

KOSMOS-G: Generating Images in Context with Multimodal Large Language Models

Xichen Pan^{1,2*} Li Dong¹ Shaohan Huang¹ Zhiliang Peng¹ Wenhui Chen³ Furu Wei¹

¹Microsoft Research ²New York University ³University of Waterloo

<https://aka.ms/GeneralAI>



Figure 1: Zero-shot subject-driven generation examples with multimodal prompts. Thanks to the advanced multimodal perception capabilities of MLLM, KOSMOS-G can generate high-fidelity subject-driven images by approaching all image inputs as a “foreign language”.

* Contribution during internship at Microsoft Research.

Abstract

Recent advancements in subject-driven image generation have made significant strides. However, current methods still fall short in diverse application scenarios, as they require test-time tuning and cannot accept interleaved multi-image and text input. These limitations keep them far from the ultimate goal of “image as a foreign language in image generation.” This paper presents KOSMOS-G, a model that leverages the advanced multimodal perception capabilities of Multimodal Large Language Models (MLLMs) to tackle the aforementioned challenge. Our approach aligns the output space of MLLM with CLIP using the textual modality as an anchor and performs compositional instruction tuning on curated data. KOSMOS-G demonstrates an impressive capability of zero-shot subject-driven generation with interleaved multi-image and text input. Notably, the score distillation instruction tuning requires no modifications to the image decoder. This allows for a seamless substitution of CLIP and effortless integration with a myriad of U-Net techniques ranging from fine-grained controls to personalized image decoder variants. We posit KOSMOS-G as an initial attempt towards the goal of “image as a foreign language in image generation.” The code can be found at <https://aka.ms/Kosmos-G>

1 Introduction

In recent studies, advancements in text-to-image (T2I) generation, particularly with diffusion models, have shown remarkable progress in producing highly photorealistic, accurate, and diverse images from textual descriptions. Building on the unprecedented success of producing highly accurate images from text descriptions, numerous studies have delved into more sophisticated subject-driven generation to integrate images into text prompts to generate new customized images.

A group of approaches [GAA⁺22, RLJ⁺22, KZZ⁺23, TGCA23, AAF⁺23, HHZW23, SHZ⁺23] propose to fine-tune the models on each set of reference images, and fail to achieve subject-driven generation through a generalized pre-trained model. [XYF⁺23, WZJ⁺23, CHL⁺23, CHSC22] inject image features into the U-Net of diffusion models. However, such injection methods segregate the guidance for text and images, thereby limiting the effectiveness of joint modeling between the two modalities. Additionally, this approach is challenging to extend to scenarios involving multiple entities. Recent work BLIP-Diffusion [LLH23] learns object representations by synthesizing images through the composition of subjects with random backgrounds. This approach effectively endows it with a zero-shot, subject-driven text-to-image generation capability. However, the specific design of its input template and training data restricts its scalability to multiple entities.

In contrast to previous methods that work with original CLIP text encoder [RKH⁺21], we propose that through leveraging Multimodal Large Language Models (MLLMs) [ADL⁺22, HSD⁺22, AHR⁺22, HDW⁺23, LLSH23], most of the challenges in subject-driven generation may be easily resolved. MLLMs have expanded the perception capabilities of language models to multimodality, enabling them to perceive diverse modalities such as images. The idea of leveraging MLLMs for subject-driven generation presents several advantages: 1) It capitalizes on the inherent vision-language alignment within the MLLM. 2) The MLLM architecture naturally supports interleaved multi-image and text input. 3) The pre-trained MLLM can effectively model multimodal input in context.

To support zero subject-driven generation with interleaved multi-image and text input, we present KOSMOS-G, which leverages the advanced multimodal perception of MLLM following an “align before instruct” manner. Specifically, we start from the multimodal language modeling stage, leading to the KOSMOS-1 [HDW⁺23] MLLM. It envisions language models as a universal task layer, perceiving free-form interleaved vision-language inputs and consolidating various task predictions into textual formats. Given the aligned vision-language representation, we then use the language modality as an anchor and align the output space of the MLLM with the CLIP text encoder. Finally, we perform instruction tuning on the curated data. KOSMOS-G accepts captions as input, where each entity is followed by its segmented image. The model is trained to faithfully reproduce all entities, render the text content, and follow the instructions. In this process, the frozen pre-trained diffusion image decoder serves as a score metric. We distill the learned data distribution to pass the differentiable gradient to the MLLM. This enables KOSMOS-G to harness rich features from the

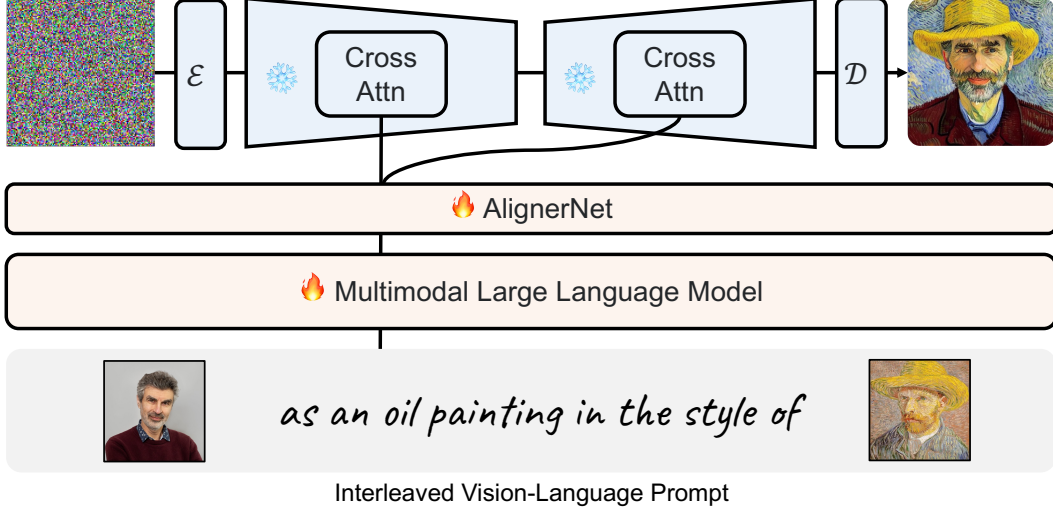


Figure 2: KOSMOS-G comprises an MLLM for multimodal perception, coupled with an AlignerNet that bridges the MLLM to the diffusion U-Net image decoder. KOSMOS-G can pass the fine concept-level guidance from interleaved input to image decoder, and offer a seamless alternative to CLIP. Orange denotes the trainable modules; Blue denotes the frozen ones.

image encoder to generate images faithfully reproducing the contents across various contexts (see Figure 1, we also present examples with more diverse interleaving in Figure 9).

Benefiting from general-purpose pre-training, KOSMOS-G approaches the objective of “image as a foreign language in image generation.” This means KOSMOS-G can capture novel concepts from input images and guide personalized creations in a zero-shot setting. Notably, KOSMOS-G also stands as the first model to master zero-shot subject-driven generation with interleaved multi-image and text input. Owing to the score distillation instruction tuning, KOSMOS-G do not need to modify any parameters of the image decoder, i.e., the diffusion U-Net and VAEs. This makes it possible for us to seamlessly substitute CLIP with KOSMOS-G in any image generation system. As a result, a plethora of applications can be unlocked in conjunction with U-Net techniques, ranging from fine-grained controls like ControlNet [ZA23] to personalized or stylized image decoder variants like amazing community contributed LoRA [HSW⁺22] checkpoints.

Overall, we propose KOSMOS-G as an initial attempt towards the objective of “image as a foreign language in image generation.” We summarize our main contributions as follows:

1. We propose to leverage the advanced multimodal perception of MLLMs for subject-driven generation with interleaved multi-image and text input.
2. We propose a compositional instruction tuning task, leading to amazing zero-shot multi-entity subject-driven generation capability.
3. Score distillation instruction tuning allows KOSMOS-G to seamlessly interface with a spectrum of U-Net techniques, indicating broad applicability and potential for integration into various frameworks.

2 KOSMOS-G: Image as a Foreign Language in Image Generation

As shown in Figure 2, KOSMOS-G is a model that can perceive interleaved multi-image and text input and generate subject-driven conditions. Specifically, the backbone of KOSMOS-G MLLM is a Transformer-based causal language model, serving as a general-purpose interface to multimodal input. We train KOSMOS-G following an “align before instruct” manner, the entire training pipeline can be divided into 3 stages:

1. **Multimodal Language Modeling:** We pre-train the MLLM from scratch on multimodal corpora, including monomodal data, cross-modal paired data, and interleaved multimodal data with language modeling loss following KOSMOS-1.

2. **Image Decoder Aligning:** We use the U-Net [RFB15] of Stable Diffusion v1.5 [RBL⁺22] as our image decoder. We trained an AlignerNet on only textual data to align the output space of KOSMOS-G to U-Net’s input space through CLIP supervision. Here, the language acts as the anchoring modality, ensuring image input is also compatible with the image decoder.
3. **Instruction Tuning:** We further fine-tune KOSMOS-G through a compositional generation task on curated data, with the differentiable gradient passed from the frozen U-Net.

In Stage 1, only the MLLM is trained. In Stage 2, AlignerNet is trained with MLLM frozen. During Stage 3, both AlignerNet and MLLM are jointly trained. The image decoder remains frozen throughout all stages.

2.1 Multimodal Language Modeling

Following KOSMOS-1, KOSMOS-G perceives general modalities in a unified way. To achieve this, we represent the input format as a single sequence using special tokens. Specifically, we use the tokens $\langle s \rangle$ and $\langle /s \rangle$ to denote start- and end-of-sequence. We also incorporate $\langle \text{image} \rangle$ and $\langle / \text{image} \rangle$ tokens to indicate the start and end of any embedded image representations within the sequence.

Our methodology involves encoding both text tokens and images into vectors, which are then fed into the decoder. For text tokens, we use a lookup table to map them into embeddings. To handle the input images, we employ a vision Transformer [DBK⁺21] as the embedding module. Furthermore, Resampler [ADL⁺22] is used as an attentive pooling mechanism to reduce the number of image embeddings. After obtaining the embeddings of an input sequence, we feed them into the Transformer-based decoder. The left-to-right causal decoder processes the sequence in an auto-regressive manner. A softmax classifier on the Transformer is used to assign probabilities to each token in the vocabulary.

KOSMOS-G is first trained using the next-token prediction task. The training objective is to maximize the log-likelihood of tokens in examples. It’s important to note that the training loss only takes into account discrete tokens, specifically text tokens. The MLLM component has 24 layers with 2,048 hidden dimensions, 8,192 FFN intermediate size, and 32 attention heads. For faster convergence, the image representation is obtained from a pre-trained CLIP ViT-L/14 model with 1,024 feature dimensions. The images are preprocessed into 224×224 resolution during training. We freeze the parameters of the CLIP model except for the last layer during training. The total number of parameters of the MLLM is about 1.6B.

2.2 Image Decoder Aligning

After undertaking multimodal language modeling, we have successfully aligned vision and language perception within MLLM. To make KOSMOS-G capable of image generation, we incorporate diffusion models [SWMG15] as our image decoder. Specifically, we adopt the widely accepted Stable Diffusion v1.5 [RBL⁺22]. It’s important to note that we only replace the CLIP text encoder [RKH⁺21] with multimodal KOSMOS-G, without making any modifications to the U-Net architecture or weight. This setup allows KOSMOS-G to effectively collaborate with techniques applied to the U-Net, like ControlNet [ZA23] and various community LoRA [HSW⁺22] variants. In this section, we will provide brief preliminaries of latent diffusion models, and then delve into the process of aligning the output space of KOSMOS-G with the image decoder after the aforementioned replacement.

Preliminaries of Latent Diffusion Models Diffusion models define a Markov chain of forward diffusion process q , adding Gaussian noise samples to the initial real data $\mathbf{z}_0 \sim q(\mathbf{z})$ over T steps. Here, \mathbf{z} denotes latent representations rather than pixel values. The efficient, low-dimensional latent space is approximately perceptually equivalent to high-dimensional RGB space, while the redundant semantically meaningless information present in the pixel domain is eliminated. Perceptual compression models (i.e., VQ-VAE) consisting of \mathcal{E} and \mathcal{D} encode the real data into the latent space and reverse, such that $\mathcal{D}(\mathcal{E}(\mathbf{x})) \approx \mathbf{x}$. Latent diffusion models use latent representations $\mathbf{z} = \mathcal{E}(\mathbf{x})$ instead of working directly with pixel values during the diffusion process. The final output can be decoded back to pixel space via $\mathcal{D}(\mathbf{z})$. The separate mild perceptual compression stage only eliminates imperceptible details, leading to competitive generation results with a much lower cost.

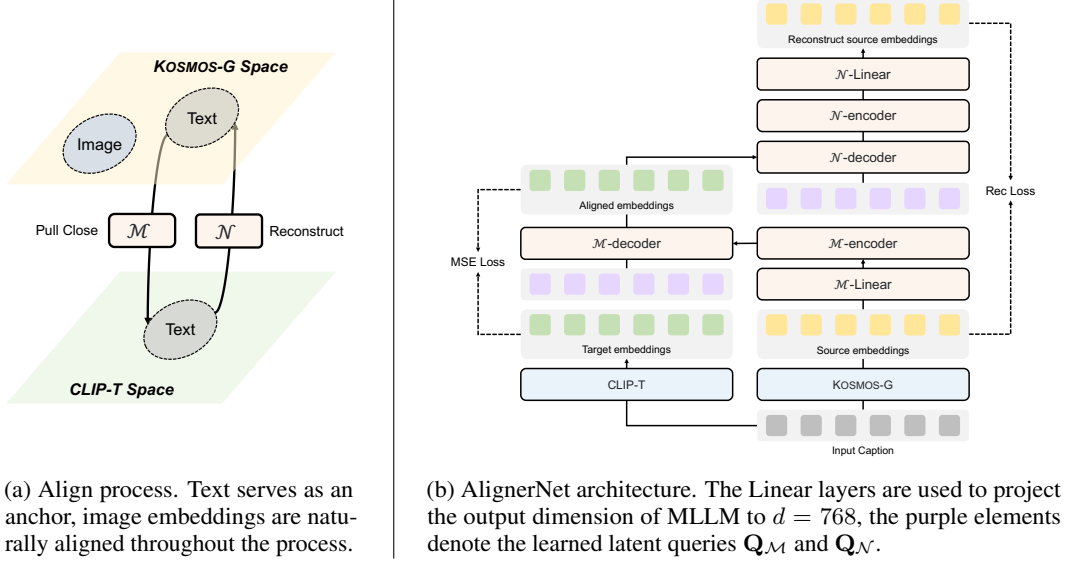


Figure 3: Overview of alignment.

The forward process $q(\mathbf{z}_t | \mathbf{z}_{t-1})$ at each time step t can be expressed as follows:

$$q(\mathbf{z}_t | \mathbf{z}_{t-1}) = \mathcal{N}(\mathbf{z}_t; \sqrt{1 - \beta_t} \mathbf{z}_{t-1}, \beta_t \mathbf{I})$$

$$q(\mathbf{z}_{1:T} | \mathbf{z}_0) = \prod_{t=1}^T q(\mathbf{z}_t | \mathbf{z}_{t-1}) \quad (1)$$

in which $\beta_t \in (0, 1)$ denotes the step size. Note $\beta_{t-1} < \beta_t$.

Diffusion models learn a U-Net [RFB15] denoted as ϵ_θ to reverse the forward diffusion process, constructing desired data samples from the noise. Let $\alpha_t = 1 - \beta_t$ and $\bar{\alpha}_t = \prod_{i=1}^t \alpha_i$. We can reparameterize the denoising process $p(\mathbf{z}_{t-1} | \mathbf{z}_t)$ also as a Gaussian distribution. This distribution can be estimated by ϵ_θ and takes the following form:

$$p_\theta(\mathbf{z}_{t-1} | \mathbf{z}_t) = \mathcal{N}(\mathbf{z}_{t-1}; \boldsymbol{\mu}_\theta(\mathbf{z}_t, t), \boldsymbol{\Sigma}_\theta(\mathbf{z}_t, t))$$

$$\text{with } \boldsymbol{\mu}_\theta(\mathbf{z}_t, t) = \frac{1}{\sqrt{\alpha_t}} \left(\mathbf{z}_t - \frac{\beta_t}{\sqrt{1 - \bar{\alpha}_t}} \epsilon_\theta(\mathbf{z}_t, t) \right) \quad (2)$$

The learning objective of diffusion models is to approximate the mean $\boldsymbol{\mu}_\theta(\mathbf{z}_t, t)$ in the reverse diffusion process. To achieve this, we can utilize the variational lower bound (ELBO) [KW14] to minimize the negative log-likelihood of $p_\theta(\mathbf{z}_0)$ [HJA20]. The simplified objective can be expressed as a denoising objective:

$$\mathcal{L}_{diff} = \mathbb{E}_{\mathbf{z}_0, \epsilon \sim \mathcal{N}(0, \mathbf{I}), t} [\|\epsilon - \epsilon_\theta(\mathbf{z}_t, t)\|^2] \quad (3)$$

During inference, [HS22] proposes to use classifier-free guidance to obtain more relevant generation results.

$$\hat{\epsilon} = w \cdot \epsilon_\theta(\mathbf{z}_t, \varphi, t) - (w - 1) \cdot \epsilon_\theta(\mathbf{z}_t, t) \quad (4)$$

where w is guidance scale, φ denotes the condition.

Align Output Space with Diffusion Models Upon replacing the previous CLIP text encoder with KOSMOS-G, the main focus is to address the misalignment issue between the KOSMOS-G and the image decoder. We discovered that simply fine-tuning KOSMOS-G using the gradient passed from the image decoder results in both trivial alignment and compromised image quality.

Inspired by [QYX⁺23], we propose the AlignerNet consisting of an encoder \mathcal{M} and a decoder \mathcal{N} to learn the alignment between the KOSMOS-G source space \mathbf{S} and CLIP text encoder target space \mathbf{T} . Given a single text-only caption \mathbf{C} , KOSMOS-G source encoder and CLIP text target encoder encode the caption into embeddings denoted as $\mathbf{s} \in \mathbb{R}^{l_s \times d_s}$ and $\mathbf{t} \in \mathbb{R}^{l_t \times d_t}$, respectively. Here, l and d indicate the length of features and embedding dimensions.

As shown in Figure 3a, we employ the encoder \mathcal{M} to minimize the distance between the text source embedding and the target embedding, aiming for a close approximation $\mathcal{M}(\mathbf{s}) \approx \mathbf{t}$ through:

$$\mathcal{L}_{mse} = \mathbb{E}_{\mathbf{s} \sim \mathbf{S}, \mathbf{t} \sim \mathbf{T}} [\|\mathbf{t} - \mathcal{M}(\mathbf{s})\|_2^2] \quad (5)$$

To mitigate the reduction in feature discrimination, we also employ a decoder \mathcal{N} to reconstruct the source embedding $\mathcal{N}(\mathcal{M}(\mathbf{s})) \approx \mathbf{s}$ through:

$$\mathcal{L}_{rec} = \mathbb{E}_{\mathbf{s} \sim \mathbf{S}} [\|\mathbf{s} - \mathcal{N}(\mathcal{M}(\mathbf{s}))\|_2^2] \quad (6)$$

Different from [QYX⁺23], KOSMOS-G is a vision-language multimodal encoder. The language modality serves as an anchor throughout the process, aligning the entire KOSMOS-G space with the image decoder input space, thus also achieving semantic alignment for the image embeddings.

To efficiently process lengthy sequences consisting of multiple images and minimize memory usage, KOSMOS-G encodes the interleaved vision-language input sequence into variable-length embeddings. However, the use of variable length embeddings makes the MLP-based GlueNet [QYX⁺23] unsuitable for learning alignment. To address this, we employ a Transformer-based architecture in AlignerNet, enabling it to effectively align the source and target spaces with mismatched sequence lengths and embedding dimensions.

As shown in Figure 3b, both \mathcal{M} and \mathcal{N} share a similar architecture design, consisting of a Transformer encoder and a Transformer decoder. The Transformer encoder and decoder in both models comprise 12 layers, with an input dimension $d = 768$ and a hidden dimension of 3072. This configuration results in approximately 225M parameters in total. In the cross attention module of Transformer decoder, we use variable length learned latent queries $\mathbf{Q}_{\mathcal{M}} \in \mathbb{R}^{l_t \times d}$ in \mathcal{M} and $\mathbf{Q}_{\mathcal{N}} \in \mathbb{R}^{l_s \times d}$ in \mathcal{N} to match sequence length. Note that as discussed in Section 4.3, we can still align MLLM with Kosmos-G through directly using diffusion loss in Equation 3 with the help of AlignerNet. While it is more costly and leads to worse performance under the same GPU days.

2.3 Instruction Tuning

After achieving a semantic alignment between KOSMOS-G and the image decoder, our model can successfully generate images following interleaved vision-language guidance. However, the multimodal language modeling and text-only alignment stage only preserve the semantic consistency between the input and output, KOSMOS-G still can not leverage rich features extracted from the image encoder to generate images faithfully reproducing the contents in various contexts.

To pursue our objective of “image as a foreign language in image generation,” we curate interleaved vision-language data and use the diffusion loss in Equation 3 to further fine-tune KOSMOS-G. Specifically, we propose a compositional generation task in which we input captions containing entities, with each of them followed by their corresponding images, like “<s> A cat <image> image embedding of the cat </image> and a dog <image> image embedding of the dog </image> sleeping in the garden <image> image embedding of the garden </image> </s>”. Our model is trained to generate images following the input instruction.

To construct the requisite data, we first caption the image, then extract the entities from the caption, and obtain the segmentation results from the image itself. A detailed introduction of the entire pipeline can be found in Section 3.1. Additionally, we leverage the data constructed by [BHE23] for InstructPix2Pix to improve KOSMOS-G’s image editing capability. This data is structured as: “<s> caption <image> embedding of the original image </image> edit instruction </s>”. We also mix some text-to-image data to preserve the language alignment already achieved.

Our goal is to leverage MLLMs to model image distributions through direct latent space sampling. In this setup, the pre-trained frozen Stable Diffusion U-Net serves as a score metric, distilling the learned data distribution. This strategy is similar to Score Distillation Sampling (SDS) [PBM22]. From the perspective of score distillation, the KL divergence between KOSMOS-G (denoted as ϕ , which encodes inputs into condition \mathcal{C}) and the score function is equivalently minimized for distilling learned probability density in the image decoder:

$$\min_{\phi} \mathcal{L}_{Diff} = \mathbb{E}_{\mathbf{z}_0, t, \mathcal{C}} [D_{KL}(q(\mathbf{z}_{t-1} | \mathbf{z}_t, \mathbf{z}_0) \parallel p_{\theta}(\mathbf{z}_{t-1} | \mathbf{z}_t; \mathcal{C}))] \quad (7)$$

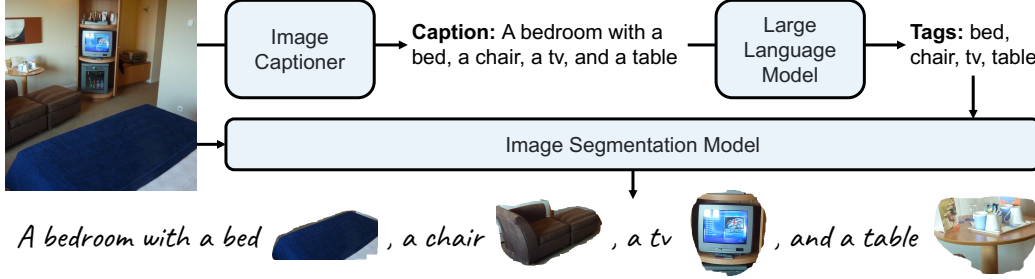


Figure 4: Overview of our data construction pipeline for compositional generation instruction tuning.

This enables KOSMOS-G to leverage rich features from the image encoder to generate an image faithfully reproducing the contents across various contexts. More details about the score distillation instruction tuning can be found in Appendix C.

3 Model Training

3.1 Multimodal Training Data

The multimodal language modeling stage in Section 2.1 using the same setting of KOSMOS-1 [HDW⁺23], where the models are trained on web-scale multimodal corpora, consisted of text corpora, image-caption pairs, and interleaved data of images and texts. For the image decoder aligning stage in Section 2.2, we only use the caption from image-caption pairs. For the instruction tuning stage in Section 2.3, we use constructed data from Open Images V7 dataset [KRA⁺20], the image-caption pairs, as well as the image editing data from InstructPix2Pix [BHE23].

Captions The image-caption pairs are sourced from multiple datasets, including English LAION-2B [SBV⁺22], LAION-400M [SVB⁺21], COYO-700M [BPK⁺22], and Conceptual Captions [SDGS18, CSDS21]. English LAION-2B, LAION-400M, and COYO-700M are collected from Common Crawl web data by extracting images and the corresponding alt-texts. Conceptual captions are also derived from web pages.

Constructed Data We use approximately 9M images from the Open Images V7 dataset [KRA⁺20] to construct our compositional generation instruction tuning data. As illustrated in Figure 4, we begin by generating captions with BLIP-2-OPT-6.7b [LLSH23]. Subsequently, we employ an LLM MPT-7B-Instruct [T⁺23] to extract entities from the captions. The original image, along with the text of each entity, is then input into the text-prompted segmentation model CLIPSeg [LE22] to derive the corresponding image of each entity.

3.2 Training Setup

Our implementation is based on the TorchScale [MWH⁺22] library, which is designed for large-scale model training. Following KOSMOS-1 [HDW⁺23], we also use MAGNETO [WMH⁺22], a Transformer variant, as the backbone architecture of our MLLM and AlignerNet. The whole training process took around four days with 256 NVIDIA V100 GPUs, i.e., one day for image decoder aligning, and three days for instruction tuning. In the instruction tuning stage, we use a blend of constructed data, InstructPix2Pix data, and caption data in a ratio of 2:2:1. For constructed data, to enhance input robustness, we randomly drop the texts of entities with a probability of 0.5 and also maintain the background of the segmented entities with a 0.5 probability. Other training configurations can be found in Appendix A.

4 Evaluation

4.1 Main Qualitative Results

As shown in Figure 5, KOSMOS-G delivers impressive zero-shot generation results across diverse settings, yielding meaningful and coherent outputs even for highly customized subjects. The visual

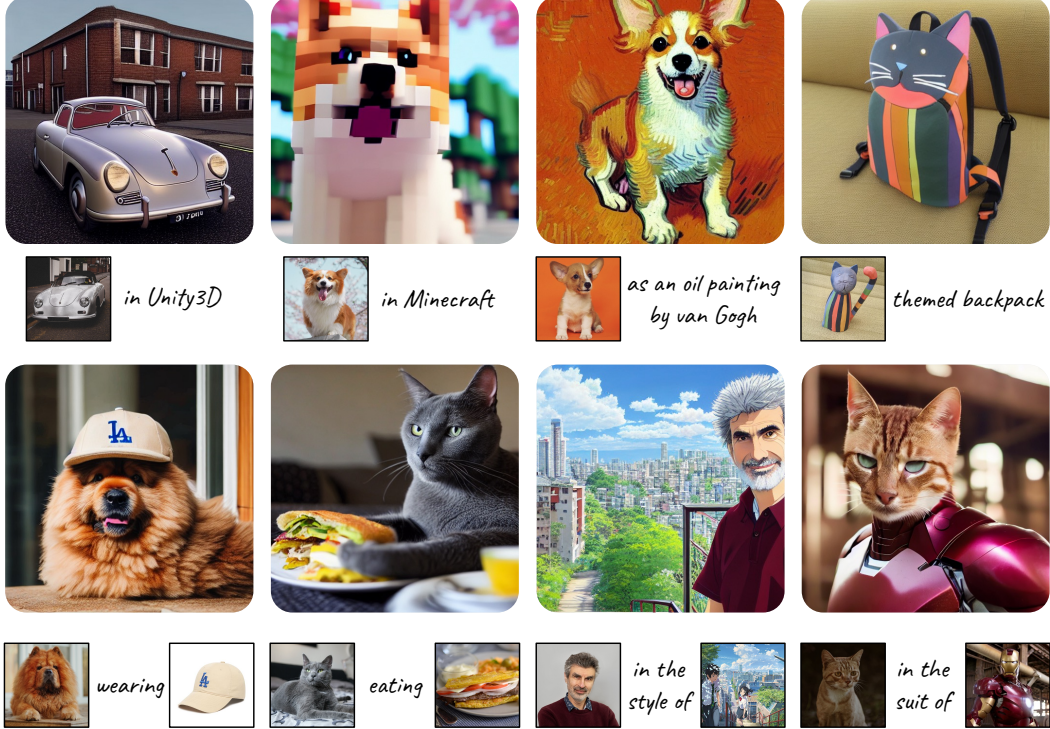


Figure 5: Zero-shot image generation examples with multimodal prompts.

Methods	DINO \uparrow	CLIP-I \uparrow	CLIP-T \uparrow	Methods	FID \downarrow
Real Images (Oracle)	0.774	0.885	-	<i>T2I Models</i>	
<i>Fine-Tuning</i>				GLIDE [NDR ⁺ 22]	12.24
Textual Inversion [GAA ⁺ 22]	0.569	0.780	0.255	Make-A-Scene [GPA ⁺ 22]	11.84
DreamBooth [RLJ ⁺ 22]	0.668	0.803	0.305	DALL-E 2 [RDN ⁺ 22]	10.39
BLIP-Diffusion [LLH23]	0.670	0.805	0.302	SD v1.5 [†] [RBL ⁺ 22]	9.34
<i>Test Time Tuning Free</i>				Imagen-3.4B [SCS ⁺ 22]	7.27
				<i>CLIP-Aligned VL2I Models</i>	
Re-Imagen* [CHSC22]	0.600	0.740	0.270	GILL-8B [KFS23]	12.20
SuTI [CHL ⁺ 23]	0.741	0.819	0.304	Emu-14B [SYC ⁺ 23]	11.66
BLIP-Diffusion* [LLH23]	0.594	0.779	0.300	KOSMOS-G-1.9B	10.99
KOSMOS-G* (single image input)	0.694	0.847	0.287		

Table 1: **Left:** Quantitative comparisons on DreamBench. * denotes zero-shot methods. **Right:** Zero-shot FID comparisons on MS-COCO. [†] indicates results evaluated by us under same settings and seed with KOSMOS-G.

samples showcase generative capabilities in re-contextualization, stylization, modification, and accessory incorporation. Notably, multi-entity subject-driven generation is very challenging even for fine-tuning methods like DreamBooth [RLJ⁺22, AAF⁺23]. While owing from the novel compositional generation instruction tuning, KOSMOS-G is the first model that is capable of achieving this in a zero-shot setting.

4.2 Quantitative Results

We do quantitative evaluations of KOSMOS-G on DreamBench [RLJ⁺22] for single-entity subject-driven generation and MS-COCO [LMB⁺14] for text-to-image generation.

The DreamBench dataset contains 30 subjects and features 25 prompt templates, resulting in 750 unique prompts covering skills like re-contextualization, modification, accessorization, etc. We follow prior work to generate 4 images for each prompt to form the 3000 images for a comprehensive evaluation. We follow DreamBooth to adopt DINO, CLIP-I to evaluate the subject fidelity,



Figure 6: Comparisons with cases presented in the second row of Figure 1.



Figure 7: Various Applications of KOSMOS-G in Conjunction with U-Net Techniques. In Figure 7b, the left image is generated using standard U-Net, the right one is produced with LoRA-tuned U-Net.

and CLIP-T to evaluate the text fidelity. We use a classifier-free guidance scale of 7.5 and 100 DPM-Solver [LZB⁺22] inference steps for sampling. As shown in Table 1, zero-shot KOSMOS-G outperforms Textual Inversion and Re-Imagen and exhibits marginally better performance than DreamBooth and BLIP-Diffusion with only a single image input. Furthermore, Our results are also comparable with SuTI, without requiring expensive apprenticeship learning supervision. KOSMOS-G accepts only a single image as input, we select a clear image from the 4-7 provided images for each subject to avoid occlusion. We slightly modify the prompt template to ensure better alignment with the instruction tuning data. The images and prompt used can be found in Appendix B.

For the text-to-image generation, We generate images using 30,000 randomly sampled captions from the MS-COCO (2014) validation set. We use a classifier-free guidance scale of 3.0 and 250 DDIM [SME21] inference steps for sampling. As shown in Table 1, KOSMOS-G surpasses other CLIP-aligned VL2I models, delivering the optimal alignment results.

4.3 Ablation Studies

We conduct ablation studies to find out the importance of the image decoder aligning and instruction tuning. Table 2 demonstrates that direct end-to-end fine-tuning or using decoder-only architecture fail to generate meaningful images. Incorporating AlignerNet and CLIP supervision, however, results in outcomes close to the original SD v1.5. End-to-end training is also feasible with AlignerNet, but it is more costly due to the additional computational of the U-Net. This leads to worse performance within the same GPU days. We also compared the generation results from KOSMOS-G before instruction tuning and the standard SD v1.5 against our final model. As illustrated in Figure 6, without instruction tuning, KOSMOS-G can only generate contents semantically aligned with the vision-language input. SD baseline also remains at the semantic level and fails to faithfully reproduce the entities in the generated images.

Methods	FID↓
SD v1.5 [RBL ⁺ 22]	9.34
E2E w/o AlignerNet	Failed
E2E w/ AlignerNet	11.30
12-Layers Decoder	Failed
12-Layers AlignerNet	9.89
24-Layers AlignerNet	9.55

Table 2: Ablation study results for image decoder aligning on MS-COCO.

4.4 Applications

As highlighted in Section 2.3, KOSMOS-G can seamlessly replace CLIP in any image generation system. This remarkable property unlocks a myriad of brand-new applications that have never

been possible before. We demonstrate its integration with ControlNet [ZA23] and LoRA variants [HSW⁺22] in Figure 7. KOSMOS-G works perfectly with these techniques. Building on the CLIP space, we believe our model will push forward the transition from text-conditioned generation toward vision-language generation, paving the way for numerous novel applications.

5 Conclusion

We propose KOSMOS-G, a model capable of high-fidelity zero-shot subject-driven generation from interleaved multi-image and text input. Our approach hinges on a unique "align before instruct" pre-training strategy. KOSMOS-G demonstrates competitive single-entity subject-driven image generation and text-to-image capability, it also stands as the first model to extend zero-shot subject-driven image generation to multi-entity scenarios. Furthermore, KOSMOS-G allows seamless replacement of CLIP, unlocking various new applications in conjunction with other U-Net techniques such as ControlNet and LoRA. In general, we present KOSMOS-G as a preliminary effort aimed at achieving the objective of "image as a foreign language in image generation."

Ethics Statement

KOSMOS-G is purely a research project. Currently, we have no plans to incorporate KOSMOS-G into a product or expand access to the public. We will also put Microsoft AI principles into practice when further developing the models.

In our research paper, we account for the ethical concerns associated with text-to-image research. To mitigate issues associated with training data, we have implemented a rigorous filtering process to purge our training data of inappropriate content, such as explicit imagery and offensive language, to minimize the likelihood of generating inappropriate content.

References

- [AAF⁺23] Omri Avrahami, Kfir Aberman, Ohad Fried, Daniel Cohen-Or, and Dani Lischinski. Break-a-scene: Extracting multiple concepts from a single image. *arXiv preprint arXiv:2305.16311*, 2023.
- [ADL⁺22] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, Roman Ring, Eliza Rutherford, Serkan Cabi, Tengda Han, Zhitao Gong, Sina Samangooei, Marianne Monteiro, Jacob Menick, Sebastian Borgeaud, Andrew Brock, Aida Nematzadeh, Sahand Sharifzadeh, Mikolaj Binkowski, Ricardo Barreira, Oriol Vinyals, Andrew Zisserman, and Karen Simonyan. Flamingo: a visual language model for few-shot learning. In *Advances in Neural Information Processing Systems*, 2022.
- [AHR⁺22] Armen Aghajanyan, Bernie Huang, Candace Ross, Vladimir Karpukhin, Hu Xu, Naman Goyal, Dmytro Okhonko, Mandar Joshi, Gargi Ghosh, Mike Lewis, and Luke Zettlemoyer. CM3: A causal masked multimodal model of the Internet. *ArXiv preprint*, abs/2201.07520, 2022.
- [BHE23] Tim Brooks, Aleksander Holynski, and Alexei A Efros. Instructpix2pix: Learning to follow image editing instructions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18392–18402, 2023.
- [BPK⁺22] Minwoo Byeon, Beomhee Park, Haecheon Kim, Sungjun Lee, Woonhyuk Baek, and Saehoon Kim. Coyo-700m: Image-text pair dataset, 2022.
- [CHL⁺23] Wenhui Chen, Hexiang Hu, Yandong Li, Nataniel Rui, Xuhui Jia, Ming-Wei Chang, and William W Cohen. Subject-driven text-to-image generation via apprenticeship learning. *ArXiv preprint*, abs/2304.00186, 2023.
- [CHSC22] Wenhui Chen, Hexiang Hu, Chitwan Saharia, and William W Cohen. Re-imagen: Retrieval-augmented text-to-image generator. *ArXiv preprint*, abs/2209.14491, 2022.

- [CSDS21] Soravit Changpinyo, Piyush Sharma, Nan Ding, and Radu Soricut. Conceptual 12m: Pushing web-scale image-text pre-training to recognize long-tail visual concepts. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*, pages 3558–3568. Computer Vision Foundation / IEEE, 2021.
- [DBK⁺21] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021.
- [DHP⁺23] Runpei Dong, Chunrui Han, Yuang Peng, Zekun Qi, Zheng Ge, Jinrong Yang, Liang Zhao, Jianjian Sun, Hongyu Zhou, Haoran Wei, et al. Dreamllm: Synergistic multimodal comprehension and creation. *ArXiv preprint*, abs/2309.11499, 2023.
- [GAA⁺22] Rinon Gal, Yuval Alaluf, Yuval Atzmon, Or Patashnik, Amit H Bermano, Gal Chechik, and Daniel Cohen-Or. An image is worth one word: Personalizing text-to-image generation using textual inversion. *ArXiv preprint*, abs/2208.01618, 2022.
- [GPA⁺22] Oran Gafni, Adam Polyak, Oron Ashual, Shelly Sheynin, Devi Parikh, and Yaniv Taigman. Make-a-scene: Scene-based text-to-image generation with human priors. *ArXiv preprint*, abs/2203.13131, 2022.
- [HDW⁺23] Shaohan Huang, Li Dong, Wenhui Wang, Yaru Hao, Saksham Singhal, Shuming Ma, Tengchao Lv, Lei Cui, Owais Khan Mohammed, Qiang Liu, et al. Language is not all you need: Aligning perception with language models. *ArXiv preprint*, abs/2302.14045, 2023.
- [HHZW23] Shaozhe Hao, Kai Han, Shihao Zhao, and Kwan-Yee K Wong. Vico: Detail-preserving visual condition for personalized text-to-image generation. *arXiv preprint arXiv:2306.00971*, 2023.
- [HJA20] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- [HS22] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *ArXiv preprint*, abs/2207.12598, 2022.
- [HSD⁺22] Yaru Hao, Haoyu Song, Li Dong, Shaohan Huang, Zewen Chi, Wenhui Wang, Shuming Ma, and Furu Wei. Language models are general-purpose interfaces. *ArXiv preprint*, abs/2206.06336, 2022.
- [HSW⁺22] Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net, 2022.
- [KFS23] Jing Yu Koh, Daniel Fried, and Ruslan Salakhutdinov. Generating images with multi-modal language models. *ArXiv preprint*, abs/2305.17216, 2023.
- [KR18] Taku Kudo and John Richardson. SentencePiece: A simple and language independent subword tokenizer and detokenizer for neural text processing. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 66–71, Brussels, Belgium, 2018. Association for Computational Linguistics.
- [KRA⁺20] Alina Kuznetsova, Hassan Rom, Neil Alldrin, Jasper Uijlings, Ivan Krasin, Jordi Pont-Tuset, Shahab Kamali, Stefan Popov, Matteo Mallocci, Alexander Kolesnikov, Tom Duerig, and Vittorio Ferrari. The open images dataset v4: Unified image classification, object detection, and visual relationship detection at scale. *IJCV*, 2020.

- [KW14] Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. In Yoshua Bengio and Yann LeCun, editors, *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*, 2014.
- [KZZ⁺23] Nupur Kumari, Bingliang Zhang, Richard Zhang, Eli Shechtman, and Jun-Yan Zhu. Multi-concept customization of text-to-image diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1931–1941, 2023.
- [LE22] Timo Lüddecke and Alexander Ecker. Image segmentation using text and image prompts. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7086–7096, 2022.
- [LLH23] Dongxu Li, Junnan Li, and Steven CH Hoi. Blip-diffusion: Pre-trained subject representation for controllable text-to-image generation and editing. *ArXiv preprint*, abs/2305.14720, 2023.
- [LLSH23] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. BLIP-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *ArXiv preprint*, abs/2301.12597, 2023.
- [LMB⁺14] 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 *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part V 13*, pages 740–755. Springer, 2014.
- [LOG⁺19] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. RoBERTa: A robustly optimized bert pretraining approach. *ArXiv preprint*, abs/1907.11692, 2019.
- [Luo22] Calvin Luo. Understanding diffusion models: A unified perspective. *arXiv preprint arXiv:2208.11970*, 2022.
- [LZB⁺22] Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. *Advances in Neural Information Processing Systems*, 35:5775–5787, 2022.
- [MWH⁺22] Shuming Ma, Hongyu Wang, Shaohan Huang, Wenhui Wang, Zewen Chi, Li Dong, Alon Benhaim, Barun Patra, Vishrav Chaudhary, Xia Song, and Furu Wei. TorchScale: Transformers at scale. *ArXiv preprint*, abs/2211.13184, 2022.
- [NDR⁺22] Alexander Quinn Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. GLIDE: towards photorealistic image generation and editing with text-guided diffusion models. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvári, Gang Niu, and Sivan Sabato, editors, *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA*, volume 162 of *Proceedings of Machine Learning Research*, pages 16784–16804. PMLR, 2022.
- [PJBM22] Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. *ArXiv preprint*, abs/2209.14988, 2022.
- [QYX⁺23] Can Qin, Ning Yu, Chen Xing, Shu Zhang, Zeyuan Chen, Stefano Ermon, Yun Fu, Caiming Xiong, and Ran Xu. Gluegen: Plug and play multi-modal encoders for x-to-image generation. *ArXiv preprint*, abs/2303.10056, 2023.
- [RBL⁺22] 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.

- [RDN⁺22] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *ArXiv preprint*, abs/2204.06125, 2022.
- [RFB15] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- [RKH⁺21] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In Marina Meila and Tong Zhang, editors, *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*, pages 8748–8763. PMLR, 2021.
- [RLJ⁺22] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. *ArXiv preprint*, abs/2208.12242, 2022.
- [SBV⁺22] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, et al. Laion-5b: An open large-scale dataset for training next generation image-text models. *ArXiv preprint*, abs/2210.08402, 2022.
- [SCS⁺22] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S Sara Mahdavi, Rapha Gontijo Lopes, et al. Photorealistic text-to-image diffusion models with deep language understanding. *ArXiv preprint*, abs/2205.11487, 2022.
- [SDGS18] Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2556–2565, Melbourne, Australia, 2018. Association for Computational Linguistics.
- [SHZ⁺23] James Seale Smith, Yen-Chang Hsu, Lingyu Zhang, Ting Hua, Zsolt Kira, Yilin Shen, and Hongxia Jin. Continual diffusion: Continual customization of text-to-image diffusion with c-lora. *arXiv preprint arXiv:2304.06027*, 2023.
- [SME21] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021.
- [SVB⁺21] Christoph Schuhmann, Richard Vencu, Romain Beaumont, Robert Kaczmarczyk, Clayton Mullis, Aarush Katta, Theo Coombes, Jenia Jitsev, and Aran Komatsuzaki. Laion-400m: Open dataset of clip-filtered 400 million image-text pairs. *ArXiv preprint*, abs/2111.02114, 2021.
- [SWMG15] Jascha Sohl-Dickstein, Eric A. Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In Francis R. Bach and David M. Blei, editors, *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6-11 July 2015*, volume 37 of *JMLR Workshop and Conference Proceedings*, pages 2256–2265. JMLR.org, 2015.
- [SYC⁺23] Quan Sun, Qiying Yu, Yufeng Cui, Fan Zhang, Xiaosong Zhang, Yueze Wang, Hongcheng Gao, Jingjing Liu, Tiejun Huang, and Xinlong Wang. Generative pre-training in multimodality. *ArXiv preprint*, abs/2307.05222, 2023.
- [T⁺23] MN Team et al. Introducing MPT-7B: A new standard for open-source, commercially usable llms, 2023.
- [TGCA23] Yoad Tewel, Rinon Gal, Gal Chechik, and Yuval Atzmon. Key-locked rank one editing for text-to-image personalization. *arXiv preprint arXiv:2305.01644*, 2023.

- [WMH⁺22] Hongyu Wang, Shuming Ma, Shaohan Huang, Li Dong, Wenhui Wang, Zhiliang Peng, Yu Wu, Payal Bajaj, Saksham Singhal, Alon Benhaim, Barun Patra, Zhun Liu, Vishrav Chaudhary, Xia Song, and Furu Wei. Foundation transformers. *ArXiv preprint*, abs/2210.06423, 2022.
- [WZJ⁺23] Yuxiang Wei, Yabo Zhang, Zhilong Ji, Jinfeng Bai, Lei Zhang, and Wangmeng Zuo. Elite: Encoding visual concepts into textual embeddings for customized text-to-image generation. *arXiv preprint arXiv:2302.13848*, 2023.
- [XYF⁺23] Guangxuan Xiao, Tianwei Yin, William T Freeman, Frédo Durand, and Song Han. Fastcomposer: Tuning-free multi-subject image generation with localized attention. *arXiv preprint arXiv:2305.10431*, 2023.
- [ZA23] Lvmin Zhang and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models, 2023.

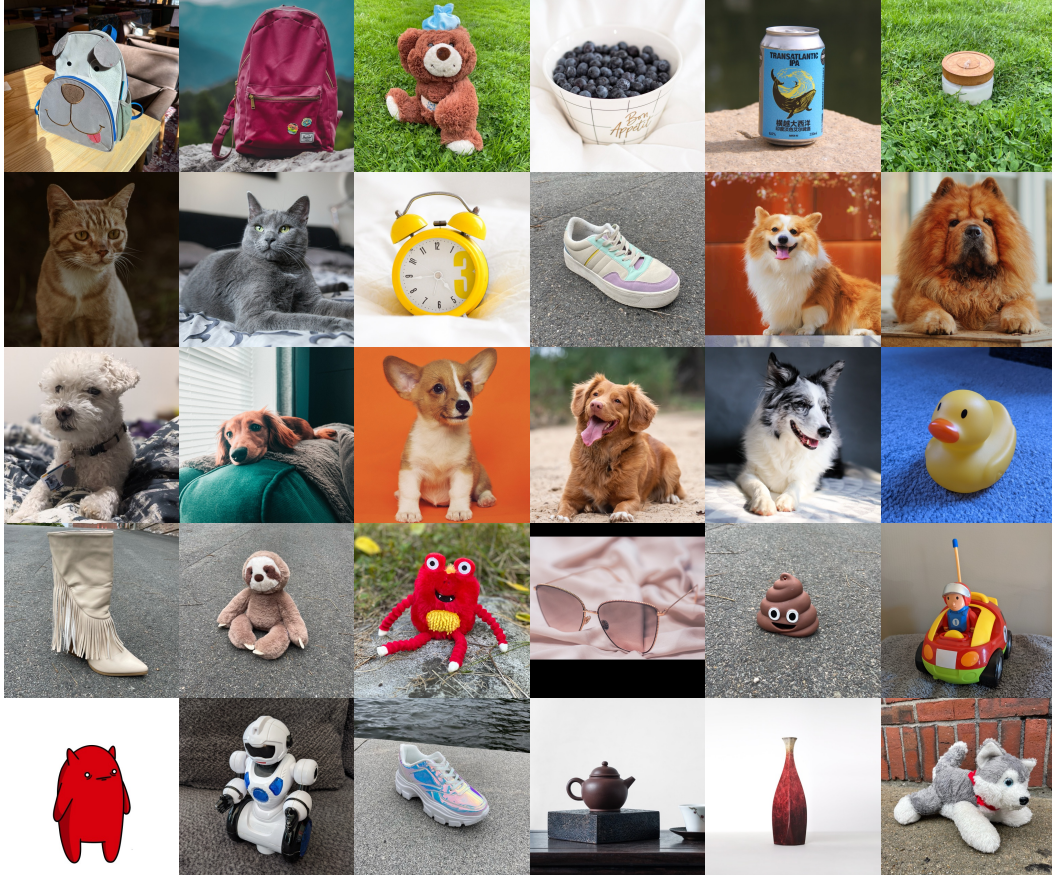


Figure 8: Selected images from DreamBench.

A Detail Training Setup

Multimodal Language Modeling We use a batch size of 1.2 million tokens which is broken down as follows: 0.5 million tokens sourced from text corpora, 0.5 million tokens derived from image-caption pairs, and 0.2 million tokens from interleaved data sets. The MLLM is trained for 300,000 steps, corresponding to about 360 billion tokens in total. We adopt the AdamW optimizer with $\beta = (0.9, 0.98)$. Furthermore, we configure the weight decay at 0.01 and the dropout rate at 0.1. The learning rate is set to escalate to $2e-4$ during the initial 375 warm-up steps and decay linearly to 0 for the rest of the training steps. For optimization stability, we initiate using Magneto. We use SentencePiece [KR18] to tokenize the text. We preprocess the data in the “full-sentence” format [LOG⁺19], where each input sequence is populated with complete sentences consecutively sampled from one or multiple documents.

Image Decoder Aligning The AlignerNet undergoes training using a batch size of 3,584 sentences for 300,000 steps, with a maximum learning rate of $1e-3$. This equates to approximately 1 billion sentences overall. The remaining configurations remain consistent with the previous stage.

Instruction Tuning The MLLM and AlignerNet are jointly trained with a batch size of 1,024 images, totaling approximately 200 million images over 200,000 steps. The learning rate peaks at $1e-3$. The rest settings are the same as in the previous stage.

B Images and Prompts for DreamBench Evaluation

The input images selected for each entity for DreamBench evaluation are displayed in Figure 8. As depicted in Table 3, We also slightly modified the original prompt to make it align better with the

training data. For the remaining prompt, we simply remove the prefix ‘‘a’’. We observe the prefix will slightly affect the performance of image editing or customized generation. This might be attributed to the high frequency of captions starting with ‘‘a photo of’’ produced by BLIP-2 in our constructed training data. Given that the compositional instruction tuning data does not contain too much editing data when a prompt starts with a prefix like ‘‘a’’, the model often refrains from altering the appearance of the input image. This can be further refined by shifting the data paradigm or by fine-tuning the alignment.

Original Prompt	Modified Prompt
a red {}	{}, red
a purple {}	{}, purple
a shiny {}	{}, shiny
a wet {}	{}, wet
a cube shaped {}	{}, cube shaped

Table 3: Prompt Modification

C Additional Details about Score Distillation Instruction Tuning

During the instruction tuning stage, we still utilize the diffusion loss for model training. However, this loss can also be regarded as score distillation, and it is effectively equivalent to optimizing the diffusion loss as the Score Distillation Sampling (SDS) [PJB22] loss.

Our objective is to distill the learned score function from the Stable Diffusion U-Net into KOSMOS-G. This process enables KOSMOS-G to encode image features into embeddings that the Stable Diffusion model can understand for subject-driven generation. Essentially, this approach is like pre-training a generalized textual inversion [GAA⁺22] model, with all conditions learnable by model-seeking. Consider the KOSMOS-G model, denoted as ϕ , which takes an input \mathbf{x} and produces an output $\mathcal{C} = \phi(\mathbf{x})$. Alongside this, we have the Diffusion U-Net, represented as θ . In our process, we optimize the KOSMOS-G model ϕ using the diffusion loss, while keeping the parameters of the Diffusion U-Net θ frozen. The diffusion loss is expressed as follows:

$$\mathcal{L}_{diff}(\phi) = \mathbb{E}_{\mathbf{z}_0, \epsilon \sim \mathcal{N}(0,1), t} \left[w(t) \|\epsilon_\theta(\mathbf{z}_t; \mathcal{C}, t) - \epsilon\|^2 \right] \quad (8)$$

Consider the gradient of \mathcal{L}_{diff} :

$$\nabla_\phi \mathcal{L}_{diff}(\phi) = \mathbb{E}_{\mathbf{z}_0, \epsilon \sim \mathcal{N}(0,1), t} \left[\underbrace{w(t) (\epsilon_\theta(\mathbf{z}_t; \mathcal{C}, t) - \epsilon)}_{\text{Noise Residual}} \underbrace{\frac{\partial \epsilon_\phi(\mathbf{z}_t; \mathcal{C}, t)}{\partial \mathcal{C}}}_{\text{U-Net Jacobian}} \underbrace{\frac{\partial \mathcal{C}}{\partial \phi}}_{\text{KOSMOS-G Jacobian}} \right] \quad (9)$$

Following the approach in Dreamfusion [PJB22], we can simplify this equation by omitting certain terms, leading to the SDS loss for KOSMOS-G:

$$\nabla_\phi \mathcal{L}_{SDS}(\phi) = \mathbb{E}_{\mathbf{z}_0, \epsilon \sim \mathcal{N}(0,1), t} \left[w(t) (\epsilon_\theta(\mathbf{z}_t; \mathcal{C}, t) - \epsilon) \frac{\partial \mathcal{C}}{\partial \phi} \right] \quad (10)$$

As established in Dreamfusion [PJB22], optimizing the SDS loss is effectively equivalent to optimizing the diffusion loss when the U-Net θ is frozen. From the perspective of score distillation, when using diffusion loss, the KL divergence defined by conditions and the pre-learned score function is equivalently minimized for distilling learned probability density in conditional image synthesis [DHP⁺23, Luo22]:

$$\min_\phi \mathcal{L}_{Diff}(\phi) = \mathbb{E}_{\mathbf{z}_0, t, \mathcal{C}} \left[D_{\text{KL}}(q(\mathbf{z}_{t-1} | \mathbf{z}_t, \mathbf{z}_0) \parallel p_\theta(\mathbf{z}_{t-1} | \mathbf{z}_t; \mathcal{C})) \right] \quad (11)$$

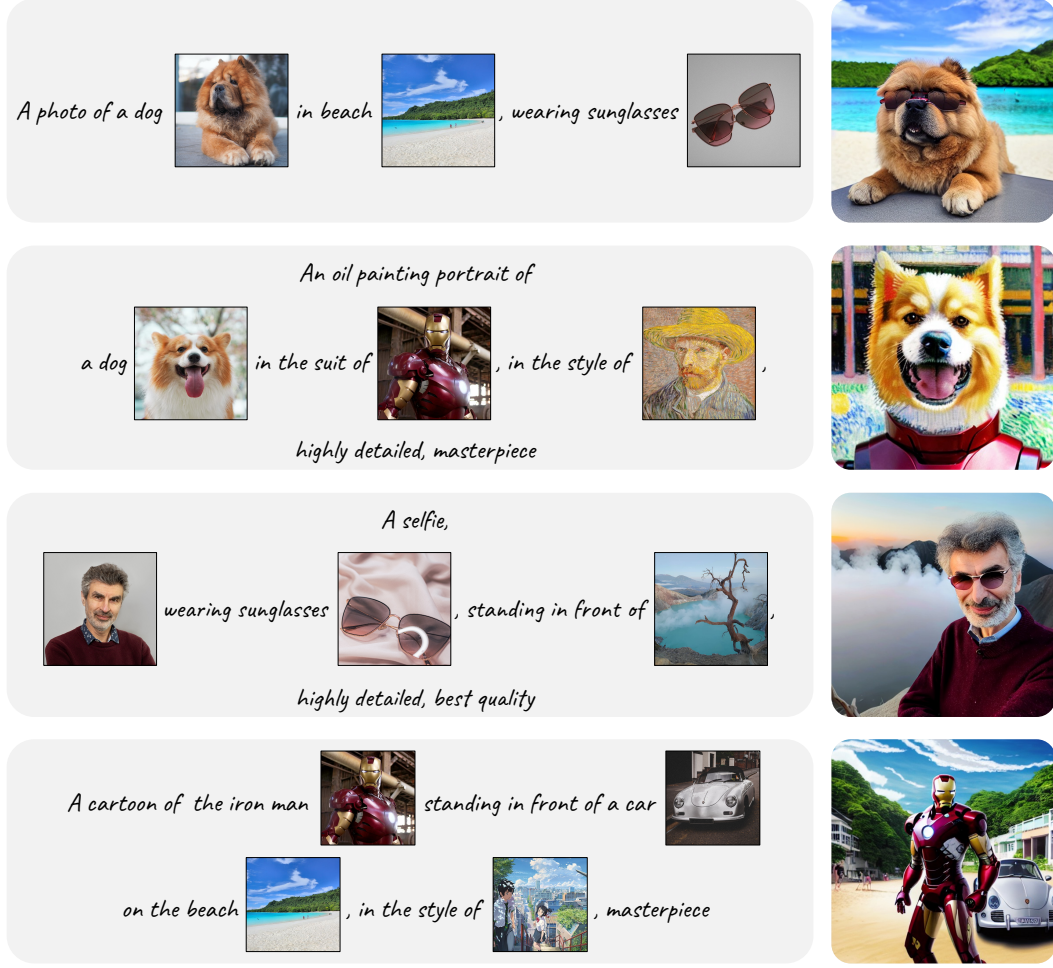


Figure 9: Additional examples under challenging multi-image and text interleaving scenario. These cases show that with KOSMOS-G, users can prompt Stable Diffusion [RBL⁺22] by approaching all image inputs as a “foreign language”.

D Additional Examples

In Figure 9, we present cases with more diverse and complex multi-image (3 to 4) and text interleaving scenarios. These examples demonstrate KOSMOS-G’s robust performance in handling these challenging cases.