

Proposing an Interactive Audit Pipeline for Visual Privacy Research

Jasmine DeHart, Chenguang Xu, Christan Grant
 School of Computer Science
 University of Oklahoma
 Norman, Oklahoma, USA

Lisa Egede
 Human Computer Interaction Institute
 Carnegie Mellon University
 Pittsburgh, PA, USA

Abstract—In an ideal world, deployed machine learning models will enhance our society. We hope that those models will provide unbiased and ethical decisions that will benefit everyone. However, this is not always the case; issues arise during the data preparation process throughout the steps leading to the models' deployment. The continued use of biased datasets and biased processes will adversely damage communities and increase the cost to fix the problem later. In this work, we walk through the decision making process that a researcher should consider before, during, and after a system deployment to understand the broader impacts of their research in the community. Throughout this paper, we discuss fairness, privacy, and ownership issues in the machine learning pipeline, assert the need for a responsible human-over-the-loop methodology to bring accountability into machine learning pipeline, and finally, reflect on the need to explore research agendas that have harmful societal impacts. We examine visual privacy research and draw lessons that can apply broadly to artificial intelligence. Our goal is to provide a systematic analysis of the machine learning pipeline for visual privacy and bias issues. With this pipeline, we hope to raise stakeholder (e.g., researchers, modelers, corporations) awareness as these issues propagate in the various machine learning phases.

Index Terms—visual privacy, fairness, human-over-the-loop

I. INTRODUCTION

As society progresses, humans are becoming more dependent on the accessibility and convenience that technology offers. Everyday a large amount of visual content is uploaded to Social Media Networks (SMNs) from billions of users across the globe, which can explain the large amounts of sensitive data that is available online. While these ecosystems have goals that revolve around helping people build connections with others; there are gaps in the methods used to protect the information for individuals and corporations who share or collect content [1]–[7]. The increased upload of images and videos emphasize the need for privacy protection and mitigation strategies for visual content. Visual privacy techniques extend from SMNs to connected networks, smart cities, lifelogging, and much more [8]. Various harms can occur as a result of sensitive information being displayed, which makes visual privacy a growing area of concern [2], [3], [9]. Existing technologies in the industry can show a disregard for protecting the information of individuals who share visual content or for individuals who are captured in the content [4], [6]. While research is being done to address these concerns, there exists a gap in understanding the overlap between

fairness, privacy, and user-feedback for issues regarding visual privacy in the machine learning (ML) pipeline.

A need for visual privacy has emerged from SMNs and the integration of Internet of Things (IoT) devices that can expose sensitive information through visual content [10]–[12]. The constant sharing and storing of videos and images bring skepticism about individual privacy and rights [13], [14]. Researchers have created datasets, models, and deployed applications that they believe will provide privacy to its' users [14]–[20]. Within these algorithms and systems, researchers should continually make decisions to assess the fairness, privacy, accessibility of the data and model in regard to the communities they serve. Bias can be curated from the data collection process, reinforced in the model's training, and systematically imposed in the deployment phase [21]. With privacy and bias issues arising throughout the ML pipeline it provokes the question: *does the impact from model development procedures outweigh the societal benefits?*

The goal of this paper is two-fold. First, the aim is to understand visual privacy and fairness as their issues intrude into the ML pipeline and potentially impact the stakeholders and community where the pipeline is deployed. Secondly, we provide a comprehensive pipeline indicating fairness and privacy issues and propose an auditing strategy to reduce these affects in visual privacyresearch. In this paper, we discuss several privacy, fairness, and ownership issues that can arise in the ML pipeline (§ II, III). We argue for the use of humanover-the-loop strategies to discover privacy and fairness issues in ML. We extend this technique to suggest two auditing processes: Fairness Forensics Auditing System (FAST) and Visual Privacy (ViP) Auditor (§ IV). Finally, reflect on the need to pursue research agendas that have harmful societal impacts (§ V). We continue our discussion on critical decisions that are often overlooked when deploying ML models in an extended version of this article [22].

II. DEFINING THE MACHINE LEARNING PIPELINE

We describe the ML pipeline as having three phases (Fig. 1). Phase 1 is the **Data Preparation** process. This phase includes considerations of (1) raw data sources, (2) data collection processes, (3) data storage, and (4) data cleaning processes that a researcher should explore before entering into the next phase. Data can come from anywhere and everywhere. With

so many data source possibilities available, the researcher should consider which sources are relevant to them. The data collection process for researchers can include using existing image datasets, social media datasets, or web scraping methods in respect to the visual privacy research task. Once a dataset is collected, a researcher could employ data cleaning tasks (e.g., crowd-sourced labeling) to derive an optimal dataset and labels.




Fig. 1. This figure illustrates the traditional ML pipeline. The pipeline has three phases: data preparation, modeling, and deployment.

In Phase 2, shown in Fig. 1, we begin the **Modeling** process. The cleaned data from Phase 1 can be divided into three datasets: training, testing, and validation. Training data is used as input for the ML algorithm. After training with the researcher’s desired ML algorithm, the researcher receives a model to run testing and validation datasets on. The model provides the output of the performance with several metrics. This new information can be used for refining the model before entering the final phase of the pipeline.

The last phase of the proposed machine learning pipeline is **Deployment**. The **Deployment** phase uses real-world data as input for the selected model. The researcher or end-user will see the real-world applications and results of their selected model from Phase 2. This phase allows the researcher to evaluate their model’s performance and impact in the communities that they serve.

A. The Guise of Pipeline Ownership

Researchers must consider who has ownership of the data and model at each phase before beginning these processes. These considerations are essential when protecting the privacy of individuals and biases that the owner could impose.

At the **Data Preparation** and **Deployment** phases, the researcher should consider who are the owners of the data and how they are receiving the content. This can explore if online images belong to the users or the corporation, if existing datasets belong to the proprietary researchers, or if agreements on volunteered data belong to companies or individuals. Furthermore, if the researcher is using online resources for data processing, the researcher should consider: how is the data stored, does it still belong to the researcher, and what information is being stored on these platforms.

In the **Modeling** and **Deployment** phases, the researchers should consider who holds ownership of the model. Considerations should be given to understand if the rights and ownership of the model are owned by the researcher, company, or from a third party [23]. The data uploaded to the model could lead to an individual perceiving that their privacy has been breached

due to the authorization and ownership of the model. If a third party owns the model, it is important to consider what information they are collecting from the use of it and who they share this information with. Throughout each phase of the ML pipeline, the stakeholders should continue to ask tough questions and make critical decisions that are ethical, fair, and in the best interest of those the technology is meant to serve.

III. EXPLORING PRIVACY AND FAIRNESS CONCERNS IN THE VISUAL PRIVACY ML PIPELINE

Efforts to implement technology that serves to mitigate harm are the motivation behind visual privacy research. In this section, we will discuss privacy and fairness issues that frequently occur with developing and deploying visual privacy systems. Examples of these issues are shown in Fig. 2 and Table I. We suggest that when evaluating visual privacy systems, researchers should consider bias issues as they arise in the ML pipeline.

A. Privacy

Visual privacy issues can arise at any point in the ML pipeline. The stakeholders must be aware of these issues and develop ways to solve them proactively as they arise.

1) *Obtaining Visual Content Consent*: Researchers can use large public image data sets [19], [32], [33] to train ML algorithms to perform various visual privacy research tasks [16], [17], [34]. Additionally, when collecting a large amount of data, many researchers question the use of *web scraping* methods to obtain this data [35]–[38]. Large data sets can be labeled using crowd-sourcing methods [32], [33], [39], [40]. While researchers’ efforts can focus on creating systems to help with visual privacy, their approach in collecting data can bring rise to privacy and ethical concerns in Phase 1 of the machine learning pipeline. The methods that researchers use to collect this data can overlook individuals’ privacy, consent, and protection. When collecting visual content or using existing datasets, researchers can un-intentionally collect private content containing minors or bystanders [41]–[45]. The topic of consent is essential to gauge participants’ willingness to participate in the study or research. For traditional studies that include people or living subjects, specific procedures and policies need to be followed according to a governing entity (i.e., institutional review board), so what excludes visual privacy research from policies and procedures when using personal data?

2) *Multiparty Conflict (MPC)*: Images and videos can be owned by multiple people [46]. Co-ownership issues can arise from several situations; a few are (1) individuals engaging in group photos, (2) a person responsible for other individuals (e.g., children, pets), (3) a person having physical possession of images of others on their device [47]. These types of conflict can affect the privacy of minors [48], [49] and bystanders [41], [44], [50]. This co-owned content can cause privacy leaks without it being the individual’s intent [51]. In the ML pipeline, the researcher should consider possible issues for MPC in all phases.

TABLE I

THIS TABLE DISPLAYS THE PRIVACY AND FAIRNESS ISSUES IN VARIOUS PHASES OF THE MACHINE LEARNING PIPELINE. THE DESCRIPTION PROVIDES A HIGH-LEVEL OVERVIEW OF WHAT THOSE ISSUES ARE. THE CHECKMARK (✓) INDICATES THAT THOSE ISSUES COULD ARISE IN THOSE PARTS OF THE PIPELINE.

			Phase 1				Phase 2				Phase 3		
			Raw Data	Data Collection	Data	Data Cleaning	Training Data	ML Training	Model	Model Analysis	Data	Model	Output
Privacy issues	Obtaining Content Consent	Exploring the ethics and privacy methodology of researchers to obtain consent for collected visual content on a public domain.		✓		✓					✓		✓
	Multiparty Conflict	Understanding privacy concerns for images and videos which are owned by multiple persons.	✓	✓	✓	✓		✓		✓			✓
	Image Removal Request	Determining when and how visual consent should be removed from the pipeline via requests.			✓	✓		✓	✓	✓	✓		✓
Fairness issues	Historical bias	The inherent bias from a biased world is absorbed by the source data. [21]	✓								✓		
	Algorithmic bias	The bias relates to the algorithm in the ML pipeline, and it could have different bias sources and types. [24] [25] [26]					✓	✓	✓	✓		✓	✓
	Software Discrimination	The output from a predictive software used to aid in decision making may lead to unfair consequences. [27]										✓	✓
	Individual fairness	Similar individuals should be treated as similarly as possible. [28]					✓	✓	✓	✓		✓	✓
	Group fairness	The groups defined by protected attributes should obtain similar treatments or equal opportunity as the privileged group. [29]					✓	✓	✓	✓		✓	✓
	Disparate treatment	Protected attributes are directed applied in the process of modeling where unfairness occurs. [30]					✓	✓	✓	✓		✓	✓
	Disparate impact	Even though the protected feature is not directly using, its relevant features still could lead a selection process to make unfair outputs. [31]					✓	✓	✓	✓		✓	✓

Considerations for content ownership and individual rights should be made early in the pipeline. When working with visual content, it can be necessary to seek permission from all parties involved. Multiparty Conflicts can enter the ML pipeline as early as the **Data Preparation** phase. In the **Deployment** phase, the real world data used for the ML task can bring rise to this issue.

3) *Image Removal Requests*: When collecting data and using existing datasets, ownership issues will arise and should be addressed early and appropriately. Instead of using “public” resources, researchers should seek participation consent from individuals. This becomes important when using data for research and in deployed systems. This raises the question, what to do if an individual’s visual content is requested to be removed from the dataset and the model’s training phase? In July 2020, MIT decided to remove the 80 Million Tiny Images dataset because of the bias and offensive labels that occurred in the dataset [52]. If researchers have used this dataset, these issues can affect the credibility of their work and the deployed system, if one exists. Image removal requests can affect all phases of the ML pipeline and should be handled accordingly.

B. Fairness

In this section, we discuss three typical fairness issues. These biases sneak into most steps of the ML pipeline and

could propagate to other parts of the pipeline. These three general issues can lead researchers to think about where or when bias can occur. Later, in the algorithmic bias section, we will discuss four more specific biases (i.e., *individual fairness* versus *group fairness*, *disparate treatment* versus *disparate impact*) that explore who is affected and how those issues arise in the pipeline.

1) *Historical Bias*: When data is generated, the inherent bias from the world could stealthily engrave into data. Historical bias can enter the ML pipeline at the start point of the **Data Preparation** phase and the **Deployment** phase. Even under ideal sampling and feature selection, historical bias could still exist and cause concern. When the historical bias proliferates through the ML pipeline, it can impact modeling and decision-making in the deployment stage [21], [53].

2) *Algorithmic Bias*: Algorithmic biases are bound together with each process in the ML pipeline. Roughly, algorithmic bias is focused in the **Modeling** phase. Because algorithms are connected with every part of ML systems, there are different bias sources and types from different components of the ML pipeline. The algorithm’s bias could be sourced from biased training data, a biased algorithm, or misinterpretation of the algorithm’s output [26]. Identifying the source of algorithmic bias contributes significantly to dissolving the fairness issue. In

addition, we must also consider the types of algorithmic bias. Usually, we can start to think about who is the victim impacted by the algorithmic bias. For example, similar individuals are treated inconsistently based on the predictions of the model, while *individual fairness* requires that each similar individual should be treated as similarly as possible [28]. As a more general example, *group fairness* considers groups defined by protected attributes (e.g., gender, race), and it requires that the protected groups should obtain similar treatments as the privileged group [29]. *Group fairness* is also referred to as statistical parity or demographic parity.

After identifying who suffers from the algorithmic bias, it becomes increasingly important to understand how fairness issues arise in the ML pipeline. *Disparate treatment*, also known as direct discrimination or intentional discrimination, occurs when protected attributes are used explicitly in the ML system. Consequently, disadvantaged groups identified by the protected attributes are deliberately treated differently. *Disparate impact* referred to indirect discrimination or unintentional discrimination, is pervasive and entrenched in our society [31]. Regarding *disparate impact* in the ML pipeline, it exists under the guise of correlated variables that implicitly correspond to the protected attributes.

3) *Software Discrimination*: Last but not least, software discrimination appears at the end of the entire ML pipeline, which is the **Deployment** phase, and bias could still exist due to the problematic model. After an ML model is passed to its end-users, the interpretability and transparency of the model can benefit from identifying and mitigating potential bias generated by the software. Researchers have developed many tools that audit fairness for a deployed ML model. Tools like IBM’s AI Fairness 360 toolkit [54] implement fairness metrics and bias mitigation algorithms. Other works have generated test suites to measure software fairness from a causality-based perspective [27].

C. Overlaps in Privacy and Fairness Issues




Fig. 2. In the ML pipeline, we indicate where privacy (green) and fairness (red) issues could arise. Possible overlaps in the system are defined in orange.

From Fig. 2, we can observe that some steps that contain both visual privacy and fairness issues. When both issues arise, researchers should be ready to deal with them; otherwise, they will affect the system’s outcome. For instance, a model builder perceives that the protected groups could be affected by the fairness issues in a facial recognition system. Consequently, the modeler strives to collect more data to make up for the disproportion. However, the privacy risk for the data collection

could be an unexpected problem and is increased for the underrepresented group [55].

IV. INTEGRATION OF INTERACTIVE AUDIT STRATEGIES FOR THE MACHINE LEARNING PIPELINE

ML models are constantly being updated once deployed to the real world; regular updates help to avoid and minimize costly errors. Differences in the time between error discovery and model correction for the deployed model are crucial. Systems should be able to respond to unexpected bias before, during, and after deployment.

It could be impossible to erase the damage caused by the aftermath of a system; however, stakeholders could start making a change now. One way to do this would be using an interactive ML approach, *human-in-the-loop* [56]–[59]. Training in the *human-in-the-loop* framework requires humans to make incremental updates to anticipate issues [60]. Traditional ML pipelines conduct training on their own without interference from humans. To debug these models, the researcher must begin a thorough investigation of the models’ predictions, parameters, and data after the learning has been completed. An interactive approach would allow a person in the model’s training process, which will reduce debugging and runtime. The human is able to check the learning for the model and coach the model to meet the desired result in a feedback cycle. Feedback cycles allow the researcher to provide positive feedback iteratively to the model after viewing the processes. This can allow the researcher to understand the possible bias and privacy issues in the model and mitigate it immediately. This approach can be extended to various ML research areas in fairness, computer vision, and privacy.

In traditional *human-in-the-loop* approaches, the human becomes a bottleneck for the feedback process. In light of this, we suggest using a **human-over-the-loop** approach [61]. *Human-over-the-loop* allows researchers to step into the pipeline as needed to perform corrections. This removes the necessity of a human approving each iteration of the model. With this feature integrated in the ML pipeline, the researchers should consider having multiple “humans” to monitor the training. This, in turn, can lower response times to resolve biases that may be imposed from “humans” during learning. Based on the *human-over-the-loop* technique, we propose the use of two interactive auditing strategies that can reduce fairness and privacy issues to allow researchers to conceptualize, develop, and deploy safer visual privacy systems.

A. Fairness Forensics Auditing System (FAST)

ML bias is a rising threat to justice, and it has been investigated in broad areas, including employee recruitment, criminal justice, and facial detection. ML research can cause unanticipated and harmful consequences on our daily life while decision-makers begin to utilize the result of the output from ML algorithms without considering fairness. Fairness forensics focuses on supporting data scientists and modelers to inspect a dataset or a new model by techniques and tools evaluating for bias.




Fig. 3. This figure illustrates the feedback loops when human-over-the-loop techniques are implemented. The green lines denote audit traces for feedback loops. The loops that are suggested to have a ViP Audit are denoted with a VP marker. The loops that are suggested to have FAST are denoted with a FF marker. Data Preparation has one feedback loop from Data to Data Cleaning. The Modeling phase has two feedback loops: (1) from model analysis to ML training, and (2) from Model Analysis to Data (in the Data Preparation phase). In the final phase of the pipeline, the deployment loops are from (1) Output to Data (in the Data Preparation phase) and (2) Output to ML training in the Modeling phase.

Fairness forensics requires overarching understandings of the types of bias, the entire pipeline of ML systems, and the analysis of bias on different pipeline stages. It is vital to understand how biases have harmful impacts on different communities of people when deploying ML systems related to visual privacy in societal domains. Fairness forensics has three major tasks: bias detection, bias interpretation, and bias mitigation. People can use fairness metrics to evaluate the input or output of ML models for bias detection. Bias report generator tools and bias visualization tools facilitate analysis and interpretation of bias for humans to understand the meaning and impact of bias detection results. Once the bias is discovered, bias mitigation strategies can be applied by the interventions to the input data, the algorithm, or the predictions. Bias mitigation algorithms can be categorized into three types: pre-processing, in-processing, and post-processing algorithms [54].

B. Visual Privacy (ViP) Auditor

For actively investigating visual privacy research, we propose the use of a human-over-the-loop technique specifically designed to handle privacy, classification, and computer vision issues. Visual privacy systems are comprised of multiple ML techniques and strategies. We envision the ViP Auditor as a comprehensive auditing tool that will enable researchers to use visual analytics [62] to understand the models' learning process. During learning, the modeler will be able to enhance the feedback process by using similar schemes as Model-Tracker [58] or Crayons Classifier [63]. The modeler will protect individual privacy in the learning process by incorporating visual privacy mitigation strategies built into the ViP Auditor. For Model Analysis, the researcher can understand the dataset attributes (e.g., number of faces, number of privacy leaks for each category), the models' classification performance, and the perceived privacy risk score of the model.

C. Understanding Pipeline Integration

From the **Data Preparation** phase in Fig. 3, researchers examine the dataset, data labels, and ownership for the content

regarding privacy concerns. This loop allows a researcher to consider the initial privacy concerns (in Section III) and develop other strategies to mitigate them. In the **Modeling** phase, the researcher should employ auditors at both feedback loops (see Fig. 3). The first feedback loop allows the researcher to conduct a privacy evaluation from the model's output. Evaluating the results from this feedback loop enables the human-over-the-loop to step in and make changes to achieve the desired level of privacy in the model. Auditing at this phase of the pipeline allows researchers to accurately correct recognition errors (bounding boxes, instance segmentation) from the models' learning. The second feedback loop conducts a privacy evaluation that allows the researcher to identify issues within the dataset from the Model Analysis. When the dataset issues are identified, the researcher can collect more data, remove the data from the pipeline, or add more tags/labels to mitigate privacy concerns that arise. The **Deployment** phase feedback loops consider the real-world output from the model. With auditors in place at this phase, the stakeholders can understand privacy issues as they arise. The stakeholders can fix issues in deployment as they arise by sanitizing the data and re-training the model. The ViP Auditor will produce a privacy risk score based on the models' performance and flag potential privacy issues.

The feedback loop from FAST is similar to the loop from the ViP Auditor. Fairness forensics system feedback occurs at different steps in all three phases of the ML pipeline (see Fig. 3), and it can encourage researchers to sanitize their data or adjust the model. The process of fairness forensics allows the human-over-the-loop to determine the need for human intervention and assess for fairness in order to achieve social justice. Imperfect fairness metrics or conflicting fairness objectives [64] means humans will need to intervene to maintain performance guarantees.

V. DISCUSSION

Being mindful of the societal impacts, evaluation methods (i.e., FAST and ViP) and monitoring strategies (i.e., human-over-the-loop) have been presented as mitigation techniques to reduce errors in the ML pipeline and in the deployed system's life cycle. However, there are no full-proof techniques for ensuring that the software is exempt from producing harm. For a stakeholder to know when to halt deployment implies that they have developed a plan for the system and require human intervention throughout the ML pipeline for proactive decision making. Monitoring for privacy and fairness issues and their potential to harm the community throughout the software's life is an essential part of this. When evaluating the fairness of a model, a researcher can explore the model's training data and performance metrics to decipher sub-trends and anomalies. From this evaluation, the researcher can generate an idea of what success can look like from their model.

It might also be helpful to pivot directions for the machine learning model to avoid going too far down a path that could prove disastrous for marginalized communities. There may be a point at which the model is beyond recognition. It may be

worth completely re-imagining the ML pipeline or abandoning the effort altogether when it has strayed far from its' intended goal. Before completely re-imaging or abandoning the model, the researcher could integrate human-over-the-loop techniques to improve the ML pipeline's consideration for privacy and fairness. The decision of which route to go ultimately involves the researcher evaluating the trade-off between the safety for the impacted communities or the potential accomplishments of producing innovative software. Success should be inspired by the ability to impact society positively, not by a system's ability to quickly solve an idea. Questioning when to stop in the ML pipeline should include prioritizing societal impact and the affect on marginalized communities.

Halting deployment on a project that has gone awry should be seen as a successful learning result, not as a failed project. If permissible, the stakeholder should consider opening up the research project or system for external review to cultivate a meaningful conversation around learning from the harms that development and deployment could have caused.

VI. CONCLUSION

Researchers should closely monitor data preparation, modeling, and deployment processes to avoid harming communities and stakeholders. The decision making process for researchers can be challenging, but it is imperative to continually evaluate to improve the model's learning process and the deployment outcomes for the communities they serve. When building a visual privacy model needing large amounts of data, it can be easy to obtain datasets that are widely distributed, but may not have been examined for discriminatory, private, or fairness issues. This work discusses privacy and fairness issues that frequently occur in the ML pipeline that could emerge at various phases. We also assert the need for responsible auditing systems to bring accountability into model training and the deployed system. To do this, we propose using human-over-the-loop strategies to introduce interactive auditing for fairness and privacy. With ML pipeline audits and engaged researchers, the evaluation and consideration given to project development and deployed systems can become a standard procedure. These proposed mitigation strategies are the first steps of a much needed effort to address privacy and fairness issues in the ML pipeline.

VII. FUTURE WORK

We believe that apart from suggesting an interactively auditable ML pipeline, future research should consider the trade-off between privacy, fairness, and model accuracy in these systems. In addition, we will develop a strategy for determining a failure rate for ML models to provide a mechanism for researchers to continually evaluate if their system is successful. These considerations warrant further investigation to determine the success and limitations of deployed systems.

VIII. ACKNOWLEDGMENTS

The researchers are partially supported by awards from the Department of Defense SMART Scholarship and the National Science Foundation under Grant No. #1952181.

REFERENCES

- [1] B. Krishnamurthy and C. E. Wills, "Characterizing privacy in online social networks," in *Proceedings of the First Workshop on Online Social Networks*. ACM, 2008, pp. 37–42.
- [2] R. Gross and A. Acquisti, "Information revelation and privacy in online social networks," in *Proceedings of the 2005 ACM Workshop on Privacy in the Electronic Society*. ACM, 2005, pp. 71–80.
- [3] D. Rosenblum, "What anyone can know: The privacy risks of social networking sites," *IEEE Security & Privacy*, vol. 5, no. 3, pp. 40–49, 2007.
- [4] M. Madejski, M. L. Johnson, and S. M. Bellovin, "The failure of online social network privacy settings," Columbia University, Tech. Rep. CUCS-010-11, 2011.
- [5] L. Van Zoonen, "Privacy concerns in smart cities," *Government Information Quarterly*, vol. 33, no. 3, pp. 472–480, 2016.
- [6] A. S. Elmaghraby and M. M. Losavio, "Cyber security challenges in smart cities: Safety, security and privacy," *Journal of advanced research*, vol. 5, no. 4, pp. 491–497, 2014.
- [7] J. DeHart, M. Stell, and C. Grant, "Social media and the scourge of visual privacy," *Information*, vol. 11, no. 2, p. 57, 2020.
- [8] J. DeHart, C. E. Baker, and C. Grant, "Considerations for designing private and inexpensive smart cities," in *ICWMC 2020-2020 IARIA The Sixteenth International Conference on Wireless and Mobile Communications (ICWMC)*. IARIA, 2020, pp. 30–33.
- [9] Y. Li, N. Vishwamitra, B. P. Knijnenburg, H. Hu, and K. Caine, "Effectiveness and users' experience of obfuscation as a privacy-enhancing technology for sharing photos," *Proceedings of the ACM on Human-Computer Interaction*, vol. 1, no. CSCW, p. 67, 2017.
- [10] M. Korayem, R. Templeman, D. Chen, D. Crandall, and A. Kapadia, "Enhancing lifelogging privacy by detecting screens," in *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, 2016, pp. 4309–4314.
- [11] R. Hoyle, R. Templeman, D. Anthony, D. Crandall, and A. Kapadia, "Sensitive lifelogs: A privacy analysis of photos from wearable cameras," in *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. ACM, 2015, pp. 1645–1648.
- [12] R. Sánchez-Corcuera, A. Nuñez-Marcos, J. Sesma-Solance, A. Bilbao-Jayo, R. Mulero, U. Zulaika, G. Azkune, and A. Almeida, "Smart cities survey: Technologies, application domains and challenges for the cities of the future," *International Journal of Distributed Sensor Networks*, vol. 15, no. 6, p. 1550147719853984, 2019.
- [13] J. M. Such, J. Porter, S. Preibusch, and A. Joinson, "Photo privacy conflicts in social media: A large-scale empirical study," in *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, 2017, pp. 3821–3832.
- [14] H. Zhong, A. Squicciarini, and D. Miller, "Toward automated multiparty privacy conflict detection," in *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*. ACM, 2018, pp. 1811–1814.
- [15] V. Arlazarov, K. Bulatov, T. Chernov, and V. Arlazarov, "Midv-500: a dataset for identity document analysis and recognition on mobile devices in video stream," *Computer Optics*, vol. 43, no. 5, p. 818–824, Oct 2019. [Online]. Available: <http://dx.doi.org/10.18287/2412-6179-2019-43-5-818-824>
- [16] A. K. Tonge and C. Caragea, "Image privacy prediction using deep features," in *Thirteenth AAAI Conference on Artificial Intelligence*, 2016.
- [17] A. Tonge and C. Caragea, "Image privacy prediction using deep neural networks," *ACM Transactions on the Web (TWEB)*, vol. 14, no. 2, pp. 1–32, 2020.
- [18] Y. Li, N. Vishwamitra, B. P. Knijnenburg, H. Hu, and K. Caine, "Blur vs. block: Investigating the effectiveness of privacy-enhancing obfuscation for images," in *2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*. IEEE, 2017, pp. 1343–1351.
- [19] S. Zerr, S. Siersdorfer, J. Hare, and E. Demidova, "Privacy-aware image classification and search," in *Proceedings of the 35th international ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM, 2012, pp. 35–44.
- [20] M. Tierney, I. Spiro, C. Bregler, and L. Subramanian, "Cryptagram: Photo privacy for online social media," in *Proceedings of the first ACM conference on Online social networks*. ACM, 2013, pp. 75–88.
- [21] H. Suresh and J. V. Guttag, "A framework for understanding unintended consequences of machine learning," *preprint arXiv:1901.10002*, 2019.

- [22] J. DeHart, C. Xu, L. Egede, and C. Grant, "Proposing an interactive audit pipeline for visual privacy research," *arXiv preprint arXiv:2111.03984*, 2021.
- [23] A. Neyaz, A. Kumar, S. Krishnan, J. Placker, and Q. Liu, "Security, privacy and steganographic analysis of faceapp and tiktok," *International Journal of Computer Science and Security (IJCSS)*, vol. 14, no. 2, p. 38, 2020.
- [24] N. Bantilan, "Themis-ml: A fairness-aware machine learning interface for end-to-end discrimination discovery and mitigation," *Journal of Technology in Human Services*, vol. 36, no. 1, pp. 15–30, 2018.
- [25] F. P. Calmon, D. Wei, B. Vinzamuri, K. N. Ramamurthy, and K. R. Varshney, "Optimized pre-processing for discrimination prevention," in *Proceedings of the 31st International Conference on Neural Information Processing Systems*, 2017, pp. 3995–4004.
- [26] D. Danks and A. J. London, "Algorithmic bias in autonomous systems," in *IJCAI*, vol. 17, 2017, pp. 4691–4697.
- [27] S. Galhotra, Y. Brun, and A. Meliou, "Fairness testing: testing software for discrimination," in *Proceedings of the 2017 11th Joint Meeting on Foundations of Software Engineering*, 2017, pp. 498–510.
- [28] C. Dwork, M. Hardt, T. Pitassi, O. Reingold, and R. Zemel, "Fairness through awareness," in *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*, 2012, pp. 214–226.
- [29] M. Hardt, E. Price, and N. Srebro, "Equality of opportunity in supervised learning," in *Proceedings of the 30th International Conference on Neural Information Processing Systems*, 2016, pp. 3323–3331.
- [30] M. B. Zafar, A. Valera, M. Gomez Rodriguez, and K. P. Gummadi, "Fairness beyond disparate treatment & disparate impact: Learning classification without disparate mistreatment," in *Proceedings of the 26th International Conference on World Wide Web*, 2017, pp. 1171–1180.
- [31] M. Feldman, S. A. Friedler, J. Moeller, C. Scheidegger, and S. Venkatasubramanian, "Certifying and removing disparate impact," in *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2015, pp. 259–268.
- [32] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, "Microsoft coco: Common objects in context," in *European Conference on Computer Vision*. Springer, 2014, pp. 740–755.
- [33] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "ImageNet: A Large-Scale Hierarchical Image Database," in *CVPR09*, 2009.
- [34] S. Zerr, S. Siersdorfer, and J. Hare, "Picalert!: A system for privacy-aware image classification and retrieval," in *Proceedings of the 21st ACM International Conference on Information and Knowledge Management*, ser. CIKM '12. New York, NY, USA: ACM, 2012, pp. 2710–2712.
- [35] M. Zimmer, "'but the data is already public': on the ethics of research in facebook," *Ethics and information technology*, vol. 12, no. 4, pp. 313–325, 2010.
- [36] M. Zimmer and K. Kinder-Kurlanda, *Internet research ethics for the social age: New challenges, cases, and contexts*. Peter Lang International Academic Publishers, 2017.
- [37] M. Mancosu and F. Vegetti, "What you can scrape and what is right to scrape: A proposal for a tool to collect public facebook data," *Social Media+ Society*, vol. 6, no. 3, p. 2056305120940703, 2020.
- [38] V. Krotov and L. Silva, "Legality and ethics of web scraping," *Twenty-fourth Americas Conference on Information Systems*, 2018.
- [39] J. Xiao, J. Hays, K. A. Ehinger, A. Oliva, and A. Torralba, "Sun database: Large-scale scene recognition from abbey to zoo," in *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2010, pp. 3485–3492.
- [40] A. Torralba, R. Fergus, and W. T. Freeman, "80 million tiny images: A large data set for nonparametric object and scene recognition," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 30, no. 11, pp. 1958–1970, 2008.
- [41] A. J. Perez, S. Zeadally, and S. Griffith, "Bystanders' privacy," *IT Professional*, vol. 19, no. 03, pp. 61–65, May 2017.
- [42] M. Dimiccoli, J. Marín, and E. Thomaz, "Mitigating bystander privacy concerns in egocentric activity recognition with deep learning and intentional image degradation," *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, vol. 1, pp. 1 – 18, 2017.
- [43] R. Hasan, D. J. Crandall, M. Fritz, and A. Kapadia, "Automatically detecting bystanders in photos to reduce privacy risks," *2020 IEEE Symposium on Security and Privacy (SP)*, pp. 318–335, 2020.
- [44] D. Darling, A. Li, and Q. Li, "Identification of subjects and bystanders in photos with feature-based machine learning," in *IEEE INFOCOM 2019 - IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, 2019, pp. 1–6.
- [45] A. Birhane, V. U. Prabhu, and E. Kahembwe, "Multimodal datasets: misogyny, pornography, and malignant stereotypes," *arXiv preprint arXiv:2110.01963*, 2021.
- [46] J. M. Such, J. Porter, S. Preibusch, and A. Joinson, "Photo privacy conflicts in social media: A large-scale empirical study," in *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, ser. CHI '17. New York, NY, USA: Association for Computing Machinery, 2017, p. 3821–3832. [Online]. Available: <https://doi.org/10.1145/3025453.3025668>
- [47] D. R. Zemmels and D. N. Khey, "Sharing of digital visual media: privacy concerns and trust among young people," *American Journal of Criminal Justice*, vol. 40, no. 2, pp. 285–302, 2015.
- [48] M. O. Lwin, A. J. Stanaland, and A. D. Miyazaki, "Protecting children's privacy online: How parental mediation strategies affect website safeguard effectiveness," *Journal of Retailing*, vol. 84, no. 2, pp. 205–217, 2008.
- [49] S. Batool, "Exploring vulnerability among children and young people who experience online sexual victimisation," Ph.D. dissertation, University of Central Lancashire, 2020.
- [50] F. Li, Z. Sun, A. Li, B. Niu, H. Li, and G. Cao, "Hideme: Privacy-preserving photo sharing on social networks," *IEEE INFOCOM 2019 - IEEE Conference on Computer Communications*, pp. 154–162, 2019.
- [51] J. DeHart and C. Grant, "Visual content privacy leaks on social media networks," *arXiv preprint arXiv:1806.08471*, 2018.
- [52] A. Torralba, R. Fergus, and B. Freeman, "80 million tiny image dataset," Jun 2020. [Online]. Available: <https://groups.csail.mit.edu/vision/TinyImages/>
- [53] T. Hellström, V. Dignum, and S. Bensch, "Bias in machine learning-what is it good for?" in *International Workshop on New Foundations for Human-Centered AI (NeHuAI) co-located with 24th European Conference on Artificial Intelligence (ECAI 2020), Virtual (Santiago de Compostela, Spain), September 4, 2020*. RWTH Aachen University, 2020, pp. 3–10.
- [54] R. K. Bellamy, K. Dey, M. Hind, S. C. Hoffman, S. Houde, K. Kannan, P. Lohia, J. Martino, S. Mehta, A. Mojsilović *et al.*, "Ai fairness 360: An extensible toolkit for detecting and mitigating algorithmic bias," *IBM Journal of Research and Development*, vol. 63, no. 4/5, pp. 4–1, 2019.
- [55] I. D. Raji, T. Gebru, M. Mitchell, J. Buolamwini, J. Lee, and E. Denton, "Saving face: Investigating the ethical concerns of facial recognition auditing," in *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, 2020, pp. 145–151.
- [56] J. A. Fails and D. R. Olsen Jr, "Interactive machine learning," in *Proceedings of the 8th International Conference on Intelligent User Interfaces*, 2003, pp. 39–45.
- [57] S. Amershi, M. Cakmak, W. B. Knox, and T. Kulesza, "Power to the people: The role of humans in interactive machine learning," *Ai Magazine*, vol. 35, no. 4, pp. 105–120, 2014.
- [58] S. Amershi, M. Chickering, S. M. Drucker, B. Lee, P. Simard, and J. Suh, "Modeltracker: Redesigning performance analysis tools for machine learning," in *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, 2015, pp. 337–346.
- [59] D. J. L. Lee, S. Macke, D. Xin, A. Lee, S. Huang, and A. G. Parameswaran, "A human-in-the-loop perspective on automl: Milestones and the road ahead," *IEEE Data Eng. Bull.*, vol. 42, no. 2, pp. 59–70, 2019.
- [60] J. Bond, C. Grant, J. Imbriani, and E. Holbrook, "A framework for interactive t-sne clustering," *International Journal of Software & Informatics*, vol. 10, no. 3, 2016.
- [61] A. Graham, Y. Liang, L. Gruenwald, and C. Grant, "Formalizing interruptible algorithms for human over-the-loop analytics," in *2017 IEEE International Conference on Big Data (Big Data)*. IEEE, 2017, pp. 4378–4383.
- [62] S. Liu, X. Wang, M. Liu, and J. Zhu, "Towards better analysis of machine learning models: A visual analytics perspective," *Visual Informatics*, vol. 1, no. 1, pp. 48–56, 2017.
- [63] J. Fails and D. Olsen, "A design tool for camera-based interaction," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2003, pp. 449–456.
- [64] S. A. Friedler, C. Scheidegger, and S. Venkatasubramanian, "The (im) possibility of fairness: Different value systems require different mechanisms for fair decision making," *Communications of the ACM*, vol. 64, no. 4, pp. 136–143, 2021.