

EchoGen: Cycle-Consistent Learning for Unified Layout-Image Generation and Understanding

Kai Zou^{1,6}, Hongbo Liu², Dian Zheng³, Jianxiong Gao⁴, Zhiwei Zhao⁵, Bin Liu^{1,6*}

¹School of Cyber Science and Technology, University of Science and Technology of China

²Tongji University

³Sun Yat-sen University

⁴Fudan University

⁵Hefei University of Technology

⁶Anhui Province Key Laboratory of Digital Security

kzou@mail.ustc.edu.cn, 1952350@tongji.edu.cn, zhengd35@mail3.sysu.edu.cn,

jxgao22@m.fudan.edu.cn, zhiweizhao@hfut.edu.cn, flowice@ustc.edu.cn

Abstract

In this work, we present EchoGen, a unified framework for layout-to-image generation and image grounding, capable of generating images with both accurate layout and high fidelity to the text description. (*e.g.*, spatial relationship), and grounding the image robustly at the same time. We believe that image grounding possesses strong text and layout understanding abilities, which can compensate for the corresponding limitations in layout-to-image generation. At the same time, images generated from layouts exhibit high diversity in content, thereby enhancing the robustness of image grounding. Jointly training both tasks within a unified model can promote performance improvements for each. However, we identify that this joint training paradigm encounters several optimization challenges and results in restricted performance. To address these issues, we propose progressive training strategies. First, the Parallel Multi-Task Pre-training (PMTP) stage equips the model with basic abilities for both tasks, leveraging shared tokens to accelerate training. Next, the Dual Joint Optimization (DJO) stage exploits task duality to sequentially integrate the two tasks, enabling unified optimization. Finally, the Cycle RL stage eliminates reliance on visual supervision by using consistency constraints as rewards, significantly enhancing the model’s unified capabilities via the GRPO strategy. Extensive experiments demonstrate state-of-the-art results on both layout-to-image generation and image grounding benchmarks, and reveal clear synergistic gains from optimizing the two tasks together.

Introduction

Applications such as poster generation (Chen et al. 2025b), interactive image editing (Liu et al. 2025), and visual question answering (Antol et al. 2015) demand both controllable image generation and accurate visual understanding. Two fundamental tasks in this domain are layout-to-image generation (Zhou et al. 2024b) and image grounding (Liu et al. 2024). The former aims to synthesize realistic images conditioned on grounding text and spatial layout, while the latter

reverses this process by identifying and describing visual elements in an image, mapping them back to structured text and layout.

Currently, researchers mainly handle the two tasks independently (Liu et al. 2024; Peng et al. 2023), lacking a unified architecture and training paradigm for joint optimization. However, they ignore the fact that joint training enables the model to naturally acquire favorable properties that are difficult to achieve in single-task settings. For example, as shown in Fig. 1, we aim to generate an image of three apples placed on a table, where their pixel positions are specified by bounding boxes, and their *spatial arrangement*—such as front-to-back or top-to-bottom—is described by the prompt. Single-task approaches struggle to accurately distinguish and interpret such spatial relationships conveyed in the text. In contrast, our unified framework benefits from the grounding task, which explicitly requires the model to understand spatial descriptions—such as identifying the nearest apple—thereby naturally enhancing alignment to both textual semantics and spatial layout specifications.

In this work, we propose EchoGen, a unified framework for layout-to-image generation and image grounding, aiming at exhibiting the synergistic effect between two tasks. In practice, we found that directly jointly training the two tasks leads to limited performance; therefore, we propose progressive training strategies. First, we perform Parallel Multi-Task Pre-training (PMTP), accelerate the learning of base capabilities by share visual token representations, leveraging the inherent duality between generation and grounding tasks. Next, we serialize the two tasks into a joint learning objective to achieve synergistic optimization and strengthen consistency along the layout–image–layout cycle mapping, which primes the model for the self-supervised Reinforcement Learning (RL) stage. Finally, by treating the grounding discrepancy as rewards, we perform self-supervised RL over the layout→image→layout loop (CycleRL) without requiring explicit supervision of the visual output, relying solely on text prompts with even random bounding boxes.

Extensive experiments on MS-COCO, LayoutSAM, and

*Corresponding authors

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

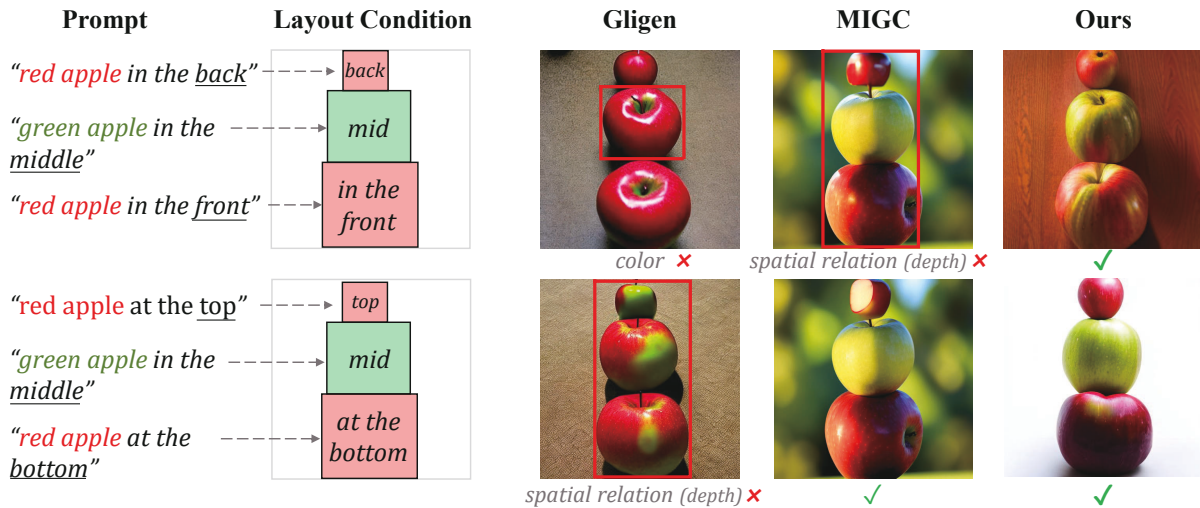


Figure 1: Our unified approach faithfully parses complex prompts—e.g., color attributes and relative spatial relations—and, combined with layout conditions, achieves more accurate semantic alignment. In contrast, specialized methods (e.g., *Gligen* and *MIGC*) often fail to handle such constraints. For instance, in the first row of *MIGC* results, the model generates a vertical spatial relation (top–middle–bottom), whereas the prompt specifies a depth-based relation (front–middle–back).

Ref-L4 demonstrate the effectiveness of our approach on both layout-to-image generation and image grounding.

The key contributions of this paper are as follows:

- We propose a unified training framework for layout-to-image generation and image grounding, named EchoGen, enabling collaborative learning and efficient training with a duality-driven joint objective for the two tasks.
- We achieve self-supervised RL over the layout→image→layout loop using cycle-consistent grounding discrepancy as rewards, eliminating the need of explicit visual supervision.
- Experiments on MS-COCO, LayoutSAM, and Ref-L4 verify state-of-the-art performance and clear synergistic gains.

Related Works

Layout Control in Image Generation

Pure text guidance in image generation lacks precise spatial control, necessitating explicit layout conditions (e.g., bounding boxes, segmentation maps). *Training-free* methods typically edit cross-attention (2023) to control text-image interactions, or apply energy-based latent updates (2023; 2023; 2024); however, these approaches often struggle with complex layouts. In contrast, *training-based* approaches (2023; 2023; 2023) introduce explicit conditioning modules. For instance, *GLIGEN* (2023) injects grounding tokens and spatial boxes; *InstanceDiffusion* (2024b) utilizes mask-guided conditioning; *MIGC* (2024b) perform optimization in a divide-and-conquer way. Methods like *IFAdapter* (2024) and *ELIGEN* (2025a) transfer control to more powerful backbones, balancing fidelity and controllability. Unified models (*PlanGen* (2025)) integrate planning, generation, and

grounding within an auto-regressive model, but they still optimize each task in isolation, resulting in suboptimal results.

Image Grounding

Image grounding aims to localise textual references in images (2022; 2024; 2023). Early open-vocabulary detectors such as GLIP (2022), Grounding DINO (2024) and the real-time YOLO-World (2024) combine detection backbones with text features to support zero-shot phrase localisation. More recently, general-purpose vision–language models like CogVLM (2024) and Qwen2.5-VL (2025) are trained on large-scale visual–text corpora with bounding-box supervision and achieve strong zero-shot results on standard grounding benchmarks(2016; 2025a). We further incorporate image grounding into a unified framework for multimodal understanding and generation, enabling joint optimization with layout-to-image synthesis.

Unified Multimodal Models

Recently, there is growing interest in unified multimodal models that can both understand and generate visual content within one architecture (2025b; 2025; 2025; 2025). Some approaches (2024a) combines an autoregressive language modeling objective with an image diffusion process to train one Transformer capable of producing either text or images from a given prompt. Janus (2025) addresses the different granularity needs between understanding and generation by decoupling its visual encoder into separate pathways, while still using a shared Transformer backbone. Building on this, Janus-Pro (2025c) yields further gains in multimodal understanding and generation. Most recently, BAGEL (2025) demonstrated that a decoder-only model pre-trained on trillions of interleaved text-image-video tokens can exhibit emergent multimodal capabilities.

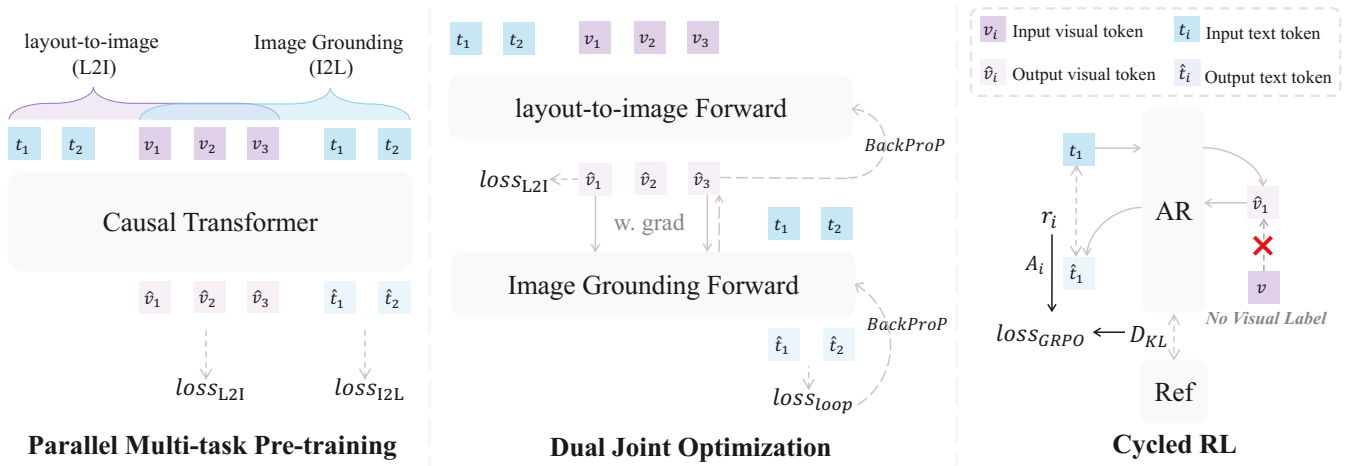


Figure 2: Overview of *EchoGen*. **Left—Parallel Multi-task Pre-training:** a unified autoregressive transformer is trained on layout-to-image and image grounding tasks in parallel, yielding fast acquisition of base capabilities. **Middle—Dual Joint Optimization:** the image tokens generated by the generation forward pass are directly reused as the input to grounding, forming a single joint objective, which strengthens layout→image→layout cycle consistency. **Right—Cycled RL:** we execute the layout→image→layout loop and treat the discrepancy between input and recovered layouts as a continuous reward, enabling self-supervised Reinforcement Learning without explicit supervision of intermediate visual outputs.

Method

Overview

To enable fully self-supervised reinforcement learning in the final stage, the model must first acquire reliable layout-to-image synthesis and image grounding capabilities and exhibit sufficient consistency along the layout→image→layout (L-I-L) loop. Accordingly, we organize training into three stages, as shown in Fig. 2. *First*, a Parallel Multi-task Pretraining phase establishes base abilities for both tasks while sharing image-token pathways to accelerate training. *Second*, a Dual Joint Optimization phase nests generation and grounding by feeding the synthesized image directly into the grounding branch, forming a single joint objective that strengthens L-I-L cycle consistency. *Finally*, a Cycled Reinforcement Learning stage executes the L-I-L loop and treats the spatial discrepancy between input and recovered layouts as a continuous reward, enabling self-supervised optimization without explicit supervision of intermediate visual outputs.

Parallel Multi-task Pre-training

Pretrained unified models often lack native layout-to-image synthesis capabilities or exhibit insufficient grounding proficiency. To prepare for subsequent stages, we introduce a multi-task learning phase that jointly acquires both layout-to-image and grounding abilities.

In both tasks, the proportion of visual tokens is overwhelmingly large—typically on the order of hundreds or even thousands—whereas the number of text tokens is usually only in the tens. Moreover, a single image can simultaneously serve as the output of layout-to-image synthesis and the input of image grounding. Leveraging this duality, we concatenate the inputs of two tasks at the to-

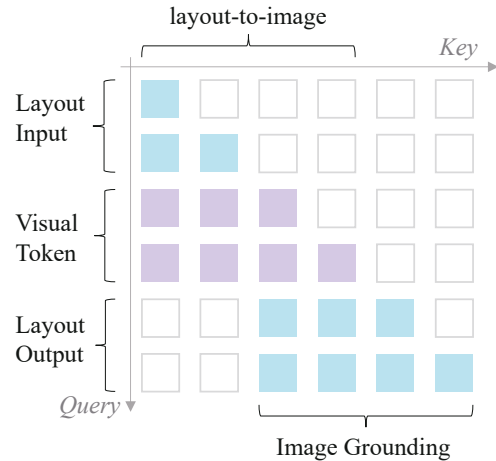


Figure 3: Attention Mask of parallel multi-task pre-training.

ken level, so that a large subset of visual tokens is shared across tasks, thereby enabling efficient parallel learning. Concretely, given a grounding–image pair (X_g, X_i) , we form the input sequence $[X_g, X_i, X_g]$ and feed it into an autoregressive Transformer, which produces the predicted sequence $[Y_g, Y_i, Y_g]$. As shown in Fig. 3, a task-aware attention-mask is designed to prevent cross-task information leakage. We then optimize the summed next-token prediction losses:

$$\mathcal{L}_{\text{pretrain}} = \text{CE}(X_i, Y_i) + \text{CE}(X_g, Y_g), \quad (1)$$

where CE denotes the cross-entropy loss function. This design efficiently enables the unified model with both layout-to-image and grounding capabilities.

Dual Joint Optimization

At this stage, we exploit the input–output duality between layout-to-image generation (L2I) and image grounding (I2L) by *nesting* them into a layout→image→layout (L–I–L) loop, yielding a joint learning objective for collaborative optimization that strengthens cycle-mapping consistency and prepares the model for the subsequent self-supervised RL stage.

Let \mathcal{L} and \mathcal{I} denote the layout- and image-token spaces, respectively. A unified model $p_\theta(\cdot | \cdot)$ with shared parameters θ implements $f_\theta : \mathcal{L} \rightarrow \mathcal{I}$ and $g_\theta : \mathcal{I} \rightarrow \mathcal{L}$. On a dataset $\mathcal{D} = \{(\ell, i)\}$, we define the supervised (token-level) cross-entropy losses

$$\begin{aligned}\mathcal{L}_{\text{I2L}}(\theta) &= \mathbb{E}_{(\ell, i) \sim \mathcal{D}}[-\log p_\theta(\ell | i)], \\ \mathcal{L}_{\text{L2I}}(\theta) &= \mathbb{E}_{(\ell, i) \sim \mathcal{D}}[-\log p_\theta(i | \ell)]\end{aligned}$$

and the loop loss

$$\mathcal{L}_{\text{loop}}(\theta) = \mathbb{E}_\ell \left[-\log p_\theta(\ell | \hat{i}) \right], \quad \text{where } \hat{i} \sim p_\theta(\cdot | \ell).$$

Remark. $\mathcal{L}_{\text{loop}}$ is also a (sequence-wise) cross-entropy objective that measures the negative log-likelihood of reconstructing ℓ from the sampled intermediate image tokens \hat{i} .

Notably, any global minimizer of $\mathcal{L}_{\text{L2I}} + \mathcal{L}_{\text{I2L}}$ on a noise-free \mathcal{D} is also a global minimizer of $\mathcal{L}_{\text{loop}}$. Specifically, at a global optimum θ^* we have $p_{\theta^*}(i | \ell) = \delta_{F(\ell)}$ and $p_{\theta^*}(\ell | i) = \delta_{G(i)}$, hence sampling $\hat{i} \sim p_{\theta^*}(\cdot | \ell)$ yields $\hat{i} = F(\ell)$ a.s., and therefore $p_{\theta^*}(\ell | \hat{i}) = 1$, making $\mathcal{L}_{\text{loop}}(\theta^*) = 0$.

However, minimizing $\mathcal{L}_{\text{loop}}$ alone does not preclude *shortcuts* that bypass faithful image synthesis (e.g., mapping ℓ to arbitrary “code” tokens that decode back to ℓ). We thus regularize the loop objective with the explicit visual supervision of L2I by the following joint loss:

$$\mathcal{J}_{\text{joint}}(\theta) = \mathcal{L}_{\text{L2I}}(\theta) + \lambda \mathcal{L}_{\text{loop}}(\theta), \quad \lambda > 0. \quad (2)$$

Both terms in $\mathcal{J}_{\text{joint}}$ are cross-entropy losses: \mathcal{L}_{L2I} enforces likelihood of visually correct image tokens given the layout, while $\lambda \mathcal{L}_{\text{loop}}$ enforces likelihood of reconstructing the original layout after passing through the sampled image-token path. Compared with the standalone grounding loss, the *nested* joint objective encourages “generate then return” consistency ($g_\theta \circ f_\theta \approx \text{Id}$) and backpropagates through the same image-token pathway as L2I, yielding better gradient alignment and smoother updates that foster synergy between the two tasks.

Because sampling intermediate image tokens is non-differentiable in an autoregressive transformer, we adopt a Gumbel–Softmax approximation to maintain gradient connectivity along the L–I–L loop. At decoding step t with logits \mathbf{z}_t , we draw Gumbel noise $\gamma_t \sim \text{Gumbel}(0, 1)$ and form the differentiable sample

$$\tilde{\mathbf{y}}_t = \text{softmax}((\mathbf{z}_t + \gamma_t)/\tau_t),$$

using the straight-through estimator in the forward pass, while backpropagating through $\tilde{\mathbf{y}}_t$. To balance exploration and discreteness, the temperature parameter follows an annealing schedule over training:

$$\tau_t \equiv \tau(k) = \max(\tau_{\min}, \tau_0 \alpha^k), \quad 0 < \alpha < 1,$$

where k indexes the update steps, τ_0 is the initial temperature, and τ_{\min} prevents premature collapse. This schedule gradually sharpens $\tilde{\mathbf{y}}_t$ toward one-hot selections, preserving stable gradients early and near-discrete decoding later. Combined with (2), this yields end-to-end differentiable, collaboratively optimized layout-to-image generation and image grounding.

Cycled Reinforcement Learning

At this stage, the model has acquired reliable layout-to-image (L2I) and image-to-layout grounding (I2L) capabilities and exhibits sufficient layout→image→layout (L–I–L) cycle consistency. We therefore perform *self-supervised* RL by using the L–I–L layout discrepancy as a continuous reward, without supervising the intermediate visual output. This is justified because the preceding stages ensure that RL samples yield meaningful visual signals, preventing degradation of the generation ability.

Concretely, let a grounding sequence be $X_g = \{(x_e^k, x_b^k)\}_{k=1}^K$, where x_e^k is the referring expression and x_b^k the ground-truth box of the k -th instance. We draw an image sample $\hat{i} \sim p_\theta(\cdot | X_g)$ via L2I, and then ground each expression on \hat{i} via I2L to obtain predicted boxes $\{\hat{y}_b^k\}_{k=1}^K$. Let $d(\cdot, \cdot)$ be a differentiable box discrepancy (e.g., 1–IoU, ℓ_1), and define the per-sample aggregate discrepancy as a continuous reward

$$r_{\text{bbox}} = \frac{1}{K} \sum_{k=1}^K d(\hat{y}_b^k, x_b^k). \quad (3)$$

We apply GRPO (Shao et al. 2024) over groups $\{r_{\text{bbox}}^i\}_{i=1}^G$, using the group-relative advantage $A_{\text{bbox}}^i = r_{\text{bbox}}^i - \frac{1}{G} \sum_{j=1}^G r_{\text{bbox}}^j$. Given an input query q , GRPO samples a set of candidate outputs $\{o^1, o^2, \dots, o^G\}$ using the current policy π_{old} . The new policy π_θ is optimized to maximize the following objective:

$$\begin{aligned}\mathcal{J}(\theta) &= \mathbb{E}_{q \sim \mathcal{D}, \{o^i\}_{i=1}^G \sim \pi_{\text{old}}(\cdot | q)} \left[\frac{1}{G} \sum_{i=1}^G \left(\right. \\ &\quad \min \left(r_{\text{bbox}}^i A_{\text{bbox}}^i, \text{clip}(r_{\text{bbox}}^i, 1-\epsilon, 1+\epsilon) A_{\text{bbox}}^i \right) \\ &\quad \left. - \beta \mathbb{D}_{\text{KL}}(\pi_\theta \| \pi_{\text{ref}}) \right) \right] \quad (4)\end{aligned}$$

where $r_{\text{bbox}}^i = \frac{\pi_\theta(o^i | q)}{\pi_{\text{old}}(o^i | q)}$ is the importance weight, A_{bbox}^i denotes the group-relative advantage, and β is a weighting factor on the regularization term. Together with the KL regularizer, the policy is optimized by the GRPO objective, yielding self-supervised, looped reinforcement learning driven solely by cycle-consistent layout rewards.

Experiments

Implementation Details

The experiments are conducted with a pretrained *Janus-Pro 1.5B* (Chen et al. 2025c) backbone on NVIDIA A100–80GB

Method	MS COCO					LayoutSAM-Eval			
	AP \uparrow	AP $_{50}$ \uparrow	AP $_{75}$ \uparrow	CLIP \uparrow	FID \downarrow	Spatial \uparrow	Color \uparrow	Texture \uparrow	Shape \uparrow
<i>Gen. only</i>									
GLIGEN	30.99	41.44	36.89	24.81	27.93	77.53	49.41	55.29	52.72
MIGC	46.16	58.29	50.21	25.04	25.35	85.66	66.97	71.24	69.06
InstanceDiffusion	49.97	61.33	52.65	25.15	25.00	87.99	69.16	72.78	71.08
IFAdapter	43.31	56.29	52.76	26.39	23.45	84.32	68.32	71.37	68.37
EliGen	45.83	59.26	49.66	27.44	21.43	85.37	72.35	75.42	74.59
<i>Und. and Gen.</i>									
PlanGen	51.39	64.70	61.10	25.33	20.44	92.21	82.69	86.53	85.36
EchoGen (Ours)	54.61	68.85	65.01	25.18	20.12	96.32	84.97	89.02	87.18

Table 1: Quantitative results on MS COCO and LayoutSAM-Eval benchmark. Method marked as *Gen. only* perform layout-to-image generation only, while *Und. and Gen.* denotes unified layout-to-image generation and grounding models.

Stage	Samples	Steps
Parallel Multi-task Pre-training	4M	125K
Dual Joint Optimization	2M	60K
Cycled RL	50K	50K

Table 2: Training recipe.

GPUs. Across all three stages, we use the AdamW optimizer with a fixed learning rate of 5×10^{-5} ; the mini-batch size is 32 for Stages 1 and 2. In the reinforcement learning stage, we adopt Group Relative Policy Optimization (GRPO) with a group size of 8. To ensure gradient propagation between the two tasks during Stage 2 (Dual Joint Optimization), we use the VQ-VAE token sequence as the visual representation for image grounding. To preserve the model’s pretrained visual understanding, we randomly retain ViT-encoded tokens with probability 50%, and we keep both the VQ-VAE encoder and the ViT encoder frozen throughout all stages.

Training data. Starting from GRIT-20M, we filter bounding boxes by confidence ≥ 0.8 and discard images containing fewer than 2 or more than 10 boxes. We then sample approximately 6M layout–image pairs for training. The number of samples and optimization steps per stage are summarized in Tab. 2.

Experiments Setup

layout-to-image Evaluation. We evaluate our layout-to-image synthesis on *MS-COCO* (Lin et al. 2014) and *LayoutSAM-Eval* (Zhang et al. 2024). For MS-COCO, we employ *Grounding-SAM* (Kirillov et al. 2023; Ren et al. 2024) to detect objects in the generated images and obtain bounding boxes; we set box threshold to 0.3 and text threshold to 0.25 without non-maximum suppression (NMS); standard COCO metrics are then computed for boxes, including AP, AP $_{50}$, and AP $_{75}$. In addition, we report the *CLIP* score (Radford et al. 2021) to assess global text–image

alignment, and the Fréchet Inception Distance (FID) (Heusel et al. 2017) to quantify distributional image quality. For *LayoutSAM-Eval*, we adopt *MiniCPM-V-2.6* (Yao et al. 2024) as the evaluation vision–language model and report four dimensions—*Spatial*, *Color*, *Texture*, and *Shape*—which probe instance-level controllability of the layout-to-image generation along spatial arrangement and appearance factors.

Image Grounding Evaluation. We assess grounding on the *Ref-L4* benchmark (2025a), tailored to modern referring expression comprehension (REC) with long, information-dense queries and broad visual diversity. Following the official protocol, we report Acc $_{0.5}$, Acc $_{0.75}$, Acc $_{0.9}$, and mAcc.

Baselines. For *layout-to-image* synthesis, we consider six representative methods based on diffusion backbones or auto-regressive transformers. *GLIGEN* (Li et al. 2023), *MIGC* (Zhou et al. 2024b), and *InstanceDiffusion* (Wang et al. 2024b) are all built upon Stable Diffusion v1.5 (Rombach et al. 2022), augmenting the backbone with task-specific modules to ingest layout signals; *IFAdapter* (Wu et al. 2024) is based on Stable Diffusion XL (Podell et al. 2023) with an adapter that explicitly learns layout control; *EliGen* (Zhang et al. 2025a) follows a *DiT* (2023) architecture and fine-tunes FLUX-dev, targeting precise *instance-level* controllability; and *PlanGen* (He et al. 2025) serves as a unified generation–understanding model based on auto-regressive transformers that supports both layout conditioning and image understanding. For *image grounding*, we evaluate three large vision–language models—*KOSMOS-2* (2023), Qwen-VL-Chat (2023), and *CogVLM-grounding* (2024a).

Layout-to-image Generation

Quantitative results. On *MS-COCO*, our approach surpasses both diffusion-based and unified-model baselines in layout-control accuracy. Relative to the previous best, we obtain absolute gains of 3.22 in AP, 4.15 in AP $_{50}$, and 3.92 in AP $_{75}$. Our method also achieves higher CLIP scores than

Prompt	Layout	Ours	Gligen	MIGC	IFAdapter	PlanGen
many dogs and cats on the grass.	<div style="display: flex; flex-wrap: wrap;"> <div style="border: 1px dashed black; padding: 2px; margin: 2px;">black dog</div> <div style="border: 1px dashed blue; padding: 2px; margin: 2px;">blue cat</div> <div style="border: 1px dashed yellow; padding: 2px; margin: 2px;">yellow dog</div> <div style="border: 1px dashed white; padding: 2px; margin: 2px;">white cat</div> </div>					
a photo showcasing a city street scene, with a person walking down the street.	<div style="display: flex; flex-wrap: wrap;"> <div style="border: 1px dashed brown; padding: 2px; margin: 2px;">wooden window</div> <div style="border: 1px dashed orange; padding: 2px; margin: 2px;">door</div> <div style="border: 1px dashed yellow; padding: 2px; margin: 2px;">wooden door</div> </div>					
a photograph showcasing a historic bridge and its surrounding buildings	<div style="display: flex; flex-wrap: wrap;"> <div style="border: 1px dashed white; padding: 2px; margin: 2px;">White clouds with blue sky</div> <div style="border: 1px dashed blue; padding: 2px; margin: 2px;">tower</div> <div style="border: 1px dashed black; padding: 2px; margin: 2px;">bridge</div> </div>					
a photo depicting a city street scene on a rainy day	<div style="display: flex; flex-wrap: wrap;"> <div style="border: 1px dashed blue; padding: 2px; margin: 2px;">car</div> <div style="border: 1px dashed black; padding: 2px; margin: 2px;">sedan</div> <div style="border: 1px dashed yellow; padding: 2px; margin: 2px;">woman with umbrella</div> <div style="border: 1px dashed black; padding: 2px; margin: 2px;">Person on bicycle</div> </div>					
a park landscape, with a tall palm tree and a lush green lawn	<div style="display: flex; flex-wrap: wrap;"> <div style="border: 1px dashed black; padding: 2px; margin: 2px;">Tall palm tree</div> <div style="border: 1px dashed blue; padding: 2px; margin: 2px;">Clear cloudless day</div> <div style="border: 1px dashed black; padding: 2px; margin: 2px;">Church with a tall tower</div> <div style="border: 1px dashed green; padding: 2px; margin: 2px;">grassy area</div> <div style="border: 1px dashed red; padding: 2px; margin: 2px;">Woman bench</div> </div>					

Figure 4: Quantitative comparison on layout-to-image generation. The first column lists the global text prompts; the second column visualizes the layout conditions, rendered as (*referring-expression*, *bounding-box*) pairs. We highlight a subset of negatives and positives with red and green boxes; note that negatives may correspond to *incorrect spatial locations or attributes* (e.g., *color*).

SD1.5-based methods and *PlanGen*, indicating stronger prompt faithfulness. The remaining CLIP gap to *IFAdapter* and *EliGen* is plausibly explained by their stronger base models—SDXL and FLUX-dev—with substantially larger parameter counts (3.5B and 12B). In terms of image quality, our method attains the best Fréchet Inception Distance (FID). Notably, *EliGen*, *PlanGen*, and our model all maintain a clear FID advantage over the other baselines; we hypothesize this stems from their Transformer-based design that treats the layout condition as contextual tokens, in contrast to approaches that add extra modules to enforce explicit spatial constraints on latent-space, which can introduce artifacts and degrade perceptual quality. On *LayoutSAM-Eval*, our approach establishes a new state of the art, improving the *Spatial*, *Color*, *Texture*, and *Shape* dimensions by 4.11, 2.28, 2.49, and 1.82, respectively, with the largest gain observed for spatial control.

Qualitative results. Figure 4 presents visual comparisons

under diverse layout conditions. For *complex* specifications (e.g., Row 1 with multiple cats and dogs of distinct colors), *GLIGEN*, *MIGC*, and *IFAdapter* tend to yield lower perceptual fidelity with overly smoothed, cartoon-like textures, whereas our method maintains consistently high image quality. Across all settings, our approach achieves precise *spatial control* (placement and scale) and accurate *attribute control* (e.g., *color*), faithfully adhering to the specified layouts.

Image Grounding

Tab. 3 summarizes the quantitative comparison on Ref-L4. Despite a smaller parameter budget, our method EchoGen achieves the best overall performance. Compared to the previous state of the art, we achieve absolute gains of 1.50 in $Acc_{0.5}$, 4.65 in $Acc_{0.75}$, and 2.37 in $mAcc$. These results highlight the synergistic effect of our unified architecture and training strategy, enabling faster performance gains with reduced training resources.

Method	Ref-L4			
	Acc _{0.5}	Acc _{0.75}	Acc _{0.9}	mAcc
<i>Und. only</i>				
KOSMOS-2	48.53	38.34	17.54	34.72
Qwen-VL-Chat	73.80	58.05	37.16	55.94
CogVLM-g.	81.70	70.77	48.35	66.09
<i>Und. and Gen.</i>				
Ours	83.20	75.42	47.32	68.46

Table 3: Quantitative results on Ref-L4. Method marked as *Und. only* perform image grounding only, while *Und. and Gen.* denotes unified layout-to-image generation and grounding models.

	Stage 1	Stage 2	Stage 3	AP \uparrow	FID \downarrow
#1	✓			47.26	21.98
#2	✓	✓		52.38	22.23
#3	✓		✓	48.01	20.32
#4	✓	✓	✓	54.61	20.12

Table 4: Ablation over training stages.

Ablation Study

Ablation for training stage. We compare different stage configurations on MS-COCO (Tab. 4). Adding *Stage 2* on top of *Stage 1* (#1 \rightarrow #2) yields a sizeable gain of 5.12 AP (from 47.26 to 52.38), and introducing *Stage 3* thereafter (#2 \rightarrow #4) provides a further 2.23 AP improvement, totaling 7.35 over #1. By contrast, skipping Stage 2 and directly applying Stage 3 (#1 \rightarrow #3) produces only a marginal AP increase (0.75), underscoring the necessity of Stage 2 for strengthening layout-image-layout consistency. Regarding image quality, FID varies modestly across stages; notably, Stage 3 consistently reduces FID, which we hypothesize arises from the additional diversity introduced by stochastic sampling.

Ablation on stage transition points. To examine how the entry point into later stages affects performance, we vary the training budget of each stage on MS-COCO and measure downstream outcomes under a fixed budget for the subsequent stage. Fig. 5 (left) plots the *post-Stage 2* AP against the AP achieved *right before* switching from Stage 1 to Stage 2. When the base capability is low (e.g., AP < 20), Stage 2 yields *negative transfer*, likely because poor layout-to-image generations induce misaligned supervision on the image grounding branch, leading to incorrect gradient signals. As the base AP improves, Stage 2 confers increasingly larger gains, indicating that collaborative training is most effective once the generator reaches a minimally reliable quality. Fig. 5 (right) shows an analogous study for the Stage 2 \rightarrow Stage 3 transition. Since Stage 3 (RL) does not provide explicit token-level supervision on visual tokens, models entering with weaker base accuracy (AP < 40) suf-




Figure 5: Ablation on stage transition points. Left: gains when switching from Stage 1 \rightarrow Stage 2; Right: gains when switching from Stage 2 \rightarrow Stage 3. The x-axis records the AP at entry to the current stage (determined by the preceding stage’s training budget), while the y-axis reports the AP improvement achieved by that stage.

Dataset	AP	AP ₅₀	AP ₇₅
From real dataset	54.61	68.85	65.01
Random	53.96	68.02	64.32
Generated by GPT-4o	54.51	68.94	64.59

Table 5: Ablation on Stage 3 (RL) training set.

fer pronounced performance drops, consistent with the difficulty of sampling sufficiently accurate generations to obtain informative rewards and gradients. Note that AP has a practical lower bound of roughly ~ 3.0 . Consequently, when the pre-switch AP is extremely low, the absolute drop is capped by this floor, so the measured degradation appears smaller; as the baseline AP increases toward ~ 30 , the potential absolute decrease correspondingly grows.

Ablation on Stage 3 training set. Stage 3 (RL) can exploit *layout text alone* as a self-supervisory signal, obviating the need for paired images. To examine whether the layout text must originate from a real paired corpus, we compare three construction strategies on MS-COCO in Tab. 5: (i) sampling layouts from the real dataset; (ii) *random* layouts with a minimal bounding-box size threshold to avoid degenerate cases; and (iii) layouts generated by GPT-4o. Across these settings, reinforcement learning yields *comparable* performance, indicating that Stage 3 has strong data robustness.

Conclusion

We present EchoGen, a unified framework for layout-to-image generation and image grounding that exploits their duality to achieve accurate layouts and robust grounding in one model. To address naïve joint-training challenges, we use a progressive scheme: Parallel Multi-Task Pre-training for base abilities with shared image-token pathways, Dual Joint Optimization to serialize generation and grounding into a unified objective that strengthens layout-image-layout loop consistency, and Cycle RL that uses grounding discrepancy as rewards, enabling self-supervised learning without visual labels.

References

- Antol, S.; Agrawal, A.; Lu, J.; Mitchell, M.; Batra, D.; Zitnick, C. L.; and Parikh, D. 2015. Vqa: Visual question answering. In *Proceedings of the IEEE international conference on computer vision*, 2425–2433.
- Avrahami, O.; Hayes, T.; Gafni, O.; Gupta, S.; Taigman, Y.; Parikh, D.; Lischinski, D.; Fried, O.; and Yin, X. 2023. Spa-Text: Spatio-Textual Representation for Controllable Image Generation. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE.
- Bai, J.; Bai, S.; Chu, Y.; Cui, Z.; Dang, K.; Deng, X.; Fan, Y.; Ge, W.; Han, Y.; Huang, F.; et al. 2023. Qwen technical report. *arXiv preprint arXiv:2309.16609*.
- Bai, S.; Chen, K.; Liu, X.; Wang, J.; Ge, W.; Song, S.; Dang, K.; Wang, P.; Wang, S.; Tang, J.; et al. 2025. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*.
- Bar-Tal, O.; Yariv, L.; Lipman, Y.; and Dekel, T. 2023. MultiDiffusion: Fusing Diffusion Paths for Controlled Image Generation. In Krause, A.; Brunskill, E.; Cho, K.; Engelhardt, B.; Sabato, S.; and Scarlett, J., eds., *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA*, volume 202 of *Proceedings of Machine Learning Research*, 1737–1752. PMLR.
- Chen, J.; Wei, F.; Zhao, J.; Song, S.; Wu, B.; Peng, Z.; Chan, S.-H. G.; and Zhang, H. 2025a. Revisiting referring expression comprehension evaluation in the era of large multimodal models. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 513–524.
- Chen, M.; Laina, I.; and Vedaldi, A. 2024. Training-free layout control with cross-attention guidance. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 5343–5353.
- Chen, S.; Lai, J.; Gao, J.; Ye, T.; Chen, H.; Shi, H.; Shao, S.; Lin, Y.; Fei, S.; Xing, Z.; et al. 2025b. PosterCraft: Rethinking High-Quality Aesthetic Poster Generation in a Unified Framework. *arXiv preprint arXiv:2506.10741*.
- Chen, X.; Wu, Z.; Liu, X.; Pan, Z.; Liu, W.; Xie, Z.; Yu, X.; and Ruan, C. 2025c. Janus-pro: Unified multimodal understanding and generation with data and model scaling. *arXiv preprint arXiv:2501.17811*.
- Cheng, T.; Song, L.; Ge, Y.; Liu, W.; Wang, X.; and Shan, Y. 2024. Yolo-world: Real-time open-vocabulary object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 16901–16911.
- Deng, C.; Zhu, D.; Li, K.; Gou, C.; Li, F.; Wang, Z.; Zhong, S.; Yu, W.; Nie, X.; Song, Z.; et al. 2025. Emerging properties in unified multimodal pretraining. *arXiv preprint arXiv:2505.14683*.
- He, R.; Cheng, B.; Ma, Y.; Jia, Q.; Liu, S.; Ma, A.; Wu, X.; Wu, L.; Leng, D.; and Yin, Y. 2025. PlanGen: Towards Unified Layout Planning and Image Generation in Auto-Regressive Vision Language Models. *arXiv preprint arXiv:2503.10127*.
- Heusel, M.; Ramsauer, H.; Unterthiner, T.; Nessler, B.; and Hochreiter, S. 2017. GANs Trained by a Two Time-Scale Update Rule Converge to a Local Nash Equilibrium. In Guyon, I.; von Luxburg, U.; Bengio, S.; Wallach, H. M.; Fergus, R.; Vishwanathan, S. V. N.; and Garnett, R., eds., *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, 6626–6637.
- Hoe, J. T.; Jiang, X.; Chan, C. S.; Tan, Y.; and Hu, W. 2023. InteractDiffusion: Interaction Control in Text-to-Image Diffusion Models. *ArXiv preprint*, abs/2312.05849.
- Kirillov, A.; Mintun, E.; Ravi, N.; Mao, H.; Rolland, C.; Gustafson, L.; Xiao, T.; Whitehead, S.; Berg, A. C.; Lo, W.-Y.; et al. 2023. Segment anything. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 4015–4026.
- Li, L. H.; Zhang, P.; Zhang, H.; Yang, J.; Li, C.; Zhong, Y.; Wang, L.; Yuan, L.; Zhang, L.; Hwang, J.-N.; et al. 2022. Grounded language-image pre-training. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10965–10975.
- Li, Y.; Liu, H.; Wu, Q.; Mu, F.; Yang, J.; Gao, J.; Li, C.; and Lee, Y. J. 2023. GLIGEN: Open-Set Grounded Text-to-Image Generation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, June 17-24, 2023*, 22511–22521. IEEE.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part V 13*, 740–755. Springer.
- Liu, S.; Zeng, Z.; Ren, T.; Li, F.; Zhang, H.; Yang, J.; Jiang, Q.; Li, C.; Yang, J.; Su, H.; et al. 2024. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. In *European conference on computer vision*, 38–55. Springer.
- Liu, Z.; Yu, Y.; Ouyang, H.; Wang, Q.; Cheng, K. L.; Wang, W.; Liu, Z.; Chen, Q.; and Shen, Y. 2025. Magicquill: An intelligent interactive image editing system. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 13072–13082.
- Peebles, W.; and Xie, S. 2023. Scalable Diffusion Models with Transformers. In *IEEE/CVF International Conference on Computer Vision, ICCV 2023, Paris, France, October 1-6, 2023*, 4172–4182. IEEE.
- Peng, Z.; Wang, W.; Dong, L.; Hao, Y.; Huang, S.; Ma, S.; and Wei, F. 2023. Kosmos-2: Grounding multimodal large language models to the world. *arXiv preprint arXiv:2306.14824*.
- Podell, D.; English, Z.; Lacey, K.; Blattmann, A.; Dockhorn, T.; Müller, J.; Penna, J.; and Rombach, R. 2023. SDXL: Improving Latent Diffusion Models for High-Resolution Image Synthesis. In *The Twelfth International Conference on Learning Representations*.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; Krueger, G.; and Sutskever, I. 2021. Learning Transferable Visual Models From Natural Language Supervision. In

- Meila, M.; and Zhang, T., eds., *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, 8748–8763. PMLR.
- Ren, T.; Liu, S.; Zeng, A.; Lin, J.; Li, K.; Cao, H.; Chen, J.; Huang, X.; Chen, Y.; Yan, F.; Zeng, Z.; Zhang, H.; Li, F.; Yang, J.; Li, H.; Jiang, Q.; and Zhang, L. 2024. Grounded SAM: Assembling Open-World Models for Diverse Visual Tasks. *ArXiv preprint*, abs/2401.14159.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-Resolution Image Synthesis with Latent Diffusion Models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, 10674–10685. IEEE.
- Shao, Z.; Wang, P.; Zhu, Q.; Xu, R.; Song, J.; Bi, X.; Zhang, H.; Zhang, M.; Li, Y.; Wu, Y.; et al. 2024. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*.
- Tian, Y.; Shao, T.; Demizu, T.; Wu, X.; and Wu, H.-T. 2024. HPE-CogVLM: Advancing Vision Language Models with a Head Pose Grounding Task. *arXiv preprint arXiv:2406.01914*.
- Wang, W.; Lv, Q.; Yu, W.; Hong, W.; Qi, J.; Wang, Y.; Ji, J.; Yang, Z.; Zhao, L.; XiXuan, S.; et al. 2024a. Cogvlm: Visual expert for pretrained language models. *Advances in Neural Information Processing Systems*, 37: 121475–121499.
- Wang, X.; Darrell, T.; Rambhatla, S. S.; Girdhar, R.; and Misra, I. 2024b. InstanceDiffusion: Instance-level Control for Image Generation.
- Wu, C.; Chen, X.; Wu, Z.; Ma, Y.; Liu, X.; Pan, Z.; Liu, W.; Xie, Z.; Yu, X.; Ruan, C.; et al. 2025. Janus: Decoupling visual encoding for unified multimodal understanding and generation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 12966–12977.
- Wu, Y.; Zhou, X.; Ma, B.; Su, X.; Ma, K.; and Wang, X. 2024. Ifadapter: Instance feature control for grounded text-to-image generation. *arXiv preprint arXiv:2409.08240*.
- Xiao, J.; Lv, H.; Li, L.; Wang, S.; and Huang, Q. 2023. R&B: Region and Boundary Aware Zero-shot Grounded Text-to-image Generation. *arXiv:2310.08872*.
- Xie, J.; Li, Y.; Huang, Y.; Liu, H.; Zhang, W.; Zheng, Y.; and Shou, M. Z. 2023. BoxDiff: Text-to-Image Synthesis with Training-Free Box-Constrained Diffusion. In *IEEE/CVF International Conference on Computer Vision, ICCV 2023, Paris, France, October 1-6, 2023*, 7418–7427. IEEE.
- Yang, Z.; Wang, J.; Gan, Z.; Li, L.; Lin, K.; Wu, C.; Duan, N.; Liu, Z.; Liu, C.; Zeng, M.; and Wang, L. 2023. ReCo: Region-Controlled Text-to-Image Generation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, June 17-24, 2023*, 14246–14255. IEEE.
- Yao, Y.; Yu, T.; Zhang, A.; Wang, C.; Cui, J.; Zhu, H.; Cai, T.; Li, H.; Zhao, W.; He, Z.; et al. 2024. Minicpm-v: A gpt-4v level mllm on your phone. *arXiv preprint arXiv:2408.01800*.
- Yu, L.; Poirson, P.; Yang, S.; Berg, A. C.; and Berg, T. L. 2016. Modeling context in referring expressions. In *European conference on computer vision*, 69–85. Springer.
- Zhang, H.; Duan, Z.; Wang, X.; Chen, Y.; and Zhang, Y. 2025a. Eligen: Entity-level controlled image generation with regional attention. *arXiv preprint arXiv:2501.01097*.
- Zhang, H.; Hong, D.; Wang, Y.; Shao, J.; Wu, X.; Wu, Z.; and Jiang, Y.-G. 2024. Creatilayout: Siamese multimodal diffusion transformer for creative layout-to-image generation. *arXiv preprint arXiv:2412.03859*.
- Zhang, X.; Guo, J.; Zhao, S.; Fu, M.; Duan, L.; Wang, G.-H.; Chen, Q.-G.; Xu, Z.; Luo, W.; and Zhang, K. 2025b. Unified multimodal understanding and generation models: Advances, challenges, and opportunities. *arXiv preprint arXiv:2505.02567*.
- Zheng, G.; Zhou, X.; Li, X.; Qi, Z.; Shan, Y.; and Li, X. 2023. Layoutdiffusion: Controllable diffusion model for layout-to-image generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 22490–22499.
- Zhou, C.; Yu, L.; Babu, A.; Tirumala, K.; Yasunaga, M.; Shamis, L.; Kahn, J.; Ma, X.; Zettlemoyer, L.; and Levy, O. 2024a. Transfusion: Predict the next token and diffuse images with one multi-modal model. *arXiv preprint arXiv:2408.11039*.
- Zhou, D.; Li, Y.; Ma, F.; Zhang, X.; and Yang, Y. 2024b. MIGC: Multi-Instance Generation Controller for Text-to-Image Synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Zou, K.; Huang, Z.; Dong, Y.; Tian, S.; Zheng, D.; Liu, H.; He, J.; Liu, B.; Qiao, Y.; and Liu, Z. 2025. Uni-MMMU: A Massive Multi-discipline Multimodal Unified Benchmark. *arXiv preprint arXiv:2510.13759*.