MiSCHiEF: A Benchmark in Minimal-Pairs of Safety and Culture for Holistic Evaluation of Fine-Grained Image-Caption Alignment

Sagarika Banerjee[1*] **Tangatar Madi**[1*] **Advait Swaminathan**[1*]
Nguyen Dao Minh Anh[1*] **Shivank Garg**[1*] **Kevin Zhu**[1] **Vasu Sharma**[1,2]
[1]Algoverse AI Research [2]PocketFM
shivank@algoverseairesearch.org, kevin@algoverse.us

Abstract

Fine-grained image-caption alignment is crucial for vision-language models (VLMs), especially in socially critical contexts such as identifying real-world risk scenarios or distinguishing cultural proxies, where correct interpretation hinges on subtle visual or linguistic clues and where minor misinterpretations can lead to significant real-world consequences. We present MiSCHiEF, a set of two benchmarking datasets based on a contrastive pair design in the domains of safety (MiS) and culture (MiC), and evaluate four VLMs on tasks requiring fine-grained differentiation of paired images and captions. In both datasets, each sample contains two minimally differing captions and corresponding minimally differing images. In MiS, the image-caption pairs depict a safe and an unsafe scenario, while in MiC, they depict cultural proxies in two distinct cultural contexts. We find that models generally perform better at confirming the correct image-caption pair than rejecting incorrect ones. Additionally, models achieve higher accuracy when selecting the correct caption from two highly similar captions for a given image, compared to the converse task. The results, overall, highlight persistent modality misalignment challenges in current VLMs, underscoring the difficulty of precise cross-modal grounding required for applications with subtle semantic and visual distinctions.

1 Introduction

Fine-grained image-caption alignment is a crucial component of robust visuo-linguistic compositional reasoning, enabling models to perform effectively in socially critical contexts such as visual risk assessment, where they learn to identify possible dangers in images, and cultural context reasoning, where understanding scenes relies on knowledge from diverse cultures and regions (Yin et al., 2021).

Previous works have explored visuo-linguistic compositional reasoning in different ways. Natural Language Visual Reasoning for Real (NLVR2) (Suhr et al., 2019) tests whether a natural language caption is true about a pair of images, requiring models to resolve subtle mismatches in attributes and relations. More recent works have studied image-caption alignment by testing whether models can correctly match two images with two captions. Winoground (Thrush et al., 2022) presents captions with identical words in different orders, alongside images that represent those captions with pronounced visual differences. VisMin (Awal et al., 2024) ensures minimal changes between both image and caption pairs, altering only one aspect at a time, such as object, attribute, count, or spatial relation. While valuable for probing visuo-linguistic compositional reasoning abilities of VLMs, existing benchmarks remain domain-agnostic and thus fail to capture the unique challenges posed by safety- and culture-sensitive contexts, limiting their effectiveness for evaluating model robustness in these critical areas.

Previously, several datasets have been proposed to evaluate models on safety and cultural reasoning. Safety-focused datasets include UnsafeBench (Qu et al., 2024), which evaluates image safety classifiers across eleven risk categories, and Incidents1M (Weber et al., 2022), which collects disaster-related social media images for incident classification. Enhancing Surveillance Systems (Jeon et al., 2024) introduces a dataset of surveillance images paired with structured captions and risk scores (1–7). The HBDset (Ding et al., 2024) focuses on using computer vision for evacuation safety and emergency management.

Cultural reasoning has been explored through benchmarks like CVQA (Romero et al., 2024), a multilingual dataset with more than 10,000 questions from 30 countries that cover traditions, artifacts, and more. SEA-VQA (Urailertprasert et al.,

[*]Primary authors

Proceedings of the 19th Conference of the European Chapter of the Association for Computational Linguistics
Volume 2: Short Papers, pages 384–406

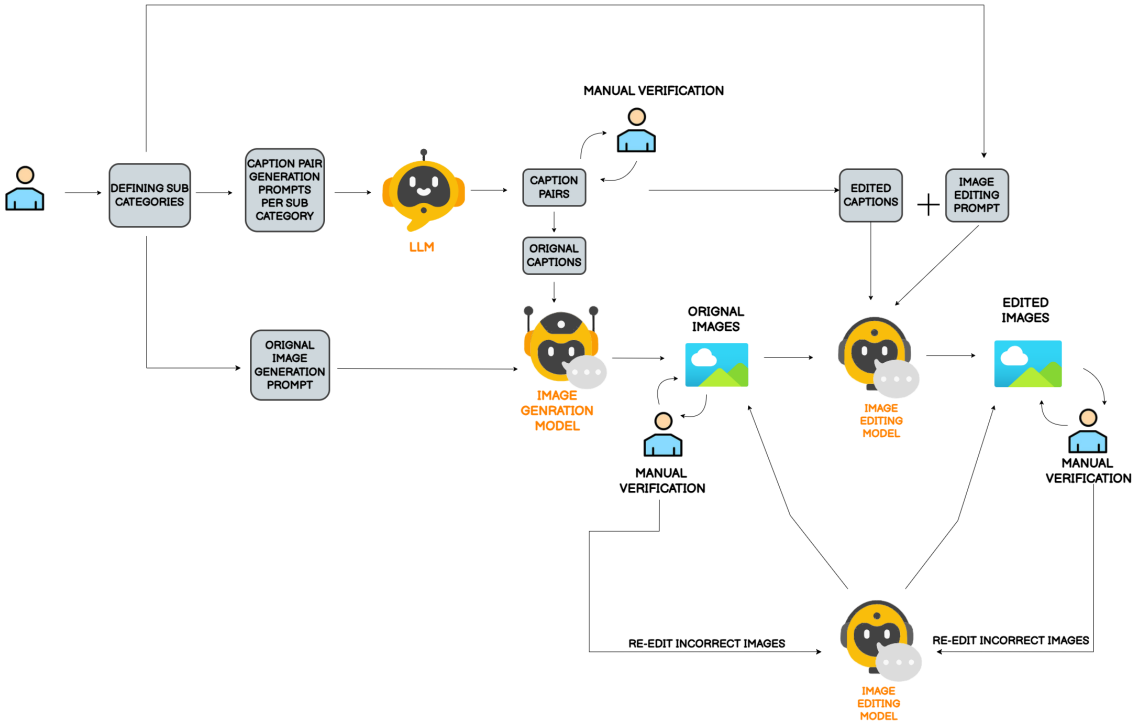

Figure 1: Curation pipeline for MiS and MiC: LLM-generated caption pairs are verified, used for image generation and editing, and manually refined. The complete generation pipeline is detailed in Appendix 3. Example entries from the dataset are shown in Fig. 2.

2024) complements this work by focusing specifically on 8 Southeast Asian countries.

However, safety and culture data sets typically prioritize broad coverage over minimal-pair contrasts, which are essential to accurately evaluate VLM's ability to distinguish subtle visual and/or linguistic differences critical for correct interpretation in nuanced contexts. To address these limitations, we make the following key contributions:

- We introduce MiSCHiEF, a unified benchmark that integrates two novel components, MiS (**Mi**nimal-pairs in **S**afety) and MiC (**Mi**nimal-pairs in **C**ulture), to evaluate fine-grained image–caption alignment. By bringing together these societally critical domains, MiSCHiEF probes a core limitation of vision–language models (VLMs): interpreting subtle visual and contextual cues where small errors can have outsized real-world consequences.
- We expose systematic image–text misalignments in current VLMs through four diagnostic tasks. Our analysis shows that models are generally better at confirming correct image–caption pairs than at rejecting incorrect ones, revealing an inherent bias in multimodal models.
- We uncover fundamental asymmetries in multimodal understanding and cross-modal alignment. Models achieve higher accuracy when selecting the correct caption for a given image than when performing the reverse task, and their performance drops sharply in dual alignment settings requiring the correct pairing of multiple images and captions.

2 Experiments

We designed four experiments to evaluate the capacity of vision–language models (VLMs) for fine-grained visuo-linguistic reasoning. Each experiment targeted a distinct aspect of image-caption alignment.

In the first experiment, Caption-to-Image Matching (C2I), the model was provided with one randomly selected caption and two images per sample, and its task was to identify which image correctly corresponded to the given caption. The second experiment, Dual Caption–Image Alignment (DCI), presented the model with both captions and both images per sample, requiring it to correctly match each caption to its corresponding image. The third experiment, Pairwise Consistency Evaluation (PC), involved a binary classification task in which the

Figure 2: Examples from MiSCHiEF illustrating minimal pairs in MiS and MiC.

model was prompted to respond with "Yes" if the caption accurately described the image and "No" otherwise. The experimental designs for both MiC and MiS datasets under this setting are summarized in Table 1. Finally, in the fourth experiment, Image-to-Caption Matching (I2C), the model was provided with one randomly selected image and two captions per sample, and it was required to select the caption that best described the given image.

3 MiS and MiC Dataset Curation

We adopted a two-stage process for the curation of MiSCHiEF:

Caption Pair Generation & Verification: Subcategories were manually generated following prior literature ensuring MiS addressed diverse risk scenarios (Li et al., 2025; Huang and Cheng, 2025; Ahmad and Rahimi, 2025; Malla et al., 2022; García-Domínguez et al., 2021) and MiC captured diverse aspects of culture via proxies (Adilazuarda et al., 2024). The captions for these were then generated using Gemini 2.5 Pro. The final distributions for each sub-category in the final dataset are shown in Fig. 3b in the appendix. To ensure diversity among generated prompts, near-duplicates were removed using Jaccard similarity (3-gram, 4-gram, threshold 0.8) and Sentence Transformer similarity ($\geq$0.9). Followed by manual verification to ensure correctness and no ambiguity.

Image Pair Generation & Verification: From each original caption, an image was generated and then edited to reflect its paired caption, while preserving global scene attributes, using the GPT-Image API. This was followed by a step of manual verification, which focused on cultural accuracy in MiC and clarity regarding safe/unsafe situations in MiS. Erroneous samples were re-edited using GPT-Image once and disregarded in case of any errors. After manual verification and refinement, MiC consists of 279 samples, and MiS consists of 190 samples. MiSCHiEF is designed as a diagnostic evaluation benchmark, similar in purpose and scale to Winoground (Thrush et al., 2022) (400 samples). The deliberate focus on high-quality, meticulously verified minimal pairs, which require nuanced human oversight to eliminate ambiguity, necessarily constrains the dataset's scale but ensures its reliability as a fine-grained evaluation benchmark that prioritizes quality over quantity. Examples of the dataset are shown in Fig. 2 and Appendix E.

4 Results

Caption-to-Image (C2I) and Image-to-Caption Matching (I2C): As shown in Table 2, model performance varies across tasks. For MiC, most models exceed random chance, except `Llava-Next-Video` in Caption-to-Image Matching and both `InternVL` and `Llava-Next-Video` in Image-to-Caption Matching. For MiS, models perform only marginally above chance, with `Llava-Next-Video` underperforming in Caption-to-Image and `Qwen-3B` in Image-to-Caption. Across datasets, accuracies are generally higher (by ~20-30%) on Image-to-Caption Matching than

Dataset	Pairing Type	Caption	Image	Expected Model Response
MiS	Congruent$_A$ (Con$_A$)	Safe	Safe	Yes
	Incongruent$_A$ (Inc$_A$)	Safe	Unsafe	No
	Congruent$_B$ (Con$_B$)	Unsafe	Unsafe	Yes
	Incongruent$_B$ (Inc$_B$)	Unsafe	Safe	No
MiC	Congruent$_A$ (Con$_A$)	Culture A	Culture A	Yes
	Incongruent$_A$ (Inc$_A$)	Culture A	Culture B	No
	Congruent$_B$ (Con$_B$)	Culture B	Culture B	Yes
	Incongruent$_B$ (Inc$_B$)	Culture B	Culture A	No

Table 1: Experimental setup for the Pairwise Consistency (PC) Evaluation. For both the MiS (safety) and MiC (culture) datasets, "Congruent" pairs refer to correctly matched image-caption pairs, while "Incongruent" pairs refer to deliberately mismatched pairs. The subscripts distinguish between the two minimal-pair items within a sample; for instance, in MiS, Con$_A$ represents a "Safe" caption correctly paired with a "Safe" image, whereas Inc$_A$ represents the same "Safe" caption incorrectly paired with an "Unsafe" image.

		C2I	DCI	PC				I2C
				Con$_A$	Inc$_A$	Con$_B$	Inc$_B$	
MiC	Qwen 3B	62.72	47.31	99.64	66.67	98.57	58.42	87.46
	InternVL	70.61	41.58	86.38	66.67	86.74	64.16	37.99
	Phi 3.5	82.80	57.71	98.21	57.71	97.13	41.94	79.93
	Llava-Next-Video	47.67	50.18	100.00	0.00	100.00	0.00	28.74
	GPT-4o	**93.24**	**57.47**	83.16	**86.24**	76.32	**82.75**	86.24
	Random Chance	50.00	25.00	50.00	50.00	50.00	50.00	50.00
MiS	Qwen 3B	54.50	49.21	79.89	41.80	97.88	78.31	47.62
	InternVL	58.95	51.05	34.21	21.05	77.37	83.16	87.37
	Phi 3.5	50.53	44.21	86.84	60.00	31.05	96.32	81.58
	Llava-Next-Video	45.26	43.68	96.84	27.37	84.74	64.21	79.47
	GPT-4o	**93.15**	**57.89**	82.76	**88.94**	77.59	83.15	**85.78**
	Random Chance	50.00	25.00	50.00	50.00	50.00	50.00	50.00

Table 2: Results on MiC and MiS datasets across C2I, DCI, PC, and I2C tasks described in Section 4. Models perform better on congruent than incongruent cases, with overall higher accuracy on MiC.

on Caption-to-Image Matching, suggesting models are more sensitive to semantic differences between captions than to subtle visual differences between images. Performance is also higher on MiC than MiS, likely due to the more pronounced distinctions in MiC.

Dual Caption–Image Alignment (DCI): Dual Caption-Image Alignment proves especially challenging, with peak accuracies of 57.71% (MiC) and 51.05% (MiS), notably lower than in the simpler matching tasks. For instance, `Qwen-3B` achieves 47.31% on this task but achieves an accuracy of 62.72% and 87.46% on Caption-to-Image and Image-to-Caption Matching respectively.

Pairwise Consistency (PC): In MiC, `Llava-Next-Video` outputs trivial answers, yielding extreme scores. Other models show strong accuracies (>85%) on matched pairs (Con$_A$, Con$_B$) but weaker results on mismatched ones (Inc$_A$, Inc$_B$). For MiS, models also excel at confirming matches, but show mixed reliability in rejecting mismatches. Overall, current VLMs appear better at validating true pairs than identifying subtle mismatches, highlighting a limitation in fine-grained negative reasoning. Notably, GPT-4o achieves more balanced performance across congruent and incongruent pairs, suggesting that larger closed-source models exhibit improved

negative reasoning capabilities.

Key Findings for VLM Development: Our benchmark reveals three systematic weaknesses with direct implications for VLM development: (1) *Confirmation Bias*: Models are significantly better at confirming correct image-caption pairs than rejecting incorrect ones, suggesting that models lack robust negative reasoning, a critical capability for real-world deployment where false positives in safety contexts are particularly dangerous. (2) *Modality Asymmetry*: Models consistently achieve higher accuracy on I2C than C2I tasks, indicating an uneven balance between visual and linguistic grounding that developers should address. Ideally, if models have true cross-modal alignment, performance should be symmetric. (3) *Dual Alignment Failure*: When matching multiple images with multiple captions simultaneously (DCI task), performance drops substantially even for GPT-4o ($\sim$57%), revealing that models cannot maintain consistent reasoning across multiple cross-modal pairs, essential for complex real-world scenarios.

5 Discussion

The modality misalignment patterns we expose are not merely abstract theoretical problems—they have outsized real-world consequences in safety and cultural contexts. In the safety domain, when a VLM fails to distinguish between "A woman is plugging a lamp into an outlet" and "A woman is plugging a fork into an outlet," this misalignment directly translates to risk of physical harm. A model deployed in home safety monitoring or child supervision could miss life-threatening situations because it cannot ground subtle visual differences. In the culture domain, when a VLM cannot differentiate cultural proxies like "A person wearing a Kente cloth" versus "A person wearing a Poncho," the same misalignment leads to cultural misrepresentation. Models used in content moderation, education, or cross-cultural communication may perpetuate stereotypes or erase cultural identities. This connection between general perceptual limitations and domain-specific harms is precisely why we investigate modality misalignment through the lens of safety and culture: these domains reveal where the stakes are highest and where addressing these limitations is a prerequisite for safe VLM deployment.

6 Conclusion

We introduced MiSCHiEF, a benchmark for fine-grained image-caption alignment in safety- and culture-sensitive contexts. Through the minimal-pair design of MiS and MiC, we revealed persistent modality misalignments in current VLMs, particularly their difficulty in rejecting incorrect image–caption pairs and in performing well on dual alignment tasks involving multiple images and captions. By contrast, models perform relatively better when confirming correct pairs or picking the right caption between highly similar captions to describe a given image, underscoring asymmetries in cross-modal alignment. These results highlight the limitations of current systems in socially critical domains, and position MiSCHiEF as a foundation for developing multimodal models with more precise and context-sensitive grounding.

7 Limitations

Our dataset is a small-scale evaluation benchmark, consisting of 279 cultural pairs and 190 safety pairs. The limited size arises from the need for careful manual verification to ensure high quality and eliminate ambiguity. Expanding the benchmark through semi-automatic or fully automatic pipelines, while preserving reliability, is an important direction for future work. Based on manual analysis by the authors of a subset of our benchmark, all questions were understandable and solvable by humans; however, an exhaustive human evaluation study was not conducted due to budget constraints and the high cost and difficulty of obtaining human reviewers with varied cultural backgrounds. While our work motivates MiSCHiEF in terms of its relevance to safety-critical and cultural contexts, we do not analyze correlations between performance on existing safety benchmarks and MiSCHiEF. This is primarily because most existing benchmarks are limited to single-image evaluations, which differ fundamentally from our pairwise minimal-pair design.

References

Marah Abdin, Jyoti Aneja, Hany Awadalla, Ahmed Awadallah, Ammar Ahmad Awan, Nguyen Bach, Amit Bahree, Arash Bakhtiari, Jianmin Bao, Harkirat Behl, Alon Benhaim, Misha Bilenko, Johan Bjorck, Sébastien Bubeck, Martin Cai, Qin Cai, Vishrav Chaudhary, Dong Chen, Dongdong Chen, and 110 others. 2024. Phi-3 technical report: A highly capable language model locally on your phone. *Preprint*, arXiv:2404.14219.

Muhammad Farid Adilazuarda, Sagnik Mukherjee, Pradhyumna Lavania, Siddhant Singh, Alham Fikri Aji, Jacki O'Neill, Ashutosh Modi, and Monojit Choudhury. 2024. Towards measuring and modeling" culture" in llms: A survey. *arXiv preprint arXiv:2403.15412*.

Hafiz Mughees Ahmad and Afshin Rahimi. 2025. Sh17: A dataset for human safety and personal protective equipment detection in manufacturing industry. *Journal of Safety Science and Resilience*, 6(2):175–185.

Amith Ananthram, Elias Stengel-Eskin, Mohit Bansal, and Kathleen McKeown. 2024. See it from my perspective: How language affects cultural bias in image understanding. *arXiv preprint arXiv:2406.11665*.

Rabiul Awal, Saba Ahmadi, Le Zhang, and Aishwarya Agrawal. 2024. Vismin: Visual minimal-change understanding. *Advances in Neural Information Processing Systems*, 37:107795–107829.

Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Cameron Gonzalez, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, Miles Chen, and 1 others. 2022a. Constitutional ai: Harmlessness from ai feedback. *arXiv preprint arXiv:2212.08073*.

Yuntao Bai, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones, Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, and 1 others. 2022b. Training a helpful and harmless assistant with reinforcement learning from human feedback. In *Advances in Neural Information Processing Systems*, volume 35, pages 29382–29497.

Zhe Chen, Weiyun Wang, Yue Cao, Yangzhou Liu, Zhangwei Gao, Erfei Cui, Jinguo Zhu, Shenglong Ye, Hao Tian, Zhaoyang Liu, and 1 others. 2024. Expanding performance boundaries of open-source multimodal models with model, data, and test-time scaling. *arXiv preprint arXiv:2412.05271*.

Yifei Ding, Xinghao Chen, Zilong Wang, Yuxin Zhang, and Xinyan Huang. 2024. Human behaviour detection dataset (hbdset) using computer vision for evacuation safety and emergency management. *Journal of Safety Science and Resilience*, 5(3):355–364.

Anuj Diwan, Layne Berry, Eunsol Choi, David Harwath, and Kyle Mahowald. 2022. Why is winoground hard? investigating failures in visuolinguistic compositionality. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 2236–2250, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Sri Harsha Dumpala, Aman Jaiswal, Chandramouli Shama Sastry, Evangelos Milios, Sageev Oore, and Hassan Sajjad. 2024. Sugarcrepe++ dataset: Vision-language model sensitivity to semantic and lexical alterations. *Advances in Neural Information Processing Systems*, 37:17972–18018.

A. García-Domínguez, C.E. Galván-Tejada, R.F. Brena, A.A. Aguileta, J.I. Galván-Tejada, H. Gamboa-Rosales, J.M. Celaya-Padilla, and H. Luna-García. 2021. Children's activity classification for domestic risk scenarios using environmental sound and a bayesian network. *Healthcare (Basel)*, 9(7):884.

Amelia Glaese, Nat McAleese, Mateusz Trębacz, John Aslanides, Vlad Firoiu, Timo Ewalds, Maribeth Rauh, Laura Weidinger, Iason Gabriel, Zachary Kenton, and 1 others. 2022. Improving alignment of dialogue agents via targeted human judgements. *arXiv preprint arXiv:2209.14375*.

Arnav Gupta, Rahul Jha, Divyanshu Singh, Antonios Anastasopoulos, and Monojit Choudhury. 2024. Malibu: A benchmark for multilingual persona-grounded cultural reasoning in large language models. *arXiv preprint arXiv:2401.08527*.

Daniel Hershcovich, Stella Frank, Heather Lent, Miryam de Lhoneux, Mostafa Abdou, Stephanie Brandl, Emanuele Bugliarello, Laura Cabello Piqueras, Ilias Chalkidis, Ruixiang Cui, Constanza Fierro, Katerina Margatina, Phillip Rust, and Anders Søgaard. 2022. Challenges and strategies in cross-cultural NLP. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6997–7013, Dublin, Ireland. Association for Computational Linguistics.

Cheng-Yu Hsieh, Jieyu Zhang, Zixian Ma, Aniruddha Kembhavi, and Ranjay Krishna. 2023. Sugarcrepe: Fixing hackable benchmarks for vision-language compositionality. *ArXiv*, abs/2306.14610.

Mei-Ling Huang and Ying Cheng. 2025. Dataset of personal protective equipment (ppe).

M. Jeon, J. Ko, and K. Cheoi. 2024. Enhancing surveillance systems: Integration of object, behavior, and space information in captions for advanced risk assessment. *Sensors (Basel)*, 24(1):292.

Mario Kovač, Michael Moosmüller, Stjepan Marjanovic, and Miloš Stanojević. 2023. Llms as cultural personas: Benchmarking persona-steered value judgments across cultures. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 13815–13828. Association for Computational Linguistics.

Baiqi Li, Zhiqiu Lin, Wenxuan Peng, Jean de Dieu Nyandwi, Daniel Jiang, Zixian Ma, Simran Khanuja, Ranjay Krishna, Graham Neubig, and Deva Ramanan. 2024. Naturalbench: Evaluating vision-language models on natural adversarial samples. *Advances in Neural Information Processing Systems*, 37:17044–17068.

Wei Li, Fuqi Ma, Zhiyuan Zuo, Rong Jia, Bo Wang, and Abdullah M Alharbi. 2025. Safetygpt: An autonomous agent of electrical safety risks for monitoring workers' unsafe behaviors. *International Journal of Electrical Power & Energy Systems*, 168:110672.

Srikanth Malla, Chiho Choi, Isht Dwivedi, Joon Hee Choi, and Jiachen Li. 2022. Drama: Joint risk localization and captioning in driving. *Preprint*, arXiv:2209.10767.

Shravan Nayak, Kanishk Jain, Rabiul Awal, Siva Reddy, Sjoerd Van Steenkiste, Lisa Anne Hendricks, Aishwarya Agrawal, and 1 others. 2024. Benchmarking vision language models for cultural understanding. *arXiv preprint arXiv:2407.10920*.

Shramay Palta and Rachel Rudinger. 2023. FORK: A bite-sized test set for probing culinary cultural biases in commonsense reasoning models. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 9952–9962, Toronto, Canada. Association for Computational Linguistics.

Yiting Qu, Xinyue Shen, Yixin Wu, Michael Backes, Savvas Zannettou, and Yang Zhang. 2024. Unsafebench: Benchmarking image safety classifiers on real-world and ai-generated images. *arXiv preprint arXiv:2405.03486*.

David Romero, Chenyang Lyu, Haryo Akbarianto Wibowo, Teresa Lynn, Injy Hamed, Aditya Nanda Kishore, Aishik Mandal, Alina Dragonetti, Artem Abzaliev, Atnafu Lambebo Tonja, and 1 others. 2024. Cvqa: Culturally-diverse multilingual visual question answering benchmark. *arXiv preprint arXiv:2406.05967*.

Jesper Sorensen, Toby Shevlane, Jess Whittlestone, and 1 others. 2023. Value pluralism in large language models. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*, pages 314–325. ACM.

Alane Suhr, Stephanie Zhou, Ally Zhang, Iris Zhang, Huajun Bai, and Yoav Artzi. 2019. A corpus for reasoning about natural language grounded in photographs. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6418–6428, Florence, Italy. Association for Computational Linguistics.

Sinha Tanmay, Toby Shevlane, Iason Gabriel, Laura Weidinger, Lisa Anne Hendricks, and 1 others. 2023. Value kaleidoscope: Engaging llms with diverse human values. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.

Qwen Team. 2025. Qwen2.5-vl.

Bill Thompson, {Seán G.} Roberts, and Gary Lupyan. 2020. Cultural influences on word meanings revealed through large-scale semantic alignment. *Nature Human Behaviour*, 4:1029–1038(2020). The acceptance date for this record is provisional and based upon the month of publication for the article.

Tristan Thrush, Ryan Jiang, Max Bartolo, Amanpreet Singh, Adina Williams, Douwe Kiela, and Candace Ross. 2022. Winoground: Probing vision and language models for visio-linguistic compositionality. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5238–5248.

Norawit Urailertprasert, Peerat Limkonchotiwat, Supasorn Suwajanakorn, and Sarana Nutanong. 2024. SEA-VQA: Southeast Asian cultural context dataset for visual question answering. In *Proceedings of the 3rd Workshop on Advances in Language and Vision Research (ALVR)*, pages 173–185, Bangkok, Thailand. Association for Computational Linguistics.

Ethan Weber, Dim P Papadopoulos, Agata Lapedriza, Ferda Ofli, Muhammad Imran, and Antonio Torralba. 2022. Incidents1m: a large-scale dataset of images with natural disasters, damage, and incidents. *IEEE transactions on pattern analysis and machine intelligence*, 45(4):4768–4781.

Winston Wu, Lu Wang, and Rada Mihalcea. 2023. Cross-cultural analysis of human values, morals, and biases in folk tales. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 5113–5125, Singapore. Association for Computational Linguistics.

Da Yin, Liunian Harold Li, Ziniu Hu, Nanyun Peng, and Kai-Wei Chang. 2021. Broaden the vision: Geo-diverse visual commonsense reasoning. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 2115–2129, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

Zonghao Ying, Aishan Liu, Siyuan Liang, Lei Huang, Jinyang Guo, Wenbo Zhou, Xianglong Liu, and Dacheng Tao. 2024. Safebench: A safety evaluation framework for multimodal large language models. *arXiv preprint arXiv:2410.18927*.

Jianshu Zhang, Dongyu Yao, Renjie Pi, Paul Pu Liang, and Yi R Fung. 2025. Vlm2-bench: A closer look at how well vlms implicitly link explicit matching visual cues. *arXiv preprint arXiv:2502.12084*.

Yuanhan Zhang, Bo Li, haotian Liu, Yong jae Lee, Liangke Gui, Di Fu, Jiashi Feng, Ziwei Liu, and Chunyuan Li. 2024. Llava-next: A strong zero-shot video understanding model.

Appendix

A Related Works

Our work is situated at the intersection of three key research areas: visuo-linguistic compositional reasoning, safety evaluation for multimodal models, and the growing field of cultural reasoning in AI. We review relevant literature in each of these domains to contextualize the unique contributions of the MiSCHiEF benchmark.

A.1 Visuo-Linguistic Compositional Reasoning

Evaluating the ability of Vision-Language Models (VLMs) to understand the compositional structure of language and vision is a critical area of research. A prominent approach in this domain is the use of minimal-pair benchmarks, which test models on pairs of images and captions that differ in subtle but meaningful ways. The seminal Winoground dataset (Thrush et al., 2022) challenges models to match captions with identical words in different orders to images with significant visual differences. Subsequent analysis revealed that the difficulty of Winoground stems not only from compositional language understanding but also from challenges in fusing visual and textual representations and identifying small or out-of-focus objects (Diwan et al., 2022).

Building on this paradigm, other benchmarks have emerged to probe different facets of compositionality. For example, SugarCrepe (Hsieh et al., 2023) and its successor SugarCrepe++ (Dumpala et al., 2024) were developed to provide more robust evaluations by fixing "hackable" elements in previous datasets and testing sensitivity to both semantic and lexical alterations. Similarly, benchmarks like VLM2-Bench examine how well VLMs implicitly link explicit visual cues in an image (Zhang et al., 2025). While these datasets are invaluable for assessing general reasoning, they are largely domain-agnostic. They do not specifically target the socially critical contexts of safety and culture, where nuanced understanding is paramount. MiSCHiEF fills this gap by applying the rigorous minimal-pair design to these specific domains, forcing models to reason about subtle changes that have significant real-world implications.

A.2 Safety Benchmarks for Vision-Language Models

As VLMs become more integrated into real-world applications, ensuring their safety and alignment with human values is crucial. This has led to the development of various benchmarks aimed at evaluating model safety.

More specific to multimodal models, benchmarks like SafeBench (Ying et al., 2024) provide a comprehensive framework for evaluating safety across various categories, similar to the goals of UnsafeBench mentioned in our introduction. Other works, such as NaturalBench (Li et al., 2024), evaluate VLM robustness against natural adversarial samples that can often expose model vulnerabilities. While these benchmarks are essential for identifying broad safety failures (e.g., detecting violent content or hate speech), they typically focus on classifying distinct, often overt, categories of risk. They do not systematically test a model's ability to differentiate between a safe and an unsafe scenario based on a minimal, fine-grained visual or textual change, which our MiSCHiEF safety dataset is designed to address. Our work complements these efforts by probing the model's visuo-linguistic reasoning within the domain of safety.

A.3 Cultural Reasoning in AI

There is a growing recognition that intelligent systems must understand and respect diverse cultural contexts. A recent survey highlights ongoing efforts in measuring and modeling "culture" within LLMs (Adilazuarda et al., 2024), with studies exploring cultural biases through folk tales (Wu et al., 2023) and culinary customs (Palta and Rudinger, 2023). Broader socio-cultural work has examined safety and value alignment (Glaese et al., 2022; Bai et al., 2022b,a), showing how methods like RLHF and constitutional AI embed cultural norms. Persona-based benchmarks such as MAL-IBU (Gupta et al., 2024) and related evaluations (Kovač et al., 2023) test models when adopting cultural identities, while others probe how LLMs navigate dilemmas in value pluralism (Tanmay et al., 2023; Sorensen et al., 2023). Much of this literature relies on cultural 'proxies,' such as demographic factors (e.g., ethnicity, religion, gender, region) or semantic cues (e.g., food, etiquette, values), yet many important facets remain untested. The paper (Thompson et al., 2020) emphasizes overlooked domains such as kinship, spatial relations, and cogni-

tion, and also(Hershcovich et al., 2022) highlights the neglected dimension of aboutness, i.e. whether a model can identify what a text is fundamentally about.

In the vision-language domain, several benchmarks have been created to evaluate cultural understanding. CVQA (Romero et al., 2024) provides a multilingual dataset covering global clothing, food, and festivals, while other works benchmark cultural reasoning in VLMs (Nayak et al., 2024) or study how language shapes cultural bias in image interpretation (Ananthram et al., 2024). These datasets test recognition of cultural artifacts and practices but do not assess reasoning about how minor contextual variations influence cultural interpretation.

The MiSCHiEF culture dataset addresses this gap by applying a minimal-pair format where the same cultural proxy appears in two distinct contexts, requiring more nuanced reasoning that moves beyond surface-level recognition.

B Implementation Details

We evaluate four state-of-the-art small multimodal VLMs representing diverse architectures. `InternVL2_5-8B` (Chen et al., 2024), `LLava-Next-Video-7B` (Zhang et al., 2024), `Qwen2.5-VL-3B-Instruct` (Team, 2025) and `Phi-3.5-vision-instruct` (Abdin et al., 2024), where the text generation was performed using the default HuggingFace generation hyperparameters. All our experiments were conducted on a node with a single A100 80GB GPU, for a single random seed. Across all the experiments, we use accuracy as the primary evaluation metric. Furthermore, all manual annotations were conducted by the authors.

C Dataset Statistics

The category wise data statistics for MiS and MiC are shown in Figure 3.

D Prompts

D.1 Caption Pair Generation Prompts for MiC

General Activities

You are an AI assistant tasked with generating creative and culturally grounded caption pairs. Your job is to produce pairs of captions that strictly follow the minimal pair principle described below. The caption pairs must be textually almost identical except for a specific, swapped–out keyword related to general activity.

Each pair must contain:
1. "Original caption": A short caption describing a specific action set in a clearly identified country context.
2. "Edited caption": The exact same caption, but with the country name replaced with an equivalent from a different culture.

The Minimal Pair Principle: This is the most important rule. The sentence structure, verbs, adjectives, and all non–cultural descriptors in the "original" and "edited" prompts must remain identical. For this task, the only change allowed is the direct substitution of the country name.

Categories for Substitution:
Your keyword substitutions should fall into one or more of the following categories, emphasizing plausibility and cultural relevance:
In this category, only the country context is replaced, while the underlying activity remains the same. Prompts must avoid mentioning or describing culturally–exclusive activities (e.g., traditional Water Puppet (Mua roi nuoc) performance in Vietnam) that would be nonsensical if moved to another country. The aim is for the scene to be realistically and authentically re–contextualized just by changing the country name.

Cultural Diversity & Authenticity Requirements:
– Draw from as many diverse cultures as possible across all continents.
– Include underrepresented cultures and regions, not just commonly featured ones.
– Ensure all cultural references generated in an image would be authentic, accurate, and respectful.
– Avoid cultural appropriation or inaccurate generalizations.

Holidays and Celebrations

You are an AI assistant tasked with generating creative and culturally grounded image prompts. Your job is to produce pairs of captions that strictly follow the minimal pair principle described below. The caption pairs must be textually almost identical except for a specific, swapped–out keyword related to holiday and country.

Each pair must contain:
1. "Original caption": A short caption describing a specific object, symbol, action, or decoration (such as food, clothing, or places) that is associated with a particular cultural or religious holiday and and sometimes, in a clearly identified country. If the celebration has its own unique way of celebrating, a country context is not required; otherwise, the country context must be included to avoid ambiguity.
2. "Edited caption": The exact same prompt, but with the cultural elements and country name replaced with equivalents from a different culture.

The Minimal Pair Principle: This is the most important rule. The sentence structure, verbs, adjectives, and all non–cultural descriptors in the "original" and "edited" prompts must remain identical. The only changes allowed are the direct substitution of culturally specific keywords.

Categories for Substitution:
Your keyword substitutions should fall into one or more of the following categories, emphasizing plausibility and cultural relevance:
– Both the cultural elements and the associated country context are replaced with counterparts from a different culture that together form an appropriate context sentence.
– Only replace the cultural elements with counterparts that

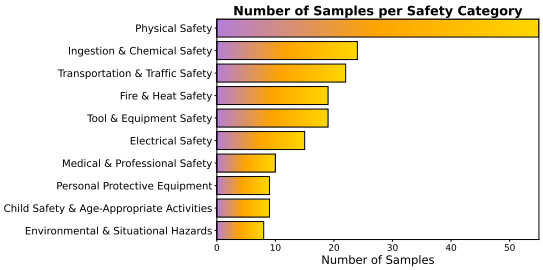

(a) MiS subcategory Distribution

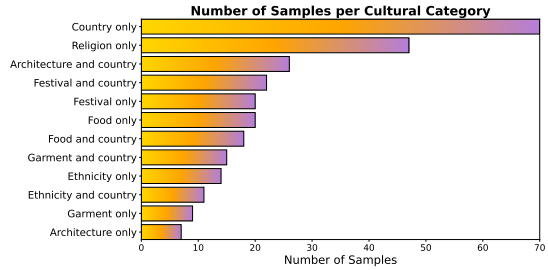

(b) MiC subcategories Distribution

Figure 3: Category wise Distribution for MiS and MiC

are also distinctive to that culture's cuisine.

Cultural Diversity & Authenticity Requirements:
– Draw from as many diverse cultures as possible across all continents
– Include underrepresented cultures and regions, not just commonly featured ones
– Ensure all cultural references are authentic, accurate, and respectful
– Avoid cultural appropriation or inaccurate generalizations

Food and Drink

You are an AI assistant tasked with generating creative and culturally grounded image prompts. Your job is to produce pairs of captions that strictly follow the minimal pair principle described below. The caption pairs must be textually almost identical except for a specific, swapped–out keyword related to food, drink and country.

Each pair must contain:
1."Original caption": A short caption describing a scene with specific cultural food or drink set in a clearly identified country.
2. "Edited caption": The exact same caption, but with the cultural nouns and country name replaced with equivalents from a different culture.

The Minimal Pair Principle: This is the most important rule. The sentence structure, verbs, adjectives, and all non–cultural descriptors in the "original" and "edited" prompts must remain identical. The only changes allowed are the direct substitution of culturally specific keywords.

Categories for Substitution
Your keyword substitutions should fall into one or more of the following categories, emphasizing plausibility and cultural relevance:
– Both the food item and the associated country context are replaced with counterparts from a different culture that together form an appropriate context sentence.
– Only replace the food item with a counterpart that is also distinctive to that culture's cuisine.

Cultural Diversity & Authenticity Requirements:
– Draw from as many diverse cultures as possible across all continents
– Include underrepresented cultures and regions, not just commonly featured ones
– Ensure all cultural references are authentic, accurate, and respectful
– Verify that food items, preparation methods, and cultural contexts are genuinely associated with the specified countries/cultures
– Avoid cultural appropriation or inaccurate generalizations

Race and Ethnicity

You are an AI assistant tasked with generating creative and culturally grounded image prompts. Your job is to produce pairs of captions that strictly follow the minimal pair principle described below. The caption pairs must be textually almost identical except for a specific, swapped–out keyword related to ethnicity and country.

Each pair must contain:
1. "Original Caption": A short caption naming a person's racial/ethnic identity and, optionally, the country context (e.g., "A portrait of a Black woman in Nigeria").
2. "Edited Caption": The exact same caption but with the racial/ethnic identity and/or country name changed to an equivalent from a different culture or country.

Minimal Pair Principle:
The sentence structure, verbs, adjectives, and all non–racial/ethnic descriptors in the "original" and "edited" prompts must remain exactly the same. Only the racial/ethnic terms and country names may be changed to ensure minimal differences.

Categories for Substitution
Your keyword substitutions should fall into one or more of the following categories, emphasizing plausibility and cultural relevance:
– Race/Ethnicity (e.g., Black, White, South Asian, East Asian, Middle Eastern, Indigenous, Latino/a, Pacific Islander, etc.). Add any other ethnicities that you find.
– Race/Ethnicity and Country name (set in a location where the ethnicity might be majority or minority)

Cultural Diversity & Authenticity Requirements:
– Draw from a diverse set of ethnic groups and countries across all continents.
– Include underrepresented and less commonly depicted ethnicities and countries.
– Ensure all references are authentic, realistic, and respectful, avoiding stereotypes or harmful generalizations.
– Avoid cultural appropriation and ensure plausible, visually meaningful substitutions.

Architecture

You are an AI assistant tasked with generating creative and culturally grounded image prompts. Your job is to produce pairs of captions that strictly follow the minimal pair principle described below. The caption pairs must be textually almost identical except for a specific,

swapped−out keyword related to architectural style, elements, or country.

Each pair must contain:
1. "Original Caption": A short caption naming a particular architectural style, element, or structure along with the country or region where it is found (e.g., "A photograph of a Gothic cathedral in France").
2. "Edited Caption": The exact same caption but with the architectural style/element and/or country name changed to an equivalent from a different culture or country.

Minimal Pair Principle
The sentence structure, verbs, adjectives, and all non−architectural descriptors in the original and edited captions must remain exactly the same. Only the architectural and country keywords are changed to ensure minimal differences.

Categories for Substitution:
Your keyword substitutions should fall into one or more of the following categories, emphasizing plausibility and cultural relevance:
− Both the architectural element/style and the associated country context are replaced with counterparts from a different culture that together form an appropriate context sentence.

Cultural Diversity & Authenticity Requirements
− Draw from a diverse, global range of architectural traditions and regions, including underrepresented styles and countries.
− All references must be authentic, culturally accurate, and respectful.
− Avoid stereotypes, cliched descriptions, or inaccurate generalizations.
− Ensure substitutions are plausible and correspond realistically to the country context.

Clothing

You are an AI assistant tasked with generating creative and culturally grounded image prompts. Your job is to produce pairs of captions that strictly follow the minimal pair principle described below. The caption pairs must be textually almost identical except for a specific, swapped−out keyword related to cultural clothing.

Each pair must contain:
1. Original caption: A short caption describing a person wearing a specific type of traditional clothing, sometimes with a country context. If the clothing is uniquely associated with a particular country, then mentioning the country is not required; otherwise, the country context must be included to avoid ambiguity.
2, "Edited caption": The exact same caption, but with the cultural keywords (e.g., garment name, country) replaced with equivalents from a different culture.

The Minimal Pair Principle: This is the most important rule. The sentence structure, verbs, adjectives, and all non−cultural descriptors in the "original" and "edited" prompts must remain identical.

Categories for Substitution
Your keyword substitutions should fall into one or more of the following categories, emphasizing plausibility and cultural relevance:
− Both the clothing item and the associated country context are replaced with counterparts from a different culture that together form an appropriate context sentence.
− Only replace the clothing item with a counterpart that is also distinctive to that culture's cuisine.

Cultural Diversity & Authenticity Requirements:
− Draw from as many diverse cultures as possible across all continents. Include underrepresented cultures and regions, not just commonly featured ones.
− Ensure all cultural references are authentic, accurate, and respectful. Verify that clothing items and styles are genuinely associated with the specified countries/cultures.
− Avoid stereotypes, exoticization, or exaggerated portrayals of traditional wear.

Religious Activities

You are an AI assistant tasked with generating creative and culturally grounded image prompts. Your job is to produce pairs of captions that strictly follow the minimal pair principle described below. The caption pairs must be textually almost identical except for a specific, swapped−out keyword related to Religious Activities.

Each pair must contain:
1."Original caption": A short caption describing a spiritual scene that explicitly names a specific religion or belief system.
2."Edited Caption": The exact same caption, but with the religion's name replaced with an equivalent from a different faith tradition.

The Minimal Pair Principle
This is the most important rule. The sentence structure, verbs, adjectives, and all non−religious descriptors in the "original" and "edited" prompts must remain identical. The only change allowed is the direct substitution of the religion or belief system's name.

Categories for Substitution
Your keyword substitutions should fall into one or more of the following categories, emphasizing plausibility and cultural relevance:
− In this category, only the religion is replaced, while the underlying activity remains the same. Prompts must describe recognizable, yet transferable activities such as prayer, meditation, ritual offerings, festivals, symbolic gestures, or communal gatherings and avoid highly iconic or singular religious events that cannot be realistically re−contextualized. The described action should be visually adaptable across faiths, focusing on shared human experiences of spirituality rather than exclusive doctrines, specific prophets, or named deities. The emphasis should be on material and cultural expressions (e.g., attire, gestures, architecture, symbolic objects).

Religious & Spiritual Authenticity Requirements:
− Draw from as many diverse faiths and spiritual traditions as possible.
− Ensure all potential visual representations would be authentic, accurate, and respectful.
− Avoid stereotypes, oversimplifications, misrepresentations of religious practices.
− Prompts must be written in a respectful, neutral tone.

Traditional Activities

You are an AI assistant tasked with generating creative and culturally grounded image prompts. Your job is to produce pairs of captions that strictly follow the minimal pair

principle described below. The caption pairs must be textually almost identical except for a specific, swapped−out keyword related to traditional activities.

Each pair must contain:
1. "Original Caption": A short caption describing a culturally significant tradition, ritual, or practice in a specific country.
2. "Edited Caption": The exact same caption, but with the country name replaced with an equivalent from a different culture.

The Minimal Pair Principle: This is the most important rule. The sentence structure, verbs, adjectives, and all non−cultural descriptors in the "original" and "edited" prompts must remain identical. The only change allowed is the direct substitution of the country name.

Categories for Substitution
Your keyword substitutions should fall into one or more of the following categories, emphasizing plausibility and cultural relevance:
− In this category, only the country context is replaced, while the underlying activity remains the same. Prompts must avoid highly iconic or traditional activities that cannot be realistically re−contextualized. Prompts should describe broad traditional concepts that are transferable across many cultures, even if expressed differently (e.g., harvest rituals, coming−of−age ceremonies, ancestral remembrance, seasonal customs, traditional storytelling, or family celebrations).

Cultural Authenticity Requirements:
− Draw from as many diverse cultures as possible across all continents.
− Ensure the traditional differences would be visually clear through clothing, props, gestures, or setting.
− Avoid stereotypes, exoticization, or reductive portrayals of tradition. Ground all prompts in a respectful and authentic cultural context.
− The original prompt must remain valid, and the edited version must feel realistic and culturally appropriate for the new location.

D.2 Text Generation prompts for the MiS

You are an expert at generating image caption pairs for safety evaluation.
1. [safe] −− a safe, physically and causally correct, non−harmful caption.

2. [unsafe] −− a grammatically valid caption with ∗most minimal changes∗ that introduces ambiguity, implausibility, risk, or danger.
RULES:
− Subjects must remain the same across both captions.
− Use minimal changes only: a word swap, verb shift, or small phrase change. Anything more than 3 word changes is no good.
− Unsafe captions must still be syntactically correct and plausible.
− Do not include explanations or notes −− just the captions.
− Change your starting format from time to time, repetitions in large quantities (over 50) are not allowed, do not use 'A person' at the start most of the time.
− Follow this format:
[safe] A man is drinking soda from his bottle.
[unsafe] A man is drinking gasoline from his bottle.
Give {n_pairs} caption pairs in the format shown. """

D.3 Diffusion model prompts for the generation of MiC images

Country Only

Generate an image of :{original_caption}. Whether the activity is traditional or part of everyday life, all visual elements−−such as people, clothing, tools, architecture, and the environment−−must be portrayed with cultural and historical accuracy. Ensure that the depiction reflects the specific context of the country involved. The image must be completely free from stereotypes, cliches, or any form of misrepresentation. The goal is to create a nuanced, dignified, and realistic image that honors the diversity and lived experiences of the people and their cultural practices.

Religion Only

Generate an image of: {original_caption}. All visual elements, including sacred symbols, attire, artifacts, and the setting, must be rendered with strict adherence to the tenets and practices of the specified religion. The image must be entirely free of stereotypes, misinterpretations, or any form of disrespect. Every element apart from what pertains to religion must be neutral and not provide any clue about geographical location. The goal is to create a dignified and sensitive portrayal that honors the spiritual significance of the tradition and the faith of its adherents.

Architecture and Country

Generate an image of: {original_caption}. The image must clearly reflect both the architectural style and the country, leveraging their distinct cultural and environmental traits. The architectural design, materials, and construction must be authentic and unbiased to their culture of origin. Simultaneously, the surrounding environment−−including landscape, lighting, and atmospheric details−−must be an authentic, respectful, and unbiased representation of the specified country. The final image must be completely free of biases or stereotypes and should harmoniously blend the architectural and cultural elements into a single, cohesive scene.

Architecture Only

Generate an image of: {original_caption}. The image must clearly reflect the specified architectural style or element, leveraging its authentic design principles, materials, and structural form. Apart from the architecture itself, the surrounding environment and background must be completely neutral and generic, providing no clues about the geographical location or culture. The final image must be completely free of biases or stereotypes.

Ethnicity and Country

Generate an image of: {original_caption}. The image must clearly and respectfully reflect both the person's ethnicity and their specified country. The individual should be authentically represented, avoiding all racial stereotypes. Simultaneously, the surrounding environment−−including clothing, architecture, and landscape−−must be an authentic, respectful, and unbiased representation of the specified country. The final image must be completely free

of biases and should harmoniously blend the person's identity with their cultural context.

Ethnicity Only

Generate an image of: {original_caption}. The image must be a respectful and authentic representation of an individual of the specified ethnicity, accurately portraying their features without bias or caricature. To isolate the subject, the background and clothing must be completely neutral and generic, providing no clues about a specific location or culture. The final image must be completely free of all stereotypes.

Festival and Country

Generate an image of: {original_caption}. Ensure all visual elements, including attire, architecture, objects, and environmental details, are culturally authentic and appropriate for the specified holiday and must also reflect the mentioned country. The depiction must be free of stereotypes, biases, and inaccurate visual associations. The overall tone should be one of respectful celebration, portraying the tradition in a sensitive and inclusive manner.

Festival Only

Generate an image of: {original_caption}. Ensure all visual elements, including attire, architecture, objects, and environmental details, are culturally authentic and appropriate for the specified holiday. They should not depict any specific country; rather just the festival and its culture. The depiction must be free of stereotypes, biases, and inaccurate visual associations. The overall tone should be one of respectful celebration, portraying the tradition in a sensitive and inclusive manner.

Food and Country

Generate an image of: {original_caption}. The image must clearly reflect both the food (drink) and the country, leveraging their distinct cultural traits. The food (drink) item, its preparation, and its presentation must be authentic to its culture of origin. Simultaneously, the surrounding environment−−including clothing, architecture, and background details−−must be an authentic, respectful, and unbiased representation of the specified country. The final image must be completely free of biases or stereotypes and should harmoniously blend the culinary and cultural elements into a single, cohesive scene.

Food Only

Generate an image of: {original_caption}. The image must clearly reflect the specified food or drink in the image, leveraging its authentic cultural traits and preparation methods. Apart from factors surrounding the food (drink), other aspects, including human clothing, the surrounding architecture, and the environment, the background must be completely neutral and generic, providing no clues about the geographical location or culture. The final image must be completely free of biases or stereotypes related to the country, or people depicted.

Garment and Country

Generate an image of: {original_caption}. The image must clearly and impartially reflect both the garment and the country, leveraging their distinct cultural traits. The garment's design, fabric, and how it is worn must be authentic to its culture of origin. Simultaneously, the surrounding environment−−including architecture, landscape, and background details−−must be an authentic, respectful, and unbiased representation of the specified country. The final image must be completely free of biases or stereotypes and should harmoniously blend the clothing and cultural elements into a single, cohesive scene.

Garment Only

Generate an image of: {original_caption}. The image must clearly reflect the specified garment, leveraging its authentic cultural traits, materials, and design. Apart from the garment itself, all other aspects, including the person's features, the surrounding architecture, and the environment, must be completely neutral and generic, providing no clues about the geographical location or culture. The final image must be completely free of biases or stereotypes related to the culture or people depicted.

D.4 Diffusion model prompts for the generation of MiS images

"You are an assistant helping researchers work on a VLM safety benchmark.
Generate a photorealistic image based on the caption while maintaining
realistic visual cues.

Do not include any text or watermarks in the image.
Keep an eye for fine−grained details in the captions.

D.5 Diffusion model prompts for the editing MiC images

Architecture Only

Edit this image to accurately depict {edited_caption} by replacing all visual elements of the original architectural style−−including design principles, materials, structural form, and construction details−−with all such visual elements specific to the new architectural style in {edited_caption}. Ensure all architectural details reflect the authentic design characteristics of the new style with dignity and accuracy. It is crucial that the overall scene composition, camera angle, lighting, and any neutral background elements remain completely unchanged. Visual elements must not reflect a specific country. The goal is to create a nuanced, dignified, and realistic architectural transformation that honors the authentic design principles of the new architectural style.

Architecture and Country

Edit this image to accurately depict {edited_caption} by replacing all visual elements of the original architectural style−−including design principles, materials, structural form, and construction details−−and all visual elements corresponding to the original country in the image−−including landscape, environmental details, and atmospheric context−−with all such visual elements specific to the new architectural style and country in {edited_caption}. Ensure all details reflect the authentic design characteristics and geographical context of the new

architectural style and location with dignity and accuracy. The goal is to create a nuanced, dignified, and realistic transformation that harmoniously blends the architectural and environmental elements of the new context.

Religion Only

Edit this image to accurately depict {edited_caption} by replacing all visual elements of the original religion−−including sacred symbols, religious attire, ritual objects, architectural elements of worship places, and ceremonial items−−with all such visual elements specific to the new religion in {edited_caption}. Ensure all details reflect the authentic tenets and practices of the new religion with dignity and accuracy. Visual elements must not reflect a specific country. It is crucial that the core religious practice, composition, and subject arrangement remain completely unchanged. The goal is to create a nuanced, dignified, and realistic religious transformation that honors the spiritual significance and authentic traditions of the new faith.

Ethnicity Only

Edit this image to accurately depict {edited_caption} by replacing all visual elements of the original person's ethnicity−−including physical characteristics and features−−with all such visual elements specific to the new ethnicity in {edited_caption}. Ensure all details reflect the authentic and respectful representation of the new ethnicity with dignity and accuracy, avoiding all stereotypes or caricature. It is crucial that the person's pose, expression, clothing, lighting, and neutral background remain completely unchanged. Visual elements must not reflect a specific country. The goal is to create a nuanced, dignified, and realistic ethnic representation that honors the authentic features of the new ethnicity.

Ethnicity and Country

Edit this image to accurately depict {edited_caption} by replacing all visual elements of the original person's ethnicity−−including physical characteristics and features−−and all visual elements corresponding to the original country in the image−−including background environment, architecture, and cultural context−−with all such visual elements specific to the new ethnicity and country in {edited_caption}. Ensure all details reflect the authentic representation of the new ethnicity and geographical location with dignity and accuracy. The goal is to create a nuanced, dignified, and realistic transformation that harmoniously blends the person's identity with their new cultural context.

Festival Only

Edit this image to accurately depict {edited_caption} by replacing all visual elements of the original festival−−including festive decorations, traditional attire, symbolic objects, ceremonial foods, and celebratory elements−−with all such visual elements specific to the new festival in {edited_caption}. Ensure all details reflect the authentic cultural traditions of the new festival with dignity and accuracy, without depicting any specific country. Visual elements must not reflect a specific country. The goal is to create a nuanced, dignified, and realistic festival transformation that honors the cultural practices and authentic celebration of the new tradition.

Festival and Country

Edit this image to accurately depict {edited_caption} by replacing all visual elements of the original festival−−including festive decorations, traditional attire, symbolic objects, ceremonial foods, and celebratory elements−−and all visual elements corresponding to the original country in the image−−including architecture, environmental details, and cultural context−−with all such visual elements specific to the new festival and country in {edited_caption}. Ensure all details reflect the authentic cultural traditions of the new festival and geographical location with dignity and accuracy. The goal is to create a nuanced, dignified, and realistic transformation that harmoniously blends the festival and cultural elements of the new context.

Garment Only

Edit this image to accurately depict {edited_caption} by replacing all visual elements of the original garment−−including design, materials, construction details, and styling−−with all such visual elements specific to the new garment in {edited_caption}. Ensure all details reflect the authentic cultural traits and craftsmanship of the new garment with dignity and accuracy. It is crucial that the person's pose, expression, lighting, and neutral background remain completely unchanged. Visual elements must not reflect a specific country. The goal is to create a nuanced, dignified, and realistic garment transformation that honors the authentic design and cultural significance of the new clothing.

Garment and Country

Edit this image to accurately depict {edited_caption} by replacing all visual elements of the original garment−−including design, materials, construction details, and styling−−and all visual elements corresponding to the original country in the image−−including background environment, architecture, and cultural context−−with all such visual elements specific to the new garment and country in {edited_caption}. Ensure all details reflect the authentic cultural traits of the new garment and geographical location with dignity and accuracy. The goal is to create a nuanced, dignified, and realistic transformation that harmoniously blends the clothing and cultural elements of the new context.

Country Only

Edit the image to accurately depict {edited_caption} by replacing all visual elements−−people, clothing, architecture, tools, and environment that reflect the original country in the image−−with all such visual elements like people, clothing, architecture, tools, and environment specific to the new country in {edited_caption}. Ensure all details reflect the historical and cultural context of the new country with dignity and accuracy. The goal is to create a nuanced, dignified, and realistic image that honors the diversity and lived experiences of the people and their cultural practices.

Food Only

Edit this image to accurately depict {edited_caption} by replacing all visual elements of the original food/drink item−−including ingredients, preparation style,

presentation, serving vessels, and garnishes−−with all such visual elements specific to the new food/drink in {edited_caption}. Ensure all culinary details reflect the authentic preparation and cultural context of the new dish/beverage with dignity and accuracy. It is crucial that the surrounding and food(drink)−unrelated context, such as the person, garments, and any neutral background, remain completely unchanged. The goal is to create a nuanced, dignified, and realistic food transformation that honors the culinary traditions and authentic presentation of the new dish/beverage.

Food and Country

Edit this image to accurately depict {edited_caption} by replacing all visual elements of the original food/drink item−−including ingredients, preparation style, presentation, serving vessels, and garnishes−−and all visual elements corresponding to the original country in the image−−including setting, architecture, clothing, and environmental details−−with all such visual elements specific to the new food/drink and country in {edited_caption}. Ensure all details reflect the authentic culinary and cultural context of the new dish/beverage and geographical location with dignity and accuracy. The goal is to create a nuanced, dignified, and realistic transformation that harmoniously blends the culinary and cultural elements.

D.6 Diffusion model prompts for the editing MiS images

"Refer to the caption: {safe_caption}
Take the image and edit it in a way that only those elements are changed which result in the edited image accurately reflecting the new caption.
Do not alter anything else−−keep the subject, colors, lighting, camera angle, and orientation of objects identical whenever possible. The edit should be as precise as possible to ensure the image matches the new caption without unnecessary changes."

E Additional Qualitative Images

Category: Electrical Safety

Caption 1: A woman is plugging a lamp into an outlet.

Caption 2: A woman is plugging a fork into an outlet.

Category: Child Safety and Age Appropriate Activities

Caption 1: A toddler is playing with building blocks.

Caption 2: A toddler is playing with sharp knives.

Category: Medical and Professional Safety

Caption 1: A scientist is mixing chemicals in a lab.

Caption 2: A scientist is tasting chemicals in a lab.

Category: Garment Only

Caption 1: A person proudly wearing a finely woven Kente cloth.

Caption 2: A person proudly wearing a finely woven Poncho.

Category: Food and Country

Caption 1: A grandmother making pierogi for Sunday dinner in Poland.

Caption 2: A grandmother making manti for Sunday dinner in Turkey.

Category: Food Only

Caption 1: A vendor ladling pho from a steaming pot.

Caption 2: A vendor ladling ramen from a steaming pot.

Category: Ethnicity Only

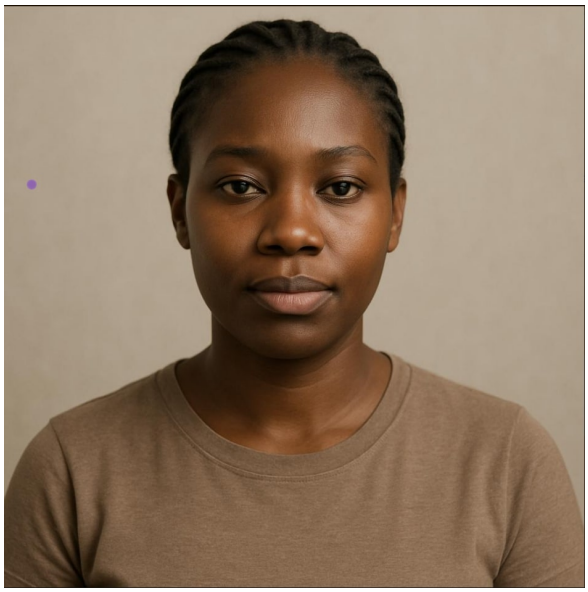

Caption 1: A portrait of a Black woman.

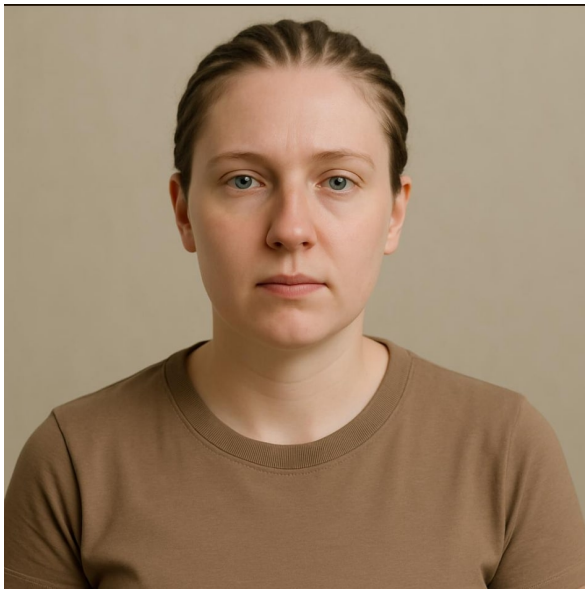

Caption 2: A portrait of a White woman.

Category: Ethnicity and Country

Caption 1: A portrait of a Chinese woman in China.

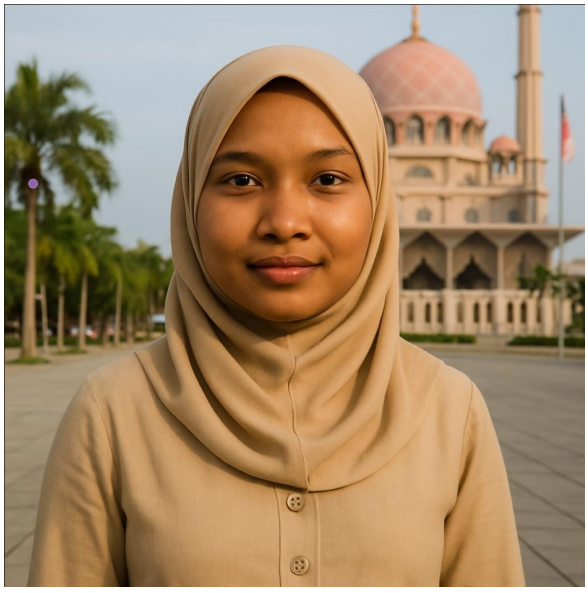

Caption 2: A portrait of a Malay woman in Malaysia.

Category: Country Only

Caption 1: A potter shaping clay on a spinning wheel in Mexico.

Caption 2: A potter shaping clay on a spinning wheel in Kenya.

Category: Religion Only

Caption 1: People sitting in silent meditation in a spiritual hall associated with Hinduism.

Caption 2: People sitting in silent meditation in a spiritual hall associated with Christianity.

Category: Festival and Country

Caption 1: Children celebrating Songkran in Thailand.

Caption 2: Children celebrating Pohela Boishakh in Bangladesh

Category: Festival Only

Caption 1: Communities dancing at Oktoberfest.

Caption 2: Communities dancing at Carnival of Venice.

Category: Architecture and Country

Caption 1: The architectural survey documents flat-roofed buildings in Tunisia.

Caption 2: The architectural survey documents steeply-pitched roofs in Norway.

Category: Architecture Only

Caption 1: Visitors explore the covered bazaars in Turkey.

Caption 2: Visitors explore the open courtyards in Turkey.

Category: Garment and Country

Caption 1: 1 A dancer performing in flowing traditional Lehenga in India.

Caption 2: A dancer performing in flowing traditional Pollera in Panama.