Practical Image Obfuscation with Provable Privacy

Liyue Fan, Department of Computer Science, University of North Carolina at Charlotte

Bibliographic reference:

L. Fan, "Practical Image Obfuscation with Provable Privacy," *2019 IEEE International Conference on Multimedia and Expo (ICME)*, Shanghai, China, 2019, pp. 784-789.
doi: 10.1109/ICME.2019.00140

Abstract:

An increasing amount of image data is being generated nowadays, thanks to the popularity of surveillance cameras and camera-equipped personal devices. While such image data can be shared widely to enable research studies, it often contains sensitive information, such as individual identities, location indications, etc. Therefore, the image data must be sanitized before sharing with untrusted parties. Current image privacy-enhancing solutions do not offer provable privacy guarantees, or sacrifice utility to achieve the standard ε-differential privacy. In this study, we propose a novel image obfuscation solution based on metric privacy, a rigorous privacy notion generalized from differential privacy. The key advantage of our solution is that our privacy model allows for higher utility by providing indistinguishability based on image visual similarity, compared to the current method with standard differential privacy. Empirical evaluation with real-world datasets demonstrates that our method provides high utility while providing provable privacy guarantees.

URL:

http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=8784836&isnumber=8784700

Draft: Image Obfuscation with Quantifiable Privacy

Liyue Fan Assistant Professor @ UNC Charlotte

UNC CHARLOTTE College of Computing and Informatics

Introduction

Image obfuscation is widely used to protect private content in photos, such as Google street view [1] and journalism [2]. Some opular obfuscation techniques:

Blurring

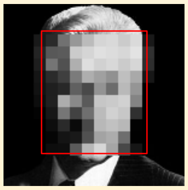

Pixelization

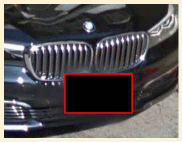

Blacking

However, machine learning models can adapt to standard obfuscation. For example:

- Hill et. al [3]

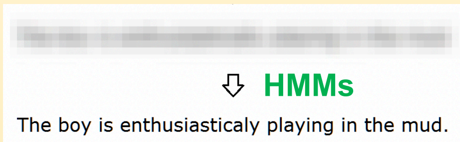

- McPherson et. al [4]

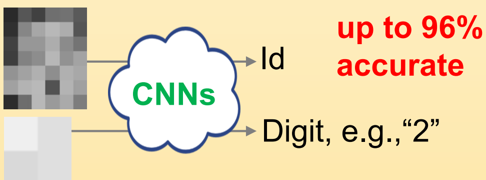

- Oh et. al [5]

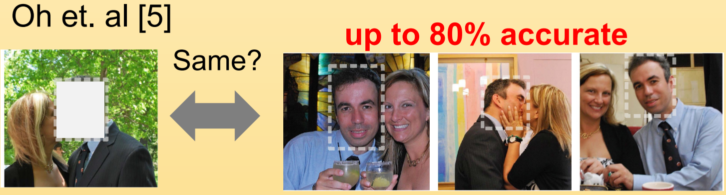

Prior Research:

- Sun et. al [6] and Ren et. al [7] adopt GANs to modify identities, but do not provide formal privacy.
- Fan [8] achieves rigorous ϵ –**Differential Privacy** but low utility, due to an *overly strong* privacy model.

Method

Objective 1: Quantifiable Privacy for ROIs

Objective 2: Privacy Utility Trade-off

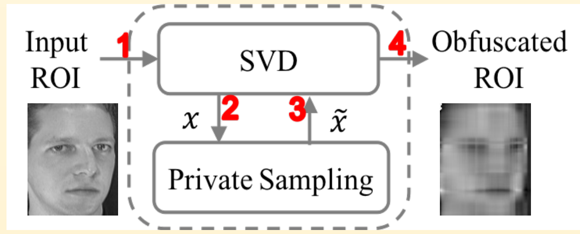

Metric Privacy ($\epsilon \times d_{\mathcal{X}}$-privacy) [9] for any secret pair x and x':

$$K(x)(Z) \le e^{\epsilon \times d_{\mathcal{X}}(x,x')} K(x')(Z), \quad \forall \text{ output } Z$$

- Privacy based on “similarity” → Utility friendly
- Standard DP is a special instance of Metric Privacy [9]

Results:

Row 1 – original images; Row 2 – our method, ϵ = 0.1; Row 3 – our method, ϵ = 0.3; Row 4 – our method, ϵ = 1; Row 5 – [8], ϵ = 1.

Acknowledgements: This research is supported in part by NSF grant CNS-1755884. Any opinions, findings, and conclusions or recommendations do not necessarily reflect the views of the National Science Foundation.

References

1. A. Frome et al., "Large-scale privacy protection in Google Street View," 2009 IEEE 12th International Conference on Computer Vision, Kyoto, 2009, pp. 2373-2380.
2. D. Aitkenhead. `I've done really bad things': The undercover cop who abandoned the war on drugs. The Guardian, 2016.
3. Hill, S., Zhou, Z., Saul, L., & Shacham, H. (2016). On the (In)effectiveness of Mosaicing and Blurring as Tools for Document Redaction, Proceedings on Privacy Enhancing Technologies, 2016(4).
4. Richard McPherson, Reza Shokri, and Vitaly Shmatikov. Defeating image obfuscation with deep learning. CoRR, abs/1609.00408, 2016.
5. Oh, Seong Joon, et al. "Faceless person recognition: Privacy implications in social media." European Conference on Computer Vision. Springer, Cham, 2016.
6. Qianru Sun, Liqian Ma, Seong Joon Oh, Luc Van Gool, Bernt Schiele, and Mario Fritz. Natural and effective obfuscation by head inpainting. In The IEEE Conference on Computer Vision and Pattern Recognition (CVPR), June 2018.
7. Zhongzheng Ren, Yong Jae Lee, and Michael S Ryoo. Learning to anonymize faces for privacy preserving action detection. In ECCV, pages 620–636, 2018.
8. Liyue Fan. Image pixelization with differential privacy. In Data and Applications Security and Privacy XXXII, pages 148–162, Springer Cham, 2018.
9. Chatzikokolakis, Konstantinos, et al. "Broadening the scope of differential privacy using metrics." International Symposium on Privacy Enhancing Technologies Symposium. Springer, Berlin, Heidelberg, 2013.