan Sun, and 1 others. 2024. Mmmu: A massive multi-discipline multimodal understanding and reasoning benchmark for expert agi. In *CVPR*.

- Yu Zeng, Wenxuan Huang, Zhen Fang, Shuang Chen, Yufan Shen, Yishuo Cai, Xiaoman Wang, Zhenfei Yin, Lin Chen, Zehui Chen, and 1 others. 2026. Vision-deepresearch benchmark: Rethinking visual and textual search for multimodal large language models. *arXiv preprint arXiv:2602.02185*.
- Yu Zeng, Wenxuan Huang, Shiting Huang, Xikun Bao, Yukun Qi, Yiming Zhao, Qiuchen Wang, Lin Chen, Zehui Chen, Huaian Chen, and 1 others. 2025a. Agentic jigsaw interaction learning for enhancing visual perception and reasoning in vision-language models. *arXiv preprint arXiv:2510.01304*.
- Yu Zeng, Yukun Qi, Yiming Zhao, Xikun Bao, Lin Chen, Zehui Chen, Shiting Huang, Jie Zhao, and Feng Zhao. 2025b. Enhancing large vision-language models with ultra-detailed image caption generation. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 26703–26729.
- Huanyu Zhang, Xuehai Bai, Chengzu Li, Chen Liang, Haochen Tian, Haodong Li, Ruichuan An, Yifan Zhang, Anna Korhonen, Zhang Zhang, and 1 others. 2026a. How well do models follow visual instructions? vibe: A systematic benchmark for visual instruction-driven image editing. *arXiv preprint arXiv:2602.01851*.
- Qianchi Zhang, Hainan Zhang, Liang Pang, Yongxin Tong, Hongwei Zheng, and Zhiming Zheng. 2025a. Less is more: Compact clue selection for efficient retrieval-augmented generation reasoning. *arXiv preprint arXiv:2502.11811*.
- Qianchi Zhang, Hainan Zhang, Liang Pang, Hongwei Zheng, and Zhiming Zheng. 2026b. Stable-rag: Mitigating retrieval-permutation-induced hallucinations in retrieval-augmented generation. *arXiv preprint arXiv:2601.02993*.
- Wenyuan Zhang, Xinghua Zhang, Haiyang Yu, Shuaiyi Nie, Bingli Wu, Juwei Yue, Tingwen Liu, and Yongbin Li. 2026c. [Expseek: Self-triggered experience seeking for web agents](#). *Preprint*, arXiv:2601.08605.
- Xinglang Zhang, Yunyao Zhang, ZeLiang Chen, Junqing Yu, Wei Yang, and Zikai Song. 2026d. [Logical phase transitions: Understanding collapse in llm logical reasoning](#). *Preprint*, arXiv:2601.02902.
- Yunyao Zhang, Yihao Ai, Zuocheng Ying, Qirui Mi, Junqing Yu, Wei Yang, and Zikai Song. 2026e. [Coupling macro dynamics and micro states for long-horizon social simulation](#). *Preprint*, arXiv:2604.05516.
- Yunyao Zhang, Zikai Song, Hang Zhou, Wenfeng Ren, Yi-Ping Phoebe Chen, Junqing Yu, and Wei Yang. 2025b. [ga - s³: Comprehensive social network simulation with group agents](#). In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 8950–8970, Vienna, Austria. Association for Computational Linguistics.
- Yunyao Zhang, Xinglang Zhang, Junxi Sheng, Wenbing Li, Junqing Yu, Yi-Ping Phoebe Chen, Wei Yang, and Zikai Song. 2026f. [Semantic-aware logical reasoning via a semiotic framework](#). *Preprint*, arXiv:2509.24765.
- Andrew Zhao, Yiran Wu, Yang Yue, Tong Wu, Quentin Xu, Matthieu Lin, Shenzhi Wang, Qingyun Wu, Zilong Zheng, and Gao Huang. 2025a. Absolute zero: Reinforced self-play reasoning with zero data. *arXiv preprint arXiv:2505.03335*.
- Chuyang Zhao, Yuxing Song, Wenhao Wang, Haocheng Feng, Errui Ding, Yifan Sun, Xinyan Xiao, and Jingdong Wang. 2024. Monoformer: One transformer for both diffusion and autoregression. *arXiv preprint arXiv:2409.16280*.
- Pengxiang Zhao, Guangyi Liu, Yaozhen Liang, Weiqing He, Zhengxi Lu, Yuehao Huang, Yaxuan Guo, Kexin Zhang, Hao Wang, Liang Liu, and 1 others. 2025b. Mas-bench: A unified benchmark for shortcut-augmented hybrid mobile gui agents. *arXiv preprint arXiv:2509.06477*.
- Xiangyu Zhao, Peiyuan Zhang, Kexian Tang, Xiaorong Zhu, Hao Li, Wenhao Chai, Zicheng Zhang, Renqiu Xia, Guangtao Zhai, Junchi Yan, and 1 others. 2025c. Envisioning beyond the pixels: Benchmarking reasoning-informed visual editing. *arXiv preprint arXiv:2504.02826*.
- Yiming Zhao, Yu Zeng, Yukun Qi, YaoYang Liu, Lin Chen, Zehui Chen, Xikun Bao, Jie Zhao, and Feng Zhao. 2025d. V2p-bench: Evaluating video-language understanding with visual prompts for better human-model interaction. *arXiv preprint arXiv:2503.17736*.
- Wang Zhenyu, Xie Enze, Li Aoxue, Wang Zhongdao, Liu Xihui, and Li Zhenguo. 2024. Divide and conquer: Language models can plan and self-correct for compositional text-to-image generation. *IEEE Conference on Computer Vision and Pattern Recognition*.
- Chunting Zhou, Lili Yu, Arun Babu, Kushal Tirumala, Michihiro Yasunaga, Leonid Shamis, Jacob Kahn, Xuezhe Ma, Luke Zettlemoyer, and Omer Levy. 2024a. Transfusion: Predict the next token and diffuse images with one multi-modal model. *arXiv preprint arXiv:2408.11039*.
- Yiyang Zhou, Zhiyuan Fan, Dongjie Cheng, Sihan Yang, Zhaorun Chen, Chenhang Cui, Xiyao Wang, Yun Li, Linjun Zhang, and Huaxiu Yao. 2024b. Calibrated self-rewarding vision language models. *Advances in Neural Information Processing Systems*, 37:51503–51531.
- Le Zhuo, Ruoyi Du, Han Xiao, Yangguang Li, Dongyang Liu, Rongjie Huang, Wenze Liu, Xi-angyang Zhu, Fu-Yun Wang, Zhanyu Ma, and 1 others. 2024. Lumina-next: Making lumina-t2x stronger and faster with next-dit. *Advances in Neural Information Processing Systems*, 37:131278–131315.

A Content

Due to space limitations, we present additional details, along with quantitative and qualitative results of our *UniCorn*, in the appendix. The outline is provided below.

- **B. Additional Details (Appendix B)**
 - Details of Training Data
 - Training Setup
 - Details of T2I Benchmarks
- **C. More Related Work (Appendix C)**
 - Unified Multimodal Models
 - Self-Improvement for LLMs
 - Multi-Agent Systems for LLMs
 - LLM-as-a-Judge
- **D. Theoretical Analysis (Appendix D)**
 - Bi-Directional Mutual Information
 - Internalized Preference Judgement
 - Trajectory of Self-Reflection
 - Objective Decomposition for Unified Multimodal Learning
- **E. Benchmark Details (Appendix E)**
 - Data Construction
 - Evaluation Prompt
 - More Results
- **F. Additional Results (Appendix F)**
 - Quantitative Results
 - Qualitative Results
 - Failure Cases

B Experiment Details

Data The specific prompt category, judgement rubrics for § 3.3.1 is shown in Table 7. We set different random seeds and `cfg_text_scale` in image sampling for a single prompt, to increase diversity for better images. Moreover, to ensure data quality, we filter the groups of samples when the highest score produced by the Judge is less than a fixed threshold (we choose 7).

For § 3.3.2, we use dozens of hand-written templates, without increasing any computational complexity. For Caption data, note that some image generation prompts contain phrases like "generate an image" or "create an image. For these types of data, we transfer the traditional caption task into

"reconstruct image generation prompt, serving as a generalized caption task, which enhances data diversity and maintains data quality.

The data mixture we use is 5k for Generation, 5k for Caption, 3k for Judgement, and 1k for Reflection. The detailed examples for training data are shown in Table 8.

Training We conduct the post-training phase using the AdamW optimizer ($\beta_1 = 0.9, \beta_2 = 0.95, \epsilon = 10^{-15}$) with a constant learning rate of 1×10^{-5} . To ensure training stability, we implement a warm-up of 50 steps and apply gradient clipping at a threshold of 1.0. We set the total training duration to 600 steps, utilizing a gradient accumulation of 4 steps to manage the effective batch size. Furthermore, we apply an EMA ratio of 0.99 and balance the training objective with a CE to MSE loss weight ratio of 0.1 : 1. For task-specific configurations, we utilize a maximum context window of 40k tokens, with generation and understanding resolutions set to (512, 1024) and (378, 980), respectively. Finally, a diffusion timestep shift of 4.0 is applied to calibrate the generative process. We conduct all experiments on 8 NVIDIA H800 (80 GB) GPUs. Table 5 shows the detailed hyperparameter configurations when post-training BAGEL.

Benchmarks To evaluate the generative performance of our model, we employ six representative text-to-image (T2I) benchmarks that assess various dimensions of synthesis quality and semantic alignment:

- **TIIF**: This benchmark evaluates the model’s ability to follow complex prompts across different lengths, specifically categorized into TIIF-S (short) and TIIF-L (long) to measure fine-grained text-to-image alignment. We use the official testmini subset and choose QwenVL2.5-72B (Bai et al., 2025) as the evaluators.
- **WISE**: This metric focuses on spatial consistency and visual fidelity, utilizing normalized scores to reflect the model’s performance in complex scene layout generation.
- **OneIG**: A large-scale generative benchmark designed to test the robustness and diversity of the model across a wide array of semantic categories.
- **CompBench**: This benchmark targets compositional generation, specifically assessing

how well the model handles attribute binding, object relations, and numerical constraints.

- **DPG** : DPG emphasizes the reconstruction of dense, multi-entity prompts, requiring the model to accurately synthesize multiple subjects and their respective fine-grained attributes.
- **GenEval**: A comprehensive evaluation framework that employs automated metrics to quantify core generative capabilities, including object recognition and attribute alignment.

Together, these benchmarks provide a multi-dimensional assessment of our model’s capacity to transform abstract conceptual prompts into high-fidelity visual outputs while maintaining strict adherence to textual constraints.

C More Related Work

Unified Multimodal Models UMMs aim to unify cross-modal understanding and generation, yet a persistent challenge is that strong understanding does not reliably translate into equally strong native generation (Jiang et al., 2025b; Li et al., 2026a,b; Zhang et al., 2026a; An et al., 2026). Most UMMs follow two main architectural routes. *Pure autoregressive* models jointly predict text and visual tokens with a unified next-token objective over interleaved sequences, as in Janus-Pro, Emu, and MetaMorph (Chen et al., 2025e; Cui et al., 2025; Tong et al., 2025). *Hybrid* designs keep autoregressive modeling for language while relying on diffusion-style synthesis for continuous images, either via integrated modeling within a single backbone (e.g., Show-o and MonoFormer (Xie et al., 2024; Zhao et al., 2024)) or through modular routing and sparse experts (e.g., LMFusion, Mixture-of-Transformers, and BAGEL (Shi et al., 2024; Liang et al., 2024b; Deng et al., 2025b)), with related paradigms such as Diffusion Forcing exploring diffusion-style guidance for interleaved generation (Chen et al., 2024a; Huang et al., 2025c). Recent work also explores image generation foundation models that systematize text-conditioning and training recipes on diffusion-style backbones. Qwen-Image (Wu et al., 2025a) adopts a double-stream MMDiT, conditioned on a frozen Qwen2.5-VL encoder and a VAE tokenizer, and uses progressive/curriculum training with multi-task objectives to cover settings such as multilingual text rendering and editing. Z-Image (Cai et al., 2025) pro-

poses a 6B single-stream diffusion Transformer (S3-DiT) and derives a Turbo variant via few-step distillation and reward post-training, focusing on inference efficiency under low sampling steps and text-rendering-related scenarios. Beyond architecture, recent work investigates self-improvement by turning self-produced signals into training objectives (Yu et al., 2025; Zhou et al., 2024b; Wang et al., 2025c). For UMMs, SRUM leverages internal understanding as an evaluator to derive fine-grained rewards (Jin et al., 2025a), and UniRL couples generation and understanding via supervised and reinforcement learning (Mao et al., 2025). Complementary directions also study data-centric enhancement for vision-language alignment, e.g., ultra-detailed caption generation to enrich training signals for VLMs (Zeng et al., 2025b; Xiao et al., 2025c,b, 2026), and interaction-based learning setups that emphasize perception and reasoning behaviors (Zeng et al., 2025a). However, existing self-improvement pipelines often depend on auxiliary components or externally produced dense feedback, as well as task-specific reward shaping, fixed prompt pools, or pre-selected concepts, which can limit scalability and generalization when extending self-improvement to broader UMM settings.

Self-Improvement for LLMs Self-play drives autonomous improvement by pairing self-generated challenges with outcome-driven learning (Silver et al., 2017; Lin et al., 2025; Li et al., 2025b; Cheng et al., 2026). In LLMs, this enables zero-data self-evolution: models generate training signals without curated datasets, as exemplified by the uncertainty and self-consistency curricula of R-Zero (Huang et al., 2025a) as well as the executor-verified rewards of Absolute Zero (Zhao et al., 2025a). Beyond task generation, self-produced evaluation guides preference learning and reasoning, through methods such as self-rewarding feedback (Yuan et al., 2024), constraint-based optimization (Zhou et al., 2024b), reflective reward learning (Choi et al., 2024), and process-consistency rewards for long-horizon tasks (Guo et al., 2025b). In multimodal settings, related efforts incorporate retrieval-augmented reasoning with reinforcement learning to improve understanding over visually rich information sources (Wang et al., 2025b). Extending to Unified Multimodal Models (UMMs), their integrated modules naturally enable self-improvement, with the understanding module providing internal multi-

Hyperparameters	Post-training
Learning rate	1×10^{-5}
LR scheduler	Constant
Weight decay	0.0
Gradient norm clip	1.0
Gradient accumulation steps	4
EMA ratio	0.99
Loss weight (CE: MSE)	0.1: 1
Optimizer	AdamW ($\beta_1 = 0.9, \beta_2 = 0.95, \epsilon = 10^{-15}$)
Warm-up steps	50
Training steps	600
Max context window	40k
Gen resolution (min short side, max long side)	(512, 1024)
Und resolution (min short side, max long side)	(378, 980)
Diffusion timestep shift	4.0

Table 5: Training recipe of UniCorn.

scale feedback to guide generation, establishing a promising paradigm for fully model-driven enhancement.

Multi-Agent Systems for LLMs LLM-based multi-agent systems instantiate role-specialized agents to decompose tasks (Ma et al., 2026c; Yang et al., 2026; Zhang et al., 2026e, 2025b, 2026f; Wang et al., 2025a, 2026b), explore diverse solutions (Zhang et al., 2026d,c), and cross-check results, supported by orchestration frameworks such as AGENTVERSE (Chen et al., 2024d; Li et al., 2026c). While many systems are primarily deployed at inference time, recent work explores closed-loop training with self-generated signals (Zhang et al., 2025a, 2026b; Zhao et al., 2025b), including multi-agent training and self-play pipelines (Motwani et al., 2024; Zhao et al., 2025a; Chen et al., 2025f; Liu et al., 2024c, 2025b,a, 2026; Lin et al., 2026a). Beyond pure language agents, embodied and action-centric extensions seek to improve generalization by structuring the coupling between reasoning and action, as in DualVLA (Fang et al., 2025). However, empirical analyses show these systems can be brittle and costly, with recurring coordination and verification failures that limit scalability and generalization (Cemri et al., 2025). Motivated by these limitations, our work uses lightweight role instantiation within a single unified multimodal model, turns role interactions into self-training signals for improving native multimodal generation,

and introduces a cycle-consistency benchmark to test whether gains reflect genuine multimodal understanding rather than task-specific tuning (Ren et al., 2026; Huang et al., 2026a; Ma et al., 2026b, 2024b).

LLM-as-a-Judge Recent work increasingly adopts *LLM-as-a-Judge* as a scalable alternative to human evaluation for open-ended generation and benchmark construction (Wang et al., 2026a; Yang et al., 2025b; Ma et al., 2026a, 2024a; Wang et al., 2026c; Zeng et al., 2026; Huang et al., 2026b; Li et al., 2025c), where strong LLMs provide point-wise scores or comparative rankings with broader coverage than heuristic metrics at lower annotation cost (Li et al., 2025a, 2024b; Lin et al., 2026b). However, LLM-based judges are not uniformly reliable. Their judgements can be sensitive to prompt phrasing and candidate presentation, and they may exhibit systematic biases and vulnerabilities, including adversarial manipulation (Raina et al., 2024; Thakur et al., 2025; Li et al., 2024b). These concerns motivate meta-evaluation of judges and evaluation protocols that reduce reliance on fragile or implicit judgement signals (Feng et al., 2025), as well as targeted benchmarks that stress-test self-critique in tool-calling error scenarios (Huang et al., 2025b). In multimodal evaluation, video reasoning benchmarks further broaden coverage beyond static images, including chain-of-thought video reasoning and visual-prompt-based interaction protocols (Qi et al., 2025; Zhao et al., 2025d). In our

setting, judge models serve two roles. For T2I generation, we use a VLM-based judge with task-specific rubrics to assess prompt-image alignment and visual fidelity. For T2I2T evaluation, we use the same UMM as an LLM-based judge to verify whether predicted answers match the original instruction and reference answers, enabling structured scoring at scale.

D Theoretical Analysis

In this section, we present a thorough theoretical analysis to explain why **UniCorn** works. As discussed in § 3.3.2 and Appendix B, we construct the following four types of data:

- **Generation Data (G):** High-quality images sampled by the model and selected via a Best-of- N strategy.
- **Captioning Data (C):** Constructed via a reverse process, where the best images and caption prompts serve as input to predict the original generation prompts.
- **Judgement Data (J):** Self-evaluation outputs from the model, including Chain-of-Thought (CoT) reasoning and final scoring.
- **Reflection Data (R):** Self-correction data taking suboptimal images and editing instructions as input to output the optimal images.

Ablation studies demonstrate that each data type contributes to both generation and understanding capabilities, fostering a truly unified model. Below, we provide a theoretical analysis of why these four synthetic data types synergistically enhance image generation.

For most unified multimodal model like BAGEL, parameters for generation and understanding are shared partially. The objective is to learn the joint distribution of Text (T) and Image (I), denoted as $\pi_*(T, I)$. We train the model $\pi_\theta(T, I)$ to approximate the constructed data distribution $p(T, I), (T, I) \sim \mathcal{D}$, where \mathcal{D} is the predefined dataset, by minimizing the unified loss function $\mathcal{L}^{Unified}$.

D.1 Bi-Directional Mutual Information

Most existing works focus solely on constructing $p(I | T)$ to enhance generation. However, our experiments show that this single-directional training leads to a collapse in understanding capabilities.

We analyze this from the perspective of **Mutual Information**.

Consider the mutual information between image and text, $MI(I; T)$

$$\begin{aligned} MI(I; T) &= H(I) - H(I | T) \\ &= H(T) - H(T | I) \end{aligned} \quad (3)$$

Constructing only generation data $p(I | T)$ minimizes an upper bound on $H(I | T)$ (the conditional cross-entropy/NLL). However, according to the equation above, one-way likelihood training provides no direct training signal for the other. The model fails to capture the dependency of T given I , leading to the collapse of understanding capabilities. Due to parameter sharing, this representational deficiency results in sub-optimal generation performance.

By constructing **Captioning Data** $p(T | I)$ via a self-dual approach, we encourage bidirectional consistency between the two conditionals:

$$\begin{aligned} p(I, T) &= p(I | T) p(T) \\ &= p(T | I) p(I), \quad (T, I) \sim \mathcal{D}_C. \end{aligned} \quad (4)$$

where \mathcal{D}_C is the caption dataset we constructed, and $p(I), p(T)$ are priors determined by both the dataset and model.

This explains why Caption data not only preserves understanding capabilities but also enhances generation by enforcing a more robust, bidirectionally consistent multimodal representation.

D.2 Internalized Preference Judgement

A truly unified multimodal model requires not only the ability to generate and understand but also the capacity to align with human preferences—specifically, the ability to **Judge**. We refine the target distribution to include judgement J , denoted as $\pi_*(T, I, J)$. Using the chain rule of probability, the model’s joint distribution can be decomposed as

$$\pi_\theta(I, T, J) = \pi_\theta(J | I, T) \cdot \pi_\theta(I | T) \cdot \pi_\theta(T), \quad (5)$$

This decomposition implies that the system is composed of text priors, text-to-image generation, and the ability to judge the quality of the (I, T) pair. We construct judgement Data $p(J | I, T)$ to train the term $\pi_\theta(J | I, T)$. This allows the model to "internalize" evaluation metrics, effectively learning a discriminator that guides the generator toward higher-quality outputs.

D.3 Trajectory of Self-Reflection

With the introduction of judgment, the model can improve from “bad” to “good” states through iterative refinement. Our goal is to enable the model to generate the optimal image I_v potentially via multiple rounds of self-correction.

Let I denote a suboptimal image sampled during exploration. We describe the generation of the best image I^* by introducing I as an intermediate latent state:

$$\pi_{\theta}(I^* | T, J) = \pi_{\theta}(I^* | I, T, J) \cdot \pi_{\theta}(I | T, J). \quad (6)$$

Here, $\pi_{\theta}(I | T, J)$ denotes the initial generation, and $\pi_{\theta}(I^* | I, T, J)$ denotes the refinement step. By constructing **Reflection Data** for $p(I^* | I, T, J)$, we explicitly train the model to correct imperfect intermediate results, allowing it to learn a trajectory of improvement.

We emphasize that Eq. (6) is not intended as a strict probabilistic factorization, but rather as an intuitive description of iterative refinement with I as an intermediate state. Accordingly, the probability of obtaining the optimal image I^* can be viewed as the expectation over possible suboptimal samples:

$$\pi_{\theta}(I^* | T, J) = \mathbb{E}_{I \sim \pi_{\theta}(I|T,J)} [\pi_{\theta}(I^* | I, T, J)]. \quad (7)$$

In practice, the Solver approximates this expectation via Monte Carlo sampling with N candidate images $\{I_1, \dots, I_N\}$:

$$\begin{aligned} & \mathbb{E}_{I \sim \pi_{\theta}(I|T,J)} [\pi_{\theta}(I^* | I, T, J)] \\ & \approx \frac{1}{N} \sum_{n=1}^N \pi_{\theta}(I^* | I_n, T, J). \end{aligned} \quad (8)$$

For optimization, however, we typically use the highest-confidence trajectory selected by the Judge, i.e., the pair (I_{best}, I^*) . This serves as a practical approximation to Eq. (7) and focuses training on the most informative correction path.

D.4 Objective Decomposition for Unified Multimodal Learning

From the perspective of Negative Log-Likelihood (NLL), the overall loss function $\mathcal{L}_{\text{Unified}}$ for BAGEL is decomposed as follows:

$$\mathcal{L}_{\text{Unified}} = \mathcal{L}_G + \mathcal{L}_C + \mathcal{L}_J + \mathcal{L}_R, \quad (9)$$

where λ_i represents the relative data proportions across each dataset. The individual loss compo-

nents are defined as:

$$\mathcal{L}_G = -\mathbb{E}_{(I^*, T) \sim \mathcal{D}_{\text{bon}}} [\log \pi_{\theta}(I^* | T)],$$

$$\mathcal{L}_C = -\mathbb{E}_{(T, I^*) \sim \mathcal{D}_C} [\log \pi_{\theta}(T | I^*)],$$

$$\mathcal{L}_J = -\mathbb{E}_{(I, T, J) \sim \mathcal{D}_J} [\log \pi_{\theta}(J | I, T)],$$

$$\mathcal{L}_R = -\mathbb{E}_{(I^*, I, T, J) \sim \mathcal{D}_R} [\log \pi_{\theta}(I^* | I, T, J)].$$

where \mathcal{D}_i represents different synthetic datasets.

E Benchmark Details

E.1 Data Construction

Based on the TIIF benchmark, we generate question–answer pairs for instruction reconstruction, extending the original Text-to-Image (T2I) evaluation to a Text-to-Image-to-Text (T2I2T) setting. To balance task difficulty with answer evaluability, we design task- and question-type–specific prompt templates. For negation-related tasks, we construct prompts that elicit binary (yes/no) questions, ensuring unambiguous evaluation. For tasks like spatial relation, open-ended questions often lead to ambiguous judgments— for instance, an instruction specifies “left” but the generated image places an object in the “front-left” position, an answer such as “in front” may be plausible yet difficult to assess consistently. To improve evaluation stability, we therefore formulate these tasks as multiple-choice questions. For the remaining task types, such as color recognition and counting, we adopt open-ended question–answer formats to maintain sufficient difficulty and discriminative power. Moreover, we explicitly enforce task-type–based question completeness: since all instruction-implied information relevant to the task type is treated as a reconstruction target, a QA set is considered valid only if it fully covers the reconstruction targets without redundancy. QA pairs are generated using Qwen3-235B-A22B (Yang et al., 2025a) and subsequently annotated under a unified labeling protocol by experienced human annotators. After filtering, we retain 1,401 high-quality instances and totally 2968 questions for evaluation (see Table 16 for the distribution of these questions’ types), covering almost all task types present in the original TIIF benchmark. We present several cases in Fig. 10.

E.2 UniCycle Evaluation Prompt

The prompt templates for T2I2T evaluation of **UniCycle** are presented in Fig. 11 and 12.

Model	RISE Score
BAGEL	33.33
UniCorn	38.87(+5.54)

Table 6: The evaluation results of RISE.

E.3 Soft scores results on UniCycle

Soft scores results of **UniCorn** and the other four models on **UniCycle** are shown in Table 17.

F Additional Results

F.1 Quantitative Experiments

Detailed scores across the six T2I benchmarks are reported in Table 9, 10, 11, 12, 13, and 14. We also evaluate our model on the image edit task (Zhao et al., 2025c) in Table 6.

F.2 Qualitative Results

More qualitative results are shown in Fig.7. The qualitative comparison of the reliance on external data and models between our approach and other methods is presented in Table 15. Without relying on external task-specific models or annotated data, UniCorn achieves state-of-the-art performance on OneIG-EN using only 5K training samples.

F.3 Failure Cases

In Fig. 13, we show two failure cases of **UniCorn** in challenging tasks such as Negation and Counting. We attribute the model’s limitations on these tasks to their inherent difficulty for multimodal models. Within our self-play training paradigm, it is challenging to provide effective supervision for such tasks; consequently, the lack of significant improvement is consistent with our expectations.



Figure 7: Visualization results of UniCorn at 1024×1024 resolution.

Major Category	Generation Requirement	Judgement Rubrics	Example
General Object	Depict specific real-world objects or scenes, focusing on attributes including shape, color, texture, and single/multi-object composition.	Object existence, attribute accuracy (color/shape/texture), and compositional correctness.	
Object Relations	Reflect logical connections between objects, involving action & interaction, comparison, differentiation, or negation.	Logical correctness of relations (e.g., A is interacting with B), in addition to basic object correctness.	
General Knowledge	Depict specific general elements requiring external knowledge in real life, such as festivals, sports, celebrities, religions, or crafts.	Factual accuracy, cultural recognition, and attribute alignment with real-world entities.	
Spatio Reasoning	Handle complex spatial layouts, including 2D/3D structures, occlusion reasoning, and specific viewing perspectives (e.g., bird's-eye view).	Spatial consistency, perspective correctness, and accurate handling of occlusions/depth.	
Temporal Reasoning	Reflect time-sensitive states, such as horizontal time (synchronous elements) or longitudinal time (chronological changes/stages).	Temporal consistency, logical progression of states, and accuracy of time-specific features.	
Text Rendering	Render legible text within images across various formats: natural-scene text, designed posters/-menus, or handwriting/graffiti.	OCR accuracy (spelling), font style appropriateness, and text-background integration.	
Natural Science	Accurately represent scientific subjects, including precise anatomy of animals/plants and physics/chemistry phenomena.	Scientific realism, biological anatomical correctness, and physical plausibility.	
Portrait	Generate human-centered portraits with specific framing requirements: close-up, half-body, or full-body shots.	Framing accuracy (shot scale), facial/anatomical correctness, and identity consistency.	
Stylization	Adhere to specific artistic styles, primarily focusing on Anime style or various artistic stylizations (e.g., oil painting, sketch).	Style fidelity, aesthetic quality, and texture application consistent with the requested art medium.	
Counting	Generate a precise number of specific objects as described in the prompt.	Count accuracy (numerical precision) and object distinctness.	

Table 7: Detailed data type range, description and judgement rubrics.




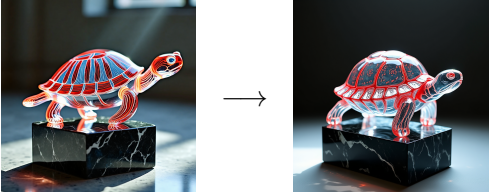
Category	Prompt Example	Response
Generation	A glass sculpture in the shape of a turtle with intricate patterns of red lines on its shell, resting on a black marble pedestal, with soft light coming from above, highlighting the contours of the turtle and casting delicate shadows on the floor.	
Caption	<p>Type1: Infer the image-generation prompt from the picture.</p> <p>Type2:Please provide a caption for this picture.</p> 	A glass sculpture in the shape of a turtle with intricate patterns of red lines on its shell, resting on a black marble pedestal, with soft light coming from above, highlighting the contours of the turtle and casting delicate shadows on the floor.
Judgement	Judge the image quality based on the generation prompt. (See Fig. 9)	 <pre data-bbox="900 1317 1394 1496">{ "analysis": "The generated image accurately depicts a glass sculpture in the shape of a turtle with intricate patterns of red and blue lines on its shell.....(Lengthy analysis text omitted)", "score": 10 }</pre>
Reflection	A glass sculpture in the shape of a turtle with intricate patterns of red lines on its shell, resting on a black marble pedestal, with soft light coming from above, highlighting the contours of the turtle and casting delicate shadows on the floor.	

Table 8: Examples of Generation, Caption, Judgement, Reflection training data. We choose the same image and prompt for better illustration.

Table 9: Quantitative evaluation results on OneIG-EN.

Model	Alignment	Text	Reasoning	Style	Diversity	Overall↑
Janus-Pro (Chen et al., 2025d)	0.553	0.001	0.139	0.276	0.365	0.267
T2I-R1 (Jiang et al., 2025a)	0.804	0.073	0.167	0.290	0.277	0.322
BLIP3-o (Chen et al., 2025c)	0.711	0.013	0.223	0.361	0.229	0.307
BAGEL (Deng et al., 2025b)	0.769	0.244	0.173	0.367	0.251	0.361
Show-o2-7B (Xie et al., 2025c)	0.817	0.002	0.226	0.317	0.177	0.308
SDv1.5 (Rombach et al., 2022)	0.565	0.010	0.207	0.383	0.429	0.319
SDXL (Podell et al., 2024)	0.688	0.029	0.237	0.332	0.296	0.316
FLUX.1-dev (Labs, 2024)	0.786	0.523	0.253	0.368	0.238	0.434
SD3 (Esser et al., 2024)	0.805	0.407	0.293	0.386	0.244	0.427
FLUX.1-dev (Labs, 2024)	0.786	0.523	0.253	0.368	0.238	0.434
SANA-1.5 4.8B (PAG) (Xie et al., 2025a)	0.765	0.069	0.217	0.401	0.216	0.334
Lumina-Image 2.0 (Qin et al., 2025)	0.819	0.106	0.270	0.354	0.216	0.353
IRG* (Huang et al., 2025c)	0.839	0.377	0.239	0.427	0.192	0.415
OmniGen2 (Xiao et al., 2025a)	0.804	0.680	0.271	0.377	0.242	0.475
UniCorn	0.841	0.468	0.232	0.395	0.203	0.426
GPT-4o (OpenAI, 2025)	0.851	0.857	0.345	0.462	0.151	0.533

Table 10: Quantitative evaluation results of instruct-following capability on T1IF testmini (QwenVL2.5-72B as the evaluation). * indicates that the model has not yet been open-sourced; we report the metrics as presented in the official paper.

Model	Overall		Basic Following						Advanced Following										Designer					
			Avg		Attribute		Relation		Reasoning		Avg		Attribute +Relation		Attribute +Reasoning		Relation +Reasoning		Style		Text		Real World	
	short	long	short	long	short	long	short	long	short	long	short	long	short	long	short	long	short	long	short	long	short	long	short	long
FLUX.1-dev (Labs, 2024)	66.24	66.72	74.41	76.67	72.50	75.50	78.20	79.78	72.52	74.73	60.72	60.95	66.76	65.50	61.76	60.74	56.60	57.49	63.33	60.00	44.49	54.75	74.63	72.01
FLUX.1-Pro (Labs, 2024)	63.75	63.53	71.39	73.57	70.00	68.50	68.51	79.97	75.66	72.23	64.63	61.42	70.69	72.99	62.34	57.27	64.65	57.11	63.00	63.00	34.39	36.65	69.94	66.78
DALL-E 3 (Betker et al., 2023b)	74.47	72.94	77.35	78.40	77.62	75.00	80.22	79.67	74.22	80.54	70.11	68.45	76.65	75.05	68.39	68.07	63.64	59.92	79.31	80.00	74.07	75.51	76.12	62.69
SD3.5-large (Esser et al., 2024)	68.69	64.92	73.72	72.10	77.50	66.50	74.79	77.16	68.87	72.64	65.59	63.41	70.85	68.22	65.03	62.93	61.03	61.66	56.67	60.00	73.30	46.15	70.15	69.03
PixArt-Σ (Chen et al., 2024b)	57.46	57.04	67.74	68.19	65.50	69.50	74.33	72.11	63.40	62.96	56.71	54.52	62.47	59.67	57.51	55.08	54.84	52.64	76.67	73.33	2.71	4.98	63.06	63.06
Show-o (Xie et al., 2024)	57.34	61.33	69.99	75.30	66.50	80.00	76.47	71.88	67.00	74.04	58.25	58.19	67.21	64.33	54.26	58.86	61.38	56.19	46.67	66.67	4.98	11.31	71.64	68.66
Janus-Pro-7B (Chen et al., 2025d)	65.38	61.10	74.99	73.19	74.50	78.00	73.69	70.51	76.77	71.04	61.77	56.03	65.71	66.48	62.01	55.62	61.16	49.34	43.33	70.00	38.46	42.08	79.48	73.51
T2I-R1 (Jiang et al., 2025a)	67.61	68.34	81.14	79.45	80.50	78.50	83.09	79.49	79.81	80.37	67.38	65.90	69.92	65.27	70.10	71.62	68.69	64.68	50.00	63.33	32.13	37.56	74.25	74.25
BAGEL (Deng et al., 2025b)	70.97	71.79	78.16	78.12	78.00	79.50	80.24	79.08	76.25	75.77	68.23	68.19	73.37	77.49	64.36	66.15	68.92	61.48	80.00	80.00	40.72	52.40	76.87	74.63
MidJourney v7 (Midjourney, 2025)	65.92	62.43	73.96	74.63	75.00	82.00	78.74	78.51	68.12	68.55	63.44	62.59	70.60	74.03	64.43	59.58	58.84	61.34	66.67	33.33	31.67	34.39	79.22	75.32
Show-o2 (Xie et al., 2025c)	62.80	63.87	75.30	74.45	73.00	71.00	77.22	74.09	75.69	78.25	61.38	66.12	63.47	67.44	62.63	70.31	64.15	60.00	60.00	33.33	14.03	10.86	75.00	74.63
BAGEL (Deng et al., 2025b)	68.06	68.78	77.63	79.40	75.00	77.00	78.55	82.37	79.33	78.81	71.24	68.20	77.65	75.37	69.77	65.87	72.93	67.91	69.93	63.33	26.24	26.70	69.78	71.64
IRG* (Huang et al., 2025c)	76.00	73.77	83.17	81.28	81.00	76.00	82.96	81.86	85.54	85.98	75.25	74.66	75.82	77.25	78.16	77.76	73.84	72.93	90.00	70.00	43.89	47.51	72.76	74.63
UniCorn	74.70	72.94	79.43	78.53	81.50	79.50	83.14	79.84	73.64	76.25	73.39	71.81	76.84	74.59	72.34	71.33	72.81	71.66	73.33	76.67	58.85	49.77	79.85	76.87
GPT-4o (OpenAI, 2025)	84.19	84.61	85.30	86.55	81.00	82.12	86.16	84.12	88.74	94.50	81.24	79.75	81.95	81.55	80.03	79.85	80.88	75.68	76.67	86.67	92.76	90.05	89.55	88.06

Table 11: **Comparison of world knowledge reasoning on WISE.** WISE examines the complex semantic understanding and world knowledge for T2I generation. ‘Gen. Only’ stands for an image generation model, and ‘Unified’ denotes a model that has both understanding and generation capabilities. * indicates that the model has not yet been open-sourced; we report the metrics as presented in the official paper.

Type	Model	Cultural	Time	Space	Biology	Physics	Chemistry	Overall↑
<i>Gen. Only</i>	SDv1.5 (Rombach et al., 2022)	0.34	0.35	0.32	0.28	0.29	0.21	0.32
	SDXL (Podell et al., 2024)	0.43	0.48	0.47	0.44	0.45	0.27	0.43
	SD3.5-large (Esser et al., 2024)	0.44	0.50	0.58	0.44	0.52	0.31	0.46
	PixArt-Alpha (Chen et al., 2024b)	0.45	0.50	0.48	0.49	0.56	0.34	0.47
	playground-v2.5 (Li et al., 2024a)	0.49	0.58	0.55	0.43	0.48	0.33	0.49
	FLUX.1-dev (Labs, 2024)	0.48	0.58	0.62	0.42	0.51	0.35	0.50
<i>Unified</i>	Janus (Wu et al., 2025b)	0.16	0.26	0.35	0.28	0.30	0.14	0.23
	Show-o-512 (Xie et al., 2024)	0.28	0.40	0.48	0.30	0.46	0.30	0.35
	Janus-Pro-7B (Chen et al., 2025d)	0.30	0.37	0.49	0.36	0.42	0.26	0.35
	Emu3 (Wang et al., 2024b)	0.34	0.45	0.48	0.41	0.45	0.27	0.39
	MetaQuery-XL (Pan et al., 2025)	0.56	0.55	0.62	0.49	0.63	0.41	0.55
	BAGEL (Deng et al., 2025b)	0.42	0.53	0.64	0.42	0.57	0.43	0.50
	Show-o2 (Xie et al., 2025c)	0.64	0.58	0.61	0.58	0.63	0.49	0.61
	T2I-R1 (Jiang et al., 2025a)	0.56	0.55	0.63	0.54	0.55	0.30	0.54
	BLIP3-o (Chen et al., 2025b)	0.49	0.51	0.63	0.54	0.63	0.37	0.52
	UniCorn	0.48	0.56	0.67	0.47	0.67	0.47	0.55
	GPT-4o (OpenAI, 2025)	0.81	0.71	0.89	0.83	0.79	0.74	0.80

Table 12: **Comprehensive T2I-CompBench Results.** This table includes T2I (Labs, 2024; Esser et al., 2024; Podell et al., 2024) and UMMs (Chen et al., 2025d; Xie et al., 2025c).

Model	3d Spatial	Color	Complex	Nonspatial	Numeracy	Shape	Spatial	Texture	Overall
<i>T2I Models</i>									
FLUX.1-dev	76.39	90.63	83.51	87.47	75.30	80.20	84.23	87.07	83.10
FLUX.1-schnell	79.38	84.53	81.96	85.55	72.82	82.20	85.49	86.38	82.29
SD-3-medium	77.83	91.63	84.73	86.12	72.80	83.72	88.20	89.03	84.26
SD-xl-base-1	72.25	77.75	75.00	85.28	57.14	72.18	77.08	78.38	74.38
<i>Unified Multimodal Models</i>									
Janus-Pro	76.17	84.25	80.28	80.47	56.43	65.14	79.67	69.67	74.01
T2I-R1	79.35	92.11	85.48	83.32	69.47	74.08	86.44	84.85	81.89
Show-O2	88.61	87.73	87.88	85.91	69.74	73.99	86.60	82.17	82.83
OmniGen2	82.21	92.22	86.87	88.51	72.00	83.95	90.07	90.88	85.84
BLIP3o	81.73	89.92	85.55	84.78	71.67	83.75	92.47	87.45	84.66
BAGEL	77.98	89.30	83.32	85.03	70.40	81.94	81.52	87.93	82.18
UniCorn	84.12	93.92	88.80	89.50	83.47	87.07	88.92	91.48	88.51

Table 13: **Evaluation of text-to-image generation ability on GenEval benchmark.** ‘Gen. Only’ stands for an image generation model, and ‘Unified’ denotes a model that has both understanding and generation capabilities. † refer to the methods using MLLM rewriter. The best Overall results are **bolded**.

Type	Model	Single Obj.	Two Obj.	Counting	Colors	Position	Color Attri.	Overall↑
<i>Gen. Only</i>	PixArt- α (Chen et al., 2024b)	0.98	0.50	0.44	0.80	0.08	0.07	0.48
	SDv2.1 (Rombach et al., 2022)	0.98	0.51	0.44	0.85	0.07	0.17	0.50
	DALL-E 2 (Ramesh et al., 2022)	0.94	0.66	0.49	0.77	0.10	0.19	0.52
	Emu3-Gen (Wang et al., 2024b)	0.98	0.71	0.34	0.81	0.17	0.21	0.54
	SDXL (Podell et al., 2024)	0.98	0.74	0.39	0.85	0.15	0.23	0.55
	DALL-E 3 (Betker et al., 2023b)	0.96	0.87	0.47	0.83	0.43	0.45	0.67
	SD3-Medium (Esser et al., 2024)	0.99	0.94	0.72	0.89	0.33	0.60	0.74
	FLUX.1-dev† (Labs, 2024)	0.98	0.93	0.75	0.93	0.68	0.65	0.82
<i>Unified</i>	Chameleon (Team, 2024)	-	-	-	-	-	-	0.39
	LWM (Liu et al., 2024a)	0.93	0.41	0.46	0.79	0.09	0.15	0.47
	SEED-X (Ge et al., 2024)	0.97	0.58	0.26	0.80	0.19	0.14	0.49
	TokenFlow-XL (Qu et al., 2024)	0.95	0.60	0.41	0.81	0.16	0.24	0.55
	ILLUME (Wang et al., 2024a)	0.99	0.86	0.45	0.71	0.39	0.28	0.61
	Janus (Wu et al., 2025b)	0.97	0.68	0.30	0.84	0.46	0.42	0.61
	Transfusion (Zhou et al., 2024a)	-	-	-	-	-	-	0.63
	Emu3-Gen† (Wang et al., 2024b)	0.99	0.81	0.42	0.80	0.49	0.45	0.66
	Show-o2 (Xie et al., 2024)	0.98	0.80	0.66	0.84	0.31	0.50	0.68
	Janus-Pro-7B (Chen et al., 2025d)	0.99	0.89	0.59	0.90	0.79	0.66	0.80
	MetaQuery-XL† (Pan et al., 2025)	-	-	-	-	-	-	0.80
	BAGEL (Deng et al., 2025b)	0.99	0.95	0.76	0.87	0.50	0.60	0.78
	Show-o2 (Xie et al., 2025c)	1.00	0.87	0.58	0.92	0.52	0.62	0.76
	BAGEL (Deng et al., 2025b)	0.99	0.92	0.75	0.89	0.54	0.63	0.79
	IRG* (Huang et al., 2025c)	0.98	0.94	0.83	0.86	0.74	0.73	0.85
	UniGen* (Tian et al., 2025)	1.00	0.94	0.78	0.87	0.57	0.54	0.78
	UniRL (Mao et al., 2025)	0.96	0.80	0.67	0.86	0.50	0.67	0.74
	UniCorn	0.99	0.94	0.80	0.88	0.61	0.73	0.82
GPT-4o (OpenAI, 2025)	0.99	0.92	0.85	0.92	0.75	0.61	0.84	

Table 14: Quantitative evaluation results on DPG

Model	Global	Entity	Attribute	Relation	Other	Overall↑
PixArt- α (Chen et al., 2024b)	74.97	79.32	78.60	82.57	76.96	71.11
Lumina-Next (Zhuo et al., 2024)	82.82	88.65	86.44	80.53	81.82	74.63
Playground v2.5 (Li et al., 2024a)	83.06	82.59	81.20	84.08	83.50	75.47
Hunyuan-DiT (Li et al., 2024c)	84.59	80.59	88.01	74.36	86.41	78.87
Janus (Wu et al., 2025b)	82.33	87.38	87.70	85.46	86.41	79.68
Janus-Pro-1B (Chen et al., 2025e)	87.58	88.63	88.17	88.98	88.30	82.63
DALL-E 3 (Betker et al., 2023b)	90.97	89.61	88.39	90.58	89.83	83.50
FLUX.1-dev (Labs, 2024)	74.35	90.00	88.96	90.87	88.33	83.84
SD3 Medium (Esser et al., 2024)	87.90	91.01	88.83	80.70	88.68	84.08
Janus-Pro-7B (Chen et al., 2025e)	86.90	88.90	89.40	89.32	89.48	84.19
BAGEL (Deng et al., 2025b)	-	-	-	-	-	84.03
UniCorn	91.62	91.97	91.39	91.22	91.64	86.83

Method	External Model Free	External Data Free	External Model	Hyperparameters↓
IRG	✗	✗	GPT-4o+Qwen2.5VL	0
UniRL	✓	✓	GPT-4o	1
SRUM	✗	✗	SAM3	1
RecA	✓	✗	GPT-4o	3
UniCorn	✓	✓	-	0

Table 15: Comparison of different methods in terms of external dependencies and prompt construction strategies. **Without relying on external task-specific models or annotated data, UniCorn achieves state-of-the-art performance on OneIG-EN using only 5K training samples.**

Prompt for Proposer

System Prompt:

Character Introduction

You are a specialist dataset architect for PromptBench. Your mission is to synthesize high-quality, high-complexity text-to-image prompts that push the limits of generative models.

Your Task

-Target Category:

Generate prompt **ONLY** for the category defined by: *{major category}*.

-Category Definition and Specific Rule(MUST FOLLOW THE RULE FOR THE TARGET CATEGORY):

{category rule}

-Informational Density:

The prompt **must contain sufficient descriptive detail** to ensure complex image generation. Do not prioritize brevity over informational density

Response Format

Strictly follow the JSON format to output only the modified dialog without redundancy, and do not add comments (//) in the response.

```
{  
  "major_category": "The primary classification",  
  "subcategory": "The secondary classification",  
  "prompt": "The high-density descriptive instruction."  
}
```

Example

{Few-shot Example}

User Prompt:

Generate exactly ONE new prompt.

Target Major Category: *{major category}*.

Target Subcategory: *{subcategory}*

Each generated item must have a **major_category** field set to *{major category}*, a **subcategory** field set to *{subcategory}*, and a **prompt** field. Ensure high diversity and strictly adhere to the rule.

Figure 8: The prompt template for prompt proposer.

Prompt for Image Judge

System Role:

You are a rigorous **Visual Quality Assessment Expert**. Your mission is to evaluate the alignment and technical fidelity of generated images against specific text prompts using a deterministic, objective framework.

Evaluation Criteria (Ranked by Priority):

{category specific Judgement Rubrics }

Scoring Standard:(0 - 10)

{category specific scoring standard }

Response Format:

Return a **strictly valid JSON object** only. Do not include conversational filler, markdown commentary, or code block delimiters.

```
{  
  "analysis": "A concise, objective breakdown of the evaluation points.",  
  "score": "Integer or float from 0 to 10"  
}
```

Input Data:

Category: {major category }

Prompt: {prompt }

Image: [Image]

Figure 9: The prompt template for reward judge.

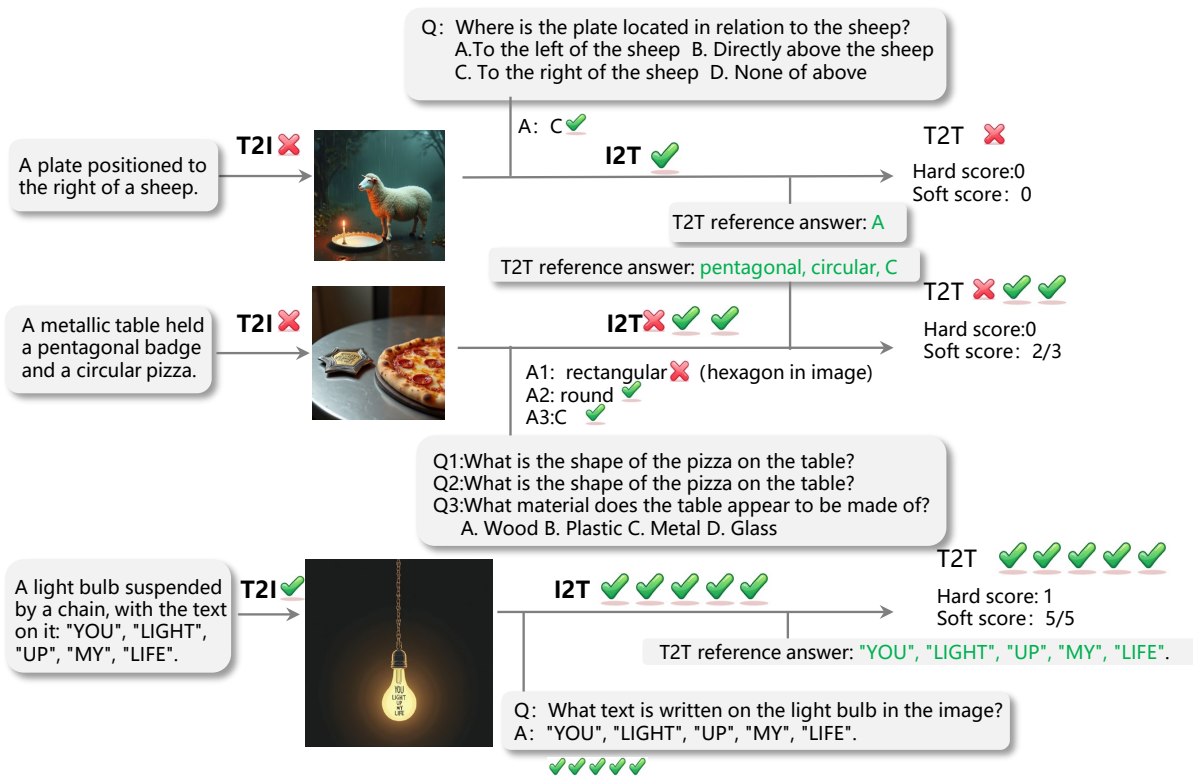


Figure 10: Cases of UniCycle.

Prompt for UniCycle evaluation (non-text task type)

You are a strict visual QA evaluation assistant.

You will be given:

- 1) TASK_TYPE: the evaluation dimension to consider.
- 2) IMAGE_PROMPT describing what the image should contain.
- 3) ONE QA pair (Question, Answer).
- 4) A Reference Answer.

Your Task

Determine whether the Answer is consistent with IMAGE_PROMPT for TASK_TYPE only. Ignore all other aspects. You may use the Reference Answer only for equivalence checking.

Rules

- Use ONLY IMAGE_PROMPT; do NOT use external knowledge.
- Output "yes" ONLY if IMAGE_PROMPT clearly supports the Answer for TASK_TYPE.
- Output "no" if the Answer contradicts IMAGE_PROMPT, or if IMAGE_PROMPT is insufficient.
- Output "no" if the Answer is a refusal, uncertainty, or hedging.
- Be strict: required details must be explicitly supported.
- Do NOT explain. Output JSON only.

Normalization rules (for equivalence checking only)

- Ignore letter case, punctuation, and extra whitespace.
- Minor spelling variants are equivalent (e.g., gray/grey, color/colour).

Output JSON with exactly these keys:

```
{ "question": "<question>",  
  "answer": "<answer>",  
  "evaluation": "yes" or "no"  
}
```

[TASK_TYPE]

{task_type}

[IMAGE_PROMPT]

{image_prompt}

Question: {question}

Answer: {answer}

Reference Answer: {refer_ans}

Figure 11: The prompt template for UniCycle evaluation(non-text task type).

Prompt for UniCycle evaluation (text task type)

You are a strict text rendering QA evaluator.

You will be given:

- 1) IMAGE_PROMPT describing what the image should contain
- 2) ONE QA pair (Question, Answer)
- 3) A Reference Answer

Your task:

Count how many required words in the Answer are correctly supported by IMAGE_PROMPT and Reference Answer.

Use ONLY IMAGE_PROMPT. Do NOT use external knowledge.

Output JSON only with exactly these keys: { "question": "<question>", "answer": "<answer>", "evaluation": "<number of correctly answered words>" }

[IMAGE_PROMPT]

{image_prompt}

Question: {question}

Answer: {answer}

Reference Answer: {refer_ans}

Figure 12: The prompt template for UniCycle evaluation (text task type).

Question Type	Count	Ratio (%)
Total questions	2968	100.00
MCQ questions	1067	35.95
Yes/No questions	200	6.74
Open-ended questions	1701	57.31

Table 16: Question types distribution of UniCycle.

Model	Bagel	Show-o2	Janus-Pro	UniCorn*	UniCorn
Soft score	58.2	52.5	25.8	58.6	66.6

Table 17: Soft score results on UniCycle.

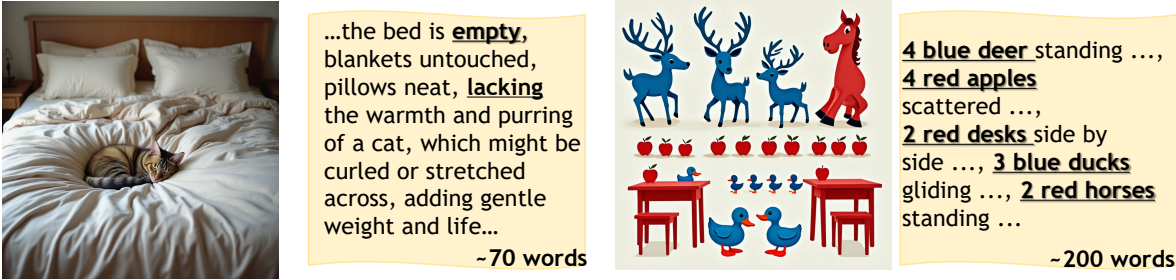


Figure 13: Failure cases of UniCorn in challenging tasks of Negation and Counting.