

# One does *not* fit all!

## On the Complementarity of Vision Encoders for Vision and Language Tasks

Gregor Geigle\*, Chen Cecilia Liu, Jonas Pfeiffer†, Iryna Gurevych  
Ubiquitous Knowledge Processing Lab (UKP Lab)  
Department of Computer Science and Hessian Center for AI (hessian.AI)  
Technical University of Darmstadt  
[www.ukp.tu-darmstadt.de](http://www.ukp.tu-darmstadt.de)

### Abstract

Current multimodal models, aimed at solving Vision and Language (V+L) tasks, predominantly repurpose Vision Encoders (VE) as feature extractors. While many VEs—of different architectures, trained on different data and objectives—are publicly available, they are not designed for the downstream V+L tasks. Nonetheless, most current work assumes that a *single* pre-trained VE can serve as a general-purpose encoder. In this work, we focus on analysis and aim to understand whether the information stored within different VEs is complementary, i.e. if providing the model with features from multiple VEs can improve the performance on a target task, and how they are combined. We exhaustively experiment with three popular VEs on six downstream V+L tasks and analyze the attention and VE-dropout patterns. Our analyses suggest that diverse VEs complement each other, resulting in improved downstream V+L task performance, where the improvements are not due to simple ensemble effects (i.e. the performance does not always improve when increasing the number of encoders). We demonstrate that future VEs, which are not *repurposed*, but explicitly *designed* for V+L tasks, have the potential of improving performance on the target V+L tasks.

### 1 Introduction

The dominant strategy for solving Vision+Language (V+L) tasks involves using Transformer models (Vaswani et al., 2017) that jointly attend over the representations of the respective modalities (Lu et al., 2019; Su et al., 2020; Li et al., 2020b; Chen et al., 2020; Huang et al., 2020, *inter alia*). While representation-learning of the text modality is comparatively straightforward using token embeddings,<sup>1</sup> image

representations are more difficult to learn. Given an image, a common approach is to use pre-trained Vision Encoders (VE), where the VE’s output features are passed as inputs, together with the text embeddings, into a Transformer model. The attention mechanism then learns a cross-modal representation space over the text and image features to solve the target V+L task.

Consequently, the success of a multimodal model builds heavily on the features extracted from a VE and is thus highly dependent on the VE’s architecture, training objectives (e.g. image classification, image encoding, object detection, etc.), and pre-training data. This dependency is further exacerbated for multimodal models that utilize VEs as static feature extractors (i.e. the weights of the VE are frozen), but also for models that are trained end-to-end, as the biases introduced by the architecture, objectives, and data of the VE remain.

Since many computer vision models can be *repurposed* as VEs for V+L tasks, a few prior works have focused on identifying *individual* VEs that perform the best on downstream tasks (Jiang et al., 2020; Shen et al., 2022; Eichenberg et al., 2022; Zhang et al., 2021). A common assumption is that a *single* pre-trained VE can perform the best for a target task or even serve as a general-purpose encoder for a wide range of V+L tasks. However, a natural question arises: to what extent is this assumption correct? Given that all VEs differ in architecture, objectives, and pre-training data, we hypothesize that the extracted features of multiple different VEs encode *complementary* information.

In this work, we focus on answering: 1) Do different VEs encode complementary features? 2) *How* are features from different VEs utilized by Transformers? We provide comprehensive analyses for multi-VE models and test whether combining VEs is beneficial over a single-VE setup under the viewpoint of feature complementarity. Similar to prior work that analyzed other components of V+L

\* Gregor is now affiliated with WüNLP & Computer Vision Lab, CAIDAS, University of Würzburg.

† Jonas is now affiliated with Google DeepMind.

<sup>1</sup>But still far from solved especially in multilingual settings (Rust et al., 2021; Clark et al., 2022; Xue et al., 2022)

Transformers (Bugliarello et al., 2021; Hendricks et al., 2021), we will not focus on improving the performance through ensembling like Yan et al. (2021b). Rather, we utilize combinations of VEs as the setting for answering our research questions.

We cover three popular classes of VEs in our experiments: 1) object detection models providing a feature representation of salient image parts containing objects (*Region*) (Anderson et al., 2018), 2) CNN models computing a feature map of the image for grid features (*Grid*), and 3) Vision Transformers (ViT) (Dosovitskiy et al., 2021) computing contextualized patch features of the image (*Patch*). As the downstream domain and task type can be heavily impacted by the different VEs, we probe all combinations of the **three** VEs on **six** different V+L tasks, covering retrieval, Q&A, and reasoning.

To investigate the VE complementarity and feature utilization, we analyze 1) the attention patterns across modalities and VEs, and 2) the dependency of specific VEs when performing VE-dropout during training and inference. While multi-VE seems to perform better than single-VE (which could partially attribute to the increased parameter count), we consistently observe performance gaps between different multi-VE configurations (e.g. a gap as large as 8.9 points for the same task) and no single winning combination for all task types. Our attention patterns analysis across the different VEs reveals that the distinctive information encoded in the VEs is important for different tasks, and the model composes the representations by enriching a dominant VE with complementary information of the other VEs.

To sum up, our results and analysis suggest that VEs trained on different objectives, architectures, and data can have a high impact on the model’s V+L task performance. We cannot rely on simple ensemble effects to improve performance; selecting and repurposing off-the-shelf VEs is non-trivial, which emphasizes the necessity to design VEs explicitly for V+L tasks in the future.

## 2 Related Work

**Multimodal Transformer Architectures.** Multimodal Transformer architectures can be divided into single-stream and dual-stream models (Bugliarello et al., 2021). The single-stream Transformer takes the concatenated visual and text tokens as input and processes them modality-agnostic, i.e. the self-attention jointly attends over

the tokens of both modalities. Dual-stream models use separate Transformers for each modality that are connected through a co-attention mechanism (Tan and Bansal, 2019; Lu et al., 2019), concatenated in a single-stream model on top (Singh et al., 2022; Kamath et al., 2021), or the image model output is used asymmetrically for cross-attention in the text model (Li et al., 2021, 2022).

The Faster R-CNN (Ren et al., 2015) object detector has been the dominant choice for multimodal models as a *Region* VE, where most methods propose to use it as a static feature extractor (Tan and Bansal, 2019; Lu et al., 2019; Su et al., 2020; Chen et al., 2020; Gan et al., 2020; Li et al., 2020b; Zhang et al., 2021; Cho et al., 2021), with the notable exception being Su et al. (2020) who backpropagate through the Faster R-CNN model. Less popular VEs are *Grid* (Huang et al., 2020; Kamath et al., 2021; Yan et al., 2021a; Shen et al., 2022; Eichenberg et al., 2022), and *Patch* (Kim et al., 2021; Wang et al., 2022; Eichenberg et al., 2022). In contrast to *Region* VEs, *Grid* and *Patch* VEs are commonly fine-tuned on the target V+L task, with the notable exception being Yan et al. (2021a). Following Bugliarello et al. (2021); Hendricks et al. (2021) we focus on single-stream models as they have been shown to perform on par with dual-stream models while being easier to extend to multi-VE setups.

**Comparing and Combining VEs.** Recently, several works aim to compare different VEs for V+L tasks. Jiang et al. (2020) compare *Region* and *Grid* for visual QA tasks, showing that training data, objectives and other factors all affect the downstream task performance. Shen et al. (2022); Eichenberg et al. (2022) compare different pre-trained *Grid* and *Patch* VEs building on CLIP (Radford et al., 2021). Zhang et al. (2021) compare different design choices for *Region* VEs with *Grid* VEs trained on the same data. Dai et al. (2023) compares different VEs in influence object hallucinations in caption generation. Closest to our work is the work by Yan et al. (2021b). While they also experiment with combining representations of *Grid*-, *Patch*-, and *Region* VEs, they only focus on the Visual Question Answering (VQA; Goyal et al., 2017) dataset and only use the combination of all three VEs. Our work provides a more in-depth evaluation of different multi-VE setups while experimenting with six diverse tasks, and shows that different combinations work best for each task.

**Analysis of Multimodal Transformers.** Our anal-

ysis methods draw inspiration from recent works that probe and analyze pre-trained multimodal Transformers for a better understanding of their different components (Bugliarello et al., 2021; Cao et al., 2020; Li et al., 2020a; Frank et al., 2021; Hendricks et al., 2021). Cao et al. (2020) propose a range of different probing tasks to understand the inner workings of multimodal models. Li et al. (2020a) analyze how accurate the attention heads of pre-trained models can perform visual grounding. Frank et al. (2021) mask parts of the text and image input and measure how the prediction performance changes for the respective other modality to test how symmetric the learned cross-modal connection is. Bugliarello et al. (2021); Hendricks et al. (2021) evaluate and disentangle which components of multimodal pre-training proposed in different works are important for their success. While previous work has only focused on models with a *Region* VE, we also experiment with *Grid* and *Patch* VEs.

In summary, our work is the first in-depth study of multimodal Transformers that use multiple VEs.

### 3 Multimodal Multi-VE Transformers

Recently, cross-modal attention is the dominant strategy to learn multimodal representations with V+L Transformers. In this work, we follow Bugliarello et al. (2021) and focus on the single-stream architecture, which shares the attention components across all modalities, i.e. the concatenated visual and text tokens are processed modality-agnostic. This architecture achieves state-of-the-art results (Bugliarello et al., 2021) while being easily extendable to multiple VEs, by concatenating all vision tokens.<sup>2</sup> Figure 1 illustrates our architecture.

**Multimodal input representations.** The raw data for a V+L task consists of either discrete tokens/characters (text-modality) or high-resolution pixel values (image-modality). To extract dense representations of the respective modalities we follow the standard pre-processing strategies: The *text* modality is tokenized using word-piece tokenization (Devlin et al., 2019) and mapped to their corresponding dense embedding representations. At the input to the first Transformer layer, positional embeddings are added to the respective token embeddings. For the *vision* modality, pre-trained VEs

<sup>2</sup>We concatenate all VE for analysis in §5. In practice, concatenation increases the sequence length and incurs a high computational cost. More efficient methods like resampling (Alayrac et al., 2022) can be explored in future work.

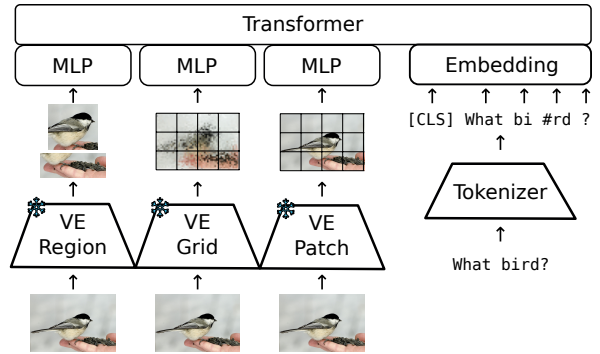


Figure 1: Our Multi-VE Architecture: Each VE produces a list of visual tokens, which are passed through MLPs and concatenated with the text embeddings. The Transformer is modality-agnostic and attends over all tokens. We freeze the VEs during training and only optimize the MLPs, embeddings, and the Transformer.

are utilized which encode the raw pixel values of the respective image into dense high-dimensional feature vectors. These VEs can either encode designated sections (e.g. *Region*), or an entire image (e.g. *Grid* and *Patch*). The extracted feature vectors are then passed through a multi-layer perceptron (MLP), and subsequently into the Transformer. This procedure can be repeated for any number of VEs of interest. In other words, the image features (from multiple VEs) and the text embeddings are concatenated and jointly passed through a shared Transformer model which learns to attend over the multimodal representations.

**V+L task training.** We place a classification head on the output of the [CLS] token (following Devlin et al. (2019)) and fine-tune the model with cross-entropy loss on the training data of the target task.

### 4 Experiments

We evaluate the impact of three different VEs on six downstream V+L tasks to assess the complementarity of different image representations. Here, we experiment with all possible combinations of the three VEs (i.e. *single VE*, *2-VE*, and *3-VE* setups).<sup>3</sup> To fairly compare the information stored in the respective VEs, we only fine-tune the multimodal models on the target V+L task in order to circumvent potentially beneficial information leaking into the multimodal model from auxiliary tasks. We therefore initialize all models with *BERT weights*

<sup>3</sup>While we report the results of a single VE, we do not aim to show that one VE outperforms others, as this would require a more controlled experimental setup, e.g. training dataset and training objectives amongst other factors (Jiang et al., 2020; Zhang et al., 2021), which is outside the scope of this work.

VE	Model	# Train	Tasks	# V. Tok.	Dim.
Region	VinVL (Zhang et al., 2021)	2.5M <sup>1</sup>	bounding box prediction, object & attribute classification	36	2054
Grid	CLIP RN50x4 (Radford et al., 2021)	400M <sup>2</sup>	image-text contrastive loss	36	2560
Patch	CLIP ViT/B-32 (Radford et al., 2021)	400M <sup>2</sup>	image-text contrastive loss	49	768

Table 1: The three VE models used in our experiments with the number of pre-training images, training objectives, the number of visual tokens (V.Tok.), and the output feature dimension. **Train Datasets:** <sup>1</sup>: Combination of multiple object detection datasets (see (Zhang et al., 2021)). <sup>2</sup>: Web-crawled & cleaned image-caption pairs (proprietary).

(Devlin et al., 2019) (base-size). We note, however, that gains can be achieved when pre-training the multimodal model on auxiliary data prior to fine-tuning on the target V+L task (Tan and Bansal, 2019; Lu et al., 2019; Chen et al., 2020, *inter alia*).

#### 4.1 Vision Encoders

We follow the standard approach and repurpose three pre-trained vision models as VEs. In a best-effort attempt for a fair setup, we use the current best publicly available models of similar sizes. Each VE has a designated, randomly initialized 2-layer perceptron (MLP) that maps the representations to the input of the Transformer and is trained on the target V+L task along with the multimodal Transformer weights. We keep the VE weights *frozen* during training. For a full summary of the VEs including pre-training data, the number of extracted tokens as well as dimensions, see Table 1.

**Region VE.** We utilize Faster R-CNN (Ren et al., 2015), an object detection model that outputs a list of bounding boxes and feature vectors for Regions of Interest—salient parts of the image that likely contain an object. Here we select the pre-trained VinVL object detector (Zhang et al., 2021),<sup>4</sup> which outperforms previous object detectors on V+L tasks. We follow Li et al. (2020b); Zhang et al. (2021) and concatenate each extracted feature vector with the corresponding normalized box coordinates and width / height. We extract the top-36 regions from the VinVL object detector.

**Grid VE.** *Grid* VEs linearize the grid feature map of a CNN (before final pooling or classification layers) to a list of visual tokens. Each visual token corresponds to a specific part of the image with image features on different scales (through different pooling operations and convolution sizes throughout the CNN) encoded in it. We use adaptive max pooling<sup>5</sup> on the feature map to reduce the number

of tokens to 36 per image. We use the CLIP CNN (RN50x4) (Radford et al., 2021) as initialization, given it’s recent success on V+L tasks (Shen et al., 2022; Eichenberg et al., 2022; Alayrac et al., 2022).

**Patch VE.** *Patch* VEs use the contextualized output representations of a Vision Transformer (ViT) (Dosovitskiy et al., 2021) as visual tokens. The ViT splits an image into uniform patches, which are used as input tokens. Different from a CNN, the ViT tokens are fixed in size throughout the model but they have a global receptive field through the ViT’s attention mechanism. We exclude the ViT’s special classification token from the Transformer input. We also utilize the CLIP-based ViT models (ViT/B-32) (Radford et al., 2021) for our *Patch*-VE. We extract all 49 tokens for the CLIP ViT due to their smaller feature dimension size.

#### 4.2 Tasks

We experiment with a set of six V+L tasks: *Image-text retrieval* (**Flickr30k** (Young et al., 2014) and **MSCOCO** (Lin et al., 2014)), *visual question answering* (**GQA** (Hudson and Manning, 2019) and **VQA2.0** (Goyal et al., 2017)), *visual entailment* (**SNLI-VE** (Xie et al., 2019)) and *memes classification* (**Hateful Memes** (Kiela et al., 2020)). For all experiments we report the mean and standard deviations over three random seeds and present training details and hyperparameters in Appendix A. We train all models with a single Nvidia V100 GPU, training a single model (all tasks, three seeds) takes approximately 10 GPU days.

#### 4.3 Results & Discussion

We report the results on the six tasks with all possible combinations of VEs in Table 2.

**No “one encoder to rule them all”.** When comparing the results of the single-VE models, there is no clear single winning VE that outperforms all other VEs across all tasks. While for QA tasks *Region* VE models perform best, for the other tasks *Grid* VE outperforms the others. We hypothesize

<sup>4</sup>Not to be confused with their Transformer.

<sup>5</sup>`torch.nn.AdaptiveMaxPool2d`

Vision Encoders	Retrieval		Question Answering		Reasoning	
	Flickr30k R@1	MSCOCO R@1	GQA Acc.	VQA Acc.	SNLI-VE Acc.	Hateful M. AUROC
Region	57.46 ±2.74	50.79 ±3.28	<b>55.32</b> ±0.33	<b>65.73</b> ±0.54	76.57 ±0.10	74.83 ±0.73
Grid	<b>66.93</b> ±3.59	<b>58.30</b> ±3.56	51.51 ±0.17	62.99 ±1.25	<b>77.32</b> ±0.11	<b>79.03</b> ±0.27
Patch	54.99 ±6.00	46.30 ±2.07	51.56 ±0.44	62.96 ±0.71	76.32 ±0.09	75.78 ±1.57
Region+Grid	63.43 ±5.85	54.87 ±6.88	55.08 ±0.44	66.30 ±1.52	77.66 ±0.11	78.68 ±1.82
Region+Patch	58.60 ±4.44	<b>58.73</b> ±4.02	<b>55.58</b> ±0.09	<b>67.05</b> ±0.42	76.60 ±0.19	75.87 ±0.63
Grid+Patch	<b>67.53</b> ±2.07	56.44 ±3.80	51.55 ±0.16	62.64 ±0.30	77.39 ±0.35	<b>79.88</b> ±0.95
Region+Grid+Patch	62.30 ±2.04	58.33 ±2.51	54.39 ±0.59	66.82 ±1.57	<b>77.87</b> ±0.24	78.81 ±0.38
With VE-Dropout Training §5.5						
Region+Grid	55.11 ±13.44	54.28 ±4.97	<b>54.91</b> ±0.26	64.72 ±3.93	77.07 ±0.21	75.75 ±0.93
Region+Patch	51.53 ±7.75	52.07 ±4.44	54.46 ±0.70	<b>65.07</b> ±0.63	76.42 ±0.15	73.57 ±0.69
Grid+Patch	<b>63.13</b> ±4.16	<b>55.56</b> ±3.92	51.80 ±0.25	60.62 ±1.38	<b>77.41</b> ±0.16	<b>77.30</b> ±0.43

Table 2: Mean and standard deviation over three seeds. Metrics: for retrieval the average recall at 1 between image-text and text-image retrieval, for Hateful Memes AUROC, and accuracy otherwise. Best single- and multi-VE setup is **bolded** and overall best score is underlined. We also report the results for VE-Dropout Training (see §5.5).

that the object-centric regions are useful for QA tasks, which focus on specific elements, while the uniform grid encoding might be useful for retrieval and other tasks that look at the entire image. Isolating *why* certain VEs are useful for specific tasks, and the role of training objectives, data, and architecture, requires a controlled setup of training VEs from scratch (Jiang et al., 2020), which we leave to future work. Interestingly, the *Patch* VE never achieves the best performance, which aligns with previous findings by Shen et al. (2022); Eichenberg et al. (2022). These single-VE results demonstrate that each VE encodes different types of information that impacts the downstream performance.

**VEs can complement each other.** When combining the representations from different VEs, we witness improvements across all V+L tasks. Interestingly, MSCOCO benefits greatly from combining the two weakest VEs (i.e. *Region* and *Patch*), surpassing their corresponding single VE results by 7.94 and 12.43 points respectively, achieving the best performance on this task. Although the *Patch* VE never achieves the best performance in single-VE setups, it provides complementary information in combination with the best performing VE, achieving the best overall performance for many tasks.<sup>6</sup> However, we see that simply using more encoders does *not* guarantee improvements as is evident by the 3-encoder model. While the 3-encoder model consistently achieves near-best

<sup>6</sup>Grid+Patch performs better than Grid for Flickr30k, Region+Patch performs better than Region for GQA, etc.

results, it is rarely the best model (only on 1 out of 6 tasks). This result shows that simply using more encoders does *not* guarantee improvement (i.e. model performance is not monotonically improving with more VEs.). Hence, it is unlikely that the improvements are due to an ensemble effect.

**One does not fit all.** In summary, we see that neither one VE alone nor a fixed combination of VEs gives the best results for the entire breadth of V+L tasks. This shows the current limitations of repurposed vision encoders and highlights the need for encoders designed specifically for V+L tasks.

## 5 Analysis

To better understand how the representations are combined in different multi-VE setups, we analyze the *flow of attention*, *phrase-to-image grounding*, and the robustness to *dropping VEs* at test time. We overload ‘cross-modality’ to include both VE-text but also VE-VE interactions for simplicity. We present the analysis for the best performing model combinations in what follows but provide a full list of results in Appendix B.

### 5.1 CLS Attention Flow

The CLS token can be seen as the fused representation of the modalities that are used for the final classification (Cao et al., 2020). We can thus estimate which VEs are important for classification by considering which VEs the CLS token attends to.<sup>7</sup>

<sup>7</sup>This is only an estimation because the modalities combine information through attention, too.

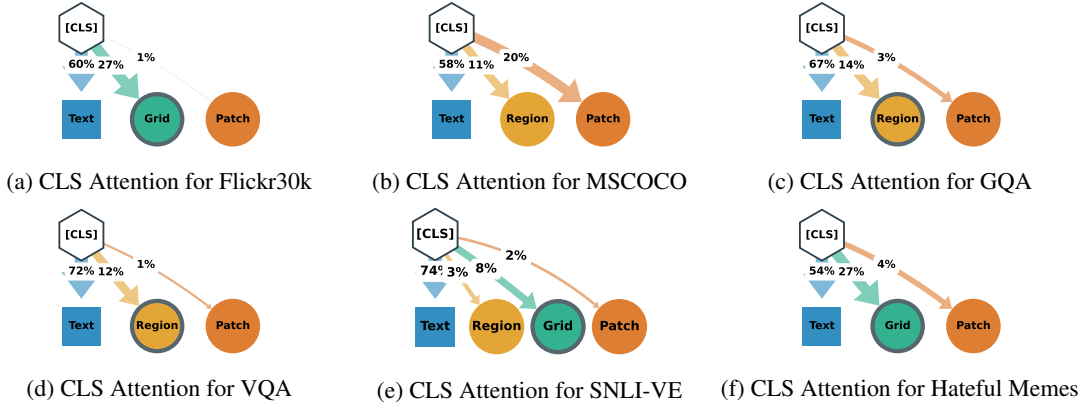


Figure 2: CLS attention (in %) to each modality/VE averaged over all heads. We add an **outline** to the VE with the best single-VE results. Numbers do not add to 100% because of CLS self-attention. We present the best multi-VE results here and all other results in Figure 8 in the Appendix.

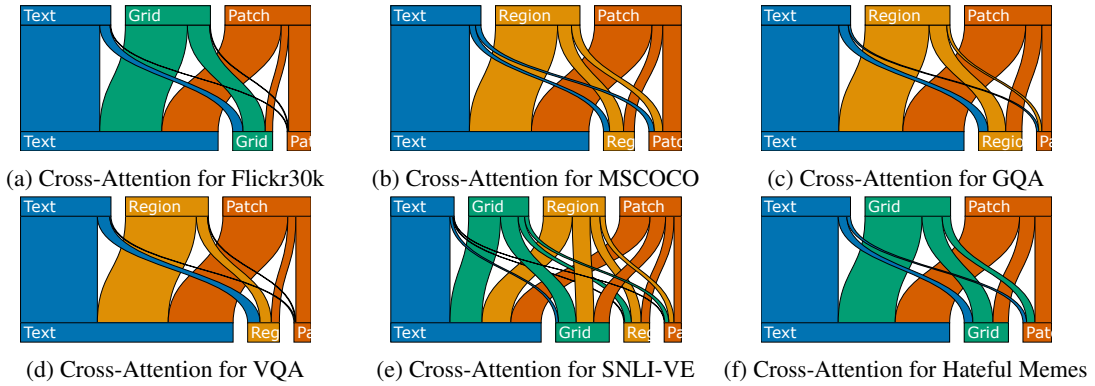


Figure 3: Cross-modal attention flow (in %) from each modality/VEs (top) to all modalities/VEs (bottom). Flow is the sum of all attention weights between two modalities, averaged over all modality tokens and attention heads. We present the best multi-VE results here and all other results in Figure 9 in the Appendix.

Following Cao et al. (2020), we compute the sum of attention from the CLS token to each modality and then average those scores over all heads.

We present the CLS attention for the best multi-VE setups in Figure 2. We see that for most tasks, the VE which performed best in the single-VE setup receives the majority of VE-attention. This suggests that one VE dominates in multi-VE setups while the others are complementary.

## 5.2 Cross-Modal Attention Flow

The attention flow between the different modalities can indicate which VEs are used by the model to reason over the input. We assume that more attention to a modality suggests that it contains useful information for others. We compute the average attention flow between two modalities  $M$  and  $N$  for an attention head as  $\frac{1}{|M|} \sum_{m \in M, n \in N} a_{m \rightarrow n}$  with  $a_{m \rightarrow n}$  as the attention weight from token  $m$  to  $n$  (excluding CLS). We average over all heads.

We present the attention flow for the best multi-

VE setups in Figure 3. Similar to §5.1, the majority of attention is paid to the VE that achieved better results in the single-VE experiment for that task.

## 5.3 Overlapping Token Surplus Attention

The attention flow between different VEs’ visual tokens that *overlap*—i.e. encode the same part of the image—can tell us if the model combines the VEs to complement the image representation. For each attention head we therefore compute the average per-token attention from a token  $t$  of *one* VE to overlapping tokens  $I|_t$  of *another* VE, and compare this—i.e. compute the surplus—to the non-overlapping tokens of that VE  $I \setminus t$  as  $(\frac{1}{|I|_t} \sum_{i \in I|_t} a_{t \rightarrow i}) - (\frac{1}{|I \setminus t|} \sum_{i \in I \setminus t} a_{t \rightarrow i})$  with  $a_{t \rightarrow i}$  being the attention weight between the tokens. We average over all tokens to get the surplus attention of an attention head for a VE pair.

We present the results for the best performing setups in Figure 4. For most VEs we can identify heads—indicated by larger dot sizes—which

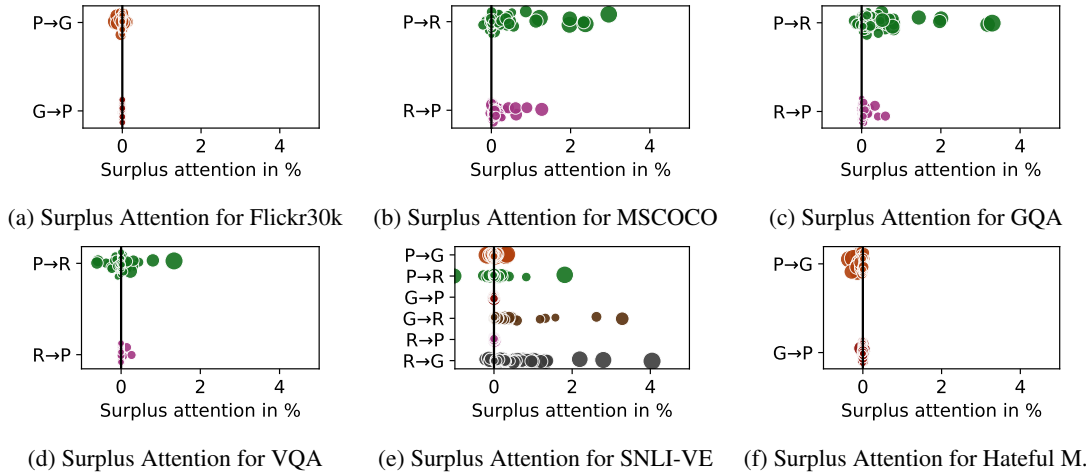


Figure 4: Surplus attention of attention heads from one VE’s tokens to another target VE’s overlapping tokens compared to the other non-overlapping tokens of the target VE. **Dot size** represents the average total attention paid to the target VE by the respective head. Positive values indicate that the respective head attributes *more* attention, negative values that *less* attention is attributed to overlapping tokens. We present the best multi-VE results here and all other results in Figure 10 in the Appendix. (Abbreviations: **Region**, **Grid**, **Patch**).




Figure 5: Visual Entity Grounding accuracy of all attention heads. An entity is grounded correctly to a VE if the attention weight from the phrase to the matching visual tokens is the highest over all the VE’s tokens. Other results can be found in Figure 11 in the Appendix.

attend particularly to *another* VE’s tokens. For most settings, these heads are also those which attend to the overlapping tokens of the respective other VE. While we witness a large surplus in attention for overlapping tokens between *Region* and *Grid/Patch*, this is not the case for *Grid-Patch*. This indicates that the complementarity of *Region* features is higher for the respective other VEs, which provides more evidence that training the VE on different data and objectives is important.

#### 5.4 Visual Entity Grounding

Visual grounding is the task of matching text phrases to their corresponding parts in the image. To analyze whether or not there are dominant VEs which learn to ground, we follow Li et al. (2020a) and count how often the highest attention weight from a text phrase is assigned to the corresponding


visual token. We use the gold phrase-to-bounding box annotations available for Flickr30k (Plummer et al., 2015) and GQA. Formally, a head correctly grounds a phrase to the gold box  $g$  if the *maximum attention* from the last phrase token  $t$  to any of a VE’s tokens  $I$  goes to any token  $I_{|g}$  overlapping with the gold box,<sup>8</sup> i.e. if  $\arg \max_{i \in I} a_{t \rightarrow i} \in I_{|g}$ . We calculate the accuracy by counting the number of correct groundings where  $I_{|g}$  is not empty.

We report the results for all heads of the best GQA and Flickr30k models in Figure 5. We can see that the accuracy of the dominant VE (*Grid* and *Region*, respectively) is generally higher than for *Patch*. While there is a clear pattern of the dominant VE achieving significantly higher accuracy, the complementary VE also achieves accuracies beyond 10%, indicating that the model learns to reason over, and utilize *both* VE’s representations.


#### 5.5 VE-Dropout

Our analyses in the previous sections suggest that there are dominant and complementary encoders, the former contributing the most to the model’s performance on the target task. To further evaluate the importance of the respective VEs we experiment with dropping all VE-specific features during test time. The results in Figure 6 show that drop-


<sup>8</sup>We consider all visual tokens with Intersection over Union between the token box and the gold box over 0.5 for region tokens (based on Li et al. (2020a)) and over 0.1 for grid and patch tokens (because a single object can be distributed over many tokens due to the uniform grid).



(a) Flickr30k VE Dropping (b) MSCOCO VE Dropping



(c) GQA VE Dropping (d) VQA VE Dropping



(e) SNLI-VE VE Dropping (f) Hateful M. VE Dropping

Figure 6: Relative performance decrease of 2-encoder models after dropping the entire first or second VE from the input compared to evaluation with both encoders in use. (Abbreviations: **R**egion, **G**rid, **P**atch).

ping the dominant encoder results in a catastrophic performance decrease, especially for the retrieval tasks. While QA and reasoning tasks have a 20% to 40% decrease in performance, for retrieval tasks R@1 decreases by almost 100%.


However, the detrimental performance of dropping out of the VE features at test time might be a result of the multi-VE models never being trained for this setting. Consequently, we train the 2-encoder models with VE-wise dropout per batch.<sup>9</sup> We hypothesize that this would force the model to take the complementary VE into account while being more robust during inference. As reported in Figure 7, the robustness in terms of dropping VEs improves, however, we see a slight drop in the final task performance, as reported in Table 2.<sup>10</sup>

## 6 Discussion and Future Directions

Our analyses demonstrate that, while combining multiple VEs consistently outperforms single VE setups, there is not a single VE or a fixed strategy on combining VEs that works best for all tasks. In particular, simply ensembling all VEs is rarely the optimal choice; consequently, best-performing combinations of VEs need to be identified for each

<sup>9</sup>Uniformly dropping the *first*, *second*, or *no* VE.

<sup>10</sup>We notice no significant changes in the attention patterns after VE-Dropout training (see Appendix C).



(a) MSCOCO VE Dropping (b) GQA VE Dropping

Figure 7: Relative performance decrease of 2-encoder models as in Figure 6 **after** VE-Dropout Training. Other tasks in Figure 12 in the Appendix.

individual task (§ 4.3). By further analyzing the attention patterns, we find a clear dominating VE (§ 5.5, Figures 6 & 7) that both the [CLS] (§ 5.1, Figure 2) and the multimodal tokens (§ 5.2, Figure 3) predominantly attend to, whereas the secondary VEs provide complementary information, supporting the model’s overall performance. The complementarity of VEs is highlighted by analyzing the cross-modal attention patterns for overlapping parts of the image (§ 5.3, Figure 4). VEs trained on different data and objectives (e.g. *Grid* and *Region*) cross-attend to the tokens of the respective *other* VE that encode the same parts of the image, aggregating their information. Further, the model learns to visually ground the text representations to all VEs, as demonstrated by their attention patterns (§ 5.4, Figure 5), and the text modality aggregates information from different VEs.

In summary, our results indicate that VEs, trained on different data and objectives, encode complementary information, resulting in improvements over approaches which only utilize a single VE. This indicates that VEs, explicitly designed for V+L tasks—e.g. by incorporating more diverse training data and objectives during pre-training—have the potential to significantly impact the performance on the target V+L tasks.

## 7 Conclusion

In this work, we investigated whether different VEs—based on repurposed pre-trained vision models—encode complementary information, which improves the performance on downstream V+L tasks. We experimented with three popular VE classes *Region*, *Grid*, and *Patch*, and trained models with all possible combinations on six different V+L tasks. While combining VEs improve over single-VE setups, our further analysis of attention patterns reveals that diverse VEs encode complementary information, which motivates future work on designing VEs explicitly for V+L tasks—e.g. by



incorporating more diverse datasets, and training objectives.

## Limitations

The main limitation of our concatenation-based multi-VE models is efficiency: The models are significantly slower than single-VE models because of the additional visual tokens in the input; the 3-VE model requires almost twice the time to train (in real time, not training steps) compared to the single-VE models. Also, in cases where images are not pre-encoded, multi-VE setups are significantly slower at inference time. However, as mentioned before, we concatenate the tokens for analysis purposes only (§5) and leave more efficient alternatives like resampling (Alayrac et al., 2022) to the future.

Several limitations could be investigated in the future, assuming access to a larger computational budget:

1. We focused on single-stream Transformers and did not take into account dual-stream or other multimodal Transformer architectures.
2. We only experimented with three popular VEs (one version per VE class). There are many other VEs we could investigate in the future.
3. We do not pre-train our multimodal models on intermediate, auxiliary multimodal tasks (Tan and Bansal, 2019; Lu et al., 2019; Chen et al., 2020, *inter alia*) as achieving state-of-the-art is not our goal.

## Acknowledgements

This work has been funded by the German Federal Ministry of Education and Research (BMBF) under the promotional reference 13N15897 (MISRIK) and by the LOEWE initiative (Hesse, Germany) within the emergenCITY center.

We thank Mert Tiftikci, Pooneh Mousavi, and Neha Warikoo for insightful feedback and suggestions on a draft of this paper.

## References

Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherin Millican, Malcolm Reynolds, Roman Ring, Eliza Rutherford, Serkan Cabi, Tengda Han, Zhitao Gong, Sina Samangooei, Marianne Monteiro, Jacob Menick, Sebastian

Borgeaud, Andrew Brock, Aida Nematzadeh, Sahand Sharifzadeh, Mikolaj Binkowski, Ricardo Barreira, Oriol Vinyals, Andrew Zisserman, and Karen Simonyan. 2022. [Flamingo: a visual language model for few-shot learning](#). In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022*.

P. Anderson, X. He, C. Buehler, D. Teney, M. Johnson, S. Gould, and L. Zhang. 2018. [Bottom-Up and Top-Down Attention for Image Captioning and Visual Question Answering](#). In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6077–6086.

Emanuele Bugliarelli, Ryan Cotterell, Naoaki Okazaki, and Desmond Elliott. 2021. [Multimodal Pretraining Unmasked: A Meta-Analysis and a Unified Framework of Vision-and-Language BERTs](#). *Transactions of the Association for Computational Linguistics*, 9:978–994.

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. 2020. [UNITER: UNiversal Image-Text Representation Learning](#). In *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings*, volume 12375 of *Lecture Notes in Computer Science*, pages 104–120. Springer.

Jaemin Cho, Jie Lei, Hao Tan, and Mohit Bansal. 2021. [Unifying vision-and-language tasks via text generation](#). 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 1931–1942. PMLR.

Jonathan H. Clark, Dan Garrette, Iulia Turc, and John Wieting. 2022. [Canine: Pre-training an Efficient Tokenization-Free Encoder for Language Representation](#). *Transactions of the Association for Computational Linguistics*, 10:73–91.

Wenliang Dai, Zihan Liu, Ziwei Ji, Dan Su, and Pascale Fung. 2023. [Plausible may not be faithful: Probing object hallucination in vision-language pre-training](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 2136–2148, Dubrovnik, Croatia. Association for Computational Linguistics.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of](#)

- Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4171–4186. Association for Computational Linguistics.
- 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. 2021. [An image is worth 16x16 words: Transformers for image recognition at scale](#). In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Constantin Eichenberg, Sidney Black, Samuel Weinbach, Letitia Parcalabescu, and Anette Frank. 2022. [MAGMA – multimodal augmentation of generative models through adapter-based finetuning](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 2416–2428, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Stella Frank, Emanuele Bugliarello, and Desmond Elliott. 2021. [Vision-and-Language or Vision-for-Language? On Cross-Modal Influence in Multimodal Transformers](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 7-11 November, 2021*, pages 9847–9857. Association for Computational Linguistics.
- Zhe Gan, Yen-Chun Chen, Linjie Li, Chen Zhu, Yu Cheng, and Jingjing Liu. 2020. [Large-Scale Adversarial Training for Vision-and-Language Representation Learning](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Y. Goyal, T. Khot, D. Summers-Stay, D. Batra, and D. Parikh. 2017. [Making the V in VQA Matter: Elevating the Role of Image Understanding in Visual Question Answering](#). In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6325–6334.
- Lisa Anne Hendricks, John Mellor, Rosalia Schneider, Jean-Baptiste Alayrac, and Aida Nematzadeh. 2021. [Decoupling the role of data, attention, and losses in multimodal transformers](#). *Transactions of the Association for Computational Linguistics*, 9:570–585.
- Zhicheng Huang, Zhaoyang Zeng, Bei Liu, Dongmei Fu, and Jianlong Fu. 2020. [Pixel-BERT: Aligning Image Pixels with Text by Deep Multi-Modal Transformers](#). *arXiv preprint*.
- Drew A. Hudson and Christopher D. Manning. 2019. [GQA: A New Dataset for Real-World Visual Reasoning and Compositional Question Answering](#). In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 6700–6709. Computer Vision Foundation / IEEE.
- Huaizu Jiang, Ishan Misra, Marcus Rohrbach, Erik G. Learned-Miller, and Xinlei Chen. 2020. [In Defense of Grid Features for Visual Question Answering](#). In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 10264–10273. Computer Vision Foundation / IEEE.
- Aishwarya Kamath, Mannat Singh, Yann LeCun, Gabriel Synnaeve, Ishan Misra, and Nicolas Carion. 2021. [MDETR - modulated detection for end-to-end multi-modal understanding](#). In *2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, October 10-17, 2021*, pages 1760–1770. IEEE.
- Douwe Kiela, Hamed Firooz, Aravind Mohan, Vedanuj Goswami, Amanpreet Singh, Pratik Ringshia, and Davide Testuggine. 2020. [The Hateful Memes Challenge: Detecting Hate Speech in Multimodal Memes](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Wonjae Kim, Bokyung Son, and Ildoo Kim. 2021. [Vilt: Vision-and-language transformer without convolution or region 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 5583–5594. PMLR.
- Junnan Li, Dongxu Li, Caiming Xiong, and Steven C. H. Hoi. 2022. [BLIP: bootstrapping language-image pre-training for unified vision-language understanding and generation](#). In *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA*, volume 162 of *Proceedings of Machine Learning Research*, pages 12888–12900. PMLR.
- 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.
- Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. 2020a. [What Does BERT with Vision Look At?](#) In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10,*

- 2020, pages 5265–5275. Association for Computational Linguistics.
- 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. 2020b. **Oscar: Object-Semantics Aligned Pre-training for Vision-Language Tasks**. In *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part XXX*, volume 12375 of *Lecture Notes in Computer Science*, pages 121–137. Springer.
- 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.
- Ilya Loshchilov and Frank Hutter. 2019. **Decoupled Weight Decay Regularization**. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- 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.
- Bryan A. Plummer, Liwei Wang, Chris M. Cervantes, Juan C. Caicedo, Julia Hockenmaier, and Svetlana Lazebnik. 2015. **Flickr30k Entities: Collecting Region-to-Phrase Correspondences for Richer Image-to-Sentence Models**. In *2015 IEEE International Conference on Computer Vision, ICCV 2015, Santiago, Chile, December 7-13, 2015*, pages 2641–2649.
- 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.
- Shaoqing Ren, Kaiming He, Ross B. Girshick, and Jian Sun. 2015. **Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks**. In *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pages 91–99.
- Phillip Rust, Jonas Pfeiffer, Ivan Vulic, Sebastian Ruder, and Iryna Gurevych. 2021. **How good is your tokenizer? on the monolingual performance of multilingual language models**. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021*, pages 3118–3135. Association for Computational Linguistics.
- Sheng Shen, Liunian Harold Li, Hao Tan, Mohit Bansal, Anna Rohrbach, Kai-Wei Chang, Zhewei Yao, and Kurt Keutzer. 2022. **How much can CLIP benefit vision-and-language tasks?** In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net.
- Amanpreet Singh, Ronghang Hu, Vedanuj Goswami, Guillaume Couairon, Wojciech Galuba, Marcus Rohrbach, and Douwe Kiela. 2022. **FLAVA: A foundational language and vision alignment model**. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pages 15617–15629. IEEE.
- Weijie Su, Xizhou Zhu, Yue Cao, Bin Li, Lewei Lu, Furu Wei, and Jifeng Dai. 2020. **VL-BERT: Pre-training of Generic Visual-Linguistic Representations**. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- 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.
- Zirui Wang, Jiahui Yu, Adams Wei Yu, Zihang Dai, Yulia Tsvetkov, and Yuan Cao. 2022. **SimVLM: Simple visual language model pretraining with weak supervision**. In *11th International Conference on Learning Representations, ICLR 2022*.
- Ning Xie, Farley Lai, Derek Doran, and Asim Kadav. 2019. **Visual Entailment: A Novel Task for Fine-Grained Image Understanding**. *arXiv preprint*.
- Linting Xue, Aditya Barua, Noah Constant, Rami Al-Rfou, Sharan Narang, Mihir Kale, Adam Roberts, and Colin Raffel. 2022. **ByT5: Towards a token-free**

future with pre-trained byte-to-byte models. *Transactions of the Association for Computational Linguistics*, 10(0):291–306.

Ming Yan, Haiyang Xu, Chenliang Li, Bin Bi, Junfeng Tian, Min Gui, and Wei Wang. 2021a. [Grid-VLP: Revisiting Grid Features for Vision-Language Pre-training](#). *arXiv preprint*.

Ming Yan, Haiyang Xu, Chenliang Li, Junfeng Tian, Bin Bi, Wei Wang, Weihua Chen, Xianzhe Xu, Fan Wang, Zheng Cao, Zhicheng Zhang, Qiyu Zhang, Ji Zhang, Songfang Huang, Fei Huang, Luo Si, and Rong Jin. 2021b. [Achieving Human Parity on Visual Question Answering](#). *arXiv preprint*.

Peter Young, Alice Lai, Micah Hodosh, and Julia Hockenmaier. 2014. [From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions](#). *Transactions of the Association for Computational Linguistics*, 2:67–78.

Pengchuan Zhang, Xiujun Li, Xiaowei Hu, Jianwei Yang, Lei Zhang, Lijuan Wang, Yejin Choi, and Jianfeng Gao. 2021. [VinVL: Revisiting Visual Representations in Vision-Language Models](#). In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5579–5588.

## A Training and Hyperparameters

We report the hyperparameters along with the task-specific training details.

### A.1 Hyperparameters

We report our hyperparameters in Table 3.4. For each task, we select the learning rate in  $\{2e - 5, 3e - 5, 5e - 5\}$  with the best validation performance for the model trained with all three VEs. We train all VE combinations for one task with the same hyperparameters. We use the training checkpoint with the best validation performance (computed each epoch) for testing.

Name	Value
Optimizer	AdamW (Loshchilov and Hutter, 2019)
Schedule	linear
Warmup steps	5%
Weight decay	0.05
Batch size	64
Max. text sequence	96

Table 3: Shared hyperparameters used during training for the different tasks.

Task	Learning Rate	Epochs
Flickr30k	2e-5	10
MSCOCO	2e-5	3
GQA	5e-5	6
VQA	5e-5	10
SNLI-VE	5e-5	10
Hateful Memes	3e-5	50

Table 4: The per-task hyperparameters used during training.

### A.2 Task Details

We describe the training details for each task. Unless noted otherwise, we use the standard splits of all tasks as described in Li et al. (2020b); Zhang et al. (2021).

**Flickr30k, MSCOCO:** We use a cross-encoder following Li et al. (2020b); Zhang et al. (2021). The model either receives the caption and the paired image or a random image (each with 50% chance). The task is to predict if the caption and image match. We use cross entropy as loss as the training objective.

During evaluation, we compute the logits for all possible image-caption pairs and use these scores as ranking to compute the recall at  $k$ . We evaluate MSCOCO on the 1k image test set.


**GQA, VQA, SNLI-VE, Hateful Memes:** We train all three tasks as standard classification task with cross entropy loss. For GQA, each class corresponds to a label appearing in the train, test, or validation set. We also use the *balanced* training data for GQA as it produces similar results to the much larger *unbalanced* training set with a fraction of the training time. For VQA, we follow Li et al. (2020b) and use the top-3000 labels for classification and we train the model with a multilabel objective using the relevance scores as soft probabilities. For testing, we use the maximum logit as the single predicted class.

## **B Full Analysis Results**


We present the full results for all VE combinations from the analysis of §5. Figure 8 shows the CLS attention, Figure 9 the attention flow, Figure 10 the surplus attention for overlapping tokens, and Figure 11 the visual grounding.

## **C Analysis Results after VE-Dropout Training**


We present the full results for all VE combinations from the analysis of §5 **after** VE-Dropout training. Figure 12 show the results for VE-Dropout at test time, Figure 14 the CLS attention, Figure 15 the attention flow, Figure 16 the surplus attention for overlapping tokens, and Figure 13 the visual grounding.




(a) CLS Attention Flickr30k




(b) CLS Attention MSCOCO




(c) CLS Attention GQA



(d) CLS Attention VQA



(e) CLS Attention SNLI-VE



(f) CLS Attention Hateful Memes

Figure 8: CLS attention weights (in %) averaged over all heads to the modalities. Numbers do not add to 100% because of CLS self-attention.




Figure 9: Attention flow (in %) from each modality (top) to all modalities (bottom). Flow is the sum of all attention weights between the modalities, averaged over all modality tokens and all attention heads.

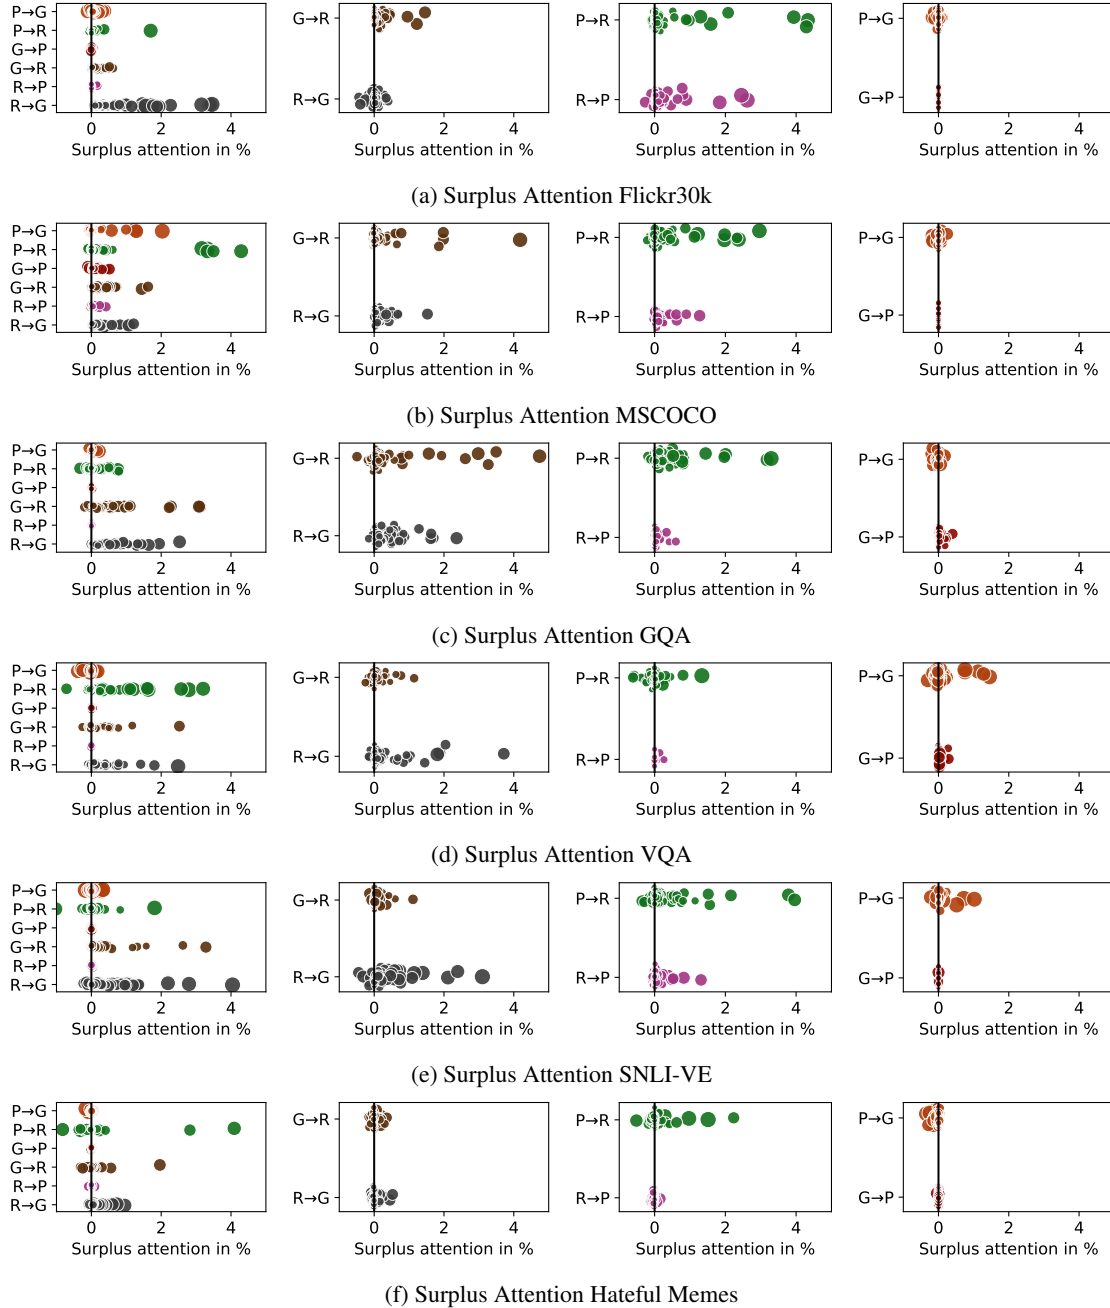


Figure 10: Surplus attention of attention heads from one VE's tokens to another target VE's overlapping tokens compared to the other non-overlapping tokens of the target VE. **Dot size** represents the average total attention paid to the target VE by each head. (Abbreviations: **R**egion, **G**rid, **P**atch).



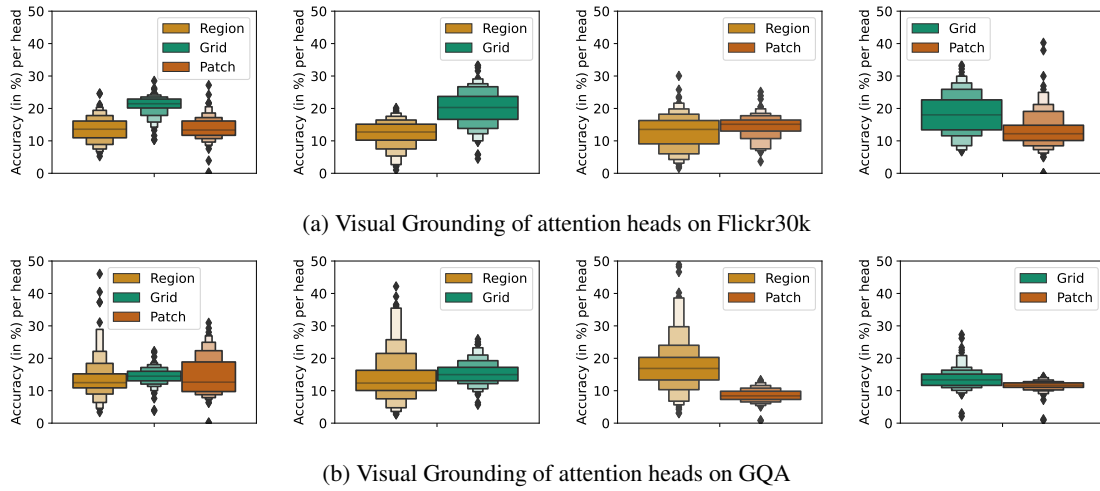


Figure 11: Visual Entity Grounding accuracy of all attention heads. An entity is grounded correctly to a VE if the attention weight from the phrase to the matching visual tokens is the highest over all the VE's tokens.

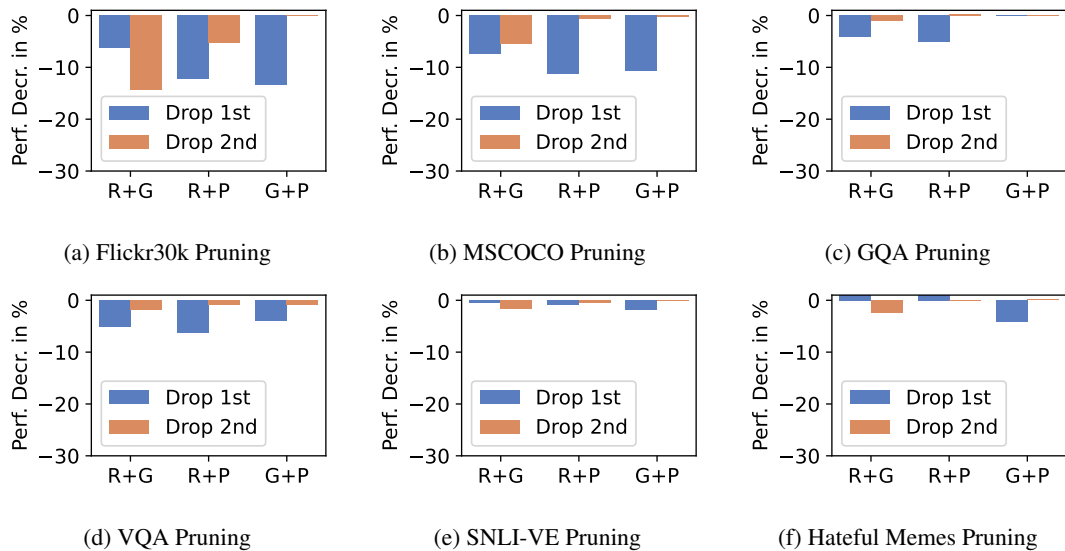




Figure 12: (After VE-Dropout Training) Relative performance decrease of 2-encoder models after dropping the entire first or second VE from the input compared to evaluation with both encoders in use. (Abbreviations: **Region**, **Grid**, **Patch**).




(a) Visual Grounding of attention heads on Flickr30k




(b) Visual Grounding of attention heads on GQA


Figure 13: (After VE-Dropout Training) Visual Entity Grounding accuracy of all attention heads. An entity is grounded correctly to a VE if the attention weight from the phrase to the matching visual tokens is the highest over all the VE's tokens.




(a) CLS Attention Flickr30k




(b) CLS Attention MSCOCO




(c) CLS Attention GQA



(d) CLS Attention VQA



(e) CLS Attention SNLI-VE



(f) CLS Attention Hateful Memes

Figure 14: (After VE-Dropout Training) CLS attention weights (in %) averaged over all heads to the modalities. Numbers do not add to 100% because of CLS self-attention.

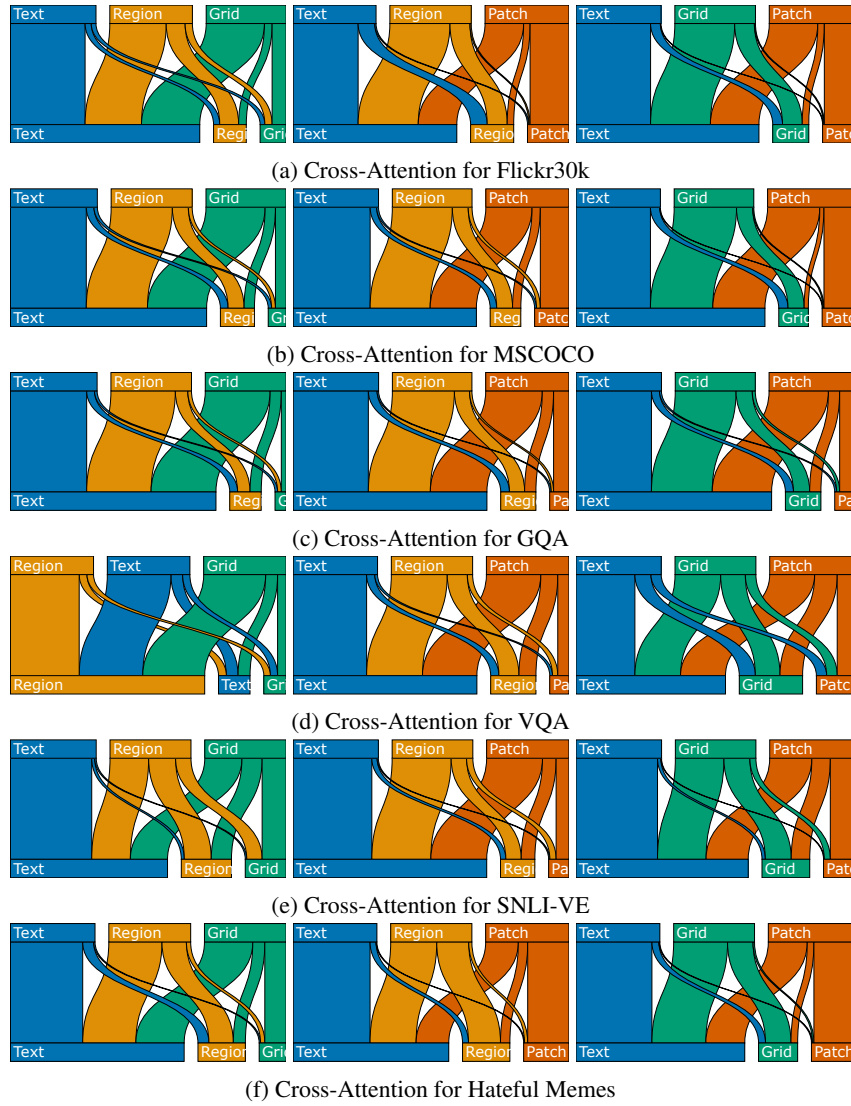


Figure 15: (After VE-Dropout Training) Attention flow (in %) from each modality (top) to all modalities (bottom). Flow is the sum of all attention weights between the modalities, averaged over all modality tokens and all attention heads.

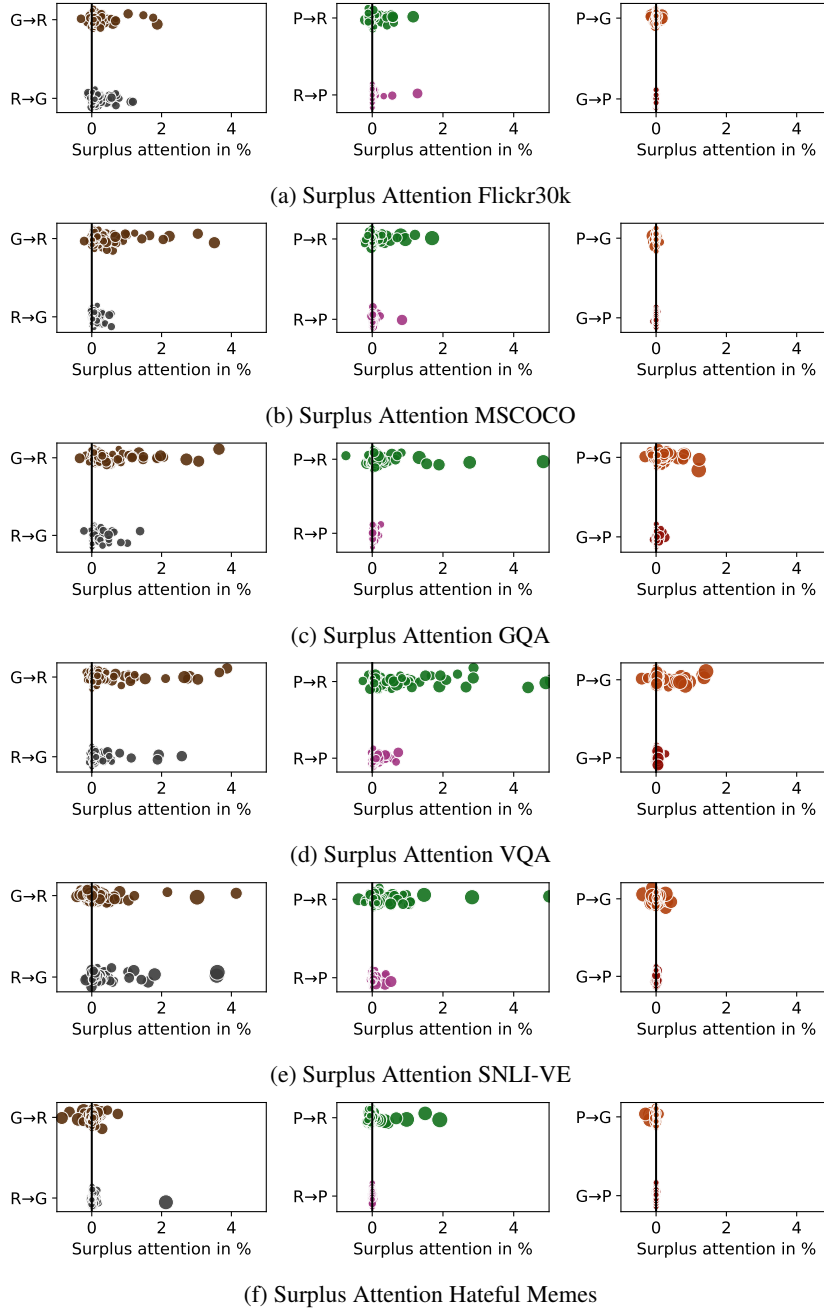


Figure 16: (After VE-Dropout Training) Surplus attention of attention heads from one VE's tokens to another target VE's overlapping tokens compared to the other non-overlapping tokens of the target VE. **Dot size** represents the average total attention paid to the target VE by each head. (Abbreviations: **R**egion, **G**rid, **P**atch).