

# MAGVLT: Masked Generative Vision-and-Language Transformer

Sungwoong Kim<sup>1,\*</sup>, Daejin Jo<sup>2,\*</sup>, Donghoon Lee<sup>2,\*</sup>, Jongmin Kim<sup>2,\*</sup>  
<sup>1</sup>Department of Artificial Intelligence, Korea University, Seoul, South Korea  
<sup>2</sup>Kakao Brain, Seongnam, South Korea

swkim01@korea.ac.kr, {daejin.jo, dhlee, jmkim}@kakaobrain.com

## Abstract

While generative modeling on multimodal image-text data has been actively developed with large-scale paired datasets, there have been limited attempts to generate both image and text data by a single model rather than a generation of one fixed modality conditioned on the other modality. In this paper, we explore a unified generative vision-and-language (VL) model that can produce both images and text sequences. Especially, we propose a generative VL transformer based on the non-autoregressive mask prediction, named **MAGVLT**, and compare it with an autoregressive generative VL transformer (ARGVLT). In comparison to ARVLT, the proposed MAGVLT enables bidirectional context encoding, fast decoding by parallel token predictions in an iterative refinement, and extended editing capabilities such as image and text infilling. For rigorous training of our MAGVLT with image-text pairs from scratch, we combine the image-to-text, text-to-image, and joint image-and-text mask prediction tasks. Moreover, we devise two additional tasks based on the step-unrolled mask prediction and the selective prediction on the mixture of two image-text pairs. Experimental results on various downstream generation tasks of VL benchmarks show that our MAGVLT outperforms ARVLT by a large margin even with significant inference speedup. Particularly, MAGVLT achieves competitive results on both zero-shot image-to-text and text-to-image generation tasks from MS-COCO by one moderate-sized model (fewer than 500M parameters) even without the use of monomodal data and networks.

## 1. Introduction

Generalizable multimodal modeling has recently made a lot of progress, especially in the field of vision-and-language (VL) modeling [3, 17, 30, 43, 45, 46, 48, 62, 63, 66, 67]. In particular, a massive amount of image-text data [9, 11, 51, 52, 54, 55] allows robust pretraining of large-scale multimodal VL models that can be easily transferred to various downstream tasks including image captioning [2, 38],

text-guided image generation [16, 17, 19, 40, 45, 46, 48, 67], visual question answering [23], and image-text retrieval [30, 38, 42, 43]. In this respect, many multimodal VL pretraining algorithms have been proposed in the literature, and these can be broadly categorized into either discriminative or generative learning algorithms. Discriminative pretraining such as contrastive learning [30, 43] aims to obtain semantic representations effective for discriminative tasks while generative pretraining learns them to reconstruct an input [4, 8, 13, 25, 32, 35, 61, 63, 64]. The recent growth of model capacity and data size has led to more interest in generative pretraining since it can provide more diverse and improved generalization ability both for VL understanding and VL generation tasks.

While generative VL pretraining has been widely exploited, most existing algorithms focus on representation learning for VL understanding tasks [7, 8, 12, 13, 33, 35, 63] or conditional generation tasks where a generation is performed on one fixed modality conditioned on the other modality [3, 13, 14, 16, 17, 19, 25, 31, 40, 45, 46, 48, 61, 64, 66, 67]. A few algorithms have tried to produce data in both modalities from a single VL model [4, 32]. If one universal model can generate both modalities, it would be beneficial in concentrating training efforts on a single model as well as resource-saving under a resource-constrained deployment. Moreover, we can expect task extension as well as synergistic performance improvement between the modalities from this multimodal generation. Therefore, in this work, we develop a unified generative VL model.

There are two prevalent paradigms of the generative modeling for image and text processing: autoregressive (AR) generative modeling [16, 32, 46, 61, 67] and non-AR generative modeling [4, 5, 45, 48]. Many previous algorithms adopt AR modeling due to its excellent generation results and high training scalability, especially with transformer networks. However, AR modeling has restrictions in unidirectional conditioning in that an image needs to be flattened into a 1D sequence by an unnatural ordering. In addition, AR sampling is performed by one-by-one predictions of elements, which incurs very slow generation for a long sequence. Recently, in order to overcome these limitations of AR modeling, non-AR generative modeling based

\*Contributed equally.

†Work done at Kakao Brain. Corresponding author.

on the mask prediction has been proposed for language [22], image [10, 36], and video processing [24, 59]. Masked modeling is usually employed for representation learning to solve understanding tasks in language, vision, and VL domains. However, with an iterative refinement-based generation and a variable mask ratio during training, it has been shown to be used as a promising generative modeling. In this regard, for our generative VL modeling, we propose **Masked Generative VL Transformer (MAGVLT)**. In contrast to AR-based generative VL transformer (ARGVLT), the proposed MAGVLT is able to exploit bidirectional conditioning and fast generation through a small number of refinement steps and parallel token predictions.

In specific, MAGVLT can generate any or both of an image and a text sequence conditioned also on any or both of them. Namely, it can perform any kind of task in a form of image-and-text-to-image-and-text (IT2IT), including image-to-text (I2T) and text-to-image (T2I) tasks. Following the previous masked generative modeling [10, 22], we conduct sampling by iterative denoising based on the masked token prediction and train MAGVLT by the masked token prediction objective with a randomly sampled mask ratio to take into account various denoising steps. Here, to perform robust training of MAGVLT especially with only image-text pairs from scratch, MAGVLT is learned basically by the composition of image-to-text, text-to-image, and joint image-and-text mask prediction objectives. We observe that our cross-modal masking (joint image-and-text mask prediction) during training helps in improving both performances of I2T and T2I tasks over single-modal masking (image-to-text + text-to-image mask predictions). Note that only masked generative modeling used in MAGVLT enables this cross-modal mask prediction during training.

In addition, we propose to use two additional tasks based on the step-unrolled mask prediction and the selective prediction on the mixture of two image-text pairs. The former one is motivated by SUNDAE [50] and is modified to perform the mask prediction on the unrolled prediction, which simulates the masked input samples encountered at the intermediate refinement steps. On the other hand, the latter one learns to reconstruct the masked tokens in accordance with a selected context between two VL contexts that are mixed as a noisy input. This selective prediction improves cross-modal attention for an accurate generation.

Through experiments on various downstream VL generation tasks, we empirically demonstrate that our MAGVLT significantly outperforms ARVLT even with greatly reduced inference time. Especially, to the best of our knowledge, MAGVLT is the first model that obtains strong performances on both zero-shot I2T and zero-shot T2I generation tasks of MS-COCO benchmark [38] by a single moderate-sized model (fewer than 500M parameters) without relying on monomodal data and networks. Previously, as unified generative VL models, L-Verse [32] and UPGen [4] have not showed zero-shot I2T results while OFA [62] has used

monomodal data and also has not showed zero-shot I2T and T2I results. Extensive ablations also validate the contribution of each component for MAGVLT.

To summarize, our main contributions are: (1) a masked generative VL transformer as a unified generative VL model that can produce both images and texts; (2) a robust training on image-text pairs by multiple training tasks that include the cross-modal mask prediction in tandem with the step-unrolled mask prediction and the selective prediction on the mixed context; and (3) an empirical validation of MAGVLT that outperforms the autoregressive model and moreover shows competitive performances on both of zero-shot I2T and T2I generation tasks for the first time without employing extra monomodal data and networks.

## 2. Related Work

### 2.1. Multimodal Vision-and-Language Modeling

There has been a lot of research on multimodal VL training. Especially, in recent years, large-scale image-text datasets have made great progress in VL pretraining with various training objectives. For example, image-text matching [39, 57, 65] and contrastive learning on image-text data [30, 43, 60, 66, 68] has been widely used for discriminative representation learning.

Since BERT [15] has shown impressive performances on many natural language processing tasks, masked language modeling has been widely adopted for VL pretraining. In particular, [13] formulated multiple VL tasks as a text generation task and applied masked token prediction objective. [6, 8, 21, 35, 63] have proposed to use a unified masked data modeling with a shared multimodal transformer while several algorithms have combined a number of objectives including image-text matching, contrastive VL loss, and masked language modeling [7, 14, 33, 71]. However, most of these masked VL pretraining algorithms have been developed for VL understanding tasks.

Meanwhile, AR generative modeling has also recently received lots of interest for VL pretraining due to its powerful generalization ability. While many algorithms have proposed to utilize the generative language modeling paradigm for VL understanding tasks [3, 25, 31, 61, 64], recent large-scale AR transformers trained on a large amount of image-text pairs have shown powerful performances for text-guided image generation [16, 17, 46, 67]. A few generative models based on AR decoding have been shown to produce both images and text sequences [32, 62], however, they have not shown competitive performances on both modalities.

### 2.2. Non-Autoregressive Generative Modeling

Non-AR generative models have been increasingly used to lift certain limitations of AR models such as unidirectional attention and slow decoding. Among them, diffusion-based models have recently shown remarkable performances on the task of text-guided image generation

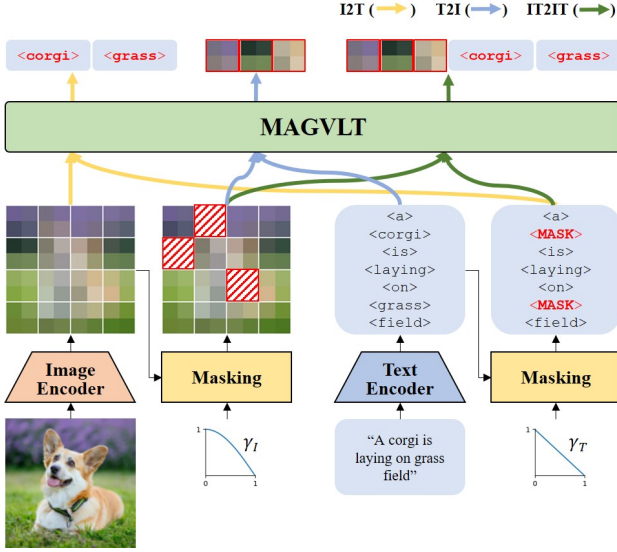


Figure 1. Masked generative VL training via three multimodal masked token prediction tasks (I2T, T2I, IT2IT). Here, we represent VQ-GAN and BPE as image encoder and text encoder. MAGVLT predicts only masked tokens according to each task.

[19, 40, 45, 47, 48]. However, for text generation, diffusion models [5, 29] are still limited in achieving competitive performances compared to AR models.

Similar to diffusion-based generative models, masked generative models also perform iterative refinements for data generation and simulate various denoising steps during training. On top of that, masked generative models often conduct fewer refinement steps leading to faster generations. Therefore, masked generative models have recently been employed for language [22], image [10, 36], and video processing [24, 59]. However, there have been almost no masked generative models that can generate both text and image data. Very recently, a concurrent work [4] has tried to combine representation and generative learning for VL tasks into a single model that is based on masked token prediction. However, their generation performances are very poor on both modalities in contrast to the strong performances of MAGVLT on both modalities, and furthermore, our MAGVLT differs in that we use multiple cross-modal tasks for robust generative training on image-text pairs.

### 3. Masked Generative Vision-and-Language Transformer

#### 3.1. Masked Image-Text Modeling

MAGVLT is based on the previous masked generative modeling for image and language processing [10, 22]. Given an image-text pair  $(I, T)$ , the image input  $I$  is mapped to latent tokens  $X = [x_i]_{i=1}^{N_I}$  by VQ-GAN [18], where  $N_I$  is the number of image tokens (e.g.,  $16 \times 16$ ), and the text sequence  $T$  is also converted to the tokenized sequence  $Y = [y_j]_{j=1}^{N_T}$  by byte pair encoding (BPE) [53],

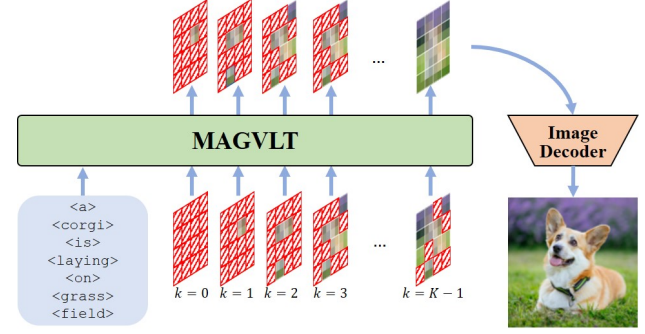


Figure 2. Iterative decoding process of MAGVLT for T2I task.

where  $N_T$  is the number of text tokens (e.g., 64). Then,  $(X, Y)$  is fed into a bidirectional transformer with full attention. In contrast to the causal attention in the AR transformer, this full attention allows to fully utilize the whole context information for decoding and therefore leads to better output predictions. Here, to indicate which modality is processed by the shared transformer, we prepend learnable special tokens representing the modality such as  $\langle \text{BOI} \rangle$  (*begin-of-image*) and  $\langle \text{BOT} \rangle$  (*begin-of-text*) to the input tokens of each modality.

For training MAGVLT on image-text pairs, we first sample a mask ratio  $\gamma(r) \in (0, 1]$  where  $r \in [0, 1]$  indicates the simulated refinement step ratio and is also uniformly sampled considering various steps during generation. Then, we uniformly sample  $\lceil \gamma \cdot N \rceil$  tokens where  $N = N_I + N_T$  and replace them with a special token  $\langle \text{MASK} \rangle$ . Here, we separately apply this masking for each modality such that  $\lceil \gamma_I \cdot N_I \rceil$  image tokens and  $\lceil \gamma_T \cdot N_T \rceil$  text tokens are masked. Let  $M_I = [m_i^I]_{i=1}^{N_I}$  and  $M_T = [m_j^T]_{j=1}^{N_T}$  be the resulting binary image mask and binary text mask, respectively, such that  $x_i$  is replaced with  $\langle \text{MASK} \rangle$  if  $m_i^I = 1$  while  $y_j$  is replaced with  $\langle \text{MASK} \rangle$  if  $m_j^T = 1$ . Note that we set  $\gamma_I(\cdot)$  and  $\gamma_T(\cdot)$  as the cosine function and the linear function, respectively, following [10] and [22].

As shown in Figure 1, MAGVLT is trained basically by the composition of three mask prediction losses:  $\mathcal{L}_{I2T}$  for the I2T task,  $\mathcal{L}_{T2I}$  for the T2I task, and  $\mathcal{L}_{IT2IT}$  for the IT2IT task. And these losses are defined by the negative log-likelihood of the masked tokens:

$$\mathcal{L}_{I2T} = - \mathbb{E}_{(X, Y) \in \mathcal{D}} \left[ \sum_{\forall j \in [1, N_T], m_j^T = 1} \log p(y_j | Y_{\bar{M}_T}, X) \right], (1)$$

$$\mathcal{L}_{T2I} = - \mathbb{E}_{(X, Y) \in \mathcal{D}} \left[ \sum_{\forall i \in [1, N_I], m_i^I = 1} \log p(x_i | X_{\bar{M}_I}, Y) \right], (2)$$

$$\mathcal{L}_{IT2IT} = - \mathbb{E}_{(X, Y) \in \mathcal{D}} \left[ \sum_{\forall j \in [1, N_T], m_j^T = 1} \log p(y_j | Y_{\bar{M}_T}, X_{\bar{M}_I}) + \sum_{\forall i \in [1, N_I], m_i^I = 1} \log p(x_i | X_{\bar{M}_I}, Y_{\bar{M}_T}) \right], (3)$$

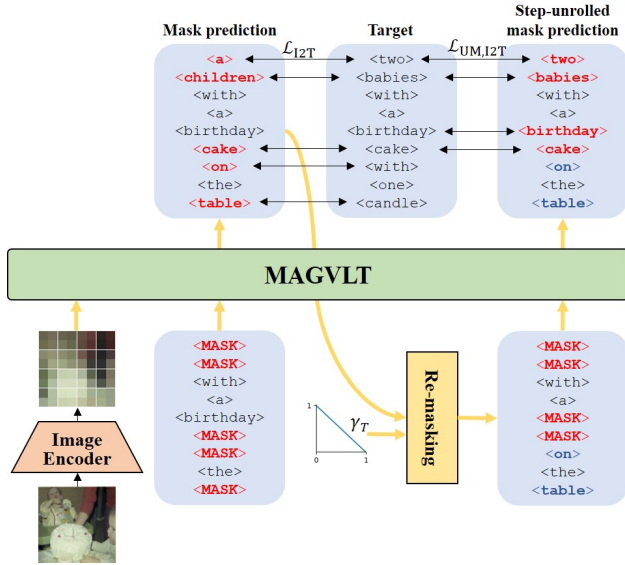


Figure 3. Step-unrolled mask prediction for I2T task. The one-step unrolled sequence is re-masked and then forwarded to the model. Throughout this process, the model faces inputs it would encounter during iterative inference.

where  $\mathcal{D}$  is the training dataset, and  $X_{\bar{M}_I}$  and  $Y_{\bar{M}_T}$  denote the masked image and the masked text, respectively.

### 3.2. Iterative Inference

During inference, the target sequence is predicted by iterative decoding. The mask ratio is defined as a function of the decoding steps as  $\gamma(\frac{k}{K})$  where  $k \in \{0, 1, \dots, K-1\}$  and  $K$  is the total number of iterations. For the first iteration ( $k=0$ ), all the tokens are masked, and the model predicts all the tokens in parallel. For the next  $k$ th iteration, the most  $\lceil \gamma(\frac{k}{K})N \rceil$  unconfident tokens are masked out and predicted again. This process is described in Figure 2. Here, it is noted that following the masking strategy in [10], unmasked image tokens in previous steps are excluded in computing confidences and accordingly will never be masked again. On the other hand, unmasked text tokens can be selected as masked tokens again in the following iterations. Since the number of refinement steps is generally small (e.g., 10), this iterative decoding with parallelizable predictions is significantly faster than the autoregressive decoding, especially when the number of tokens is very large.

**Target Length Prediction.** Since the length of text sequence is varied in contrast to the fixed number of image tokens, non-AR models like MAGVLT need to perform target length prediction for text generation. We follow the length predictor proposed in [22] where the output of  $\langle \text{BOT} \rangle$  that is located between  $X$  and  $Y$  is the predicted text length. At test time, the text sequence is generated after the length is predicted. The loss of the auxiliary target length predictor,  $\mathcal{L}_{\text{TL}}(N_T, \hat{N}_T)$ , is defined by the cross entropy loss between the ground-truth length  $N_T$  and the predicted length  $\hat{N}_T$




Figure 4. An example of a paired data with masked caption (Left), and biased statistics of word compositions (Right).

given the maximum possible number of text tokens.

### 3.3. Step-Unrolled Mask Prediction

Although a variable mask ratio during training reflects various intermediate refinement steps, there still exists a gap between a corruption on the target tokens at training time and a corruption on the partially predicted tokens at test time. SUNDAE [50] tries to resolve this issue by optimizing the model conditioned on a corrupted target sequence which is sampled through one step generative unrolling during training and achieves significant performance improvements of the non-AR autoencoder for text generation.

Here, we adopt and modify this step-unrolled denoising as an additional training task for MAGVLT. Since MAGVLT is based on the masked token prediction, we re-mask the one-step predicted sequence where the mask ratio is reduced from the previous mask ratio and then predict the re-masked tokens by MAGVLT. We call this task as step-unrolled mask prediction, dubbed UnrollMask. We apply UnrollMask only to the I2T and T2I training tasks to maintain the uncorrupted cross-modal context. Figure 3 visually depicts this UnrollMask especially for the I2T task. We denote the UnrollMask loss as  $\mathcal{L}_{\text{UM}}$ , and for example  $\mathcal{L}_{\text{UM}}$  on the I2T task can be defined as

$$\mathcal{L}_{\text{UM,I2T}} = -\mathbb{E}_{(X,Y) \in \mathcal{D}} \left[ \sum_{\forall j \in [1, N_T], m_j^{T^{(+1)}}=1} \log p(y_j | \hat{Y}_{\bar{M}_T}^{(+1)}, X) \right], (4)$$

where  $\hat{Y}_{\bar{M}_T}^{(+1)}$  indicates the re-masked one-step unrolled prediction of  $Y_{\bar{M}_T}$ .

### 3.4. Selective Prediction on Mixed Context

In this multimodal generative modeling, the model often ignores the cross-modal context and produces an output that is biased to the within-modal statistics. For example, in the I2T task of Figure 4, the model should predict the masked word token as ‘cat’ by the given image, however, the model often rather outputs ‘dog’ since ‘dog jumps’ is more likely occurred than ‘cat jumps’ before the text of ‘on grass field’ in the set of training text sequences.

Thus, in order to reduce such bias, we propose a simple yet effective additional learning task, named selective prediction on the mixed context (MixSel), which is described

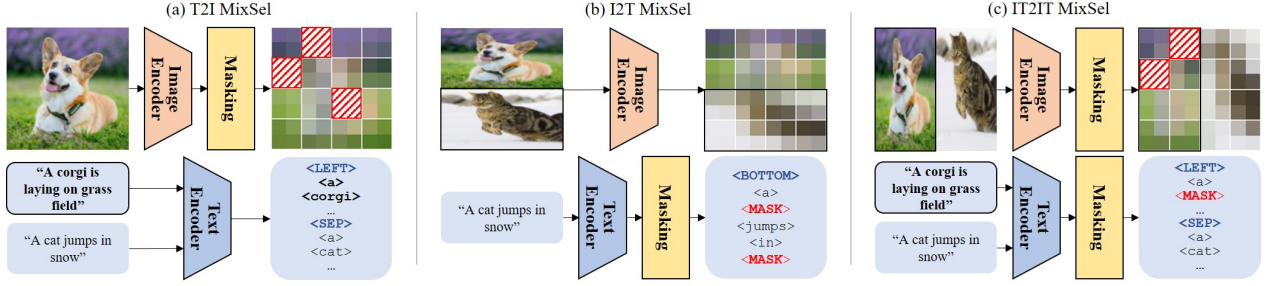


Figure 5. *MixSel* learning tasks corresponding to three multimodal tasks.

in Figure 5. As shown in the figure, two different input contexts are mixed in a half-and-half concatenated manner, and one of them is randomly selected to be the target context in generation. Here, a special token is appended to inform the selected context, for instance `<LEFT>` or `<RIGHT>` is used for the horizontally combined image or the concatenated text sequence while `<TOP>` or `<BOTTOM>` is used for the vertically combined image. Also, when two different text sequences are concatenated, another special token, `<SEP>`, is inserted between them. The *MixSel* objective is denoted as  $\mathcal{L}_{MS}$ , and for instance  $\mathcal{L}_{MS}$  on the I2T task can then be defined as

$$\mathcal{L}_{MS, I2T} = -\mathbb{E}_{(X, Y) \in \mathcal{D}} \left[ \sum_{\forall j \in [1, N_T], m_j^T = 1} \log p(y_j^\ell | \hat{Y}_{M_T}^\ell, \phi(X^1, X^2)) \right], \quad (5)$$

where  $\phi$  is the mixture function on the two images  $X^1$  and  $X^2$ , and  $\ell \in \{1, 2\}$  represents the selected context.

From this *MixSel* training task, the model is able to attend more carefully to the appropriate span of the cross-modal context and improve the accuracy of the cross-modal attention by mixing the original cross-modal content with randomly unrelated one. This could make the model to utilize the cross-modal context more trustfully and hence more often in generation as the training progresses. Therefore, *MixSel* indeed helps in reducing the overlooking of the cross-modal context and circumventing the within-modal bias problem in test-time generation. Note that it is different from the previous mix-based data augmentation techniques [26, 69, 70] in that we retain the information of the original contexts entirely and randomly select the target one for generation. Moreover, although *MixSel* training is relevant to classifier-free guidance (CFG) [28] in that both try to strengthen the effect of the condition, *MixSel* does not have to perform the forward processing twice at test time. Also, we can adapt CFG along with *MixSel* training.

### 3.5. Multitask Pretraining

As we mentioned above, MAGVLT is basically trained via three types of multimodal tasks: I2T, T2I, and IT2IT.

During training, a task  $\tau \in \{I2T, T2I, IT2IT\}$  is sampled from the categorical distribution with the predefined sampling probability  $p_\tau$  for each iteration (batch-wise), and then apply the associated mask prediction loss  $\mathcal{L}_{mask, \tau} \in \{\mathcal{L}_{I2T}, \mathcal{L}_{T2I}, \mathcal{L}_{IT2IT}\}$ . Along with this mask prediction loss, we also add the target length prediction loss  $\mathcal{L}_{TL, \tau}$ , the UnrollMask loss  $\mathcal{L}_{UM, \tau}$ , and the *MixSel* loss  $\mathcal{L}_{MS, \tau}$  according to the selected task  $\tau$ . Overall, the final objective of MAGVLT with respect to  $\tau$  is

$$\mathcal{L}_\tau = \mathcal{L}_{mask, \tau} + \lambda_{TL} \mathcal{L}_{TL, \tau} \cdot \mathbb{I}[\tau \neq T2I] + \lambda_{UM} \mathcal{L}_{UM, \tau} \cdot \mathbb{I}[\tau \neq IT2IT] + \lambda_{MS} \mathcal{L}_{MS, \tau}, \quad (6)$$

where  $\mathbb{I}[\cdot]$  is the indicator function, and  $\lambda_{TL}$ ,  $\lambda_{UM}$ , and  $\lambda_{MS}$  are relative loss weights. Here, we fix  $\lambda_{TL} = 0.01$ ,  $\lambda_{UM} = 1.0$ , and  $\lambda_{MS} = 0.5$  for all our experiments.

## 4. Experiment

In this section, we elaborate the experiments on the VL generation tasks with extensive ablation studies to identify the contribution of each factor in the proposed algorithm. Official codes will be available<sup>†</sup>.

### 4.1. Experimental Setup

**Model.** VLTs (*i.e.* ARGVLT and MAGVLT) have 447M parameters (24 layers, 1024 hidden dimension, and 8 attention heads) including VQ-GAN in total. We also perform experiments about scaling of VLTs, and the results are presented in Appendix. As an image encoder, VQ-GAN [18] converts a  $256 \times 256$  image into  $16 \times 16$  tokens with 16,384 codebook size. For text sequence, we adopt the BPE tokenizer [53] used in CLIP [43] with 49,408 vocabulary size. We fix the text sequence length to 64.

**Dataset.** We pretrain ARGVLT and MAGVLT from scratch using paired image-text datasets. Our pretraining data consists of Conceptual Captions 3M (CC3M) [54], Conceptual Captions 12M (CC12M) [11], SBU Caption [41], and Visual Genome [34] datasets. Together, there are about 17M image-text pairs.

**Pretraining.** There are many options to train VLTs. Note that T2I and I2T are available for both ARGVLT and

<sup>†</sup><https://github.com/kakaobrain/magvlt>




Figure 6. Text-to-Image samples on MS-COCO captions. The images in the second column are sampled from MAGVLT trained without *UnrollMask* and *MixSel*. MAGVLT generated more appropriate images on the corresponding caption. More samples will be found in Appendix.

MAGVLT while IT2IT is only available for MAGVLT. We experiment various subsets of the multimodal tasks in order to investigate the effectiveness of each task. In specific, we pretrain ARGVLT on three different subsets which consist of *T2I only*, *I2T only*, and *T2I & I2T*. Likewise, we pretrain MAGVLT on three different subsets which consist of *T2I only*, *I2T only*, and *T2I & I2T & IT2IT*. More details of pre-training will be found in Appendix.

**Evaluation.** We compare generative VL models based on cross-modality generation tasks especially in *zero-shot* settings in order to evaluate the generalization ability of the proposed method. We also provide more results including finetuning VLTs on downstream tasks in Appendix.

**Sampling.** For text-to-image generation, following [32], we obtain 32 samples from each trained VLT (*i.e.* ARGVLT and MAGVLT) and calculate similarity scores between the sampled images and the conditioning text by CLIP [43] to select a top ranked image (*clip reranking*). Likewise, for image-to-text generation (image captioning), we produce 64 samples from each trained VLT and select a top ranked text by CLIP scores. The number of refinement steps is set to  $K = 10$  and  $K = 12$  for image generation tasks and text generation tasks, respectively.

## 4.2. Image Generation

**Text-to-Image.** We evaluate the zero-shot generalization capability of MAGVLT under text-to-image generation. We measure quantitative metrics of quality of generated images by Fréchet Inception Distance (FID) [27] and Inception Score (IS) [49] on MS-COCO [38] validation. In addition, we compare the relative decoding speed of MAGVLT against ARGVLT. The results are shown in Table 1, where the models are grouped according to: (1) whether they are AR-based or not, and (2) the modality they can generate; the models in the first two groups are able to generate image

Model	FID (↓)	IS (↑)	Speed
<b>AR based</b>			
CM3-Medium (2.7B) [1]	36.78	-	-
DALL-E (12B) [46]	27.5	17.9	-
CogView (4B) [16]	27.1	18.2	-
CogView2 (6B) [17]	24.0	22.4	-
Parti-350M (350M) [67]	14.10	-	-
Make-A-Scene (4B) [20]	<b>11.84</b>	-	-
ARGVLT ( <i>T2I only</i> ) (447M)	21.80	19.27	1.00×
<b>Non-AR based</b>			
GLIDE (3.5B) [40]	12.24	-	-
DALL-E-2 (6.5B) [45]	10.39	-	-
Imagen (4.9B) [48]	7.27	-	-
ERNIE-ViLG 2.0 (24B) [19]	<b>6.75</b>	-	-
<b>MAGVLT (<i>T2I only</i>) (447M)</b>	10.74	23.94	8.12×
<b>Available for both <i>T2I</i> &amp; <i>I2T</i></b>			
UPGen (307M) [4]	65.25	-	-
L-Verse (500M) [32]	37.2	-	-
ARGVLT (447M)	16.93	22.50	1.00×
<b>MAGVLT (447M)</b>	<b>12.08</b>	<b>22.75</b>	8.12×

Table 1. *Zero-shot* T2I results on MS-COCO validation set. Here, we compute FID and IS on a subset of 30,000 captions sampled from COCO validation.

Model	FID (↓)	IS (↑)
ARGVLT	5.62	29.21
<b>MAGVLT</b>	<b>3.17</b>	<b>30.79</b>

Table 2. *Zero-shot* image inpainting results on MS-COCO validation.

only while the models in the last group can generate both image and text. MAGVLTs significantly outperform AR-based methods including ARGVLTs as well as obtain comparable scores to the state-of-the-art diffusion-based methods. Note that all the other models in the non-AR group have more than 1B parameters while MAGVLTs have less than 500M parameters. And the performance gap between the task-specific MAGVLT (*T2I only*) and the universal MAGVLT is small. Moreover, MAGVLTs generate an image more than eight times faster than ARGVLTs. We provide qualitative samples in Figure 6.

**Image Inpainting.** One of the key advantages of MAGVLT over ARGVLT is that it enables bidirectional encoding of conditional information. MaskGIT [10] already demonstrated this advantage. To reconfirm it on MAGVLT, we conduct similar image inpainting experiments. In detail, the central  $8 \times 8$  image tokens corresponding to the central 50% of the whole  $16 \times 16$  image tokens are masked out, and then replaced with newly-generated tokens conditioning on the unmasked image tokens and the ground-truth text tokens. The output images are blended with the input images along the mask boundary following [10]. Quantitatively, as shown in Table 2, MAGVLT outperforms ARGVLT, which is also observed in qualitative samples of Figure 11 in Appendix.

Model	B-4	M	C	S	Speed
<b>with external language model</b>					
ZeroCap (345M) [58]	2.6	11.5	14.6	5.5	-
MAGIC (1.5B) [56]	12.9	17.4	49.3	11.3	-
VLKD <sub>VIT-B/16</sub> (406M) [14]	<b>16.7</b>	<b>19.7</b>	58.3	<b>13.4</b>	-
Flamingo-3B (3B) [3]	-	-	<b>73.0</b>	-	-
<b>without external language model</b>					
SimVLM <sub>huge</sub> (632M) [64]	11.2	14.7	32.2	8.5	-
ARGVLT (I2T only) (447M)	11.4	15.1	47.4	11.4	1.00×
ARGVLT (447M)	10.9	14.9	45.5	11.2	1.00×
<b>MAGVLT (I2T only)</b> (447M)	12.9	17.1	53.5	12.9	1.56×
<b>MAGVLT</b> (447M)	<b>14.6</b>	<b>19.0</b>	<b>60.4</b>	<b>14.3</b>	1.56×

Table 3. Zero-shot I2T results on MS-COCO Karpathy test.

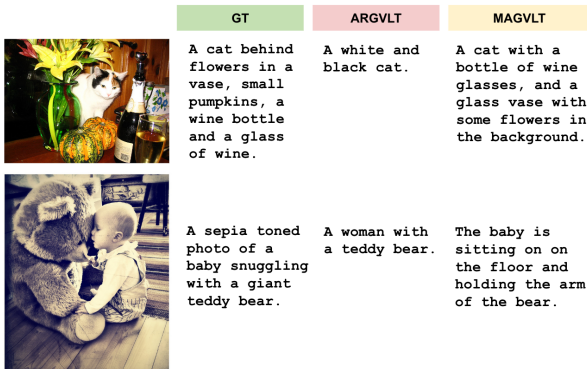


Figure 7. Image-to-text samples on MS-COCO images. MAGVLT generated more proper captions on the corresponding images. More samples will be found in Appendix.

### 4.3. Text Generation

**Image Captioning.** We evaluate the zero-shot image caption generation on MS-COCO Caption [38], and NoCaps [2]. We measure quantitative metrics of the quality of the generated caption compared to the ground truth by BLEU-4 (B-4), METEOR (M), CIDEr (C), and SPICE (S). The evaluation results on MS-COCO are shown in Table 3. Likewise as in text-to-image generation, we can verify that MAGVLTs improve performances over ARGVLTs significantly. Moreover, MAGVLT outperforms some baselines that leverage external language models in CIDEr and SPICE which are specifically designed for the captioning task. Note that MAGVLT has about six times smaller parameters than Flamingo-3B [3], and even can generate images by a single model. In addition, the universal MAGVLT performs better than the task-specific MAGVLT (I2T only) maybe due to the synergetic improvement between the modalities. Regarding the marginal speedup (1.56×) by MAGVLT over ARGVLT for I2T, compared to the fixed number of image tokens (256), the numbers of text tokens are quite small and generally less than 32, thus, the relative speedup by parallel predictions for 12 steps is reduced in text generation. Here, MAGVLT generates a text sequence given the pre-predicted

Model	CIDEr	SPICE
<b>with MS-COCO in training</b>		
FewVLM <sub>large</sub> (740M) [31]	47.7	<b>9.1</b>
MetaLM (545M) [25]	<b>58.7</b>	8.6
<b>without MS-COCO in training</b>		
VLKD <sub>RN50x16</sub> (406M) [14]	54.0	<b>9.6</b>
SimVLM <sub>huge</sub> (632M) [64]	<b>101.4</b>	-
ARGVLT (I2T only) (447M)	34.8	6.5
ARGVLT (447M)	33.4	6.4
<b>MAGVLT (I2T only)</b> (447M)	37.7	7.2
<b>MAGVLT</b> (447M)	46.3	8.7

Table 4. Zero-shot I2T results on NoCaps validation.

Model	B-4	M	C	S
ARGVLT	39.5	32.2	119.8	23.8
<b>MAGVLT</b>	<b>42.7</b>	<b>35.2</b>	<b>135.3</b>	<b>26.3</b>

Table 5. Zero-shot text infilling results on MS-COCO Karpathy test.

target length. We provide qualitative samples in Figure 7.

Here, since L-Verse [32], which can generate both modalities, provides I2T results only by scratch training on MS-COCO, for more comparison to L-Verse, we also perform the training from scratch on MS-COCO alone and obtain a much better FID score for T2I but a slightly lower CIDEr score for I2T compared to L-Verse (18.49 vs. 45.8 in FID, 85.3 vs. 102.2 in CIDEr).

The evaluation results on NoCaps are shown in Table 4. Basically, including MS-COCO in training is beneficial in performing on NoCaps since the interface of the caption collection for NoCaps closely resembles that used for the collection of the MS-COCO captions. Yet, MAGVLT shows comparable performances compared to FewVLM [31]. Also in this task, MAGVLT significantly outperforms ARGVLT. MAGVLTs underperform in comparisons to VLKD [14] and SimVLM [64]. This might be due to that VLKD (RN50x16) has larger parameters ( $\approx 700M$ ) than MAGVLT and also leverages an external language model (BART [37]) unlike MAGVLT. SimVLM uses ALIGN dataset [30] which contains 1.8B image-text pairs as well as C4 [44] dataset which contains 360M text-only instances while MAGVLT uses only 17M image-text pairs for pretraining.

**Text Infilling.** Similar to image inpainting, we conduct the text infilling experiment where the central 50% part of the text is erased by a mask and then replaced with generated tokens from the trained model. The qualitative samples are shown in Figure 13 in Appendix, where we can see that the infilled words generated by MAGVLT are better aligned with surrounding context words, compared to ARGVLT. Also, as shown in Table 5, MAGVLT quantitatively outperforms ARGVLT.

Task sample weights	CIDEr ( $\uparrow$ )	FID ( $\downarrow$ )
T2I:I2T:IT2IT		
1:0:0 ( <i>T2I only</i> )	-	<b>10.74</b>
0:1:0 ( <i>I2T only</i> )	53.5	-
0:0:1 ( <i>IT2IT only</i> )	55.3	12.06
8:2:0 ( <i>T2I &amp; I2T</i> )	59.7	13.09
2:1:1	<b>61.7</b>	15.17
6:1:1	60.7	12.65
8:1:1*	60.4	12.08
10:1:1	59.2	12.07

Table 6. Variants of MAGVLT.

#### 4.4. Ablation Studies

**Variants of MAGVLT.** Here, we investigate the effects of sampling weights for three multimodal tasks corresponding to  $p_\tau$  in learning MAGVLT, and the results are shown in Table 6. We denote the three weights in a form of T2I:I2T:IT2IT in the table. The trained model by T2I only produces the best performance for text-to-image generation. However, in contrast to this, the trained model by I2T only shows the worst performance for captioning. We observe that the training loss of the mask prediction of this I2T only model is the lowest compared to the other models. This may suggest that the bias issue we mentioned in subsection 3.4 can be more serious, especially in non-AR methods. As shown in the result, learning T2I along with I2T resolves the issue to some extent. The ratio of 8:2:0 shows better I2T performance in CIDEr but inferior T2I performance in FID than that of 0:0:1. We observe that the inclusion of the uncorrupted cross-modal context is necessary. The model trained with 2:1:1 weights shows the best captioning performance but the worst T2I performance at a time. As the portion of T2I training is getting larger, the model performs better in T2I but worsens in I2T. It means that there is a trade-off between T2I and I2T according to the sampling ratio. We use the most balanced one (8:1:1) super-scripted by \* as the default setting for MAGVLT.

It is noted that regarding the performance drop by IT2IT modeling rather than T2I-only modeling for the T2I task, compared to I2T, in T2I the benefit of focusing on the cross-modal context is less significant, especially in the later refinement steps. Therefore, improved cross-modal attention by IT2IT training could be less effective for the T2I task. Having said that, the performance drop by MAGVLT for T2I is small even though it can also perform I2T. And we observe that in terms of CLIP scores, our IT2IT training is slightly better than T2I-only training (0.3176 vs. 0.3145) due to the enhanced cross-modal alignment. Moreover, the capacity of our moderate-sized model would be still limited in generating both modalities. Under this limited capacity, there exist some performance trade-offs, and here it would be more leaned to I2T. Table 9 in Appendix shows that when we increase the model capacity about two times, the large model with IT2IT performs better than the medium model

Model	CIDEr ( $\uparrow$ )	FID ( $\downarrow$ )
MAGVLT ( <i>T2I only</i> )	-	<b>10.74</b>
w/o MixSel	-	10.97
w/o UnrollMask and MixSel	-	11.72
MAGVLT ( <i>I2T only</i> )	<b>53.5</b>	-
w/o MixSel	51.3	-
w/o UnrollMask and MixSel	48.0	-
MAGVLT	<b>60.4</b>	<b>12.08</b>
w/o MixSel	58.9	<b>12.07</b>
w/o UnrollMask	56.5	13.26
w/o UnrollMask and MixSel	53.8	14.12

Table 7. Effectiveness of additional training tasks.

with T2I-only training on the T2I task.

**Effectiveness of Additional Tasks.** Here, we investigate the effects of the additional tasks (*i.e.* *UnrollMask* and *MixSel*), and the results are shown in Table 7. MAGVLTs without the additional tasks are clearly underperformed in total. Applying *UnrollMask* in pretraining significantly improves the performances of the base models on both T2I and I2T tasks. Including *MixSel* along with *UnrollMask* also improves the performances of the T2I-only model and especially the models on I2T. Similar to IT2IT training, this improved cross-modal attention by *MixSel* could be less effective for T2I task compared to I2T task. We also experimentally found that the performance gain by *MixSel* is somewhat marginal for ARGVLT. We hypothesize that the causal attention of AR models would be hard to encode the randomly changing span of context.

## 5. Conclusion

In this work, we propose MAGVLT as a unified generative VL model that can produce both image and text data. MAGVLT is robustly trained on image-text pairs by multiple cross-modal tasks and significantly outperforms ARGVLT achieving strong performances on both of zero-shot I2T and T2I tasks. In future work, we first need to scale-up MAGVLT in terms of both model and data for better generalizability. Also, we aim to leverage a pretrained language model to enhance natural language processing and eventually to enable in-context VL learning. We will also try to exploit an encoder-decoder architecture to robustly perform on both understanding and generation tasks.

## Acknowledgements

We would like to thank Brain Cloud Team at Kakao Brain for their support. This work was also supported by Korea University Grant (K2304351) and Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No. 2022-0-00612, Geometric and Physical Commonsense Reasoning based Behavior Intelligence for Embodied AI).



## References

- [1] Armen Aghajanyan, Bernie Huang, Candace Ross, Vladimir Karpukhin, Hu Xu, Naman Goyal, Dmytro Okhonko, Mandar Joshi, Gargi Ghosh, Mike Lewis, et al. Cm3: A causal masked multimodal model of the internet. *arXiv preprint arXiv:2201.07520*, 2022. **6**
- [2] Harsh Agrawal, Karan Desai, Yufei Wang, Xinlei Chen, Rishabh Jain, Mark Johnson, Dhruv Batra, Devi Parikh, Stefan Lee, and Peter Anderson. nocaps: novel object captioning at scale. In *ICCV*, 2019. **1, 7**
- [3] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katie Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. *arXiv preprint arXiv:2204.14198*, 2022. **1, 2, 7**
- [4] Anonymous. Connecting representation and generation via masked vision-language transformer. In *Submitted to The Eleventh International Conference on Learning Representations*, 2023. under review. **1, 2, 3, 6**
- [5] Jacob Austin, Daniel D. Johnson, Jonathan Ho, Daniel Tarlow, and Rianne van den Berg. Structured denoising diffusion models in discrete state-spaces. *arXiv preprint arXiv:2107.03006*, 2021. **1, 3**
- [6] Roman Bachmann, David Mizrahi, Andrei Atanov, and Amir Zamir. Multimaec: Multi-modal multi-task masked autoencoders. In *ECCV*, 2022. **2**
- [7] Hangbo Bao, Wenhui Wang, Li Dong, Qiang Liu, Owais Khan Mohammed, Kriti Aggarwal, Subhojit Som, and Furu Wei. Vlmo: Unified vision-language pre-training with mixture-of-modality-experts. *arXiv preprint arXiv:2111.02358*, 2021. **1, 2**
- [8] Hangbo Bao, Wenhui Wang, Li Dong, and Furu Wei. Vl-beit: Generative vision-language pretraining. *arXiv preprint arXiv:2206.01127*, 2022. **1, 2**
- [9] Minwoo Byeon, Beomhee Park, Haecheon Kim, Sungjun Lee, Woonhyuk Baek, and Saehoon Kim. Coyo-700m: Image-text pair dataset. <https://github.com/kakaobrain/coyo-dataset>, 2022. **1**
- [10] Huiwen Chang, Han Zhang, Lu Jiang, Ce Liu, and William T. Freeman. Maskgit: Masked generative image transformer. In *CVPR*, 2022. **2, 3, 4, 6**
- [11] 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 *CVPR*, 2021. **1, 5**
- [12] Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. Uniter: Universal image-text representation learning. In *ECCV*, 2020. **1**
- [13] Jaemin Cho, Jie Lei, Hao Tan, and Mohit Bansal. Unifying vision-and-language tasks via text generation. In *ICML*, 2021. **1, 2**
- [14] Wenliang Dai, Lu Hou, Lifeng Shang, Xin Jiang, Qun Liu, and Pascale Fung. Enabling multimodal generation on clip via vision-language knowledge distillation. *arXiv preprint arXiv:2203.06386*, 2022. **1, 2, 7, 13**
- [15] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. **2**
- [16] Ming Ding, Zhuoyi Yang, Wenyi Hong, Wendi Zheng, Chang Zhou, Da Yin, Junyang Lin, Xu Zou, Zhou Shao, Hongxia Yang, et al. Cogview: Mastering text-to-image generation via transformers. In *NeurIPS*, 2021. **1, 2, 6**
- [17] Ming Ding, Wendi Zheng, Wenyi Hong, and Jie Tang. Cogview2: Faster and better text-to-image generation via hierarchical transformers. *arXiv preprint arXiv:2204.14217*, 2022. **1, 2, 6**
- [18] Patrick Esser, Robin Rombach, and Björn Ommer. Taming transformers for high-resolution image synthesis. In *CVPR*, 2021. **3, 5**
- [19] Zhida Feng, Zhenyu Zhang, Xintong Yu, Yewei Fang, Lanxin Li, Xuyi Chen, Yuxiang Lu, Jiayang Liu, Weichong Yin, Shikun Feng, Yu Sun, Hao Tian, Hua Wu, and Haifeng Wang. Ernie-vilg 2.0: Improving text-to-image diffusion model with knowledge-enhanced mixture-of-denoising-experts. *arXiv preprint arXiv:2210.15257*, 2022. **1, 3, 6**
- [20] 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. In *ECCV*, 2022. **6**
- [21] Xinyang Geng, Hao Liu, Lisa Lee, Dale Schuurmans, Sergey Levine, and Pieter Abbeel. Multimodal masked autoencoders learn transferable representations. *arXiv preprint arXiv:2205.14204*, 2022. **2**
- [22] Marjan Ghazvininejad, Omer Levy, Yinhan Liu, and Luke Zettlemoyer. Mask-predict: Parallel decoding of conditional masked language models. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 6112–6121, Hong Kong, China, Nov. 2019. Association for Computational Linguistics. **2, 3, 4**
- [23] Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In *CVPR*, 2017. **1, 13**
- [24] Agrim Gupta, Stephen Tian, Yunzhi Zhang, Jiajun Wu, Roberto Martín-Martín, and Li Fei-Fei. Maskvit: Masked visual pre-training for video prediction. *arXiv preprint arXiv:2206.11894*, 2022. **2, 3**
- [25] 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 arXiv:2206.06336*, 2022. **1, 2, 7, 13**
- [26] Dan Hendrycks, Norman Mu, Ekin D. Cubuk, Barret Zoph, Justin Gilmer, and Balaji Lakshminarayanan. Augmix: A simple data processing method to improve robustness and uncertainty. In *ICLR*, 2020. **5**
- [27] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *NeurIPS*, page 6629–6640, Red Hook, NY, USA, 2017. Curran Associates Inc. **6**
- [28] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022. **5**
- [29] Emiel Hoogeboom, Didrik Nielsen, Priyank Jaini, Patrick Forré, and Max Welling. Argmax flows and multinomial dif-

- fusion: Learning categorical distributions. In *NeurIPS*, 2021. [3](#)
- [30] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc V. Le, Yunhsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *ICML*, 2021. [1](#), [2](#), [7](#)
- [31] Woojeong Jin, Yu Cheng, Yelong Shen, Weizhu Chen, and Xiang Ren. A good prompt is worth millions of parameters? low-resource prompt-based learning for vision-language models. *arXiv preprint arXiv:2110.08484*, 2021. [1](#), [2](#), [7](#)
- [32] Taehoon Kim, Gwangmo Song, Sihaeng Lee, Sangyun Kim, Yewon Seo, Soonyoung Lee, Seung Hwan Kim, Honglak Lee, and Kyunghoon Bae. L-verse: Bidirectional generation between image and text. In *CVPR*, 2022. [1](#), [2](#), [6](#), [7](#)
- [33] Wonjae Kim, Bokyung Son, and Ildoo Kim. Vilt: Vision-and-language transformer without convolution or region supervision. In *ICML*, 2021. [1](#), [2](#)
- [34] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, Michael Bernstein, and Li Fei-Fei. Visual genome: Connecting language and vision using crowdsourced dense image annotations. 2016. [5](#)
- [35] Gukyeong Kwon, Zhaowei Cai, Avinash Ravichandran, Erhan Bas, Rahul Bhotika, and Stefano Soatto. Masked vision and language modeling for multi-modal representation learning. *arXiv preprint arXiv:2208.02131*, 2022. [1](#), [2](#)
- [36] Doyup Lee, Chiheon Kim, Saehoon Kim, Minsu Cho, and Wook-Shin Han. Draft-and-revise: Effective image generation with contextual rq-transformer. In *NeurIPS*, 2022. [2](#), [3](#)
- [37] Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In *ACL*, pages 7871–7880, Online, July 2020. Association for Computational Linguistics. [7](#)
- [38] Tsung-Yi Lin, Michael Maire, Serge Belongie, Lubomir Bourdev, Ross Girshick, James Hays, Pietro Perona, Deva Ramanan, C. Lawrence Zitnick, and Piotr Dollár. Microsoft coco: Common objects in context. *arXiv preprint arXiv:1405.0312*, 2014. [1](#), [2](#), [6](#), [7](#)
- [39] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In *NeurIPS*, 2019. [2](#)
- [40] Alex 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. *arXiv preprint arXiv:2112.10741*, 2021. [1](#), [3](#), [6](#)
- [41] Vicente Ordonez, Girish Kulkarni, and Tamara Berg. Im2text: Describing images using 1 million captioned photographs. In J. Shawe-Taylor, R. Zemel, P. Bartlett, F. Pereira, and K.Q. Weinberger, editors, *NeurIPS*, volume 24. Curran Associates, Inc., 2011. [5](#)
- [42] Bryan A. Plummer, Liwei Wang, Chris M. Cervantes, Juan C. Caicedo, Julia Hockenmaier, and Svetlana Lazebnik. Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. In *ICCV*, 2015. [1](#)
- [43] 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. *arXiv preprint arXiv:2103.00020*, 2021. [1](#), [2](#), [5](#), [6](#)
- [44] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21(140):1–67, 2020. [7](#)
- [45] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022. [1](#), [3](#), [6](#)
- [46] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *ICML*, 2021. [1](#), [2](#), [6](#), [12](#)
- [47] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022. [3](#)
- [48] 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, Tim Salimans, Jonathan Ho, David J Fleet, and Mohammad Norouzi. Photorealistic text-to-image diffusion models with deep language understanding. *arXiv preprint arXiv:2205.11487*, 2022. [1](#), [3](#), [6](#)
- [49] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. *Advances in neural information processing systems*, 29, 2016. [6](#)
- [50] Nikolay Savinov, Junyoung Chung, Mikolaj Binkowski, Erich Elsen, and Aaron van den Oord. Step-unrolled denoising autoencoders for text generation. In *ICLR*, 2022. [2](#), [4](#)
- [51] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, Patrick Schramowski, Srivatsa Kundurthy, Katherine Crowson, Ludwig Schmidt, Robert Kaczmarczyk, and Jenia Jitsev. Laion-5b: An open large-scale dataset for training next generation image-text models. *arXiv preprint arXiv:2210.08402*, 2022. [1](#)
- [52] 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 arXiv:2111.02114*, 2021. [1](#)
- [53] Rico Sennrich, Barry Haddow, and Alexandra Birch. Neural machine translation of rare words with subword units. In *ACL*, 2016. [3](#), [5](#)
- [54] Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricu. Conceptual captions: A cleaned, hypemymed, image alt-text dataset for automatic image captioning. In *ACL*, 2018. [1](#), [5](#)

- [55] Krishna Srinivasan, Karthik Raman, Jiecao Chen, Michael Bendersky, and Marc Najork. Wit: Wikipedia-based image text dataset for multimodal multilingual machine learning. *arXiv preprint arXiv:2103.01913*, 2021. [1](#)
- [56] Yixuan Su, Tian Lan, Yahui Liu, Fangyu Liu, Dani Yogatama, Yan Wang, Lingpeng Kong, and Nigel Collier. Language models can see: Plugging visual controls in text generation. *arXiv preprint arXiv:2205.02655*, 2022. [7](#)
- [57] Hao Tan and Mohit Bansal. Lxmert: Learning cross-modality encoder representations from transformers. In *EMNLP*, 2019. [2](#)
- [58] Yoad Towel, Yoav Shalev, Idan Schwartz, and Lior Wolf. Zerocap: Zero-shot image-to-text generation for visual-semantic arithmetic. In *CVPR*, 2022. [7](#)
- [59] Ruben Villegas, Mohammad Babaeizadeh, Pieter-Jan Kindermans, Hernan Moraldo, Han Zhang, Mohammad Taghi Saffar, Santiago Castro, Julius Kunze, and Dumitru Erhan. Phenaki: Variable length video generation from open domain textual description. *arXiv preprint arXiv:2210.02399*, 2022. [2](#), [3](#)
- [60] Jianfeng Wang, Xiaowei Hu, Zhe Gan, Zhengyuan Yang, Xiyang Dai, Zicheng Liu, Yumao Lu, and Lijuan Wang. Ufo: A unified transformer for vision-language representation learning. *arXiv preprint arXiv:2111.10023*, 2021. [2](#)
- [61] Jianfeng Wang, Zhengyuan Yang, Xiaowei Hu, Linjie Li, Kevin Lin, Zhe Gan, Zicheng Liu, Ce Liu, and Lijuan Wang. Git: A generative image-to-text transformer for vision and language. *arXiv preprint arXiv:2205.14100*, 2022. [1](#), [2](#)
- [62] Peng Wang, An Yang, Rui Men, Junyang Lin, Shuai Bai, Zhikang Li, Jianxin Ma, Chang Zhou, Jingren Zhou, and Hongxia Yang. Ofa: Unifying architectures, tasks, and modalities through a simple sequence-to-sequence learning framework. In *ICML*, 2022. [1](#), [2](#)
- [63] Wenhui Wang, Hangbo Bao, Li Dong, Johan Bjorck, Zhiliang Peng, Qiang Liu, Kriti Aggarwal, Owais Khan Mohammed, Saksham Singhal, Subhojit Som, and Furu Wei. Image as a foreign language: Beit pretraining for all vision and vision-language tasks. *arXiv preprint arXiv:2208.10442*, 2022. [1](#), [2](#)
- [64] Zirui Wang, Jiahui Yu, Adams Wei Yu, Zihang Dai, Yulia Tsvetkov, and Yuan Cao. Simvlm: Simple visual language model pretraining with weak supervision. In *ICLR*, 2022. [1](#), [2](#), [7](#), [13](#)
- [65] Haiyang Xu, Ming Yan, Chenliang Li, Bin Bi, Songfang Huang, Wenming Xiao, and Fei Huang. E2e-vlp: End-to-end vision-language pre-training enhanced by visual learning. In *ACL*, 2021. [2](#)
- [66] Jiahui Yu, Zirui Wang, Vijay Vasudevan, Legg Yeung, Mojtaba Seyedhosseini, and Yonghui Wu. Coca: Contrastive captioners are image-text foundation models. *arXiv preprint arXiv:2205.01917*, 2022. [1](#), [2](#)
- [67] Jiahui Yu, Yuanzhong Xu, Jing Yu Koh, Thang Luong, Gunjan Baid, Zirui Wang, Vijay Vasudevan, Alexander Ku, Yinfei Yang, Burcu Karagol Ayan, Ben Hutchinson, Wei Han, Zarana Parekh, Xin Li, Han Zhang, Jason Baldridge, and Yonghui Wu. Scaling autoregressive models for content-rich text-to-image generation. *arXiv preprint arXiv:2206.10789*, 2022. [1](#), [2](#), [6](#)
- [68] Lu Yuan, Dongdong Chen, Yi-Ling Chen, Noel Codella, Xiyang Dai, Jianfeng Gao, Houdong Hu, Xuedong Huang, Boxin Li, Chunyuan Li, Ce Liu, Mengchen Liu, Zicheng Liu, Yumao Lu, Yu Shi, Lijuan Wang, Jianfeng Wang, Bin Xiao, Zhen Xiao, Jianwei Yang, Michael Zeng, Luowei Zhou, and Pengchuan Zhang. Florence: A new foundation model for computer vision. *arXiv preprint arXiv:2111.11432*, 2021. [2](#)
- [69] Sangdoon Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *ICCV*, 2019. [5](#)
- [70] Hongyi Zhang, Moustapha Cisse, Yann N. Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. In *ICLR*, 2018. [5](#)
- [71] Pengchuan Zhang, Xiujun Li, Xiaowei Hu, Jianwei Yang, Lei Zhang, Lijuan Wang, Yejin Choi, and Jianfeng Gao. Vinvl: Revisiting visual representations in vision-language models. *arXiv preprint arXiv:2101.00529*, 2021. [2](#)