







Scene-Centric vs. Object-Centric Image-Text Cross-Modal Retrieval: A Reproducibility Study

Mariya Hendriksen¹, Svitlana Vakulenko², Ernst Kuiper³,
and Maarten de Rijke⁴

¹ AIRLab, University of Amsterdam, Amsterdam, The Netherlands

m.hendriksen@uva.nl

² Amazon, Madrid, Spain

svvakul@amazon.com

³ Bol.com, Utrecht, The Netherlands

ekuiper@bol.com

⁴ University of Amsterdam, Amsterdam, The Netherlands

m.derijke@uva.nl

Abstract. Most approaches to (CMR) focus either on object-centric datasets, meaning that each document depicts or describes a single object, or on scene-centric datasets, meaning that each image depicts or describes a complex scene that involves multiple objects and relations between them. We posit that a robust CMR model should generalize well across both dataset types. Despite recent advances in CMR, the reproducibility of the results and their generalizability across different dataset types has not been studied before. We address this gap and focus on the reproducibility of the state-of-the-art CMR results when evaluated on object-centric and scene-centric datasets. We select two state-of-the-art CMR models with different architectures: (i) CLIP; and (ii) X-VLM. Additionally, we select two scene-centric datasets, and three object-centric datasets, and determine the relative performance of the selected models on these datasets. We focus on reproducibility, replicability, and generalizability of the outcomes of previously published CMR experiments. We discover that the experiments are not fully reproducible and replicable. Besides, the relative performance results partially generalize across object-centric and scene-centric datasets. On top of that, the scores obtained on object-centric datasets are much lower than the scores obtained on scene-centric datasets. For reproducibility and transparency we make our source code and the trained models publicly available.

1 Introduction

Cross-modal retrieval (CMR) is the task of finding relevant items across different modalities. For example, given an image, find a text or vice versa. The main challenge in CMR is known as *the heterogeneity gap* ([5, 22]). Since items from

Research conducted while the author was at the University of Amsterdam.

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2023
J. Kamps et al. (Eds.): ECIR 2023, LNCS 13982, pp. 68–85, 2023.

https://doi.org/10.1007/978-3-031-28241-6_5

different modalities have different data types, the similarity between them cannot be measured directly. Therefore, the majority of CMR methods published to date attempt to bridge this gap by learning a latent representation space, where the similarity between items from different modalities can be measured [57].

In this work, we specifically focus on *image-text* CMR, which uses textual and visual data. The retrieval task is performed on *image-text pairs*. In each image-text pair, the text (often referred to as *caption*) describes the corresponding image it is aligned with. For image-text CMR we use either an image or a text as a query [57]. Hence, the CMR task that we address in this paper consists of two subtasks: (i) *text-to-image retrieval*: given a text that describes an image, retrieve all the images that match this description; and (ii) *image-to-text retrieval*: given an image, retrieve all texts that can be used to describe this image.

Scene-centric vs. Object-centric vs. Datasets. Existing image datasets can be grouped into *scene-centric* and *object-centric* datasets [48, 62]. The two types of dataset are typically used for different tasks, viz. the tasks of scene and object understanding, respectively. They differ in important ways that are of interest to us when evaluating performance and generalization abilities of CMR models.

Scene-centric images depict complex scenes that typically feature multiple objects and relations between them. These datasets contain image-text pairs, where, in each pair, an image depicts a complex scene of objects and the corresponding text describes the whole scene, often focusing on *relations and activities*.

Images in object-centric image datasets are usually focused on a single object of interest that they primarily depict. This object is often positioned close to the center of an image with other objects, optionally, in the background. Object-centric datasets contain image-text pairs, where, in each pair, an image depicts an object of interest and the corresponding text describes the depicted *object and its (fine-grained) attributes*.

To illustrate the differences between the two dataset types in CMR, we consider the examples provided in Fig. 1 with an object-centric image-caption (left) and a scene-centric image-caption (right). Note how the pairs differ considerably in terms of the visual style and the content of the caption. The pair on the left focuses on a single object (“pants”) and describes its fine-grained visual attributes (“multicolor,” “boho,” “batic”). The pair on the right captures a scene describing multiple objects (“seagulls,” “pier,” “people”) and relations between them (“sitting,” “watching”).

Research Goals. We focus on (traditional) CMR methods that extract features from each modality and learn a common representation space. Recent years have seen extensive experimentation with such CMR methods, mostly organized into two groups: (i) contrastive experiments on object-centric datasets [17], and (ii) contrastive experiments on scene-centric datasets [35]. In this paper, we consider representative state-of-the-art CMR methods from both groups. We select two pre-trained models which demonstrate state-of-the-art performance



Multicolor boho batic pants



Seagulls sitting on the ledge of a pier with people watching

Fig. 1. An object-centric (left) and a scene-centric (right) image-text pair. Sources: Fashion200k (left); MS COCO (right).

on CMR task and evaluate them in a zero-shot setting. In line with designs used in prior reproducibility work on CMR [3] we select two models for the study. Following the ACM terminology [1], we focus on *reproducibility* (different team, same experimental setup) and *replicability* (different team, different experimental setup) of previously reported results. And following Voorhees [55], we focus on relative (a.k.a. comparative) performance results. In addition, for the reproducibility experiment, we consider the absolute difference between the reported scores and the reproduced scores.

We address the following research questions: (RQ1) Are published relative performance results on CMR reproducible? This question matters because it allows us to confirm the validity of reported results. We show that the relative performance results are not fully reproducible. Specifically, the results are reproducible for one dataset, but not for the other dataset).

We then shift to replicability and examine whether lessons learned on scene-centric datasets transfer to object-centric datasets: (RQ2) To what extent are the published relative performance results replicable? That is, we investigate the validity of the reported results when evaluated in a different setup. We find that relative performance results are partially replicable, using other datasets.

After investigating the reproducibility and replicability of the results, we consider the generalizability of the results. We contrastively evaluate the results on object-centric and scene-centric datasets: (RQ3) Do relative performance results for state-of-the-art CMR methods generalize from scene-centric datasets to object-centric datasets? We discover that the relative performance results only partially generalize across the two dataset types.

Main Contributions. Our main contributions are: (i) We are one of the first to consider reproducibility in the context of CMR and reproduce scene-centric CMR experiments from two papers [44,61] and find that the results are only partially

reproducible. (ii) We perform a replicability study and examine whether relative performance differences reported for CMR methods generalize from scene-centric datasets to object-centric datasets. (iii) We investigate the generalizability of obtained results and analyze the effectiveness of pre-training on scene-centric datasets for improving the performance of CMR on object-centric datasets, and vice versa. And, finally, (iv) to facilitate the reproducibility of our work, we provide the code and the pre-trained models used in our experiments.¹

2 Related Work

Cross-modal Retrieval. CMR methods attempt to construct a multimodal representation space, where the similarity of concepts from different modalities can be measured. Some of the earliest approaches in CMR utilised canonical correlation analysis [15, 26]. They were followed by a dual encoder architecture equipped with a recurrent and a convolutional component, a hinge loss [12, 58] and hard-negative mining [11]. Later on, several attention-based architectures were introduced such as architectures with dual attention [39], stacked cross-attention [31], bidirectional focal attention [36].

Another line of work proposed to use transformer encoders [54] for CMR task [38], and adapted the BERT model [8] as a backbone [13, 67]. Some other researchers worked on improving CMR via modality-specific graphs [56], or image and text generation modules [16].

There is also more domain-specific work that focused on CMR in fashion [14, 28–30], e-commerce [19, 20], cultural heritage [49] and cooking [56].

In contrast to the majority of prior work on the topic, we focus on the reproducibility, replicability, and generalizability of CMR methods. In particular, we explore the state-of-the-art models designed for the CMR task by examining their performance on scene-centric and object-centric datasets.

Scene-centric and Object-centric Datasets. The majority of prior work related to object-centric and scene-centric datasets focuses on computer vision tasks such as object recognition, object classification, and scene recognition. Herranz et al. [21] investigated biases in a CNN when trained on scene-centric versus object-centric datasets and evaluated on the task of object classification.

In the context of object detection, prior work focused on combining feature representations learned from object-centric and scene-centric datasets to improve the performance when detecting small objects [48], and using object-centric images to improve the detection of objects that do not appear frequently in complex scenes [62]. Finally, for the task of scene recognition, Zhou et al. [66] explored the quality of feature representations learned from both scene-centric and object-centric datasets and applied to the task of scene recognition.

Unlike prior work on the topic, in this paper, we focus on both scene-centric and object-centric datasets for evaluation on CMR task. In particular, we explore

¹ <https://github.com/mariyahendriksen/ecir23-object-centric-vs-scene-centric-CMR>.

how state-of-the-art(SOTA) CMR models perform on object-centric and scene-centric datasets.

Reproducibility in Cross-modal Retrieval. To the best of our knowledge, despite the popularity of the CMR task, there are very few papers that focus on reproducibility of research in CMR. Some rare (recent) examples include [3], where the authors survey metric learning losses used in computer vision and explore their applicability for CMR. Rao et al. [45] analyze contributing factors that affect the performance of the state-of-the-art CMR models. However, all prior work focuses on exploring model performance only on two popular scene-centric datasets: Microsoft COCO (MS COCO) and Flickr30k.

In contrast, in this work, we take advantage of the diversity of the CMR datasets and specifically focus on examining how the state-of-the-art CMR models perform across different dataset types: scene-centric and object-centric datasets.

3 Task Definition

We follow the same notation as in previous work [4, 53, 65]. An image-caption cross-modal dataset consists of a set of images \mathcal{I} and texts \mathcal{T} where the images and texts are aligned as image-text pairs: $\mathcal{D} = \{(\mathbf{x}_{\mathcal{I}}^1, \mathbf{x}_{\mathcal{T}}^1), \dots, (\mathbf{x}_{\mathcal{I}}^n, \mathbf{x}_{\mathcal{T}}^n)\}$.

The *cross-modal retrieval* (CMR) task is defined analogous to the standard information retrieval task: given a query \mathbf{q} and a set of m candidates $\Omega_{\mathbf{q}} = \{\mathbf{x}^1, \dots, \mathbf{x}^m\}$ we aim to rank all the candidates w.r.t. their relevance to the query \mathbf{q} . In CMR, the query can be either a text $\mathbf{q}_{\mathcal{T}}$ or an image $\mathbf{q}_{\mathcal{I}}$: $\mathbf{q} \in \{\mathbf{q}_{\mathcal{T}}, \mathbf{q}_{\mathcal{I}}\}$. Similarly, the set of candidate items can be either visual $\mathcal{I}_{\mathbf{q}} \subset \mathcal{I}$, or textual $\mathcal{T}_{\mathbf{q}} \subset \mathcal{T}$ data: $\Omega \in \{\mathcal{I}_{\mathbf{q}}, \mathcal{T}_{\mathbf{q}}\}$.

The CMR task is performed across modalities, therefore, if the query is a text then the set of candidates are images, and vice versa. Hence, the task comprises effectively two subtasks: (i) *text-to-image retrieval*: given a textual query $\mathbf{q}_{\mathcal{T}}$ and a set of candidate images $\Omega \subset \mathcal{I}$, we aim to rank all instances in the set of candidate items Ω w.r.t. their relevance to the query $\mathbf{q}_{\mathcal{T}}$; (ii) *image-to-text retrieval*: given an image as a query $\mathbf{q}_{\mathcal{I}}$ and a set of candidate texts $\Omega \subset \mathcal{T}$, we aim to rank all instances in the set of candidate items Ω w.r.t. their relevance to the query $\mathbf{q}_{\mathcal{I}}$.

4 Methods

In this section, we give an overview of the models included in the study, of the models which were excluded, and provide justification for it. All the approaches we focus on belong to the traditional CMR framework and comprise two stages. First, we extract textual and visual features. The features are typically extracted with a textual encoder and a visual encoder. Next, we learn a latent representation space where the similarity of items from different modalities can be measured directly.

4.1 Methods Included for Comparison

We focus on CMR in *zero-shot setting*, hence, we only consider pre-trained models. Therefore, we focus on the models that are released for public use. Besides, as explained in Sect. 1, we follow prior reproducibility work to inform our experimental choices regarding the number of models. Given the above-mentioned requirements, we selected two methods that demonstrate state-of-the-art performance on the CMR task: CLIP and X-VLM.

Contrastive Language-Image Pretraining (CLIP) [44]. This model is a dual encoder that comprises an image encoder, and a text encoder. The model was pre-trained in a contrastive manner using a symmetric loss function. It is trained on 400 million image-caption pairs scraped from the internet. The text encoder is a transformer [54] with modification from [43]. For the image encoder, the authors present two architectures. The first one is based on ResNet [18] and it is represented in five variants in total. The first two options are ResNet-50, ResNet-101; the last three options are variants of ResNet scaled up in the style of EfficientNet [51]. The second image encoder architecture is a Vision Transformer (ViT) [9]. It is presented in three variants: ViT-B/32, a ViT-B/16, and a ViT-L/14. The CMR results reported in the original paper are obtained with a model configuration where vision transformer ViT-L/14 is used as an image encoder, and the text transformer is a text encoder. Hence, we use this configuration in our experiments.

X-VLM [61]. This model consists of three encoders: an image encoder, a text encoder, and a cross-modal encoder. The image and text encoder take an image and text as inputs and output their visual and textual representations. The cross-modal encoder fuses the output of the image encoder and the output of the text encoder. The fusion is done via a cross-attention mechanism. For CMR task, the model is fine-tuned via a contrastive learning loss and a matching loss. All encoders are transformer-based. The image encoder is a ViT initialised with Swin Transformer_{base} [37]. Both the text encoder and the cross-modal encoder are initialised using different layers of BERT [8]: the text encoder is initialized using the first six layers, whereas the cross-modal encoder is initialised using the last six layers.

4.2 Methods Excluded from Comparison

While selecting the models for the experiments, we considered other architectures with promising performance on the MS COCO and the Flickr30k datasets. Below, we outline the architectures we considered and explain why they were not included.

Several models such as Visual N-Grams [32], Unicoder-VL [33], ViLT-B/32 [25], UNITER [6] were excluded because they were consistently outperformed by CLIP on the MS COCO and Flickr30k datasets by large margins.

Besides, we excluded ImageBERT [42] because it was outperformed by CLIP on the MS COCO dataset. ALIGN [23], ALBEF [34], VinVL [64], METER [10] were not included because X-VLM consistently outperformed them. UNITER [6] was beaten by both CLIP and X-VLM. We did not include other well-performing models such as ALIGN [23], Flamingo [2], CoCa [60] because the pre-trained models were not publicly available.

5 Experimental Setup

In this section, we discuss our experimental design including the choice of datasets, subtasks, metrics, and implementation details.

5.1 Datasets

We run experiments on two scene-centric and three object-centric datasets. Below, we discuss each of the datasets in more detail.

Scene-centric Datasets. We experiment with two scene-centric datasets: (i) Microsoft COCO (MS COCO) [35] contains 123,287 images depicting regular scenes from everyday life with multiple objects placed in their natural contexts. There are 91 different object types such as “person”, “bicycle”, “apple”. (ii) Flickr30k contains 31,783 images of regular scenes from everyday life, activities, and events. For both scene-centric datasets, we use the splits provided in [24]. The MS COCO dataset is split into 113,287 images for training, 5,000 for testing and 5,000 for validation; the Flickr30k dataset has 29,783 images for training, 1,000 for testing and 1,000 for validation. In both datasets, every image was annotated with five captions using Amazon Mechanical Turk. Besides, we select one caption per image randomly, and use the test set for our experiments.

Object-centric Datasets. We consider three object-centric datasets in our experiments: (i) Caltech-UCSD Birds 200 (CUB-2000) [59] contains 11,788 images of 200 birds species. Each image is annotated with a fine-grained caption from [46]. We selected one caption per image randomly. Each caption is at least 10 words long and does not contain any information about the birds’ species or actions. (ii) Fashion200k contains 209,544 images that depict various fashion items in five product categories (dress, top, pant, skirt, jacket) and their corresponding descriptions. (iii) Amazon Berkley Objects(ABO) [7] contains 147,702 product listings associated with 398,212 images. This dataset was derived from Amazon.com product listings. We selected one image per listing and used the associated product description as its caption. The majority of images depict a single product on a white background. The product is located in the center of the image and takes at least 85% of the image area. For all object-centric datasets, we use the splits provided by the dataset authors and use the test split for our experiments.

5.2 Subtasks

Our goal is to assess and compare the performance of the CMR methods (described in Sect. 5) across the object-centric and scene-centric datasets described in the previous subsection. We design an experimental setup that takes into account two CMR subtasks and two dataset types. It can be summarized using a tree with branches that correspond to different configurations (see Fig. 2). We explain how we cover the branches of this tree in the next subsection.

The tree starts with a root (“Image-text CMR” with label 0) that has sixteen descendants, in total. The root node has two children corresponding to the two image-text CMR subtasks: a text-to-image retrieval (node 1) and image-to-text retrieval (node 2). Since we want to evaluate each of these subtasks on both object-centric and scene-centric datasets, nodes 1 and 2 also have two children each, i.e., the nodes {3, 4, 5, 6}. Finally, every object-centric node has three children: CUB-200, Fashion200k, and ABO datasets {7, 8, 9, 12, 13, 14}; and every scene-centric node has two children: MS COCO and Flickr30k datasets {10, 11, 15, 16}.

5.3 Experiments

To answer the research questions introduced in Sect. 1, we conduct two experiments. In all the experiments, we use CLIP and X-VLM models in a zero-shot setting. Following [55], we focus on relative performance results. In each experiment, we consider different subtrees from Fig. 2. Following [25, 32, 33, 44, 61], we use Recall@K where $K = \{1, 5, 10\}$ to evaluate the model performance in all our experiments. In addition, following [50, 52, 63], we calculate the sum of recalls (rsum) for text-to-image, and image-to-text retrieval tasks as well as the total sum of recalls for both tasks.

For text-to-image retrieval, we first obtain representations for all the candidate images by passing them through the image encoder of the model. Then we




Fig. 2. Our experimental design for evaluating CMR methods across object-centric and scene-centric datasets. The blue colour indicates parts of the tree used in Experiment 1, the green colour indicates parts of the tree used in Experiment 2, and the red colour indicates parts used in all experiments. (Best viewed in color.) (Color figure online)

pass each textual query through the text encoder of the model and retrieve the top- k candidates ranked by cosine similarity w.r.t. the query.

For image-to-text retrieval, we do the reverse, using the texts as candidates and images as queries. More specifically, we start by obtaining representations of the candidate captions by passing them through the text encoder. Afterwards, for each of the visual queries, we pass the query through the image encoder and retrieve top- k candidates ranked by cosine similarity w.r.t. the query.

In *Experiment 1* we evaluate the reproducibility of the CMR results reported in the original publications (RQ1). Both models we consider (CLIP and X-VLM) were originally evaluated on two scene-centric datasets, viz. MS COCO and Flickr30k. Therefore, for our reproducibility study, we also evaluate these models on these two datasets. We evaluate both text-to-image and image-to-text retrieval. That is, we focus on the two sub-trees $0 \leftarrow 1 \leftarrow 4 \leftarrow \{10, 11\}$ and $0 \leftarrow 2 \leftarrow 6 \leftarrow \{15, 16\}$ (the red and blue parts of the tree) from Fig. 2. In addition to relative performance results, we consider absolute differences between the reported scores and the reproduced scores. Following Petrov and Macdonald [41], we assume that the score is reproduced if we obtain a score value equal to the reported score given a relative tolerance of $\pm 5\%$.

In *Experiment 2* we focus on the replicability of the reported results on object-centric datasets (RQ2). Thus, we evaluate CLIP and X-VLM on the CUB-200, Fashion200k, and ABO datasets. This experiment covers the subtrees $0 \leftarrow 1 \leftarrow 3 \leftarrow \{7, 8, 9\}$ and $0 \leftarrow 2 \leftarrow 5 \leftarrow \{12, 13, 14\}$ (the red and green parts of the tree) in Fig. 2.

After obtaining the results from Experiment 1 and 2, we examine the generalizability of the obtained scores (RQ3). We do so by comparing the relative performance results the models achieve on the object-centric versus scene-centric datasets. More specifically, we compare the relative performance of CLIP and X-VLM on CUB-200, Fashion200k, ABO with their relative performance on MS COCO and Flickr30k. Thus, this experiment captures the complete tree in Fig. 2.

Table 1. Results of experiment 1 (reproducibility study), using the MS COCO and Flickr30k datasets. ‘‘Orig.’’ indicates the scores from the original publications. ‘‘Repr.’’ indicates the scores that we obtained.

	Model	Text-to-image			Image-to-text			Rsum		
		R@1	R@5	R@10	R@1	R@5	R@10	t2i	i2t	total
MS COCO (5k)										
Orig.	CLIP [44]	37.80	62.40	72.20	58.40	81.50	88.10	172.40	228.00	400.40
	X-VLM [61]	55.60	82.70	90.00	70.80	92.10	96.50	228.30	259.40	487.70
Repr	CLIP	21.59	40.22	49.80	24.36	44.13	53.41	111.61	121.90	233.51
	X-VLM	42.79	67.61	67.64	64.60	84.48	84.50	178.04	233.58	411.62
Flickr30k (1 k)										
Orig.	CLIP [44]	68.70	90.60	95.20	88.00	98.70	99.40	254.50	286.10	540.60
	X-VLM [61]	71.90	93.30	96.40	85.30	97.80	99.60	261.60	282.70	544.30
Repr.	CLIP	74.95	93.09	96.15	77.02	94.18	96.84	264.19	268.04	532.23
	X-VLM	37.82	82.36	82.48	63.30	91.10	91.10	202.66	245.50	448.16

6 Results

We focus on the reproducibility (different team, same setup) and replicability (different team, different setup) of the CMR experiments reported in the original papers devoted to CLIP [44] and X-VLM [61]. To organize our result presentation, we refer to the tree in Fig. 2. We traverse the tree bottom up, from the leaves to the root.

6.1 RQ1: Reproducibility

To address RQ1, we report on the outcomes of Experiment 1. We investigate to what extent the CMR results reported in the original papers devoted to CLIP [44] and X-VLM [61] are reproducible. Given that both methods were originally evaluated on two scene-centric datasets, viz. MS COCO and Flickr30k, we evaluate the models on the text-to-image and image-to-text tasks on these two datasets. Therefore, we focus on the two blue sub-trees $0 \leftarrow 1 \leftarrow 4 \leftarrow \{10, 11\}$ and $0 \leftarrow 2 \leftarrow 6 \leftarrow \{15, 16\}$ from Fig. 2.

Results. The results of Experiment 1 are shown in Table 1. We recall the scores obtained in the original papers [44, 61] (“Orig.”) and the scores that we obtained (“Repr.”), on the MS COCO and Flickr30k datasets. Across the board, the scores that we obtained (the “reproduced scores”) tend to be lower than the scores obtained in the original publications (the “original scores”).

On the MS COCO dataset, X-VLM consistently outperforms CLIP, both in the original publications and in our setup, for both the text-to-image and the image-to-text tasks. Moreover, this holds for all $R@n$ metrics, and, hence, for the Rsum metrics. Interestingly, the relative gains that we obtain tend to be larger than the ones obtained in the original publications. For example, our biggest relative difference is for the image-to-text task in terms of the $R@1$ metric: according to the scores reported in [44, 61], X-VLM outperforms CLIP by 21%, whereas in our experiments the relative gain is 165%.

On average, the original CLIP scores are as much as $\sim 70\%$ higher than the reproduced scores; the original scores for X-VLM are $\sim 20\%$ higher than the reproduced ones. When considering the absolute differences between the original scores and the reproduced scores and assuming a relative tolerance of $\pm 5\%$, we see that, on the MS COCO dataset, the scores are not reproducible for both models.

On the Flickr30k dataset, we see a different pattern. For the text-to-image task, the original results indicate that X-VLM consistently outperforms CLIP, on all $R@n$ metrics, but according to our results, the relative order is consistently reversed. For the image-to-text task, we obtained mixed outcomes: for $R@1$ and $R@5$, the original order (CLIP outperforms X-VLM) is confirmed, but for $R@10$ the order is swapped. According to our experimental results, however, CLIP consistently outperforms X-VLM on all tasks, and on all $R@n$ metrics (and hence also on the Rsum metrics).

On the Flickr30k dataset, the CLIP scores are reproduced on the text-to-image and image-to-text retrieval tasks when the model is evaluated on R@5 and R@10. On the text-to-image task, the reproduced R@5 score is 2.7% higher than the original score; the reproduced R@10 score is 1% higher than the original score. For the image-to-text retrieval task, the reproduced R@5 score is 4% lower than the original score; the reproduced R@10 score is 2% lower than the original score.

Answer to RQ1. In the case of the CLIP model, the obtained *absolute* scores were reproducible only on the Flickr30k dataset for the text-to-image and the image-to-text tasks when evaluated on R@5 and R@10. For X-VLM, we did not find the absolute scores obtained when evaluating the model on the MS COCO and Flickr20k datasets to be reproducible, neither for the text-to-image nor the image-to-text tasks.

The *relative* outcomes on the MS COCO dataset could be reproduced, for all tasks and metrics, whereas on the Flickr30k dataset they could only partially be reproduced, that is, only for the image-to-text task on the R@1 and R@5 metrics; for the text-to-image task, X-VLM outperforms CLIP according to the original scores, but CLIP outperforms X-VLM according to our reproduced scores.

Upshot. As explained in Sect. 4, in this paper we focus on CMR in a zero-shot setting. This implies that the differences that we observed between the original scores and the reproduced scores must be due to differences in text and image data (pre-)processing and loading. We, therefore, recommend that the future work includes (as much as is practically possible) tools and scripts used in these stages of the experiment with the publication of its implementations.

6.2 RQ2: Replicability

To answer RQ2, we replicate the originally reported text-to-image and image-to-text retrieval experiments in a different setup, i.e., by evaluating CLIP and X-VLM using object-centric datasets instead of scene-centric datasets. Thus, we evaluate CLIP and X-VLM on the CUB-200, Fashion200k, and ABO datasets and focus on the green subtrees $0 \leftarrow 1 \leftarrow 3 \leftarrow \{7, 8, 9\}$ and $0 \leftarrow 2 \leftarrow 5 \leftarrow \{12, 13, 14\}$ from Fig. 2.

Results. The results of Experiment 2 (aimed at answering RQ2) can be found in Table 2. On the CUB-200 dataset, CLIP consistently outperforms X-VLM. The biggest relative increase is 124% for image-to-text in terms of R@10, while the smallest relative increase is 1% for text-to-image in terms of R@1. Overall, on the text-to-image retrieval task, CLIP outperforms X-VLM by 38%, and on the image-to-text retrieval task, the relative gain is 70%.

Table 2. Results of Experiment 2 (replicability study), using the CUB-200, Fashion200k, and ABO datasets.

Model	Text-to-image			Image-to-text			Rsum		
	R@1	R@5	R@10	R@1	R@5	R@10	t2i	i2t	total
CUB-200									
CLIP	0.71	2.38	4.42	1.23	3.40	5.48	7.51	10.11	17.62
X-VLM	0.70	2.28	2.45	1.16	2.35	2.45	5.43	5.96	11.39
Fashion200k									
CLIP	3.05	8.56	12.85	3.43	9.82	14.56	24.46	27.81	52.27
X-VLM	2.80	6.62	6.70	1.84	3.96	4.04	16.12	09.84	25.96
ABO									
CLIP	6.25	13.90	18.50	7.99	18.96	25.57	38.65	52.52	91.17
X-VLM	3.10	6.48	6.56	3.20	7.42	7.50	16.14	18.12	34.26

On Fashion200k, CLIP outperforms X-VLM, too. The smallest relative increase is 9% for text-to-image in terms of R@1, the biggest relative increase is 260% for image-to-text in terms of R@10. In general, on the text-to-image retrieval task, CLIP outperforms X-VLM by 52%; on the image-to-text retrieval task, the relative gain is 83%.

Finally, on the ABO dataset, CLIP outperforms X-VLM again. The smallest relative increase is 101% for text-to-image in terms of R@1, the biggest relative increase is 241% for image-to-text again in terms of R@10. In general, on the text-to-image retrieval task, CLIP outperforms X-VLM by 139%; on the image-to-text retrieval task, the relative gain is 190%. All in all, CLIP outperforms X-VLM on all three scene-centric datasets. The overall relative gain on CUB-200 dataset is 55%, on Fashion200k dataset -101% . The biggest relative gain of 166% is obtained on the ABO dataset.

Answer to RQ2. The outcome of Experiment 2 is clear. The original relative performance results obtained on the MS COCO and Flickr30k (Table 1) are only partially replicable to the CUB-200, Fashion200k, and ABO datasets. On the latter datasets CLIP consistently outperforms X-VLM by a large margin, whereas the original scores obtained on the former datasets indicate that X-VLM mostly outperforms CLIP.

Upshot. We hypothesize that the failure to replicate the relative results originally reported for scene-centric datasets (viz. X-VLM outperforms CLIP) is due to CLIP being pre-trained on more and more diverse image data. We, therefore, recommend that future work aimed at developing large-scale CMR models quantifies and reports the diversity of the training data used.

6.3 RQ3: Generalizability

To answer RQ3, we compare the relative performance of the selected models on object-centric and scene-centric data. Thus, we compare the relative performance of CLIP and X-VLM on CUB-200, Fashion200k, ABO with their relative performance on MS COCO and Flickr30k. We focus on the complete tree from Fig. 2.

Results. The results of our experiments on the scene-centric datasets are in Table 1; the results that we obtained on the object-centric datasets are in Table 2. On object-centric datasets, CLIP consistently outperforms X-VLM. However, the situation with scene-centric results is partially the opposite. There, X-VLM outperforms CLIP on the MS COCO dataset.

Answer to RQ3. Hence, we answer RQ3 by stating that the relative performance results for CLIP and X-VLM that we obtained in our experiments only partially generalize from scene-centric to object-centric datasets. The MS COCO dataset is the odd one out.²

Upshot. Given the observed differences in relative performance results for CLIP and X-VLM on scene-centric vs. object-centric datasets, we recommend that CMR be trained in both scene-centric and object-centric datasets to help improve the generalizability of experimental outcomes.

7 Discussion and Conclusions

We have examined two SOTA image-text CMR methods, CLIP and X-VLM, by contrasting their performance on two scene-centric datasets (MS COCO and Flickr30k) and three object-centric datasets (CUB-200, Fashion200k, ABO) in a zero-shot setting.

We focused on the *reproducibility* of the CMR results reported in the original publications when evaluated on the selected scene-centric datasets. The reported scores were not reproducible for X-VLM when evaluated on the MS COCO and the Flickr30k datasets. For CLIP, we were able to reproduce the scores on the Flickr30k dataset when evaluated using R@5 and R@10. Conversely, the relative results were reproducible on the MS COCO dataset, for all metrics and tasks, and partially reproducible on the Flickr30k dataset only for image-to-text task when evaluated on R@1 and R@5. We also examined the *replicability* of the CMR results using three object-centric datasets. We discovered that the relative results are replicable when we compare the relative performance on the object-centric datasets with the relative scores on the Flickr30k dataset. However, for

² On the GitHub repository for CLIP, several issues have been posted related to the performance of CLIP on the MS COCO dataset. See, e.g., <https://github.com/openai/CLIP/issues/115>.

the MS COCO dataset, the relative outcomes were not replicable. And, finally, we explored the generalizability of the obtained results by comparing the models' performance on scene-centric vs. object-centric datasets. We observed that the absolute scores obtained when evaluating models on object-centric datasets are much lower than the scores obtained on scene-centric datasets.

Our findings demonstrate that the reproducibility of CMR methods on scene-centric datasets is an open problem. Besides, we show that while the majority of CMR methods are evaluated on the MS COCO and the Flickr30k datasets, the object-centric datasets represent a challenging and relatively unexplored set of benchmarks.

A limitation of our work is the relatively small number of scene-centric and object-centric datasets used for the evaluation of the models. Another limitation is that we only considered CMR in a zero-shot setting, ignoring, e.g., few-shot scenarios; this limitation did, however, come with the important advantage of reducing the number of experimental design decisions to be made for contrastive experiments.

A promising direction for future work is to include further datasets when contrasting the performance of CMR models, both scene-centric and object-centric. In particular, it would be interesting to investigate the models' performance on datasets, e.g., Conceptual Captions [47], the Flower [40], and the Cars [27] datasets. A natural step after that would be to consider few-shot scenarios.

Acknowledgements. We thank Paul Groth, Andrew Yates, Thong Nguyen, and Maurits Bleeker for helpful discussions and feedback.

This research was supported by Ahold Delhaize, and the Hybrid Intelligence Center, a 10-year program funded by the Dutch Ministry of Education, Culture and Science through the Netherlands Organisation for Scientific Research, <https://hybrid-intelligence-centre.nl>.

All content represents the opinion of the authors, which is not necessarily shared or endorsed by their respective employers and/or sponsors.

References

1. ACM (2020) Artifact Review and Badging - Current. <https://www.acm.org/publications/policies/artifact-review-and-badging-current> Accessed Aug 7 2022
2. Alayrac, J.B., et al.: Flamingo: a visual language model for few-shot learning. arXiv preprint [arXiv:2204.14198](https://arxiv.org/abs/2204.14198) (2022)
3. Bleeker, M., de Rijke, M.: Do lessons from metric learning generalize to image-caption retrieval? In: Hagen, M., et al. (eds.) ECIR 2022. LNCS, vol. 13185, pp. 535–551. Springer, Cham (2022). https://doi.org/10.1007/978-3-030-99736-6_36
4. Brown, A., Xie, W., Kalogeiton, V., Zisserman, A.: Smooth-AP: smoothing the path towards large-scale image retrieval. In: Vedaldi, A., Bischof, H., Brox, T., Frahm, J.-M. (eds.) ECCV 2020. LNCS, vol. 12354, pp. 677–694. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-58545-7_39
5. Carvalho, M., Cadène, R., Picard, D., Soulier, L., Thome, N., Cord, M.: Cross-modal retrieval in the cooking context: Learning semantic text-image embeddings. In: The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval, pp. 35–44 (2018)

6. Chen, Y.C., et al.: Uniter: Learning universal image-text representations. In: Computer Vision - ECCV 2020, Springer International Publishing, pp. 104–120 (2020)
7. Collins, J., et al.: Abo: Dataset and benchmarks for real-world 3d object understanding. CVPR (2022)
8. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint (2018). [arXiv:1810.04805](https://arxiv.org/abs/1810.04805)
9. Dosovitskiy, A., et al. (2020) An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint [arXiv:2010.11929](https://arxiv.org/abs/2010.11929)
10. Dou, Z.Y., et al.: An empirical study of training end-to-end vision-and-language transformers. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 18166–18176 (2022)
11. Faghri, F., Fleet, D.J., Kiros, J.R., Fidler, S.: Vse++: Improving visual-semantic embeddings with hard negatives (2017). arXiv preprint [arXiv:1707.05612](https://arxiv.org/abs/1707.05612)
12. Frome, A.: Devise: A deep visual-semantic embedding model. In: Proceedings of the 26th International Conference on Neural Information Processing Systems - Volume 2, Curran Associates Inc., NIPS'13, pp 2121–2129 (2013)
13. Gao, D.: Fashionbert: Text and image matching with adaptive loss for cross-modal retrieval. In: Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 2251–2260 (2020)
14. Goei, K., Hendriksen, M., de Rijke, M.: Tackling attribute fine-grainedness in cross-modal fashion search with multi-level features. In: SIGIR 2021 Workshop on eCommerce, ACM (2021)
15. Gong, Y., Wang, L., Hodosh, M., Hockenmaier, J., Lazebnik, S.: Improving image-sentence embeddings using large weakly annotated photo collections. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) ECCV 2014. LNCS, vol. 8692, pp. 529–545. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-10593-2_35
16. Gu, J., Cai, J., Joty, S.R., Niu, L., Wang, G.: Look, imagine and match: Improving textual-visual cross-modal retrieval with generative models. In: Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 7181–7189 (2018)
17. Han, X., et al.: Automatic spatially-aware fashion concept discovery. In: Proceedings of the IEEE international Conference on Computer Vision, pp. 1463–1471 (2017)
18. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 770–778 (2016)
19. Hendriksen, M.: Multimodal retrieval in e-commerce. In: Hagen, M., et al. (eds.) ECIR 2022. LNCS, vol. 13186, pp. 505–512. Springer, Cham (2022). https://doi.org/10.1007/978-3-030-99739-7_62
20. Hendriksen, M., Bleeker, M., Vakulenko, S., van Noord, N., Kuiper, E., de Rijke, M.: Extending CLIP for category-to-image retrieval in e-commerce. In: Hagen, M., et al. (eds.) ECIR 2022. LNCS, vol. 13185, pp. 289–303. Springer, Cham (2022). https://doi.org/10.1007/978-3-030-99736-6_20
21. Herranz, L., Jiang, S., Li, X.: Scene recognition with cnns: objects, scales and dataset bias. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 571–579 (2016)
22. Hu, P., Zhen, L., Peng, D., Liu, P.: Scalable deep multimodal learning for cross-modal retrieval. In: Proceedings of the 42nd international ACM SIGIR conference on research and development in information retrieval, pp. 635–644 (2019)

23. Jia, C., et al.: Scaling up visual and vision-language representation learning with noisy text supervision. In: International Conference on Machine Learning, PMLR, pp. 4904–4916 (2021)
24. Karpathy, A., Fei-Fei, L.: Deep visual-semantic alignments for generating image descriptions. In: Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 3128–3137 (2015)
25. Kim, W., Son, B., Kim, I.: Vilt: Vision-and-language transformer without convolution or region supervision. In: International Conference on Machine Learning, PMLR, pp. 5583–5594 (2021)
26. Klein, B., Lev, G., Sadeh, G., Wolf, L.: Fisher vectors derived from hybrid gaussian-laplacian mixture models for image annotation. (2014) arXiv preprint [arXiv:1411.7399](https://arxiv.org/abs/1411.7399)
27. Krause, J., Stark, M., Deng, J., Fei-Fei, L.: 3d object representations for fine-grained categorization. In: 4th International IEEE Workshop on 3D Representation and Recognition (3dRR-13), Sydney, Australia (2013)
28. Laenen, K.: Cross-modal representation learning for fashion search and recommendation. PhD thesis, KU Leuven (2022)
29. Laenen, K., Zoghbi, S., Moens, M.F.: Cross-modal search for fashion attributes. In: Proceedings of the KDD 2017 Workshop on Machine Learning Meets Fashion, ACM, vol 2017, pp. 1–10 (2017)
30. Laenen, K., Zoghbi, S., Moens, M.F.: Web search of fashion items with multimodal querying. In: Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining, pp. 342–350 (2018)
31. Lee, K.H., Chen, X., Hua, G., Hu, H., He, X.: Stacked cross attention for image-text matching. In: Proceedings of the European Conference on Computer Vision (ECCV), pp. 201–216 (2018)
32. Li, A., Jabri, A., Joulin, A., Van Der Maaten, L.: Learning visual n-grams from web data. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 4183–4192 (2017)
33. Li, G., Duan, N., Fang, Y., Gong, M., Jiang, D.: Unicoder-vl: a universal encoder for vision and language by cross-modal pre-training. Proc. AAAI Conf. Artif. Intell. **34**, 11336–11344 (2020)
34. Li, J., Selvaraju, R., Gotmare, A., Joty, S., Xiong, C., Hoi, S.C.H.: Align before fuse: vision and language representation learning with momentum distillation. Adv. Neural. Inf. Process. Syst. **34**, 9694–9705 (2021)
35. Lin, T.-Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft COCO: common objects in context. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) ECCV 2014. LNCS, vol. 8693, pp. 740–755. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-10602-1_48
36. Liu, C., Mao, Z., Liu, A.A., Zhang, T., Wang, B., Zhang, Y.: Focus your attention: A bidirectional focal attention network for image-text matching. In: Proceedings of the 27th ACM International Conference on Multimedia, pp. 3–11 (2019)
37. Liu, Z.: Swin transformer: Hierarchical vision transformer using shifted windows. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 10012–10022 (2021)
38. Messina, N., Amato, G., Esuli, A., Falchi, F., Gennaro, C., Marchand-Maillet, S.: Fine-grained visual textual alignment for cross-modal retrieval using transformer encoders. ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM) **17**(4), 1–23 (2021)

39. Nam, H., Ha, J.W., Kim, J.: Dual attention networks for multimodal reasoning and matching. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 299–307 (2017)
40. Nilsback, M.E., Zisserman, A.: Automated flower classification over a large number of classes. In: Indian Conference on Computer Vision, Graphics and Image Processing (2008)
41. Petrov, A., Macdonald, C.: A systematic review and replicability study of bert4rec for sequential recommendation. In: Proceedings of the 16th ACM Conference on Recommender Systems, pp. 436–447 (2022)
42. Qi, D., Su, L., Song, J., Cui, E., Bharti, T., Sacheti, A.: Imagebert: Cross-modal pre-training with large-scale weak-supervised image-text data. (2020) arXiv preprint [arXiv:2001.07966](https://arxiv.org/abs/2001.07966)
43. Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I., et al.: Language models are unsupervised multitask learners. OpenAI blog **1**(8), 9 (2019)
44. Radford, A., et al.: Learning transferable visual models from natural language supervision. In: International Conference on Machine Learning, PMLR, pp. 8748–8763 (2021)
45. Rao, J, et al.: Where does the performance improvement come from?: - A reproducibility concern about image-text retrieval. In: SIGIR '22: The 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, Madrid, Spain, July 11–15, 2022, ACM, pp. 2727–2737 (2022)
46. Reed, S., Akata, Z., Lee, H., Schiele, B.: Learning deep representations of fine-grained visual descriptions. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 49–58 (2016)
47. Sharma, P., Ding, N., Goodman, S., Soricut, R.: Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In: Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pp. 2556–2565 (2018)
48. Shen, Z.Y., Han, S.Y., Fu, L.C., Hsiao, P.Y., Lau, Y.C., Chang, S.J.: Deep convolution neural network with scene-centric and object-centric information for object detection. *Image Vis. Comput.* **85**, 14–25 (2019)
49. Sheng, S., Laenen, K., Van Gool, L., Moens, M.F.: Fine-grained cross-modal retrieval for cultural items with focal attention and hierarchical encodings. *Computers* **10**(9), 105 (2021)
50. Song, J., Choi, S. Image-text alignment using adaptive cross-attention with transformer encoder for scene graphs (2021)
51. Tan, M., Le, Q.: Efficientnet: Rethinking model scaling for convolutional neural networks. In: International Conference on Machine Learning, PMLR, pp. 6105–6114 (2019)
52. Ueki, K.: Survey of visual-semantic embedding methods for zero-shot image retrieval. In: 2021 20th IEEE International Conference on Machine Learning and Applications (ICMLA), IEEE, pp. 628–634 (2021)
53. Varamesh, A., Diba, A., Tuytelaars, T., Van Gool, L.: Self-supervised ranking for representation learning. (2020) arXiv preprint [arXiv:2010.07258](https://arxiv.org/abs/2010.07258)
54. Vaswani, A.: Attention is all you need. In: Advances in Neural Information Processing Systems, vol. 30 (2017)
55. Voorhees, E.M.: The philosophy of information retrieval evaluation. In: Evaluation of Cross-Language Information Retrieval Systems, pp. 355–370. Springer, Berlin Heidelberg (2002)

56. Wang, H., et al.: Cross-modal food retrieval: learning a joint embedding of food images and recipes with semantic consistency and attention mechanism. *IEEE Trans. Multimedia* **24**, 2515–2525 (2021)
57. Wang, K., Yin, Q., Wang, W., Wu, S., Wang, L.: A comprehensive survey on cross-modal retrieval. *arXiv preprint* (2016). [arXiv:1607.06215](https://arxiv.org/abs/1607.06215)
58. Wang, L., Li, Y., Lazebnik, S.: Learning deep structure-preserving image-text embeddings. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 5005–5013 (2016)
59. Welinder, P.: Caltech-UCSD Birds 200. Tech. Rep. CNS-TR-2010-001, California Institute of Technology (2010)
60. Yu, J., Wang, Z., Vasudevan, V., Yeung, L., Seyedhosseini, M., Wu, Y.: Coca: Contrastive captioners are image-text foundation models. (2022). *arXiv preprint* [arXiv:2205.01917](https://arxiv.org/abs/2205.01917)
61. Zeng, Y., Zhang, X., Li, H.: Multi-grained vision language pre-training: Aligning texts with visual concepts. In: Chaudhuri K, Jegelka S, Song L, Szepesvári C, Niu G, Sabato S (eds) *International Conference on Machine Learning, ICML 2022*, 17–23 July 2022, Baltimore, Maryland, USA, PMLR, *Proceedings of Machine Learning Research*, vol 162, pp. 25994–26009 (2022)
62. Zhang, C., et al.: Mosaicos: a simple and effective use of object-centric images for long-tailed object detection. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 417–427 (2021)
63. Zhang, K., Mao, Z., Wang, Q., Zhang, Y.: Negative-aware attention framework for image-text matching. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 15661–15670 (2022)
64. Zhang, P., et al.: Vinvl: Revisiting visual representations in vision-language models. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 5579–5588 (2021)
65. Zhang, Y., Jiang, H., Miura, Y., Manning, C.D., Langlotz, C.P.: Contrastive learning of medical visual representations from paired images and text (2020). *arXiv preprint* [arXiv:2010.00747](https://arxiv.org/abs/2010.00747)
66. Zhou, B., Lapedriza, A., Xiao, J., Torralba, A., Oliva, A.: Learning deep features for scene recognition using places database. In: *Advances in Neural Information Processing Systems* vol. 27 (2014)
67. Zhuge, M., et al.: Kaleido-bert: Vision-language pre-training on fashion domain. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 12647–12657 (2021)