
More effort is needed to protect pedestrian privacy in the era of AI

Xingchen Zhang*

Fusion Intelligence Laboratory
Department of Computer Science
University of Exeter
Exeter, UK, EX4 4RN
x.zhang12@exeter.ac.uk

Zixian Zhao

Fusion Intelligence Laboratory
Department of Computer Science
University of Exeter
Exeter, UK, EX4 4RN
zz541@exeter.ac.uk

Abstract

In the era of artificial intelligence (AI), pedestrian privacy is increasingly at risk. In research areas such as autonomous driving, computer vision, and surveillance, large datasets are often collected in public spaces, capturing pedestrians without consent or anonymization. These datasets are used to train systems that can identify, track, and analyze individuals, often without their knowledge. Although various technical methods and regional regulations have been proposed to address this issue, existing solutions are either insufficient to protect privacy or compromise data utility, thereby limiting their effectiveness for research. In this paper, we argue that **more effort is needed to protect pedestrian privacy in the era of AI while maintaining data utility**. We call on the AI and computer vision communities to take pedestrian privacy seriously and to rethink how pedestrian data are collected and anonymized. Collaboration with experts in law and ethics will also be essential for the responsible development of AI. Without stronger action, it will become increasingly difficult for individuals to protect their privacy, and public trust in AI may decline.

1 Introduction

In the era of AI, large amounts of data are collected and used to train deep learning models for various applications. Pedestrians in public or semi-public spaces such as streets, parks, or campuses are frequently recorded by cameras. In some research and commercial projects, additional sensors such as LiDARs, thermal cameras, or event cameras are also used to capture pedestrian data, often without the individuals' awareness. These recordings are then analyzed by neural networks and stored in datasets for AI training and deployment. For example, video data is widely used in pedestrian crossing prediction Zhang et al. [2022b], Chen et al. [2025b], trajectory forecasting Yao et al. [2024], Chen et al. [2025a], object tracking Wang et al. [2025], Wu et al. [2024], and various benchmarks Hind et al. [2024], Fu et al. [2025].

However, many public datasets do not sufficiently address pedestrian privacy. Faces and other personal details are often clearly visible, and some datasets do not mention ethics approval or describe any privacy protection measures. As shown in Fig. 1, pedestrians are not anonymized in many datasets used to train AI models for various tasks. These datasets are publicly available and have been widely used and cited in research, meaning that identifiable individuals can be easily viewed by anyone. Sharing personal information in this way raises serious privacy concerns.

In addition to existing datasets, real-world AI systems that rely on sensors such as cameras, LiDARs, thermal sensors, and event cameras continuously collect and analyze pedestrian data in public spaces.

*Corresponding author

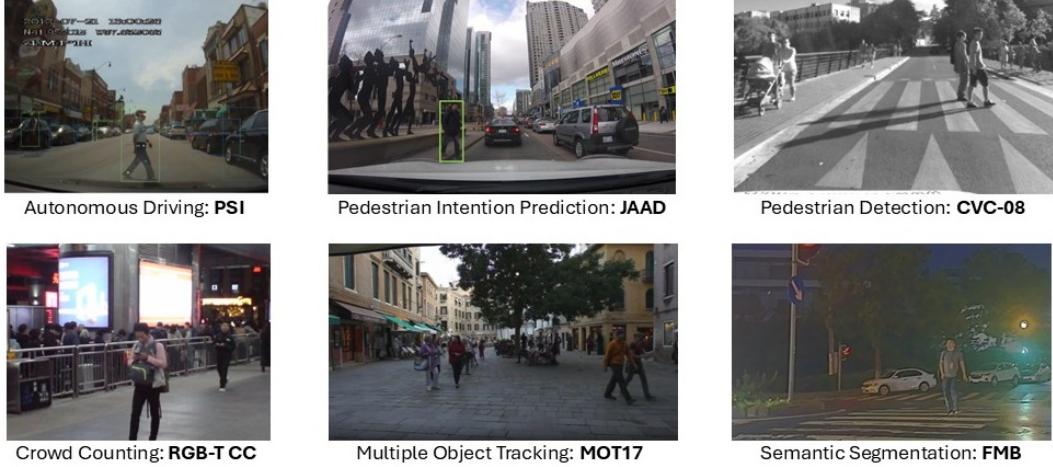


Figure 1: Pedestrian identities are often visible in datasets used for various computer vision tasks, such as autonomous driving (PSI dataset Chen et al. [2021]), pedestrian intention prediction (JAAD Rasouli et al. [2017]), pedestrian detection (CVC-08 Alzate et al. [2015]), crowd counting (RGB-T CC Liu et al. [2021]), multi-object tracking (MOT17 Milan et al. [2016]), and semantic segmentation (FMB Liu et al. [2023a]). These datasets are publicly accessible and have been used in a large number of papers, raising important privacy concerns for the individuals depicted. Images are taken from the respective datasets, and some may have been cropped for clarity.

Many pedestrians are unaware that they are being recorded or analyzed, and they are not given the opportunity to consent or opt out. This issue deserves more attention from the research community.

To protect privacy, certain rules must be followed in some countries or regions, such as the GDPR in Europe Voigt and Von dem Bussche [2017]. Many universities in Europe, including those in the UK, require a strict ethics approval process to ensure that the privacy of individuals, such as pedestrians in public spaces, is respected. However, since it is often infeasible to obtain explicit consent from every pedestrian in public settings, establishing strong default privacy protections remains a challenge. Moreover, **the belief that face anonymization alone is sufficient should be reconsidered**, given the rapid advancement of AI systems that can identify individuals from other features such as body shape, clothing, or gait. Additionally, although regulations exist in many regions, it is not uncommon for companies to violate them. For example, Amazon was fined \$25 million for violating children’s privacy in 2023 Wright [2023]. Therefore, robust pedestrian anonymization in videos is essential for protecting privacy in AI-driven applications.

Various pedestrian anonymization methods have been proposed, such as blurring, masking, face-swapping, and full-body anonymization. However, most of these methods significantly reduce the utility of the anonymized data, limiting their effectiveness in practical applications. In this paper, utility refers to whether the data remains useful for a specific task. In particular, we are concerned with whether the anonymized data still preserves its usefulness for the target task. It should be noted that utility is task-dependent. When discussing utility, the corresponding task should be specified, as anonymization may affect one task while leaving another unaffected.

In this paper, we argue that **more effort is needed to protect pedestrian privacy in the era of AI while maintaining data utility**. So far, the AI community has not paid sufficient attention to this issue. If no meaningful action is taken, it will become increasingly difficult for pedestrians to protect their privacy, leading to serious privacy risks. This, in turn, may reduce public trust in AI systems.

Position: Current methods for protecting pedestrian privacy are insufficient. More effort is needed to achieve effective and utility-preserving privacy protection in the era of AI.

2 Limitations of existing privacy protection methods

Various methods have been used to anonymize pedestrians in images and videos. Simple techniques such as cropping, Gaussian blurring, and masking appear in some studies, while more advanced approaches have also been proposed. Table 1 lists several representative methods. In this section, we discuss the general limitations of existing anonymization techniques.

Table 1: An overview of representative anonymization methods. ‘NA’ means this information is not available from the corresponding papers.

Reference	Year	Face	Full body	Real-time	Image/Video
Gafni et al. [2019]	2019	✓	-	✓	Video
DeepPrivacy Hukkelås et al. [2019]	2019	✓	-	NA	Image
Cho et al. [2020]	2020	✓	-	NA	Image
Zhu et al. [2020]	2020	✓	-	NA	Video
Li et al. [2021]	2021	✓	-	NA	Image
Wilson et al. [2022]	2022	✓	-	NA	Video
DeepPrivacy2 Hukkelås and Lindseth [2023]	2023	-	✓	-	Image
SG-GAN Hukkelås et al. [2023]	2023	-	✓	-	Image
LDFA Klemp et al. [2023]	2023	✓	-	NA	Image
Xue et al. [2023]	2023	✓	-	NA	Image
3PFS Zhao et al. [2024]	2024	✓	-	-	Video
Maximov et al. [2024]	2024	-	✓	NA	Video

2.1 Traditional anonymization methods

It is common to use masking or blurring to anonymize videos or images, as seen in some map or street-view services. Some public datasets also contain images where faces are blurred Grauman et al. [2022], as shown in Fig. 2 (a). However, this approach is not sufficient. Traditional anonymization methods can reduce the utility of videos for downstream tasks Lee and You [2024], i.e., make the anonymized dataset less useful for specific tasks. Moreover, face anonymization alone is inadequate, as will be discussed in detail in Section 2.3. Additionally, with the rapid progress in low-level vision tasks, traditional anonymization techniques are increasingly vulnerable to reversal using methods such as denoising Wu and Gao [2021], Ilesanmi and Ilesanmi [2021], Elad et al. [2023], deblurring Zhang et al. [2022a], Xiang et al. [2024], and inpainting Jam et al. [2021], Zhang et al. [2023], Xu et al. [2023].

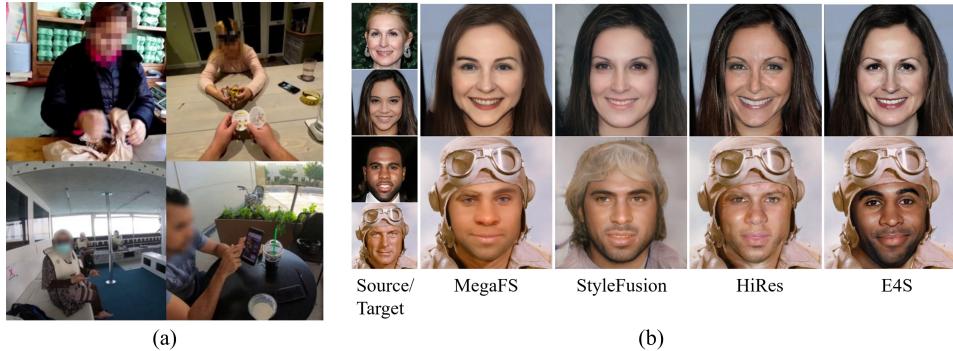


Figure 2: (a) Sample images from the Ego4D dataset Grauman et al. [2022], which uses pixelation or blurring to anonymize pedestrians and protect their privacy. However, these traditional methods often distort facial regions and remove important cues, such as head pose direction, that could be useful for model training. (b) Existing StyleGAN-based face-swapping methods Zhu et al. [2021], Kafri et al. [2021], Xu et al. [2022], Liu et al. [2023b] perform well when high-resolution target faces are available. However, in real-world settings, such as those involving cameras on robots or autonomous vehicles, pedestrian faces are often low-resolution, which limits the effectiveness of such techniques.

2.2 Generating fake faces or bodies

Some studies have used generative AI to generate fake faces or bodies for anonymization. For example, Hukkelås et al. Hukkelås and Lindseth [2023] generated full bodies using Generative

Adversarial Networks (GANs), showing promise for anonymization. However, as reported in Nature Nature [2023], generating fake faces or bodies may cause copyright issues. It is important to ensure that measures are taken to avoid such legal risks. Meanwhile, most existing anonymization efforts primarily emphasize model architectures and pipeline design, but offer limited theoretical analysis and formal privacy guarantees to rigorously quantify disclosure risk. Furthermore, most generative models are designed to operate on individual frames without temporal awareness, which leads to a lack of temporal consistency. As a result, the appearance of a generated face or body may fluctuate across consecutive frames, causing unnatural motion artifacts. This inconsistency reduces the visual coherence of pedestrian action, reducing the utility of the anonymized data in downstream tasks. Moreover, uncontrollable generative process may raise serious ethical concerns. For example, if a system unintentionally generates faces or bodies that resemble individuals from specific religions or ethnic groups, it may lead to unintended identity implications or cultural misrepresentation.

2.3 Face anonymization

Most studies have focused on face anonymization, with face swapping-based anonymization being one of the most commonly used techniques. For example, Mahajan et al. Mahajan et al. [2017] developed a face swapping application for privacy protection. Wilson et al. Wilson et al. [2022] proposed a method to protect patient privacy via face swapping, but it was designed for close-up medical images and is not suitable for pedestrians captured by vehicle-mounted cameras. In general, most existing methods are designed for high-resolution and near-distance face images, as shown in Fig. 2 (b). Additionally, they often have face detection before face anonymization, which means their anonymization effectiveness is highly dependent on the accuracy and robustness of the face detector.

In addition to the limitations discussed above, face anonymization methods often fail to generalize to full-body scenarios, which are common in real-world settings such as autonomous driving and surveillance. For example, DeepPrivacy Hukkelås et al. [2019] and LDFA Klemp et al. [2023] perform well on face anonymization but do not extend effectively to full-body cases. This raises a broader concern: pedestrians can still be identified through non-facial visual cues. In many real-world cases, a person can still be identified through their overall appearance, such as clothing style and body shape. For example, if another camera records the same person without anonymization, it is possible to match the clothing and figure across videos, even if the face is hidden in one of them. Full-body anonymization can better protect privacy by hiding more identifying features. We therefore **argue that full-body anonymization is necessary**.

2.4 Full-body anonymization

In recent years, several full-body anonymization methods have been proposed Hukkelås et al. [2023], Hukkelås and Lindseth [2023], Maximov et al. [2024]. These methods typically combine generative models with pedestrian segmentation techniques. However, they still face important technical limitations. First, for video data, these methods often fail to maintain temporal consistency, resulting in flickering or unstable appearance across frames. Second, they frequently do not preserve key pedestrian attributes such as age and gender, making the generated appearances somewhat random and less useful for downstream tasks. Moreover, even if full-body anonymization is achieved, gait patterns may still reveal individual identity. We therefore **argue that gait anonymization is also necessary** Hirose et al. [2022].

3 What is a good pedestrian privacy protection method?

A good privacy protection method should meet several key requirements. For example, Zhao et al. Zhao et al. [2024] suggest that such a method must effectively conceal pedestrian identity. In addition, the anonymized data should remain useful for downstream tasks. Based on these criteria and our review of related work, we summarize our view on the essential qualities of good pedestrian privacy protection in the box below.

Our opinion: What is a good pedestrian anonymization method?

- Can protect privacy for multiple pedestrians
- Do not significantly reduce the utility of anonymized data
- The anonymized video should be consistent and smooth temporally
- Resistance to model attack

As mentioned in Section 1, utility here refers to whether the data remains useful for a specific task. In practice, it is both important and challenging to evaluate the utility of anonymized data. The design of utility metrics is application-dependent because anonymized datasets may affect one task but not another. For example, in our previous work Zhao et al. [2024], we evaluated utility with several metrics, such as face detection and attribute reusability.

4 Alternative Views

In this section, we discuss some alternative views and then emphasize our perspectives.

4.1 Face anonymization is sufficient to protect pedestrian privacy

As mentioned earlier, many studies focus solely on face anonymization, with the underlying assumption that it is sufficient for protecting pedestrian privacy. In some cases, researchers have released datasets where only the face is masked or blurred.

Our perspective. We emphasize that **face anonymization alone is not sufficient**. This is because other visual features, such as body shape and gait, can also reveal identity, as demonstrated by pedestrian re-identification methods Ye et al. [2021] and gait recognition models Sepas-Moghaddam and Etemad [2022].

4.2 Anonymization reduces utility significantly

Some researchers argue that anonymization reduces utility. This is partly correct, but our opinion is that **utility-preserving anonymization is possible and necessary**. For instance, in our previous work Zhao et al. [2024], we show that the face detection performance based on anonymized images is only slightly worse than that of using original images. Specifically, in the experimental results, 99.5% faces remain detectable after their anonymization. This means that by using a proper anonymization method, it is possible to minimize the effect of anonymization on utility.

Our perspective. Utility-preserving anonymization is possible and necessary. We believe it should be treated as a key research direction in pedestrian privacy protection.

4.3 Using generative AI to generate fake images and videos does not have copyright issues

Generative AI has been widely used to create synthetic images and videos. Due to its ability to produce high-quality content, some researchers have adopted generative models for privacy protection, for example by generating fake faces or bodies. Typical examples include DeepPrivacy Hukkelås et al. [2019] and DeepPrivacy2 Hukkelås and Lindseth [2023]. However, since many pre-trained generative models are trained on datasets that may include unlicensed or copyrighted materials, the outputs they produce could also raise copyright concerns.

Our perspective. We argue that using synthetic images generated by generative models, especially pre-trained ones, **may cause copyright issues**. This risk should not be overlooked when applying generative AI for anonymization.

4.4 Federated Learning is sufficient for privacy protection

Federated learning Kairouz et al. [2021] is often considered a promising solution to privacy concerns, as it allows model training without directly sharing raw data Andrew et al. [2024]. However, it also faces several limitations Bak et al. [2024]. A key limitation is that it only begins after data collection. If privacy is not considered at the data collection stage, sensitive information, such as faces, body

shapes, or movement patterns, may already be captured and stored, even without being shared with a central server. This creates potential risks of leakage or misuse before any federated training takes place.

Our perspective. Federated learning alone is not sufficient to protect pedestrian privacy. To truly protect pedestrian privacy, we need to incorporate privacy-preserving strategies at every stage of the AI pipeline, including data collection, storage, and model training.

5 Challenges in pedestrian anonymization

5.1 Real-time performance

An advanced anonymization method is usually time-consuming. For example, the several methods mentioned earlier, such as DeepPrivacy2 Hukkelås and Lindseth [2023] and SG-GAN Hukkelås et al. [2023], cannot meet the real-time requirement. Moreover, when a method aims to preserve utility while protecting privacy, the method is usually more time-consuming because more modules are needed in the method. A typical example is the 3PFS proposed by Zhao et al. Zhao et al. [2024], where several seconds are needed to anonymize a single image.

5.2 Scaling to many pedestrians

Most methods are designed for handling one or several pedestrians. However, these methods may not be able to scale up to images with many pedestrians. When there are many pedestrians in one image, occlusion will become a significant issue. This will significantly affect the anonymization performance of existing deep learning-based pedestrian anonymization methods. Moreover, for those methods that require a detection or segmentation step prior to anonymization Zhao et al. [2024], crowded scenes also pose a significant challenge to the detectors and segmentation models.

5.3 Temporal consistency in anonymized videos

Another challenge in video anonymization lies in achieving strong temporal consistency. SG-GAN Hukkelås et al. [2023] is a representative approach method. However, the temporal consistency of its anonymized videos is limited, as shown in Fig. 3. In contrast, a temporally consistent anonymized video better aligns with the temporal coherence of human behavior, thereby preserving visual naturalness. This consistency also enhances the reliability of the anonymized data in downstream tasks such as pedestrian tracking and action recognition. Moreover, maintaining temporal consistency for both the face and body is essential. Facial consistency ensures stable facial attributes such as eye direction and head pose, while body consistency preserves coherent actions and overall appearance. Consistency in both the face and body is critical for downstream tasks such as pedestrian crossing intention prediction and age or gender recognition. However, existing anonymization methods often fail to maintain such consistency over time. Therefore, ensuring temporal consistency across both regions remains a key challenge for future anonymization research.



Figure 3: The inconsistency performance of SG-GAN Hukkelås et al. [2023], where the appearance of generated faces and bodies fluctuates across consecutive frames. This temporal inconsistency reduces the usefulness of the anonymized video for downstream computer vision tasks.

5.4 Avoid copyright and ethical issues of generative AI

The generated faces or bodies from generative AI have the risk of exposing identities present in the training data Nature [2023]. Since the training data of generative models may contain unlicensed materials, there are potential copyright issues associated with the generated identities

in generative AI-based anonymization methods. Moreover, since most generative models used in pedestrian anonymization are uncontrollable, the generated faces or bodies used for anonymization may unintentionally resemble individuals of specific religions or ethnicities, leading to unintended identity implications or cultural misrepresentation. Therefore, how to avoid copyright and ethical issues of generative AI is also a challenge that must be addressed in future pedestrian anonymization.

5.5 Consider individual difference of pedestrians

To maintain the utility of anonymized data, we usually need to replace original pedestrians with alternative representations, which may be artificially generated by AI Hukkelås et al. [2019] or obtained by swapping with real individuals Zhao et al. [2024]. During this process, it is important to preserve meaningful individual differences among pedestrians, such as age and gender. These characteristics are not only important in some downstream vision tasks but also essential for other fields like pedestrian demographic analysis Zhang et al. [2020]. How to maintain these characteristics for different pedestrians existing in the original video is a challenging problem.

5.6 Reversibility under attack

One of the key challenges for the future development of anonymization methods lies in evaluating and mitigating the risk of reversibility under adversarial conditions. Specifically, anonymized data may be vulnerable to reconstruction attacks. For example, when noise-based anonymization is employed, advances in denoising algorithms can potentially recover the original data. Similarly, when more advanced anonymization (e.g., generating fake faces or fake bodies) is employed, we also need to consider the risk of de-anonymization or model inversion attacks, in which attackers may leverage generative priors to reconstruct the original identity or other sensitive information.

6 Future work and call to actions

Based on our discussions, we propose several future directions for pedestrian anonymization. We call on AI researchers and policymakers to take coordinated action and develop a comprehensive strategy for protecting pedestrian privacy at every stage, from data collection to model deployment.

6.1 Technical Development

We believe a key technical direction in pedestrian privacy protection is to develop utility-preserving anonymization methods that ensure downstream task performance (e.g., detection or navigation) while hiding pedestrian identity. We call on researchers to work on pedestrian privacy protection from the following aspects.

6.1.1 From face anonymization to full-body anonymization and gait anonymization

As introduced earlier, face anonymization is not sufficient to protect pedestrian privacy. Therefore, it is crucial to recognize that more research efforts should be devoted to full-body anonymization in the future. Additionally, as gaits can be used to identify humans Eddine and Dugelay [2022], more research efforts should also be devoted to gait anonymization. Overall, **our perspective is that pedestrian privacy protection methods will go from face anonymization to full-body anonymization, and then go to gait anonymization**, as illustrated in Fig. 4. A possible way to achieve gait anonymization is to add noise to gait data so that the resulting motions remain similar but no longer reveal the pedestrian's identity.

6.1.2 Temporal consistency in anonymized videos

Most existing pedestrian anonymization methods fail to achieve temporal consistency in anonymized videos. Improving temporal consistency will be an important research direction for pedestrian privacy protection, as it can enhance both visual realism and data utility. Therefore, we call on researchers to further explore this direction and push the boundaries of pedestrian privacy protection.

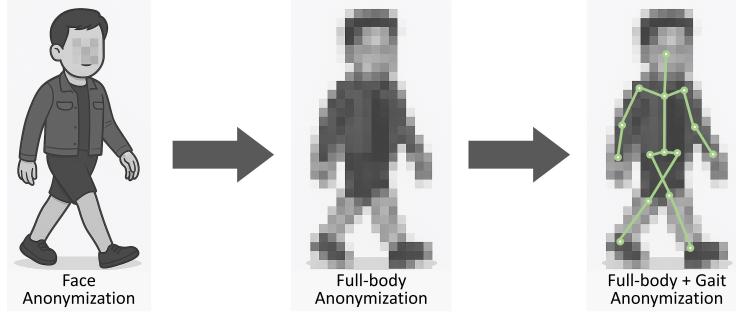


Figure 4: Pedestrian privacy protection will go from face anonymization to full-body anonymization, and then go to gait anonymization.

6.1.3 Protecting the privacy of multiple pedestrians

Another promising direction is protecting the privacy of multiple pedestrians, for example, scenes containing more than 20 individuals. To the best of our knowledge, no existing pedestrian anonymization method has demonstrated effective performance in such crowded scenarios. How to handle occlusions and overlapping pedestrians remains an open question that deserves further research attention.

6.1.4 Avoiding demographic bias

Existing deep learning-based pedestrian anonymization methods are often trained on datasets that may be biased toward certain demographic groups Paproki et al. [2024], Dehdashtian et al. [2024]. As a result, these methods may perform better on those groups and worse on others. This bias raises concerns about fairness in privacy protection. It remains an open challenge and calls for more attention and action from the research community.

6.1.5 Real-time performance

There has been very limited research on deploying pedestrian anonymization modules in real-time perception systems, particularly in recent works that rely on generative models. Existing methods often focus on offline processing, which limits their applicability in time-sensitive scenarios. Future research should explore the development of real-time privacy-preserving systems that can efficiently suppress identity-revealing information in RGB images.

Our perspective. If more research efforts are devoted to real-time pedestrian anonymization, it could become possible to collect high-quality anonymized pedestrian data directly, rather than collecting raw data first and anonymizing it offline. This would enhance pedestrian privacy protection by addressing potential risks from the data collection stage itself.

6.1.6 Model security

Measures should also be taken to ensure the security of anonymization models, so that they are not vulnerable to attacks such as adversarial examples. This is another underexplored direction in pedestrian privacy protection and deserves more attention from the research community.

6.1.7 Theoretical guarantees

Most existing pedestrian anonymization methods emphasize model and pipeline design but lack formal theoretical analysis of privacy protection. A systematic theoretical analysis would strengthen the interpretability and robustness of anonymization methods. For example, incorporating frameworks with formal guarantees, such as differential privacy, could enable anonymization methods to provide quantifiable and theoretically grounded privacy protection. Therefore, we call on researchers to complement model and pipeline design with formal theoretical guarantees so that methods are not only effective in practice but also provably safe, thereby advancing pedestrian privacy protection.

6.1.8 Protecting pedestrian privacy in non-RGB and multimodal data

Most existing pedestrian anonymization methods rely on RGB images or videos. There is still limited research on other modalities, such as thermal images and event camera data, as illustrated in Fig. 5(a) and Fig. 5(b), respectively. However, these data can still pose privacy risks, as body heat patterns or motion signatures may reveal a person’s identity through attributes such as body shape or gait. For example, some researchers argue that thermal data cannot be used to identify individuals. However, recent thermal-to-RGB image translation methods Luo et al. [2022], Wadsworth et al. [2024] challenge this assumption, showing that identity information may still be recoverable.

Beyond single modalities, multimodal systems combining RGB with thermal or event data further complicate privacy protection. While multimodal fusion improves perception robustness, it also increases the diversity of captured personal information. For example, fusing RGB and thermal data may reveal both visual appearance and body heat patterns. These multimodal privacy risks deserve greater attention as multimodal sensing becomes increasingly common in robotics and intelligent systems.



Figure 5: Pedestrians captured by thermal and event cameras. These data can still pose privacy risks.
(a) A thermal image containing pedestrians (image taken from the FMB dataset Liu et al. [2023a]).
(b) Pedestrians captured by a Prophesee event camera (image taken from Moore [2020]).

6.2 Evaluation and benchmarks

To the best of our knowledge, there is no well-recognized benchmark in the field of pedestrian anonymization. Zhao et al. Zhao et al. [2024] attempted to create an anonymization benchmark based on public datasets. However, that work focuses only on face anonymization. For full-body and gait anonymization, benchmarks are still missing, leading to inconsistent performance evaluation across studies. Therefore, it is important to develop standardized benchmarks and propose suitable evaluation metrics. Moreover, when developing such benchmarks, it is necessary to establish shared datasets and protocols for evaluating the privacy–utility trade-off in downstream tasks and the robustness of anonymization against potential privacy attacks.

Our perspective. In addition to accuracy-based metrics, we encourage competitions in computer vision and multimodal learning to introduce privacy scores. This would promote models that balance task performance and privacy preservation, especially in scenarios involving human data.

6.3 Policy and Governance

6.3.1 Regulatory frameworks and governance

Because current regulations often lag behind technological developments and remain too general, they cannot keep up with the technical details of emerging AI systems, especially those involving multimodal data. For example, while the GDPR provides guidance on traditional RGB images, it offers no clear rules for non-standard modalities such as thermal images or event camera data, which are increasingly used in modern perception systems and robotics. Additionally, existing legal definitions of personal data, which focus mainly on faces, may no longer be sufficient in the era of AI.

We argue that interdisciplinary collaboration is urgently needed. Researchers should work with policymakers to redefine and continuously evolve privacy regulations in public spaces, keeping pace

with advances in AI. These regulations should take into account new forms of identifiable information, such as body shape and gait, and new modalities, such as thermal images and event camera data. Moreover, privacy must be considered from the data collection stage, not just during model training or deployment. For example, thermal images or event camera data may require specific justification for collection.

To support these efforts, we encourage the development of standardized protocols to assess anonymization methods and quantify privacy risks. A clear consent framework should also be established through visible notices or designated collection zones with proper oversight. These measures can improve transparency and public trust while reducing privacy risks and ensuring that privacy protection keeps pace with advances in AI.

6.3.2 Global ethics and policy alignment

Mitigating the inconsistencies of ethics policies across different regions and countries is a challenging yet important task. Currently, there is no universal standard for ethics policies of pedestrian data. One practical approach is to follow the ethics policies of the region where the dataset is collected. For example, datasets collected in Europe should follow GDPR, while those collected in China should follow local regulations. In the long term, the research community and policymakers need to work together to establish shared ethical standards to ensure global consistency in privacy protection.

7 Societal Implications of Failing to Act

There could be negative effects on society in the long run if we do not take pedestrian privacy seriously. People may feel uncomfortable walking in public if they know they are being recorded frequently or used to train AI models. This can lead to a loss of trust in AI systems. Some groups of people may also be more affected than others, especially those who are already underrepresented or more vulnerable. This could make privacy protection unfair. We believe that ignoring these issues may be harmful for both individuals and the future development of AI. Therefore, we emphasize that pedestrian privacy should require more effort and be treated as a core part of responsible AI.

8 Conclusions

In this paper, we discussed pedestrian privacy protection in the age of AI. We showed that many datasets include pedestrians without proper anonymization, and that current anonymization methods are often insufficient to fully protect pedestrian privacy. We also discussed the challenges in pedestrian anonymization and pointed out that existing policies may not fully cover new modalities of pedestrian data, such as data collected using non-RGB sensors. Moreover, as AI systems become more capable and pervasive, the potential risks to pedestrian privacy are increasing. The AI community must take responsibility for protecting pedestrian privacy, as privacy is important not only for human dignity but also for building public trust in AI. Based on our discussions, **we emphasize that pedestrian privacy needs to be protected through greater efforts in both technology and policy.**

Although we mainly talk about pedestrian privacy in this paper, the privacy concerns can extend to other domains such as patient privacy in healthcare. As Prof. Fei-Fei Li has emphasized, we should aim for human-centered AI [2023]. We believe privacy protection is a crucial step in human-centered AI.

We call on NeurIPS and the broader AI community to treat privacy as a core part of AI research. To support this vision, NeurIPS and other major AI conferences can play important roles. They can encourage the development of benchmarks to evaluate anonymization methods and promote privacy-aware datasets and model design. In addition, they can host workshops that bring together researchers from AI and policy to address privacy challenges. These actions will help make privacy protection a fundamental principle of future AI research.

Acknowledgment

This study has received funding from the Royal Society Research Grant (No. RG\R1\251462).

References

Alejandro Gonzalez Alzate, Sebastian Ramos, David Vazquez, Antonio Lopez, and Jaume Amores. Spatiotemporal stacked sequential learning for pedestrian detection. In *Proceedings of the 7th Iberian Conference on Pattern Recognition and Image Analysis*, pages 3–12, 2015.

Galen Andrew, Peter Kairouz, Sewoong Oh, Alina Oprea, H Brendan McMahan, and Vinith M Suriyakumar. One-shot empirical privacy estimation for federated learning. In *International Conference on Learning Representations*, 2024.

Marieke Bak, Vince I Madai, Leo Anthony Celi, Georgios A Kaassis, Ronald Cornet, Menno Maris, Daniel Rueckert, Alena Buyx, and Stuart McLennan. Federated learning is not a cure-all for data ethics. *Nature Machine Intelligence*, 6(4):370–372, 2024.

Tina Chen, Taotao Jing, Renran Tian, Yaobin Chen, Joshua Domeyer, Heishiro Toyoda, Rini Sherony, and Zhengming Ding. PSI: A pedestrian behavior dataset for socially intelligent autonomous car. *arXiv preprint arXiv:2112.02604*, 2021.

Wangxing Chen, Haifeng Sang, Jinyu Wang, and Zishan Zhao. DSTIGCN: Deformable Spatial-Temporal Interaction Graph Convolution Network for Pedestrian Trajectory Prediction. *IEEE Transactions on Intelligent Transportation Systems*, 26(5):6923 – 6935, 2025a.

Xiaobo Chen, Wei Xu, Shilin Zhang, and Yingfeng Cai. Pedestrian crossing intention prediction via progressive multimodal token fusion for autonomous driving. *IEEE Transactions on Intelligent Transportation Systems*, 26(9):12959 – 12973, 2025b.

Durkhyun Cho, Jin Han Lee, and Il Hong Suh. CLEANIR: Controllable attribute-preserving natural identity remover. *Applied Sciences*, 10(3):1120, 2020.

Sepehr Dehdashtian, Ruozhen He, Yi Li, Guha Balakrishnan, Nuno Vasconcelos, Vicente Ordóñez, and Vishnu Naresh Boddeti. Fairness and bias mitigation in computer vision: A survey. *arXiv preprint arXiv:2408.02464*, 2024.

Mohamed Jamel Eddine and Jean-Luc Dugelay. Gait3: An event-based, visible and thermal database for gait recognition. In *2022 International Conference of the Biometrics Special Interest Group*, pages 1–5. IEEE, 2022.

Michael Elad, Bahjat Kawar, and Gregory Vaksman. Image denoising: The deep learning revolution and beyond—a survey paper. *SIAM Journal on Imaging Sciences*, 16(3):1594–1654, 2023.

Teng Fu, Yuwen Chen, Zhuofan Chen, Mengyang Zhao, Bin Li, and Xiangyang Xue. Crowdtrack: A benchmark for difficult multiple pedestrian tracking in real scenarios. *arXiv preprint arXiv:2507.02479*, 2025.

Oran Gafni, Lior Wolf, and Yaniv Taigman. Live face de-identification in video. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9378–9387, 2019.

Kristen Grauman, Andrew Westbury, Eugene Byrne, Zachary Chavis, Antonino Furnari, Rohit Girdhar, Jackson Hamburger, Hao Jiang, Miao Liu, Xingyu Liu, et al. Ego4d: Around the world in 3,000 hours of egocentric video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18995–19012, 2022.

Sam Hind, Fernando N van der Vlist, and Max Kanderske. Challenges as catalysts: how waymo’s open dataset challenges shape ai development. *AI & SOCIETY*, pages 1–17, 2024.

Yuki Hirose, Kazuaki Nakamura, Naoko Nitta, and Noboru Babaguchi. Anonymization of human gait in video based on silhouette deformation and texture transfer. *IEEE Transactions on Information Forensics and Security*, 17:3375–3390, 2022.

Håkon Hukkelås and Frank Lindseth. Deepprivacy2: Towards realistic full-body anonymization. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1329–1338, 2023.

Håkon Hukkelås, Rudolf Mester, and Frank Lindseth. Deepprivacy: A generative adversarial network for face anonymization. In *International symposium on visual computing*, pages 565–578. Springer, 2019.

Håkon Hukkelås, Morten Smebye, Rudolf Mester, and Frank Lindseth. Realistic Full-Body Anonymization with Surface-Guided GANs. In *2023 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 1430–1440, 2023.

Ademola E Ilesanmi and Taiwo O Ilesanmi. Methods for image denoising using convolutional neural network: a review. *Complex & Intelligent Systems*, 7(5):2179–2198, 2021.

Jireh Jam, Connah Kendrick, Kevin Walker, Vincent Drouard, Jison Gee-Sern Hsu, and Moi Hoon Yap. A comprehensive review of past and present image inpainting methods. *Computer vision and image understanding*, 203:103147, 2021.

Omer Kafri, Or Patashnik, Yuval Alaluf, and Daniel Cohen-Or. Stylefusion: A generative model for disentangling spatial segments. *arXiv preprint arXiv:2107.07437*, 2021.

Peter Kairouz, H Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Kallista Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, et al. Advances and open problems in federated learning. *Foundations and Trends® in Machine Learning*, 14(1–2):1–210, 2021.

Marvin Klemp, Kevin Rösch, Royden Wagner, Jannik Quehl, and Martin Lauer. Ldfa: Latent diffusion face anonymization for self-driving applications. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3198–3204, 2023.

Jun Ha Lee and Su Jeong You. Balancing privacy and accuracy: Exploring the impact of data anonymization on deep learning models in computer vision. *IEEE Access*, 2024.

Fei-Fei Li. *The worlds I see: Curiosity, exploration, and discovery at the dawn of AI*. Flatiron books: a moment of lift book, 2023.

Yongxiang Li, Qianwen Lu, Qingchuan Tao, Xingbo Zhao, and Yanmei Yu. Sf-gan: face de-identification method without losing facial attribute information. *IEEE Signal Processing Letters*, 28:1345–1349, 2021.

Jinyuan Liu, Zhu Liu, Guanyao Wu, Long Ma, Risheng Liu, Wei Zhong, Zhongxuan Luo, and Xin Fan. Multi-interactive feature learning and a full-time multi-modality benchmark for image fusion and segmentation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 8115–8124, 2023a.

Lingbo Liu, Jiaqi Chen, Hefeng Wu, Guanbin Li, Chenglong Li, and Liang Lin. Cross-modal collaborative representation learning and a large-scale rgbt benchmark for crowd counting. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4823–4833, 2021.

Zhian Liu, Maomao Li, Yong Zhang, Cairong Wang, Qi Zhang, Jue Wang, and Yongwei Nie. Fine-grained face swapping via regional gan inversion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8578–8587, 2023b.

Yi Luo, Dechang Pi, Yue Pan, Lingqiang Xie, Wen Yu, and Yufei Liu. ClawGAN: Claw connection-based generative adversarial networks for facial image translation in thermal to RGB visible light. *Expert Systems with Applications*, 191:116269, 2022.

Sachit Mahajan, Ling-Jyh Chen, and Tzu-Chieh Tsai. Swapitup: A face swap application for privacy protection. In *Proceedings of the IEEE 31st International Conference on Advanced Information Networking and Applications*, pages 46–50. IEEE, 2017.

Maxim Maximov, Tim Meinhardt, Zoe Papakipos, Caner Hazirbas, Cristian Canton, and Laura Leal-Taixé. Data-driven but privacy-conscious: Pedestrian dataset de-identification via full-body person synthesis. In *Proceedings of the IEEE 18th International Conference on Automatic Face and Gesture Recognition*, pages 1–10. IEEE, 2024.

Anton Milan, Laura Leal-Taixé, Ian Reid, Stefan Roth, and Konrad Schindler. MOT16: A benchmark for multi-object tracking. *arXiv preprint arXiv:1603.00831*, 2016.

Samuel K. Moore. Prophesee’s event-based camera reaches high resolution, 2 2020. URL <https://spectrum.ieee.org/prophesee-eventbased-camera-reaches-high-resolution>.

Nature. Why nature will not allow the use of generative ai in images and videos. *Nature*, 2023. Editorial, no listed author.

Anthony Paproki, Olivier Salvado, and Clinton Fookes. Synthetic data for deep learning in computer vision & medical imaging: A means to reduce data bias. *ACM Computing Surveys*, 56(11):1–37, 2024.

Amir Rasouli, Iuliia Kotseruba, and John K Tsotsos. Are they going to cross? a benchmark dataset and baseline for pedestrian crosswalk behavior. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 206–213, 2017.

Alireza Sepas-Moghaddam and Ali Etemad. Deep gait recognition: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 45(1):264–284, 2022.

Paul Voigt and Axel Von dem Bussche. The EU General Data Protection Regulation (GDPR). *A practical guide*, 1st ed., Cham: Springer International Publishing, 10(3152676):10–5555, 2017.

Emma Wadsworth, Advait Mahajan, Raksha Prasad, and Rajesh Menon. Deep learning for thermal-rgb image-to-image translation. *Infrared Physics & Technology*, 141:105442, 2024.

Junke Wang, Zuxuan Wu, Dongdong Chen, Chong Luo, Xiyang Dai, Lu Yuan, and Yu-Gang Jiang. OmniTracker: Unifying Visual Object Tracking by Tracking-with-Detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2025.

Ethan Wilson, Frederick Shic, Jenny Skytta, and Eakta Jain. Practical digital disguises: Leveraging face swaps to protect patient privacy. *arXiv preprint arXiv:2204.03559*, 2022.

George Wright. Amazon to pay \$25m over child privacy violations. <https://www.bbc.com/news/technology-65772154>, May 2023. Accessed: 2025-10-23.

Changhe Wu and Tianhan Gao. Image denoise methods based on deep learning. In *Journal of Physics: Conference Series*, volume 1883, page 012112. IOP Publishing, 2021.

Zongwei Wu, Jilai Zheng, Xiangxuan Ren, Florin-Alexandru Vasluiianu, Chao Ma, Danda Pani Paudel, Luc Van Gool, and Radu Timofte. Single-model and any-modality for video object tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19156–19166, 2024.

Yawen Xiang, Heng Zhou, Chengyang Li, Fangwei Sun, Zhongbo Li, and Yongqiang Xie. Deep learning in motion deblurring: current status, benchmarks and future prospects. *The Visual Computer*, pages 1–27, 2024.

Yangyang Xu, Bailin Deng, Junle Wang, Yanqing Jing, Jia Pan, and Shengfeng He. High-resolution face swapping via latent semantics disentanglement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7642–7651, 2022.

Zishan Xu, Xiaofeng Zhang, Wei Chen, Minda Yao, Jueting Liu, Tingting Xu, and Zehua Wang. A review of image inpainting methods based on deep learning. *Applied Sciences*, 13(20):11189, 2023.

Hanyu Xue, Bo Liu, Xin Yuan, Ming Ding, and Tianqing Zhu. Face image de-identification by feature space adversarial perturbation. *Concurrency and Computation: Practice and Experience*, 35(5):e7554, 2023.

Pengfei Yao, Yinglong Zhu, Huikun Bi, Tianlu Mao, and Zhaoqi Wang. Trajclip: Pedestrian trajectory prediction method using contrastive learning and idempotent networks. *Advances in Neural Information Processing Systems*, 37:77023–77037, 2024.

Mang Ye, Jianbing Shen, Gaojie Lin, Tao Xiang, Ling Shao, and Steven CH Hoi. Deep learning for person re-identification: A survey and outlook. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(6):2872–2893, 2021.

Kaihao Zhang, Wenqi Ren, Wenhan Luo, Wei-Sheng Lai, Björn Stenger, Ming-Hsuan Yang, and Hongdong Li. Deep image deblurring: A survey. *International Journal of Computer Vision*, 130(9):2103–2130, 2022a.

Xiaobo Zhang, Donghai Zhai, Tianrui Li, Yuxin Zhou, and Yang Lin. Image inpainting based on deep learning: A review. *Information Fusion*, 90:74–94, 2023.

Xingchen Zhang, Panagiotis Angeloudis, and Yiannis Demiris. St crossingpose: A spatial-temporal graph convolutional network for skeleton-based pedestrian crossing intention prediction. *IEEE Transactions on Intelligent Transportation Systems*, 23(11):20773–20782, 2022b.

Yunchang Zhang, Yu Qiao, and Jon D Fricker. Investigating pedestrian waiting time at semi-controlled crossing locations: Application of multi-state models for recurrent events analysis. *Accident Analysis & Prevention*, 137:105437, 2020.

Zixian Zhao, Xingchen Zhang, and Yiannis Demiris. 3pfs: Protecting pedestrian privacy through face swapping. *IEEE Transactions on Intelligent Transportation Systems*, 2024.

Bingquan Zhu, Hao Fang, Yanan Sui, and Luming Li. Deepfakes for medical video de-identification: Privacy protection and diagnostic information preservation. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, pages 414–420, 2020.

Yuhao Zhu, Qi Li, Jian Wang, Cheng-Zhong Xu, and Zhenan Sun. One shot face swapping on megapixels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4834–4844, 2021.