

WHEN AND WHY VISION-LANGUAGE MODELS BEHAVE LIKE BAGS-OF-WORDS, AND WHAT TO DO ABOUT IT?

Anonymous authors

Paper under double-blind review

ABSTRACT

Despite the success of large vision and language models (VLMs) in many downstream applications, it is unclear how well they encode the compositional relationships between objects and attributes. Here, we create the Attribution, Relation, and Order (ARO) benchmark to systematically evaluate the ability of VLMs to understand different types of relationships, attributes, and order information. ARO consists of *Visual Genome Attribution*, to test the understanding of objects’ properties; *Visual Genome Relation*, to test for relational understanding; and *COCO-Order & Flickr30k-Order*, to test for order sensitivity in VLMs. ARO is orders of magnitude larger than previous benchmarks of compositionality, with more than 50,000 test cases. We present the settings where state-of-the-art VLMs behave like bags-of-words—i.e. when they have poor relational understanding, can blunder when linking objects to their attributes, and demonstrate a severe lack of order sensitivity. VLMs are predominantly trained and evaluated on large scale datasets with rich compositional structure in the images and captions. Yet, training on these datasets has not been enough to address the lack of compositional understanding, and evaluating on these datasets has failed to surface this deficiency. To understand why these limitations emerge and are not represented in the standard tests, we zoom into the evaluation and training procedures. We demonstrate that it is possible to perform well on image-text retrieval over existing datasets without using the composition and order information. This further motivates the value of using ARO to benchmark VLMs. Given that contrastive pretraining optimizes for retrieval on large datasets with similar shortcuts, we hypothesize that this can explain why the models do not need to learn to represent compositional information. This finding suggests a natural solution: composition-aware hard negative mining. We show that a simple-to-implement modification of contrastive learning significantly improves the performance on tasks requiring understanding of order and compositionality.

1 INTRODUCTION

Vision and language models (VLMs) have demonstrated high performance on dozens of well-established benchmarks (Radford et al., 2021; Li et al., 2022; Singh et al., 2022; Alayrac et al., 2022; Wang et al., 2022a;b; Zhai et al., 2022). Yet it is unclear whether high performance on these benchmarks indicates rich *compositional understanding* of either text or images. For example, does CLIP distinguish between “the man is eating the sandwich” and “the sandwich is eating the man”? Natural scenes are complex, composed of many objects and attributes, in relationships with one another. While there have been important efforts to test compositional representations of objects, attributes, and relations (Thrush et al., 2022), such efforts are based on small sets of hand-crafted examples, often combined with testing many other types of knowledge. This makes it hard to evaluate the role of relational and attributional knowledge in isolation and lacks the statistical power to quantify how well VLMs perform on granular subtypes of compositions. Here, we provide a large-scale test bed to evaluate VLMs’ attribution, relation, and order understanding. Using the test bed we create, we find significant deficiencies: many models fail to perform beyond chance-level at simple tasks requiring compositional understanding.

Many VLMs are pretrained and tested on large datasets with complex scenes and detailed captions with rich compositional structure. Yet, training on these datasets has not been enough to address

the lack of compositional understanding, and evaluating on these datasets has failed to surface this deficiency. In the recent literature, the dominant VLM training paradigm is image-text contrastive pretraining (Jia et al., 2021; Radford et al., 2021; Zhang et al., 2020) over these large pretraining datasets. Contrastive pretraining optimizes for the task of image-text retrieval, and naturally many VLMs are tested in the retrieval task. In this work, we provide an analysis of retrieval, as an evaluation and objective. We propose experiments to analyze how these models are evaluated and trained, to understand the underlying issues.

Our main contributions are three-fold:

1. **Introducing the Attribution, Relation, and Order benchmark (ARO) for fine-grained evaluation of VLMs’ relation, attribution, and order understanding.** We present four new tasks: *Visual Genome Attributions* and *Visual Genome Relations* to test the understanding of objects’ attributes and relations in complex natural scenes; and *COCO(Order)* and *Flickr30k(Order)*, to test the models’ ability to identify the correct ordering of the words in a caption (Section 2). Using these evaluations, we show that state-of-the-art VLMs fail to represent simple relations such as “to the right of” and “behind”, fail to represent the attributive difference between “the black jacket and the blue sky” versus “the blue jacket and the black sky”, and fail to represent the difference between correct and permuted captions. We provide fine-grained insights into the types of attributions and relations that models most frequently fail to understand.
2. **A critique of retrieval and contrastive pretraining.** Given we find VLMs exhibit poor compositional understanding, why have these issues not surfaced in many previous evaluations? Existing retrieval datasets are equipped with complex scenes and detailed descriptions as captions, likely full of rich compositional structure. Intriguingly, the models can perform well on retrieval without having a good compositional understanding. Our experiments (Section 3) show that models can achieve a high performance on retrieval even when the order and composition cues are removed from captions or images. Hence, it is natural that models with compositional deficiencies can still perform well on the standard evaluations. This suggests that *standard retrieval tasks are limited in their ability to assess compositional understanding of the model*, further motivating the need for our comprehensive ARO benchmark. Since contrastive pretraining optimizes for retrieval, our findings also show that models can perform well on contrastive pretraining without learning compositional information. Given our results, we argue that not learning the compositional information is a valid shortcut strategy (Geirhos et al., 2020), and VLMs have little incentive to learn to encode compositionality during contrastive pretraining.
3. **Composition-aware hard negatives can go a long way.** We propose a simple fix: mining of composition-aware hard negatives (Section 4). First, we introduce hard negatives consisting of nearest neighboring images into each batch, to force models to represent fine-grained differences between very similar scenes. Second, we introduce hard negative captions into each batch, consisting of the true captions with word order perturbed, to force models to distinguish between correct and incorrect order. Finally, we show that this simple finetuning modification provides significant improvements in model understanding of attributes and relations.

2 ATTRIBUTION, RELATION, AND ORDER (ARO) BENCHMARK: WHEN DO MODELS BEHAVE LIKE A BAG-OF-WORDS?

Whereas humans effortlessly parse natural scenes containing rich objects in relation to one another, it is unclear whether machines understand the complexity in these scenes. To do so, models must be able to correctly represent objects, their attributes, and the relations between objects. Recent research has started to probe VLMs’ for such information. Thrush et al. (2022) proposed Winoground, a dataset of hand-curated test cases that document a clear lack of compositional and pragmatic understanding in VLMs. The dataset is high quality but relatively small scale; its 400 test cases cover a wide range of linguistic phenomena (e.g., relation, attribution, pragmatics, world knowledge), making it hard to render statistically significant results about relational and attributive abilities. Here, we introduce a large dataset for testing VLMs’ ability to encode relations and attributes of objects along with order information. Our large focused dataset facilitates fine-grained and statistically strong analyses of model understanding.

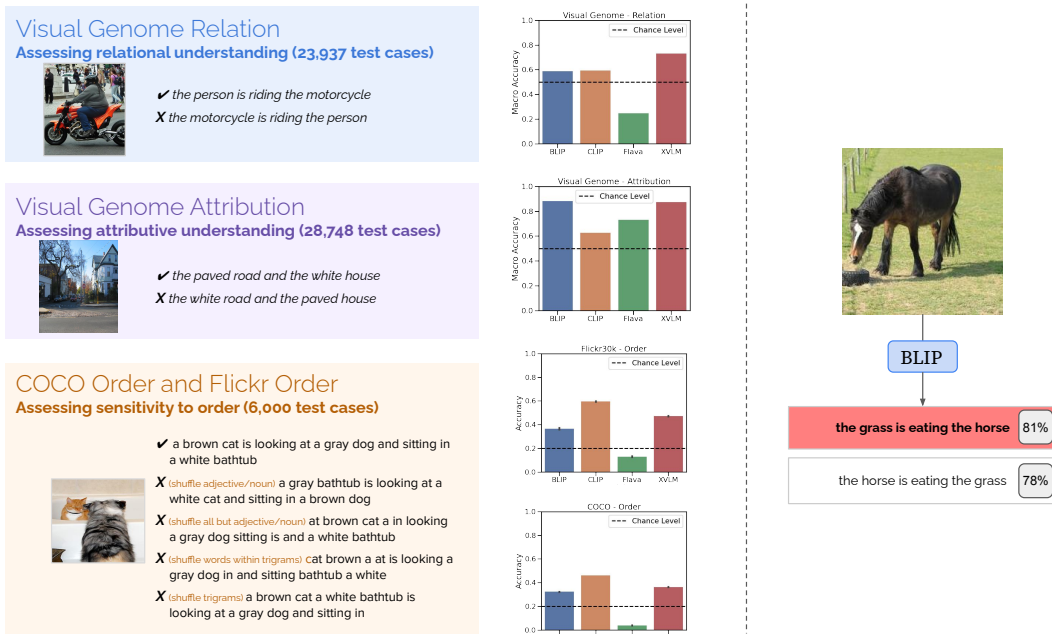


Figure 1: ARO(Attribution, Relation and Order) a benchmark to test composition and order understanding. We present four large scale tasks to test the model’s relational, attributive, and order understanding. These datasets probe the models’ ability to pick the correct ordering of the constituents of a relation, e.g. by asking the model to pick between ‘the horse is eating the grass’ vs ‘the grass is eating the horse’. Existing VLMs exhibit intriguing deficiencies at these simple tasks: several models remain at or below chance level. For example, BLIP chooses ‘the grass is eating the horse’, with 81% probability.

2.1 NEW BENCHMARKS FOR ASSESSING RELATIONAL AND ATTRIBUTIVE UNDERSTANDING

We leverage Visual Genome (VG) (Krishna et al., 2017) – a large-scale dataset with over 100,000 images, annotated with objects, attributes, and relations – and the high quality GQA annotations (Hudson & Manning, 2019) established in past work. Building upon these, we generate two novel datasets for probing relation and attribution understanding:

- **Visual Genome Relation.** Given an image and a constituent relation of the form X relation Y , we test whether the model can pick the correct order. Specifically, we probe models to pick between X relation Y and Y relation X with test cases from various relations, such as prepositional relations (e.g. “the **man** is behind the **tree**” vs “the **tree** is behind the **man**”) and verbs (e.g. “the **man** eating the **sandwich**” vs “the **sandwich** is eating the **man**”).
- **Visual Genome Attribution.** We test the ability to attribute the properties of objects appropriately. For instance, we ask the model to pick between “the **crouched man** and the **open door**” and “the **open man** and the **crouched door**”.

We extract images with 48 relations including *eating*, *sitting on*, *inside*, and *below*, with 23, 937 test cases in total; and 117 unique attribute pairs including ‘gray vs wood’, ‘open vs white’, and ‘small vs brown’, with 28, 748 test cases in total. The details of the dataset generation procedure are presented in Appendix A, along with the full list of relations with the count statistics. In brief, we mined Visual Genome for cases in which both of the constituent objects for the relations/attributes took a meaningfully large space in the image, and extracted the smallest bounding box containing both of the constituents, presenting this cropped image alongside the correct and swapped relation/attribute annotation. Each test case is thus made of an image (e.g., the image of a man eating a sandwich) a correct caption (e.g., “the man is eating the sandwich”) and a grammatically correct, but swapped, caption (e.g., “the sandwich is eating the man”). For each of the test cases in these datasets, we quantify the performance of each model in identifying the correct caption from the two alternatives, i.e. chance level performance here is 50%. Examples of these tests can be seen in Figure 1.

2.2 NEW BENCHMARKS FOR ASSESSING ORDER SENSITIVITY

Whereas Visual Genome Relation and Visual Genome Attribution assess the model’s ability to understand order and compositionality *related to attributes and relations*, we also wish to discern whether this is connected to a broader inability to represent word order *in general*. With this motivation, we specifically want to test the order sensitivity of models. Do models broadly exhibit any preference towards the correct ordering of the words in a scene description, or are they indifferent towards any permutation, even for those that are unreasonable?

We propose an additional stress test, to test the models’ ability to pick the correct ordering of the words within a caption. Given an image, we probe the models to pick between the correct ordering of a caption versus alternatives where the words are reordered in systematic ways. Given four systematic permutations of a given caption and the caption itself, can the model ‘pick the right caption’? We augment existing retrieval datasets to derive *COCO - Order*, and *Flickr30k - Order* (Lin et al., 2014; Young et al., 2014). To generate these datasets, we utilize four different perturbations of a caption, provided in Table 1. These perturbations largely follow prior work that evaluates language models (O’Connor & Andreas, 2021). We use Spacy (Honnibal & Montani, 2017) to perform part-of-speech tagging, when necessary for performing the perturbation.

2.3 EVALUATING VLMS ON ARO RESULTS

We evaluate four state-of-the-art VLMS: CLIP (Radford et al., 2021), BLIP (Li et al., 2022), Flava (Singh et al., 2022), and X-VLM (Zeng et al., 2022b). More details on these models can be found in Appendix B.

Models exhibit deficiencies in compositional understanding: In Figure 1, we present model performance on the Visual Genome Relation and Visual Genome Attribution evaluations. In relation tests, we observe that most models are near or below chance level, indicating severe deficiencies in relational understanding. In Appendix Table 2, we provide the performance on each relation separately. We observe interesting trends. For instance, while BLIP is relatively accurate at understanding positional relations, its performance is generally near chance level for verbs, such as ‘eating’ or ‘watching’; whereas CLIP generally performs at chance level on positional relations. Quantitatively, while BLIP obtains 66% macro accuracy for spatial relations, it obtains 56% accuracy in verbs. In contrast, CLIP achieves 56% in spatial relations, and 61% in verbs. In attribution tests, although BLIP (88%) and XLVM (87%) perform remarkably well, CLIP (62%) is again close to chance level. While Flava is reasonably good at attribution (73%), its performance is below-chance for relations(25%). Overall, the VLMS exhibit significant deficiencies in compositional understanding, particularly relational understanding. These deficiencies motivate our interest in probing whether these models are failing to represent compositional information *in particular* — e.g. failing to represent order of relation constituents — or whether models are in fact failing to represent word order more broadly.

Table 1: List of perturbations used in order sensitivity experiments.

Perturbation Type	Example
Original Caption	remarkable scene with a blue ball behind a green chair
Shuffle nouns and adjectives	green ball with a remarkable chair behind a blue scene
Shuffle everything but nouns and adjectives	remarkable scene behind a blue ball with a green chair
Shuffle trigrams	a green chair remarkable scene with a blue ball behind
Shuffle words within each trigram	scene with remarkable a ball blue a green behind chair

Models have little to no preference toward correctly formed sentences: In Figure 1, we present the COCO/Flickr30k order task and the performance of the tested VLMS. Given an image, the VLM must pick between the original caption and four alternative captions, augmented with the 4 perturbations listed in Table 1 respectively; thus, chance-level performance is 20%. Given the randomness over the permutations of captions, we repeat the experiment with 5 different seeds, and present the mean performance and error bars. All numerical values can be found in Appendix 5. Overall, models exhibit different levels of preference towards the correct ordering. For instance, while BLIP performs relatively well on the earlier tasks compared to CLIP, we observe here that its

performance is much closer to chance level. Similarly, while Flava obtains a good performance with the Attribution task, its performance is below chance-level for the COCO-Flickr30k Order task.

Connection to prior evaluations with text-condition image generation: We observe that CLIP largely cannot identify the correct ordering of constituents in a relation. We take the view that this is in line with prior observations, where it is shown that text-conditioned image generators that use CLIP as the text encoder, struggle with generating images that are faithful to the relations in the provided description (Conwell & Ullman, 2022; Ramesh et al., 2022; Saharia et al., 2022). Our results provide a potential explanation, indicating the problem may lie with CLIP’s inability to encode order.

3 WHY DO MODELS BEHAVE LIKE BAG-OF-WORDS? A CRITIQUE OF RETRIEVAL AND CONTRASTIVE PRETRAINING

Given that VLMs exhibit poor compositional and order understanding, why have these issues not surfaced in many previous evaluations? Most VLMs are consistently evaluated over image-to-text retrieval, and demonstrate impressive performance. In this section, we demonstrate why retrieval can be problematic, both as an evaluation and as an objective. We first show that models can perform well in the text-image retrieval evaluations on the existing large-scale datasets, without using order or composition information. It is thus natural that the issues of lack of compositional and order information have been masked by the celebrated performance. Next, we discuss the connection between contrastive pretraining on large datasets to the task of retrieval, and argue that models may not have an incentive to learn composition and order.

3.1 LIMITATIONS OF RETRIEVAL AS AN EVALUATION

Even though existing retrieval datasets are equipped with complex scenes and detailed descriptions, it is unclear how much complexity the models need to understand to perform well on this task. In particular, do the models have to use compositional information to perform well in a large-scale retrieval task? With the purpose of understanding what it does and does not take to perform well on the task of retrieval, we propose evaluations with modified datasets. First, we propose augmentations to the existing datasets, where we remove the order and composition cues. Next, we evaluate models on these augmented versions of these datasets to understand whether it is possible to perform well without these cues.

Datasets: We zoom into two of the standard cross-modal text-image retrieval datasets, namely COCO (Lin et al., 2014) and Flickr30k (Young et al., 2014). Following prior work (Li et al., 2022), we use the Karpathy split (Karpathy & Fei-Fei, 2015) for both of the datasets. We test the models on the test splits; where COCO contains 5k images, and Flickr30k contains 1k images. We report Recall@1 and Recall@5 for both datasets.

Experimental Protocol: Our goal is to understand whether models need order information / compositional understanding to perform well on existing text-image retrieval datasets. We test our hypothesis on the augmented versions of the existing datasets. We propose two augmented retrieval evaluations (see Figure 2):

1. **Perturbing the order and composition information in the captions:** To understand if models need the compositional information in captions to perform well, we aim to remove these, and test whether models can perform well without this information. In this case, we take the permutations of the words in a given caption, using the strategies specified in Table 1. For instance, we take the COCO dataset, and shuffle all of the words in all of the captions, and compute the retrieval performance over this caption-modified dataset. Given that all words are shuffled, the compositional structure within the captions is generally altered, e.g. “the sandwich is eating the man” becomes “eating the sandwich man a”.
2. **Perturbing the order and composition information in the images:** Similar to captions, to understand if models need the compositional information in the images to perform well, we aim to test the models in the absence of these. As it is harder to manipulate entities in an image, we resort to a more severe method, and take the permutations over the patches of an image. For instance, we split images into 9 equally sized patches and take a permutation of those patches to form the new images. Again, this should alter most compositional structure within an image; for instance

an object that is “below” another can move into an arbitrary point in an image. We then compute the retrieval performance with this image-modified dataset. We explore three such strategies: splitting the image into 4 equally-sized rows and shuffling, splitting into 4 equally-sized columns and shuffling, and splitting into 9 equally-sized patches and shuffling.

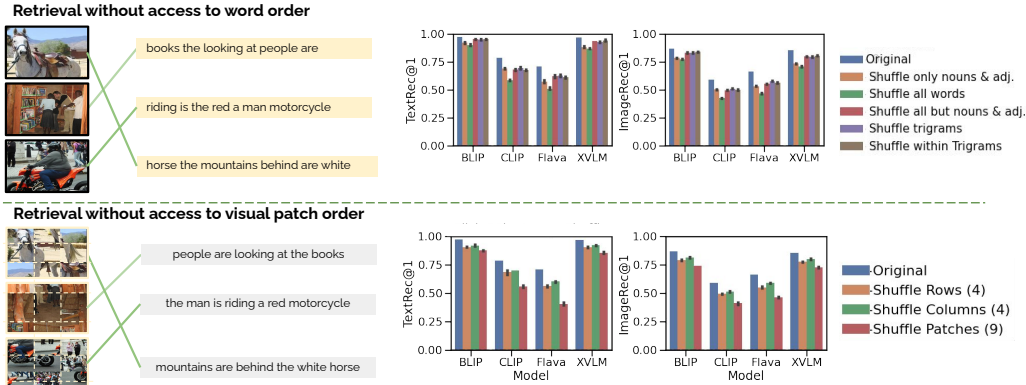


Figure 2: **Retrieval without access to order information.** We show that models can achieve substantially high performance even when order information is destroyed. In particular, in datasets where the captions are augmented with order perturbations, models show marginal performance degradation compared to the original dataset.

Models can achieve high performance even when order information is inaccessible In Figure 2, we present the retrieval performance of existing models under different augmentation strategies, and a detailed table can be found in Appendix 3.4. For each perturbation, we augment the dataset and compute the retrieval performance 3 times with different random seeds, and report standard errors bars. Notably, across all of the mentioned perturbation strategies, most models lose marginal performance in performing the retrieval task with the perturbed caption, or the perturbed image. These results demonstrate that it is possible to obtain high performance in a retrieval task without utilizing the compositional structure, even in the case of these arguably large-scale datasets.

3.2 LIMITATIONS OF RETRIEVAL AND CONTRASTIVE PRETRAINING AS AN OBJECTIVE

To understand why these deficiencies emerge in the first place, we discuss the training procedure of VLMs. Most state-of-the-art VLMs, at their core, are trained with a contrastive loss (Radford et al., 2021; Chen et al., 2020; Zeng et al., 2022a; Li et al., 2021; 2022; Zhang et al., 2020) on a large pretraining dataset scraped from the web. We hypothesize that the lack of compositional understanding can be attributed to the way the models are trained. Here, we discuss the connection between the contrastive pretraining to retrieval, to better understand the underlying phenomenon.

Quoting from Radford et al. (2021): “CLIP pre-trains for the task of image-text retrieval on our noisy web-scale dataset”. The goal in contrastive text-image pretraining is to optimize the model to identify the matching pairs of caption and text; namely *retrieval*. Simple as it is, training models with a retrieval objective, on large pretraining datasets has demonstrated remarkable off-the-shelf performance, for instance in tasks requiring single object recognition.

Although the existing large retrieval/pretraining datasets can have long and detailed captions, our results demonstrate that it is possible to perform well on these datasets without using compositional structure. Accordingly, these results provide evidence that *models can achieve high performance in retrieval objectives, thus also obtaining low contrastive loss, without order information, unless the datasets are carefully designed.* These datasets are designed to cover a large conceptual and semantic space, in order to make models broadly useful for downstream tasks; yet these datasets are not designed to contain many images with captions containing similar words that must be differentiated. Without such alternatives in the dataset, the task can be solved without taking order information into account — and behaving like a bag-of-words becomes a high-reward strategy. There is a large body of evidence showing that neural networks are prone to exploiting shortcut strategies (Geirhos et al., 2020; 2019), and our findings demonstrate that not learning the order information is a valid shortcut strategy, for general purpose retrieval/captioning datasets. We thus argue that it is unclear what should

incentivize models trained with a contrastive loss to learn to pay attention to order structure, unless the datasets or algorithms are carefully designed with this consideration.

4 A SIMPLE FIX: COMPOSITION-AWARE HARD NEGATIVES

Our analysis of the retrieval objective and datasets leads to a natural solution: generating hard negatives for contrastive learning (Robinson et al., 2021; Kalantidis et al., 2020). We propose a natural extension of CLIP’s contrastive objective to alleviate some of the issues discovered in the previous sections. To make CLIP sensitive to the word ordering and better at capturing compositions, we introduce strong alternatives:

1. **Generation of negative captions:** For each image-caption pair, we generate a negative caption by swapping different linguistic elements: *noun phrases, nouns, adjectives, adverbs, verb phrases*. For example, the caption “The man is eating the sandwich and the woman is watching the television” either becomes “The woman is eating the sandwich and the man is watching the television” (noun swapping) or “The man is watching the sandwich and the woman is eating the television” (verb phrase swapping).
2. **Sampling strong alternative images:** To generate images that are strong alternatives to the images in a batch, we first use CLIP to compute the pairwise similarity between all images in the dataset. During training, for each image in the batch we sample one of the $K = 3$ nearest neighbors as the strong alternative image. The sampled alternative images (and the respective captions and negative captions) are added to the batch.

We first briefly describe the standard contrastive methodology used in CLIP: Let $f_i : \mathcal{X}_{\text{image}} \rightarrow \mathbb{R}^d$ be the image encoder, $f_t : \mathcal{X}_{\text{text}} \rightarrow \mathbb{R}^d$ be the text encoder for a VLM. In the CLIP objective, given a batch of N image-caption pairs, the goal is to predict the actual pairings that occurred. Given a batch of N images $\mathbf{I}_N = \{I_1, I_2, \dots, I_N\}$ and N captions $\mathbf{T}_N = \{T_1, T_2, \dots, T_N\}$, CLIP first computes the matrix of cosine similarities, denoted by $\mathbf{S} \in \mathbb{R}^{N \times N}$, where each item $S_{j,k}$ computes the cosine similarity between the image j and caption k , namely $S_{j,k} = \frac{f_i(I_j) \cdot f_t(T_k)}{\|f_i(I_j)\| \cdot \|f_t(T_k)\|}$. Using these similarities, the row-wise and column-wise cross-entropy losses are computed.

For composition-aware mining of hard negatives, we propose a simple modification to the CLIP objective. In Figure 3 we give an overview of the procedure. Given a batch of images \mathbf{I}_N , captions \mathbf{T}_N , we first generate negative captions \mathbf{T}_N^- , and then concatenate the two matrices to obtain $(\tilde{\mathbf{T}}_{2N})$. Next, we again compute the similarity matrix $\tilde{\mathbf{S}} \in \mathbb{R}^{N \times 2N}$. Here, the row-wise and column-wise cross-entropy losses are computed as in CLIP, with the difference that we do not compute the loss for the negative captions column-wise (as there is no matching image for a negative caption).

Experimental protocol: Due to the computational cost of training CLIP from scratch, we only focused on finetuning experiments. Specifically, we finetune the ViT-B/32 variant of CLIP on the COCO dataset with hard negatives (NegCLIP). As an ablation, we also perform finetuning on COCO, without the sampled hard negatives to disentangle the effect of finetuning. The details of finetuning, hyperparameter selection and the ablation are provided in Appendix C.

Evaluation: We propose two main sets of evaluations. First, we evaluate models on four of the order and composition sensitive tasks, namely *Visual Genome Relation*, *Visual Genome Attribution*, *COCO & Flickr30k Order*. In addition to these, to ensure that the model is still comparable to the original CLIP, we perform evaluations on five downstream tasks: CIFAR10, 100 (Krizhevsky et al., 2009) and ImageNet (Deng et al., 2009) for image classification; and Flickr30k and COCO for cross-modal retrieval.

Results: In Figure 3, we provide a comparison of CLIP to NegCLIP with a radarplot for an overview; numerical values can be found in Appendix Table 6 with the additional ablation on a CLIP model fine-tuned on MSCOCO without negative samples. NegCLIP does not suffer in downstream tasks, and it improves the performance on VG-Relation from 63% to 81%, on VG-Attribution from 62% to 71%, on COCO-Order from 46% to 86%, and on Flickr30k-Order from 59% to 91%. Further, in VG-Relations, NegCLIP becomes the best model, in comparison to all other models, and in VG-Attribution it becomes comparable to X-VLM and BLIP.

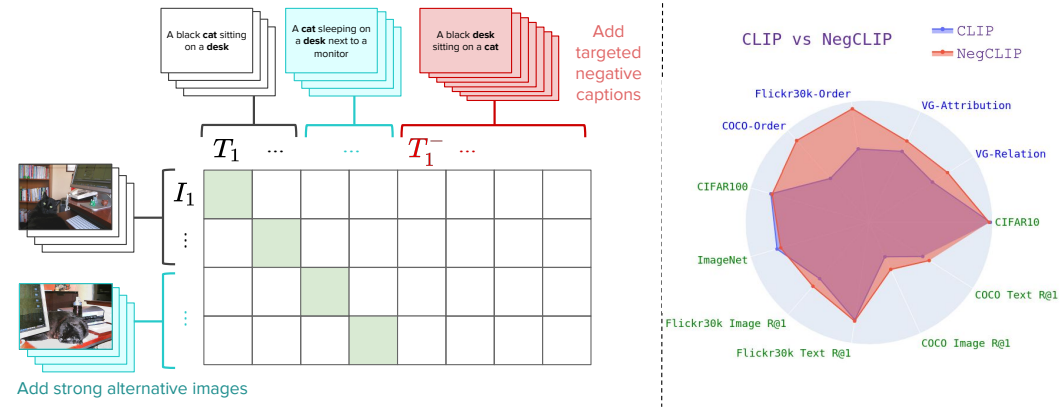


Figure 3: **Finetuning CLIP with targeted alternatives.** We propose a straightforward extension of CLIP. For each image, we sample strong alternatives among the dataset using nearest neighbors, and we create targeted negative captions to enhance order sensitivity. This method improves CLIP in a suite of compositional tasks, while not substantially hurting performance in important downstream tasks. Text in blue: tasks we proposed. Text in green: standard downstream evaluations.

Overall, we observe that NegCLIP does not have a substantial loss in performance on the downstream tasks, yet provides substantial gains on the order sensitive tasks. Our results highlight that mining targeted and cheap negatives can lead to substantial improvements in compositional tasks, without causing a loss of performance in important downstream tasks. While contrastive learning provided substantial improvements in representation learning, it is unclear if simply scaling the size of pretraining datasets will be as efficient as exploring algorithmic improvements, especially in learning the compositional structure. We believe our results provide evidence that seeking such modifications can bring various improvements. We do not suggest that ours is the best possible model for encoding relations; rather, in order to represent these bits of compositional information, our addition of targeted negatives should be considered as a promising option when training models with the contrastive objective.

5 RELATED WORK

Visio-linguistic compositionality: Understanding what aspects of language and vision VLMs capture is the main objective of several recent papers. For example, Frank et al. (2021) ablations suggest that information sharing between the text and vision modalities is not balanced, as the representations from the text encoder are more influenced by the vision modality than the other way round. Parcalabescu et al. (2021) show that VLMs have difficulties in counting objects in images. In terms of the evaluation part of our paper, Winoground (Thrush et al., 2022) presents the nearest neighbor to our work. Winoground is a carefully curated dataset that aims to evaluate compositional and pragmatics language understanding of VLMs. Results on Winoground show that VLMs perform around chance level over a set of image-text matching tasks. Our main goal is to propose a fine-grained evaluation of relation and attribution understanding with a large scale dataset; hence we believe our dataset complements Winoground (See Table 2 in the Appendix). Namely, while Winoground proposes a great breadth; here, we attempt to focus on the depth of testing for relation and attributive understanding. Similar to our methodology, Bogin et al. (2021) utilizes Visual Genome to test VQA systems for compositional generalization; in comparison here our purpose is to implement tasks that can probe any VLM.

Closer in spirit to our Visual Genome - Relations dataset, Conwell & Ullman (2022) demonstrate the lack of competency of a text-to-image model, DALL-E, in generating images faithful to the relationships described in the textual prompt. A Benchmark for Compositional Visual Reasoning (Zerroug et al., 2022) designs a benchmark to evaluate vision models’ understanding of simple compositional attributes of abstract shapes. GQA (Hudson & Manning, 2019) is a visual question-answering dataset derived from Visual Genome (Krishna et al., 2017) to test visual question answering systems for scene understanding with compositional question answering. Recently, Saharia et al. (2022) released a set of

prompts, collectively called DrawBench, as a benchmark for text-to-image models’ ability to generate images faithful to the given challenging prompts, involving a set of compositional tests. Previous evaluations, while informative, were limited by small data size; Winoground has 400 examples, limiting the statistical strength of the fine-grained conclusions. Our ARO dataset (50,000 test cases) is 2 orders of magnitude larger, enabling us to characterize the performance of VLM different types of compositions. Moreover, previous works did not quantify how using composition-aware hard negatives in contrastive learning could affect the ability of VLM to understand composition.

Order information in language and vision: Several works have highlighted the lack of word order sensitivity in large language models [Hessel & Schofield \(2021\)](#); [O’Connor & Andreas \(2021\)](#); [Pham et al. \(2021\)](#); [Sinha et al. \(2021\)](#). [Sinha et al. \(2021\)](#) show that pre-training BERT with sentences with shuffled words, marginally affect the performance on down stream tasks. [Pham et al. \(2021\)](#) show that some of the tasks in GLUE ([Wang et al., 2018](#)) can be solved even when disregarding word order. Our analysis of image retrieval leans in a very similar direction, suggesting the need for more careful benchmarks. [O’Connor & Andreas \(2021\)](#) show that for long-range contexts, models use content words and local co-occurrence statistics to make predictions. [Ettinger \(2020\)](#) uses a set of psycholinguistic tasks to evaluate the linguistic and contextual information used by BERT, showing BERT’s insensitivity to elements like negation. On the vision side, [Brendel & Bethge \(2019\)](#) show that a bag-of-local-features model performs almost as well as their state-of-the-art counterparts. Closer to our experiments, the work of [Tejankar et al. \(2021\)](#) shows that training contrastive vision language models using only bag-of-words in place of the caption does not significantly hurt performance on zero-shot classification. Our work generalize these results, showcasing the general limits of vision language models when dealing with relations, attributes and shuffled captions.

Negative mining and contrastive learning: The use of hard negatives has been successful in metric and representation learning ([Harwood et al., 2017](#); [Wu et al., 2017](#); [Ge, 2018](#)). Furthermore, hard negatives have also been shown to improve contrastive learning ([Kalantidis et al., 2020](#); [Robinson et al., 2021](#)). Our work differs in that we propose to explore hard negatives with contrastive learning, particularly in the context of vision-language models and compositional abilities. Note that [Li et al. \(2021\)](#) and [Li et al. \(2022\)](#) use negative mining by selecting pairs of items with high similarity. However, in light of our results where we test BLIP, this strategy alone does not seem to be enough to effectively train the model to deal with relationships and word order.

6 CONCLUSION

In this work, we evaluate VLMs for the ability to encode composition and order structure, introducing large-scale test beds to generate fine-grained and statistically strong insights. We show that models struggle with relation, attribution, and order understanding, and our datasets revealed different limitations of models. We show that models can achieve high performance on the task of crossmodal retrieval without needing to learn order and composition information. Given that contrastive pretraining optimizes models for the task of retrieval, we argue that this can explain why VLMs need not learn to encode order and compositional cues. Using these insights, we proposed a simple modification to the training procedure, namely composition-aware hard negative mining. Through several evaluations, we demonstrate that by generating composition-aware hard negatives during model training, we can improve the compositional and order understanding of VLMs.

Overall, our work demonstrates the importance of the interaction between the pretraining objective and large datasets VLMs are trained on. In this work, we focused on finetuning of VLMs for demonstrating the use of the composition-aware negative mining. For future work, we are interested in further exploring composition-aware contrastive pretraining of VLMs. Given that VLMs are trained with large datasets with rich text corpora, the limited language understanding of these models are intriguing. While here we specifically focused on contrastive learning, in light of our findings, studying the interaction between different pretraining objectives and compositional understanding is an interesting future avenue. Our results further highlight the importance of rich evaluations of VLMs. We hope future VLMs will release results on these fine-grained evaluations in addition to standard tasks, and more fine-grained evaluations will be developed.

ETHICAL STATEMENT

We are aware that contrastive models like CLIP and BLIP are biased (see, for example, the work by Wolfe & Caliskan (2022)). Our work uses standard datasets in the literature and we do not foresee additional issues created by our fine-tuning or methodology. Nonetheless, our dataset, based on Visual Genome, inherits stereotypical patterns present in the original data (e.g., the original dataset has more images of men than women (Krishna et al., 2017)).

REPRODUCIBILITY STATEMENT

The code to reproduce all experiments, along with the code to generate the datasets and tasks we propose will be released at `redacted` for anonymity. Experiments on caption perturbation have been run with three different seeds to take into account the randomness of the permutation methods. All of the models that are used were obtained from the checkpoints released with the respective paper. Particularly, we obtained the checkpoints released at BLIP¹, X-VLM², CLIP³, and Flava⁴ repositories.

Experiments with CLIP and NegCLIP have not been run multiple times due to the computational requirements. However, the experiments have been run using fixed seed, so they can be replicated by other researchers. In addition to this, our NegCLIP is implemented as a fork of the open clip project⁵. The code will be released with the same license.

REFERENCES

- Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katie Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. *arXiv preprint arXiv:2204.14198*, 2022. 1
- Ben Bogin, Shivanshu Gupta, Matt Gardner, and Jonathan Berant. Covr: A test-bed for visually grounded compositional generalization with real images. *arXiv preprint arXiv:2109.10613*, 2021. 8
- Wieland Brendel and Matthias Bethge. Approximating CNNs with bag-of-local-features models works surprisingly well on imagenet. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=SkfMWhAqYQ>. 9
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pp. 1597–1607. PMLR, 2020. 6
- Colin Conwell and Tomer D Ullman. Testing relational understanding in text-guided image generation. *arXiv preprint arXiv:2208.00005*, 2022. 5, 8
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pp. 248–255. Ieee, 2009. 7
- Allyson Ettinger. What bert is not: Lessons from a new suite of psycholinguistic diagnostics for language models. *Transactions of the Association for Computational Linguistics*, 8:34–48, 2020. 9
- Stella Frank, Emanuele Bugliarello, and Desmond Elliott. 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)*. Association for Computational Linguistics, nov 2021. URL <https://arxiv.org/abs/2109.04448>. 8

¹<https://github.com/salesforce/BLIP>

²<https://github.com/zengyan-97/X-VLM>

³<https://github.com/openai/CLIP>

⁴<https://huggingface.co/facebook/flava-full>

⁵https://github.com/mlfoundations/open_clip/

- Weifeng Ge. Deep metric learning with hierarchical triplet loss. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 269–285, 2018. 9
- Robert Geirhos, Patricia Rubisch, Claudio Michaelis, Matthias Bethge, Felix A. Wichmann, and Wieland Brendel. Imagenet-trained CNNs are biased towards texture; increasing shape bias improves accuracy and robustness. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=Bygh9j09KX>. 6
- Robert Geirhos, Jörn-Henrik Jacobsen, Claudio Michaelis, Richard Zemel, Wieland Brendel, Matthias Bethge, and Felix A Wichmann. Shortcut learning in deep neural networks. *Nature Machine Intelligence*, 2(11):665–673, 2020. 2, 6
- Ben Harwood, Vijay Kumar BG, Gustavo Carneiro, Ian Reid, and Tom Drummond. Smart mining for deep metric learning. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 2821–2829, 2017. 9
- Jack Hessel and Alexandra Schofield. How effective is bert without word ordering? implications for language understanding and data privacy. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pp. 204–211, 2021. 9
- Matthew Honnibal and Ines Montani. spaCy 2: Natural language understanding with Bloom embeddings, convolutional neural networks and incremental parsing. To appear, 2017. 4
- Drew A Hudson and Christopher D Manning. Gqa: A new dataset for real-world visual reasoning and compositional question answering. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 6700–6709, 2019. 3, 8, 14
- Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *International Conference on Machine Learning*, pp. 4904–4916. PMLR, 2021. 2
- Yannis Kalantidis, Mert Bulent Sariyildiz, Noe Pion, Philippe Weinzaepfel, and Diane Larlus. Hard negative mixing for contrastive learning. *Advances in Neural Information Processing Systems*, 33: 21798–21809, 2020. 7, 9
- Andrej Karpathy and Li Fei-Fei. Deep visual-semantic alignments for generating image descriptions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3128–3137, 2015. 5
- Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International journal of computer vision*, 123(1):32–73, 2017. 3, 8, 10
- Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 7
- Junnan Li, Ramprasaath Selvaraju, Akhilesh Gotmare, Shafiq Joty, Caiming Xiong, and Steven Chu Hong Hoi. Align before fuse: Vision and language representation learning with momentum distillation. *Advances in neural information processing systems*, 34:9694–9705, 2021. 6, 9
- Junnan Li, Dongxu Li, Caiming Xiong, and Steven C. H. Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *ICML*, 2022. 1, 4, 5, 6, 9, 14
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pp. 740–755. Springer, 2014. 4, 5

- Joe O’Connor and Jacob Andreas. What context features can transformer language models use? In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pp. 851–864, Online, August 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.acl-long.70. URL <https://aclanthology.org/2021.acl-long.70>. 4, 9
- Letitia Parcalabescu, Albert Gatt, Anette Frank, and Iacer Calixto. Seeing past words: Testing the cross-modal capabilities of pretrained V&L models on counting tasks. In *Proceedings of the 1st Workshop on Multimodal Semantic Representations (MMSR)*, pp. 32–44, Groningen, Netherlands (Online), June 2021. Association for Computational Linguistics. URL <https://aclanthology.org/2021.mmsr-1.4>. 8
- Thang Pham, Trung Bui, Long Mai, and Anh Nguyen. Out of order: How important is the sequential order of words in a sentence in natural language understanding tasks? In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pp. 1145–1160, Online, August 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.findings-acl.98. URL <https://aclanthology.org/2021.findings-acl.98>. 9
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pp. 8748–8763. PMLR, 2021. 1, 2, 4, 6, 14
- Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022. 5
- Joshua Robinson, Ching-Yao Chuang, Suvrit Sra, and Stefanie Jegelka. Contrastive learning with hard negative samples. *ArXiv*, abs/2010.04592, 2021. 7, 9
- Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S Sara Mahdavi, Rapha Gontijo Lopes, et al. Photorealistic text-to-image diffusion models with deep language understanding. *arXiv preprint arXiv:2205.11487*, 2022. 5, 8
- Amanpreet Singh, Ronghang Hu, Vedanuj Goswami, Guillaume Couairon, Wojciech Galuba, Marcus Rohrbach, and Douwe Kiela. Flava: A foundational language and vision alignment model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 15638–15650, 2022. 1, 4, 14
- Koustuv Sinha, Robin Jia, Dieuwke Hupkes, Joëlle Pineau, Adina Williams, and Douwe Kiela. Masked language modeling and the distributional hypothesis: Order word matters pre-training for little. In *EMNLP*, 2021. 9
- Ajinkya Tejankar, Bichen Wu, Saining Xie, Madian Khabsa, Hamed Pirsiavash, and Hamed Firooz. A fistful of words: Learning transferable visual models from bag-of-words supervision. *arXiv preprint arXiv:2112.13884*, 2021. 9
- Tristan Thrush, Ryan Jiang, Max Bartolo, Amanpreet Singh, Adina Williams, Douwe Kiela, and Candace Ross. Winoground: Probing vision and language models for visio-linguistic compositionality. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 5238–5248, 2022. 1, 2, 8
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman. Glue: A multi-task benchmark and analysis platform for natural language understanding. *arXiv preprint arXiv:1804.07461*, 2018. 9
- Junke Wang, Dongdong Chen, Zuxuan Wu, Chong Luo, Luowei Zhou, Yucheng Zhao, Yujia Xie, Ce Liu, Yu-Gang Jiang, and Lu Yuan. Omnivl: One foundation model for image-language and video-language tasks. *arXiv preprint arXiv:2209.07526*, 2022a. 1
- Wenhui Wang, Hangbo Bao, Li Dong, Johan Bjorck, Zhiliang Peng, Qiang Liu, Kriti Aggarwal, Owais Khan Mohammed, Saksham Singhal, Subhojit Som, et al. Image as a foreign language: Beit pretraining for all vision and vision-language tasks. *arXiv preprint arXiv:2208.10442*, 2022b. 1

- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pp. 38–45, Online, October 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-demos.6. URL <https://aclanthology.org/2020.emnlp-demos.6>. 14
- Robert Wolfe and Aylin Caliskan. American == white in multimodal language-and-image ai. In *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society, AIES '22*, pp. 800–812, New York, NY, USA, 2022. Association for Computing Machinery. ISBN 9781450392471. doi: 10.1145/3514094.3534136. URL <https://doi.org/10.1145/3514094.3534136>. 10
- Chao-Yuan Wu, R Manmatha, Alexander J Smola, and Philipp Krahenbuhl. Sampling matters in deep embedding learning. In *Proceedings of the IEEE international conference on computer vision*, pp. 2840–2848, 2017. 9
- Peter Young, Alice Lai, Micah Hodosh, and Julia Hockenmaier. 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, 2014. 4, 5
- Yan Zeng, Xinsong Zhang, and Hang Li. Multi-grained vision language pre-training: Aligning texts with visual concepts. *ArXiv*, abs/2111.08276, 2022a. 6
- Yan Zeng, Xinsong Zhang, and Hang Li. Multi-grained vision language pre-training: Aligning texts with visual concepts. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato (eds.), *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 25994–26009. PMLR, 2022b. 4, 14
- Aimen Zerroug, Mohit Vaishnav, Julien Colin, Sebastian Musslick, and Thomas Serre. A benchmark for compositional visual reasoning. *arXiv preprint arXiv:2206.05379*, 2022. 8
- Xiaohua Zhai, Xiao Wang, Basil Mustafa, Andreas Steiner, Daniel Keysers, Alexander Kolesnikov, and Lucas Beyer. Lit: Zero-shot transfer with locked-image text tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 18123–18133, 2022. 1
- Yuhao Zhang, Hang Jiang, Yasuhide Miura, Christopher D Manning, and Curtis P Langlotz. Contrastive learning of medical visual representations from paired images and text. *arXiv preprint arXiv:2010.00747*, 2020. 2, 6

A PROPOSED DATASETS

A.1 GENERATING THE DATASETS

We follow the below steps while generating the datasets:

1. We go through the scene graphs annotated in GQA (Hudson & Manning, 2019).
2. First, we identify candidate objects in all scenes. Namely, we require the objects to be large enough such that they are recognizable in an image. We enforce a heuristical criterion, where we discard all objects that have a width lower than $\frac{1}{4}$ of the width of the entire image, or with a height lower than $\frac{1}{4}$ of the height of the entire image.
3. For VG-Relation, we identify all pairings of objects, where we make sure that the pair of objects are not from the same category (say, both are not simultaneously “dogs”). Similarly, for VG-Attribution, we identify all pairs of objects that are modified by at least 1 attribute, where both objects and attributes are different from each other.
4. After identifying the pairs of objects, we extract the smallest bounding box containing both of the objects from the scene. This is to minimize the distraction in the rich scenes in Visual Genome.
5. Finally, for each identified pairs of objects, we fill the preset templates. For relations, we fill the templates of “the [object 1] is [relation] [object 2]” for the true caption, and “the [object 2] is [relation] [object 1]” for the false caption. For attributes, we fill the templates of “the [attribute 1] [object 1] and the [attribute 2] [object 2]” and “the [attribute 2] [object 1] and the [attribute 1] [object 2]”.
6. For Visual Genome Relations, we post-process the relations to remove symmetric relations; such as “near” or “next to”.

Overall, this process results in a set of 23, 937 test cases for relations, and 28, 748 test cases for attributes.

B MODELS

1. CLIP (Radford et al., 2021): We use the ‘ViT-B/32’ variant of CLIP, released at <https://github.com/openai/CLIP/>.
2. BLIP (Li et al., 2022): We report results for the ‘Base’ variants of BLIP. For experiments with COCO, we use the version finetuned on COCO; and for experiments with Flickr30k, we use the version finetuned on Flickr. For Visual Genome experiments, we use the variant that gives better results, that is the one finetuned on COCO. We use the checkpoints released at <https://github.com/salesforce/BLIP>.
3. X-VLM (Zeng et al., 2022b): Similar to BLIP, we report results with the ‘Base’ variants of X-VLM. For experiments with COCO, we use the version finetuned on COCO; and for experiments with Flickr30k, we use the version finetuned on Flickr. We use the checkpoints released at <https://github.com/zengyan-97/X-VLM/>.
4. FLAVA (Singh et al., 2022): We use the model released at the Huggingface Transformers (Wolf et al., 2020) library. We follow the tutorial shared by authors, and use the “flava-full” variant⁶.

C NEGATIVE MINING

C.1 NEGATIVE TEXT MINING

Starting from the MSCOCO image dataset, we use spacy to swap the position of two elements of the caption. These elements can be either nouns, adjectives, adverbs, verb phrases and noun phrases (only if the noun phrase is composed by three tokens or more, to not overlap with noun swapping).

⁶<https://github.com/apsdehal/flava-tutorials/blob/main/winoground-flava-example.ipynb>

Table 2: Fine-grained results in Visual Genome Relation dataset.

	CLIP	NegCLIP	CLIP-FT	XVLM	BLIP	Flava	Freq
Accuracy	0.59	0.8	0.64	0.73	0.59	0.24	
Spatial Relationships							
Accuracy	0.56	0.66	0.57	0.74	0.66	0.34	
above	0.48	0.60	0.54	0.80	0.64	0.55	269
at	0.59	0.93	0.71	0.72	0.49	0.15	75
behind	0.56	0.29	0.34	0.82	0.77	0.28	574
below	0.56	0.46	0.48	0.74	0.69	0.44	209
beneath	0.80	0.70	0.70	0.80	0.70	0.40	10
in	0.63	0.89	0.63	0.73	0.72	0.09	708
in front of	0.54	0.75	0.70	0.66	0.55	0.78	588
inside	0.50	0.91	0.67	0.69	0.72	0.12	58
on	0.52	0.86	0.58	0.86	0.76	0.12	1684
on top of	0.43	0.75	0.58	0.85	0.79	0.19	201
to the left of	0.49	0.50	0.50	0.52	0.51	0.50	7741
to the right of	0.49	0.50	0.50	0.52	0.49	0.51	7741
under	0.64	0.43	0.54	0.86	0.73	0.27	132
Verbs							
Accuracy	0.61	0.86	0.66	0.73	0.56	0.2	
carrying	0.33	0.83	0.75	0.75	0.67	0.08	12
covered by	0.47	0.36	0.36	0.61	0.58	0.56	36
covered in	0.79	0.50	0.50	0.14	0.29	0.14	14
covered with	0.56	0.56	0.50	0.56	0.50	0.19	16
covering	0.39	0.58	0.45	0.67	0.55	0.06	33
cutting	0.75	0.83	0.83	0.67	0.25	0.00	12
eating	0.57	1.00	0.67	0.62	0.52	0.00	21
feeding	0.90	0.80	0.80	0.60	0.30	0.20	10
grazing on	0.10	0.90	0.30	0.60	0.40	0.50	10
hanging on	0.79	1.00	0.93	0.93	0.79	0.00	14
holding	0.58	0.97	0.79	0.67	0.44	0.27	142
leaning on	0.67	1.00	1.00	0.75	0.58	0.08	12
looking at	0.84	1.00	0.68	0.68	0.55	0.26	31
lying in	0.47	1.00	0.60	0.87	0.67	0.00	15
lying on	0.60	0.88	0.50	0.93	0.75	0.17	60
parked on	0.67	0.86	0.38	0.76	0.86	0.00	21
reflected in	0.64	0.71	0.57	0.50	0.43	0.43	14
resting on	0.38	0.85	0.23	0.92	0.54	0.15	13
riding	0.71	0.98	0.78	0.82	0.41	0.02	51
sitting at	0.62	1.00	0.88	0.88	0.46	0.00	26
sitting in	0.57	0.96	0.78	0.87	0.83	0.30	23
sitting on	0.58	0.97	0.78	0.94	0.73	0.14	175
sitting on top of	0.50	0.90	0.80	1.00	0.80	0.10	10
standing by	0.67	0.92	0.67	0.83	0.67	0.67	12
standing in	0.73	0.98	0.69	0.69	0.49	0.05	59
standing on	0.60	1.00	0.63	0.83	0.73	0.06	52
surrounded by	0.64	0.71	0.64	0.71	0.64	0.79	14
using	0.84	1.00	1.00	0.68	0.58	0.00	19
walking in	0.70	1.00	0.70	0.60	0.50	0.00	10
walking on	0.79	1.00	0.79	0.84	0.42	0.05	19
watching	0.45	0.55	0.27	0.59	0.68	0.36	22
wearing	0.47	0.99	0.88	0.68	0.48	0.64	949

For each caption we thus build a set of 5 possible negative captions (note that, a negative caption might have less than 5 negative captions if there are not enough elements to swap). If a caption has zero negative captions, we remove it from the dataset. During training, at each epoch for each caption we sample one of its negative captions as additional element of the batch.

Table 3: Performance in the text shuffled retrieval task. Performance is averaged over 3 seeds, and standard deviations are reported next to each mean.

Strategy	Model	Text Recall@1	Text Recall@5	Image Recall@1	Image Recall@5
COCO Text Shuffle					
No Shuffling	BLIP	0.814	0.953	0.635	0.855
Shuffle only nouns and adj.	BLIP	0.710 ± 0.006	0.907 ± 0.001	0.513 ± 0.001	0.768 ± 0.001
Shuffle all words	BLIP	0.690 ± 0.004	0.899 ± 0.002	0.505 ± 0.002	0.763 ± 0.001
Shuffle all but nouns and adj.	BLIP	0.767 ± 0.002	0.935 ± 0.001	0.579 ± 0.000	0.817 ± 0.000
Shuffle trigrams	BLIP	0.762 ± 0.001	0.933 ± 0.000	0.581 ± 0.000	0.816 ± 0.000
Shuffle within Trigrams	BLIP	0.767 ± 0.003	0.934 ± 0.000	0.585 ± 0.000	0.820 ± 0.000
No Shuffling	CLIP	0.503	0.748	0.301	0.557
Shuffle only nouns and adj.	CLIP	0.420 ± 0.003	0.684 ± 0.007	0.244 ± 0.002	0.480 ± 0.001
Shuffle all words	CLIP	0.341 ± 0.001	0.608 ± 0.005	0.205 ± 0.001	0.422 ± 0.003
Shuffle all but nouns and adj.	CLIP	0.415 ± 0.002	0.671 ± 0.001	0.248 ± 0.002	0.483 ± 0.001
Shuffle trigrams	CLIP	0.411 ± 0.006	0.673 ± 0.004	0.251 ± 0.001	0.490 ± 0.002
Shuffle within Trigrams	CLIP	0.404 ± 0.002	0.661 ± 0.004	0.243 ± 0.001	0.478 ± 0.001
No Shuffling	Flava	0.454	0.788	0.388	0.682
Shuffle only nouns and adj.	Flava	0.335 ± 0.003	0.645 ± 0.002	0.287 ± 0.001	0.566 ± 0.000
Shuffle all words	Flava	0.338 ± 0.006	0.631 ± 0.001	0.260 ± 0.001	0.526 ± 0.002
Shuffle all but nouns and adj.	Flava	0.392 ± 0.005	0.692 ± 0.003	0.317 ± 0.002	0.601 ± 0.001
Shuffle trigrams	Flava	0.385 ± 0.006	0.698 ± 0.007	0.333 ± 0.002	0.621 ± 0.001
Shuffle within Trigrams	Flava	0.394 ± 0.005	0.698 ± 0.000	0.324 ± 0.001	0.606 ± 0.001
No Shuffling	XVLM	0.791	0.947	0.610	0.848
Shuffle only nouns and adj.	XVLM	0.655 ± 0.005	0.884 ± 0.000	0.462 ± 0.001	0.731 ± 0.000
Shuffle all words	XVLM	0.633 ± 0.006	0.879 ± 0.001	0.450 ± 0.002	0.723 ± 0.002
Shuffle all but nouns and adj.	XVLM	0.734 ± 0.004	0.928 ± 0.002	0.547 ± 0.003	0.803 ± 0.000
Shuffle trigrams	XVLM	0.727 ± 0.004	0.920 ± 0.002	0.544 ± 0.001	0.799 ± 0.004
Shuffle within Trigrams	XVLM	0.739 ± 0.007	0.929 ± 0.004	0.554 ± 0.005	0.808 ± 0.001
Flickr30k Text Shuffle					
No Shuffling	BLIP	0.972	0.999	0.869	0.974
Shuffle only nouns and adj.	BLIP	0.919 ± 0.008	0.993 ± 0.004	0.786 ± 0.003	0.949 ± 0.001
Shuffle all words	BLIP	0.902 ± 0.008	0.988 ± 0.003	0.770 ± 0.002	0.939 ± 0.001
Shuffle all but nouns and adj.	BLIP	0.950 ± 0.003	0.996 ± 0.002	0.829 ± 0.006	0.961 ± 0.001
Shuffle trigrams	BLIP	0.948 ± 0.005	0.997 ± 0.001	0.828 ± 0.001	0.965 ± 0.002
Shuffle within Trigrams	BLIP	0.953 ± 0.004	0.997 ± 0.001	0.838 ± 0.002	0.964 ± 0.001
No Shuffling	CLIP	0.784	0.950	0.591	0.835
Shuffle only nouns and adj.	CLIP	0.690 ± 0.005	0.909 ± 0.001	0.501 ± 0.003	0.770 ± 0.003
Shuffle all words	CLIP	0.587 ± 0.007	0.854 ± 0.011	0.423 ± 0.006	0.694 ± 0.002
Shuffle all but nouns and adj.	CLIP	0.678 ± 0.010	0.904 ± 0.007	0.493 ± 0.003	0.764 ± 0.002
Shuffle trigrams	CLIP	0.698 ± 0.017	0.910 ± 0.005	0.509 ± 0.004	0.775 ± 0.002
Shuffle within Trigrams	CLIP	0.680 ± 0.006	0.903 ± 0.011	0.498 ± 0.006	0.766 ± 0.001
No Shuffling	Flava	0.707	0.941	0.664	0.900
Shuffle only nouns and adj.	Flava	0.573 ± 0.021	0.869 ± 0.009	0.532 ± 0.001	0.816 ± 0.002
Shuffle all words	Flava	0.504 ± 0.002	0.817 ± 0.009	0.467 ± 0.006	0.754 ± 0.006
Shuffle all but nouns and adj.	Flava	0.622 ± 0.016	0.888 ± 0.005	0.553 ± 0.004	0.823 ± 0.002
Shuffle trigrams	Flava	0.626 ± 0.016	0.895 ± 0.003	0.578 ± 0.002	0.849 ± 0.003
Shuffle within Trigrams	Flava	0.613 ± 0.007	0.889 ± 0.003	0.564 ± 0.006	0.834 ± 0.003
No Shuffling	XVLM	0.967	1.000	0.855	0.968
Shuffle only nouns and adj.	XVLM	0.881 ± 0.013	0.987 ± 0.001	0.733 ± 0.002	0.920 ± 0.002
Shuffle all words	XVLM	0.869 ± 0.003	0.982 ± 0.003	0.708 ± 0.001	0.908 ± 0.002
Shuffle all but nouns and adj.	XVLM	0.937 ± 0.003	0.994 ± 0.003	0.798 ± 0.003	0.949 ± 0.003
Shuffle trigrams	XVLM	0.924 ± 0.005	0.995 ± 0.002	0.795 ± 0.008	0.948 ± 0.001
Shuffle within Trigrams	XVLM	0.942 ± 0.013	0.996 ± 0.002	0.807 ± 0.003	0.953 ± 0.001

C.2 NEGATIVE IMAGE MINING

Starting from the MSCOCO image dataset, we compute the pairwise similarity between all the images. Then, for each image we collect the 3 most similar images. During training, at each epoch for each image we sample one of its negative images as additional element of the batch.

Table 4: Performance in the image shuffled retrieval task. Performance is averaged over 3 seeds, and standard deviations are reported next to each mean.

Strategy	Model	Text Recall@1	Text Recall@5	Image Recall@1	Image Recall@5
COCO Image Shuffle					
No Shuffling	BLIP	0.814	0.953	0.635	0.855
Shuffle Rows (4)	BLIP	0.692 ± 0.005	0.880 ± 0.002	0.546 ± 0.001	0.790 ± 0.000
Shuffle Columns (4)	BLIP	0.705 ± 0.004	0.880 ± 0.005	0.553 ± 0.003	0.795 ± 0.001
Shuffle Patches (9)	BLIP	0.594 ± 0.003	0.817 ± 0.001	0.488 ± 0.002	0.741 ± 0.000
No Shuffling	CLIP	0.503	0.748	0.301	0.557
Shuffle Rows (4)	CLIP	0.402 ± 0.006	0.657 ± 0.003	0.252 ± 0.004	0.488 ± 0.002
Shuffle Columns (4)	CLIP	0.410 ± 0.002	0.661 ± 0.002	0.256 ± 0.001	0.493 ± 0.000
Shuffle Patches (9)	CLIP	0.284 ± 0.002	0.531 ± 0.012	0.209 ± 0.003	0.428 ± 0.002
No Shuffling	Flava	0.454	0.788	0.388	0.682
Shuffle Rows (4)	Flava	0.305 ± 0.007	0.629 ± 0.004	0.311 ± 0.004	0.588 ± 0.003
Shuffle Columns (4)	Flava	0.325 ± 0.000	0.650 ± 0.001	0.318 ± 0.000	0.603 ± 0.002
Shuffle Patches (9)	Flava	0.182 ± 0.004	0.449 ± 0.001	0.243 ± 0.001	0.500 ± 0.001
No Shuffling	XVLM	0.791	0.947	0.610	0.848
Shuffle Rows (4)	XVLM	0.664 ± 0.006	0.876 ± 0.005	0.524 ± 0.002	0.781 ± 0.003
Shuffle Columns (4)	XVLM	0.687 ± 0.003	0.883 ± 0.001	0.539 ± 0.000	0.790 ± 0.002
Shuffle Patches (9)	XVLM	0.463 ± 0.005	0.760 ± 0.003	0.471 ± 0.001	0.731 ± 0.001
Flickr30k Image Shuffle					
No Shuffling	BLIP	0.972	0.999	0.869	0.974
Shuffle Rows (4)	BLIP	0.905 ± 0.002	0.980 ± 0.005	0.791 ± 0.004	0.943 ± 0.004
Shuffle Columns (4)	BLIP	0.918 ± 0.007	0.991 ± 0.003	0.812 ± 0.004	0.954 ± 0.001
Shuffle Patches (9)	BLIP	0.875 ± 0.004	0.973 ± 0.002	0.739 ± 0.001	0.920 ± 0.003
No Shuffling	CLIP	0.784	0.950	0.591	0.835
Shuffle Rows (4)	CLIP	0.687 ± 0.020	0.895 ± 0.001	0.493 ± 0.002	0.763 ± 0.002
Shuffle Columns (4)	CLIP	0.699 ± 0.001	0.892 ± 0.006	0.511 ± 0.004	0.769 ± 0.005
Shuffle Patches (9)	CLIP	0.559 ± 0.008	0.814 ± 0.008	0.411 ± 0.011	0.683 ± 0.008
No Shuffling	Flava	0.707	0.941	0.664	0.900
Shuffle Rows (4)	Flava	0.562 ± 0.008	0.849 ± 0.008	0.550 ± 0.008	0.823 ± 0.001
Shuffle Columns (4)	Flava	0.602 ± 0.006	0.882 ± 0.003	0.588 ± 0.002	0.852 ± 0.004
Shuffle Patches (9)	Flava	0.404 ± 0.011	0.725 ± 0.006	0.463 ± 0.004	0.753 ± 0.007
No Shuffling	XVLM	0.967	1.000	0.855	0.968
Shuffle Rows (4)	XVLM	0.905 ± 0.005	0.982 ± 0.002	0.775 ± 0.001	0.935 ± 0.001
Shuffle Columns (4)	XVLM	0.921 ± 0.005	0.990 ± 0.003	0.799 ± 0.005	0.948 ± 0.002
Shuffle Patches (9)	XVLM	0.855 ± 0.007	0.971 ± 0.004	0.728 ± 0.006	0.908 ± 0.003

Table 5: Performance in the Pick the Right Caption task. Performance is averaged over 3 seeds, and standard deviations are reported next to each mean.

Dataset	Model	Accuracy
Flickr30k-PRC	BLIP	0.369 ± 0.009
Flickr30k-PRC	CLIP	0.595 ± 0.006
Flickr30k-PRC	Flava	0.129 ± 0.005
Flickr30k-PRC	XVLM	0.473 ± 0.003
Flickr30k-PRC	Chance-level	0.2
COCO-PRC	BLIP	0.321 ± 0.001
COCO-PRC	CLIP	0.460 ± 0.001
COCO-PRC	Flava	0.039 ± 0.001
COCO-PRC	XVLM	0.362 ± 0.003
COCO-PRC	Chance-level	0.2

C.3 FINE-TUNING DETAILS

For finetuning models, we build our code on https://github.com/mlfoundations/open_clip/. We finetune all models on the training split of the COCO dataset, and validate on the validation split. We finetune both CLIP-FT and NegCLIP for 5 epochs, with sweeping learning

Table 6: Here, we report the result of fine-tuning CLIP with hard-negatives. Metrics reported: Accuracy for CIFAR10, CIFAR100, ImageNet; Recall@1 for Flickr30k and COCO text-to-image and image-to-text retrieval; Macro Accuracy for VG-Relations and VG-Attribution; Accuracy for Flickr30k and COCO pick-the-right-caption tasks.

	CLIP	CLIP-FT	NegCLIP
Compositional Tasks			
VG-Relation	0.63	0.63	0.81
VG-Attribution	0.62	0.65	0.71
Flickr30k-PRC	0.59	0.50	0.91
COCO-PRC	0.46	0.36	0.86
Downstream Tasks			
CIFAR10	0.95	0.95	0.94
CIFAR100	0.80	0.80	0.79
ImageNet	0.75	0.74	0.72
Flickr30k Image R@1	0.59	0.67	0.67
Flickr30k Text R@1	0.78	0.83	0.79
COCO Image R@1	0.30	0.42	0.41
COCO Text R@1	0.50	0.59	0.56

rates in $\{1e-5, 5e-6, 1e-6\}$ and picking the models based on the retrieval performance in the COCO validation set. We use 50 steps of warmup and AdamW optimizer with a cosine-annealing learning rate schedule