

Probing Cross-modal Semantics Alignment Capability from the Textual Perspective

Zheng Ma Shi Zong Mianzhi Pan Jianbing Zhang*

Shujian Huang Xinyu Dai Jiajun Chen

Nanjing University

{maz, panmz}@smail.nju.edu.cn

{szong, zjb, daixinyu, huangsj, chenjj}@nju.edu.cn

Abstract

In recent years, vision and language pre-training (VLP) models have advanced the state-of-the-art results in a variety of cross-modal downstream tasks. Aligning cross-modal semantics is claimed to be one of the essential capabilities of VLP models. However, it still remains unclear about the inner working mechanism of alignment in VLP models. In this paper, we propose a new probing method that is based on image captioning to first empirically study the cross-modal semantics alignment of VLP models. Our probing method is built upon the fact that given an image-caption pair, the VLP models will give a score, indicating how well two modalities are aligned; maximizing such scores will generate sentences that VLP models believe are of good alignment. Analyzing these sentences thus will reveal in what way different modalities are aligned and how well these alignments are in VLP models. We apply our probing method to five popular VLP models, including UNITER, ROSITA, ViLBERT, CLIP, and LXMERT, and provide a comprehensive analysis of the generated captions guided by these models. Our results show that VLP models (1) focus more on just aligning objects with visual words, while neglecting global semantics; (2) prefer fixed sentence patterns, thus ignoring more important textual information including fluency and grammar; and (3) deem the captions with more visual words are better aligned with images. These findings indicate that VLP models still have weaknesses in cross-modal semantics alignment and we hope this work will draw researchers' attention to such problems when designing a new VLP model.¹

1 Introduction

Vision-language pre-training (VLP) models are designed to understand visual information, textual

* Corresponding author.

¹Our code is publicly available at https://github.com/aaronma2020/probing_vlp

Which caption matches the image better?



- (a) a bunch of bananas on a table
- (b) a bunch of the bananas table
- (c) a closeup image that is a a bunch but a a bananas but a a table
- (d) bunch that the bananas at the table
- (e) a image image of a bunch near a bananas near a table
- (f) a bunch of bananas of a table

Models	Choice
UNITER	(b)
ROSITA	(c)
ViLBERT	(d)
CLIP	(e)
LXMERT	(f)
Human	(a)

Table 1: Given an image, VLP models are asked to choose from a reasonable caption (a) generated by a captioning model and five unreadable captions (b) - (f) modified from (a). We observe that all VLP models pick unreadable captions.

semantics, and cross-modal relationships. To align cross-modal semantics, most of these VLP models follow two structures: The single-stream architecture (Chen et al., 2019; Cui et al., 2021) aligns the cross-modal information as they start being fed into the model; while alignment happens in later layers for the two-stream architecture (Lu et al., 2019; Radford et al., 2021; Tan and Bansal, 2019). Besides, both types of models are generally trained with a cross-modality matching task during pre-training (Chen et al., 2019) or a contrastive loss for better alignment (Li et al., 2021; Yao et al., 2021). These models have achieved state-of-the-art in many cross-modal tasks, such as image-text retrieval (Lee et al., 2018) and visual question answering (Antol et al., 2015).

Although VLP models have excellent performance on a large panoply of cross-modal tasks, it remains unclear which information VLP models have aligned between different modalities. Some studies have investigated this problem and designed a series of probing tasks to explore it (Parcalabescu et al., 2020; Lindström et al., 2020). However,

these probing tasks almost probe VLP models based on the classification task, such as classifying or counting objects and recognizing consistencies between regions and phrases. We argue that the simple classified probing tasks fail to explicitly explore VLP models’ inner alignment mechanism, as they only require models to use partial information, such as a region, a phrase, or both, to make a decision. To illustrate this point, in Table 1, we present a reasonable caption and five unreadable captions. VLP models are asked which caption fits the image better. Surprisingly, no model chooses reasonable caption (a). It is thus reasonable to question the capability of VLP models for cross-modal alignment and motivates us to study which kind of captions VLP models deem match images better.

Without doubt, it is too difficult and nearly impossible to find out all unreasonable captions the models favor, by first generating some manually and then testing them, as in fact, we do not know the model’s preferences beforehand. However, if we can use the matching score from VLP models as a signal to train a cross-modal generated model, then maximizing such scores will amplify the models’ preferences and reflect them in the generated sentences. Inspired by this idea, in this paper we present the first work of using a novel probing method to empirically study cross-modal alignment via the image captioning task. The generated captions from our probing method have higher matching scores, which are considered to fit images better and contain important alignment information that the VLP models focus on.

We then apply our probing method to five powerful and representative VLP models, including UNITER, ROSITA, ViLBERT, CLIP, and LXMERT, which cover two mainstream architectures and two mainstream alignment tasks. Our approach explicitly maps the alignment information to the generated captions. Through analyzing these captions, we discover that objects and visual words tend to receive more attention in VLP models’ inner alignment mechanism, and the more nouns the captions contain, the more VLP models deem they match images. It indicates these VLP models are overwhelmingly dependent on partial information (objects in images and visual words in captions), rather than the whole semantics of images and captions when they judge whether image-caption pairs are aligned. Moreover, we find that these models favor certain sentence patterns (details in Section 4).

Captions that follow these patterns will be considered more consistent with images. It suggests when deciding whether image-caption pairs are aligned, VLP models ignore more important textual information, such as fluency and grammar.

2 Related Work

There has been an increasing number of papers studying how different modalities are aligned in VLP models. We now review recent representative studies. Cao et al. (2020) report the dominance of textual modality with VALUE, a comprehensive framework they introduce including multiple probing tasks. Lindström et al. (2020) design three probing tasks to analyze linguistic properties of multi-modal embeddings, such as estimating the number of object instances in the image. All these tasks employ a simple neural network classifier to probe the ability of VLP models in certain aspects or reveal the importance of textual compared with visual information.

Parcalabescu et al. (2020) evaluate VLP models on count, a task requiring the model to correctly predict the number of objects in an image, and find several prevailing VLP models that fail to identify entities in an image. Moreover, Parcalabescu et al. (2022) propose VALSE to test VLP models for their vision-linguistic grounding capabilities on specific linguistic phenomena and find that VLP models struggle to ground their interdependence and relationships in visual scenes when forced to respect linguistic indicators. It is also suggested that more targeted investigations are needed to probe the cross-modal alignment capacity of VLP models. Instead of the previous simple classification-based method, in this work, we introduce a new probing method based on the image captioning model, targeting for the cross-modal alignment. As we will discuss in the following sections, our approach is able to generate sentences that enable a more transparent analysis towards the alignment of visual and textual modality.

3 Our Probing Method

3.1 Probing Method Overview

VLP models are trained to capture the relationship between different modalities. They are able to score the matching of an image-caption pair, which can reflect whether the cross-modal semantics are aligned. Motivated by this, in this paper we

propose a new probing method to probe the cross-modal semantics alignment capability. Specifically, we first train a captioning model to generate a caption S for a given image I . This image-caption pair (I, S) is then fed into a certain VLP model to get an alignment score r (referred to as **VLP score**), which indicates the degree of alignment between images and captions from the VLP model. Finally, the captioning model is adjusted according to this feedback score. By continuously updating the model based on the above process, we are able to gradually generate captions with higher alignment scores from VLP models. We can then probe the alignment capability of VLP models, by simply evaluating the quality of these generated captions.

Our probing method needs to feed the scores back to the captioning model, guiding it to generate captions with higher scores. However, because of the non-differentiable problem, these scores can not be used directly to optimize the captioning model. To address this problem, we use self-critical sequence training (SCST; Rennie et al. (2017)), which is a standard method in image captioning task and has been widely used (Anderson et al., 2018; Huang et al., 2019).

Self-critical Sequence Training (SCST). We now briefly introduce self-critical sequence training (Rennie et al., 2017), a two-stage training method based on reinforcement learning. Given an image-text pair (I, S) , with $S = (s_0, s_1, \dots, s_t)$, in the first stage, the base model tries to minimize the cross-entropy loss (CE training):

$$L(\theta) = - \sum_{t=1}^T \log(p_{\theta}(s_t | I, s_1, \dots, s_{t-1})). \quad (1)$$

In the second state, SCST adopts two search strategies (greedy and sample) to generate sentences and computes the difference of a particular metric between two sentences as a reward to optimize the model. The goal of training is to minimize the negative expected reward:

$$L(\theta) = -\mathbb{E}_{S_{\text{sample}} \sim p_{\theta}} [r(S_{\text{sample}}) - r(S_{\text{greedy}})], \quad (2)$$

where S_{sample} is the sample sentence, and S_{greedy} is the greedy sentence. In Rennie et al. (2017), r is a type of image captioning metric. In our task, r is the matching score of VLP models.

3.2 Evaluated VLP Models

We conduct analyses of the following five VLP models, all of which have lifted the state-of-the-art

results across various vision-language tasks.

UNITER. Chen et al. (2019) propose word-region alignment via optimal transport. This task and the use of a conditional masking strategy during pre-training greatly enhance the fine-grained alignment capacity of UNITER.

ROSITA. Cui et al. (2021) adopt an elaborate pre-training task for fine-grained alignment of different modalities. It modifies commonly-used masked language modeling and masked region modeling to structural knowledge masking, an innovative masking strategy based on the unified vision-language scene graph.

ViLBERT. Lu et al. (2019) first introduce the two-stream architecture where image and text are encoded by two independent transformers and further fused by a co-attention mechanism.

LXMERT. Tan and Bansal (2019) also explore two-stream architecture. Compared with ViLBERT, LXMERT modifies cross-modal co-attention layers and introduces extra pre-training tasks, like ROI-feature regression and image question answering.

CLIP. Radford et al. (2021) use contrastive learning to fuse visual and textual features after they are encoded separately. This simple designed task renders powerful zero-shot transfer ability to CLIP across a wide range of downstream tasks, such as optical character recognition, action recognition, and text retrieval.

3.3 Experimental Setup

3.3.1 Training an Image Captioning Model

Dataset. To train an image captioning model, we use MSCOCO dataset (Lin et al., 2014). We follow the data split in Karpathy and Fei-Fei (2015) and divide the dataset into 113,287 images for training, 5,000 for validation, and 5,000 for test. Each image has at least five reference captions. We count all words in captions, drop the words with a frequency less than or equal to five, and finally keep 9,487 words to build a vocabulary. In Section 4 and Section 5, we will use this test set for analysis.

Captioning models. An image captioning model is essential in our probing method, as we rely on the generated captions from the image caption model to analyze the potential problems of the VLP models. Current prevailing approaches are based on deep neural networks. For example, Vinyals et al. (2015) feeds image features to LSTM-based lan-

Model	VLP Model Score					Image Captioning Metrics					
	UNITER	ROSITA	ViLBERT	CLIP	LXMERT	Bleu1	Bleu4	METEOR	ROUGE	CIDEr	SPICE
UNITER	90.4↑	75.4	89.5↑	27.7	88.2↑	44.5	6.0	14.3	35.8	34.9	9.7
ROSITA	85.7↑	97.3↑	89.2↑	26.8	89.1↑	25.2	1.0	13.8	26.7	5.3	9.1
ViLBERT	59.8	80.2	97.3↑	26.0	75.4	26.0	2.5	14.4	24.8	7.5	12.0
CLIP	55.0	86.0	85.5↑	32.1↑	79.7	31.2	3.9	16.4	31.0	10.3	10.3
LXMERT	73.9↑	90.7↑	84.6	27.2	93.6↑	30.0	3.3	16.1	31.2	9.3	11.3
CE	71.6	86.5	84.6	27.8	80.1	72.2	28.7	24.4	52.4	92.0	17.4

Table 2: Results on scores of all VLP models and image captioning metrics. CE means training the FC model with the cross-entropy loss and other models are trained under the SCST framework. Symbol \uparrow means training with a certain VLP model has an improvement on other VLP models than using the cross-entropy loss. Because CLIP is different from other VLP models in measuring an image-caption pair, the scale of its score is also different.

guage models, and various attention mechanisms are incorporated to generate better captions (Xu et al., 2015; Lu et al., 2017; Anderson et al., 2018). In this work, we use FC model (Rennie et al., 2017) as our image captioning model, which is similar to Vinyals et al. (2015).²

We train the FC model with cross-entropy loss (referred to as CE) for 30 epochs, using Adam (Kingma and Ba, 2015) optimizer with the learning rate of $5e-4$. We anneal the learning rate by 0.8 every three epochs and increase the probability of feeding back a sample of the word posterior by 0.05 every five epochs. We evaluate the model on the development set every 3,000 steps and select the model with the best CIDEr score as the initialized caption model.

3.3.2 Probing Method Setup

Probing process. In Section 3.3.1, we have trained the image captioning model with cross-entropy loss (CE model). We then follow the probing process in Section 3.1 and further train this model using VLP models matching scores as rewards for extra 20 epochs and collect generated captions from all test images in MSCOCO dataset for further analysis.

Evaluation metrics for generated captions. We evaluate the quality of the generated captions using the following automatic metrics. BLEU (Papineni et al., 2002), METEOR (Denkowski and Lavie, 2014), ROUGE (Lin, 2004), and CIDEr (Vedantam et al., 2015) evaluate captions based on n -gram overlap. SPICE (Anderson et al., 2016) measures scene graph similarity between candidates and reference captions. We use the publicly released code

²We also experiment with other captioning models, such as BUTD (Anderson et al., 2018) and vanilla Transformers (Vaswani et al., 2017). In pilot studies, we observe that the constructions of generated captions are similar. We choose the FC model which has a smaller number of parameters.

to compute all metrics.³

3.4 Probing Results

We present our probing results of five VLP models in Table 2. We observe the following trends.

From the left of Table 2, we observe that the scores on the diagonal all improve, which shows that a certain VLP score treated as the reward will improve after SCST training compared to previous cross-entropy training. It means that VLP models consider the generated sentences are more consistent with images. These results are in line with our expectations. We also compute scores of generated sentences using a certain VLP score on other VLP models. We observe that using a certain score does not necessarily improve the scores of other VLP models. It indicates that the preferred patterns of captions of those models differ from each other.

However, as shown in the right of Table 2, we observe that for captions generated from our probing method, all image captioning metrics have dropped sharply compared to the CE model. It indicates that although the captioning model can generate sentences that obtain higher scores from VLP models, which have a potential better cross-modality alignment, the quality of these sentences may not as good as we expect. In Table 3, we provide three examples of the generated sentences. We observe one notable issue that captions generated by the CE model are normal and fit images, but captions after SCST training become unreadable. For example, the visual words in some captions are grouped together, e.g., “*elephant elephant*”, and some captions seem to follow certain sentence patterns, e.g., “*a a motor bike but a a motor but a a road*”. It motivates us to take a closer investigation around this abnormal phenomenon, which we will discuss in detail in Section 4.

³<https://github.com/tylin/coco-caption>



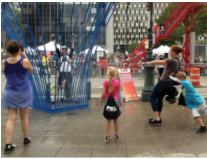
Image	Generated Captions
	<p>CE a motorcycle parked on a dirt road next to a fence</p> <p>UNITER a motorcycle bike motorcycle bike motorcycle bike</p> <p>ROSITA a image that is a a motor bike but a a motor but a a road but a a dirt</p> <p>ViLBERT a motorcycle that was in the metal at this park at the bike across the parking meter at this party</p> <p>CLIP a image image of a row of motorcycles parked near a row area and a wire area near a background</p> <p>LXMERT a bike parked parked of a motorcycle parked parked of a ground near a ground near a sand sand shore</p>
	<p>CE a group of elephants standing next to each other</p> <p>UNITER a elephant of elephants elephant</p> <p>ROSITA a image image that is a a elephants elephants elephants touching a a zoo exhibit but a a zoo pin</p> <p>ViLBERT a baby elephant that across the dirt towards the elephant at this enclosure behind the dirt behind his enclosure at</p> <p>CLIP a group image of people petting a elephant near a fenced area and a baby near a background area area</p> <p>LXMERT these elephants elephants elephants elephants interacting near a elephant elephant touching a touching a elephant near a dirt ground near</p>
	<p>CE a group of people standing on a sidewalk</p> <p>UNITER a group playing a court playing the court</p> <p>ROSITA a young people a a young person a a object that is a a court match but a a court</p> <p>ViLBERT a girl that a frisbee across the basketball towards the basketball at this team at the basketball at the basketball</p> <p>CLIP a group image of people petting a frisbee near a camera area and a neck shirt near a background area</p> <p>LXMERT these young young people standing young standing playing a holding a walking a holding a playing a holding a sky</p>

Table 3: The examples of sentences generated by captioning models trained with different VLP models. All captions are unreadable except the CE model’s.

4 Issues of Generated Captions with High VLP Scores

In Section 3, we have raised the issue that there exists a mismatch between increased scores from VLP models and decreased scores in standard image captioning metrics. In this section, we provide an in-depth and systematic analysis of these generated sentences with high VLP scores.

4.1 Method

We use 5,000 images from the test set of MSCOCO dataset to generate captions with high VLP scores and then perform an analysis on these sentences. We first quantitatively evaluate the problem in these captions, by calculating a set of statistics at the sentence level and token level. For the sentence level, we count the length (Avg. Leng.) and perplexity (PPL) of captions, which reflects the influence of captions. The perplexity is calculated by using SRILM,⁴ a language modeling toolkit. We use it to train a tri-gram language model on the MSCOCO corpus. For the token level, we run a part-of-speech tagger to count the number of nouns.⁵ We count the number of nouns (Noun) and non-repeated nouns (Uni. Noun) in each sentence and the top 10 uni-grams of all captions, which can reflect the prefer-

⁴<https://www.speech.sri.com/projects/srilm>

⁵<https://spacy.io>

ences of selection on tokens. The above statistics of generated captions are reported in Table 4, and the top 10 uni-grams are reported in Appendix A.

VLP models	Sentence Level		Token Level	
	PPL↓	Avg. Leng.	Noun	Uni. Noun
UNITER	134.6	6.7	3.2	2.5
ROSITA	505.2	19.9	6.6	4.2
ViLBERT	174.8	20.0	6.8	4.9
CLIP	131.2	19.7	8.7	6.5
LXMERT	176.6	19.0	7.2	3.2
CE	7.4	9.5	3.4	3.2

Table 4: Statistics of the generated sentences. CE means training FC model with the cross-entropy loss. At the sentence level, we calculate perplexity (PPL) and average length (Avg. Leng.). At the token level, we calculate the average number of nouns (Noun), and the average number of non-repeating nouns (Uni. Noun).

We are also interested in the sentence patterns in the generated captions. To do this, we first run a part-of-speech tagger for each sentence and count the n -grams ($n = 1, 2, 3, 4, 5$). Then we summarize the patterns using regular expressions. Figure 1 illustrates such process: Each token is first associated with a POS tag; we then extract its prefix “a image that is” by 4-grams and merge the nouns phrases DET DET NOUN NOUN; we finally sum up into a regular expression $\text{Prefix} + ((\text{CONJ}) ? + (\text{NOUN.P}) *) *$. All sen-

tence patterns are listed in Table 5.

4.2 Common Issues

We observe two common issues of generated captions with high matching scores, regardless of the choice of VLP models: one is these captions are not fluent, and another one is the captions contain more nouns (visual words), which are grouped together.

Fluency issue. As shown in Table 4, generated captions using SCST training have higher perplexities (more than 100), while the perplexity is only 7.4 in cross-entropy training mode. It suggests that the generated sentences are not fluent or even unreadable (also can be seen from examples in Table 3).

Noun issue. From Table 4, we observe that generated captions with higher VLP scores tend to contain more nouns (except UNITER, as its sentences are shorter). Also, these sentences contain many repeating nouns, while it does not happen in sentences that the CE model generates (only 0.2 difference). Sentences in Table 3 demonstrate this abnormal phenomenon.

4.3 Specific Issues

We observe that VLP models may have some specific issues, which we will discuss below.

Sentence prefix. We summarize prefixes of the generated captions on the left of Table 5. We observe that ROSITA and CLIP have a special fondness for prefixes related to the word “*image*” (e.g., “*a image image that is*” or “*a image image of*”), which accounts for 56.7% and 74.6% respectively.⁶ The difference is that ROSITA prefers prefixes ending with “*of*”, but CLIP prefers prefixes ending with “*that is*”. We also observe prefixes that UNITER, ViLBERT, and LXMERT prefer are similar, which begin with $a + (\text{NOUN})^*$. The difference is that UNITER and LXMERT favor prefixes ending with a preposition `PREP` or verb `VERB`, but ViLBERT prefers prefixes ending with “*that*”.

Sentence pattern. The summarized patterns of the generated sentences are shown in the right of Table 5. One notable observation is that all VLP models prefer to pack noun phrases together, while each VLP model has its own favorite pattern. Single-stream architecture models (UNITER and ROSITA) hardly use a preposition or a verb to connect noun phrases. UNITER

⁶“*a image*” is directly taken from the generated captions by the model and is not an accidental typo.

favors packing the nouns together without any connections, e.g., “*table table*”. ROSITA prefers to pack `determiner+determiner+NOUN` together and they are connected with “*but*”, e.g., “*a a cabinets but a a cabinets*”. Two-stream architecture models (ViLBERT, CLIP, and LXMERT) frequently pack `determiner+(NOUN)` together and they are connected with a preposition or verb. In particular, CLIP prefers to use “*and*” to connect two noun phrases.

Uni-Grams. We observe that top uni-grams of generated sentences are different between VLP models, and far away from those from model trained with cross-entropy loss (top 10 uni-grams are in Appendix A). We observe that single-stream architecture models pay more attention to nouns (6 in UNITER’s uni-grams and 5 in ROSITA’s), emphasizing on different aspects: UNITER focuses more aspects, e.g., animal (“*cat*”, “*dog*”), place (“*room*”, “*road*” and “*apartment*”) and person (“*people*”), while ROSITA focuses more on person (“*person*”, “*man*” and so on). The uni-grams of two-stream architecture models’ captions are relatively uniform in types of words.

5 Limitations of Current VLP Models

The above issues reveal that VLP models tend to prefer the captions with fixed sentence patterns and more visual words. In this section, we design a set of experiments to verify this phenomenon and further discuss the problems of current VLP models. We argue that these limitations have potentially hindered VLP models from aligning visual and textual modalities authentically.

5.1 Setup

The main idea of our verification experiments is to first construct captions, for instance to replace certain tokens, so that these captions carry characteristics that we want to test VLP models. These captions with their corresponding images are then fed to VLP models. Finally, we draw conclusions by comparing the changes in the matching scores. All experiments are done using the test set of MSCOCO dataset, containing 5,000 images.

5.2 Findings and Discussions

VLP models excessively rely on objects and visual words in their inner alignment mechanism, ignoring global semantics. To test if VLP models

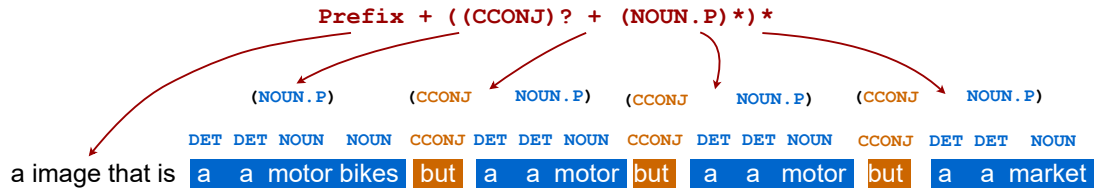


Figure 1: An example of the process of summarizing sentence patterns. The sentence is generated by the model training with ROSITA. (·) * represents a component repeated one to more times, (·) ? represents a component repeated zero or one time

VLP Models	Prefix		Pattern	
	Top Prefix	Ratio	Top Pattern	Ratio
UNITER	Prefix: (NOUN . P) *+REL Example: a table table of	89.7%	Pattern: Prefix+(NOUN . P) * Example: a table table of the table tools	87.2%
ROSITA	Pattern1: a+IMAGE . P+that+is Example1: a image image that is	56.7%	Pattern: Prefix+ ((CONJ) ?+(NOUN . P) *) * Example1: a image image that is a a cabinets but a a cabinets but a a sink Example2: a man man that is a a person bike a a dirt but a a rural dirt	60.7%
	Pattern2: (NOUN . P) *+that+is Example2: a man man that is	28.9%		
ViLBERT	Pattern: (NOUN . P) *+that Example: a bird that	91.9%	Pattern: Prefix+ ((AUX) ?+(REL) ?+(NOUN . P)) * Example: a bird that the kite above the sky above the beach towards this water	82.9%
CLIP	Pattern1: a+IMAGE . P+of Example1: a image image of	74.6%	Pattern: Prefix+ ((CONJ) ?+(REL) ?+(NOUN . P) *) * Example: a image image of a white bathroom with a brown shower curtain near a corner area and seat floor holder	54.0%
LXMERT	Pattern: (NOUN . P) *+REL Example: a man man posing	99.6%	Pattern: Prefix+ ((REL) ?+NOUN . P) * Example: a man man posing a bike riding a bike bike riding a dirt near a road near a dirt ground	94.1%

Table 5: The top prefixes and sentence patterns in generated captions. NOUN . P represents the nouns phrase, REL represents the relationship word (including preposition and verb), CONJ represents the conjunction and AUX represents the copula (including “is” and “was”). IMAGE . P represents a phrase related to the word “image”, such as “image image” and “closeup image”. (·) * represents a component repeated one to more times, (·) ? represents a component repeated zero or one time.



Image	Constructed Captions
	CE a man riding a horse in a field Visual Words “man”, “horse”, “field” UNITER a man of the horse field ROSITA a image image that a a man but a a horse but a a field ViLBERT a man that the horse at the field CLIP a image image of a man near a horse near a field LXMERT a man of a horse of a field
	Visual Words “tracks”, “person”, “guy”, “background”, “bicycle”, “man”, “bike”, “train” UNITER a tracks of the person guy background bicycle ROSITA a image image that a a tracks but a a person but a a guy but a a background but a a bicycle ViLBERT a tracks that the person at the guy at the background at the bicycle CLIP a image image of a tracks near a person near a guy near a background near a bicycle LXMERT a tracks of a person of a guy of a background of a bicycle

Table 6: The example of constructed captions by different templates. (Top) exhibits the constructed captions using different sentence templates, where the visual words are from the CE caption. (Bottom) exhibits the constructed caption with five visual words, where the visual words are from the ground truth.

only align objects and visual words, we first prepare three kinds of captions: (1) captions generated by the CE model (CE captions), (2) the captions that replace visual words with the wrong ones, and (3) the captions that keep the visual words but replace other words with random ones. We then feed these captions into five VLP models and reported our results in Figure 2.

We observe that replacing visual words leads to a sharp drop in scores, but keeping them and replacing other words have a little effect on scores. Although the meaning of two kinds of replaced captions has changed a lot, the captions kept visual words make the VLP models deem they still match images. It shows that VLP model will directly consider a caption matches an image as long as the

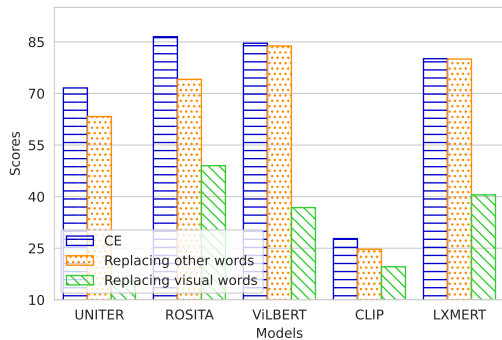


Figure 2: VLP scores of three kinds of captions. CE represents the captions generated by the model trained with cross-entropy loss. Replacing others represents the CE captions replaced all words except visual words with random words. Replacing nouns represents the CE captions replaced visual words with the wrong ones.

Model	CE	Recons.	Template
UNITER	71.6	80.5	a+NOUN+of+the+(NOUN)*
ROSITA	86.5	93.1	a+image+image+that+is+a+a+NOUN (but+a+a+NOUN)*
ViLBERT	72.5	87.7	a+NOUN+that+the+NOUN+(at+the+NOUN)*
CLIP	27.8	28.6	a+image+image+of+NOUN+(near+a+NOUN)*
LXMERT	80.1	81.1	a+NOUN+(of+a+NOUN)*

Table 7: VLP model scores of reconstructed captions on various VLP models. We also present five templates that are used to reconstruct these captions.

caption contains visual words related to the image.

VLP models prefer certain sentence patterns, thus ignoring more important textual information, such as fluency and grammar. We also experiment with reconstructing the captions based on the generated captions using cross-entropy loss, by injecting visual words into sentence templates. We design different templates for different VLP models and list these templates in Table 7. We choose “*but*” for the CCONJ in the template of ROSITA, because “*but*” is the conjunction that appears most frequently in top 10 uni-grams of ROSITA. Similarly, we choose “*at*”, “*near*”, “*of*” for the preposition in the templates of ViLBERT, CLIP and LXMERT respectively. Table 6 (Top) exhibits complete captions for these templates.

Our experimental results are in Table 7. As we just extract visual words from the original CE model and put them in fixed sentence patterns, such reconstructed captions are normally unreadable. However, all VLP models deem the reconstructed unreadable captions fit images better than the captions from cross-entropy loss. It means VLP mod-

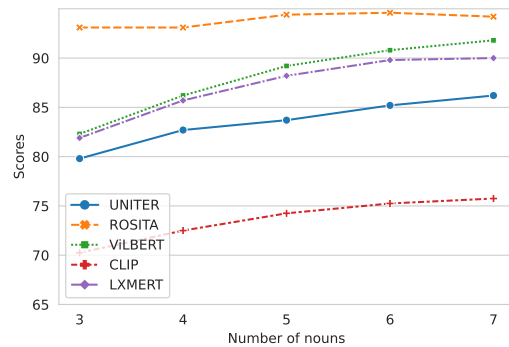


Figure 3: Scores from VLP models with varying number of nouns k in the captions. For clarity, we multiply the score of CLIP by 2.5 according to Hessel et al. (2021).

els have preferences for fixed sentence patterns, while ignoring fluency and grammar issues.

VLP models tend to judge captions with many visual words match images better, which might weaken the significance of key objects in images. We finally evaluate the role of visual tokens in VLP’s cross-modal alignment. To do this, we follow the sentence templates in Table 7 to construct many captions containing visual words with different numbers. As the CE captions contain a few visual words (3.4 on average), we further extract visual words from the ground truth (each image with five captions at least) and merge different visual words into a set. Table 6 (Bottom) exhibits the constructed captions with five visual words.

We vary the number of visual words k ($k = 3, 4, 5, 6, 7$) and report our experimental results in Figure 3. It is observed that with visual words increasing, VLP models deem the captions containing more visual words to be more consistent with images (UNITER does not change much because it gives a high score of the captions containing three visual words). It indicates VLP models deem the captions containing more visual words are more consistent with images.

6 Conclusion

In this paper, we empirically study the cross-modal semantics alignment capability of VLP models, using a newly proposed probing method via an image captioning model. By analyzing the issues of generated captioning guided by the VLP models, we find that VLP models have particular weaknesses in cross-modal semantics alignment, including paying more attention to aligning objects and visual words, while neglecting global semantics; preferring fixed

sentence patterns; and considering captions with more visual words are better aligned.

We hope our work sheds light on promoting better architecture or pre-training tasks for cross-modal semantics alignment to overcome these limitations. Researchers can also use our probing method to discover potential problems when designing new VLP models.

Limitations

One limitation of our work is that our experiment does not cover all VLP models, as some are not open-sourced currently. As a result, we carefully select five VLP models that are powerful and representative of both two mainstream architectures. We plan to experiment with more VLP models in the future to generalize our findings.

Acknowledgements

We would like to thank the anonymous reviewers for their constructive comments. This work was supported by NSFC No. 62176115.

References

- Peter Anderson, Basura Fernando, Mark Johnson, and Stephen Gould. 2016. [SPICE: semantic propositional image caption evaluation](#). In *Computer Vision - ECCV 2016 - 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part V*, volume 9909 of *Lecture Notes in Computer Science*, pages 382–398. Springer.
- Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. 2018. [Bottom-up and top-down attention for image captioning and visual question answering](#). In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 6077–6086. Computer Vision Foundation / IEEE Computer Society.
- Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C. Lawrence Zitnick, and Devi Parikh. 2015. [VQA: visual question answering](#). In *2015 IEEE International Conference on Computer Vision, ICCV 2015, Santiago, Chile, December 7-13, 2015*, pages 2425–2433. IEEE Computer Society.
- Jize Cao, Zhe Gan, Yu Cheng, Licheng Yu, Yen-Chun Chen, and Jingjing Liu. 2020. [Behind the scene: Revealing the secrets of pre-trained vision-and-language models](#). In *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part VI*, volume 12351 of *Lecture Notes in Computer Science*, pages 565–580. Springer.
- Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. 2019. [UNITER: learning universal image-text representations](#). *CoRR*, abs/1909.11740.
- Yuhao Cui, Zhou Yu, Chunqi Wang, Zhongzhou Zhao, Ji Zhang, Meng Wang, and Jun Yu. 2021. [ROSITA: enhancing vision-and-language semantic alignments via cross- and intra-modal knowledge integration](#). In *MM '21: ACM Multimedia Conference, Virtual Event, China, October 20 - 24, 2021*, pages 797–806. ACM.
- Michael J. Denkowski and Alon Lavie. 2014. [Meteor universal: Language specific translation evaluation for any target language](#). In *Proceedings of the Ninth Workshop on Statistical Machine Translation, WMT@ACL 2014, June 26-27, 2014, Baltimore, Maryland, USA*, pages 376–380. The Association for Computer Linguistics.
- Jack Hessel, Ari Holtzman, Maxwell Forbes, Ronan Le Bras, and Yejin Choi. 2021. [Clipscore: A reference-free evaluation metric for image captioning](#). *CoRR*, abs/2104.08718.
- Lun Huang, Wenmin Wang, Jie Chen, and Xiaoyong Wei. 2019. [Attention on attention for image captioning](#). In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, pages 4633–4642. IEEE.
- Andrej Karpathy and Li Fei-Fei. 2015. [Deep visual-semantic alignments for generating image descriptions](#). In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015, Boston, MA, USA, June 7-12, 2015*, pages 3128–3137. IEEE Computer Society.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Kuang-Huei Lee, Xi Chen, Gang Hua, Houdong Hu, and Xiaodong He. 2018. [Stacked cross attention for image-text matching](#). In *Computer Vision - ECCV 2018 - 15th European Conference, Munich, Germany, September 8-14, 2018, Proceedings, Part IV*, volume 11208 of *Lecture Notes in Computer Science*, pages 212–228. Springer.
- Junnan Li, Ramprasaath R. Selvaraju, Akhilesh Gotmare, Shafiq R. Joty, Caiming Xiong, and Steven Chu-Hong Hoi. 2021. [Align before fuse: Vision and language representation learning with momentum distillation](#). In *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 9694–9705.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81.

- Tsung-Yi Lin, Michael Maire, Serge J. Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. 2014. **Microsoft COCO: common objects in context**. In *Computer Vision - ECCV 2014 - 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part V*, volume 8693 of *Lecture Notes in Computer Science*, pages 740–755. Springer.
- Adam Dahlgren Lindström, Johanna Björklund, Suna Bensch, and Frank Drewes. 2020. **Probing multimodal embeddings for linguistic properties: the visual-semantic case**. In *Proceedings of the 28th International Conference on Computational Linguistics, COLING 2020, Barcelona, Spain (Online), December 8-13, 2020*, pages 730–744. International Committee on Computational Linguistics.
- Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. 2019. **Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks**. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 13–23.
- Jiasen Lu, Caiming Xiong, Devi Parikh, and Richard Socher. 2017. **Knowing when to look: Adaptive attention via a visual sentinel for image captioning**. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 3242–3250. IEEE Computer Society.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. **Bleu: a method for automatic evaluation of machine translation**. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics, July 6-12, 2002, Philadelphia, PA, USA*, pages 311–318. ACL.
- Letitia Parcalabescu, Michele Cafagna, Lilitta Muradjan, Anette Frank, Iacer Calixto, and Albert Gatt. 2022. **VALSE: A task-independent benchmark for vision and language models centered on linguistic phenomena**. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2022, Dublin, Ireland, May 22-27, 2022*, pages 8253–8280. Association for Computational Linguistics.
- Letitia Parcalabescu, Albert Gatt, Anette Frank, and Iacer Calixto. 2020. **Seeing past words: Testing the cross-modal capabilities of pretrained v&l models**. *CoRR*, abs/2012.12352.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. **Learning transferable visual models from natural language supervision**. In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pages 8748–8763. PMLR.
- Steven J. Rennie, Etienne Marcheret, Youssef Mroueh, Jerret Ross, and Vaibhava Goel. 2017. **Self-critical sequence training for image captioning**. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 1179–1195. IEEE Computer Society.
- Hao Tan and Mohit Bansal. 2019. **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 2019, Hong Kong, China, November 3-7, 2019*, pages 5099–5110. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. **Attention is all you need**. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008.
- Ramakrishna Vedantam, C. Lawrence Zitnick, and Devi Parikh. 2015. **Cider: Consensus-based image description evaluation**. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015, Boston, MA, USA, June 7-12, 2015*, pages 4566–4575. IEEE Computer Society.
- Oriol Vinyals, Alexander Toshev, Samy Bengio, and Dumitru Erhan. 2015. **Show and tell: A neural image caption generator**. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015, Boston, MA, USA, June 7-12, 2015*, pages 3156–3164. IEEE Computer Society.
- Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron C. Courville, Ruslan Salakhutdinov, Richard S. Zemel, and Yoshua Bengio. 2015. **Show, attend and tell: Neural image caption generation with visual attention**. In *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6-11 July 2015*, volume 37 of *JMLR Workshop and Conference Proceedings*, pages 2048–2057. JMLR.org.
- Lewei Yao, Runhui Huang, Lu Hou, Guansong Lu, Minzhe Niu, Hang Xu, Xiaodan Liang, Zhenguo Li, Xin Jiang, and Chunjing Xu. 2021. **FILIP: fine-grained interactive language-image pre-training**. *CoRR*, abs/2111.07783.

A Top 10 Uni-Grams

We provide top 10 uni-grams of generated captions. Figure 4 (b) and Figure 4 (c) are top 10 uni-grams of UNITER and ROSITA respectively, which contain more visual words. Figure 4 (d), Figure 4 (e) and Figure 4 (f) are top 10 uni-grams of ViLBERT, CLIP and LXMERT, which are relatively uniform in types of words. Especially, uni-grams of LXMERT' sentences contain four verbs and other models' sentences hardly have.

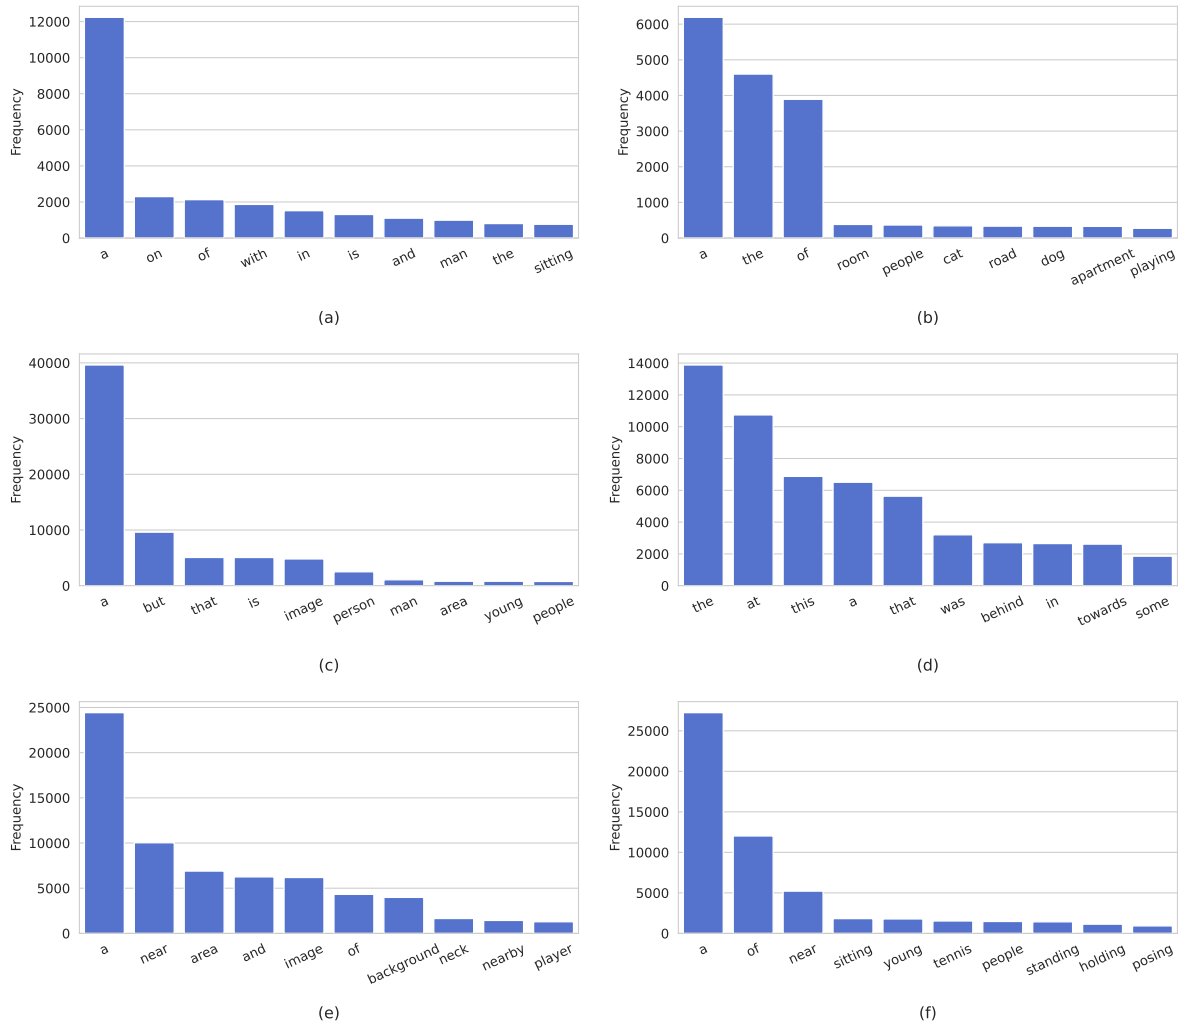


Figure 4: Top 10 uni-grams of various models. Figure (a) is from the model training with cross-entropy, and Figures (b) - (f) are from UNITER, ROSITA, ViLBERT, CLIP, and LXMERT, respectively.