

UniAlignment: Semantic Alignment for Unified Image Generation, Understanding, Manipulation and Perception

Xinyang Song^{1, 2*}, Libin Wang³, Weining Wang^{2†}, Shaozhen Liu^{1, 2}, Dandan Zheng³, Jingdong Chen³, Qi Li^{1, 2}, Zhenan Sun^{1, 2}

¹School of Artificial Intelligence, University of Chinese Academy of Sciences

²Institute of Automation, Chinese Academy of Sciences

³AntGroup

{xinyang.song, shaozhen.liu}@cripac.ia.ac.cn, {libin.wlb, yuandan.zdd, jingdongchen.cjd}@antgroup.com, {weining.wang, qli, znsun}@nlpr.ia.ac.cn

Abstract

The remarkable success of diffusion models in text-to-image generation has sparked growing interest in expanding their capabilities to a variety of multi-modal tasks, including image understanding, manipulation, and perception. These tasks require advanced semantic comprehension across both visual and textual modalities, especially in scenarios involving complex semantic instructions. However, existing approaches often rely heavily on vision-language models (VLMs) or modular designs for semantic guidance, leading to fragmented architectures and computational inefficiency. To address these challenges, we propose UniAlignment, a unified multimodal generation framework within a single diffusion transformer. UniAlignment introduces a dual-stream diffusion training strategy that incorporates both intrinsic-modal semantic alignment and cross-modal semantic alignment, thereby enhancing the model’s cross-modal consistency and instruction-following robustness. Additionally, we present SemGen-Bench, a new benchmark specifically designed to evaluate multimodal semantic consistency under complex textual instructions. Extensive experiments across multiple tasks and benchmarks demonstrate that UniAlignment outperforms existing baselines, underscoring the significant potential of diffusion models in unified multimodal generation.

Introduction

The emergence of unified models for image generation and understanding (Team 2024; Xie et al. 2024a; Chen et al. 2025c; Pan et al. 2025) reflects the growing demand for general-purpose generative frameworks in multimodal learning. Recent models, such as GPT-4o-Image (OpenAI 2025) and BLIP-3o (Chen et al. 2025b) have demonstrated impressive capabilities in both visual comprehension and generation, suggesting the feasibility of integrating diverse vision-language tasks within a single architecture.

Despite these advances, achieving a fully unified multimodal framework remains challenging. AR-based unified

*Work done during internship at Ant Group.

†Corresponding Author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

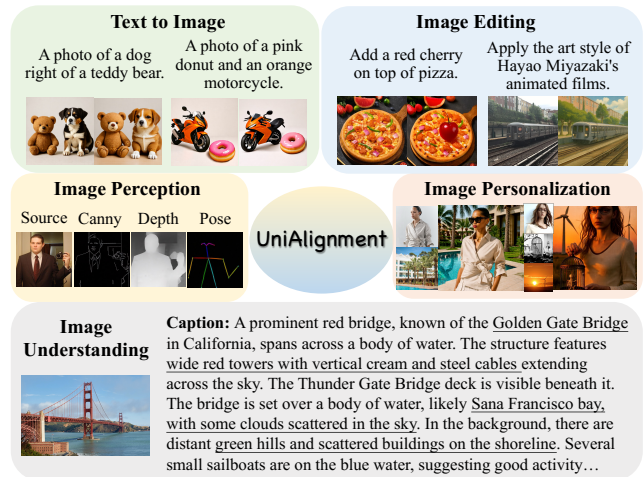


Figure 1: **Showcase of UniAlignment’s capabilities.** Our approach enables a single lightweight DiT to handle diverse multimodal tasks, including image understanding, generation, editing, perception, and personalization, achieving versatile capabilities within a unified framework.

models (Wu et al. 2025a; Xie et al. 2024a; Zhou et al. 2024), built upon transformer architectures adapted from large language models (Bai et al. 2023; Touvron et al. 2023), exhibit strong compositional reasoning and instruction-following capabilities. However, these models typically depend on discrete token-based decoding, which inherently constrained their ability to capture fine-grained visual fidelity and continuous pixel-level manipulation necessary for high-quality image synthesis and editing. In contrast, Diffusion-based models (Rombach et al. 2022; Peebles and Xie 2023; Esser et al. 2024; Black Forest Labs 2024) excel in generating high-fidelity images through iterative denoising, but often lack robust semantic perception.

Several methods employing hybrid architectures, such as MetaQueries (Pan et al. 2025), Step1X-Edit (Liu et al. 2025), UniWorld-V1 (Lin et al. 2025), and OmniGen2 (Wu et al. 2025b), which enhance diffusion with VLM guid-

ance but remain structurally fragmented, involving modular designs and specialized connectors, thus diminishing their generalizability and scalability. These approaches fail to construct a unified model architecture, resulting in redundant parameters that degrade both training and inference efficiency. Recent efforts (Li et al. 2025; Swerdlow et al. 2025; Yang et al. 2025) attempt to construct a unified model purely based on diffusion models. However, DualDiffusion (Li et al. 2025) despite employing a unified diffusion transformer, suffers from task interference and weak semantic grounding. UniDisc (Swerdlow et al. 2025) introduces a purely discrete diffusion approach, which enhances token-level alignment but limits pixel-level precision essential for fine-grained editing. MMaDA (Yang et al. 2025) integrates reinforcement learning to align reasoning process, but depends on separate encoders, which increases complexity architectural and undermines end-to-end optimization.

To address the aforementioned challenges, we propose UniAlignment, a unified multimodal generative framework based entirely on a single shared diffusion transformer architecture, designed to harmonize image understanding, generation, editing, and perception tasks. The model adopts a dual-stream diffusion design, where a shared DiT backbone jointly parameterizes a continuous diffusion branch for image synthesis and a discrete diffusion branch for text-based understanding. This design enables a symmetric and scalable modeling of both modalities within the same generative space. By avoiding reliance on external vision-language models (VLMs) and modular components during inference, UniAlignment maintains high computational efficiency and structural simplicity. To further enhance multimodal coherence, UniAlignment introduces two lightweight yet highly effective semantic alignment strategies. The first is cross-modal semantic alignment between text-to-image and image-to-text branches, incorporating contrastive learning into dual stream training. The second is intrinsic-modal semantic alignment via intermediate feature matching, enriching the latent semantic during denoising process. As shown in Fig. 2, these mechanisms significantly improve semantic consistency of instruction following and robustness in vision-language semantic grounding.

To enable collaborative optimization across multiple tasks, we adopt a progressive multi-stage training strategy on large-scale datasets to achieve balanced performance. Furthermore, given the limitations of existing benchmarks, which typically adopt simple semantics and rigid instruction formats, we construct a new benchmark, SemGen-Bench, specifically designed to evaluate models’ semantic fidelity and multimodal alignment under complex, compositional instructions. Extensive experiments demonstrate that UniAlignment achieves remarkable performance across diverse multimodal tasks, outperforming existing unified models in both general scenarios and challenging generation and editing tasks. Our contributions can be summarized as follows:

- We propose UniAlignment, a unified multimodal generative model based on a single Diffusion Transformer, demonstrating outstanding performance while maintain-

ing lightweight design and computational efficiency.

- We introduce two complementary semantic alignment mechanisms that significantly enhances image-text semantic consistency and instruction-following robustness.
- We construct a rigorous new benchmark SemGen-Bench for evaluating multimodal semantic alignment under complex, compositional instructions, establishing a high-standard baseline for future research.

Related Works

Semantic Representation Alignment

Semantic representation alignment plays a pivotal role in vision-language representation learning, significantly influencing downstream tasks such as visual understanding and generation. Early works (Radford et al. 2021; Sun et al. 2023; Zhai et al. 2023) leverage large-scale image-text pairs and contrastive learning to construct multimodal pre-training frameworks that align visual and textual modalities. Other approaches (Li et al. 2022, 2023) introduce Q-former (Zhang et al. 2024a) to bridge vision encoders and language models, achieving parameter-efficient multimodal adaptation. Building on these foundations, several methods focus on optimizing generative models through semantic representation alignment. (Hudson et al. 2024; Wang et al. 2025b; Ma et al. 2025) utilize diffusion feedback to refine CLIP representations, aligning pretrained VLMs with diffusion-based generation through self-supervised reconstruction. Others aim to enhance perceptual quality by aligning internal representations between diffusion and vision models. Early works (Xiang et al. 2023; Wei et al. 2023; Tian et al. 2024) explore unified visual representation learning via self-supervised diffusion training. VAVAE (Yao, Yang, and Wang 2025) aligns high-dimensional visual tokens with backbone encoder features to push the limits of reconstruction quality, while REPA (Yu et al. 2025b) introduces intermediate-layer alignment between DiTs and pretrained vision encoders (Caron et al. 2021), improving convergence and generation fidelity. Extending this idea, Soft-REPA (Lee et al. 2025) employs contrastive learning with learnable soft tokens to enhance text-image consistency. In this work, we propose two semantic alignment strategies, jointly optimizing semantic consistency and visual-language grounding within a unified generative framework.

Unified Multi-Modal Generation Models

Recent breakthroughs in vision-language models (Google 2025; OpenAI 2025; Chen et al. 2025b) have driven a new wave of innovation in multimodal large models, accelerating the pursuit of unified frameworks for general-purpose generation. Early explorations (Wang et al. 2024; Team 2024; Xie et al. 2024a; Wu et al. 2025a) focus on augmenting LLMs with visual encoders, using autoregressive decoding over pixel-level tokens to enable vision synthesis. Follow-up works (Wu et al. 2025c; Qu et al. 2025; Xie et al. 2024b) investigate the impact of discrete or continuous visual tokenization on unified vision-language modeling, still grounded within autoregressive transformer architectures. In contrast, methods (Li et al. 2025; Swerdlow et al. 2025;

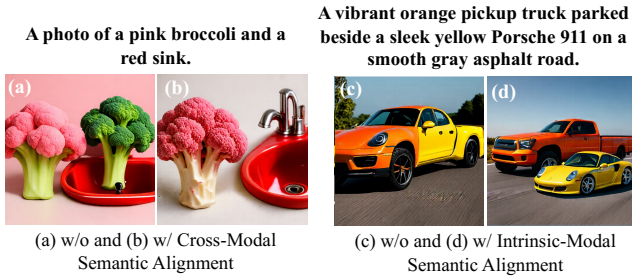


Figure 2: **Analysis of the proposed semantic representation alignment.** (a–d) present the generated images with and without the two semantic alignment mechanisms.

Yang et al. 2025) attempt to unify image generation and understanding via diffusion models, proposing dual-stream diffusion frameworks that jointly model denoising processes in discrete or continuous latent spaces.

However, achieving general-purpose multimodal generative models requires going beyond text-to-image synthesis. A line of recent work (Brooks, Holynski, and Efros 2023; Zhang et al. 2023; Yu et al. 2025a; Shi, Wang, and Huang 2024; Wei et al. 2024; Wang et al. 2025a; Liu et al. 2025; Mao et al. 2025; Xiao et al. 2025; Zhang et al. 2025; Pan et al. 2025; AI et al. 2025; Huang et al. 2025; Lin et al. 2025) aim to build multitask, generalizable visual generation frameworks. Among them, (Pan et al. 2025; AI et al. 2025; Huang et al. 2025; Lin et al. 2025) adopt a hybrid architecture, using connectors to bridge VLMs and diffusion backbones, thereby combining semantic understanding with pixel-level synthesis. Although representing a step towards integrating visual understanding, generation, and manipulation, these methods remain structurally fragmented, resulting in redundant computation and a lack of architectural unification.

Method

In this section, we present UniAlignment, a unified framework for multimodal generation tasks. We begin by introducing the overall architecture, which jointly addresses both image generation and understanding within a single diffusion-based backbone. To meet the semantic alignment demands posed by various vision-language generation scenarios, we propose two lightweight yet effective semantic alignment strategies: cross-modal semantic alignment and intrinsic-modal semantic alignment, implicitly enhancing the semantic grounding capability of diffusion models without additional parameters or learnable queries. To further support balanced learning across diverse generation tasks, a multi-stage training strategy is introduced to improve multi-task performance by progressively optimization.

Unified Generation & Understanding Framework

Autoregressive models have demonstrated limited efficacy in visual generation tasks, struggling with fine-grained spatial coherence and semantic fidelity. In contrast, diffu-

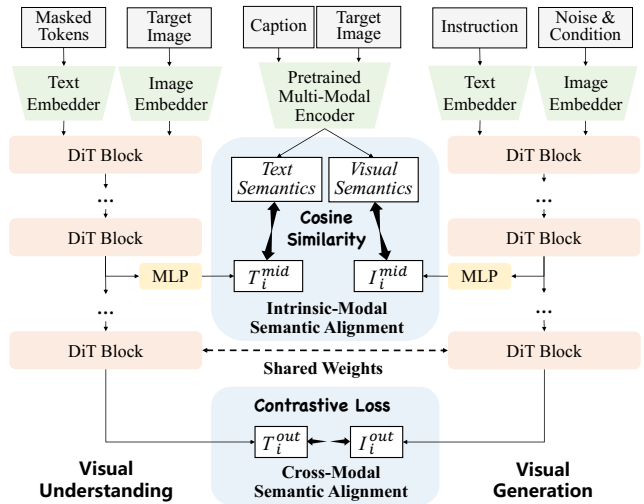


Figure 3: **The training pipeline of UniAlignment.** Text diffusion (left branch) and image diffusion (right branch) are both unified within a single Transformer with shared weights. The central part illustrates the proposed cross-modal semantic alignment and intrinsic-modal semantic alignment. The image condition is set absent during T2I generation. Only parameters within the DiT blocks and the MLP heads are optimized throughout the entire training.

sion models have shown remarkable potential for general-purpose generation. Motivated by this, we leverage a single diffusion transformer (DiT) to unify both visual generation and visual understanding, without relying on VLMs or external visual encoders during inference.

As illustrated in Fig. 3, UniAlignment adopts a dual-stream diffusion architecture, parameterizing both continuous and discrete diffusion processes within a shared DiT backbone. Inspired by DualDiffusion (Li et al. 2025), this architecture simultaneously models image distribution $p(\mathbf{x}_{\text{img}})$ and text distribution $p(\mathbf{x}_{\text{txt}})$. The continuous diffusion branch targets image generation following the flow matching principle (Lipman et al. 2022), with its training objective defined as follows:

$$\mathcal{L}_{\text{img}} = \mathbb{E}_{t, q_{\text{img}}} \left\| \mathbf{v}_{\theta}(\mathbf{x}_{\text{img}}^t, t, \mathbf{x}_{\text{txt}}) - (\epsilon - \mathbf{x}_{\text{img}}) \right\|_2^2, \quad (1)$$

where $\mathbf{x}_{\text{img}}^t$ denotes the noisy image at timestep t , \mathbf{v}_{θ} is the predicted velocity field, and ϵ represents the noise.

The discrete diffusion branch, masked token prediction is adopted to model the reverse process based on the masked state of discrete tokens. Specifically, the model learns to iteratively denoise masked tokens, thereby approximating the underlying data distribution in the discrete space:

$$\mathcal{L}_{\text{txt}} = \mathbb{E}_{q_{\text{txt}}} \left[-\frac{1}{K} \sum_{i=1}^K \log [\mathbf{x}_{\theta}(\mathbf{x}_{\text{txt}}^{t_i}, \mathbf{x}_{\text{img}}) \cdot \mathbf{x}] / t_i \right], \quad (2)$$

where timestep t_i is determined based on the text sequence length. The understanding and generation branches share the same parameters throughout training, allowing both modalities to be jointly optimized within a unified architecture.

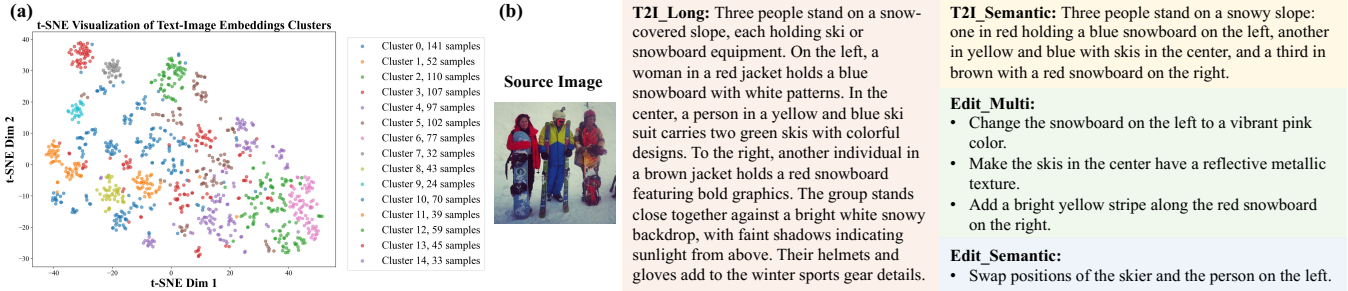


Figure 4: **The proposed SemGen-Bench.** (a) Visualization of clustered embeddings of image data. (b) Samples from the proposed SemGen-Bench.

To support general-purpose multimodal generation, the continuous diffusion branch extends beyond text-to-image synthesis to include a broad range of conditional image generation tasks, including image manipulation (e.g., instruction-based editing, subject customization, stylization) and image perception (e.g., depth estimation, edge detection, and human pose prediction). The source images are first encoded via a pretrained VAE to obtain conditional tokens, which are then concatenated with the noisy latents and fed into the DiT blocks. During training, the continuous diffusion branch receives both the conditional image and the textual instruction as input, while the discrete diffusion branch is conditioned on the target image and its corresponding caption. This design enables the incorporation of visual conditions without architectural modifications, additional encoders, or query tokens. The unified DiT framework of UniAlignment allows seamless extension to a variety of multimodal generation tasks, maintaining lightweight design and high efficiency, avoiding the semantic gap between autoregressive and diffusion-based paradigms.

Semantic Representation Alignment

Cross-modal Semantic Alignment. Although jointly training bi-directional diffusion processes within a single transformer allows for synchronized learning of generation and understanding, the inherent differences in their optimization objectives can lead to conflicting gradients and suboptimal convergence. As illustrated in Fig. 2(a-b), these conflicts often degrade instruction-following capabilities, reducing the model’s text-image consistency. This highlights the challenge of achieving balanced optimization in unified multimodal frameworks.

To alleviate the conflict arising from bi-directional diffusion training, we propose a solution that adjusts the direction of joint optimization. Our key insight is to leverage the benefits of joint vision-language training to improve text-image consistency through representation alignment within the transformer. As contrastive learning has proven effective in aligning multimodal representations in prior works (Radford et al. 2021; Li et al. 2022; Jia et al. 2021), we introduce a contrastive objective over the output embeddings from the dual diffusion branches. This objective encourages semantic similarity between matched image-text pairs while suppressing that of unmatched pairs. Specifically, given a batch

of N image-text pairs, we optimize the contrastive loss over the corresponding embeddings to strengthen cross-modal semantic alignment as follows:

$$\mathcal{L}_{\text{Cross}} = -\frac{1}{N} \sum_{i=1}^N \log \frac{\exp(\text{sim}(\mathbf{I}_i^o, \mathbf{T}_i^o) / \tau)}{\sum_{j=1}^N \exp(\text{sim}(\mathbf{I}_i, \mathbf{T}_j) / \tau)}, \quad (3)$$

where \mathbf{I}_i^o and \mathbf{T}_i^o denote the output embeddings from the image and text diffusion branches, respectively. τ is the temperature parameter.

Intrinsic-modal Semantic Alignment. Unified multimodal generation poses high requirements on a model’s representational capacity, particularly in tasks such as image editing and visual perception. These tasks require the model to not only comprehend textual semantics but also accurately perceive high-level visual attributes from images. However, the T5 text encoder and the VAE used in DiT models are pretrained independently on unimodal data, lacking the ability to capture the inherent semantic dependencies between modalities. As a result, the model struggles with aligning text and image semantics effectively. As shown in Fig. 2(c-d), current DiT-based model exhibit limited semantic understanding, often failing to detect critical attributes embedded in the textual instruction and image content.

To address this limitation, we propose an intrinsic-modal semantic alignment strategy that enhances the model’s visual-linguistic perception without introducing additional inference-time parameters. Inspired by REPA (Yu et al. 2025b), we regularize intermediate hidden states of the DiT blocks using semantic embeddings extracted from a pretrained vision-language encoder. Specifically, for each image-text pair, the target image and caption are passed through a pretrained vision-language encoder to obtain semantic embeddings \mathbf{h}_I and \mathbf{h}_T . These embeddings are then aligned with the hidden states \mathbf{I}_i^m and \mathbf{T}_i^m from specific layers of the model’s image and text diffusion branches, respectively. Two Multilayer perceptrons (MLPs) are introduced as projection heads M to mitigate the dimension discrepancy. The intrinsic-modal semantic alignment loss is defined as:

$$\mathcal{L}_{\text{Intrinsic}} = -\frac{1}{N} \sum_{n=1}^N [\text{sim}(\mathbf{h}_I, M(\mathbf{I}_i^m)) + \text{sim}(\mathbf{h}_T, M(\mathbf{T}_i^m))], \quad (4)$$

Type	Method	Params	T2I Generation		Image Editing		Understanding			SemGen-Bench			
			GenEval	DPG-Bench	ImgEdit-Bench	GEdit-Bench-EN	MMB	MMMU	MM-Vet	T2I.Long	T2I.Sem	Edit.Multi	Edit.Sem
Gen. Only	SDXL	-	0.55	74.7	-	-	-	-	-	-	-	-	-
	SD3-medium	-	0.62	84.1	-	-	-	-	-	-	-	-	-
	FLUX.1-dev	-	0.66	84.0	-	-	-	-	-	-	-	-	-
	DualDiffusion	-	0.65	81.3	-	-	-	-	-	7.71	7.29	-	-
Edit. Only	Instruct-P2P	-	-	-	1.88	3.68	-	-	-	-	-	-	-
	MagicBrush	-	-	-	1.90	1.86	-	-	-	-	-	-	-
	AnyEdit	-	-	-	2.45	3.21	-	-	-	-	-	-	-
	Step1X-Edit	-	-	-	3.06	6.70	-	-	-	-	-	6.24	6.10
	IC-Edit	-	-	-	3.05	4.84	-	-	-	-	-	-	-
Und. Only	LLaVA-1.5	-	-	-	-	-	36.4	67.8	36.3	-	-	-	-
	LLaVA-NeXT	-	-	-	-	-	79.3	51.1	57.4	-	-	-	-
Unified	Show-o	-	0.68	67.27	-	-	-	27.4	-	-	-	-	-
	Janus-Pro	-	0.80	84.19	-	-	75.5	36.3	39.8	-	-	-	-
	Emu3	-	0.54	80.60	-	-	58.5	31.6	37.2	-	-	-	-
	BAGEL	7B+7B*	0.82	<u>85.07</u>	3.20	6.52	85.0	55.3	67.2	-	-	-	-
	Uniworl-V1	7B+12B*	0.80	81.38	<u>3.26</u>	4.85	<u>83.5</u>	58.6	<u>67.1</u>	6.42	5.58	4.76	5.77
	OmniGen2	3B+4B*	0.77	83.57	3.44	6.42	79.1	53.1	61.8	<u>7.95</u>	<u>7.90</u>	<u>6.40</u>	5.61
UniAlignment		2B	<u>0.81</u>	85.64	3.25	<u>6.57</u>	80.6	<u>61.3</u>	63.0	8.07	7.97	6.85	<u>5.98</u>

Table 1: **Comparison across Generation, Editing and Understanding tasks.** *: The first term and the second term represent the number of parameters for text generation and image generation, respectively.

where sim denotes the cosine similarity. Notably, both the pretrained vision-language encoder and the MLPs are only used during training and do not introduce any additional parameters during inference.

Based on aforementioned designs, the overall optimization objective is the weighted sum of above loss functions:

$$\mathcal{L} = \mathcal{L}_{img} + \lambda_1 \mathcal{L}_{txt} + \lambda_2 \mathcal{L}_{Cross} + \lambda_3 \mathcal{L}_{Intrinsic} \quad (5)$$

The integration of semantic alignment facilitates fine-grained semantic understanding of the lightweight DiT framework and improves the model’s ability to capture cross-modal dependencies across diverse generation tasks.

Multi-Stage Training

To effectively train a unified generative model across diverse tasks and modalities, we design a multi-stage training strategy that promotes task balance and gradual performance improvement. In the first stage, we conduct image-text pretraining using 2 million image-text pairs from Text-to-Image 2M (Jacky Hate 2024) and an additional 10 million curated internal samples, covering both captioning and generation tasks. This stage focuses on building foundational alignment for both captioning and generation tasks. The second stage is multi-task joint pretraining. We incorporate heterogeneous datasets, including image editing and visual perception examples collected from previous works (Yu et al. 2025a; Zhang et al. 2023; Shi, Wang, and Huang 2024; Wei et al. 2024; Yu et al. 2024; Lin et al. 2025), as well as GPT-4o-generated samples. To adapt these datasets for text-image joint training, we use vision-language models to generate captions for target images lacking textual descriptions. To prevent task-specific overfitting and catastrophic forgetting,

T2I samples are interleaved with conditional task data during this stage. A hybrid batch sampling strategy is adopted, where multiple data types are alternated within training iterations, enabling balanced optimization across tasks. The final stage is supervised finetuning using high-quality datasets including BLIP-3o (Chen et al. 2025b), ShareGPT-4o (Chen et al. 2025a), and LLaVA-OneVision (Li et al. 2024). The progressive multi-stage training encourages the model to incrementally acquire multimodal competencies while mitigating optimization imbalance across diverse tasks.

Experiments

Implementation Details

We build UniAlignment upon the open-sourced SD 3.0 backbone (Esser et al. 2024). A frozen T5 text encoder is employed, and image tokenization is performed using a pretrained VAE. The core architecture, a multimodal DiT, contains only 2 million trainable parameters. The vision encoder of Qwen2.5-VL-7B is utilized for intrinsic-modal semantic alignment instead of CLIP or SigLIP due to their limited input length for text tokens. The DiT embeddings are sourced from a transformer block at depth 8, which we empirically found to yield optimal performance.

Training images are resized to 512×512 , and text sequences are truncated to 256 tokens. The model is trained for 80K steps in the first two stages and 30K steps in the finetuning stage, totaling 25M training samples. During training, we adopt a constant learning rate of $3e-5$ and weight decay schedule of $1e-2$. Gradient accumulation is employed throughout all training stages, allowing for a total batch size of 256. The overall loss function parameters are set as $\lambda_1 = 0.2$, $\lambda_2 = 0.05$ and $\lambda_3 = 0.1$.

Method	GenEval							DPG-Bench					
	Single obj	Two obj	Counting	Colors	Position	Color attr	Overall	Global	Entity	Attribute	Relation	Other	Overall
DualDiffusion	0.97	0.80	0.54	0.76	0.32	0.50	0.65	87.33	89.48	86.72	89.95	86.32	81.32
SDXL	0.98	0.74	0.39	0.85	0.15	0.23	0.55	83.27	82.43	80.91	86.76	80.41	74.65
SD3-medium	0.99	0.94	0.72	0.89	0.33	0.60	0.74	87.90	91.01	88.83	80.70	88.68	84.08
FLUX.1-dev	0.99	0.81	0.79	0.74	0.20	0.47	0.67	82.10	89.50	88.70	91.10	89.40	84.00
LUMINA-Next	0.92	0.46	0.48	0.70	0.09	0.13	0.46	82.82	88.65	86.44	80.53	81.82	74.63
OmniGen	0.98	0.84	0.66	0.74	0.40	0.43	0.68	87.90	88.97	88.47	87.95	83.56	81.16
Show-o	0.98	0.80	0.66	0.84	0.31	0.50	0.68	79.33	75.44	78.02	84.45	60.80	67.27
Janus	0.97	0.68	0.30	0.84	0.46	0.42	0.61	82.33	87.38	87.70	85.46	86.41	79.68
Janus-Pro	0.99	0.89	0.59	0.90	0.79	0.66	0.80	86.90	88.90	89.40	89.32	89.48	84.19
Emu3	0.99	0.81	0.42	0.80	0.49	0.45	0.66	85.21	86.68	86.84	90.22	83.15	80.60
TokenFlow-XL	0.95	0.60	0.41	0.81	0.16	0.24	0.55	78.72	79.22	81.29	85.22	71.20	73.38
BAGEL	0.99	0.94	0.81	0.88	0.64	0.63	0.82	88.94	90.37	91.29	90.82	88.67	<u>85.07</u>
Uniworl-V1	0.99	0.93	0.79	0.89	0.49	0.70	0.80	83.64	88.39	88.44	89.27	87.22	81.38
OmniGen2	1	0.94	0.64	0.88	0.45	0.70	0.77	88.81	88.83	90.18	89.37	90.27	83.57
UniAlignment	0.99	0.95	0.76	0.88	0.56	0.78	<u>0.81</u>	91.84	90.16	90.44	91.58	89.55	85.64

Table 2: Evaluation of text-to-image generation ability on **GenEval benchmark and DPG-Bench**.

Method	Gedit-Bench-EN			ImgEdit-Bench									
	SC	PQ	O	Add	Adjust	Extract	Replace	Remove	Background	Style	Hybrid	Action	Overall
Instruct-P2P	3.58	5.49	3.68	2.45	1.83	1.44	2.01	1.50	1.44	3.55	1.2	1.46	1.88
MagicBrush	4.68	5.66	4.52	2.84	1.58	1.51	1.97	1.58	1.75	2.38	1.62	1.22	1.90
AnyEdit	3.18	5.82	3.21	3.18	2.95	1.88	2.47	2.23	2.24	2.85	1.56	2.65	2.45
OmniGen	5.96	5.89	5.06	3.47	3.04	1.71	2.94	2.43	3.21	4.19	2.24	3.38	2.96
Step1X-Edit	7.09	6.76	6.70	3.88	3.14	1.76	3.40	2.41	3.16	4.63	2.64	2.52	3.06
ICEdit	5.11	6.85	4.84	3.58	3.39	1.73	3.15	2.93	3.08	3.84	2.04	3.68	3.05
BAGEL	7.36	6.83	6.52	3.56	3.31	1.70	3.30	2.62	3.24	4.49	2.38	4.17	3.20
Uniworl-V1	4.93	7.43	4.85	3.82	3.64	2.27	3.47	3.24	2.99	4.21	2.96	2.74	<u>3.26</u>
OmniGen2	7.16	6.77	6.41	3.57	3.06	1.77	3.74	3.20	3.57	4.81	2.52	4.68	3.44
UniAlignment	7.25	6.81	<u>6.57</u>	3.66	3.21	2.07	3.45	2.51	3.28	4.63	2.50	3.95	3.25

Table 3: **Quantitative comparison on GEdit-Bench-EN and ImgEdit-Bench**. For GEdit-Bench, SC (Semantic Consistency) evaluates instruction following, and PQ (Perceptual Quality) assesses image naturalness and artifacts. Higher scores are better.

Evaluation Settings

To comprehensively evaluate UniAlignment, we assess its performance across multiple standard and newly proposed benchmarks. Specifically, text-to-image generation is assessed using GenEval (Ghosh, Hajishirzi, and Schmidt 2023) and DPG-Bench (Hu et al. 2024), while image editing capabilities are evaluated on GEdit-Bench-EN (Liu et al. 2025) and ImgEdit-Bench (Ye et al. 2025). To further investigate multimodal semantic fidelity under complex compositional instructions, we introduce SemGen-Bench. Qualitative assessments on tasks such as image perception and personalization further indicate the versatility of our approach.

Main Results

As shown in Table 1, we conduct a comprehensive comparison of our proposed UniAlignment with state-of-the-art models across text-to-image generation, image editing, visual understanding, and the proposed SemGen-Bench. Despite its compact scale (2B parameters), UniAlignment achieves comparable overall performance across all tasks.

These results demonstrate the effectiveness of our unified DiT architecture and semantic alignment strategies in enabling a lightweight, general-purpose multimodal framework. Below, we detail the evaluation results for each task.

Detailed Task-specific Evaluation

New Evaluation Benchmark: SemGen-Bench A critical challenge in evaluating unified generative models lies in assessing their ability of modeling multimodal semantics, particularly under complex semantic instructions. Although taking visual attributes into account, existing visual generation benchmarks primarily focus on single image synthesis tasks, often under simplistic and narrowly defined settings. For instance, Geneval (Ghosh, Hajishirzi, and Schmidt 2023) evaluates compositional subject attributes within images, but it restricts instruction length and format to predefined templates, limiting its applicability to arbitrary textual inputs. GEdit-Bench (Liu et al. 2025) and ImgEdit-Bench (Ye et al. 2025) target real-world image editing but confined to single-turn instructions with limited semantic

Method	Time	T2I-Long			T2I-Sem			Edit-Multi			Edit-Sem		
		IC	PQ	O	IC	PQ	O	IC	PQ	O	IC	PQ	O
DualDiffusion	2s	7.86	7.65	7.71	7.49	7.54	7.29	-	-	-	-	-	-
StepIX-Edit	91s	-	-	-	-	-	-	6.16	6.65	6.24	6.43	<u>6.84</u>	6.10
UniWorld	21s	5.86	7.82	6.42	5.13	7.59	5.58	4.22	7.16	4.76	5.98	6.95	5.77
OmniGen2	37s	<u>8.29</u>	7.72	<u>7.95</u>	<u>8.29</u>	7.77	<u>7.90</u>	<u>6.42</u>	6.72	<u>6.40</u>	5.97	6.70	5.61
UniAlignment	2s	8.53	<u>7.80</u>	8.07	8.48	<u>7.69</u>	7.97	6.94	<u>7.03</u>	6.85	<u>6.23</u>	6.80	<u>5.98</u>

Table 4: **Overall comparison on our proposed SemGen-Bench.** Instruction Consistency (IC), Perceptual Quality (PQ) and Overall Scores (O) are represented. “Time” denotes the average generation time of each image.



Figure 5: **Comparison of text to image generation with opensource approaches.** The underlined parts emphasize the details in the text instructions.

depth. Prompts in these benchmarks are often direct and lack the complexity required to evaluate fine-grained semantic perception. To address these limitations, we introduce **SemGen-Bench (Semantic Generation Benchmark)**, a large-scale benchmark designed to measure model’s visual generation performance in semantically challenging scenarios. It spans both image generation and editing tasks under diverse and semantically rich instructions, providing a comprehensive assessment of multimodal generative models.

As illustrated in Fig. 4, SemGen-Bench comprises four task categories: (1)*T2I-Long* (long-form text-to-image generation) includes instructions exceeding 200 tokens, testing the ability to process extended contextual informa-

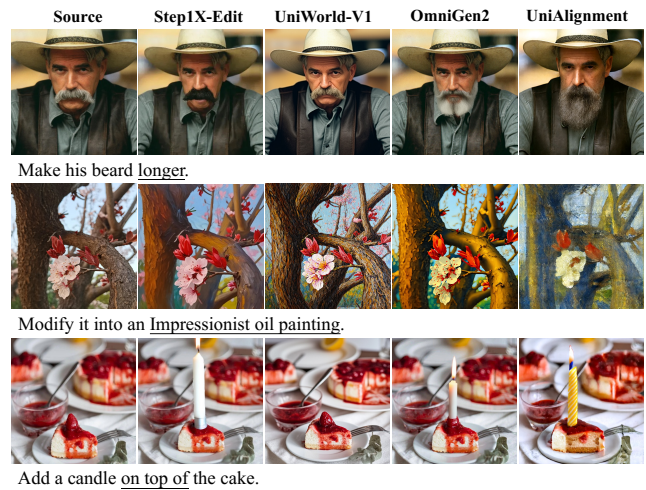


Figure 6: **Comparison of instruction-guided image editing.**

tion. (2)*T2I-Sem* (semantically complex text-to-image generation) demands descriptions incorporating object count, color, spatial relationships, background attributes and logical reasoning. (3)*Edit-Multi* (multi-turn image editing) involves three parallel editing instructions applied to a single image, evaluating the ability to execute coordinated transformations. (4)*Edit-Sem* (semantically complex image editing) focuses on modifying object-level attributes while requiring semantic reasoning beyond surface-level changes. Each category contains 150 examples covering a wide range of entities and scenes.

To construct SemGen-Bench, we adopt a hybrid pipeline combining MLLMs with human annotation. Raw images from MSCOCO (Lin et al. 2014) are filtered to remove samples with low quality and extreme aspect ratios. Next, a multimodal embedding model (Zhang et al. 2024b) is employed to extract vector representations of text-image pairs, followed by hierarchical clustering to obtain 15 semantically diverse categories. Image samples are selected from each category to ensure diversity and balance in the test set. Prompt instructions are first generated using Qwen2.5-VL-72B (Bai et al. 2025), then manually refined to ensure broad diversity in both semantic content and syntactic structure.

To ensure fair evaluation under complex semantic conditions, we adopt an MLLM-based protocol to assess model

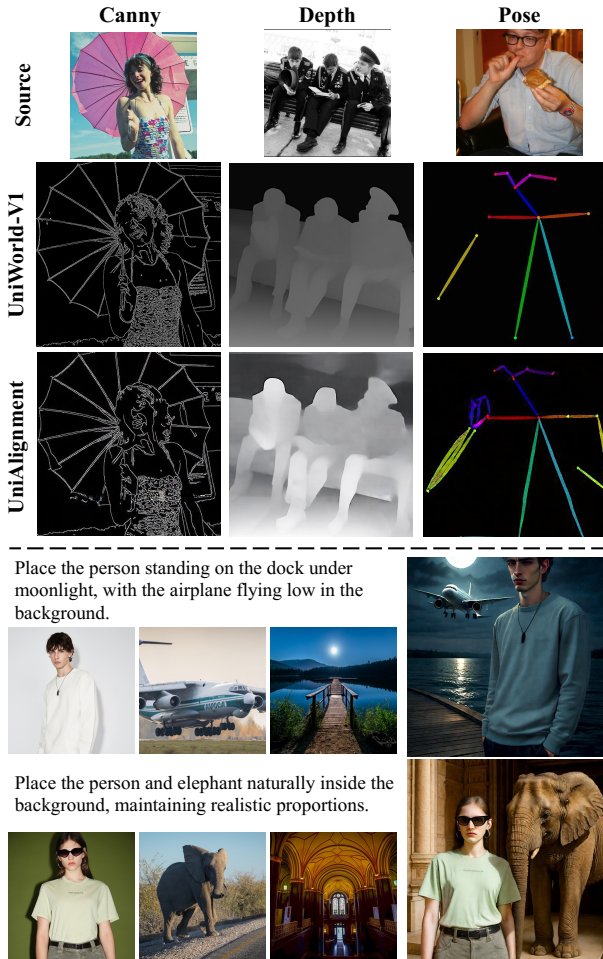
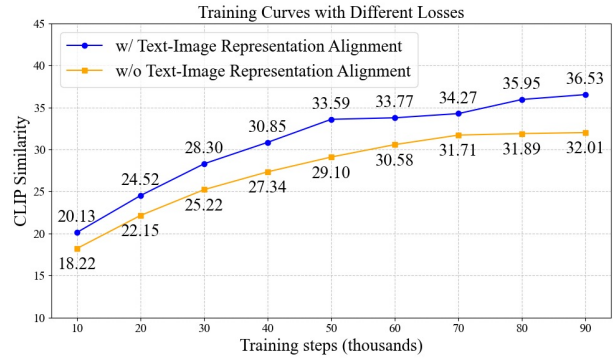


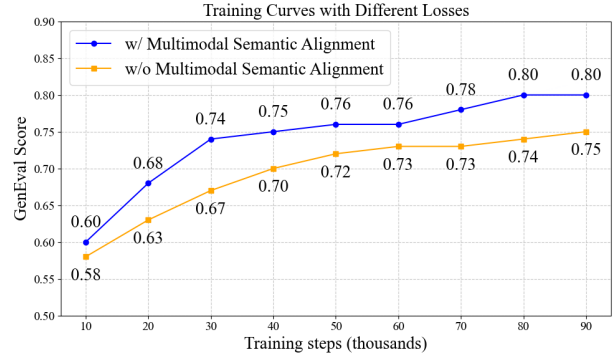
Figure 7: **Showcases of image perception and personalization.** UniAlignment is capable of predicting canny, depth, pose as well as single and multiple subject personalization. Image perception results are compared with UniWorld.

outputs across three dimensions: Instruction Consistency (IC), Perceptual Quality (PQ), and an Overall Score. Following the VIEScore framework (Ku et al. 2024), outputs are rated on a 0–10 scale using MLLMs. The SemGen-Bench will serve as a valuable benchmark for advancing research in unified multimodal generation.

Text-to-Image Generation. As shown in Table 2, UniAlignment achieves highly competitive performance with state-of-the-art methods across all evaluated subcategories. Notably, UniAlignment achieves the highest overall score on DPG-Bench demonstrating superior capability in handling dense textual prompts. Additionally, UniAlignment secures the second-highest score on GenEval with a marginal gap of only 0.01, underscoring its robustness in semantically understanding visual attributes. Qualitative comparisons with three unified generation methods (Li et al. 2025; Lin et al. 2025; Wu et al. 2025b) are presented in Fig. 5. The visualization results exhibit stronger consistency with textual prompts of UniAlignment compared to existing



(a) CLIP score of generated captions and images



(b) GenEval score of generated images

Figure 8: **Ablation results of semantic alignment.**

methods, while also achieving superior visual fidelity.

Image Editing. As shown in Table 3, UniAlignment ranks second only to the specialized editing method StepIX-Edit (Liu et al. 2025) on GEdit-Bench, outperforming UniWorld (1.72) and OmniGen2 (0.16). Notably, UniAlignment achieves the highest score in Semantic Consistency, highlighting its strong capability for joint visual-linguistic semantic grounding. On ImgEdit-Bench, UniAlignment consistently ranks in the top three across nearly all metrics, demonstrating robust and consistent performance in instruction-based image editing tasks. Qualitative visualizations with open-sourced approaches are provided in Fig. 6. Without relying on VLMs or additional visual encoders, UniAlignment demonstrates superior editing performance compared to other methods, effectively balancing fidelity to the original image and adherence to the instructions.

Evaluation on SemGen-Bench. To rigorously benchmark generation performance under complex semantic instructions, we evaluate different models on our proposed SemGen-Bench, which includes tasks involving long-form, complex and multi-step instructions, as shown in Table 4. UniAlignment achieves top-2 performance across nearly all evaluated categories. Importantly, since our framework does not rely on external VLMs at inference, it achieves significantly faster generation time, highlighting the efficiency of its unified diffusion-based architecture. These results further

validate UniAlignment’s strength in visual-language semantic understanding and its effectiveness in handling semantically challenging multimodal tasks.

Image Perception and Personalization. Besides common tasks such as text-to-image generation and image editing, UniAlignment is also capable of realizing image perception as well as subject personalization. As shown in Fig. 7, UniAlignment is capable of generating more refined and visually distinctive image-perception results compared to UniWorld, while also excelling at multi-subject customized image generation tasks.

Ablation Study

We conduct ablation study of our proposed two semantic alignment mechanisms. As illustrated in Fig. 8, training with the cross-modal semantic alignment encourages the joint modelling of visual and textual modalities within a shared latent space, mitigating the optimization conflict and enhancing instruction consistency. Meanwhile, the intrinsic-modal semantic alignment enriches the latent semantics during the diffusion denoising process, leading to improved semantic grounding for image generation.

Conclusion

In this work, we propose UniAlignment, a unified generative model for multimodal generation. Built upon a single lightweight diffusion transformer, UniAlignment is capable of handling a broad spectrum of multimodal tasks, including image generation, understanding, editing, perception, and personalization. Two semantic alignment mechanisms are introduced to enhance the semantic grounding and image-text consistency, without introducing extra parameters. Furthermore, a multi-stage training strategy is implemented to enable progressive task specialization. Experimental results across standard benchmarks as well as our proposed SemGen-Bench demonstrate the superior performance of UniAlignment, offering a promising direction toward general-purpose multimodal intelligence.

Acknowledgments

This work was supported in part by the National Key Research and Development Program of China under Grant No. 2024YFC3210803, in part by the National Natural Science Foundation of China (Grant Nos. U23B2054, 62531026, 62276263), in part by the Youth Innovation Promotion Association CAS (Grant No. Y2023143). This work was supported by Ant Group Research Intern Program.

References

AI, I.; Gong, B.; Zou, C.; Zheng, D.; Yu, H.; Chen, J.; Sun, J.; Zhao, J.; Zhou, J.; Ji, K.; et al. 2025. Ming-lite-uni: Advancements in unified architecture for natural multimodal interaction. *arXiv preprint arXiv:2505.02471*.

Bai, J.; Bai, S.; Chu, Y.; Cui, Z.; Dang, K.; Deng, X.; Fan, Y.; Ge, W.; Han, Y.; Huang, F.; et al. 2023. Qwen technical report. *arXiv preprint arXiv:2309.16609*.

Bai, S.; Chen, K.; Liu, X.; Wang, J.; Ge, W.; Song, S.; Dang, K.; Wang, P.; Wang, S.; Tang, J.; et al. 2025. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*.

Black Forest Labs. 2024. Flux. <https://github.com/black-forest-labs/flux>.

Brooks, T.; Holynski, A.; and Efros, A. A. 2023. Instruct-pix2pix: Learning to follow image editing instructions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 18392–18402.

Caron, M.; Touvron, H.; Misra, I.; Jégou, H.; Mairal, J.; Bojanowski, P.; and Joulin, A. 2021. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 9650–9660.

Chen, J.; Cai, Z.; Chen, P.; Chen, S.; Ji, K.; Wang, X.; Yang, Y.; and Wang, B. 2025a. ShareGPT-4o-Image: Aligning Multimodal Models with GPT-4o-Level Image Generation. *arXiv preprint arXiv:2506.18095*.

Chen, J.; Xu, Z.; Pan, X.; Hu, Y.; Qin, C.; Goldstein, T.; Huang, L.; Zhou, T.; Xie, S.; Savarese, S.; et al. 2025b. Blip3-o: A family of fully open unified multimodal models-architecture, training and dataset. *arXiv preprint arXiv:2505.09568*.

Chen, X.; Wu, Z.; Liu, X.; Pan, Z.; Liu, W.; Xie, Z.; Yu, X.; and Ruan, C. 2025c. Janus-pro: Unified multimodal understanding and generation with data and model scaling. *arXiv preprint arXiv:2501.17811*.

Esser, P.; Kulal, S.; Blattmann, A.; Entezari, R.; Müller, J.; Saini, H.; Levi, Y.; Lorenz, D.; Sauer, A.; Boesel, F.; et al. 2024. Scaling rectified flow transformers for high-resolution image synthesis. In *Forty-first International Conference on Machine Learning*.

Ghosh, D.; Hajishirzi, H.; and Schmidt, L. 2023. Geneval: An object-focused framework for evaluating text-to-image alignment. *Advances in Neural Information Processing Systems*, 36: 52132–52152.

Google. 2025. Gemini 2.0 flash. <https://developers.googleblog.com/en/experiment-with-gemini-20-flash-native-image-generation>.

Hu, X.; Wang, R.; Fang, Y.; Fu, B.; Cheng, P.; and Yu, G. 2024. Ella: Equip diffusion models with llm for enhanced semantic alignment. *arXiv preprint arXiv:2403.05135*.

Huang, R.; Wang, C.; Yang, J.; Lu, G.; Yuan, Y.; Han, J.; Hou, L.; Zhang, W.; Hong, L.; Zhao, H.; et al. 2025. Illume+: Illuminating unified mllm with dual visual tokenization and diffusion refinement. *arXiv preprint arXiv:2504.01934*.

Hudson, D. A.; Zoran, D.; Malinowski, M.; Lampinen, A. K.; Jaegle, A.; McClelland, J. L.; Matthey, L.; Hill, F.; and Lerchner, A. 2024. Soda: Bottleneck diffusion models for representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 23115–23127.

Jacky Hate. 2024. Text-to-image-2m dataset. <https://huggingface.co/datasets/jackyhate/text-to-image-2M>.

- Jia, C.; Yang, Y.; Xia, Y.; Chen, Y.-T.; Parekh, Z.; Pham, H.; Le, Q.; Sung, Y.-H.; Li, Z.; and Duerig, T. 2021. Scaling up visual and vision-language representation learning with noisy text supervision. In *International Conference on Machine Learning*, 4904–4916. PMLR.
- Ku, M.; Jiang, D.; Wei, C.; Yue, X.; and Chen, W. 2024. VEScore: Towards Explainable Metrics for Conditional Image Synthesis Evaluation. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 12268–12290.
- Lee, J.-Y.; Cha, B.; Kim, J.; and Ye, J. C. 2025. Aligning text to image in diffusion models is easier than you think. *arXiv preprint arXiv:2503.08250*.
- Li, B.; Zhang, Y.; Guo, D.; Zhang, R.; Li, F.; Zhang, H.; Zhang, K.; Zhang, P.; Li, Y.; Liu, Z.; et al. 2024. Llava-onevision: Easy visual task transfer. *arXiv preprint arXiv:2408.03326*.
- Li, J.; Li, D.; Savarese, S.; and Hoi, S. 2023. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International Conference on Machine Learning*, 19730–19742. PMLR.
- Li, J.; Li, D.; Xiong, C.; and Hoi, S. 2022. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International Conference on Machine Learning*, 12888–12900. PMLR.
- Li, Z.; Li, H.; Shi, Y.; Farimani, A. B.; Kluger, Y.; Yang, L.; and Wang, P. 2025. Dual diffusion for unified image generation and understanding. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 2779–2790.
- Lin, B.; Li, Z.; Cheng, X.; Niu, Y.; Ye, Y.; He, X.; Yuan, S.; Yu, W.; Wang, S.; Ge, Y.; et al. 2025. Uniworld: High-resolution semantic encoders for unified visual understanding and generation. *arXiv preprint arXiv:2506.03147*.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *European Conference on Computer Vision*, 740–755. Springer.
- Lipman, Y.; Chen, R. T.; Ben-Hamu, H.; Nickel, M.; and Le, M. 2022. Flow Matching for Generative Modeling. In *The Eleventh International Conference on Learning Representations*.
- Liu, S.; Han, Y.; Xing, P.; Yin, F.; Wang, R.; Cheng, W.; Liao, J.; Wang, Y.; Fu, H.; Han, C.; et al. 2025. Step1x-edit: A practical framework for general image editing. *arXiv preprint arXiv:2504.17761*.
- Ma, S.; Ge, Y.; Wang, T.; Guo, Y.; Ge, Y.; and Shan, Y. 2025. GenHancer: Imperfect generative models are secretly strong vision-centric enhancers. *arXiv preprint arXiv:2503.19480*.
- Mao, C.; Zhang, J.; Pan, Y.; Jiang, Z.; Han, Z.; Liu, Y.; and Zhou, J. 2025. Ace++: Instruction-based image creation and editing via context-aware content filling. *arXiv preprint arXiv:2501.02487*.
- OpenAI. 2025. Gpt-4o. <https://openai.com/index/introducing-4o-image-generation>.
- Pan, X.; Shukla, S. N.; Singh, A.; Zhao, Z.; Mishra, S. K.; Wang, J.; Xu, Z.; Chen, J.; Li, K.; Juefei-Xu, F.; et al. 2025. Transfer between modalities with metaqueries. *arXiv preprint arXiv:2504.06256*.
- Peebles, W.; and Xie, S. 2023. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 4195–4205.
- Qu, L.; Zhang, H.; Liu, Y.; Wang, X.; Jiang, Y.; Gao, Y.; Ye, H.; Du, D. K.; Yuan, Z.; and Wu, X. 2025. Tokenflow: Unified image tokenizer for multimodal understanding and generation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 2545–2555.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, 8748–8763.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10684–10695.
- Shi, Y.; Wang, P.; and Huang, W. 2024. Seedit: Align image re-generation to image editing. *arXiv preprint arXiv:2411.06686*.
- Sun, Q.; Fang, Y.; Wu, L.; Wang, X.; and Cao, Y. 2023. Evalclip: Improved training techniques for clip at scale. *arXiv preprint arXiv:2303.15389*.
- Swerdlow, A.; Prabhudesai, M.; Gandhi, S.; Pathak, D.; and Fragkiadaki, K. 2025. Unified multimodal discrete diffusion. *arXiv preprint arXiv:2503.20853*.
- Team, C. 2024. Chameleon: Mixed-modal early-fusion foundation models. *arXiv preprint arXiv:2405.09818*.
- Tian, C.; Tao, C.; Dai, J.; Li, H.; Li, Z.; Lu, L.; Wang, X.; Li, H.; Huang, G.; and Zhu, X. 2024. ADDP: Learning General Representations for Image Recognition and Generation with Alternating Denoising Diffusion Process. In *ICLR*.
- Touvron, H.; Lavril, T.; Izacard, G.; Martinet, X.; Lachaux, M.-A.; Lacroix, T.; Rozière, B.; Goyal, N.; Hambro, E.; Azhar, F.; et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Wang, S.; Li, W.; Wang, Q.; Zhao, S.; and Zhang, J. 2025a. MIND-Edit: MLLM Insight-Driven Editing via Language-Vision Projection. *arXiv preprint arXiv:2505.19149*.
- Wang, W.; Sun, Q.; Zhang, F.; Tang, Y.; Liu, J.; and Wang, X. 2025b. Diffusion Feedback Helps CLIP See Better. In *The Thirteenth International Conference on Learning Representations*.
- Wang, X.; Zhang, X.; Luo, Z.; Sun, Q.; Cui, Y.; Wang, J.; Zhang, F.; Wang, Y.; Li, Z.; Yu, Q.; et al. 2024. Emu3: Next-token prediction is all you need. *arXiv preprint arXiv:2409.18869*.
- Wei, C.; Mangalam, K.; Huang, P.-Y.; Li, Y.; Fan, H.; Xu, H.; Wang, H.; Xie, C.; Yuille, A.; and Feichtenhofer, C. 2023. Diffusion models as masked autoencoders. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 16284–16294.

- Wei, C.; Xiong, Z.; Ren, W.; Du, X.; Zhang, G.; and Chen, W. 2024. Omniedit: Building image editing generalist models through specialist supervision. In *The Thirteenth International Conference on Learning Representations*.
- Wu, C.; Chen, X.; Wu, Z.; Ma, Y.; Liu, X.; Pan, Z.; Liu, W.; Xie, Z.; Yu, X.; Ruan, C.; et al. 2025a. Janus: Decoupling visual encoding for unified multimodal understanding and generation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 12966–12977.
- Wu, C.; Zheng, P.; Yan, R.; Xiao, S.; Luo, X.; Wang, Y.; Li, W.; Jiang, X.; Liu, Y.; Zhou, J.; et al. 2025b. OmniGen2: Exploration to Advanced Multimodal Generation. *arXiv preprint arXiv:2506.18871*.
- Wu, Y.; Zhang, Z.; Chen, J.; Tang, H.; Li, D.; Fang, Y.; Zhu, L.; Xie, E.; Yin, H.; Yi, L.; et al. 2025c. Vila-u: a unified foundation model integrating visual understanding and generation. In *The Thirteenth International Conference on Learning Representations*.
- Xiang, W.; Yang, H.; Huang, D.; and Wang, Y. 2023. Denoising diffusion autoencoders are unified self-supervised learners. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 15802–15812.
- Xiao, S.; Wang, Y.; Zhou, J.; Yuan, H.; Xing, X.; Yan, R.; Li, C.; Wang, S.; Huang, T.; and Liu, Z. 2025. Omnigen: Unified image generation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 13294–13304.
- Xie, J.; Mao, W.; Bai, Z.; Zhang, D. J.; Wang, W.; Lin, K. Q.; Gu, Y.; Chen, Z.; Yang, Z.; and Shou, M. Z. 2024a. Show-o: One single transformer to unify multimodal understanding and generation. *arXiv preprint arXiv:2408.12528*.
- Xie, R.; Du, C.; Song, P.; and Liu, C. 2024b. Muse-vl: Modeling unified vlm through semantic discrete encoding. *arXiv preprint arXiv:2411.17762*.
- Yang, L.; Tian, Y.; Li, B.; Zhang, X.; Shen, K.; Tong, Y.; and Wang, M. 2025. Mmada: Multimodal large diffusion language models. *arXiv preprint arXiv:2505.15809*.
- Yao, J.; Yang, B.; and Wang, X. 2025. Reconstruction vs. generation: Taming optimization dilemma in latent diffusion models. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 15703–15712.
- Ye, Y.; He, X.; Li, Z.; Lin, B.; Yuan, S.; Yan, Z.; Hou, B.; and Yuan, L. 2025. Imgedit: A unified image editing dataset and benchmark. *arXiv preprint arXiv:2505.20275*.
- Yu, Q.; Chow, W.; Yue, Z.; Pan, K.; Wu, Y.; Wan, X.; Li, J.; Tang, S.; Zhang, H.; and Zhuang, Y. 2025a. Anyedit: Mastering unified high-quality image editing for any idea. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 26125–26135.
- Yu, S.; Kwak, S.; Jang, H.; Jeong, J.; Huang, J.; Shin, J.; and Xie, S. 2025b. Representation Alignment for Generation: Training Diffusion Transformers Is Easier Than You Think. In *The Thirteenth International Conference on Learning Representations*.
- Yu, Y.; Zeng, Z.; Hua, H.; Fu, J.; and Luo, J. 2024. Prompt-Fix: you prompt and we fix the photo. In *Proceedings of the 38th International Conference on Neural Information Processing Systems*, 40000–40031.
- Zhai, X.; Mustafa, B.; Kolesnikov, A.; and Beyer, L. 2023. Sigmoid loss for language image pre-training. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 11975–11986.
- Zhang, H.; Duan, Z.; Wang, X.; Zhao, Y.; Lu, W.; Di, Z.; Xu, Y.; Chen, Y.; and Zhang, Y. 2025. Nexus-gen: A unified model for image understanding, generation, and editing. *arXiv preprint arXiv:2504.21356*.
- Zhang, K.; Mo, L.; Chen, W.; Sun, H.; and Su, Y. 2023. Magicbrush: A manually annotated dataset for instruction-guided image editing. *Advances in Neural Information Processing Systems*, 36: 31428–31449.
- Zhang, Q.; Zhang, J.; Xu, Y.; and Tao, D. 2024a. Vision transformer with quadrangle attention. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(5): 3608–3624.
- Zhang, X.; Zhang, Y.; Xie, W.; Li, M.; Dai, Z.; Long, D.; Xie, P.; Zhang, M.; Li, W.; and Zhang, M. 2024b. GME: Improving Universal Multimodal Retrieval by Multimodal LLMs. *arXiv preprint arXiv:2412.16855*.
- Zhou, C.; Yu, L.; Babu, A.; Tirumala, K.; Yasunaga, M.; Shamis, L.; Kahn, J.; Ma, X.; Zettlemoyer, L.; and Levy, O. 2024. Transfusion: Predict the next token and diffuse images with one multi-modal model. *arXiv preprint arXiv:2408.11039*.