

Seeing What You Miss: Vision-Language Pre-training with Semantic Completion Learning

Yatai Ji^{1*} Rongcheng Tu^{2*} Jie Jiang^{2*} Weijie Kong² Chengfei Cai²
Wenzhe Zhao² Hongfa Wang² Yujiu Yang^{1†} Wei Liu^{2†}

¹Tsinghua University ²Tencent

jyt21@mails.tsinghua.edu.cn turongcheng@gmail.com

{zeus, jacobkong, fletchercai, carsonzhao, hongfawang}@tencent.com

yang.yujiu@sz.tsinghua.edu.cn wl2223@columbia.edu

Abstract

Cross-modal alignment is essential for vision-language pre-training (VLP) models to learn the correct corresponding information across different modalities. For this purpose, inspired by the success of masked language modeling (MLM) tasks in the NLP pre-training area, numerous masked modeling tasks have been proposed for VLP to further promote cross-modal interactions. The core idea of previous masked modeling tasks is to focus on reconstructing the masked tokens based on visible context for learning local-to-local alignment. However, most of them pay little attention to the global semantic features generated for the masked data, resulting in a limited cross-modal alignment ability of global representations. Therefore, in this paper, we propose a novel *Semantic Completion Learning* (SCL) task, complementary to existing masked modeling tasks, to facilitate global-to-local alignment. Specifically, the SCL task complements the missing semantics of masked data by capturing the corresponding information from the other modality, promoting learning more representative global features which have a great impact on the performance of downstream tasks. Moreover, we present a flexible vision encoder, which enables our model to perform image-text and video-text multimodal tasks simultaneously. Experimental results show that our proposed method obtains state-of-the-art performance on various vision-language benchmarks, such as visual question answering, image-text retrieval, and video-text retrieval.

1. Introduction

Our real-world contains a wide variety of information, such as texts, images, sounds, etc. For a powerful general artificial intelligence system, it is necessary to capture the se-

*Equal contribution.

†Corresponding Author.



Figure 1. (a) The comparisons between previous masked modeling tasks and our proposed Semantic Completion Learning (SCL), which is composed of “MVSC” and “MLSC”. (b) The cross-modal attention map visualization of the text global representation ([CLS]) on the input image for our model pre-trained with or without SCL.

mantic association from different modality sources. Towards this goal, multimodal representation learning is a crucial technique to bridge the heterogeneity gap between different modalities [6, 36]. In this area, vision-language pre-training models [14, 21, 30, 35, 43] have shown an impressive semantic alignment ability, which brings substantial advances on various downstream tasks, for instance, visual question answering, image-text retrieval, etc.

Recently, numerous self-supervised vision-language pre-training models [3, 10, 14, 17, 31, 32, 43] have been proposed. These methods model the interactions between vision and

language features mainly by using various masked modeling tasks, such as masked language modeling (MLM) and masked vision modeling (MVM). As shown in Fig. 1(a), the basic idea of MLM and MVM is self-reconstructing the masked tokens via leveraging informative visible tokens to realize local-to-local alignment. Specifically, MLM adopted by BERT [24] is to predict the original vocabulary IDs of the masked words. Inspired by the success of MLM in pre-training, there is a flourishing trend to extend it to visual pre-training tasks. Generally, by masking some visual patches, MVM tasks predict their original pixels [13, 17], corresponding discrete tokens [4, 10, 46] generated by the VQ-VAE variants, or Histograms of Oriented Gradients (HOG) features [11], etc.

These masked modeling tasks only focus on reconstructing the local masked tokens, and pay little attention to recovering the missing global semantic information caused by data corruption. The token-level reconstruction may lead to inadequate learning of global representations for cross-modal information. As illustrated in Fig. 1(b), in the situation of token-level reconstructions, the global representation is disordered in its attention on the other modality. It implies that the global-to-local alignment ability of the pre-training model is limited, leading to a degraded global representation. However, the global semantic features have a great impact on the performance of the pre-training model as they are usually used to deal with downstream tasks. Therefore, it is crucial to ensure the global semantic features to learn more accurate global-to-local alignment.

Intuitively, considering that the paired vision and text data are two views of the same semantic information, the missing semantics of masked data can be completed by capturing information from the other modality. From this point of view, we propose a novel pre-training task called Semantic Completion Learning (SCL). Specifically, SCL is composed of dual parts: masked vision semantic completion (MVSC) and masked language semantic completion (MLSC). As shown in Fig. 1(a), MVSC (MLSC) exploits information of complete text (vision) data to recover the global semantic representations of masked vision (text) data. In this way, the model can generate representative global features with accurate global-to-local alignment. For example, as illustrated in Fig. 1(b), compared with the model pre-trained without SCL, the attention maps with SCL pre-training are more discriminative and reasonable.

For the architecture of the vision-language pre-training model, we adopt a general framework that consists of two uni-modal encoders and a fusion encoder. Moreover, we present a flexible vision encoder to enable our model to perform image-text and video-text multimodal tasks simultaneously. Specifically, for video inputs, the vision encoder only adds a few additional learning parameters, and the [CLS] feature of each frame is treated as a bridge associating spatial

modeling within the frame and temporal modeling among frames. Inspired by curriculum learning [3], we train the model with image-text and video-text datasets successively to transfer visual knowledge from images to videos.

In a nutshell, our contributions are three-fold. (1) To enhance the global-to-local alignment of global representations, we propose a new pre-training task called Semantic Completion Learning (SCL), which recovers missing semantic information from unmasked data, promoting learning more representative global features. (2) We design an adaptive vision encoder, which can transfer multimodal pre-training knowledge between images and videos readily. (3) We conduct multiple vision-language downstream tasks to demonstrate the generalization of semantic completion learning, and the vision encoder, including visual question answering, visual reasoning, image-text retrieval, and video-text retrieval. Our model SCL achieves state-of-the-art performance based on a similar pre-training data scale. Our code is available at <https://github.com/IIGROUP/SCL>.

2. Related Works

2.1. Vision-Language Pre-training

Existing vision-language pre-training works can be divided into two categories: dual-tower and cross-fusion architecture.

The dual-tower architecture based methods [1–3, 12, 22, 39, 49] employ two individual encoders to separately extract the features for the visual data (images or videos) and textual data, and then map these features into a common semantic space. Among them, CLIP [39] exploits contrastive learning with a huge quantity of noisy image-text pairs directly collected from the Internet, achieving remarkable results on plenty of vision-language tasks. Similarly, FROZEN [3] proposes a curriculum learning schedule to train the vision-language model on both image-text and video-text datasets by treating an image as a single-frame video. Although these two-stream architecture based methods perform well on cross-modal retrieval tasks with high efficiency, their performances on the more complex multimodal downstream tasks are not inspirational due to the insufficient interaction between local vision and text features.

To overcome this limitation, the cross-fusion architecture based methods [5, 9, 30, 31, 43] have been proposed, which employ a cross-modal fusion encoder to enhance the interactions between vision and text features. For example, ALBEF [1] not only aligns the image and text features with contrastive learning but also feeds them into a cross-modal attention-based encoder to obtain the fused features. Clover [18] improves cross-modal feature alignment and fusion via a tri-modal alignment pre-training task. Our model also conducts multimodal feature fusion to achieve encouraging performance on more downstream tasks.

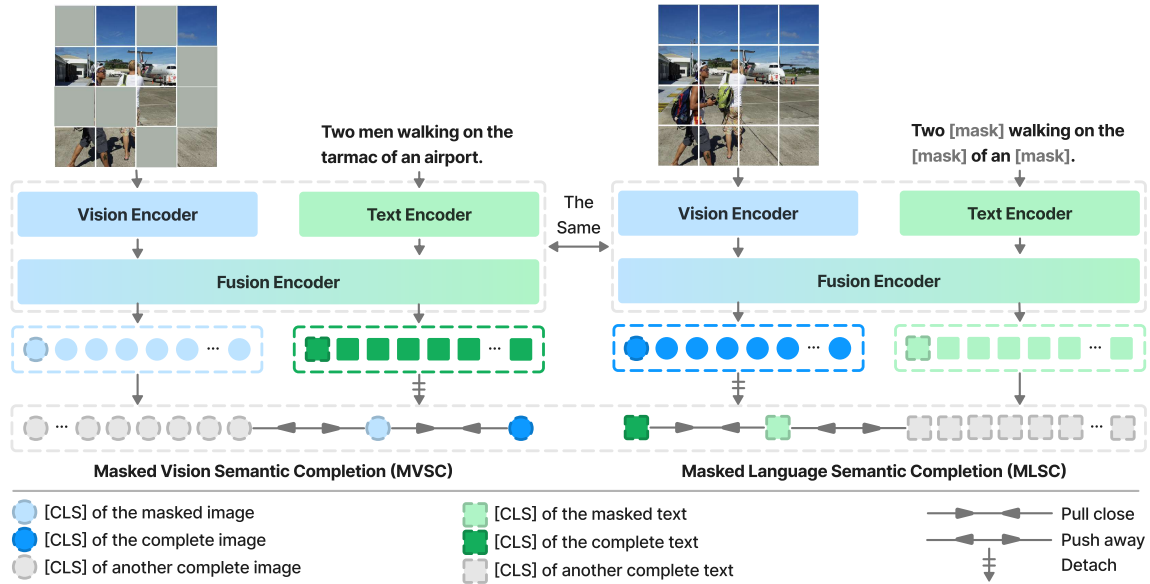


Figure 2. The overview of our proposed Semantic Completion Learning (SCL). The two versions of an image-text pair are forward propagated, respectively, to perform masked vision/language semantic completion.

2.2. Masked Modeling Tasks

Recently, various masked modeling tasks have been proposed, whose strategy is self-reconstructing the masked data. Masked Language modeling (MLM) adopted by BERT [24] is the most classical one. It randomly masks some tokens of the input and then predicts the original vocabulary IDs of the masked words based on their context. By pre-training with the MLM, BERT achieves state-of-the-art results on eleven natural language processing (NLP) tasks. Inspired by the success of MLM in NLP, some works extend it into the visual domain and propose masked vision modeling (MVM). For example, VLMAE [17] proposes the Regional Masked Image Modeling (RMIM) task to facilitate the fusion of multimodal features. The RMIM masks some patches of an input image and then reconstructs the original pixels depending on the visible patches and the corresponding text. Similarly, VIOLET [10] proposes a masked visual-token modeling task, which first maps the original video frame patches into discrete visual tokens and then recovers the corresponding visual tokens of masked patches to train a joint encoder for the vision-language fusion. However, these tasks focus on reconstructing local masked tokens, ignoring the recovery of global semantic information of the masked data after cross-modal interactions. Hence, we propose a novel semantic completion learning (SCL) task.

3. Approaches

In this section, we first introduce our pre-training objectives in Sec. 3.1, and then describe the model architecture in Sec. 3.2. Please refer to Appendix A for the whole architec-

ture figure and the details of previous pre-training tasks.

3.1. Pre-training Tasks

3.1.1 Previous Pre-training Tasks

Contrastive Learning (CL). The input images and texts are projected into vision and language embedding spaces with two uni-modal encoders, respectively. We utilize contrastive learning to adjust the positions of semantic features, enforcing the paired image-text features close and negative samples far apart. Then the token-wise fusion is employed for features of different modalities in the unified semantic space.

Vision-Text Matching (VTM). VTM aims to determine the correspondence of an image-text pair. The model conducts a binary classification on the concatenation of vision and text global representations generated by the fusion encoder, which contributes to the overall alignment of different modalities.

Masked Language Modeling (MLM). MLM was first used as a pretext task in natural language processing and was later introduced to multi-modal pre-training. Following the text tokens masked out with a probability of 15%, the model attempts to predict the original words based on visual information and textual context. The token-level reconstruction task plays an important role in the way that the model learns to associate linguistic words and visual entities, realizing local-to-local semantic alignment.

3.1.2 Semantic Completion Learning (SCL)

It is significant for the model to learn multi-modal information fusion, that is, to extract knowledge from the other modality. Instead of local information reconstruction in former masked modeling tasks, we expect that the model can also recover the global semantics of masked images or texts after cross-modal interaction.

As shown in Fig. 2, for each data pair, we first randomly mask the image and text separately to get $\{I_{mask}, T\}$ and $\{I, T_{mask}\}$, so that the masked one manages to learn semantic information from the other complete modality. Then the two couples of data are sent to the model respectively. The recovered features of masked data are obtained by leveraging information from the other modality to complete its missing semantic information:

$$\begin{aligned} I_{Re}, T_{Co} &= \text{Model}(I_{mask}, T), \\ I_{Co}, T_{Re} &= \text{Model}(I, T_{mask}), \end{aligned} \quad (1)$$

where I_{Re}, T_{Re} are recovered global features of masked data, and I_{Co}, T_{Co} refer to global features of complete data. Then we conduct masked vision semantic completion (MVSC) and masked language semantic completion (MLSC) simultaneously. Specifically, we bridge the gap between recovered global features and the complete ones in the form of contrastive learning. The InfoNCE loss is adopted to maximize the mutual information (MI) between two versions of the input data pair, $\{I_{mask}, T\}$ and $\{I, T_{mask}\}$:

$$\begin{aligned} \text{NCE}_V &= -\frac{1}{N} \sum_{i=1}^N \log \frac{\exp(s(I_{Re}^i, I_{Co}^i)/\tau)}{\sum_{n=1}^N \exp(s(I_{Re}^i, I_{Co}^n)/\tau)}, \\ \text{NCE}_L &= -\frac{1}{N} \sum_{i=1}^N \log \frac{\exp(s(T_{Re}^i, T_{Co}^i)/\tau)}{\sum_{n=1}^N \exp(s(T_{Re}^i, T_{Co}^n)/\tau)}, \end{aligned} \quad (2)$$

where s denotes cosine similarity and τ serves as the temperature hyper-parameter. The negative samples are global features of other complete images or texts in a batch. Note that I_{Co} and T_{Co} are detached for gradient backward, which makes the model more focused on the learning of recovering global features. Finally, the SCL loss is defined as:

$$\mathcal{L}_{SCL} = \text{NCE}_V + \text{NCE}_L. \quad (3)$$

By minimizing Eq.(3), it will make the global feature I_{Re} of the masked image similar to I_{Co} of the complete image (T_{Re} similar to T_{Co}). To recover the semantic information of masked data, the global representations will learn supplementary knowledge from the corresponding tokens of the other modality, i.e., accurate global-to-local alignment.

3.2. Model Architecture

Our model consists of three components: Vision Encoder, Text Encoder, and Fusion Encoder.

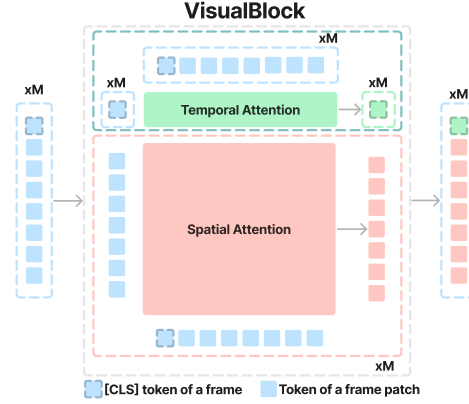


Figure 3. The architecture of a VisualBlock.

3.2.1 Vision Encoder

Input. The vision encoder takes visual data (a video or image) $I \in \mathcal{R}^{M \times 3 \times H \times W}$ containing M frame(s) of resolution $H \times W$ as input, and when I is an image, $M = 1$. The visual data I is first split into $M \times N$ patches $x \in \mathcal{R}^{M \times 3 \times N \times P \times P}$, where $P \times P$ is the size of patches and $N = HW/P^2$. Then, the patches x are transformed into M sequences of vision tokens $\mathbf{V} = \{v_i\}_{i=1}^M \in \mathcal{R}^{M \times N \times D}$, where $v_i \in \mathcal{R}^{N \times D}$ denotes the sequence of tokens for the i^{th} frame in the visual data and D denotes the dimension of vision tokens. Next, a learnable [CLS] token is concatenated to every token sequence v_i , and we obtain $\mathbf{V} = \{v_i\}_{i=1}^M \in \mathcal{R}^{M \times (N+1) \times D}$. Finally, the tokens \mathbf{V} are summed with learnable spatial positional embeddings $\mathbf{E}^s \in \mathcal{R}^{(N+1) \times D}$ and temporal positional embeddings $\mathbf{E}^t \in \mathcal{R}^{M \times D}$:

$$\mathbf{g}_{ij}^0 = v_{ij} + \mathbf{E}_i^t + \mathbf{E}_j^s, \quad (4)$$

where all patches in the same spatial location of different frames are given the same spatial positional embedding \mathbf{E}_j^s , and all patches in the same frame share the same temporal positional embedding \mathbf{E}_i^t .

VisualBlock. The pre-processed vision tokens $\mathbf{G}^0 = \{\mathbf{g}_{ij}^0\}_{i=1}^M \in \mathcal{R}^{M \times (N+1) \times D}$ are fed into the vision encoder which can process image and video data. The vision encoder is a modified ViT [8], containing a stack of VisualBlocks.

The detail of each VisualBlock is shown in Fig. 3. Specifically, each VisualBlock will perform temporal attention to exploit the global temporal information of the visual data, and conduct the spatial attention to capture sufficient local spatial semantic information. For temporal attention, we perform multi-head attention for the [CLS] tokens $\{\mathbf{g}_{i0}^{l-1}\}_{i=1}^M$ of all frames through attending to all $M \times (N+1)$ tokens to produce [CLS] tokens $\{\mathbf{g}_{i0}^l\}_{i=1}^M$. Regarding spatial attention, it is the multi-head attention within each frame. Taking the i^{th} frame as an example, we use $\{\mathbf{g}_{ij}^{l-1}\}_{j=1}^N$ without [CLS] token as queries and all the $N+1$ tokens in the frame as keys and values to conduct attention and obtain the output

Model	VQA2.0		NLVR2	
	test-dev	test-std	dev	test-p
<i>Pre-trained with >10M images</i>				
ALBEF(14M) [29]	75.84	76.04	82.55	83.14
SimVLM [48]	77.87	78.14	81.72	81.77
OFA [45]	78.0	78.1	-	-
BLIP [28]	78.25	78.32	82.15	82.24
<i>Pre-trained with <10M images</i>				
Oscar [31]	73.16	73.44	78.07	78.36
UNITER [5]	72.70	72.91	77.18	77.85
ViLT [25]	71.26	-	75.70	76.13
TCL [52]	74.90	74.92	80.54	81.33
VLMo [47]	76.64	76.89	82.77	83.34
METER [9]	77.68	77.64	82.33	83.05
Ours	78.72	78.78	83.63	84.27

Table 1. Performance comparison on VQA2.0 and NLVR2.

tokens $\{g_{ij}^l\}_{j=1}^N$. After the temporal attention and spatial attention are conducted, we concatenate the M [CLS] tokens $\{g_{i0}^l\}_{i=1}^M$ with the $M \times N$ tokens $\{\{g_{ij}^l\}_{j=1}^N\}_{i=1}^M$ of frame patches as the output of the VisualBlock, denoted as $G^l = \{g_i^l\}_{i=1}^M$.

3.2.2 Text Encoder

Given the input text T , we first tokenize it into word embeddings $\{t_i\}_{i=1}^K$, where K is the total number of words. Then, the text encoder, which consists of a stack of bidirectional transformers [44], maps $\{t_i\}_{i=1}^K$ into token features $W = \{w_i\}_{i=1}^K$ by modeling the contextual relationships.

3.2.3 Fusion Encoder

Similar to [9, 43], we adopt a two-stream architecture for the fusion encoder, each layer of which consists of two modality-specific self-attention blocks and two cross-attention blocks. Specifically, taking the vision features as an example, following intra-modal interactions in the visual self-attention block, we conduct cross-modal interactions in the language-to-vision cross-attention block, which takes the vision tokens $G = \{g_i\}_{i=1}^M$ as queries and the text tokens $W = \{w_i\}_{i=1}^K$ as keys and values. The text features are conducted with similar operations. In the end, we use the mean pooling of all the frame [CLS] tokens yielded by the fusion encoder as the global representation for visual data and the [CLS] token of text as the global representation for text data.

4. Experiments

4.1. Implementation Details

Following a recent line of works, we use COCO [33], Visual Genome (VG) [26], Conceptual Captions (CC3M) [41],

¹Pre-trained on 14M images.

Model	Flicker30K-ZS					
	IR@1	IR@5	IR@10	TR@1	TR@5	TR@10
<i>Evaluate pre-trained models directly</i>						
UNITER [5]	66.16	88.40	92.94	80.70	95.70	98.00
ViLT [25]	55.0	82.5	89.8	73.2	93.6	96.5
ALIGN [23]	75.70	93.80	96.80	88.60	98.70	99.70
METER [9]	79.60	94.96	97.28	90.90	98.30	99.50
Ours	79.74	95.46	97.86	91.70	99.30	99.90
<i>Evaluate models fine-tuned on COCO</i>						
ALBEF [29]	76.8	93.7	96.7	90.5	98.8	99.7
ALBEF ¹ [29]	82.8	96.3	98.1	94.1	99.5	99.7
TCL [52]	79.6	95.1	97.4	93.0	99.1	99.6
Ours	81.74	96.72	98.54	94.80	99.60	100.00

Table 2. Performance comparison of zero-shot image-text retrieval on Flickr30K.

and SBU Captions [37] for image-text pre-training, which contain 4M images in total. Then the pre-trained model is applied to image-text downstream tasks and the initialization for video-text pre-training. For the following video-text pre-training, we utilize WebVid [3] with 2.5M videos as the pre-training corpus. We employ an extensive set of evaluation benchmarks on a wide variety of vision-language understanding and retrieval tasks, including visual question answering (VQA2.0 [16]), visual reasoning (NLVR2 [42]), image-text retrieval (Flicker30K [38], COCO [33]), and video-text retrieval (MSRVTT [51], LSMDC [40]). We utilize CLIP-ViT-224/16 [39] and RoBERTa [34] to initialize vision and language encoders following METER [9]. For our proposed SCL, we adopt high mask ratios, 80% for images and 40% for texts. Other details of pre-training settings can be found in Appendix B.

4.2. Evaluation Results

4.2.1 Image-Text Understanding

We conduct multimodal understanding tasks on VQA2.0 and NLVR2, which require the model to exploit vision and language semantic fusion. As the results shown in Table 1, SCL achieves new state-of-the-art performance compared with previous models, implying that cross-modal fusion benefits from our new pre-training task. Specifically, when pre-trained with fewer than 10M images, our model outperforms METER [9] by +1.04 and +1.14 scores on VQA2.0 test-dev and test-std. On NLVR2, we gain +0.86 and +0.93 score improvements over VLMo [47], respectively. Moreover, our model pre-trained with 4M images also surpasses some models with more than 10M images, for instance, SimVLM [48], BLIP [28].

4.2.2 Image-Text Retrieval

We evaluate image-text retrieval in both zero-shot and fine-tuning scenarios. Our model achieves substantial perfor-

Model	COCO						Flickr30K					
	IR@1	IR@5	IR@10	TR@1	TR@5	TR@10	IR@1	IR@5	IR@10	TR@1	TR@5	TR@10
<i>Pre-trained with > 10M images</i>												
ALIGN [23]	59.9	83.3	89.8	77.0	93.5	96.9	84.9	97.4	98.6	95.3	99.8	100.0
ALBEF(14M) [29]	60.7	84.3	90.5	77.6	94.3	97.2	85.6	97.5	98.9	95.9	99.8	100.0
<i>Pre-trained with < 10M images</i>												
UNITER [5]	52.93	79.93	87.95	65.68	88.56	93.76	75.56	94.08	96.76	87.30	98.00	99.20
PixelBERT [20]	50.1	77.6	86.2	63.6	87.5	93.6	71.5	92.1	95.8	87.0	98.9	99.5
VinVL [53]	58.1	83.2	90.1	74.6	92.6	96.3	-	-	-	-	-	-
ALBEF(4M) [29]	56.80	81.50	89.20	73.10	91.40	96.00	82.80	96.70	98.40	94.30	99.40	99.80
SOHO [19]	50.6	78.0	86.7	66.4	88.2	93.8	72.5	92.7	96.1	86.5	98.1	99.3
TCL [52]	59.0	83.2	89.9	75.6	92.8	96.7	84.0	96.7	98.5	94.9	99.5	99.8
METER [9]	57.08	82.66	90.07	76.16	93.16	96.82	82.22	96.34	98.36	94.30	99.60	99.90
Ours	60.14	84.56	91.45	77.70	94.10	97.44	84.56	97.42	98.94	95.90	99.80	100.00

Table 3. Performance comparison of fine-tuned image-text retrieval on Flickr30K and COCO datasets.

Model	MSRVTT			LSMDC		
	R@1	R@5	R@10	R@1	R@5	R@10
<i>Fine-tune</i>						
VideoCLIP [49]	30.9	55.4	66.8	-	-	-
Frozen [3]	31.0	59.5	70.5	15.0	30.8	39.8
VIOLET [10]	34.5	63.0	73.4	16.1	36.6	41.2
ALPRO [27]	33.9	60.7	73.2	-	-	-
MCQ [14]	37.6	64.8	75.1	17.9	35.4	44.5
MILES [15]	37.7	63.6	73.8	17.8	35.6	44.1
Clover [18]	38.6	67.4	76.4	22.7	42.0	52.6
Ours	43.2	76.0	86.7	32.8	62.5	72.9
<i>Zero-shot</i>						
VideoCLIP [49]	10.4	22.2	30.0	-	-	-
Frozen [3]	18.7	39.6	51.6	9.3	22.0	30.1
VIOLET [10]	25.9	49.5	59.7	-	-	-
ALPRO [27]	24.1	44.7	55.4	-	-	-
MCQ [14]	26.0	46.4	56.4	12.2	25.9	32.2
MILES [15]	26.1	47.2	56.9	11.1	24.7	30.6
Clover [18]	25.8	49.6	60.1	13.8	28.1	38.3
Ours	30.9	54.4	65.0	17.2	32.4	39.1

Table 4. Performance comparison of text-to-video retrieval on MSRVTT and LSMDC.

mance improvements on Flickr30K and COCO datasets with similar model sizes and pre-training data scales. The results are also competitive with models pre-trained on larger datasets, such as ALIGN [23] and ALBEF [29]. In the fine-tuning phase, the model is trained with CL and VTM losses. During inference, for the sake of efficiency, we first filter top-k candidates with vision and language encoders and then compute VTM scores for ranking.

To investigate the generalization ability of our model, we conduct zero-shot experiments on the Flickr30K dataset. As shown in Table 2, SCL achieves the best performance in both settings of zero-shot retrieval on Flickr30K. When we evaluate with the pre-trained model directly, SCL gains

a comprehensive boost from previous methods, reaching 79.74% and 91.7% in terms of IR@1 and TR@1. When evaluated with the model fine-tuned on COCO, SCL outperforms models pre-trained on datasets of similar sizes, including ALBEF [29] and TCL [52]. Moreover, compared with ALBEF pre-trained on 14M images, SCL also has a more impressive performance in five out of six recall metrics, further demonstrating the effectiveness of our proposed strategies.

For fine-tuning experiments, our model surpasses previous models by a large margin, as shown in Table 3. TCL [52] has a distinguished retrieval performance with triple contrastive learning, which is cross-modal, intra-modal, and global-local. Compared with TCL, our method brings +1.14%/ + 0.56% IR@1 boost and +2.10%/ + 1.00% TR@1 on COCO and Flickr30K, respectively. It is worth noting that our model also has higher scores than ALIGN [23] with 1.8B image-text pairs pre-trained. Thanks to semantic completion learning, the global features capture more cross-modal information, leading to an encouraging performance on retrieval.

4.2.3 Video-Text Retrieval

Due to our adaptable vision encoder, the image-text pre-trained model can be readily transferred to video-text pre-training. We evaluate text-to-video retrieval on two popular datasets, MSRVTT and LSMDC, to prove the performance of the video pre-training model. Table 4 summarizes results under both fine-tuning and zero-shot settings. In the fine-tuning situation, compared with the previous SOTA model, SCL achieves notable performance improvements with +4.6% and +10.1% in R@1 on MSRVTT and LSMDC. When doing zero-shot retrieval, SCL also gains remarkable improvements over the existing methods with +4.8% and +3.4% in R@1 on MSRVTT and LSMDC, respectively. These results demonstrate that the knowledge of

Pre-training tasks	Flicker30K-ZS						VQA2.0	NLVR2	
	IR@1	IR@5	IR@10	TR@1	TR@5	TR@10	test-dev	dev	test-p
Ours (MLM+VTM+CL+SCL)	70.4	91.1	95.2	86.2	97.3	99.2	77.59	81.01	81.89
w/o MLM	67.6	89.5	94.4	84.7	97.5	99.3	75.46	78.93	79.41
w/o VTM	49.7	78.5	87.2	61.1	86.4	94.7	76.96	77.86	79.54
w/o CL	67.8	90.7	94.9	82.5	97.3	99.1	77.66	81.10	82.18
w/o SCL	67.0	89.4	94.3	80.0	96.5	99.4	77.31	80.30	81.46

Table 5. Ablation study of each pre-training task. Note that the model pre-trained without VTM can only conduct zero-shot retrieval with vision and text encoders without feature fusion, so the recall metrics are not comparable to the others.

Pre-training tasks	Flicker30K-ZS				VQA2.0	NLVR2
	IR@1	IR@5	TR@1	TR@5	test-dev	test-p
MLM+VTM+CL	67.0	89.4	80.0	96.5	77.31	81.46
+MLSC	67.0	90.0	80.5	97.1	77.60	81.88
+MVSC	69.5	90.5	85.5	98.3	77.50	81.69
+SCL	70.4	91.1	86.2	97.3	77.59	81.89

Table 6. Ablation study of MLSC and MVSC.

mask ratio		Flicker30K-ZS				VQA2.0
image	text	IR@1	IR@5	TR@1	TR@5	test-dev
0.7	0.4	68.8	90.6	84.0	97.7	77.51
0.8	0.3	69.5	90.5	83.8	97.6	77.63
0.8	0.4	70.4	91.1	86.2	97.3	77.59
0.8	0.5	68.8	90.5	81.9	97.4	77.63
0.9	0.4	68.6	90.6	82.8	97.3	77.57

Table 7. Effect of different image and text mask ratios.

Vision Encoder	Video Retrieval			Text Retrieval		
	R@1	R@5	R@10	R@1	R@5	R@10
Mean Pooling	23.0	45.6	56.5	22.5	46.4	54.9
Global CLS	22.5	45.9	55.1	22.4	45.7	55.5
Frame CLS	23.3	48.3	56.8	24.0	46.5	55.7

Table 8. Ablation study of the vision encoder on zero-shot retrieval of MSRVTT.

our SCL model learned from image-text data can be used to improve the performance of video-text retrieval tasks.

4.3. Ablation Studies

We conduct empirical ablation experiments on pre-training tasks and the vision encoder. Since pre-training is time-consuming, we use COCO and VG as pre-training datasets, which is also a common setting in previous works [19, 43, 50].

4.3.1 Different Pre-training Tasks

There are four pretext tasks in our method, including CL, VTM, MLM, and SCL. As summarized in Table 5, we explore the impact of each task on both retrieval and understanding datasets. The first row shows the results of our model with all pre-training tasks, and the second to fifth

rows reflect the effect of removing each task separately. According to the chart, we observe that the retrieval performance drops most due to the lack of SCL when conducting retrieval with feature fusion. Specifically, SCL brings +3.38% and +6.20% boost in IR@1 and TR@1 on F30K-ZS. The model without MLM loses 2.13% in the accuracy of VQA2.0, which indicates that MLM has a great effect on multimodal understanding tasks. As for NLVR2, VTM has a relatively large impact. However, contrastive learning is only effective for retrieval in our model, which is perhaps because the other three pre-training tasks have already learned cross-modal fusion sufficiently for understanding tasks. Overall, comparing the first row with the fifth row, the model with SCL makes progress on all downstream tasks, which demonstrates that the model learns more accurate cross-modal alignment to generate representative global features.

Furthermore, SCL comprises MVSC and MLSC, whose effects we showcase in Table 6. According to the first three rows, either MVSC or MLSC can improve the performance of downstream tasks. We find that MVSC has a superior impact on retrieval tasks, which is probably because it improves the robustness of visual information understanding. In VQA2.0 and NLVR2, MLSC plays a more important role. Additionally, when combining the two sub-tasks, our model performs better in most metrics, which indicates that they are in synergy.

4.3.2 Mask Ratio in SCL

As shown in Table 7, we observe that the mask ratios of image and text affect downstream tasks, especially on zero-shot retrieval. VQA2.0 is less sensitive to the mask ratio because the model has been fine-tuned with a large amount of data. Considering the second to fourth rows, when the image mask ratio is fixed, the model with a text mask ratio of 0.4 has almost the best performance. Moreover, when the text mask ratio is set to 0.4, the results of the image mask ratio of 0.8 are the highest. We speculate that when the mask ratio is lower, semantic completion will rely more on intra-modal information and lack learning across modalities, leading to inferior performance. When the mask ratio is too high, the

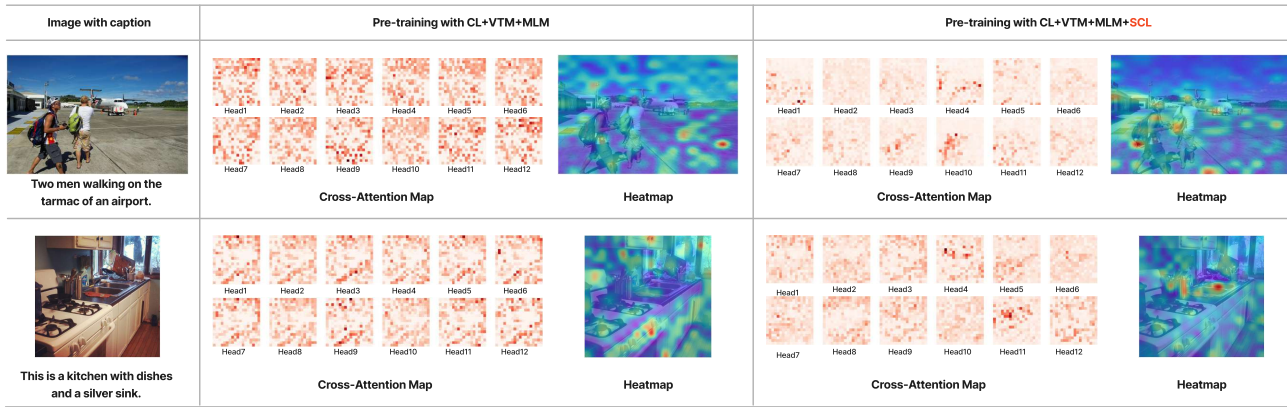


Figure 4. The cross-attention visualization of text [CLS] on the whole image for the model pre-trained with or without SCL. The cross-modal attention maps of 12 heads are from the last layer of the fusion encoder. Then we depict heatmaps by max-pooling the attention maps.

small number of remaining tokens can only perform very limited cross-modal interactions. In conclusion, we choose 0.4 and 0.8 as text and image mask ratios, respectively.

4.3.3 Vision Encoder Design

To investigate the effectiveness of our designed vision encoder in processing video data, we compare it with two other variants: (1) The first variant, termed Mean Pooling, directly treats a video as M separate images and then uses the mean pooling of M [CLS] tokens as the video representation. (2) The second variant, termed Global CLS, is the vision encoder proposed by MCQ [14]. In this experiment, we pre-train the vision encoder and text encoder on the WebVid [3] dataset via contrastive learning and then conduct zero-shot cross-modal retrieval on the MSRVT dataset. The experimental results are shown in Table 8, where the Frame CLS denotes our designed vision encoder. It can be found that Frame CLS achieves the best performance on both video-to-text and text-to-video retrieval tasks, which demonstrates the outstanding capability of video temporal modeling.

4.4. Visualization Analysis

To demonstrate that SCL boosts cross-modal alignment for global representations, we visualize cross-attention maps between text [CLS] and the whole image in the last layer of the fusion encoder. We conduct max-pooling on attention maps of 12 heads to draw heatmaps, as shown in Fig. 1(b) and Fig. 4. Compared with pre-trained by CL, VTM and MLM, the model with SCL can recognize relevant regions more precisely. For example, in the first image of Fig 1(b), the attention distribution of global text representation to the image is scattered without SCL, while after semantic completion learning, [CLS] pays attention to the fish, lemons, asparagus in the image. Taking the second image of Fig 4 as another example, the model pre-trained with SCL identifies

the dishes and sink in the kitchen, which indicates a desirable global-to-local alignment ability.

Observing attention maps of 12 heads in Fig. 4, we find that the attention maps without SCL are basically the same for an image, but for the model pre-trained with SCL, the attention maps of different heads are distinctive, which means that each head learns various information from the image. Overall, semantic completion learning encourages global representations to learn cross-modal interactions, extracting useful knowledge from the other modality. There are more visualization cases in Appendix C.

5. Conclusion

In this paper, we proposed a new vision-language pre-training task called Semantic Completion Learning (SCL). Different from previous pre-training tasks that reconstruct masked local tokens, SCL leverages cross-modal interactions to recover global semantic information of masked data, promoting cross-modal alignment for global representations. Ablation studies and visualization analysis demonstrate the effectiveness of SCL. Moreover, we introduced a flexible vision encoder, which adapts to image-text and video-text multimodal tasks readily. We conducted image-text and video-text pre-training sequentially and applied our model to various challenging downstream tasks. The extensive evaluations validate the great superiority of our SCL method.

6. Acknowledgements

This work was partly supported by the National Natural Science Foundation of China (Grant No. 61991450) and the Shenzhen Science and Technology Program (JSGG20220831093004008; ZDSYS20200811142605016).

References

- [1] Hassan Akbari, Liangzhe Yuan, Rui Qian, Wei-Hong Chuang, Shih-Fu Chang, Yin Cui, and Boqing Gong. VATT: transformers for multimodal self-supervised learning from raw video, audio and text. In *NeurIPS*, 2021. 2
- [2] Alex Andonian, Shixing Chen, and Raffay Hamid. Robust cross-modal representation learning with progressive self-distillation. In *CVPR*, 2022. 2
- [3] Max Bain, Arsha Nagrani, Gül Varol, and Andrew Zisserman. Frozen in time: A joint video and image encoder for end-to-end retrieval. In *ICCV*, 2021. 1, 2, 5, 6, 8
- [4] Hangbo Bao, Wenhui Wang, Li Dong, and Furu Wei. Vl-beit: Generative vision-language pretraining. *CoRR*, abs/2206.01127, 2022. 2
- [5] Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. Uniter: Universal image-text representation learning. In *ECCV*, 2020. 2, 5, 6
- [6] Xing Cheng, Hezheng Lin, Xiangyu Wu, Fan Yang, and Dong Shen. Improving video-text retrieval by multi-stream corpus alignment and dual softmax loss. *CoRR*, abs/2109.04290, 2021. 1
- [7] Ekin Dogus Cubuk, Barret Zoph, Jonathon Shlens, and Quoc Le. Randaugment: Practical automated data augmentation with a reduced search space. In *NeurIPS*, 2020. 12
- [8] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021. 4
- [9] Zi-Yi Dou, Yichong Xu, Zhe Gan, Jianfeng Wang, Shuohang Wang, Lijuan Wang, Chenguang Zhu, Pengchuan Zhang, Lu Yuan, Nanyun Peng, Zicheng Liu, and Michael Zeng. An empirical study of training end-to-end vision-and-language transformers. In *CVPR*, 2022. 2, 5, 6, 11
- [10] Tsu-Jui Fu, Linjie Li, Zhe Gan, Kevin Lin, William Yang Wang, Lijuan Wang, and Zicheng Liu. VIOLET : End-to-end video-language transformers with masked visual-token modeling. *CoRR*, abs/2111.12681, 2021. 1, 2, 3, 6, 12
- [11] Tsu-Jui Fu, Linjie Li, Zhe Gan, Kevin Lin, William Yang Wang, Lijuan Wang, and Zicheng Liu. An empirical study of end-to-end video-language transformers with masked visual modeling. *CoRR*, abs/2209.01540, 2022. 2
- [12] Valentin Gabeur, Chen Sun, Karteek Alahari, and Cordelia Schmid. Multi-modal transformer for video retrieval. In *ECCV*, 2020. 2
- [13] Hongyu Gao, Chao Zhu, Mengyin Liu, Weibo Gu, Hongfa Wang, Wei Liu, and Xu-Cheng Yin. Calic: Accurate and efficient image-text retrieval via contrastive alignment and visual contexts modeling. In *ACM Multimedia*, 2022. 2
- [14] Yuying Ge, Yixiao Ge, Xihui Liu, Dian Li, Ying Shan, Xiaohu Qie, and Ping Luo. Bridging video-text retrieval with multiple choice questions. In *CVPR*, 2022. 1, 6, 8, 12
- [15] Yuying Ge, Yixiao Ge, Xihui Liu, Alex Jinpeng Wang, Jianping Wu, Ying Shan, Xiaohu Qie, and Ping Luo. MILES: visual BERT pre-training with injected language semantics for video-text retrieval. *CoRR*, abs/2204.12408, 2022. 6
- [16] Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. Making the V in VQA matter: Elevating the role of image understanding in Visual Question Answering. In *Proc. of CVPR*, 2017. 5
- [17] Sunan He, Taian Guo, Tao Dai, Ruizhi Qiao, Chen Wu, Xijun Shu, and Bo Ren. VLMAE: vision-language masked autoencoder. *CoRR*, abs/2208.09374, 2022. 1, 2, 3
- [18] Jingjia Huang, Yinan Li, Jiashi Feng, Xiaoshuai Sun, and Rongrong Ji. Clover: Towards A unified video-language alignment and fusion model. *CoRR*, abs/2207.07885, 2022. 2, 6, 12
- [19] Zhicheng Huang, Zhaoyang Zeng, Yupan Huang, Bei Liu, Dongmei Fu, and Jianlong Fu. Seeing out of the box: End-to-end pre-training for vision-language representation learning. In *CVPR*, 2021. 6, 7
- [20] Zhicheng Huang, Zhaoyang Zeng, Bei Liu, Dongmei Fu, and Jianlong Fu. Pixel-bert: Aligning image pixels with text by deep multi-modal transformers. *CoRR*, 2020. 6
- [21] Yatai Ji, Junjie Wang, Yuan Gong, Lin Zhang, Yanru Zhu, Hongfa Wang, Jiaying Zhang, Tetsuya Sakai, and Yujiu Yang. MAP: modality-agnostic uncertainty-aware vision-language pre-training model. *CoRR*, abs/2210.05335, 2022. 1
- [22] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc V. Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *ICML*, 2021. 2
- [23] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc V. Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *ICML*, 2021. 5, 6
- [24] Jacob Devlin Ming-Wei Chang Kenton and Lee Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*, 2019. 2, 3, 11
- [25] Wonjae Kim, Bokyung Son, and Ildoo Kim. Vilt: Vision-and-language transformer without convolution or region supervision. In *ICML*, 2021. 5
- [26] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A. Shamma, Michael S. Bernstein, and Li Fei-Fei. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *Int. J. Comput. Vis.*, 2017. 5
- [27] Dongxu Li, Junnan Li, Hongdong Li, Juan Carlos Niebles, and Steven C. H. Hoi. Align and prompt: Video-and-language pre-training with entity prompts. In *CVPR*, 2022. 6
- [28] Junnan Li, Dongxu Li, Caiming Xiong, and Steven C. H. Hoi. BLIP: bootstrapping language-image pre-training for unified vision-language understanding and generation. In *ICML*, 2022. 5
- [29] 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. *NeurIPS*, 2021. 5, 6

- [30] Wei Li, Can Gao, Guocheng Niu, Xinyan Xiao, Hao Liu, Jiachen Liu, Hua Wu, and Haifeng Wang. UNIMO: towards unified-modal understanding and generation via cross-modal contrastive learning. In *ACL*, 2021. 1, 2
- [31] Xiujun Li, Xi Yin, Chunyuan Li, Pengchuan Zhang, Xiaowei Hu, Lei Zhang, Lijuan Wang, Houdong Hu, Li Dong, Furu Wei, Yejin Choi, and Jianfeng Gao. Oscar: Object-semantics aligned pre-training for vision-language tasks. In *ECCV*, 2020. 1, 2, 5
- [32] Kevin Qinghong Lin, Alex Jinpeng Wang, Mattia Soldan, Michael Wray, Rui Yan, Eric Zhongcong Xu, Difei Gao, Rong-Cheng Tu, Wenzhe Zhao, Weijie Kong, Chengfei Cai, Hongfa Wang, Dima Damen, Bernard Ghanem, Wei Liu, and Mike Zheng Shou. Egocentric video-language pretraining. *CoRR*, abs/2206.01670, 2022. 1
- [33] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014. 5, 12
- [34] 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. *CoRR*, 2019. 5, 11
- [35] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In *NeurIPS*, 2019. 1
- [36] Jiasen Lu, Christopher Clark, Rowan Zellers, Roozbeh Motlaghi, and Aniruddha Kembhavi. Unified-io: A unified model for vision, language, and multi-modal tasks. *CoRR*, abs/2206.08916, 2022. 1
- [37] Vicente Ordonez, Girish Kulkarni, and Tamara L. Berg. Im2text: Describing images using 1 million captioned photographs. In *NeurIPS*, 2011. 5
- [38] Bryan A. Plummer, Liwei Wang, Chris M. Cervantes, Juan C. Caicedo, Julia Hockenmaier, and Svetlana Lazebnik. Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. In *ICCV*, 2015. 5, 12
- [39] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *ICML*, 2021. 2, 5, 11
- [40] Anna Rohrbach, Marcus Rohrbach, Niket Tandon, and Bernt Schiele. A dataset for movie description. In *CVPR*, 2015. 5, 12
- [41] Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *ACL*, 2018. 5
- [42] Alane Suhr, Stephanie Zhou, Ally Zhang, Iris Zhang, Huajun Bai, and Yoav Artzi. A corpus for reasoning about natural language grounded in photographs. In *ACL*, 2019. 5
- [43] Hao Tan and Mohit Bansal. LXMERT: learning cross-modality encoder representations from transformers. In *EMNLP*, 2019. 1, 2, 5, 7
- [44] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017. 5
- [45] Peng Wang, An Yang, Rui Men, Junyang Lin, Shuai Bai, Zhikang Li, Jianxin Ma, Chang Zhou, Jingren Zhou, and Hongxia Yang. OFA: unifying architectures, tasks, and modalities through a simple sequence-to-sequence learning framework. In *ICML*, 2022. 5
- [46] Wenhui Wang, Hangbo Bao, Li Dong, Johan Bjorck, Zhiliang Peng, Qiang Liu, Kriti Aggarwal, Owais Khan Mohammed, Saksham Singhal, Subhojit Som, and Furu Wei. Image as a foreign language: Beit pretraining for all vision and vision-language tasks. *CoRR*, abs/2208.10442, 2022. 2
- [47] Wenhui Wang, Hangbo Bao, Li Dong, and Furu Wei. Vlmo: Unified vision-language pre-training with mixture-of-modality-experts. *CoRR*, 2021. 5
- [48] Zirui Wang, Jiahui Yu, Adams Wei Yu, Zihang Dai, Yulia Tsvetkov, and Yuan Cao. Simvlm: Simple visual language model pretraining with weak supervision. In *ICLR*, 2022. 5
- [49] Hu Xu, Gargi Ghosh, Po-Yao Huang, Dmytro Okhonko, Armen Aghajanyan, Florian Metze, Luke Zettlemoyer, and Christoph Feichtenhofer. Videoclip: Contrastive pre-training for zero-shot video-text understanding. In *EMNLP*, 2021. 2, 6
- [50] Haiyang Xu, Ming Yan, Chenliang Li, Bin Bi, Songfang Huang, Wenming Xiao, and Fei Huang. E2E-VLP: end-to-end vision-language pre-training enhanced by visual learning. In *ACL*, 2021. 7
- [51] Jun Xu, Tao Mei, Ting Yao, and Yong Rui. MSR-VTT: A large video description dataset for bridging video and language. In *CVPR*, 2016. 5, 12
- [52] Jinyu Yang, Jiali Duan, Son Tran, Yi Xu, Sampath Chanda, Liquan Chen, Belinda Zeng, Trishul Chilimbi, and Junzhou Huang. Vision-language pre-training with triple contrastive learning. In *CVPR*, 2022. 5, 6
- [53] 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 *CVPR*, 2021. 6