

Behind the Scene: Revealing the Secrets of Pre-trained Vision-and-Language Models

Jize Cao^{*1}, Zhe Gan², Yu Cheng², Licheng Yu^{*3}
Yen-Chun Chen², and Jingjing Liu²

¹ University of Washington
caojize@cs.washington.edu

² Microsoft Dynamics 365 AI Research

{zhe.gan,yu.cheng,yen-chun.chen,jingjl}@microsoft.com

³ Facebook AI
lichengyu@fb.com

Abstract. Recent Transformer-based large-scale pre-trained models have revolutionized vision-and-language (V+L) research. Models such as ViLBERT, LXMERT and UNITER have significantly lifted state of the art across a wide range of V+L benchmarks. However, little is known about the inner mechanisms that destine their impressive success. To reveal the secrets behind the scene, we present VALUE (**V**ision-**A**nd-**L**anguage **U**nderstanding **E**valuation), a set of meticulously designed probing tasks (*e.g.*, Visual Coreference Resolution, Visual Relation Detection) generalizable to standard pre-trained V+L models, to decipher the inner workings of multimodal pre-training (*e.g.*, implicit knowledge garnered in individual attention heads, inherent cross-modal alignment learned through contextualized multimodal embeddings). Through extensive analysis of each archetypal model architecture via these probing tasks, our key observations are: (*i*) Pre-trained models exhibit a propensity for attending over text rather than images during inference. (*ii*) There exists a subset of attention heads that are tailored for capturing cross-modal interactions. (*iii*) Learned attention matrix in pre-trained models demonstrates patterns coherent with the latent alignment between image regions and textual words. (*iv*) Plotted attention patterns reveal visually-interpretable relations among image regions. (*v*) Pure linguistic knowledge is also effectively encoded in the attention heads. These are valuable insights serving to guide future work towards designing better model architecture and objectives for multimodal pre-training.⁴

1 Introduction

Recently, Transformer-based [35] large-scale pre-trained models [32,22,6,19,21,27] have prevailed in Vision-and-Language (V+L) research, an important area that

^{*} This work was done when Jize and Licheng worked at Microsoft.

⁴ Code is available at <https://github.com/JizeCao/VALUE>.

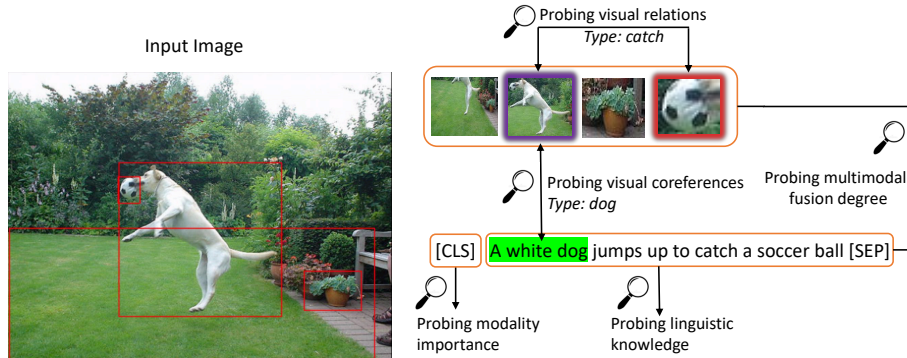


Fig. 1: Illustration of the proposed VALUE framework for investigating pre-trained vision-and-language models. VALUE consists of a set of well-designed probing tasks that unveil the inner mechanisms of V+L pre-trained models across: (i) Multimodal Fusion Degree; (ii) Modality Importance; (iii) Cross-modal Interaction via probing visual coreferences; (iv) Image-to-image Interaction via probing visual relations; and (v) Text-to-text Interaction via probing learned linguistic knowledge.

sits at the nexus of computer vision and natural language processing (NLP). Inspired by BERT [9], a common practice for pre-training V+L models is to first encode image regions and sentence words into a common embedding space, then use multiple Transformer layers to learn image-text contextualized joint embeddings through well-designed pre-training tasks. There are two main schools of model design: (i) single-stream architecture, such as VLBERT [27] and UNITER [27], where a single Transformer is applied to both image and text modalities; and (ii) two-stream architecture, such as LXMERT [32] and ViLBERT [22], in which two Transformers are applied to images and text independently, and a third Transformer is stacked on top for later fusion. When finetuned on downstream tasks, these pre-trained models have achieved new state of the art on image-text retrieval [18], visual question answering [4,11], referring expression comprehension [37], and visual reasoning [13,28,38]. This suggests that substantial amount of visual and linguistic knowledge has been encoded in the pre-trained models.

There has been several studies that investigate latent knowledge encoded in pre-trained language models [33,7,25,16]. However, analyzing multimodal pre-trained models is still an unexplored territory. It remains unclear how the inner mechanisms of cross-modal pre-trained models induce their empirical success on downstream tasks. Motivated by this, we present VALUE (**V**ision-**A**nd-**L**anguage **U**nderstanding **E**valuation), a set of well-designed probing tasks that aims to reveal the secrets of these pre-trained V+L models. To investigate both single- and two-stream model architectures, we select one model from each category (LXMERT for two-stream and UNITER for single-stream, because of their superb performance across many V+L tasks). As illustrated in Figure 1, VALUE is designed to provide insights on: (a) Multimodal Fusion Degree; (b) Modal-

ity Importance; (c) Cross-modal Interaction (Image-to-text/Text-to-image); (d) Image-to-image Interaction; and (e) Text-to-text Interaction.

For (a) Multimodal Fusion Degree, clustering analysis between image and text representations shows that in single-stream models like UNITER, as the network layers go deeper, the fusion between two modalities becomes more intertwined. However, the opposite phenomenon is observed in two-stream models like LXMERT. For (b) Modality Importance, by analyzing the attention trace of the [CLS] token, which is commonly considered as containing the intended fused multimodal information and often used as the input signal for downstream tasks, we find that the final predictions tend to depend more on textual input rather than visual input.

To gain deeper insights into how pre-trained models drive success in downstream tasks, we look into three types of interactions between modalities. For (c) Cross-modal Interaction, we propose a Visual Coreference Resolution task to probe its encoded knowledge. For (d) Image-to-image Interaction, we conduct analysis via Visual Relation Detection between two image regions. For (e) Text-to-text Interaction, we evaluate the linguistic knowledge encoded in each layer of the tested model with SentEval toolkit [8], and compare with the original BERT [9]. Experiments show that both single- and two-stream models, especially the former, can well capture cross-modal alignment, visual relations, and linguistic knowledge.

To the best of our knowledge, this is the first known effort on thorough analysis of pre-trained V+L models, to gain insights from different perspectives about the latent knowledge encoded in self-attention weights, and to distill the secret ingredients that drive the empirical success of prevailing V+L models.

2 Related Work

For image-text representation learning, ViLBERT [22] and LXMERT [32] used two-stream architecture for pre-training, while B2T2 [2], VisualBERT [21], Unicoder-VL [19], VL-BERT [27] and UNITER [6] adopted single-stream architecture. VLP [39] proposed a unified pre-trained model for both image captioning and VQA. Multi-task learning [23] and adversarial training in VILLA [10] have been studied to boost performance. On video+language side, VideoBERT [30] applied BERT to learn joint embeddings of video frame tokens and linguistic tokens from video-text pairs. CBT [29] introduced contrastive learning to handle real-valued video frame features, and HERO [20] proposed hierarchical Transformer architectures to leverage both global and local temporal visual-textual alignments. However, except for some simple visualization of the learned attention maps [32], no existing work has systematically analyzed these pre-trained models.

There has been some recent studies on assessing the capability of BERT in capturing structural properties of language [34,15,31]. Multi-head self-attention has been analyzed for machine translation in [36,24], which observed that only a small subset of heads is important, and the other heads can be pruned without affecting model performance. [33] reported that BERT can rediscover the

classical NLP pipeline, where basic syntactic information appears in lower layers, while high-level semantic information appears in higher layers. Analysis on BERT self-attention [7,12,16] showed that BERT can learn syntactic relations, and a limited set of attention patterns are repeated across different heads. [25,5] and [40] demonstrated that BERT has surprisingly strong ability to recall factual relational knowledge and perform commonsense reasoning. A layer-wise analysis of Transformer representations in [1] provided insights to the reasoning process of how BERT answers questions. All these studies have focused on the analysis of BERT, while investigating pre-trained V+L models is still an uncharted territory. Given their empirical success and unique multimodal nature, we believe it is instrumental to conduct an in-depth analysis to understand these models, to provide useful insights and guidelines for future studies. New probing tasks such as visual coreference resolution and visual relation detection are proposed for this purpose, which can lend insights to other evaluation tasks as well.

3 VALUE: Probing Pre-trained V+L Models

Key curiosities this study aims to unveil include:

- (i) What is the correlation between multimodal fusion and the number of layers in pre-trained models? (Sec. 3.1)
- (ii) Which modality plays a more important role that drives the pre-trained model to make final predictions? (Sec. 3.2)
- (iii) What knowledge is encoded in pre-trained models that supports cross-modal interaction and alignment? (Sec. 3.3)
- (iv) What knowledge has been learned for image-to-image (intra-modal) interaction (*i.e.*, visual relations)? (Sec. 3.4)
- (v) Compared with BERT, do pre-trained V+L models effectively encode linguistic knowledge for text-to-text (intra-modal) interaction? (Sec. 3.5)

To answer these questions, we select one model from each archetypal model architecture for dissection: UNITER-base [6] (12 layers, 768 hidden units per layer, 12 attention heads) for single-stream model, and LXMERT [32] for two-stream model.⁵ Single-stream model (UNITER) shares the same structure as BERT [9], except that the input now becomes a mixed sequence of two modalities. Two-stream model (LXMERT) first performs self-attention through several layers on each modality independently, then fuses the inputs through a stack of cross-self-attention layers (first cross-attention, then self-attention). Therefore, the attention pattern in two-stream models is fixed, as one modality is only allowed to attend over either itself or the other modality at any time (there is no such constraint in single-stream models). A more detailed model description is provided in the Appendix.

Two datasets are selected for our probing experiments:

- (i) **Visual Genome (VG)** [17]: image-text dataset with annotated dense captions and scene graphs.

⁵ Our probing analysis can be readily extended to other pre-trained models as well.

(ii) **Flickr30k Entities** [26]: image-text dataset with annotated visual co-reference links between image regions and noun phrases in the captions.

Each data sample consists of three components: (i) an input image; (ii) a set of detected image regions;⁶ and (iii) a caption. In Flickr30k Entities, the caption is relatively long, describing the whole image; while in VG, a few short captions are provided (called dense captions), each describing an image region.

The number of annotated image regions in an image is relatively small (5-6 for Flickr30k Entities, and 2-4 for VG dense annotated graph); while in pre-trained V+L models, the number of image regions fed to the model is typically 36 [3]. Therefore, we extract an additional set of image regions from a Faster R-CNN pre-trained on the VG dataset [3], and combine them with the original image regions provided in the dataset. The initial image representation is obtained from the Faster R-CNN as well. Finally, we feed both image regions and textual tokens into the pre-trained model for probing analysis. The following sub-sections describe analysis results and key observations from each proposed probing task.

3.1 Deep to Profound: Deeper Layers Lead to More Intertwined Multimodal Fusion

By nature, visual features in images and linguistic clues in text have distinctive characteristics. However, it is unknown whether the semantic gap between their corresponding representations narrows (*i.e.*, the contextualized representations of the two modalities become less differentiable) through the cross-modal fusion between intermediate layers.

3.1.1 Probing Task To answer this question, we design a probing task to test the multimodal fusion degree of a model. First, we extract all the embedding features of both image regions and textual tokens from aforementioned two datasets (VG and Flickr30k Entities). For single-stream model (UNITER-base), we extract the output representation from each layer, as multimodal fusion is performed through all the layers. For two-stream model (LXMERT), we only consider the output embeddings from the last 5 layers (*i.e.*, the layers constituting the cross-modality encoder), because the two modalities do not have any interaction prior to that. To gain quantitative measurement, for each data sample, we apply the k -means algorithm (with $k = 2$) on the representations from each layer to partition them into two clusters, and measure the difference between the formed clusters and ground-truth visual/textual clusters via Normalized Mutual Information (NMI), an unsupervised metric for evaluating differences between clusters. A larger NMI value implies that the distinction between two clusters is more significant, indicating a lower fusion degree [14]. For example, when NMI is equal to 1.0, the two clusters represent the original visual tokens and textual tokens, respectively. We use the mean value of NMI for all the data samples to measure the level of multimodal fusion.

⁶ An image region is also called a visual token in this paper; these two terms will be used interchangeable throughout the paper.

Layer	0	1	2	3	4	5	6	7	8	9	10	11
<i>single-stream</i>												
Flickr30k	0.36	0.38	0.39	0.41	0.38	0.38	0.38	0.38	0.32	0.20	0.26	0.20
Visual Genome	0.25	0.25	0.24	0.24	0.22	0.22	0.21	0.21	0.20	0.17	0.16	0.16
<i>two-stream (cross)</i>												
Flickr30k	0.42	0.48	0.67	0.75	0.43	–	–	–	–	–	–	–
Visual Genome	0.43	0.56	0.68	0.78	0.57	–	–	–	–	–	–	–

Table 1: NMI scores on multimodal fusion probing. A smaller NMI value indicates a higher fusion degree. Note that the two-stream model (LXMERT) only has 5 layers in its cross-modality encoder. A larger layer number corresponds to an upper layer.

3.1.2 Results Table 1 summarizes the probing results on multimodal fusion degree. For single-stream model (UNITER-base), the NMI scores gradually decrease, indicating that the representations from the two modalities fuse together deeper and deeper from lower to higher layers. This observation matches our intuition that the embedding of a modality from a higher layer better attends over the other modality than a lower layer. However, in two-stream model (LXMERT), as the layers go deeper (except for the last layer), the representations deviate from each other. One possible explanation is that single-stream model applies the same set of parameters to both image and text modalities, while two-stream model uses two separate sets of parameters (as part of its network design) to model the attention on the image and text modality independently. The latter makes it relatively easier to distinguish the two representations, leading to a higher NMI score. An t-SNE visualization is provided in the Appendix.

3.2 Who Pulls More Strings: Textual Modality Is More Dominant than Image

Following BERT [9], pre-trained V+L models add a special [CLS] token at the beginning of a sequence, and a special [SEP] token at the end. In practice, the [CLS] token representation from the last layer is used as the fused representation for both modalities in downstream tasks. Since the [CLS] token absorbs information from both modalities through self-attention, the degree of attention of the [CLS] token over each modality could be regarded as evidence to answer the following question: *which modality is more dominant during inference?* Note that in the two-stream model design, the [CLS] token is not allowed to attend to both image and textual modality simultaneously. Therefore, the probing experiments for modality importance is focused on single-stream model.

3.2.1 Probing Task Formally, to quantitatively analyze the [CLS] attention trace, the *Modality Importance* (MI) of a head j is defined as the sum of the attention values that the [CLS] token spent on the modality M (visual or textual) at j for the whole sequence $S = ([CLS], t_1, \dots, t_m, [SEP], v_1, \dots, v_n)$, where

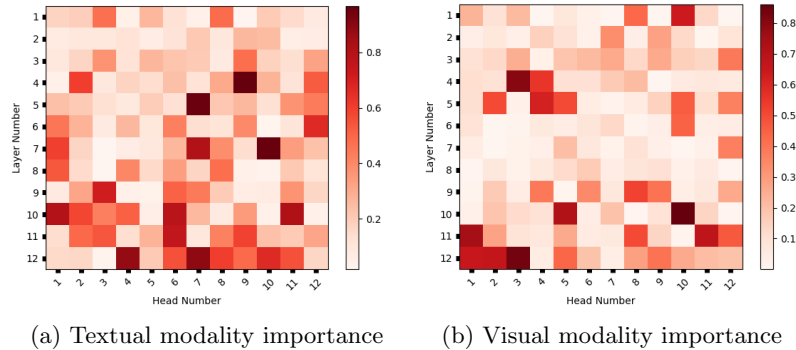


Fig. 2: Visualization of the modality importance score for all the 144 attention heads.

t_1, \dots, t_m and v_1, \dots, v_n are textual and visual tokens, respectively. That is,

$$I_{M,j} = \sum_{i \in S} \mathbb{1}(i \in M) \cdot \alpha_{ij}, \quad (1)$$

where α_{ij} represents the attention weight that the [CLS] token attends on token i at head j , and $\mathbb{1}(\cdot)$ is the indicator function.

Since the attention heads in the same layer perform attention on the same representation, it is natural to expand the above head-level analysis to a layer-level analysis, by considering the mean of MI scores of all the 12 heads as the MI measurement of that layer. In addition, we calculate an overall MI score via summing up the MI scores of all the 144 heads. The mean MI value of all the data samples is reported as the score for probing modality importance. Note that the MI score for visual/textual modality is calculated based on the attention weights on the visual/textual tokens; while the attention weights on the special [CLS] and [SEP] tokens are not considered. Therefore, the two mean MI scores from the two modalities do not sum to one.

3.2.2 Results Experiments are conducted on the Flickr30k Entity dataset. Figure 2 provides the MI scores for all the 144 attention heads of UNITER-base. The heatmap on the textual MI is denser than that of the visual MI, showing that more attention heads are learning useful knowledge from the textual modality than the image modality. Figure 3a further shows the layer-level MI scores for each modality. The average MI score on the text modality is higher than that on the image modality, especially for intermediate layers, suggesting that the pre-trained model relies more on the textual modality for making decisions during inference time.

3.3 Winner Takes All: A Subset of Heads is Specialized for Cross-modal Interaction

The key difference between single-modal Transformer (such as BERT) and two-modal Transformer (such as UNITER and LXMERT) is that two-modal Trans-

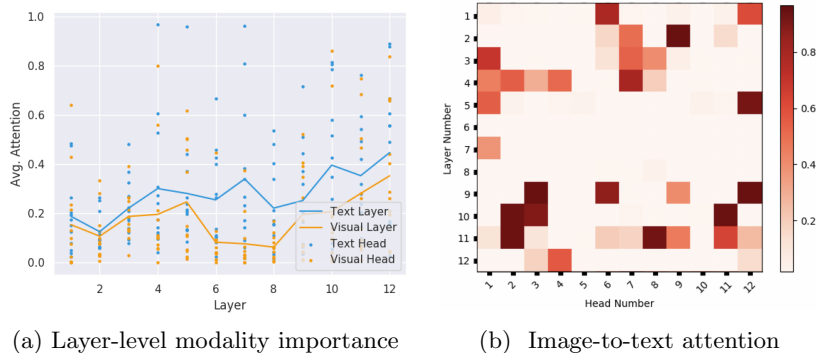


Fig. 3: Layer-level modality importance scores and visualization of the image-to-text attention for all the 144 attention heads.

former requires extra cross-modal interaction. To gain an in-depth understanding of cross-modal attention heads, which is instructive to prompt better model design and enhance model interpretability, we look into two special types of head: (i) image-to-text head; and (ii) visual-coreference head.

3.3.1 Probing Image-to-text Head We first analyze whether there exists any head specialized in learning cross-modal interaction. Formally, for a given image-text pair, visual and textual tokens are denoted as V and T , respectively. We define a head as *image-to-text head* if:

$$\exists v \in V \sum_{t \in T} \alpha_{v \rightarrow t} > 0.5, \quad (2)$$

where $\alpha_{v \rightarrow t}$ denotes the attention weight from a visual token to a textual token. This defines whether a visual token pays more attention to the text modality. Specifically, if there exists one visual token that has higher attention weight on text than other tokens, we regard the corresponding head as performing cross-modal attention from image to text.

Based on the above definition, we count the number of occurrences of head as an *image-to-text head* for all data samples in the Flickr30k Entities dataset, and report the empirical probability of each head being an *image-to-text head*. Note that in the two-stream model, the image-to-text head is by design, therefore, we only conduct this analysis on single-stream model.

Figure 3b shows there is a specific set of heads in UNITER-base that perform cross-modal interaction. The maximum probability of a head performing image-to-text attention is 0.92, the minimum probability is 0, and only 15% heads have more than 0.5 probability to pay the majority attention weight on the image-to-text part. Interestingly, by training single-stream model, the attention heads are automatically learned to exhibit a “two-stream” pattern, where some heads control the message sharing from the visual modality to the textual modality.

<i>Results on single-stream model</i>				
Coref Type	$V \rightarrow T$	$T \rightarrow V$	$V \rightarrow T$ (Rand.)	$T \rightarrow V$ (Rand.)
people	0.165 (9-12)	0.060 (3-1)	0.106 (9-12)	0.035 (11-8)
bodyparts	0.108 (9-3)	0.051 (11-8)	0.084 (9-12)	0.038 (11-8)
scene	0.151 (9-12)	0.048 (3-1)	0.111 (9-12)	0.035 (11-8)
clothing	0.157 (9-3)	0.040 (3-1)	0.092 (9-12)	0.040 (10-2)
animals	0.285 (9-12)	0.137 (3-1)	0.139 (9-12)	0.047 (9-12)
instruments	0.244 (11-8)	0.042 (9-12)	0.091 (9-12)	0.031 (9-12)
vehicles	0.194 (9-12)	0.065 (3-1)	0.112 (9-12)	0.039 (9-12)
<i>Results on two-stream model</i>				
Coref Type	$V \rightarrow T$	$T \rightarrow V$	$V \rightarrow T$ (Rand.)	$T \rightarrow V$ (Rand.)
people	0.145 (2-8)	0.063 (2-6)	0.056 (1-7)	0.057 (3-9)
bodyparts	0.079 (5-7)	0.059 (1-7)	0.041 (1-2)	0.060 (3-9)
scene	0.076 (5-7)	0.059 (1-7)	0.038 (1-7)	0.060 (3-9)
clothing	0.062 (5-7)	0.062 (2-6)	0.040 (1-7)	0.061 (3-9)
animals	0.235 (3-4)	0.106 (4-7)	0.075 (1-7)	0.072 (3-9)
instruments	0.144 (5-7)	0.040 (1-7)	0.055 (1-11)	0.058 (3-9)
vehicles	0.097 (1-1)	0.046 (2-6)	0.097 (1-7)	0.062 (3-9)

Table 2: Results on visual coreference resolution. Each number represents the maximum attention weights between two linked tokens, averaged across all data samples. $V \rightarrow T$ records the attention trace where a visual token attends to the linked noun phrase; $T \rightarrow V$ records the attention trace on the other direction. The head that achieves the maximum attention weight is shown in the bracket.

3.3.2 Visual Coreference Resolution One straightforward way to investigate the visual-linguistic knowledge encoded in the model is to evaluate whether the model is able to match an image region to its corresponding textual phrase in the sentence. Thus, we design a *Visual Coreference Resolution* task (similar to coreference resolution) to predict whether there is a link between an image region and a noun phrase in the sentence that describes the image. In addition, each coreference link in the dataset is annotated with a label (8 in total).

Through this task, we can find out whether the coreference knowledge can be captured by the attention trace. To achieve this goal, for each data sample in the Flickr30k Entity dataset, we extract the encoder’s attention weights for all the 144 heads. Note that noun phrases typically consist of two or more tokens in the sequence. Thus, we extract the maximum attention weight between the image region and each word of the noun phrase for each head. The maximum weight is then used to evaluate which head identifies visual coreference (*i.e.*, performing multimodal alignment).

Results are summarized in Table 2. The columns labeled with “(Rand.)” are considered as ablation groups to identify whether the high attention weight of a certain coreference relationship is triggered by the relation between the image-text pair rather than the effect of one specific image/text token. For a link $V \rightarrow T$, we measure the maximum attention weight of the visual token V

to a random noun phrase, to obtain the results for $V \rightarrow T$ (Rand.). Similarly, we use the maximum attention weight of a noun phrase T to a random visual token to obtain the results for $T \rightarrow V$ (Rand.).

Results show that the relation between a noun phrase and its linked visual token is encoded in the attention pattern, especially for $V \rightarrow T$. The heads (9-3)⁷, (9-12) and (3-1) have captured richer coreference knowledge than other heads, indicating that there exists a subset of heads in the pre-trained model that is specialized in coreference linking between the two modalities. Furthermore, for single-stream model, we observe that some heads encode the coreference information in both directions, serving as additional evidence that these heads perform cross-modal alignment specifically. On the other hand, the amount of learned coreference knowledge is limited in the text modality attention trace ($T \rightarrow V$), which provides indirect evidence that the text modality does not incorporate much visual information, even with the forced cross-attention design in the two-stream model.

3.3.3 Probing Combinations of Heads Previous analysis mainly investigates whether cross-modal knowledge can be captured through individual attention head. It is also possible that such knowledge can be induced via the cooperation of multiple heads. To quantitatively analyze this, we further examine visual coreference through probing over combinations of heads.

1) Attention Prober To reveal the learned knowledge across different attention heads, we use a linear classifier based on the combination of attention weights, following [7]. Specifically,

$$p(c|i, j) \propto \sum_{k=1}^N (w_k \alpha_{ij}^k + \mu_k \alpha_{ji}^k), \quad (3)$$

where given tokens i and j in the sequence, $p(c|i, j)$ is the probability of the link label between these two tokens being c . α_{ij}^k is the attention weight for token i attending to token j at head k , w_k and μ_k are two learnable scalars, and N is the number of attention heads.

2) Layer-wise Embedding Prober Similarly, to further examine the knowledge encoded in the model, we can naturally extend the above attention-based prober into an embedding-based prober:

$$p_k(c|i, j) \propto w_{cls}^{k\top} (W_w^k e_i^k \odot W_\mu^k e_j^k), \quad (4)$$

where $p_k(c|i, j)$ is the same as defined in (3), e_i^k and e_j^k are the embeddings of token i and j at the k -th layer, and $W_w^k, W_\mu^k \in \mathbb{R}^{768 \times 768}, w_{cls}^k \in \mathbb{R}^{768 \times 1}$. By training a linear classifier on top of extracted embeddings, this prober provides a way to probe the encoded knowledge between each pair of tokens on each layer.

For visual coreference resolution, we probe the model on two sub-tasks: (i) **Visual Coref Detection (VCD)**: determine whether a noun phrase is coreferenced to a specific visual token (binary classification); and (ii) **Visual Coref**

⁷ Head (i - j) means the j -th head at the i -th layer.

Classifier Input	VCD (SS)	VCD (TS)	VCC (SS)	VCC (TS)
144 Attention Heads	52.04	53.68	75.10	54.47
Random Guess	50.00	50.00	12.50	12.50
Layer 1	56.86	53.68	93.51	93.35
Layer 5	59.12	52.59	94.05	92.62
Layer 12	58.40	/	93.44	/

Table 3: Results of attention and layer-wise embedding probers on Visual Coreference Detection and Classification (VCD and VCC). SS: single-stream; TS: two-stream.

Classification (VCC): classify the label of the coreference relation between a noun phrase and a visual token (multi-class classification). With these tasks, we can examine whether certain type of knowledge is encoded, as well as the granularity of the knowledge encoded in the probed feature space⁸.

Results are shown in Table 3. Neither model performs well on the VCD task, suggesting that there is no attention pattern or embedding feature that can handle all the coreference relations. However, the results on VCC are encouraging. This aligns with our observation from Table 2 that some attention heads of the two models are significantly effective in certain coreference relations.⁹

3.4 Secret Liaison Revealed: Cross-Modality Fusion Registers Visual Relations

To evaluate the encoded knowledge learned from the image modality, we adopt the visual relation detection task, which requires a model to identify and classify the relation between two image regions. This task can be viewed as examining whether the model captures visual relations between image regions.

The VG dataset is used for this task, which contains 1,531,448 first-order object-object relations. First-order object-object relation can be determined simply by the visual representations of two objects, independent to other objects in the image or text annotation. To reduce the imbalance in the number of relations per relation type, we randomly select at most 15,000 subject-object relation pairs (s, o) per relation type. Furthermore, to de-duplicate cases where the same type of relation comes from the same text annotation, we select at most 5 same relation types (s, o) from the same annotation. This probing task is performed on 32 most frequent relation pairs in the dataset.

⁸ Since noun phrase may contain several tokens, we use the maximum attention weight among the tokens in that phrase over an image region as the attention weight between the noun phrase and the image region. The embedding of the noun phrase is the mean of all the representations of its textual tokens.

⁹ Though both models’ embedding probers achieve higher than 94% accuracy on the VCC task, it is worth noting that text embedding input can potentially leak the link information. For instance, the phrase “A guard with a white hat” may already provide coreference information between **person** and the corresponding image region.

Relation Type	SS	SS (Rand.)	TS	TS (Rand.)
on	0.154 (10-1)	0.055 (1-8)	0.157 (3-12)	0.063 (5-9)
standing in	0.107 (2-8)	0.051 (1-8)	0.173 (7-4)	0.064 (3-1)
wearing	0.311 (10-1)	0.049 (1-8)	0.230 (7-4)	0.055 (3-1)
playing	0.135 (4-1)	0.050 (1-8)	0.103 (7-10)	0.062 (3-1)
eating	0.138 (10-1)	0.056 (1-8)	0.142 (7-4)	0.067 (3-1)
holds	0.200 (10-1)	0.055 (1-8)	0.173 (7-4)	0.066 (3-1)
covering	0.151 (7-2)	0.053 (1-8)	0.173 (3-1)	0.061 (3-6)

Table 4: Results on Visual Relation Identification/Classification using maximum attention weight between two visual tokens. SS: single stream; TS: two stream.

1) Probing Individual Attention Head We apply the same analysis here similar to the visual coreference resolution task. The only difference is that we do not consider the directions of attention (*i.e.*, $s \rightarrow o$, or $o \rightarrow s$), as both directions correspond to the same visual modality. The average of maximum attention value in both $s \rightarrow o$ and $o \rightarrow s$ directions is reported for visual relation analysis.

Results are summarized in Table 4. We report the maximum attention weights of 7 of 30 relations. The columns with “(Rand.)” are ablation groups. As shown in the comparison, the learned attention heads encode rich knowledge about the relations between visual tokens with much higher attention weights. Moreover, similar to the observation in visual coreference resolution task, specific heads (10-1 head in the single-stream model, 7-4 head in the two-stream model) have captured richer visual relations than others.


2) Probing Combinations of Attention Heads The above analysis only reveals the behavior of individual attention heads. Similar to Sec. 3.3, we further examine visual relations by training a linear classifier on top of a combination of attention heads over two sub-tasks: (*i*) **Relation Identification**: determine whether two image regions have a relationship; and (*ii*) **Relation Classification**: classify the relation label between two image regions.

Results on probing a combination of heads are summarized in Figure 4. Two baselines are considered in this task. (*i*) Original visual embeddings from Faster R-CNN [3]. This setup evaluates how much correlation between two related visual representations has elevated or diluted by the V+L models. (*ii*) Mismatched image-text representation. In this baseline, we construct a dataset where an image is associated with an unrelated dense annotation instead of a related one. This baseline evaluates how the correlation between the image regions changes based on the text modality. Note that for the two-stream model, there are 10 layers involved in visual relation reasoning. The first five layers perform self-attention across the visual modality only, and in the last five layers the visual representation interacts with the text modality.

As shown, both models perform much better than the baseline with related annotations (the dashed black line in Figure 4), indicating there is a substantial amount of visual-relation knowledge encoded in these representations. On the

Classifier Input	VRI (SS)	VRC (SS)	VRI (TS)	VRC (TS)
144 Attention Heads	69.81	24.67	67.53	18.89
144 Attention Heads (mismatch)	64.66	23.64	64.72	18.42
Random Guess	50.00	3.33	50.00	3.33

(a) Accuracies (%) of the attention probers on VRI and VRC



(b) Layer-wise emb. prober on VRC (c) Layer-wise emb. prober on VRI

Fig. 4: Results of attention and layer-wise embedding probers on Visual Relation Identification and Classification (VRI and VRC).

other hand, the specific visual relation knowledge degrades a lot with mismatched image regions and dense annotations (the solid red and blue lines vs. the dashed lines in Figure 4). For single-stream model, the visual-relation knowledge decayed to that of the original visual embedding. This makes sense because no visual relation information can be obtained from the mismatched caption.

The visual relation knowledge of the two-stream model is greatly influenced by the unpaired caption, leading to a huge performance drop after Layer 5 in the VRC task. This result may contribute to different inductive bias between the two models. For two-stream model, since the visual modality has to attend to the text modality during cross attention, the visual representation will be greatly influenced by the text modality even when the two modalities are totally unrelated. On the other hand, because the visual representation in the single-stream model is capable of selectively choosing whether to attend over the text modality, the effect of unrelated caption on the visual representation is negligible.

3.5 No Lost in Translation: Pre-trained V+L Models Encode Rich Linguistic Knowledge

Besides looking into the knowledge learned from the visual modality, we are also interested in the encoded knowledge learned from the text modality. To achieve this goal, we probe the pre-trained models over nine tasks defined in the SentEval toolkit [8]. Descriptions about the tasks are provided in the Appendix.

First, we extract contextualized word representations from the pre-trained models. For single-stream model, the input is the sequence of tokens with text

Model	SentLen (Surface)	TreeDepth (Syntactic)	TopConst (Syntactic)	BShift (Syntactic)	Tense (Semantic)	SubjNum (Semantic)	ObjNum (Semantic)	SOMO (Semantic)	CoordInv (Semantic)
SS	88.8 (L2)	36.4 (L5)	79.0 (L7)	81.6 (L10)	86.6 (L11)	83.4 (L7)	78.8 (L9)	57.1 (L12)	62.1 (L11)
TS	83.8 (L7)	34.0 (L8)	67.2 (L8)	64.9 (L8)	75.6 (L6)	78.8 (L9)	76.8 (L8)	51.4 (L8)	58.7 (L8)
BERT	96.2 (L3)	41.3 (L6)	84.1 (L7)	87.0 (L9)	90.0 (L9)	88.1 (L6)	82.2 (L7)	65.2 (L12)	78.7 (L9)

Table 5: Results on the linguistic probing tasks. L_n : Layer n .

only. For two-stream model, since the last 5 cross-attention layers require visual input, which is not covered by these linguistic benchmarks, we only evaluate the first 9 layers that are performing self-attention over pure text input.

We use the Google-pretrained BERT-base model as the baseline. For each task, we obtain the results from all the layers and report the best number. The results in Table 5 show that pre-trained V+L model generally performs worse than the original BERT-base model on these linguistic benchmarks, which is as expected. A full table with results from all the layers is provided in the Appendix. As shown in the table, the single-stream model performs better than the two-stream model across all the tasks. One possible reason is that LXMERT does not initialize the parameters of the language Transformer encoder from BERT-base, whereas UNITER does. Thus, using the parameters of BERT as initialization is potentially useful for the model to acquire rich linguistic knowledge, and helpful for tasks involving complex text-based reasoning.

To measure the gains due to learning, we have conducted all the above experiments (Secs. 3.1-3.5) on untrained baselines with random weights. These additional results are provided in the Appendix.

4 Conclusion and Key Takeaways

Intrigued by the five questions presented at the beginning of Section 3, we have provided a thorough analysis of UNITER-base and LXMERT models as a deep dive into Vision+Language pre-training. To summarize our key findings:

(i) In single-stream model, deeper layers lead to more intertwined multimodal fusion; while the opposite trend is observed in two-stream model.

(ii) Textual modality plays a more important role than image in making final decisions, consistent across both single- and two-stream models.

(iii) In single-stream model, a subset of heads organically evolves to pivot on cross-modal interaction and alignment, which on the other hand is enforced by model design in two-stream model.

(iv) Visual relations are inherently registered in both single- and two-stream pre-trained models.

(v) Rich linguistic knowledge is naturally encoded, even though the models are specifically designed for multimodal pre-training.

We provide additional guidelines in the Appendix. For future work, we plan to perform model compression via pruning attention heads based on the analysis and observations in this work.

References

1. van Aken, B., Winter, B., Löser, A., Gers, F.A.: How does bert answer questions? a layer-wise analysis of transformer representations. In: CIKM (2019) 4
2. Alberti, C., Ling, J., Collins, M., Reitter, D.: Fusion of detected objects in text for visual question answering. In: EMNLP (2019) 3
3. Anderson, P., He, X., Buehler, C., Teney, D., Johnson, M., Gould, S., Zhang, L.: Bottom-up and top-down attention for image captioning and visual question answering. In: CVPR (2018) 5, 12
4. Antol, S., Agrawal, A., Lu, J., Mitchell, M., Batra, D., Lawrence Zitnick, C., Parikh, D.: Vqa: Visual question answering. In: ICCV (2015) 2
5. Bouraoui, Z., Camacho-Collados, J., Schockaert, S.: Inducing relational knowledge from bert. In: AAAI (2020) 4
6. Chen, Y.C., Li, L., Yu, L., Kholy, A.E., Ahmed, F., Gan, Z., Cheng, Y., Liu, J.: Uniter: Learning universal image-text representations. arXiv preprint arXiv:1909.11740 (2019) 1, 3, 4
7. Clark, K., Khandelwal, U., Levy, O., Manning, C.D.: What does bert look at? an analysis of bert’s attention. arXiv preprint arXiv:1906.04341 (2019) 2, 4, 10
8. Conneau, A., Kiela, D.: Senteval: An evaluation toolkit for universal sentence representations. arXiv preprint arXiv:1803.05449 (2018) 3, 13
9. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. In: NAACL (2019) 2, 3, 4, 6
10. Gan, Z., Chen, Y.C., Li, L., Zhu, C., Cheng, Y., Liu, J.: Large-scale adversarial training for vision-and-language representation learning. arXiv preprint arXiv:2006.06195 (2020) 3
11. Goyal, Y., Khot, T., Summers-Stay, D., Batra, D., Parikh, D.: Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In: CVPR (2017) 2
12. Htut, P.M., Phang, J., Bordia, S., Bowman, S.R.: Do attention heads in bert track syntactic dependencies? arXiv preprint arXiv:1911.12246 (2019) 4
13. Hudson, D.A., Manning, C.D.: Gqa: a new dataset for compositional question answering over real-world images. In: CVPR (2019) 2
14. Jawahar, G., Sagot, B., Seddah, D.: What does bert learn about the structure of language? In: ACL (2019) 5
15. Jiang, Z., Xu, F.F., Araki, J., Neubig, G.: How can we know what language models know? arXiv preprint arXiv:1911.12543 (2019) 3
16. Kovaleva, O., Romanov, A., Rogers, A., Rumshisky, A.: Revealing the dark secrets of bert. In: EMNLP (2019) 2, 4
17. Krishna, R., Zhu, Y., Groth, O., Johnson, J., Hata, K., Kravitz, J., Chen, S., Kalantidis, Y., Li, L.J., Shamma, D.A., et al.: Visual genome: Connecting language and vision using crowdsourced dense image annotations. IJCV (2017) 4
18. Lee, K.H., Chen, X., Hua, G., Hu, H., He, X.: Stacked cross attention for image-text matching. In: ECCV (2018) 2
19. Li, G., Duan, N., Fang, Y., Jiang, D., Zhou, M.: Unicoder-vl: A universal encoder for vision and language by cross-modal pre-training. In: AAAI (2020) 1, 3
20. Li, L., Chen, Y.C., Cheng, Y., Gan, Z., Yu, L., Liu, J.: Hero: Hierarchical encoder for video+ language omni-representation pre-training. arXiv preprint arXiv:2005.00200 (2020) 3

21. Li, L.H., Yatskar, M., Yin, D., Hsieh, C.J., Chang, K.W.: Visualbert: A simple and performant baseline for vision and language. arXiv preprint arXiv:1908.03557 (2019) [1](#), [3](#)
22. Lu, J., Batra, D., Parikh, D., Lee, S.: Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In: NeurIPS (2019) [1](#), [2](#), [3](#)
23. Lu, J., Goswami, V., Rohrbach, M., Parikh, D., Lee, S.: 12-in-1: Multi-task vision and language representation learning. In: CVPR (2020) [3](#)
24. Michel, P., Levy, O., Neubig, G.: Are sixteen heads really better than one? In: NeurIPS (2019) [3](#)
25. Petroni, F., Rocktäschel, T., Lewis, P., Bakhtin, A., Wu, Y., Miller, A.H., Riedel, S.: Language models as knowledge bases? In: EMNLP (2019) [2](#), [4](#)
26. Plummer, B.A., Wang, L., Cervantes, C.M., Caicedo, J.C., Hockenmaier, J., Lazebnik, S.: Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. IJCV (2015) [5](#)
27. Su, W., Zhu, X., Cao, Y., Li, B., Lu, L., Wei, F., Dai, J.: Vl-bert: Pre-training of generic visual-linguistic representations. In: ICLR (2020) [1](#), [2](#), [3](#)
28. Suhr, A., Zhou, S., Zhang, A., Zhang, I., Bai, H., Artzi, Y.: A corpus for reasoning about natural language grounded in photographs. In: ACL (2019) [2](#)
29. Sun, C., Baradel, F., Murphy, K., Schmid, C.: Contrastive bidirectional transformer for temporal representation learning. arXiv preprint arXiv:1906.05743 (2019) [3](#)
30. Sun, C., Myers, A., Vondrick, C., Murphy, K., Schmid, C.: Videobert: A joint model for video and language representation learning. In: ICCV (2019) [3](#)
31. Talmor, A., Elazar, Y., Goldberg, Y., Berant, J.: olympics—on what language model pre-training captures. arXiv preprint arXiv:1912.13283 (2019) [3](#)
32. Tan, H., Bansal, M.: Lxmert: Learning cross-modality encoder representations from transformers. In: EMNLP (2019) [1](#), [2](#), [3](#), [4](#)
33. Tenney, I., Das, D., Pavlick, E.: Bert rediscovered the classical nlp pipeline. In: ACL (2019) [2](#), [3](#)
34. Tenney, I., Xia, P., Chen, B., Wang, A., Poliak, A., McCoy, R.T., Kim, N., Van Durme, B., Bowman, S.R., Das, D., et al.: What do you learn from context? probing for sentence structure in contextualized word representations. In: ICLR (2019) [3](#)
35. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. In: NeurIPS (2017) [1](#)
36. Voita, E., Talbot, D., Moiseev, F., Sennrich, R., Titov, I.: Analyzing multi-head self-attention: Specialized heads do the heavy lifting, the rest can be pruned. In: ACL (2019) [3](#)
37. Yu, L., Poirson, P., Yang, S., Berg, A.C., Berg, T.L.: Modeling context in referring expressions. In: ECCV (2016) [2](#)
38. Zellers, R., Bisk, Y., Farhadi, A., Choi, Y.: From recognition to cognition: Visual commonsense reasoning. In: CVPR (2019) [2](#)
39. Zhou, L., Palangi, H., Zhang, L., Hu, H., Corso, J.J., Gao, J.: Unified vision-language pre-training for image captioning and vqa. In: AAAI (2020) [3](#)
40. Zhou, X., Zhang, Y., Cui, L., Huang, D.: Evaluating commonsense in pre-trained language models. In: AAAI (2020) [4](#)