A World Leading SFI Research Centre

Science Foundation Ireland SFI For what's next

How important are faces for person re-identification?

Insight

SFI RESEARCH CENTRE FOR DATA ANALYTICS

Julia Dietlmeier, Joseph Antony,
Kevin McGuinness and Noel E. O'Connor

Dublin City University, Ireland

HOST INSTITUTIONS

DCU

NUI Galway
OÉ Gaillimh

UCC
University College Cork, Ireland
Coláiste na hOllscoile Corcaigh

UCD
DUBLIN

PARTNER INSTITUTIONS

Maynooth University
National University of Ireland Maynooth

Tyndall
National Institute

Trinity College Dublin
Coláiste na Tríonóide, Baile Átha Cliath
The University of Dublin

UNIVERSITY OF LIMERICK
OLLSCOIL LUIMNIGH

FUNDED BY:

European Union
European Regional Development Fund

Science Foundation Ireland SFI For what's next

Person re-identification (person re-ID)

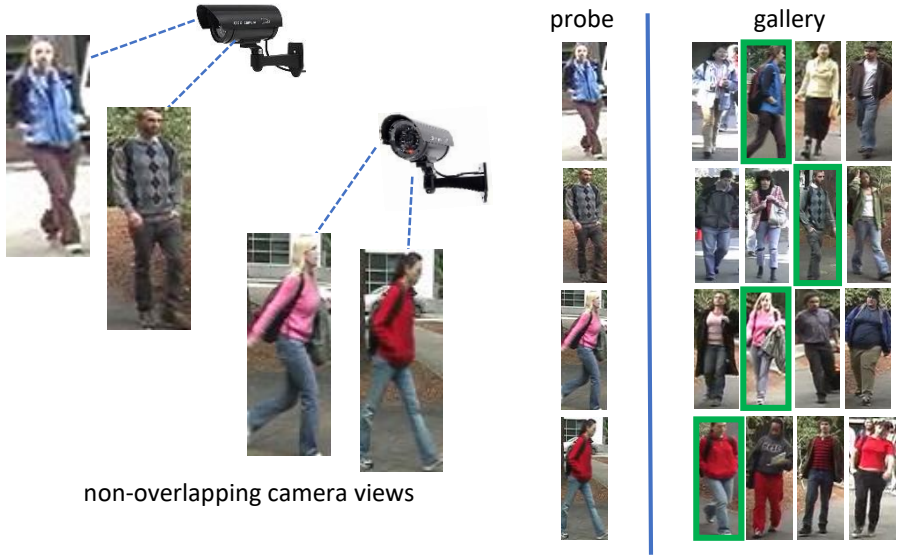

Figure 1 Person re-ID

The task of person re-identification (re-ID) is to retrieve images of a specified individual in a large non-overlapping multi-camera database, also called gallery, given a query person of interest.

It is an important component in intelligent video surveillance systems, which can be used to improve public safety with the increasing number of surveillance cameras in University campuses, theme parks, streets, train stations, and airports.

Motivation

- The release of new person re-identification datasets can raise legitimate privacy concerns.
- Sensitive data and its use is protected under law.
- Effective anonymization can allow data to be released to benefit society without compromising the identities of the individuals that appear in the data.
- Data protection laws such as GDPR do not apply to personal data that have been anonymized.
- Therefore, in the context of person re-ID, the anonymization process of a dataset could involve de-identification of individual faces.

Privacy preserving face anonymization

- Face anonymization or de-identification aims at removing privacy-sensitive information from detected faces.
- Existing approaches to face de-identification can be categorized into naive, the k-same family of algorithms, and the methods based on Generative Adversarial Networks (GANs).
- To our best knowledge, there have not been any prior studies on face anonymization in the context of person re-ID.
- We use a pre-trained state-of-the-art **TinyFaces** model to detect faces and locate a bounding box for each face. A crop of the bounding box is then extracted and blurred using a Gaussian blur filter with radius equal to 1/8 the width of the size of the bounding box.

Privacy preserving face anonymization

Figure 2 Examples of original (first row) and our anonymized (second row) images from the Market1501 re-ID dataset. For face de-identification, we apply the **TinyFaces** face detector first and then blur the detected region with a large kernel Gaussian to remove all privacy-sensitive information.

Image-based person re-ID datasets

Table 1 Statistics of the image-based person re-ID datasets used in the experimental part of this paper.

Dataset	Released in	Cameras	Identities	Training images	Gallery images	Query images
Market1501 [22]	2015	6	1,501	12,936	19,732	3,368
DukeMTMC-reID [23]	2017	8	1,404	16,522	17,661	2,228
CUHK03 (detected) [24]	2014	6	1,360	7,365	1,400	5,332
Airport [26], [36]	2018	6	770	3,493	2,420	1,003
VIPeR [25]	2008	2	632	316	316	316

We evaluate the effect of anonymization on five popular person re-identification datasets: Market1501, DukeMTMCreID, CUHK03, VIPeR, and Airport. **Table 1** summarizes statistics for these datasets.

Person re-ID models

Table 2 Reported performance of selected state-of-the-art person re-ID models.

Models	Venue	Market1501		DukeMTMC-reID	
		mAP	Rank1	mAP	Rank1
BoT [27]	CVPRW 2019	94.2	95.4	89.1	90.3
PCB [28]	ECCV 2018	81.6	93.8	69.2	83.3
MLFN [29]	CVPR 2018	74.3	90.0	62.8	81.0
HACNN [30]	CVPR 2018	75.7	91.2	63.8	80.5
Resnet50Mid [31]	arXiv 2017	75.6	89.9	63.9	80.4

Best performing →

- Approaches to person re-identification can be categorized into traditional and deep learning-based methods. Traditional methods aim to design hand-crafted features and learn an effective distance metric. Convolutional neural networks (CNNs), recurrent neural networks (RNNs), and GANs form the basis for deep learning-based person re-ID.

Performance evaluation results

Table 3 Market1501 performance evaluation.

	Trained and tested on original		Trained on original, tested on anonymized		Trained and tested on anonymized	
Models	mAP	Rank1	mAP	Rank1	mAP	Rank1
PCB	72.8	87.8	71.7	87.4	72.9	88.2
MLFN	71.4	87.4	70.9	86.9	71.3	87.5
HACNN	67.4	85.5	66.5	85.7	66.6	85.2
Resnet50Mid	72.9	88.1	71.5	87.5	72.7	87.9
MuDeep	44.9	70.4	43.7	70.0	44.8	69.6
BoT	94.1	95.5	93.7	95.0	94.0	95.2

Table 4 DukeMTMC-reID performance evaluation.

	Trained and tested on original		Trained on original, tested on anonymized		Trained and tested on anonymized	
Models	mAP	Rank1	mAP	Rank1	mAP	Rank1
PCB	66.3	80.3	64.4	79.3	65.4	80.2
MLFN	60.4	78.1	57.9	76.8	59.7	77.2
HACNN	57.4	74.0	55.5	72.2	56.7	73.1
Resnet50Mid	62.9	80.1	60.1	78.7	61.8	80.8
MuDeep	36.0	56.6	34.4	54.4	34.8	54.7
BoT	88.8	90.4	87.8	89.7	88.6	90.3

We report the Rank1 accuracy and mean Average Precision (mAP) as standard performance evaluation metrics.

Performance evaluation results

- We observe that the performance results follow the same pattern for each dataset. The performance is marginally diminished when we train on the original and evaluate all models on the anonymized versions of all datasets. The accuracy is, however, recovered by training on the anonymized versions of all datasets.

- The average mAP across all models and datasets before anonymization is 58.07 and 56.68 after anonymization (-1.39), and 57.44 after retraining on the anonymized data (-0.63).
  Anonymization and retraining does have a statistically significant effect on mAP
  (one-sided paired t-test, $p < 0.01$) but the effect size is small (0.63).
  Similar conclusions are found for rank-1 accuracy: average rank-1 accuracy before anonymization
  is 66.69, and after retraining is 66.25 (-0.44).
  Again, the difference in mean is significant ($p < 0.05$), but very small (less than half a percentage point).

Effect of different anonymization techniques

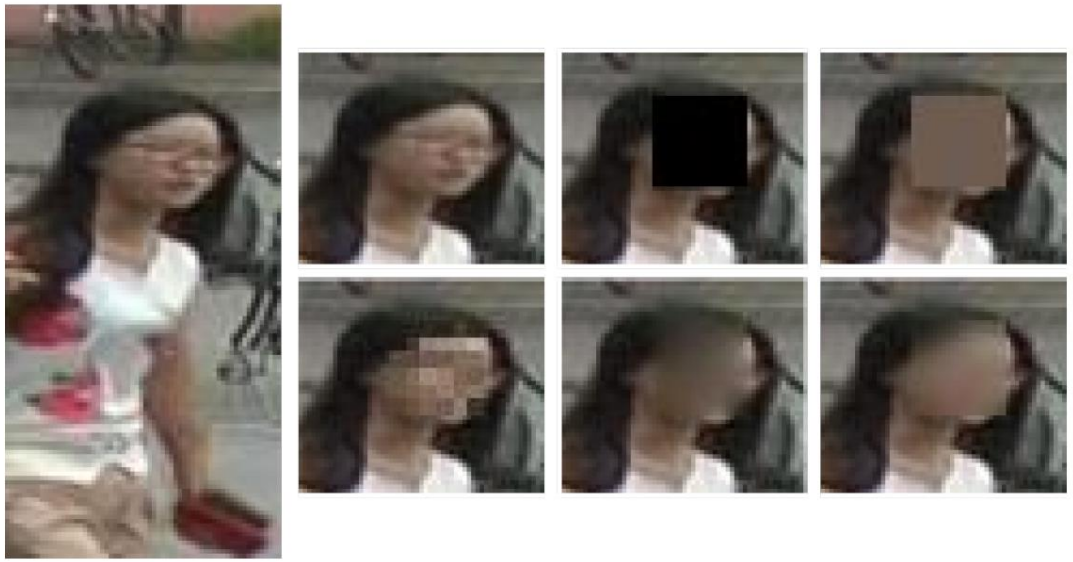

Figure 3 Example of different face anonymization techniques on a 64x128 sample image from the image-based Market1501 dataset. **Top row**: original face image, blackout on the face detected with the TinyFaces detector, blank; **Bottom row**: pixelated, inpainted, blurred with a large-kernel Gaussian.

Effect of different anonymization techniques

Table 5 Performance evaluation of the BoT model on Market1501 dataset under different face anonymization techniques.

	Trained on original, tested on anonymized		Trained and tested on anonymized	
Face anon. tech.	mAP	Rank1	mAP	Rank1
Blur	93.7	95.0	94.0	95.2
Zero	93.4	94.6	93.8	94.5
Blank	93.7	95.0	93.8	95.7
Pixelate	94.0	95.3	94.0	95.4
Inpaint	93.6	95.0	93.6	95.0

- We observe that most methods provide comparable results, with blur and pixelate being slightly superior. The results are somewhat better when trained and tested on the anonymized versions of the Market1501 dataset.
- Zeroing the faces (replacing all face pixels with zero) is only 0.25% on average below blur. Therefore, the complete zeroing of the face can be used instead of blurring in cases when extra security is needed.

Discussion

- Existing SoA for deblurring faces are focused on either removing a type of structured face blurring that occurs due to motion or relatively minor blurring that can occur due to camera defocus and similar effects.
- The structure present in motion blur offers the opportunity to achieve very good results in removing this type of distortion, but these techniques do not work on strong (large kernel, high variance) Gaussian blurs, which simply remove all high frequency components in the Fourier spectrum.
- **Deblurring techniques cannot recover faces that have been blurred using strong Gaussian blurs, since the information about the detail in the face is effectively completely removed.**
- As an example, Google uses similar image blurring technology in their Google Street View system to blur faces and license plates and that this technology has been deployed since 2009 without any major incidents in terms of deblurring the images.

Conclusions

- In this study, we empirically observed that **the effect of blurring faces on the person re-identification performance is surprisingly small**.
- We also find that the relative performance of different state-of-the-art methods is preserved after anonymization, meaning that **new approaches can be safely compared using anonymized data**.
- This finding could potentially pave the way for the **release of larger anonymized datasets**.

Thank you!

ACKNOWLEDGMENT

This publication has emanated from research conducted with the financial support of Science Foundation Ireland (SFI) under grant number SFI/15/SIRG/3283 and SFI/12/RC/2289 P2. We thank Bhartendu Sharma for his assistance in running several experiments. We would also like to thank Dublin Airport Authority for their contributions to this work.