

How important are faces for person re-identification?

Julia Dietlmeier, Joseph Antony,
Kevin McGuinness and Noel E. O'Connor
Dublin City University, Ireland



SFI RESEARCH CENTRE FOR DATA ANALYTICS

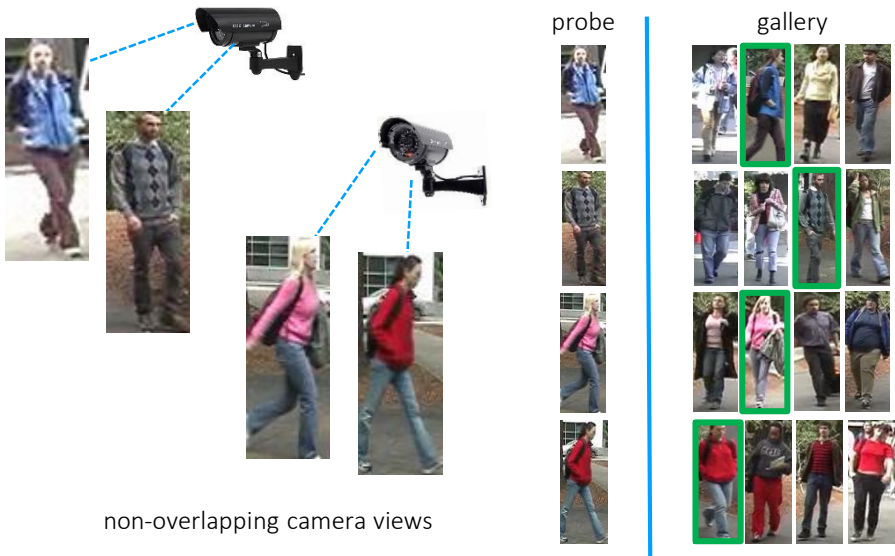
A World
Leading SFI
Research
Centre



Overview

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.

Face anonymization is needed to preserve privacy and to comply with GDPR.



Models and Datasets

We apply a face detection and blurring algorithm to create **anonymized versions of several popular person re-identification datasets** including Market1501, DukeMTMC-reID, CUHK03, Viper, and Airport.

We benchmark several state-of-the-art person re-ID models such as PCB, MLFN, HACNN, Resnet50Mid, MuDeep, and BoT.

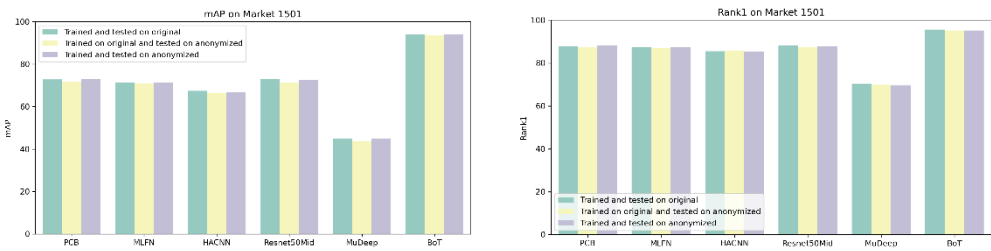


The figure above shows selected examples of original (first row) and our anonymized (second row) images from the Market1501 re-ID dataset.

To anonymize the data we first apply the **TinyFaces** face detector and then blur the detected region with a large kernel Gaussian to remove all privacy-sensitive information. The effectiveness of this approach was manually verified.

Experimental Results

Extensive experiments on multiple datasets showed that face anonymization **surprisingly only has a very slight impact on performance**, and that this slight impact can be mostly recovered by training on anonymized data. The graphic below shows mAP and Rank1 for Market1501.



We also investigate how **different types of face de-identification** impact re-ID performance. The table below shows the results. This experiment focused on the state-of-the-art BoT model and the Market1501 dataset. **All methods gave similar results**, with blur and pixelate being slightly superior.

Face anon. tech.	Trained on original, tested on anonymized		Trained and tested on anonymized	
	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



Example of the different face anonymization techniques used

Conclusion

The study empirically showed that the effect of blurring faces on the person re-identification performance is surprisingly small and that the relative performance of different state-of-the-art methods is preserved after anonymization. This implies that that **new Re-ID algorithms can be safely compared using anonymized data**, which could potentially pave the way for the release of larger anonymized datasets.

Acknowledgement

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 would also like to thank Dublin Airport Authority for their contributions to this work.

HOST INSTITUTIONS



PARTNER INSTITUTIONS



FUNDED BY:

