

SenseNova-U1: Unifying Multimodal Understanding and Generation with NEO-unify Architecture

Abstract

Recent large vision–language models (VLMs) remain fundamentally constrained by a persistent dichotomy: understanding and generation are treated as distinct problems, leading to fragmented architectures, cascaded pipelines, and misaligned representation spaces. We argue that this divide is not merely an engineering artifact, but a structural limitation that hinders the emergence of native multimodal intelligence. Hence, we introduce **SenseNova-U1**, a native unified multimodal paradigm built upon **NEO-unify** [112], in which understanding and generation evolve as synergistic views of a single underlying process.

We launch two native unified variants, **SenseNova-U1-8B-MoT** and **SenseNova-U1-A3B-MoT**, built on dense (8B) and mixture-of-experts (30B-A3B) understanding baselines, respectively. Designed from first principles, they rival top-tier understanding-only VLMs across text understanding, vision-language perception, knowledge reasoning, agentic decision-making, and spatial intelligence. Meanwhile, they deliver strong semantic consistency and visual fidelity, excelling in conventional or knowledge-intensive any-to-image (X2I) synthesis, complex text-rich infographic generation, and interleaved vision-language generation, with or without think patterns. Beyond performance, we show detailed model design, data preprocessing, pre-/post-training, and inference strategies to support community research.

Last but not least, preliminary evidence demonstrates that our models extend beyond perception and generation, performing strongly in vision–language–action (VLA) and world model (WM) scenarios. This points toward a broader roadmap where models do not translate between modalities, but think-and-act across them in a native manner. Multimodal AI is no longer about connecting separate systems, but about building a unified one and trusting the necessary capabilities to emerge from within.



Official Demo: <https://unify.light-ai.top/>

GitHub Code: <https://github.com/OpenSenseNova/SenseNova-U1>

HuggingFace Model: <https://huggingface.co/collections/sensenova/sensenova-u1>

NEO-unify Blog: <https://huggingface.co/blog/sensenova/neo-unify> (March 5, 2026)



Figure 1 Showcases of SenseNova-U1-8B-MoT in infographics and human generation.




Figure 2 Showcases of SenseNova-U1-8B-MoT in image editing and interleaved generation.

Contents

1	Introduction	5
2	Related Works	5
2.1	Native Multimodal Models	5
2.2	Native Multimodal Unified Models	6
3	Methodology	7
3.1	Near-Lossless Visual Interface	7
3.2	Native Multimodal Unified Modeling	8
3.3	Joint Training Objective	8
3.4	Training Procedure	9
3.5	Inference Infrastructure	12
4	Data Construction	13
4.1	Understanding Data Organization	13
4.2	Generation Data Organization	14
5	Experiments	16
5.1	Main Results	16
5.1.1	Image Understanding	16
5.1.2	Text Understanding	17
5.1.3	Image Generation	18
5.1.4	Image Editing	25
5.1.5	Interleaved Generation	27
5.2	Ablation Studies	29
5.2.1	Native Encoder-Free Design Preserves Both Semantic and Pixel Representations	29
5.2.2	Understanding and Generation Synergize with Native MoT Backbone	31
5.2.3	Native Multimodal Architecture Shows High Data-Scaling Efficiency	32
5.3	Visualization Results	32
6	Conclusion	34
7	Contributors	34

1 Introduction

Recent advances in multimodal foundation models [2, 60, 134] have markedly enhanced both perception and generation across vision and language. Yet these capabilities have largely evolved in isolation. This divide stems from the underlying system design: understanding is typically mediated by pretrained vision encoders (VEs) [107, 118, 163], whereas generation relies on latent variational autoencoders (VAEs) [58, 152]. These choices impose distinct learning objectives [58, 107] and training pipelines [1, 67, 79, 109], resulting in divergent feature representations that bifurcate multimodal modeling into separate regimes. Consequently, early unified multimodal models (UMMs) [14, 18, 28, 75, 139, 140] remain loosely integrated, with perception and generation connected through different tokenizers, latent spaces, or auxiliary modules rather than being learned jointly within a truly unified system.

Against this backdrop, native vision–language models (VLMs) have emerged along two distinct directions. One casts multimodality as an extension of language, mapping all modalities into discrete tokens within a unified autoregressive framework [23, 73, 78, 90, 116, 122, 125, 135]. While enabling seamless cross-modal reasoning, this discretization inevitably compresses non-linguistic signals into lossy representations, constraining both high-level semantics and visual fidelity. The other instead pursues a unified continuous visual interface spanning understanding and generation [35, 84, 127, 152, 168, 170], seeking to reconcile conceptual structure with high-fidelity reconstruction within a shared representation space — but often with trade-offs. Yet neither resolves the fundamental tension between semantic abstraction and pixel-level granularity. This leaves open a central question: can multimodal intelligence be unified in a truly native form, breaking free from latent bottlenecks and intermediate representations?

We return to the first principles: building a model that directly engages with native inputs (*i.e.* pixels and words), and steps beyond representation arguments or pre-trained priors. Crucially, we dispense with both pretrained vision encoders and deep decoder heads, yielding a unified architecture that supports concise and scalable training. Hence, we introduce **SenseNova-U1**, a native unified multimodal paradigm built on the **NEO-unify** [112] model. As a first step toward truly end-to-end unification, it learns directly from lossless inputs and self-organizes potential representation spaces tailored to diverse application scenarios. Specifically, it incorporates: (i) a near-lossless visual interface that simultaneously preserves semantic structure and fine-grained pixel detail without any pretrained VEs or VAEs; (ii) a unified end-to-end modeling over raw inputs that jointly couples autoregressive cross-entropy for language with pixel-space flow matching for vision; (iii) a native mixture-of-transformers (MoT) architecture that synergizes understanding and generation in an intrinsically multimodal system with minimal objective interference and powerful scaling efficiency.

We launch two variants, **SenseNova-U1-8B-MoT** and **SenseNova-U1-A3B-MoT**, built upon dense (8B) and mixture-of-experts (30B-A3B) multimodal understanding backbones, respectively. Both models adopt a native MoT architecture, enabling efficient scaling while reducing interference across heterogeneous multimodal objectives. Empirically, SenseNova-U1 rivals top-tier understanding-only VLMs across text understanding, vision-language perception, knowledge reasoning, agentic decision-making, and spatial intelligence, while simultaneously achieving strong any-to-image (X2I) generation performance under a $32\times$ compression ratio across conventional, knowledge-intensive, and text-rich scenarios. Beyond them, it supports visual-centric reasoning and coherent interleaved generation across modalities, enabling applications such as illustrated guides, visual storytelling, presentations, posters, comics, resumes, and other information-dense visual formats requiring structured layout generation and high-fidelity rendering. Overall, SenseNova-U1 sets a brand-new paradigm for unified multimodal understanding and generation, outperforming prior open-source models across a wide range of understanding, reasoning, and generation benchmarks.

Preliminary experiments further suggest promising capabilities in vision–language–action (VLA) and world modeling (WM), indicating that our models can reason and act natively across modalities without relying on external adapters or modular bridges. More broadly, SenseNova-U1 points toward a shift in multimodal AI: from connecting separate modality-specific systems to learning perception, reasoning, and generation within a natively unified architecture.

2 Related Works

2.1 Native Multimodal Models

Recently, vision–language models (VLMs) [2, 47, 52, 57, 106, 115, 134] have rapidly advanced multimodal understanding by coupling visual encoders (VEs) with large language models (LLMs), through either staged pretraining or joint optimization. Despite their success, such designs inherit pretrained semantic biases and introduce additional

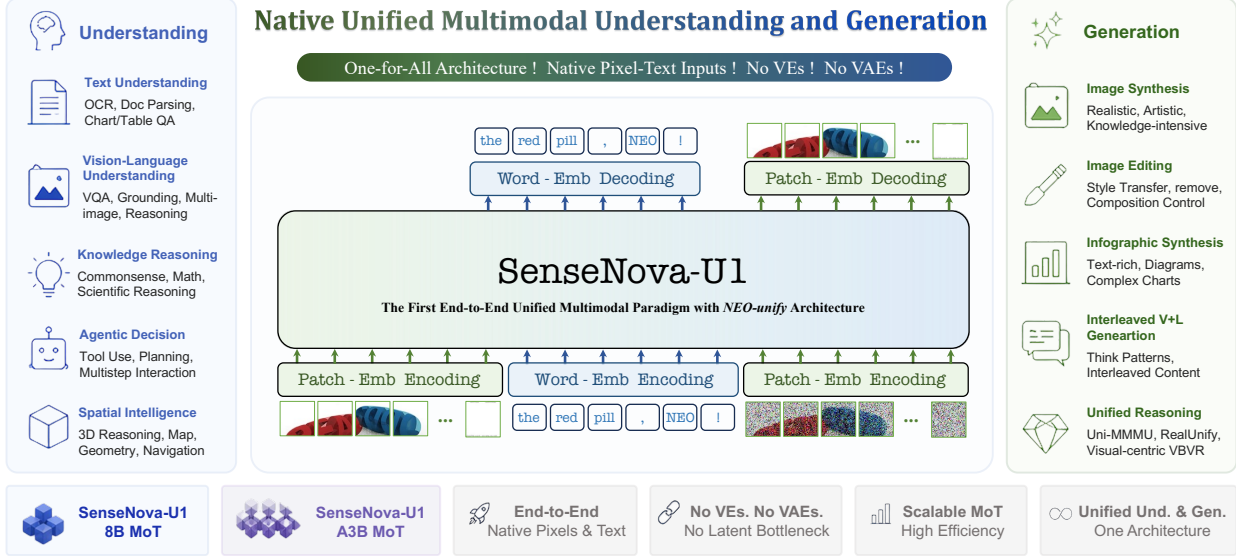


Figure 3 Overview of SenseNova-U1. With extremely lightweight encoding and decoding interfaces, SenseNova-U1 enables efficient and deeply correlated pixel-word correspondence within a single end-to-end architecture. As a native unified multimodal paradigm, it jointly supports diverse application scenarios, including perception, synthesis, and interleaved vision-language generation.

complexity, along with inherent capacity trade-offs across components. This has motivated a shift toward native multimodal backbones without VEs, as exemplified by Fuyu [4] and EVE [29]. Subsequent works push further by efficiently constructing visual perception while mitigating vision-language conflicts through distillation [29, 70, 131], data mixing [19, 63], shared modules [121, 151], and modality decomposition [31, 88, 89]. Notably, NEO [30] advances this line by exploring a native pixel-word primitive, substantially narrowing the gap with leading modular VLMs over diverse understanding tasks. For years, visual generation has been dominated by low-dimensional VAE or VQ-VAE latents [58, 128], with heavy compression limiting semantic expressivity under reconstruction-driven objectives. Although recent efforts [64, 152] enrich these latents with pretrained representations or auxiliary objectives, they remain fundamentally constrained by the compression bottleneck and fragmented training pipelines. In parallel, emerging works [17, 20, 69, 158] validate that direct pixel-space modeling can rival or even surpass latent diffusion, pointing toward a fundamentally new direction via fully end-to-end optimization from raw pixels.

2.2 Native Multimodal Unified Models

Early efforts to unify multimodal understanding and generation have largely converged on shared backbones, as exemplified by Show-o [145, 146], Janus [18, 92, 140], OmniGen [141, 144], and BAGEL [28]. While these systems demonstrate that perception and synthesis can coexist within a single model, they remain split across fundamentally different tokenizers, diffusion heads, or decoupled pathways, reflecting a deeper mismatch between understanding and generation. A complementary line of work shifts the focus to the visual interface itself, including shared discrete tokenizers [76, 90, 105, 143, 167] or continuous representation-based autoencoder [13, 35, 84, 113, 127, 161, 168]. These approaches partially reconcile perception and synthesis, yet remain fundamentally constrained by intermediate representations, where semantic structure and visual fidelity must be traded against each other.

Native multimodal modeling is increasingly diverging along two distinct directions. Discrete unified models [23, 66, 73, 78, 122, 125, 135] recast multimodal learning as token-level autoregression, achieving architectural unification while sacrificing visual fidelity and expressivity under discrete tokenization. In parallel, continuous native approaches pursue end-to-end modeling without explicit tokenizers or latent bottlenecks. NEO-unify [112] takes a first step toward this direction by learning directly from near-lossless inputs, achieving strong performance across diverse understanding and generation tasks. Tuna-2 [85] further demonstrates that pixel-space modeling can match latent-space methods, reinforcing the view that high-fidelity generation need not rely on compressed representations. Notably, SenseNova-U1 builds on NEO-unify [112] by scaling this paradigm across data corpus, model capacity, and application scenarios, moving toward a truly unified foundation in which multimodal intelligence emerges natively.

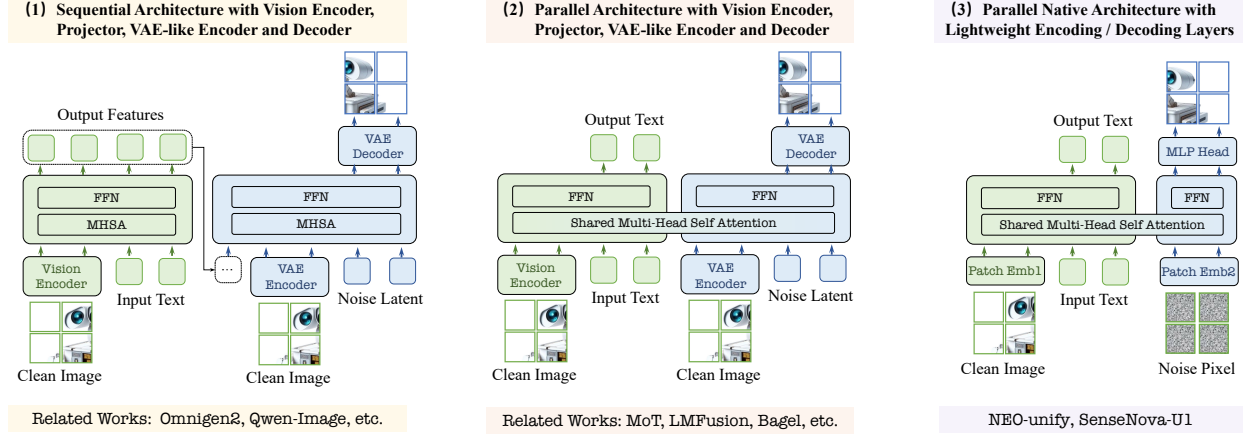


Figure 4 SenseNova-U1 built on NEO-unify: one native paradigm for multimodal understanding and generation. SenseNova-U1 operates directly on native pixel and text inputs without relying on separate VEs or VAEs. The framework combines a near-lossless visual interface, implemented with two-layer convolutional encoding and MLP-like decoding layers, together with a native Mixture-of-Transformers (MoT) main architecture. Despite using a $32\times$ compression ratio, it delivers strong performance across a broad range of understanding and generation tasks, while substantially simplifying system design and improving computational efficiency.

3 Methodology

For years, multimodal models have relied on a vision encoder (VE) for perception and a variational autoencoder (VAE) for generation. Recent efforts attempt to unify these components through shared tokenizers, yet remain constrained by representational trade-offs. SenseNova-U1 returns to first principles, introducing a native, unified, end-to-end framework that operates directly on pixels and words, eliminating reliance on pretrained encoder priors and the scaling limitations imposed by fixed representations. The overall framework is illustrated in Figure 4.

3.1 Near-Lossless Visual Interface

Patch Encoding Layer. We follow NEO [30] to construct lightweight patch encoding layers. Given an input image or noise, we map it into a sequence of visual tokens using two convolutional layers with GELU activation and 2D sinusoidal positional encoding. The convolutional strides are set to 16 and 2, so that each token corresponds to a 32×32 image patch. Two special $\langle \text{img} \rangle$ and $\langle / \text{img} \rangle$ tokens are used to delimit visual content. Besides, text words are encoded using the original tokenizer of the underlying language model without modification. After that, visual and textual tokens are projected into a shared embedding space and processed jointly within a unified backbone.

Patch Decoding Layer. The understanding stream uses a linear projection head to map tokens to the word vocabulary for text prediction. The generation stream directly predicts pixel patches via a multi-layer perceptron (MLP) head, bypassing deep diffusion heads and VAE decoders. This design enables fully end-to-end learning of the representation space, free from the inductive biases and representational constraints imposed by intermediate modules.

Dynamic Noise Scale. Because the generation stream operates over varying resolutions, a naive unit-variance prior $\mathbf{z}_1 \sim \mathcal{N}(0, \mathbf{I})$ becomes mismatched to the signal scale, leading to inconsistent signal-to-noise ratios (SNRs) across resolutions at the same flow timestep. To address this, we introduce a resolution-adaptive noise scale σ_R . Let $N(H, W) = (H \cdot W)/32^2$ denote the number of generation tokens for an image of size $H \times W$, and let N_0 be a reference token count. We define: $\sigma_R(H, W) = \sigma_0 \sqrt{N(H, W)/N_0}$, where σ_0 is a base noise scale. During training, terminal noise is sampled from a Gaussian distribution scaled by σ_R , which also initializes the flow ordinary differential equation (ODE) at inference. Intuitively, the square-root scaling preserves approximately constant per-token noise energy from low to high resolutions, ensuring a consistent SNR distribution for flow matching.

Noise-Scale Conditioning. Since σ_R varies with image resolutions, we explicitly feed it to the denoiser. We normalize the scale as $\bar{\sigma} = \sigma_R/\sigma_{\max} \in [0, 1]$ and encode it using a dedicated sinusoidal MLP embedder $\text{NSEmb}(\cdot)$. The resulting embedding is combined with the timestep embedding τ_t to form the conditioning signal: $\mathbf{s}_t = \tau_t + \text{NSEmb}(\bar{\sigma}(H, W))$, where \mathbf{s}_t denotes the joint time and noise-scale conditioning applied to the input image tokens.

Configuration	SenseNova-U1-8B-MoT	SenseNova-U1-A3B-MoT
Patch Size	32×32	32×32
Pre-Buffer	✓	✗
# Num Layers	42	48
# Num Heads (Q / KV)	32 / 8	32 / 4
Head Size (T / H / W)	64 / 32 / 32	64 / 32 / 32
Hidden Size	4,096	2,048
# Und / Gen Experts	1 / 1	128 / 32 (A8)
# Und / Gen Parameters	8.2B / 8.2B	30.0B / 8.2B (A3B)

Table 1 Architectural configurations of SenseNova-U1 variants: SenseNova-U1-8B-MoT & SenseNova-U1-A3B-MoT.

3.2 Native Multimodal Unified Modeling

Improved Native Primitive. We refine the native VLM primitive from NEO [30] as the base transformer block. Its native rotary position embedding (Native RoPE) unifies temporal and spatial encoding within a single representation. Text tokens evolve along the temporal axis T with $H = W = 0$, while image tokens additionally carry spatial indices along height H and width W . The new design *reallocates* pretrained LLM head dimensions across the T , H , and W axes, each associated with independent frequency bases and incurring no additional parameters. It is applied to the Query and Key projections, along with their corresponding normalizations, all initialized from the understanding backbone. Besides, we maintain native multimodal attention that jointly supports language and vision modeling.

Native Mixture-of-Transformers. At the core of SenseNova-U1 is a native Mixture-of-Transformers (MoT) backbone that unifies understanding and generation within a monolithic framework. The understanding stream processes clean image and text inputs, while the generation stream operates on noise-conditioned inputs. All modalities are represented within a single sequence and processed under a shared self-attention mechanism, enabling perception and synthesis to interact natively at every layer. Here, text tokens attend causally to preceding tokens only. Image tokens within the same block attend bidirectionally to one another while remaining causally conditioned on all preceding context. Noise tokens within each image block also attend bidirectionally, with full access to clean inputs, whereas clean tokens are prevented from attending to any noise tokens. Crucially, we adopt full parameter decoupling between the two streams, with separate projections, normalizations, and feedforward blocks dynamically routed by token type at each layer.

Model Variants. SenseNova-U1 is instantiated at two scales (detailed model configurations are provided in Table 1):

- **SenseNova-U1-8B-MoT.** The shallow Pre-Buffer layers map raw pixel and text inputs into a unified representation, while the Post-LLM layers retain the linguistic proficiency and reasoning capabilities of a pretrained LLM. Besides, both streams are instantiated as dense 8B networks in a symmetric parallel configuration.
- **SenseNova-U1-A3B-MoT.** To scale efficiently, we extend the MoT framework with stream-wise mixture-of-experts (MoE) without Pre-Buffer layers. The understanding stream employs 128 experts with a total of 30B parameters, while the generation stream uses 32 experts totaling 8B parameters. A top- k routing strategy activates 8 experts per token in each stream, resulting in approximately 3B active parameters during inference.

3.3 Joint Training Objective

SenseNova-U1 is optimized end-to-end with text and visual generation objectives weighted by λ_1 and λ_2 :

$$\mathcal{L}_{\text{total}} = \lambda_1 \mathcal{L}_{\text{Und}} + \lambda_2 \mathcal{L}_{\text{Gen}} \quad (1)$$

Autoregressive Text Loss. For understanding tasks, we employ standard next-token prediction as follows:

$$\mathcal{L}_{\text{Und}} = -\frac{1}{N} \sum_{n=1}^N \log p_{\theta}(x_n \mid x_{<n}, \mathbf{c}) \quad (2)$$

where x_n denotes the n -th text token, $x_{<n}$ the preceding tokens, and \mathbf{c} the multimodal context tokens.

Pixel-Space Flow Matching. For visual generation, we follow the former JiT [69] with \mathbf{x} -predict and \mathbf{v} -loss, operating directly in the pixel-level space. Given a clean image $\mathbf{x} \in \mathbb{R}^{3 \times H \times W}$ and a Gaussian sample $\epsilon \sim \mathcal{N}(0, \mathbf{I})$, we form the noisy sample along the rectified-flow interpolant, formulated as follows:

$$\mathbf{z}_t = t\mathbf{x} + (1-t)\sigma_R\epsilon, \quad t \in [0, 1], \quad (3)$$

where $t = 0$ corresponds to pure noise ($\mathbf{z}_0 = \sigma_R\epsilon$), and $t = 1$ corresponds to the clean image ($\mathbf{z}_1 = \mathbf{x}$). Note that σ_R denotes the resolution-adaptive noise scale. The unified framework directly regresses the clean signal $\hat{\mathbf{x}}_\theta$, which is then converted into a velocity term for \mathbf{v} -loss computation as follows:

$$\mathbf{v}_\theta(\mathbf{z}_t, t) = \frac{\hat{\mathbf{x}}_\theta(\mathbf{z}_t, t, \mathbf{s}_t) - \mathbf{z}_t}{1-t}, \quad (4)$$

where \mathbf{s}_t is joint time-and-noise-scale conditioning. We adopt mean squared error (MSE) for velocity-space loss,

$$\mathcal{L}_{\text{Gen}} = \mathbb{E}_{t, \mathbf{x}, \epsilon, (H, W)} \left[\|\mathbf{v}_\theta(\mathbf{z}_t, t) - \mathbf{v}^*\|_2^2 \right], \quad \mathbf{v}^* = \frac{\mathbf{x} - \mathbf{z}_t}{1-t}. \quad (5)$$

Classifier-Free Guidance. For generation tasks, including text-to-image synthesis, image editing, and interleaved image-text generation, we adopt a unified classifier-free guidance formulation that independently modulates the influence of textual and visual conditions. Let \mathbf{c}_{txt} denote the text condition and \mathbf{c}_{img} the visual context. During training, we randomly drop the text condition with probability 10%, and drop both text and image conditions with an additional probability of 10%, enabling the model to learn conditional, image-only, and unconditional generation within a single framework. During inference, the guided score is formulated as:

$$\begin{aligned} \nabla_{\mathbf{x}} \log p(\mathbf{x} \mid \mathbf{c}_{\text{img}}, \mathbf{c}_{\text{txt}}) &= \gamma \left(\nabla_{\mathbf{x}} \log p(\mathbf{x} \mid \mathbf{c}_{\text{img}}, \mathbf{c}_{\text{txt}}) - \nabla_{\mathbf{x}} \log p(\mathbf{x} \mid \mathbf{c}_{\text{img}}) \right) \\ &+ \gamma_{\text{img}} \left(\nabla_{\mathbf{x}} \log p(\mathbf{x} \mid \mathbf{c}_{\text{img}}) - \nabla_{\mathbf{x}} \log p(\mathbf{x}) \right) + \nabla_{\mathbf{x}} \log p(\mathbf{x}). \end{aligned} \quad (6)$$

Here, γ controls text guidance and γ_{img} controls image-context guidance. Empirically, $\gamma = 4$ and $\gamma_{\text{img}} = 1$ consistently yield the best performance across X2I tasks, suggesting that explicit image-context guidance plays a comparatively minor role. This observation implies that the model already captures visual conditioning effectively, while stronger guidance is primarily needed to enforce textual alignment. In practice, this guidance is applied to the predicted flow velocity used for generation. Note that we apply a timestep shift of 3.0 and global CFG renormalization strategies.

3.4 Training Procedure

SenseNova-U1 is trained via progressive stages in Table 2 that incrementally build native multimodal capabilities.

Stage 1: Understanding Warmup. We initialize from a pretrained NEO [30] and perform two efficiency-oriented adaptations: an attention-fusion phase that simplifies original QK projections and normalization, followed by a full-model continuation phase that re-equilibrates the network under the enhanced attention modules.

(i) *Attention-Fusion Phase.* We unify NEO’s QK projections and normalization across the temporal and spatial axes into a single shared set, halving the QK parameter footprint while preserving the native RoPE multi-axis structure and maintaining separate frequency scaling for temporal (rope theta = 5,000,000) and spatial dimensions (rope theta = 10,000). To mitigate the short-term performance drop, we freeze the rest of the network and fine-tune only the attention layers, including Q, K, V, and output projections as well as QK normalization, until the model recovers its pre-fusion accuracy. Training data for this phase is drawn from an updated mid-training corpus described in Sec. 4.1.

(ii) *Full-Model Continuation Phase.* We then unfreeze the entire understanding branch and continue training on the same updated mid-training corpus, using a learning rate of 2×10^{-5} . The resulting model forms the understanding backbone of SenseNova-U1, providing rich contextual conditioning for subsequent generation phases.

Stage 2: Generation Pre-Training. With the understanding branch frozen, we pretrain the generation branch on text-to-image data. It learns to synthesize pixel patches directly via pixel-space flow matching, conditioned on the semantic context from the frozen understanding branch. This stage establishes a stable generative foundation before joint optimization. Specifically, we conduct the overall generation pre-training processes in three phases.

	Stage 1 : Understanding Warmup	Stage 2 : Generation Pre-Training			Stage 3 : Unified Mid-Training	Stage 4 : Unified SFT
		Phase I	Phase II	Phase III		
Hyperparameters						
Peak learning rate	2×10^{-5}	2×10^{-4}	1×10^{-4}	1×10^{-4}	2×10^{-5}	2×10^{-5}
Min learning rate	–	–	–	2×10^{-5}	–	0
LR scheduler	Constant	Constant	Constant	Cosine decay	Constant	Cosine decay
Optimizer	AdamW ($\beta_1 = 0.9, \beta_2 = 0.95, \epsilon = 10^{-8}$)					
Weight decay	0.0	0.0	0.0	0.0	0.0	0.0
Gradient norm clip	1.0	1.0	1.0	1.0	1.0	1.0
EMA ratio	–	0.9999	0.9999	0.9999	0.999	0.999
Training steps	120K	120K	60K	120K	84K	9K
Warmup Steps	–	2000	2000	2000	2000	100
Loss weight (CE : MSE)	1 : 0	0 : 1	0 : 1	0 : 1	0.1 : 1	0.1 : 1
Und resolution	$256^2 \rightarrow 4096^2$	–	–	–	$256^2 \rightarrow 4096^2$	$256^2 \rightarrow 4096^2$
Gen resolution	–	$256^2 \rightarrow 512^2$	$512^2 \rightarrow 2048^2$	$512^2 \rightarrow 2048^2$	$512^2 \rightarrow 2048^2$	$512^2 \rightarrow 2048^2$
Seq Length	32768	8192	16384	16384	32768	32768
Time shift	–	$\mu = -0.8, \sigma = 0.8$ in the logit-normal t -sampler				
Noise scale	–	$\sigma_0 \sqrt{N/N_0}, \sigma_0 = 1, N_0 = 64$ ($\sigma \in [1, 8]$ for $N \in [64, 4096]$)				
# Training tokens	0.75T	0.25T	0.25T	0.88T	1.19T	0.13T
Data sampling ratio						
Understanding data	1.00	0.00	0.00	0.00	0.33	0.33
Generation data	0.00	1.00	1.00	0.56	0.37	0.37
Editing data	0.00	0.00	0.00	0.37	0.24	0.24
Interleave data	0.00	0.00	0.00	0.07	0.06	0.06

Table 2 Training recipe of SenseNova-U1 from Stage 1 to Stage 4. Stage 2 is divided into three phases for generation pre-training.

In Phase I, we train on text-to-image data with resolutions ranging from 256×256 to 1024×1024 pixels, resizing images larger than 512×512 to 512×512 while preserving aspect ratios. This phase runs for 120K steps with a constant learning rate of 2×10^{-4} . In Phase II, we continue training on samples with resolution no smaller than 512^2 pixels, resizing images larger than 2048×2048 to 2048×2048 . This phase lasts for 60K steps, using a learning rate of 1×10^{-4} . In Phase III, we introduce image-editing, reasoning, and interleaved image-text generation data for an additional 120K steps, progressively expanding the model’s generative capabilities over diverse application scenarios. The entire training uses a cosine learning-rate schedule that decays from 1×10^{-4} to 2×10^{-5} . For the A3B variant, we apply a MoE balance loss coefficient of 5×10^{-3} to ensure balanced expert utilization in the generation branch.

Stage 3: Unified Mid-Training. Both branches are jointly trained end-to-end on a curated mixture of understanding and generation data, allowing the MoT backbone to develop coherent shared representations across various tasks.

The training mixture includes text-only & multimodal understanding, text-to-image generation, image editing, and interleaved image-text data, with sampling ratios of 0.33 : 0.37 : 0.24 : 0.06, respectively. We train the full model framework for 84K steps (nearly converges within 40K steps with <80M data) with a constant learning rate of 2×10^{-5} . For joint optimization, we set the loss weights in Eq.(1) to $\lambda_1 = 0.1$ and $\lambda_2 = 1.0$. Specifically for the A3B variant, we apply a MoE balance loss coefficient of 1×10^{-3} to both the generation and understanding branches.

Stage 4: Unified Supervised Fine-Tuning. The full model is fine-tuned on high-quality, instruction-following data spanning both understanding and generation tasks, including multimodal dialogue, image generation, editing, and interleaved data. This stage sharpens instruction alignment and task-specific performance across modalities.

We use the same data mixture as in Stage 3, covering understanding, text-to-image generation, image editing, and interleaved image-text data. The full model is further fine-tuned for 9k steps with a cosine learning-rate schedule decaying from 2×10^{-5} to 0. We retain the same loss weights as in Stage 3, with $\lambda_1 = 0.1$ and $\lambda_2 = 1.0$ in Eq.(1).

Stage 5: Post Training for T2I Generation. We present the post-training recipe for SenseNova-U1 via an initial round of text-to-image generation training, which leverages reinforcement learning (RL) following Flow-GRPO [80] to improve generation quality, and employs Distribution Matching Distillation [157] to enhance efficiency.

Dynamic Resolution Warmup. Since different resolutions exhibit significant reward variance, we introduce a warmup strategy to improve training stability. The candidate resolution set is constructed from aspect ratios $\{1:1, 16:9, 9:16, 3:2, 2:3\}$ and target image areas $\{1536^2, 2048^2\}$, each assigned a base sampling probability p_i . We assign each resolution a difficulty score $d_i \in [0, 1]$ based on its aspect ratio and pixel count, and gate the sampling probabilities as:

$$\hat{p}_i = p_i \cdot \text{clamp}\left(\frac{\min(e/E_{\text{warm}}, 1) - d_i}{\delta} + 1, 0, 1\right), \quad (7)$$

where e is current epoch, E_{warm} is warmup duration, and $\delta = 0.3$ is a smoothing margin. The gated probabilities are then renormalized for sampling, starting from easier configurations and progressively incorporating more challenging ones.

Reward Function Design. We employ three reward components in our RL training pipeline as follows:

- **Text Rendering Reward.** For text-rendering tasks, we use PaddleOCR [22] to extract text \hat{T} from generated images and compare it with the ground-truth T^* . The reward is based on Intersection-over-Union (IoU):

$$r_{\text{ocr}} = \frac{|\mathcal{C}(\hat{T}) \cap \mathcal{C}(T^*)|}{|\mathcal{C}(\hat{T}) \cup \mathcal{C}(T^*)|}, \quad (8)$$

where $\mathcal{C}(\cdot)$ denotes the Counter of texts, and \cap, \cup represent multi-set intersection and union, respectively.

- **Style Following Reward.** For prompts with explicit style constraints, we use a VLM judge [47] to assess whether the generated image follows the specified style. The judge assigns a discrete score $s \in \{1, 2, 3, 4\}$, linearly mapped to a style reward $r_{\text{sty}} \in [0, 1]$, where 0 indicates a complete mismatch and 1 denotes a perfect match.
- **Aesthetic Reward.** For aesthetic and preference alignment, we use Human Preference Score (HPSv3) [93] as an aesthetic quality reward. Given a generated image \mathbf{x} and its prompt p , the aesthetic reward $r_{\text{aes}} = H_{\text{v3}}(\mathbf{x}, p)$, where r_{aes} denotes the image-text preference score and H_{v3} is its corresponding scorer.

(i) *Text Rendering RL Training.* The training is conducted using text-rendering prompts in both English and Chinese, with the text rendering reward r_{ocr} serving as the optimization signal. The model is trained for 600 epochs with a learning rate of 1×10^{-5} and a KL coefficient $\beta = 0.01$. In each epoch, we sample $N = 48$ prompts, and for each prompt generate $K = 16$ images (i.e., 48×16 samples in total) using 10-step flow matching, with a guidance scale of 4.0 and a noise level of 0.7. A dynamic-resolution warmup is applied for the first $E_{\text{warm}} = 200$ epochs.

(ii) *Unified General RL Training.* This stage further improves image generation quality through interleaved multi-reward training: text rendering, style following, and general visual quality. To balance these heterogeneous objectives, we organize the training data into two reward groups and adopt an interleaved training strategy, where the active reward group alternates every training epoch. The two reward groups are defined as follows:

- **Group 1: Text Rendering & Style Following.** This group uses text-rendering prompts with style constraints. The composite reward combines text accuracy and overall style, following:

$$r = r_{\text{ocr}} + \lambda_{\text{sty}} \cdot r_{\text{sty}}, \quad (9)$$

where λ_{sty} controls the relative weight of the style following reward.

- **Group 2: Human Preference and Aesthetics.** This group uses general image generation prompts and adopts r_{aes} to assess aesthetics. Notably, it tends to favor darker backgrounds, with limited benefit to OCR performance.

We interleave the two reward groups at every epoch. The coefficient λ_{sty} is set to 0.5. The 8B variant is trained for 1,600 epochs, with all other hyperparameters consistent with the previous stage. The A3B variant is trained for 200 epochs and leaves room for further improvement. Here, we freeze the understanding branch, the last three transformer layers, and the MLP head of the generation branch to mitigate grid artifacts. This issue likely arises because the final FFN layer and the MLP head model disjoint 32×32 pixel patches independently. *A promising direction for future work is to replace the MLP head with PixelShuffle modules followed by two convolutional layers to further alleviate this issue.*

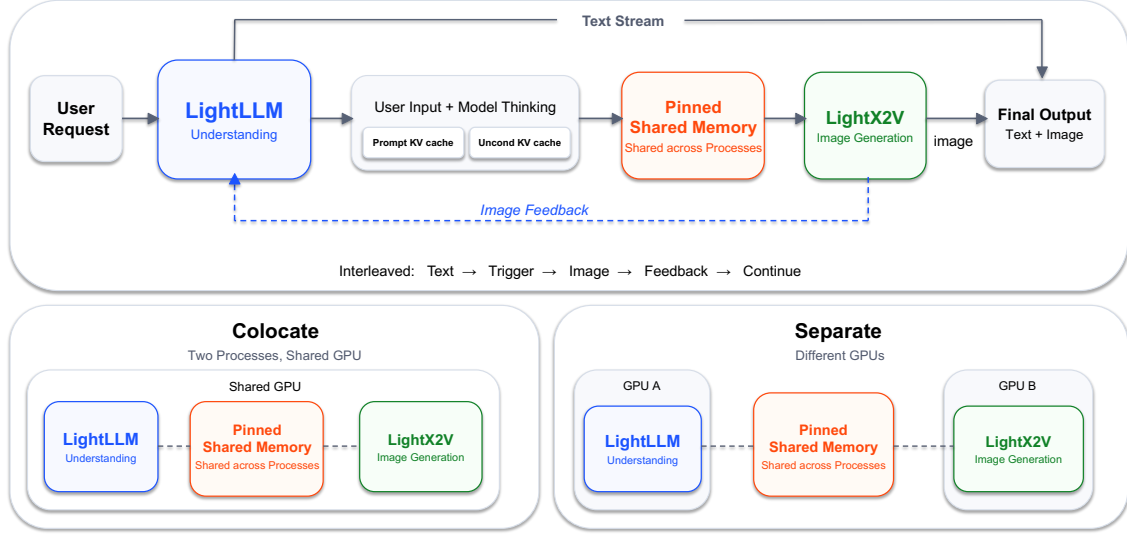


Figure 5 Disaggregated inference architecture of SenseNova-U1. LIGHTLLM serves multimodal understanding, text streaming, and control flow, while LIGHTX2V serves image generation. The two engines exchange generation state through pinned shared memory, enabling independent scheduling, parallelism, and resource allocation.

Stage 6: CFG & Step Distillation. We employ distribution matching distillation (DMD2) [157] to reduce the number of function evaluations (NFE) for image synthesis from 100 to 8 for impressive efficiency. It involves three models: a generator G to be distilled, a fake flow model F that estimates the score of the evolving generative distribution, and a teacher T that models the score of the target data distribution. All three are initialized from the teacher model.

During distillation, only the generation branch is optimized, including the MoT parameters as well as the patch encoding and decoding layers. The distillation process is performed in a unified setting, where text-to-image, editing, and interleaved data are jointly used for training. For editing and interleaved generation, backward simulation uses ground-truth images as references. We follow the hyperparameter settings of Phased DMD [36]. The generator G is updated once every five updates of F . G is optimized with AdamW optimizer at a learning rate of 2×10^{-6} , betas of (0.0, 0.999), and weight decay of 0.01, while F uses a learning rate of 4×10^{-7} with the same optimizer settings. Backward simulation employs an Euler solver with a timestep shift of 3.0, and the CFG scale is set to 4.0.

3.5 Inference Infrastructure

Although SenseNova-U1 is exposed to users as a unified multimodal model, its understanding and generation pathways have different inference characteristics. The understanding path is dominated by multimodal prefill, autoregressive text decoding, streaming, batching, and control-flow management, while the image generation path is dominated by iterative pixel-space denoising with different parallelism and memory-access patterns. Serving both paths in a single monolithic runtime would unnecessarily couple their scheduling policies, parallelization strategies, and resource budgets.

Disaggregated Deployment. We adopt a disaggregated inference architecture using two specialized open-source engines: LIGHTLLM [44, 97] for multimodal understanding, text streaming, and request orchestration, and LIGHTX2V [98] for image generation. The two engines exchange generation state through pinned shared memory and optimized transfer kernels, preserving a unified API abstraction while allowing each execution path to be independently optimized.

This design brings three practical benefits. First, it enables different parallelization strategies for different workloads: the understanding engine uses LLM-oriented Tensor Parallelism (TP), while the generation engine uses diffusion-oriented strategies, e.g., Classifier-Free Guidance Parallelism (CFG Parallelism) or Sequence Parallelism (SP). Second, it supports independent resource allocation, including separate GPU groups, memory budgets, and batching policies. Third, it improves operational isolation, so text-heavy and image-heavy traffic can be scaled, profiled, and tuned independently.

The infrastructure supports both separate and colocate deployments. In separate mode, LIGHTLLM and LIGHTX2V run on distinct GPU groups, which is preferred in production because it provides clear bottleneck attribution and independent




Figure 6 Hybrid attention pattern for unified multimodal prefill. Text rows follow the standard causal mask, while image rows can attend to the full preceding text prefix and the entire image span. The serving kernel preserves the causal fast path for pure-text blocks and only expands the key range for blocks that contain image tokens.

scaling. In colocate mode, the two engines run as separate processes on the same GPU group, which is useful for lightweight validation, smaller hardware configurations, or deployment scenarios where the image generation workload is substantially higher than the understanding workload. For 2048×2048 image generation with SenseNova-U1-8B-MoT, both modes support a TP2+CFG2 configuration. In separate mode, the per-step latencies on 5090 and L40S GPUs are 0.415 and 0.443 seconds, respectively. Because the generation-stage key-value cache is provided by the understanding module, text-to-image generation and image editing share similar runtime characteristics.

Hybrid Attention Kernel. A key systems challenge in unified multimodal prefill is the hybrid attention pattern: text rows remain causal, while image rows attend to the full text prefix and the full image span. This preserves standard autoregressive behavior for text tokens while allowing bidirectional interaction within image tokens.

To support this efficiently, we introduce an optional `b_image_token_end` in the attention kernel. The kernel makes the masking decision at the M-block level. If an M-block contains no image token, it keeps the standard causal key range. Otherwise, its key range is extended to the image-span end, so image rows can attend to the text prefix and the full image span. This design preserves the causal fast path for pure-text blocks and only introduces extra computation for blocks containing image tokens. We implement this mechanism in both a Triton kernel and a modified FlashAttention3 backend. The Triton version is easier to integrate, while the FlashAttention3 version provides higher throughput.

4 Data Construction

4.1 Understanding Data Organization

Pre-training Stage. The training corpus comprises large-scale web text, image-text pairs, and interleaved multimodal documents, organized into four categories: image-text pairs (32%), captions (17%), infographic understanding (14%), and pure text (37%). The data curation pipeline includes four stages: cross-source deduplication, content and safety filtering, image quality filtering, and CLIP-ratio-balanced re-captioning to ensure balanced alignment across the corpus.

Mid-training Stage. This stage is primarily drawn from internal SenseNova V6.5 datasets, spanning four categories: General (39.2%), Agent and Spatial (22.3%), Knowledge Reasoning (19.3%), and Pure Text (19.2%). The General category is further divided into general visual question answering (26.6%), multi-turn dialogue (26.4%), captioning (20.3%), OCR (18.6%), and multi-image understanding (8.2%), while Knowledge Reasoning includes knowledge-oriented (12.0%) and reasoning-oriented (7.2%) data. To ensure quality and diversity, we adopt a three-stage curation pipeline in Figure 7: distribution-balanced sampling, prompt augmentation, and multi-criteria filtering.

(i) **Distribution-Balanced Sampling.** We adopt a two-stage process to extract a diverse subset from the initial pool. First, CLIP-based diversity sampling clusters visual embeddings via K -means and samples uniformly across clusters to improve long-tail coverage. This is followed by attribute profiling, which evaluates each sample along perceptual and semantic dimensions and applies stratified sampling to ensure balanced representation across attribute tiers.

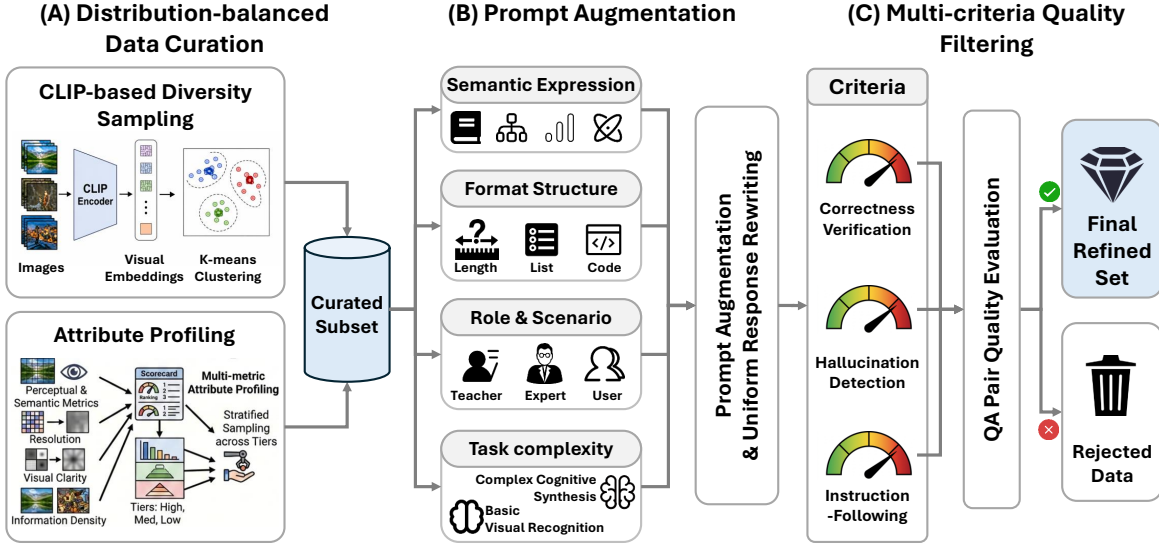


Figure 7 Data processing pipeline for the understanding corpus. Large-scale multimodal instruction data are curated across ten vertical domains through a systematic process consisting of distribution-balanced data curation, prompt augmentation, and multi-criteria quality filtering, producing a high-quality and diverse data corpus for the midtraining process.

(ii) **Prompt Augmentation.** To improve the diversity and complexity of training instructions, we augment the initial prompts along four dimensions: semantic expression, format and structural constraints, role and scenario, and task complexity, ranging from perceptual recall to compositional reasoning. After prompt augmentation, all answers are uniformly regenerated to ensure consistent quality and stylistic coherence across the whole corpus.

(iii) **Multi-Criteria Filtering.** To ensure dataset fidelity, we employ an automated model-based scoring pipeline to evaluate each question-answering (QA) pair across three dimensions: correctness verification against ground-truth annotations, hallucination detection to penalize visually unsupported fabrications, and instruction-following assessment to measure alignment with specified constraints such as formatting and persona.

Supervised Fine-Tuning. The final SFT corpus is organized along fine-grained, capability-atomic dimensions to enable precise control over the supervision mixture. The distribution spans spatial intelligence ($\sim 15\%$), general multimodal understanding ($\sim 13\%$), reasoning ($\sim 12\%$), general NLP ($\sim 11\%$), OCR and document analysis ($\sim 11\%$), agentic function calling ($\sim 10\%$), long-context conversation ($\sim 8\%$), code ($\sim 6\%$), multi-turn dialogue ($\sim 4\%$), complex compositional understanding ($\sim 4\%$), and supplementary data covering additional capabilities for the remaining proportion.

Rather than recollecting data from scratch, we refine the midtraining candidate pool with a dual emphasis on quality and difficulty. For quality-oriented selection, we reuse the multi-criteria filtering framework from midtraining, scoring each candidate across visual fidelity, instruction clarity, response correctness, reasoning quality, and safety, while increasing the sampling proportion of high-scoring examples relative to midtraining. For difficulty-oriented reconstruction, we rebalance supervision along three axes: composing longer and structurally richer instances by concatenating short samples into long-context, multi-image, and multi-turn settings; applying rejection sampling for reasoning-intensive domains to retain examples in the intermediate difficulty regime where learning is most effective; and rewriting under-specified queries to inject explicit constraints on output format, stylistic attributes, and target granularity.

4.2 Generation Data Organization

To balance broad coverage with high fidelity, we curate a large-scale generation corpus spanning text-to-image and image-editing data, organized into four domains: Nature, Design, People, and Synthetic. As shown in Figure 8, the distribution is carefully balanced across domains while preserving a pronounced long tail of natural, synthetic, and text-rich content. All samples are processed through the unified pipeline in Figure 9, combining low-level filtering, deduplication, VLM-based captioning, and quality-aware filtering to ensure a consistent standard of quality.




Figure 8 Data distribution of SenseNova-U1’s training corpus. From left to right, the four sunburst charts depict the hierarchical composition of the *Understanding*, *Text-to-Image*, *Editing*, and *Interleaved* datasets. The inner ring shows top-level categories and their proportions, while the outer ring breaks them down into fine-grained subclasses. Together, they illustrate the diversity of natural images, synthetic content, and text-rich samples across all four capabilities.

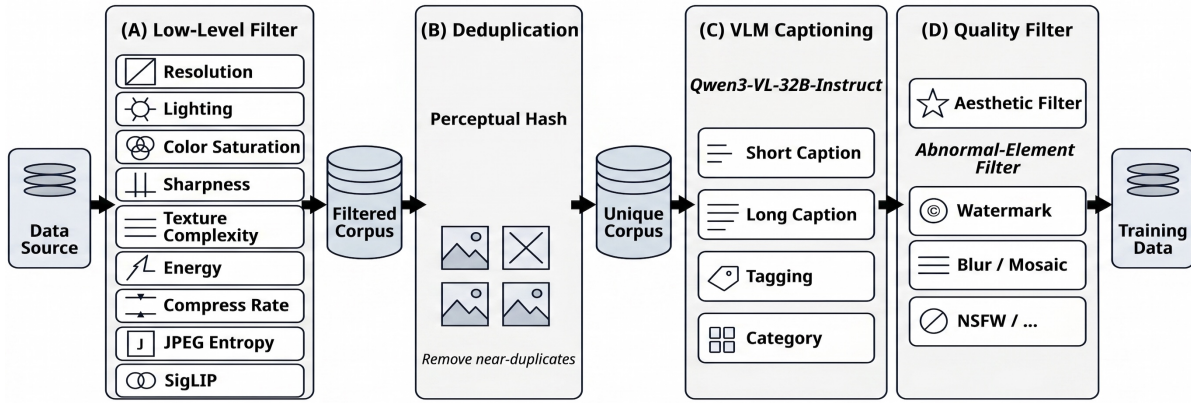


Figure 9 Data processing pipeline for the generation corpus. The same four-stage flow, *i.e.*, low-level filtering, deduplication, VLM captioning, and quality filtering, is applied to T2I and image editing to ensure high-quality and diverse generation data.

Text-to-Image Data. The entire corpus is composed of Nature ($\sim 40.5\%$), People ($\sim 26.7\%$), and Design ($\sim 20.7\%$), and further enriched with complex infographics and bilingual text rendering data, along with a long tail of fine-grained subclasses such as posters, charts, and cityscapes. This diversity provides broad visual coverage, structured layouts, and text-intensive scenarios, which foster strong visual priors, aesthetic quality, and robust text-rendering ability.

Image Editing Data. The editing corpus is primarily sourced from web-scale data. For in-domain data, it reflects similar diversity at content and operation levels. For content level, natural scenes ($\sim 52.3\%$) and human subjects ($\sim 14.7\%$) dominate real-world coverage, with the remainder consisting of infographic and synthetic edits. For operation level, the data spans subject addition and removal, background and color changes, identity transfer, motion manipulation, portrait editing, compositing, and reasoning-driven transformations. Beyond standard filtering, each editing pair is further validated by decomposing its instruction into dynamic objectives that specify what should change and what must remain unchanged. These objectives are jointly verified with a static physical-consistency constraint against source image.

Interleaved Data. To strengthen interleaved reasoning and generation, we construct a compact vision–text corpus in which sequences alternate between text and images to form coherent multimodal narratives [23, 148]. The corpus spans four complementary categories: Video, Lifestyle, Infographics, and Reasoning domains, each targeting distinct capabilities, and is built under a unified pipeline of preprocessing, task-specific synthesis, and post-processing that jointly verifies text quality, image quality, image–text consistency, and trajectory-level correctness. As shown in Figure 8, lifestyle data dominates at roughly 44%, including tutorials (26%), daily-life scenarios (14%), and picture books (4%). Infographics account for about 29%, providing dense supervision for text-rich page synthesis. Video contributes around 19%, capturing temporal continuity and world dynamics. Reasoning comprises approximately 8% and represents the most reasoning-intensive subset, as each sample includes an explicit chain-of-thought trace.

Benchmark	SenseNova-U1 8B-Think	Qwen3VL 8B-Think	Qwen3.5 9B	SenseNova-U1 30BA3B-Think	Qwen3VL 30BA3B-Think	Qwen3.5 35BA3B	Gemma4 26BA4B	LongCat-Next 68BA3B
<i>STEM & Reasoning</i>								
MMMU [159]	74.78	74.10	78.40	80.55	76.00	81.40	76.56	70.60
MMMU-Pro [160]	67.69	60.40	70.10	72.83	63.00	75.10	73.80	60.30
MathVista _{mini} [87]	84.20	81.40	85.70	85.30	81.90	86.20	72.70	83.10
MathVision [132]	75.82	62.70	78.90	79.63	65.70	83.90	68.95	64.70
<i>General VQA</i>								
MMBench-EN [82]	90.25	87.50	90.10	91.59	88.90	91.50	91.68	–
MMStar [16]	78.27	75.30	79.70	80.92	75.50	91.90	76.93	69.30
<i>OCR</i>								
InfoVQA [94]	82.46	86.00	90.76	83.04	86.00	94.22	–	83.30
OCRBench-v2 [37]	61.30	61.55	66.54	68.64	61.50	73.71	–	58.90
AI2D [51]	91.74	84.90	90.20	92.23	86.90	92.60	86.04	–
OCRBench [83]	82.10	81.90	89.20	91.90	83.90	91.00	86.30	86.50
<i>Hallucination</i>								
HallusionBench [49]	67.75	65.40	69.30	68.95	66.00	67.90	–	–
<i>Visual Reasoning</i>								
BabyVision [15]	25.00	17.78	25.80	31.70	18.60	29.60	11.34	–
TiR [68]	28.15	22.30	31.90	29.30	22.50	42.30	24.19	–
<i>Spatial Intelligence</i>								
VSI-Bench [150]	62.66	56.61*	55.67*	56.90	51.56*	58.10*	32.91	–
ViewSpatial [65]	56.19	47.25	48.19	58.52	47.37	50.78	41.68	–
MindCube-Tiny [156]	62.01	43.17	57.59	70.86	40.86	63.46	48.84	–
3DSR-Bench [91]	64.88	54.48	56.77	62.96	55.55	66.60	53.61	–

Table 3 Quantitative evaluation results on multimodal understanding benchmarks. For spatial intelligence, we adopt EASI [9] as the standard evaluation, using 32 input frames on VSI-Bench for all models. We observe that Qwen variants require 128-frame inputs to reach their best performance; we report these results separately and mark them with an asterisk (*).

5 Experiments

5.1 Main Results

5.1.1 Image Understanding

We evaluate SenseNova-U1 on various multimodal understanding and reasoning benchmarks, covering perception-centric understanding, multimodal reasoning, OCR recognition, visual-centric reasoning, and spatial intelligence. Note that we evaluate multimodal and pure-text understanding using EvalScope, following the LLM-as-a-judge paradigm with the `llm_recall` strategy and `gpt-4o-mini-2024-07-18` as the judge model. All experiments use a standardized inference setup with `temperature = 0.6`, `top_p = 0.95`, `top_k = 20`, and `repetition_penalty = 1.00`. To support long-context multimodal reasoning, we set the maximum sequence length to 40,960 tokens and the request timeout to 600 seconds, and enable internal reasoning via `enable_thinking: true`.

General Understanding. We report multimodal tasks including multimodal reasoning (MMMU [159], MMMU-Pro [160], MathVista [87], MathVision [132]), general VLM understanding (MMBench-EN [82], MMStar [16]), OCR (InfoVQA [94], OCRBench [83], OCRBench-v2 [37], AI2D [51]), hallucination detection (HallusionBench [49]), and advanced visual reasoning (BabyVision [15], TiR [68]). For fair comparison, we prioritize official results when available and otherwise re-evaluate baselines using vLLM under their recommended settings within the EvalScope framework.

In Table 3, our native SenseNova-U1 achieves strong performance, even without particular reinforcement learning for understanding domains. It consistently outperforms strong baselines such as Qwen3VL-8B [2] built on the same LLM [149] on multimodal reasoning benchmarks, while showing clear advantages in mathematical reasoning, highlighting the effectiveness of the encoder-free architecture. In text-rich understanding, our models obtain notable gains over both similarly sized and larger competitors [2, 27, 125]. On general vision–language benchmarks and hallucination evaluation, SenseNova-U1 matches or surpasses leading models while preserving robust and grounded predictions. Overall, compared to Qwen3.5 [106], SenseNova-U1 can achieve competitive performance across a wide range of tasks while establishing a new frontier in efficient training and modeling through encoder-free architectures.

Benchmark	SenseNova-U1 8B-Think	Qwen3VL 8B-Think	Qwen3.5 9B	SenseNova-U1 30BA3B-Think	Qwen3VL 30BA3B-Think	Qwen3.5 35BA3B	Gemma4 26BA4B
<i>Knowledge</i>							
MMLU-Pro [136]	81.44	77.30	82.50	84.04	80.50	85.30	84.89
MMLU-Redux [41]	87.61	88.80	91.10	89.44	90.90	93.30	91.91
C-Eval [55]	84.40	83.88	88.20	85.89	87.29	90.20	82.39
SuperGPQA [33]	49.67	51.20	58.20	59.71	56.40	63.40	61.88
<i>Instruction Following</i>							
IFEval [171]	91.13	83.20	91.50	92.39	81.70	91.90	–
IFBench [164]	67.01	29.93	64.50	79.79	34.69	70.20	25.51
<i>Agent</i>							
Tau2 [153]	71.70	31.65	79.10	75.39	46.40	81.20	68.20
Claw eval [154]	58.90	21.70	65.40	58.50	22.10	36.50	60.60

Table 4 Quantitative evaluation results on language understanding benchmarks.

Spatial Intelligence. We employ EASI [9] toolkit to evaluate performance across key spatial intelligence benchmarks, covering capabilities such as metric measurement, spatial relations, perspective-taking, and comprehensive reasoning [10]. SenseNova-U1 demonstrates remarkable performance on VSI-Bench [150], ViewSpatial [65], MindCube-Tiny [156], and 3DSR-Bench [91], highlighting strong spatial intelligence across both high-level reasoning and low-level geometric representation. These results suggest that native end-to-end multimodal modeling not only benefits semantic and compositional spatial understanding, but also preserves fine-grained structural and geometric perception essential for robust spatial reasoning. Note that the EASI protocol uses 32 input frames on VSI-Bench; we adopt this standard for all models except the Qwen variants, which require 128-frame input for the best performance.

5.1.2 Text Understanding

General Understanding. We evaluate pure-text capabilities across two categories. For knowledge-intensive reasoning, we report results on MMLU-Pro [136], MMLU-Redux [41], C-Eval [55], and SuperGPQA [33], covering broad academic knowledge and advanced reasoning across disciplines and languages. For instruction following, we evaluate on IFEval [171] and IFBench [164], which measure adherence to complex and constraint-heavy instructions.

As shown in Table 4, SenseNova-U1 demonstrates particularly strong gains in instruction following, consistently outperforming the Qwen3.5 series on both IFEval and IFBench. These results indicate that native encoder-free architectures can effectively handle complex, constraint-heavy instructions. It surpasses the Qwen3VL series and narrows the gap with the Qwen3.5 series on MMLU-Pro and SuperGPQA, reflecting strong academic knowledge and professional-level reasoning capabilities. Besides, strong text knowledge capability is further evidenced by consistent trends on C-Eval and MMLU-Redux. These results demonstrate that our native models using fewer training resources can rival the top-tier encoder-based architecture, further proving its effectiveness with efficient data-scaling capability.

Agentic Function. We evaluate the agentic capabilities of SenseNova-U1 on two complementary benchmarks: τ^2 -Bench [153] and Claw-Eval [154]. Specifically, τ^2 -Bench focuses on end-to-end task completion across domains such as Retail, Airline, and Telecom, requiring sustained interactions, adaptive tool use, and environment-aware decision making. Claw-Eval instead emphasizes trustworthy agent behavior through trajectory-aware evaluation over task completion, safety, and robustness, using execution traces, audit logs, and repeated-trial verification.

As shown in Table 4, SenseNova-U1 demonstrates powerful multi-turn reasoning and agent capabilities across both τ^2 -Bench and Claw-Eval. Notably, our A3B variant consistently outperforms existing multimodal baselines [2] and approaches the performance of substantially larger reasoning-oriented models [27, 106] despite using fewer active parameters. In particular, SenseNova-U1 exhibits strong trajectory-level reliability, coherent long-horizon interaction, and robust tool-use behavior in complex multi-step environments. We attribute the remaining gap to dense reasoning-focused baselines partly to deliberate training trade-offs that prioritize high-fidelity multimodal generation and interleaved vision-language capabilities. Nevertheless, these results highlight the effectiveness of our native multimodal framework in sustaining trustworthy and capable agent behavior across diverse real-world scenarios.

Model	# Params	Single Object	Two Object	Counting	Colors	Position	Attribute Binding	Overall↑
<i>Closed-source Models</i>								
GPT-Image-1 [100]	-	0.99	0.92	0.85	0.92	0.75	0.61	0.84
Seedream 4.0 [111]	-	0.99	0.92	0.72	0.91	0.76	0.74	0.84
Seedream 3.0 [38]	-	0.99	0.96	0.91	0.93	0.47	0.80	0.84
<i>Open-source Models</i>								
SenseNova-U1	8BA3B	1.00	0.96	0.89	0.91	0.92	0.77	0.91
SenseNova-U1	8B	1.00	0.96	0.92	0.92	0.91	0.76	0.91
Tuna [84]	7B	1.00	0.97	0.81	0.91	0.88	0.83	0.90
OneCAT [66]	9BA3B	1.00	0.96	0.84	0.94	0.84	0.80	0.90
NEO-unify [112]	8B	1.00	0.96	0.90	0.91	0.91	0.77	0.90
Mogao [74]	7B	1.00	0.97	0.83	0.93	0.84	0.80	0.89
Lumina-DiMOO [147]	8B	1.00	0.94	0.85	0.89	0.85	0.76	0.88
Qwen-Image [139]	20B	0.99	0.92	0.89	0.88	0.76	0.77	0.87
NEO-unify [112]	2B	0.99	0.92	0.89	0.86	0.77	0.76	0.87
Tuna-2 [85]	7B	0.99	0.96	0.80	0.91	0.84	0.76	0.87
InternVL-U [126]	1.7B	0.99	0.94	0.74	0.91	0.77	0.74	0.85
LongCat-Next [125]	68BA3B	-	-	-	-	-	-	0.84
Z-Image [7]	6B	1.00	0.94	0.78	0.93	0.62	0.77	0.84
BLIP3-o [14]	1.4B	-	-	-	-	-	-	0.84
X-Omni [42]	12B	0.98	0.95	0.75	0.91	0.71	0.68	0.83
BAGEL [28]	7B	0.99	0.94	0.81	0.88	0.64	0.63	0.82
Janus-Pro [18]	7B	0.99	0.89	0.59	0.90	0.79	0.66	0.80
OmniGen2 [141]	4B	1.00	0.95	0.64	0.88	0.55	0.76	0.80
UniWorld-V1 [75]	12B	0.99	0.93	0.79	0.89	0.49	0.70	0.80
Show-o2 [146]	7B	1.00	0.87	0.58	0.92	0.52	0.62	0.76
SD3-Medium [34]	2B	0.99	0.94	0.72	0.89	0.33	0.60	0.74
Emu3.5 [23]	32B	-	-	-	-	-	-	0.73
FLUX.1-dev [60]	12B	0.98	0.81	0.74	0.79	0.22	0.45	0.66

Table 5 Quantitative evaluation results on GenEval. The parameters of the generation component are denoted as # Params; A in this column denotes activated parameters, e.g., 8BA3B means 8B total generation parameters with 3B activated during inference.

5.1.3 Image Generation

We evaluate SenseNova-U1 on text-to-image generation from complementary perspectives: general generation, text-centric generation, complex infographic generation, and reasoning-centric generation. Together, these evaluations cover object-level composition, prompt following, long-text rendering, knowledge-informed generation, structured professional visual content creation, and more tightly coupled understanding-generation behaviors. Given our 32×32 downsampling ratios, we generate 2K images and downsample them to 1K for evaluation under comparable computational budgets.

General Generation. For general text-to-image generation, we adopt GenEval [43], DPG-Bench [53], OneIG-Bench [12], and TIIF-Bench [138]. These benchmarks examine object-level compositional generation, dense prompt following, and fine-grained overall capability from complementary perspectives. Across them, SenseNova-U1 remains highly competitive, showing that the native unified modeling paradigm does not sacrifice fundamental generation quality.

GenEval. GenEval [43] mainly evaluates compositional generation ability across object co-occurrence, counting, color, position, and attribute binding. As shown in Table 5, SenseNova-U1-A3B-MoT and SenseNova-U1-8B-MoT both achieve an overall score of 0.91, outperforming representative open-source models such as Qwen-Image at 0.87, Lumina-DiMOO at 0.88, and BAGEL at 0.82. More specifically, our model maintains consistently strong performance on Single Object, Two Object, Counting, Colors, and Position, which indicates stable object-level control and compositional consistency. Although the Attribute Binding score remains slightly below the 0.80 achieved by OneCAT and Mogao, SenseNova-U1 still attains the best overall result through a more balanced performance profile. These results show that SenseNova-U1 maintains strong and balanced compositional generation ability across object-level factors.

Model	# Params	Global	Entity	Attribute	Relation	Other	Overall↑
<i>Closed-source Models</i>							
Seedream 4.5 [6]	-	89.24	94.30	92.14	92.23	93.83	88.63
Nano-Banana-Pro [24]	-	91.00	92.85	91.56	92.39	89.93	87.16
GPT-Image-1 [100]	-	88.89	88.94	89.84	92.63	90.96	85.15
<i>Open-source Models</i>							
Qwen-Image [139]	20B	91.32	91.56	92.02	94.31	92.73	88.32
SenseNova-U1	8BA3B	94.19	92.05	91.05	93.22	93.44	88.14
Z-Image [7]	6B	93.39	91.22	93.16	92.22	91.52	88.14
JoyAI-Image [117]	16B	-	-	-	-	-	88.05
SenseNova-U1	8B	88.74	90.90	92.43	92.43	92.50	87.78
X-Omni [42]	12B	-	-	-	-	-	87.65
Tuna [84]	7B	90.42	91.68	90.94	91.87	90.73	86.76
NEO-unify [112]	8B	91.00	91.53	92.06	94.14	90.43	86.71
Tuna-2 [85]	7B	89.50	91.40	92.07	91.91	88.81	86.54
NEO-unify [112]	2B	89.49	92.87	91.26	92.29	92.13	86.54
Show-o2 [146]	7B	-	-	-	-	-	86.14
Lumina-DiMOO [147]	8B	81.46	92.08	88.98	94.31	82.00	86.04
InternVL-U [126]	1.7B	90.39	90.78	90.68	90.29	88.77	85.18
BAGEL [28]	7B	88.94	90.37	91.29	90.82	88.67	85.07
LongCat-Next [125]	68BA3B	-	-	-	-	-	84.66
OneCAT [66]	9BA3B	-	-	-	-	-	84.53
Mogao [74]	7B	-	-	-	-	-	84.33
Janus-Pro [18]	7B	86.90	88.90	89.40	89.32	89.48	84.19
SD3-Medium [34]	2B	87.90	91.01	88.83	80.70	88.68	84.08
FLUX.1-dev [60]	12B	74.35	90.00	88.96	90.87	88.33	83.84
Ovis-U1 [130]	1.2B	82.37	90.08	88.68	93.35	85.20	83.72
OmniGen2 [141]	4B	88.81	88.83	90.18	89.37	90.27	83.57
BLIP3-o [14]	1.4B	-	-	-	-	-	81.60
UniWorld-V1 [75]	12B	83.64	88.39	88.44	89.27	87.22	81.38

Table 6 Quantitative evaluation results on DPG-Bench. The parameters of the generation component are denoted as # *Params*; A in this column denotes activated parameters, e.g., 8BA3B means 8B total generation parameters with 3B activated during inference.

DPG-Bench. DPG-Bench [53] evaluates fine-grained instruction following using dense prompts across five dimensions: Global, Entity, Attribute, Relation, and Other. As shown in Table 6, SenseNova-U1 ranks among the top-performing models on this benchmark. Both SenseNova-U1-A3B-MoT and SenseNova-U1-8B-MoT remain highly competitive with leading specialized generation systems such as closed-source Seedream 4.5, open-source Qwen-Image, and Z-Image, despite being trained within a unified multimodal framework rather than optimized solely for image synthesis. Notably, SenseNova-U1-A3B-MoT achieves the highest Global score, underscoring the strong semantic planning, holistic scene composition, and long-range instruction consistency enabled by native end-to-end multimodal modeling.

OneIG-Bench. OneIG-Bench [12] provides a fine-grained evaluation of image generation quality across Alignment, Text, Reasoning, Style, and Diversity in both English and Chinese. As shown in Table 7 and Table 8, SenseNova-U1 remains highly competitive across both language settings, demonstrating strong multilingual generation capability within a unified framework. In particular, our models exhibit clear strengths in Alignment and Text understanding. SenseNova-U1-A3B-MoT achieves a leading Alignment score on the English benchmark, while SenseNova-U1-8B-MoT attains the best Text performance on both the English and Chinese tracks among all compared methods. These results underscore the effectiveness of native end-to-end multimodal modeling in preserving fine-grained text-image alignment, robust multilingual text rendering, and precise instruction-following under complex generation scenarios.

TIIF-Bench. TIIF-Bench [138] systematically evaluates image generation spanning attributes, relations, reasoning, style, and text. As shown in Table 9 and Table 10, SenseNova-U1 achieves consistently strong performance under

Model	# Params	Alignment	Text	Reasoning	Style	Diversity	Overall↑
<i>Closed-source Models</i>							
Gemini-2.5-Flash-Image [26]	-	0.878	0.894	0.346	0.450	0.182	0.550
GPT-Image-1 [100]	-	0.851	0.857	0.345	0.462	0.151	0.533
Seedream 3.0 [38]	-	0.818	0.865	0.275	0.413	0.277	0.530
Imagen4 [45]	-	0.857	0.805	0.338	0.377	0.199	0.515
Recraft V3 [108]	-	0.810	0.795	0.323	0.378	0.205	0.502
Kolors 2.0 [59]	-	0.820	0.427	0.262	0.360	0.300	0.434
Imagen3 [3]	-	0.843	0.343	0.313	0.359	0.188	0.409
<i>Open-source Models</i>							
Emu3.5 [23]	32B	0.902	0.994	0.345	0.427	0.151	0.564
SenseNova-U1	8B	0.882	0.969	0.330	0.396	0.166	0.549
SenseNova-U1	8BA3B	0.887	0.861	0.317	0.458	0.194	0.543
Qwen-Image [139]	20B	0.882	0.891	0.306	0.418	0.197	0.539
HiDream-I1-Full [8]	17B	0.829	0.707	0.317	0.347	0.186	0.477
SD3.5 Large [34]	8B	0.809	0.629	0.294	0.353	0.225	0.462
FLUX.1 [Dev] [60]	12B	0.786	0.523	0.253	0.368	0.238	0.434
BAGEL [28]	7B	0.769	0.244	0.173	0.367	0.251	0.361
BLIP3-o [14]	1.4B	0.711	0.013	0.223	0.361	0.229	0.307
Janus-Pro [18]	7B	0.553	0.001	0.139	0.276	0.365	0.267

Table 7 Quantitative evaluation results on OneIG-EN. The parameters of the generation component are denoted as *# Params*; *A* in this column denotes activated parameters, e.g., 8BA3B means 8B total generation parameters with 3B activated during inference.

Model	# Params	Alignment	Text	Reasoning	Style	Diversity	Overall↑
<i>Closed-source Models</i>							
Seedream 3.0 [38]	-	0.793	0.928	0.281	0.397	0.243	0.528
GPT-Image-1 [100]	-	0.812	0.650	0.300	0.449	0.159	0.474
Kolors 2.0 [59]	-	0.738	0.502	0.226	0.331	0.333	0.426
Gemini-2.5-Flash-Image [26]	-	0.825	0.276	0.298	0.427	0.198	0.337
<i>Open-source Models</i>							
Qwen-Image [139]	20B	0.825	0.963	0.267	0.405	0.279	0.548
SenseNova-U1	8BA3B	0.847	0.906	0.301	0.446	0.202	0.540
SenseNova-U1	8B	0.826	0.977	0.303	0.392	0.176	0.535
Emu3.5 [23]	32B	0.853	0.941	0.300	0.386	0.166	0.529
BAGEL [28]	7B	0.672	0.365	0.186	0.357	0.268	0.370
HiDream-I1-Full [8]	17B	0.620	0.205	0.256	0.304	0.300	0.337
BLIP3-o [14]	1.4B	0.608	0.092	0.213	0.369	0.233	0.303
Janus-Pro [18]	7B	0.324	0.148	0.104	0.264	0.358	0.240

Table 8 Quantitative evaluation results on OneIG-ZH. The parameters of the generation component are denoted as *# Params*; *A* in this column denotes activated parameters, e.g., 8BA3B means 8B total generation parameters with 3B activated during inference.

both short and long instruction settings. In particular, SenseNova-U1-8B-MoT attains the best overall results among all compared methods, while SenseNova-U1-A3B-MoT also remains highly competitive. These results suggest that SenseNova-U1 extends beyond accurate text rendering to more challenging text-centric generation scenarios that require jointly satisfying fine-grained textual constraints, compositional reasoning, and global instruction consistency.

Text-centric Generation. We evaluate text-centric generation, with a focus on long-text rendering, multi-region text generation, and complex text-conditioned instruction following. We conduct experiments on CVTG-2K [32] and

Model	# Params	Overall↑	Basic				Advanced						Design
			Avg	Attr	Rel	Rsn	Avg	ARel	ARsn	RRsn	Style	Text	RealW
Closed-source Models													
GPT-Image-1 [100]	-	89.15	90.75	91.33	84.57	96.32	88.55	87.07	87.22	85.59	90.00	89.83	89.73
Seedream 3.0 [38]	-	86.02	87.07	90.50	89.85	80.86	79.16	79.76	77.23	75.64	100.00	97.17	83.21
DALL-E 3 [5]	-	74.96	78.72	79.50	80.82	75.82	73.39	73.45	72.01	63.59	89.66	66.83	72.93
MidJourney v7 [95]	-	68.74	77.41	77.58	82.07	72.57	64.66	67.20	81.22	60.72	83.33	24.83	68.83
FLUX.1 [Pro] [60]	-	67.32	79.08	78.83	82.82	75.57	61.10	62.32	69.84	65.96	63.00	35.83	71.80
Open-source Models													
SenseNova-U1	8B	89.74	90.38	91.50	88.46	91.19	85.21	85.37	83.78	84.10	100.00	89.59	93.66
Emu3.5 [23]	32B	89.48	87.05	90.50	89.80	80.85	84.65	82.91	83.76	83.45	100.00	100.00	94.03
SenseNova-U1	8BA3B	89.25	88.28	90.50	86.83	87.50	87.60	87.48	86.85	87.03	96.67	90.05	90.30
Qwen-Image [139]	20B	86.14	86.18	90.50	88.22	79.81	79.30	79.21	78.85	75.57	100.00	92.76	90.30
FLUX.1 [dev] [60]	12B	71.09	83.12	87.05	87.25	75.01	65.79	67.07	73.84	69.09	66.67	43.83	70.72
SD 3 [34]	8B	67.46	78.32	83.33	82.07	71.07	61.46	61.07	68.84	50.96	66.67	59.83	63.23
Janus-Pro [18]	7B	66.50	79.33	79.33	78.32	80.32	59.71	66.07	70.46	67.22	60.00	28.83	65.84
Infinity [50]	8B	62.07	73.08	74.33	72.82	72.07	56.64	60.44	74.22	60.22	80.00	10.83	54.28
Show-o [145]	1.3B	59.72	73.08	74.83	78.82	65.57	53.67	60.95	68.59	66.46	63.33	3.83	55.02

Table 9 Quantitative evaluation results on TIIF testmini (short). Abbrev.: Avg = Average, Attr = Attribute, Rel = Relation, Rsn = Reasoning, ARel = Attribute+Relation, ARsn = Attribute+Reasoning, RRsn = Relation+Reasoning, RealW = Real World.

Model	# Params	Overall↑	Basic				Advanced						Design
			Avg	Attr	Rel	Rsn	Avg	ARel	ARsn	RRsn	Style	Text	RealW
Closed-source Models													
GPT-Image-1 [100]	-	88.29	89.66	87.08	84.57	97.32	88.35	89.44	83.96	83.21	93.33	86.83	93.46
Seedream 3.0 [38]	-	84.31	84.93	90.00	85.94	78.86	80.60	81.82	78.85	78.64	93.33	87.78	83.58
DALL-E 3 [5]	-	70.81	78.50	79.83	78.82	76.82	67.27	67.20	71.34	60.72	86.67	54.83	60.99
FLUX.1 [Pro] [60]	-	69.89	78.91	81.33	83.82	71.57	65.37	65.57	71.47	67.72	63.00	55.83	68.80
MidJourney v7 [95]	-	65.69	76.00	81.83	76.82	69.32	60.53	62.70	71.59	64.59	80.00	20.83	63.61
Open-source Models													
SenseNova-U1	8B	89.17	91.02	94.50	90.49	88.06	85.34	87.23	84.14	83.14	100.00	82.81	92.16
Emu3.5 [23]	32B	88.18	88.41	92.50	90.78	81.94	84.04	83.08	85.73	81.09	90.00	95.93	92.54
SenseNova-U1	8BA3B	87.36	87.95	92.50	89.45	81.90	84.39	83.88	86.52	81.44	96.67	82.81	91.04
Qwen-Image [139]	20B	86.83	87.22	91.50	90.78	79.38	80.88	78.94	81.69	78.59	100.00	89.14	91.42
FLUX.1 [dev] [60]	12B	71.78	78.65	83.17	80.39	72.39	68.54	73.69	73.34	71.59	66.67	52.83	71.47
SD 3 [34]	8B	66.09	77.75	79.83	78.82	74.07	59.56	64.07	70.34	57.84	76.67	20.83	67.34
Janus-Pro [18]	7B	65.02	78.25	82.33	73.32	79.07	58.82	56.20	70.84	59.97	70.00	33.83	60.25
Infinity [50]	8B	62.32	75.41	76.83	77.57	71.82	54.98	55.57	64.71	59.71	73.33	23.83	56.89
Show-o [145]	1.3B	58.86	75.83	79.83	78.32	69.32	50.38	56.82	68.96	56.22	66.67	2.83	50.92

Table 10 Quantitative evaluation results on TIIF testmini (long). Abbrev.: Avg = Average, Attr = Attribute, Rel = Relation, Rsn = Reasoning, ARel = Attribute+Relation, ARsn = Attribute+Reasoning, RRsn = Relation+Reasoning, RealW = Real World.

LongText-Bench [42]. Across these benchmarks, SenseNova-U1 shows consistently strong performance, indicating that it has reached a text-centric generation capability comparable to the strongest current text-to-image models.

CVTG-2K. CVTG-2K [32] evaluates complex text-centric generation with multiple text regions. As shown in Table 11, SenseNova-U1 performs well on this benchmark. Impressively, SenseNova-U1-8B-MoT achieves the best average word accuracy of 0.940, with consistently strong results across settings from 2 to 5 text regions. These results demonstrate that SenseNova-U1 can accurately render textual content under dense multi-region settings.

Model	# Params	NED	CLIPScore	Word Accuracy				Average↑
				2 regions	3 regions	4 regions	5 regions	
Closed-source Models								
Seedream 4.5 [6]	-	0.948	0.807	0.878	0.895	0.908	0.901	0.899
GPT-Image-1 [100]	-	0.948	0.798	0.878	0.866	0.873	0.822	0.857
Nano-Banana-Pro [24]	-	0.875	0.737	0.737	0.775	0.786	0.793	0.779
Open-source Models								
SenseNova-U1	8B	0.972	0.825	0.945	0.954	0.944	0.936	0.940
Emu3.5 [23]	32B	-	-	-	-	-	-	0.912
SenseNova-U1	8BA3B	0.944	0.824	0.884	0.883	0.883	0.875	0.881
JoyAI-Image [117]	16B	0.937	0.799	-	-	-	-	0.874
Z-Image [7]	6B	0.937	0.797	0.901	0.872	0.865	0.851	0.867
Qwen-Image [139]	20B	0.912	0.802	0.837	0.836	0.831	0.816	0.829
LongCat-Next [125]	68BA3B	-	-	-	-	-	-	0.764
InternVL-U [126]	1.7B	0.804	0.816	0.729	0.660	0.618	0.549	0.623
Lumina-DiMOO [147]	8B	0.805	0.831	0.723	0.646	0.571	0.505	0.590
FLUX.1-dev [60]	12B	0.688	0.740	0.609	0.553	0.466	0.432	0.497
BAGEL [28]	7B	0.657	0.779	0.498	0.391	0.332	0.291	0.356
Ovis-U1 [130]	1.2B	0.477	0.725	0.133	0.109	0.091	0.065	0.093

Table 11 Quantitative evaluation results on CVTG-2K. The parameters of the generation component are denoted as # Params.

Model	# Params	LongText-Bench-EN↑	LongText-Bench-ZH↑
<i>Closed-source Models</i>			
Seedream 4.5 [6]	-	0.989	0.987
Nano-Banana-Pro [24]	-	0.981	0.949
GPT-Image-1 [100]	-	0.956	0.619
<i>Open-source Models</i>			
SenseNova-U1	8B	0.979	0.962
Emu3.5 [23]	32B	0.976	0.928
JoyAI-Image [117]	16B	0.963	0.963
SenseNova-U1	8BA3B	0.950	0.955
Qwen-Image [139]	20B	0.943	0.946
Z-Image [7]	6B	0.935	0.936
LongCat-Next [125]	68BA3B	0.932	0.891
X-Omni [42]	12B	0.900	0.814
InternVL-U [126]	1.7B	0.738	0.860
NEO-unify [112]	2B	0.748	0.495
FLUX.1-dev [60]	12B	0.607	0.005
OmniGen2 [141]	4B	0.561	0.059
Lumina-DiMOO [147]	8B	0.437	0.047
BAGEL [28]	7B	0.373	0.310

Table 12 Quantitative evaluation results on LongText-Bench. A in # Params denotes activated parameters during inference.

LongText-Bench. LongText-Bench [42] primarily evaluates the accuracy and stability of long-text generation in both English and Chinese. As shown in Table 12, SenseNova-U1-8B-MoT achieves 0.979 on LongText-Bench-EN and 0.962 on LongText-Bench-ZH, while SenseNova-U1-A3B-MoT also reaches 0.950 and 0.955, despite it not being fully converged. These results indicate that our model can accurately render long-form text in both languages while maintaining high readability and semantic accuracy as text length and structural complexity increase.



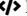
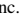



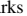

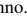
Model	# Params	Question Type										Overall	
		 Comp.	 Enc.	 Order	 Marks	 Anno.	 Axes	 Leg.	 Chart	 Title	 Deco.	Q-ACC↑	I-ACC
Closed-source Models													
Nano-Banana-Pro [24]	-	0.84	0.86	0.90	0.87	0.93	0.93	0.96	0.92	0.98	0.94	0.90	0.49
Seedream-4.5 [6]	-	0.34	0.37	0.47	0.48	0.70	0.70	0.81	0.68	0.95	0.84	0.61	0.06
GPT-Image-1.5 [101]	-	0.38	0.48	0.44	0.57	0.50	0.54	0.57	0.68	0.60	0.80	0.55	0.12
Nano-Banana [26]	-	0.18	0.31	0.27	0.44	0.54	0.57	0.52	0.60	0.65	0.81	0.48	0.02
P-Image [104]	-	0.08	0.15	0.19	0.27	0.36	0.28	0.54	0.43	0.58	0.68	0.34	0.00
Image-01 [96]	-	0.01	0.05	0.04	0.10	0.10	0.14	0.03	0.22	0.14	0.47	0.13	0.00
Open-source Models													
SenseNova-U1	8B	0.27	0.23	0.49	0.45	0.54	0.61	0.70	0.65	0.74	0.82	0.51	0.04
SenseNova-U1	8BA3B	0.17	0.22	0.33	0.41	0.36	0.49	0.56	0.60	0.55	0.78	0.42	0.02
Qwen-Image [139]	20B	0.10	0.13	0.19	0.29	0.43	0.37	0.51	0.48	0.56	0.78	0.36	0.01
Z-Image-Turbo [7]	6B	0.10	0.16	0.16	0.25	0.38	0.31	0.58	0.42	0.61	0.73	0.35	0.00
HiDream-I1 [8]	17B	0.01	0.03	0.03	0.10	0.07	0.14	0.10	0.26	0.19	0.20	0.11	0.00
FLUX.1-dev [60]	12B	0.00	0.03	0.01	0.08	0.06	0.06	0.01	0.24	0.09	0.39	0.10	0.00

Table 13 Quantitative evaluation results on IGenBench. The parameters of the generation component are denoted as # Params; A in this column denotes activated parameters, e.g., 8BA3B means 8B total generation parameters with 3B activated during inference.

Model	# Params	Layout	Attribute	Text	Knowledge	Average↑
<i>Closed-source Models</i>						
GPT-Image-2 [102]	-	88.5 / 95.3	81.9 / 91.0	83.9 / 92.6	74.2 / 90.9	82.1 / 92.5
Nano-Banana-Pro [24]	-	72.2 / 91.2	65.6 / 92.2	86.4 / 95.0	82.6 / 96.2	76.7 / 93.7
Nano-Banana-2.0 [46]	-	68.4 / 91.0	57.4 / 91.6	83.4 / 94.6	64.6 / 93.0	68.5 / 92.5
Seedream-5.0 [110]	-	67.6 / 89.0	42.4 / 77.2	43.4 / 75.6	41.8 / 75.2	48.8 / 79.2
GPT-Image-1.5 [101]	-	51.6 / 84.8	25.8 / 75.2	40.4 / 82.8	26.0 / 83.6	35.9 / 81.6
Seedream-4.5 [6]	-	35.4 / 71.6	22.4 / 62.8	41.4 / 72.4	21.4 / 58.2	30.1 / 66.2
Wan2.6-T2I [129]	-	46.4 / 80.6	16.6 / 60.6	12.6 / 52.6	12.2 / 41.0	21.9 / 58.7
Seedream-4.0 [111]	-	27.6 / 73.4	11.4 / 59.2	11.4 / 52.8	6.8 / 54.8	14.3 / 60.1
GPT-Image-1 [100]	-	21.4 / 60.2	6.8 / 48.6	8.6 / 41.0	7.8 / 60.0	11.2 / 52.4
<i>Open-source Models</i>						
SenseNova-U1	8B	61.6 / 81.6	47.5 / 72.8	46.3 / 74.6	3.5 / 17.9	39.7 / 61.7
SenseNova-U1	8BA3B	50.9 / 76.7	35.5 / 60.9	24.5 / 58.7	2.0 / 11.5	28.2 / 51.9
Emu3.5 [23]	32B	30.4 / 63.4	14.2 / 52.6	7.0 / 33.6	1.2 / 11.0	13.2 / 40.2
HunyuanImage-3.0 [11]	80BA13B	27.8 / 65.0	13.8 / 53.6	10.2 / 39.6	0.0 / 2.0	13.0 / 40.1
Z-Image [7]	6B	26.8 / 69.2	2.6 / 47.6	2.8 / 45.0	0.6 / 13.2	8.2 / 43.8
Qwen-Image-2512 [139]	20B	22.2 / 70.6	1.2 / 47.8	1.8 / 39.2	0.0 / 6.4	6.3 / 41.0
FLUX.2-dev [61]	32B	17.2 / 67.8	1.2 / 49.2	1.0 / 43.0	0.0 / 8.2	4.9 / 42.0
Qwen-Image [139]	20B	10.4 / 51.2	0.2 / 22.2	0.6 / 17.6	0.0 / 4.4	2.8 / 23.8
GLM-Image [162]	7B	5.4 / 43.2	0.0 / 13.4	0.2 / 4.4	0.0 / 0.4	1.4 / 15.3
LongCat-Image [124]	6B	2.4 / 35.8	0.2 / 11.6	0.0 / 4.4	0.0 / 0.0	0.7 / 13.0
X-Omni [42]	12B	2.0 / 22.8	0.0 / 5.6	0.0 / 8.0	0.0 / 1.4	0.5 / 9.4
BAGEL [28]	7B	0.6 / 12.8	0.0 / 1.6	0.0 / 0.0	0.0 / 0.2	0.2 / 3.7

Table 14 Quantitative evaluation results on BizGenEval. Each cell reports hard / easy testset scores.

Complex Infographic Generation. Here we evaluate complex infographic and commercial visual content generation on IGenBench [120] and BizGenEval [71]. Compared with general text-to-image generation, these infographic tasks are substantially more challenging, since the model must not only generate correct text and visual elements, but also satisfy structured layouts, chart expression, and multiple semantic constraints at the same time.

Model	# Params	Cultural	Time	Space	Biology	Physics	Chemistry	Overall↑
<i>Closed-source Models</i>								
Nano-Banana-Pro [24]	-	0.89	0.80	0.89	0.88	0.86	0.85	0.87
GPT-Image-1 [100]	-	0.81	0.71	0.89	0.83	0.79	0.74	0.80
Seedream 4.0 [111]	-	0.78	0.73	0.85	0.79	0.84	0.67	0.78
<i>Open-source Models</i>								
SenseNova-U1-SFT (w/ CoT)	8BA3B	0.81	0.77	0.84	0.81	0.84	0.82	0.81
SenseNova-U1-SFT (w/ CoT)	8B	0.78	0.73	0.82	0.80	0.85	0.77	0.78
SenseNova-U1-SFT	8BA3B	0.73	0.69	0.80	0.73	0.82	0.69	0.74
NEO-unify (w/ CoT) [112]	8B	0.73	0.67	0.79	0.70	0.75	0.66	0.72
BAGEL (w/ CoT) [28]	7B	0.76	0.69	0.75	0.65	0.75	0.58	0.70
SenseNova-U1-SFT	8B	0.65	0.66	0.82	0.68	0.81	0.66	0.69
Qwen-Image [139]	20B	0.63	0.62	0.76	0.60	0.72	0.39	0.63
BLIP3-o [14]	1.4B	-	-	-	-	-	-	0.62
NEO-unify (w/ CoT) [112]	2B	0.59	0.54	0.68	0.57	0.69	0.50	0.59
InternVL-U (w/ CoT) [126]	1.7B	0.55	0.57	0.74	0.51	0.72	0.46	0.58
Emu3.5 [23]	32B	-	-	-	-	-	-	0.58
LongCat-Next [125]	68BA3B	-	-	-	-	-	-	0.57
UniWorld-V1 [75]	12B	0.53	0.55	0.73	0.45	0.59	0.41	0.55
FLUX.1-dev [60]	12B	0.48	0.58	0.62	0.42	0.51	0.35	0.50
BAGEL [28]	7B	0.44	0.52	0.65	0.42	0.62	0.41	0.49
NEO-unify [112]	8B	-	-	-	-	-	-	0.47
InternVL-U [126]	1.7B	0.37	0.51	0.68	0.39	0.62	0.39	0.46
SD3-Medium [34]	2B	0.43	0.50	0.52	0.41	0.53	0.33	0.45
Ovis-U1 [130]	1.2B	0.36	0.46	0.64	0.35	0.52	0.28	0.42
NEO-unify [112]	2B	-	-	-	-	-	-	0.41
Lumina-DiMOO [147]	8B	0.35	0.43	0.59	0.31	0.49	0.34	0.40
Janus-Pro [18]	7B	0.30	0.37	0.49	0.36	0.42	0.26	0.35

Table 15 Quantitative evaluation results on WISE. The parameters of the generation component are denoted as # *Params*; A in this column denotes activated parameters, e.g., 8BA3B means 8B total generation parameters with 3B activated during inference.

IGenBench. IGenBench [120] evaluates the reliability of text-to-infographic generation, requiring models to jointly satisfy textual, chart, data, and structural constraints. As shown in Table 13, SenseNova-U1 achieves the strongest performance among open-source models, substantially outperforming Qwen-Image and Z-Image-Turbo while remaining competitive with several closed-source systems. These results demonstrate the strong infographic generation reliability of SenseNova-U1, although truly robust infographic synthesis remains challenging for current models.

BizGenEval. BizGenEval [71] evaluates visual generation in real-world commercial scenarios across dimensions, including Layout, Attribute, Text, and Knowledge. As shown in Table 14, SenseNova-U1 achieves the best hard-split average among all open-source models while remaining competitive on the easy split. In particular, our models exhibit strong layout planning, attribute control, and text rendering capabilities, highlighting the potential of native unified multimodal modeling for complex professional visual content generation under multi-constraint settings.

Reasoning-centric Generation. We further evaluate reasoning-driven WISE [99]. It examines whether a model can effectively utilize internal world knowledge and combine it with reasoning during image generation, thereby handling tasks involving cultural commonsense, temporal understanding, spatial understanding, and scientific knowledge.

WISE. As shown in Table 15, SenseNova-U1 displays a clear advantage on WISE. Even without chain-of-thought (CoT), SenseNova-U1-A3B-MoT-SFT substantially outperforms representative open-source models such as Qwen-Image, BAGEL, and InternVL-U, while CoT further boosts performance to a level competitive with closed-source systems. A similar trend is observed for the 8B variant, indicating that reasoning-enhanced generation scales consistently across model sizes. Notably, our models excel in Cultural, Biology, Physics, and Chemistry, particularly in science-oriented tasks requiring explicit knowledge retrieval and multi-step reasoning. This suggests that SenseNova-U1 effectively translates structured semantic reasoning into accurate and knowledge-consistent visual generation.

Model	# Params	Add	Adjust	Extract	Replace	Remove	Background	Style	Hybrid	Action	Overall↑
<i>Closed-source Models</i>											
UniWorld-V2 [72]	-	4.29	4.44	4.32	4.69	4.72	4.41	4.91	3.83	4.83	4.49
Nano-Banana-Pro [24]	-	4.44	4.62	3.42	4.60	4.63	4.32	4.97	3.64	4.69	4.37
Seedream 4.5 [6]	-	4.57	4.65	2.97	4.66	4.46	4.37	4.92	3.71	4.56	4.32
Seedream 4.0 [111]	-	4.33	4.38	3.89	4.65	4.57	4.35	4.22	3.71	4.61	4.30
Nano-Banana [26]	-	4.62	4.41	3.68	4.34	4.39	4.40	4.18	3.72	4.83	4.29
GPT-Image-1 [100]	-	4.61	4.33	2.90	4.35	3.66	4.57	4.93	3.96	4.89	4.20
FLUX.1 Kontext [Pro] [62]	-	4.25	4.15	2.35	4.56	3.57	4.26	4.57	3.68	4.63	4.00
<i>Open-source Models</i>											
Qwen-Image-Edit-2511 [139]	20B	4.54	4.57	4.13	4.70	4.46	4.36	4.89	4.16	4.81	4.51
LongCat-Image-Edit [124]	6B	4.44	4.53	3.83	4.80	4.60	4.33	4.92	3.75	4.82	4.45
Emu3.5 [23]	32B	4.61	4.32	3.96	4.84	4.58	4.35	4.79	3.69	4.57	4.41
FLUX.2 [Dev] [61]	32B	4.50	4.18	3.83	4.65	4.65	4.31	4.88	3.46	4.70	4.35
Qwen-Image-Edit-2509 [139]	20B	4.32	4.36	4.04	4.64	4.52	4.37	4.84	3.39	4.71	4.35
Z-Image-Edit [7]	6B	4.40	4.14	4.30	4.57	4.13	4.14	4.85	3.63	4.50	4.30
Qwen-Image-Edit [139]	20B	4.38	4.16	3.43	4.66	4.14	4.38	4.81	3.82	4.69	4.27
Ovis-U1 [130]	1.2B	3.99	3.73	2.66	4.38	4.15	4.05	4.86	3.43	4.68	3.97
SenseNova-U1	8BA3B	4.03	4.10	2.73	4.27	3.91	4.06	4.92	2.87	4.29	3.91
SenseNova-U1	8B	3.83	4.15	3.12	4.32	3.26	4.18	4.85	3.03	4.41	3.90
InternVL-U (w/ CoT) [126]	1.7B	4.24	3.80	2.58	4.36	3.51	3.92	4.69	3.00	4.31	3.82
FLUX.1 Kontext [Dev] [62]	12B	4.12	3.80	2.04	4.22	3.09	3.97	4.51	3.35	4.25	3.71
InternVL-U [126]	1.7B	4.13	3.40	2.27	4.13	3.39	3.84	4.77	3.03	4.05	3.67
OmniGen2 [141]	4B	3.57	3.06	1.77	3.74	3.20	3.57	4.81	2.52	4.68	3.44
UniWorld-V1 [75]	12B	3.82	3.64	2.27	3.47	3.24	2.99	4.21	2.96	2.74	3.26
BAGEL [28]	7B	3.56	3.31	1.70	3.30	2.62	3.24	4.49	2.38	4.17	3.20
Step1X-Edit [81]	12B	3.88	3.14	1.76	3.40	2.41	3.16	4.63	2.64	2.52	3.06
ICEdit [165]	12B	3.58	3.39	1.73	3.15	2.93	3.08	3.84	2.04	3.68	3.05
OmniGen [144]	3.8B	3.47	3.04	1.71	2.94	2.43	3.21	4.19	2.24	3.38	2.96

Table 16 Quantitative evaluation results on ImgEdit. The parameters of the generation component are denoted as # Params; A in this column denotes activated parameters, e.g., 8BA3B means 8B total generation parameters with 3B activated during inference.

5.1.4 Image Editing

We evaluate SenseNova-U1 on image editing across three distinct settings: single-image editing, multi-image editing, and reasoning-driven editing. Collectively, these evaluations probe the model’s ability to follow complex instructions, perform realistic human-guided modifications, and synthesize knowledge-aware visual content.

General Editing. We further evaluate SenseNova-U1 on general image editing. Compared with text-to-image generation, image editing requires the model not only to follow textual instructions, but also to preserve the original image content, structure, and visual style while performing precise local or global modifications. This places substantially higher demands on instruction following, content preservation, and fine-grained controllability.

ImgEdit. ImgEdit [155] provides a fine-grained evaluation of image editing across diverse dimensions, including Add, Adjust, Replace, Remove, Background, Style, Hybrid, and Action. As shown in Table 16, both SenseNova-U1-A3B-MoT and SenseNova-U1-8B-MoT achieve decent overall performance, outperforming existing open-source unified editing systems while remaining competitive with several specialized editing models. A performance gap nevertheless remains relative to the strongest dedicated editing approaches, particularly on complex hybrid edits and scenarios requiring precise content preservation under substantial transformations. *We attribute this gap primarily to limitations in the current editing data, which remains dominated by open-source resources and lacks sufficiently diverse editing pipelines and large-scale preference-aligned optimization.* Despite these limitations, the results indicate that SenseNova-U1 already provides a strong general-purpose foundation for image editing within a native unified framework. We expect future improvements to arise naturally from richer editing data, stronger editing-oriented supervision, and reinforcement learning strategies more closely aligned with human editing preferences and long-horizon editing objectives.

Model	# Params	GEdit-Bench-EN		
		G_SC	G_PQ	G_O \uparrow
Closed-source Models				
UniWorld-V2 [72]	-	8.39	8.02	7.83
Seedream 4.5 [6]	-	8.27	8.17	7.82
Nano-Banana-Pro [24]	-	8.10	8.34	7.74
Seedream 4.0 [111]	-	8.14	8.12	7.70
GPT-Image-1 [100]	-	7.85	7.62	7.53
Nano-Banana [26]	-	7.40	8.45	7.29
FLUX.1 Kontext [Pro] [62]	-	7.02	7.60	6.56
Open-source Models				
Qwen-Image-Edit-2511 [139]	20B	8.30	8.20	7.88
Longcat-Image-Edit [124]	6B	8.13	8.18	7.75
Emu3.5 [23]	32B	8.11	7.70	7.59
Z-Image-Edit [7]	6B	8.11	7.72	7.57
Qwen-Image-Edit [139]	20B	8.00	7.86	7.56
Qwen-Image-Edit-2509 [139]	20B	8.15	7.86	7.54
SenseNova-U1	8B	8.27	7.49	7.47
FLUX.2 [Dev] [61]	32B	7.84	8.06	7.41
SenseNova-U1	8BA3B	8.07	7.36	7.32
Step1X-Edit [81]	12B	7.66	7.35	6.97
BAGEL [28]	7B	7.36	6.83	6.52
OmniGen2 [141]	4B	7.16	6.77	6.41
FLUX.1 Kontext [Dev] [62]	12B	6.52	7.38	6.00
OmniGen [144]	3.8B	5.96	5.89	5.06
UniWorld-V1 [75]	12B	4.93	7.43	4.85

Table 17 Quantitative evaluation results on GEdit-Bench.

GEdit-Bench. GEdit-Bench [81] evaluates general instruction-based image editing with a stronger emphasis on overall editing quality and prompt consistency. As shown in Table 17, both SenseNova-U1-A3B-MoT and SenseNova-U1-8B-MoT achieve competitive performance against representative editing models such as Emu-3.5, Z-Image-Edit, and Qwen-Image-Edit. While specialized editing systems still retain an advantage in highly optimized editing workflows, SenseNova-U1 demonstrates strong generalization across diverse editing instructions, maintaining coherent semantics and stable visual consistency under a unified native modeling framework.

Reasoning-centric Editing. Beyond general editing ability, we further evaluate our SenseNova-U1 on reasoning-driven image editing. Compared with standard editing tasks, these scenarios are substantially more challenging because the model must first infer implicit temporal, causal, spatial, or logical relationships from the prompt instruction, and then translate the reasoning outcome into precise and visually consistent modifications.

RISEBench. RISEBench [166] mainly evaluates reasoning-informed image editing, covering four types of reasoning-centric editing tasks: Temporal, Causal, Spatial, and Logical, together with auxiliary metrics such as IR, AC, and VP. As shown in Table 18, SenseNova-U1-A3B-MoT-SFT achieves an overall score of 25.3 without CoT, substantially outperforming most open-source unified editing models. With CoT enabled, SenseNova-U1-A3B-MoT-SFT further improves to 30.0, reaching the best level among the open-source methods in our comparison. A similar trend also appears for SenseNova-U1-8B-MoT-SFT, which improves from 23.9 to 26.9, indicating that this pattern remains stable across model scales. A closer look shows that the gains from CoT are especially notable on dimensions that rely more heavily on explicit reasoning, such as Causal and Logical. For example, the Logical score of SenseNova-U1-A3B-MoT-SFT improves from 7.1 to 20.0. This suggests that the model can use an explicit reasoning process to better decompose complex editing goals. Overall, these results indicate that the advantage of SenseNova-U1 lies not only in executing edits, but also in performing the necessary understanding and inference before editing, which leads to stronger performance on reasoning-centric image editing than existing unified multimodal models.

Model	# Params	Temporal	Causal	Spatial	Logical	Overall↑	IR	AC	VP
<i>Closed-source Models</i>									
GPT-Image-1.5 [101]	-	54.1	60.0	62.0	21.2	50.0	69.7	92.5	94.9
Nano-Banana-Pro [24]	-	41.2	61.1	48.0	37.6	47.2	77.0	85.5	94.4
Nano-Banana [26]	-	25.9	47.8	37.0	18.8	32.8	61.2	86.0	91.3
GPT-Image-1 [100]	-	34.1	32.2	37.0	10.6	28.9	62.8	80.2	94.9
GPT-Image-1-mini [100]	-	24.7	28.9	33.0	9.4	24.4	54.1	71.5	93.7
Gemini-2.0-Flash-exp [25]	-	8.2	15.5	23.0	4.7	13.3	48.9	68.2	82.7
Seedream 4.0 [111]	-	12.9	12.2	11.0	7.1	10.8	58.9	67.4	91.2
Gemini-2.0-Flash-pre [25]	-	10.6	13.3	11.0	2.3	9.4	49.9	68.4	84.9
<i>Open-source Models</i>									
SenseNova-U1-SFT (w/ CoT)	8BA3B	24.7	46.7	28.0	20.0	30.0	63.2	84.1	87.4
SenseNova-U1-SFT (w/ CoT)	8B	31.8	33.3	27.0	15.3	26.9	60.8	86.6	88.2
SenseNova-U1-SFT	8BA3B	25.9	41.1	26.0	7.1	25.3	57.4	82.6	85.4
SenseNova-U1-SFT	8B	22.4	33.3	27.0	11.8	23.9	58.2	84.1	82.4
Qwen-Image-Edit-2511 [139]	20B	21.2	18.9	31.0	4.7	19.4	49.9	71.0	91.5
BAGEL (w/ CoT) [28]	7B	5.9	17.8	21.0	1.2	11.9	45.9	73.8	80.1
InternVL-U (w/ CoT) [126]	1.7B	4.7	7.8	1.8	5.9	9.4	43.9	64.4	79.7
Qwen-Image-Edit-2509 [139]	20B	4.7	10.0	17.0	2.4	8.9	37.2	66.4	86.9
BAGEL [28]	7B	2.4	5.6	14.0	1.2	6.1	36.5	53.5	73.0
FLUX.1-Kontext-Dev [62]	12B	2.3	5.5	13.0	1.2	5.8	26.0	71.6	85.2
InternVL-U [126]	1.7B	3.5	2.2	5.0	3.5	3.6	35.6	52.7	75.9
Ovis-U1 [130]	1.2B	1.2	3.3	4.0	2.4	2.8	33.9	52.7	72.9
Lumina-DiMOO [147]	8B	2.4	1.1	4.0	1.2	2.2	34.0	50.7	72.3
Step1X-Edit [81]	12B	0.0	2.2	2.0	3.5	1.9	25.1	41.5	73.5
OmniGen [144]	3.8B	1.2	1.0	0.0	1.2	0.8	22.0	32.6	55.3
Emu2 [119]	37B	1.2	1.1	0.0	0.0	0.5	22.6	38.2	78.3

Table 18 Quantitative evaluation results on RISEBench. The parameters of the generation component are denoted as # Params; A in this column denotes activated parameters, e.g., 8BA3B means 8B total generation parameters with 3B activated during inference.

5.1.5 Interleaved Generation

We further evaluate the model on interleaved generation and unified reasoning, examining whether understanding and generation can reinforce each other within a single framework. These evaluations cover open-ended interleaved generation, generation-assisted multimodal reasoning, and the bidirectional synergy between understanding and generation.

Interleaved Generation. This is no longer to generate a single image in one shot, but to alternately produce text and images in an open-ended output process while maintaining semantic coherence, cross-modal consistency, and overall completeness across multiple generation steps. Here we adopt OpenING [172] and VBVR-Image (Preview) [133].

OpenING. OpenING [172] evaluates the overall quality of open-ended interleaved image-text generation across dimensions including completeness, quality, richness, correctness, human alignment, image-text coherence, and multi-step consistency. As shown in Table 19, SenseNova-U1 demonstrates consistently strong performance under this challenging setting. In particular, SenseNova-U1-A3B-MoT-SFT with CoT achieves the best overall score of 9.16, while SenseNova-U1-8B-MoT-SFT with CoT reaches 9.07, outperforming representative systems such as Nano Banana, Wan-Weaver, and GPT-4o+DALL-E3. These results suggest that SenseNova-U1 not only produces high-quality unimodal outputs, but also maintains strong semantic coherence, long-range consistency, and instruction fidelity across interleaved multimodal generation trajectories. More importantly, the strong performance under multi-step image-text interaction indicates that our unified framework can effectively coordinate generation and reasoning within a shared native modeling space, rather than treating image synthesis and language generation as isolated processes.

VBVR-Image (Preview). VBVR-Image [133] is a recently introduced benchmark derived from VBVR that evaluates reasoning behaviors emerging through visual generation. It extends this setting from video generation to interleaved image generation, where models must generate images to solve visual reasoning tasks such as maze navigation, pattern

Model	# Params	Complete	Quality	Richness	Correct	Human Align.	IT Coherency	Multi-step	Overall↑
<i>Closed-source Models</i>									
Nano-Banana [26]	-	9.34	8.58	8.00	9.17	8.88	9.27	8.70	8.85
Wan-Weaver [148]	-	9.41	8.32	8.03	8.90	8.69	8.78	8.56	8.67
GPT-4o [56]+DALL-E3 [5]	-	8.66	8.01	7.42	7.98	8.77	8.15	8.38	8.20
Gemini [123]+Flux [60]	12B	7.58	7.26	6.48	7.03	7.98	6.98	7.33	7.23
<i>Open-source Models</i>									
SenseNova-U1-SFT (w/ CoT)	8BA3B	9.27	9.11	8.45	9.16	9.40	9.55	9.21	9.16
SenseNova-U1-SFT (w/ CoT)	8B	9.14	9.03	8.43	9.08	9.35	9.40	9.09	9.07
SEED-X [40]	17B	5.65	6.07	4.92	5.77	7.03	5.72	5.72	5.84
Emu3 [135]	8B	5.90	5.96	5.52	5.43	6.47	5.66	5.37	5.76
Anole [21]	7B	6.27	6.02	5.28	5.06	6.91	4.90	5.81	5.75
SEED-LLaMA [39]	14B	5.59	5.50	4.61	4.59	6.5	4.43	5.13	5.19
VILA-U [143]	7B	5.60	5.14	4.68	4.78	5.69	4.74	4.79	5.06
Show-o [145]	1.3B	4.37	4.79	3.83	3.76	5.78	4.04	4.33	4.41
MiniGPT-5 [169]	0.86B	3.91	4.5	3.61	3.63	5.51	3.56	4.10	4.12
NExT-GPT [142]	1.3B	3.89	4.25	3.35	3.61	5.35	3.32	3.85	3.95

Table 19 Quantitative evaluation results on OpenING. The parameters of the generation component are denoted as *# Params*; A in this column denotes activated parameters, e.g., 8BA3B means 8B total generation parameters with 3B activated during inference.

Model	# Params	Overall↑	In-Domain by Category						Out-of-Domain by Category					
			Avg.	Abst.	Know.	Perc.	Spat.	Trans.	Avg.	Abst.	Know.	Perc.	Spat.	Trans.
Closed-source Models														
Nano-Banana-2 [46]	-	62.3	61.1	64.4	49.6	78.6	51.1	53.5	63.6	83.2	63.1	61.4	61.0	54.2
GPT-Image-2 [102]	-	60.1	57.9	62.0	46.9	70.1	45.8	58.1	62.3	82.9	63.3	61.9	48.2	55.4
Open-source Models														
SenseNova-U1-SFT	8BA3B	68.9	73.9	76.1	72.2	80.8	66.5	68.9	64.0	89.2	61.9	66.2	75.5	41.8
SenseNova-U1-SFT	8B	68.8	70.8	73.9	60.5	81.6	65.8	66.3	66.8	83.7	65.4	69.9	68.0	51.9
VBVR-ThinkMorph [48]	7B	63.0	64.7	67.3	61.3	66.1	64.1	62.0	61.4	84.9	55.5	60.4	69.5	48.0
ThinkMorph [48]	7B	38.7	37.0	45.8	26.8	36.3	36.7	33.3	40.5	57.3	46.0	42.5	55.6	19.3
VBVR-BAGEL [28]	7B	36.5	37.0	44.1	28.6	36.7	32.2	38.8	36.0	51.9	40.3	35.6	31.9	26.6
BAGEL [28]	7B	29.1	32.1	35.5	26.6	33.1	27.3	36.1	26.0	45.3	20.1	20.5	22.0	25.9

Table 20 Quantitative evaluation results on VBVR-Image (Preview). VBVR- prefix is the model tuned on VBVR training set.

discovery, and spatial inference. As the benchmark is currently available only in a preview version, we report results on this subset in Table 20. The results show that SenseNova-U1 exhibits strong reasoning capability within the generation process itself, outperforming both competitive in-domain trained baselines and several powerful proprietary systems. These findings suggest that the unified native framework of SenseNova-U1 supports not only high-quality generation, but also reasoning behaviors that can emerge and be executed directly through multimodal generation trajectories.

Unified Reasoning. Beyond one-way generation ability, we further investigate whether the model can achieve genuine bidirectional synergy between understanding and generation. Unlike conventional evaluations that assess these capabilities in isolation, the Generation-aids-Understanding (GaU) components of Uni-MMMU [173] and RealUnify [114] explicitly examine whether generation can enhance understanding (GEU), and conversely, whether understanding can improve generation (UEG) within a unified multimodal framework.

Uni-MMMU. Uni-MMMU [173] (GaU) evaluates whether generation can actively assist multimodal understanding and reasoning. As shown in Table 21, SenseNova-U1 achieves strong performance under this setting. SenseNova-U1-8B-MoT-SFT attains a GaU average of 35.0, substantially outperforming unified baselines such as BAGEL, OmniGen2, and Ovis-U1, while SenseNova-U1-A3B-MoT-SFT also achieves a competitive score of 32.6. This suggests that the generation branch of SenseNova-U1 can provide meaningful support for multimodal reasoning, highlighting the synergistic interaction between generation and understanding within our unified framework.

Model	# Params	Jigsaw-T	Maze-T	Sliding-T	Geometry-T	Avg↑
<i>Closed-source Models</i>						
Nano-Banana [26]	-	57.0	4.7	0.0	47.8	27.4
<i>Open-source Models</i>						
SenseNova-U1-SFT	8B	87.3	28.6	0.0	24.2	35.0
SenseNova-U1-SFT	8BA3B	88.0	34.0	1.2	7.1	32.6
BAGEL [28]	7B	48.0	0.0	1.2	32.8	20.5
Ovis-U1 [130]	1.2B	53.0	0.0	0.0	3.5	14.1
OmniGen2 [141]	4B	48.0	0.0	0.0	5.7	13.4
Qwen-Image-Edit [139]	20B	43.3	0.7	0.0	8.5	13.1

Table 21 Quantitative evaluation results on UniMMM (Generation aids Understanding, GaU). For all models, we report text accuracy (T) for all tasks. For multi-step tasks (Maze, Sliding Puzzle), we report sample-level accuracy.

RealUnify. RealUnify [114] further evaluates bidirectional synergy through both Understanding Enhances Generation (UEG) and Generation Enhances Understanding (GEU). As shown in Table 22, SenseNova-U1 demonstrates clear advantages under both settings. We can observe that SenseNova-U1-8B-MoT-SFT achieves the best overall average of 52.4, including 55.7 on Avg-UEG and 47.5 on Avg-GEU, while SenseNova-U1-A3B-MoT-SFT also attains a strong overall average of 50.5. These results suggest that SenseNova-U1 achieves genuine synergy between understanding and generation, rather than merely colocating the two capabilities within a shared backbone.

Model	# Params	Understanding Enhances Generation							Generation Enhances Understanding					Avg↑
		WK	CR	MR-I	LR	SR	C2I	Avg-UEG	MR-II	MT	AF	CN	Avg-GEU	
SenseNova-U1-SFT	8B	88	68	33	45	54	46	55.7	36	63	51	40	47.5	52.4
SenseNova-U1-SFT	8BA3B	81	56	41	40	55	44	52.8	30	65	58	35	47.0	50.5
BAGEL [28]	7B	74	80	26	37	29	40	47.7	38	25	52	28	35.8	42.9
Ovis-U1 [130]	1.2B	59	71	30	34	17	25	39.3	38	25	31	24	29.5	35.4
OneCAT [66]	9BA3B	64	65	20	27	31	27	39.0	29	26	26	36	29.2	35.1
UniPic2 [137]	2B	62	72	30	38	26	15	40.5	28	24	27	16	23.8	33.8
UniWorld-V1 [75]	12B	56	59	26	37	24	9	35.2	33	25	36	20	28.5	32.5
OmniGen2 [141]	4B	55	60	26	28	20	6	32.5	42	24	38	19	30.8	31.8
ILLUME+ [54]	3B	52	62	22	25	26	7	32.3	27	20	38	25	27.5	30.4

Table 22 Quantitative evaluation results on RealUnify. We report the step-wise inference results; for SenseNova-U1, the GEU results are obtained via a single interleaved process, while the UEG results are obtained via text-to-image generation inference.

5.2 Ablation Studies

We conduct a series of ablation studies about SenseNova-U1, focusing on three key questions: whether the encoder-free design preserves both semantic and pixel-level representations, whether it synergizes effectively with the MoT backbone while minimizing intrinsic modality conflict, and whether it exhibits strong data-scaling efficiency.

5.2.1 Native Encoder-Free Design Preserves Both Semantic and Pixel Representations

Image Reconstruction. As reported in Table 23, our previously released NEO-unify (2B) [112] attains 31.56 PSNR and 0.85 SSIM on MS-COCO 2017 [77] after only 90K pretraining steps, approaching the 31.56 PSNR and 0.93 SSIM achieved by the FLUX.1-dev VAE. This result suggests that the native near-lossless interface is capable of retaining both high-level semantic information and fine-grained visual details without depending on pretrained vision encoders or latent autoencoders. Representative reconstruction examples are presented in Figure 10.

Image Editing. For editing tasks, NEO-unify (2B) [112] routes all conditional contexts through the understanding branch, while the generation branch directly synthesizes the target images. Despite freezing the understanding branch throughout training, the model still exhibits strong editing capability, together with substantially improved token

Method	Downsampling Ratio	Resolution	PSNR \uparrow	SSIM \uparrow
SDXL VAE [103]	8	256	25.76	0.76
SD3 VAE [86]	8	256	29.47	0.86
FLUX.1-dev VAE [62]	8	256	30.43	0.93
RAE (DINOv2-B) [168]	14	256	18.36	0.47
UniFlow (DINOv2-L) [161]	14	256	30.66	0.94
UAE (DINOv2-L) [35]	14	256	32.74	0.94
FLUX.1-dev VAE [62]	8	512	31.56	0.93
Neo-unify (2B)	32	512	31.56	0.85

Table 23 Reconstruction performance with a frozen understanding branch on MS-COCO 2017.




Figure 10 Reconstructing out-domain images with 2B NEO-unify under a frozen understanding branch.

efficiency. Using only public text-to-image and editing datasets, it achieves an ImgEdit score of 3.32 after an initial 60K-step mixed training process. Representative editing examples on ImgEdit prompts are shown in Figure 11.

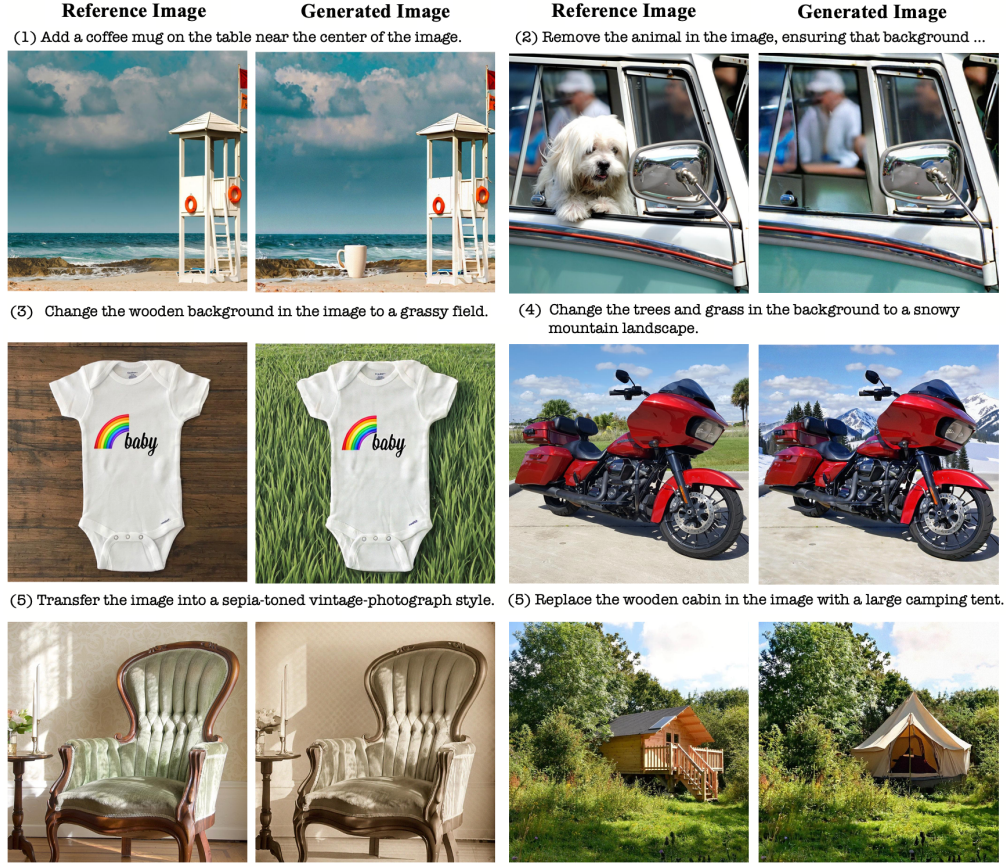


Figure 11 Validating ImgEdit prompts with 2B NEO-unify under a frozen understanding branch.






Figure 12 Understanding-generation co-training with 8B-MoT backbone. GEdit-Bench scores are normalized to 0–100 scale.

5.2.2 Understanding and Generation Synergize with Native MoT Backbone


Starting from pretrained dual branches, we jointly optimize all components during mid-training. Even with low data ratios and small understanding loss weights, understanding remains stable while generation converges rapidly. As shown in Figure 12, the two capabilities co-evolve effectively within the MoT backbone with minimal intrinsic conflict.




(a) Scores on the image generation task (DPG-Bench).



(b) Scores on the reasoning image generation task (WISE).



(c) Scores on the image editing task (GEdit-Bench).



(d) Scores on the reasoning image editing task (RISEBench).

Figure 13 Data-scaling curves of the 8B-MoT backbone. We set the resolution as 512 and 1,024 for DPG-Bench and others.

5.2.3 Native Multimodal Architecture Shows High Data-Scaling Efficiency

We begin with web-scale pretraining, followed by mid-training and supervised fine-tuning using diverse, high-quality data corpora spanning both understanding and generation tasks. As shown in Figure 13, the model delivers strong data-scaling efficiency, with both generation quality and understanding-generation synergy improving steadily as training data scale. Overall, these results further validate the advantages of the native MoT design from three complementary perspectives: preserving both semantic structure and pixel-level fidelity, reducing intrinsic conflict between understanding and generation, and enabling efficient scaling across data scale and multimodal generation tasks.

5.3 Visualization Results

In addition to quantitative evaluations, we provide qualitative visualizations to illustrate the behaviors of SenseNova-U1 in complex multimodal scenarios. We focus on representative cases spanning text-to-image generation, infographic generation, image editing, interleaved image-text generation, visual understanding, and agentic multimodal interaction, covering both general and reasoning-intensive settings. Additional showcases are available at <https://github.com/OpenSenseNova/SenseNova-U1/blob/main/docs/showcases.md>.

Vision-Language-Action. We first present representative video-based action reasoning examples in Figure 14. For each example, four frames are uniformly sampled from the input video and arranged in a single row to illustrate the temporal progression of the manipulation process. These cases show that SenseNova-U1 can capture action-relevant visual dynamics across time, maintain coherent visual understanding under embodied settings, and reason about object states, manipulation trajectories, and task progression from sparse temporal observations.




Figure 14 Visualizing vision-language-action behaviors of SenseNova-U1 on robotic manipulation videos.

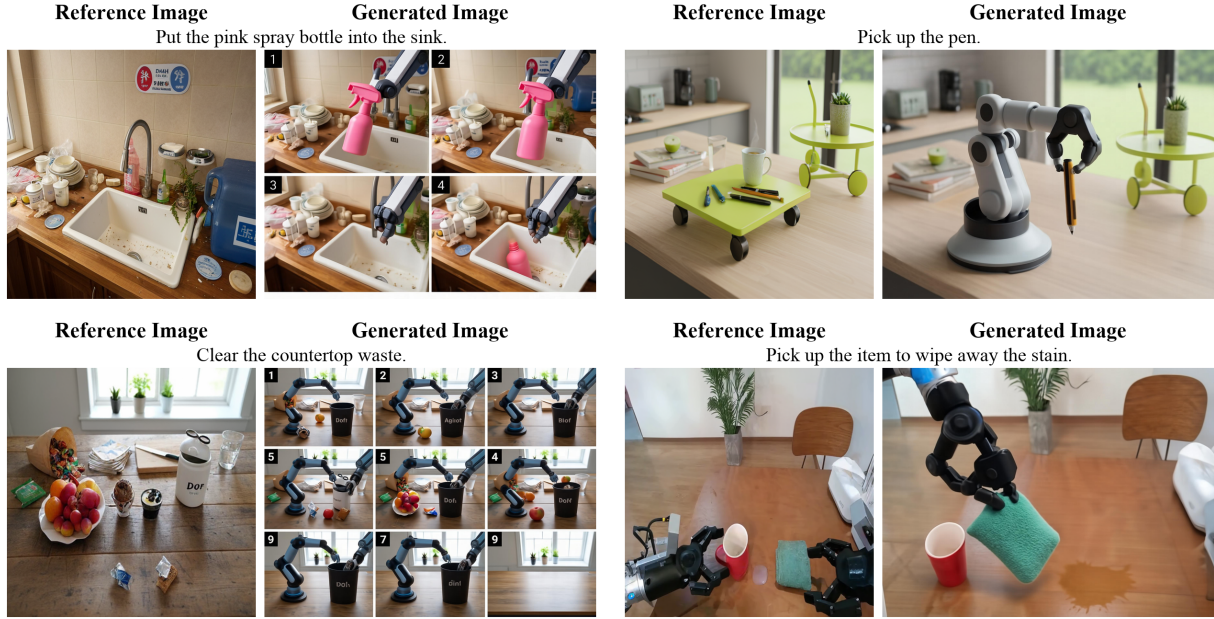


Figure 15 Visualizing world-modeling predictions from the robotic arm view with SenseNova-U1.

World Modeling. We further visualize the world modeling capability of SenseNova-U1 in Figure 15. Given an input image and an action-oriented instruction, the model is required to predict the corresponding visual outcome. For readability, simplified instructions are shown in the figure, while the original prompts are used during inference. The selected examples demonstrate that SenseNova-U1 can translate structured action instructions into plausible visual state transitions while preserving overall scene consistency and object coherence.

6 Conclusion

We present a unified multimodal foundation model in which understanding, generation, and reasoning emerge within a single native architecture rather than through the coordination of separate systems. Across a broad range of tasks, the model exhibits strong capabilities in vision-language perception, semantic reasoning, high-fidelity generation, and interleaved multimodal interaction, suggesting that a shared representation can simultaneously support analytical and creative intelligence. More fundamentally, our results point toward a broader transition in multimodal AI. Rather than merely aligning isolated modalities, unified models begin to internalize a coherent abstraction of the world itself, enabling perception, imagination, and decision-making to arise within a shared latent space. Early advances in vision-language-action models and world modeling further indicate a path from passive understanding toward embodied, goal-directed intelligence. We believe the next generation of AI will emerge not from increasingly complex collections of modular components, but from unified architectures grounded in a single underlying intelligence.

7 Contributors

The list is organized by contribution role, with individuals listed alphabetically by their first name within each category.

Project Sponsor and Advisor:

Dahua Lin

Senior Project Lead:

Lei Yang, Lewei Lu, Quan Wang, Ruihao Gong, Wenxiu Sun, Ziwei Liu

Project Lead:

Haiwen Diao

Core Contributor:

Hanming Deng, Jiahao Wang, Penghao Wu, Shihao Bai, Silei Wu, Weichen Fan, Wenjie Ye, Wenwen Tong, Xiangyu Fan, Yan Li, Yubo Wang, Zhijie Cao, Zhiqian Lin, Zhitao Yang, Zhongang Cai

Contributor:

Bo Liu, Chengguang Lv, Haojia Yu, Haozhe Xie, Hongli Wang, Jianan Fan, Jiaqi Li, Jiefan Lu, Jingcheng Ni, Junxiang Xu, Kaihuan Liang, Lianqiang Shi, Linjun Dai, Linyan Wang, Oscar Qian, Peng Gao, Pengfei Liu, Qingping Sun, Rui Shen, Ruisi Wang, Shengnan Ma, Shuang Yang, Siyi Xie, Siying Li, Tianbo Zhong, Xiangli Kong, Xuanke Shi, Yang Gao, Yongqiang Yao, Yue Zhu, Yuwei Niu, Yves Wang, Zhengqi Bai, Zhengyu Lin, Zixin Yin

Acknowledgement:

We would like to thank Bo Yang, Boxuan Li, Chen Feng, Chen Wei, Chenyang Gu, Fanyi Pu, Fanzhou Wang, Guanzhou Chen, Haoge Deng, Hongyu Liang, Houyuan Chen, Huaping Zhong, Huchuan Lu, Jiawei Hong, Jinkun Xie, Jinwei Liang, Mingxuan Li, Mutian Xu, Ruize Ma, Siqi Luo, Tiankuo Yao, Tongxi Zhou, Wangqi Yin, Xiaotong Li, Yinfei Zeng, Yong Xien Chng, Yuhao Dong, Yukang Cao, Zheng Ma, Ziming Wu, Zongpu Zhang, and Zukai Chen for their valuable support and contributions to this project, including data preparation, model evaluation, infrastructure support, architecture analysis, and helpful discussions.

References

- [1] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. *Advances in Neural Information Processing Systems*, 35:23716–23736, 2022.
- [2] Shuai Bai, Yuxuan Cai, Ruizhe Chen, Keqin Chen, Xionghui Chen, Zesen Cheng, Lianghao Deng, Wei Ding, Chang Gao, Chunjiang Ge, et al. Qwen3-vl technical report. *arXiv preprint arXiv:2511.21631*, 2025.
- [3] Jason Baldridge, Jakob Bauer, Mukul Bhutani, Nicole Brichtova, Andrew Bunner, Lluís Castrejon, Kelvin Chan, Yichang Chen, Sander Dieleman, Yuqing Du, et al. Imagen 3. *arXiv preprint arXiv:2408.07009*, 2024.
- [4] Rohan Bavishi, Erich Elsen, Curtis Hawthorne, Maxwell Nye, Augustus Odena, Arushi Somani, and Sağnak Taşlılar. Introducing our multimodal models, 2023. URL <https://www.adept.ai/blog/fuyu-8b>.
- [5] James Betker, Gabriel Goh, Li Jing, Tim Brooks, Jianfeng Wang, Linjie Li, Long Ouyang, Juntang Zhuang, Joyce Lee, Yufei Guo, Wesam Manassra, Prafulla Dhariwal, Casey Chu, Yunxin Jiao, and Aditya Ramesh. Improving image generation with better captions, 2023. URL <https://cdn.openai.com/papers/dall-e-3.pdf>.
- [6] ByteDance. Seedream 4.5, 2025. URL https://seed.bytedance.com/en/seedream4_5.
- [7] Huanqia Cai, Sihan Cao, Ruoyi Du, Peng Gao, Steven Hoi, Shijie Huang, Zhaohui Hou, Dengyang Jiang, Xin Jin, Liangchen Li, et al. Z-image: An efficient image generation foundation model with single-stream diffusion transformer. *arXiv preprint arXiv:2511.22699*, 2025.
- [8] Qi Cai, Jingwen Chen, Yang Chen, Yehao Li, Fuchen Long, Yingwei Pan, Zhaofan Qiu, Yiheng Zhang, Fengbin Gao, Peihan Xu, et al. Hidream-1l: An open-source high-efficient image generative foundation model. *arXiv preprint arXiv:2505.22705*, 2025.
- [9] Zhongang Cai, Yubo Wang, Qingping Sun, Ruisi Wang, Chenyang Gu, Wanqi Yin, Zhiqian Lin, Zhitao Yang, Chen Wei, Xuanke Shi, Kewang Deng, Xiaoyang Han, Zukai Chen, Jiaqi Li, Xiangyu Fan, Hanming Deng, Lewei Lu, Bo Li, Ziwei Liu, Quan Wang, Dahua Lin, and Lei Yang. Holistic evaluation of multimodal llms on spatial intelligence. *arXiv preprint arXiv:2508.13142*, 2025.
- [10] Zhongang Cai, Ruisi Wang, Chenyang Gu, Fanyi Pu, Junxiang Xu, Yubo Wang, Wanqi Yin, Zhitao Yang, Chen Wei, Qingping Sun, Tongxi Zhou, Jiaqi Li, Hui En Pang, Oscar Qian, Yukun Wei, Zhiqian Lin, Xuanke Shi, Kewang Deng, Xiaoyang Han, Zukai Chen, Xiangyu Fan, Hanming Deng, Lewei Lu, Liang Pan, Bo Li, Ziwei Liu, Quan Wang, Dahua Lin, and Lei Yang. Scaling spatial intelligence with multimodal foundation models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2026.
- [11] Siyu Cao, Hangting Chen, Peng Chen, Yiji Cheng, Yutao Cui, Xincheng Deng, Ying Dong, Kipper Gong, Tianpeng Gu, Xiuse Gu, et al. Hunyuanimage 3.0 technical report. *arXiv preprint arXiv:2509.23951*, 2025.
- [12] Jingjing Chang, Yixiao Fang, Peng Xing, Shuhan Wu, Wei Cheng, Rui Wang, Xianfang Zeng, Gang Yu, and Hai-Bao Chen. Oneig-bench: Omni-dimensional nuanced evaluation for image generation. *arXiv preprint arXiv:2506.07977*, 2025.
- [13] Bowei Chen, Sai Bi, Hao Tan, He Zhang, Tianyuan Zhang, Zhengqi Li, Yuanjun Xiong, Jianming Zhang, and Kai Zhang. Aligning visual foundation encoders to tokenizers for diffusion models. *arXiv preprint arXiv:2509.25162*, 2025.
- [14] Jiu-hai Chen, Zhiyang Xu, Xichen Pan, Yushi Hu, Can Qin, Tom Goldstein, Lifu Huang, Tianyi Zhou, Saining Xie, Silvio Savarese, Le Xue, Caiming Xiong, and Ran Xu. Blip3-o: A family of fully open unified multimodal models-architecture, training and dataset. *arXiv preprint arXiv:2505.09568*, 2025.
- [15] Liang Chen, Weichu Xie, Yiyan Liang, Hongfeng He, Hans Zhao, Zhibo Yang, Zhiqi Huang, Haoning Wu, Haoyu Lu, Yiping Bao, et al. Babyvision: Visual reasoning beyond language. *arXiv preprint arXiv:2601.06521*, 2026.
- [16] Lin Chen, Jinsong Li, Xiaoyi Dong, Pan Zhang, Yuhang Zang, Zehui Chen, Haodong Duan, Jiaqi Wang, Yu Qiao, Dahua Lin, et al. Are we on the right way for evaluating large vision-language models? *Advances in Neural Information Processing Systems*, 37:27056–27087, 2024.
- [17] Shoufa Chen, Chongjian Ge, Shilong Zhang, Peize Sun, and Ping Luo. Pixelflow: Pixel-space generative models with flow. *arXiv preprint arXiv:2504.07963*, 2025.
- [18] Xiaokang Chen, Zhiyu Wu, Xingchao Liu, Zizheng Pan, Wen Liu, Zhenda Xie, Xingkai Yu, and Chong Ruan. Janus-pro: Unified multimodal understanding and generation with data and model scaling. *arXiv preprint arXiv:2501.17811*, 2025.

- [19] Yangyi Chen, Xingyao Wang, Hao Peng, and Heng Ji. A single transformer for scalable vision-language modeling. Transactions on Machine Learning Research, 2024.
- [20] Zhennan Chen, Junwei Zhu, Xu Chen, Jiangning Zhang, Xiaobin Hu, Hanzhen Zhao, Chengjie Wang, Jian Yang, and Ying Tai. Dip: Taming diffusion models in pixel space. arXiv preprint arXiv:2511.18822, 2025.
- [21] Ethan Chern, Jiadi Su, Yan Ma, and Pengfei Liu. Anole: An open, autoregressive, native large multimodal models for interleaved image-text generation. arXiv preprint arXiv:2407.06135, 2024.
- [22] Cheng Cui, Ting Sun, Manhui Lin, Tingquan Gao, Yubo Zhang, Jiaxuan Liu, Xueqing Wang, Zelun Zhang, Changda Zhou, Hongen Liu, et al. Paddleocr 3.0 technical report. arXiv preprint arXiv:2507.05595, 2025.
- [23] Yufeng Cui, Honghao Chen, Haoge Deng, Xu Huang, Xinghang Li, Jirong Liu, Yang Liu, Zhuoyan Luo, Jinsheng Wang, Wenxuan Wang, et al. Emu3. 5: Native multimodal models are world learners. arXiv preprint arXiv:2510.26583, 2025.
- [24] Google DeepMind. Gemini 3 pro image model card, Nov 2025. URL <https://storage.googleapis.com/deepmind-media/Model-Cards/Gemini-3-Pro-Image-Model-Card.pdf>.
- [25] Google DeepMind. Gemini 2.0 flash, 2025. URL <https://developers.googleblog.com/en/experiment-with-gemini-2.0-flash-native-image-generation>.
- [26] Google DeepMind. Gemini 2.5 flash & gemini 2.5 flash image model card, 2025. URL <https://storage.googleapis.com/deepmind-media/Model-Cards/Gemini-2-5-Flash-Model-Card.pdf>.
- [27] Google DeepMind. Gemma 4: Byte for byte, the most capable open models, 2026. URL <https://blog.google/innovation-and-ai/technology/developers-tools/gemma-4/>.
- [28] Chaorui Deng, Deyao Zhu, Kunchang Li, Chenhui Gou, Feng Li, Zeyu Wang, Shu Zhong, Weihao Yu, Xiaonan Nie, Ziang Song, Guang Shi, and Haoqi Fan. Emerging properties in unified multimodal pretraining. arXiv preprint arXiv:2505.14683, 2025.
- [29] Haiwen Diao, Yufeng Cui, Xiaotong Li, Yueze Wang, Huchuan Lu, and Xinlong Wang. Unveiling encoder-free vision-language models. Advances in Neural Information Processing Systems, 37:52545–52567, 2024.
- [30] Haiwen Diao, Mingxuan Li, Silei Wu, Linjun Dai, Xiaohua Wang, Hanming Deng, Lewei Lu, Dahua Lin, and Ziwei Liu. From pixels to words—towards native vision-language primitives at scale. arXiv preprint arXiv:2510.14979, 2025.
- [31] Haiwen Diao, Xiaotong Li, Yufeng Cui, Yueze Wang, Haoge Deng, Ting Pan, Wenxuan Wang, Huchuan Lu, and Xinlong Wang. Evev2: Improved baselines for encoder-free vision-language models. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pages 21014–21025, 2025.
- [32] Nikai Du, Zhennan Chen, Shan Gao, Zhizhou Chen, Xi Chen, Zhengkai Jiang, Jian Yang, and Ying Tai. Textcrafter: Accurately rendering multiple texts in complex visual scenes. arXiv preprint arXiv:2503.23461, 2025.
- [33] Xinrun Du, Yifan Yao, Kaijing Ma, Bingli Wang, Tianyu Zheng, King Zhu, Minghao Liu, Yiming Liang, Xiaolong Jin, Zhenlin Wei, et al. Supergpqa: Scaling llm evaluation across 285 graduate disciplines. arXiv preprint arXiv:2502.14739, 2025.
- [34] Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers for high-resolution image synthesis. In International Conference on Machine Learning, 2024.
- [35] Weichen Fan, Haiwen Diao, Quan Wang, Dahua Lin, and Ziwei Liu. The prism hypothesis: Harmonizing semantic and pixel representations via unified autoencoding. arXiv preprint arXiv:2512.19693, 2025.
- [36] Xiangyu Fan, Zesong Qiu, Zhuguanyu Wu, Fanzhou Wang, Zhiqian Lin, Tianxiang Ren, Dahua Lin, Ruihao Gong, and Lei Yang. Phased dmd: Few-step distribution matching distillation via score matching within subintervals. arXiv preprint arXiv:2510.27684, 2025.
- [37] Ling Fu, Zhebin Kuang, Jiajun Song, Mingxin Huang, Biao Yang, Yuzhe Li, Linghao Zhu, Qidi Luo, Xinyu Wang, Hao Lu, et al. Ocrbench v2: An improved benchmark for evaluating large multimodal models on visual text localization and reasoning. arXiv preprint arXiv:2501.00321, 2024.
- [38] Peiyuan Gao, Le Zhuo, Ziyi Lin, Chris Liu, Junsong Chen, Ruoyi Du, Enze Luo, Long Zhang, Guo Chen, Shengyu Zhang, et al. Seedream 3.0 technical report. arXiv preprint arXiv:2504.11346, 2025.
- [39] Yuying Ge, Sijie Zhao, Ziyun Zeng, Yixiao Ge, Chen Li, Xintao Wang, and Ying Shan. Making llama see and draw with seed tokenizer. arXiv preprint arXiv:2310.01218, 2023.

- [40] Yuying Ge, Sijie Zhao, Jinguo Zhu, Yixiao Ge, Kun Yi, Lin Song, Chen Li, Xiaohan Ding, and Ying Shan. Seed-x: Multimodal models with unified multi-granularity comprehension and generation. *arXiv preprint arXiv:2404.14396*, 2024.
- [41] Aryo Pradipta Gema, Joshua Ong Jun Leang, Giwon Hong, Alessio Devoto, Alberto Carlo Maria Mancino, Rohit Saxena, Xuanli He, Yu Zhao, Xiaotang Du, Mohammad Reza Ghasemi Madani, et al. Are we done with mmlu? In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 5069–5096, 2025.
- [42] Zigang Geng, Yibing Wang, Yeyao Ma, Chen Li, Yongming Rao, Shuyang Gu, Zhao Zhong, Qinglin Lu, Han Hu, Xiaosong Zhang, et al. X-omni: Reinforcement learning makes discrete autoregressive image generative models great again. *arXiv preprint arXiv:2507.22058*, 2025.
- [43] Dhruva Ghosh, Hannaneh Hajishirzi, and Ludwig Schmidt. Geneval: An object-focused framework for evaluating text-to-image alignment. *Advances in Neural Information Processing Systems*, 36:52132–52152, 2023.
- [44] Ruihao Gong, Shihao Bai, Siyu Wu, Yunqian Fan, Zaijun Wang, Xiuhong Li, Hailong Yang, and Xianglong Liu. Past-future scheduler for llm serving under sla guarantees. In *Proceedings of the 30th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2*, pages 798–813, 2025. ISBN 9798400710797.
- [45] Google. Imagen 4 model card, 2025. URL <https://storage.googleapis.com/deepmind-media/Model-Cards/Imagen-4-Model-Card.pdf>.
- [46] Google. Nano banana 2: Combining pro capabilities with lightning-fast speed, 2026. URL <https://blog.google/innovation-and-ai/technology/ai/nano-banana-2/>.
- [47] Google DeepMind. Gemini 3 Pro Model Card, November 2025. URL <https://storage.googleapis.com/deepmind-media/Model-Cards/Gemini-3-Pro-Model-Card.pdf>.
- [48] Jiawei Gu, Yunzhuo Hao, Huichen Will Wang, Linjie Li, Michael Qizhe Shieh, Yejin Choi, Ranjay Krishna, and Yu Cheng. Thinkmorph: Emergent properties in multimodal interleaved chain-of-thought reasoning. *arXiv preprint arXiv:2510.27492*, 2025.
- [49] Tianrui Guan, Fuxiao Liu, Xiyang Wu, Ruiqi Xian, Zongxia Li, Xiaoyu Liu, Xijun Wang, Lichang Chen, Furong Huang, Yaser Yacoob, et al. Hallusionbench: an advanced diagnostic suite for entangled language hallucination and visual illusion in large vision-language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14375–14385, 2024.
- [50] Jian Han, Jinlai Liu, Yi Jiang, Bin Yan, Yuqi Zhang, Zehuan Yuan, Bingyue Peng, and Xiaobing Liu. Infinity: Scaling bitwise autoregressive modeling for high-resolution image synthesis. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 15733–15744, 2025.
- [51] Tuomo Hiippala, Malihe Alikhani, Jonas Haverinen, Timo Kalliokoski, Evanfiya Logacheva, Serafina Orekhova, Aino Tuomainen, Matthew Stone, and John A Bateman. Ai2d-rst. *Language Resources and Evaluation*, 55(3):661–688, 2021.
- [52] Wenyi Hong, Wenmeng Yu, Xiaotao Gu, Guo Wang, Guobing Gan, Haomiao Tang, Jiale Cheng, Ji Qi, Junhui Ji, Lihang Pan, et al. Glm-4.5 v and glm-4.1 v-thinking: Towards versatile multimodal reasoning with scalable reinforcement learning. *arXiv preprint arXiv:2507.01006*, 2025.
- [53] Xiwei Hu, Rui Wang, Yixiao Fang, Bin Fu, Pei Cheng, and Gang Yu. Ella: Equip diffusion models with llm for enhanced semantic alignment. *arXiv preprint arXiv:2403.05135*, 2024.
- [54] Runhui Huang, Chunwei Wang, Junwei Yang, Guansong Lu, Yunlong Yuan, Jianhua Han, Lu Hou, Wei Zhang, Lanqing Hong, Hengshuang Zhao, et al. Illume+: Illuminating unified mllm with dual visual tokenization and diffusion refinement. *arXiv preprint arXiv:2504.01934*, 2025.
- [55] Yuzhen Huang, Yuzhuo Bai, Zhihao Zhu, Junlei Zhang, Jinghan Zhang, Tangjun Su, Junteng Liu, Chuancheng Lv, Yikai Zhang, Yao Fu, et al. C-eval: A multi-level multi-discipline chinese evaluation suite for foundation models. *Advances in Neural Information Processing Systems*, 36:62991–63010, 2023.
- [56] Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, et al. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*, 2024.
- [57] Kimi Team. Kimi K2: Open agentic intelligence. *arXiv preprint arXiv:2507.20534*, 2025.
- [58] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- [59] Kuaishou Kolors Team. Kolors 2.0, 2025. URL <https://kolors.kuaishou.com/>.

- [60] Black Forest Labs. Flux, 2024. URL <https://github.com/black-forest-labs/flux>.
- [61] Black Forest Labs. FLUX.2: Frontier Visual Intelligence, 2025. URL <https://bfl.ai/blog/flux-2>.
- [62] Black Forest Labs, Stephen Batifol, Andreas Blattmann, Frederic Boesel, Saksham Consul, Cyril Diagne, Tim Dockhorn, Jack English, Zion English, Patrick Esser, et al. Flux. 1 kontext: Flow matching for in-context image generation and editing in latent space. *arXiv preprint arXiv:2506.15742*, 2025.
- [63] Weixian Lei, Jiacong Wang, Haochen Wang, Xiangtai Li, Jun Hao Liew, Jiashi Feng, and Zilong Huang. The scalability of simplicity: Empirical analysis of vision-language learning with a single transformer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 20758–20769, 2025.
- [64] Xingjian Leng, Jaskirat Singh, Yunzhong Hou, Zhenchang Xing, Saining Xie, and Liang Zheng. Repa-e: Unlocking vae for end-to-end tuning of latent diffusion transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 18262–18272, 2025.
- [65] Dingming Li, Hongxing Li, Zixuan Wang, Yuchen Yan, Hang Zhang, Siqi Chen, Guiyang Hou, Shengpei Jiang, Wenqi Zhang, Yongliang Shen, et al. Viewspatial-bench: Evaluating multi-perspective spatial localization in vision-language models. *arXiv preprint arXiv:2505.21500*, 2025.
- [66] Han Li, Xinyu Peng, Yaoming Wang, Zelin Peng, Xin Chen, Rongxiang Weng, Jingang Wang, Xunliang Cai, Wenrui Dai, and Hongkai Xiong. Onecat: Decoder-only auto-regressive model for unified understanding and generation. *arXiv preprint arXiv:2509.03498*, 2025.
- [67] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International conference on machine learning*, pages 19730–19742. PMLR, 2023.
- [68] Ming Li, Jike Zhong, Shitian Zhao, Haoquan Zhang, Shaoheng Lin, Yuxiang Lai, Chen Wei, Konstantinos Psounis, and Kaipeng Zhang. Tir-bench: A comprehensive benchmark for agentic thinking-with-images reasoning. *arXiv preprint arXiv:2511.01833*, 2025.
- [69] Tianhong Li and Kaiming He. Back to basics: Let denoising generative models denoise. *arXiv preprint arXiv:2511.13720*, 2025.
- [70] Tianle Li, Yongming Rao, Winston Hu, and Yu Cheng. Breen: bridge data-efficient encoder-free multimodal learning with learnable queries. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 5384–5395, 2026.
- [71] Yan Li, Zezi Zeng, Ziwei Zhou, Xin Gao, Muzhao Tian, Yifan Yang, Mingxi Cheng, Qi Dai, Yuqing Yang, Lili Qiu, et al. Bizgeneval: A systematic benchmark for commercial visual content generation. *arXiv preprint arXiv:2603.25732*, 2026.
- [72] Zongjian Li, Zheyuan Liu, Qihui Zhang, Bin Lin, Feize Wu, Shenghai Yuan, Zhiyuan Yan, Yang Ye, Wangbo Yu, Yuwei Niu, et al. Uniworld-v2: Reinforce image editing with diffusion negative-aware finetuning and mllm implicit feedback. *arXiv preprint arXiv:2510.16888*, 2025.
- [73] Weixin Liang, Lili Yu, Liang Luo, Srinivasan Iyer, Ning Dong, Chunting Zhou, Gargi Ghosh, Mike Lewis, Wen-tau Yih, Luke Zettlemoyer, et al. Mixture-of-transformers: A sparse and scalable architecture for multi-modal foundation models. *arXiv preprint arXiv:2411.04996*, 2024.
- [74] Chao Liao, Liyang Liu, Xun Wang, Zhengxiong Luo, Xinyu Zhang, Wenliang Zhao, Jie Wu, Liang Li, Zhi Tian, and Weilin Huang. Mogao: An omni foundation model for interleaved multi-modal generation. *arXiv preprint arXiv:2505.05472*, 2025.
- [75] Bin Lin, Zongjian Li, Xinhua Cheng, Yuwei Niu, Yang Ye, Xianyi He, Shenghai Yuan, Wangbo Yu, Shaocong Wang, Yunyang Ge, et al. Uniworld: High-resolution semantic encoders for unified visual understanding and generation. *arXiv preprint arXiv:2506.03147*, 2025.
- [76] Haokun Lin, Teng Wang, Yixiao Ge, Yuying Ge, Zhichao Lu, Ying Wei, Qingfu Zhang, Zhenan Sun, and Ying Shan. Toklip: Marry visual tokens to clip for multimodal comprehension and generation. *arXiv preprint arXiv:2505.05422*, 2025.
- [77] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755, 2014.
- [78] Xi Victoria Lin, Akshat Shrivastava, Liang Luo, Srinivasan Iyer, Mike Lewis, Gargi Ghosh, Luke Zettlemoyer, and Armen Aghajanyan. Moma: Efficient early-fusion pre-training with mixture of modality-aware experts. *arXiv preprint arXiv:2407.21770*, 2024.

- [79] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *Advances in Neural Information Processing Systems*, 36:34892–34916, 2023.
- [80] Jie Liu, Gongye Liu, Jiajun Liang, Yangguang Li, Jiaheng Liu, Xintao Wang, Pengfei Wan, Di Zhang, and Wanli Ouyang. Flow-grpo: Training flow matching models via online rl. *arXiv preprint arXiv:2505.05470*, 2025.
- [81] Shiyu Liu, Yucheng Han, Peng Xing, Fukun Yin, Rui Wang, Wei Cheng, Jiaqi Liao, Yingming Wang, Honghao Fu, Chunrui Han, et al. Step1x-edit: A practical framework for general image editing. *arXiv preprint arXiv:2504.17761*, 2025.
- [82] Yuan Liu, Haodong Duan, Yuanhan Zhang, Bo Li, Songyang Zhang, Wangbo Zhao, Yike Yuan, Jiaqi Wang, Conghui He, Ziwei Liu, et al. Mmbench: Is your multi-modal model an all-around player? In *Proceedings of the European Conference on Computer Vision*, pages 216–233. Springer, 2024.
- [83] Yuliang Liu, Zhang Li, Mingxin Huang, Biao Yang, Wenwen Yu, Chunyuan Li, Xu-Cheng Yin, Cheng-Lin Liu, Lianwen Jin, and Xiang Bai. Ocrbench: on the hidden mystery of ocr in large multimodal models. *Science China Information Sciences*, 67(12):220102, 2024.
- [84] Zhiheng Liu, Weiming Ren, Haozhe Liu, Zijian Zhou, Shoufa Chen, Haonan Qiu, Xiaoke Huang, Zhaochong An, Fanny Yang, Aditya Patel, et al. Tuna: Taming unified visual representations for native unified multimodal models. *arXiv preprint arXiv:2512.02014*, 2025.
- [85] Zhiheng Liu, Weiming Ren, Xiaoke Huang, Shoufa Chen, Tianhong Li, Mengzhao Chen, Yatai Ji, Sen He, Jonas Schult, Tao Xiang, Wenhu Chen, Ping Luo, Luke Zettlemoyer, and Yuren Cong. Tuna-2: Pixel embeddings beat vision encoders for unified understanding and generation. *arXiv preprint arXiv:2604.24763*, 2026.
- [86] Joshua Lopez. Stable diffusion 3: Research paper, 2025. URL <https://stability.ai/news/stable-diffusion-3-research-paper>. Stability AI.
- [87] Pan Lu, Hritik Bansal, Tony Xia, Jiacheng Liu, Chunyuan Li, Hannaneh Hajishirzi, Hao Cheng, Kai-Wei Chang, Michel Galley, and Jianfeng Gao. Mathvista: Evaluating mathematical reasoning of foundation models in visual contexts. *arXiv preprint arXiv:2310.02255*, 2023.
- [88] Gen Luo, Wenhan Dou, Wenhao Li, Zhaokai Wang, Xue Yang, Changyao Tian, Hao Li, Weiyun Wang, Wenhao Wang, Xizhou Zhu, et al. Mono-intervl-1.5: Towards cheaper and faster monolithic multimodal large language models. *arXiv preprint arXiv:2507.12566*, 2025.
- [89] Gen Luo, Xue Yang, Wenhan Dou, Zhaokai Wang, Jiawen Liu, Jifeng Dai, Yu Qiao, and Xizhou Zhu. Mono-intervl: Pushing the boundaries of monolithic multimodal large language models with endogenous visual pre-training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 24960–24971, 2025.
- [90] Chuofan Ma, Yi Jiang, Junfeng Wu, Jihan Yang, Xin Yu, Zehuan Yuan, Bingyue Peng, and Xiaojuan Qi. Unitok: A unified tokenizer for visual generation and understanding. *arXiv preprint arXiv:2502.20321*, 2025.
- [91] Wufei Ma, Haoyu Chen, Guofeng Zhang, Yu-Cheng Chou, Celso M de Melo, and Alan Yuille. 3dsrbench: A comprehensive 3d spatial reasoning benchmark. *arXiv preprint arXiv:2412.07825*, 2024.
- [92] Yiyang Ma, Xingchao Liu, Xiaokang Chen, Wen Liu, Chengyue Wu, Zhiyu Wu, Zizheng Pan, Zhenda Xie, Haowei Zhang, Xingkai Yu, et al. Janusflow: Harmonizing autoregression and rectified flow for unified multimodal understanding and generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7739–7751, 2025.
- [93] Yuhang Ma, Xiaoshi Wu, Keqiang Sun, and Hongsheng Li. Hpsv3: Towards wide-spectrum human preference score. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15086–15095, 2025.
- [94] Minesh Mathew, Viraj Bagal, Rubèn Tito, Dimosthenis Karatzas, Ernest Valveny, and CV Jawahar. Infographicvqa. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1697–1706, 2022.
- [95] Midjourney. Midjourney v7, 2025. URL <https://www.midjourney.com/home>.
- [96] Minimax. Image-01, 2025. URL <https://www.minimax.io/news/image-01>.
- [97] ModelTC. LightLLM: A python-based llm inference and serving framework, 2025. URL <https://github.com/ModelTC/LightLLM>.
- [98] ModelTC. LightX2V: A lightweight video and image generation inference framework, 2025. URL <https://github.com/ModelTC/LightX2V>.

- [99] Yuwei Niu, Munan Ning, Mengren Zheng, Weiyang Jin, Bin Lin, Peng Jin, Jiaqi Liao, Chaoran Feng, Kunpeng Ning, Bin Zhu, et al. Wise: A world knowledge-informed semantic evaluation for text-to-image generation. arXiv preprint arXiv:2503.07265, 2025.
- [100] OpenAI. Gpt-image-1, 2025. URL <https://platform.openai.com/docs/guides/image-generation>.
- [101] OpenAI. Gpt-image-1.5, 2025. URL <https://openai.com/index/new-chatgpt-images-is-here/>.
- [102] OpenAI. Introducing chatgpt images 2.0, 2026. URL <https://openai.com/index/introducing-chatgpt-images-2-0/>.
- [103] Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. Sdxl: Improving latent diffusion models for high-resolution image synthesis. arXiv preprint arXiv:2307.01952, 2023.
- [104] Pruna AI. Pruna ai, 2025. URL <https://www.pruna.ai>.
- [105] Liao Qu, Huichao Zhang, Yiheng Liu, Xu Wang, Yi Jiang, Yiming Gao, Hu Ye, Daniel K Du, Zehuan Yuan, and Xinglong Wu. Tokenflow: Unified image tokenizer for multimodal understanding and generation. In Proceedings of the Computer Vision and Pattern Recognition Conference, pages 2545–2555, 2025.
- [106] Qwen Team. Qwen3.5: Towards native multimodal agents, February 2026. URL <https://qwen.ai/blog?id=qwen3.5>.
- [107] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In International conference on machine learning, pages 8748–8763, 2021.
- [108] Recraft. Recraft v3, 2024. URL <https://www.recraft.ai/docs/recreate-models/recreate-V3>.
- [109] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 10684–10695, 2022.
- [110] Bytedance Seed. Seedream 5.0 lite. https://seed.bytedance.com/en/seedream5_0_lite, 2026.
- [111] Team Seedream, Yunpeng Chen, Yu Gao, Lixue Gong, Meng Guo, Qiushan Guo, Zhiyao Guo, Xiaoxia Hou, Weilin Huang, Yixuan Huang, et al. Seedream 4.0: Toward next-generation multimodal image generation. arXiv preprint arXiv:2509.20427, 2025.
- [112] SenseNova. Neo-unify: Building native multimodal unified models end to end, 2026. URL <https://huggingface.co/blog/sensenova/neo-unify>.
- [113] Minglei Shi, Haolin Wang, Wenzhao Zheng, Ziyang Yuan, Xiaoshi Wu, Xintao Wang, Pengfei Wan, Jie Zhou, and Jiwen Lu. Latent diffusion model without variational autoencoder. arXiv preprint arXiv:2510.15301, 2025.
- [114] Yang Shi, Yuhao Dong, Yue Ding, Yuran Wang, Xuanyu Zhu, Sheng Zhou, Wenting Liu, Haochen Tian, Rundong Wang, Huanqian Wang, et al. Realunify: Do unified models truly benefit from unification? a comprehensive benchmark. arXiv preprint arXiv:2509.24897, 2025.
- [115] Aaditya Singh, Adam Fry, Adam Perelman, Adam Tart, Adi Ganesh, Ahmed El-Kishky, Aidan McLaughlin, Aiden Low, AJ Ostrow, Akhila Ananthram, et al. Openai gpt-5 system card. arXiv preprint arXiv:2601.03267, 2025.
- [116] Wei Song, Yuran Wang, Zijia Song, Yadong Li, Haoze Sun, Weipeng Chen, Zenan Zhou, Jianhua Xu, Jiaqi Wang, and Kaicheng Yu. Dualtoken: Towards unifying visual understanding and generation with dual visual vocabularies. arXiv preprint arXiv:2503.14324, 2025.
- [117] JD Open Source. Joyai-image: Awakening spatial intelligence in unified multimodal understanding and generation, 2026. URL <https://github.com/jd-opensource/JoyAI-Image>.
- [118] Quan Sun, Yuxin Fang, Ledell Wu, Xinlong Wang, and Yue Cao. Eva-clip: Improved training techniques for clip at scale. arXiv preprint arXiv:2303.15389, 2023.
- [119] Quan Sun, Yufeng Cui, Xiaosong Zhang, Fan Zhang, Qiying Yu, Yueze Wang, Yongming Rao, Jingjing Liu, Tiejun Huang, and Xinlong Wang. Generative multimodal models are in-context learners. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 14398–14409, 2024.
- [120] Yinghao Tang, Xueding Liu, Boyuan Zhang, Tingfeng Lan, Yupeng Xie, Jiale Lao, Yiyao Wang, Haoxuan Li, Tingting Gao, Bo Pan, et al. Igenbench: Benchmarking the reliability of text-to-infographic generation. arXiv preprint arXiv:2601.04498, 2026.

- [121] Chenxin Tao, Shiqian Su, Xizhou Zhu, Chenyu Zhang, Zhe Chen, Jiawen Liu, Wenhai Wang, Lewei Lu, Gao Huang, Yu Qiao, et al. Hovle: Unleashing the power of monolithic vision-language models with holistic vision-language embedding. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 14559–14569, 2025.
- [122] Chameleon Team. Chameleon: Mixed-modal early-fusion foundation models. *arXiv preprint arXiv:2405.09818*, 2024. doi: 10.48550/arXiv.2405.09818. URL <https://github.com/facebookresearch/chameleon>.
- [123] Gemini Team, Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, Katie Millican, et al. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023.
- [124] Meituan LongCat Team, Hanghang Ma, Haoxian Tan, Jiale Huang, Junqiang Wu, Jun-Yan He, Lishuai Gao, Songlin Xiao, Xiaoming Wei, Xiaoqi Ma, et al. Longcat-image technical report. *arXiv preprint arXiv:2512.07584*, 2025.
- [125] Meituan LongCat Team, Bin Xiao, Chao Wang, Chengjiang Li, Chi Zhang, Chong Peng, Hang Yu, Hao Yang, Haonan Yan, Haoze Sun, et al. Longcat-next: Lexicalizing modalities as discrete tokens. *arXiv preprint arXiv:2603.27538*, 2026.
- [126] Changyao Tian, Danni Yang, Guanzhou Chen, Erfei Cui, Zhaokai Wang, Yuchen Duan, Penghao Yin, Sitao Chen, Ganlin Yang, Mingxin Liu, et al. Internvl-u: Democratizing unified multimodal models for understanding, reasoning, generation and editing. *arXiv preprint arXiv:2603.09877*, 2026.
- [127] Shengbang Tong, David Fan, John Nguyen, Ellis Brown, Gaoyue Zhou, Shengyi Qian, Boyang Zheng, Théophane Vallaëys, Junlin Han, Rob Fergus, et al. Beyond language modeling: An exploration of multimodal pretraining. *arXiv preprint arXiv:2603.03276*, 2026.
- [128] Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *Advances in Neural Information Processing Systems*, 30, 2017.
- [129] Team Wan, Ang Wang, Baole Ai, Bin Wen, Chaojie Mao, Chen-Wei Xie, Di Chen, Fei Wu Yu, Haiming Zhao, Jianxiao Yang, et al. Wan: Open and advanced large-scale video generative models. *arXiv preprint arXiv:2503.20314*, 2025.
- [130] Guo-Hua Wang, Shanshan Zhao, Xinjie Zhang, Liangfu Cao, Pengxin Zhan, Lunhao Duan, Shiyin Lu, Minghao Fu, Xiaohao Chen, Jianshan Zhao, et al. Ovis-u1 technical report. *arXiv preprint arXiv:2506.23044*, 2025.
- [131] Han Wang, Yongjie Ye, Bingru Li, Yuxiang Nie, Jinghui Lu, Jingqun Tang, Yanjie Wang, and Can Huang. Vision as lora. *arXiv preprint arXiv:2503.20680*, 2025.
- [132] Ke Wang, Junting Pan, Weikang Shi, Zimu Lu, Houxing Ren, Aojun Zhou, Mingjie Zhan, and Hongsheng Li. Measuring multimodal mathematical reasoning with math-vision dataset. *Advances in Neural Information Processing Systems*, 37: 95095–95169, 2024.
- [133] Maijunxian Wang, Ruisi Wang, Juyi Lin, Ran Ji, Thaddäus Wiedemer, Qingying Gao, Dezhi Luo, Yaoyao Qian, Lianyu Huang, Zelong Hong, Jiahui Ge, Qianli Ma, Hang He, Yifan Zhou, Lingzi Guo, Lantao Mei, Jiachen Li, Hanwen Xing, Tianqi Zhao, Fengyuan Yu, Weihang Xiao, Yizheng Jiao, Jianheng Hou, Danyang Zhang, Pengcheng Xu, Boyang Zhong, Zehong Zhao, Gaoyun Fang, John Kitaoka, Yile Xu, Hua bureau Xu, Kenton Blacutt, Tin Nguyen, Siyuan Song, Haoran Sun, Shaoyue Wen, Linyang He, Runming Wang, Yanzhi Wang, Mengyue Yang, Ziqiao Ma, Raphaël Millièvre, Freda Shi, Nuno Vasconcelos, Daniel Khashabi, Alan Yuille, Yilun Du, Ziming Liu, Dahua Lin, Ziwei Liu, Vikash Kumar, Yijiang Li, Lei Yang, Zhongang Cai, and Hokin Deng. A very big video reasoning suite. *arXiv preprint arXiv:2602.20159*, 2026.
- [134] Weiyun Wang, Zhangwei Gao, Lixin Gu, Hengjun Pu, Long Cui, Xingguang Wei, Zhaoyang Liu, Linglin Jing, Shenglong Ye, Jie Shao, et al. Internvl3. 5: Advancing open-source multimodal models in versatility, reasoning, and efficiency. *arXiv preprint arXiv:2508.18265*, 2025.
- [135] Xinlong Wang, Xiaosong Zhang, Zhengxiong Luo, Quan Sun, Yufeng Cui, Jinsheng Wang, Fan Zhang, Yueze Wang, Zhen Li, Qiyang Yu, et al. Emu3: Next-token prediction is all you need. *arXiv preprint arXiv:2409.18869*, 2024.
- [136] Yubo Wang, Xueguang Ma, Ge Zhang, Yuansheng Ni, Abhranil Chandra, Shiguang Guo, Weiming Ren, Aaran Arulraj, Xuan He, Ziyang Jiang, et al. Mmlu-pro: A more robust and challenging multi-task language understanding benchmark. *Advances in Neural Information Processing Systems*, 37:95266–95290, 2024.
- [137] Hongyang Wei, Baixin Xu, Hongbo Liu, Size Wu, Jie Liu, Yi Peng, Peiyu Wang, Zexiang Liu, Jingwen He, Yidan Xietian, et al. Skywork unipic 2.0: Building kontext model with online rl for unified multimodal model. *arXiv preprint arXiv:2509.04548*, 2025.
- [138] Xinyu Wei, Jinrui Zhang, Zeqing Wang, Hongyang Wei, Zhen Guo, and Lei Zhang. Tiif-bench: How does your t2i model follow your instructions? *arXiv preprint arXiv:2506.02161*, 2025.

- [139] Chenfei Wu, Jiahao Li, Jingren Zhou, Junyang Lin, Kaiyuan Gao, Kun Yan, Sheng-ming Yin, Shuai Bai, Xiao Xu, Yilei Chen, et al. Qwen-image technical report. [arXiv preprint arXiv:2508.02324](#), 2025.
- [140] Chengyue Wu, Xiaokang Chen, Zhiyu Wu, Yiyang Ma, Xingchao Liu, Zizheng Pan, Wen Liu, Zhenda Xie, Xingkai Yu, Chong Ruan, et al. Janus: Decoupling visual encoding for unified multimodal understanding and generation. [arXiv preprint arXiv:2410.13848](#), 2024.
- [141] Chenyuan Wu, Pengfei Zheng, Ruiran Yan, Shitao Xiao, Xin Luo, Yuezhe Wang, Wanli Li, Xiyan Jiang, Yexin Liu, Junjie Zhou, et al. Omnigen2: Exploration to advanced multimodal generation. [arXiv preprint arXiv:2506.18871](#), 2025.
- [142] Shengqiong Wu, Hao Fei, Leigang Qu, Wei Ji, and Tat-Seng Chua. NEX-T-GPT: Any-to-Any Multimodal LLM. In [Proceedings of the International Conference on Machine Learning](#), 2024.
- [143] Yecheng Wu, Zhuoyang Zhang, Junyu Chen, Haotian Tang, Dacheng Li, Yunhao Fang, Ligeng Zhu, Enze Xie, Hongxu Yin, Li Yi, et al. Vila-u: a unified foundation model integrating visual understanding and generation. [arXiv preprint arXiv:2409.04429](#), 2024.
- [144] Shitao Xiao, Yuezhe Wang, Junjie Zhou, Huaying Yuan, Xingrun Xing, Ruiran Yan, Shuting Wang, Tiejun Huang, and Zheng Liu. Omnigen: Unified image generation. [arXiv preprint arXiv:2409.11340](#), 2024.
- [145] Jinheng Xie, Weijia Mao, Zechen Bai, David Junhao Zhang, Weihao Wang, Kevin Qinghong Lin, Yuchao Gu, Zhijie Chen, Zhenheng Yang, and Mike Zheng Shou. Show-o: One single transformer to unify multimodal understanding and generation. [arXiv preprint arXiv:2408.12528](#), 2024.
- [146] Jinheng Xie, Weijia Mao, Zechen Bai, David Junhao Zhang, Weihao Wang, Kevin Qinghong Lin, Yuchao Gu, Zhijie Chen, Zhenheng Yang, Mike Zheng Shou, et al. Show-o2: Improved native unified multimodal models. [arXiv preprint arXiv:2510.23218](#), 2025.
- [147] Yi Xin, Qi Qin, Siqi Luo, Kaiwen Zhu, Juncheng Yan, Yan Tai, Jiayi Lei, Yuewen Cao, Keqi Wang, Yibin Wang, et al. Lumina-dimoo: An omni diffusion large language model for multi-modal generation and understanding. [arXiv preprint arXiv:2510.06308](#), 2025.
- [148] Jinbo Xing, Zeyinzi Jiang, Yuxiang Tuo, Chaojie Mao, Xiaotang Gai, Xi Chen, Jingfeng Zhang, Yulin Pan, Zhen Han, Jie Xiao, et al. Wan-weaver: Interleaved multi-modal generation via decoupled training. [arXiv preprint arXiv:2603.25706](#), 2026.
- [149] An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, et al. Qwen3 technical report. [arXiv preprint arXiv:2505.09388](#), 2025.
- [150] Jihan Yang, Shusheng Yang, Anjali W Gupta, Rilyn Han, Li Fei-Fei, and Saining Xie. Thinking in space: How multimodal large language models see, remember, and recall spaces. In [Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition](#), pages 10632–10643, 2025.
- [151] Rui Yang, Lin Song, Yicheng Xiao, Runhui Huang, Yixiao Ge, Ying Shan, and Hengshuang Zhao. Haplovl: A single-transformer baseline for multi-modal understanding. [arXiv preprint arXiv:2503.14694](#), 2025.
- [152] Jingfeng Yao, Bin Yang, and Xinggang Wang. Reconstruction vs. generation: Taming optimization dilemma in latent diffusion models. In [Proceedings of the Computer Vision and Pattern Recognition Conference](#), pages 15703–15712, 2025.
- [153] Shunyu Yao, Karthik Narasimhan, and Jian Cao. τ -bench: A benchmark for tool-agent-user interaction in real-world domains. [arXiv preprint arXiv:2406.12045](#), 2024.
- [154] Bowen Ye, Rang Li, Qibin Yang, Yuanxin Liu, Linli Yao, Hanglong Lv, Zhihui Xie, Chenxin An, Lei Li, Lingpeng Kong, Qi Liu, Zhifang Sui, and Tong Yang. Claw-eval: Toward trustworthy evaluation of autonomous agents. [arXiv preprint arXiv:2604.06132](#), 2026.
- [155] Yang Ye, Xianyi He, Zongjian Li, Bin Lin, Shenghai Yuan, Zhiyuan Yan, Bohan Hou, and Li Yuan. Imgedit: A unified image editing dataset and benchmark. [arXiv preprint arXiv:2505.20275](#), 2025.
- [156] Baiqiao Yin, Qineng Wang, Pingyue Zhang, Jianshu Zhang, Kangrui Wang, Zihan Wang, Jieyu Zhang, Keshigeyan Chandrasegaran, Han Liu, Ranjay Krishna, et al. Spatial mental modeling from limited views. In [Proceedings of the IEEE/CVF International Conference on Computer Vision Workshop](#), 2025.
- [157] Tianwei Yin, Michaël Gharbi, Taesung Park, Richard Zhang, Eli Shechtman, Fredo Durand, and William T Freeman. Improved distribution matching distillation for fast image synthesis. [Advances in Neural Information Processing Systems](#), 37: 47455–47487, 2024.

- [158] Yongsheng Yu, Wei Xiong, Weili Nie, Yichen Sheng, Shiqiu Liu, and Jiebo Luo. Pixeldit: Pixel diffusion transformers for image generation. arXiv preprint arXiv:2511.20645, 2025.
- [159] Xiang Yue, Yuansheng Ni, Kai Zhang, Tianyu Zheng, Ruoqi Liu, Ge Zhang, Samuel Stevens, Dongfu Jiang, Weiming Ren, Yuxuan Sun, et al. Mmmu: A massive multi-discipline multimodal understanding and reasoning benchmark for expert agi. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 9556–9567, 2024.
- [160] Xiang Yue, Tianyu Zheng, Yuansheng Ni, Yubo Wang, Kai Zhang, Shengbang Tong, Yuxuan Sun, Botao Yu, Ge Zhang, Huan Sun, et al. Mmmu-pro: A more robust multi-discipline multimodal understanding benchmark. In Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pages 15134–15186, 2025.
- [161] Zhengrong Yue, Haiyu Zhang, Xiangyu Zeng, Boyu Chen, Chenting Wang, Shaobin Zhuang, Lu Dong, KunPeng Du, Yi Wang, Limin Wang, et al. Uniflow: A unified pixel flow tokenizer for visual understanding and generation. arXiv preprint arXiv:2510.10575, 2025.
- [162] Z.ai. Glm-image: Auto-regressive for dense-knowledge and high-fidelity image generation, 2026. URL <https://z.ai/blog/glm-image/>.
- [163] Xiaohua Zhai, Basil Mustafa, Alexander Kolesnikov, and Lucas Beyer. Sigmoid loss for language image pre-training. In Proceedings of the IEEE/CVF international conference on computer vision, pages 11975–11986, 2023.
- [164] Tao Zhang, Yuyang Hong, Yang Xia, Kun Ding, Zeyu Zhang, Ying Wang, Shiming Xiang, and Chunhong Pan. If-bench: Benchmarking and enhancing mllms for infrared images with generative visual prompting. arXiv preprint arXiv:2512.09663, 2025.
- [165] Zechuan Zhang, Ji Xie, Yu Lu, Zongxin Yang, and Yi Yang. In-context edit: Enabling instructional image editing with in-context generation in large scale diffusion transformer. arXiv preprint arXiv:2504.20690, 2025.
- [166] Xiangyu Zhao, Peiyuan Zhang, Kexian Tang, Xiaorong Zhu, Hao Li, Wenhao Chai, Zicheng Zhang, Renqiu Xia, Guangtao Zhai, Junchi Yan, et al. Envisioning beyond the pixels: Benchmarking reasoning-informed visual editing. arXiv preprint arXiv:2504.02826, 2025.
- [167] Yue Zhao, Fuzhao Xue, Scott Reed, Linxi Fan, Yuke Zhu, Jan Kautz, Zhiding Yu, Philipp Krähenbühl, and De-An Huang. Qlip: Text-aligned visual tokenization unifies auto-regressive multimodal understanding and generation. arXiv preprint arXiv:2502.05178, 2025.
- [168] Boyang Zheng, Nanye Ma, Shengbang Tong, and Saining Xie. Diffusion transformers with representation autoencoders. arXiv preprint arXiv:2510.11690, 2025.
- [169] Kaizhi Zheng, Xuehai He, and Xin Eric Wang. MiniGPT-5: Interleaved vision-and-language generation via generative vokens. arXiv preprint arXiv:2310.02239, 2023.
- [170] Chunting Zhou, Lili Yu, Arun Babu, Kushal Tirumala, Michihiro Yasunaga, Leonid Shamis, Jacob Kahn, Xuezhe Ma, Luke Zettlemoyer, and Omer Levy. Transfusion: Predict the next token and diffuse images with one multi-modal model. arXiv preprint arXiv:2408.11039, 2024.
- [171] Jeffrey Zhou, Tianjian Lu, Swaroop Mishra, Siddhartha Brahma, Sujoy Basu, Yi Luan, Denny Zhou, and Le Hou. Instruction-following evaluation for large language models. arXiv preprint arXiv:2311.07911, 2023.
- [172] Pengfei Zhou, Xiaopeng Peng, Jiajun Song, Chuanhao Li, Zhaopan Xu, Yue Yang, Ziyao Guo, Hao Zhang, Yuqi Lin, Yefei He, et al. Opening: A comprehensive benchmark for judging open-ended interleaved image-text generation. In Proceedings of the Computer Vision and Pattern Recognition Conference, pages 56–66, 2025.
- [173] Kai Zou, Ziqi Huang, Yuhao Dong, Shulin Tian, Dian Zheng, Hongbo Liu, Jingwen He, Bin Liu, Yu Qiao, and Ziwei Liu. Uni-mmmu: A massive multi-discipline multimodal unified benchmark. arXiv preprint arXiv:2510.13759, 2025.