DIPA: An Image Dataset with Cross-cultural Privacy Concern Annotations

Anran Xu
Interactive Intelligent Systems Lab.,
The University of Tokyo
Tokyo, Japan
anran@iis-lab.org

Zhongyi Zhou
Interactive Intelligent Systems Lab.,
The University of Tokyo
Tokyo, Japan
zhongyi@iis-lab.org

Kakeru Miyazaki
Interactive Intelligent Systems Lab.,
The University of Tokyo
Tokyo, Japan
kakeru-miyazaki@iis-lab.org

Ryo Yoshikawa
Interactive Intelligent Systems Lab.,
The University of Tokyo
Tokyo, Japan
ryo@iis-lab.org

Simo Hosio
Center for Ubiquitous Computing,
University of Oulu
Oulu, Finland
simo.hosio@oulu.fi

Koji Yatani
Interactive Intelligent Systems Lab.,
The University of Tokyo
Tokyo, Japan
koji@iis-lab.org

ABSTRACT

Image privacy protection is an important topic in Human-Computer Interaction and usable security. Researchers have examined different aspects of image privacy by collecting samples by themselves. However, there does not exist a publicly-available dataset on image privacy, which prevents these efforts from sharing common technical foundations. We introduce DIPA, an open source dataset that provides content-level annotations that specifically focus on image privacy. We include 1,495 images from two existing datasets in DIPA, and augment them with 5,671 annotations. Each annotation includes reasons why the associated visual content can be privacy-threatening, a rating of how informative annotators thought the associated content is to threaten privacy, and another rating of how broadly the image could be shared. We also collected annotations from people living in Japan and UK to enable researchers and developers to perform analysis from the perspective of cultural differences. In this paper, we present the construction procedure of DIPA and report high-level statistics of the data we obtained. We hope that DIPA would accelerate various future research, including quantitative understandings of cultural differences on perceptions of image privacy and the development of robust recognition models for image privacy protection.

CCS CONCEPTS

• **Security and privacy** → **Privacy protections**; **Usability in security and privacy**.

KEYWORDS

Image privacy, usable security

ACM Reference Format:
Anran Xu, Zhongyi Zhou, Kakeru Miyazaki, Ryo Yoshikawa, Simo Hosio, and Koji Yatani. 2023. DIPA: An Image Dataset with Cross-cultural Privacy Concern Annotations. In *28th International Conference on Intelligent User Interfaces (IUI '23 Companion), March 27–31, 2023, Sydney, NSW, Australia.* ACM, New York, NY, USA, 8 pages. https://doi.org/10.1145/3581754.3584176

IUI '23 Companion, March 27–31, 2023, Sydney, NSW, Australia

ACM ISBN 979-8-4007-0107-8/23/03.
https://doi.org/10.1145/3581754.3584176

1 INTRODUCTION

Image privacy protection is a critical topic in usable security and Human-Computer Interaction. Although the use of social network software, crowdsourcing platforms, and other image sharing platforms increase the possibility of exposing privacy through images, there exist rarely open-sourced resources [14, 23] that help researchers understand privacy concerns on images from users' perspective. Defining privacy in images is ambiguous and subjective because it involves multiple factors, such as personal preferences, sharing scenarios and social relationships. This leads to different definitions and perspectives of image privacy in existing work. For example, researchers developed different recognition methods for detecting privacy-threatening images [2, 5, 11, 17, 19, 20, 23], but each of them had different criteria of what content can be privacy-threatening. To discover commonalities of privacy, researchers attempted to summarize privacy-threatening content under their given sharing contexts [10, 13]. Other work reckoned that privacy is a user-specific issue, and they built personalized recommendation models to mitigate users' concerns [22, 24]. While prior work successfully broadened the research landscape of image privacy, researchers would need to perform their own surveys and data collection, which is a large obstacle in this field. A publicly-available data resource about image privacy would thus contribute to the advance of artificial intelligence technology for future usable security. Furthermore, it would greatly accelerate the open science of usable security research by providing opportunities to deeply understand user perception of and attitude toward image privacy. For instance, social scientists could take the advantage of such a dataset to quantitatively examine the impact of cultural backgrounds on image privacy. We, therefore, argue that a dataset on image privacy would bring an impact not only on the computer science field but also on other disciplines related to people's privacy, security, and media use.

The primary contribution of this work is DIPA[1] (a Dataset with Image Privacy Annotations), a publicly-available dataset that contains annotations of privacy concerns (i.e., what visual contents in an image can be considered privacy-threatening) on images already available in public datasets (OpenImages [12] and LVIS [9]). We collected such annotations through asking crowdworkers to perform object annotation tasks along with how privacy-threatening they consider each of the annotated objects. Our data collection was executed in two crowdsourcing platforms (CrowdWorks [8] and Prolific [15]) so that our dataset includes annotations by people in Japan and UK. The annotations in DIPA thus cover cultural diversity, which is one critical factor for image privacy perception [6, 21]. More specifically, DIPA has the following features.

- DIPA provides 1,495 images containing 5,671 content-level annotations about privacy concerns. Each annotation accompanies reasons why the associated visual content can be privacy-threatening, a rating of how informative annotators thought the associated content is to threaten privacy, and another rating of how broadly the image could be shared.
- Annotations available in DIPA are provided by annotators in Japan and UK recruited from two different crowdsourcing platforms. This allows researchers and developers to perform quantitative comparisons on perceived image privacy from a cross-cultural perspective.
- DIPA also provides annotators' demographic information and Big Five personality test results with anonymization. Researchers and developers may use these data to study relationships between personality traits and image privacy preferences.

2 DATASET CONSTRUCTION

Building a public dataset on image privacy concerns is a challenging research topic. Using users' own photos would reflect the most realistic settings, but it is not feasible from a privacy perspective. A dataset would also need a large number of images for statistical analysis and machine learning applications. We, therefore, decided to utilize existing image datasets (OpenImage [12] and LVIS [9]) for our dataset. A challenge associated with this approach is that generic image datasets like OpenImages contain many non-privacy-threatening visual contents. Performing annotations in every image in such datasets would be inefficient to achieve our purpose.

We employed a two-stage data collection method to avoid privacy leakage from the participants in our study and collect annotated images efficiently. In the first stage, we designed a formative study to derive common categories of privacy-threatening contents in images. In the second stage, we chose images from existing open-sourced image datasets [9, 12] that include at least one privacy-threatening type of contents in the categories we previously derived. In this manner, we are able to use images that are already publicly-available and efficiently collect annotations on contents our participants consider privacy-threatening. All the data collection procedures received approval from our institutional review board prior to the execution.

[1]The dataset is available at: https://dipa-download.s3.ap-northeast-1.amazonaws.com/dataset.zip

2.1 Stage I – Deriving Common Privacy-Threatening Contents

This part of the data collection involves the derivation of common privacy-threatening contents in images. Li et al. conducted a similar study [13], and we expand it by incorporating a hypothetical scenario where aggressive attackers would try to identify the owners' identity while considering cultural differences. In particular, Li et al.'s study was conducted in the U.S., and we obtained annotations from people in Japan and UK in our study. This expansion would increase the generalizability of the dataset by mitigating issues that may stem from cultural specificity (e.g., guns were one privacy-threatening content category in Li et al.'s study, but they are prohibited in many countries).

In this part of the study, participants were required to input their basic demographic information (age, gender, and nationality) and complete a brief Big-five personality inventory [16]. We next required them to choose 10 photos that were taken by themselves and had been shared online. We asked participants to assume that malicious privacy attackers would try their best to identify you from each of their photos. They were then instructed to perform annotations on any visual content that could be exploited by such attackers and provide reasons as well as brief descriptions of the content. In addition, we offered a blurring tool in case participants wanted to hide specific content before submitting.

To highlight the differences by cultural background, we conducted the data collection described above in Japan. We recruited 200 participants through CrowdWorks [8], one of the largest crowdsourcing platforms in Japan. The interface and instructions were provided in Japanese. We paid approximately 3 USD at the completion of the task.

2.1.1 Results. 171 of 200 participants (69 males, 78 females, and 4 prefer not to say) successfully provided valid annotated images. Participants uploaded 1,632 photos with 1,894 annotations, many of which included rich descriptions. We carefully examined each annotation and performed categorization. The first author conducted the initial categorization. We then collaboratively iterated our categorization, ultimately summarized in 25 categories (Table 1). Besides direct privacy leakages from appearance or identity, our participants also expressed their concerns that malicious people might exploit various indirect information in images, such as backgrounds that revealed the photo owner's living places or professional books indicated careers. We also observed that our participants expressed different levels of severity in their reasoning.

2.2 Stage II – Collecting Privacy Concern Annotations

This part of the data collection focused to collect privacy concern annotations on images from existing datasets. We first extracted images that contain at least one category of privacy-threatening contents we identified from two large-scale datasets – OpenImages [12] and LVIS [9]. We aimed to extract 50 images from each image dataset for each category though images for several categories are not immediately available. For the categories of "Cosmetics" and "Photo", there are only 46 and 44 images, respectively.

Table 1: Twenty-five common categories of privacy-threatening content observed in the first study. We excluded "Nudity," "Accident," and "Bystander" for the second study because contents in these categories are either unavailable or indistinguishable.

Category [% (count)]	Description
Person (except bystanders) [43% (809)]	People who are intended to be photographed
Place Identifier [17% (328)]	Road signs, building signs, maps, and other environmental hints that indicate locations
Identity [8% (155)]	Identity information on tickets, passports, nameplates, etc
Home Interior [5% (91)]	Furniture and home decoration that may imply people's habits and locations
Vehicle Plate [4% (84)]	Identifiable information for cars but also for their owners
Bystander [4% (72)]	People who are photographed without permission.
Food [4% (68)]	Close-up views of food or party scenes.
Printed Materials [3% (62)]	Various printed materials containing private information.
Screen [3% (57)]	Computer monitors, smartphone screens, electronic information boards, etc
Clothing [2% (46)]	Clothing that may imply personal identity, habits, and occupations.
Scenery [2% (42)]	Backgrounds which may imply locations or personal information
Pet [1% (23)]	Pet ownership can be private information to some people.
Book [1% (12)]	Books that may imply personal information and preferences
Photo [1% (11)]	Other photos in the captured image
Machine [0% (9)]	Machines used in a workplace or specific areas
Table [0% (5)]	Tables with many personal items
Electronic Devices [0% (5)]	Electronic devices that photo owners' regard as private
Cosmetics [0% (4)]	Personal care products that may reveal owners' habits or locations
Toy [0% (4)]	Toys for children or photo owners
Finger [0% (2)]	Finger close-up that can be used to infer a person's fingerprint
Cigarettes [0% (1)]	Cigarettes or smoking scenes
Accident [0% (1)]	Accident scenes
Musical Instrument [0% (1)]	Instruments or playing scenes.
Nudity [0% (1)]	Naked upper body
Accessory [0% (1)]	Accessories worn by people photographed
Total [100% (1894)]	

Images for the categories of "Nudity" and "Accident' are not available in these datasets. Furthermore, the category of "Bystander" is not readily identifiable. As a result, we extracted 2,090 images for 22 categories (all without "Nudity", "Accident" and "Bystander"). Please note that these images also contain various visual contents including those that are not in the 22 categories.

We recruited participants through two crowdsourcing platforms: CrowdWorks [8] and Prolific [15]. 200 participants who resided in Japan and UK were recruited through CrowdWorks and Prolific, respectively. We provided the annotation interface and task descriptions in Japanese and English for CrowdWorks and Prolific participants, respectively. We paid them approximately 3 dollars in their local currency for their successful task completion.

Participants were first asked to provide their basic demographic information and fill out the Big-five personality questionnaire similar to our formative study. Participants were then asked to annotate privacy-threatening content in ten randomly-given images with our custom annotation interface (Figure 1). This interface highlighted all the visual contents for which object annotations are already available in the existing datasets. In addition, participants were asked to add object annotations in case they found anything in images that can be considered privacy-threatening but not highlighted yet. Participants were then asked to provide the following information for each highlighted object if they agreed that it was privacy-threatening.

- **Information Type.** We asked participants to clarify what kind of information about a photo owner they were able to infer from the given content. Participants were provided with four default options of "personal identity", "location of shooting", "personal habits", and "social circle" as well as a free-form textbox.
- **Informativeness.** Participants were also asked to rate the perceived degree of severity of privacy threats of each object on a 7-Likert scale (1: Extremely uninformative – 4: Neutral – 7: Extremely informative) [5, 19, 20, 23]. In these responses, we used the terms "uninformative" and "informative" to mean that content conveys a small or large amount of information related to privacy, respectively.
- **Maximum Sharing Scope** We also asked participants to judge how broadly they would be willing to share the given image if they were the owner [3, 18].

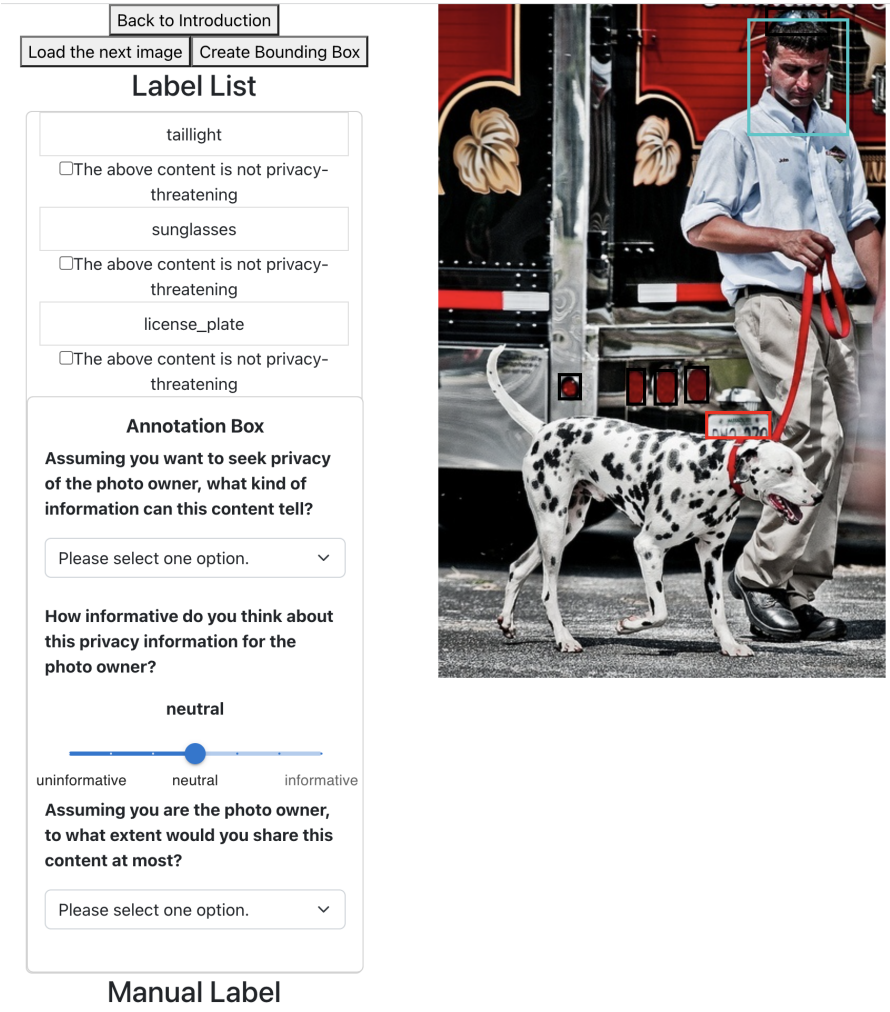

Figure 1: The annotation interface used in our second study. The red bounding box (surrounding the vehicle plate) represents the current annotation area. The light blue bounding box (surrounding the person's face) represents a manually-highlighted privacy-threatening content by participants. For each highlighted object (including those created by participants themselves), if participants considers the corresponding content to be privacy-threatening, they were asked to provide more details. Otherwise, they simply ticked a checkbox and moved on to the next object.

We derived the following four choices based on the previous studies [1, 4, 7, 13] – "I won't share it at all", "Family or friend", "Public", and "Broadcast programs" as well as a free-form text box.

3 DATASET ANALYSIS

We filtered participants who did not provide any annotation in the task. As a result, 177 and 183 participants recruited from CrowdWorks and Prolific successfully completed their tasks, respectively. Table 2 details the demographic information of our participants. Out of 2,090 images in our initial collection (Section 2.2), participants from CrowdWorks and Prolific identified 1,244 and 949 images that they claim contain at least one privacy-threatening content, respectively. After combining overlaps, we obtained 1,495 unique images that contain one or more annotated privacy-threatening contents by at least one participant, all of which are included in our DIPA dataset. Images without any annotation of privacy-threatening contents were excluded from this study and our dataset.

We obtained 5,671 annotations for visual contents that are considered privacy-threatening in the 1,495 images.

One visual content can be annotated up to twice as we assigned two participants to annotate each image. The 5,671 annotations contain 3,613 provided by one participant for the same amount of visual contents. Two participants provided their annotations for 1,029 visual contents, resulting in 2,058 annotations.

Among the 5,671 annotations, 3,170 belong to the 22 privacy-threatening categories we derived. 2,252 annotations were made for visual contents that are not in our 22 categories but where object

Table 2: Demographic information of participants in our second study.

	Age Range					All
	18–24	25–34	35–44	45–54	55–	
CrowdWorks						**177**
Male	3	15	31	34	11	94
Female	3	21	29	18	7	78
Others	0	0	3	2	0	5
Prolific						**183**
Male	21	23	21	12	7	84
Female	11	19	29	9	15	83
Others	0	5	2	5	4	16

Table 3: Category-wise distribution in DIPA. 249 extra annotations manually added by participants were categorized into "other categories". The binomial test confirmed that privacy-threatening content was more likely to appear in 22 categories (excluding "Nudity", "Accident", and "Bystander") of privacy-threatening content we identified.

Category	CrowdWorks (Japan)		Prolific (UK)	
	Privacy-threatening	Not privacy-threatening	Privacy-threatening	Not privacy-threatening
Person (except bystanders)	1, 034	584	644	937
Place Identifier	68	38	32	63
Identity	9	5	7	6
Home Interior	36	71	12	90
Vehicle Plate	74	24	74	25
Food	25	50	10	54
Printed Materials	40	42	30	50
Screen	103	58	77	74
Clothing	158	174	72	251
Scenery	44	40	24	60
Pet	56	44	33	59
Book	56	59	32	70
Photo	13	16	6	20
Machine	21	36	13	43
Table	70	129	15	169
Electronic Devices	18	80	15	85
Cosmetics	15	11	7	17
Toy	23	33	6	47
Finger	22	83	9	83
Cigarettes	20	57	13	60
Musical Instrument	30	58	14	73
Accessory	57	68	33	89
Sum of the above	1, 992	1, 760	1, 178	2, 425
Other categories	1, 643	2, 701	858	3, 238

annotations exist in the original image datasets. The remaining 249 annotations were newly added by those participants by hand.

Table 3 details the distributions of annotations for visual contents according to crowdsourcing platforms. Our binomial test to

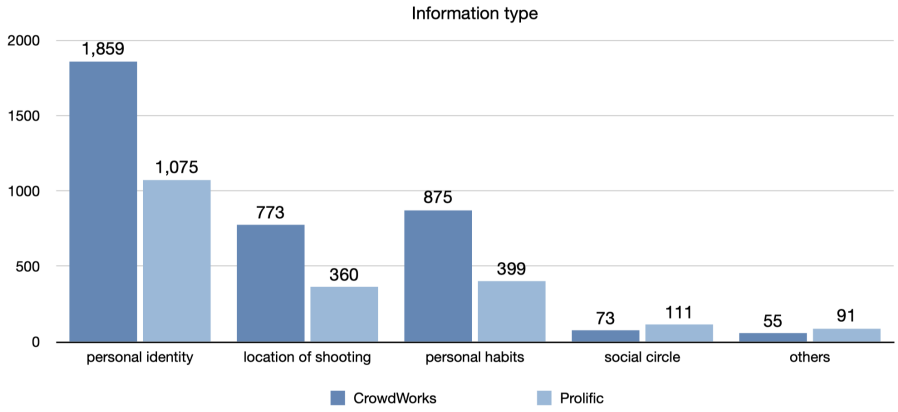

Figure 2: The distributions of information type of privacy-threatening contents annotated by participants in both crowdsourcing platforms.

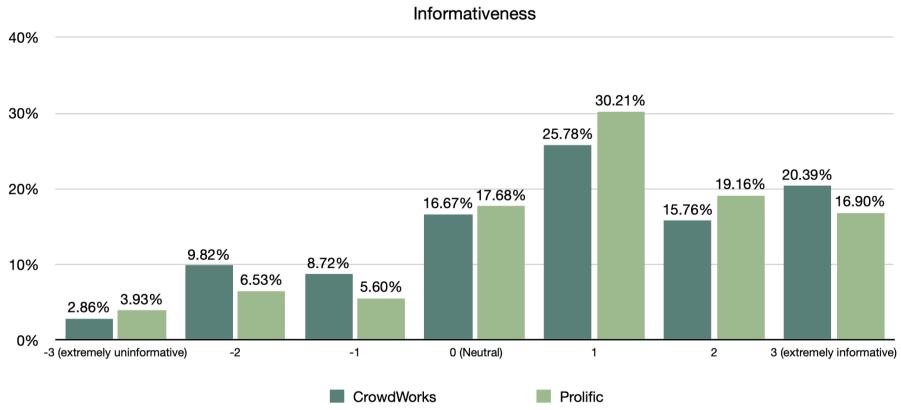

Figure 3: The distribution of informativeness scores provided by participants in both crowdsourcing platforms.

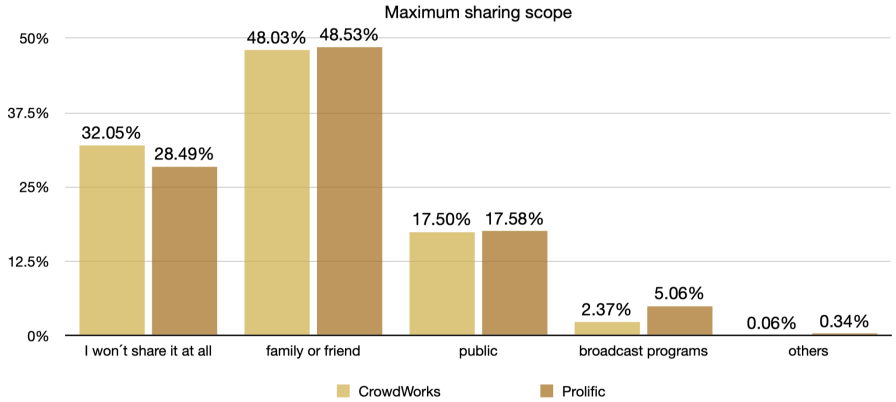

Figure 4: The distributions of maximum sharing scopes provided by participants in both crowdsourcing platforms.

compare the probability of annotations by Prolific participants for privacy-threatening contents against those by CrowdWorks participants revealed a significant difference (p<.001, 95%CI: [0.31, 0.34]).

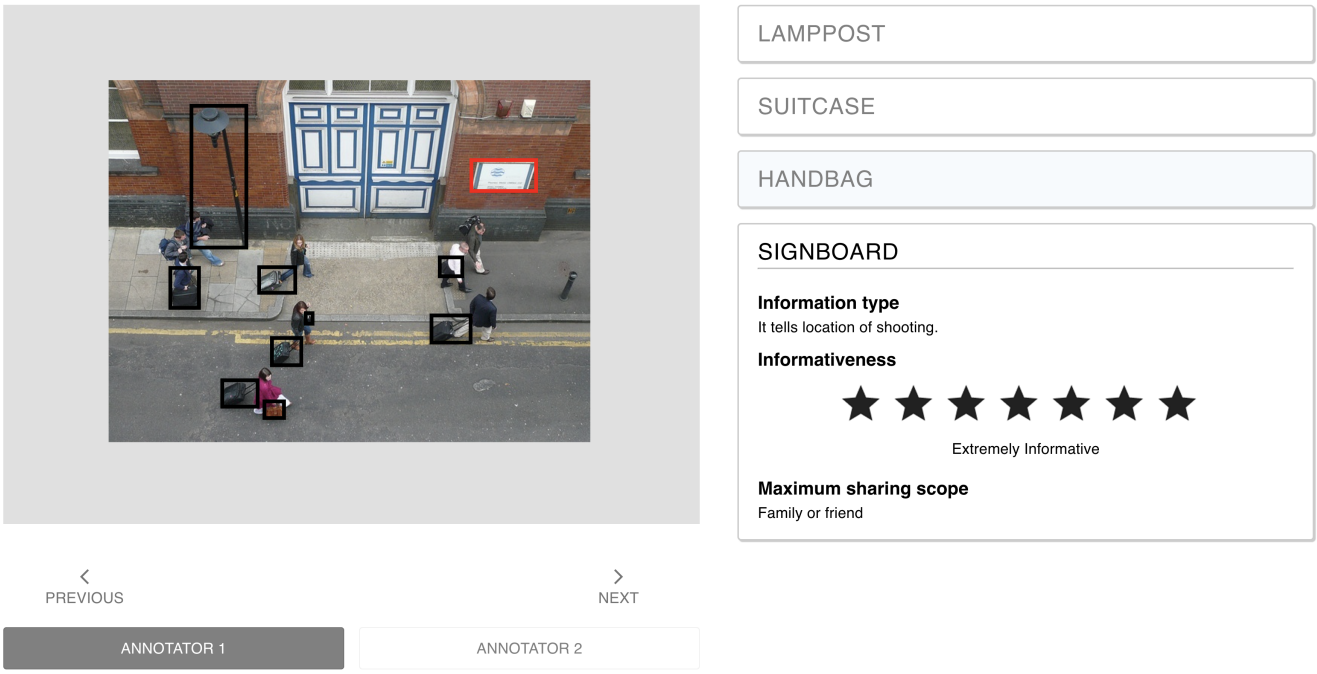

Figure 5: The online interface for DIPA. It presents an image and its annotation, allowing users to examine our dataset without needing technical skills.

This result confirms that CrowdWorks participants had a tendency of flagging up more contents as privacy-threatening than Prolific participants.

Figure 2 shows the distributions of the information types in the annotations. The category of personal identity was most frequent. There are also differences in the occurrences of annotations by participants from CrowdWorks and Prolific. Our Chi-square tests for each information type revealed significant differences except personal identity (personal identity: $\chi^2(1)=1.37$, $p=.242$, location of shooting: $\chi^2(1)=10.26$, $p<.001$, personal habits: $\chi^2(1)=14.74$, $p<.001$, social circle: $\chi^2(1)=48.21$, $p<.001$, others: $\chi^2(1)=46.72$, $p<.001$). These results also suggest that Crowdworks annotators considered more visual contents as privacy-threatening than Prolific annotators.

Figure 3 shows the distributions of the informativeness scores rated by CrowdWorks and Prolific participants in percentage. The two distributions show seemingly similar patterns. Figure 4 presents the distributions of the maximum sharing scopes rated by CrowdWorks and Prolific participants in percentage. For the majority of the images, our participants considered that they would be willing to share with up to families and friends. Future work should investigate what differences exist in the annotations in our dataset and how cultural backgrounds could contribute to such differences.

We also provide an online interface to see each image in our dataset (Figure 5). The interface presents detailed information about each annotation (information type, informativeness, and maximum sharing scope)[2]. This interface would help non-technical users to review our dataset.

4 CONCLUSION

We present DIPA – an open-source image dataset that provides content-level annotations that focus on how these contents can be privacy-threatening. Our dataset includes 1,495 images from OpenImages [12] and LVIS [9] with augmentation of annotations about how their corresponding visual contents are perceived as privacy-threatening. The dataset contains 5,671 annotations that include three kinds of information about the perceived threat of privacy: reasons for the threat of privacy, informativeness of the corresponding content with respect to privacy, and the broadest possible scope for sharing. DIPA would enable various quantitative research in image privacy. We thus hope that DIPA would stimulate open science for usable security and broader research.

REFERENCES

[1] Shane Ahern, Dean Eckles, Nathaniel S Good, Simon King, Mor Naaman, and Rahul Nair. 2007. Over-exposed? Privacy patterns and considerations in online and mobile photo sharing. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. 357–366.

[2] Rawan Alharbi, Mariam Tolba, Lucia C Petito, Josiah Hester, and Nabil Alshurafa. 2019. To mask or not to mask? balancing privacy with visual confirmation utility in activity-oriented wearable cameras. *Proceedings of the ACM on interactive, mobile, wearable and ubiquitous technologies (IMWUT '19)* 3, 3 (2019), 1–29.

[3] Mary Jean Amon, Rakibul Hasan, Kurt Hugenberg, Bennett I Bertenthal, and Apu Kapadia. 2020. Influencing photo sharing decisions on social media: A case

[2]The online interface for DIPA is available at: https://anranxu.github.io/DIPA_visualization/

of paradoxical findings. In *2020 IEEE Symposium on Security and Privacy (SP)*. IEEE, 1350–1366.
[4] K Caine. 2008. Linking Studies of HCI to Psychological Theories of Privacy. *Georgia Institute of Technology* (2008).
[5] Daniel Castro, Steven Hickson, Vinay Bettadapura, Edison Thomaz, Gregory Abowd, Henrik Christensen, and Irfan Essa. 2015. Predicting daily activities from egocentric images using deep learning. In *proceedings of the 2015 ACM International symposium on Wearable Computers*. 75–82.
[6] Hichang Cho, Milagros Rivera-Sánchez, and Sun Sun Lim. 2009. A multinational study on online privacy: global concerns and local responses. *New media & society* 11, 3 (2009), 395–416.
[7] Sunny Consolvo, Ian E Smith, Tara Matthews, Anthony LaMarca, Jason Tabert, and Pauline Powledge. 2005. Location disclosure to social relations: why, when, & what people want to share. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. 81–90.
[8] CrowdWorks. 2022. https://crowdworks.jp. Accessed: 2022-10-31.
[9] Agrim Gupta, Piotr Dollar, and Ross Girshick. 2019. Lvis: A dataset for large vocabulary instance segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 5356–5364.
[10] Hana Habib, Neil Shah, and Rajan Vaish. 2019. Impact of contextual factors on snapchat public sharing. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–13.
[11] Mohammed Korayem, Robert Templeman, Dennis Chen, David Crandall, and Apu Kapadia. 2016. Enhancing lifelogging privacy by detecting screens. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. 4309–4314.
[12] Alina Kuznetsova, Hassan Rom, Neil Alldrin, Jasper Uijlings, Ivan Krasin, Jordi Pont-Tuset, Shahab Kamali, Stefan Popov, Matteo Malloci, Alexander Kolesnikov, et al. 2020. The open images dataset v4. *International Journal of Computer Vision* 128, 7 (2020), 1956–1981.
[13] Yifang Li, Nishant Vishwamitra, Hongxin Hu, and Kelly Caine. 2020. Towards a taxonomy of content sensitivity and sharing preferences for photos. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI '20)*. 1–14.
[14] Tribhuvanesh Orekondy, Bernt Schiele, and Mario Fritz. 2017. Towards a visual privacy advisor: Understanding and predicting privacy risks in images. In *Proceedings of the IEEE international conference on computer vision*. 3686–3695.
[15] Prolific. 2022. https://www.prolific.co/. Accessed: 2022-10-31.
[16] Beatrice Rammstedt and Oliver P John. 2007. Measuring personality in one minute or less: A 10-item short version of the Big Five Inventory in English and German. *Journal of research in Personality* 41, 1 (2007), 203–212.
[17] Julian Steil, Marion Koelle, Wilko Heuten, Susanne Boll, and Andreas Bulling. 2018. Privaceye: Privacy-preserving first-person vision using image features and eye movement analysis. *arXiv preprint arXiv:1801.04457* (2018).
[18] Jose M Such, Joel Porter, Sören Preibusch, and Adam Joinson. 2017. Photo privacy conflicts in social media: A large-scale empirical study. In *Proceedings of the 2017 CHI conference on human factors in computing systems*. 3821–3832.
[19] Ashwini Tonge and Cornelia Caragea. 2020. Image privacy prediction using deep neural networks. *ACM Transactions on the Web (TWEB '20)* 14, 2 (2020), 1–32.
[20] Lam Tran, Deguang Kong, Hongxia Jin, and Ji Liu. 2016. Privacy-cnh: A framework to detect photo privacy with convolutional neural network using hierarchical features. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 30.
[21] Ho Keung Tsoi and Li Chen. 2011. From privacy concern to uses of social network sites: A cultural comparison via user survey. In *2011 IEEE Third International Conference on Privacy, Security, Risk and Trust and 2011 IEEE Third International Conference on Social Computing*. IEEE, 457–464.
[22] Nishant Vishwamitra, Yifang Li, Hongxin Hu, Kelly Caine, Long Cheng, Ziming Zhao, and Gail-Joon Ahn. 2022. Towards Automated Content-based Photo Privacy Control in User-Centered Social Networks. In *Proceedings of the Twelveth ACM Conference on Data and Application Security and Privacy*. 65–76.
[23] Sergej Zerr, Stefan Siersdorfer, Jonathon Hare, and Elena Demidova. 2012. Privacy-aware image classification and search. In *Proceedings of the 35th international ACM SIGIR conference on Research and development in information retrieval*. 35–44.
[24] Haoti Zhong, Anna Cinzia Squicciarini, David J Miller, and Cornelia Caragea. 2017. A Group-Based Personalized Model for Image Privacy Classification and Labeling.. In *International Joint Conference on Artificial Intelligence (IJCAI '17)*, Vol. 17. 3952–3958.