

GABInsight: Exploring Gender-Activity Binding Bias in Vision-Language Models

Ali Abdollahi^{a,1}, Mahdi Ghaznavi^{a,1}, Mohammad Reza Karimi Nejad^{a,d}, Arash Mari Oriyad^a, Reza Abbasi^a, Ali Salesi^a, Melika Behjati^{b,c}, Mohammad Hossein Rohban^a and Mahdiah Soleymani Baghshah^{a,*}

^aSharif University of Technology

^becole polytechnique fédérale de lausanne (EPFL)

^cIDIAP research institute

^dShamsipour Technical and Vocational College

Abstract. Vision-language models (VLMs) are intensively used in many downstream tasks, including those requiring assessments of individuals appearing in the images. While VLMs perform well in simple single-person scenarios, in real-world applications, we often face complex situations in which there are persons of different genders doing different activities. We show that in such cases, VLMs are biased towards identifying the individual with the expected gender (according to ingrained gender stereotypes in the model or other forms of sample selection bias) as the performer of the activity. We refer to this bias in associating an activity with the gender of its actual performer in an image or text as the Gender-Activity Binding (GAB) bias and analyze how this bias is internalized in VLMs. To assess this bias, we have introduced the GAB dataset with approximately 5500 AI-generated images that represent a variety of activities, addressing the scarcity of real-world images for some scenarios. To have extensive quality control, the generated images are evaluated for their diversity, quality, and realism. We have tested 12 renowned pre-trained VLMs on this dataset in the context of text-to-image and image-to-text retrieval to measure the effect of this bias on their predictions. Additionally, we have carried out supplementary experiments to quantify the bias in VLMs' text encoders and to evaluate VLMs' capability to recognize activities. Our experiments indicate that VLMs experience an average performance decline of about 13.2% when confronted with gender-activity binding bias.

1 Introduction

Nowadays large multi-modal models have shown tremendous potential in various tasks, from information retrieval systems to image captioning, visual question answering, image generation, and visual reasoning. As a branch of multi-modal models, Vision-Language Models (VLMs) [23] that provide a shared cross-modal embedding space between text and image modalities, have been intensively investigated by researchers recently and are used in many real-world applications [41, 25, 11].

Despite their wide range of success in downstream tasks, VLMs are subject to many critical biases, which can consequently affect their performance, especially in sensitive applications. For example, Agarwal et al. [3] provides a comprehensive analysis of biases and

their implications in the CLIP Models, revealing how this model is exposed to gender biases and how challenging it is to trust it in real-world applications. As Srinivasan and Bisk [29] and Lee et al. [18] state, the bias is introduced in visual-linguistic pre-training due to unfairness in the training data, and additionally, at inference time, the visual and linguistic contexts that is used for few-shot applications can also promote bias.

One important type of bias, which is highly discussed in the literature, is the gender bias. Gender bias demonstrates the undesirable association of a factor with a gender according to the model. For instance, due to the high occurrence of data in which men are repairing devices, this activity is usually associated with men by deep models, which leads to failure when faced with women repairing devices. Hall et al. [12] suggests that different VLMs do not perform equally well on determining gender for an occupation, and do not assign equal retrieval likelihood to images of male and female professionals. Also, Lee et al. [18] and Wang et al. [33] show that CLIP [23] shows biased behavior in retrieval based on gender-neutral queries. Please note that, in line with common practice in the literature, our focus is on two genders, masculine and feminine. This approach streamlines the process of generating the dataset and analyzing the experiments.

In this research, we delve into the issue of gender-activity binding in retrieval tasks. We aim to retrieve the corresponding caption or image from two given captions or images, each depicting an activity performed by a different gender, based on a provided image or caption that shows one of the genders performing the activity. We refer to this bias in associating gender and activity as the *Gender-Activity Binding (GAB)* bias. This bias arises due to ingrained gender stereotypes in the model or other types of sample selection biases [5]. Our research indicates that VLMs do not display a substantial bias in retrieval tasks when both text and image modalities depict only one gender. This is because the identification of the performer's gender can directly lead to the correct output, eliminating the need to bind the performed activity with the performer's gender. However, text encoders have shown a considerable bias towards the expected gender. The bias becomes more pronounced when the scene is more complex, i.e., when two individuals of different genders are present in the image. We observe a drop in retrieval accuracy in scenarios where the activity performer is the unexpected gender. For example, in the context of the "repairing" activity, the retrieval accuracy of VLMs in identifying the performer decreases when a woman is re-

* Corresponding Author. Email: soleymani@sharif.edu.

¹ Equal contribution.

pairing a device and a man is also present in the scene. This is in comparison with scenarios where the performer is a man, or there is no man present in the scene to associate the activity with him.

The significance of this bias escalates when we consider its underlying implications. Not only does it perpetuate societal biases and fairness concerns related to gender stereotypes, but it can also lead to serious complications if incorporated into judgement and decision-making systems. The potential for such bias to inadvertently influence outcomes underscores the importance of addressing it.

We have created a dataset known as the *Gender-Activity Binding (GAB)* dataset as shown in Figure 1(a). This dataset comprises images depicting various activities performed by men or women in different contexts (with or without the presence of an individual of the opposite gender), each accompanied by a descriptive caption. Given that our scenarios are uncommon due to biases and stereotypes, suitable real-world images are scarce for evaluation purposes. To address this issue, all images are generated using DALL-E 3². We employ extensive prompt enhancement techniques to ensure the diversity and quality of the generated images. We also evaluated the generated images on diversity, quality, and realism. We performed multiple quantitative assessments consisting of Fréchet Inception Distance (FID), Structural Similarity Index (SSIM), Learned Perceptual Image Patch Similarity (LPIPS), and coverage score[21] for quality and diversity assessment, and human feedback for a final qualitative assessment to ensure the final results accurately represent the intended activities and their related contexts. The activities included in GAB dataset are sourced from various places to encapsulate societal gender stereotypes and other activities that are statistically biased towards a specific gender, as observed in the LAION-400M dataset [26].

We benchmark the performance of 12 renowned multi-modal foundation models and assess their performance in terms of text-to-image retrieval, image-to-text retrieval, and recognizing activities on GAB dataset as shown in Figure 1(b). Specifically, in the image-to-text retrieval task, we observed that while an unexpected gender is doing a stereotypically biased activity, the average performance of VLMs declined by approximately 33.2% due to the presence of the expected gender in the scene. Moreover, VLMs undergo an average accuracy reduction of approximately 13.2% when encountering gender-activity binding bias. On the other hand, in the text-to-image retrieval task, most VLMs achieve an accuracy of approximately 50%, indicating that their performance is nearly random, and they fail to recognize the performer of the activity based on the text.

As far as we are aware, this research is the first to delve into the origins of the association between genders and activities in vision-language models by conducting a thorough analysis both image-to-text and text-to-image retrieval accuracy of VLMs. It encompasses not only straightforward scenarios involving single individuals, where VLMs demonstrate high accuracy for both genders [15, 4, 30], but also more complex situations with individuals of different genders [12]. This comprehensive approach allows us to isolate the impact of the complexity of the scenario on performance decline [32] and more accurately measure the Gender-Activity Binding bias. Additionally, the study carries out supplementary experiments to investigate the bias in VLMs’ text encoders and their proficiency in recognizing activities.^{3,4}

To summarize, our contributions are as follows:

- We have created a novel dataset, known as the gender-activity binding dataset. This dataset includes AI-generated images of various activities that are typically associated with a specific gender, and are excellent in terms of diversity, quality, and realism (Section 3).
- We conduct an intensive performance benchmark of well-known vision-language models in retrieval tasks to assess their robustness against the Gender-Activity Binding bias. Our findings reveal that when two individuals of different genders are present, VLMs exhibit a bias towards binding the activity with the gender that is expected to perform it, resulting in a 13.2% drop in retrieval accuracy. However, in scenarios involving a single individual, these models demonstrate high image-to-text retrieval accuracy. We also demonstrate that VLMs lack the ability to bind gender and activity in text-to-image retrieval tasks (Section 4).
- We carry out additional experiments to investigate the binding bias in VLMs’ capability to comprehend activities and the bias present in their text encoders (Section 4 and the Supplementary Material [1]). We delve into the insights that can be derived about how VLMs internalize this bias and its impact on the shared embedding space of VLMs (Section 5).

2 Gender-Activity Binding Bias

Despite the recent increase in large-scale data collection, it’s important to note that these datasets are not necessarily devoid of biases. These biases can manifest in large datasets derived from real-world distributions, primarily due to sample selection bias. The issue causes VLMs to have varying levels of accuracy in recognizing the performer of an activity depending on the gender of the individual involved [40]. We refer to this bias in associating an activity with the gender of its actual performer in an image or text as *Gender-Activity Binding Bias*.

Consider an image-to-text retrieval task, where we have two pieces of text C_1 and C_2 , each describing an activity performed by a different gender. We also have an image I_1 depicting a person (either a man or a woman) performing an action. The task is to match the image with the correct piece of text (Figure 1-(b) left). Also, consider the text-to-image retrieval task, where we have two images I_1 and I_2 , each showing an activity being performed by a specific gender, and we want to retrieve the correct image based on a given piece of text C_1 (Figure 1-(b) right). The Gender-Activity Binding bias manifests as a decline in the retrieval accuracy of VLMs when they incorrectly identify the gender of an individual performing activities typically associated with a different gender.

Let’s denote the probability of a VLM correctly identifying the performer of an activity in a data i (image or text) as p , the activity as a , and the gender as g . Ideally, the probability $p(g|a, i)$ should be equal for all g . However, because of the co-occurrence of gender g of the performer and the activity in the dataset, which is referred to as sample selection bias [5], the gender of the performer becomes correlated with the activity. Thus, we often find that $p(g_1|a, i) > p(g_2|a, i)$, where g_1 and g_2 represent different genders and g_1 shows the expected gender for the performer of a . In simple terms, VLMs tend to incorrectly select the text or image because they are biased towards associating the activity with a specific gender. This is not always accurate, leading to errors in scenarios where the activity is performed by the less expected gender.

The accuracy of models in image and text retrieval tasks is assessed in two scenarios: one where only a person of an unexpected gender (in terms of the bias) is depicted or described in the text or im-

² <https://openai.com/dall-e3>

³ Supplementary materials are available on Arxiv [1]: <https://arxiv.org/abs/2407.21001>

⁴ The dataset and codes for reproducing results are available on Github [2]: <https://github.com/sharif-ml-lab/GABInsight>

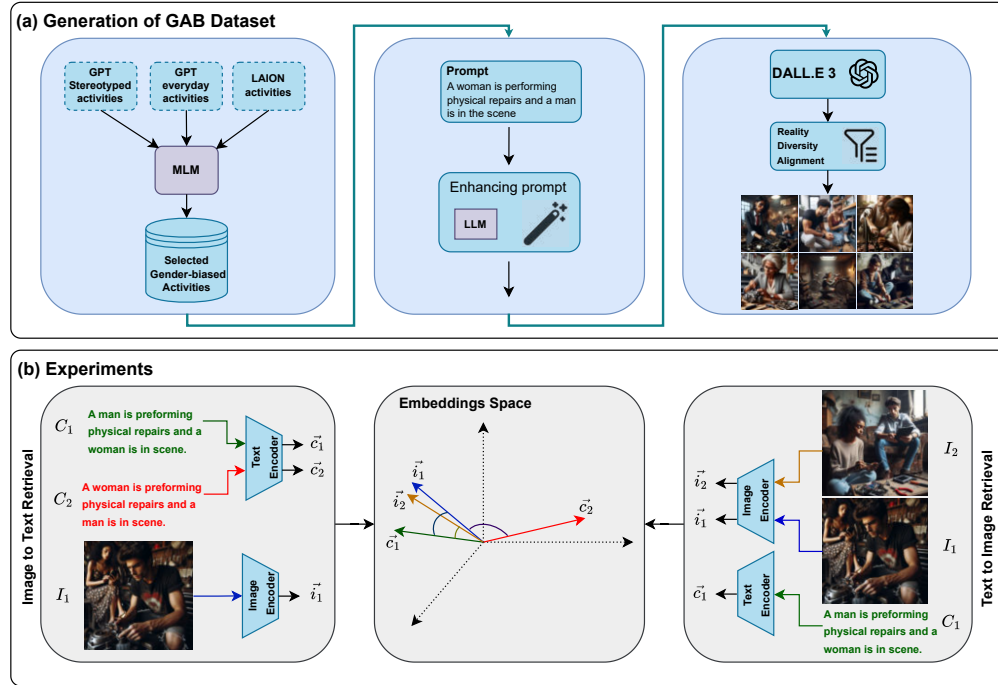


Figure 1. Overview of (a) the creation process of the introduced dataset and (b) the empirical tests conducted to assess the gender-activity binding bias in retrieval tasks within vision-language models. (a) Left: we gather three sets of activities that show bias, including stereotypical, everyday activities, and those that exhibit gender bias in the captions of LAION-400M [26]. Middle: we employ prompt enhancement techniques to develop a diverse, descriptive, and detailed prompt from a basic initial one, aiding us in generating a wider range of images with superior quality and realism. Right: we utilize DALL-E 3 to construct our dataset based on the enhanced prompts. The generated images are selected to align with the activity and scenario and are evaluated for diversity, quality, and realness to achieve a high score based on standard metrics. (b) Middle: joint embedding space of text and images in vision-language models. Left/Right: an overview of image-to-text/text-to-image retrieval tasks. The caption/image with the highest cosine similarity to the input image/caption is retrieved.

age, and another where both genders are present in the data, but the activity is performed by only one of them. Consequently, the impact of the bias can be observed as a decrease in accuracy in two situations: when the expected gender is absent or present, and when the activity is performed by the unexpected or expected gender. These cases together provide better insights into how these models handle gender-activity associations under different conditions. The first scenario allows us to understand how well the models can recognize and correctly associate an activity with an unexpected gender when the other gender is not present. This helps us gauge the models' ability to break away from societal stereotypes and biases. The second scenario, on the other hand, tests the models' ability to correctly identify the performer of an activity when both genders are present, which is a more complex and realistic situation. This indicates these biases affect the models' performance in real-world scenarios. Together, these cases provide a comprehensive understanding of the models' strengths and weaknesses in handling gender-activity binding, paving the way for targeted improvements to mitigate the gender-activity binding bias.

3 Dataset

This section details the creation process of the Gender-Activity Binding (GAB) Dataset, which serves to evaluate various VLMs against the gender-activity binding bias we have identified and their ability to associate activities with performers. This dataset comprises AI-

generated images that illustrate a variety of activities with potential subjects of both genders.

3.1 Overview of the dataset

The GAB dataset divides images into four distinct groups:

1. In two groups, both genders are present in the scene. In one of these groups, the expected gender is performing the activity while the other gender is just present in the image. In the other group, the unexpected gender is performing the activity while the other gender is also in the scene.
2. In the next two groups, only one gender is present in the scene. In one of these groups, the expected gender is performing the action. In the other group, the scenario is reversed, with the unexpected gender acting.

For each activity, with respect to these groupings, we can consider four groups based on the performer of the activity (expected (E) or unexpected (U)) and the number of persons (with different genders) in the scene (1 or 2): E1, E2, U1, U2.

For each image in the dataset, we have templates that are replaced with activity or gender identifiers ("man" and "woman") in accordance with the designed experiments. Details of the experiments and these replacements can be found in Section 4. The template for images containing two genders is: "a <man/woman> is <doing activity> and a <woman/man> is in the scene," and for images containing one

gender, the template is: "a <man/woman> is <doing activity>." This specific template can also be used for images with two genders.

Each of the mentioned groups contains at least 26 images for each activity, resulting in a total of 5500 images.

This dataset structure facilitates the evaluation of gender-activity binding bias and the ability of VLMs to bind and recognize activities by defining different retrieval tasks, as you can see in Section 4.

The pipeline of the dataset creation is shown in Figure 1(a).

3.2 Selection of Activities

We proposed three methods to identify the biased activities, all utilizing Masked Language Modeling (MLM) as part of their approach to detecting biased activities. We utilized the 'RoBERTa-large' [20] model for its superior performance in masked token prediction, which was attributed to its objective function. The sentences are formatted as "a <mask> is doing <activity>".⁵ We then calculate the probabilities associated with "man" and "woman" (as well as related masculine and feminine identifiers like "boy" and "girl"), from which we calculate the log ratios as follows.

$$\text{bias}(\text{sentence}) = \log \left(\frac{P(\text{mask} = \text{man} | \text{sentence})}{P(\text{mask} = \text{woman} | \text{sentence})} \right) \quad (1)$$

3.2.1 Stereotyped Activities from GPT

Our first method involved using GPT-4 to request activities stereotypically associated with men or women in society. We provided over 100 activities from GPT-4 and formed sentences by replacing the <activity> token in the template of 'a <mask> is doing <activity>'. We then input these sentences into the RoBERTa model and calculated the log ratio.

We applied this methodology separately for activities stereotypically associated with each gender. For men, this resulted in a normal distribution⁶ with a positive mean ($\mu > 0$), from which we chose the sentences at the rightmost end. Similarly, for women, it also resulted in a normal distribution but with a negative mean ($\mu < 0$), from which we then chose the sentences at the leftmost end.

3.2.2 Everyday Activities from GPT

In our second approach, we used GPT-4 to compile a list exceeding 1,000 sentences related to hobbies, jobs, and everyday activities. We then masked the subjects of the sentences just like in the first method and asked the RoBERTa model to fill the mask. We stored the log ratios of predicted probabilities. The log ratios formed a normal distribution, from which we selected sentences at each end.

3.2.3 Activities from LAION-400M Dataset

In the third method, we used the LAION-400M [26] dataset, which consists of 400 million image-caption pairs. We filtered out the NSFW⁷ pairs and sampled 20 million pairs randomly. We then extracted the verbs, subjects, and objects of all caption sentences using spaCy [14], a natural language processing toolkit. We counted the number of masculine and feminine subjects for each verb-object pair

and calculated their log ratios. This also formed a normal distribution.

We then used the Z-test to determine whether each verb-object pair had a significantly different ratio compared to the mean of all verb-object pairs. This test is appropriate due to the large dataset size of LAION-400M. By checking if differences are statistically significant, we identify biased activities. This method enhances the reliability of our activities by ensuring they are not a result of random variance or outliers. Then, we created sentences with the biased verb-object pairs, masked the subject, and submitted them to the RoBERTa model to fill the mask, conducting the same experiment as in the previous methods to select the biased activities.

3.3 Generation

The generation process consisted of two phases. The first phase involved generating a good prompt that helped the image generation system meet the criteria given below, derived from the base sentence. The second phase involved generating the image from the prompt. The generated images should accurately represent the original activity, look natural and high quality, and be diverse.

We used a Large Language Model (LLM) to generate diverse prompts while also keeping the original activity intact (Details about the utilized LLM are provided in the Supplementary Material [1]). We then created a list of predefined settings like the environment of the house, the skin color of the person, etc. We created a list of different combinations of these settings and encoded them using MiniLM [34] sentence transformer [24]. We then clustered them into K^8 clusters and sampled $\frac{N}{K}$ samples from each cluster to enforce more diversity semantically among N selected prompts. We then generated a set of diverse prompts based on the selected prompts by instructing Llama2 to generate diverse prompts that maintain the original activity while describing the environment and the settings. Using these prompts, the generated images were significantly improved in quality and detail. Among the various image generation options we explored, we ultimately selected DALL-E3 [22] as our final solution for image generation (detailed in the Supplementary Material [1]).

3.4 Image Filtering Methodology

Our image filtering methodology employs a combination of automated metrics and human evaluation to ensure the generated images are of high quality, relevant, and diverse.

3.4.1 Quality Assessment

The quality of generated images was assessed using the Fréchet Inception Distance (FID) [13]. The FID score aims to measure the similarity between the generated images and a set of reference images (in this work, we selected images from the COCO [13] dataset that contain at least one human in the scene). To achieve high-quality images, we initially filtered out images with high Fréchet Inception Distance (FID) scores. Furthermore, for an image to pass our quality assessment, individuals must constitute at least 15% of the detected objects and must have a detectable face, as confirmed by the MTCNN model. The FID score for images in the whole dataset is 11.9. It's worth mentioning that the FID scores for both expected (E1 and E2) and unexpected (U1 and U2) categories of images are

⁸ $K = 6$ which resulted in sufficient diversity in our experiments

⁵ <mask> is the special mask token of the RoBERTa model

⁶ All the tests, including log ratios, formed normal distributions that we tested using the Shapiro-Wilk test for normality

⁷ 'Not Safe For Work'. In this context, it refers specifically to content involving nudity, vulgarity, or violence.

12.1 and 11.7, respectively, signifying an acceptable level of quality across all scenarios.

3.4.2 Diversity Metrics

To ensure image diversity, we employed the Structural Similarity Index (SSIM) [35], the Learned Perceptual Image Patch Similarity (LPIPS) [39], and the coverage [21] metrics.

SSIM compares images on a per-region basis. Lower SSIM indicates greater diversity. Our target mean SSIM threshold across all image pairs is 0.046. It’s also worth mentioning that the SSIM for the expected (E1 and E2) and unexpected (U1 and U2) categories of images are 0.046 and 0.045, respectively. This indicates a satisfactory level of diversity in both image categories.

LPIPS measures the perceptual similarity between two images, using deep network features to closely mimic human visual perception. To indicate greater diversity, we aim for higher LPIPS values, setting the mean target for all pairs at 0.66. Also, the LPIPS scores for the expected (E1 and E2) and unexpected (U1 and U2) categories of images are 0.66 and 0.65, respectively.

In our analysis of the GAB dataset’s diversity, we also report the coverage [21], which measures how well the GAB dataset’s distribution aligns with that of the COCO dataset. This is done by calculating the proportion of real samples from the COCO dataset for which there exists a similar sample in the GAB dataset (in this work, the used samples of COCO contain at least one human in the scene). Using the ViT-base-16-224 [36, 9] model, the coverage scores for the entire dataset, as well as for the expected (E1 and E2) and unexpected (U1 and U2) categories, are 0.71, 0.73, and 0.75, respectively. These scores indicate that the GAB dataset effectively represents the diversity found in the COCO dataset.

3.4.3 Human Feedback

After quantitative evaluations, we performed a qualitative analysis with human feedback. This involved scrutinizing images for alignment with the desired activities and contexts. Images that did not align with the activities or did not depict the intended subjects were removed. Additionally, images with unrealistic or nonsensical elements, and those appearing abnormal to an external observer, were excluded.

Following this step, a list of images that have passed both quantitative and qualitative evaluations is obtained. As a result, the images are diverse, high-quality, and realistic while also meeting our criteria for the generation phase.

4 Experiments

We designed several experiments on our proposed GAB dataset to evaluate Vision-Language Models in text-to-image and image-to-text retrieval tasks, as shown in Figure 1(b). Based on these experiments, we assessed the defined gender-activity binding bias and the ability of VLMs to recognize gender-biased activities. Finally, we evaluated the bias in the text encoders and image encoders of different VLMs separately.

4.1 Setup

4.1.1 Task Definitions

The experiments we have designed to assess the performance of VLMs on the aforementioned aspects are grounded in their zero-shot

performance on both text-to-image and image-to-text retrieval tasks.

More formally, consider the space of images as \mathcal{I} and the space of captions as \mathcal{C} . Also consider a VLM with an image encoder $E_I : \mathcal{I} \rightarrow \mathbb{R}^D$ and a text encoder $E_C : \mathcal{C} \rightarrow \mathbb{R}^D$, that map images and texts respectively to a shared embedding space, which is shown in the middle box in Figure 1(b). We can define a matching score function $s : \mathcal{I} \times \mathcal{C} \rightarrow \mathbb{R}^+$, which measures the similarity between a caption and an image, as the cosine similarity between the embedding of an image I and the embedding of a caption C :

$$s(I, C) = \frac{E_I(I) \cdot E_C(C)}{\|E_I(I)\| \|E_C(C)\|}. \quad (2)$$

Text-to-image retrieval involves finding the image that best matches a given text from a collection of images. Conversely, image-to-text retrieval is about finding the text that best describes a given image. We compute the similarity using the previously defined similarity measure in the upcoming experiments. For the text-to-image retrieval task, we have a caption C and make comparisons between two images: I , which shows an image corresponding to the caption C , and I^R , which corresponds to C^R (obtained by reversing the gender of subjects in C). Similarly, for the image-to-text retrieval task, we have an image I and make comparisons between two captions C and C^R .

4.1.2 Selected Vision-Language Models

In our experiments, we benchmarked 12 VLMs and presented the results of various tests. To compare the effects of patch size and backbone size, we reported the results for 4 CLIP models released by OpenAI [23], each with different backbones and patch sizes.

Another model selected for evaluation is NegCLIP [38], a version of the CLIP-ViT-B-32 model by OpenAI [23] which is fine-tuned on a modified subset of the Visual Genome Dataset [17] annotations. NegCLIP is fine-tuned to purportedly enhance the base model’s performance on tasks requiring relational and activity understanding.

We also included results for several recently proposed VLMs, including Eva01 [10] and Eva02 [31], FLAVA [28], ALIGN [16], COCA [37], and AltCLIP [7].

4.2 Results

4.2.1 Image-to-Text Retrieval

We assess the performance of the models discussed in Section 4.1.2 on the GAB dataset and calculate their accuracy in retrieving the correct caption as outlined in Section 4.1. We employ a caption template of the form: “*a <man/woman> is <doing activity> and a <woman/man> is in the scene*”. The results for activities that are stereotypically biased are presented in Figure 2, while the results for the other two groups of activities are provided in the Supplementary Material [1].

In this experiment, we assessed image-to-text retrieval across three distinct scenarios. The first scenario involves images where an individual of an unexpected gender is performing a typically biased activity with no other individuals present in the scene. The second scenario comprises images where both genders are present, and the activity is performed by the gender typically associated with it. The third scenario mirrors the second, but in this case, the activity is performed by the gender not typically associated with it.

Performance Drop Due to Presence of Expected Gender: As depicted in Figure 2, there is a noticeable drop in the accuracy of the




Figure 2. Average retrieval accuracy of VLMs on the image-to-text retrieval task across various scenarios. The chart highlights the performance drop between these scenarios for each model. The purple bar represents the accuracy in the scenario where the unexpected gender is performing the activity in the reference image, while the expected gender is also present in the scene. The red bar corresponds to the accuracy in the reverse scenario, where the expected gender is performing the activity. The blue bar denotes the scenario where the unexpected gender is performing the activity and is the only one present in the scene.




Figure 3. Average retrieval accuracy of VLMs on the image-to-text retrieval task for the E1 and U1 categories of images on stereotypical activities. The purple (left-hand) bar represents the accuracy in the scenario where the unexpected gender is performing the activity in the reference image and is alone in the scene. The blue (right-hand) bar denotes the scenario where the expected gender is performing the activity and is the only one present in the scene. In this task, the template is “*a <man/woman> is <doing activity>*”.

models when the scene includes two genders compared to scenarios where only the unexpected gender is present and performing the activity associated with bias. The chart reveals that VLMs experience an average performance decline of approximately 33.2%. However, a few models demonstrate only a slight decrease in performance.

Performance Drop Due to Gender-Activity Binding Bias: In the scenarios where there are two individuals present at the scene, the majority of models display a substantial decline in retrieval accuracy when the action is performed by an individual of the unexpected gender. As indicated in Figure 2, the VLMs undergo an average accuracy reduction of approximately 13.2% when encountering gender activity binding bias.

Note that, as illustrated in Figure 3, the models demonstrate satisfactory performance in the Image-to-Text retrieval task on both image scenarios of U1 and E1, indicating that the gender bias adversely affects the performance of the models only when there is more than one individual in the image. Additionally, the accuracy of the models remains consistent, regardless of whether the expected or unexpected

gender is performing the activity. The models’ average performance in both scenarios is approximately 80%.

4.2.2 Text Encoder Bias

Given an activity a , consider $e, u \in \{\text{man, woman}\}$ where e is the expected gender of the individual performing a and u is the unexpected one. To separately assess the bias of the text encoder, we determine the frequency at which the embedding of a gender-neutral sentence “a person is <doing activity>” is closer to the embedding of “a e is <doing activity>” than the embedding of “a u is <doing activity>”. The outcomes are presented in Table 1. It can be observed that in all categories of the gathered activities, the majority of activities exhibit a bias towards the expected gender.

Table 1. Bias of the text encoder for each model. Taking into account the template “a <mask> is <doing activity>”, the table shows the proportion of activities where substituting <mask> with the gender-neutral term “person” results in a higher matching score (Eq. 2) in the embedding space compared to the instance in which <mask> is replaced by the expected gender reference (woman/man) compared to the unexpected one. The columns denote the collected activities across different categories as outlined in Section 3.2.

Model	LAION-400M Biased Activities	GPT Gathered Biased Activities	Biased Stereotype Activities
AltCLIP	0.80	0.73	0.91
EVA01-g-14	0.80	0.80	0.92
EVA02-L-14	0.80	0.80	0.92
CLIP-RN50x64	0.50	0.60	0.59
CLIP-ViT-B-16	0.70	0.60	0.83
NegCLIP-ViT-B-32	0.70	0.60	0.75
CLIP-ViT-B-32	0.60	0.66	0.59
CLIP-ViT-L-14	0.60	0.73	0.83
COCA-ViT-B-32	0.90	0.93	0.92
COCA-ViT-L-14	1.00	0.93	0.92
Flava	0.70	0.53	0.75
Align	0.90	1.00	1.00

4.2.3 Text-to-Image Retrieval

Figure 4 illustrates the accuracy of VLMs in the task of text-to-image retrieval. For each activity, we form random pairs of images from E2 and U2⁹. Each model is then evaluated for its ability to retrieve the image that best matches captions formatted as “a <man/woman> is <doing activity>” and a “<woman/man> is in the scene.” As can be seen in Figure 4, VLMs achieve an accuracy of approximately 50%, indicating that their performance is nearly random in this task. In essence, the benchmarked models do not seem to comprehend the properties of the given images that would aid in retrieving the correct one given embedded information from the caption, i.e., they fail to recognize the performer of the activity based on the text.

4.2.4 Activity Recognition

To investigate whether gender bias in associating certain activities with specific genders arises from a misunderstanding of the activities themselves, we evaluated the model’s comprehension of these activities. We collected several alternative activities performed in the same environment as the reference activity for each entry in the GAB dataset. For each image, we generated captions based on both the reference activity and the collected alternative activities. These captions were then ranked by their matching scores with the image.

⁹ Pairing is in the form of one-to-one correspondence.




Figure 4. Average accuracy of models in the text-to-image retrieval tasks. The images are sourced from the E2 and U2 groups, which are images that feature two individuals of different genders. The performance of models on the U2 and E2 groups is represented by the purple (left-hand) and blue (right-hand) bars, respectively. Additionally, the red (upper) and black (lower) dashed lines depict the average performance for the U2 and E2 groups, respectively.

We report the mean reciprocal rank (MRR), recall@1, and recall@3 for the true caption (reference activity) in the context of everyday activities (Section 3.2.2) in Table 2 (see Supplementary Material [1] for results for other sets of activities). The results indicate that the model is capable, to some extent, of correctly identifying the true activity from a list of alternatives, suggesting that VLMs can recognize the performed activity in image-to-text retrieval.

Table 2. Models’ ability to retrieve the correct activity from a list of activities. For every image, we formulate a collection of captions derived from a list of activities and arrange them according to their matching score (Eq. 2) with the image. We report Mean Reciprocal Rank (MRR), recall@1, and recall@3, which indicate that VLMs can to some degree recognize the correct captions in image-to-text retrieval tasks.

Model	MRR		Recall@1		Recall@3	
	One Gender Presence	Two Genders Presence	One Gender Presence	Two Genders Presence	One Gender Presence	Two Genders Presence
ALIGN	0.677	0.593	0.553	0.449	0.731	0.659
EVA01-g-14	0.585	0.512	0.425	0.363	0.672	0.564
EVA02-L-14	0.688	0.563	0.564	0.415	0.760	0.635
CLIP-RN50x64	0.601	0.495	0.465	0.334	0.648	0.557
CLIP-ViT-B-16	0.483	0.456	0.333	0.310	0.531	0.486
NegCLIP-ViT-B-32	0.347	0.332	0.159	0.153	0.414	0.385
CLIP-ViT-B-32	0.296	0.269	0.128	0.110	0.316	0.274
CLIP-ViT-L-14	0.606	0.550	0.469	0.404	0.672	0.618
COCA-ViT-B-32	0.319	0.292	0.137	0.128	0.382	0.316
COCA-ViT-L-14	0.459	0.418	0.280	0.243	0.529	0.489
Flava	0.474	0.469	0.283	0.282	0.554	0.538
ALIGN	0.534	0.497	0.375	0.331	0.620	0.589

5 Discussion

5.1 Image and Text Embedding in VLMs

To gain a deeper understanding of how texts and images are encoded into the shared embedding space, consider these points:

1. Given that the image retrieval accuracy is nearly 50% for both expected and unexpected caption groups (as shown in Figure 4), we can infer that the cosine similarity (as per Eq. 2) between both sets of captions and images (for expected and unexpected groups) is almost identical.
2. As the average text retrieval accuracy exceeds 60% for E2 and falls below 50% for U2, it suggests that text embeddings are somewhat closer to the image embeddings in E2.

A plausible interpretation of these observations is that VLMs encode images more closely to the encoding of the more expected text. However, the image encoders of VLMs embed images from both groups

closely together, resulting in nearly equal cosine similarity with any caption embedding.

5.2 Bias Mitigation

Multiple works have previously discussed approaches for mitigating bias from VLMs, such as orthogonal projection [8], making unbiased datasets [19, 15], representation correction [27, 33], and prompt tuning [6]. Several of these approaches could be applied also for robustness to gender-activity binding bias. For more detail on these works, refer to the Supplementary Material [1].

5.3 Future Works

Study Other Social Biases. In this work, we focused solely on gender bias; however, the experimental approaches described can be extended to other social biases, such as race and age. In future research, we plan to explore how VLMs perform across these dimensions using the experiments outlined in this study.

Study the Source of Bias in the Training Data. Due to limited accessibility to the training sets, especially for models with unpublished datasets, this study could not investigate the source of bias in the training data directly. We plan to address this limitation in future work by examining the sources of bias in publicly available training datasets for VLMs and analyzing the origins of these biases within those datasets.

Exploring the Influence of More Attributes on Image Retrieval Accuracy Our image generation prompts took into account various factors such as clothing type (casual wear, formal attire, traditional attire, activewear, and business casual) and age (teenager, young adult, adult, middle-aged, and senior). However, it remains unexplored whether specific attributes, like gendered clothing or pregnancy, have a greater impact on retrieval accuracy than the physical characteristics of the individuals depicted. We recognize the potential to investigate these factors separately in the future to understand their influence on activity binding by VLMs.

6 Conclusions

In this study, we have explored a significant yet overlooked bias in VLMs known as the gender-activity binding bias. We used a unique dataset named GAB, which includes about 5500 images. We found that the capacity of VLMs to link an activity with the gender of its performer notably decreases when another person of a different gender is in the scene. We discovered that while the bias is evident in image-to-text retrieval tasks and in the text encoder alone, the models perform randomly in text-to-image retrieval tasks. This suggests that the gender-activity binding bias is primarily absorbed by the text encoder rather than the image encoder in VLMs. Furthermore, we noted that while VLMs have difficulty with gender-activity binding, they do have some ability to recognize activities. We believe that the ability of VLMs to comprehend activities and perform compositional reasoning in complex scenes are other crucial factors contributing to the gender-activity binding bias that could be investigated in future research.

Acknowledgements

We would like to express our sincere gratitude to Ms. Fahimeh Hosseini Noohdani and Mr. Feraiidoon Mehri for their invaluable assistance in completing this paper, as well as their constructive and insightful comments.

References

- [1] A. Abdollahi, M. Ghaznavi, M. R. Karimi Nejad, A. Mari Oriyad, R. Abbasi, A. Salehi, M. Behjati, M. H. Rohban, and M. Soleymani Baghshah. Gabinsight: Exploring gender-activity binding bias in vision-language models. *ArXiv*, abs/2407.21001, 2024.
- [2] A. Abdollahi, M. Ghaznavi, M. R. Karimi Nejad, A. Mari Oriyad, R. Abbasi, A. Salehi, M. Behjati, M. H. Rohban, and M. Soleymani Baghshah. Gabinsight: Exploring gender-activity binding bias in vision-language models-code and data, 2024. URL <https://github.com/sharif-ml-lab/GABInsight>.
- [3] S. Agarwal, G. Krueger, J. Clark, A. Radford, J. W. Kim, and M. Brundage. Evaluating clip: Towards characterization of broader capabilities and downstream implications. *ArXiv*, abs/2108.02818, 2021.
- [4] J. Aneja, A. Deshpande, and A. G. Schwing. Convolutional image captioning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5561–5570, 2018.
- [5] E. Bareinboim, J. Tian, and J. Pearl. Recovering from selection bias in causal and statistical inference. *Probabilistic and Causal Inference*, 2014.
- [6] H. Berg, S. Hall, Y. Bhalgat, H. Kirk, A. Shtedritski, and M. Bain. A prompt array keeps the bias away: Debiasing vision-language models with adversarial learning. In *Proceedings of the 2nd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 12th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 806–822, Online only, 2022. Association for Computational Linguistics.
- [7] Z. Chen, G. Liu, B.-W. Zhang, Q. Yang, and L. Wu. AltCLIP: Altering the language encoder in CLIP for extended language capabilities. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 8666–8682, 2023.
- [8] C.-Y. Chuang, V. Jampani, Y. Li, A. Torralba, and S. Jegelka. Debiasing vision-language models via biased prompts. *ArXiv*, 2302.00070, 2023.
- [9] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.
- [10] Y. Fang, W. Wang, B. Xie, Q. Sun, L. Wu, X. Wang, T. Huang, X. Wang, and Y. Cao. Eva: Exploring the limits of masked visual representation learning at scale. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19358–19369, 2023.
- [11] Z. Gan, L. Li, C. Li, L. Wang, Z. Liu, J. Gao, et al. Vision-language pre-training: Basics, recent advances, and future trends. *Foundations and Trends® in Computer Graphics and Vision*, 14(3–4):163–352, 2022.
- [12] S. M. Hall, F. Gonçalves Abrantes, H. Zhu, G. Sodunke, A. Shtedritski, and H. R. Kirk. Visogender: A dataset for benchmarking gender bias in image-text pronoun resolution. *Advances in Neural Information Processing Systems*, 36, 2024.
- [13] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, and S. Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017.
- [14] M. Honnibal, I. Montani, S. Van Landeghem, and A. Boyd. spaCy: Industrial-strength Natural Language Processing in Python. 2020.
- [15] P. Howard, A. Madasu, T. Le, G. L. Moreno, A. Bhiwandiwala, and V. Lal. Probing and mitigating intersectional social biases in vision-language models with counterfactual examples. *CVPR*, 2024.
- [16] C. Jia, Y. Yang, Y. Xia, Y.-T. Chen, Z. Parekh, H. Pham, Q. Le, Y.-H. Sung, Z. Li, and T. Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 4904–4916. PMLR, 2021.
- [17] R. Krishna, Y. Zhu, O. Groth, J. Johnson, K. Hata, J. Kravitz, S. Chen, Y. Kalantidis, L.-J. Li, D. A. Shamma, et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International journal of computer vision*, 123:32–73, 2017.
- [18] N. Lee, Y. Bang, H. Lovenia, S. Cahyawijaya, W. Dai, and P. Fung. Survey of social bias in vision-language models. *CoRR*, abs/2309.14381, 2023.
- [19] Y. Li and N. Vasconcelos. Debias your VLM with counterfactuals: A unified approach, 2024. URL <https://openreview.net/forum?id=xx05gm7oQw>.
- [20] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov. Roberta: A robustly optimized bert pretraining approach. *ArXiv*, abs/1907.11692, 2019.
- [21] M. F. Naeem, S. J. Oh, Y. Uh, Y. Choi, and J. Yoo. Reliable fidelity and diversity metrics for generative models. In H. D. III and A. Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 7176–7185. PMLR, 13–18 Jul 2020.
- [22] OpenAI. Dall-e 3. <https://openai.com/dall-e-3/>, 2023.
- [23] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, G. Krueger, and I. Sutskever. Learning transferable visual models from natural language supervision. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 8748–8763. PMLR, 2021.
- [24] N. Reimers and I. Gurevych. Sentence-bert: Sentence embeddings using siamese bert-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 2019.
- [25] E. Salin, S. Ayache, and B. Favre. Towards an exhaustive evaluation of vision-language foundation models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops*, pages 339–352, 2023.
- [26] C. Schuhmann, R. Vencu, R. Beaumont, R. Kaczmarczyk, C. Mullis, A. Katta, T. Coombes, J. Jitsev, and A. Komatsuzaki. Laion-400m: Open dataset of clip-filtered 400 million image-text pairs. *ArXiv*, abs/2111.02114, 2021.
- [27] A. Seth, M. Hemani, and C. Agarwal. Dear: Debiasing vision-language models with additive residuals. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6820–6829, 2023.
- [28] A. Singh, R. Hu, V. Goswami, G. Couairon, W. Galuba, M. Rohrbach, and D. Kiela. Flava: A foundational language and vision alignment model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15638–15650, 2022.
- [29] T. Srinivasan and Y. Bisk. Worst of both worlds: Biases compound in pre-trained vision-and-language models. In *Proceedings of the 4th Workshop on Gender Bias in Natural Language Processing (GeBNLP)*, pages 77–85. Association for Computational Linguistics, 2022.
- [30] T. Srinivasan and Y. Bisk. Worst of both worlds: Biases compound in pre-trained vision-and-language models. In *Proceedings of the 4th Workshop on Gender Bias in Natural Language Processing (GeBNLP)*, pages 77–85, 2022.
- [31] Q. Sun, Y. Fang, L. Wu, X. Wang, and Y. Cao. Eva-clip: Improved training techniques for clip at scale. *ArXiv*, arXiv:2303.15389, 2023.
- [32] T. Thrush, R. Jiang, M. Bartolo, A. Singh, A. Williams, D. Kiela, and C. Ross. Winoground: Probing vision and language models for visiolinguistic compositionality. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5228–5238, 2022.
- [33] J. Wang, Y. Liu, and X. E. Wang. Are gender-neutral queries really gender-neutral? mitigating gender bias in image search. *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, page 1995–2008, 2021.
- [34] W. Wang, F. Wei, L. Dong, H. Bao, N. Yang, and M. Zhou. Minilm: Deep self-attention distillation for task-agnostic compression of pre-trained transformers. *ArXiv*, abs/2002.10957, 2020.
- [35] Z. Wang, A. Bovik, H. Sheikh, and E. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, 2004.
- [36] B. Wu, C. Xu, X. Dai, A. Wan, P. Zhang, Z. Yan, M. Tomizuka, J. Gonzalez, K. Keutzer, and P. Vajda. Visual transformers: Token-based image representation and processing for computer vision. *ArXiv*, abs/2006.03677, 2020.
- [37] J. Yu, Z. Wang, V. Vasudevan, L. Yeung, M. Seyedhosseini, and Y. Wu. Coca: Contrastive captioners are image-text foundation models. *Trans. Mach. Learn. Res.*, 2022, 2022.
- [38] M. Yuksekgonul, F. Bianchi, P. Kalluri, D. Jurafsky, and J. Zou. When and why vision-language models behave like bags-of-words, and what to do about it? In *The Eleventh International Conference on Learning Representations*, 2023.
- [39] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 586–595, 2018. doi: 10.1109/CVPR.2018.00068.
- [40] Y. Zhang, J. Wang, and J. Sang. Counterfactually measuring and eliminating social bias in vision-language pre-training models. In *Proceedings of the 30th ACM International Conference on Multimedia*, MM '22, page 4996–5004, New York, NY, USA, 2022. Association for Computing Machinery. ISBN 9781450392037. doi: 10.1145/3503161.3548396.
- [41] X. Zhou, M. Liu, E. Yurtsever, B. L. Zagar, W. Zimmer, H. Cao, and A. C. Knoll. Vision language models in autonomous driving: A survey and outlook. *IEEE Transactions on Intelligent Vehicles*, pages 1–20, 2024. doi: 10.1109/TIV.2024.3402136.