Protecting Cross-camera Person Re-identification Data with Image Differential Privacy

Lucas Maris
Nara Institute of Science and Technology, Japan
lucas.maris.lo3@is.naist.jp

Yuki Matsuda
Okayama University, Japan
RIKEN Center for Advanced Intelligence Project (AIP)
yukimat@okayama-u.ac.jp

Keiichi Yasumoto
Nara Institute of Science and Technology, Japan
RIKEN Center for Advanced Intelligence Project (AIP)
yasumoto@is.naist.jp

***Abstract*—To achieve smart cities, leveraging data from cameras, which are often readily-installed IoT devices, can offer precious insights on the behavior of pedestrians and play a crucial role in designing and maintaining efficient transportation, appropriate infrastructure, or attractive tourism facilities. Such pedestrian flow data collection is often achieved through cross-camera person re-identification. This task is heavily privacy-invasive task by design, benefiting from rich visual data, which thus carries highly sensitive personal details about individuals. We here study how image data can be protected upfront, and introduce a novel image differential privacy mechanism leveraging both pixelization and color quantization for this purpose. Our extensive experiments show that through its random noise additions, our mechanism can obfuscate data more effectively than standard image obfuscation methods while retaining high utility for cross-camera re-identification, preserving reasonable re-identification metrics and demographic information even under low privacy budgets.**

***Index Terms*—smart cities, image differential privacy, person re-identification, demographic predictions, anonymization**

I. Introduction

Cities keep growing, creating logistic challenges in terms of traffic management, city planning, and tourism. Access to comprehensive data regarding human flows within a city would help mitigate these issues, and such data can be collected through cameras, due to their prevalent presence within our lives, making them prime candidates for smart city applications. In this context, cross-camera person re-identification has been extensively studied over the past decades [1], as the computer vision task of matching individuals from different camera perspectives. Despite promising applications in security, planning, or tourism areas, little concern has been paid to the glaring privacy concerns this raises and, thus, the public acceptance of such systems.

To address these concerns, this study investigates the possibility of protecting stored image representations of individuals against sporadic data breaches. Assuming a CCTV-based re-identification system operating within a smart city, where each camera stores and transmits a large amount of sensitive visual pedestrian images, is it possible to store these images in such a way that they are of minimal value to attackers, yet still retain utility for their intended purpose, in case of a data leak at an arbitrary point in the system's operation?

To answer this question, we examine whether cross-camera person re-identification data can be protected through differential privacy [2], a formal data privacy model which characterizes the privacy loss that results from having one's data published. This privacy loss, or budget, can be affected and decreased through random data distortions, which data aggregators can define according to their privacy requirements.

This ability to quantify and control the privacy leakage of smart systems is key to their social acceptance, which requires building trustful relationships with citizens, customers, or individuals at large. Extending differential privacy guarantees to images is expected to move forward towards this goal.

Our contributions are as follows:

- We formulate a new strict image differential privacy mechanism leveraging both pixelization and color quantization, which significantly reduces data sensitivity and allows for lower privacy budgets.
- We extensively evaluate its effect on the public Market1501 dataset, in terms of re-identification performances, gender classification ability and visual image quality, and identify privacy-utility trade-off points.
- We highlight the robustness of centroid-based re-identification models against differentially private noise.
- We show our differential privacy mechanism is able to provide a quantifiable privacy guarantee by reducing general utility and visual quality of images, while preserving reasonable re-identification performances.

II. Related Work

This section briefly overviews existing studies on person re-identification and differential privacy.

A. Cross-camera person re-identification (reID)

The recognition of individuals across different visual snapshots from different points of view has become a traditional computer vision task over the last 20 years [1]. It is usually formalized as an image-retrieval task, where the aim is to train a model on a *training* set such that it can rank images from the *gallery* set in order of their similarity to a given image from the *query* set. Recent state-of-the-art reID systems achieve maximum performances on traditional datasets such as Market1501 [3] by combining the now conventional CNN feature extractors [4] and triplet loss [5] with a variety of other techniques [6], e.g., by aggregating information from multiple person images to increase model robustness to outliers [7].

Privacy-preserving re-identification has also been explored, mostly by substituting the classic RGB camera with less privacy-invasive sensors, such as continuous-wave radars [8], event cameras [9], or LiDARs [10]. Other approaches apply classic anonymization methods to reID data, such as face blurring [11]. Some studies have focused on anonymizing visual data, e.g., by replacing pedestrians with synthetic objects [12] or wireframe representations [13], but do not evaluate the usability of protected data for practical computer vision tasks. Privacy based on cryptographic data encryption [14] and federated learning [15] has also been considered in the context of person re-identification. To our best knowledge, the effect of applying formal differential privacy directly to image data collected from RGB cameras on reID performances has not yet been studied.

B. Differential privacy (DP)

Over the last decade, differential privacy [2] has become one of the most popular ways of modeling formal data privacy, due to its ability to provide quantifiable protection against arbitrary risks, with implementations by, e.g., Google, Apple, Microsoft, the U.S. Census Bureau and SAP. By perturbing computations over statistical databases, it promises indistinguishability between data records and plausible deniability to every individual composing such databases, providing them with roughly the same privacy that would result from having their data removed.

Recent work regarding differential privacy has aimed at extending its desirable properties to other forms of data. With unstructured data making up most of today's data landscape, a line of work focused on image differential privacy has also emerged, with studies leveraging pixelization [16], singular value decomposition [17], data streams [18], or generative models [19], [20]. As noted in a recent survey on differential privacy for unstructured data [21], the common approach is to vectorize unstructured data into a structured form, which can then be obfuscated with conventional DP methods.

These past studies have focused on analyzing the effect of differential privacy mechanisms on adverse attacks [16], [17], [20], on the visual quality of output images [18], [19], or on their similarity to the unprotected images [16], [17], [19], [20], but usually provide limited insights regarding the concrete usability of differentially private image data for practical computer vision tasks. We here study how such data can be used specifically for person re-identification under strict differential privacy conditions, without assuming the size of sensitive areas in images [16] nor the image sensitivity based on empirical measures [19], and without producing outputs similar to a given training set [19], [20].

III. Method

This section goes over the threat model we consider, the differential privacy mechanism we propose, the re-identification model we use, the way we evaluate gender predictions, and the baseline methods we compare to.

A. Threat model

Cross-camera person re-identification is a task where rich (and thus privacy-sensitive) data is collected in order to produce outputs that are actually much less privacy-sensitive. While the inputs to this task are a large number of image-representations of individuals, its output can be viewed as answering questions such as "is person x the same person as person y?". The answer to these questions is evidently still privacy-sensitive information; however, given that x's and y's image-representations are stored in such a way that they cannot be linked to the individuals' real-world visual appearance, this has the system store much less private information overall.

In a smart city system operators may wish to collect information such as the percentage of people that visit point B shortly after point A. However, for their purposes, there is no value in knowing whether this percentage includes a specific visually recognizable individual. Therefore, if the input data to the reID system can be distorted to be minimally useful for other applications, the privacy footprint of the system is reduced to what is strictly necessary to its operation. We believe this further useful for the public release of image datasets, which have been under scrutiny due to privacy concerns, and believe differential privacy a promising approach to strip datasets of sensitive information and restrict their fitness for undesirable, unforeseen uses.

B. ε-Image Differential Privacy

To provide a quantifiable privacy guarantee on images, we extend the grayscale pixel-level differential privacy definition first introduced as DP-Pix [16]. We further restrict its definition; instead of providing indistinguishability between *grayscale* same-sized images differing by at most m pixels, we here aim for indistinguishability between ***RGB*** same-sized images differing in *any* amount of pixels.

Definition 1. *ε-Image Differential Privacy*: a randomized mechanism $\mathcal{M}$ gives ε-image differential privacy if for any two images i and j of same dimension, and for any possible output $R \subseteq \text{Range}(\mathcal{M})$,

$$\Pr[\mathcal{M}(i) \in R] \leq \exp(\varepsilon)\ \Pr[\mathcal{M}(j) \in R] \tag{1}$$

As shown in [2], the Laplace mechanism can achieve such a guarantee, provided $\mathcal{M}$ is defined as the noisy function:

$$\mathcal{M}(x) = f(x) + n, \text{ where } n \sim Laplace\left(0, \frac{\Delta f}{\varepsilon}\right) \tag{2}$$

The exact amount of noise n is to be calibrated to the sensitivity Δf of function f, which we define as the identity function. Our broader image neighborhood definition, which does not assume the size of the area containing private information in images, leads to a much higher sensitivity than existing image differential privacy mechanisms; this then calls for more noise addition to achieve low privacy budgets ε.

To reduce the magnitude of this sensitivity value, we propose to use not only pixelization, but also color quantization as a means to reduce the dimensionality of images prior to

differentially private noise addition. As such, we generalize and extend the DP-Pix definition of function f; instead of the pixelization of grayscale image x, we define f as the identity function applied to RGB image x, with optional pixelization and color quantization parameters b and c. The sensitivity of this function then becomes:

$$\Delta f = \frac{wh}{4^b}\left(\frac{2^8}{2^c} - 1\right)^3 \tag{3}$$

where w and h are the width and height of images. The amount of pixelization to be applied to images is defined through b, such that 4^b pixels are reduced to a single pixel. The range of available color values in images is characterized through c, where each color channel is reduced from its original 8 bits to $(8 - c)$ bits. If $b = 0$ and $c = 0$, no pixelization nor quantization is applied, and the sensitivity is equivalent to that of the identity function.

By protecting image data upfront with a quantifiable privacy guarantee in the form of a privacy budget ε, its vulnerability can be mitigated. The re-identification system, or any other unforeseen use of the images after applying ε-Image DP, is a form of post-processing, which has no effect on the privacy guarantee offered by the mechanism [2].

To thoroughly evaluate how dimensionality reduction parameters b and c affect privacy budgets ε and reID performances, we here present results for 4 parameter combinations, as introduced in Table I. Among the combinations we experimented with, these parameters were chosen for their good performances and ability to illustrate the importance of color quantization in our mechanism. Privacy budgets ε were made to vary between $\{1, 2.5, 5\} \times 10^x$, where $x \in \{0, ..., 12\}$.

TABLE I: Dimensionality reduction parameter settings. Sensitivity Δf is calculated for images with $w = 64$ and $h = 128$.

	b	c	Δf
(A) High color quantization	0	6	221,184
(B) Mixed pix. and color quant.	1	5	702,464
(C) High pixelization	2	4	1,728,000
(D) No pix. nor color quant.	0	0	135,834,624,000

C. Re-identification system

We opt for a simple but effective reID model in the form of the now common Bag of Tricks (BOT) model [6], and increase the robustness of this model against differentially private noise by combining it with the centroid-based approach introduced by Wieczorek *et al* [7]. By averaging training samples into centroids, i.e., aggregated class representations, the task is shifted from ranking specific identity-instances to classifying into actual identities, which arguably also makes more sense in practical applications. An illustration is provided in Fig. 1.

We expect the use of these averaged individual representations to be able to magnify identity-specific latent features that remain underneath the noise. We test this hypothesis by training and testing both regular and centroid-based models directly onto noised images. Performances are reported in terms of Mean Average Precision (mAP), the mean of the average precision score as evaluated for each image within the query set, and Rank-1, the percentage of images from the query set for which the gallery match predicted with highest confidence is a true match. Experiments are repeated three times and we report the average metric values.

(a) Regular reID model. The query is wrongly assigned to the nearest image's identity.

(b) Centroid-based reID model. The query is correctly assigned to the nearest mean centroid representation's identity.

Fig. 1: Difference between regular and centroid reID models.

In terms of dataset, we use Market1501 [3], perhaps the most common and widely used public re-identification dataset, which contains 32,668 pedestrian images from 1501 different individuals collected across 6 cameras. Market1501 images are composed of 64x128 RGB pixels (using 24 bits per pixel). Train, gallery and query sets are kept consistent with the standard splits for this dataset.

D. Gender predictions

To evaluate the effect of our privacy mechanism beyond the reID task, we additionally consider its effect on gender classification, a common attribute classification task. This gives an idea of how feasible attribute extraction remains after image DP-obfuscation, and can assist system operators in choosing a suitable privacy budget ε. Market1501 is annotated in terms of binary gender labels (male, female) [22]. Gender classification is implemented as a simple fully connected layer on top of a pretrained ResNet50v2 [4] backbone.

Performances are reported in terms of F1-scores as evaluated on the test set, composed of both the gallery and the query sets. All experiments are repeated three times, and we here report the average metric values. To present a fairer comparison with the centroid-based reID model, which leverages multiple image representations of the same identity, we additionally report the classification F1-scores on a per-identity basis, by aggregating all the predictions obtained on the image representations of a given identity through a majority vote, for each of the identities in the dataset.

E. Baselines

We compare our method to existing image obfuscation methods illustrated in Fig. 2. Blurring is widely used as a simple privacy-preserving method, and applies a Gaussian kernel to modify each pixel based on neighboring pixels. We here present results for blurring with kernel size $k = 25$. Replacing faces with black boxes is also a common way to increase privacy. We here do this using a CE2P human parsing model [23] to segment images into 20 different areas, merge

those labeled as *face*, *hair*, *hat* and *sunglasses*, increase the resulting area's size to form a rectangular shape, and then zero all pixels in this area. We also compare our results to simple pixelization, with kernel size $b = 2$, and color quantization, reducing color richness by a factor of $c = 6$.

Fig. 2: Visual effect of the considered baseline methods.

IV. Results

In this section, we report reID and gender classification performances on ε-Image DP-protected data, compare our proposed method to existing baselines, and provide examples of ε-Image DP-protected images.

A. ReID

Figure 3 shows the performance of both a regular and a centroid-based reID model on Market1501 obfuscated through ε-Image DP. For the sake of brevity, we only include this comparison for the mixed pixelization & color quantization parameter setting (B), as all other parameter settings exhibit the same behavior. As one would expect, both mAP and Rank-1 metrics degrade as the privacy budget ε decreases, i.e., as privacy increases. It is however striking from our experiments that centroid-based reID models (dark lines) offer much higher robustness to noise than regular reID models (light lines). Their observed reID performance after applying ε-Image DP is sensibly higher and therefore can withstand lower privacy budgets ε before performances degrade significantly.

Having confirmed that averaged individual representations can help decrease privacy budgets for reID, we now discuss and compare the effects of dimensionality reduction parameters b and c and privacy budget ε. Figure 4 shows the performances of centroid-based reID models on Market1501 data noised with different sets of privacy parameters.

At a glance, it appears Market1501 suffers very little from dimensionality reduction parameters b and c, with all four

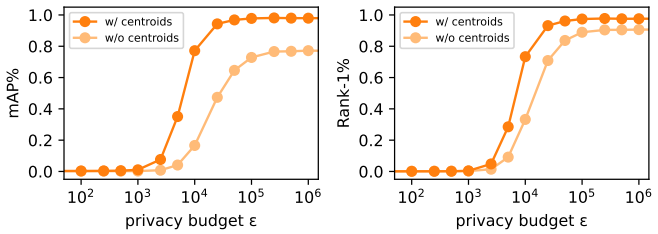

Fig. 3: ReID performances (mAP left, Rank-1 right) of centroid-based (dark line) and regular (light line) models on Market1501 after ε-Image DP, in parameter setting (B).

settings displaying similar behaviors when little noise is added (right-hand sides of the graphs). The order of parameter sets in which performances degrade to chance-level is directly related to dimensionality reduction parameters b and c, with c seemingly having the largest impact. Not using any dimensionality reduction (red line, D) has performances degrade at much higher privacy budgets ε than using either pixelization or color quantization. High pixelization (green line, C) can decrease privacy budgets while retaining good reID performances, but not as much as high color quantization (blue line, A), which can achieve privacy budgets as low as $\varepsilon = 2500$ despite retaining a mAP $> 90\%$.

This behavior is due to the lower sensitivity values Δf of parameter settings (A)-(C), as previously indicated in Table I. This lower sensitivity is achieved by reducing the amount of information in images prior to differentially private obfuscation. Due to the nature of image data, raw images contain a lot of redundancy, which pixelization and color quantization help eliminate. The more noticeable effect of the latter is likely inherent to the RGB color model, which defines many more possible color values than appears to be necessary for reID models to distinguish people.

It is also quite apparent that there exists a clear cutoff point where reID goes from feasible to nearly chance-level. Considering pixel color channels have a very limited range of values (256 for each channel), every pixel color channel is likely to be forced into its minimum or maximum value at random the instant the privacy budget becomes low enough, essentially erasing all information within images. We mark these cut-off points with $\star$ in Fig. 4, as the lowest privacy budgets ε where centroid-reID performances are kept reasonably high, and further discuss these in Section IV-C below.

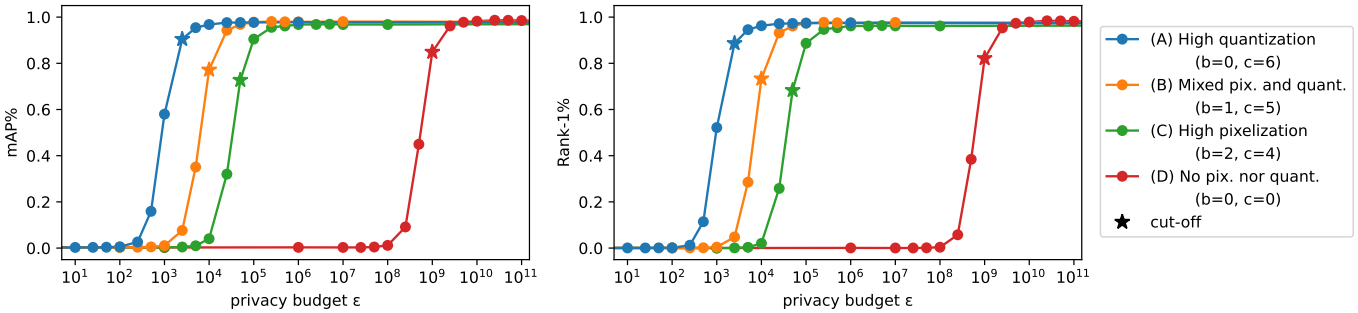

Fig. 4: ReID performances of centroid-based models on Market1501 after ε-Image DP.

B. Gender predictions

Using the same noised Market1501 images to perform gender classification gives the results shown in Fig. 5. As with reID, leveraging multiple image representations per identity (dark lines) yields higher performances than using single images (light lines). The x-values at which points are marked with $\star$ are kept consistent with Fig. 4. At a glance, gender performances appear to degrade around the same privacy budgets ε as reID performances, but the usability drops more gradually than for reID, with softer slopes around $\star$-marked points. It is also interesting to observe that the effect of dimensionality reduction alone, as shown in the right-hand parts of each graph, which have very little added noise, is nearly nonexistent, with all settings on both datasets displaying very similar performance. This can be attributed to gender classification being a simpler task overall, thus suffering less from pixelization and color quantization.

C. Baseline comparison

Table II summarizes the cut-off points identified and marked with $\star$ in Fig. 4 and Fig. 5 for our proposed method in each parameter setting A-D. This table compares the behavior of our models at these cutoff points on ε-Image DP-protected datasets with state-of-the-art performances on an unprotected dataset, and with normal reID and gender classification performances on datasets obfuscated through traditional methods. To further quantify the relationship between the original images and their obfuscated counterparts, this table also reports the mean Structural Similarity Index Measure (SSIM) [24], a common measure for assessing the perceived similarity between images. The larger this measure, computed between two images, the more said images are visually similar; a good obfuscation mechanism is expected to minimize this metric.

From Table II, it appears our method can generally achieve reID metrics closer to SOTA than traditional obfuscation methods, on top of providing a quantifiable privacy measure in the form of a privacy budget ε, which traditional methods can not. While replacing individual's faces with black boxes does achieve comparably high reID performance, it does not protect any other private body features; this is confirmed by the high associated SSIM value. It clearly appears our proposed method achieves lower SSIM values than all the other baselines, thus introducing more distortion, all the while preserving similar or better reID and gender classification performances.

D. Visual image quality

Figure 6 illustrates our proposed mechanism with an example image from the Market1501 dataset. Images at the cut-off points discussed above are lined in red. As can be observed, both the dimensionality reduction (pixelization & color quantization) as well as the actual noising mechanism affect the output image. A higher level of pixelization appears to suffer more rapidly from a decreasing privacy budget ε, thus requiring less Laplacian noise to obfuscate identity-specific attributes (e.g., face), at least on a visual level, while a higher level of color quantization seems to streamline color regions, making them more robust against low privacy budgets ε. This aligns with our results in Section IV.

While the images at the selected cutoff points bear some form of resemblance to the original, it is hard to make out much details of the person's appearance. As such, the selected comparison points appear like good trade-off points between low general-purpose visual quality and reasonable reID performances. Visually, the mixed setting (B), strikes us as perhaps the best trade-off, sacrificing performance to an extent that makes images hardly recognizable yet still reasonably useful for person re-identification. When compared to the images obfuscated using traditional methods, as in Fig. 2, our proposed method appears to achieve a more thorough obfuscation, erasing more high-level details of peoples' appearances.

TABLE II: Comparison of our proposed method (A-D) on Market1501 with SOTA and traditional obfuscation baselines. Best values are highlighted in bold.

		reID		Gender	
Parameters	**Privacy ε**	**mAP**	**Rank-1**	**F1**	**SSIM**
SOTA [7]	none	*98.3%*	*98.0%*		
Blurring (k=25)	none	71.5%	87.3%	82.8%	0.469
No face	none	82.1%	**92.4%**	85.1%	0.914
Pix. (b=2)	none	67.6%	85.2%	83.2%	0.661
Quant. (c=6)	none	71.3%	87.4%	84.4%	0.785
(A) b=0, c=6	2500	**90.5%**	88.6%	**85.6%**	0.220
(B) b=1, c=5	10000	77.2%	73.3%	79.2%	0.232
(C) b=2, c=4	50000	72.7%	68.3%	83.5%	0.330
(D) b=0, c=0	10^9	84.9%	82.1%	79.5%	**0.144**

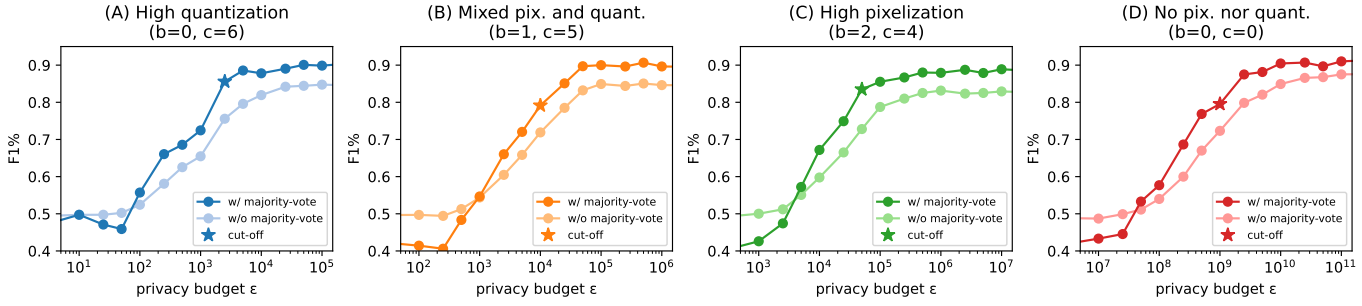

Fig. 5: Gender classification performances on Market1501 after ε-Image DP, per-identity (dark line) and per-image (light line).

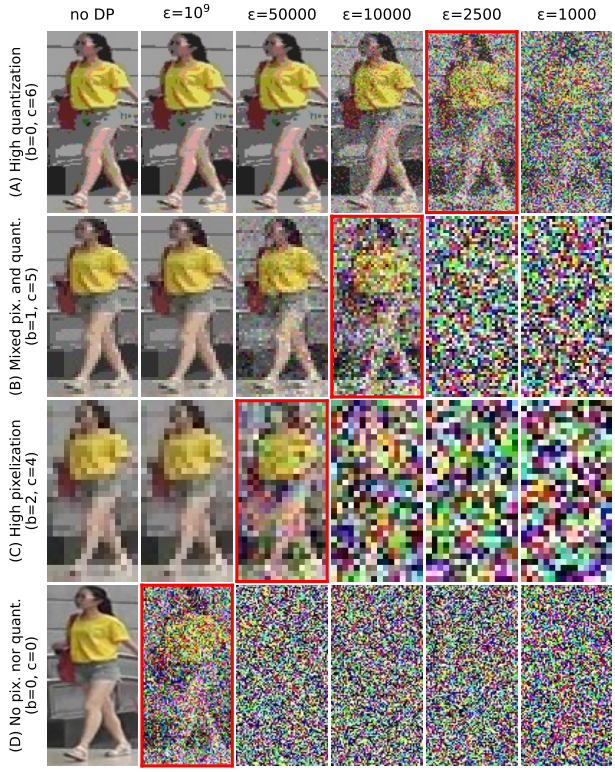

Fig. 6: Visual effect of image differential privacy on an example images from Market1501. Images obtained at the cutoff points selected in Table II are lined with a red frame.

V. Conclusion

Cameras are one of the most promising sensors for achieving IoT-enabled smart cities, but exploiting their data in large-scale projects raises legitimate privacy concerns. To address these, we introduce a novel, strict pixel-level image differential privacy mechanism, which allows smart city system operators to store pedestrian image data with quantifiable privacy guarantees. We show that by applying our privacy mechanism, one can significantly reduce the visual quality of an image dataset and limit the personal information it holds, thus lessening the gravity of potential data leaks. Despite this obfuscation, we identify sensible privacy-utility trade-off points where the images can still reasonably be used for person re-identification through centroid-based models, which we have shown to exhibit high robustness to differentially private noise addition. We expect these results to be useful for building privacy-compliant camera-based pedestrian flow information systems in smart cities, able to link together highly noised person representations without compromising pedestrians' privacy.

Acknowledgment

This study was supported in part by JSPS Grant-in-Aid for Scientific Research JP21H03431.

References

[1] M. Ye, J. Shen, G. Lin, T. Xiang, L. Shao, and S. C. H. Hoi, "Deep learning for person re-identification: A survey and outlook," *IEEE Transactions on Pattern Analysis & Machine Intelligence*, vol. 44, no. 06, pp. 2872–2893, 2022.

[2] C. Dwork and A. Roth, "The algorithmic foundations of differential privacy," *Foundations and Trends® in Theoretical Computer Science*, vol. 9, no. 3–4, pp. 211–407, 2014.

[3] L. Zheng, L. Shen, L. Tian, S. Wang, J. Wang, and Q. Tian, "Scalable person re-identification: A benchmark," in *2015 IEEE International Conference on Computer Vision*, 2015, pp. 1116–1124.

[4] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *2016 IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 770–778.

[5] A. Hermans, L. Beyer, and B. Leibe, "In defense of the triplet loss for person re-identification," *arXiv preprint arXiv:1703.07737*, 2017.

[6] H. Luo, Y. Gu, X. Liao, S. Lai, and W. Jiang, "Bag of tricks and a strong baseline for deep person re-identification," in *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 2019, pp. 1487–1495.

[7] M. Wieczorek, B. Rychalska, and J. Dabrowski, "On the unreasonable effectiveness of centroids in image retrieval," in *Neural Information Processing*, vol. 13111, 2021, pp. 212–223.

[8] L. Fan, T. Li, R. Fang, R. Hristov, Y. Yuan, and D. Katabi, "Learning longterm representations for person re-identification using radio signals," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020.

[9] S. Ahmad, G. Scarpellini, P. Morerio, and A. Del Bue, "Event-driven re-id: A new benchmark and method towards privacy-preserving person re-identification," in *IEEE/CVF Winter Conference on Applications of Computer Vision Workshops*, 2022, pp. 459–468.

[10] M. Ohno, R. Ukyo, T. Amano, H. Rizk, and H. Yamaguchi, "Privacy-preserving pedestrian tracking using distributed 3d lidars," in *2023 IEEE International Conference on Pervasive Computing and Communications*, 2023, pp. 43–52.

[11] J. Dietlmeier, J. Antony, K. McGuinness, and N. E. O'Connor, "How important are faces for person re-identification?" in *2020 25th International Conference on Pattern Recognition*, 2021, pp. 6912–6919.

[12] H. Wang, Y. Hong, Y. Kong, and J. Vaidya, "Publishing video data with indistinguishable objects," *Advances in database technology : proceedings. International Conference on Extending Database Technology*, vol. 2020, pp. 323–334, 2020.

[13] A. Kunchala, M. Bouroche, and B. Schoen-Phelan, "Towards a framework for privacy-preserving pedestrian analysis," in *2023 IEEE/CVF Winter Conference on Applications of Computer Vision*, 2023, p. 4370-4380.

[14] H. Cheng, H. Wang, X. Liu, Y. Fang, M. Wang, and X. Zhang, "Person re-identification over encrypted outsourced surveillance videos," *IEEE Transactions on Dependable and Secure Computing*, vol. 18, no. 3, pp. 1456–1473, 2021.

[15] G. Wu and S. Gong, "Decentralised learning from independent multi-domain labels for person re-identification," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, no. 4, pp. 2898–2906, 2021.

[16] L. Fan, "Image pixelization with differential privacy," in *Data and Applications Security and Privacy XXXII*, vol. 10980, 2018, p. 148-162.

[17] ——, "Practical image obfuscation with provable privacy," in *2019 IEEE International Conference on Multimedia and Expo*, 2019, pp. 784–789.

[18] C. Liu, J. Yang, W. Zhao, Y. Zhang, J. Li, and C. Mu, "Face image publication based on differential privacy," *Wireless Communications and Mobile Computing*, vol. 2021, pp. 1–20, 2021.

[19] T. Li and C. Clifton, "Differentially private imaging via latent space manipulation," in *IEEE Symposium on Security and Privacy*, 2021.

[20] W. Croft, J.-R. Sack, and W. Shi, "Obfuscation of images via differential privacy: From facial images to general images," *Peer-to-Peer Networking and Applications*, vol. 14, p. 1705–1733, 2021.

[21] Y. Zhao and J. Chen, "A survey on differential privacy for unstructured data content," *ACM Computing Surveys*, vol. 54, no. 10s, pp. 1–28, 2022.

[22] Y. Lin, L. Zheng, Z. Zheng, Y. Wu, Z. Hu, C. Yan, and Y. Yang, "Improving person re-identification by attribute and identity learning," *Pattern Recognition*, 2019.

[23] T. Ruan, T. Liu, Z. Huang, Y. Wei, S. Wei, and Y. Zhao, "Devil in the details: Towards accurate single and multiple human parsing," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, 2019, pp. 4814–4821.

[24] Z. Wang, A. Bovik, H. Sheikh, and E. Simoncelli, "Image quality assessment: from error visibility to structural similarity," *IEEE Transactions on Image Processing*, vol. 13, no. 4, pp. 600–612, 2004.