

WHICH PRIVATE ATTRIBUTES DO VLMS AGREE ON AND PREDICT WELL?

Olena Hrynenko^{1,2}, Darya Baranouskaya^{1,2}, Alina Elena Baia², Andrea Cavallaro^{1,2}

¹EPFL, Switzerland, ²Idiap Research Institute, Switzerland

ABSTRACT

Visual Language Models (VLMs) are often used for zero-shot detection of visual attributes in the image. We present a zero-shot evaluation of open-source VLMs for privacy-related attribute recognition. We identify the attributes for which VLMs exhibit strong inter-annotator agreement, and discuss the disagreement cases of human and VLM annotations. Our results show that when evaluated against human annotations, VLMs tend to predict the presence of privacy attributes more often than human annotators. In addition to this, we find that in cases of high inter-annotator agreement between VLMs, they can complement human annotation by identifying attributes overlooked by human annotators. This highlights the potential of VLMs to support privacy annotations in large-scale image datasets.

Index Terms— VLMs, attributes recognition, privacy

1. INTRODUCTION

Human labels are the reference for privacy attributes in images. Multi-label annotation of a large number of images (e.g., 103 attributes and approximately 7,000 images per annotator¹ in the VISPR dataset [1]) causes a high cognitive load and leads to annotator fatigue [2]. On average, a person recalls up to seven concepts in their active memory [3], hence there is a high chance that while the annotated image attributes, or concepts, are correct, they are not the only attributes that are present (see Fig. 1), leading to false negatives.

Recent research has demonstrated that Large Language Models (LLMs) and Visual Language Models (VLMs) are emerging as viable substitutes for human annotators [4, 5]. For example, using VLMs in image annotation tasks, such as those conducted on the CelebA [6] dataset, shows that automated annotation outputs can achieve high agreement with humans on objective binary classification tasks (e.g., "Is this person wearing eyeglasses? Answer with only yes or no."), while significantly reducing cost [5]. Motivated by the limitations of human annotation and the capabilities of VLMs, in this work, we investigate whether VLMs can recognise a large set of privacy attributes and how their predictions align with human annotations. Our study focuses on the privacy attributes defined in the VISPR dataset [1], which were derived from regulations, social network policies, and the personal judgment of the authors [1]. Examples of attributes include *Gender*, *Signature*, *Full Name*, and *Tattoo*.

Similar to Chiang and Lee [4], to ensure that the experimental setup is similar for both human and VLM annotators, we replicate the annotation instructions provided to the annotators when prompting VLMs. However, we adapt the setup to a multimodal context. Building on prior work that assesses the ability of VLMs to detect or extract private information in images [7, 8, 9, 10, 11], we extend the evaluation to the more fine-grained task of private attribute recognition. Unlike related works [7, 8, 9, 10], we use the entire VISPR [1]

¹In the VISPR dataset [1], three annotators performed an image multi-labelelling for a dataset of 22,167 images from Flickr, one annotator per image.



Annotated attributes:

Age Group
Weight Group
Height Group
Gender
Eye Color
Hair Color ...



Annotated attributes:

Age Group
Weight Group
Height Group
Gender
Hair Color
Face (Complete) ...

Fig. 1. Examples of images from the VISPR dataset [1]. Attributes annotated as *present* are shown in green, and those annotated as *absent* are shown in red. While both images show a group of people, the *Weight Group* and *Height Group* attributes have been omitted by a human annotator for the image in the first row.

image test set for the evaluation of VLMs' zero-shot performance, consider the full list of privacy attributes provided by VISPR, and use the annotation instructions provided in the VISPR dataset to prompt the models. Only a subset of attributes or images is usually considered [7, 8, 10]. In addition, differently from prior work, we also consider potential human errors in the annotation of the VISPR dataset. While [5] evaluated personal appearance attributes of celebrities, we consider attributes and images that describe not only humans, but also the objects, or locations, widening the set of possible attributes.

In summary, we evaluate the ability of VLMs to recognise private attributes and analyse annotation disagreements between the VLMs and human annotators in the VISPR dataset.

2. RELATED WORK

Privacy-related studies examine the inference-time privacy risks of VLMs, such as identifying the presence [7, 8] or extracting private information [8, 9, 10, 11].

The Multi- P^2A benchmark [7] includes measuring VLMs' ability to recognise the presence of any private attribute in the image. For this, Zhang et al. [7] build a balanced dataset with an equal number of images that contain and do not contain privacy-sensitive content. They select 23 from the original 67 VISPR [1] privacy attributes, excluding some categories on which VLMs can fail, such as *Eye Color*. The authors [7] evaluate whether at least one privacy attribute is present, rather than differentiating among specific privacy attributes. REVAL [9] tests VLMs on various personally identifiable information (PII), such as *Home Address*, *Credit Card Number*, *Telephone Number*, and observes that these models struggle to accurately recognise specific privacy attributes in images. Based on this, Zhang et al. [9] propose to reformulate the evaluation as a binary task that assesses only the presence or absence of any privacy

attribute. Similarly, MultiTrust [8] prompts VLMs to make a binary decision about the presence of private information, using a small set of images derived from VISPR [1] and VizWiz-Priv [12] covering a limited range of PII categories.

Tömekçe et al. [11] show that private attribute inference from social media images is already feasible and that safeguards can be bypassed with simple prompt engineering. Schultenkämper and Bäumer [10] construct a vision-question answering dataset with prompts crafted for extracting private information. The authors [10] then assess the ability of VLMs to extract such information, focusing only on directly visible personal attributes, while excluding attributes like *Full Name* or *Place of Birth*. MultiTrust [8] evaluates the ability of VLMs to extract private information considering a limited set of PIIs (such as *Email*, *Name*, *Address*, *Credit Card*, and *ID Number*) and a small set of images derived from VISPR. In addition, MMDT [13] and GeoLocator [14] highlight that VLMs can accurately predict locations from landmarks and street cues, exposing geographic privacy. RTVLM [15] shows that many open-source VLMs fail to refuse privacy-sensitive queries, often disclosing personal or celebrity data.

While several studies [7, 8, 9, 10, 13] reuse VISPR attributes to construct privacy-sensitive probes and test VLMs, they focus on a limited number of private attributes/PIIs and do not evaluate fine-grained recognition of specific privacy attributes. Our work addresses this limitation by systematically evaluating VLMs on all VISPR privacy attributes using the instructions as for human annotators, and by analysing the discrepancies between VLMs and human annotators.

3. EXPERIMENTAL SETUP

We use three open-source instruction-following VLMs, namely, Gemma-3-4b-it [16], Qwen2.5-VL-7B-Instruct [17], and Llama-3.2-11B-Vision-Instruct [18]. The experiments are run on a single RTX3090. For each image in the VISPR test dataset (8000 images), we run a prompt for each of the 67 privacy attributes in a zero-shot setup. The prompt is formulated as follows:

User prompt:
Here is information about the attribute:
Attribute: {attribute}
Attribute's definition: {description}

By using this definition for any subject in the image – either in the foreground or background, answer the following question: is attribute present or absent in the image?
Please answer [Present] or [Absent].

Since the response to the prompt varies from a single word (e.g., "Present") to one or more sentences (e.g., "The image shows two nail polish designs, but there is no signature present in the image."), we use a two-step process to parse the response into a label. As in [19], we use the same model for parsing the response. We first use the following prompt:

System prompt:
You are provided with an open-ended reply.
The possible closed-ended answers are:
[present], [absent]. Map open-ended reply to the closed-ended answer. You can reply with: [present], [absent] or [NONE] if none of the closed-ended replies is suitable. You MUST begin your reply with: 'THE ANSWER: '

User Prompt:
Open-ended reply: {reply}

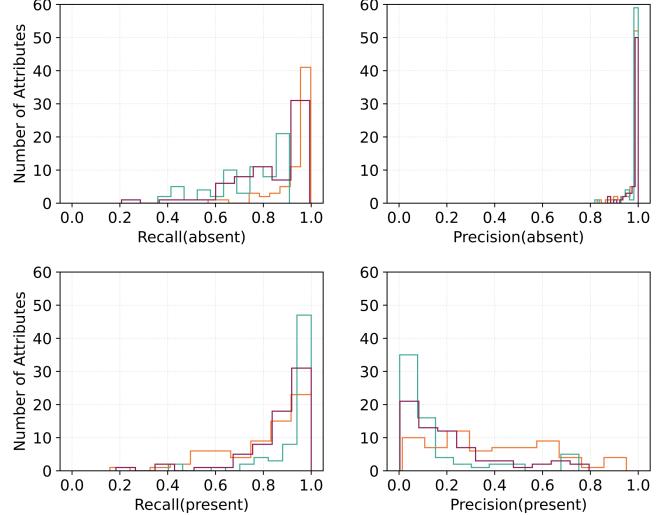


Fig. 2. Distribution of precision and recall for present and absent labels for zero-shot recognition of Qwen2.5-VL-7B-Instruct (orange), Gemma-3-4b-it (green), Llama-3.2-11B-Vision-Instruct (red) for 67 attributes of the VISPR test set [1].

We then apply logic-based substring parsing to obtain the present, absent, or NONE labels. We included the NONE case to account for the inconclusive model's responses (e.g., "The image shows a ticket for a Metal Hammer Radio event in Barcelona on June 6, 2002. ... However, there is no indication of Sexual Orientation present or absent in the image."). We enable stochastic decoding (set `do_sample=True` in the model's generation parameters), which allows the model to sample from the output distribution and generate more diverse predictions, thereby reducing the risk of consistently predicting the same class when multiple outcomes are considered plausible.

4. EVALUATION

4.1. VLMs recognition of private attributes

We analyse the distributions of the precision and recall scores for present and absent classes (see Fig. 2). Additionally, we evaluate how well VLMs recognise privacy attributes by considering the balanced accuracy (see Fig. 3).

Models have high recall for both present and absent labels, indicating that they rarely miss attributes when they are actually present in the image and correctly identify truly absent attributes. While the precision for the absent class is consistently close to 1, the precision for the present class is notably lower, with models often detecting attributes that are not annotated in the dataset. This discrepancy may not solely reflect model error, given the high cognitive load of labelling a large set of attributes for human annotators. As a result, for the majority of attributes and models, the balanced accuracy is above 0.75. However, low precision negatively affects the F1-macro scores, which, for the worst-performing model Gemma-3-4b-it, range from 0.4 to 0.6 across the majority of attributes (57 out of 67). Across the three models, Qwen2.5-VL-7B-Instruct significantly outperforms Gemma-3-4b-it and Llama-3.2-11B-Vision-Instruct, reaching the highest F1-macro score for the majority of attributes (64 out of 67) and F1-macro over 0.6 for 55 attributes. Overall, despite some limitations in precision for the present class, the performance of Qwen2.5-VL-7B-Instruct aligns the most with annotators' labels, while Gemma-3-4b-it and Llama-3.2-11B-Vision-Instruct perform noticeably worse.

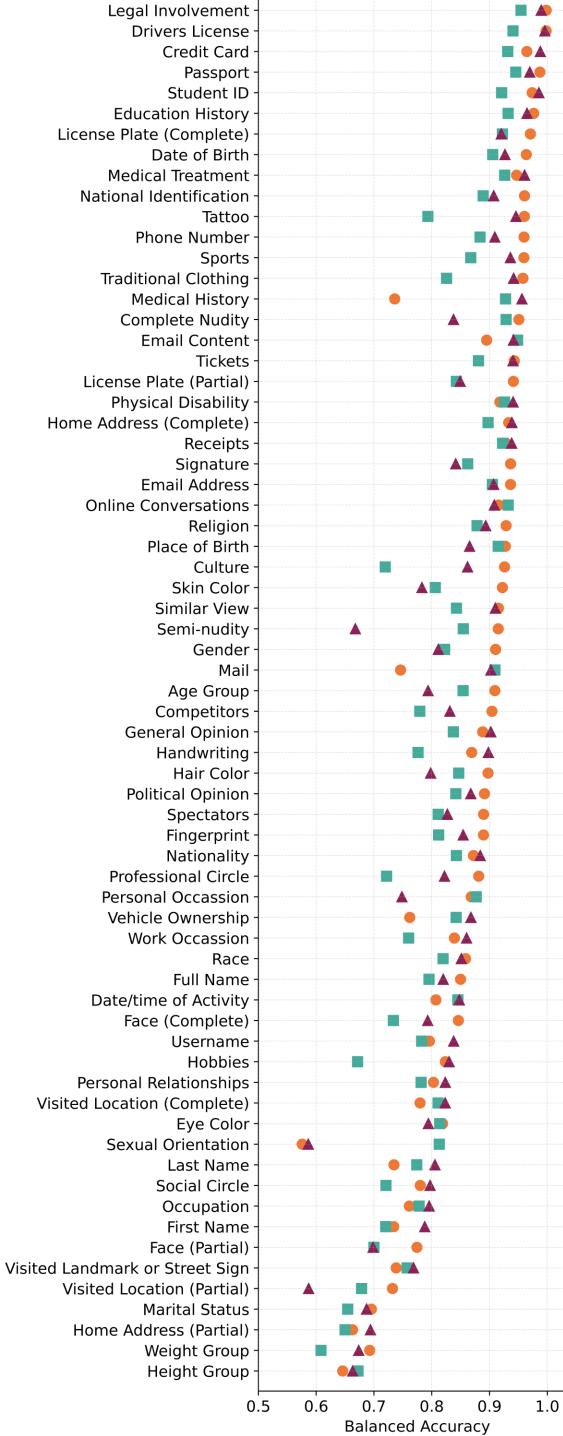


Fig. 3. Balanced accuracy of zero-shot privacy attribute recognition for Qwen2.5-VL-7B-Instruct (●), Gemma-3-4b-it (■), and Llama-3.2-11B-Vision-Instruct (▲) on the VISPR test set [1].

4.2. VLMs and human disagreements

To analyse the disagreements between the VLMs and the human annotators, we take the top-five attributes with the highest Fleiss kappa [20] among the VLM annotators (see Tab. 1). The Fleiss kappa measures how well the annotators are aligned in their responses, and ranges from -1 (agreement worse than chance),

Table 1. Top-five attributes with the highest VLM annotation agreement according to Fleiss kappa (1 – perfect agreement, 0 – annotation agreement by chance). For each attribute, we detail the cases in which VLMs identify it as present, whereas human annotators as absent. We manually inspect 50 images per attribute, report the percentage of cases in which VLMs are correct, and list the image categories in which VLMs are incorrect, from most common to least.

Attribute	VLM label: present		VLM incorrect (categories)
	Fleiss kappa	VLM correct	
Age Group	0.693	64%	animals; drawings of people; statues; dolls
Gender	0.680	18%	statues; drawings of people; dolls; person visible but gender indistinguishable; helmet; toy
Hair Color	0.650	22%	drawings of people; statues; animals; doll; person visible but hair color indistinguishable
Spectators	0.641	70%	a group of people is present, but no clear spectating of an event; only participants of an event
Face (Partial)	0.580	40%	visible full face; animal; drawings of people; part of the body; mask; fireworks

to 1 (perfect agreement). Taking attributes with high Fleiss kappa allows us to avoid selecting images where the models agreed by chance. For each attribute, we consider the images on which models agree. This reduces the number of images per attribute to 5875 for *Age Group*, 5810 for *Gender*, 5636 for *Hair Color*, 5755 for *Spectators*, and 5407 for *Face (Partial)* attributes. We refer to these collective VLM annotations as VLM labels, and to human annotations as human labels. We compute the confusion matrices for VLM and human labels for each attribute (see Fig. 4) and discuss them below.

Age Group, Gender, and Hair Color. We now analyse the disagreement cases for *Age Group*, *Gender*, and *Hair Color* attributes. We investigate which other human-labelled attributes were present in the image. Out of 67 VISPR attributes, we manually select the attributes that correspond to a person being present in the image, the human-defining attributes: *Age Group*, *Weight Group*, *Height Group*, *Gender*, *Eye Color*, *Hair Color*, *Face (Complete)*, *Face (Partial)*, *Semi-nudity*, *Complete Nudity*, *Race*, *Skin Color*. If many of the human-defining attributes appear in the image annotation, it is highly likely that either *Age Group*, *Gender*, or *Hair Color* attribute is also present, but was omitted by the human annotator.

For *Age Group*, *Gender*, or *Hair Color*, for the cases when the human label is absent, and the VLM label is present, we plot the percentage of images that contain at least N other human-defining attributes (see Fig. 5, left). We notice an elbow drop at $N = 6$, which means that at least six other human-defining attributes are present in these images. We assume that the presence of the other six human-defining attributes is sufficient for *Age Group*, *Gender*, and *Hair Color* attributes to be present in the image. For example, if *Face (Partial)*, *Race*, *Skin Color*, *Gender*, *Hair Color*, and *Face (Complete)* attributes are present, the *Age Group* is likely to

		Age Group		Gender		Hair Color		Spectators		Face (Partial)		100%
		absent	present									
Human label	absent	2440 (41.5%)	523 (8.9%)	2401 (41.3%)	382 (6.6%)	2287 (40.6%)	395 (7.0%)	3982 (69.2%)	1251 (21.7%)	1558 (28.8%)	1819 (33.6%)	
	present	17 (0.3%)	2895 (49.3%)	47 (0.8%)	2980 (51.3%)	21 (0.4%)	2933 (52.0%)	0 (0.0%)	522 (9.1%)	8 (0.1%)	2022 (37.4%)	

Fig. 4. The disagreements between the human labels and VLM labels. We note that there are numerous disagreements when the VLMs predict the presence of an attribute, when, according to the human annotators, the attribute is absent. Integers denote image counts, with proportions shown in parentheses.

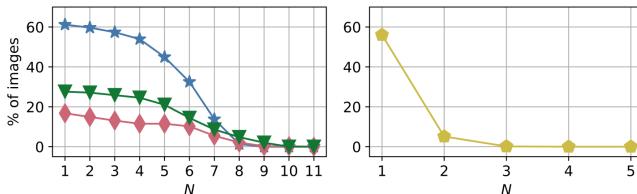


Fig. 5. Left: Percentage of images with at least N other human-defining attributes present according to human annotators for *Age Group* (—★—), *Gender* (—◆—), *Hair Color* (—▼—) attributes. Right: Percentage of images with at least N other relationship-defining attributes present according to human annotators for *Spectators* (—◆—) attribute. The percentage is computed out of the cases when the VLM label is present, and the human label is absent (i.e., 523 for the *Age Group* attribute).

be present as well. Under this assumption, at least 167 images are missing a human label for the *Age Group*: out of the images where all three VLMs annotated this attribute as present, but the human-annotated label is absent (523 images), 32% of them include at least six other human-identifying attributes. Similarly, 10% of the 382 images for the *Gender* attribute, and 14% of 395 images for the *Hair Color* attribute are likely to have been omitted by human annotators. **Spectators.** We now analyse the cases of disagreement for the *Spectators* attribute. From the original list of attributes, we take a subset of the relationship-defining attributes: *Personal Relationships*, *Social Circle*, *Professional Circle*, *Competitors*, *Spectators*, and *Similar View*. We consider the cases when the VLM label is present for *Spectators*, and the human label is absent. In 56% of such cases, there exists at least one other relationship-defining attribute according to the human annotator (see Fig. 5, right). This means that while the models are able to detect the presence of relationship attributes, they might be in disagreement with the human annotators on the types of relationships that are present.

Face (Partial). To analyse the disagreement cases for this attribute, we firstly define the group of attributes that corresponds to the presence of the human face, face-defining attributes: *Face (Complete)*, *Face (Partial)*, and *Eye Color*. Out of 1819 disagreement images where *Face (Partial)* is present according to the VLM, and absent according to a human, 59% of the images contain at least one other face-defining attribute annotated by a human annotator, and 46% contain both of these attributes. While this does not guarantee that the face was partially visible, it guarantees the presence of a face.

We compare the difference in annotation for *Face (Partial)* with *Face (Complete)* attributes. The annotation instructions for *Face (Partial)* and *Face (Complete)* are as follows: "Less than 70% of the face is visible or there is occlusion, such as when the subject

is wearing sunglasses" and "A face is completely visible. Also includes photographs of faces on identity cards, documents or billboards" [1]. For human annotation, both *Face (Partial)* and *Face (Complete)* have the same label in 75% of the cases. A similar pattern holds for Qwen2.5-VL-7B-Instruct (77%), Gemma-3-4b-it (84%), and Llama-3.2-11B-Vision-Instruct (75%).

When either one of the attributes is present, but not both, the human annotates *Face (Complete)* more often than the *Face (Partial)*, which corresponds to 15% and 10% of the cases, respectively. The tendency is reversed for the Qwen2.5-VL-7B-Instruct and Llama-3.2-11B-Vision-Instruct, since when either one of the attributes is present, but not both, *Face (Partial)* is annotated more often than *Face (Complete)*. For Qwen2.5-VL-7B-Instruct, 21% of the cases include a present label for *Face (Partial)* and absent for *Face (Complete)*, and only 1% of the cases when *Face (Complete)* is present but the *Face (Partial)* is absent. Instead, for Llama-3.2-11B-Vision-Instruct, it is 19% for *Face (Partial)* being present and *Face (Complete)* absent, and 5% vice versa. For Gemma-3-4b-it, the model predicts only *Face (Complete)*, and only *Face (Partial)* 7% and 9% of the cases accordingly.

Manual inspection. For each of the five attributes, we randomly select 50 images where the VLM label is present and the human label is absent, and manually analyse them. We re-annotate the selected images with two annotators per attribute, following the VISPR definition of attributes. We define a subject in annotation instructions as a living person; therefore, we do not consider attributes related to drawings, statues, and dolls as present. The disagreements were resolved further in the discussion. As a result, we found that, for the majority of such images related to *Age Group* and *Spectators* attributes, VLMs correctly identify attributes that were initially overlooked by VISPR human annotators. However, VLMs erroneously detect attributes in photos of statues, drawings of people, religious illustrations, and animals.

5. CONCLUSION

We analysed and evaluated the ability of VLMs to recognise privacy attributes in images. We identified the attributes for which VLMs are in strong agreement when annotating, such as *Age Group*, *Gender*, *Hair Color*, *Spectators*, *Face (Partial)*, and analysed disagreement between model predictions and human annotations. We found that VLMs can help to detect attributes such as *Age Group*, *Gender*, and *Hair Color*, which are at times omitted by human annotators. Finally, we observed that VLMs often tend to annotate images with a *Face (Partial)* rather than a *Face (Complete)* attribute and, for the *Spectators* attribute, either correctly identify it or predict related relationship attributes in the image. For these attributes, the disagreements with human annotations often stem from human annotation errors or the presence of related attributes in the image.

6. REFERENCES

- [1] T. Orekondy, B. Schiele, and M. Fritz, “Towards a Visual Privacy Advisor: Understanding and predicting privacy risks in images,” in *Proc. IEEE International Conference on Computer Vision*, Venice, Oct. 2017, pp. 3706–3715.
- [2] R. G. Kost and J. Correa Da Rosa, “Impact of survey length and compensation on validity, reliability, and sample characteristics for ultrashort-, short-, and long-research participant perception surveys,” *Clinical and Translational Science*, vol. 2, no. 1, pp. 31–37, Feb. 2018.
- [3] G. A. Miller, “The magical number seven, plus or minus two: Some limits on our capacity for processing information,” *Psychological Review*, vol. 63, no. 2, pp. 81–97, Mar. 1956.
- [4] C.-H. Chiang and H.-Y. Lee, “Can Large Language Models be an alternative to human evaluations?,” in *Annual Meeting of the Association for Computational Linguistics*, July 2023.
- [5] H. Lu and F. Zhong, “Can Vision-Language Models replace human annotators: A case study with CelebA dataset,” *arXiv preprint arXiv:2410.09416*, 2024.
- [6] Z. Liu, P. Luo, X. Wang, and X. Tang, “Deep learning face attributes in the wild,” in *Proc. IEEE International Conference on Computer Vision*, USA, 2015, p. 3730–3738.
- [7] J. Zhang, X. Cao, Z. Han, S. Shan, and X. Chen, “Multi-P²A: A multi-perspective benchmark on privacy assessment for Large Vision-Language Models,” *arXiv preprint arXiv:2412.19496*, 2024.
- [8] Y. Zhang, Y. Huang, Y. Sun, C. Liu, Z. Zhao, Z. Fang, Y. Wang, H. Chen, X. Yang, X. Wei, H. Su, Y. Dong, and J. Zhu, “Multi-Trust: A comprehensive benchmark towards trustworthy Multimodal Large Language Models,” in *Proc. International Conference on Neural Information Processing Systems*, Red Hook, NY, USA, 2025.
- [9] J. Zhang, Z. Yuan, Z. Wang, B. Yan, S. Wang, X. Cao, Z. Guo, S. Shan, and X. Chen, “REVAL: A comprehension evaluation on reliability and values of large vision-language models,” *arXiv preprint arXiv:2503.16566*, 2025.
- [10] S. Schultenkämper and F. S. Bäumer, “Pixels versus privacy: Leveraging Vision-Language Models for sensitive information extraction,” *Advances in Security*, vol. 17, no. 1-2, pp. 1–10, 2024.
- [11] B. Tömekçe, M. Vero, R. Staab, and M. Vechev, “Private attribute inference from images with Vision-Language Models,” in *Proc. Advances in Neural Information Processing Systems*, 2024.
- [12] D. Gurari, Q. Li, C. Lin, Y. Zhao, A. Guo, A. Stangl, and J. P. Bigham, “VizWiz-Priv: A dataset for recognizing the presence and purpose of private visual information in images taken by blind people,” in *Proc. of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, June 2019.
- [13] C. Xu, J. Zhang, Z. Chen, C. Xie, M. Kang, Y. Potter, Z. Wang, Z. Yuan, A. Xiong, Z. Xiong, C. Zhang, L. Yuan, Y. Zeng, P. Xu, C. Guo, A. Zhou, J. Z. Tan, X. Zhao, F. Pinto, Z. Xiang, Y. Gai, Z. Lin, D. Hendrycks, B. Li, and D. Song, “MMDT: Decoding the trustworthiness and safety of multimodal foundation models,” in *Proc. International Conference on Learning Representations*, 2025.
- [14] Y. Yang, S. Wang, D. Li, S. Sun, and Q. Wu, “GeoLocator: A location-integrated Large Multimodal Model (LMM) for inferring geo-privacy,” *Applied Sciences*, vol. 14, no. 16, 2024.
- [15] M. Li, L. Li, Y. Yin, M. Ahmed, Z. Liu, and Q. Liu, “Red teaming Visual Language Models,” in *Findings of the Association for Computational Linguistics*, Bangkok, Thailand, Aug. 2024, pp. 3326–3342.
- [16] Gemma Team, A. Kamath, J. Ferret, S. Pathak, N. Vieillard, R. Merhej, S. Perrin, T. Matejovicova, A. Ramé, M. Rivière, et al., “Gemma 3 technical report,” *arXiv preprint arXiv:2503.19786*, 2025.
- [17] S. Bai, K. Chen, X. Liu, J. Wang, W. Ge, S. Song, K. Dang, P. Wang, S. Wang, J. Tang, et al., “Qwen2.5-VL technical report,” *arXiv preprint arXiv:2502.13923*, 2025.
- [18] A. Grattafiori, A. Dubey, A. Jauhri, A. Pandey, A. Kadian, A. Al-Dahle, A. Letman, A. Mathur, A. Schelten, A. Vaughan, et al., “The Llama 3 herd of models,” *arXiv preprint arXiv:2407.21783*, 2024.
- [19] H. Li, W. Hu, H. Jing, Y. Chen, Q. Hu, S. Han, T. Chu, P. Hu, and Y. Song, “PrivaCI-bench: Evaluating privacy with contextual integrity and legal compliance,” in *Proc. Annual Meeting of the Association for Computational Linguistics*, July 2025.
- [20] J. L. Fleiss, “Measuring nominal scale agreement among many raters.,” *Psychological bulletin*, vol. 76, no. 5, pp. 378, 1971.