

DU-VLG: Unifying Vision-and-Language Generation via Dual Sequence-to-Sequence Pre-training

Luyang Huang Guocheng Niu Jiachen Liu Xinyan Xiao and Hua Wu

Baidu Inc., Beijing, China

{huangluyang, niuguocheng, liujiachen, xiaoxinyan, wu_hua}@baidu.com

Abstract

Due to the limitations of the model structure and pre-training objectives, existing vision-and-language generation models cannot utilize pairwise images and text through bi-directional generation. In this paper, we propose **DU-VLG**, a framework which unifies vision-and-language generation as sequence generation problems. DU-VLG is trained with novel dual pre-training tasks: multi-modal denoising autoencoder tasks and modality translation tasks. To bridge the gap between image understanding and generation, we further design a novel commitment loss. We compare pre-training objectives on image captioning and text-to-image generation datasets. Results show that DU-VLG yields better performance than variants trained with uni-directional generation objectives or the variant without the commitment loss. On the image captioning task, our model reaches better performance than other pre-trained systems. On text-to-image generation datasets, our model achieves better or comparable results than previous state-of-the-art models. In addition, human judges further confirm that our model generates real and relevant images as well as faithful and informative captions.

1 Introduction

Pre-trained models for vision-and-language tasks have made remarkable progress recently (Lu et al., 2019; Su et al., 2020; Chen et al., 2020). Existing pre-trained models either focus on text-to-image synthesis or image-to-text generation (Ramesh et al., 2021; Cho et al., 2021). These models are often pre-trained with image-text pairs which are aligned in semantics. However, due to the limitations of model structure, existing models cannot be adapted to each other. In addition, pre-training objectives are designed either for text generation conditioned on the image or image generation conditioned on the text, limiting the model to learn better semantic alignment from bi-directional generation (Xu et al., 2021; Ding et al., 2021).

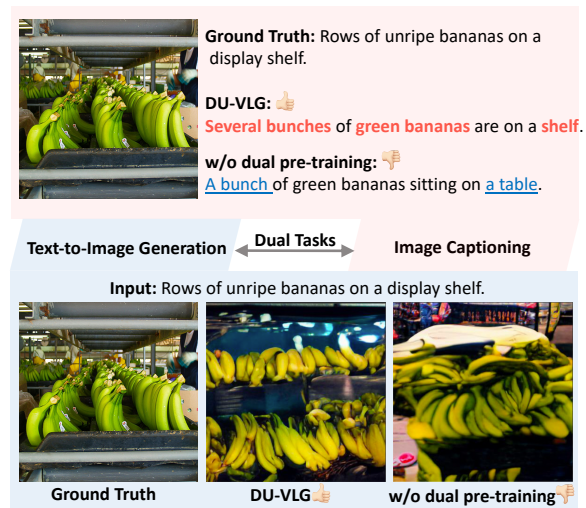


Figure 1: An example from COCO dataset. For image captioning, our system generates informative captions, with key words highlighted in **bold**. Incorrect information is underlined. For text-to-image generation, our system synthesizes vivid images aligned with captions.

We argue that *image-to-text and text-to-image generation appear as dual tasks, which both require strong visual and textual representations aligned in the same semantic space*. Images and text descriptions are of different information quantity and density. The images often contain more information, but are with heavy redundancy, while text descriptions are semantically condensed, but may neglect details. Uni-directional generation paradigm may induce the model to amplify this property. Take Fig.1 as an example, the uni-directional model may fail in capturing details. Inspired by this observation, we propose to utilize bi-directional generation objectives to learn better generalization of image and text representations.

To this end, we present **DU-VLG**, a framework with **DUal** sequence-to-sequence pre-training for **Vision-and-Language Generation**. Under the encoder-decoder Transformer framework, our model takes text and raw images as inputs and gen-

erate text and images autoregressively. Concretely, images are represented as continuous patch features in the encoder and discrete visual tokens in the decoder. With the hybrid image embedding schema, DU-VLG is able to unify vision-and-language generation in a single model.

In order to utilize dualities of image-text pairs, we further propose **two pairs of dual pre-training tasks**: multi-modal denoising autoencoder task and modality translation task. For the multi-modal denoising autoencoder task, our model takes image-text pairs with some image patches or words randomly masked as inputs and learns image-text alignment through reconstruction of the corrupted modality. For modality translation tasks, we form image captioning and text-to-image generation as dual pre-training tasks, which further enhance model ability of semantic alignment. Different from existing multi-modal pre-trained models, our model learns image-text alignment through bi-directional generation objectives.

Moreover, we propose **a novel commitment loss** to drive the model to acquire better image representation. Concretely, the commitment loss is designed to connect visual embeddings in the decoder to patch-based features in the encoder. In tandem with our model design, the commitment loss aims to unify image understanding and generation in a single model, which allows for better utilization of bi-directional generation objectives.

We conduct experiments on various vision-and-language generation tasks. We first study effects of dual pre-training tasks and the commitment loss. On both image captioning and text-to-image generation tasks, DU-VLG outperforms its variant without commitment loss or the variants that only learns uni-directional generation objectives. For image captioning, we achieve better BLEU-4 and CIDER than existing pre-trained models on COCO dataset (Lin et al., 2014). For text-to-image generation, our model achieves better results than both Transformer-based and GAN-based methods on both COCO and CUB dataset (Welinder et al., 2010). Human judges confirm that our model generates captions and images with high-quality. Importantly, we test our model on a challenging vision-and-language generation task: visual commonsense reasoning (Park et al., 2020). Results demonstrate that our model is able to handle challenging multi-modal generation tasks effectively.

The main contributions of DU-VLG are as fol-

lows:

- We unify vision-and-language generation tasks with a single model, DU-VLG. With an encoder-decoder Transformer, DU-VLG is able to handle various vision-and-language generation tasks.
- DU-VLG is pre-trained with novel dual pre-training tasks, which utilizes dualities of image-text pairs. DU-VLG yields better or comparable results than existing state-of-the-art methods on three vision-and-language generation tasks.
- We further propose a new commitment loss, which aims to bridge the gap between image understanding and generation inner with our proposed dual paradigm. Experimental results show that the ability of dual tasks is further enhanced.

The rest of the paper is organized as follows. We describe our model in § 2 and introduce our proposed pre-training task and commitment loss in § 3. Training details are presented in § 4. In § 5, we discuss experimental results. Related work is listed in § 6 and we finally draw our conclusion in § 7.

2 Model

In this section, we describe our proposed model. Overall, our model design is mainly inspired by two observations: (1) sharing parameters that play the same role boosts model performance (Xia et al., 2018) and (2) image understanding and generation require representing image features in different granularity (Cho et al., 2020). Hence, we use a standard Transformer with the encoder-decoder structure (Vaswani et al., 2017), as illustrated in Fig.2. Our model takes images and text as inputs and treats image and text generation as sequence generation problems. Importantly, we propose to use a hybrid image embedding schema in the encoder and the decoder.

2.1 Encoder

In the encoder, images and text are first passed to embedding layers to obtain text embeddings \mathbf{x}_{text} and image embeddings $\mathbf{x}_{\text{image}}$. For text embedding, we follow RoBERTa and tokenize inputs into BPEs (Liu et al., 2020). Each BPE token is represented as the summation of word embedding and position embedding. Unlike text, Images are represented as pixels in a continuous semantic space. However, using pixels as image tokens results in a huge amount of computational cost since model needs to process long sequences. In order to main-

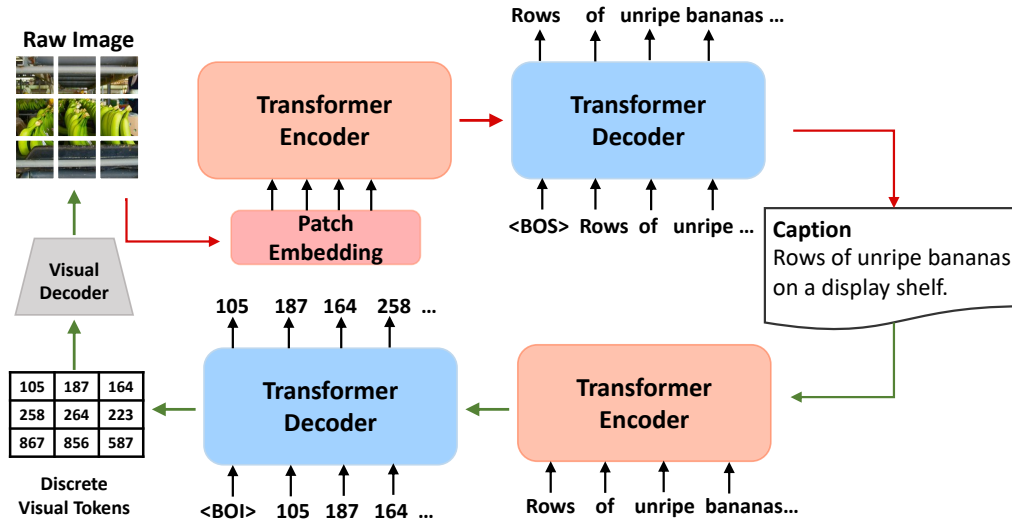


Figure 2: An overview of DU-VLG. Our model is able to take images and text as inputs and generates images and text recurrently. In order to adapt image inputs to the Transformer-based model, we use a hybrid image embedding schema in encoder and decoder. The same color indicates that model parameters are shared for both images and text. The visual decoder weights are not used during training. The symmetric structure is designed for learning better representations from dual pre-training tasks.

tain semantic information as well as reduce the computational cost, we split raw images into a grid of patches.

Image Embedding for Encoder. In the encoder, image inputs are flattened to a sequence of patches, with each patch represents the feature of $p \times p$ pixels. To obtain patch embedding, we pass input images to a trained Vision Transformer (ViT) (Dosovitskiy et al., 2021) and take hidden states of the last layer $\mathbf{x}_{\text{image}}$ as image patch embeddings.

Image and text embeddings are then concatenated and fed into the encoder self-attention layers. If either image or text is missing in the input, we use a [IMAGEPAD] or [TEXTPAD] token as the placeholder.

2.2 Decoder

In the decoder, we use two embeddings: the text embedding which shares weights with the text embedding in the encoder and the image embedding which maps discrete visual tokens to embedding vectors. To enable autoregressive generation, we add [BOI] and [EOI] token to denote the start and the end of the image sequence.

Discrete Visual Tokens for Decoder. In the decoder, the model generates a sequence of discrete visual tokens recurrently. During training, ground truth visual tokens are obtained by a Vector Quantised Variational Autoencoder (VQ-VAE) (van den Oord et al., 2017). The VQ-VAE contains two mod-

ules, an image tokenizer and a visual decoder. The image tokenizer first extracts grid features from raw images and maps into discrete tokens $\mathbf{y}_{\text{image}}$. The visual decoder reconstructs the original image from discrete visual tokens. The image tokenizer represents each $p \times p$ pixels as a visual token, with a vocabulary size of $|\mathcal{V}|$. Therefore, the number of decoder visual tokens is the same as the number of encoder patch tokens. We refer to the original paper for more details. Importantly, during testing, model first generates a sequence of image tokens recurrently and reconstruct the image with the visual decoder.

3 Dual Pre-training Tasks and Pre-training Objectives

Next, we introduce our pre-training method. Pre-training corpus consists of millions of aligned image-text pairs. In order to effectively learn vision-and-language understanding and generation, we propose dual pre-training tasks. Dual pre-training tasks drive the model to learn from reconstruction of the image or text description based on given context. We propose two pairs of pre-training tasks: (1) multi-modal denoising autoencoder task (§ 3.1) and (2) modality translation task (§ 3.2), as shown in Fig.3. In § 3.3, we formulate a commitment loss to connect image understanding and generation.

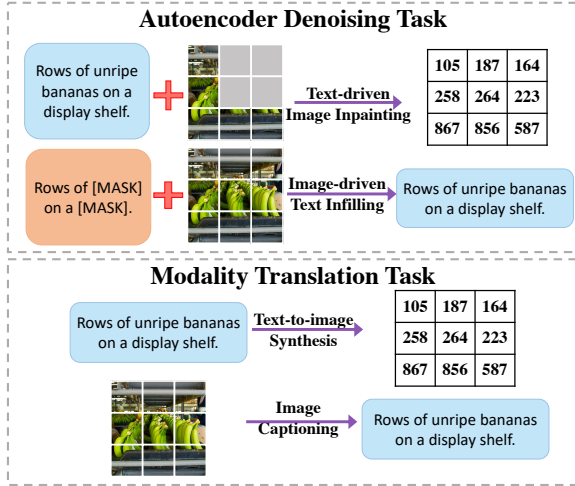


Figure 3: An illustration of our proposed dual pre-training tasks. The model reconstructs the image or text conditioned on its visual and textual context.

3.1 Multi-modal Denoising Autoencoder Task

Given an image-text pair (V, W) from the training set D , we first obtain image patch embeddings $\mathbf{x}_{\text{image}}$ computed by ViT layers and attain text embeddings \mathbf{x}_{text} . To encourage the model to learn cross-modal contextualized embeddings, we propose two dual tasks: 1) text-driven image inpainting task which aims to reconstruct the original image and 2) image-driven text infilling task which aims to reconstruct the original text.

Text-Driven Image Inpainting. Given image patch embeddings $\mathbf{x}_{\text{image}}$, we replace 50 percent of image patches with the same number of trainable [MASK] embeddings, producing masked image sequences $\tilde{\mathbf{x}}_{\text{image}}$. We use blockwise masking algorithm (Bao et al., 2021) to randomly select patches. Meanwhile, we feed the input image to the image tokenizer and produce a sequence of visual tokens $\mathbf{y}_{\text{image}}$. The model is trained to reconstruct the image by optimizing negative log likelihood loss of the ground-truth visual tokens:

$$\mathcal{L}_{\text{image}}^{\text{DAE}} = - \sum_{(V,W) \in D} \log p(\mathbf{y}_{\text{image}} | \tilde{\mathbf{x}}_{\text{image}}, \mathbf{x}_{\text{text}}) \quad (1)$$

Image-Driven Text Infilling. Inspired by text infilling (Lewis et al., 2020), we randomly sample a number of text spans from a Poisson distribution ($\lambda = 3$) and replace with a single [MASK]. Different from text infilling, we randomly mask 50 percent of tokens since we additionally include image as visual context. The model is trained to

optimize negative log likelihood loss of original text tokens:

$$\mathcal{L}_{\text{text}}^{\text{DAE}} = - \sum_{(V,W) \in D} \log p(\mathbf{x}_{\text{text}} | \tilde{\mathbf{x}}_{\text{text}}, \mathbf{x}_{\text{image}}) \quad (2)$$

where $\tilde{\mathbf{x}}_{\text{text}}$ represents the corrupted text sequence.

3.2 Modality Translation Task

In addition to the denoising autoencoder task, we further enhance the model with the modality translation task. The modality translation task drives the model to learn mapping from a modality to the other. Given an image-text pair, we form the modality translation task as two dual tasks: 1) image captioning and 2) text-to-image synthesis.

Image Captioning. Given an image as input, model first produces image patch embeddings $\mathbf{x}_{\text{image}}$ from ViT and encodes image features with encoder self-attentions. The decoder is trained to generate text based on image features. The loss function can be defined as:

$$\mathcal{L}_{\text{text}}^{\text{MT}} = - \sum_{(V,W) \in D} \log p(\mathbf{x}_{\text{text}} | \mathbf{x}_{\text{image}}) \quad (3)$$

Text-to-Image Synthesis. Given a visual description as input, model encodes the input with the encoder and the decoder generates discrete visual tokens $\mathbf{y}_{\text{image}}$ recurrently. During training, the ground truth visual tokens are computed by the image tokenizer. The loss function can be defined as:

$$\mathcal{L}_{\text{image}}^{\text{MT}} = - \sum_{(V,W) \in D} \log p(\mathbf{y}_{\text{image}} | \mathbf{x}_{\text{text}}) \quad (4)$$

3.3 Connecting Image Embedding between Encoder and Decoder.

In the encoder-decoder structure, text embedding is often shared among the encoder, the decoder and the token generation layer (Paulus et al., 2018). This allows the model to learn better syntactic and semantic information. For image embedding, since we use a hybrid embedding schema in the encoder and the decoder, we propose a commitment loss to connect image understanding and generation during training. Intuitively, decoder visual token embeddings $\mathbf{y}_{\text{image}}$ should commit to corresponding

patch embeddings $\mathbf{x}_{\text{image}}$ in encoder. Therefore, the commitment loss uses a square loss to connect the encoder and the decoder:

$$\mathcal{L}_{\text{com}} = - \sum_{(V) \in D} \|\text{sg}[\mathbf{x}_{\text{image}}] - \mathbf{y}_{\text{image}}\|^2 \quad (5)$$

where sg means stopgradient operator which is identity at forward computation but has zero partial derivatives at backward computation. The commitment loss is applied to the text-driven image inpainting objective and the text-to-image synthesis objective.

During training, for each instance, we randomly select a couple of objectives from denoising autoencoder and modality translation. We set probability of denoising autoencoder as 0.6 for all experiments. Therefore, for each batch, the pre-training loss is a combination of three losses:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{text}} + \alpha \mathcal{L}_{\text{image}} \quad (6)$$

$$\mathcal{L}_{\text{image}} = \mathcal{L}_{\text{image}}^{\text{DAE}} + \mathcal{L}_{\text{image}}^{\text{MT}} + \beta \mathcal{L}_{\text{com}} \quad (7)$$

$$\mathcal{L}_{\text{text}} = \mathcal{L}_{\text{text}}^{\text{DAE}} + \mathcal{L}_{\text{text}}^{\text{MT}} \quad (8)$$

where α and β are hyperparameters to control the scale of image loss and commitment loss.

4 Experimental Setup

4.1 Pre-training

Pre-training Corpus. We train our model on four existing datasets that consist of image-text pairs. Our pre-training datasets include 1) Common Objects in Context (COCO) (Lin et al., 2014), 2) Conceptual Captions (CC) (Sharma et al., 2018), 3) SBU Captioned Photo (SBU) (Ordonez et al., 2011) and 4) Visual Genome (VG) (Krishna et al., 2016). For Visual Genome dataset, since captions are collected for image regions, we use image regions and captions as pairs. We additionally filter captions which are fewer than five words. We end up with a collection of about 5 million image-text pairs.

Implementation Detail. We report results on two model sizes: 1) a base version with 6 layers for the encoder and decoder and 2) a large version with 12 layers for the encoder and decoder. For each model size, we report results with two different input image resolutions: 224×224 and 384×384 . Following ViT, we use a patch size of $p = 16$ for all the experiments. For VQ-VAE, we take the off-the-shelf VQ-GAN (Esser et al., 2021), which

is a variant of VQ-VAE. The VQ-GAN maps each 16×16 pixels as a discrete visual token, with a vocabulary size of $|\mathcal{V}| = 16384$.

For base and large model, we use ViT-base and ViT-large with a patch size of $p = 16$ to extract image patch embeddings. ViT weights are set frozen during pre-training. Since image sequences are longer than text sequences, we set $\alpha = 0.05$ and $\beta = 1$ for all experiments. For model optimization, we utilize Adam optimizer with a gradient clipping of 1.0 and a batch size equivalent of 1024.

4.2 Fine-tuning on Downstream Tasks

In order to evaluate model capability of vision-and-language generation tasks, we test on three downstream tasks: 1) text-to-image generation, 2) image captioning and 3) visual commonsense reasoning. Here we mainly introduce evaluation metrics. For additional fine-tuning details, we refer to the appendices.

Text-to-Image Generation. We experiment with two popular text-to-image generation datasets: the Caltech-UCSD Birds 200 dataset (CUB) and Common Objects in Context dataset (COCO).

The CUB dataset contains 200 bird categories with 11,788 images. Each image has ten text descriptions. We follow the standard split which uses 150 categories with 8,855 images for training and the remaining 50 categories with 2,933 images for testing. The COCO dataset contains 82,784 images for training and 40,505 for testing. Each image has five text descriptions.

We fine-tune on the pre-trained model with a learning rate of $1e-4$ for 300 epoches on both datasets. Similar to Ramesh et al. (2021), we sample 16 images per caption with nucleus sampling strategy (Holtzman et al., 2020). During testing, we first sample 16 images per caption and rerank the generated images with a CLIP model (Radford et al., 2021). The CLIP model selects the best image based on its correlation with the text description.

We include two widely used evaluation metrics: 1) Inception Score (IS) (Salimans et al., 2016) and 2) Fréchet Inception Distance (FID) (Heusel et al., 2017). The IS score computes the KL-divergence between the conditional class distribution and the marginal class distribution obtained by a pre-trained Inception v3 model (Szegedy et al., 2016). The FID computes the Fréchet distance between ground-truth images and generated images

based on the features obtained by the Inception v3 model. Higher IS scores and lower FID scores denote that images synthesized by the model are of better quality. Previous work (Li et al., 2019b) reports that the IS score fails in evaluating the quality of images on COCO dataset. Hence, we do not report the IS score on COCO dataset. For fair comparison, we resize our model outputs to 256×256 and calculate FID and IS scores.

Image Captioning. For image captioning, we test our model on COCO dataset. We report four metrics based on word overlapping on COCO dataset: 1) BLEU-4 (Papineni et al., 2002), 2) METEOR (Lavie and Agarwal, 2007), 3) CIDEr (Vedantam et al., 2015) and 4) SPICE (Johnson et al., 2020).

For COCO dataset, we follow the Karparthy split (Karparthy and Fei-Fei, 2015) which has 113,287, 5000 and 5000 images for training, validation and test. Each image has 5 human-written captions. During inference, we generate a caption for each image and evaluate against five references.

We fine-tune on COCO dataset with a learning rate of $3e-5$. Vision Transformer layers are trainable during fine-tuning. Following Li et al. (2020), we add object labels detected by the object detection model as additional text inputs. We find object labels improve CIDEr and BLEU scores for at least 1 point and 0.3 points. During testing, we use beam search with a beam size of 5.

Visual Commonsense Reasoning. Besides image captioning and text-to-image generation, which only requires model to encode one modality, we further test our model on a more challenging dataset, VisualCOMET (Park et al., 2020). VisualCOMET is a visual commonsense reasoning task which provides the model with an image and the event that happens at present. The model is required to infer what may happen next, before and the people’s intents at present. VisualCOMET requires the model to jointly comprehend image and text and generate reasonable inference. Similar to image captioning, we use BLEU-2, METEOR and CIDEr as metrics.

5 Results

In this section, we start with comparing our proposed pre-training objectives in § 5.1. We then conduct automatic evaluation on three vision-and-language generation tasks (§ 5.2) and further report human evaluation on both caption and synthesized

Image -> Text System	COCO Caption			
	BLEU-4	CIDEr	METEOR	SPICE
DU-VLGB-224	38.8	124.8	29.2	22.0
w/o L_{image}	36.9	118.8	28.4	20.5
w/o L_{text}	35.2	112.8	27.4	19.6
w/o L_{com}	38.4	123.1	28.8	21.7
Text -> Image		CUB	COCO	
System	IS \uparrow	FID \downarrow	FID \downarrow	
DU-VLGB-224	5.14	23.78	26.82	
w/o L_{image}	4.84	25.28	36.59	
w/o L_{text}	5.03	24.68	29.64	
w/o L_{com}	5.08	24.44	27.92	

Table 1: Ablation study on pre-training tasks and objectives. The best result per metric per dataset is **bolded**. DU-VLGB-224 yields significantly higher scores than other comparisons with approximate randomization test ($p < 0.0005$).

image quality (§ 5.2). Finally, we investigate inference speed of our proposed model (§ 5.3).

5.1 Comparing Pre-training Objectives

Comparisons. We first investigate whether our proposed dual pre-training tasks and commitment loss improve generation quality. We fine-tune on two downstream tasks: image captioning and text-to-image generation. We report our base model with an input image resolution of 224×224 (DU-VLGB-224). We compare our base model with three variants: 1) the model trained without text-driven image inpainting and text-to-image synthesis tasks (w/o L_{image}), 2) the model trained without image-driven text infilling and image captioning tasks (w/o L_{text}) and 3) the model trained without commitment loss (w/o L_{com}).

Results. As displayed in Tab.1, *our model with dual pre-training tasks performs the best on both image captioning and text-to-image generation tasks*. This demonstrates the benefit of dual pre-training tasks and the commitment loss. For image captioning, comparing with the variant without image generation objectives, our model with dual pre-training tasks significantly improves automatic metrics, which indicates that image generation objectives can boost visual understanding. For text-to-image generation, our model yields better FID and IS scores than the variant without text generation objectives on both CUB and COCO dataset. This demonstrates that using text generation objectives can guide better semantic interpretation of text content.

Moreover, our model outperforms the variant

trained without the commitment loss on two downstream tasks. This further illustrates that the commitment loss improves model performance on both image understanding and generation.

5.2 Automatic Evaluation

Comparisons. We then compare our model with other vision-and-language models. For image captioning, we include state-of-the-art vision-and-language pre-trained models: (1) object-semantics aligned pre-training (OSCAR) (Li et al., 2020), (2) unified modal understanding and generation pre-training (UNIMO) (Li et al., 2021), (3) improving visual representations for vision-and-language pre-training (VINVL) (Zhang et al., 2021b) and (4) end-to-end vision-and-language pre-training (E2E-VLP) (Xu et al., 2021). For OSCAR and VINVL, we report their results with cross-entropy optimization for fair comparison.

For text-to-image generation, we include four Transformer-based models: (1) X-LXMERT, which has 228 million parameters and is trained on 9 million image-text pairs, (2) DALLE, which has 12 billion parameters and is trained on 250 million text-image pairs (Ramesh et al., 2021), (3) COGVIEW, which has 4 billion parameters and is trained on 30 million data (Ding et al., 2021) and (4) NUWA, which has 870 million parameters and is trained on a mixture of text-image pairs and text-video pairs (Wu et al., 2021). We further compare our model with three traditional methods based on generative adversarial network (GAN): (1) DM-GAN (Zhu et al., 2019), (2) DF-GAN (Tao et al., 2020) and (3) XMC-GAN (Zhang et al., 2021a).

For visual commonsense reasoning, we include Vision-Language Transformer (V-L TRANSFORMER) (Park et al., 2020) as a baseline, which fuses region-based visual features into a pre-trained GPT-2 (Radford et al., 2019).

Results. For image captioning, our model achieves better scores than both end-to-end method and two-stage methods. In Tab.2, DU-VLG outperforms previous state-of-the-art pre-trained model VINVL, e.g., improving BLEU-4 and CIDEr by more than 1 and 3 points.

Moreover, for text-to-image generation tasks, our model achieves state-of-the-art IS and FID on CUB dataset, as displayed in Tab.3, outperforming traditional GAN-based methods. Compared with Transformer-based methods, our model yields better or comparable FID scores on COCO datasets.

Image -> Text System	CoCo Caption			
	BLEU-4	CIDER	METEOR	SPICE
OSCAR _B	36.5	123.7	30.7	23.5
UNIMO _B	38.8	124.4	29.8	22.1
VINVL _B	38.2	129.3	30.3	23.6
E2E-VLP	36.2	117.3	–	–
DU-VLG _{B-224}	38.8	124.8	29.2	22.0
DU-VLG _{B-384}	40.0	133.0	30.2	23.8
OSCAR _L	37.4	127.8	30.7	23.5
UNIMO _L	39.6	127.7	29.5	22.4
VINVL _L	38.5	130.8	30.4	23.4
DU-VLG _{L-224}	39.2	128.1	29.8	22.8
DU-VLG _{L-384}	40.1	135.8	30.8	23.9

Table 2: Automatic evaluation on Image Captioning datasets. We report our model and comparisons with two model sizes: the base version (B) and the large version (L) and two input image resolution: 224×224 and 384×384 . Our base and large models have comparable number of parameters compared to other comparisons. The best metric of each model size is **bolded**.

It is worth to note that our models are with fewer parameters and less training data compared with DALLE, COGVIEW and NUWA. This demonstrates the effectiveness of our proposed framework.

In addition, we study the effect of different input image resolutions. We compare two different resolutions of the input images: 224×224 and 384×384 . In Tab.2 and Tab.3, we find higher resolution as inputs leads to better results on both image-to-text and text-to-image generation tasks. This observation remarks the importance of fine-grained image representation.

We then evaluate our model on a more challenging vision-and-language task, visual commonsense reasoning. As shown in Tab.4, our model significantly outperforms V-L TRANSFORMER, which is fine-tuned based on a language model, GPT-2. This demonstrates that our model is able to jointly comprehend image and text inputs and generate informative inference.

5.3 Human Evaluation

We conduct human evaluation to analyze generation quality of images and text. For both image captioning and text-to-image generation, we select 100 samples from COCO test set and hire three annotators to rate captions and images. For image captioning, we include three systems: (1) best performed pre-trained model VINVL (2) our model that removes dual pre-training DU-VLG w/o L_{image} and

Text -> Image System	CUB		COCO
	IS \uparrow	FID \downarrow	FID \downarrow
DM-GAN	4.75	16.09	32.64
DF-GAN	5.10	14.81	21.42
XMC-GAN	–	–	9.33
X-LXMERT	–	–	37.40
DALLE	–	–	27.50
COGVIEW	–	–	26.00
NUWA	–	–	12.90
DU-VLGB ₂₂₄	5.14	23.78	26.82
DU-VLGB ₃₈₄	5.26	14.60	22.41
DU-VLGL ₂₂₄	5.18	21.50	23.25
DU-VLGL ₃₈₄	5.28	14.15	14.48

Table 3: Automatic evaluation on Text-to-Image Generation datasets. For fair comparison, we resize generated images to 256×256 pixels before calculating IS and FID scores.

System	VisualCOMET		
	BLEU-2	CIDER	METEOR
V-L TRANSFORMER	13.5	18.2	11.5
DU-VLGB ₃₈₄	21.5	36.6	25.6
DU-VLGL ₃₈₄	23.9	41.9	27.1

Table 4: Automatic evaluation on visual commonsense reasoning. Our model generates informative inference compared to the baseline.

(3) our best performed model DU-VLG. For text-to-image generation, we compare three models: (1) Transformer-based model pre-trained on about 9 million data X-LXMERT, (2) our model trained without text generation objectives DU-VLG w/o L_{text} and (3) DU-VLG. For our model, we use the large version with the input image resolution of 384×384 .

For image captioning, human judges are asked to rate on two aspects: **informativeness**—whether the caption covers important objects from the image and **faithfulness**—whether the caption correctly describes the image. For text-to-image generation, we consider two aspects: **fidelity**—whether the image is realistic and **relevance**—whether the image matches with the caption. All aspects are rated on a Likert scale from 1 (poor) to 5 (good).

Results. From Fig.4, we find our DU-VLG model obtains better scores in relevance, fidelity, informativeness and faithfulness than the variant that removes dual pre-training tasks. This confirms our claim that bi-directional generation objectives improve semantic alignment between images and text. Meanwhile, compared with well-performed model




Figure 4: Human evaluation on COCO dataset. DU-VLG yields significantly higher scores than other systems on fidelity, relevance, informativeness and faithfulness ($p < 0.05$).

VINVL and X-LXMERT, our model yields better scores on four aspects. This implies that our model generates more informative captions committed to the input images and synthesizes more realistic images aligned with the captions compared to the state-of-the-art pre-trained models.

Interestingly, image captioning models yield higher scores than text-to-image generation models, closer to 5 (perfect). After inspection, we find that our model yields near-perfect captions compared to human written ones, while the generated images sometimes fail in synthesizing details. For example, the shape of a banana may be distorted, limiting the fidelity of the image.

5.4 Inference Efficiency

Next, we compare the inference speed and the number of model parameters with existing models. For image captioning, we compare our model with two best performed pre-trained models: the base version of UNIMO and VINVL. For text-to-image generation, we compare with two transformer-based large models DALLE and Cogview. For our model, we report the base version. We test speed on COCO test set with one 32GB NVIDIA TESLA V100. We include the visual decoder when calculating the inference speed.

In Tab.5, we find our model is roughly $7\times$ faster than two-stage methods on image captioning. This is mainly because extracting image features with ViT is much faster than object detection. Importantly, our model has comparable parameters compared with UNIMO and VINVL.

For text-to-image generation, our model is roughly $400\times$ faster than large model Cogview and has only 5 percent of parameters. This further confirms the importance of dual pre-training tasks.

System	Time(s)	# Param. (M)
Image Captioning		
UNIMOB	0.88+0.12	172
VINVL _B	0.90+0.12	187
DU-VLG _{B-224}	0.14	228
Text-to-Image Generation		
DALLE	–	12,000
COGVIEW	300	4,000
DU-VLG _{B-224}	0.76	228

Table 5: Comparing inference speed (time) and number of parameters (# Param.) on different tasks. For two-stage methods UNIMO and VINVL, we report image feature extraction and beam search time respectively.

6 Related Work

Vision-and-Language Pre-training for Image-to-Text Generation Tasks. Transformer backbones have achieved great success in language pre-training (Devlin et al., 2019; Lewis et al., 2020; Liu et al., 2020). In order to adapt Transformers to multi-modal pre-training, previous work mainly focuses on (1) better image features and (2) designing pre-training tasks (Lu et al., 2019; Li et al., 2019a). To obtain high-quality image features, Image region features extracted from an object detection model are widely adopted in multi-modal pre-training (Zhou et al., 2020; Li et al., 2020; Zhang et al., 2021b). Kim et al. (2021) points out that two-stage method is time-consuming and the trained object detector may fail in the unlabeled domain (Jiang et al., 2021). To that end, Huang et al. (2020) feeds raw images to convolutional backbones such as ResNets (He et al., 2016) and takes its outputs as image features. Kim et al. (2021) uses linear projection to obtain patch-based image features. However, currently, end-to-end image feature extraction methods cannot yield comparable results compared to two-stage methods on image captioning.

To learn image-text alignment, masked token prediction, which masks a portion of text or image tokens and predicts masked positions conditioned on the context, is widely used as the pre-training task (Xia et al., 2020). Li et al. (2020) designs image-text matching task, which predicts whether the image and the text are paired or not. Li et al. (2021) proposes special self-attention masks to unify text understanding and generation. Xu et al. (2021) includes image captioning and object detection as pre-training objectives to enhance the decoder. However, current methods for generation

tasks are limited to text generation and are struggled to learn fine-grained image-text alignment.

In this paper, we introduce a hybrid image embedding schema to connect image understanding and generation, which unifies image and text generation via sequence-to-sequence pre-training. Concretely, we enhance image-text alignment with novel dual pre-training tasks. Our model outperforms state-of-the-art pre-trained systems on image captioning.

Vision-and-Language Pre-training for Text-to-Image Generation Tasks. To generate images autoregressively, images are represented as discrete tokens. X-LXMERT (Cho et al., 2020) partitions image grid features into clusters and obtains visual tokens via nearest-neighbor search. However, X-LXMERT needs to train an image generator from scratch to synthesize images from visual tokens, which accumulates errors during training. Ding et al. (2021); Ramesh et al. (2021) use discrete visual tokens from a trained vector-quantised variational autoencoder (VQ-VAE) (van den Oord et al., 2017) for text-to-image generation. However, their models consist of billions of parameters and require a huge corpus to pre-train (more than 100 million image-text pairs). In this paper, we present a relative small model (about 200M parameters), with better generation quality on COCO dataset. In particular, we offer a detailed analysis on the inference speed and the model size in the appendices.

7 Conclusion

We presented a novel framework, DU-VLG, which unifies vision-and-language generation tasks with an encoder-decoder Transformer. We propose to use a hybrid image embedding schema in the encoder and decoder. In addition, we pre-train the model with novel dual pre-training tasks, along with a new commitment loss, to guide better image and text understanding and generation. Experiments show that our proposed dual pre-training objectives significantly improve performance on three vision-and-language generation tasks. Human evaluation further confirms that our model with dual pre-training tasks improves generation quality on image captioning and text-to-image generation.

8 Acknowledgments

This work was supported by the National Key Research and Development Project of China (No. 2018AAA0101900)

9 Ethics Statement

Large models that are pre-trained on heterogeneous data can be potentially harmful to marginalized populations. Along with the improved controllability, we also recognize that our system might be misused to create offensive or fabricated content. We therefore advocate cautious usage in real-world deployment.

References

- Hangbo Bao, Li Dong, and Furu Wei. 2021. [Beit: BERT pre-training of image transformers](#). *CoRR*, abs/2106.08254.
- Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. 2020. Uniter: Universal image-text representation learning. In *Computer Vision – ECCV 2020*, pages 104–120, Cham. Springer International Publishing.
- Jaemin Cho, Jie Lei, Hao Tan, and Mohit Bansal. 2021. [Unifying vision-and-language tasks via text generation](#). In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 1931–1942. PMLR.
- Jaemin Cho, Jiasen Lu, Dustin Schwenk, Hannaneh Hajishirzi, and Aniruddha Kembhavi. 2020. [X-LXMERT: Paint, Caption and Answer Questions with Multi-Modal Transformers](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8785–8805, Online. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Ming Ding, Zhuoyi Yang, Wenyi Hong, Wendi Zheng, Chang Zhou, Da Yin, Junyang Lin, Xu Zou, Zhou Shao, Hongxia Yang, and Jie Tang. 2021. [Cogview: Mastering text-to-image generation via transformers](#). *CoRR*, abs/2105.13290.
- 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. 2021. [An image is worth 16x16 words: Transformers for image recognition at scale](#). In *International Conference on Learning Representations*.
- Patrick Esser, Robin Rombach, and Bjorn Ommer. 2021. [Taming transformers for high-resolution image synthesis](#). In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12873–12883.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. [Deep residual learning for image recognition](#). In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778.
- Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. 2017. [Gans trained by a two time-scale update rule converge to a local nash equilibrium](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2020. [The curious case of neural text de-generation](#). In *International Conference on Learning Representations*.
- Zhicheng Huang, Zhaoyang Zeng, Bei Liu, Dongmei Fu, and Jianlong Fu. 2020. [Pixel-bert: Aligning image pixels with text by deep multi-modal transformers](#). *CoRR*, abs/2004.00849.
- Junguang Jiang, Baixu Chen, Jianmin Wang, and Ming-sheng Long. 2021. [Decoupled adaptation for cross-domain object detection](#).
- Khia A. Johnson, Molly Babel, Ivan Fong, and Nancy Yiu. 2020. [SpiCE: A new open-access corpus of conversational bilingual speech in Cantonese and English](#). In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 4089–4095, Marseille, France. European Language Resources Association.
- Andrej Karpathy and Li Fei-Fei. 2015. [Deep visual-semantic alignments for generating image descriptions](#). In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Wonjae Kim, Bokyung Son, and Ildoo Kim. 2021. [Vilt: Vision-and-language transformer without convolution or region supervision](#). In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 5583–5594. PMLR.
- Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A. Shamma, Michael S. Bernstein, and Li Fei-Fei. 2016. [Visual genome: Connecting language and vision using crowdsourced dense image annotations](#). *International Journal of Computer Vision*, 123:32–73.
- Alon Lavie and Abhaya Agarwal. 2007. [METEOR: An automatic metric for MT evaluation with high levels of correlation with human judgments](#). In *Proceedings of the Second Workshop on Statistical Machine Translation*, pages 228–231, Prague, Czech Republic. Association for Computational Linguistics.

- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. [BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. 2019a. [Visualbert: A simple and performant baseline for vision and language](#). *CoRR*, abs/1908.03557.
- Wei Li, Can Gao, Guocheng Niu, Xinyan Xiao, Hao Liu, Jiachen Liu, Hua Wu, and Haifeng Wang. 2021. [UNIMO: Towards unified-modal understanding and generation via cross-modal contrastive learning](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 2592–2607, Online. Association for Computational Linguistics.
- Wenbo Li, Pengchuan Zhang, Lei Zhang, Qiuyuan Huang, Xiaodong He, Siwei Lyu, and Jianfeng Gao. 2019b. Object-driven text-to-image synthesis via adversarial training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Xiujun Li, Xi Yin, Chunyuan Li, Pengchuan Zhang, Xiaowei Hu, Lei Zhang, Lijuan Wang, Houdong Hu, Li Dong, Furu Wei, Yejin Choi, and Jianfeng Gao. 2020. Oscar: Object-semantics aligned pre-training for vision-language tasks. In *Computer Vision – ECCV 2020*, pages 121–137, Cham. Springer International Publishing.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. 2014. Microsoft coco: Common objects in context. In *Computer Vision – ECCV 2014*, pages 740–755, Cham. Springer International Publishing.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Ro{bert}a: A robustly optimized {bert} pretraining approach](#).
- Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. 2019. [Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks](#). In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Vicente Ordonez, Girish Kulkarni, and Tamara L. Berg. 2011. Im2text: Describing images using 1 million captioned photographs. In *Neural Information Processing Systems (NIPS)*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Weijing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Jae Sung Park, Chandra Bhagavatula, Roozbeh Motlaghi, Ali Farhadi, and Yejin Choi. 2020. Visualcomet: Reasoning about the dynamic context of a still image. In *Proceedings of the European Conference on Computer Vision (ECCV)*.
- Romain Paulus, Caiming Xiong, and Richard Socher. 2018. [A deep reinforced model for abstractive summarization](#). In *International Conference on Learning Representations*.
- 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. 2021. [Learning transferable visual models from natural language supervision](#). In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 8748–8763. PMLR.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.
- Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. 2021. [Zero-shot text-to-image generation](#). *CoRR*, abs/2102.12092.
- Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, Xi Chen, and Xi Chen. 2016. [Improved techniques for training gans](#). In *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc.
- Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. 2018. [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. Association for Computational Linguistics.
- Weijie Su, Xizhou Zhu, Yue Cao, Bin Li, Lewei Lu, Furu Wei, and Jifeng Dai. 2020. [Vi-bert: Pre-training of generic visual-linguistic representations](#). In *International Conference on Learning Representations*.
- C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna. 2016. [Rethinking the inception architecture for computer vision](#). In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2818–2826, Los Alamitos, CA, USA. IEEE Computer Society.

- Ming Tao, Hao Tang, Songsong Wu, Nicu Sebe, Fei Wu, and Xiao-Yuan Jing. 2020. [Df-gan: Deep fusion generative adversarial networks for text-to-image synthesis](#). *CoRR*, abs/2008.05865.
- Aaron van den Oord, Oriol Vinyals, and Koray Kavukcuoglu. 2017. Neural discrete representation learning. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS'17*, page 6309–6318, Red Hook, NY, USA. Curran Associates Inc.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Ramakrishna Vedantam, C. Lawrence Zitnick, and Devi Parikh. 2015. Cider: Consensus-based image description evaluation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- P. Welinder, S. Branson, T. Mita, C. Wah, F. Schroff, S. Belongie, and P. Perona. 2010. Caltech-UCSD Birds 200. Technical Report CNS-TR-2010-001, California Institute of Technology.
- Chenfei Wu, Jian Liang, Lei Ji, Fan Yang, Yuejian Fang, Daxin Jiang, and Nan Duan. 2021. [Nüwa: Visual synthesis pre-training for neural visual world creation](#). *CoRR*, abs/2111.12417.
- Qiaolin Xia, Haoyang Huang, Nan Duan, Dongdong Zhang, Lei Ji, Zhifang Sui, Edward Cui, Taroon Bharti, Xin Liu, and Ming Zhou. 2020. [XGPT: cross-modal generative pre-training for image captioning](#). *CoRR*, abs/2003.01473.
- Yingce Xia, Xu Tan, Fei Tian, Tao Qin, Nenghai Yu, and Tie-Yan Liu. 2018. [Model-level dual learning](#). In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 5383–5392. PMLR.
- Haiyang Xu, Ming Yan, Chenliang Li, Bin Bi, Songfang Huang, Wenming Xiao, and Fei Huang. 2021. [E2E-VLP: End-to-end vision-language pre-training enhanced by visual learning](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 503–513, Online. Association for Computational Linguistics.
- Han Zhang, Jing Yu Koh, Jason Baldridge, Honglak Lee, and Yinfei Yang. 2021a. Cross-modal contrastive learning for text-to-image generation.
- Pengchuan Zhang, Xiujun Li, Xiaowei Hu, Jianwei Yang, Lei Zhang, Lijuan Wang, Yejin Choi, and Jianfeng Gao. 2021b. [Vinvl: Revisiting visual representations in vision-language models](#). In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5579–5588.
- Luowei Zhou, Hamid Palangi, Lei Zhang, Houdong Hu, Jason Corso, and Jianfeng Gao. 2020. [Unified vision-language pre-training for image captioning and vqa](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(07):13041–13049.
- Minfeng Zhu, Pingbo Pan, Wei Chen, and Yi Yang. 2019. [Dm-gan: Dynamic memory generative adversarial networks for text-to-image synthesis](#). In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.

A Additional Evaluation

We include 5 examples on COCO dataset for image captioning and text-to-image generation tasks. In Fig.5 and Fig.6, we find that DU-VLG generates captions and images of high quality.

B Human Evaluation Guideline

In human evaluation, each annotator is presented with 100 model generated images and 100 model generated captions from 3 systems (in random order). For text-to-image generation, the human judges are asked to evaluate on fidelity and informativeness on a scale of 1 to 5 (1 being good and 5 being poor). Here are descriptions of two aspects:

- **Fidelity:** Whether the image is realistic and looks like a real photo.

- **Relevance:** Whether the image provides necessary content coverage from the text description.

For image captioning, the human annotators are asked to evaluate on faithfulness and informativeness on a scale of 1 to 5 (1 being good and 5 being poor). Here are detailed descriptions of two aspects:

- **Faithfulness:** Whether the caption correctly describes main objects in the image.

- **Informativeness:** Whether the caption covers enough information from the image.

The definition of four aspects can be found in Tab.6.

Image Captioning	
Informativeness:	
1	Not relevant to the image.
3	Relevant, but misses the main objects of the image.
5	Successfully captures the main point of the image.
Faithfulness:	
1	The caption is full of fabricated content.
3	The caption is overall relevant to the image, but contains some fake details.
5	The caption matches with the image.
Text-to-Image Generation	
Fidelity:	
1	The image is unreal, distorted or blurred.
3	The image is overall realistic, but some details are blurred or distorted.
5	The image is vivid and looks like a real photo.
Relavance:	
1	The image does not match with the caption.
3	The image is related to the caption, but some details are hallucinated.
5	The image clearly reflects the caption.

Table 6: The definition of four aspects in human evaluation.

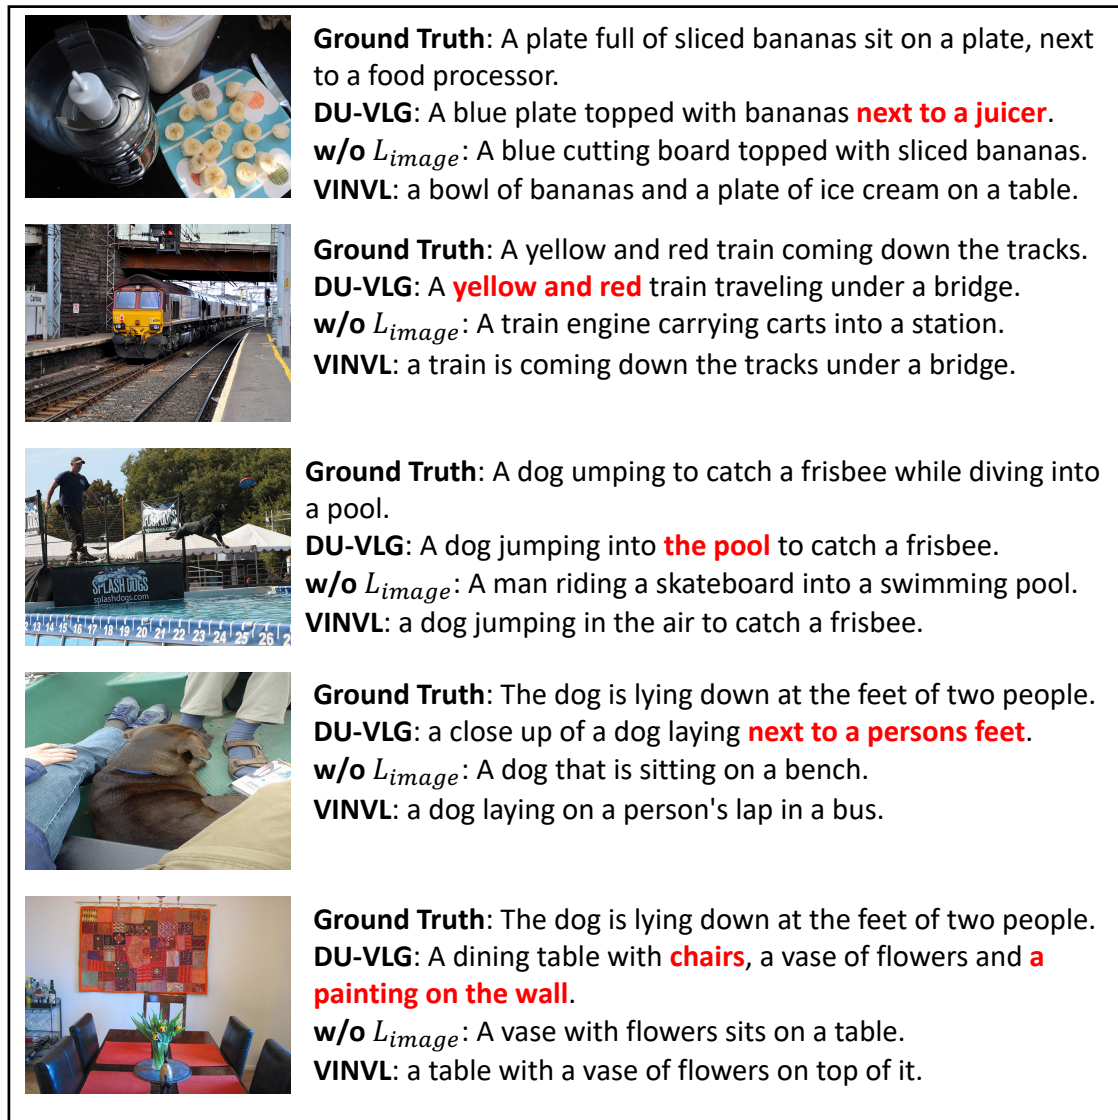


Figure 5: Samples on image captioning from COCO dataset. DU-VLG generates faithful and informative captions, highlighted in **red**.



Figure 6: Samples on text-to-image generation from COCO dataset. DU-VLG generates vivid and relevant images.