

MASS: Overcoming Language Bias in Image-Text Matching

Jiwan Chung, Seungwon Lim Sangkyu Lee, Youngjae Yu

Yonsei University, 50 Yonsei-ro, Seodaemun-gu, Seoul, South Korea
jiwan.chung.research@gmail.com

Abstract

Pretrained visual-language models have made significant advancements in multimodal tasks, including image-text retrieval. However, a major challenge in image-text matching lies in language bias, where models predominantly rely on language priors and neglect to adequately consider the visual content. We thus present Multimodal Association Score (MASS), a framework that reduces the reliance on language priors for better visual accuracy in image-text matching problems. It can be seamlessly incorporated into existing visual-language models without necessitating additional training. Our experiments have shown that MASS effectively lessens language bias without losing an understanding of linguistic compositionality. Overall, MASS offers a promising solution for enhancing image-text matching performance in visual-language models.

code — https://github.com/JiwanChung/mass_aaai

1 Introduction

Recently, pretrained visual-language models have showcased impressive performance across a broad spectrum of multimodal tasks (Li et al. 2022; Alayrac et al. 2022; Wang et al. 2022b), including simple image-text retrieval (Chen et al. 2015; Young et al. 2014). In particular, CLIP (Radford et al. 2021) has emerged as one of the most popular models for image-text matching. The image-text similarity score derived from CLIP is shown to be effective for assessing image caption quality (Hessel et al. 2021), providing reward signals to train multimodal models (Yu et al. 2023), and building a large-scale image search system (Beaumont 2022).

However, CLIP falters when the given task requires accurate language modeling. Recent benchmarks have unveiled CLIP’s deficiency in modeling linguistic compositionality (Thrush et al. 2022; Nikolaus et al. 2022; Yuksekgonul et al. 2022). Furthermore, current visual-language models also have demonstrated inconsistent outcomes on image-text matching that demand comprehension of linguistic constructs, such as the existence and quantity of objects or coreference (Parcalabescu et al. 2022; Shekhar et al. 2017).

As an alternative to the contrastive objective in CLIP, recent studies (Petryk et al. 2024; Tschannen et al. 2024) revealed that log-likelihood induced from the autoregressive



Figure 1: Captions retrieved with each method given the image, where only MASS succeeds in ruling out the failure modes. Models trained Image-Text Contrastive (ITC) objectives such as CLIP (Radford et al. 2021) fail to model linguistic structure. Token Likelihood (TL) of image captioning models including OFA (Wang et al. 2022b) shows overreliance on its language prior. Our MASS amends the language bias of image captioning models for accurate image-text matching capability.

image captioning objective features a better understanding of linguistic structures, including compositionality. The objective fosters language generation capability in the trained models, inducing stronger linguistic understanding.

Still, there is a caveat in directly using the image captioning models to assess image-text similarity. While training for the captioning objective, the models inherently build a prior for language distribution on what text sequence is more likely *regardless of the given image*. As shown in Figure 1, overreliance on this prior leads to weaker visual conditioning and then to incorrect image-text matching. Drawing inspiration from previous research (Niu et al. 2020), we refer to such phenomenon as *language bias*. Language bias refers to the propensity of visual-language models to rely heavily on language priors in the training data instead of properly conditioning their output on the given images. It has been repeatedly reported that language bias is a major bottleneck in visual question-answering (Agrawal, Batra, and Parikh 2016; Zhang et al. 2016; Agrawal et al. 2018; Si et al. 2022), and other vision-language tasks (Srinivasan and Bisk 2021; Salin et al. 2022).

In this paper, we present Multimodal Association Score (MASS) as an inference-time framework designed to reduce language bias in image-text matching. MASS aims to measure the strength of the association between image and text modality without involving the textual prior. Our approach




Figure 2: (a) Given an image of a girl playing tennis, the visual-language model falsely retrieves captions describing the subject as male by relying on language bias. (b) On the other hand, MASS reduces such gender bias by adopting pointwise mutual information which suppresses the text-only marginal likelihood. We provide the corresponding experimental results in section 4.3.

uniquely functions as follows: Given a pretrained visual-language model, we first extract the image-conditional and text-only likelihood per text token. We then compute the pointwise mutual information of each image and text token and aggregate them for the text sequence to create a debiased scalar similarity score. Importantly, MASS can be computed for any off-the-shelf visual-language model that outputs image-conditioned text likelihood. Also, it does not require any additional training.

We first demonstrate that MASS can mitigate language bias in visual-language models by evaluating it with multimodal color (section 4.1), number (section 4.2), and gender (section 4.3) bias benchmarks. Our results indicate that MASS significantly improves performance over CLIP (Radford et al. 2021) and raw token likelihood (TL) in the datasets with clear language bias (color), explicit bias in labels (number), and practical image search problems (gender) about reliance on language bias. We further evaluate MASS’s capability to maintain linguistic understanding against two visual-linguistic compositionality benchmarks; Winoground (Thrush et al. 2022) and SVO-Probe (Hendricks and Nematzadeh 2021). MASS greatly enhances the performance of the backbone visual-language model in both experiments and outperforms the strong baselines.

The main contributions of this paper are threefold. Firstly, we present MASS as an effective framework for language bias reduction in image-text matching tasks. Secondly, we demonstrate that MASS can both reduce language bias and maintain linguistic understanding with various biases and compositionality benchmarks. Lastly, our approach does not necessitate additional training and can be directly applied to a wide variety of visual-language models for robust image-text matching.

2 Image-Text Similarity Functions

Our objective is to identify a similarity function that accurately pairs image and text, even in the presence of language bias in the pretrained model. Given an image $\mathbf{c} \in \mathcal{C}$ and a text $\mathbf{x} \in \mathcal{X}$, the similarity function $\mathcal{S}(\mathbf{c}, \mathbf{x})$ evaluates the closeness of alignment such that parallel image-text pairs have high scores. We divide the similarity functions by the granularity of the training objectives: Image-Text Con-

trastive Learning (ITC) and Image-Text Matching (ITM) provide scores on *sequence-level* (section 2.1), while Token-Level supervision (TL) such as supervised image captioning applies *token-level* (section 2.2) feedback.

2.1 Sequence-Level Similarity Functions

The sequence-level similarity functions are defined on top of dual vision-language encoders. The dual encoders may be either *separate* (i.e., architectures with independent encoders per modality as CLIP (Radford et al. 2021) and the lower modules of ALBEF (Li et al. 2021)) or *intertwined* (e.g., the upper module of ALBEF) in terms of architecture.

Image-Text Contrastive Learning. ITC is a training objective for learning multimodal joint embedding space based on the Noise Contrastive Estimation (Oord, Li, and Vinyals 2018). We use *separate* image and text encoders f_ϕ, g_ψ as often used with the ITC objective (Radford et al. 2021).

ITC requires the aligned image and text pair (\mathbf{c}, \mathbf{x}) and a set of misaligned text $\tilde{\mathbf{x}}$ typically sourced from other text within the minibatch. ITC first obtains vector embeddings of each image and text and computes the cosine similarity of them to build image-text pair logits, which are then used to calculate the cross-entropy loss. We use the term ITC both to denote the training objective and corresponding similarity function \mathcal{S}_{ITC} induced from the trained model with parameters $\bar{\psi}$ and $\bar{\phi}$. In experiments, we include CLIP (Radford et al. 2021) as the baseline similarity function of this type.

$$\mathcal{S}_{\text{ITC}}(\mathbf{c}, \mathbf{x}) := \frac{f_{\bar{\phi}}(\mathbf{c})^T \cdot g_{\bar{\psi}}(\mathbf{x})}{\|f_{\bar{\phi}}(\mathbf{c})\| \cdot \|g_{\bar{\psi}}(\mathbf{x})\|} \quad (1)$$

Image-Text Matching. ITM is a binary classification objective that decides whether a pair of image and text (\mathbf{c}, \mathbf{x}) is aligned or not (Li et al. 2021; Wang et al. 2022a,b; Li et al. 2022). While both ITM and ITC require negative samples for training, ITC typically compares a positive sample with a large number of negatives, while ITM determines the correctness of a single sample. In practice, ITM is often used to train *intertwined* models (Li et al. 2021, 2022), which we adopt for explanation here.

For each pair of images and text (\mathbf{c}, \mathbf{x}) , an *intertwined* image-text encoder $f_{\phi, \psi}$ processes both to yield a multi-

modal vector representation. Then, a linear classifier h_ω is applied on top of the representation to build a binomial logit for the cross-entropy loss. At inference time, the probability of the image-text pair being true is used as the similarity function output $\mathcal{S}_{\text{ITM}}(\mathbf{c}, \mathbf{x})$. In our setting, all baselines marked ITM fall into this category. Note that OFA uses the text generation head to obtain the ITM score instead of a separate linear classifier. Refer to the corresponding paper (Wang et al. 2022b) for details.

$$\mathcal{S}_{\text{ITM}}(\mathbf{c}, \mathbf{x}) = \frac{\exp(h_\omega(f_{\bar{\phi}, \bar{\psi}}(\mathbf{c}, \mathbf{x})))}{1 + \exp(h_\omega(f_{\bar{\phi}, \bar{\psi}}(\mathbf{c}, \mathbf{x})))} \quad (2)$$

2.2 Token-Level Similarity Functions

Token-Level Likelihood. Let us divide a sample of text into a sequence of *tokens* $\mathbf{x} = \{x_1, \dots, x_l\}$ of length l . The token-level functions require models to output the image-text alignment score for each text token. A popular training objective that admits such constraint is image captioning. Given a pair of image-text (\mathbf{c}, \mathbf{x}) , an image captioning model with parameter θ is trained to maximize the log-likelihoods $\log p_\theta$ of the text sequence \mathbf{x} factorized in an autoregressive manner. With trained parameter $\bar{\theta}$, the model likelihood for the probability of the next token conditioned on both the image and previously generated text $p_{\bar{\theta}}(x_t|x_{<t}, \mathbf{c})$ can be treated as an approximation to the true conditional likelihood $p(x_t|x_{<t}, \mathbf{c})$.

Aggregation. The token-level likelihood of an image captioning model can be directly used as the token-level similarity function (Petryk et al. 2024). For all token-level similarity scores, we compute the mean of all text token scores in the sequence to build the sequence-level image-text alignment metric. We denote this type of similarity function with TL in the experiments.

$$\mathcal{S}_{\text{TL}}(\mathbf{c}, \mathbf{x}) := \frac{1}{l} \sum_{t < l} p_{\bar{\theta}}(x_t|x_{<t}, \mathbf{c}) \quad (3)$$

3 Multimodal ASsociation Score

Directly using likelihood from autoregressive visual-language models as an image-text similarity function is problematic since this does not represent *pure* image-text similarity. For instance, image captioning models are also required to generate *linguistically plausible* captions. Whether the token x_t matches the given image \mathbf{c} or not, the model can assign a high likelihood if it was exposed frequently with the prefix $x_{<t}$ in the training process. To address this ignorance of the context, *pointwise mutual information* $\text{PMI}(\mathbf{x}; \mathbf{c})$ has been utilized in NLP tasks such as the conversation model (Li et al. 2015) and text summarization (Van der Poel, Cotterell, and Meister 2022). However, its application to the visual-language model has yet to be investigated, and we propose it as an effective method for reducing language bias in image-text matching.

$$\text{PMI}(\mathbf{x}; \mathbf{c}) := \log \frac{p(\mathbf{x}, \mathbf{c})}{p(\mathbf{x})p(\mathbf{c})} = \log \frac{p(\mathbf{x}|\mathbf{c})}{p(\mathbf{x})} \quad (4)$$

3.1 Language Bias Reduction

Formally, consider the log-likelihood of a given image and text pair (\mathbf{c}, \mathbf{x}) from trained parameter $\bar{\theta}$, processed in an autoregressive manner. According to Bayes' rule, we can decompose the log-likelihood to two distinct terms based on whether image \mathbf{c} is related to the text \mathbf{x} or not. We label the term not related to the \mathbf{c} as the *linguistic* log-likelihood and the term that is related to the \mathbf{c} as the *association* log-likelihood.

$$\log p_{\bar{\theta}}(\mathbf{x}|\mathbf{c}) = \underbrace{\log p_{\bar{\theta}}(\mathbf{x})}_{\text{linguistic}} + \underbrace{\log \frac{p_{\bar{\theta}}(\mathbf{c}|\mathbf{x})}{p_{\bar{\theta}}(\mathbf{c})}}_{\text{association}} \quad (5)$$

Under this perspective, the linguistic log-likelihood represents the linguistic plausibility of text \mathbf{x} . While the term is independent of the image context, it can overrule the similarity between image \mathbf{c} and text \mathbf{x} measured by the *association* term when the model relies on text bias. As such, a simple remedy for the likelihood-based similarity functions would be to regularize the log-likelihood by subtracting the linguistic log-likelihood. For the token-level similarity function, we can deduct the *token-level* linguistic log-likelihood $\log p_{\bar{\theta}}(x_t|x_{<t})$ from token-level log-likelihood $\log p_{\bar{\theta}}(x_t|x_{<t}, \mathbf{c})$ for each token x_t .

It is noteworthy that if $p_{\bar{\theta}}(\mathbf{x})$ and $p_{\bar{\theta}}(\mathbf{x}|\mathbf{c})$ are accurately estimating the ground-truth marginal density $p(\mathbf{x})$ and conditional density $p(\mathbf{x}|\mathbf{c})$, this approach equates to estimating the pointwise mutual information. In this sense, we denote this similarity function as Multimodal ASsociation Score (MASS), which is calculated by averaging each token's pointwise mutual information over the total length l .

$$\mathcal{S}_{\text{MASS}}(\mathbf{c}, \mathbf{x}) := \frac{1}{l} \sum_{t < l} \log \frac{p_{\bar{\theta}}(x_t|x_{<t}, \mathbf{c})}{p_{\bar{\theta}}(x_t|x_{<t})} \quad (6)$$

3.2 Estimating the Marginal Likelihood

The autoregressive visual-language model inherently assumes an image input \mathbf{c} , as it is trained to estimate the *conditional* log-likelihood $\log p(\mathbf{x}|\mathbf{c})$. As a result, we cannot directly obtain the token-level marginal log-likelihood $\log p_{\bar{\theta}}(x_t|x_{<t})$, which is required to calculate the MASS score. One possible solution is to use a Monte Carlo approximation by uniformly sampling N random image $\tilde{\mathbf{c}}$. For each token x_t , we can approximate the marginal log-likelihood $\log p_{\bar{\theta}}(x_t|x_{<t})$ by taking the average of the conditional log-likelihood $\log p_{\bar{\theta}}(x_t|x_{<t}, \tilde{\mathbf{c}})$ over random images $\tilde{\mathbf{c}}_i$.


$$\log p_{\bar{\theta}}(x_t|x_{<t}) = \int_{\mathbf{c}} \log p_{\bar{\theta}}(x_t|x_{<t}, \mathbf{c}) d\mathbf{c} \quad (7)$$

$$\approx \frac{1}{N} \sum_{i=1}^N \log p_{\bar{\theta}}(x_t|x_{<t}, \tilde{\mathbf{c}}_i) \quad (8)$$

However, this approach requires N more forward passes to compute the score for one image-text pair, and it is also challenging to ensure the accuracy of the approximation. As an alternative, we discovered that using image input as a *null image* \mathbf{c}_\emptyset , which is a 'black-filled' image, is a good alternative for estimating $\log p_{\bar{\theta}}(x_t|x_{<t})$. For the rest of the



(a) Top: A sample from our color debiasing experiment. A model without language bias would say the tomato is gray. Bottom: A sample from the counting benchmark.



(b) Results on Natural Colors Dataset (Anwar et al. 2020). MASS effectively debiases TL both in sample and type-level.

Figure 3: Data samples and experimental results from our color debiasing experiment.

paper, we use $\log p_{\bar{\theta}}(x_t|x_{<t}, \mathbf{c}_{\theta})$ as the approximation of $\log p_{\bar{\theta}}(x_t|x_{<t})$ to calculate MASS.

$$S_{\text{MASS}}(\mathbf{c}, \mathbf{x}) \approx \frac{1}{l} \sum_{t < l} \log \frac{p_{\bar{\theta}}(x_t|x_{<t}, \mathbf{c})}{p_{\bar{\theta}}(x_t|x_{<t}, \mathbf{c}_{\theta})} \quad (9)$$

4 Experiments: Language Debiasing

We show that MASS effectively reduces language bias in three different domains; natural color association (section 4.1), object number counting (section 4.2), and gender balance in image search (section 4.3). The improvement is especially substantial when the tested data have a different language bias from the training data of the base visual-language model.

Backbones. We use OFA (Wang et al. 2022b) as the backbone VL model for both ITM, TL, and MASS similarity functions. Also, we experiment with alternative backbones of BLIP-2 (Li et al. 2023) and LLaVA (Liu et al. 2024). Note that ITM is not applicable for LLaVA since it only has the language modeling head.

4.1 Debiasing Colors

When there is linguistic bias in the data, the visual-language models trained on them would likely reflect the

bias (Birhane, Prabhu, and Kahembwe 2021; Ross, Katz, and Barbu 2021). However, evaluating the model for bias is not a simple problem: how do we know when the model is not relying on the text context without looking at the image?

Here, we devise a toy experiment where the model should perform perfectly if it truly looks at the image. The Natural Color Dataset (NCD) (Anwar et al. 2020) is a dataset of various fruits colored either in the natural color or in gray. Given a gray image of a fruit, we asked the model if the color is gray. Surprisingly, the token likelihood score said no for many fruits and vegetables.

Approach and metrics. We use only the grayscale images from the NCD dataset, which leaves 723 images consisting of 20 different fruits. Given an image, we build a true caption "The fruit is gray." and an adversarial caption "The fruit is color.", in which *fruit* and *color* are variables replaced with the type of the fruit and its natural color. Then, we compute the difference between the score of the true caption and that of the adversarial caption given the image. Intuitively, the score should always be greater than zero for an image-text matching algorithm without severe language bias.

Results. Fig 3b compares MASS and token likelihood of the captioning model (TL) on the color debiasing experiment. We report the ratio of the biased samples and the biased types (fruit or vegetable types of mean value lower than zero) here. TL prefers its language bias over the image context in a third of the categories. In contrast, MASS effectively reduces language bias, resulting in mean differences that are consistently above zero for all but one category.

4.2 Debiasing Numbers

VL models are known to prefer specific numbers such as 0, 1, and 2 (Parcalabescu et al. 2021), regardless of the actual number of objects in an image. Here, we test MASS on the task of counting the number of visual entities.

Approach. We adopt the counting benchmark in VALSE dataset (Parcalabescu et al. 2022). VALSE formulates counting tasks as foils. Given an image, a model is asked to differentiate the true caption with the correct number from the false caption with the wrong number (foil). The benchmark consists of three sets: a set with balanced number distribution in the foils (*balanced*), a set that only uses small numbers (*small*), and a set in which true captions have only large numbers and foils have only small numbers (*adversarial*). Note that models typically understand small numbers better since most images in training data do not have many objects. Hence, the *adversarial* set is especially hard for models with language bias.

Metric and baselines. We report pairwise ranking accuracy metric, which formulates the counting task as an image-text alignment scoring problem. Other baselines here include the multitask models such as LXMERT (Tan and Bansal 2019) and 12-in-1 (Lu et al. 2020).

Results. Table 1 shows that our MASS outperforms all baselines except one in all subsets. Importantly, the score excels at the *adversarial* set and performs on par with the best model 12-in-1. As the *adversarial* set is designed to be hard for models with language bias, this result further verifies that the improvement of MASS comes from language debiasing.

Models	balanced			Small			Adv.		
	ITM/C	TL	MASS	ITM/C	TL	MASS	ITM/C	TL	MASS
Chance	16.7			16.7			16.7		
CLIP	62.1	-	-	62.5	-	-	57.5	-	-
LXMERT	62.2	-	-	69.2	-	-	42.6	-	-
12-in-1	76.7	-	-	80.2	-	-	77.3	-	-
OFA _{large}	66.8	65.4	70.0 (\uparrow 3.2)	71.7	71.1	73.2 (\uparrow 1.5)	72.0	66.8	76.7 (\uparrow 4.7)

Table 1: Counting experiment in VALSE dataset (Parcalabescu et al. 2022). Adv. denotes the adversarial set with explicit language bias. The best numbers given the same backbone are **bolded**.

	Model	Method	Gender Bias \downarrow			Recall \uparrow		
			Bias@1	Bias@5	Bias@10	Recall@1	Recall@5	Recall@10
T2I	SCAN	ITC	13.8	21.3	24.8	25.4	54.1	67.8
	FairSample	ITC	11.3	19.2	22.9	26.8	55.3	68.5
	CLIP-clip	ITC	6.7	14.7	16.1	27.3	50.8	62.0
	CLIP	ITC	8.8	15.7	<u>18.9</u>	36.5	61.0	71.6
	OFA _{large}	ITM	13.2	19.5	23.3	35.3	67.1	76.9
		TL	<u>8.2</u>	<u>15.4</u>	19.2	38.0	65.1	74.4
MASS		9.2	15.9	<u>18.3</u>	55.6	75.2	79.5	
I2T	CLIP	ITC	4.0	11.4	15.4	56.9	80.0	87.0
	OFA _{large}	ITM	6.3	12.7	15.7	42.1	79.7	88.9
		TL	6.6	12.4	15.3	33.4	70.7	84.8
		MASS	3.5	9.9	13.4	57.3	84.7	89.8

Table 2: Gender bias reduction experiment on MS-COCO test split. *T2I* denotes text-to-image retrieval and *I2T* denotes image-to-text. We compare our image-text matching score (MASS) with the baselines on both the gender bias metric Bias@K and the retrieval metric Recall@K. The best numbers are **bolded** and the second-best numbers are underlined.

4.3 Debiasing Genders

We now turn to a more practical problem of balancing gender bias in image-text matching. Here, we adopt the conventional experiment setup (Wang, Liu, and Wang 2021) to evaluate the gender bias of a text-based image search algorithm. The hypothesis is simple: given a gender-neutral text query, the retrieved images should contain a similar number of men and women in them.

Approach. We evaluate our method in both text-to-image and image-to-text retrieval tasks using the test split of MS COCO dataset (Chen et al. 2015).¹ Our text-to-image experiment follows Wang *et al.*'s setup. First, we *neutralize* each caption in the dataset by replacing any gender-specific word with the corresponding gender-neutral alternatives. Then, we perform text-to-image retrieval on the COCO dataset using neutralized captions.

For the image-to-text retrieval task, we follow a similar approach; however, neither captions nor images are neutralized to align with the definition of the bias score described below. Instead, retrieval is performed using original captions. Neutralization of images (e.g., masking gender-related features) was avoided as it led to a loss of critical contextual information necessary for valid retrieval comparisons. We utilize CLIP (Radford et al. 2021) to retrieve the top 20 candidates, followed by re-ranking the list using image-text matching algorithms. This two-stage retrieval strategy is

¹We use the Karpathy split (Karpathy and Fei-Fei 2015) with 5000 test images.

commonly adopted to improve evaluation efficiency (Wang et al. 2022a; Li et al. 2023).

For the gender debiasing experiments, we include additional baselines from prior debiasing methods: SCAN (Lee et al. 2018), FairSample (Lee et al. 2018), and CLIP-clip (Wang, Liu, and Wang 2021). It should be noted that our results for the CLIP model differ from those reported by Wang *et al.* due to the use of a larger model scale (ViT-Large in our study versus ViT-Base in theirs).

Metrics. The *bias score* suggested in Wang *et al.* compares the proportions of masculine and feminine images or captions in the search results. When N_m and N_f are respectively the number of retrieved images or captions depicting male and female in dataset D , the bias score for the top K retrieved results ($Bias@K$) is defined as follows:

$$f(x) := \begin{cases} 0 & \text{if } N_m + N_f = 0 \\ \frac{N_m - N_f}{N_m + N_f} & \text{otherwise} \end{cases} \quad (10)$$

$$Bias@K := \frac{1}{|D|} \sum_{x \in D} f(x) \quad (11)$$

Note that the query x is an image to retrieve a list of captions in image-to-text retrieval, and vice versa for text-to-image. We use the words in captions to decide whether an image or caption contains gender-specific information: refer to Wang *et al.* (Wang, Liu, and Wang 2021) for more details.

We also evaluate the standard *recall* metric as an indicator of correctness: for each query, recall is one when the re-

trieved samples contain any of the ground-truth image-text pairs and zero in the other case. A good algorithm should be able to minimize bias and keep the recall high.

Results. Table 2 summarizes the results of our gender debiasing experiments. In the text-to-image setting, MASS outperforms all other baselines by a wide margin in recall, while maintaining the bias score close to the best non-debiasing method (TL) of the baselines. Also, while a previous debiasing-specific method (CLIP-clip) effectively ameliorates gender biases, it comes at an expense of high degradation of recall accuracy.

On image-to-text retrieval, the captions retrieved by MASS contained much more balanced gender keywords than the baselines. Further, MASS still shows the best recall among others, proving that its social bias reduction does not come at an expense of correctness. In conclusion, MASS is the best option here for gender-debiased retrieval.

5 Experiments: Linguistic Complexity

As MASS is built on captioning models, it could suffer less from the well-known linguistic insensitivity defect (Yuksekgonul et al. 2022) of ITC-based models such as CLIP (Radford et al. 2021). However, we still need to verify whether the debiasing effect of MASS also comes at a loss of this linguistic capability. Thus we investigate whether MASS can outperform the baselines in a visual-language compositionality benchmark (section 5.1) and a benchmark to test the capability to distinguish subject, noun, and verb in a caption (section 5.2). Both tasks require understanding linguistic structures beyond basic bag-of-words representations.

5.1 Winoground

The Winoground benchmark (Thrush et al. 2022) evaluates a VL model’s capability to understand compositionality. A sample in Winoground consists of two images and two captions with two correct and two wrong image-text pairs. Importantly, the two captions contain the same set of words but in a different order. This way, a model that perceives captions as bag-of-words (Yuksekgonul et al. 2022) without understanding visual-linguistic compositionality cannot do well on the task.

Metrics and baselines. We use the suggested set of metrics (Thrush et al. 2022): *Textscore* which compares captions given the images, *Imagescore* which compares images given the captions, and *Groupscore* that is true only if both are true. Note that *Groupscore* is the preferred metric since the other metrics can overestimate model performance by separating score computations for captions in the same sample (Elazar et al. 2021). Baselines include popular VL models (VinVL (Zhang et al. 2021), LXMERT (Tan and Bansal 2019), and CLIP (Radford et al. 2021)).

Results. Table 3 shows MASS outperforms all baselines by a wide margin in the main metric (*Groupscore*). The same OFA_{large} model saw more than twofold improvement by using MASS over the ITM score. Also, note that MASS with OFA_{large} is the only option that surpasses random chance in *Groupscore*. Finally, MASS’s contribution is better represented in *Imagescore* than in *Textscore*, indicating that

MASS can balance the unusually low *Imagescore* of the baselines with the *Textscore*.

We also observe a similar trend of results with an alternative backbone model (BLIP-2 (Li et al. 2023)) as MASS improves the base token likelihood method (TL) across all scores. However, there are considerable gaps in Textscore between ITM and TL. We suspect that the advanced training techniques used to optimize the ITM head in BLIP-2, such as hard negative mining (Li et al. 2021), contributed to its better performance over the OFA model. Still, MASS helps close the performance gap in the Groupscore.

5.2 SVO-Probes

SVO-Probes (Hendricks and Nematzadeh 2021) is another benchmark testing VL models’ sensitivity to linguistic alterations. Here, the models are given a true caption and a false caption with either the subject, verb, or object changed to break the image-caption alignment.

Metrics and baselines. Unlike the original dataset paper (Hendricks and Nematzadeh 2021), we here use Winoground-style metrics for evaluation to prevent overvaluing a VL model’s response as in the previous section (section 5.1). Note that the original accuracy metric does not accept an arbitrary similarity function since it requires a binary decision on the image-text pair match. We include the vanilla CLIP model as the baseline here.

Results. As in the Winoground experiment, Table 4 also shows that MASS can improve image-text matching in data with linguistic complexities. MASS outperforms all baselines on each metric, showing that it distinguishes the linguistic aspects better than other methods.

6 Related Work

Visual-Language Model. Early visual-language models like LXMERT (Tan and Bansal 2019) established visual-language alignment based on pretrained text encoder representations. Later, the success of CLIP (Radford et al. 2021) introduced contrastive learning as a core method for visual-language models, leading to further research like ALBEF (Li et al. 2021), which incorporates contrastive loss as a part of its multi-task losses. BLIP-2 (Li et al. 2023) demonstrated that frozen large language models can enhance image captioning performance. OFA (Wang et al. 2022b) further extend visual-language models to a wider range of tasks including image generation. Recent visual instruction tuning models allow prompt-based control of visually conditioned knowledge. LLaVA (Liu et al. 2024) extend the pretrained language model LLAMA to multimodal inputs. A prior study (Lin et al. 2024) proposes a framework for mitigating language priors in vision-language models by introducing ‘null images’ and leveraging Pointwise Mutual Information (PMI), techniques similar to those presented in this work. Although both works share methodological elements, they were developed independently and address different objectives: (Lin et al. 2024) focuses on theoretical contributions to vision-language alignment, while this work targets societal bias reduction and introduces new evaluation metrics, such as bias-aware Text-to-Image retrieval.

Models	Text			Image			Group			
	ITM/C	TL	MASS	ITM/C	TL	MASS	ITM/C	TL	MASS	
Human		89.5			88.5			85.5		
Chance		25.0			25.0			16.7		
VinVL	37.8	-	-	17.8	-	-	14.5	-	-	
LXMERT	19.3	-	-	7.0	-	-	4.0	-	-	
CLIP	30.8	-	-	10.5	-	-	8.0	-	-	
OFA _{large}	26.4	26.8	32.0 (↑5.2)	16.5	29.8	31.5 (↑2.7)	9.5	15.8	20.3 (↑4.5)	
base	26.8	23.8	30.3 (↑3.5)	10.8	25.3	21.3 (↓4.0)	6.5	12.5	15.3 (↑2.8)	
tiny	22.8	15.5	18.0 (↓4.8)	7.8	14.8	17.5 (↑2.7)	4.5	5.0	8.0 (↑3.0)	
BLIP-2 _{T5}	42.5	22.8	26.8 (↓15.7)	19.5	21.5	32.0 (↑10.5)	15.5	11.0	18.0 (↑2.5)	
OPT	42.5	23.3	22.8 (↓19.7)	19.5	24.8	29.8 (↑5.1)	15.5	12.5	15.0 (↓0.5)	
LLaVA	-	27.0	32.3 (↑5.3)	-	24.7	31.8 (↑7.1)	-	14.0	19.0 (↑5.0)	

Table 3: Results on Winoground (Thrush et al. 2022). The best matching algorithm for each model is marked **bold**.

Models	Text			Image			Group			
	ITM/C	TL	MASS	ITM/C	TL	MASS	ITM/C	TL	MASS	
Chance		25.0			25.0			16.7		
CLIP	61.1	-	-	38.3	-	-	31.8	-	-	
OFA _{large}	59.0	31.8	66.0 (↑7.0)	47.5	21.5	50.1 (↑2.6)	39.7	12.5	43.0 (↑4.7)	

Table 4: Result on SVO-Probes (Hendricks and Nematzadeh 2021). The best numbers are **bolded**.

Bias in Multimodal Models. Visual-language models often suffer from language bias induced by their training dataset. This deficiency has been explored as visual priming bias (Zhang et al. 2016) and language prior (Agrawal et al. 2018). Language bias is also closely connected to socially offensive bias, such as gender bias (de Vassimon Manela et al. 2021) and racial bias (Davidson, Bhattacharya, and Weber 2019) in language models. Social biases are also observed in the VLMs (Hendricks et al. 2018). Other works explore diverse categories (Garcia et al. 2023) and metric extension (Ross, Katz, and Barbu 2021) for social bias understanding.

7 Conclusion

We introduce MASS (Multimodal ASSociation Score) as an inference-time framework to address the issue of language bias in image-text matching. MASS builds on image captioning models to extract image-text matching capability devoid of their linguistic biases. Specifically, it computes pointwise mutual information between the image and each text token and aggregates them to generate a scalar similarity score. The proposed method significantly reduces language bias without additional training and demonstrates improved performance in debiasing tasks and visuo-linguistic compositionality tests, outperforming existing models such as CLIP and token-level likelihood.

We hope that our research sparks interest in 1. investigating linguistic bias inherent in the image-text alignment mechanism of the recent visual-language models and 2. devising a method to reduce such language bias in off-the-shelf visual-language models without incorporating the computationally overwhelming training process.

Broader Impact

The social bias reduction effect of MASS has two important limitations. First, it relies on the output of visual-language models, which may encode hidden biases. Since MASS is an inference-time method, it cannot audit the information stored in the models’ parameters. Hence, we acknowledge that while our framework can reduce social bias, the range and intensity of its effect may vary with respect to the base VL model used. Second, we only explore social bias in terms of language bias. Social bias could be manifested as visual bias or even in the form of bias that is identified only as combinations of image and text (Yamada, Tang, and Yildirim 2022). We hope this research serves as a simple but effective baseline to initiate a search for more robust approaches to social bias reduction in image-to-text alignment problems.

Acknowledgements

This work was partly supported by an IITP grant funded by the Korean Government (MSIT) (No.RS-2020-II201361, Artificial Intelligence Graduate School Program (Yonsei University) and RS-2024-00353131) and the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No. RS-2024-00354218).

References

- Agrawal, A.; Batra, D.; and Parikh, D. 2016. Analyzing the Behavior of Visual Question Answering Models. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, 1955–1960*.
- Agrawal, A.; Batra, D.; Parikh, D.; and Kembhavi, A. 2018. Don’t just assume; look and answer: Overcoming priors for visual question answering. In *Proceedings of the IEEE con-*

- ference on computer vision and pattern recognition, 4971–4980.
- Alayrac, J.-B.; Donahue, J.; Luc, P.; Miech, A.; Barr, I.; Hasson, Y.; Lenc, K.; Mensch, A.; Millican, K.; Reynolds, M.; et al. 2022. Flamingo: a visual language model for few-shot learning. *Advances in neural information processing systems*, 35: 23716–23736.
- Anwar, S.; Tahir, M.; Li, C.; Mian, A.; Khan, F. S.; and Muzaffar, A. W. 2020. Image colorization: A survey and dataset. *arXiv preprint arXiv:2008.10774*.
- Beaumont, R. 2022. Clip retrieval: Easily compute clip embeddings and build a clip retrieval system with them. *GitHub*.
- Birhane, A.; Prabhu, V. U.; and Kahembwe, E. 2021. Multimodal datasets: misogyny, pornography, and malignant stereotypes. *arXiv preprint arXiv:2110.01963*.
- Chen, X.; Fang, H.; Lin, T.-Y.; Vedantam, R.; Gupta, S.; Dollár, P.; and Zitnick, C. L. 2015. Microsoft coco captions: Data collection and evaluation server. *arXiv preprint arXiv:1504.00325*.
- Davidson, T.; Bhattacharya, D.; and Weber, I. 2019. Racial bias in hate speech and abusive language detection datasets. *arXiv preprint arXiv:1905.12516*.
- de Vassimon Manela, D.; Errington, D.; Fisher, T.; van Breugel, B.; and Minervini, P. 2021. Stereotype and skew: Quantifying gender bias in pre-trained and fine-tuned language models. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, 2232–2242.
- Elazar, Y.; Zhang, H.; Goldberg, Y.; and Roth, D. 2021. Back to Square One: Artifact Detection, Training and Commonsense Disentanglement in the Winograd Schema. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 10486–10500.
- Garcia, N.; Hirota, Y.; Wu, Y.; and Nakashima, Y. 2023. Uncurated image-text datasets: Shedding light on demographic bias. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6957–6966.
- Hendricks, L. A.; Burns, K.; Saenko, K.; Darrell, T.; and Rohrbach, A. 2018. Women also snowboard: Overcoming bias in captioning models. In *Proceedings of the European conference on computer vision (ECCV)*, 771–787.
- Hendricks, L. A.; and Nematzadeh, A. 2021. Probing Image-Language Transformers for Verb Understanding. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, 3635–3644.
- Hessel, J.; Holtzman, A.; Forbes, M.; Le Bras, R.; and Choi, Y. 2021. CLIPScore: A Reference-free Evaluation Metric for Image Captioning. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 7514–7528.
- Karpathy, A.; and Fei-Fei, L. 2015. Deep visual-semantic alignments for generating image descriptions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 3128–3137.
- Lee, K.-H.; Chen, X.; Hua, G.; Hu, H.; and He, X. 2018. Stacked cross attention for image-text matching. In *Proceedings of the European conference on computer vision (ECCV)*, 201–216.
- Li, J.; Galley, M.; Brockett, C.; Gao, J.; and Dolan, B. 2015. A diversity-promoting objective function for neural conversation models. *arXiv preprint arXiv:1510.03055*.
- Li, J.; Li, D.; Savarese, S.; and Hoi, S. 2023. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International conference on machine learning*, 19730–19742. PMLR.
- Li, J.; Li, D.; Xiong, C.; and Hoi, S. 2022. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International conference on machine learning*, 12888–12900. PMLR.
- Li, J.; Selvaraju, R.; Gotmare, A.; Joty, S.; Xiong, C.; and Hoi, S. C. H. 2021. Align before fuse: Vision and language representation learning with momentum distillation. *Advances in neural information processing systems*, 34: 9694–9705.
- Lin, Z.; Chen, X.; Pathak, D.; Zhang, P.; and Ramanan, D. 2024. Revisiting the Role of Language Priors in Vision-Language Models. In *Forty-first International Conference on Machine Learning*.
- Liu, H.; Li, C.; Wu, Q.; and Lee, Y. J. 2024. Visual instruction tuning. *Advances in neural information processing systems*, 36.
- Lu, J.; Goswami, V.; Rohrbach, M.; Parikh, D.; and Lee, S. 2020. 12-in-1: Multi-Task Vision and Language Representation Learning. In *The IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Nikolaus, M.; Salin, E.; Ayache, S.; Fourtassi, A.; and Favre, B. 2022. Do Vision-and-Language Transformers Learn Grounded Predicate-Noun Dependencies? In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*.
- Niu, Y.; Tang, K.; Zhang, H.; Lu, Z.; Hua, X.; and rong Wen, J. 2020. Counterfactual VQA: A Cause-Effect Look at Language Bias. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 12695–12705.
- Oord, A. v. d.; Li, Y.; and Vinyals, O. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*.
- Parcalabescu, L.; Cafagna, M.; Muradjan, L.; Frank, A.; Calixto, I.; and Gatt, A. 2022. VALSE: A Task-Independent Benchmark for Vision and Language Models Centered on Linguistic Phenomena. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 8253–8280.
- Parcalabescu, L.; Gatt, A.; Frank, A.; and Calixto, I. 2021. Seeing past words: Testing the cross-modal capabilities of pretrained V&L models on counting tasks. *IWCS 2021*, 32.
- Petryk, S.; Whitehead, S.; Gonzalez, J. E.; Darrell, T.; Rohrbach, A.; and Rohrbach, M. 2024. Simple token-level confidence improves caption correctness. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 5742–5752.

- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; Krueger, G.; and Sutskever, I. 2021. Learning Transferable Visual Models From Natural Language Supervision. In *International Conference on Machine Learning*.
- Ross, C.; Katz, B.; and Barbu, A. 2021. Measuring Social Biases in Grounded Vision and Language Embeddings. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 998–1008.
- Salin, E.; Farah, B.; Ayache, S.; and Favre, B. 2022. Are vision-language transformers learning multimodal representations? a probing perspective. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 11248–11257.
- Shekhar, R.; Pezzelle, S.; Klimovich, Y.; Herbelot, A.; Nabi, M.; Sangineto, E.; and Bernardi, R. 2017. FOIL it! Find One mismatch between Image and Language caption. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 255–265.
- Si, Q.; Liu, Y.; Meng, F.; Lin, Z.; Fu, P.; Cao, Y.; Wang, W.; and Zhou, J. 2022. Towards Robust Visual Question Answering: Making the Most of Biased Samples via Contrastive Learning. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, 6650–6662.
- Srinivasan, T.; and Bisk, Y. 2021. Worst of both worlds: Biases compound in pre-trained vision-and-language models. *arXiv preprint arXiv:2104.08666*.
- Tan, H.; and Bansal, M. 2019. Lxmert: Learning cross-modality encoder representations from transformers. *arXiv preprint arXiv:1908.07490*.
- Thrush, T.; Jiang, R.; Bartolo, M.; Singh, A.; Williams, A.; Kiela, D.; and Ross, C. 2022. Winoground: Probing vision and language models for visio-linguistic compositionality. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5238–5248.
- Tschannen, M.; Kumar, M.; Steiner, A.; Zhai, X.; Houlsby, N.; and Beyer, L. 2024. Image captioners are scalable vision learners too. *Advances in Neural Information Processing Systems*, 36.
- Van der Poel, L.; Cotterell, R.; and Meister, C. 2022. Mutual information alleviates hallucinations in abstractive summarization. *arXiv preprint arXiv:2210.13210*.
- Wang, J.; Liu, Y.; and Wang, X. 2021. Are Gender-Neutral Queries Really Gender-Neutral? Mitigating Gender Bias in Image Search. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 1995–2008.
- Wang, J.; Yang, Z.; Hu, X.; Li, L.; Lin, K.; Gan, Z.; Liu, Z.; Liu, C.; and Wang, L. 2022a. Git: A generative image-to-text transformer for vision and language. *arXiv preprint arXiv:2205.14100*.
- Wang, P.; Yang, A.; Men, R.; Lin, J.; Bai, S.; Li, Z.; Ma, J.; Zhou, C.; Zhou, J.; and Yang, H. 2022b. Ofa: Unifying architectures, tasks, and modalities through a simple sequence-to-sequence learning framework. In *International Conference on Machine Learning*, 23318–23340. PMLR.
- Yamada, Y.; Tang, Y.; and Yildirim, I. 2022. When are lemons purple? the concept association bias of clip. *arXiv preprint arXiv:2212.12043*.
- Young, P.; Lai, A.; Hodosh, M.; and Hockenmaier, J. 2014. From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. *Transactions of the Association for Computational Linguistics*, 2: 67–78.
- Yu, Y.; Chung, J.; Yun, H.; Hessel, J.; Park, J. S.; Lu, X.; Zellers, R.; Ammanabrolu, P.; Le Bras, R.; Kim, G.; et al. 2023. Fusing pre-trained language models with multimodal prompts through reinforcement learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10845–10856.
- Yuksekgonul, M.; Bianchi, F.; Kalluri, P.; Jurafsky, D.; and Zou, J. 2022. 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*.
- Zhang, P.; Goyal, Y.; Summers-Stay, D.; Batra, D.; and Parikh, D. 2016. Yin and yang: Balancing and answering binary visual questions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 5014–5022.
- Zhang, P.; Li, X.; Hu, X.; Yang, J.; Zhang, L.; Wang, L.; Choi, Y.; and Gao, J. 2021. Vinvl: Revisiting visual representations in vision-language models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 5579–5588.