

# Towards Unifying Medical Vision-and-Language Pre-training via Soft Prompts

Zhihong Chen<sup>1,2\*</sup> Shizhe Diao<sup>3\*</sup>

Benyou Wang<sup>1,2†</sup> Guanbin Li<sup>4†</sup> Xiang Wan<sup>2</sup>

<sup>1</sup>The Chinese University of Hong Kong, Shenzhen <sup>2</sup>Shenzhen Research Institute of Big Data

<sup>3</sup>The Hong Kong University of Science and Technology <sup>4</sup>Sun Yat-sen University

zhihongchen@link.cuhk.edu.cn sdiaoaa@connect.ust.hk

wangbenyou@cuhk.edu.cn liguanbin@mail.sysu.edu.cn wanxiang@sribd.com

## Abstract

*Medical vision-and-language pre-training (Med-VLP) has shown promising improvements on many downstream medical tasks owing to its applicability to extracting generic representations from medical images and texts. Practically, there exist two typical types, i.e., the fusion-encoder type and the dual-encoder type, depending on whether a heavy fusion module is used. The former is superior at multi-modal tasks owing to the sufficient interaction between modalities; the latter is good at uni-modal and cross-modal tasks due to the single-modality encoding ability. To take advantage of these two types, we propose an effective yet straightforward scheme named PTUnifier to unify the two types. We first unify the input format by introducing visual and textual prompts, which serve as DETR-like queries that assist in extracting features when one of the modalities is missing. By doing so, a single model could serve as a foundation model that processes various tasks adopting different input formats (i.e., image-only, text-only, and image-text-pair). Furthermore, we construct a prompt pool (instead of static ones) to improve diversity and scalability, enabling queries conditioned on different input instances. Experimental results show that our approach achieves competitive results on a broad range of tasks, spanning uni-modal tasks (i.e., image/text classification and text summarization), cross-modal tasks (i.e., image-to-text generation and image-text/text-image retrieval), and multi-modal tasks (i.e., visual question answering), demonstrating the effectiveness of our approach. Note that the adoption of prompts is orthogonal to most existing Med-VLP approaches and could be a beneficial and complementary extension to these approaches.<sup>1</sup>*

\*Equal contributions.

†Corresponding authors.

<sup>1</sup>The source code is available at <https://github.com/zjohnchan/ptunifier>.

## 1. Introduction

Medical data is multi-modal in general, among which vision and language are two critical modalities. It includes visual data (e.g., radiography, magnetic resonance imaging, and computed tomography) and textual data (e.g., radiology reports, and medical texts). More importantly, such images and texts are *pair-collected* in routine clinical practice (e.g., X-ray images and their corresponding radiology reports). Medical vision-and-language pre-training (Med-VLP) aims to learn generic representation from large-scale medical image-text pairs and then transfer it to various medical tasks, which is believed to be beneficial in addressing the data scarcity problem in the medical field.

Recently, substantial progress has been made toward research on Med-VLP [65, 32, 20, 43, 8]. In general, most existing Med-VLP models can be classified into two types: the dual-encoder type and the fusion-encoder type, where the former encodes images and texts separately to learn uni-modal/cross-modal representations following a shallow interaction layer (i.e., an image-text contrastive layer), and the latter performs an early fusion of the two modalities through the self-attention/co-attention mechanisms to learn multi-modal representations.<sup>2</sup> For dual-encoders, the purpose of existing studies [65, 20, 45, 59, 56, 60, 3] is to develop label-efficient algorithms to learn effective uni-modal/cross-modal representations since large-scale manually labeled datasets are difficult and expensive to obtain for medical images. The learned representations can improve the *effectiveness* of uni-modal (i.e., vision-only or language-only) tasks<sup>3</sup> and the *efficiency* of cross-modal (i.e., image-to-text or text-to-image) retrieval tasks significantly. For fusion-encoders,

<sup>2</sup>Although the terminologies “cross-modal” and “multi-modal” have been used interchangeably in the literature, we treat them as terms with different meanings in this paper.

<sup>3</sup>It is worth noting that most existing studies only conduct the evaluation on the vision-only tasks and disregard the language-only tasks although the text representations are simultaneously learned.

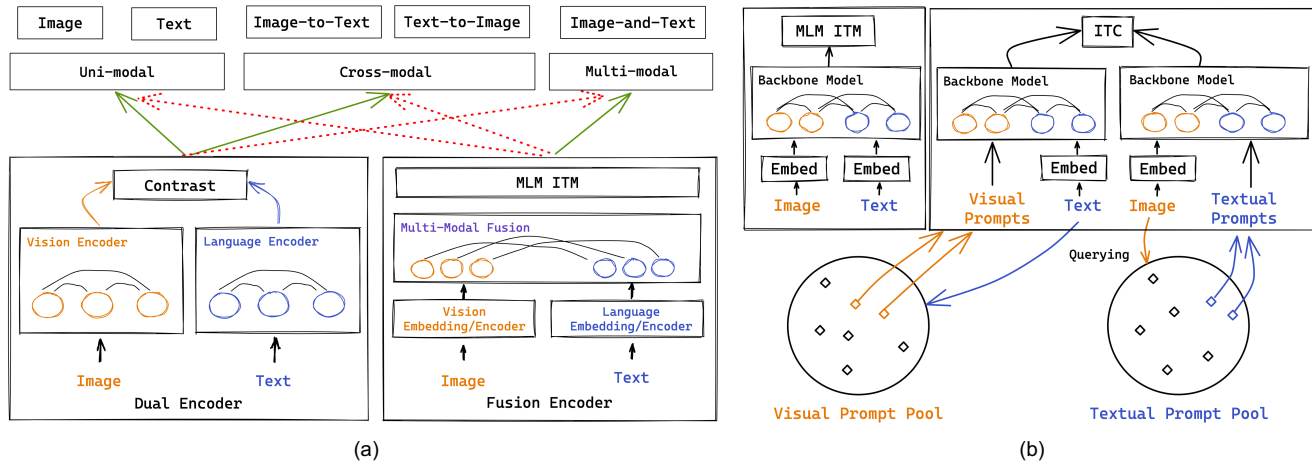


Figure 1: (a) Illustrations of two Med-VLP paradigms and their advantages (pointed by green arrows) and disadvantages (pointed by red arrows) in downstream tasks; (b) The overall architecture of our proposed approach, where the backbone models share the same parameters, and we duplicate them for illustration.

existing studies [32, 25, 43, 8, 9] aim to jointly process these two modalities with an early interaction to learn multi-modal representations to solve those tasks requiring multi-modal reasoning (e.g., medical visual question answering and medical image-text classification). However, it seems that “*you can’t have your cake and eat it, too.*”: the fusion-encoders can not perform uni-modal tasks effectively and cross-modal tasks efficiently due to the lack of single-modal encoding, while the dual-encoders underperform on multi-modal tasks owing to the insufficient interaction between modalities as shown in Figure 1(a).

In this paper, we aim to learn a unified medical vision-and-language pre-trained model. Although there exist some solutions [4, 51] to achieve a similar goal in the general domain, we propose an architecture- and task-agnostic approach named PTUnifier, which is much simpler and lighter-weighted. Technically, we develop the designs from the following perspectives: (i) *Compatibility*: we introduce visual and textual prompts to make the Med-VLP model compatible with different kinds of inputs (i.e., image-only inputs, text-only inputs, and image-text pairs); (ii) *Scalability*: we improve the diversity of the prompts by constructing prompt pools for different modalities from which different inputs are able to select their corresponding prompts, which enhances the capacity and makes it scalable to larger-scale Med-VLP. *In intuition, the introduced soft prompts serve as a query bank that stores the queries for extracting the most representative uni-modal features, the working mechanism of which is similar to the query vectors in DETR [5].* As a result, the proposed approach can be employed in unifying Med-VLP with many existing VLP model architectures (e.g., classic one [29] or even a single vanilla Transformer model) and does not require extra modality-dependent ar-

chitectures, resulting in better applicability. We perform the pre-training on three large-scale medical image-text datasets, i.e., ROCO [47], MedCaT [53], and MIMIC-CXR [24]. To verify the effectiveness of our approach and facilitate further research, we construct a medical vision-language benchmark including uni-modal tasks (i.e., image classification (IC) for vision and text classification (TC) and text summarization (TS) for language), cross-modal tasks (i.e., image-to-text retrieval (ITR), text-to-image retrieval (TIR), and image-to-text generation<sup>4</sup> (ITG)), and multi-modal tasks (i.e., visual question answering (VQA)). The proposed PTUnifier achieves excellent performance on all datasets, demonstrating its effectiveness.

## 2. Related Work

**Vision-and-Language Pre-training (VLP)** Motivated by the success of the self-supervised pre-training recipe in natural language processing (NLP) (e.g., BERT [14]) and computer vision (CV) (e.g., SimCLR [6] and MoCo [18]), there has been an increasing interest in developing VLP methods to address a wide range of vision-and-language-related tasks. In general, VLP methods can be classified into two categories according to the vision-and-language interaction, i.e., dual-encoders and fusion-encoders. Existing dual-encoder methods can be summarized according to the following aspects: (i) using medium-scale curated image-text data [48], (ii) using large-scale noisy image-text data [23], (iii) designing more fine-grained image-text contrast [62], (iv) adopting extra single modal contrastive learning [44]. For fusion-encoder approaches, existing studies can be further catego-

<sup>4</sup>Medical image-to-text generation refers to medical/radiology report generation in previous studies [11, 10].

alized with respect to these three perspectives: (i) Uni-modal encoders: different methods adopt different image features (*e.g.*, region features [29, 40], patch embeddings [26], and grid features [21]) and distinct text features (*e.g.*, statistic embeddings [26], and dynamic embeddings [17]); (ii) Multi-modal fusion modules: existing studies adopted the single-stream fusion scheme [52, 30] or dual-stream fusion scheme [54, 63]; (iii) Pretext tasks: existing studies explore a variety of pre-training tasks, including masked language modeling [29], masked image modeling [40, 7], image-text matching [64]. This paper adopts the model architecture of fusion encoders and the pretext tasks from both dual-encoder and fusion-encoder types.

### Medical Vision-and-Language Pre-Training (Med-VLP)

Being one of the applications and extensions of VLP to the medical domain, Med-VLP aims to understand the content of medical images and texts, which can be traced back to [65] for dual-encoders and [32] for fusion-encoders. For dual-encoders, the follow-up studies [20, 45, 56] explored the global-local image-text contrastive learning to capture more fine-grained information among medical images and texts and have achieved state-of-the-art results in the medical image classification task. For fusion-encoders, [25, 43, 8] performed pre-training to improve the multi-modal reasoning ability of the vision-and-language models for the downstream task (*e.g.*, Medical VQA). Besides, [9] integrated medical knowledge into the pre-training procedure to improve the performance on downstream medical tasks.

**Unified Vision-and-Language Pre-training** To unify the dual and fusion encoders, existing studies mainly adopted/designed specific model architectures to accommodate different pretext tasks. The most common scheme is to add an extra multi-modal fusion module to the dual encoders and perform the cross-modal pretext task (*i.e.*, image-text contrast) before the fusion and multi-modal pretext tasks (*e.g.*, MLM and ITM) after the fusion [28, 51]. Another line of research [12, 57] resorts to multi-tasking on various downstream supervised vision-language tasks by formulating them as sequence-to-sequence tasks. Besides, [4, 58] proposed a mixture-of-modality experts (MoME) Transformer to unify vision-and-language models by employing a set of modality experts to replace the feed-forward networks (FFN) in the standard Transformer. More recently, [15] proposed an encoder-decoder generative model learned from prefix language modeling and prefix image modeling. However, the aforementioned studies are either architecture-dependent or task-dependent, and they perform the unifying through training different parts of the models when applying different types of VLP objectives. Therefore, it is expected to unify the existing Med-VLP types in an *architecture- and*

*task-agnostic* fashion to improve the generalization and extensibility ability of Med-VLP methods.

## 3. Bridging the Gap

In this section, we introduce the PTUnifier framework for unifying the fusion-encoder and dual-encoder types. §3.1 details the problem to be addressed. This work proposes to unify inputs using prompts (in §3.2). Thus one could jointly train various tasks even with different input formats (in §3.3) in either pre-training or fine-tuning.

### 3.1. Problem Definition

We adopt the general problem formulation for pre-training following existing studies [29, 54]. Formally, given a medical image  $I$  and its corresponding description text  $T$ , the representation learning process can be formulated as

$$\theta^*, \theta_1^*, \dots, \theta_S^* = \arg \min_{\theta, \theta_1, \dots, \theta_S} \sum_{s=1}^S \mathcal{L}_s(Y_s, \mathcal{H}_{\theta_s}(\mathcal{M}_\theta(\mathbf{X}))), \quad (1)$$

where  $S$  refers to the number of pretext tasks;  $\mathcal{L}_s$  are the loss functions of pretext tasks;  $Y_s$  are the corresponding ground-truth labels;  $\mathcal{H}_{\theta_s}$  are the prediction heads with their parameters  $\theta_s$ ;  $\mathcal{M}_\theta$  is the backbone model which is parameterized by  $\theta$ ;  $\mathbf{X}$  represents the input to the backbone model, which could be one of the following cases:

$$\mathbf{X} = \begin{cases} (\mathbf{X}^v) & \text{if } \textit{image-only} \\ (\mathbf{X}^l) & \text{if } \textit{text-only} \\ (\mathbf{X}^v, \mathbf{X}^l) & \text{if } \textit{image-text} \end{cases} \quad (2)$$

where we suppose that we have embedded a medical image  $I$  as  $\mathbf{X}^v \in \mathbb{R}^{D_v \times N_v}$  or a medical text  $T$  as  $\mathbf{X}^l \in \mathbb{R}^{D_l \times N_l}$  when dealing with vision and language modalities. The challenge of the problem is to make the backbone model  $\mathcal{M}_\theta$  deal with such variable-size and heterogeneous input. After overcoming this challenge, we can perform different types of downstream vision-language tasks (*i.e.*, uni-modal, cross-modal, and multi-modal tasks).

### 3.2. Unifying Inputs via Prompts

To unify inputs, we propose to unify the inputs via prompts so as to perform different types of tasks. The working mechanism of the soft prompts is similar to the query vectors in DETR [5]. In specific, we design two solutions, *i.e.*, a basic solution for *compatibility* and an advanced solution for *scalability*. In this work, we use the advanced solution in default if not specified.

#### 3.2.1 Compatibility using Soft Prompts

To make the backbone model compatible with variable-size and heterogeneous input, this work proposes a simple yet

effective approach, namely using Prompt (PT) as a placeholder for missing modality.  $\mathcal{M}_\theta$  naturally accepts two inputs (visual and textual embeddings  $\mathbf{X}^v, \mathbf{X}^l$ ), which is by definition compatible to inputs with image-text pairs. For image-only/text-only inputs, we propose to introduce visual/textual prompts to enable the backbone model to perceive the missing input in a specific modality:

$$\mathbf{X} = \begin{cases} (\mathbf{X}^v, \mathbf{PT}^l) & \text{if } \textit{image-only} \\ (\mathbf{PT}^v, \mathbf{X}^l) & \text{if } \textit{text-only} \\ (\mathbf{X}^v, \mathbf{X}^l) & \text{if } \textit{image-text} \end{cases} \quad (3)$$

where  $\mathbf{PT}^v \in \mathbb{R}^{D_v \times k}$  and  $\mathbf{PT}^l \in \mathbb{R}^{D_l \times k}$  are the visual and textual prompts, respectively.

### 3.2.2 Scalability of Soft Prompts

The above solution adopts a static fashion to introduce prompts, which might have limited diversity and therefore harm its capacity. Hence, we construct a pool of visual/textual prompts *instead of static prompts*. Importantly, the selection of prompts is *conditioned on the input embeddings*.

Formally, we define a visual prompt pool  $\mathbf{V} \in \mathbb{R}^{D_v \times N_v}$  and a textual prompt pool  $\mathbf{T} \in \mathbb{R}^{D_l \times N_l}$ .  $N_v$  and  $N_l$  are the size of the visual/textual prompt pool, respectively. Given the image-only input with its visual embedding sequence  $\mathbf{X}^v$  or language-only input with its textual embedding sequence  $\mathbf{X}^l$ , we conduct a pooling operation (e.g., average/max pooling) to obtain a *query vector* for existing modality (denoted as  $\mathbf{q}^v$  or  $\mathbf{q}^l$ ), namely,  $\mathbf{q}^v = \text{pooling}(\mathbf{X}^v)$  and  $\mathbf{q}^l = \text{pooling}(\mathbf{X}^l)$ , respectively. To get the prompts of the missing modality, the selection of prompts is based on the similarity scores between the query vector and all prompts in the pool from the missing modality:

$$\begin{aligned} \mathbf{PT}^l &= \text{top-}k \left[ \mathbf{w}^T \mathbf{q}^v \right]_{\mathbf{w} \in \mathbf{V}}, \\ \mathbf{PT}^v &= \text{top-}k \left[ \mathbf{w}^T \mathbf{q}^l \right]_{\mathbf{w} \in \mathbf{T}}, \end{aligned} \quad (4)$$

where  $\mathbf{w}$  is an embedding vector in the prompt pool, and we select  $k$  closest prompts as the input embedding sequence of the missing modality.

**Intuitive Explanation** Without loss of any generality, we take a text-only scenario as an example, but it also holds for the image-only scenario. To select the best visual prompts for the text-only input, the proposed method chooses the input-conditioned ones compared to the given textual query vector. As an intuitive explanation, one could treat the visual prompt pool as a query bank that stores the queries for extracting uni-modal features when one of the modalities is missing.

**Linking to Prompts** We find that the PTUnifier (especially the static one in §3.2.1) is quite similar to the prompt tuning [31, 36]. They both introduce special tokens or vectors as a certain signal for training or inference. One notable difference is that in a special version of PTUnifier using prompts pools (see §3.2.2), the selection of additional tokens/vectors is conditioned on the input, while prompts are generally static and constant to input.

### 3.3. Unifying Multiple Pre-training Objectives

Owing to the unified image and/or text input formulation, we can adopt pretext tasks of both fusion-encoders and dual-encoders (see Eq. 1). Following previous studies [29, 54, 65, 48], we develop two commonly used pretext tasks (i.e., masked language modeling (MLM) and image-text matching (ITM)) for fusion-encoders and the image-text contrast (ITC) pretext task for dual-encoders. To produce the prediction for the aforementioned MLM and ITM tasks, we use two independent prediction heads  $\mathcal{H}_{\text{MLM}}$  and  $\mathcal{H}_{\text{ITM}}$  (i.e., two two-layer multilayer perceptrons (MLP)).

**Masked Language Modeling (MLM)** Following BERT [14], we randomly mask 15% of the words (denoted as  $Y_{\text{MLM}}$ ) of the input text  $T$  and recover them according to the remaining text ( $T_{\text{M}}$ ) and the input  $I$ . The MLM objective is given by:

$$\mathcal{L}_{\text{MLM}} = - \sum_{(I, T)} \log p_{\text{MLM}}(Y_{\text{MLM}} | I, T_{\text{M}}), \quad (5)$$

where  $p_{\text{MLM}}$  is obtained by applying  $\mathcal{H}_{\text{MLM}}$  followed by a softmax operation on the corresponding representations of [MASK] in  $\mathbf{Z}^l$ .

**Image-Text Matching (ITM)** aims to distinguish whether an image-text pair is a match. In detail, a positive image-text pair and a randomly sampled negative pair are fed into  $\mathcal{M}_\theta$  and the concatenation of  $\mathbf{z}_{[\text{CLS}]}^v$  and  $\mathbf{z}_{[\text{CLS}]}^l$  is processed by  $\mathcal{H}_{\text{MLM}}$  followed by a softmax layer to output a binary probability  $p_{\text{ITM}}$ . Therefore, the ITM objective is given by

$$\mathcal{L}_{\text{ITM}} = - \sum_{(I, T)} \log p_{\text{ITM}}(Y_{\text{ITM}} | I, T). \quad (6)$$

**Image-Text Contrast (ITC)** aims to learn better uni-modal/cross-modal representation from the instance-level contrast. In this work, given an image-text pair, we use two different forward procedures on the image-only input  $I$  and the text-only input  $T$ , respectively, to obtain the image-only representation (denoted as  $\mathbf{z}^v$ ) and text-only representation (denoted as  $\mathbf{z}^l$ ). Afterward, we adopt the similarity function  $s(I, T) = \mathbf{z}^{v \top} \mathbf{z}^l$  to compute the image-to-text similarity

and text-to-image similarity between  $z^v$  and  $z^l$ . Subsequently, the similarities are normalized as follows:

$$p_n^{i2t} = \frac{\exp(s(I, T_n)/\tau)}{\sum_{n=1}^N \exp(s(I, T_n)/\tau)}, \quad (7)$$

$$p_n^{t2i} = \frac{\exp(s(I_n, T)/\tau)}{\sum_{n=1}^N \exp(s(I_n, T)/\tau)}, \quad (8)$$

where  $N$  is the size of the mini-batch. The ground-truth labels  $Y^{i2t}$  and  $Y^{t2i}$  are two  $N \times N$  one-hot matrices, where negative pairs have a probability of 0 and the positive pair has a probability of 1. Therefore, the ITC objective is given by

$$\mathcal{L}_{ITC} = -\frac{1}{2} \sum_{(I,T)} \log p^{i2t}(Y^{i2t}|I, T) \quad (9)$$

$$-\frac{1}{2} \sum_{(I,T)} \log p^{t2i}(Y^{t2i}|I, T). \quad (10)$$

## 4. The Model Architecture

The previous section documents the unification at the input and task levels. This section will introduce the overall architecture of our work. As a pipeline, we first map visual and textual tokens into embeddings space ( $\mathbf{X}^v$  and  $\mathbf{X}^l$  as specified in §4.1). Such token embeddings with or without prompts will be jointly processed by an identical backbone model  $\mathcal{M}_\theta$  (§4.2). An overview of the proposed approach is shown in Figure 1(b).

### 4.1. Visual and Textual Embeddings

**Visual embedding** For an input image  $I$ , it is first segmented into patches following [16]. Then the patches are linearly projected into patch embeddings  $\mathbf{X}^v = (\mathbf{x}_1^v, \mathbf{x}_2^v, \dots, \mathbf{x}_{N_v}^v), \mathbf{x}_i^v \in \mathbb{R}^{D_v}$  through a linear transformation and a special learnable token embedding  $\mathbf{x}_{[\text{CLS}]}^v$  is prepended for the aggregation of visual information. Therefore, the image embedding sequence is obtained by summing up the patch embeddings and learnable 1D position embeddings  $\mathbf{E}_{pos}^v \in \mathbb{R}^{D_v \times (N_v+1)}$ :

$$\mathbf{X}^v = [\mathbf{x}_{[\text{CLS}]}^v; \mathbf{x}_1^v; \mathbf{x}_2^v; \dots; \mathbf{x}_{N_v}^v] + \mathbf{E}_{pos}^v, \quad (11)$$

where  $[\cdot; \cdot]$  represents the column concatenation.<sup>5</sup>

**Textual embedding** Similarly, for an input text  $T$ , we follow BERT [14] to tokenize the input text to subword tokens by WordPiece [61]. Afterwards, the tokens are linearly projected into embeddings  $\mathbf{X}^l = (\mathbf{x}_1^l, \mathbf{x}_2^l, \dots, \mathbf{x}_{N_l}^l), \mathbf{x}_i^l \in \mathbb{R}^D$  through a linear transformation with a start-of-sequence token embedding  $\mathbf{x}_{[\text{CLS}]}^l$ , and a special boundary token embedding  $\mathbf{x}_{[\text{SEP}]}^l$  added. Therefore, the text embedding sequence

<sup>5</sup>We overload the notation  $\mathbf{X}^v$  for simplicity (same for  $\mathbf{X}^l$ ).

Datasets	Image #	Text #	Avg. Len.	Avg. Sent. #
ROCO	81k	81k	20.42	1.46
MedICaT	124k	321k	40.88	2.82
MIMIC-CXR	232k	367k	36.49	5.07

Table 1: The statistics of the three pre-training datasets including the numbers of images, texts, the average word-based length (Avg. Len.) of texts, and the average number of sentences (Avg. Sent. #).

is obtained by summing up the sub-word token embeddings and text position embeddings  $\mathbf{E}_{pos}^l \in \mathbb{R}^{D \times (N_l+2)}$ :

$$\mathbf{X}^l = [\mathbf{x}_{[\text{CLS}]}^l; \mathbf{x}_1^l; \dots; \mathbf{x}_{N_l}^l; \mathbf{x}_{[\text{SEP}]}^l] + \mathbf{E}_{pos}^l. \quad (12)$$

### 4.2. The Backbone Model

Since the input image and/or text are represented as a unified image-text sequence, the backbone model can be any model for sequential modeling. In this work, we adopt an attention-based Med-VLP model with the multi-modal interaction, which can be an effective model (including uni-modal encoders and a multi-modal fusion module) or an efficient one (*i.e.*, a single Transformer model), where the attention mechanism is defined as

$$\text{ATTN}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{softmax}\left(\frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{D_k}}\right) \mathbf{V}, \quad (13)$$

where  $\mathbf{Q}$ ,  $\mathbf{K}$ , and  $\mathbf{V}$  are the query, key, and value matrix linearly transformed from the input embedding sequence, respectively, and  $D_k$  is the dimension of  $\mathbf{K}$ . Formally, for a given input (defined in Eq. 3), the whole representation process can be formulated as

$$\mathbf{Z}^v, \mathbf{Z}^l = \mathcal{M}_\theta(\mathbf{X}), \quad (14)$$

where  $\mathbf{Z}^v = (z_{[\text{CLS}]}^v, z_1^v, z_2^v, \dots, z_{N_v}^v)$  and  $\mathbf{Z}^l = (z_{[\text{CLS}]}^l, z_1^l, \dots, z_{N_l}^l, z_{[\text{SEP}]}^l)$  are the image and text representations from the backbone model.

## 5. Experimental Settings

### 5.1. Pre-training Datasets

In our experiments, we perform the pre-training on three datasets, which are described as follows:

- **ROCO** [47]: a dataset of radiology figure-caption pairs from PubMed Central, an open-access biomedical literature database.
- **MedICaT** [53]: a dataset of medical figure-caption pairs also extracted from PubMed Central. Different from ROCO, 75% of its figures are compound figures, including several sub-figures.

- **MIMIC-CXR** [24]: the largest radiology dataset to date consisting of chest X-ray images (in frontal or lateral views) and their reports from the Beth Israel Deaconess Medical Center.

For all the datasets, we exclude those samples with a text length of less than 3. For ROCO and MedICaT, we filter non-radiology samples, and for MIMIC-CXR, we only keep images in the frontal view. As for the dataset split, we adopt the official splits of ROCO and MIMIC-CXR. For MedICaT, we randomly sample 1,000 image-text pairs for validation and 1,000 for testing, and the remaining image-text pairs are used for training. Different from the texts in general-domain VLP, medical texts are long narratives consisting of multiple sentences. To deal with this case, we randomly sample a sentence from the input text in each iteration. Table 1 shows the statistics of the pre-training datasets.

## 5.2. Medical Vision-Language Benchmark

To evaluate the performance, we construct a medical vision-language evaluation benchmark including three types of tasks, *i.e.*, uni-modal, cross-modal, and multi-modal evaluations. All the adopted datasets are related to radiology.

**Uni-modal Evaluation** requires the model to process a single modality with vision-only or language-only inputs. For vision-only tasks, we conduct the image classification (IC) experiments on CheXpert [22] and RSNA Pneumonia [50]. For language-only tasks, we perform both the understanding task (*i.e.*, text classification (TC)) and the generation task (*i.e.*, text summarization (TS)) on the RadNLI [49, 42] and MIMIC-CXR datasets, respectively.

**Cross-modal Evaluation** requires the model to align the vision and language modalities. We conduct experiments on three kinds of tasks (*i.e.*, image-to-text retrieval (ITR), text-to-image retrieval (TIR), and image-to-text generation (ITG)). For ITR and TIR, we adopt the ROCO dataset and measure both zero-shot and fine-tuned performance. During the evaluation, we sample 2,000 image-text pairs from the ROCO test set and report the results on the 2,000 sampled image-text pairs due to the large time complexity of the ranking process. For ITG, we conduct experiments on the MIMIC-CXR dataset to evaluate its ability for radiology report generation.

**Multi-modal Evaluation** requires the model to reason over both the image and text inputs through the multi-modal interaction. We conduct the experiments on the medical visual question answering (VQA) task, which requires the model to answer natural language questions about a medical image. We adopt three publicly available Med-VQA datasets (*i.e.*, VQA-RAD [27], SLAKE [35], and MedVQA-2019

[1]), where VQA-RAD consists of 3,515 image-question pairs, SLAKE contains 14,028 image-question pairs and MedVQA-2019 contains 15,292 image-question pairs.

## 5.3. Implementation Details

**Pre-training** We adopt the classical VLP model as the backbone model, including a vision encoder, a language encoder, and a multi-modal fusion module. For the vision and language encoders, we adopt base-size Transformer encoders with 12 layers initialized from CLIP-ViT-B [48] RoBERTa-base [38] and their hidden dimension is set to 768. For the multi-modal fusion module, we set the number of Transformer layers to 6, the dimension of the hidden states to 768, and the number of heads to 12. For the visual/textual prompt pools, the dimension and the pool size is set to 768 and 1,024, respectively, by default. For optimization, the pre-training takes 100,000 steps with AdamW optimizer [39] with a weight decay of 0.01. The learning rates for the vision and language encoders and the remaining parameters are set to  $1e-5$  and  $5e-5$ , respectively. We use the warm-up strategy during the first 10% of the total number of steps, and the learning rate is linearly decayed to 0 after warm-up. For data augmentation, we use center-crop to resize each image to the size of  $288 \times 288$ .

**Fine-tuning** For all downstream tasks, we use the AdamW optimizer with the learning rate set to  $5e-6$  and  $2.5e-4$  for the backbone model and task-specific layers, respectively. The fine-tuning strategies can be divided into three categories according to the type of tasks. Specifically, for the classification tasks (*i.e.*, IC, TC, and VQA), we feed the concatenation of the image/visual prompt and text/textual prompt representations to a randomly initialized two-layer MLP to predict the labels. For the retrieval tasks (*i.e.*, ITR and TIR), we adopt the prediction head for the image-text contrast pre-text task and test its zero-shot and fine-tuned performance. For the generation tasks (*i.e.*, TS and ITG), we feed the concatenation of the sequence of image/visual prompt and text/textual prompt representations to a Transformer decoder with its parameters (except for the parameters of cross-attention layers) initialized from the pre-trained language encoder. For the evaluation metrics, we follow the previous studies to adopt AUROC for IC, accuracy for TC and VQA, Recall@K ( $K=1, 5, 10$ ) for ITR and TIR, and natural language generation (NLG) metrics (*i.e.*, BLEU [46], METEOR [13], CIDEr [55], and ROUGE [33]) for TS and ITG.

All pre-training and fine-tuning experiments are conducted on 80GB NVIDIA A100 GPUs with mixed-precision [41] to accelerate training and save memory. To demonstrate the effectiveness of the proposed approach, we compare it with previous studies, including ConVIRT [65], GLORIA [20], ClinicalBERT [2], IFCC [42], TransABS [37], WG-Sum [19], R2Gen [11], R2GenCMN [10], ViLT [26], ME-

Methods	Uni-Modal				Cross-Modal				Multi-Modal						
	Image		Text		Image-to-Text		Text-to-Image								
	CheXpert AUROC	PNAS AUROC	RadNLI Acc	MIMIC RL	MIMIC BL4	ROCO R@1	ROCO R@1	VQA-RAD Acc	SLAKE Acc	MedVQA-2019 Acc					
Study <sub>1</sub>	ConVIRT [65]	87.3	PNAS 81.3	ClinicalBERT [2]	72.6	TransABS [37]	43.8	R2Gen [11]	8.0	11.9	ViLT [26]	9.8	VQA-RAD 72.7	SLAKE 82.1	MedVQA-2019 -
Study <sub>2</sub>	GLoRIA [20]	88.1	88.6	IFCC [42]	77.8	WGSum [19]	45.1	M2Trans [42]	10.5	14.5	METER [17]	11.3	VQA-RAD 72.0	SLAKE -	MedVQA-2019 77.9
PTUnifier (ours)	90.1	90.6	80.0	46.2	10.7	21.0	20.8	78.3	85.2	79.3					

Table 2: Comparisons of our proposed method with previous studies on three types of evaluations (*i.e.*, uni-modal, cross-modal, and multi-modal evaluations). Study<sub>1</sub> and Study<sub>2</sub> denote two state-of-the-art approaches of each type of tasks, respectively. BL-4 denotes BLEU score using 4-grams and RG-L denotes ROUGE-L (same below). Dark and light grey colors highlight the top and second best results on each metric (same below). Note that the results of image-to-text generation are replicated using the front images only with the original reports as the reference instead of the tokenized ones.

TER [17], CPRD [34], and MMBERT [25]. Note that owing to the fact that there are few studies researching the medical cross-modal retrieval task, we replicate two state-of-the-art general-domain approaches on the medical image-text dataset. Besides, for image-to-text and text-to-text generation tasks, since we only use the chest X-ray images of the front views and different computation methods of metrics, we replicate existing studies that release their codes.<sup>6</sup>

## 6. Results and Analyses

### 6.1. Main Results

To demonstrate the effectiveness of the proposed approach, we conduct experiments on the aforementioned medical vision-language benchmark. The results of the main experiments are reported in Table 2. There are several observations. First, our approach achieves the best performance on all tasks. It outperforms previous studies on uni-modal image classification (+2.0% AUROC), text classification (+2.2% Accuracy), text summarization (+1.1% Rouge-L), image-to-text generation (+0.2% BLEU-4), image-to-text retrieval (+7.5% Recall@1), text-to-image retrieval (+9.5% Recall@1), and multi-modal VQA (+3.4% Accuracy), which confirms the validity of the proposed approach. Second, the proposed approach outperforms those complicated methods designed for specific tasks. For example, R2Gen introduced recurrent memory networks to the Transformer decoder to augment its decoding ability; WGSum used extra word graphs to improve the ability to detect keywords in the findings section of radiology reports. CPRD adopted representation distillation to alleviate the data scarcity problem in the Med-VQA task. These observations show that

<sup>6</sup>We use the original reports as the reference instead of the tokenized ones.

different pre-training ways can enable distinct abilities of the model, and it is possible to design an appropriate approach to exploit the knowledge shared across different tasks and perform various tasks using a unified model. Note that *the existing studies are only designed for a single task*, while our approach generally targets all vision- and/or language-related tasks, namely, without any tailored adaptations to a specific task.

### 6.2. Ablation Study

To further illustrate the effectiveness of our proposed approach, we perform an ablation study on the pre-training objectives, including the ones from fusion-encoders (*i.e.*, MLM and ITM) and the one from dual-encoders (*i.e.*, ITC).

There are several observations drawn from different aspects. First, the objectives of fusion encoders (*i.e.*, MLM and ITM) guide the models (*i.e.*, ID 3 and 5) to the more powerful multi-modal representations than other models without them, which could be observed from the performance on the downstream Med-VQA task. Second, the image-text contrast objective of dual encoders assists the models (*i.e.*, ID 4 and 5) in learning the uni-modal image representations and the cross-modal representations, and the models pre-trained with the ITC objective outperform those pre-trained without the ITC objective. More importantly, the models pre-trained with the ITC objective (*i.e.*, ID 4 and 5) demonstrate their great transfer ability where the pre-trained models can achieve high performance with very little data (*e.g.*, 1% and 10%). Third, it is interesting to note that the ITC objective does not promote the performance of the uni-modal text classification task. We can explain this phenomenon by the reason that images and texts are abstracted at different levels, where pixels of images have a lower semantic level than tokens of texts. Therefore, in the ITC process, the texts can be

ID	Objectives			Uni-Modal Image			Text	Cross-Modal Image-to-Text		Multi-Modal		
	MLM	ITM	ITC	CheXpert			RadNLI	MIMIC		VQA-RAD		Overall
				1% AUROC	10% AUROC	100% AUROC		Acc	BL4	CDr	Open Acc	
1	✓			66.1	79.1	81.1	77.2	6.9	11.1	57.5	79.5	70.8
2		✓		56.9	83.0	85.8	77.5	10.0	18.2	23.5	82.8	59.3
3	✓	✓		74.5	87.2	88.4	78.3	9.9	17.1	67.0	84.6	77.7
4			✓	88.0	88.9	89.3	76.5	10.3	19.0	64.8	81.0	74.6
5	✓	✓	✓	88.7	89.0	90.1	80.0	10.7	21.0	68.7	84.6	78.3

Table 3: Ablation studies on the different types of objectives, including the fusion-encoders ones (*i.e.*, masked language modeling (MLM) and image-text matching (ITM)) and the dual-encoders one (*i.e.*, image-text contrast (ITC)). 1%, 10%, and 100% represent the different portion of training data.

Pool Size	Pool Para.	MLM	ITM	ITC	Total
0	0	1.055	0.232	1.901	3.188
512	393.2k	1.053	0.215	1.787	3.055
1024	784.4k	1.049	0.211	1.778	3.038
2048	1.5m	1.057	0.229	1.861	3.147

Table 4: Pre-training losses (including MLM, ITM, and ITC) of our approach against different pool size, where the parameters of the pool (Pool Para.) are also shown.

treated as a kind of “supervision signals” for the learning of image encoding, yet, it is harder for the images to play such a role in contrast. This can be observed from previous studies [48, 23, 44], where the dual-encoders were only evaluated on the uni-modal vision tasks or cross-modal tasks. Fourth, performing both types of objectives promotes the model (*i.e.*, ID 5) to achieve the best performance across all the tasks, which confirms the feasibility of the research direction on unifying the fusion-encoders and dual-encoders.

### 6.3. Effects of Soft Prompt Pools

To analyze the impacts of soft prompt pools, we perform the pre-training with different pool sizes (ranging from 0 to 2048) with the results shown in Table 4, where we show the pre-training losses since they directly reflect how well the models perform the designed pretext tasks.

We have several observations: (i) Although enlarging pool size leads to increasing parameter numbers, it is demonstrated that there are not too many parameters (less than 0.5%) introduced compared with the total parameters (350M); (ii) All models with soft prompt pools have a better convergence (with a lower convergence loss) than the one without prompt pools (*i.e.*, pool size equal to 0), which demonstrates the effectiveness of introducing the prompt

pools; (iii) It is found that setting a proper pool size is important, where the model achieves the best convergence when the pool size is set to 1024. This might owe to the fact that the pool size controls how much the querying information is stored during the pre-training procedure, and a large pool size with a large capacity might “absorb” too much noise in the pre-training corpus.

## 7. Conclusion

In this paper, we proposed a simple yet effective scheme to take advantage of both fusion encoders and dual encoders, where visual and textual prompt pools are used to make our model compatible with different kinds of inputs (*i.e.*, image-only, text-only, and image-text-pair), and thus different types of objectives (*e.g.*, MLM and ITM for fusion-encoders and ITC for dual-encoders) can be adopted for pre-training. It is worth noting that our proposed approach is complementary to most of the existing Med-VLP models. Experimental results confirm the validity of our approach, where it achieves state-of-the-art performance on the downstream tasks. Further analyses investigate the effects of different types of objectives. Such empirical studies might provide a valuable reference for future research in this area.

**Limitation** The proposed vision-and-language approach is orthogonal to the domains (*e.g.*, the general domain, and the medical domain). However, limited by GPU resources, we do not perform the pre-training in the general domain. Instead, we conduct experiments in the important medical domain, which allows us pre-training the models with an academic budget. Nonetheless, we agree that it would be better to evaluate domain-agnostic approaches in the general domain to verify their generalization.



## Acknowledgements

This work was supported in part by the Chinese Key-Area Research and Development Program of Guangdong Province (2020B0101350001), in part by the Guangdong Basic and Applied Basic Research Foundation (NO. 2020B1515020048), in part by the National Natural Science Foundation of China (NO. 61976250), in part by the Shenzhen Science and Technology Program (NO. JCYJ20220530141211024, NO. JCYJ20220818103001002), in part by Shenzhen Key Research Project (C10120230151), in part by Shenzhen Doctoral Startup Funding (RCBS20221008093330065), in part by the Fundamental Research Funds for the Central Universities under Grant 22lgqb25 and in part by the Guangdong Provincial Key Laboratory of Big Data Computing, The Chinese University of Hong Kong, Shenzhen.

## References

- [1] Asma Ben Abacha, Sadid A. Hasan, Vivek Datla, Joey Liu, Dina Demner-Fushman, and Henning Müller. Vqa-med: Overview of the medical visual question answering task at imageclef 2019. In *CLEF*, 2019. 6
- [2] Emily Alsentzer, John Murphy, William Boag, Wei-Hung Weng, Di Jindi, Tristan Naumann, and Matthew McDermott. Publicly available clinical bert embeddings. In *Proceedings of the 2nd Clinical Natural Language Processing Workshop*, pages 72–78, 2019. 6, 7
- [3] Shruthi Bannur, Stephanie Hyland, Qianchu Liu, Fernando Perez-Garcia, Maximilian Ilse, Daniel C Castro, Benedikt Boecking, Harshita Sharma, Kenza Bouzid, Anja Thieme, et al. Learning to exploit temporal structure for biomedical vision-language processing. *arXiv preprint arXiv:2301.04558*, 2023. 1
- [4] Hangbo Bao, Wenhui Wang, Li Dong, Qiang Liu, Owais Khan Mohammed, Kriti Aggarwal, Subhojit Som, Songhao Piao, and Furu Wei. Vlmo: Unified vision-language pre-training with mixture-of-modality-experts. In *Advances in Neural Information Processing Systems*, 2022. 2, 3
- [5] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part I 16*, pages 213–229. Springer, 2020. 2, 3
- [6] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020. 2
- [7] Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. Uniter: Universal image-text representation learning. In *European conference on computer vision*, pages 104–120. Springer, 2020. 3
- [8] Zhihong Chen, Yuhao Du, Jinpeng Hu, Yang Liu, Guanbin Li, Xiang Wan, and Tsung-Hui Chang. Multi-modal masked autoencoders for medical vision-and-language pre-training. In *Medical Image Computing and Computer Assisted Intervention—MICCAI 2022: 25th International Conference, Singapore, September 18–22, 2022, Proceedings, Part V*, pages 679–689. Springer, 2022. 1, 2, 3
- [9] Zhihong Chen, Guanbin Li, and Xiang Wan. Align, reason and learn: Enhancing medical vision-and-language pre-training with knowledge. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 5152–5161, 2022. 2, 3
- [10] Zhihong Chen, Yaling Shen, Yan Song, and Xiang Wan. Cross-modal memory networks for radiology report generation. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 5904–5914, 2021. 2, 6
- [11] Zhihong Chen, Yan Song, Tsung-Hui Chang, and Xiang Wan. Generating radiology reports via memory-driven transformer. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*, pages 1439–1449, 2020. 2, 6, 7
- [12] Jaemin Cho, Jie Lei, Hao Tan, and Mohit Bansal. Unifying vision-and-language tasks via text generation. In *International Conference on Machine Learning*, pages 1931–1942. PMLR, 2021. 3
- [13] Michael Denkowski and Alon Lavie. Meteor 1.3: Automatic metric for reliable optimization and evaluation of machine translation systems. In *Proceedings of the sixth workshop on statistical machine translation*, pages 85–91, 2011. 6
- [14] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, 2019. 2, 4, 5
- [15] Shizhe Diao, Wangchunshu Zhou, Xinsong Zhang, and Jiawei Wang. Prefix language models are unified modal learners. *arXiv preprint arXiv:2206.07699*, 2022. 3
- [16] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2020. 5
- [17] Zi-Yi Dou, Yichong Xu, Zhe Gan, Jianfeng Wang, Shuohang Wang, Lijuan Wang, Chenguang Zhu, Zicheng Liu, Michael Zeng, et al. An empirical study of training end-to-end vision-and-language transformers. *arXiv preprint arXiv:2111.02387*, 2021. 3, 7
- [18] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9729–9738, 2020. 2
- [19] Jinpeng Hu, Jianling Li, Zhihong Chen, Yaling Shen, Yan Song, Xiang Wan, and Tsung-Hui Chang. Word graph guided

- summarization for radiology findings. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 4980–4990, 2021. 6, 7
- [20] Shih-Cheng Huang, Liyue Shen, Matthew P Lungren, and Serena Yeung. Gloria: A multimodal global-local representation learning framework for label-efficient medical image recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3942–3951, 2021. 1, 3, 6, 7
- [21] Zhicheng Huang, Zhaoyang Zeng, Bei Liu, Dongmei Fu, and Jianlong Fu. Pixel-bert: Aligning image pixels with text by deep multi-modal transformers. *arXiv preprint arXiv:2004.00849*, 2020. 3
- [22] Jeremy Irvin, Pranav Rajpurkar, Michael Ko, Yifan Yu, Silvana Ciurea-Ilcus, Chris Chute, Henrik Marklund, Behzad Haghgoo, Robyn Ball, Katie Shpanskaya, et al. Chexpert: A large chest radiograph dataset with uncertainty labels and expert comparison. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pages 590–597, 2019. 6
- [23] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *International Conference on Machine Learning*, pages 4904–4916. PMLR, 2021. 2, 8
- [24] Alistair EW Johnson, Tom J Pollard, Nathaniel R Greenbaum, Matthew P Lungren, Chih-ying Deng, Yifan Peng, Zhiyong Lu, Roger G Mark, Seth J Berkowitz, and Steven Horng. Mimic-cxr-jpg, a large publicly available database of labeled chest radiographs. *arXiv preprint arXiv:1901.07042*, 2019. 2, 6
- [25] Yash Khare, Viraj Bagal, Minesh Mathew, Adithi Devi, U Deva Priyakumar, and CV Jawahar. Mmbert: Multimodal bert pretraining for improved medical vqa. In *2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI)*, pages 1033–1036. IEEE, 2021. 2, 3, 7
- [26] Wonjae Kim, Bokyung Son, and Ildoo Kim. Vilt: Vision-and-language transformer without convolution or region supervision. In *International Conference on Machine Learning*, pages 5583–5594. PMLR, 2021. 3, 6, 7
- [27] Jason J Lau, Soumya Gayen, Asma Ben Abacha, and Dina Demner-Fushman. A dataset of clinically generated visual questions and answers about radiology images. *Scientific data*, 5(1):1–10, 2018. 6
- [28] Junnan Li, Ramprasaath Selvaraju, Akhilesh Gotmare, Shafiq Joty, Caiming Xiong, and Steven Chu Hong Hoi. Align before fuse: Vision and language representation learning with momentum distillation. *Advances in Neural Information Processing Systems*, 34, 2021. 3
- [29] Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. Visualbert: A simple and performant baseline for vision and language. *arXiv preprint arXiv:1908.03557*, 2019. 2, 3, 4
- [30] Xiujun Li, Xi Yin, Chunyuan Li, Pengchuan Zhang, Xiaowei Hu, Lei Zhang, Lijuan Wang, Houdong Hu, Li Dong, Furu Wei, et al. Oscar: Object-semantics aligned pre-training for vision-language tasks. In *European Conference on Computer Vision*, pages 121–137. Springer, 2020. 3
- [31] Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4582–4597, 2021. 4
- [32] Yikuan Li, Hanyin Wang, and Yuan Luo. A comparison of pre-trained vision-and-language models for multimodal representation learning across medical images and reports. In *2020 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, pages 1999–2004. IEEE, 2020. 1, 2, 3
- [33] Chin-Yew Lin. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81, 2004. 6
- [34] Bo Liu, Li-Ming Zhan, and Xiao-Ming Wu. Contrastive pre-training and representation distillation for medical visual question answering based on radiology images. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 210–220. Springer, 2021. 7
- [35] Bo Liu, Li-Ming Zhan, Li Xu, Lin Ma, Yan Yang, and Xiao-Ming Wu. Slake: A semantically-labeled knowledge-enhanced dataset for medical visual question answering. In *2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI)*, pages 1650–1654. IEEE, 2021. 6
- [36] Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *ACM Computing Surveys*, 55(9):1–35, 2023. 4
- [37] Yang Liu and Mirella Lapata. Text summarization with pre-trained encoders. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3730–3740, 2019. 6, 7
- [38] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019. 6
- [39] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019. 6
- [40] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. *Advances in neural information processing systems*, 32, 2019. 3
- [41] Paulius Micikevicius, Sharan Narang, Jonah Alben, Gregory Diamos, Erich Elsen, David Garcia, Boris Ginsburg, Michael Houston, Oleksii Kuchaiev, Ganesh Venkatesh, et al. Mixed precision training. In *International Conference on Learning Representations*, 2017. 6
- [42] Yasuhide Miura, Yuhao Zhang, Emily Tsai, Curtis Langlotz, and Dan Jurafsky. Improving factual completeness and consistency of image-to-text radiology report generation. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5288–5304, 2021. 6, 7

- [43] Jong Hak Moon, Hyungyung Lee, Woncheol Shin, Young-Hak Kim, and Edward Choi. Multi-modal understanding and generation for medical images and text via vision-language pre-training. *IEEE Journal of Biomedical and Health Informatics*, 26(12):6070–6080, 2022. 1, 2, 3
- [44] Norman Mu, Alexander Kirillov, David Wagner, and Saining Xie. Slip: Self-supervision meets language-image pre-training. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXVI*, pages 529–544. Springer, 2022. 2, 8
- [45] Philip Müller, Georgios Kaissis, Congyu Zou, and Daniel Rückert. Joint learning of localized representations from medical images and reports. *ArXiv preprint*, abs/2112.02889, 2021. 1, 3
- [46] 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*, pages 311–318, 2002. 6
- [47] Obioma Pelka, Sven Koitka, Johannes Rückert, Felix Nensa, and Christoph M Friedrich. Radiology objects in context (roco): a multimodal image dataset. In *Intravascular Imaging and Computer Assisted Stenting and Large-Scale Annotation of Biomedical Data and Expert Label Synthesis*, pages 180–189. Springer, 2018. 2, 5
- [48] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021. 2, 4, 6, 8
- [49] Alexey Romanov and Chaitanya Shivade. Lessons from natural language inference in the clinical domain. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1586–1596, 2018. 6
- [50] George Shih, Carol C Wu, Safwan S Halabi, Marc D Kohli, Luciano M Prevedello, Tessa S Cook, Arjun Sharma, Judith K Amorosa, Veronica Arteaga, Maya Galperin-Aizenberg, et al. Augmenting the national institutes of health chest radiograph dataset with expert annotations of possible pneumonia. *Radiology: Artificial Intelligence*, 1(1):e180041, 2019. 6
- [51] Amanpreet Singh, Ronghang Hu, Vedanuj Goswami, Guillaume Couairon, Wojciech Galuba, Marcus Rohrbach, and Douwe Kiela. Flava: A foundational language and vision alignment model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15638–15650, 2022. 2, 3
- [52] Weijie Su, Xizhou Zhu, Yue Cao, Bin Li, Lewei Lu, Furu Wei, and Jifeng Dai. Vi-bert: Pre-training of generic visual-linguistic representations. In *International Conference on Learning Representations*, 2019. 3
- [53] Sanjay Subramanian, Lucy Lu Wang, Ben Bogin, Sachin Mehta, Madeleine van Zuylen, Sravanthi Parasa, Sameer Singh, Matt Gardner, and Hannaneh Hajishirzi. Medicat: A dataset of medical images, captions, and textual references. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 2112–2120, 2020. 2, 5
- [54] Hao Tan and Mohit Bansal. Lxmert: Learning cross-modality encoder representations from transformers. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5100–5111, 2019. 3, 4
- [55] 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*, pages 4566–4575, 2015. 6
- [56] Fuying Wang, Yuyin Zhou, Shujun Wang, Varut Vardhanabhuti, and Lequan Yu. Multi-granularity cross-modal alignment for generalized medical visual representation learning. In *Advances in Neural Information Processing Systems*, 2022. 1, 3
- [57] 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 *International Conference on Machine Learning*, pages 23318–23340. PMLR, 2022. 3
- [58] Wenhui Wang, Hangbo Bao, Li Dong, Johan Bjorck, Zhiliang Peng, Qiang Liu, Kriti Aggarwal, Owais Khan Mohammed, Saksham Singhal, Subhojit Som, et al. Image as a foreign language: Beit pretraining for all vision and vision-language tasks. *arXiv preprint arXiv:2208.10442*, 2022. 3
- [59] Zifeng Wang, Zhenbang Wu, Dinesh Agarwal, and Jimeng Sun. Medclip: Contrastive learning from unpaired medical images and text. *arXiv preprint arXiv:2210.10163*, 2022. 1
- [60] Chaoyi Wu, Xiaoman Zhang, Ya Zhang, Yanfeng Wang, and Weidi Xie. Medklip: Medical knowledge enhanced language-image pre-training. *medRxiv*, pages 2023–01, 2023. 1
- [61] Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, et al. Google’s neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*, 2016. 5
- [62] Lewei Yao, Runhui Huang, Lu Hou, Guansong Lu, Minzhe Niu, Hang Xu, Xiaodan Liang, Zhenguo Li, Xin Jiang, and Chunjing Xu. Filip: Fine-grained interactive language-image pre-training. In *International Conference on Learning Representations*, 2022. 2
- [63] Fei Yu, Jiji Tang, Weichong Yin, Yu Sun, Hao Tian, Hua Wu, and Haifeng Wang. Ernie-vil: Knowledge enhanced vision-language representations through scene graphs. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 3208–3216, 2021. 3
- [64] 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*, pages 5579–5588, 2021. 3
- [65] Yuhao Zhang, Hang Jiang, Yasuhide Miura, Christopher D Manning, and Curtis P Langlotz. Contrastive learning of medical visual representations from paired images and text. *arXiv preprint arXiv:2010.00747*, 2020. 1, 3, 4, 6, 7