Faceless Person Recognition; Privacy Implications in Social Media

Seong Joon Oh, Rodrigo Benenson, Mario Fritz, Bernt Schiele

Max Planck Institute for Informatics, Saarland Informatics Campus, Saarbrücken, Germany

Motivation

- How much private information can be exposed from social photos via computer vision?
- How robust are the state of the art person recognisers to head blur?
- Which actions can users take to protect their privacy?

Challenges in Analysis

- Can only lower bound on the performance of the best corporate systems, due to a limited access to the large scale private user databases.
- How to simulate users with varying degrees of privacy sensitivity?
- How to aggregate personal information spread across multiple photos?

Setup for Analysis

- Person recognition in social media.
- Closed world assumption: Recognise from a finite set of identities (200~600).
- GT head boxes are given on all the instances.
- Fuse information from non-tagged instances in the same album and < 10 tagged instances per identity.
- Consider multiple identity protection scenarios.
- **Dataset: Person In Photo Albums (PIPA) [1]**

Who is this person inside an album?

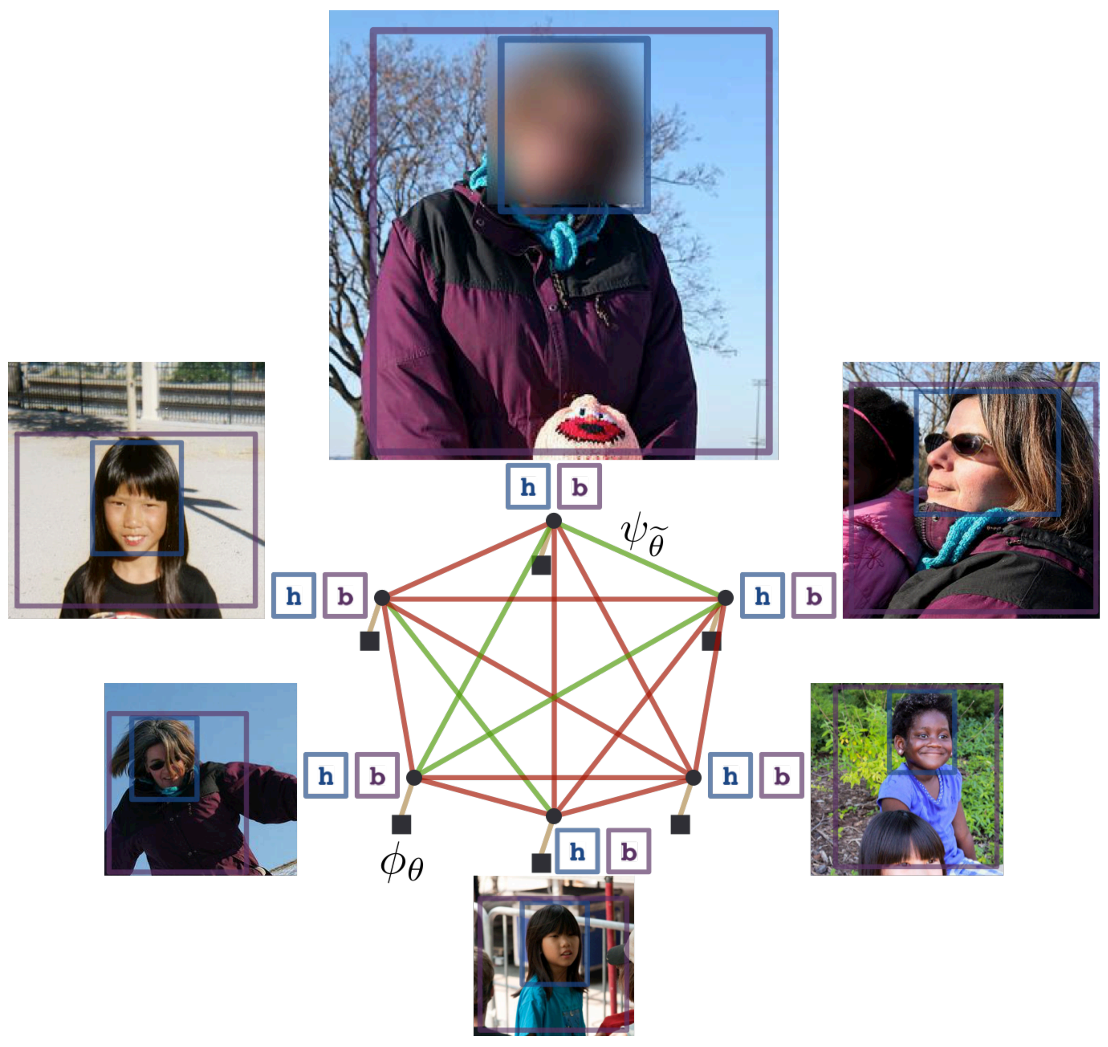

… given some tagged images?

Only 4 tagged images.

+

Other tagged people.

Conclusion in a Nutshell

1. State of the art person recognisers are robust to common identity protection measures.
2. Further performance boost from 1) adapting system to obfuscation patterns and 2) jointly reasoning across photos.
3. Even in the most protective scenario (no identity tag in the same event photos, all heads obfuscated), achieve 12x above naïve guess.

Faceless Person Recognition

$$\arg\max_{Y} \frac{1}{|V|}\sum_{i\in V}\phi_\theta(Y_i|X_i) + \frac{\alpha}{|E|}\sum_{(i,j)\in E} 1_{[Y_i=Y_j]}\psi_{\tilde{\theta}}(X_i, X_j)$$

Unary: single person recogniser. ϕ_θ

- Identity probabilities for a single person.
- State of the art CNN full-body recogniser [2].
- Fine-tuned for obfuscation patterns.

Pairwise: person pair matcher. $\psi_{\tilde{\theta}}$

- Match probabilities for person pairs.
- Siamese network trained for matching.
- Fine-tuned for obfuscation patterns.

Identity Protection Scenarios

Number of tagged photos & amount of head obfuscation

Who? | In the same album | Tagged examples

$S_1^{\tau=2}$ Many tagged heads

$S_1^{\tau=1}$ Few tagged heads

S_2 Obfuscate query head

S_3 Obfuscate every head

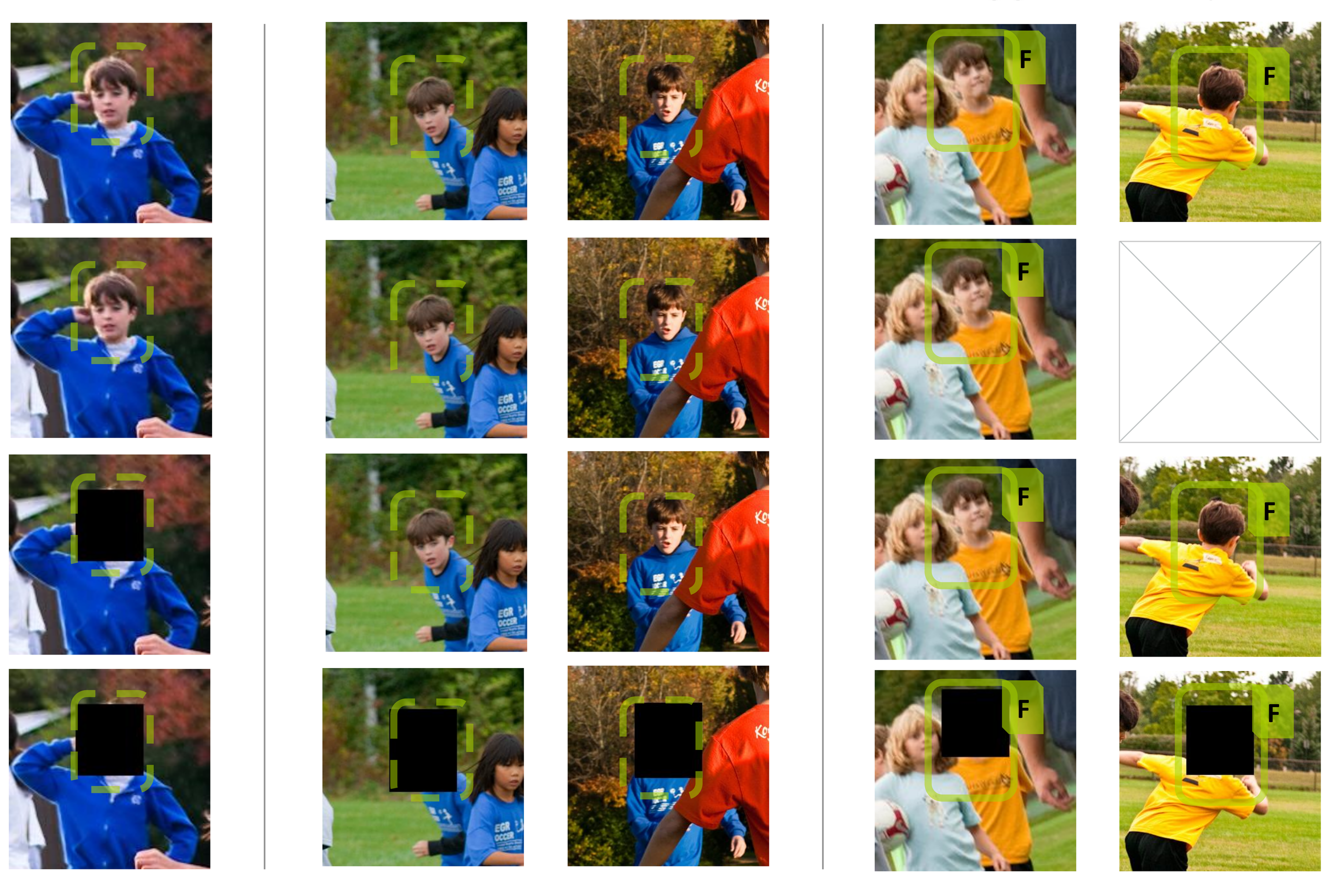

Head obfuscation types

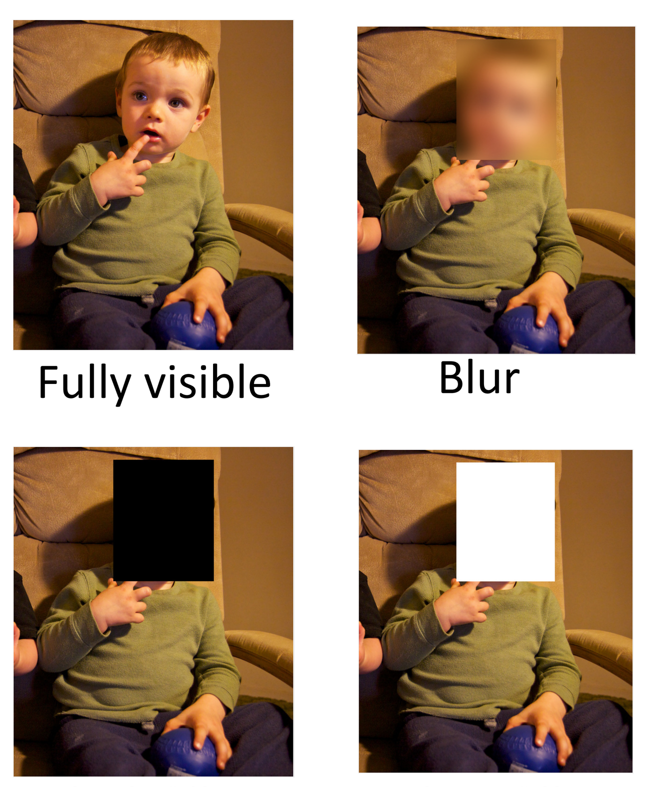

Fully visible | Blur | Black fill-in | White fill-in

Domain shift [2]

- Within events: Similar clothing.
- Across events: Changed clothing.

Quantitative Results

Identification accuracy versus tag rate.

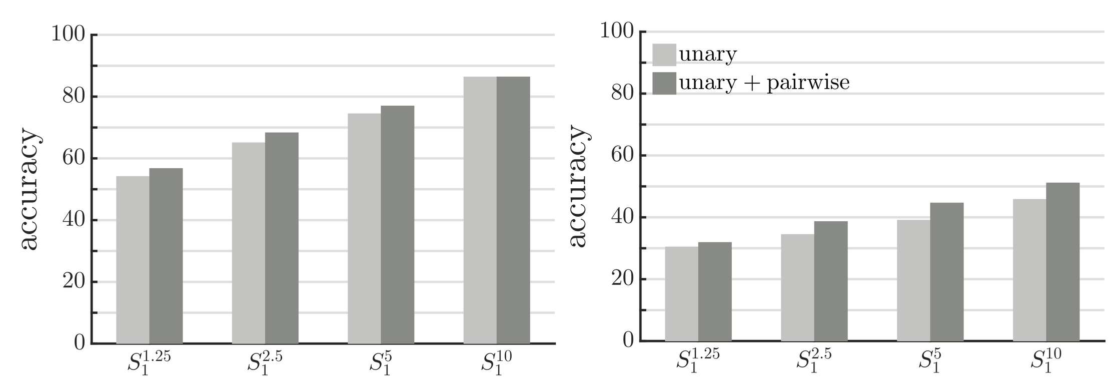

… versus obfuscation type & amount.

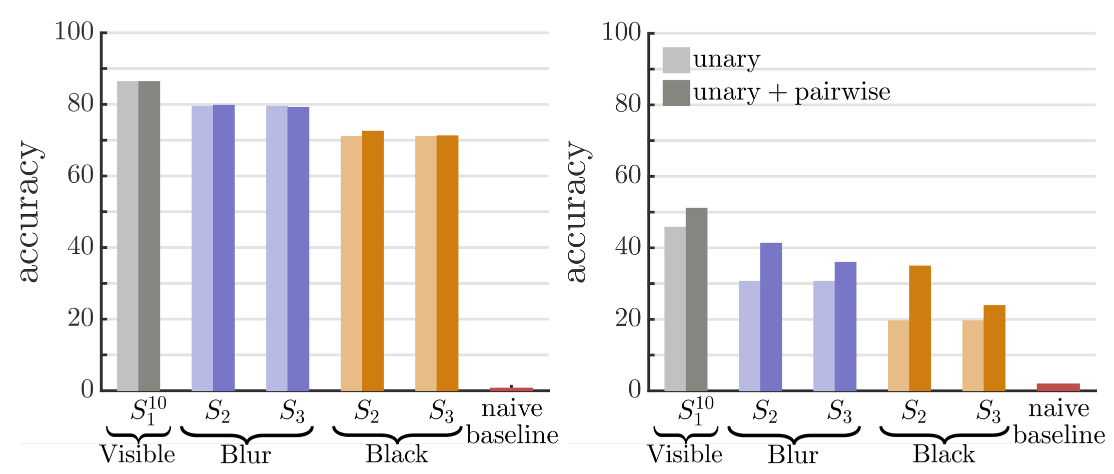

- **Number of tagged photos:**
  - 1.25 tags / person → still far better than naïve baseline.
- **Amount of head obfuscation:**
  - Within events: ineffective way of protection.
  - Across events: most effective if all heads are blacked out.
- **Head obfuscation types:**
  - Black ≈White >Blur >Visible.
- **Domain shift:**
  - Recognition system struggles more across events.
  - **Take-away**: Make sure no tagged heads exist for the event where you want protection.

Qualitative Results

S_2 Blur | S_3 Blur | S_3 Black

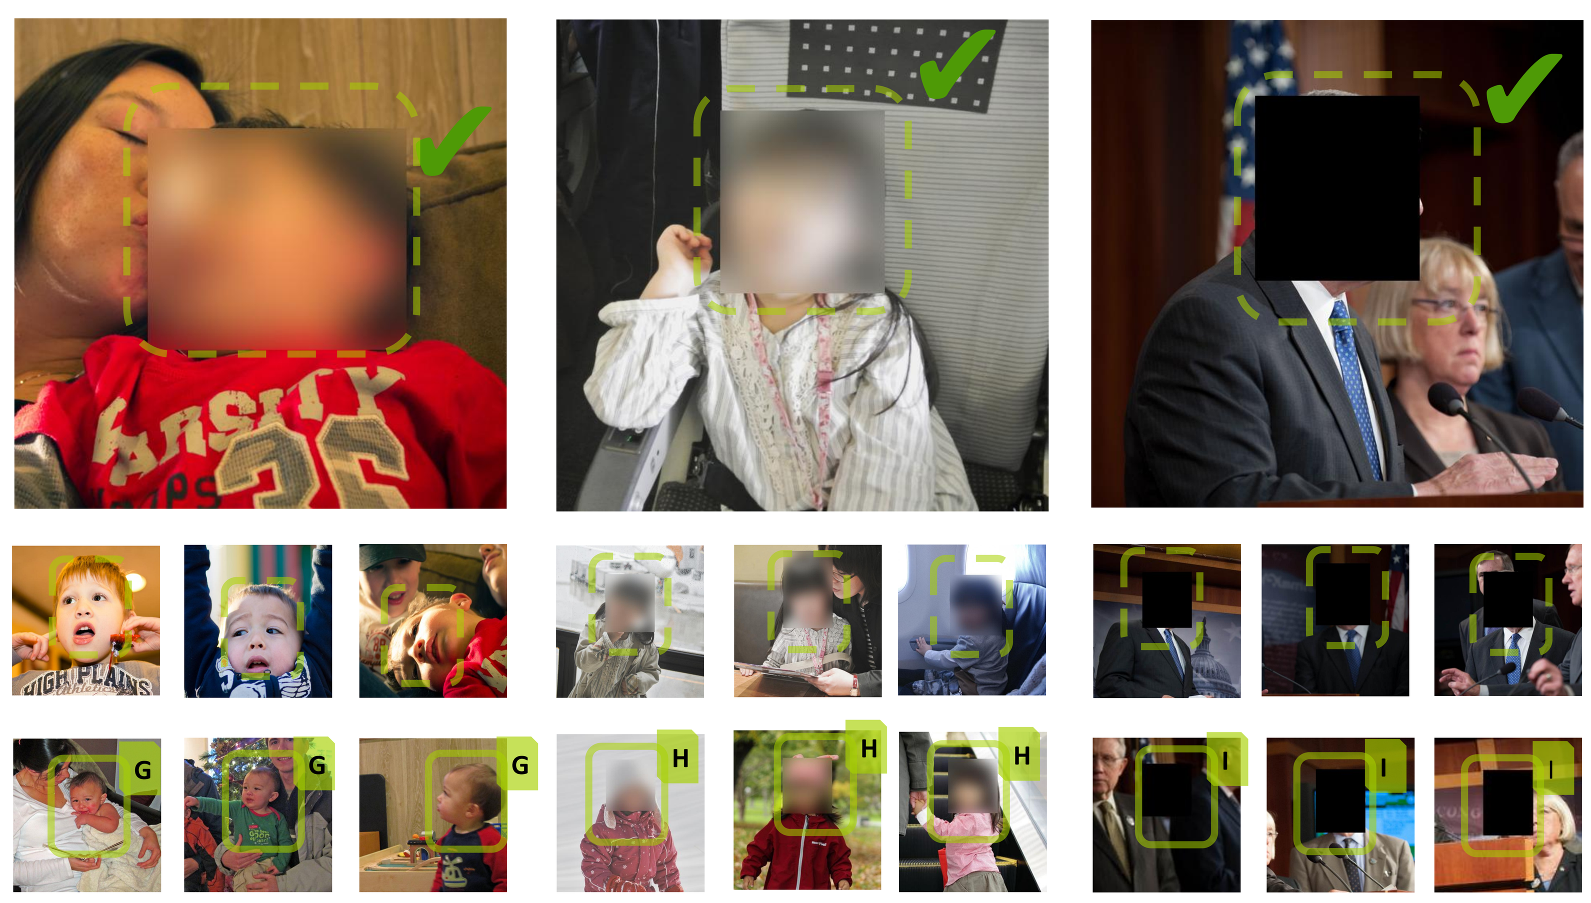

In the same album

Tagged positive examples

Acknowledgements

This research was supported by the German Research Foundation (DFG CRF 1223).

References

[1] Zhang et al. Beyond frontal faces: improving person recognition using multiple cues. CVPR'15.
[2] Oh et al. Person Recognition in Personal Photo Collections. ICCV'15.