

Exposing and Mitigating Spurious Correlations for Cross-Modal Retrieval

Jae Myung Kim¹, A. Sophia Koepke¹, Cordelia Schmid², Zeynep Akata^{1,3}

¹University of Tübingen ²Inria, Ecole normale supérieure, CNRS, PSL Research University
³MPI for Intelligent Systems

Abstract

Cross-modal retrieval methods are the preferred tool to search databases for the text that best matches a query image and vice versa. However, image-text retrieval models commonly learn to memorize spurious correlations in the training data, such as frequent object co-occurrence, instead of looking at the actual underlying reasons for the prediction in the image. For image-text retrieval, this manifests in retrieved sentences that mention objects that are not present in the query image. In this work, we introduce *ODmAP@k*, an object decorrelation metric that measures a model’s robustness to spurious correlations in the training data. We use automatic image and text manipulations to control the presence of such object correlations in designated test data. Additionally, our data synthesis technique is used to tackle model biases due to spurious correlations of semantically unrelated objects in the training data. We apply our proposed pipeline, which involves the finetuning of image-text retrieval frameworks on carefully designed synthetic data, to three state-of-the-art models for image-text retrieval. This results in significant improvements for all three models, both in terms of the standard retrieval performance and in terms of our object decorrelation metric. The code is available at https://github.com/ExplainableML/Spurious_CM_Retrieval.

1. Introduction

Quickly growing image and text databases necessitate an urgent need to search data efficiently. The cross-modal text-image retrieval task considers the setting in which query images are used to retrieve matching text from a text database and vice versa. For instance, for a query image that shows dogs playing with a frisbee, an image-to-text retrieval model is expected to retrieve a sentence describing the scene (e.g. “Two dogs fighting over a frisbee”).

Commonly, image-text retrieval models are trained on paired image-text data. It has been shown that text biases and correlations in the training data translate to vision-

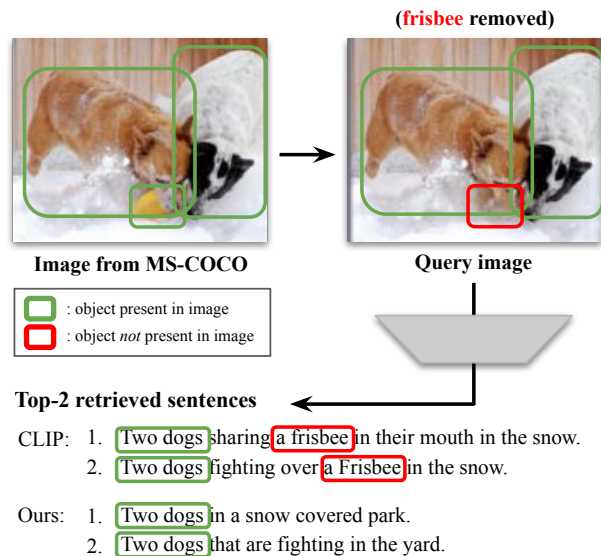


Figure 1. **Image-to-text retrieval results with CLIP and our proposed pipeline.** The query image is generated by removing an object (i.e. the frisbee). The top-2 text retrieval results show that CLIP wrongly retrieves sentences that mention the frisbee. Correct and incorrect words in the retrieved sentences are marked green and red respectively.

language models (e.g. for image captioning [30]). In this work, we investigate the impact of spurious correlations in the training data, such as frequent co-occurrences of semantically unrelated objects, on cross-modal retrieval models trained with it. We indeed observe spurious correlations in retrieved sentences describing objects not present in the query image although they tend to be somewhat related to the objects actually present in the image in the sense that they frequently appear together in the dataset. For instance, when an image of dogs *without* a frisbee is the query, the CLIP model [58] nevertheless retrieves the sentence “Two dogs sharing a frisbee in their mouth in the snow” which contains the word ‘frisbee’ based on common co-occurrences (see Fig. 1).

It is easy for humans to identify failure cases where retrieved sentences mention unrelated objects. However, the

standard retrieval evaluation protocols do not specifically measure these errors. To quantify those, we propose the object decorrelation metric ODmAP@k which captures a model’s robustness to semantically unrelated object correlations in the training data. Being able to explicitly measure this specific type of error is a first step towards mitigating it since those errors also affect the retrieval task performance.

Our object decorrelation metric uses a designated test set with carefully designed synthetic images. For this, object detections are used to remove commonly co-occurring objects with an inpainting framework [72]. Our ODmAP@k then measures if the retrieved text i) contains the objects that occur in the synthetic image, and ii) if the retrieved text does not mention any of the objects that have been removed from the original image and are not present in the synthetic image. This enables us to quantify if a model has memorized commonly co-occurring objects in the training data, or if it is actually able to retrieve text that matches the objects in the query images.

Recently, [1] identified spurious correlations in the context of Visual Question Answering (VQA) and alleviated the biases due to spurious correlations in the training data by using synthetic training data designed for the VQA task. Similarly, we propose a finetuning pipeline that mitigates the impact of frequent co-occurrences of semantically unrelated objects in the training data on the trained models. While data augmentation is commonly used for computer vision tasks, cross-modal augmentation is more challenging. In particular, we aim at formulating automatic augmentation strategies in both the image and text domains, whilst ensuring that the image-text pairs match and that our synthetically generated data challenges the memorization of spurious correlations in the training data. Our proposed finetuning pipeline improves debiasing retrieval models while having a competitive performance in a standard evaluation protocol on the MSCOCO [42] and Flickr30k [71] datasets.

To summarise, we make the following contributions: 1) We reveal that the performance of existing cross-modal retrieval models suffers from the existence of spurious object correlations in the training data. We propose the object decorrelation metric ODmAP@k to measure this correlation bias. 2) We propose a finetuning pipeline for mitigating the impact of spurious object correlations in the training data, which uses carefully designed synthesized data. 3) We demonstrate that our finetuning pipeline mitigates the model to learn the spuriousness while having a competitive performance on the standard retrieval evaluation compared with the model trained on the original dataset.

2. Related work

Cross-modal retrieval. Commonly, cross-modal retrieval methods use a learnt shared latent space to relate different

modalities to one another. This has been investigated for different modalities paired with text, such as for text-image retrieval [9, 12, 19–22, 27, 32–34, 36, 38, 40, 41, 45, 58, 63, 64, 66, 67, 73], text-video retrieval [3, 4, 13, 17, 23, 24, 49, 54, 68, 69], and text-audio retrieval [37, 44, 54, 56]. Furthermore, [29, 31, 53] have explored audio-visual-text and audio-visual retrieval. In this paper, we focus on image-text retrieval.

Image-text retrieval. Text-image retrieval models are usually trained to align representations across the text and image modalities for matching image-text pairs. Several methods have proposed different ways of quantifying the similarity between learnt visual and text embeddings [9, 12, 20–22, 27, 34, 63, 64]. In particular, [21] use a triplet loss, and [12] consider a probabilistic formulation. A different line of works has developed specialized network components which allow the modeling of relations across modalities [19, 38, 40, 67, 73]. Differently from the aforementioned works that aim at learning better cross-modal embeddings, we focus specifically on addressing the problem of biased cross-modal models due to spurious correlations in the training data.

Image-text representations can be learned by using millions of image-text pairs sourced from the internet for training [32, 59]. The contrastive alignment of images and text in the two-stream CLIP [59] and ALIGN [32] frameworks combined with (noisy) large-scale training data has resulted in impressive generalization capabilities. Consequently, the success of the CLIP model [59] has influenced domains far beyond text-image retrieval, as CLIP embeddings have been used for varied tasks, such as semantic segmentation [61], image generation [14], and image video retrieval [46], to name just a few. Furthermore, several works have built on CLIP for learning strong and generalizable vision-language representations in a dual-stream fashion which allows efficient retrieval [2, 25, 39, 51, 70]. In this paper, we apply our proposed method to CLIP and to the more recent BLIP [39] which outperforms CLIP on zero-shot image-text retrieval.

Biases in vision-language models. Exposing and mitigating biases in vision-language models is of increasing research interest. Recent works consider, for instance, societal biases [5, 65, 75], the missing correspondence of annotation [11], the language bias in VQA [8, 55], hubness in cross-modal retrieval [6], multimodal spurious correlations in VQA [1], the spuriousness in captioning [30], object hallucination in captioning [60], large-scale vision-language pretraining [74], or cross-modal retrieval in a unique e-commerce setting [47]. To reduce the multimodal spurious correlations and language bias in VQA, [1, 8, 28, 55] proposed to create additional data to balance the training dataset. Inspired by the success of data augmentation in the VQA setting, we design a setup for synthesizing data that allows us to identify and mitigate biases in trained image-text retrieval models that arise from spurious correlations in

the training data.

3. Object decorrelation framework

In this work, we propose a framework for measuring and mitigating the bias in retrieval models due to spurious object correlations in the training data. To examine the spuriousness of the retrieval models, we propose the object decorrelation metric ODmAP@k that is measured using synthetic images. In Section 3.1, we describe the process of generating synthetic images. We then explain our proposed ODmAP@k metric in Section 3.2. In Section 3.3, we provide details about our finetuning pipeline for mitigating the negative impact of object correlations in the training data.

3.1. Synthetic image generation

We consider a dataset \mathcal{D} consisting of image-text pairs $(\mathbf{x}, \mathbf{y}) \in \mathcal{D}$. We examine whether the image-to-text retrieval model retrieves a sentence based on the object in a query image. Concretely, we consider an image \mathbf{x} that contains several objects $\mathcal{O} = \{(b_i, c_i) \mid i = 1, \dots, p\}$, where b_i and c_i refer to a box region and the class name for the i -th object. We synthesize an additional input image \mathbf{x}' where objects of the class c , $\mathcal{O}_c = \{(b_i, c_i) \mid c_i = c\}$, are removed from the original image \mathbf{x} . The purpose of removing the region related to class c is to examine if the retrieved sentence for this query image contains the word related to class c . If the retrieved sentence describes the query image \mathbf{x}' well by not mentioning the class c , this model would be robust to spuriousness. The removed regions $\{b_i \mid c_i = c\}$ are filled in by an inpainting model [72] to avoid the data-distribution shift which occurs when these regions are filled in by a constant value [7]. In the following, we describe the details of using inpainting for synthesizing images.

Multiple object classes in the original image. To generate a synthetic image, we select a reference image \mathbf{x} that contains objects from multiple classes, $|\text{set}(\{c_i\}_1^p)| \geq 2$. Our aim is to synthesize data that can expose and fix a model’s bias towards frequently co-occurring objects. By removing objects of a specific class (sometimes it could be multiple classes, which we discuss below), we increase the number of images in which correlated classes do not co-occur. We propose two scenarios for synthesizing synthetic images: removing a single class and removing multiple classes.

Removing a single class. Assume we select a class c^r which we want to remove from the image \mathbf{x} . In this scenario, we want only the objects of class c^r to be removed while others should remain in the synthesized image \mathbf{x}' . Let us define \mathcal{C}^r as all classes in the image \mathbf{x} with the exception of class c^r , i.e. $\mathcal{C}^r = \text{set}(\{c_i\}_1^p) - \{c^r\}$, with the box regions of objects for the class c^r and $c^r \in \mathcal{C}^r$ defined as $\mathcal{B}^r = \{b_i \mid c_i = c^r\}$ and $\mathcal{B}^{\setminus r} = \{b_i \mid c_i = c^{\setminus r}\}$ respectively. To remove only the objects of the class c^r and keep

the objects of all the classes in $\mathcal{C}^{\setminus r}$, the intersection between object regions should be small enough,

$$\frac{\text{area}(\mathcal{B}^r \cap \mathcal{B}^{\setminus r})}{\text{area}(\mathcal{B}^{\setminus r})} < \alpha_1, \quad \text{for all } c^{\setminus r} \in \mathcal{C}^{\setminus r}, \quad (1)$$

where the function $\text{area}(\cdot)$ denotes the area of the given argument. If the LHS of Eq. 1 is big enough, the region \mathcal{B}^r will overlap with a large part of the region $\mathcal{B}^{\setminus r}$. If this is the case, the objects that we want to preserve in the synthetic image would also be removed when removing the region \mathcal{B}^r . For instance, in Figure 2 (left), we select the class ‘person’ as the class to be removed in the reference image \mathbf{x} . Since the ‘person’ region is not covering the ‘horse’ region, the inpainting model restores the removed horse region well, while the people are fully removed. In practice, we select $\alpha_1 = 0.4$.

Removing multiple classes. As before, we select a class c^r which we would like to remove from the image \mathbf{x} . However, this time we want to remove objects of multiple classes, including the class c^r . Let us define \mathcal{C}^{r+} as the class set where the objects of the class $c^{r+} \in \mathcal{C}^{r+}$ largely overlap with the region \mathcal{B}^r ,

$$\frac{\text{area}(\mathcal{B}^r \cap \mathcal{B}^{r+})}{\text{area}(\mathcal{B}^{r+})} > \alpha_2, \quad \text{for all } c^{r+} \in \mathcal{C}^{r+}, \quad (2)$$

where $\mathcal{B}^{r+} = \{b_i \mid c_i = c^{r+}\}$. When we try to remove the region \mathcal{B}^r in the image \mathbf{x} , the object information of class c^{r+} would almost be lost since the region \mathcal{B}^{r+} highly overlaps with the region \mathcal{B}^r . Therefore, instead of removing only the region \mathcal{B}^r , we remove object regions of multiple classes that satisfy Eq. 2, i.e. $\{b_i \mid c_i \in \{c^r\} \cup \mathcal{C}^{r+}\}$. We select $\alpha_2 = 0.8$.

Size of the removed region. If the removed region largely covers the image, the inpainting model would struggle to generate a realistic image. Therefore, we only consider synthetic images that satisfy:

$$\frac{\text{area}(\mathcal{B}')}{\text{area}(\mathbf{x})} < \alpha_3, \quad (3)$$

where \mathcal{B}' is the removed region, i.e. $\mathcal{B}' = \mathcal{B}^r$ when objects of a single class are removed and $\mathcal{B}' = \{b_i \mid c_i \in \{c^r\} \cup \mathcal{C}^{r+}\}$ when objects of multiple classes are removed. We use $\alpha_3 = 0.7$.

3.2. Object decorrelation metric ODmAP@k

To evaluate the bias of cross-modal retrieval models due to spurious object correlations in the training dataset, we propose the object decorrelation metric **ODmAP@k**. It

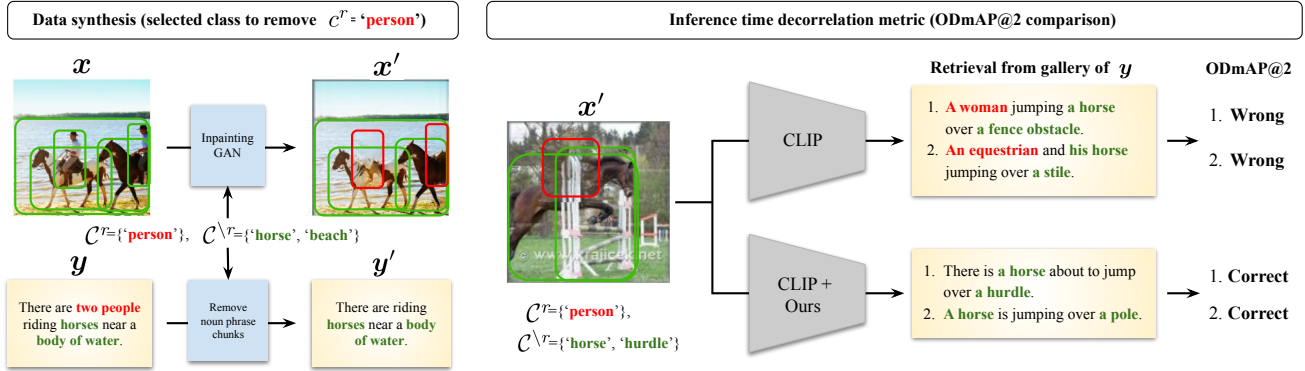


Figure 2. **Pipeline of synthesizing the image/text pair.** The left figure shows the pipeline of how the synthetic data is generated. The right figure shows how the proposed object decorrelation metric works. Existing and nonexisting words/objects in the sentence and image are colored green and red respectively.

measures the model’s ability to capture the semantic correspondence between the query image and retrieved sentences. To measure ODMAP@k, we use synthetic test images x' as queries and retrieve sentences from the gallery of text captions in the dataset, $\mathcal{G}_y = \{y \mid (x, y) \in \mathcal{D}\}$. Our objective is to examine whether the retrieved sentence describes the class c^r that is removed and therefore does not exist in the query image x' .

The best scenario would be that we know the matching sentence $y' \in \mathcal{G}_y$ for the synthetic image x' since this implies that y' would not describe the class c^r that does not exist in the image x' , and we can observe whether the model retrieves the sentence y' . However, it is expensive to manually pair the synthetic image with matching sentences. Instead, we focus on the correspondence between the object class in the image and the noun phrases in the sentence. Assume the caption y is composed of several noun phrases $\mathcal{N} = \{n_i \mid i = 1, \dots, q\}$. We then measure ODMAP@k by (1) checking if the sentence y retrieved by the model does not have any of the noun phrases that is related to the class c^r , i.e. $\text{match}(c^r, n_i) = 0$ for all $n_i \in \mathcal{N}$, and (2) determining whether the retrieved sentence y contains the noun phrases that are related with the class existing in the query image x' , i.e. $\text{match}(c^{\setminus r}, n_i) = 1$ if $c^{\setminus r} \in \mathcal{C}^{\setminus r}$ and $n_i \in \mathcal{N}$, where the function $\text{match}(\cdot, \cdot)$ is 1 when two arguments are related and 0 if not. We assume the retrieved sentence is correct if both conditions (1) and (2) are satisfied, and then measure the accuracy using the mean average precision at k (mAP@k) [52]. For instance, in Figure 2 (right) the query image does not contain the class ‘person’ but contains the classes ‘horse’ and ‘hurdle’. If the retrieved sentence contains a noun phrase that is related to the class ‘person’, e.g. ‘A woman’ or ‘An equestrian’, we regard this retrieved sentence as wrong. To reduce the possibility of no relevant and correct sentences in the gallery that describe the query synthetic image x' well, we form the gallery \mathcal{G} with sentences

not only from the test dataset but also from the training and validation datasets. Also, the design of the matching function $\text{match}(\cdot, \cdot)$ differs on the dataset which is described in Section 4.1.

3.3. Finetuning pipeline

In this section, we introduce a data augmentation method \mathcal{D}' that helps mitigate the spuriousness of the dataset and leads to debiasing the model when \mathcal{D}' is used for finetuning. We use the synthetic image x' generated by the method introduced in Section 3.1 using the training dataset. To form the synthetic dataset $(x', y') \in \mathcal{D}'$, we should make a caption y' that describes the synthetic image x' . One solution would be to manually describe each image, which requires huge annotation costs. Instead, we propose three methods that automatically generate sentence with the pre-existing resources: prompts, a pre-trained captioning model, and noun phrase chunking.

Prompts. Language prompts have recently gained attention for boosting the performance of large language and vision-language models [48, 57, 58]. For instance, CLIP [58] achieved an impressive zero-shot classification performance on ImageNet [15]. To predict the class label in CLIP, a prompt, e.g. “A photo of [classname]”, passes through the CLIP text encoder to form a text embedding and measure the distance with an image embedding. In our case, we use the prompt for describing the synthetic image x' . Assume $\mathcal{C}^{\setminus r}$ is the set of classes for which objects are present in the synthetic image x' . We synthesize the caption y' with the prompt that includes classes in $\mathcal{C}^{\setminus r}$, e.g. $y' = \text{“A photo of person and dog”}$ when $\mathcal{C}^{\setminus r} = \{\text{'person'}, \text{'dog'}\}$. In practice we use 80 prompts and randomly select a single prompt to generate the caption y' .

Pre-trained captioning model. Another way to easily generate a caption for a synthetic image is to use a pre-trained

captioning model. We use the ClipCap [50] image captioning model trained on the MSCOCO dataset [42] to generate the caption y' .

Removing noun phrase chunks. We can synthesize the caption y' by removing relevant noun phrases from the original caption y . We assume that the caption y is composed of several noun phrases $\mathcal{N} = \{n_i | i = 1, \dots, q\}$, and the synthetic image x' is obtained by removing the objects from classes in \mathcal{C}^r from the original image x . Similar to how the ODmAP@k is measured, we select the noun phrases, $\mathcal{N}^r \subset \mathcal{N}$, that are related to the class in \mathcal{C}^r , i.e. $\text{match}(c^{\setminus r}, n_i) = 1$ for $c^{\setminus r} \in \mathcal{C}^{\setminus r}$ and $n_i \in \mathcal{N}^r$, and remove the selected noun phrases \mathcal{N}^r from the caption y to generate y' . For instance, the noun phrases for $y = \text{“Two dogs fighting over a frisbee”}$ are $\mathcal{N} = \{\text{‘Two dogs’}, \text{‘a frisbee’}\}$. Removing the class ‘frisbee’ gives the synthetic caption $y' = \text{“Two dogs fighting over”}$.

Even though this process of synthesizing the sentence is similar to the method used in ODmAP@k in the sense that they both use the noun phrase that matches the class, we argue that it is not designed to explicitly improve the ODmAP@k. In the process of generating y' , the noun phrase is erased from the original sentence y , and therefore y' might be grammatically incorrect. Instead, ODmAP@k is basically retrieving the sentence from the gallery formed with the original sentence y , and the noun phrase is used only to check whether the retrieved sentence describes a specific class when measuring the spuriousness.

Training with the synthetic dataset. With the use of both datasets \mathcal{D} and \mathcal{D}' , the training dataset becomes more balanced in terms of the co-occurrence of visual objects and semantics in the image/text pair. This balanced dataset can be used to mitigate the bias that arises from memorizing object co-occurrence.

4. Experiment

We describe our experimental setup in Section 4.1. We then show the evaluation results on the spuriousness in Section 4.2 and the standard cross-modal retrieval in Section 4.3. Finally, we provide qualitative results and an analysis of our method in Section 4.4.

4.1. Experimental setup

Cross-modal retrieval datasets. We conduct experiments on two datasets: MS-COCO [42] and Flickr30k [71]. **MS-COCO** contains 123,287 images and each image is manually annotated with 5 sentences. Following existing works on this benchmark, we adopted the standard evaluation split to test the general cross-modal retrieval performance, using 113,287 images for training, 5,000 images for validation and 5,000 images for testing. **Flickr30k** contains 31,783 images from the Flickr30k website and each image is an-

notated with 5 sentences. We used the standard evaluation split that contains 29,000 images for training, 1,000 images for validation and 1,000 images for testing.

Compared methods. We adopt three cross-modal retrieval methods for comparison: VSE++ [21], CLIP [58], and BLIP [39]. The methods we compare to are selected to cover different architectures (CNN, RNN, or Transformer-based) and loss functions (triplet loss, contrastive loss, or matching loss). Specifically, we consider the VGG19 [62] and Bi-GRU [10] architectures for the image encoder and the text encoder in VSE++, and ViT-B [18], and BERT [16] in both CLIP and BLIP. Following the original papers, VSE++ is trained with a triplet loss, CLIP with a contrastive loss, and BLIP with a contrastive loss and a matching loss. We refer to the zero-shot application of the CLIP and BLIP models as CLIP (zeroshot) and BLIP (zeroshot), respectively, and to the model that was fine-tuned on the original dataset \mathcal{D} as CLIP (finetune) and BLIP (finetune).

Implementation details. We finetune the compared methods with both the original and synthetic datasets, $\mathcal{D} \cup \mathcal{D}'$, on a single Quadro RTX 6000 GPU. To finetune VSE++ and CLIP, we run 10 epochs with a batch size of 256. We use the Adam [35] optimizer with a learning rate in $1e-4$ for VSE and $2e-6$ for CLIP, and decay the learning rate linearly with a rate of 0.5 every 2 epochs. To finetune the BLIP, we run 2 epochs with a batch size of 16. We use AdamW [43] optimizer with a learning rate of $1e-5$, and decay the learning rate linearly with a rate of 0.5 every half epoch. If not mentioned otherwise, we use the method of removing noun phrase chunks to synthesize the text dataset as default. When synthesizing the text by removing noun phrase chunks, we match the noun phrase with the class name of the to be removed class. The mechanism for matching these depends on the dataset. On the Flickr30k dataset, the (class of object in image, noun phrase in text) pair is annotated. On the MS-COCO dataset, we manually create a list of words that are related to the class, and regard the noun phrase and the class as a matching pair when the noun phrase contains a word that is related to the class. The manual list is provided in the Appendix. The number of classes in MS-COCO and Flickr30k is 80 and 4, respectively.

Synthetic dataset. With the data augmentation pipeline described in Section 3.3, we synthesized 45,467 image/text pairs for the MS-COCO training dataset and 4,650 pairs for the Flickr30k training dataset. These are obtained by considering a single caption among 5 possible captions for each image in the original dataset \mathcal{D} to generate a synthetic image/text pair. Using all 5 (which results in about 5 times more synthetic image/text pairs) did not yield a noticeable performance improvement (discussed in Section 4.4). Therefore, we use 45,467 and 4,650 synthetic pairs for finetuning the models on MS-COCO and Flickr30k respectively.

	MS-COCO (5K test set)			Flickr30k (1K test set)		
	ODmAP@1	ODmAP@5	ODmAP@10	ODmAP@1	ODmAP@5	ODmAP@10
VSE++ [21]	56.0	49.3	45.2	59.7	53.2	49.3
VSE++ [21] + Ours	59.8	52.8	48.3	62.5	54.6	51.5
CLIP [58] (zeroshot)	58.6	51.6	47.1	59.0	52.0	48.4
CLIP [58] (finetune)	59.8	53.2	48.8	60.2	52.2	48.4
CLIP [58] + Ours	70.1	62.3	56.8	61.4	52.3	47.6
BLIP [39] (zeroshot)	60.2	52.3	47.4	62.9	54.9	51.2
BLIP [39] (finetune)	65.3	58.3	53.8	58.3	52.5	50.3
BLIP [39] + Ours	71.6	63.7	58.5	62.5	55.4	52.4

Table 1. **Evaluation of object decorrelation of cross-modal retrieval on the MS-COCO (left) and Flickr30k (right) datasets.** We evaluate three different models that use different architectures (CNN, RNN, or Transformer-based) and loss functions for training the model (triplet loss, contrastive loss, or matching loss).

	MS-COCO (5K test set)						Flickr30k (1K test set)					
	image → text			text → image			image → text			text → image		
	R@1	R@5	R@10	R@1	R@5	R@10	R@1	R@5	R@10	R@1	R@5	R@10
VSE++ [21]	32.8	62.0	74.6	24.1	52.9	66.3	40.9	68.7	78.5	31.8	59.9	70.9
VSE++ [21] + Ours	35.0	64.3	76.2	24.7	53.5	66.6	42.5	69.5	77.7	31.8	60.7	71.8
CLIP [58] (zeroshot)	50.6	75.1	83.6	30.1	55.7	66.8	79.0	94.3	98.2	58.0	82.9	89.9
CLIP [58] (finetune)	65.5	87.4	93.3	48.6	75.7	84.7	84.2	96.1	98.1	68.6	90.4	95.1
CLIP [58] + Ours	65.6	87.2	93.1	48.4	75.7	84.4	85.0	96.5	99.0	69.9	90.9	95.4
BLIP [39] (zeroshot)	71.2	90.1	94.6	55.0	79.3	86.9	85.5	97.9	99.0	77.7	94.2	96.9
BLIP [39] (finetune)	78.0	93.8	97.0	61.0	84.2	90.8	96.1	99.8	99.9	85.8	97.2	98.8
BLIP [39] + Ours	78.7	94.5	97.6	61.3	84.4	90.8	96.2	99.6	100.0	86.2	97.5	98.8

Table 2. **Evaluation of standard cross-modal retrieval on the MS-COCO (left) and Flickr30k (right) datasets.** We evaluate three different models that use different architectures (CNN, RNN, or Transformer-based) and loss functions for training the model (triplet loss, contrastive loss, or matching loss).

4.2. Evaluating object decorrelation

We evaluate the object decorrelation of our method and four compared methods using the ODmAP@K metric on the MS-COCO and Flickr30k datasets in Table 1. We observe that our method outperforms the other frameworks that we compare to. For instance, CLIP (finetune) gives 59.8% and 60.2% scores in ODmAP@1 on the MS-COCO and Flickr30k datasets, respectively, while our method yields 70.1% and 61.4%, outperforming CLIP (finetune) by a margin of 10.3% and 1.2%. Similar trends are observed across different datasets and compared methods except for BLIP on Flickr30k where BLIP (zeroshot) is 0.4% better than ours. These results imply that our method has the ability to get better retrieval results than the compared methods based on the correct object cue present in the image. Also, we observe that the baseline model tends to have a better ability to debias the spurious correlation when its standard retrieval ability is better. For instance on the

MS-COCO dataset, BLIP (finetune) has the best score on ODmAP@1 at 65.3%, followed by CLIP (finetune) with 59.8% and VSE++ at 56%. Finally, we observe that for the large-scale models the fine-tuned model gives a better score than the zero-shot model. The ODmAP@1 score improves from 60.2% to 65.3% for BLIP and from 58.6% to 59.8% for CLIP when the zero-shot model is fine-tuned only on the original dataset.

4.3. Evaluating standard cross-modal retrieval

Standard metric. We evaluate the cross-modal retrieval performance using the recall at K (**R@K**) which measures the fraction of queries for which at least one correct sample is in the top K retrieved items.

Quantitative results for the standard retrieval evaluation. We evaluate the cross-modal retrieval results using the standard retrieval metric. The results are reported in Table 2. Our method shows competitive performance across

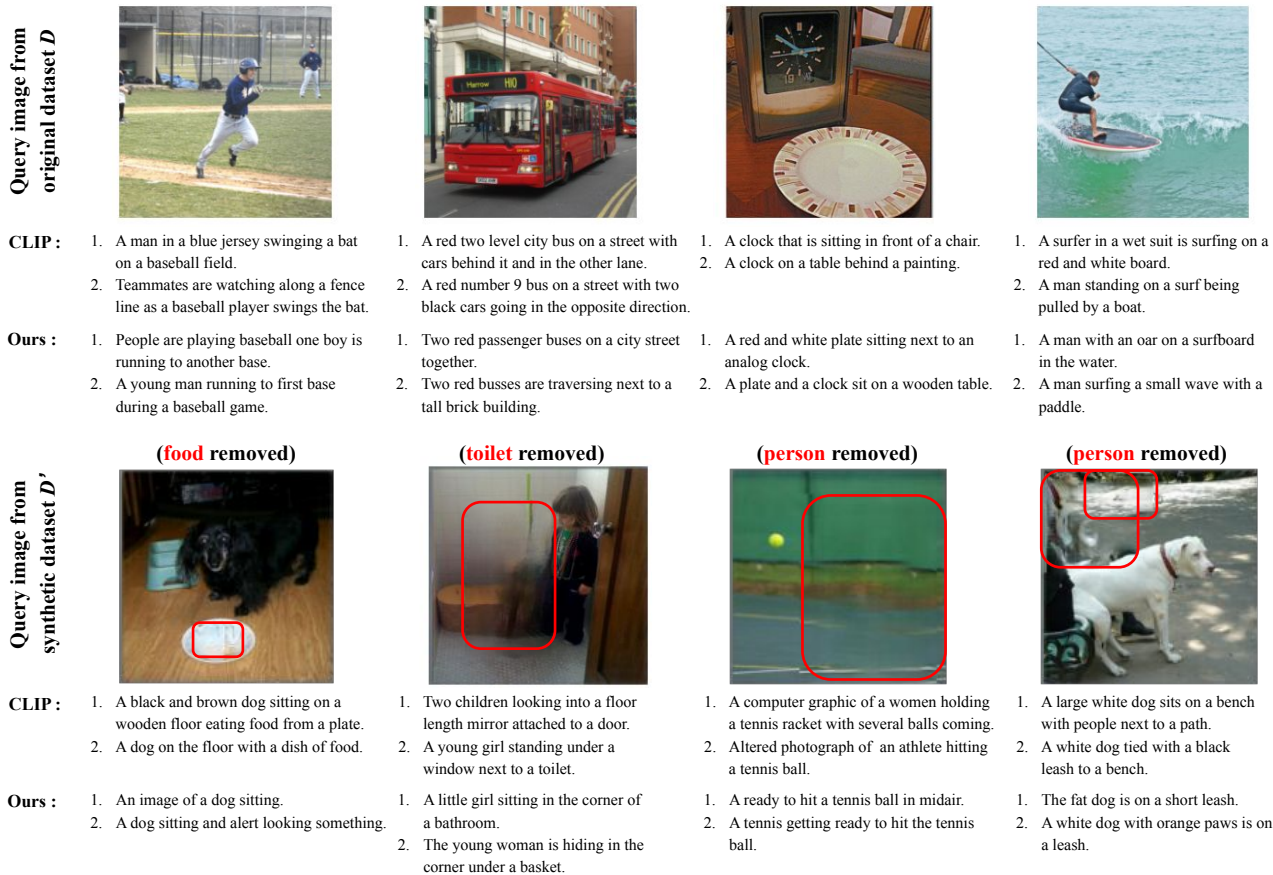


Figure 3. **Qualitative results of image-to-text retrieval.** The first (resp. second) row shows the image-to-text retrieval results when the query image is from the original dataset (resp. synthetic dataset). For the synthetic dataset, we indicate which class has been removed from the original dataset in red.

different datasets and baseline models. For instance, CLIP (finetune) yields 65.5% and 48.6% for image-to-text and text-to-image R@1 respectively on the MS-COCO dataset, while our method gives 65.6% and 48.4%. In both cases, the difference is less than 0.5%. Similar trends are observed across different datasets and compared methods, where our method shows competitive or marginally better results. Overall, these results suggest that our method addressed the co-occurrence bias in the model without harming the overall retrieval performance.

4.4. Analysis

Qualitative result. Figure 3 shows a qualitative comparison between CLIP [58] and CLIP fine-tuned on our augmented dataset. The first row shows the top-2 retrieved sentences given the original image as a query. It can be observed that the retrieved sentence describes an object that is not present in the query image. For instance, the first retrieved sentence by CLIP for the first image as a query

contains the word ‘bat’ which is not visible in the image but is related to other objects like the man with a blue jersey or a baseball field. This can be observed more clearly in the second row of Figure 3 where the synthetic image is given as a query. While objects of the selected class (or classes) are removed from the original image to generate the synthetic image, we observe that CLIP still retrieves the sentence that describes the object that is removed. For instance, the top-1 retrieved sentence by CLIP for the first image as a query contains the word ‘food’ which is the class that has been removed from the original image. Our method retrieves the sentence that describes the visible objects without mentioning other objects that are not present in the query image.

Impact of different methods for synthetically generating data. Here, we explore the impact of different methods for synthetically generating text and images in D' . For the image synthesis, we consider three different methods that are commonly used for removing the information from the image: zero padding, mean padding, and blur padding. Zero

padding refers to the method where the removed region \mathcal{B}^r in the original image is filled in with zeroes. Mean padding and blur padding refer to the compared methods when \mathcal{B}^r is filled in with the average pixel values of \mathcal{B}^r and the Gaussian blur of \mathcal{B}^r , respectively. For text synthesis, we consider two additional methods: prompts and the captioning model to generate the text (as described in Section 3.3).

The results for different data synthesis methods are shown in Table 3. We observe that except for using prompts to synthesize text and inpainting for image synthesis, using the synthetic dataset generated by noun phrase removal and inpainting method appear to be suited best for debiasing the model. When the image is generated using blur-, zero-, or average padding, the synthesized image would be considered out-of-distribution which would result in the model learning other biases. Furthermore, using prompts for text synthesis shows the best result on ODMAP@1 but drastically decreases on ODMAP@5 and ODMAP@10. We discovered that this model suffers from the hubness issue, i.e. the retrieved sentence tends to be the same for many different queries. We conjecture that this happens because the prompts use a similar text format, and the model learns to match this specific text format to the synthetic images. These results suggest that the careful design of data augmentation is needed to debias the model.

Impact of varying the ratio of original and synthetic data. To determine the best experimental configuration for improving the model’s performance with the synthetic dataset \mathcal{D}' , we experiment with different ratios of synthetic data included for training of the CLIP model in the COCO dataset. Concretely, we first generate a synthetic dataset using all 5 original captions per image, which is then $|\mathcal{D}'|/|\mathcal{D}| \approx 0.4$. We use the full original dataset \mathcal{D} , but vary the subset from the synthetic dataset $\mathcal{D}'_{sub} \subset \mathcal{D}'$ and use both \mathcal{D} and \mathcal{D}'_{sub} for training the model. We show the R@1 and ODMAP@1 as we change the amount of synthetic data used for training in Figure 4. We observe that the ODMAP@1 increases drastically until $|\mathcal{D}'_{sub}|/|\mathcal{D}| \approx 0.1$, and then slowly increases and saturates after the ratio is larger than 0.2. Given that using one out of five original captions per image to generate a synthetic dataset would have $|\mathcal{D}'_{sub}|/|\mathcal{D}| \approx 0.08$ which is near 0.1, we conclude that using synthetic data for a single caption per image is sufficient. We also observe that the R@1 score does not vary much as we change the ratio of the datasets.

5. Conclusion

In this paper, we studied the learnt bias in the image-text retrieval models that arises from spurious correlations in the training data. We discovered that existing methods tend to retrieve samples based on cues that might not be semantically related. To address this issue, we trained the model with additional synthetic data which eliminates frequent

Method	Text Synth.	Image Synth.	Object decorrelation evaluation		
			ODmAP@1	ODmAP@5	ODmAP@10
\mathcal{D}	-	-	59.8	53.2	48.8
$\mathcal{D} + \mathcal{D}'$	NP removal	blur pad.	67.9	60.9	55.8
	NP removal	zero pad.	68.9	61.0	55.7
	NP removal	avg pad.	69.1	62.0	57.1
	capt. model	inpaint	61.9	55.6	51.6
	prompts	inpaint	73.8	56.2	49.5
	NP removal	inpaint	70.1	63.7	58.5

Table 3. Comparing the impact of different methods for synthetically generating text and images in \mathcal{D}' from the MS-COCO dataset on CLIP (finetuned) in terms of its spuriousness.




Figure 4. Impact of varying the ratio of original and synthetic samples used for training the model.

object co-occurrences in the original training data. Additionally, we proposed a new object decorrelation metric, ODMAP@k, which measures how well the model retrieves samples based on the right cue. Applying our method shows significant improvements on ODMAP@k for a variety of image-text retrieval models, without harming the standard retrieval performance.

Limitation and future work. Our method focuses on data augmentation for the image-text retrieval task. It alleviates the bias caused by spurious correlations between objects in the training data, but does not analyze and solve other biases, e.g. texture bias [26]. Finding and analyzing other biases in the retrieval model would be an interesting topic. Another promising future direction would be to further explore the spuriousness of other cross-modal retrieval tasks such as text-video retrieval or text-audio retrieval.

Acknowledgements

This work was supported by DFG project number 276693517, by BMBF FKZ: 01IS18039A, by the ERC (853489 - DEXIM), and by EXC number 2064/1 – project number 390727645. Jae Myung Kim thanks the European Laboratory for Learning and Intelligent Systems (ELLIS) PhD program and the International Max Planck Research School for Intelligent Systems (IMPRS-IS) for support.

References

- [1] Vedika Agarwal, Rakshith Shetty, and Mario Fritz. Towards causal vqa: Revealing and reducing spurious correlations by invariant and covariant semantic editing. In *CVPR*, 2020. 2
- [2] 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. 2
- [3] Yusuf Aytar, Mubarak Shah, and Jiebo Luo. Utilizing semantic word similarity measures for video retrieval. In *CVPR*, 2008. 2
- [4] Max Bain, Arsha Nagrani, Gül Varol, and Andrew Zisserman. Frozen in time: A joint video and image encoder for end-to-end retrieval. In *ICCV*, 2021. 2
- [5] Hugo Berg, Siobhan Mackenzie Hall, Yash Bhalgat, Wonsuk Yang, Hannah Rose Kirk, Aleksandar Shtedritski, and Max Bain. A prompt array keeps the bias away: Debiasing vision-language models with adversarial learning. *arXiv preprint arXiv:2203.11933*, 2022. 2
- [6] Simion-Vlad Bogolin, Ioana Croitoru, Hailin Jin, Yang Liu, and Samuel Albanie. Cross modal retrieval with querybank normalisation. In *CVPR*, 2022. 2
- [7] Chun-Hao Chang, Elliot Creager, Anna Goldenberg, and David Duvenaud. Explaining image classifiers by counterfactual generation. In *ICLR*, 2018. 3
- [8] Long Chen, Xin Yan, Jun Xiao, Hanwang Zhang, Shiliang Pu, and Yueting Zhuang. Counterfactual samples synthesizing for robust visual question answering. In *CVPR*, 2020. 2
- [9] Tianlang Chen, Jiajun Deng, and Jiebo Luo. Adaptive offline quintuplet loss for image-text matching. In *ECCV*, 2020. 2
- [10] Kyunghyun Cho, Bart Van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. On the properties of neural machine translation: Encoder-decoder approaches. *arXiv preprint arXiv:1409.1259*, 2014. 5
- [11] Sanghyuk Chun, Wonjae Kim, Song Park, Minsuk Chang, and Seong Joon Oh. Eccv caption: Correcting false negatives by collecting machine-and-human-verified image-caption associations for ms-coco. In *ECCV*, 2022. 2
- [12] Sanghyuk Chun, Seong Joon Oh, Rafael Sampaio De Rezende, Yannis Kalantidis, and Diane Larlus. Probabilistic embeddings for cross-modal retrieval. In *ICCV*, 2021. 2
- [13] Ioana Croitoru, Simion-Vlad Bogolin, Marius Leordeanu, Hailin Jin, Andrew Zisserman, Samuel Albanie, and Yang Liu. Teachtext: Crossmodal generalized distillation for text-video retrieval. In *ICCV*, 2021. 2
- [14] Katherine Crowson, Stella Biderman, Daniel Kornis, Dashiell Stander, Eric Hallahan, Louis Castriaco, and Edward Raff. Vqgan-clip: Open domain image generation and editing with natural language guidance. In *ECCV*, 2022. 2
- [15] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009. 4
- [16] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. 2018. 5
- [17] Jianfeng Dong, Xirong Li, and Cees GM Snoek. Predicting visual features from text for image and video caption retrieval. *IEEE Transactions on Multimedia*, 2018. 2
- [18] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *iclr*, 2020. 5
- [19] Aviv Eisenschat and Lior Wolf. Linking image and text with 2-way nets. In *CVPR*, 2017. 2
- [20] Martin Engilberge, Louis Chevallier, Patrick Pérez, and Matthieu Cord. Finding beans in burgers: Deep semantic-visual embedding with localization. In *CVPR*, 2018. 2
- [21] Fartash Faghri, David J Fleet, Jamie Ryan Kiros, and Sanja Fidler. Vse++: Improving visual-semantic embeddings with hard negatives. In *BMVC*, 2018. 2, 5, 6
- [22] Andrea Frome, Greg S Corrado, Jon Shlens, Samy Bengio, Jeff Dean, Marc’Aurelio Ranzato, and Tomas Mikolov. Devise: A deep visual-semantic embedding model. In *NeurIPS*, 2013. 2
- [23] Valentin Gabeur, Arsha Nagrani, Chen Sun, Karteek Alahari, and Cordelia Schmid. Masking modalities for cross-modal video retrieval. 2022. 2
- [24] Valentin Gabeur, Chen Sun, Karteek Alahari, and Cordelia Schmid. Multi-modal transformer for video retrieval. In *ECCV*, 2020. 2
- [25] Yuting Gao, Jinfeng Liu, Zihan Xu, Jun Zhang, Ke Li, and Chunhua Shen. Pyramidclip: Hierarchical feature alignment for vision-language model pretraining. *arXiv preprint arXiv:2204.14095*, 2022. 2
- [26] 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 *ICLR*, 2019. 8
- [27] Yunchao Gong, Liwei Wang, Micah Hodosh, Julia Hockenmaier, and Svetlana Lazebnik. Improving image-sentence embeddings using large weakly annotated photo collections. In *ECCV*, 2014. 2
- [28] Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In *CVPR*, 2017. 2
- [29] David Harwath, Adria Recasens, Dídac Surís, Galen Chuang, Antonio Torralba, and James Glass. Jointly discovering visual objects and spoken words from raw sensory input. In *ECCV*, 2018. 2
- [30] Lisa Anne Hendricks, Kaylee Burns, Kate Saenko, Trevor Darrell, and Anna Rohrbach. Women also snowboard: Overcoming bias in captioning models. In *ECCV*, 2018. 1, 2
- [31] Sungeun Hong, Woobin Im, and Hyun S Yang. Deep learning for content-based, cross-modal retrieval of videos and music. *arXiv preprint arXiv:1704.06761*, 2017. 2

- [32] 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. 2021. [2](#)
- [33] Andrej Karpathy and Li Fei-Fei. Deep visual-semantic alignments for generating image descriptions. In *CVPR*, 2015. [2](#)
- [34] Andrej Karpathy, Armand Joulin, and Li F Fei-Fei. Deep fragment embeddings for bidirectional image sentence mapping. 2014. [2](#)
- [35] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. [5](#)
- [36] Ryan Kiros, Ruslan Salakhutdinov, and Richard S Zemel. Unifying visual-semantic embeddings with multimodal neural language models. *arXiv preprint arXiv:1411.2539*, 2014. [2](#)
- [37] A Sophia Koepke, Andreea-Maria Oncescu, Joao Henriques, Zeynep Akata, and Samuel Albanie. Audio retrieval with natural language queries: A benchmark study. *IEEE Transactions on Multimedia*, 2022. [2](#)
- [38] Kuang-Huei Lee, Xi Chen, Gang Hua, Houdong Hu, and Xiaodong He. Stacked cross attention for image-text matching. In *ECCV*, 2018. [2](#)
- [39] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *ICML*, 2022. [2](#), [5](#), [6](#)
- [40] Kunpeng Li, Yulun Zhang, Kai Li, Yuanyuan Li, and Yun Fu. Visual semantic reasoning for image-text matching. In *ICCV*, 2019. [2](#)
- [41] Xiujun Li, Xi Yin, Chunyuan Li, Pengchuan Zhang, Xiaowei Hu, Lei Zhang, Lijuan Wang, Houdong Hu, Li Dong, Furu Wei, et al. Oscar: Object-semantic aligned pre-training for vision-language tasks. In *ECCV*, 2020. [2](#)
- [42] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014. [2](#), [5](#)
- [43] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *ICLR*, 2018. [5](#)
- [44] Siyu Lou, Xuenan Xu, Mengyue Wu, and Kai Yu. Audio-text retrieval in context. In *ICASSP*, 2022. [2](#)
- [45] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. *NeurIPS*, 2019. [2](#)
- [46] Huaishao Luo, Lei Ji, Ming Zhong, Yang Chen, Wen Lei, Nan Duan, and Tianrui Li. Clip4clip: An empirical study of clip for end to end video clip retrieval and captioning. *Neurocomputing*, 2022. [2](#)
- [47] Haoyu Ma, Handong Zhao, Zhe Lin, Ajinkya Kale, Zhangyang Wang, Tong Yu, Jiuxiang Gu, Sunav Choudhary, and Xiaohui Xie. Ei-clip: Entity-aware interventional contrastive learning for e-commerce cross-modal retrieval. In *CVPR*, 2022. [2](#)
- [48] Sachit Menon and Carl Vondrick. Visual classification via description from large language models. In *ICLR*, 2023. [4](#)
- [49] Niluthpol Chowdhury Mithun, Juncheng Li, Florian Metze, and Amit K Roy-Chowdhury. Learning joint embedding with multimodal cues for cross-modal video-text retrieval. In *ACM International Conference on Multimedia Retrieval*, 2018. [2](#)
- [50] Ron Mokady, Amir Hertz, and Amit H Bermano. Clip-cap: Clip prefix for image captioning. *arXiv preprint arXiv:2111.09734*, 2021. [5](#)
- [51] Norman Mu, Alexander Kirillov, David Wagner, and Saining Xie. Slip: Self-supervision meets language-image pre-training. In *ECCV*, 2022. [2](#)
- [52] Kevin Musgrave, Serge Belongie, and Ser-Nam Lim. A metric learning reality check. In *ECCV*, 2020. [4](#)
- [53] Arsha Nagrani, Samuel Albanie, and Andrew Zisserman. Learnable pins: Cross-modal embeddings for person identity. In *ECCV*, 2018. [2](#)
- [54] Arsha Nagrani, Paul Hongsuck Seo, Bryan Seybold, Anja Hauth, Santiago Manen, Chen Sun, and Cordelia Schmid. Learning audio-video modalities from image captions. *arXiv preprint arXiv:2204.00679*, 2022. [2](#)
- [55] Yulei Niu, Kaihua Tang, Hanwang Zhang, Zhiwu Lu, Xian-Sheng Hua, and Ji-Rong Wen. Counterfactual vqa: A cause-effect look at language bias. In *CVPR*, 2021. [2](#)
- [56] Andreea-Maria Oncescu, A Koepke, Joao F Henriques, Zeynep Akata, and Samuel Albanie. Audio retrieval with natural language queries. In *INTERSPEECH*, 2021. [2](#)
- [57] Sarah Pratt, Rosanne Liu, and Ali Farhadi. What does a platypus look like? generating customized prompts for zero-shot image classification. *arXiv preprint arXiv:2209.03320*, 2022. [4](#)
- [58] 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. 2021. [1](#), [2](#), [4](#), [5](#), [6](#), [7](#)
- [59] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019. [2](#)
- [60] Anna Rohrbach, Lisa Anne Hendricks, Kaylee Burns, Trevor Darrell, and Kate Saenko. Object hallucination in image captioning. In *EMNLP*, 2018. [2](#)
- [61] Gyungin Shin, Weidi Xie, and Samuel Albanie. Reco: Retrieve and co-segment for zero-shot transfer. In *NeurIPS*, 2022. [2](#)
- [62] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. [5](#)
- [63] Yale Song and Mohammad Soleymani. Polysemous visual-semantic embedding for cross-modal retrieval. In *CVPR*, 2019. [2](#)
- [64] Christopher Thomas and Adriana Kovashka. Preserving semantic neighborhoods for robust cross-modal retrieval. In *ECCV*, 2020. [2](#)
- [65] Jialu Wang, Yang Liu, and Xin Wang. Are gender-neutral queries really gender-neutral? mitigating gender bias in image search. In *EMNLP*, 2021. [2](#)

- [66] Liwei Wang, Yin Li, and Svetlana Lazebnik. Learning deep structure-preserving image-text embeddings. In *CVPR*, 2016. [2](#)
- [67] Zihao Wang, Xihui Liu, Hongsheng Li, Lu Sheng, Junjie Yan, Xiaogang Wang, and Jing Shao. Camp: Cross-modal adaptive message passing for text-image retrieval. In *ICCV*, 2019. [2](#)
- [68] Michael Wray, Diane Larlus, Gabriela Csurka, and Dima Damen. Fine-grained action retrieval through multiple parts-of-speech embeddings. In *ICCV*, 2019. [2](#)
- [69] Ran Xu, Caiming Xiong, Wei Chen, and Jason Corso. Jointly modeling deep video and compositional text to bridge vision and language in a unified framework. In *AAAI*, 2015. [2](#)
- [70] Lewei Yao, Runhui Huang, Lu Hou, Guansong Lu, Minzhe Niu, Hang Xu, Xiaodan Liang, Zhenguo Li, Xin Jiang, and Chunjing Xu. Filip: Fine-grained interactive language-image pre-training. *arXiv preprint arXiv:2111.07783*, 2021. [2](#)
- [71] 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*, 2014. [2](#), [5](#)
- [72] Jiahui Yu, Zhe Lin, Jimei Yang, Xiaohui Shen, Xin Lu, and Thomas S Huang. Generative image inpainting with contextual attention. In *CVPR*, 2018. [2](#), [3](#)
- [73] Qi Zhang, Zhen Lei, Zhaoxiang Zhang, and Stan Z Li. Context-aware attention network for image-text retrieval. In *CVPR*, 2020. [2](#)
- [74] Shengyu Zhang, Tan Jiang, Tan Wang, Kun Kuang, Zhou Zhao, Jianke Zhu, Jin Yu, Hongxia Yang, and Fei Wu. De-
vlbert: Learning deconfounded visio-linguistic representations. In *ACM MM*, 2020. [2](#)
- [75] Dora Zhao, Angelina Wang, and Olga Russakovsky. Understanding and evaluating racial biases in image captioning. In *ICCV*, 2021. [2](#)