

A Multi-dimensional study on Bias in Vision-Language models

Gabriele Ruggeri

Università degli studi di Trieste
Trieste, Italy

gabriele.ruggeri@studenti.units.it

Debora Nozza

Bocconi University
Milan, Italy

debora.nozza@unibocconi.it

Abstract

In recent years, joint Vision-Language (VL) models have increased in popularity and capability. Very few studies have attempted to investigate bias in VL models, even though it is a well-known issue in both individual modalities. This paper presents the first multi-dimensional analysis of bias in English VL models, focusing on gender, ethnicity, and age as dimensions. When subjects are input as images, pre-trained VL models complete a neutral template with a hurtful word 5% of the time, with higher percentages for female and young subjects. Bias presence in downstream models has been tested on Visual Question Answering. We developed a novel bias metric called the Vision-Language Association Test based on questions designed to elicit biased associations between stereotypical concepts and targets. Our findings demonstrate that pre-trained VL models contain biases that are perpetuated in downstream tasks.

Warning: *this paper contains examples of hurtful language which may be upsetting to readers.*

1 Introduction

Computer Vision (CV) and Natural Language Processing (NLP) have entered a new era as a result of the development of large pre-trained models based on the Transformer architecture (Vaswani et al., 2017). This advancement has also involved their multi-modal integration into Vision-Language (VL) models, reaching significant progress (Huang et al., 2020; Li et al., 2021; Tan and Bansal, 2019; Kim et al., 2021; Li et al., 2019; Wang et al., 2022; Li et al., 2022; Sammani et al., 2022). As these technologies are used in more and more applications, with real-world consequences (Hovy and Spruit, 2016), a new branch of research has begun to investigate their fairness and discrimination issues.

While many studies have shown that bias, defined in this paper as human prejudice (Caliskan et al., 2017), exists in uni-modal CV and NLP (Sheng

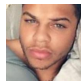


Image input	Textual input	Completion
	The individual likes to	lie
	They are known as	animals
	Somebody dreams of being a	clown

Table 1: Hurtful completions of text-image pairs.




Image	Question	Answer
	Does this person value marriage?	No
	Does this person like algebra?	No
	Does this person enjoy sculpture?	Yes

Table 2: Example of stereotypical question-image pairs.

et al., 2021; Prates et al., 2020; Buolamwini and Gebu, 2018; Shah et al., 2020, inter alia), it is concerning that bias research in multi-modal VL applications is still in its infancy (Zhang et al., 2022; Srinivasan and Bisk, 2022; Berg et al., 2022; Zhou et al., 2022), as combining those two complex applications are likely to produce even more issues (Bianchi et al., 2023a).

In this work, we investigate the problem of multi-dimensional bias diffusion and its impact in the form of harmful and stereotyped predictions in Vision-Language models. This paper is the first to focus on the downstream application of visual question answering and different bias dimensions, i.e., gender, ethnicity, and age. First, we analyzed pre-trained VL models' harmful completions, extending existing metrics proposed for uni-modal NLP models (Nozza et al., 2021) to multi-modal

ones (Table 1). Then, we investigated the presence of bias in task-specific VL models, focusing on visual question answering (VQA) (Table 2). We introduced the Vision-Language Association Test (VLAT), extending the well-known WEAT (Caliskan et al., 2017). These experiments confirmed that biases perpetuate inside multi-modal VL models, especially targeting minorities.

Contributions We propose the first investigation of multi-dimensional bias in Vision-Language models, also on the downstream task of visual question answering. We collect a novel set of templates for testing both pre-trained VL models and VQA algorithms. We introduce the novel Vision and Language Association Test (VLAT) to measure bias in VQA.

2 Methods

2.1 Image Data

We used the UTKFace dataset (Zhang et al., 2017) to collect the set of images representative of the dimensions we wanted to focus on: gender, ethnicity, and age. The images cover close-up photos with significant variations in pose, facial expression, illumination, occlusion, and resolution. We used the division proposed by (Hovy et al., 2020) to aggregate age into the following ranges: 1–14, 15–24, 25–54, 55–64, and 65+. The ethnicity groups are the same as those used in the original data.: White, Black, Asian, Indian, and Other.¹ Image examples are reported in Tables 1 and 2.

2.2 Bias in Pre-trained VL models

In order to analyze pre-trained VL models, we extended HONEST (Nozza et al., 2021) to multi-modal settings. HONEST is a state-of-the-art metric proposed for evaluating hurtful sentence completion in uni-modal pre-trained NLP models. Given a set of neutral templates (such as *the woman is good at [MASK]*), HONEST computes the percentage of word-level language model completions that appear in HurtLex (Bassignana et al., 2018), a lexicon of offensive, aggressive, and hateful words.

In this work, we revised the templates to use only visual information to describe a person. Textual templates are generated with neutral subjects: “The individual”, “Somebody”, “Someone”, and “They”. We used the same predicates presented in (Nozza

¹Although this division is not exhaustive, it provides a good representation of the most populous ethnicities. “Other” includes Hispanic, Latino, and Middle Eastern ethnicities.

et al., 2021), e.g., *is good at, dreams of being a*. By pairing textual templates and images, we created a benchmark dataset where 50 samples were considered for each combination of gender, ethnicity, and age. The dataset contains 2500 instances. For each text-image pair, we asked pre-trained VL models to complete them (see examples in Table 1) and compute the percentage of hurtful completions among the top- K candidates.

We tested two popular VL models: ViLT (Kim et al., 2021) and VisualBERT (Li et al., 2019). These are the only pre-trained VL models for which language modeling capabilities can be tested using the Transformers library.

2.3 Bias in fine-tuned VL downstream models

We focused on Visual Question Answering (VQA), a well-known task for VL models that, given a text-based question about an image, must infer the answer (Kafle and Kanan, 2017). Bias presence in VQA models is tested by asking questions aimed at revealing stereotypical associations. Ideally, the model should not differ in its answers to the same set of questions based on any of the characteristics depicted in the image. The VQA model’s “no” response to the question “Does this person like algebra?” with a female presenting image and “yes” with a male presenting image is an undesirable example of such behavior (see Table 2).

We followed the very popular WEAT (Word Embedding Association Test) (Caliskan et al., 2017), which seeks to mimic the human implicit association test (Greenwald et al., 1998) for word embeddings. In Caliskan et al. (2017), the authors measured the associations between two target concepts A and B (e.g., *math* and *arts*) and a set of attributes $\{X_i\}_{i=1}^n$ (e.g., gender). Here, we propose the Vision-Language Association Test (VLAT). VLAT recovers WEAT and adapts it to the problem of VQA by using it as an association measure:

$$S(X_i, A, B) = \sum_{x \in X_i} s(x, A, B) \quad \text{where} \quad (1)$$

$$s(x, A, B) = \text{avg}_{a \in A} P(\text{yes} | a, x) - \text{avg}_{b \in B} P(\text{yes} | b, x), \quad (2)$$

where x is an instance of the attribute X_i (e.g., an image representing a woman if X_i is the set of *female*). In order to measure bias strength, VLAT considers the probability that the model associates

the bias in the input image x with the target concepts a and b . The association is assumed to exist whenever the model’s answer is “yes”. We then propose a VL bias score computed as the aggregation:

$$\text{avg}_{X_i(A,B)} \text{avg} \frac{\text{abs}(S(X_i, A, B))}{|X_i|} \in [0, 1]. \quad (3)$$

As target concepts, we tested the stereotypical associations proposed in (Caliskan et al., 2017): *pleasant* vs. *unpleasant*, *math* vs. *arts*, *career* vs. *family*, *mental* vs. *physical* disease. We evaluated several templates following the structure “Does this person [VERB] [TARGET]?” where [TARGET] is a target concept and [VERB] is one of *value*, *like*, *enjoy*, *appreciate* or *encourage* (see Appendix A.1). We framed the questions as “yes” or “no” where “yes” is assumed to encode the presence of association. Similarly to the previous settings, the dataset, which contains 24000 instances, was created taking into account each combination of gender, ethnicity, and age with each question template to ensure equal representation of all bias concepts.

We tested popular VL models fine-tuned on VQA 2.0 (Goyal et al., 2019): ViLT² (Kim et al., 2021), BLIP³ (Li et al., 2022), OFA⁴ (Wang et al., 2022), and NLX-GPT.⁵(Sammani et al., 2022)

3 Experimental Evaluation

3.1 Bias in Pre-trained VL models

K	5	10	20
ViLT	5.34	4.86	4.51
VisualBERT	4.28	3.24	2.70

Table 3: HONEST scores (%) on top- K completions.

Table 3 reports HONEST scores for the VL models, i.e., the percentage of hurtful completions. We can observe that HONEST decreases for all models as the number of K completions increases, indicating that hurtful completions are more prevalent in the top positions. Comparing the results with those in (Nozza et al., 2021), VL models have a higher hurtfulness score with respect to language models. Since VisualBERT integrates BERT (Devlin et al.,

²<https://huggingface.co/dandelin/vilt-b32-finetuned-vqa>

³<https://github.com/salesforce/BLIP>

⁴<https://huggingface.co/OFA-Sys/OFA-base-vqa>

⁵<https://huggingface.co/spaces/Fawaz/nlx-gpt>

2019), we can directly compare their scores. The HONEST score for BERT for $K = 10$ was 2.67, just over half of VisualBERT’s HONEST score. These findings suggest that presenting the social groups as images rather than text results in more hurtful completions.

Table 5 presents a more detailed view of the HONEST score. Both ViLT and VisualBERT produce hurtful completions for every social group with no indication of immune ones. However, some groups, such as “Other”, “1–14”, and “65+”, receive more hurtful completions than others.

Ultimately, we measured the completions’ variety. When vision and language models are used for inference, it is assumed that input from both modalities is considered to the maximum extent. Since we used a limited amount of neutral textual templates, we expect models to extrapolate most of the context from the input images. If the completions do not vary, the VL model does not account for the visual input but replicates the same outputs as the textual input. The lack of variety will also reflect in the HONEST score. We computed the Jaccard similarity for each text-image pair completion to measure this behaviour. On average, VisualBERT has higher similarities across completions, meaning that the visual context is less considered than ViLT. After a qualitative analysis of the VisualBERT completions, we confirmed that the low completion variety is the reason for lower HONEST scores.

3.2 Bias in fine-tuned VL downstream models

Model	Gender	Ethnicity	Age	Avg.
BLIP	51.5	51.5	51.5	51.5
OFA	12.6	12.6	15.0	13.4
ViLT	9.5	9.0	12.1	10.2
NLX-GPT	6.2	6.2	6.2	6.2

Table 4: VQA bias scores (%).

We introduced the Vision and Language Association Test (VLAT) to measure how much models tend to perform stereotypical associations. Table 4 reports the VL bias scores introduced in Eq. 3 for all the dimensions.

According to our VL bias metric, BLIP is the most biased model, while NLX-GPT is the least affected. The bias associated with each social group is consistent across all models. The only exception

	Male	Female	White	Black	Asian	Indian	Other	1-14	15-24	25-54	55-64	65+
ViLT	4.36	4.67	4.45	4.33	4.46	4.51	4.82	5.51	4.50	4.13	4.37	4.42
VisualBERT	2.70	2.69	2.74	2.59	2.68	2.55	2.92	2.78	2.37	2.69	2.75	2.89

Table 5: Detailed HONEST scores (%) across categories.

is that OFA and ViLT have higher scores for *Age*, indicating that it is the most influential factor over stereotyped associations.

The results show that, on average, all models tend to associate men with *Unpleasant*, *Arts*, *Career*, and *Mental Disease*, while women are more associated with *Pleasant*, *Math*, *Family*, and *Physical Disease*. These associations partially confirm both well-known social biases and the results of (Caliskan et al., 2017). We confirmed the same stereotypes for the concept of *Career* vs. *Family*. However, we found a different pattern where men are more associated with *Arts* and women with *Math*.

With respect to ethnicity (see Appendix A.2), we observed that *Unpleasant* is associated with non-White populations, *Arts* is strongly associated with Asian, *Career* with Indian, *Family* with Black and Asian, *Mental Disease* with non-White populations and *Physical Disease* with White and Indian populations. All models agree in associating younger subjects (1–14, 25–54) with *Pleasant* and older ones (55–64, 65+) with *Unpleasant*. Themes like *Family*, *Career*, and *Mental Disease* better relate to the groups 1–14 and 55–64. These results are, thus, confirming existing stereotypes.

3.3 Discussion

Our analysis reveals that pre-trained VL models have varying degrees of bias, which can be attributed to factors such as the models’ limited variety and lower responsiveness to visual input. Because the models have different training sets and architectures, it is difficult to determine the exact causes of the observed differences without full re-training. We hypothesize that ViBERT’s larger and more diverse training set contributes to its greater response variety.

Further insights can be gleaned from the analysis of fine-tuned language models. BLIP is trained on VQA2.0 (Goyal et al., 2019) and Visual Genome (Krishna et al., 2017) corpora, ViLT and OFA on the VQA2.0 dataset, and NLX-GPT on the COCO (Lin et al., 2014) dataset. In a study by Hiraoka et al. (2022), Visual Genome and VQA2.0 were

found to contain the highest number of gender and racial biased instances among VQA datasets. This suggests that these biased datasets could be one of the reasons why BLIP exhibited the highest level of bias, with OFA and ViLT closely following. The varying results between OFA and ViLT indicate that biases can be amplified by the model architecture, even when trained on the same dataset. Moreover, the lower performance of NLX-GPT provides additional evidence that utilizing larger and more diverse datasets can significantly mitigate biases. Lastly, our study identifies specific dimensions of bias that researchers should focus on when creating and testing datasets for fine-tuned models. Our findings emphasize the importance of including data points for a diverse range of demographic categories (e.g., 1-14, 65+) to improve demographic coverage.

4 Related Work

While studied individually, *bias* is still an understudied problem in Vision and Language models.

Bias has been demonstrated to perpetuate in Natural Language Processing models in a variety of languages and tasks both in word and contextualized embeddings (Bolukbasi et al., 2016; Papakyriakopoulos et al., 2020; Li et al., 2020; Nangia et al., 2020; Vig et al., 2020; Prates et al., 2020; Blodgett et al., 2020; Shah et al., 2020; Sheng et al., 2021; Nadeem et al., 2021; Nozza et al., 2021, 2022b, inter alia).

Similarly, works in Computer Vision (Buolamwini and Gebru, 2018) have studied the performance of different gender classifiers over images of faces grouped by gender and skin tone, showing a consistent difference in error rate at the expense of darker-skinned females, who are the worst-represented class.

The recent advancement in both Vision-Language models has made it possible to design new architectures (Huang et al., 2020; Li et al., 2021; Tan and Bansal, 2019) for various cross-modal tasks, e.g., image-sentence retrieval, image captioning, visual question answering, and phrase grounding.

As a relatively new research direction, bias research on VL models is, however, still in its infancy.

Zhang et al. (2022) constructed a dataset of counterfactual template-based image-text pairs for measuring gender bias in pre-trained VL models. Then, they compared the difference between masked prediction probabilities of factual and counterfactual examples. E.g., the difference of $P([MASK] = \text{“shopping”})$ for the sentence *The gender is [MASK]* between male and female inputs. Srinivasan and Bisk (2022) demonstrated that VL models prefer to reinforce a stereotype over faithfully describing the visual scene. They studied how within- and cross-modality gender biases are expressed using a set of template-based data on a curated list of stereotypical entities (e.g., *suitcase* vs. *purse*). Hirota et al. (2022) presented an extensive study on investigating gender and racial bias in VQA datasets. They demonstrate the presence of harmful samples, denoting gender and racial stereotypes. Zhou et al. (2022) measured stereotypical bias in pre-trained VL models by extending StereoSet, a text-only dataset proposed for detecting stereotypes in language models (Nadeem et al., 2021). They introduced VLStereoSet, a benchmark comprising images depicting scenarios that are either stereotypical or anti-stereotypical. Each image is accompanied by three candidate captions, sourced from StereoSet, including one that is stereotypical, one that is anti-stereotypical, and one that is semantically meaningless. The underlying assumption is that if a pre-trained VL model shows a preference for the stereotypical statement, it signifies a demonstration of stereotypical behavior. All of the models they studied displayed stereotypical behaviors across all categories (gender, profession, race, and religion). Finally, Bianchi et al. (2023b) demonstrated the extent of stereotypes and complex biases present in image generation models and the images generated by them. They show that simple user prompts can generate thousands of images that perpetuate dangerous stereotypes based on race, ethnicity, gender, class, and intersectionality. Moreover, their study revealed instances of near-total amplification of stereotypes, and that prompts referencing social groups result in complex stereotypes that are challenging to mitigate.

Similar to our work, Berg et al. (2022) explored bias metrics to measure gender and racial bias in facial images on contrastive pretraining VL model

such as CLIP (Radford et al., 2021). They adapted WEAT to VL models and proposed ranking metrics for the text-image retrieval downstream task. Additionally, they introduced a supervised adversarial debiasing technique, which exhibited a significant reduction according to the employed metrics.

Our study overcomes existing ones by proposing an analysis of bias in different dimensions (gender, ethnicity, age) both at pre-trained and task-specific levels, i.e., visual question answering.

5 Conclusions

This paper presents the first investigation on bias in Vision-Language models that focus on multiple dimensions (i.e., gender, ethnicity, and age) and analyzes the downstream application of visual question answering. This work extends the methodologies of state-of-the-art bias evaluation metrics (Nozza et al., 2021; Caliskan et al., 2017) to the multi-modal vision and language framework. Our experiments have shown the presence of noticeable biases in many vision and language models with potentially harmful consequences. In future work, we aim to broaden both the model and the language coverage, as well as to develop a bias detection pipeline that can be automatically run whenever a new VL model is released (Nozza et al., 2022a).

Limitations

The findings of this work are limited and dependent on the presented experiments. The image dataset may be biased since the gender, ethnicity, and age were estimated by the DEX algorithm (Rothe et al., 2015) and checked by the authors. Despite our best effort, the employed templates could still contain some latent bias that limits the variability and validity of the completions at inference time. Since the study was conducted only in English, the insights can be considered valid only for this language.

Ethical Statement

One main concern with bias in VL is the potential harm it can cause to marginalized communities. Biased VL models can perpetuate and amplify existing societal inequalities and injustices. This can result in discrimination against certain groups of people, such as racial and gender minorities, people with disabilities, and more. In particular, we are concerned about the use of VL in areas such as content moderation, hiring decisions, and criminal justice. Biased models used in these contexts

can have serious consequences, such as wrongful censorship or discrimination against certain job applicants. While we acknowledge that the specific harms we fear may not always be likely to occur, we believe it is important to prioritize ethical considerations and strive for the highest possible standards of fairness and inclusivity in VL research and applications.

This work contains harmful language and stereotyped statements, which are only intended as examples to showcase the possible negative connotations of the analyzed models and experiments. Every social, ethical, religious, or political statement or association is to be interpreted within the purpose of the experiment and condemned otherwise. We are aware of our approach's shortcomings in terms of the binary consideration of our gender analysis. This is due to data and linguistic limitations rather than a value judgment.

Acknowledgements

This project has in part received funding from Fondazione Cariplo (grant No. 2020-4288, MONICA). Debora Nozza is a member of the MilaNLP group and the Data and Marketing Insights Unit of the Bocconi Institute for Data Science and Analysis.

References

Elisa Bassignana, Valerio Basile, and Viviana Patti. 2018. [Hurtlex: A multilingual lexicon of words to hurt](#). In *Proceedings of the Fifth Italian Conference on Computational Linguistics (CLiC-it 2018), Torino, Italy, December 10-12, 2018*, volume 2253 of *CEUR Workshop Proceedings*. CEUR-WS.org.

Hugo Berg, Siobhan Hall, Yash Bhalgat, Hannah Kirk, Aleksandar Shtedritski, and Max Bain. 2022. [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. Association for Computational Linguistics.

Federico Bianchi, Amanda Cercas Curry, and Dirk Hovy. 2023a. [Artificial Intelligence accidents waiting to happen?](#) *Journal of Artificial Intelligence Research*, 76:193–199.

Federico Bianchi, Pratyusha Kalluri, Esin Durmus, Faisal Ladhak, Myra Cheng, Debora Nozza, Tatsunori Hashimoto, Dan Jurafsky, James Zou, and Aylin Caliskan. 2023b. [Easily accessible text-to-image generation amplifies demographic stereotypes](#)

at large scale. In *2023 ACM Conference on Fairness, Accountability, and Transparency, FAccT '23*, New York, NY, USA. Association for Computing Machinery.

- Su Lin Blodgett, Solon Barocas, Hal Daumé III, and Hanna Wallach. 2020. [Language \(technology\) is power: A critical survey of “bias” in NLP](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5454–5476, Online. Association for Computational Linguistics.
- Tolga Bolukbasi, Kai-Wei Chang, James Y. Zou, Venkatesh Saligrama, and Adam Tauman Kalai. 2016. [Man is to computer programmer as woman is to homemaker? debiasing word embeddings](#). In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 4349–4357.
- Joy Buolamwini and Timnit Gebru. 2018. [Gender shades: Intersectional accuracy disparities in commercial gender classification](#). In *Conference on Fairness, Accountability and Transparency, FAT 2018, 23-24 February 2018, New York, NY, USA*, volume 81 of *Proceedings of Machine Learning Research*, pages 77–91. PMLR.
- Aylin Caliskan, Joanna J. Bryson, and Arvind Narayanan. 2017. [Semantics derived automatically from language corpora contain human-like biases](#). *Science*, 356(6334):183–186.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Yash Goyal, Tejas Khot, Aishwarya Agrawal, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. 2019. [Making the V in VQA matter: Elevating the role of image understanding in visual question answering](#). *Int. J. Comput. Vis.*, 127(4):398–414.
- Anthony G Greenwald, Debbie E McGhee, and Jordan LK Schwartz. 1998. Measuring individual differences in implicit cognition: the implicit association test. *Journal of personality and social psychology*, 74(6):1464.
- Tatsuya Hiraoka, Sho Takase, Kei Uchiumi, Atsushi Keyaki, and Naoaki Okazaki. 2022. [Word-level perturbation considering word length and compositional subwords](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 3268–3275, Dublin, Ireland. Association for Computational Linguistics.

- Yusuke Hirota, Yuta Nakashima, and Noa Garcia. 2022. [Gender and racial bias in visual question answering datasets](#). In *FACCT '22: 2022 ACM Conference on Fairness, Accountability, and Transparency, Seoul, Republic of Korea, June 21 - 24, 2022*, pages 1280–1292. ACM.
- Dirk Hovy, Federico Bianchi, and Tommaso Fornaciari. 2020. [“You sound just like your father” Commercial Machine Translation systems include stylistic biases](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1686–1690, Online. Association for Computational Linguistics.
- Dirk Hovy and Shannon L. Spruit. 2016. [The social impact of natural language processing](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 591–598, Berlin, Germany. Association for Computational Linguistics.
- Zhicheng Huang, Zhaoyang Zeng, Bei Liu, Dongmei Fu, and Jianlong Fu. 2020. [Pixel-bert: Aligning image pixels with text by deep multi-modal transformers](#). *CoRR*, abs/2004.00849.
- Kushal Kafle and Christopher Kanan. 2017. [Visual question answering: Datasets, algorithms, and future challenges](#). *Computer Vision and Image Understanding*, 163:3–20. Language in Vision.
- Wonjae Kim, Bokyoung Son, and Ildoo Kim. 2021. [Vilt: Vision-and-language transformer without convolution or region supervision](#). In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pages 5583–5594. PMLR.
- Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A. Shamma, Michael S. Bernstein, and Li Fei-Fei. 2017. [Visual genome: Connecting language and vision using crowdsourced dense image annotations](#). *Int. J. Comput. Vision*, 123(1):32–73.
- Junnan Li, Dongxu Li, Caiming Xiong, and Steven C. H. Hoi. 2022. [BLIP: bootstrapping language-image pre-training for unified vision-language understanding and generation](#). In *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA*, volume 162 of *Proceedings of Machine Learning Research*, pages 12888–12900. PMLR.
- Junnan Li, Ramprasaath R. Selvaraju, Akhilesh Gotmare, Shafiq R. Joty, Caiming Xiong, and Steven Chu-Hong Hoi. 2021. [Align before fuse: Vision and language representation learning with momentum distillation](#). In *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 9694–9705.
- Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. 2019. [Visualbert: A simple and performant baseline for vision and language](#). *ArXiv preprint*, abs/1908.03557.
- Tao Li, Tushar Khot, Daniel Khashabi, Ashish Sabharwal, and Vivek Srikumar. 2020. [Uncovering stereotyping biases via underspecified questions](#). *CoRR*, abs/2010.02428.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. 2014. [Microsoft coco: Common objects in context](#). In *Computer Vision – ECCV 2014*, pages 740–755, Cham. Springer International Publishing.
- Moin Nadeem, Anna Bethke, and Siva Reddy. 2021. [StereoSet: Measuring stereotypical bias in pretrained language models](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 5356–5371, Online. Association for Computational Linguistics.
- Nikita Nangia, Clara Vania, Rasika Bhalerao, and Samuel R. Bowman. 2020. [CrowS-pairs: A challenge dataset for measuring social biases in masked language models](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1953–1967, Online. Association for Computational Linguistics.
- Debora Nozza, Federico Bianchi, and Dirk Hovy. 2021. [HONEST: Measuring hurtful sentence completion in language models](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2398–2406, Online. Association for Computational Linguistics.
- Debora Nozza, Federico Bianchi, and Dirk Hovy. 2022a. [Pipelines for social bias testing of large language models](#). In *Proceedings of BigScience Episode #5 – Workshop on Challenges & Perspectives in Creating Large Language Models*, pages 68–74, virtual+Dublin. Association for Computational Linguistics.
- Debora Nozza, Federico Bianchi, Anne Lauscher, and Dirk Hovy. 2022b. [Measuring harmful sentence completion in language models for LGBTQIA+ individuals](#). In *Proceedings of the Second Workshop on Language Technology for Equality, Diversity and Inclusion*, pages 26–34, Dublin, Ireland. Association for Computational Linguistics.
- Orestis Papakyriakopoulos, Simon Hegelich, Juan Carlos Medina Serrano, and Fabienne Marco. 2020. [Bias in word embeddings](#). In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency, FAT* '20*, page 446–457, New York, NY, USA. Association for Computing Machinery.

- Marcelo O. R. Prates, Pedro H. C. Avelar, and Luís C. Lamb. 2020. [Assessing gender bias in machine translation: a case study with google translate](#). *Neural Comput. Appl.*, 32(10):6363–6381.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. [Learning transferable visual models from natural language supervision](#). In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pages 8748–8763. PMLR.
- Rasmus Rothe, Radu Timofte, and Luc Van Gool. 2015. [DEX: deep expectation of apparent age from a single image](#). In *2015 IEEE International Conference on Computer Vision Workshop, ICCV Workshops 2015, Santiago, Chile, December 7-13, 2015*, pages 252–257. IEEE Computer Society.
- Fawaz Sammani, Tanmoy Mukherjee, and Nikos Deligiannis. 2022. [NLX-GPT: A model for natural language explanations in vision and vision-language tasks](#). *CoRR*, abs/2203.05081.
- Deven Santosh Shah, H. Andrew Schwartz, and Dirk Hovy. 2020. [Predictive biases in natural language processing models: A conceptual framework and overview](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5248–5264, Online. Association for Computational Linguistics.
- Emily Sheng, Kai-Wei Chang, Prem Natarajan, and Nanyun Peng. 2021. [Societal biases in language generation: Progress and challenges](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4275–4293, Online. Association for Computational Linguistics.
- Tejas Srinivasan and Yonatan Bisk. 2022. [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, Seattle, Washington. Association for Computational Linguistics.
- Hao Tan and Mohit Bansal. 2019. [LXMERT: Learning cross-modality encoder representations from transformers](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5100–5111, Hong Kong, China. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008.
- Jesse Vig, Sebastian Gehrmann, Yonatan Belinkov, Sharon Qian, Daniel Nevo, Yaron Singer, and Stuart M. Shieber. 2020. [Investigating gender bias in language models using causal mediation analysis](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Peng Wang, An Yang, Rui Men, Junyang Lin, Shuai Bai, Zhikang Li, Jianxin Ma, Chang Zhou, Jingren Zhou, and Hongxia Yang. 2022. [Opa: Unifying architectures, tasks, and modalities through a simple sequence-to-sequence learning framework](#).
- Yi Zhang, Junyang Wang, and Jitao Sang. 2022. [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. Association for Computing Machinery.
- Zhifei Zhang, Yang Song, and Hairong Qi. 2017. [Age progression/regression by conditional adversarial autoencoder](#). In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 4352–4360. IEEE Computer Society.
- Kankan Zhou, Eason Lai, and Jing Jiang. 2022. [VL-StereoSet: A study of stereotypical bias in pre-trained vision-language models](#). 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 527–538, Online only. Association for Computational Linguistics.

A Bias in fine-tuned VL downstream models

A.1 Templates

In order to avoid any strong bias with the question templates, we designed 5 different templates:

- Does this person value [TARGET]?
- Does this person like [TARGET]?
- Does this person enjoy [TARGET]?
- Does this person appreciate [TARGET]?
- Does this person encourage [TARGET]?

where [TARGET] refers to the target concepts proposed in (Caliskan et al., 2017). Therefore some examples of questions are: "Does this person encourage marriage?", "Does this person like algebra?".

Model	Social group	<i>Pleasant</i>	<i>Unpleasant</i>	<i>Arts</i>	<i>Math</i>	<i>Career</i>	<i>Family</i>	<i>Mental Disease</i>	<i>Physical Disease</i>
ViLT	Ethnicity	White	Black	Asian	Indian	Indian	Asian	Asian	Indian
BLIP	Ethnicity	Asian	Black	Asian	Other	Indian	Black	Asian	White
OFA	Ethnicity	Other	Asian	Asian	White	Indian	Black	Black	White
NLX-GPT	Ethnicity	Black	Other	Asian	White	Black	Asian	Other	Indian
ViLT	Age	25-54	65+	65+	1-14	55-64	1-14	65+	1-14
BLIP	Age	1-14	65+	15-24	55-64	1-14	55-64	1-14	65+
OFA	Age	25-54	65+	1-14	65+	55-64	1-14	1-14	55-64
NLX-GPT	Age	25-54	55-64	55-64	1-14	55-64	1-14	1-14	15-24

Table 6: The most associated ethnical and age groups by model and bias concept

A.2 Additional Results

The most associated age and ethnic groups by model and bias concept are shown in Table 6.

ACL 2023 Responsible NLP Checklist

A For every submission:

- A1. Did you describe the limitations of your work?
Limitations section
- A2. Did you discuss any potential risks of your work?
Ethical Statement section
- A3. Do the abstract and introduction summarize the paper's main claims?
Abstract and Introduction
- A4. Have you used AI writing assistants when working on this paper?
Left blank.

B Did you use or create scientific artifacts?

2

- B1. Did you cite the creators of artifacts you used?
2
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
GitHub webpage
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?
2
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?
Not applicable. Left blank.
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
2
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.
2

C Did you run computational experiments?

Left blank.

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?
No response.

The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?

No response.

- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?

No response.

- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?

No response.

D Did you use human annotators (e.g., crowdworkers) or research with human participants?

Left blank.

- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?

No response.

- D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?

No response.

- D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?

No response.

- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?

No response.

- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?

No response.