

ZeroNLG: Aligning and Autoencoding Domains for Zero-Shot Multimodal and Multilingual Natural Language Generation

Bang Yang*, Fenglin Liu*, Yuexian Zou, Xian Wu, Yaowei Wang, and David A. Clifton

Abstract—Natural Language Generation (NLG) accepts input data in the form of images, videos, or text and generates corresponding natural language text as output. Existing NLG methods mainly adopt a supervised approach and rely heavily on coupled data-to-text pairs. However, for many targeted scenarios and for non-English languages, sufficient quantities of labeled data are often not available. As a result, it is necessary to collect and label data-text pairs for training, which is both costly and time-consuming. To relax the dependency on labeled data of downstream tasks, we propose an intuitive and effective zero-shot learning framework, ZeroNLG, which can deal with multiple NLG tasks, including image-to-text (image captioning), video-to-text (video captioning), and text-to-text (neural machine translation), across English, Chinese, German, and French within a unified framework. ZeroNLG does not require any labeled downstream pairs for training. During training, ZeroNLG (i) projects different domains (across modalities and languages) to corresponding coordinates in a shared common latent space; (ii) bridges different domains by aligning their corresponding coordinates in this space; and (iii) builds an unsupervised multilingual auto-encoder to learn to generate text by reconstructing the input text given its coordinate in shared latent space. Consequently, during inference, based on the data-to-text pipeline, ZeroNLG can generate target sentences across different languages given the coordinate of input data in the common space. Within this unified framework, given visual (imaging or video) data as input, ZeroNLG can perform zero-shot visual captioning; given textual sentences as input, ZeroNLG can perform zero-shot machine translation. We present the results of extensive experiments on twelve NLG tasks, showing that, without using any labeled downstream pairs for training, ZeroNLG generates high-quality and “believable” outputs and significantly outperforms existing zero-shot methods. Our code and data are available at <https://github.com/yangbang18/ZeroNLG>.

Index Terms—Zero-shot Learning, Natural Language Generation, Multimodal Language Generation, Multilingual Language Generation, Visual Captioning, Neural Machine Translation.

1 INTRODUCTION

NATURAL language generation (NLG), also known as the data-to-text generation, aims to comprehend the content of provided data, which may come in various forms such as text [1], image [2], and video [3], and produce coherent text in natural language automatically. NLG has a wide range of applications, including image and video captioning and machine translation. Due to its broad usage scenarios, NLG has been receiving extensive research interests [4], [5], [6], [7]. Existing NLG approaches usually adopt an encoder-decoder framework [1], where the encoder calculates vector representations for the input data, and the decoder employs RNNs [8] or Transformers [9] to generate the target sentences using the encoded representation. Such approaches have demonstrated state-of-the-art performance in various

natural language generation tasks [6], [10], [11], [12].

Most existing encoder-decoder-based approaches are purely data-driven and their performance is heavily reliant on the volume and quality of available labeled data-text pairs. However the acquisition of paired data can be time-consuming and costly in real-world situations, and its scarcity can prohibit the scale of models; with less restrictive approaches, we could dramatically increase the size of available data used to train, and thereby substantially enlarge model scalability to these increased dataset sizes. Although numerous datasets of data-English pairs have been made available publicly, the availability of data-text pairs in non-English languages is often relatively uncommon or may even be unavailable. Consequently, the lack of labeled training data poses a significant challenge to developing NLG models for non-English languages - contributing to the highly uneven representation of less commonly-employed languages, which can in itself be a barrier to “fair AI” that would be usable by otherwise marginalized and under-represented communities. For instance, to generate video captions in Chinese/French/German or to translate Chinese text into French/German, it is necessary to collect video-Chinese/French/German sentence pairs or Chinese-French/German sentence pairs, respectively. This problem scales in difficulty when considering even less commonly-encountered languages, many of which correspond directly to less privileged communities. To relax the dependence on

- * Equal Contributions. Ordered by a coin toss.
- Bang Yang and Yuexian Zou are with ADSPLAB, School of ECE, Peking University Shenzhen Graduate School, Shenzhen 518055, China, and also with Peng Cheng Laboratory, Shenzhen 518052, China. E-mail: {yangbang, zouyx}@pku.edu.cn.
- Fenglin Liu and David A. Clifton are with the Department of Engineering Science, University of Oxford, OX3 7DQ Oxford, U.K.. DAC is also with the Oxford-Suzhou Centre for Advanced Research, Suzhou, China. E-mail: {fenglin.liu, david.clifton}@eng.ox.ac.uk.
- Xian Wu is with Tencent AI Lab, China. E-mail: kevinxwu@tencent.com.
- Yaowei Wang is with Peng Cheng Laboratory, Shenzhen 518052, China. E-mail: wangyw@pcl.ac.cn.

Manuscript received XXX; revised XXX. (Corresponding author: Yuexian Zou.)




Fig. 1. During training, ZeroNLG first (i) projects different data across modalities and languages to corresponding coordinates in a shared common latent space; (ii) aligns their coordinates to bridge different domains; Here S_i and S_j refer to the text in non-English text, e.g. Chinese and German; (iii) performs unsupervised auto-encoding to learn to generate/reconstruct text given the coordinate of input text in this space. During inference, ZeroNLG encodes the input data acquiring its coordinate in this space, which can be directly used to perform zero-shot data-to-text generation (i.e., visual captioning and machine translation) without the need for downstream labeled pairs.

labeled data for downstream NLG tasks, in this paper we propose the novel ZeroNLG framework, which is particularly useful for non-English languages where data-text pairs are limited in availability.

To conduct zero-shot NLG across different modalities and languages, the core objective is to bridge the gap between various domains, e.g., vision and language domains, English and Chinese domains, and Chinese and German domains. To this end, as shown in Figure 1, we propose the Cross-Domain Alignment pre-training objective, which adopts Mean Square Error (MSE) and Info Noise Contrastive Estimation (InfoNCE) [13] losses, to pre-train a vision encoder and a multilingual encoder. In this way, the aligned and bridged domains can be used for zero-shot multimodal and multilingual natural language generation. In implementations, we introduce a vision encoder, an English encoder, and a multilingual encoder. The motivation for introducing an English encoder comes from the fact that there are lots of English-centric corpus, e.g., image-to-English datasets [14], [15], video-to-English datasets [16], [17], and non-English-to-English translation datasets [9], [18], [19]. Therefore, we introduce the English encoder to make use of the existing English resources to pre-train the vision encoder and the multilingual encoder. We first exploit the vision-English pairs (D_1) to pre-train the vision encoder and English encoder to *align the vision and English domains*. Then, we fix the parameters of the English encoder and exploit the English - non-English, e.g., the English-Chinese pairs (D_2), the English-German pairs (D_3), and the English-French pairs (D_4), to pre-train the multilingual encoder to *align and bridge domains between English and any language*. It is worth noting that the used datasets D_1 , D_2 , D_3 , and D_4 are independent and can have no overlap, i.e., English text in D_1 , the English text in D_2 , the English text in D_3 , and the English text in D_4 , are separate sets and have no overlap. It means that there are no pairs of vision and Chinese/German/French; Chinese and German/French; German and French. Considering that we have aligned the vision and English domains, our method can *align and bridge the vision and non-English domains* without training on the pairs of vision and non-English text. Meanwhile, due to various non-English domains being aligned with the English domain, *the non-English domains, i.e., Chinese, German, and French, are aligned with each other*. As a result, the ZeroNLG can *align and bridge different domains*

across modalities and languages in a shared common latent space without the training on downstream data-text pairs.

After aligning and bridging various domains, as shown in Figure 1, we further propose an unsupervised training objective Denoising Language Reconstruction (DLR) to learn to conduct zero-shot NLG. Here we present a multilingual auto-encoder, including the pre-trained multilingual encoder and a randomly initialized multilingual decoder. The DLR aims to reconstruct the input sentences across various languages. During training, our method samples the English/Chinese/German/French sentences $S_1/S_2/S_3/S_4$ from $D_1/D_2/D_3/D_4$, as input to learn to reconstruct the input sentences in the $S_i \rightarrow S_i$ ($i = 1, 2, 3, 4$) auto-encoding pipeline. In the prediction stage, due to we have aligned and bridged various domains, we can directly replace the S_i with images/videos V as input to generate the zero-shot visual captions for different languages in the $V \rightarrow S_i$ ($i = 1, 2, 3, 4$) pipeline. Meanwhile, we can perform the zero-shot Chinese \leftrightarrow German, Chinese \leftrightarrow French, and German \leftrightarrow French machine translation by inputting S_j to the model to generate the translation in the $S_j \rightarrow S_i$ ($j \neq i$) pipeline. Overall, the proposed ZeroNLG can perform zero-shot multimodal and multilingual natural language generation without the requirements of any downstream data-text pairs for training. Besides, we can find that our method has the potential to be easily extended to other languages (e.g., Swedish and Italian) by aligning and bridging the English and target language domains. The extensive experiments on various NLG tasks, including image captioning, video captioning, and machine translation, across English, Chinese, German, and French, significantly prove the effectiveness of the proposed ZeroNLG.

Overall, the contributions of this work are:

- We propose an effective method ZeroNLG to make the first attempt to perform zero-shot multimodal and multilingual natural language generation in a unified framework, where the downstream training pairs are not available.
- Our method bridges different domains across modalities and languages by aligning them in a common latent space; then learns to perform zero-shot language generation by auto-encoding/reconstructing the sentences in different languages.
- The extensive experiments and analyses on twelve natural language generation tasks across multiple languages show that our ZeroNLG can generate desirable sentences without using any labeled downstream data-text pairs for training and significantly outperforms existing state-of-the-art zero-shot learning methods.

2 RELATED WORKS

We introduce the related works from two aspects: i) natural language generation and ii) zero-shot learning.

2.1 Natural Language Generation (NLG)

The goal of NLG is to automatically generate fluent and accurate natural language text based on given input data such as text [1], images [2], and video [3].

This is typically achieved using an encoder-decoder framework where the encoder computes intermediate representations of the input data and the decoder uses RNN [8] to generate the final output. Attention mechanisms [20], [21], [22], [23] have been proposed to provide the decoder with full access to the source information, resulting in more efficient use of the input data. In particular, fully attentive models such as the Transformer [24] have been successful in achieving state-of-the-art results in multiple NLG tasks such as image captioning [25], [26], video captioning [27], [28], and neural machine translation [18], [19]. However, to efficiently train the data-driven models, most existing works rely on pairs of input data and corresponding output text, which could be difficult to obtain in the real world. Additionally, there has been relatively little research concerning zero-shot learning.

2.2 Zero-shot Learning

Recently, few-shot learning [29] has received growing research interests [30], [31], [32], [33], [34], [35]. Inspired by the success of few-shot learning, several works [36], [37] explored such an approach for data-to-text NLG tasks, which mainly include text-to-text and vision-to-text tasks.

In recent years, lots of zero-shot text-to-text machine translation models have been proposed [36], [37], [38], [39], [40], [41]. Typically, the source language and the target language are mapped into a common latent space, where sentences with the same semantic meaning are well aligned, thus the zero-shot text-to-text translation can be carried out. Another line of research focuses on the “prompt-based learning” [42] of large language models (LLMs) that acquire impressive sentence completion ability from massive pre-training text data. By providing LLMs with a textual template that consists of task-specific information [43], [44], several in-context examples [12], [45], or a chain of thoughts [11], [46], [47], prompt-based methods can exploit the potentials of LLMs to perform zero-shot or few-shot NLG.

For the vision-to-text task (i.e., visual captioning), zero-shot learning is particularly difficult because of the great disparities between the vision and the language domains, as well as the distinct characteristics of each modality. As a result, the zero-shot vision-to-text works [48], [49], [50], [51], [52], [53], [54], [55], [56], [57] are relatively much less and the overall frameworks are more complex than those used for text-to-text tasks. For example, [48] proposed a method to generate captions in a central language (Chinese) and subsequently translate them into the target language (English), without the training on downstream vision-text pairs. However, this proposed method can not adopt visual information from images to generate more robust captions containing accurate visual details. [49] aligned and bridged the vision and language domains with visual objects (e.g., *girl*, *umbrella*). To achieve it, it was necessary to employ complex models and strategies to obtain higher-quality vision captions, e.g. object detection model [58], image reconstruction [49] and adversarial learning [59], in which the detector is limited to a pre-defined set of objects. Nevertheless, it does not incorporate other visual information, e.g., attributes (*wooden*), relations (*holding*), and color (*red*) to include more visual details to generate captions. [50], [60]

used the scene graph to bridge the gap between the vision and language domains. However, in order to construct an accurate and reliable scene graph, they were obliged to utilize Faster-RCNN [61] as the object detector, MOTIFS [62] as the relationship detector, and an additional classifier for attribute identification [60]. [56], [63] require the adversarial learning [59], [64] for training. Most recently, several works [53], [54], [57] adopt LLMs like GPT [45], [65] for zero-shot visual captioning. While these techniques are effective, they have some drawbacks including the excessive parameterization of LLMs and a lack of adaptability to multilingual environments. Overall, although the existing methods for zero-shot vision-to-text have shown considerable progress, they are hard to implement and still far from real-world applications.

To this end, we propose the ZeroNLG framework. The unique advantages of our method are i) it is simple but highly effective, outperforming all existing zero-shot methods; ii) it can utilize the full information of input data to perform zero-shot generation; iii) it can deal with various NLG problems across modalities and languages within a unified framework; iv) it could easily be extended to other languages - we prove its effectiveness on Chinese, German, and French. As a result, our method could have the potential to promote the application of NLG, especially vision-to-text, for various low-resource language applications.

3 APPROACH

We first formulate how ZeroNLG tackles zero-shot NLG and then introduce two key components: cross-domain alignment and denoising language reconstruction.

3.1 Formulation

As shown in Figure 2, our ZeroNLG first aligns and bridges different domains in a shared common latent space, then performs unsupervised auto-encoding to learn to generate text by reconstructing the input text. In implementations, we choose English, Chinese, German, and French to evaluate our approach. Therefore, we denote the vision data as V , the English sentence as S_1 , the Chinese sentence as S_2 , the German sentence as S_3 , and the French sentence as S_4 . As shown in Figure 2, ZeroNLG includes four modules, i.e., a vision encoder $\mathbb{E}_v(\cdot)$, an English text encoder $\mathbb{E}_e(\cdot)$, a multilingual encoder $\mathbb{E}_m(\cdot)$, and a multilingual decoder $\mathbb{D}_m(\cdot)$. Our ZeroNLG is defined as follows:

$$\begin{array}{l}
 \text{Pre-training} \left\{ \begin{array}{l}
 \text{Vision-English Alignment} : \mathbb{E}_v(V) \xrightarrow{\text{Align}} \mathbb{E}_e(S_1) \\
 \text{Cross-Lingual Alignment} : \mathbb{E}_e(S_1) \xleftarrow{\text{Align}} \mathbb{E}_m(S_i) \\
 \text{Denoising Reconstruction} : \mathbb{E}_m(S_i) \xrightarrow[\mathbb{D}_m(\cdot)]{\text{Reconstruct}} S_i
 \end{array} \right. \\
 \\
 \text{Zero-shot Inference} \left\{ \begin{array}{l}
 \text{Zero-shot Vision-to-Text} : \mathbb{E}_v(V) \xrightarrow[\mathbb{D}_m(\cdot)]{\text{Generate}} S_j \\
 \text{Zero-shot Text-to-Text} : \mathbb{E}_m(S_i) \xrightarrow[\mathbb{D}_m(\cdot)]{\text{Generate}} S_j
 \end{array} \right. \quad (1)
 \end{array}$$

where $i, j = 1, 2, 3, 4, j \neq i$, S_j denotes the target sentences. It is worth noting that introducing the English text encoder $\mathbb{E}_e(\cdot)$ has several merits: 1) we can make full use of existing English-centric resources, e.g., the pairs of vision and English, and the pairs of non-English and English; 2) we can adopt the English-centric pre-trained model, e.g., CLIP

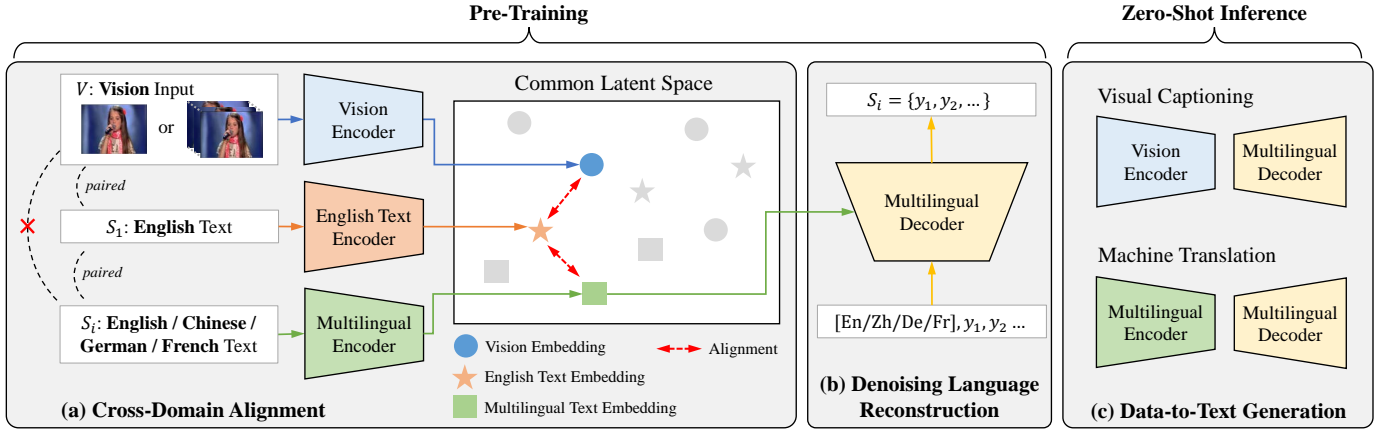


Fig. 2. The illustration of our proposed ZeroNLG, including two components: cross-domain alignment and denoising language reconstruction, where the former aims to align and bridge different data in a shared common latent space and the latter aims to reconstruct the input sentences across different languages, learning to generate sentences based on the embeddings in the common latent space. We rely on English-centric pairs for training, i.e., vision-English, English-Chinese, English-German, and English-French, where the English texts in different sets have no overlap. During inference, we can perform zero-shot natural language generation, including vision-to-Chinese/German/French captioning, and Chinese \leftrightarrow German, Chinese \leftrightarrow French, and German \leftrightarrow French machine translation.

[66], as our basic model to boost the alignment of vision and English domains. Such a continual learning scheme is a critical step towards sustainable AI [67]. After pre-training, our ZeroNLG can directly perform zero-shot vision-to-text and text-to-text in the $\mathbb{D}_m(\mathbb{E}_v(V)) \rightarrow S_j$ and $\mathbb{D}_m(\mathbb{E}_m(S_i)) \rightarrow S_j$ pipelines, respectively. Therefore, the ZeroNLG can deal with zero-shot NLG across modalities and languages within a single framework without the requirement of downstream labeled data-text pairs, which are not easy to acquire in the real world, especially for non-English languages.

3.2 Cross-Domain Alignment

We introduce the Info Noise Contrastive Estimation (InfoNCE) [13] and Mean Square Error (MSE) losses to align and bridge different domains. In particular, InfoNCE loss is a type of contrastive loss function used for self-supervised learning and has shown success in bridging the gap between the visual and textual modalities [66], [68], [69]; MSE loss minimizes the distance between two different distributions.

Given a batch of K English-centric training samples, including K pairs of English text and Data, we denote the encoded English text and Data of k^{th} training sample as (s_k, d_k) , where $s = \mathbb{E}_e(S_1)$. We take the vision data and the text data in i^{th} language as an example, if the Data is vision data, $d = \mathbb{E}_v(V)$; if the Data is multilingual text data in i^{th} language, $d = \mathbb{E}_m(S_i)$. Therefore, the InfoNCE loss can be formulated as follows:

$$\begin{aligned} \mathcal{L}_{\text{InfoNCE}}^{s \rightarrow d} &= -\frac{1}{K} \sum_{k=1}^K \log \frac{\exp(\langle s_k, d_k \rangle / \tau)}{\sum_{l=1}^K \exp(\langle s_k, d_l \rangle / \tau)} \\ \mathcal{L}_{\text{InfoNCE}}^{d \rightarrow s} &= -\frac{1}{K} \sum_{k=1}^K \log \frac{\exp(\langle d_k, s_k \rangle / \tau)}{\sum_{l=1}^K \exp(\langle d_k, s_l \rangle / \tau)} \\ \mathcal{L}_{\text{InfoNCE}} &= \frac{1}{2} \left(\mathcal{L}_{\text{InfoNCE}}^{s \rightarrow d} + \mathcal{L}_{\text{InfoNCE}}^{d \rightarrow s} \right), \end{aligned} \quad (2)$$

where the $\langle \cdot, \cdot \rangle$ and τ denote the cosine similarity and a temperature hyper-parameter [70], respectively.

The MSE loss can be defined as follows:

$$\mathcal{L}_{\text{MSE}} = \frac{1}{2K} \sum_{k=1}^K \|s_k - d_k\|_2^2, \quad (3)$$

where $\|\cdot\|_2$ denotes L2-norm.

By combing the InfoNCE and MSE losses, we obtain the final training loss of the Cross-Domain Alignment (CDA):

$$\mathcal{L}_{\text{CDA}} = \lambda_1 \cdot \mathcal{L}_{\text{InfoNCE}} + \lambda_2 \cdot \mathcal{L}_{\text{MSE}}, \quad (4)$$

where $\lambda_1, \lambda_2 \in [0, 1]$ are hyper-parameters that control the strength of each loss item. Through the above equation, our method can align and bridge different domains in a shared latent space, which provides a solid bias for zero-shot natural language generation.

3.3 Denoising Language Reconstruction

After aligning multimodal and multilingual domains, the next key step is to enable ZeroNLG to learn to generate multilingual text. Here, we introduce an unsupervised objective dubbed Denoising Language Reconstruction (DLR) to train ZeroNLG. Specifically, we adopt Transformer [24] as the decoder to reconstruct S_i ($i = 1, 2, 3, 4$) in the $S_i \rightarrow S_i$ auto-encoding pipeline. In implementations, we first randomly sample the sentences in i^{th} language S_i , then adopt the multilingual encoder $\mathbb{E}_m(\cdot)$ to acquire the coordinates (i.e., embeddings) of input sentences in the latent space $\mathbb{E}_m(S_i)$, and finally adopt the multilingual decoder $\mathbb{D}_m(\cdot)$ to reconstruct the input sentences, defined as:

$$S_i = \mathbb{D}_m(\mathbb{E}_m(S_i)), \quad (5)$$

where $\mathbb{E}_m(S_i)$ denotes the coordinates of input text in the latent space. To train our method, we take the input sentence as the ground truth to be generated, i.e., $S_i = \{y_0, y_1, y_2, \dots, y_{|S_i|}\}$, where y_0 and $|S_i|$ denote the begin-of-sentence token and the number of tokens, respectively, and utilize the cross-entropy loss, which is widely used in natural language generation problems:

$$\mathcal{L}_{\text{DLR}} = -\sum_{l=1}^{|S_i|} \log p(y_l | y_{0:l-1}; \mathbb{E}_m(S_i)), \quad (6)$$

where we implement y_0 as a language-specific token following [71], [72] so that the decoder can be aware of which language to be generated.

Data Corruption For vision-to-text, due to (i) the large variations of images and videos caused by different object attributes, occlusion, motion blur etc [73]; (ii) the great disparities between the vision and the language domains [74], we propose two data corruption strategies to further improve the performance and robustness of our ZeroNLG.

In implementations, we simultaneously consider the input and feature corruptions. For input corruption, we adopt the masking strategy as in BERT [75] to randomly mask $r\%$ tokens of the input sentences S_i , obtaining the corrupted input sentences S'_i . As a result, the DLR process becomes:

$$S_i = \mathbb{D}_m(\mathbb{E}_m(S'_i)), \quad (7)$$

For the feature corruption, we propose to add Gaussian noise $n \sim \mathcal{N}(0, \epsilon)$ into the text features $\mathbb{E}_m(S_i)$ (i.e., the coordinates) of input sentences S_i , acquiring the corrupted features of input sentences $\mathbb{E}'_m(S_i) = \mathbb{E}_m(S_i) + n$. Therefore, the DLR process is re-defined as follows:

$$S_i = \mathbb{D}_m(\mathbb{E}'_m(S_i)) = \mathbb{D}_m(\mathbb{E}_m(S_i) + n), \quad (8)$$

Through data corruption, we can encourage the model to learn more robust latent representations, achieving strong performance on zero-shot natural language generation.

4 EXPERIMENTS

In this section, we conduct experiments on multiple NLG tasks, i.e., vision-to-text image captioning and video captioning, and text-to-text neural machine translation. We first describe public datasets for pre-training and evaluation. Then, we present the performance of our approach on zero-shot multimodal and multilingual natural language generation across modalities and languages.

4.1 Experimental Setups

4.1.1 Datasets and Downstream Tasks

Pre-Training Datasets The WebImageText [66], [76] dataset (WIT) is used for vision-English alignment. WIT consists of 400 million image-English text pairs collected from the internet. For ease of experimentation, we directly use CLIP [66] pre-trained on WIT. Besides, the CC3M dataset [77] is used for cross-lingual alignment and denoising language reconstruction. For CC3M consisting of 3.3M English (En) sentences, as we consider three non-English languages: Chinese (Zh), German (De), and French (Fr), we split the corpus into three non-overlapping splits, each of which contains 1.1M English sentences and is translated to the corresponding language via Google Translator [78]. As a result, we can acquire 1.1M En-Zh pairs, 1.1M En-De pairs, and 1.1M En-Fr pairs. Moreover, we randomly sample a subset of 1.1M English sentences from CC3M to ensure that the data (En, En-Zh, En-De, En-Fr) for pre-training is balanced.

It is worth noting that more strong and robust performance could be achieved by using large-scale human-annotated translation datasets. In our experiments, we find

TABLE 1

Pre-training and testing datasets used for experiments. It is worth noting that the four English corpora used for pre-training are independent and can have no overlap. For English application scenarios, as the training data and testing data are from different domains and we also do not use any downstream data for training, we can still consider the evaluation on the English corpora as zero-shot NLG, as in the existing works [53], [54], [55], [57], [79]. As a result, we evaluate ZeroNLG on twelve tasks across modalities and languages. As there are no human-annotated datasets available for the video to German/French tasks, we report the qualitative results in Figure 4.

	Data-Text Pairs	Pre-training	Testing
English Corpora	Image-English	WIT [66], [76]	MS-COCO [14]
	English-Chinese	CC3M-Zh [77]	En-Zh [80]
	English-German	CC3M-De [77]	WMT16 [81]
	English-French	CC3M-Fr [77]	WMT17 [82]
	Video-English	-	MSR-VTT [16]
Non-English Corpora	Video-Chinese	-	VATEX-Zh [17]
	Image-Chinese	-	Flickr30k-Zh [80]
	Image-German	-	Flickr30k-De [83]
	Image-French	-	Flickr30k-Fr [83]
	Chinese-German	-	WMT16 [81]
	Chinese-French	-	WMT17 [82]
German-French	-	WMT17 [82]	

that our approach can achieve state-of-the-art zero-shot results with the machine-translated sentences, thus we do not attempt to utilize more human-labeled high-quality datasets for pre-training. Therefore, our model is not limited to the currently used pre-training data.

Downstream Tasks and Datasets As shown in Table 1, we consider two types of natural language generation tasks: visual captioning and machine translation. We focus on four languages, i.e., Chinese, German, French, and English. For visual non-English captioning, we report performance on video-to-Chinese, image-to-Chinese, image-to-German and image-to-French directions using VATEX-Zh [17], Flickr30k-Zh [80], Flickr30k-De [83], and Flickr30k-Fr [83] datasets, respectively. For visual English captioning, we utilize the widely-adopted MSR-VTT [16] and MS-COCO [14] datasets to measure the performance of video-to-English and image-to-English, respectively. For evaluating the performance of our method on machine translation, we only adopt humanly translated and annotated data. We obtain the translation pairs used for evaluation from the WMT16 [81] and WMT17 [82] machine translation competitions and English-Chinese dataset [80], including (a) English-Chinese pairs, (b) English-German pairs, (c) English-French pairs, (d) Chinese-German pairs, (e) Chinese-French pairs, and (f) German-French pairs. For data preparation, we adopt the official splits to split the datasets and only adopt the testing set to evaluate the zero-shot performance of our approach.

4.1.2 Metrics

For visual captioning in English, we follow the common practice in the literature to report BLEU-4 [84], METEOR [85], ROUGE-L [86] and CIDEr [87]. We also include the SPICE metric [88] for the image-to-English generation. For visual non-English captioning, METEOR and SPICE metrics are excluded because they consider synonym matching and named entity recognition in English by default. All metrics are computed by the widely-used Microsoft COCO

TABLE 2

Performance of zero-shot vision-to-text visual captioning across three non-English Languages, i.e., Chinese, German, and French. B-4, R-L, and C are short for BLEU-4, ROUGE-L, and CIDEr, respectively. Higher is better in all columns. All previous works can not deal with zero-shot captioning for non-English languages. For comparison, we re-implement three state-of-the-art zero-shot English captioning models equipped with Google Translator and a state-of-the-art machine translation model NLLB-200. As we can see, our ZeroNLG can simultaneously generate desirable visual captioning across different languages in a single unified framework and achieves the best zero-shot results.

Methods	Year	Video-to-Text			Image-to-Text								
		Chinese			Chinese			German			French		
		B-4	R-L	C	B-4	R-L	C	B-4	R-L	C	B-4	R-L	C
CoCa [55] + Google Translator	2022	1.4	15.3	4.3	2.9	19.0	9.8	3.5	20.3	11.0	2.4	15.1	19.8
CoCa [55] + NLLB-200 [7]	2022	0.0	11.6	1.2	0.8	13.2	2.6	3.3	20.5	10.5	2.3	15.3	18.8
CLIP-Re [54] + Google Translator	2022	2.7	20.8	9.6	2.9	23.7	15.2	2.1	21.6	13.1	1.8	15.5	21.9
CLIP-Re [54] + NLLB-200 [7]	2022	0.8	16.1	2.8	1.1	17.5	4.6	2.2	21.7	12.8	1.9	15.6	21.4
CapDec [52] + Google Translator	2022	2.9	22.3	5.1	4.7	26.6	13.9	5.4	26.7	16.9	2.5	18.2	23.2
CapDec [52] + NLLB-200 [7]	2022	0.9	17.1	1.9	1.8	19.3	4.7	5.2	27.0	16.9	2.5	18.5	23.6
ZeroNLG	Ours	7.1	29.6	9.8	8.4	31.8	18.0	5.7	27.2	17.1	2.8	18.6	24.8

Evaluation Server [14], where we use different toolkits to segment Chinese, German, and French sentences (as introduced next). For machine translation, we report the widely adopted BLEU [84] measured by the SacreBLEU toolkit [89].

4.1.3 Implementation Details

As shown in Eq. 1, ZeroNLG includes a vision encoder, an English text encoder, a multilingual encoder, and a multilingual decoder. We implement them as follows.

- Vision Encoder $\mathbb{E}_v(\cdot)$ and English text encoder $\mathbb{E}_e(\cdot)$: following existing works [52], [53], [54], [57], [66], we adopt the pre-trained and frozen CLIP [66], which is composed of a ViT-B/32 model [90] and a decoder-only text encoder [91], to implement our vision encoder and English text encoder. Such practice can substantially save computing resources and energy, which is a critical step towards sustainable AI [67].
- Multilingual encoder $\mathbb{E}_m(\cdot)$: we adopt the pre-trained multilingual DistilBERT [92] as the multilingual encoder, which adopts WordPiece embeddings [78] and has a vocabulary of size 119,547.
- Multilingual decoder $\mathbb{D}_m(\cdot)$: we implement it as Transformer decoder [24] with 768 model dimensions, 12 attention heads, 3 layers, and the same word embeddings as that of the multilingual encoder.

For the English text encoder, we extract global features from the position of the end-of-sentence token following CLIP. For the multilingual encoder, we truncate sentences into a maximum length of 128 and follow sBERT [93] to obtain mean pooled global features. We use AdamW [94] with L2 weight decay of 0.01 to train models for 3 epochs. We set the learning rate fixed to $2e-5$ after 5K warm-up iterations. The batch size is 128 for cross-domain alignment and 32 for language reconstruction. Based on the validation performance, we set $r = 0$ and $\epsilon = 0.1$ for visual captioning and $r = 5$ and $\epsilon = 0.01$ for machine translation. $\{\lambda_1, \lambda_2\} = \{1, 0\}$ (Eq. 1) is used in CLIP for vision-English alignment. For cross-lingual alignment, we set $\{\lambda_1, \lambda_2\} = \{0, 1\}$ (see Table 6). We use the Jieba toolkit¹ to segment Chinese sentences, and use the CoreNLP toolkit [95] for German and French sentences. When processing

1. <https://github.com/fxsjy/jieba>

videos, we uniformly sample 8 frames. We use beam search with a beam size of 3 to generate texts. Our ZeroNLG can be trained on *only* an NVIDIA T4 card within 18 hours.

4.2 Zero-shot Vision-to-Text (Visual Captioning)

Vision-to-Non-English We conduct the experiments on three non-English languages, i.e., Chinese (Zh), German (De), and French (Fr). In detail, we adopt the VATEX-Zh [17] to report the performance on video-to-Chinese; adopt the Flickr30k-Zh [80], Flickr30k-De [83], and Flickr30k-Fr [83] to report the performance on image-to-Chinese, image-to-German, image-to-French, respectively.

Since all previous works can not deal with zero-shot visual captioning tasks for non-English languages, we re-implement three existing state-of-the-art zero-shot English captioning methods, i.e., CoCa [55], CLIP-Re [54], CapDec [52], and equip them with strong machine translation models to generate non-English captions. In implementations, we use CoCa (ViT-B/32 variant) pre-trained on LAION-2B [96] and released by OpenCLIP [97], and re-implement the CLIP-Re and CapDec on the same corpus as ours. For translation, we adopt a commercial translation product Google Translator [78] and a recent state-of-the-art machine translation model NLLB-200 [7].

Table 2 reports the zero-shot visual captioning results across Chinese, German, and French. As we can see, our approach achieves the best zero-shot results across all metrics and languages. For video captioning, our proposed ZeroNLG consistently outperforms the existing zero-shot captioning methods. In detail, it achieves up to 4.2%, 7.3%, and 0.2% absolute improvements compared to previous best results in terms of BLEU-4, ROUGE-L, and CIDEr, respectively. For image captioning, our ZeroNLG successfully surpasses all baselines across all languages by up to 18.4%, 1.2%, and 5.1% relative improvements on the CIDEr metric in terms of Chinese, German, and French, respectively. In brief, the above results significantly prove the effectiveness of our approach ZeroNLG in dealing with zero-shot multimodal and multilingual natural language generation within a unified framework.

Vision-to-English Considering that several zero-shot English captioning methods have been proposed [53], [54], [55], [57], [79], for comparison with these existing works,

TABLE 3

Performance of zero-shot vision-to-text visual captioning in English. ‡: Our re-implementations. *: Pre-trained on LAION-2B [96] that consists of 2 billion image-English text pairs. †: SMs [57] needs to call GPT-3 API which is quite time-consuming. Therefore SMs [57] only reported their performance on 100 randomly sampled test instances. To conduct a fair comparison, we report the result of ZeroNLG on the same 100 test instances used by SMs. As we can see, our approach outperforms previous state-of-the-art zero-shot methods on most metrics.

Methods	Year	Pre-trained Backbone	Video-to-Text (English)				Image-to-Text (English)				
			B-4	M	R-L	C	B-4	M	R-L	C	S
CapDec [‡] [52]	2022	CLIP	7.3	14.1	33.5	8.4	8.8	13.5	33.0	25.0	7.9
ZeroCap [‡] [53]	2022	CLIP + GPT-2	2.3	10.8	23.1	7.3	1.8	9.1	19.7	14.4	5.0
MAGIC [54]	2022	CLIP + GPT-2	5.5	13.3	35.4	7.4	5.2	12.5	30.7	18.3	5.7
EPT [79]	2022	CLIP + GPT-2	3.0	14.6	27.7	11.3	-	-	-	-	-
CoCa [‡] [55]	2022	CLIP-like*	3.4	10.2	21.5	5.6	5.2	11.2	24.2	16.7	6.8
ZeroNLG	Ours	CLIP	8.7	15.0	35.4	9.9	9.6	14.4	34.9	29.9	8.7
SMs [†] [57]	2023	CLIP + GPT-3	-	-	-	-	10.0	16.2	36.1	50.1	10.8
ZeroNLG [†]	Ours	CLIP	-	-	-	-	12.0	15.7	38.1	51.6	11.1

TABLE 4

Performance of zero-shot machine translation across English (En), Chinese (Zh), German (De), and French (Fr). We report the BLEU metric measured by the SacreBLEU toolkit [89]. Higher is better in all columns. → and ← denote the translation direction. (·) is calculated by comparing with our method in terms of the sentence pairs used for pre-training. Our approach achieves encouraging performance with less number of parameters and pre-training pairs. More importantly, these listed works can not deal with vision-to-text multimodal NLG tasks.

Methods	Year	#Params	#Sentence Pairs	En-Zh		En-De		En-Fr		Average English	Zh-De		Zh-Fr		De-Fr		Average Non-English
				→	←	→	←	→	←		→	←	→	←			
mBART-50 [71]	2020	610.9M	203.7M (62x)	18.9	12.5	32.4	34.0	30.4	41.1	28.2	6.9	0.3	4.2	1.7	7.6	17.9	6.4
M2M-100 [72]	2021	483.9M	7.5B (2,000x)	16.4	10.5	24.5	30.2	30.7	36.4	24.8	8.5	13.3	6.8	14.9	22.6	23.5	14.9
NLLB-200 [7]	2022	617.2M	18B (5,000x)	6.3	12.8	37.5	39.8	49.8	46.8	32.2	10.7	4.1	5.7	4.9	34.2	30.8	14.7
ZeroNLG	Ours	165.0M	3.3M (1x)	14.7	8.8	20.5	21.1	22.0	24.6	18.6	7.3	11.9	5.2	16.2	16.7	18.5	12.6

we further conduct the experiments on zero-shot vision-to-English to verify the effectiveness of our approach. In particular, for a fair comparison, we follow previous works to conduct experiments on the MSR-VTT video captioning [16] and MSCOCO image English captioning [14] datasets. It is worth noting that the training data and testing data are from different domains, therefore, like previous works [53], [54], [55], [57], [79], we can consider this evaluation as zero-shot English captioning.

As shown in Table 3, our method ZeroNLG substantially outperforms existing state-of-the-art zero-shot methods. For example, on image captioning, our approach can outperform all previous works, several of which incorporate a large language modeling model – GPT [45], [65]. On video captioning, although EPT [79] achieves a better CIDEr score than ours (11.3 vs. 9.9), it needs more time to generate a caption for each video (> 1 minute in EPT vs. < 1 second in ZeroNLG). The vision-to-English experiments further prove the effectiveness of our approach, which achieves state-of-the-art zero-shot results across different languages.

4.3 Zero-shot Text-to-Text (Machine Translation)

To further prove the effectiveness of our approach, we conduct six zero-shot machine translation tasks across English (En), Chinese (Zh), German (De), and French (Fr), i.e., En-Zh, En-De, En-Fr, Zh-De, Zh-Fr, De-Fr. We report the results on Table 4. We compare our ZeroNLG with three pre-trained large language models (LLMs) designed for neural machine translation: mBART-50 [71], M2M-100 [72], and NLLB-200 [7]. In detail, M2M-100 and NLLB-200 are respectively pre-trained on around 7.5B and 18B many-to-many training

pairs, including En-Zh, En-De, En-Fr, Zh-De, Zh-Fr, and De-Fr, while mBART-50 and our ZeroNLG are only pre-trained on English-centric training pairs, i.e., En-Zh, En-De, and En-Fr, without any Zh-De, Zh-Fr, De-Fr training pairs. Nevertheless, due to the training set and testing set being from different domains, we can consider all these LLMs and our ZeroNLG as zero-shot machine translation models.

As shown in Table 4, our approach shows potential when compared with the existing state-of-the-art LLMs, which adopt more model parameters and training pairs. Especially, without the downstream pairs for training, our ZeroNLG can significantly outperform the mBART-50 model, across all non-English translation tasks. For the Zh←Fr task, our ZeroNLG achieves the best results. It further proves the effectiveness of our method in dealing with zero-shot machine translation, where the downstream pairs are not available.

Overall Combining the results of zero-shot vision-to-text and text-to-text results, we can conclude that our proposed ZeroNLG can perform zero-shot multimodal and multilingual natural language generation in a single framework and outperform previous state-of-the-art zero-shot methods. The advantages under the scenarios without any downstream labeled data-text pairs show that ZeroNLG might be applied to other low-resource languages (Swedish, Italian, etc.).

5 ANALYSIS

To understand the effect of each component in our framework, we conduct several analyses in this section. We focus on the more challenging task, i.e., vision-to-text visual captioning, to perform the analysis.

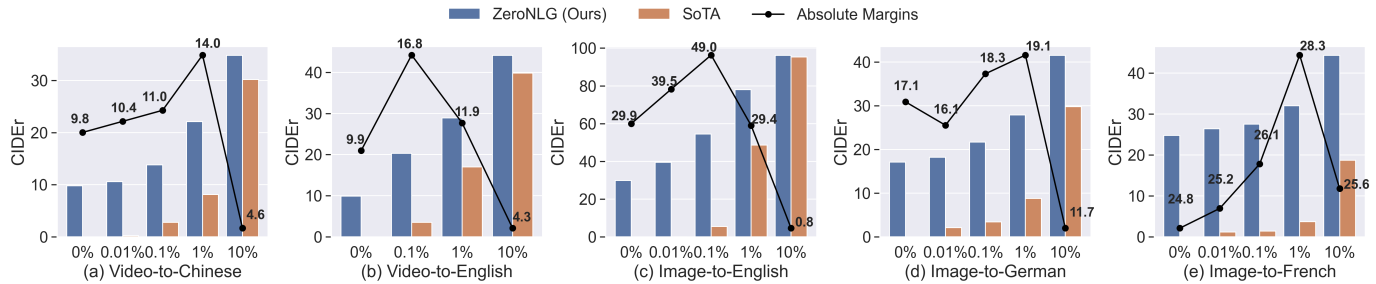


Fig. 3. Results of vision-to-text visual captioning with respect to different ratios of downstream data used for training. The absolute margins between our model and the state-of-the-art (SoTA) model ClipCap [4] are shown with the polyline. Our method consistently and significantly outperforms the SoTA under the very limited pairs setting (i.e., 0.01%, 0.1%, and 1%).

TABLE 5

Quantitative analysis of our ZeroNLG, which consists of the cross-domain alignment (CDA) and denoising language reconstruction (DLR).

Setting	DLR		Video-to-Text												Image-to-Text															
	CDA	Data Corruption	Languages				English (En)				Chinese (Zh)				English (En)				Chinese (Zh)				German (De)				French (Fr)			
			En	Zh	De	Fr	B-4	M	R-L	C	B-4	R-L	C	B-4	M	R-L	C	S	B-4	R-L	C	B-4	R-L	C	B-4	R-L	C			
Full	✓	✓	✓	✓	✓	8.7	15.0	35.4	9.9	7.1	29.6	9.8	9.6	14.4	34.9	29.9	8.7	8.4	31.8	18.0	5.7	27.2	17.1	2.8	18.6	24.8				
(a)	✓	-	✓	✓	✓	1.5	10.2	24.8	3.7	7.9	18.6	2.9	1.0	7.7	19.0	5.7	3.0	0.7	17.7	3.5	0.0	15.2	3.5	0.5	10.1	8.0				
(b)	-	✓	✓	✓	✓	0.6	7.5	22.7	0.5	0.0	16.8	0.5	0.7	5.3	15.8	1.0	1.0	0.0	14.9	0.7	0.0	11.4	0.7	0.0	7.2	2.8				
(c)	-	-	✓	✓	✓	Fail																								
(d)	✓	✓	✓			7.3	14.6	34.0	9.0	-	-	-	9.3	14.2	34.3	27.5	8.3	-	-	-	-	-	-	-	-	-				
(e)	✓	✓		✓		-	-	-	-	6.4	28.1	8.6	-	-	-	-	-	8.0	30.4	17.5	-	-	-	-	-					
(f)	✓	✓			✓	-	-	-	-	-	-	-	-	-	-	-	-	-	-	3.6	25.3	13.8	-	-	-					
(g)	✓	✓			✓	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	2.3	17.8	21.2					

TABLE 6

Impact of encoder choices, the number of decoder layers (L), the corpora for cross-lingual alignment (CLA), and loss designs (λ_1, λ_2) of CLA. We perform the analysis on the image-to-text visual captioning under different languages and report the CIDEr scores. Default settings are highlighted in a gray background. As we can see, our ZeroNLG can benefit from more advanced pre-trained multilingual encoders, decoders with larger capacities, and proper pre-training corpora.

Encoder	En	Zh	De	Fr	L	En	Zh	De	Fr	Corpora	En	Zh	De	Fr	λ_1	λ_2	En	Zh	De	Fr
mDistilBERT	29.9	18.0	17.1	24.8	3	29.9	18.0	17.1	24.8	(1) CC3M _{4L}	29.9	18.0	17.1	24.8	0	1	29.9	18.0	17.1	24.8
mBERT	32.0	19.4	18.0	26.9	6	31.0	19.0	17.9	26.2	(2) WMT _{4L}	21.2	14.7	9.3	13.2	0.1	1	28.6	18.9	15.0	23.6
XLNet-R	28.7	18.5	16.5	23.9	12	32.1	19.2	19.3	28.8	(2) \rightarrow (1)	31.6	19.5	17.0	26.4	1	1	27.5	16.3	12.6	22.8
sBERT*	32.7	20.1	18.0	27.1	24	34.4	19.8	21.2	30.9	(1) \rightarrow (2)	29.4	18.0	14.2	21.4	1	0.1	24.9	15.4	13.1	20.7
										None	1.0	0.7	0.7	2.8	1	0	23.1	14.2	10.3	19.1

*: paraphrase-multilingual-mpnet-base-v2

\rightarrow : sequential training

5.1 Semi-Supervised Learning

To further prove the effectiveness of our method, we propose to utilize a few labeled downstream data-text pairs for fine-tuning. To this end, in Figure 3, we evaluate the performance of our approach on visual captioning across English, Chinese, German, and French with respect to the increasing amount of paired data. For a fair comparison, we also re-train the state-of-the-art (SoTA) model ClipCap [4] using the same amount of downstream data-text pairs and the same architecture as our model. As we can see, our ZeroNLG can be significantly boosted with few downstream labeled pairs and outperforms the SoTA under all ratios of data used for training. More importantly, i) without any downstream pairs for training (0%), our ZeroNLG can even significantly surpass the SoTA trained with 10% downstream pairs by 6.1% absolute CIDEr score on image-to-French; ii) under very limited pairs settings, e.g., 0.01% and 0.1%, our approach can outperform the SoTA trained with 1% downstream pairs on all vision-to-text tasks. The strong

performance of our ZeroNLG under the very limited pairs setting proves the effectiveness of our approach in relaxing the reliance on the downstream pairs to provide a solid basis for natural language generation, which is particularly useful for low-resource language applications, where the labeled data-text pairs are scarce and hard to collect.

Overall, with very limited labeled downstream training pairs, our method can be efficiently fine-tuned and deployed on low-resource language applications. Such capacity could improve the practicality of natural language generation in real-world applications, and contribute to marginalized and under-represented communities.

5.2 Quantitative Analysis

In this section, we conduct a quantitative analysis to understand the contributions of each component in our approach.

Ablation study As shown in Table 5, removing our components significantly degrades the performance across different languages, which demonstrates the effectiveness of



Fig. 4. The examples of visual captions generated by the state-of-the-art zero-shot model [52] and our ZeroNLG model for different languages, i.e., English, Chinese, German, and French, under the zero-shot setting. For better understanding, we add English translations below the non-English captions in brackets. We highlight accurate keywords and wrong details. As we can see, ZeroNLG can generate accurate and vivid descriptions with more visual details across languages.

our proposed methods for zero-shot natural language generation (NLG). In particular, both settings (b) and (c) show that when removing the cross-domain alignment, the model fails to accurately perform zero-shot NLG, which indicates the importance of bridging different domains for zero-shot NLG. By comparing Full and setting (a), we can find that data corruption strategy can bring significant improvements. It can be explained that the data corruption strategy in language reconstruction can encourage the model to efficiently bridge the domains and learn more robust latent representations. We verify it in the following visualization. Overall, the ablation study proves our arguments and the effectiveness of the proposed components.

Effect of the number of languages By comparing the performance of Full and settings (d-g) in Table 5, we can find that adding more languages successfully enables our ZeroNLG to not only deal with multiple language application scenarios within a single unified framework but also consistently boost the performance of each language application scenario across all metrics. The improved results show that the different knowledge from different language texts can be unified in our ZeroNLG to achieve a better language understanding, producing an overall improvement across all metrics regardless of the downstream language application scenarios.

Effect of encoder choices We compare different pre-trained multilingual encoders, including multilingual DistilBERT [92] (*mDistilBERT*), multilingual BERT [75] (*mBERT*), XLM-RoBERTa [98] (*XLM-R*), and sentence BERT [93] (*sBERT*). All these models are base-size. As we can observe from Table 6, mBERT outperforms mDistilBERT due to a larger capacity. Surprisingly, XLM-R obtains inferior performance, possibly because it can not measure semantic textual similarity well². By contrast, sBERT — the model that has been pre-trained on a sentence similarity task, performs the best among all variants. Thus, our ZeroNLG is not limited to the default encoder choice (i.e., mDistilBERT) and can benefit from a more sophisticated encoder like sBERT.

2. <https://www.sbert.net/examples/training/multilingual>.

Effect of decoder capacities The decoder of our ZeroNLG is shallow (i.e., $L = 3$ layers) by default. We here ablate different L in Table 6, where we can observe gradually boosted performance as L increases. Therefore, our ZeroNLG can be improved by a larger decoder capacity. It not only proves that more model parameters or stronger decoders can lead to further improvements, but also shows the potential of our ZeroNLG, which could be further improved by directly scaling up the model parameters.

Effect of corpora Apart from the cross-lingual alignment corpus introduced in Section 4.1.1 (abbreviated as $CC3M_{4L}$), we here consider translation data that focuses on news text for comparison and construct a corpus dubbed WMT_{4L} of the same data scale as $CC3M_{4L}$. From Table 6, we can observe that although using WMT_{4L} is not as competitive as $CC3M_{4L}$ due to substantial distribution shifts, it still achieves obvious performance gains compared with not using any corpora. Besides, switching the training orders of WMT_{4L} and $CC3M_{4L}$ can produce encouraging results, proving the robustness of our model to training. More importantly, sequentially training on WMT_{4L} and $CC3M_{4L}$ (i.e., “(2) \rightarrow (1)”) yields generally the best results. This demonstrates the importance of pre-training on corpora with gradually decreasing distribution shifts.

Effect of loss designs In Table 6, we evaluate the effect of the InfoNCE loss (λ_1) and the MSE loss (λ_2) when performing cross-lingual alignment. We can observe that $\lambda_1 > 0$ suffers from performance degradation compared with $\lambda_1 = 0$. We speculate that this is because contrastive learning requires a large batch size to include sufficient (hard) negative samples [99], which is hard to realize in our case due to resource restrictions. In contrast, $\{\lambda_1, \lambda_2\} = \{0, 1\}$ is compute-friendly and performs the best.

5.3 Visualization

To better understand our method, in Figure 5, we adopt t-SNE [100] to visualize vision and multilingual embeddings. For comparison, we consider (a) the Base model (i.e., without our proposed CDA and DLR), (b) the Base model with CDA, and (c) our ZeroNLG. As we can see,




Fig. 5. We show the t-SNE visualization [100] of vision and multilingual embeddings. We plot the scatter diagrams with 200 samples for each modality and language. For comparison, we show the embeddings learned by (a) the Base model (i.e., without our CDA and DLR), (b) the Base model with CDA (Eq. 1), and (c) our full model ZeroNLG.

for the Base model, although the embeddings belonging to the same domain are well-clustered, there is a clear gap between different domains. With CDA, texts in different languages are well-aligned with each other whereas vision and text domains still exhibit disparities, possibly because contrastive learning converges to local minima [101]. Finally, by considering data corruption during DLR, our ZeroNLG can align and bridge different domains across modalities and languages well.

5.4 Qualitative Analysis

In this section, we conduct a qualitative analysis in Figure 4 to intuitively understand our proposed ZeroNLG.

Case study In Figure 4, we give four examples to compare our proposed ZeroNLG with the state-of-the-art (SoTA) zero-shot model CapDec [52]. Both models are trained on the same pre-training corpus and have no access to downstream vision-text data. As we can observe, under the zero-shot setting, the SoTA can not well describe the vision content and generates several wrong descriptions (Red-colored text), failing to depict some important visual objects, e.g., “dog” in (a), “girl” and “flowers” in (b), and “bride and groom” in (d), while our approach can generate fluent and “believable” captions containing these important objects. More encouragingly, the captions generated by our ZeroNLG are well supported by accurate visual details, e.g., “beautiful” and “cute” in (b), “happy” in (c), and “wedding” in (d). Besides, considering that there are no public human-annotated datasets available for the video to German/French tasks, the capacity of our ZeroNLG in generating reasonable German and French captions is encouraging. Overall, our approach can generate high-quality and desirable outputs for different languages within a unified framework. It further proves our arguments and the effectiveness of our proposed approach.

Error analysis We further perform the error analysis to analyze our ZeroNLG. As we can see, our model suffers from several common drawbacks: i) generating repeated or incomplete sentences, e.g., “in the sand on the bench” and the German sentence in (a); and ii) misunderstanding objects and scenes in some cases, e.g., the French sentence in (d). They can be attributed to the lack of detailed visual relationships and accurate visual information. We may alleviate these drawbacks by introducing a visual object extractor and a scene graph, where the former predicts a set of visual objects and the latter models the relationships between objects. Both of them are widely used in previous

zero-shot methods [49], [50], [60]. However, it is unlikely to be avoided completely, as these drawbacks are common in natural language generation [102].

6 CONCLUSIONS AND FUTURE WORKS

In this work, we make the first attempt to achieve zero-shot multimodal and multilingual natural language generation in a unified framework. To this end, we propose the intuitive ZeroNLG approach, which first exploits English-centric data to align and bridge different domains across modalities and languages, and then auto-encodes languages to learn to perform zero-shot NLG. Our experiments demonstrate that, without any available downstream data for training, ZeroNLG can produce desirable outputs given various forms of input data, i.e. images, video, and text. Extensive investigation of performance with twelve NLG tasks, including image captioning, video captioning, and neural machine translation, demonstrate the effectiveness of our approach, and where we conclude that ZeroNLG significantly outperforms previous state-of-the-art zero-shot methods within a single framework.

Future works: Substantial research avenues exist in i) further aligning with various types of images, such as artwork and 3D renderings, to perform the image-to-image translation and text-to-image translation (i.e., image generator) within a unified framework; ii) further boosting performance by learning and unifying knowledge from various different languages (and images). We note in passing that our proposed method can be improved by incorporating more languages, which has been shown in Table 5. We emphasize the applicability of this method in improving the manner in which we offer data-to-data mappings for under-represented languages and communities, which is an area of particular future opportunity in promoting “fair AI”.

REFERENCES

- [1] D. Bahdanau, K. Cho, and Y. Bengio, “Neural machine translation by jointly learning to align and translate,” in *ICLR*, 2015. 1, 2
- [2] O. Vinyals, A. Toshev, S. Bengio, and D. Erhan, “Show and tell: A neural image caption generator,” in *CVPR*, 2015. 1, 2
- [3] S. Venugopalan, H. Xu, J. Donahue, M. Rohrbach, R. J. Mooney, and K. Saenko, “Translating videos to natural language using deep recurrent neural networks,” in *HLT-NAACL*, 2015. 1, 2
- [4] R. Mokady, A. Hertz, and A. H. Bermano, “Clipcap: Clip prefix for image captioning,” *arXiv preprint arXiv:2111.09734*, 2021. 1, 8
- [5] B. Yang, Y. Zou, F. Liu, and C. Zhang, “Non-autoregressive coarse-to-fine video captioning,” in *AAAI*, 2021. 1
- [6] X. Hu, Z. Gan, J. Wang, Z. Yang, Z. Liu, Y. Lu, and L. Wang, “Scaling up vision-language pre-training for image captioning,” in *CVPR*, 2022. 1
- [7] M. R. Costa-jussà, J. Cross, O. Çelebi, M. Elbayad, K. Heffernan, K. Heffernan, E. Kalbassi, J. Lam, D. Licht, J. Maillard *et al.*, “No language left behind: Scaling human-centered machine translation,” *arXiv preprint arXiv:2207.04672*, 2022. 1, 6, 7
- [8] S. Hochreiter and J. Schmidhuber, “Long short-term memory,” *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997. 1, 3
- [9] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, “Attention is all you need,” in *NIPS*, 2017. 1, 2
- [10] F. Stahlberg, “Neural machine translation: A review,” *Journal of Artificial Intelligence Research*, vol. 69, pp. 343–418, 2020. 1
- [11] A. Chowdhery, S. Narang, J. Devlin, M. Bosma, G. Mishra, A. Roberts, P. Barham, H. W. Chung, C. Sutton, S. Gehrmann *et al.*, “Palm: Scaling language modeling with pathways,” *arXiv preprint arXiv:2204.02311*, 2022. 1, 3

- [12] J.-B. Alayrac, J. Donahue, P. Luc, A. Miech, I. Barr, Y. Hasson, K. Lenc, A. Mensch, K. Millican, M. Reynolds *et al.*, "Flamingo: a visual language model for few-shot learning," in *NeurIPS*, 2022. 1, 3
- [13] A. v. d. Oord, Y. Li, and O. Vinyals, "Representation learning with contrastive predictive coding," *arXiv preprint arXiv:1807.03748*, 2018. 2, 4
- [14] X. Chen, H. Fang, T.-Y. Lin, R. Vedantam, S. Gupta, P. Dollár, and C. L. Zitnick, "Microsoft coco captions: Data collection and evaluation server," *arXiv preprint arXiv:1504.00325*, 2015. 2, 5, 6, 7
- [15] P. Young, A. Lai, M. Hodosh, and J. Hockenmaier, "From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions," *TACL*, 2014. 2
- [16] J. Xu, T. Mei, T. Yao, and Y. Rui, "MSR-VTT: A large video description dataset for bridging video and language," in *CVPR*, 2016. 2, 5, 7
- [17] X. Wang, J. Wu, J. Chen, L. Li, Y.-F. Wang, and W. Y. Wang, "Vatex: A large-scale, high-quality multilingual dataset for video-and-language research," in *ICCV*, 2019. 2, 5, 6
- [18] M. Ott, S. Edunov, A. Baevski, A. Fan, S. Gross, N. Ng, D. Grangier, and M. Auli, "fairseq: A fast, extensible toolkit for sequence modeling," in *NAACL-HLT (Demonstrations)*, 2019. 2, 3
- [19] F. Wu, A. Fan, A. Baevski, Y. N. Dauphin, and M. Auli, "Pay less attention with lightweight and dynamic convolutions," in *ICLR*, 2019. 2, 3
- [20] T. Luong, H. Pham, and C. D. Manning, "Effective approaches to attention-based neural machine translation," in *EMNLP*, 2015. 3
- [21] F. Liu, X. Ren, X. Wu, S. Ge, W. Fan, Y. Zou, and X. Sun, "Prophet attention: Predicting attention with future attention," in *NeurIPS*, 2020. 3
- [22] K. Xu, J. Ba, R. Kiros, K. Cho, A. Courville, R. Salakhudinov, R. Zemel, and Y. Bengio, "Show, attend and tell: Neural image caption generation with visual attention," in *ICML*, 2015. 3
- [23] Y. Li, B. Yang, X. Cheng, S. Zhu, H. Li, and Y. Zou, "Unify, align and refine: Multi-level semantic alignment for radiology report generation," in *ICCV*, 2023, pp. 2863–2874. 3
- [24] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," in *NIPS*, 2017. 3, 4, 6
- [25] M. Cornia, M. Stefanini, L. Baraldi, and R. Cucchiara, "Meshed-memory transformer for image captioning," in *CVPR*, 2020. 3
- [26] F. Liu, Y. Liu, X. Ren, X. He, and X. Sun, "Aligning visual regions and textual concepts for semantic-grounded image representations," in *NeurIPS*, 2019. 3
- [27] F. Liu, X. Ren, X. Wu, B. Yang, S. Ge, and X. Sun, "O2NA: an object-oriented non-autoregressive approach for controllable video captioning," in *ACL/IJCNLP (Findings)*, 2021. 3
- [28] B. Yang, T. Zhang, and Y. Zou, "Clip meets video captioning: Concept-aware representation learning does matter," in *PRCV*. Springer, 2022, pp. 368–381. [Online]. Available: https://doi.org/10.1007/978-3-031-18907-4_29 3
- [29] Y. Wang, Q. Yao, J. T. Kwok, and L. M. Ni, "Generalizing from a few examples: A survey on few-shot learning," *ACM Computing Surveys*, 2020. 3
- [30] G. S. Dhillon, P. Chaudhari, A. Ravichandran, and S. Soatto, "A baseline for few-shot image classification," in *ICLR*, 2020. 3
- [31] Y. Tian, Y. Wang, D. Krishnan, J. B. Tenenbaum, and P. Isola, "Rethinking few-shot image classification: A good embedding is all you need?" in *ECCV*, 2020. 3
- [32] E. Perez, D. Kiela, and K. Cho, "True few-shot learning with language models," in *NeurIPS*, 2021. 3
- [33] Y. Gu, X. Han, Z. Liu, and M. Huang, "PPT: pre-trained prompt tuning for few-shot learning," in *ACL*, 2022. 3
- [34] T. Gao, A. Fisch, and D. Chen, "Making pre-trained language models better few-shot learners," in *ACL/IJCNLP*, 2021. 3
- [35] M. Tsimpoukelli, J. Menick, S. Cabi, S. M. A. Eslami, O. Vinyals, and F. Hill, "Multimodal few-shot learning with frozen language models," in *NeurIPS*, 2021. 3
- [36] J. Philip, A. Berard, M. Gallé, and L. Besacier, "Monolingual adapters for zero-shot neural machine translation," in *EMNLP*, 2020. 3
- [37] G. Chen, S. Ma, Y. Chen, D. Zhang, J. Pan, W. Wang, and F. Wei, "Towards making the most of cross-lingual transfer for zero-shot neural machine translation," in *ACL*, 2022. 3
- [38] M. Johnson, M. Schuster, Q. V. Le, M. Krikun, Y. Wu, Z. Chen, N. Thorat, F. Viégas, M. Wattenberg, G. Corrado, M. Hughes, and J. Dean, "Google's multilingual neural machine translation system: Enabling zero-shot translation," *TACL*, vol. 5, pp. 339–351, 2017. 3
- [39] M. Artetxe, G. Labaka, E. Agirre, and K. Cho, "Unsupervised neural machine translation," in *ICLR*, 2018. 3
- [40] G. Lample, A. Conneau, L. Denoyer, and M. Ranzato, "Unsupervised machine translation using monolingual corpora only," in *ICLR*, 2018. 3
- [41] G. Lample, M. Ott, A. Conneau, L. Denoyer, and M. Ranzato, "Phrase-based & neural unsupervised machine translation," in *EMNLP*, 2018. 3
- [42] P. Liu, W. Yuan, J. Fu, Z. Jiang, H. Hayashi, and G. Neubig, "Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing," *ACM Computing Surveys*, vol. 55, no. 9, pp. 1–35, 2023. 3
- [43] C. Raffel, N. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, and P. J. Liu, "Exploring the limits of transfer learning with a unified text-to-text transformer," *The Journal of Machine Learning Research*, vol. 21, no. 1, pp. 5485–5551, 2020. 3
- [44] X. L. Li and P. Liang, "Prefix-tuning: Optimizing continuous prompts for generation," in *ACL/IJCNLP*, 2021. 3
- [45] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell *et al.*, "Language models are few-shot learners," in *NeurIPS*, 2020. 3, 7
- [46] J. Wei, X. Wang, D. Schuurmans, M. Bosma, F. Xia, E. H. Chi, Q. V. Le, D. Zhou *et al.*, "Chain-of-thought prompting elicits reasoning in large language models," in *NeurIPS*, 2022. 3
- [47] T. Kojima, S. S. Gu, M. Reid, Y. Matsuo, and Y. Iwasawa, "Large language models are zero-shot reasoners," in *NeurIPS*, 2022. 3
- [48] J. Gu, S. R. Joty, J. Cai, and G. Wang, "Unpaired image captioning by language pivoting," in *ECCV*, 2018. 3
- [49] Y. Feng, L. Ma, W. Liu, and J. Luo, "Unsupervised image captioning," in *CVPR*, 2019. 3, 10
- [50] J. Gu, S. R. Joty, J. Cai, H. Zhao, X. Yang, and G. Wang, "Unpaired image captioning via scene graph alignments," in *ICCV*, 2019. 3, 10
- [51] F. Liu, M. Gao, T. Zhang, and Y. Zou, "Exploring semantic relationships for image captioning without parallel data," in *ICDM*, 2019. 3
- [52] D. Nukrai, R. Mokady, and A. Globerson, "Text-only training for image captioning using noise-injected CLIP," in *Findings of EMNLP*, 2022. 3, 6, 7, 9, 10
- [53] Y. Tewel, Y. Shalev, I. Schwartz, and L. Wolf, "Zerocap: Zero-shot image-to-text generation for visual-semantic arithmetic," in *CVPR*, 2022. 3, 5, 6, 7
- [54] Y. Su, T. Lan, Y. Liu, F. Liu, D. Yogatama, Y. Wang, L. Kong, and N. Collier, "Language models can see: Plugging visual controls in text generation," *arXiv preprint arXiv:2205.02655*, 2022. 3, 5, 6, 7
- [55] J. Yu, Z. Wang, V. Vasudevan, L. Yeung, M. Seyedhosseini, and Y. Wu, "Coca: Contrastive captioners are image-text foundation models," *TMLR*, pp. 1–20, 2022. 3, 5, 6, 7
- [56] F. Liu, X. Wu, C. You, S. Ge, Y. Zou, and X. Sun, "Aligning source visual and target language domains for unpaired video captioning," *TPAMI*, vol. 44, no. 12, pp. 9255–9268, 2022. 3
- [57] A. Zeng, M. Attarian, brian ichter, K. M. Choromanski, A. Wong, S. Welker, F. Tombari, A. Purohit, M. S. Ryoo, V. Sindhwani, J. Lee, V. Vanhoucke, and P. Florence, "Socratic models: Composing zero-shot multimodal reasoning with language," in *ICLR*, 2023. 3, 5, 6, 7
- [58] J. Huang, V. Rathod, C. Sun, M. Zhu, A. Korattikara, A. Fathi, I. Fischer, Z. Wojna, Y. Song, S. Guadarrama *et al.*, "Speed/accuracy trade-offs for modern convolutional object detectors," in *CVPR*, 2017. 3
- [59] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. C. Courville, and Y. Bengio, "Generative adversarial nets," in *NIPS*, 2014. 3
- [60] X. Yang, H. Zhang, and J. Cai, "Auto-encoding and distilling scene graphs for image captioning," *IEEE Trans. Pattern Anal. Mach. Intell. (TPAMI)*, 2020. 3, 10
- [61] S. Ren, K. He, R. B. Girshick, and J. Sun, "Faster R-CNN: towards real-time object detection with region proposal networks," in *NIPS*, 2015. 3
- [62] R. Zellers, M. Yatskar, S. Thomson, and Y. Choi, "Neural motifs: Scene graph parsing with global context," in *CVPR*, 2018. 3
- [63] I. Laina, C. Rupprecht, and N. Navab, "Towards unsupervised image captioning with shared multimodal embeddings," in *ICCV*, 2019. 3

- [64] J. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired image-to-image translation using cycle-consistent adversarial networks," in *ICCV*, 2017. 3
- [65] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, I. Sutskever *et al.*, "Language models are unsupervised multitask learners," *OpenAI blog*, vol. 1, no. 8, p. 9, 2019. 3, 7
- [66] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark *et al.*, "Learning transferable visual models from natural language supervision," in *CIML*, 2021. 4, 5, 6
- [67] S. Budenny, V. Lazarev, N. Zakharenko, A. Korovin, O. Plosskaya, D. Dimitrov, V. Akhrikin, I. Pavlov, I. Oseledets, I. Barsola *et al.*, "Eco2ai: carbon emissions tracking of machine learning models as the first step towards sustainable ai," in *Doklady Mathematics*, 2023, pp. 1–11. 4, 6
- [68] C. Jia, Y. Yang, Y. Xia, Y.-T. Chen, Z. Parekh, H. Pham, Q. Le, Y.-H. Sung, Z. Li, and T. Duerig, "Scaling up visual and vision-language representation learning with noisy text supervision," in *ICML*, 2021. 4
- [69] J. Yang, C. Li, P. Zhang, B. Xiao, C. Liu, L. Yuan, and J. Gao, "Unified contrastive learning in image-text-label space," in *CVPR*, 2022. 4
- [70] T. Chen, S. Kornblith, M. Norouzi, and G. E. Hinton, "A simple framework for contrastive learning of visual representations," in *ICML*, 2020. 4
- [71] Y. Tang, C. Tran, X. Li, P.-J. Chen, N. Goyal, V. Chaudhary, J. Gu, and A. Fan, "Multilingual translation with extensible multilingual pretraining and finetuning," *arXiv preprint arXiv:2008.00401*, 2020. 5, 7
- [72] A. Fan, S. Bhosale, H. Schwenk, Z. Ma, A. El-Kishky, S. Goyal, M. Baines, O. Celebi, G. Wenzek, V. Chaudhary *et al.*, "Beyond english-centric multilingual machine translation," *JMLR*, vol. 22, no. 1, pp. 4839–4886, 2021. 5, 7
- [73] L. Meng, H. Li, B.-C. Chen, S. Lan, Z. Wu, Y.-G. Jiang, and S.-N. Lim, "Adavit: Adaptive vision transformers for efficient image recognition," in *CVPR*, 2022. 5
- [74] W. Liang, Y. Zhang, Y. Kwon, S. Yeung, and J. Zou, "Mind the gap: Understanding the modality gap in multi-modal contrastive representation learning," *arXiv preprint arXiv:2203.02053*, 2022. 5
- [75] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," *arXiv preprint arXiv:1810.04805*, 2018. 5, 9
- [76] C. Schuhmann, R. Kaczmarczyk, A. Komatsuzaki, A. Katta, R. Vencu, R. Beaumont, J. Jitsev, T. Coombes, and C. Mullis, "Laion-400m: Open dataset of clip-filtered 400 million image-text pairs," in *NeurIPS Workshop*, 2021. 5
- [77] P. Sharma, N. Ding, S. Goodman, and R. Soicrut, "Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning," in *ACL*, 2018. 5
- [78] Y. Wu, M. Schuster, Z. Chen, Q. V. Le, M. Norouzi, W. Macherey, M. Krikun, Y. Cao, Q. Gao, K. Macherey, J. Klingner, A. Shah, M. Johnson, X. Liu, L. Kaiser, S. Gouws, Y. Kato, T. Kudo, H. Kazawa, K. Stevens, G. Kurian, N. Patil, W. Wang, C. Young, J. Smith, J. Riesa, A. Rudnick, O. Vinyals, G. Corrado, M. Hughes, and J. Dean, "Google's neural machine translation system: Bridging the gap between human and machine translation," *arXiv preprint arXiv:1609.08144*, 2016. 5, 6
- [79] Y. Tewel, Y. Shalev, R. Nadler, I. Schwartz, and L. Wolf, "Zero-shot video captioning with evolving pseudo-tokens," *arXiv preprint arXiv:2207.11100*, 2022. 5, 6, 7
- [80] W. Lan, X. Li, and J. Dong, "Fluency-guided cross-lingual image captioning," in *ACM MM*, 2017. 5, 6
- [81] "Wmt16 shared task: Multimodal machine translation," in *Proceedings of the First Conference on Machine Translation*, 2016. [Online]. Available: <https://www.statmt.org/wmt16/multimodal-task.html> 5
- [82] "Wmt17 shared task: Multimodal machine translation," in *Proceedings of the Second Conference on Machine Translation*, 2017. [Online]. Available: <https://www.statmt.org/wmt17/multimodal-task.html> 5
- [83] D. Elliott, S. Frank, K. Sima'an, and L. Specia, "Multi30k: Multilingual english-german image descriptions," in *ACL workshop*, 2016. 5, 6
- [84] K. Papineni, S. Roukos, T. Ward, and W. Zhu, "BLEU: a Method for automatic evaluation of machine translation," in *ACL*, 2002. 5, 6
- [85] S. Banerjee and A. Lavie, "METEOR: An automatic metric for MT evaluation with improved correlation with human judgments," in *ACL*, 2005. 5
- [86] C.-Y. Lin, "ROUGE: A package for automatic evaluation of summaries," in *ACL*, 2004. 5
- [87] R. Vedantam, C. L. Zitnick, and D. Parikh, "Cider: Consensus-based image description evaluation," in *CVPR*, 2015. 5
- [88] P. Anderson, B. Fernando, M. Johnson, and S. Gould, "SPICE: Semantic propositional image caption evaluation," in *ECCV*, 2016. 5
- [89] M. Post, "A call for clarity in reporting bleu scores," in *WMT*, 2018, pp. 186–191. 6, 7
- [90] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby, "An image is worth 16x16 words: Transformers for image recognition at scale," in *ICLR*, 2021. 6
- [91] A. Radford, K. Narasimhan, T. Salimans, and I. Sutskever, "Improving language understanding by generative pre-training," *OpenAI, Tech. Rep.*, 2018. 6
- [92] V. Sanh, L. Debut, J. Chaumond, and T. Wolf, "Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter," *arXiv preprint arXiv:1910.01108*, 2019. 6, 9
- [93] N. Reimers and I. Gurevych, "Sentence-bert: Sentence embeddings using siamese bert-networks," in *EMNLP-IJCNLP*, 2019, pp. 3982–3992. 6, 9
- [94] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," in *ICLR*, 2019. 6
- [95] C. D. Manning, M. Surdeanu, J. Bauer, J. R. Finkel, S. Bethard, and D. McClosky, "The stanford corenlp natural language processing toolkit," in *ACL (Demonstrations)*, 2014. 6
- [96] C. Schuhmann, R. Beaumont, R. Vencu, C. W. Gordon, R. Wightman, M. Cherti, T. Coombes, A. Katta, C. Mullis, M. Wortsman, P. Schramowski, S. R. Kundurthy, K. Crowson, L. Schmidt, R. Kaczmarczyk, and J. Jitsev, "LAION-5b: An open large-scale dataset for training next generation image-text models," in *NeurIPS (Datasets and Benchmarks Track)*, 2022. 6, 7
- [97] G. Ilharco, M. Wortsman, R. Wightman, C. Gordon, N. Carlini, R. Taori, A. Dave, V. Shankar, H. Namkoong, J. Miller, H. Hajishirzi, A. Farhadi, and L. Schmidt, "Openclip," *Zenodo*, Jul. 2021. [Online]. Available: <https://doi.org/10.5281/zenodo.5143773> 6
- [98] A. Conneau, K. Khandelwal, N. Goyal, V. Chaudhary, G. Wenzek, F. Guzmán, E. Grave, M. Ott, L. Zettlemoyer, and V. Stoyanov, "Unsupervised cross-lingual representation learning at scale," in *ACL*, 2020, pp. 8440–8451. 9
- [99] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, "A simple framework for contrastive learning of visual representations," in *ICML*. PMLR, 2020, pp. 1597–1607. 9
- [100] L. van der Maaten and G. Hinton, "Visualizing data using t-sne," *JMLR*, vol. 9, no. 86, pp. 2579–2605, 2008. 9, 10
- [101] P. Shi, M. C. Welle, M. Björkman, and D. Kragic, "Towards understanding the modality gap in clip," in *ICLR Workshop*, 2023, pp. 1–8. 10
- [102] A. See, P. J. Liu, and C. D. Manning, "Get to the point: Summarization with pointer-generator networks," in *ACL*, R. Barzilay and M. Kan, Eds., 2017. 10



Bang Yang received the B.E. degree from Sun Yat-sen University in 2018 and the M.S. degree from Peking University in 2021, where he is currently pursuing the Ph.D. degree. He is engaged in a joint training program with the Peng Cheng Laboratory. His research interests include multimodal learning and AI in healthcare.



Fenglin Liu is a PhD student at the University of Oxford. His research interests include Vision-and-Language Processing, Machine Learning, and their applications to healthcare. He has published papers at premier journals and conferences, e.g., TPAMI, NeurIPS, CVPR, ACL, EMNLP, NAACL. He has served as a senior program committee member for IJCAI and was awarded as the Distinguished/Outstanding Reviewer of CVPR, AAAI, and IJCAI.



David A. Clifton is the Royal Academy of Engineering Chair of Clinical Machine Learning at the University of Oxford, and OCC Fellow in AI & Machine Learning at Reuben College, Oxford. He was the first AI scientist to be appointed to an NIHR Research Professorship, which is the UK medical research community's "flagship Chair programme". He is a Fellow of the Alan Turing Institute, Research Fellow of the Royal Academy of Engineering, Visiting Chair in AI for Healthcare at the University of Manchester, and

a Fellow of Fudan University, China. He studied Information Engineering at Oxford's Department of Engineering Science, supervised by Prof. Lionel Tarassenko CBE, Chair of Electrical Engineering. His research focuses on the development of machine learning for tracking the health of complex systems. His previous research resulted in patented systems for jet-engine health monitoring, used with the engines of the Airbus A380, the Boeing 787 "Dreamliner", and the Eurofighter Typhoon. Since graduating from his DPhil in 2009, he has focused mostly on the development of AI-based methods for healthcare. Patents arising from this collaborative research have been commercialised via university spin-out companies OBS Medical, Oxehhealth, and Sensyne Health, in addition to collaboration with multinational industrial bodies. He was awarded a Grand Challenge award from the UK Engineering and Physical Sciences Research Council, which is an EPSRC Fellowship that provides long-term strategic support for "future leaders in healthcare". His research has been awarded over 35 academic prizes; in 2018, he was joint winner of the inaugural "Vice-Chancellor's Innovation Prize", which identifies the best interdisciplinary research across the entirety of the University of Oxford.



Yuexian Zou (Senior Member, IEEE) is currently a Full Professor with Peking University and the Director of the Advanced Data and Signal Processing Laboratory in Peking University and serves as the Deputy Director of Shenzhen Association of Artificial Intelligence (SAAI). She was a recipient of the award Leading Figure for Science and Technology by Shenzhen Municipal Government in 2009 and now is the adjunct professor in Peng Cheng Laboratory. She conducted more than 20 research projects including

NSFC and 863 projects. She has published more than 280 academic papers in famous journals and flagship conferences, and issued nine invention patents. Her research interests are mainly in machine learning and scene understanding.



Xian Wu is now a Principal Researcher in Tencent. Before joining Tencent, he worked as a Senior Scientist Manager and a Staff Researcher in Microsoft and IBM Research. Xian Wu received his PhD degree from Shanghai Jiao Tong University. His research interests include Medical AI, Natural Language Processing and Multi-Modal modeling. Xian Wu has published papers in CVPR, NeurIPS, ACL, WWW, AAAI, IJCAI etc. He also served as PC member of TKDE, TKDD, TOIS, TIST, CVPR, ICCV, AAAI etc.



Yaowei Wang (Member, IEEE) received the Ph.D. degree in computer science from the Graduate University of Chinese Academy of Sciences in 2005. He is currently an Associate Professor with the Peng Cheng Laboratory, Shenzhen, China. He was a Professor at the National Engineering Laboratory for Video Technology Shenzhen (NELVT), Peking University Shenzhen Graduate School, in 2019. From 2014 to 2015, he worked as an Academic Visitor at the Vision Laboratory, Queen Mary University

of London. He worked at the Department of Electronics Engineering, Beijing Institute of Technology, from 2005 to 2019. His research interests include machine learning, multimedia content analysis, and understanding. He is the author or coauthor of over 70 refereed journals and conference papers. He was a recipient of the second prize of the National Technology Invention in 2017 and the first prize of the CIE Technology Invention in 2015. His team was ranked as one of the best performers in the TRECVID CCD/SED tasks from 2009 to 2012 and PETS in 2012. He is a member of CIE, CCF, and CSIG.