

SRUM: Fine-Grained Self-Rewarding for Unified Multimodal Models

Weiyang Jin^{1,2,5*}, Yuwei Niu^{3*}, Jiaqi Liao², Chengqi Duan^{1,2}, Aoxue Li⁴,
Shenghua Gao², and Xihui Liu^{1,2†}

¹ HKU MMLab

² The University of Hong Kong

³ Peking University

⁴ Noah's Ark Lab, Huawei

⁵ Shenzhen Loop Area Institute

`weiyangjin@connect.hku.hk`

Abstract. Recently, remarkable progress has been made in Unified Multimodal Models (UMMs), which integrate vision-language generation and understanding capabilities within a single framework. However, a model's strong visual understanding often fails to transfer to visual generation: it may correctly judge prompt-image alignment while failing to generate a faithful image from the same prompt. This raises a compelling question: Can a model improve itself by using its understanding module to reward its generation module? We introduce SRUM, a self-rewarding post-training framework directly applicable to existing UMMs of various designs. SRUM creates a feedback loop where the model's own understanding module acts as an internal "evaluator", providing corrective signals to improve generation without additional human-labeled data or external reward models. To provide comprehensive feedback, SRUM uses a global-local dual reward system: a **global reward** ensures overall visual semantics and layout, while a **local reward** refines fine-grained, object-level fidelity. SRUM shows strong generalization, boosting performance on T2I-CompBench from 82.18 to **88.37** and on T2I-ReasonBench from 43.82 to **46.75**. Overall, our work establishes a powerful paradigm for enabling a UMM's understanding module to guide and enhance its own generation via self-rewarding.

Keywords: UMMs · Self-improvement · Post-training

1 Introduction

Text-to-Image (T2I) models have achieved remarkable progress in generating high-quality and diverse images from given prompts [28, 32, 33]. However, they often fail to accurately interpret instructions involving world knowledge, complex spatial relationships, detailed attribute binding, or compositional reasoning [16].

*Equal contribution. †Corresponding author: xihuiliu@eee.hku.hk. Project page: https://waynejin0918.github.io/srum_web/.

These limitations point to a fundamental lack of deep semantic understanding in standard T2I models. To address this challenge, researchers have developed Unified Multimodal Models (UMMs), which integrate both understanding and generation capabilities within a single framework [7, 44, 45, 47]. By sharing a common backbone, UMMs possess the inherent potential for synergy, offering a promising path to resolve the comprehension challenges that plague traditional T2I models.

Despite their advanced architecture, a fundamental paradox plagues current UMMs: their capacity to generate falls far behind their ability to understand [4, 27, 37, 41, 48]. For instance, a model can often correctly judge the alignment between a detailed prompt and a complex image, yet be incapable of generating a faithful image from that same prompt (Figure 1). This persistent gap between understanding and generation suggests that the key to unlocking better generation lies within the model itself.

To address this challenge, we propose bridging this module gap through self-rewarding. We introduce **Self-Rewarding for Unified Multimodal Models** (SRUM), a novel post-training framework designed to create a synergistic feedback loop within the model itself. Our core insight is that the solution lies within the UMMs’ own architecture. By treating the generation module as a “student” and the more capable understanding module as an internal “teacher” or “evaluator,” we establish a practical self-rewarding system for improvement, without external reward models, human labels, or additional image data during the training phase.

Furthermore, to effectively guide the generation of complex scenes, a reward signal should provide multi-scale feedback. As our ablation studies confirm, a single, holistic score is insufficient because it fails to provide the fine-grained corrective signals needed for detailed improvement. Therefore, we propose a global-local dual reward framework. The global reward evaluates high-level compositional coherence to ensure overall scene plausibility. Concurrently, the local reward targets object-level details, optimizing attribute binding and spatial arrangements. This synergistic design enables SRUM to enhance the performance of the base model on complex generation tasks.

Through extensive experiments, we demonstrate that our approach significantly improves the composition, reasoning, and visual fidelity of UMMs, showing strong generalization across in-domain and out-of-domain settings. SRUM achieves SOTA results on T2I-CompBench and T2I-ReasonBench, improving the overall score of a strong baseline model from 82.18 to 88.37 in composition and from 43.82 to 46.75 in reasoning. Our key contributions can be summarized as follows:

- We are the first to propose a comprehensive self-rewarding framework for UMMs at the post-training stage, successfully bridging the gap between their understanding and generation.
- We introduce a novel dual reward design that combines global compositional assessment with local object-level feedback, providing solid and multi-scale guidance during model training.




Fig. 1: The example on the left suggests that the current UMMs’ understanding module has exceeded the capability of its generation module: the generation module is prone to producing incorrect candidate images based on a given prompt in relevant scenarios, a situation which the understanding module can reasonably identify. This not only highlights a gap between understanding and generation but also reveals the potential for understanding to guide generation. Inspired by this insight, we propose **SRUM** to bridge this gap, particularly in complex generation domains.

- We achieve better performance on complex compositional generation and demonstrate strong generalization. Ultimately, SRUM establishes a powerful paradigm for a UMM’s understanding module to guide its own generation module toward self-improvement.

2 Related Works

2.1 Architectures for Unified Multimodal Models

Unified Multimodal Models (UMMs) have emerged as a prominent research direction, aiming to integrate diverse tasks like visual understanding and generation within a single, end-to-end trained architecture. This consolidation seeks to foster synergy across modalities and reduce systemic complexity. Recent architectural paradigms can be broadly categorized. The **Purely Autoregressive (AR)** approach extends the next-token prediction paradigm of LLMs to visual data, treating images as a sequence of discrete tokens [36, 41]. A key refinement in this area involves decoupling the visual encoders, using a semantic encoder for understanding tasks while retaining a reconstruction-based tokenizer for generation, as demonstrated by Janus [44]. Show-O further refines this by integrating a discrete-diffusion schedule to improve token prediction [47]. More prevalent are hybrid architectures that combine the strengths of AR and diffusion models. One major category consists of **Sequential AR-Diffusion** models,

where an AR component generates an intermediate representation that conditions a diffusion-based decoder. In some variants, a pre-trained MLLM is kept frozen for reasoning, and its features are routed via learnable queries or hidden states to an external image generator [19, 34, 37]. This cascaded design effectively leverages powerful existing models. A more integrated approach uses a **Unified Transformer Backbone** [2, 54], where both AR and diffusion objectives are optimized simultaneously within a single transformer. To improve scalability, the **Mixture-of-Transformers (MoT)** paradigm has been introduced [6, 18]. This approach, exemplified by Bagel, employs a sparse, modular design where specialized experts handle different modalities but share information through a common attention mechanism. Our work, SRUM, inherits this flexible MoT framework, demonstrating the versatility of our method on advanced UMMs.

2.2 Post-Training Stage in UMMs

In addition to architectural innovations, considerable research has focused on the post-training stage to enhance the generative abilities of UMMs. Methods such as Chain-of-Thought (CoT) and test-time verification introduce explicit reasoning steps or iterative output validation [8, 9, 13]. However, these often depend on external models and do not fundamentally improve the native generative capacity of UMMs. Reinforcement learning techniques, including Direct Preference Optimization (DPO), leverage human or automated feedback to refine generation policies. While effective, these require carefully curated paired data and delicate advantage-function tuning with text-dependent rewards [12, 29, 31]. Reconstruction Alignment (RecA) introduces a post-training method based on reconstruction loss, yielding improved semantic understanding [46]. Some work has also attempted to use rule-level rewards for guidance, but such rewards are not universal and must be redesigned for different tasks [14, 15, 25]. In contrast, SRUM leverages the model’s inherent understanding to score self-generated samples, thereby enhancing performance without external reward models or human preference labels.

2.3 Self-Rewarding in Understanding Models

Self-rewarding mechanisms have emerged as a significant paradigm for enhancing the understanding and reasoning capabilities of MLLMs. These approaches aim to reduce reliance on external preference data by enabling models to generate their own reward signals, thereby facilitating continuous self-improvement. For instance, CSR [55] achieves zero-cost self-enhancement through iterative online DPO with visual constraint rewards. SRPO [5] introduces a two-stage reflective reward mechanism, significantly improving the quality of reflection and answer accuracy in complex reasoning tasks. R1-Reward leverages process consistency rewards and stable reinforcement learning algorithms to enhance long-range reasoning stability [12]. Collectively, these works signal a paradigm shift from external rewards to self-criticism and optimization. However, they tend to focus

on a single dimension of feedback. Our SRUM framework proposes a more holistic approach. It distinguishes itself by incorporating a global-local dual reward system designed to provide a more comprehensive training signal.

3 SRUM: Self-Rewarding for Unified Multimodal Models

To drive self-improvement where the model’s understanding capabilities guide its generation abilities, we establish a multi-stage self-rewarding process. First, the Unified Multimodal Models (UMMs) generate high-quality candidate images with corresponding bounding boxes (as detailed in Section 3.1). Next, these candidates are evaluated by a global-local judgment framework that assesses both overall composition and fine-grained details, producing a holistic reward signal (Section 3.2). Finally, the cached rewards directly inform a reward-weighted training process, which enables targeted, region-specific optimization and reduces reward hacking (Section 3.3).




Fig. 2: Showcase of the SRUM pipeline. It consists of two main stages: Self-Rewarding Data Generation and Reward-Weighted Training. The first stage generates high-quality data and scores it to produce a reward for the next training stage for self-improvement.

3.1 Image Candidates and Bounding Box Generation

We developed a self-data generation pipeline that enables our model to create its own training data, removing the need for external image resources. It starts with the UMM using its “think” mode (a form of CoT) to generate semantically rich images [6, 42]. For Bagel, an external segmentation model proposes spatial supports for grounding [17]; for BLIP3o, native grounding can be used instead. In both cases, the UMM verifies boxes, matches prompt-relevant regions, and assigns all semantic rewards, so the external component is only a localization aid rather than an external reward model.

3.2 Rewarding Process

Self-Judgment for Reliable Rewarding. A cornerstone of self-improvement is enabling the model’s internal understanding module to serve as a stable and re-

liable “evaluator”. To ensure the scores it generates are consistently trustworthy, we designed a comprehensive self-judgment mechanism to assess image quality and prompt alignment [11, 20, 49, 53]. This dual-level judgment is key to guaranteeing a thorough assessment. First, a local judgment evaluates object-level fidelity and artifacts on a strict $[-1.0, 1.0]$ scoring scale. A mandatory “Reason” field elicits an interpretable rationale for the score, akin to chain-of-thought prompting [9, 13], which further bolsters the reliability of the process. We enforce semantic grounding by verifying that identified objects correspond to prompt keywords, and a non-linear penalty maps severe distortions to a high-penalty negative range (e.g., -0.9 to -0.5) to better reflect human visual sensitivity. Subsequently, a global judgment evaluates the holistic composition and spatial alignment with the prompt’s intent. Crucially, for prompts lacking specific compositional directives (e.g., “a picture of a tree”), a neutral score range (e.g., -0.4 to 0.4) is applied, ensuring a fair assessment.

Reward Generation for Training. To serve as the core learning signal for self-improvement, the self-judgment scores are converted into a dense reward map. This step leverages the UMM’s grounding capabilities to produce fine-grained local rewards for prompt-relevant regions and a single global reward for the whole image. The global score is normalized to $[0, 1]$ to avoid spurious positive signals from multiplying two negative values. During reward generation, the understanding module is frozen and receives no gradients. SRUM is therefore offline: it generates candidates, caches rewards, and then updates the trainable generation/flow parameters.

3.3 Reward-Weighted Training

The reward-weighted training stage is where the model achieves self-improvement through training with rewards. The core objective is to translate the capabilities of the understanding module directly into the functionality of the generation module. By using fine-grained local rewards and layout-aware global rewards to weight the training objective, we guide the generator to learn more detailed and accurate patterns from the original data. This process is the key to bridging the gap between the model’s understanding and generation components, enabling the generator to benefit from the insights of the evaluator. The mechanism for this goal is a reward-weighted training objective, centered on the loss term $\mathcal{L}_r$. This term operates on the model’s velocity prediction v_θ , a standard practice in flow-based frameworks [21, 22]. The loss is modulated by two cached feedback signals from the frozen understanding module: a regional reward map $R \in [-1, 1]$ for localized refinement and a global scalar α for overall compositional quality. The product of these signals, $\alpha \cdot R$, weights the squared error between the predicted velocity v_θ and the target velocity derived from the original latent x_0^{gt} . This allows for fine-grained control, encouraging preservation where feedback is positive ($\alpha \cdot R > 0$) and promoting change where it is negative ($\alpha \cdot R < 0$):

$$\mathcal{L}_r = \mathbb{E} \left[\alpha \cdot R \odot (v_\theta - (x_0^{\text{gt}}))^2 \right] \quad (1)$$

Second, to ensure the model’s output conforms to the desired overall structure and to prevent reward hacking, we introduce a reference constraint term, $\mathcal{L}_{\text{ref}}$. This term acts as a regularizer, penalizing the squared ℓ_2 distance to the target velocity of the artifact-free latent x_0^{gt} :

$$\mathcal{L}_{\text{ref}} = \mathbb{E} \left[\|v_\theta - (\epsilon - x_0^{\text{gt}})\|^2 \right] \quad (2)$$

The final training objective is a weighted sum of these two losses, balanced by a hyperparameter λ_c . This is the complete objective used by SRUM; no hidden online reward loss or evaluator update is applied. The composite design enables targeted local refinement while maintaining global coherence:

$$\mathcal{L}_{\text{Total}} = \mathcal{L}_r + \lambda_c \cdot \mathcal{L}_{\text{ref}} \quad (3)$$

4 Analysis of Self-Rewarding

We validate our Self-Rewarding for Unified Multimodal Models (SRUM) method across various unified multimodal models (UMMs) and evaluation benchmarks. In particular, we investigate the following aspects:

- **Generality and Performance:** SRUM achieves better performance on composition and delivers consistent gains across frameworks. (Table 1)
- **Component Efficacy:** Ablation studies confirm each component of the SRUM framework makes a critical contribution to the overall performance. (Figure 3)
- **Generalization:** SRUM demonstrates in-domain and out-of-domain generalization, indicating that improvements in generation can be transferred from understanding. (Table 5, Figure 6, Table 6)

4.1 Experimental Setup

Model Architectures. We evaluate SRUM on two powerful open-source UMMs. All experiments are conducted as a post-training phase, starting from the official pre-trained weights. Bagel [6] is a versatile UMM that serves as our primary model for comprehensive analysis, including main results, ablation studies, and generalization tests. We evaluate both its standard and CoT inference modes. BLIP3o [3,30,43] is another current SOTA UMM used to validate the generality and effectiveness of SRUM with frozen MLLM training. We focus on these two families because some AR-type models, such as Show-O and Janus, can exhibit biases stemming from under-leveraged generation or understanding capabilities.

Datasets and Benchmarks. Our experiments leverage several specialized datasets for training and evaluation to ensure a thorough and multi-faceted analysis. For consistent and objective scoring across all generation benchmarks, we employ QwenVL-2.5-32B and QwenVL-2.5-72B [1] as the designated multimodal evaluators. It is crucial to clarify that these external models serve strictly as objective evaluation metrics (akin to accuracy in classification or mIoU in segmentation) and are completely excluded from our core self-rewarding training

pipeline. We also report traditional metrics following prior work [3]. Our experiment begins with instruction data sourced from the T2I-CompBench training set [16]. For our primary evaluation, we use the standard split of the same benchmark to compare SRUM-enhanced models against leading T2I and UMM baselines. To assess generalization, we evaluate the model’s in-domain transferability on GenEval [11], which includes similar compositional challenges, and WISE [26], which evaluates knowledge-informed generation. Furthermore, we evaluate broader out-of-domain reasoning-informed capabilities on T2I-Reason Bench [35], a benchmark containing complex prompts that require knowledge beyond the training distribution.

4.2 Main Results

Model	3D	Spat.	Color	Compl.	Nonspat.	Num.	Shape	Spatial	Text.	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	
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	
+SRUM	83.78	90.22	86.57	85.10	74.52	85.44	93.88	86.52	85.75	
Bagel	77.98	89.30	83.32	85.03	70.40	81.94	81.52	87.93	82.18	
+SRUM	83.10	92.90	88.69	88.47	78.52	84.23	86.92	89.57	86.55	
Bagel _(CoT)	84.66	88.85	86.10	85.64	75.36	84.33	82.71	88.07	84.46	
+SRUM	88.60	92.90	91.31	90.48	80.12	84.47	89.93	89.15	88.37	

Table 1: Comprehensive T2I-CompBench Results. This table includes T2I and UMMs models. Models incorporating the SRUM are denoted with +SRUM. **Bold** indicate the highest score in each column. Green values indicate the improvements.

As shown in Table 1, our proposed method, **SRUM**, achieves consistent and substantial performance gains across various compositional generation tasks. Specifically, when evaluated in CoT mode, **Bagel+SRUM** attains an overall score of 88.37, ranking first among current UMM baselines. This marks a significant improvement of +3.91 points over the baseline Bagel with CoT, demonstrating the efficacy of our approach. The advantages of SRUM are particularly pronounced in categories demanding spatial and complex reasoning as well as numeracy. For instance, our method sets new SOTA scores in **Spatial (93.88)**,

Model	Attribute Binding ($\uparrow$,%)			Object Relationship ($\uparrow$,%)		Complex ($\uparrow$,%)
	Color	Shape	Texture	Spatial	Non-Spat.	
Unified Multimodal Models						
Show-o	56.00	41.00	46.00	20.00	30.00	29.00
Janus-Pro-7B	63.59	35.28	49.36	20.61	30.85	35.59
Bagel	70.62	73.55	74.21	72.15	58.20	43.10
+SRUM	72.88	76.86	77.48	74.40	60.91	44.41
Bagel _(CoT)	80.95	83.92	84.35	71.80	57.55	43.65
+SRUM	82.14	85.17	89.47	74.03	59.96	44.51

Table 2: Detailed T2I-CompBench Evaluation via Traditional Metrics. This table compares various UMMs models using specialized traditional evaluators. Models incorporating the SRUM are denoted with +SRUM. **Bold** indicates the highest score in each column. All values are reported in percentage (%) (higher is better $\uparrow$).

3D Spatial (88.60) and **Complex (91.31)** reasoning, including 3D and action-based prompts. Although we observe a slight drop in texture and color categories in some cases, the overall trend remains positive, likely because our algorithm does not over-optimize low-level information for certain objects. Additionally, we report traditional metrics in Table 2 to provide a multi-angle comparison. While these evaluators confirm the consistent gains of **SRUM**, the improvements appear less pronounced than those in the LLM-based evaluation (Table 1). This discrepancy stems from the inherent limitations of traditional models (e.g., CLIP or BLIP), which often suffer from “bag-of-words” effects and lack fine-grained logical parsing. In contrast, our **Qwen-based evaluation** offers a more solid and nuanced assessment by leveraging superior multimodal reasoning, capturing complex spatial and attribute relationships that traditional metrics tend to overlook.

4.3 Empirical Study

We primarily employ three Bagel variants for analysis: **Base Model**, where Bagel’s open-source weights are used directly for inference; **SFT Model**, where Bagel generates images from training instructions and is then trained on this self-generated data; and **SRUM Model**, where the same self-generated data is trained with SRUM’s reward-weighted objective.

Computational Efficiency. A key advantage of SRUM is its practical efficiency. By parallelizing the understanding module, scoring a batch of 6K candidate images requires fewer than 4 GPU-hours on an NVIDIA H100. Furthermore, the loss computation introduces negligible overhead compared to standard SFT (12.5 GPU-hours for SRUM vs. 12.4 GPU-hours for SFT). Crucially, this self-rewarding mechanism eliminates the need for laborious manual data cleaning, making it highly scalable and practical given the current scarcity of high-quality image-text data.

Ablation Results. To further verify the effectiveness of our proposed reward configuration, we perform an ablation study on Bagel results on T2I-CompBench by systematically modifying the reward scheme. As shown in Figure 3 (Left), our full SRUM model achieves the highest overall accuracy, with the ablation results confirming the critical role of each component. Specifically, omitting the **local reward** leads to a performance drop of 0.76 in CoT mode and 1.04 in standard inference mode, confirming the necessity of fine-grained spatial feedback that bounding boxes provide. Omitting the **global reward** leads to a notable decrease in performance, underscoring its importance for capturing the overarching coherence and compositional structure of the generated images. Removing the **reference constraint** also results in a significant drop, proving its value in ensuring training stability. This aligns with conclusions from post-training methods like DPO [31], where a constraint is essential to prevent significant policy deviation due to reward hacking. Furthermore, using a simple **sparse reward** leads to substantial performance degradation, reinforcing the necessity of a continuous, dense reward signal for richer gradient information. This is particularly evident as sparse reward schemes, such as Dance-GRPO [50], are ill-suited for providing granular regional feedback, which highlights the value of our dense reward design. Overall, this ablation confirms that the efficacy of our framework stems from the synergistic contributions of each component.




Fig. 3: Left: Module Evaluation. We report the accuracy drop ($\Delta \text{Acc. (\%)}$) from our SRUM. Specifically, 0-1 Reward represents the sparse reward. **Right:** Hyperparameter evaluation on T2I-CompBench. We report the accuracy under different λ_c values in two modes: CoT and non-CoT.

In Figure 3 (Right), we analyze the effect of different reference-constraint ratios on the experimental outcomes. Across both Bagel with CoT and non-CoT configurations, the results consistently indicate that $\lambda_c = 0.5$ is the most effective choice. Consequently, we keep this hyperparameter fixed in subsequent experiments.

Comparison with External Rewards and Post-Training Methods. To further validate the necessity and superiority of the *self-rewarding* mechanism, we compare SRUM against popular post-training baselines including Reconstruction Alignment (ReCA) and the RL-based UnifiedReward, as well as

Method	UR	ReCA	Qwen	Qwen	Qwen	Qwen	SRUM (Ours)
Size	7B	-	7B	7B	32B	32B	7B
Config.	-	-	w/o GR	w/ GR	w/o GR	w/ GR	Self-Rewarding
Overall Score	85.15	86.45	85.78	87.52	85.53	88.01	88.37

Table 3: Comparison with external reward models and post-training methods on T2I-CompBench. UR and GR denote UnifiedReward and Global Reward, respectively. Our SRUM achieves the highest overall score.

external VLM reward models (QwenVL2.5-7B and 32B). As shown in Table 3, SRUM significantly outperforms ReCA and UnifiedReward. Counter-intuitively, replacing our internal understanding module with powerful external VLMs does not yield stable improvements. On saved image/prompt pairs, the internal scorer achieves stronger calibration with the final evaluator ($\rho = .58$ vs. $.43$) and better accept/reject separation (0.67 vs. -0.31), explaining why it can be a better training signal even when external VLMs are larger. This shows that SRUM’s effectiveness comes from the intrinsic multi-scale design rather than external parameter scale.

Further Analysis. For a more granular investigation, we leverage the same powerful MLLM, QwenVL-2.5-72B, from our primary evaluation to conduct a deeper analysis of our method and the baseline. Specifically, we employ the MLLM to perform a step-by-step scoring of the inference process.




Fig. 4: Score per step during inference in Bagel with its ablation models.

The evaluation is divided into two metrics: (1) layout, which assesses the concordance of the overall structure and quality, and (2) detail, which measures

the fidelity of the generated fine-grained details. Our ablation study, visualized in Figure 4, systematically isolates the effects of each component. We observe that the “think” mode primarily bolsters the initial layout generation by improving the high-level reasoning process. The global reward component of SRUM then further refines this layout during the early stages of inference. In contrast, a baseline using only this global reward (labeled “sample reward”) yields negligible improvements in detail fidelity. This highlights a crucial finding: the fine-grained, local rewards are essential for the subsequent optimization of details, with their benefits becoming most apparent in the later inference steps. Collectively, these results demonstrate that our dual global-local reward mechanism provides a multi-stage optimization path: first establishing a coherent layout and then progressively refining the details. This synergistic approach allows SRUM to significantly outperform standard SFT on the same self-generated data.

Impact on Understanding Module. As shown in Table 4, our method has a minimal impact on the model’s core understanding capabilities. On prevalent benchmarks such as MME [10], MM-Vet [51], MMBench [23], MMMU [52], and MathVista [24], the results exhibit only marginal fluctuations compared to the base version. Notably, performance on MMVP [38] even improves, which is consistent with prior works [37, 39, 40]. This indicates that our method holds significant potential for further iterative enhancement. In Figure 5, we track the activation dynamics of two distinct functional clusters, Understanding and Generation, across the Base, SFT, and SRUM models. In Bagel’s inference, the parameters primarily activated during understanding tasks are defined as the understanding cluster, and those primarily activated during generation tasks are defined as the generation cluster. Our analysis reveals two distinct finetuning paradigms. Conventional SFT exhibits a narrowing effect, achieving specialization by suppressing irrelevant functional clusters. In contrast, SRUM demonstrates an enhancing and orchestrating effect, strengthening the primary task-relevant cluster while maintaining supportive activation in secondary clusters. This promotes robust and generalizable representations.




Fig. 5: Functional cluster activation patterns of the different models (Bagel, SFT and SRUM) on different understanding and generation tasks. The average activation strength of Understanding and Generation clusters is shown.

	Base	SFT	SRUM
MME-P	1687	1682	1673
MME-C	701	683	677
MMBench	85.0	84.6	84.8
MM-Vet	67.2	66.5	67.0
MMMU	55.3	55.0	55.2
MathVista	73.1	72.8	73.0
MMVP	69.3	68.7	70.0

Table 4: Comparison with the results of different models (Bagel, SFT and SRUM) on understanding benchmarks. MME-P and MME-C represent the perception and cognition parts, respectively.

In-Domain Generalization. We then investigate the in-domain generalization capability of our model. We posit that the compositional abilities learned from the T2I-CompBench training set should be transferable to other benchmarks with similar evaluation perspectives. To test this hypothesis, we evaluate SRUM on the GenEval benchmark without any further fine-tuning. The comparative results are summarized in Table 5.

As shown in the table, our evaluation on GenEval further validates the strengths of SRUM, particularly in the challenging domain of object counting. SRUM attains the highest score of 0.83 in Counting, surpassing both the base model and the SFT baseline. Crucially, this superior performance in numerical generation aligns with our previous results on T2I-CompBench. This consistency across benchmarks underscores our method’s reliable improvement in processing quantitative information. By excelling at a complex task like counting while retaining proficiency in simpler ones, the model demonstrates strong in-domain generalization. This confirms that the targeted enhancements by SRUM are transferable.

Model	Single obj.	Two obj.	Counting	Colors	Position	Color attr.
Bagel	0.99	0.94	0.81	0.88	0.64	0.82
Bagel+SFT	0.96	0.94	0.79	0.92	0.59	0.78
Bagel+SRUM	0.98	0.94	0.83	0.90	0.64	0.83

Table 5: Results on key visual attributes at GenEval. For brevity, some model names are shortened; abbreviations are defined at the beginning of Section 4.3. **Bold values** are the best in each column.

In-Domain Knowledge-based Generalization. Following this, we explore whether our method holds a distinct advantage for reasoning-based generation, a current area of focus in the community. We train the model on one category of prompts from the WISE benchmark and perform in-domain evaluations on the remaining two categories. This protocol yields three distinct evaluation sets for analyzing the model’s generalization capabilities.

As illustrated in Figure 6, selecting any single group for training generally enhances the image generation performance of the other two groups. This improvement is consistent across both standard and CoT reasoning paradigms. It shows that SRUM can promote knowledge-domain generalization, enabling generation to better fit instruction semantics at the knowledge level.

Out-of-Domain Generalization. We directly evaluate the model trained on T2I-CompBench prompts on **T2I-ReasonBench**, which tests reasoning-informed image generation. We report accuracy for high-level semantic alignment and use Bagel without CoT to avoid self-rewrite effects that could leak hidden meanings in categories such as **Idiom**.




Fig. 6: Results of Bagel on WISE. We train on one of the three WISE tasks and evaluate on the other two, showing SRUM’s knowledge-domain generalization.

As shown in Table 6, SRUM outperforms both Base and SFT, with steadier gains on image-based prompts, demonstrating stronger adaptation to unseen reasoning tasks.

Model	Entity	Idiom	Scientific	Textual	Overall
Bagel	49.70	34.46	47.52	43.59	43.82
Bagel _{+SFT}	50.53	39.43	47.45	44.08	45.37
Bagel _{+SRUM}	52.85	40.51	47.83	45.83	46.75

Table 6: Performance comparison on T2I-ReasonBench. **Bold values** indicate the best performance in each column.

5 Conclusion

This paper introduces SRUM, a fine-grained post-training framework that lets a model’s understanding module reward its generation module. By combining local object-level feedback with global compositional assessment, SRUM transfers visual understanding into generation without external reward models or human-labeled preference data. Experiments on compositional and reasoning benchmarks such as T2I-CompBench and T2I-ReasonBench show consistent gains, robust in-domain and out-of-domain generalization, and the value of dense multi-scale self-rewarding.

SRUM is an initial step toward self-rewarding UMMs. Future work will improve scorer prompts, scale to larger datasets, and explore self-play questions and answers for a more closed-loop training system where understanding and generation improve together.

Acknowledgements

This work is supported by the Shenzhen Loop Area Institute under grant FPF 10120250006, and partially supported under grant FPF10120250002.

References

1. Bai, S., Chen, K., Liu, X., Wang, J., Ge, W., Song, S., Dang, K., Wang, P., Wang, S., Tang, J., Zhong, H., Zhu, Y., Yang, M., Li, Z., Wan, J., Wang, P., Ding, W., Fu, Z., Xu, Y., Ye, J., Zhang, X., Xie, T., Cheng, Z., Zhang, H., Yang, Z., Xu, H., Lin, J.: Qwen2.5-vl technical report. arXiv preprint arXiv:2502.13923 (2025)
2. Chen, B., Martí Monsó, D., Du, Y., Simchowitz, M., Tedrake, R., Sitzmann, V.: Diffusion forcing: Next-token prediction meets full-sequence diffusion. In: NeurIPS (2024)
3. Chen, J., Xu, Z., Pan, X., Hu, Y., Qin, C., Goldstein, T., Huang, L., Zhou, T., Xie, S., Savarese, S., et al.: Blip3-o: A family of fully open unified multimodal models-architecture, training and dataset. arXiv preprint arXiv:2505.09568 (2025)
4. Chen, X., Wu, C., Wu, Z., Ma, Y., Liu, X., Pan, Z., Liu, W., Xie, Z., Yu, X., Ruan, C., Luo, P.: Janus-pro: Unified multimodal understanding and generation with data and model scaling. arXiv preprint arXiv:2501.17811 (2025)
5. Choi, E., Ahmadian, A., Geist, M., Pietquin, O., Azar, M.G.: Self-improving robust preference optimization. arXiv preprint arXiv:2406.01660 (2024)
6. Deng, C., Zhu, D., Li, K., Gou, C., Li, F., Wang, Z., Zhong, S., Yu, W., Nie, X., Song, Z., et al.: Emerging properties in unified multimodal pretraining. arXiv preprint arXiv:2505.14683 (2025)
7. Dong, R., Han, C., Peng, Y., Qi, Z., Ge, Z., Yang, J., Zhao, L., Sun, J., Zhou, H., Wei, H., et al.: Dreamllm: Synergistic multimodal comprehension and creation. In: ICLR (2024)
8. Duan, C., Fang, R., Wang, Y., Wang, K., Huang, L., Zeng, X., Li, H., Liu, X.: Got-r1: Unleashing reasoning capability of mllm for visual generation with reinforcement learning. arXiv preprint arXiv:2505.17022 (2025)
9. Fang, R., Duan, C., Wang, K., Huang, L., Li, H., Yan, S., Tian, H., Zeng, X., Zhao, R., Dai, J., et al.: Got: Unleashing reasoning capability of multimodal large language model for visual generation and editing. arXiv preprint arXiv:2503.10639 (2025)
10. Fu, C., Chen, P., Shen, Y., Qin, Y., Zhang, M., Lin, X., Yang, J., Zheng, X., Li, K., Sun, X., et al.: Mme: A comprehensive evaluation benchmark for multimodal large language models. arXiv preprint arXiv:2306.13394 (2023)
11. Ghosh, D., Hajishirzi, H., Schmidt, L.: Geneval: An object-focused framework for evaluating text-to-image alignment. In: NeurIPS (2023)
12. Guo, D., Yang, D., Zhang, H., Song, J., Zhang, R., Xu, R., Zhu, Q., Ma, S., Wang, P., Bi, X., et al.: Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. arXiv preprint arXiv:2501.12948 (2025)
13. Guo, Z., Zhang, R., Tong, C., Zhao, Z., Huang, R., Zhang, H., Zhang, M., Liu, J., Zhang, S., Gao, P., et al.: Can we generate images with cot? let's verify and reinforce image generation step by step. arXiv preprint arXiv:2501.13926 (2025)
14. Han, Y., Chen, H., Han, A., Wang, Z., Lin, X., Zhang, Y., Zhang, S., Zou, D.: Self-contradiction as self-improvement: Mitigating the generation-understanding gap in mllms. arXiv preprint arXiv:2507.16663 (2025)

15. Hong, J., Zhang, Y., Wang, G., Liu, Y., Wen, J.R., Yan, R.: Reinforcing multimodal understanding and generation with dual self-rewards. *arXiv preprint arXiv:2506.07963* (2025)
16. Huang, K., Sun, K., Xie, E., Li, Z., Liu, X.: T2i-compbench: A comprehensive benchmark for open-world compositional text-to-image generation. *Advances in Neural Information Processing Systems* **36**, 78723–78747 (2023)
17. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., et al.: Segment anything. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 4015–4026 (2023)
18. Liang, W., YU, L., Luo, L., Iyer, S., Dong, N., Zhou, C., Ghosh, G., Lewis, M., tau Yih, W., Zettlemoyer, L., Lin, X.V.: Mixture-of-transformers: A sparse and scalable architecture for multi-modal foundation models. *TMLR* (2025)
19. Lin, B., Li, Z., Cheng, X., Niu, Y., Ye, Y., He, X., Yuan, S., Yu, W., Wang, S., Ge, Y., et al.: Uniworld: High-resolution semantic encoders for unified visual understanding and generation. *arXiv preprint arXiv:2506.03147* (2025)
20. Lin, Z., Pathak, D., Li, B., Li, J., Xia, X., Neubig, G., Zhang, P., Ramanan, D.: Evaluating text-to-visual generation with image-to-text generation. In: *European Conference on Computer Vision*. pp. 366–384. Springer (2024)
21. Lipman, Y., Chen, R.T., Ben-Hamu, H., Nickel, M., Le, M.: Flow matching for generative modeling. In: *ICLR* (2023)
22. Liu, X., Gong, C., et al.: Flow straight and fast: Learning to generate and transfer data with rectified flow. In: *ICLR* (2023)
23. Liu, Y., Duan, H., Zhang, Y., Li, B., Zhang, S., Zhao, W., Yuan, Y., Wang, J., He, C., Liu, Z., et al.: Mmbench: Is your multi-modal model an all-around player? In: *ECCV* (2024)
24. Lu, P., Bansal, H., Xia, T., Liu, J., Li, C., Hajishirzi, H., Cheng, H., Chang, K.W., Galley, M., Gao, J.: Mathvista: Evaluating mathematical reasoning of foundation models in visual contexts. In: *NeurIPS Workshop on Mathematical Reasoning and AI* (2023)
25. Mao, W., Yang, Z., Shou, M.Z.: Unirl: Self-improving unified multimodal models via supervised and reinforcement learning. *arXiv preprint arXiv:2505.23380* (2025)
26. Niu, Y., Ning, M., Zheng, M., Lin, B., Jin, P., Liao, J., Ning, K., Zhu, B., Yuan, L.: Wise: A world knowledge-informed semantic evaluation for text-to-image generation. *arXiv preprint arXiv:2503.07265* (2025)
27. Pan, X., Shukla, S.N., Singh, A., Zhao, Z., Mishra, S.K., Wang, J., Xu, Z., Chen, J., Li, K., Juefei-Xu, F., Hou, J., Xie, S.: Transfer between modalities with meta-queries. *arXiv preprint arXiv:2504.06256* (2025)
28. Podell, D., English, Z., Lacey, K., Blattmann, A., Dockhorn, T., Müller, J., Penna, J., Rombach, R.: Sdxl: Improving latent diffusion models for high-resolution image synthesis. In: *ICLR* (2024)
29. Qu, L., Li, H., Wang, W., Liu, X., Li, J., Nie, L., Chua, T.S.: Silmm: Self-improving large multimodal models for compositional text-to-image generation. In: *Proceedings of the Computer Vision and Pattern Recognition Conference*. pp. 18497–18508 (2025)
30. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: *International conference on machine learning*. pp. 8748–8763. PmLR (2021)
31. Rafailov, R., Sharma, A., Mitchell, E., Manning, C.D., Ermon, S., Finn, C.: Direct preference optimization: Your language model is secretly a reward model. *Advances in neural information processing systems* **36**, 53728–53741 (2023)

32. Ramesh, A., Pavlov, M., Goh, G., Gray, S., Voss, C., Radford, A., Chen, M., Sutskever, I.: Zero-shot text-to-image generation. In: ICML (2021)
33. Saharia, C., Chan, W., Saxena, S., Li, L., Whang, J., Denton, E.L., Ghasemipour, K., Gontijo Lopes, R., Karagol Ayan, B., Salimans, T., et al.: Photorealistic text-to-image diffusion models with deep language understanding. In: NeurIPS (2022)
34. Shi, W., Han, X., Zhou, C., Liang, W., Lin, X.V., Zettlemoyer, L., Yu, L.: Llamafusion: Adapting pretrained language models for multimodal generation. arXiv preprint arXiv:2412.15188 (2024)
35. Sun, K., Fang, R., Duan, C., Liu, X., Liu, X.: T2i-reasonbench: Benchmarking reasoning-informed text-to-image generation. arXiv preprint arXiv:2508.17472 (2025)
36. Team, C.: Chameleon: Mixed-modal early-fusion foundation models. arXiv preprint arXiv:2405.09818 (2024)
37. Tong, S., Fan, D., Zhu, J., Xiong, Y., Chen, X., Sinha, K., Rabbat, M., LeCun, Y., Xie, S., Liu, Z.: Metamorph: Multimodal understanding and generation via instruction tuning. arXiv preprint arXiv:2412.14164 (2024)
38. Tong, S., Liu, Z., Zhai, Y., Ma, Y., LeCun, Y., Xie, S.: Eyes wide shut? exploring the visual shortcomings of multimodal llms. In: CVPR. pp. 9568–9578 (2024)
39. Wang, H., Zheng, A., Zhao, Y., Wang, T., Ge, Z., Zhang, X., Zhang, Z.: Reconstructive visual instruction tuning. arXiv preprint arXiv:2410.09575 (2024)
40. Wang, W., Sun, Q., Zhang, F., Tang, Y., Liu, J., Wang, X.: Diffusion feedback helps clip see better. arXiv preprint arXiv:2407.20171 (2024)
41. Wang, X., Zhang, X., Luo, Z., Sun, Q., Cui, Y., Wang, J., Zhang, F., Wang, Y., Li, Z., Yu, Q., et al.: Emu3: Next-token prediction is all you need. arXiv preprint arxiv:2409.18869 (2024)
42. Wang, Y., Liu, M., He, W., Zhang, L., Huang, Z., Zhang, G., Shu, F., Tao, Z., She, D., Yu, Z., et al.: Mint: Multi-modal chain of thought in unified generative models for enhanced image generation. arXiv preprint arXiv:2503.01298 (2025)
43. Wang, Z., Li, Y., Chen, X., Lim, S.N., Torralba, A., Zhao, H., Wang, S.: Detecting everything in the open world: Towards universal object detection (2023)
44. Wu, C., Chen, X., Wu, Z., Ma, Y., Liu, X., Pan, Z., Liu, W., Xie, Z., Yu, X., Ruan, C., Luo, P.: Janus: Decoupling visual encoding for unified multimodal understanding and generation. arXiv preprint arXiv:2410.13848 (2024)
45. Wu, Y., Zhang, Z., Chen, J., Tang, H., Li, D., Fang, Y., Zhu, L., Xie, E., Yin, H., Yi, L., Han, S., Lu, Y.: Vila-u: A unified foundation model integrating visual understanding and generation. arXiv preprint arXiv:2409.04429 (2024), <https://arxiv.org/abs/2409.04429>
46. Xie, J., Darrell, T., Zettlemoyer, L., Wang, X.: Reconstruction alignment improves unified multimodal models (2025), <https://arxiv.org/abs/2509.07295>
47. Xie, J., Mao, W., Bai, Z., Zhang, D.J., Wang, W., Lin, K.Q., Gu, Y., Chen, Z., Yang, Z., Shou, M.Z.: Show-o: One single transformer to unify multimodal understanding and generation. arXiv preprint arxiv:2408.12528 (2024)
48. Xie, J., Yang, Z., Shou, M.Z.: Show-o2: Improved native unified multimodal models. arXiv preprint arXiv:2506.15564 (2025)
49. Xu, J., Liu, X., Wu, Y., Tong, Y., Li, Q., Ding, M., Tang, J., Dong, Y.: Imagereward: Learning and evaluating human preferences for text-to-image generation. Advances in Neural Information Processing Systems **36**, 15903–15935 (2023)
50. Xue, Z., Wu, J., Gao, Y., Kong, F., Zhu, L., Chen, M., Liu, Z., Liu, W., Guo, Q., Huang, W., et al.: Dancegrpo: Unleashing grpo on visual generation. arXiv preprint arXiv:2505.07818 (2025)

51. Yu, W., Yang, Z., Li, L., Wang, J., Lin, K., Liu, Z., Wang, X., Wang, L.: Mm-vet: Evaluating large multimodal models for integrated capabilities. In: ICML (2024)
52. Yue, X., Ni, Y., Zhang, K., Zheng, T., Liu, R., Zhang, G., Stevens, S., Jiang, D., Ren, W., Sun, Y., et al.: Mmmu: A massive multi-discipline multimodal understanding and reasoning benchmark for expert agi. In: CVPR (2024)
53. Zhang, W., Zhai, G., Wei, Y., Yang, X., Ma, K.: Blind image quality assessment via vision-language correspondence: A multitask learning perspective. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 14071–14081 (2023)
54. Zhao, C., Song, Y., Wang, W., Feng, H., Ding, E., Sun, Y., Xiao, X., Wang, J.: Monoformer: One transformer for both diffusion and autoregression. arXiv preprint arXiv:2409.16280 (2024), <https://arxiv.org/abs/2409.16280>
55. Zhou, Y., Fan, Z., Cheng, D., Yang, S., Chen, Z., Cui, C., Wang, X., Li, Y., Zhang, L., Yao, H.: Calibrated self-rewarding vision language models. *Advances in Neural Information Processing Systems* **37**, 51503–51531 (2024)