

# Large-Scale Bidirectional Training for Zero-Shot Image Captioning

Taehoon Kim<sup>\*†</sup>    Mark Marsden<sup>\*‡</sup>    Pyunghwan Ahn<sup>†</sup>    Sangyun Kim<sup>†</sup>  
Sihaeng Lee<sup>†</sup>    Alessandra Sala<sup>‡</sup>    Seung Hwan Kim<sup>†</sup>

LG AI Research<sup>†</sup>

Shutterstock<sup>‡</sup>

## Abstract

*When trained on large-scale datasets, image captioning models can understand the content of images from a general domain but often fail to generate accurate, detailed captions. To improve performance, pretraining-and-finetuning has been a key strategy for image captioning. However, we find that large-scale bidirectional training between image and text enables zero-shot image captioning. In this paper, we introduce Bidirectional Image Text Training in largER Scale, BITTERS, an efficient training and inference framework for zero-shot image captioning. We also propose a new evaluation benchmark which comprises of high quality datasets and an extensive set of metrics to properly evaluate zero-shot captioning accuracy and societal bias. We additionally provide an efficient finetuning approach for keyword extraction. We show that careful selection of large-scale training set and model architecture is the key to achieving zero-shot image captioning.*

## 1. Introduction

Trained with billions of image-text pairs, large-scale vision-language models (LVLMs) are outperforming previous state-of-the-art approaches in visual understanding [25, 30, 58, 65] and text-to-image generation [9, 10, 13, 39, 43, 44, 46, 63]. Recent advancements in text-to-image generation models now allow us to generate high-quality, unseen images from a single text prompt. Unlike text-to-image generation which focuses on zero-shot capability, image-to-text generation (image captioning) mainly relies on the pretraining-and-finetuning strategy. Even after pretraining on millions of image-text pairs [25, 58], previous captioning models are observed to lack zero-shot capability as they struggle to generalize to a given domain without finetuning. On the other hand, unified image-text training methods [27, 59] achieve notable results on various tasks.

Inspired by the bidirectional training strategy introduced in [27], we propose a large-scale training, inference, and evaluation framework for zero-shot image captioning. We find the main cause of poor zero-shot capability from the quality of training data. Texts in large-scale web-crawled datasets [5, 8, 48, 50, 55] have a wide variety of tones and manners. However, the *scale* is also what makes it difficult to filter out slang, inappropriate languages, or poor descriptions. To gain control over generated captions, we train our model with a new collection of 100 million image-text pairs, specifically curated for zero-shot image captioning. As shown in Figure 1, our **Bidirectional Image Text Training in largER Scale, BITTERS**, generates detailed captions over a diverse set of image categories and styles. Our contribution for the training and evaluation of a zero-shot image captioning model are as follows:

- We propose a 2D discrete wavelet transform (DWT) based cross-level feature augmentation method with a new architecture for AugVAE [27]. Our improved AugVAE (WaveVAE) shows enhanced image reconstruction performance on the ImageNet1K [7] validation set. With WaveVAE and other changes to the BiART [27] parameter configuration, BITTERS shows 32% less GPU memory usage and 18% faster sampling speed compared to L-Verse [27].
- We provide training (TIP100M) and evaluation (ICE-A and B) datasets for zero-shot image captioning. We also suggest various metrics to assess a given model’s zero-shot image captioning performance and societal bias. We show how the distribution of the training set affects the quality and bias of generated captions.
- We further introduce an adapter-based [23, 24, 33] finetuning approach for vision-language transformers. With mere 0.02% increase in number of parameters and sampling speed to BITTERS, we also enable keyword extraction from images.

Section 2 reviews previous works related to zero-shot image captioning. Section 3 introduces the BITTERS architecture as well as its training and sampling procedures. Sections 4 and 5 explains details of our datasets and evaluation metrics.

<sup>\*</sup>Equal contribution.

<sup>†</sup>{taehoon.kim, p.ahn, syun, sihaeng.lee, sh.kim}@lgresearch.ai

<sup>‡</sup>{mmarsden, asala}@shutterstock.com




Figure 1. Example captions generated with BITTERS. Our models shows zero-shot capability of image captioning by covering a wide variety of categories including photos, stock images, illustrations, vectors, as well as images of people from various demographic groups.

## 2. Related Work

In this section, we introduce large vision language models (LVLMs) for image-text generation tasks. Specifically, we introduce: (i) VQ-VAE for image encoding and decoding, (ii) LVLMs for image  $\leftrightarrow$  text generation, and (iii) a brief introduction on zero-shot image captioning. We also introduce more previous works related to BITTERS in our supplementary material.

### 2.1. Vector Quantized Variational Autoencoder

Vector quantized variational autoencoder, VQ-VAE [56], is widely used to encode and decode images to and from a series of quantized vectors. VQ-VAEs consist of an encoder, a decoder, and a vector quantizer with a visual codebook for learning discrete representations of images. As the vector quantizer factorizes the continuous representation of an image into a set of vectors within the limited possibilities of a visual codebook, information loss is inevitable. Reconstruction Fréchet Inception Distance (rFID) [11, 18] is widely used to measure the difference between original and reconstructed images.

From the assumption that the rFID of a VQ-VAE also affects the generation performance of a transformer it is attached to [11, 27, 44, 62], various methods have been proposed to improve rFID. Razavi *et al.* [45] and Kim *et al.* [27] focus on hierarchical feature representations to update visual codebook with diverse features. Gumbel-softmax relaxation [44] and  $L_2$  normalization [62] are proposed to improve codebook usage. Combinations of  $L_1$ ,  $L_2$ , logit-laplace [44], LPIPS [66], or adversarial [14] losses are proposed in [11, 27, 44, 62] to directly reduce rFID.

### 2.2. Image $\leftrightarrow$ Text Generation

In the early stage of text-to-image generation, a combination of a VQ-VAE and an auto-regressive transformer were widely used [9, 13, 44, 63] and showed promising results on zero-shot text-to-image generations. Recent works [39, 43, 47] further improve the generation quality by replacing the transformer with a denoising diffusion probabilistic model (DDPM) [20].

Unlike text-to-image, previous works [6, 25, 30, 38, 58, 65] on image-to-text generation (image captioning) focus on a pretraining-and-finetuning strategy. In accordance with LLMs, scaling up the pretraining data and the model size is important to build LVLMs. To control the generation result, the model is usually finetuned with a smaller dataset for each downstream-task.

To learn a cross-modal relationship between image and text, unified image-text training approaches [27, 59] have also been proposed. The bidirectional training strategy [27] shows data and parameter-efficient results compared to unidirectional models. OFA [59] unifies modalities (vision, language, multimodal) and tasks for pretraining.

### 2.3. Zero-Shot Image Captioning

Tewel *et al.* [54] proposes ZeroCap, a combination of pretrained multimodal (CLIP [41]) and language (GPT-2 [42]) models to enable zero-shot image captioning. Unlike previous works mentioned above which require additional finetuning, ZeroCap performs inference time gradient updates and optimization of GPT-2 to generate a caption that matches the given image. This requires a CLIP text encoder forward per each word of the caption. Concurrently to this work, Li *et al.* [29] decodes CLIP latents with a text-only trained language model for zero-shot image captioning.

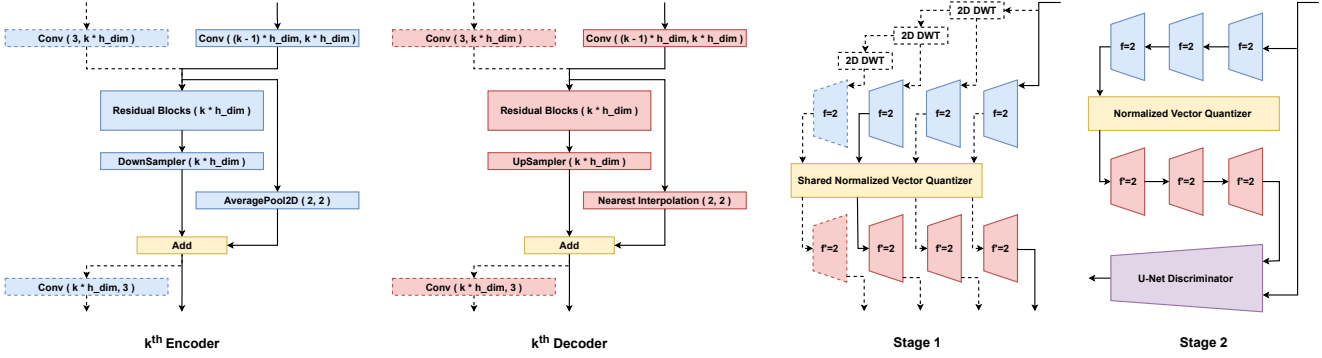


Figure 2. Proposed WaveVAE architecture. We remove unnecessary components (dotted line) and run short calibration before Stage 2.

### 3. Method

#### 3.1. Preliminary

Previous auto-regressive transformer [4] based methods follow a two-stage training procedure proposed by Ramesh *et al.* [44] for image-to-text [27, 58] or text-to-image [9, 10, 13, 27, 44] generation:

- **Stage 1:** Train a vector-quantized variational autoencoder (VQ-VAE) [11, 27, 28, 44, 45] to compress each RGB image into a sequence of image tokens with each element of  $d_Z$  possible values.
- **Stage 2:** Concatenate BPE-encoded text tokens and image tokens before training an auto-regressive transformer [4] to model the joint distribution over text and image tokens. According to the order of text and image tokens, the transformer learns to generate an image from given text or vice-versa.

For zero-shot image captioning, we bring the bidirectional image-text training concept of L-Verse [27] to a larger scale with 100 million image-text pairs. We also propose a new VQ-VAE, WaveVAE, which shows better image reconstruction quality with 75% fewer parameters compared to AugVAE [27].

#### 3.2. WaveVAE

According to Kim *et al.* [27], similar patterns in various patch sizes can appear throughout the training. Kim *et al.* [27] proposes a two stage training theme for a VQ-VAE which utilizes this cross-level patch similarity:

- **Stage 1:** Train a hierarchical AugVAE (AugVAE-ML) with cross-level feature augmentation. Four latent feature maps of different sizes are extracted to update a single vector quantizer.
- **Stage 2:** Remove unnecessary components from AugVAE-ML and finetune the model into a single-level AugVAE (AugVAE-SL) of  $32 \times 32$  latent map.

Following the notation from Kim *et al.* [27], we define the  $k^{th}$  encoder as  $z = E_k(x, f, d_{in}, d_{out})$ , where  $x$  is an  $n \times n \times d_{in}$  tensor,  $f$  is a downsampling factor, and  $z$  is an  $\frac{n}{f} \times \frac{n}{f} \times d_{out}$  tensor. The vector quantizer is  $z_q = VQ(z, d_Z)$ , where  $z$  is an  $n \times n \times d$  tensor with continuous  $d$ -size vectors and  $z_q$  is a quantized version of  $z$  with  $d_Z$  possible values. The  $k^{th}$  decoder is  $\hat{x} = G_k(\hat{z}, f', d'_{in}, d'_{out})$ , where  $\hat{z}$  is an  $n \times n \times d'_{in}$  tensor,  $f'$  is an upsampling factor, and  $\hat{x}$  is an  $n \times n \times d'_{out}$  tensor. We further improve cross-level feature augmentation method in both theoretical and architectural aspects.

While Kim *et al.* [27] directly use the output of  $k^{th}$  encoder as the input for  $k + 1^{th}$  encoder for pretraining (Stage 1), we utilize 2D discrete wavelet transform (DWT) [2] to generate an input for each encoder. The 2D DWT decomposes an image into three high-pass filtered images, each describing local changes in details, and one low-pass filtered image. The low-pass filtered image is a downscaled approximation of the original image which can replace the output of previous encoder. Specifically, our DWT applied AugVAE (WaveVAE) is trained as follows:

- **Stage 1:** Pair each  $E_k(2, 3, k * hidden\_dim)$  with  $G_k(2, k * hidden\_dim, 3)$ . For each pair, apply DWT decomposition  $k - 1$  times to generate input  $x_k$ . All pairs of encoder and decoder shares one  $VQ(8192)$  for cross-level feature augmentation.
- **Stage 2:** Integrate encoders into a single encoder  $E(8, 3, 3 * hidden\_dim)$  and decoders into a single decoder  $G(8, 3 * hidden\_dim, 3)$ . After removing redundant components, calibrate the single-level encoder - vector quantizer - decoder architecture for few iterations and finetune the model with various loss terms.

Along with DWT-based cross-level feature augmentation, we optimized the architecture of proposed WaveVAE for better image reconstruction quality with fewer parameters. We reduced the number of parameters to 25 million in total, a 75% reduction compared to AugVAE-SL.

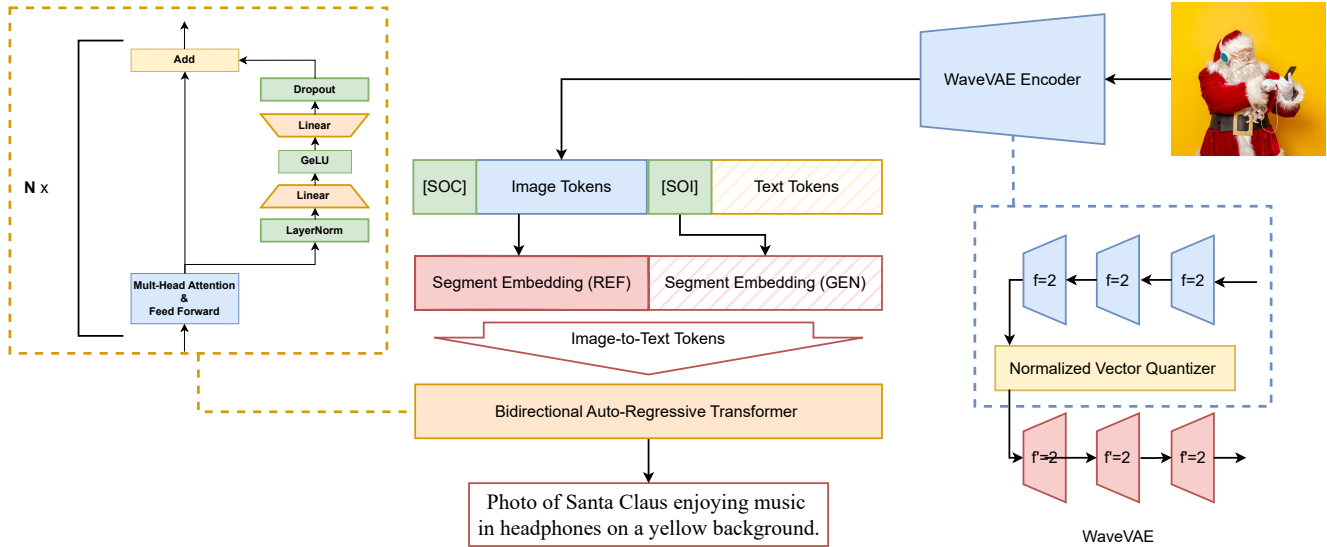


Figure 3. Proposed BITTERS architecture with the adapter for finetuning. Adapter is attached to each layer of Bidirectional Auto-Regressive Transformer (BiART) for parameter-efficient finetuning.

**Architecture** The architecture of the proposed WaveVAE is depicted in Figure 2. The encoder and decoder of WaveVAE are both bottleneck-style residual network similar to AugVAE [27]. Unlike AugVAE, we use Parametric Rectified Linear Unit (PReLU) [15] as the activation function. We also use PixelShuffle [51] for upsampling. We use encoders and decoders with 8 residual blocks and the hidden dimension (*hidden\_dim*) of 64. For the vector quantizer, we use a visual codebook with the embedding dimension of 64 and the codebook size of 8192. We do not use an exponential moving average (EMA) vector quantizer [45] as in [27]. Instead, we use an  $L_2$  normalized vector quantizer as proposed in [62]. Precise details including the source code are provided in our supplementary material.

### 3.3. BiART

**Architecture** Along with WaveVAE, we also modify the parameter setting of BiART [27] according to the scaling law proposed in [21]. BITTERS uses a 24-layer BiART [27] with 1280 dimensional states and 10 masked self-attention heads. We use 64 BPE-encoded [49] text tokens with 49408 possibilities and 1024 encoded image tokens with 8192 possibilities. BITTERS has 650 million parameters in total. Compared to L-Verse [27], BITTERS shows 32% less memory usage and 18% faster inference speed on a single NVIDIA A100 GPU. While BiART is initially designed for bidirectional training between image and text, we adapt the bidirectional training concept of BiART *only for stable training*. As stated in Kim *et al.* [27], bidirectional training relieves the heterogeneity between image and text, and prevents overflow of gradients. We also tried to train a model only in image-to-text direction but the model diverged.

**Parameter-Efficient Finetuning** Unlike previous methods [6, 25, 30, 58, 65], we don't update all weights during finetuning as this risks modifying the learned joint distribution between image and text. Among the various finetuning approaches [23, 24, 33–35] for transformers, we take an adapter-based approach [23, 24, 33]. As shown in Figure 3, we attach an adapter with bottleneck dimension of 320 to each layer of BiART. During finetuning, we only update adapters with hyperparameters used for pretraining. Using adapters brings only 0.02% increase in number of parameters and latency.

### 3.4. Sampling

**Image Captioning** As we focus on zero-shot image captioning, we modify the text sampling process proposed in [27] to delicately control the generated caption. We sample 32 text tokens with pretrained BITTERS model to generate a caption for each image. For each token selection, we first select 10% of logits with the highest probabilities (*top-k* sampling) [12] and apply *top-p* sampling [22] with  $p = 0.95$ . We sample 64 captions in total and calculate CLIPScore [17] to select a Top-1 caption.

**Keyword Extraction** To extract keywords from image, we sample 48 text tokens with finetuned BITTERS model to generate a list of keywords for each image. We sample 64 lists of keywords in total and calculate CLIPScore [17] to select a Top-8 lists of keywords. From Top-8 lists of keywords, we further select keywords that appeared 3 or more times in the Top-8 lists. Other details are identical to the sampling method for image captioning.

## 4. Dataset

In this section, we briefly introduce our newly curated training and evaluation sets for zero-shot image captioning. More details including license and level of publicity for each dataset are provided in our supplementary material.

### 4.1. Training

As stated in [27, 58], it is difficult to control the generated caption from a model trained with a large number of web-harvested image-text pairs. This is due to the wide variety of text styles, inappropriate language, and intrinsic biases present. For this reason, we train BITTERS using a new, quality-controlled 100 million image dataset which we will refer to as Text Image Pairs 100 Million (TIP100M).

All images are random sampled from Shutterstock’s<sup>1</sup> image catalog. Images are 500px on the longest side with a single ground truth caption provided. A list of keywords is also included for each image to enable keyword extraction model training. All captions and keywords are in English and were moderated for hate speech, slurs, and expletives.

This dataset contains a very broad set of concepts and scenarios (indoor and outdoor, with or without people, with or without animals). The majority of images are photographs (69%) while the rest are either illustrations or vector graphics (31% combined). Roughly 1/3 (24/69%) of included photographs contain at least one person. For photographs containing people: 75% are Caucasian, 53% are 20-30 years of age and 67% are female. 10% of images are editorial use content (containing logos, celebrities, news content).

### 4.2. Evaluation

Evaluation is carried out using two newly curated image datasets. These collections are produced to better assess the true limits of zero-shot image captioning models by identifying where a given model succeeds or fails to produce accurate and fair image captions. We refer to each dataset as Image Captioning Evaluation Accuracy (ICE-A) and Image Captioning Evaluation Bias (ICE-B) respectively.

**ICE-A** This collection consists of 17 image categories from 4 high-level groups (People, Animals, Objects, Other). There is a varied set of 300 curated images per category, 5100 images in total. To avoid redundancy, the full set of category names are outlined in Table 4. All image categories consist of photos unless otherwise stated (i.e. vector graphics, illustrations). In this context, ‘Stocky Setting’ refers to photographs taken in a very controlled setting such as a studio. In comparison, ‘Authentic’ refers to photographs taken in a more natural, candid setting, typically with natural lighting.

<sup>1</sup>[www.shutterstock.com](http://www.shutterstock.com)

**ICE-B** This collection was generated to expand upon the ‘Various Demographics’ category in ICE-A, allowing for an in depth assessment of societal bias in a manner similar to [19]. During the creation of this dataset, the following constraints were enforced: (i) each image must be a photograph containing exactly one person and (ii) demographic metadata (gender, ethnicity) must be available for each image. Enforcing both constraints allows us to directly assess societal biases and split the dataset into demographic subgroups. The dataset consists of 14K female subject images and 8K of male subjects<sup>2</sup>. The following ethnicity labels breakdown is observed: Southeast Asian (5686), Caucasian (4782), Hispanic (3965), African American (3878), East Asian (2136), South Asian (1240), Middle Eastern (704)<sup>3</sup>. Beyond diversity of human subjects, this dataset also contains a variety of scenarios and concepts including indoor, outdoor, facial masks being worn, hand gestures, medical settings, exercising, workplace, and family home.

## 5. Metrics

In this section, we briefly introduce quantitative measures that we use to assess the zero-shot image captioning performance of our model. More details on each evaluation metric are provided in our supplementary material.

**Caption Accuracy** We evaluate caption accuracy using a set of commonly employed metrics: SPICE [1], CIDEr [57], METEOR [3], ROGUE [31] and BLEU-4 [40]. For qualitative assessment, we also conduct a human evaluation on BITTERS generated captions along with generations from other models.

**Bias Assessment** To prevent potential negative impacts of zero-shot image captioning, we also put emphasis on detecting gender or racial prejudice in generated captions. We assess the societal bias of a given model using three distinct approaches: Gender Error and Term Ratio [16], VADER Sentiment Score [67], and Leakage for Image Captioning (LIC) [19].

**Keyword Extraction** We evaluate the keyword extraction performance of BITTERS on image-keywords ground truth pairs from ICE-A using the following two metrics: Normalized Keyword Overlap and CLIP Cosine Similarity. To show the relevance between image captioning and keyword extraction, we use the same set of example images for generated captions and extracted keywords in our supplementary material.

<sup>2</sup>We use a binary simplification of gender. We acknowledge that this is not inclusive or representative and should be addressed in future work.

<sup>3</sup>We use a fixed set of 7 race/ethnicity groups. We acknowledge that this set of groups is far from comprehensive and that these attributes are much more complex in reality.



Figure 4. Examples of input images (*top*) and reconstructions from WaveVAE (*bottom*). The resolution of each image is  $256 \times 256$ .

Model	Training Data	rFID
AugVAE-SL [27]	ImageNet1K	3.28
WaveVAE	ImageNet1K	<b>2.70</b>
AugVAE-SL [27]	TIP100M	3.94
WaveVAE	TIP100M	<b>2.35</b>

Table 1. Reconstruction Fréchet Inception Distance (rFID) on ImageNet1K validation set. Our WaveVAE shows better (lower) rFID in all cases compared to AugVAE-SL with 75% less parameters.

Metric	MS-COCO		ICE-A	
	ZeroCap [54]	BITTERS	ZeroCap [54]	BITTERS
BLEU-4	2.6	<b>6.1</b>	0.24	<b>2.9</b>
METEOR	11.5	<b>13.7</b>	3.5	<b>11.7</b>
ROUGE	-	27.8	8.1	<b>18.3</b>
CIDEr	14.6	<b>31.2</b>	5.7	<b>35.3</b>
SPICE	5.5	<b>9.5</b>	4.8	<b>13.2</b>

Table 2. Zero-shot image captioning accuracy of ZeroCap and BITTERS on MS-COCO Captions *karpathy test split* and ICE-A.

## 6. Experiment

### 6.1. Image Reconstruction

From the result in Table 1, WaveVAE shows better (lower) reconstruction Fréchet Inception Distance [18] (rFID) on the ImageNet1K [7] validation set when trained in every case with different training datasets. It is also notable that WaveVAE shows improvement in rFID when trained with a larger dataset (TIP100M), while AugVAE-SL shows performance degradation. This suggests that WaveVAE is more suitable than AugVAE-SL for large-scale training.

### 6.2. Zero-Shot Image Captioning

**Result on Public Benchmark** We first compare BITTERS with ZeroCap [54], current state-of-the-art in zero-shot image captioning, on MS-COCO Captions [32] *karpathy test split* and ICE-A. From the results in Table 2, BITTERS outperforms ZeroCap on MS-COCO Captions in all metrics. Due to gradient update and optimization over the context cache, ZeroCap shows 30 times slower inference speed compared to BITTERS. From the results in Tables 2 and 3, ZeroCap shows poor result on ICE-A even compared to ClipCap [38] and GIT [58]. ZeroCap fails to generate proper captions for images in categories such as vectors, stylized, illustrations, and medical equipment, *which is not included in MS-COCO*. This demonstrates the importance of a new evaluation set with diverse image categories (ICE-A) to assess a given model’s zero-shot image captioning performance.

**Experimental Setup** We perform quantitative evaluation on four model architectures. Two model architectures trained on TIP100M, L-Verse [27] and the proposed BITTERS, are set as an experimental group. We also include ClipCap [38] and GIT (BASE\_MODEL) [58] as a control group. These models are chosen for comparison as they are trained on large-scale datasets and perform well on other captioning benchmarks.

**Caption Accuracy** Table 3 shows the overall caption accuracy metrics for ICE-A. L-Verse performs best in terms of BLEU-4, CIDEr and SPICE. We also use SPICE score to compare caption accuracy over different categories. In Table 4, L-Verse performs best overall (highest SPICE score in 9/17 categories) with BITTERS a close second. Regarding computational efficiency as mentioned in Section 3.3, BITTERS shows comparable performance to L-Verse in all metrics. We find that the distribution of contents in a training set highly influences generated captions. While GIT performs noticeably better on ‘Food’, ClipCap performs best on several categories in ‘Animals’ group. On the other hand, L-Verse and BITTERS lead on ‘Vector Graphics’, ‘Illustrations’, and ‘Stylized’ by a large margin. Some content is likely out of domain for ClipCap and GIT (i.e. ‘Stylized’) hence the large difference in SPICE score.

For ICE-B, SPICE score per gender group is in Table 5 and SPICE score per ethnicity group is in Table 6. L-Verse and BITTERS again outperform the approaches in the control group. While the internal performance gap between

Metric	ClipCap [38]	GIT [58]	L-Verse [27]	BITTERS
BLEU-4	0.8	0.7	<b>3.6</b>	2.9
METEOR	7.2	5.4	<b>11.7</b>	<b>11.7</b>
ROUGE	14.2	12.1	18.1	<b>18.3</b>
CIDEr	21.4	19.9	<b>37.8</b>	35.3
SPICE	11.9	10.5	<b>13.8</b>	13.2

Table 3. Zero-shot image captioning accuracy on ICE-A.

Category	ClipCap [38]	GIT [58]	L-Verse [27]	BITTERS
P: AuI	<b>13.6</b>	7.6	13.2	12.3
P: AuO	<b>13.4</b>	13.5	11.2	13.1
P: StS	11.9	13.6	16.3	<b>20.1</b>
P: VaD	12.7	9.1	<b>16.5</b>	13.1
A: AuI	<b>17.2</b>	10.3	10.8	14.6
A: AuO	9.9	10.9	<b>12.2</b>	11.5
A: StS	<b>22.2</b>	17.4	21.2	21.3
Ob: CoS	10.5	8.2	<b>13.2</b>	12.3
Ob: CoU	14.1	13.2	<b>16.3</b>	14.9
Ob: Ele	10.7	10.0	<b>14.4</b>	11.8
Ob: Food	13.7	<b>20.2</b>	18.8	15.3
Ob: Fur	12.5	9.4	10.9	<b>13.2</b>
Ob: MeE	<b>8.6</b>	4.9	8.2	7.8
Ob: RdS	<b>10.4</b>	9.1	8.9	8.3
Ot: Ill	8.2	8.6	<b>15.1</b>	10.9
Ot: Sty	5.8	5.5	<b>17.8</b>	16.8
Ot: VeG	7.0	6.5	<b>12.1</b>	10.3

\* **P**: People **A**: Animals **Ob**: Objects **Ot**: Other

\*\* **AuI**: Authentic Indoors **AuO**: Authentic Outdoors **StS**: Stocky Setting **VaD**: Various Demographics **CoS**: Construction Site **CoU**: Cooking Utensils **Ele**: Electronics **Fur**: Furniture **MeE**: Medical Equipment **RdS**: Road Signs **Ill**: Illustrations **Sty**: Stylized **VeG**: Vector Graphics

Table 4. SPICE score per image category on ICE-A.

gender groups is minimal for all models, there is a much larger variation between ethnicity groups. It is also observed that models which perform better in terms of overall SPICE score (L-Verse, BITTERS) have much larger variation in SPICE score between ethnicity groups.

**Human Evaluation** We further conduct a human evaluation over generated captions on ICE-A. We use a web-based human evaluation tool as shown in our supplementary material. 100 anonymous evaluators participated in this research. For each participant, we provided 20 images randomly sampled from ICE-A with a caption generated from each of the four models (ClipCap, GIT, L-Verse, BITTERS). We asked them to choose the most appropriate caption for each image. We also allowed each participant to choose “None of the captions well describes the image.” Evaluation results are provided in Table 7. With the question “Which caption best describes the given image?” asked, BITTERS generated captions received **29.9%** of the votes among other

Group	ClipCap [38]	GIT [58]	L-Verse [27]	BITTERS
Male	12.2	9.8	<b>16.3</b>	15.4
Female	12.5	10.4	<b>16.9</b>	16.1
All	12.3	10.2	<b>16.7</b>	15.9

Table 5. SPICE score per gender group on ICE-B.

Group	ClipCap [38]	GIT [58]	L-Verse [27]	BITTERS
AF-AM	11.6	9.4	<b>18.6</b>	17.2
C	12.5	10.5	<b>14.6</b>	13.9
E-A	12.1	9.6	12.7	<b>13.0</b>
H	11.7	10.9	<b>22.8</b>	20.8
M-E	11.3	9.2	13.5	<b>13.5</b>
S-A	17.1	13.4	<b>27.9</b>	24.4
SE-A	12.4	9.7	12.4	<b>12.9</b>
All	12.3	10.2	<b>16.7</b>	15.9

\* **AF-AM**: African American **C**: Caucasian **E-A**: East Asian **H**: Hispanic **M-E**: Middle Eastern **S-A**: South Asian **SE-A**: Southeast Asian

Table 6. SPICE score per ethnicity group on ICE-B.

Percentage (%)	ClipCap [38]	GIT [58]	L-Verse [27]	BITTERS
w/ None	18.2	16.9	20.7	<b>29.9</b>
w/o None	21.2	19.7	24.2	<b>34.9</b>

\* **w/ None**: Percentage including votes for “None of the captions well describes the image” (None, 14.3%) **w/o None**: Percentage without votes for None.

Table 7. Human evaluation results on ICE-A. 100 anonymous participants are asked the question “Which caption best describes the given image?” for 20 random images sampled from ICE-A.

choices including “None of the captions well describes the image.” (14.3%). BITTERS generated captions also had higher preference (**34.9%**) by a large margin compared to the other three models.

**Bias Assessment** Table 8 presents gender error and ratio results on ICE-B. GIT has the lowest gender error rate while the proposed BITTERS has the highest. All four models have a gender term ratio similar to the true ratio of female to male subjects (1.75). While GIT and ClipCap generated captions are less prone to gender error, they are also observed to be shorter and less semantically accurate (i.e. SPICE score). There appears to be an accuracy-bias trade-off for image captioning models. As shown in [19], other image captioning models also have similar gender error rates of 2-4%. Overall results suggest that gender bias is a common issue in image captioning.

Neutral sentiment rate (%) per ethnicity group and gender group are shown in Table 9 and Table 10. Models trained on TIP100M (L-Verse, BITTERS) contain more sentimental terms in their captions compared to ClipCap

Metric	ClipCap [38]	GIT [58]	L-Verse [27]	BITTERS
Error (%) ↓	3.1	<b>1.8</b>	4.7	4.8
Term Ratio	2.01	1.83	1.82	1.88

Table 8. Gender error (%) and term ratio on ICE-B. Gender term ratio for ground truth captions is 1.75.

Group	ClipCap [19]	GIT [58]	L-Verse [27]	BITTERS
AF-AM	67.6	89.7	15.2	14.8
C	75.5	91.4	29.3	26.5
E-A	82.1	91.8	27.9	28.3
H	59.7	85.8	10.7	10.9
M-E	75.8	95.5	21.4	22.4
S-A	69.3	95.5	20.2	20.9
SE-A	83.8	91.5	32.4	29.9

\* **AF-AM**: African American **C**: Caucasian **E-A**: East Asian  
**H**: Hispanic **M-E**: Middle Eastern **S-A**: South Asian **SE-A**: Southeast Asian

Table 9. Neutral sentiment rate (%) per ethnicity group on ICE-B.

Group	ClipCap [38]	GIT [58]	L-Verse [27]	BITTERS
Male	77.6	91.0	26.4	25.7
Female	71.5	90.2	21.8	20.3

Table 10. Neutral sentiment rate (%) per gender group on ICE-B.

Metric	ClipCap [38]	GIT [58]	L-Verse [27]	BITTERS
LIC	<b>44.9</b>	45.3	49.9	52.0

Table 11. Gender bias LIC scores on ICE-B. LIC score for ground truth captions is 49.2.

and GIT, which both show a much higher neutral sentiment rate. All models tend to generate more sentimentally neutral captions for images of male subjects, while they use more emotional languages for female subjects. Images of Hispanic subjects have the lowest neutral sentiment rate for all captioning models. The use of sentimental terms (i.e. ‘sad’, ‘angry’, ‘happy’) in image captions is not inherently good or bad as this language is often required to properly describe the image contents. However, large differences between gender or ethnicity groups are undesirable (i.e. showing strong tendency to use emotive language to describe certain gender or ethnicity of a person being photographed).

Table 11 presents gender bias LIC scores for all models on ICE-B. The LIC score for the human-provided ground truth captions is 49.2. ClipCap performs best in terms of LIC. Both ClipCap and GIT captions are less biased than human-provided captions. This result suggests that certain language is being used for gender groups (i.e. terms such as ‘attractive’, ‘young’, ‘beautiful’ only being used for images of female subjects). Once again the models (L-Verse,

Metric	BITTERS
Normalized Keyword Overlap (%) ↑	35.7
CLIP Cosine Similarity (%) ↑	37.6

Table 12. Keyword extraction performance on ICE-A.

BITTERS) with more accurate captions semantically contain more gender bias. LIC scores in this range (40-50) have been observed for other captioning models [19], suggesting that most of image captioning approaches show comparable levels of gender bias.

### 6.3. Zero-Shot Keyword Extraction

Table 12 presents keyword extraction results on ICE-A. As described in Section 3.4, we currently sample 48 tokens from each image. Our BITTERS extracts a median of 15 keywords per image while the ground truth median is 41. Despite the lack of overlap with the ground truth, the model extracted keywords are observed to be highly relevant and successfully describe the contents of the image. We provide examples of extracted keywords along with generated captions in our supplementary material.

## 7. Conclusion

In this paper, we present a novel model architecture and a set of evaluation data and metrics for zero-shot image captioning. As we bring the bidirectional image-text training strategy [27] to large-scale, we find that a well-curated dataset of image-text pairs enables zero-shot image captioning. BITTERS is parameter-efficient architecture specially designed for zero-shot image captioning and keyword extraction. The proposed WaveVAE uses 2D DWT [2] for cross-level feature augmentation [27]. Our WaveVAE shows improved image reconstruction performance and adaptability to large-scale training compared to AugVAE. [27]. We further assess the accuracy and bias of zero-shot generated captions with BITTERS in both a quantitative and qualitative manner.

## 8. Discussion

**Limitation** Future work in this area will need to address the trade-off between societal bias and semantic accuracy. This can be accomplished through training set audits to mitigate bias and improve representation. The development of improved bias detection and mitigation techniques for caption generation will also be required. Even when trained in a bidirectional manner, our model lacks text-to-image generation capability compared to other state-of-the-art models. We provide further discussion on *zero-shot* and *large-scale bidirectional training* in our supplementary material.



## References

- [1] Peter Anderson, Basura Fernando, Mark Johnson, and Stephen Gould. Spice: Semantic propositional image caption evaluation. In *Proceedings of the European Conference on Computer Vision*, 2016. 5
- [2] M. Antonini, M. Barlaud, P. Mathieu, and I. Daubechies. Image coding using wavelet transform. *IEEE Transactions on Image Processing*, 1992. 3, 8
- [3] Satantjeev Banerjee and Alon Lavie. Meteor: An automatic metric for mt evaluation with improved correlation with human judgments. In *Proceedings of the acl workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*, 2005. 5
- [4] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. In *Advances in Neural Information Processing Systems*, 2020. 3
- [5] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021. 1
- [6] Jaemin Cho, Jie Lei, Hao Tan, and Mohit Bansal. Unifying vision-and-language tasks via text generation. In *Proceedings of the International Conference on Machine Learning*, 2021. 2, 4
- [7] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2009. 1, 6, 12
- [8] Karan Desai, Gaurav Kaul, Zubin Aysola, and Justin Johnson. RedCaps: Web-curated image-text data created by the people, for the people. In *NeurIPS Datasets and Benchmarks*, 2021. 1
- [9] Ming Ding, Zhuoyi Yang, Wenyi Hong, Wendi Zheng, Chang Zhou, Da Yin, Junyang Lin, Xu Zou, Zhou Shao, Hongxia Yang, and Jie Tang. Cogview: Mastering text-to-image generation via transformers. In *Advances in Neural Information Processing Systems*, 2021. 1, 2, 3, 12, 14
- [10] Ming Ding, Wendi Zheng, Wenyi Hong, and Jie Tang. Cogview2: Faster and better text-to-image generation via hierarchical transformers, 2022. 1, 3, 12, 14
- [11] Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming transformers for high-resolution image synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021. 2, 3, 12
- [12] Angela Fan, Mike Lewis, and Yann Dauphin. Hierarchical neural story generation. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2018. 4, 14
- [13] 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, 2022. 1, 2, 3, 12, 14
- [14] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, 2014. 2
- [15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of International Conference on Computer Vision*, 2015. 4
- [16] Lisa Anne Hendricks, Kaylee Burns, Kate Saenko, Trevor Darrell, and Anna Rohrbach. Women also snowboard: Overcoming bias in captioning models. In *Proceedings of the European Conference on Computer Vision*, 2018. 5, 13
- [17] Jack Hessel, Ari Holtzman, Maxwell Forbes, Ronan Le Bras, and Yejin Choi. CLIPScore: A reference-free evaluation metric for image captioning. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 2021. 4, 14
- [18] 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 *Advances in Neural Information Processing Systems*, 2017. 2, 6
- [19] Yusuke Hirota, Yuta Nakashima, and Noa Garcia. Quantifying societal bias amplification in image captioning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. 5, 7, 8, 13
- [20] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems*, 2020. 2
- [21] Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, Tom Hennigan, Eric Noland, Katie Millican, George van den Driessche, Bogdan Damoc, Aurelia Guy, Simon Osindero, Karen Simonyan, Erich Elsen, Jack W. Rae, Oriol Vinyals, and Laurent Sifre. Training compute-optimal large language models, 2022. 4
- [22] Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. The curious case of neural text degeneration. In *Proceedings of the International Conference on Learning Representations*, 2020. 4, 14
- [23] Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin de Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. Parameter-efficient transfer learning for NLP. In *Proceedings of the International Conference on Machine Learning*, 2019. 1, 4, 12
- [24] Edward Hu, Yelong Shen, Phil Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. In *Proceedings of the International Conference on Learning Representations*, 2022. 1, 4, 12
- [25] Xiaowei Hu, Zhe Gan, Jianfeng Wang, Zhengyuan Yang, Zicheng Liu, Yumao Lu, and Lijuan Wang. Scaling up vision-language pre-training for image captioning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. 1, 2, 4

- [26] Clayton Hutto and Eric Gilbert. Vader: A parsimonious rule-based model for sentiment analysis of social media text. In *Proceedings of the international AAAI conference on web and social media*, 2014. [13](#)
- [27] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [7](#), [8](#), [12](#), [14](#)
- [28] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. In *Proceedings of the International Conference on Learning Representations*, 2014. [3](#)
- [29] Wei Li, Linchao Zhu, Longyin Wen, and Yi Yang. Decap: Decoding CLIP latents for zero-shot captioning via text-only training. In *International Conference on Learning Representations*, 2023. [2](#)
- [30] 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. Oscar: Object-semantics aligned pre-training for vision-language tasks. In *Proceedings of the European Conference on Computer Vision*, 2020. [1](#), [2](#), [4](#)
- [31] Chin-Yew Lin. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, 2004. [5](#)
- [32] 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. In *Proceedings of the European Conference on Computer Vision*, 2014. [6](#)
- [33] Zhaojiang Lin, Andrea Madotto, and Pascale Fung. Exploring versatile generative language model via parameter-efficient transfer learning. In *Findings of the Association for Computational Linguistics: EMNLP*, 2020. [1](#), [4](#), [12](#)
- [34] Xiao Liu, Kaixuan Ji, Yicheng Fu, Zhengxiao Du, Zhilin Yang, and Jie Tang. P-tuning v2: Prompt tuning can be comparable to fine-tuning universally across scales and tasks, 2021. [12](#)
- [35] Xiao Liu, Yanan Zheng, Zhengxiao Du, Ming Ding, Yujie Qian, Zhilin Yang, and Jie Tang. Gpt understands, too, 2021. [4](#), [12](#)
- [36] Ilya Loshchilov and Frank Hutter. Fixing weight decay regularization in adam. In *Proceedings of the International Conference on Learning Representations*, 2018. [12](#)
- [37] Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization for generative adversarial networks. In *Proceedings of International Conference on Learning Representations*, 2018. [12](#)
- [38] Ron Mokady, Amir Hertz, and Amit H Bermano. Clipcap: Clip prefix for image captioning, 2021. [2](#), [6](#), [7](#), [8](#)
- [39] 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, 2021. [1](#), [2](#), [12](#), [14](#)
- [40] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, 2002. [5](#)
- [41] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *Proceedings of the International Conference on Machine Learning*, 2021. [2](#), [12](#), [14](#)
- [42] Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners, 2019. [2](#)
- [43] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents, 2022. [1](#), [2](#), [12](#), [14](#)
- [44] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *Proceedings of the International Conference on Machine Learning*, 2021. [1](#), [2](#), [3](#), [12](#), [14](#)
- [45] Ali Razavi, Aaron van den Oord, and Oriol Vinyals. Generating diverse high-fidelity images with vq-vae-2. In *Advances in Neural Information Processing Systems*, 2019. [2](#), [3](#), [4](#)
- [46] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models, 2021. [1](#)
- [47] 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, 2022. [2](#), [12](#), [14](#)
- [48] 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, 2021. [1](#)
- [49] Rico Sennrich, Barry Haddow, and Alexandra Birch. Neural machine translation of rare words with subword units. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, 2016. [4](#)
- [50] Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, 2018. [1](#), [14](#)
- [51] Wenzhe Shi, Jose Caballero, Ferenc Huszár, Johannes Totz, Andrew P. Aitken, Rob Bishop, Daniel Rueckert, and Zehan Wang. Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network, 2016. [4](#)
- [52] Leslie N. Smith and Nicholay Topin. Super-convergence: Very fast training of residual networks using large learning rates, 2017. [12](#)
- [53] Ming Tao, Hao Tang, Songsong Wu, Nicu Sebe, Xiao-Yuan Jing, Fei Wu, and Bingkun Bao. Df-gan: Deep fusion generative adversarial networks for text-to-image synthesis, 2021. [14](#)
- [54] Yoav Tewel, Yoav Shalev, Idan Schwartz, and Lior Wolf. Zerocap: Zero-shot image-to-text generation for visual-

- semantic arithmetic. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 17918–17928, June 2022. [2](#), [6](#)
- [55] Bart Thomee, David A. Shamma, Gerald Friedland, Benjamin Elizalde, Karl Ni, Douglas Poland, Damian Borth, and Li-Jia Li. YFCC100m. *Communications of the ACM*, 2016. [1](#)
- [56] Aaron van den Oord, Oriol Vinyals, and Koray Kavukcuoglu. Neural discrete representation learning. In *Advances in Neural Information Processing Systems*, 2017. [2](#)
- [57] Ramakrishna Vedantam, C. Lawrence Zitnick, and Devi Parikh. Cider: Consensus-based image description evaluation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2015. [5](#)
- [58] 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, 2022. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [7](#), [8](#)
- [59] 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 *Proceedings of the International Conference on Machine Learning*, 2022. [1](#), [2](#)
- [60] Xintao Wang, Liangbin Xie, Chao Dong, and Ying Shan. Real-esrgan: Training real-world blind super-resolution with pure synthetic data. In *International Conference on Computer Vision Workshops (ICCVW)*, 2021. [12](#)
- [61] Tao Xu, Pengchuan Zhang, Qiuyuan Huang, Han Zhang, Zhe Gan, Xiaolei Huang, and Xiaodong He. Attngan: Fine-grained text to image generation with attentional generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018. [14](#)
- [62] Jiahui Yu, Xin Li, Jing Yu Koh, Han Zhang, Ruoming Pang, James Qin, Alexander Ku, Yuanzhong Xu, Jason Baldridge, and Yonghui Wu. Vector-quantized image modeling with improved VQGAN. In *Proceedings of the International Conference on Learning Representations*, 2022. [2](#), [4](#)
- [63] 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, 2022. [1](#), [2](#), [12](#), [14](#)
- [64] Han Zhang, Jing Yu Koh, Jason Baldridge, Honglak Lee, and Yinfei Yang. Cross-modal contrastive learning for text-to-image generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021. [14](#)
- [65] 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. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021. [1](#), [2](#), [4](#)
- [66] Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018. [2](#), [12](#)
- [67] Dora Zhao, Angelina Wang, and Olga Russakovsky. Understanding and evaluating racial biases in image captioning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021. [5](#), [13](#)
- [68] Minfeng Zhu, Pingbo Pan, Wei Chen, and Yi Yang. 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*, 2019. [14](#)