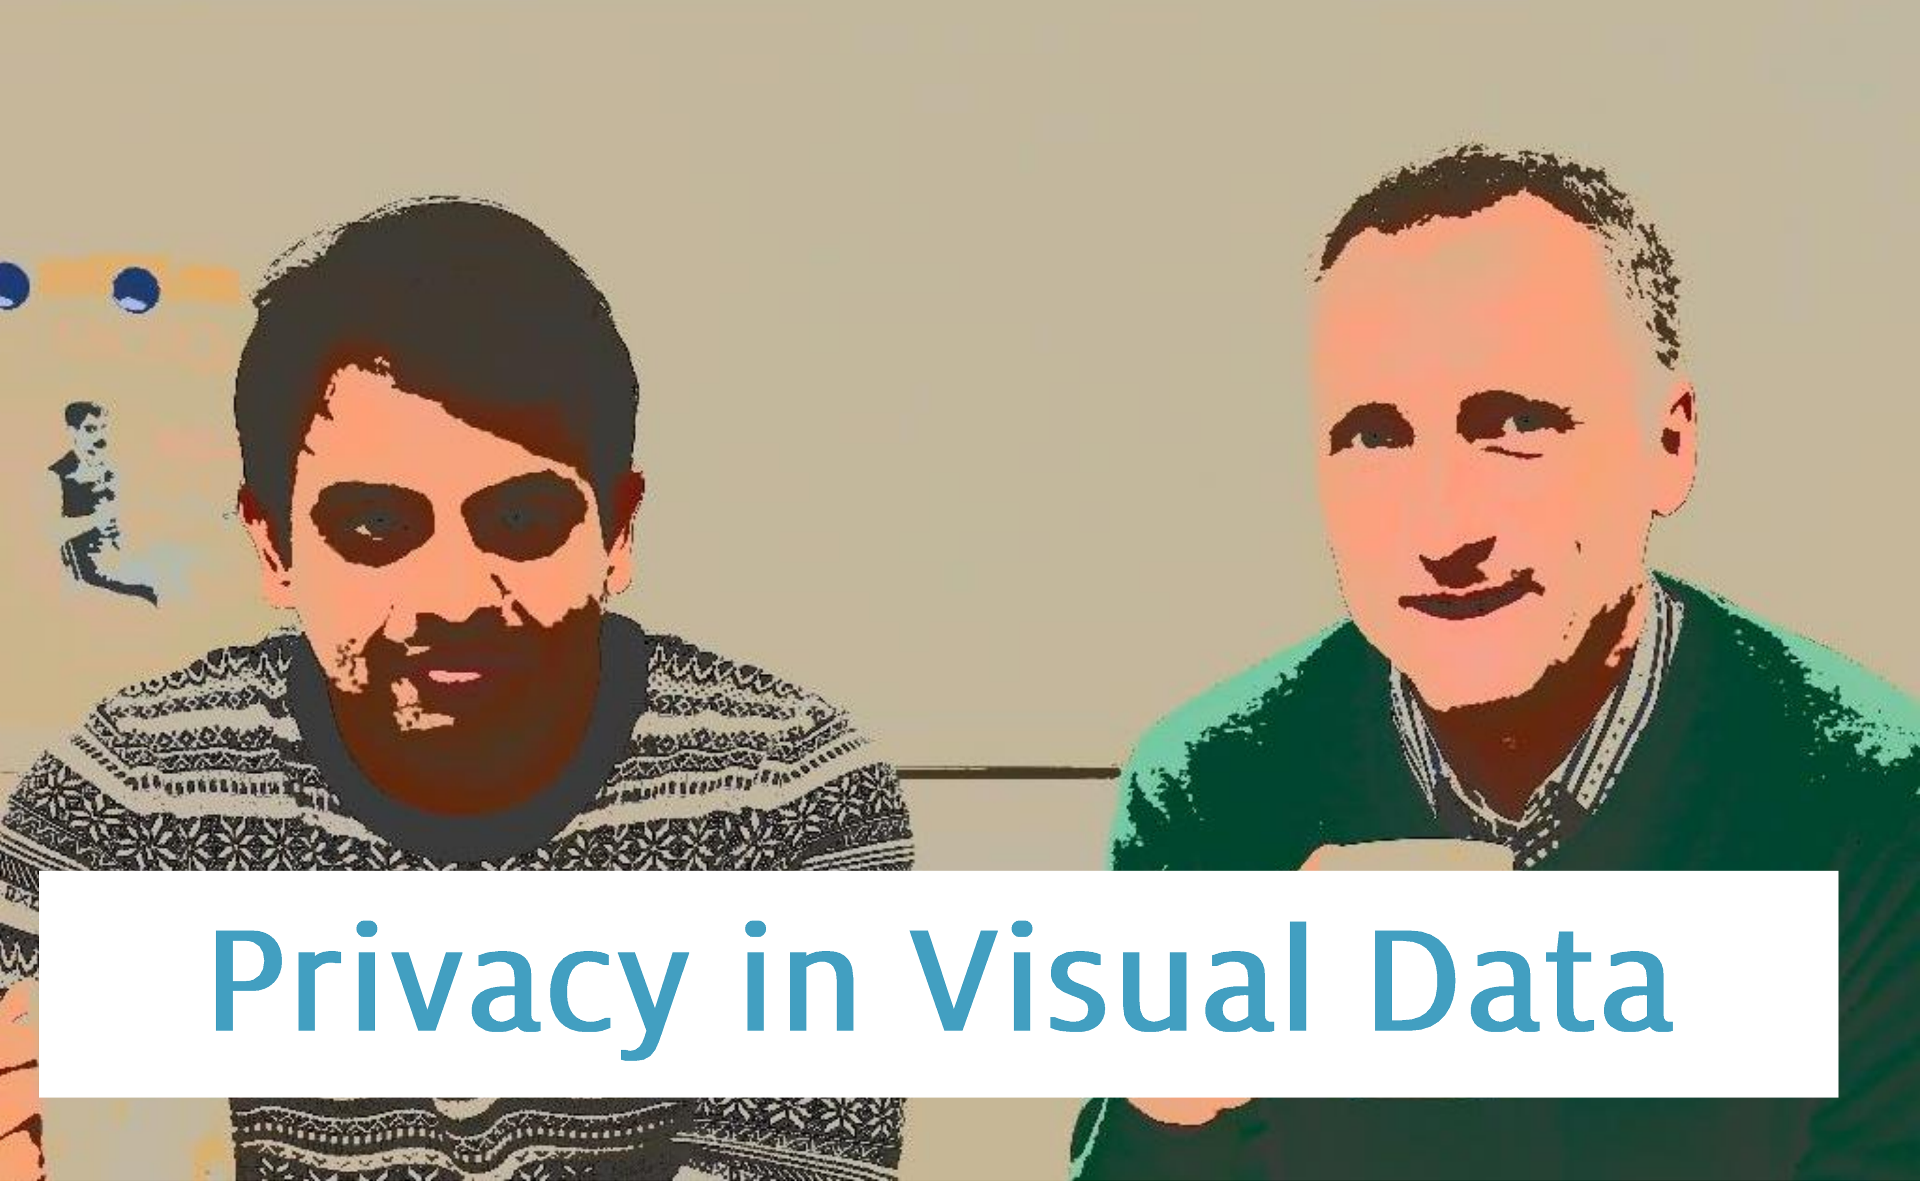

Privacy in Visual Data

Bernhard Rinner

Ljubljana, September 20, 2017

ALPEN-ADRIA UNIVERSITÄT KLAGENFURT | WIEN GRAZ
Institute of Networked and Embedded Systems

Privacy and its Protection

- Privacy is related to “the ability to seclude themselves, or information about themselves”
  - highly subjective and context dependent
- Privacy has a significant impact on society
  - addressed in numerous fields
  - controversially discussed
- Privacy is increasingly at risk
  - Technological advances, limited awareness, change in politics

[catphi.wordpress.com]

Ubiquity of Cameras

- We are surrounded by billions of cameras in public, private and business
- Huge amounts of image/video data is endlessly captured and shared
- Analysis and networking capabilities advance at astonishing rates
- Limited awareness about privacy threats

[spiegel.de; givenimaging.com]

Privacy in Data(bases)

- Draw conclusions for the entire population (or parts of) but avoid linkage of sensitive information to individuals

Name	SSN	Age	ZIP	Sex	Disease
[redacted]		[30,39]	9***	female	Flu
[redacted]		[40,49]	9***	male	Cancer
[redacted]		[30,39]	9***	female	Flu
[redacted]		[40,49]	9***	male	Flu
…	…	…		…	…

Explicit identifier (Name, SSN) — Quasi identifier (Age, ZIP, Sex) — Sensitive information (Disease)

- Anonymization as key protection method
- Modify quasi identifier to achieve k anonymity

Privacy in Visual Data

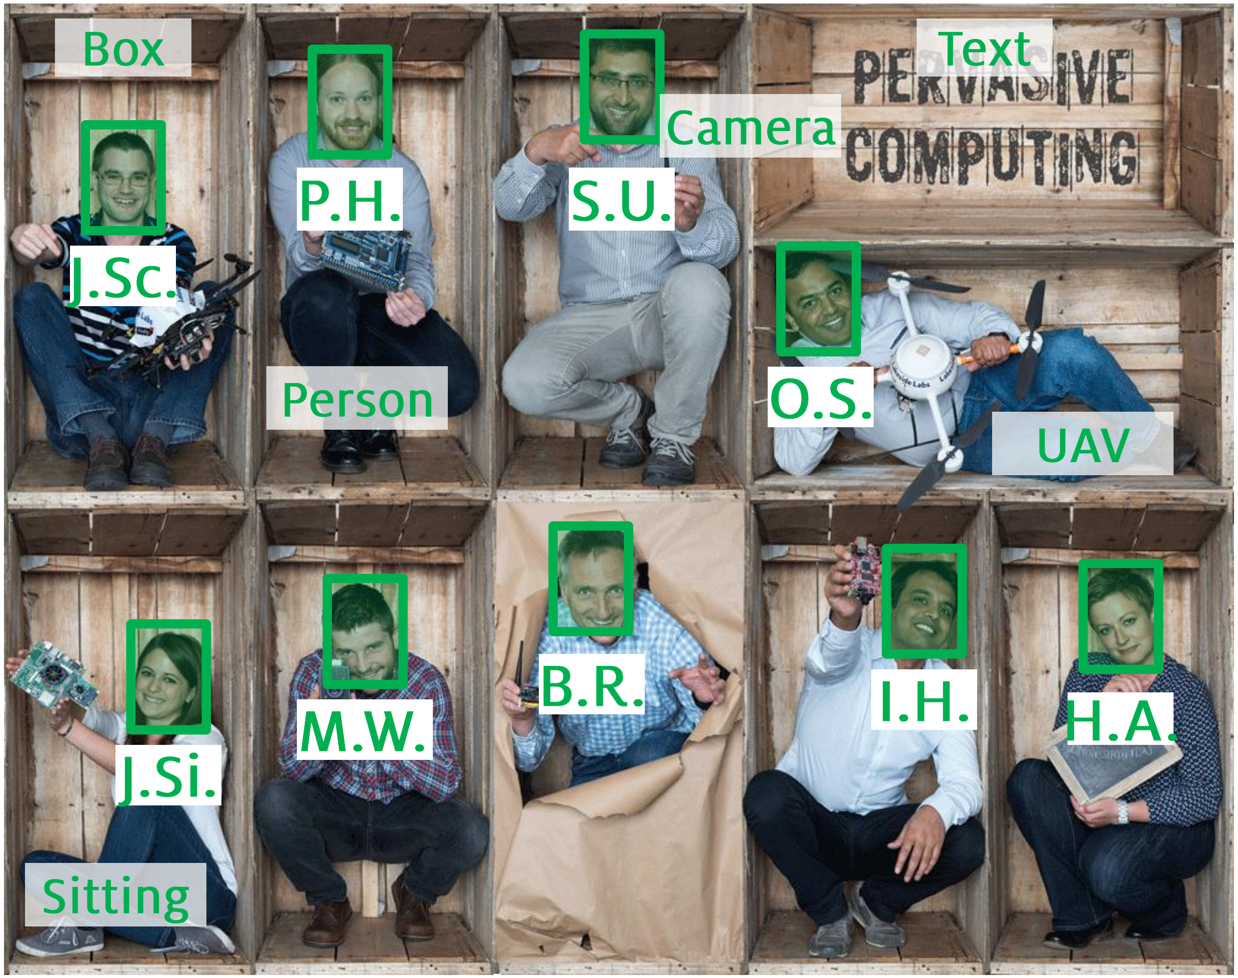

Who is there?

- (Quasi-) Identifiers
- Body or face regions

What is shown?

- Sensitive information
- Presence,
  „show an object"
  „captured in a box"

How to avoid linkage of sensitive information to individuals?

Utility and Privacy Tradeoff

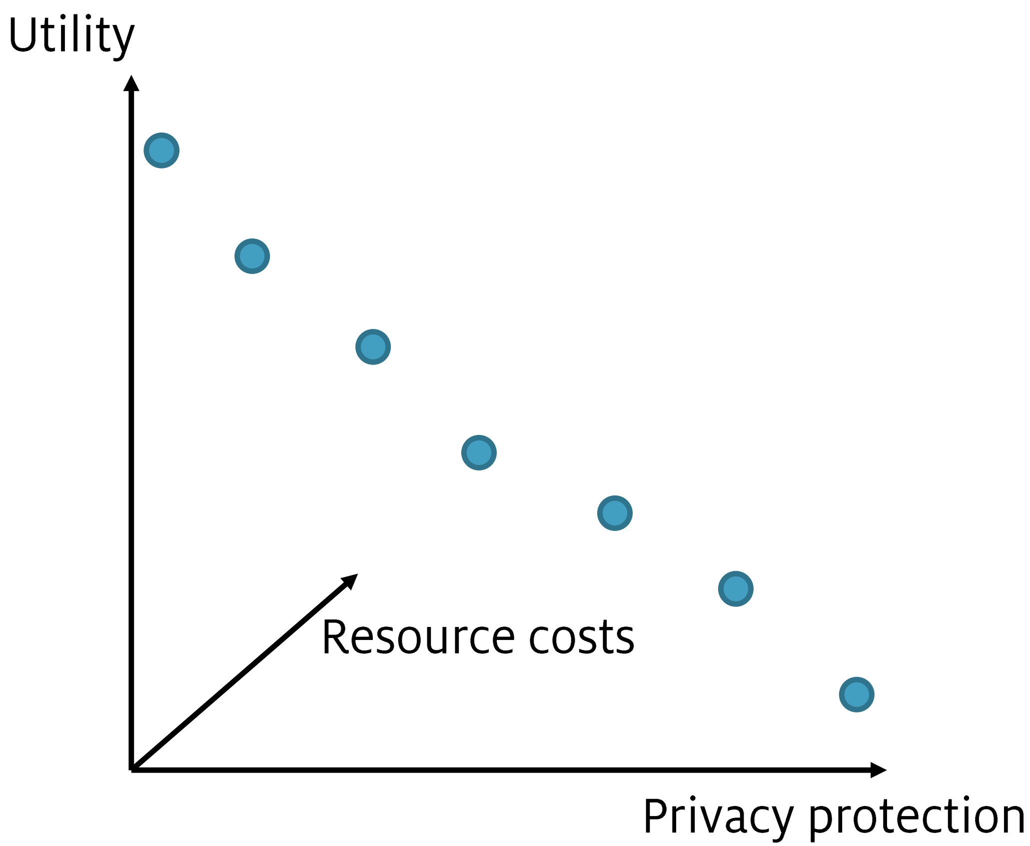

Distortion as key protection method

- Blanking
- Pixelation
- Bluring
- Cartooning

Utility dependent on level of distortion

- Similarity
- Appearance
- Detectability

No single best protection method available

Agenda

1. What distortion method to use in aerial imagery?
   - Explore utility/privacy/cost design space
   - Adapt filter strength for recreational images
   - Measure achieved privacy protection and utility

[www.radiogong.de]

2. How to securely implement privacy protection?
   - Apply security methods at sensory edge
   - Rely on hardware-supported protection

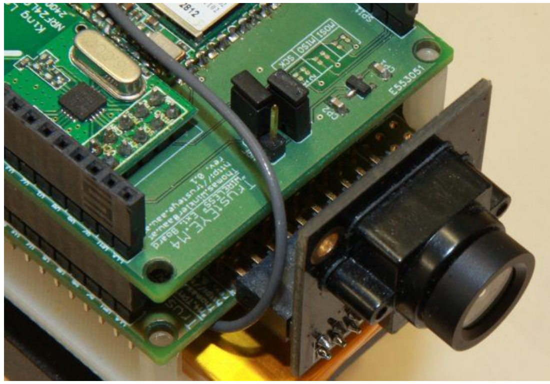

[Winkler, Rinner. Security and Privacy Protection in Visual Sensor Networks: A Survey. ACM Computing Surveys. 2014.]

#1 Privacy Protection in Recreational Aerial Images

Recreational Airborne Cameras

- Micro Aerial Vehicles (MAVs) are becoming common in public places

  for recreational and business video capturing

  with high-resolution cameras

- How can we protect privacy while maintaining high fidelity of visual data?

- Exploring the privacy design space
  - When is protection necessary at all?
- Configuring an adaptive privacy filter
  - What is the minimal protection?

www.hexaplus.com

www.airdog.com

www.kickstarter.com

Adapt Blur to Target Resolution

- Privacy design space exploration with adaptive filtering
  - Determine target‘s pixel density based on camera pose
  - Decide whether target is inherently protected
  - Configure privacy protection filter
  - Perform adaptive filtering
- Studied for aerial images

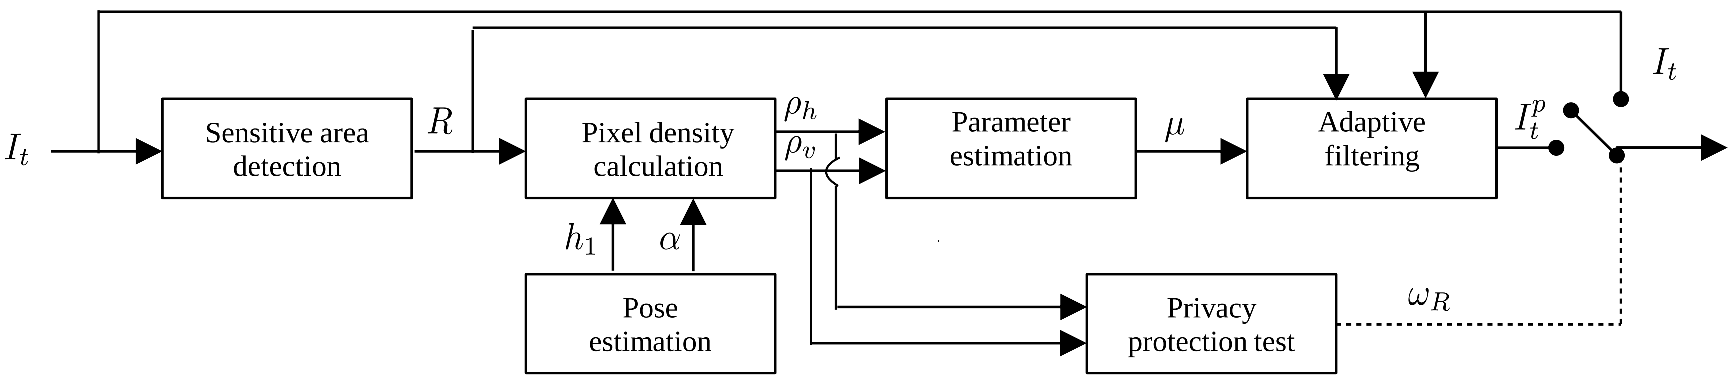

[Sawar, Rinner, Cavallaro. Design Space Exploration for Adaptive Privacy Protection in Airborne Images. In Proc. AVSS 2016.]

Pixel Density Estimation

- Horizontal and vertical density at target center

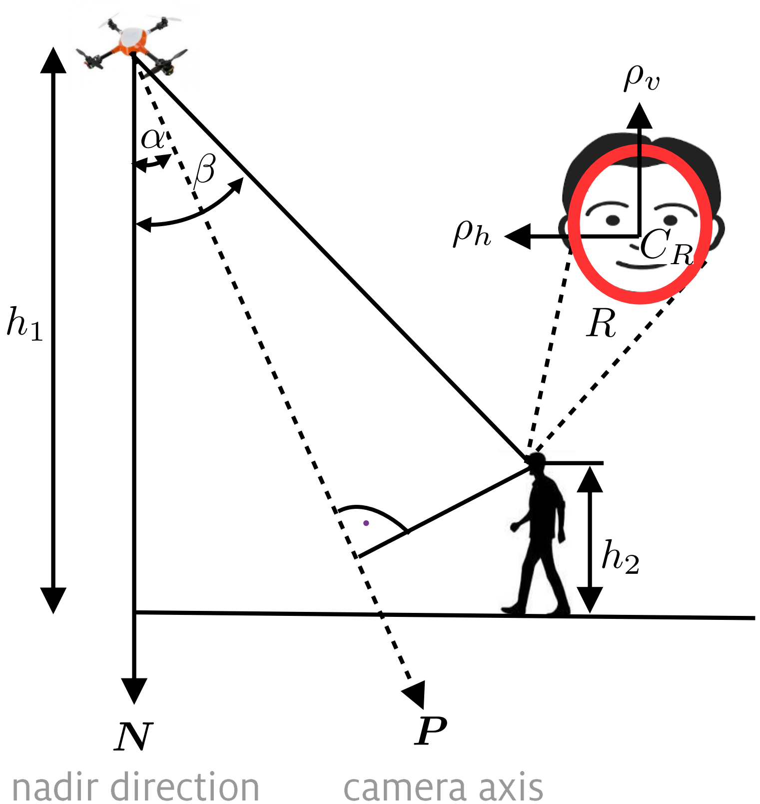

focal length

$$\rho_h = \frac{f\cos(\beta)}{p_h(h_1 - h_2)}$$

horizontal pixel size

$$\rho_v \approx \frac{f\cos(\beta)\sin(\beta)}{p_v(h_1 - h_2)}$$

vertical pixel size

Privacy Design Space

- Region protected (ω_R=0), if horizontal or vertical density is below threshold

$$\omega_R = \begin{cases} 1 & if \quad \rho_h > \rho_h^0 \quad \& \quad \rho_v > \rho_v^0 \\ 0 & otherwise \end{cases}$$

- Pixel density values for different heights (3-150 m), focal lengths (10-200 mm) and viewing angles (0-90 degrees)
  - For Canon EOS 5 MkII camera

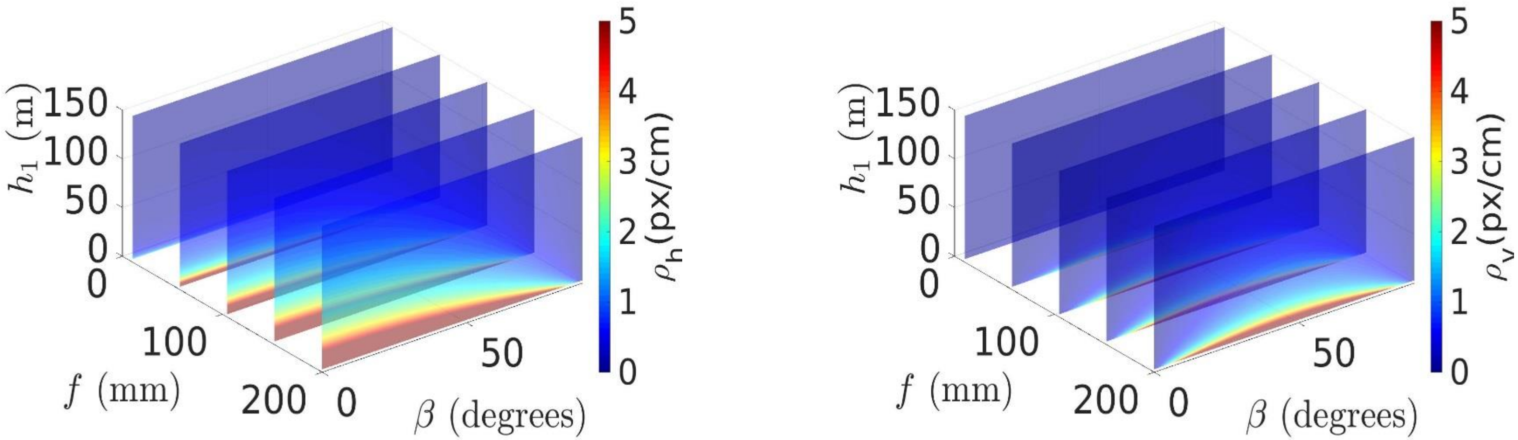

Privacy Design Space

- Region protected (ω_R=o), if horizontal or vertical density is below threshold

$$\omega_R = \begin{cases} 1 & if \quad \rho_h > \rho_h^0 \quad \& \quad \rho_v > \rho_v^0 \\ 0 & otherwise \end{cases}$$

- Separation between privacy sensitive and inherently protected space
  - For given threshold values (shown for $\rho_h^0 = \rho_v^0 = 1$ px/cm)

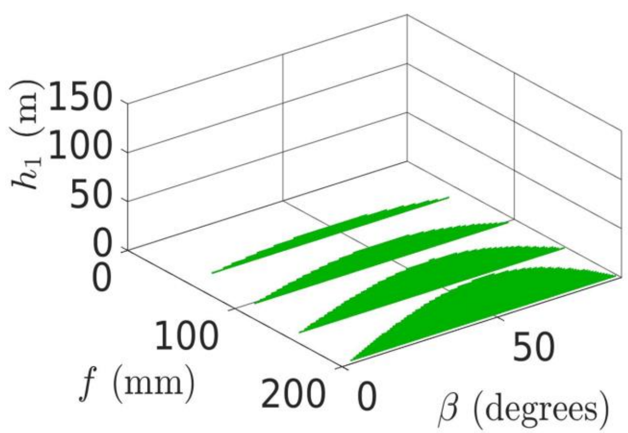

Adaptive Privacy Filter

- Configure filter G so that privacy protection is increased while fidelity is maintained

$$I_t^p = \mathcal{G}(I_t, R, \mu)$$

- filter strength
- face region
- unprotected frame at time t
- filter operator

- Determine filter strength μ such that the pixel resolution in the protected image is just below the threshold

Gaussian Blur as Privacy Filter

- Approximated anisotropic Gaussian kernel

$$g(v,h) = \frac{1}{2\pi\sigma_v\sigma_h} e^{-\left(\frac{v^2}{2\sigma_v^2}+\frac{h^2}{2\sigma_h^2}\right)}$$

with

$$\sigma_i = \frac{3\rho_i}{\pi\rho_i^0} \quad where \quad i \in \{v,h\}$$

- Filtering with kernel size

$$\mu_i = 2\lceil 3\sigma_i \rceil + 1$$

useful information in I_t^p is reduced to the threshold ρ_i^o

Adaptive Gaussian Blur Example

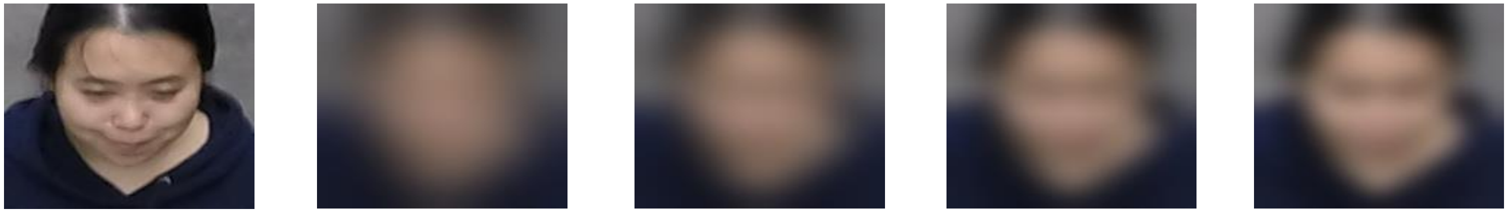

ρ: (5.03, 3.88) μ: (121, 105) μ: (99, 77) μ: (75, 57) μ: (59, 47)

Original Fixed Over Optimal Under

Gaussian blur for LDA face recognizer
Fixed: w.r.t. highest pixel density image in the data

Adaptive Gaussian Blur Example

Original	Fixed*	Over*	Optimal*	Under*
ρ: (5.03, 3.88)	μ: (121, 105)	μ: (99, 77)	μ: (75, 57)	μ: (59, 47)
ρ: (3.96, 2.87)	μ: (121, 105)	μ: (75, 57)	μ: (59, 43)	μ: (47, 35)
ρ: (3.06, 2.28)	μ: (121, 105)	μ: (61, 45)	μ: (45, 35)	μ: (37, 29)

*Gaussian Blur for LDA face recognizer
Fixed: w.r.t. highest pixel density image in the data

Measuring Privacy & Utility

- Subjective methods based on user studies
  - Predefined criteria
  - Crowd approaches
- Objective methods exploit CV algorithms
  - Detectors, classifiers, recognizers etc.
  - Metric based on performance difference between protected and unprotected input
  - Do not consider context or side-channel information

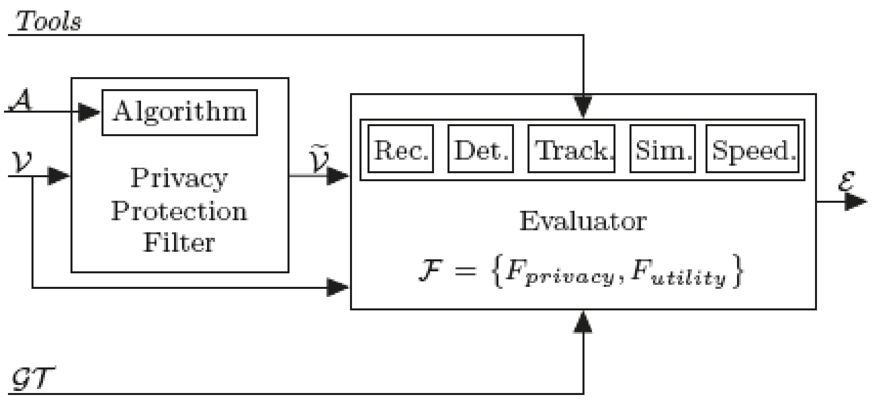

[Erdelyi, Winkler, Rinner. Privacy Protection vs. Utility in Visual Data: An Objective Evaluation Framework. Multimedia Tools and Applications, 2017.]

Experimental Setup

- Dataset from [Hsu, 2015]
  - Population size: 11 persons
  - Test data: 693 (63 x 11) images collected from 63 different positions.
  - Training data: 121 images i.e. 11 images of each person.

- Popular face recognizers for privacy measurement:
  - Linear Discriminant Analysis (LDA) [Belhumeur, 1997]
  - Local Binary Patterns Histograms (LBPH) [Ahonen, 2006]

- Fidelity measurement:
  - Peak Signal to Noise Ratio (PSNR)
  - Structural Similarity Index metric (SSIM) [Wang 2004]

Privacy of adaptively blurred Faces

LDA face recognizer

Thresholds: 0,6 & 0.4 px/cm

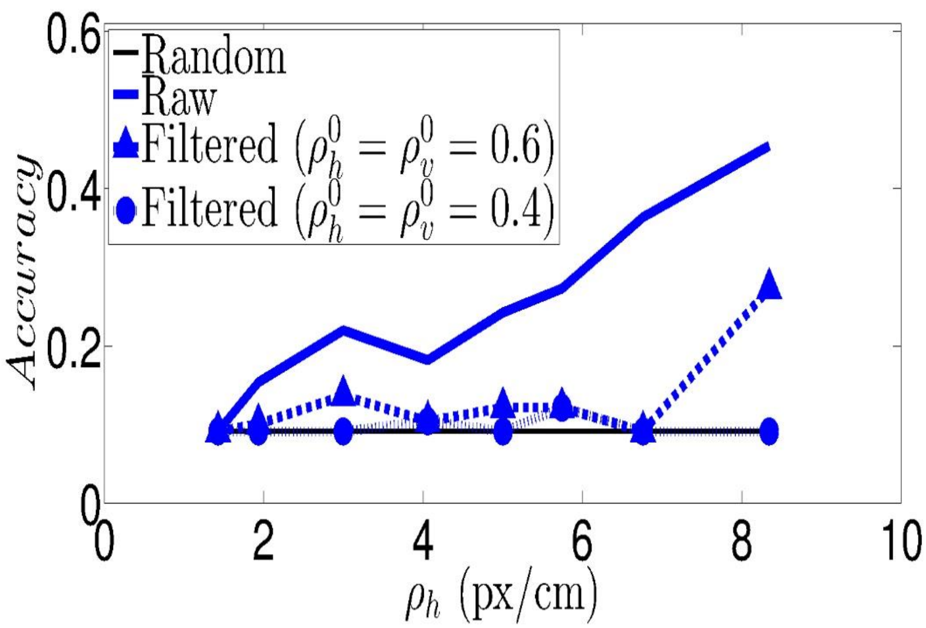

LBPH face recognizer

Thresholds: 0.4 & 0.2 px/cm

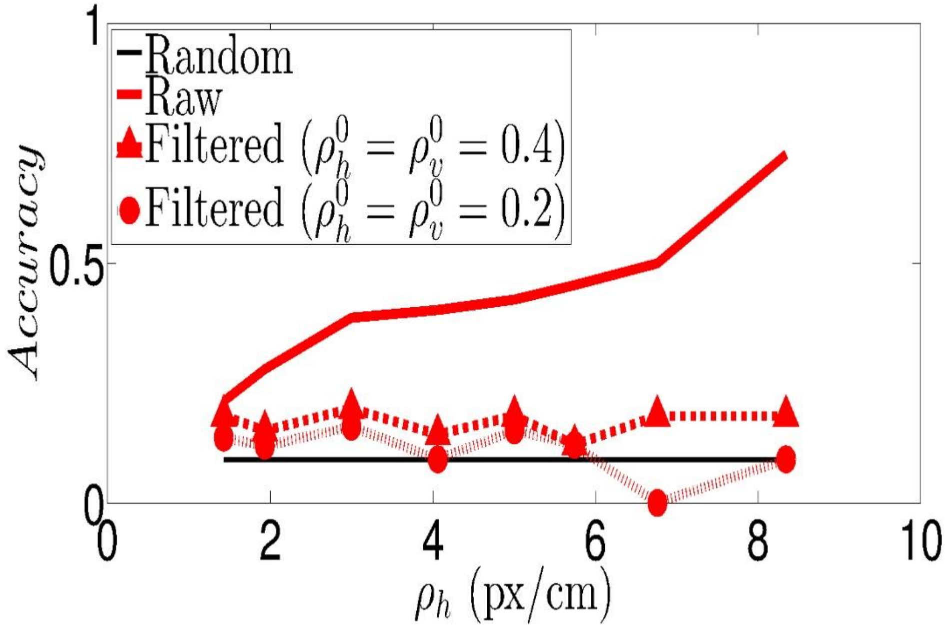

Fidelity Comparison

Peak Signal to Noise Ratio

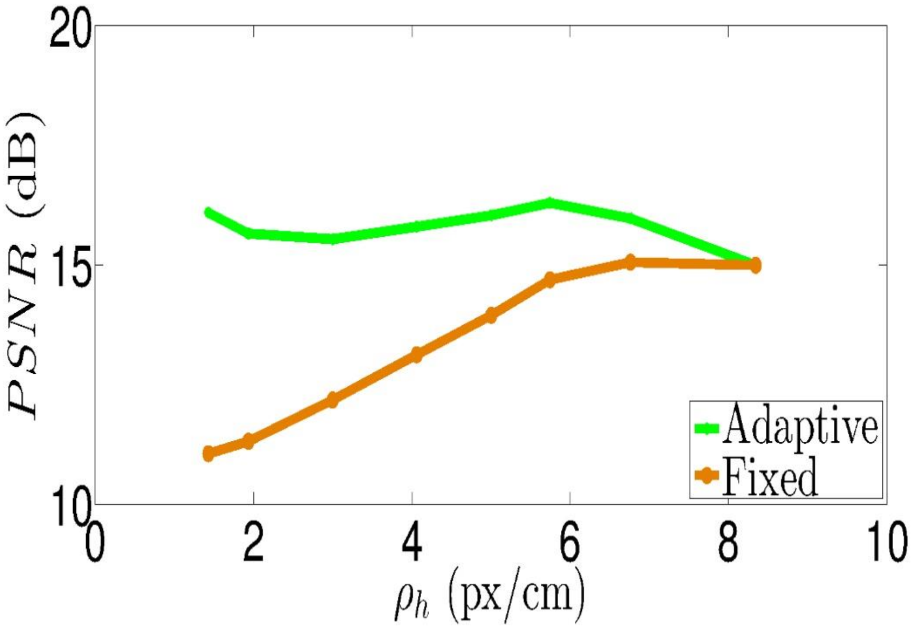

Structural Similarity Index

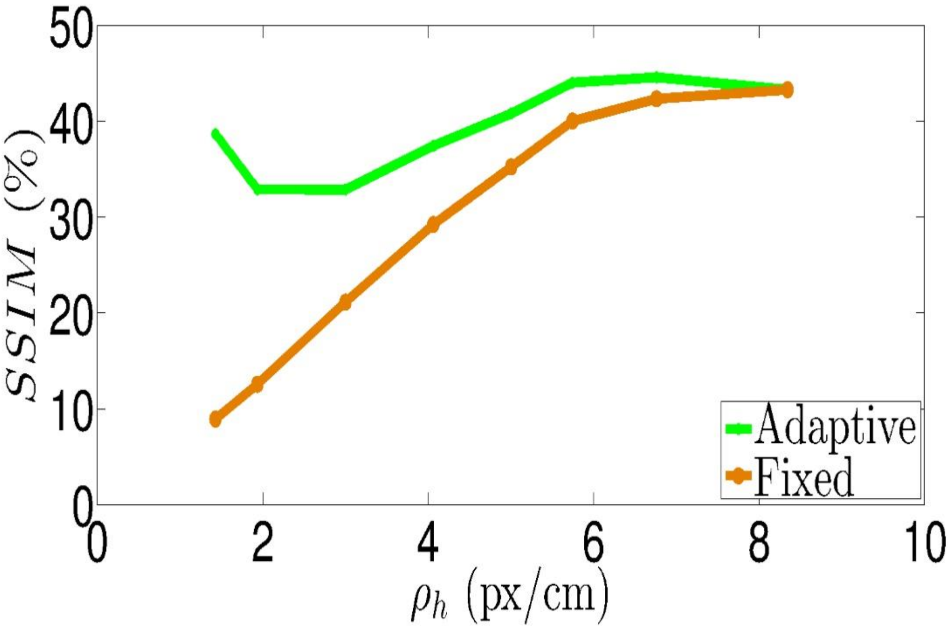

Privacy Attacks

Modelling privacy protection systems

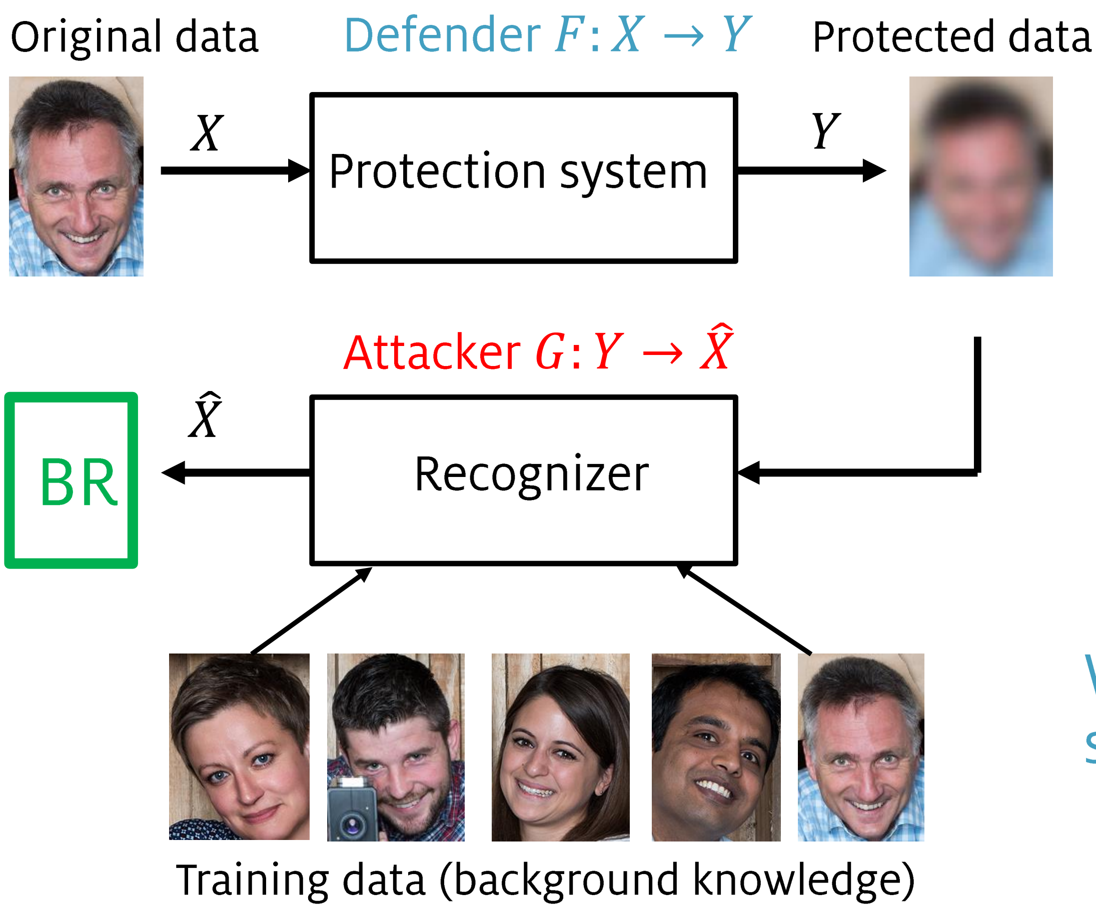

Distortion (utility)

$$D = \lambda(X; Y)$$

Information leakage (privacy protection)

$$L = \lambda(X; \hat{X})$$

What if the attacker has some knowledge about F?

Parrot Attacks

Attacker knows (learns) the protection filter (eg. blurring filter)

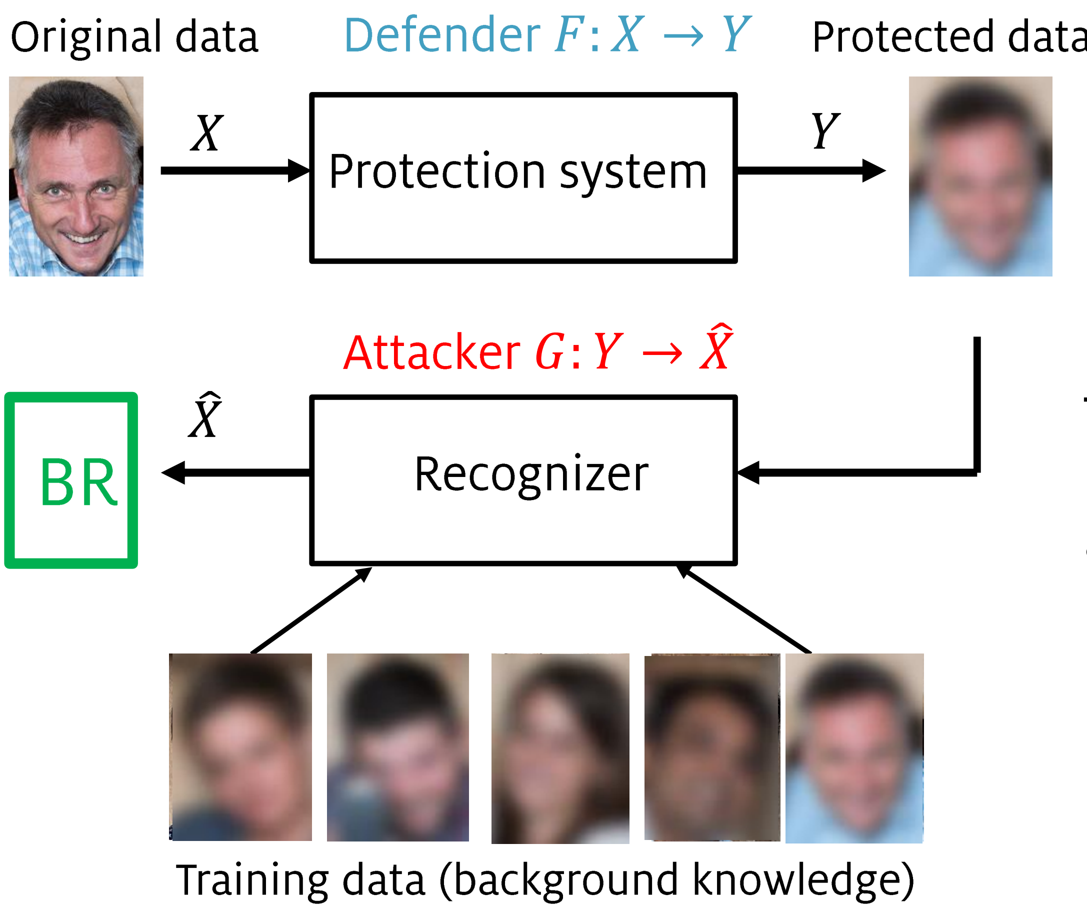

Train the recognizer in protected domain

- increase of information leakage

Reconstruction Attacks

Attacker knows (learns) how to reconstruct original data

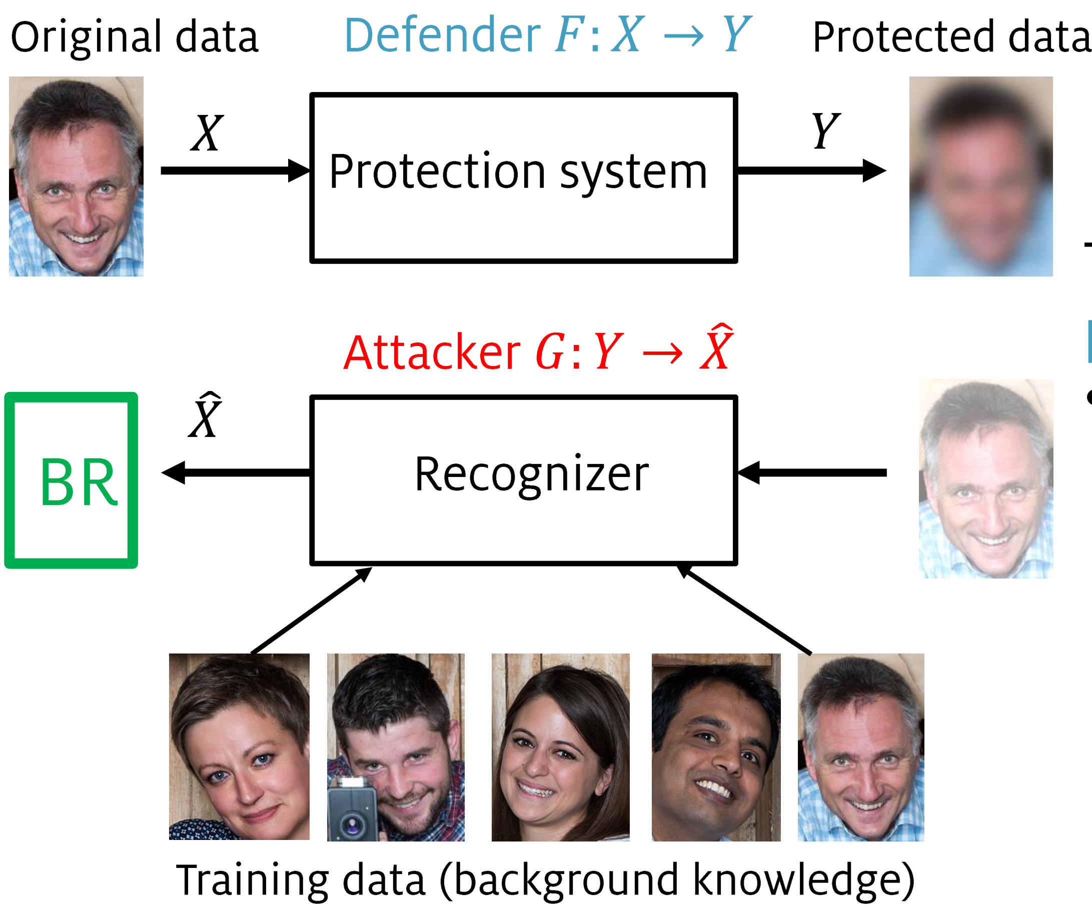

Train reconstruction of protected data

- Eg., superresolution

Adaptive Blurring with Spatial Hopping (AHGMM)

Pseudo-randomly change filter parameters for small patches to hinder

- Estimation of filter parameter
- Reconstruction of original image

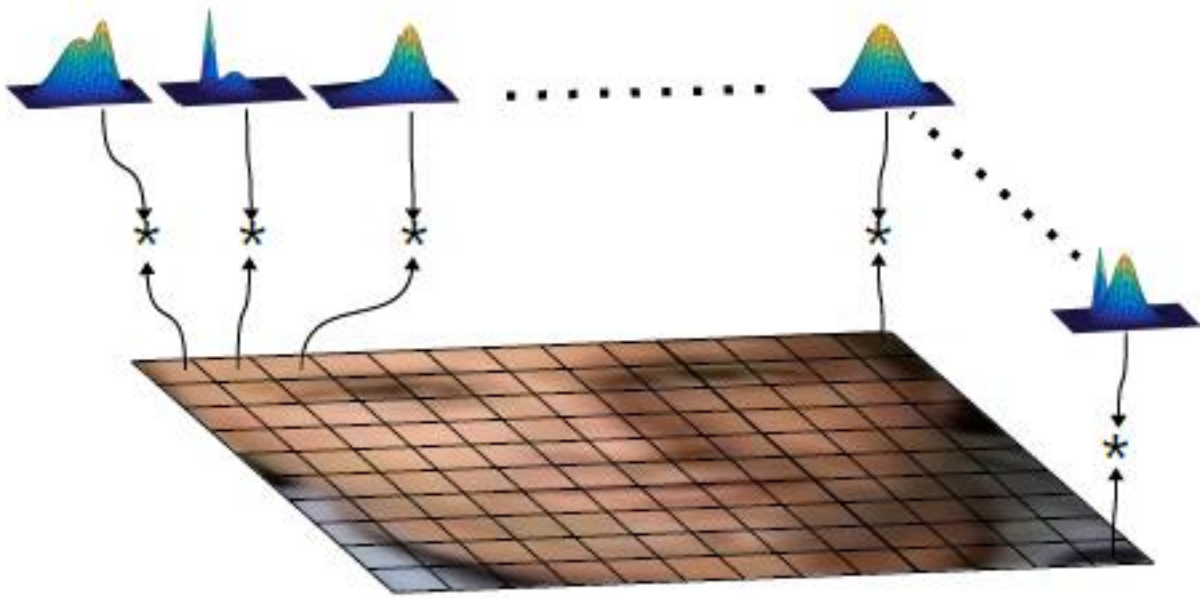

[Sawar, Rinner, Cavallaro. Adaptive Hopping Gaussian Mixture Model for Privacy-Preserving Aerial Photography. Under review 2017.]

Experimental Setup

- Labelled Faces in the Wild Dataset
  - Population size: 5749 persons
  - Expanded for aerial imagery
    40 instances for each person (variation in pitch angle and resolution)

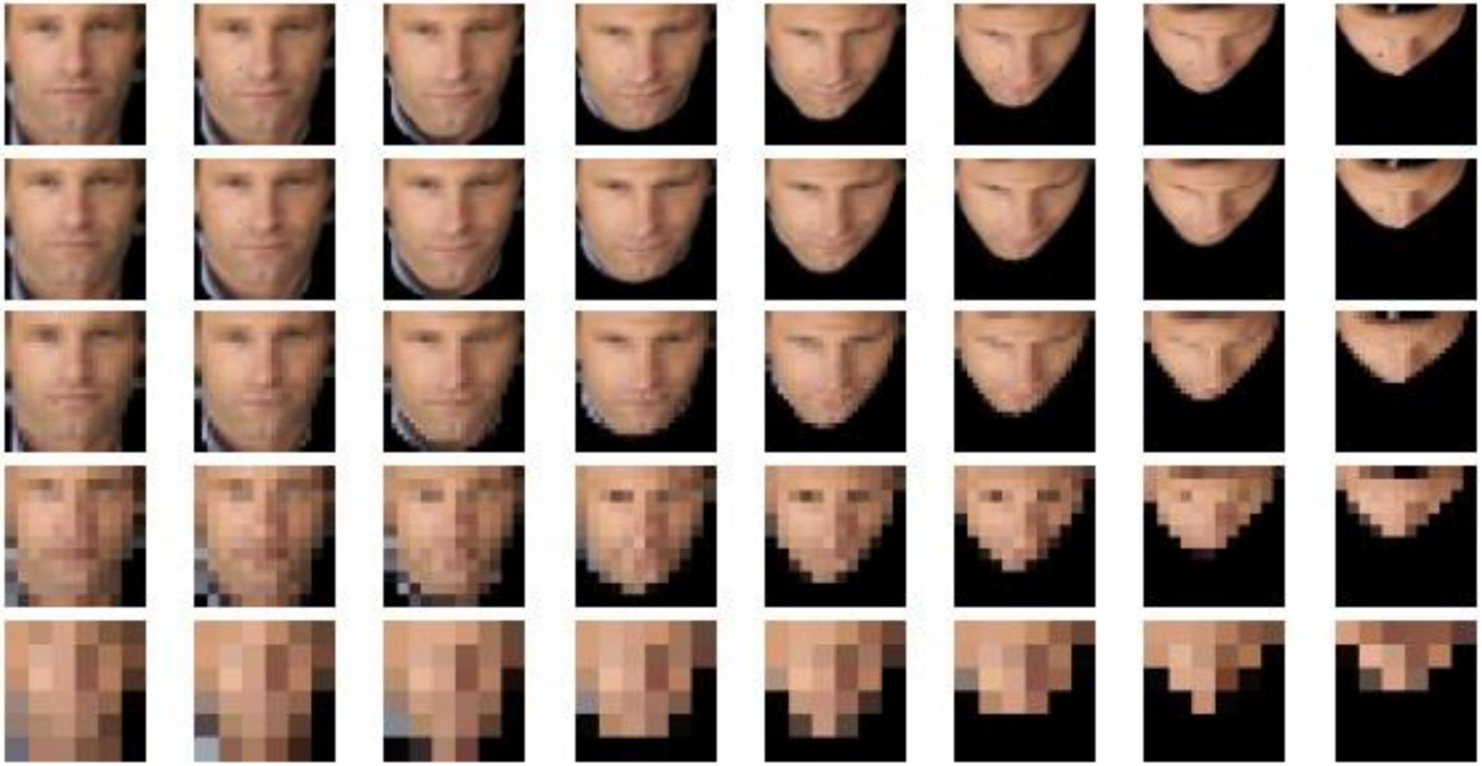

Experimental Setup (2)

- Privacy attack scenarios
  - Naïve: training with raw data
  - Parrot: training with AHGMM filtered data (3 variants)
  - Pitch angle is known by attacker as background
  - Tested with 380000 face images in total

- OpenFace recognizer for privacy measurement:
  - Verificiation test (600 persons with 10x cross validation)

- Fidelity measurement:
  - Peak Signal to Noise Ratio (PSNR)
  - Structural Similarity Index metric (SSIM) [Wang 2004]

Privacy Evaluation

- Comparison with 3 state-of-the-art privacy filters (-AHGMM)
  - Charts: privacy level η vs. pich angle; rows: different filter thresholds

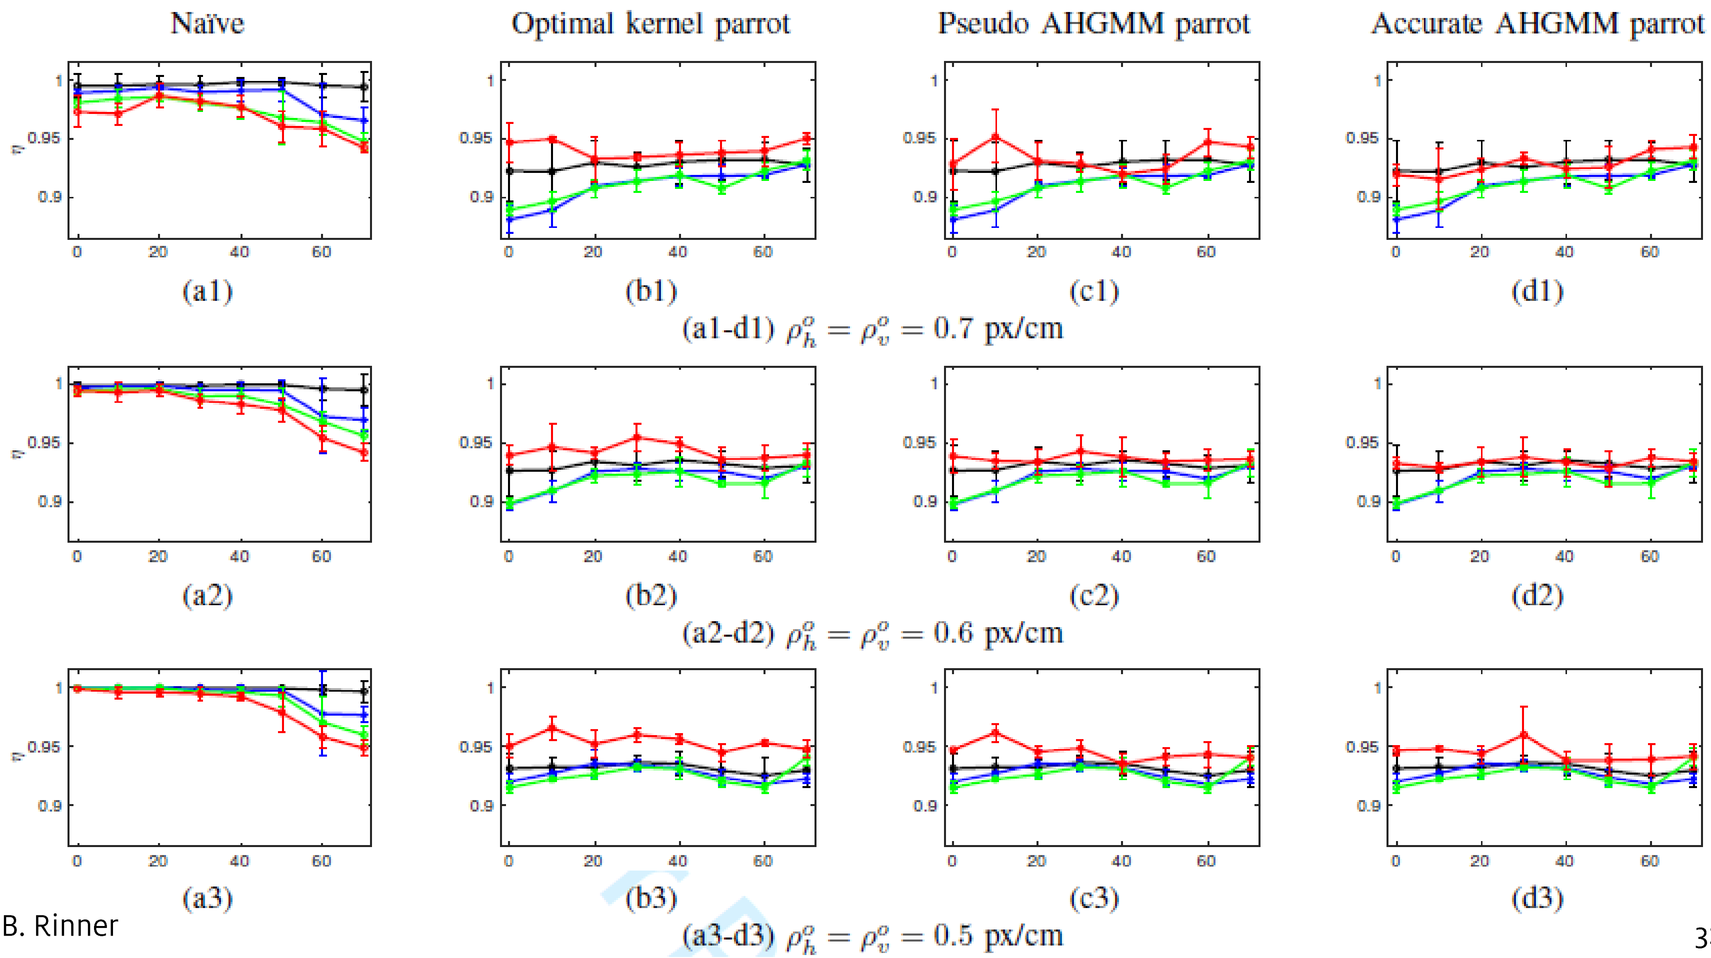

Privacy/Utility Tradeoff

- Privacy level vs. utility compared with 3 privacy filters (-AHGMM)

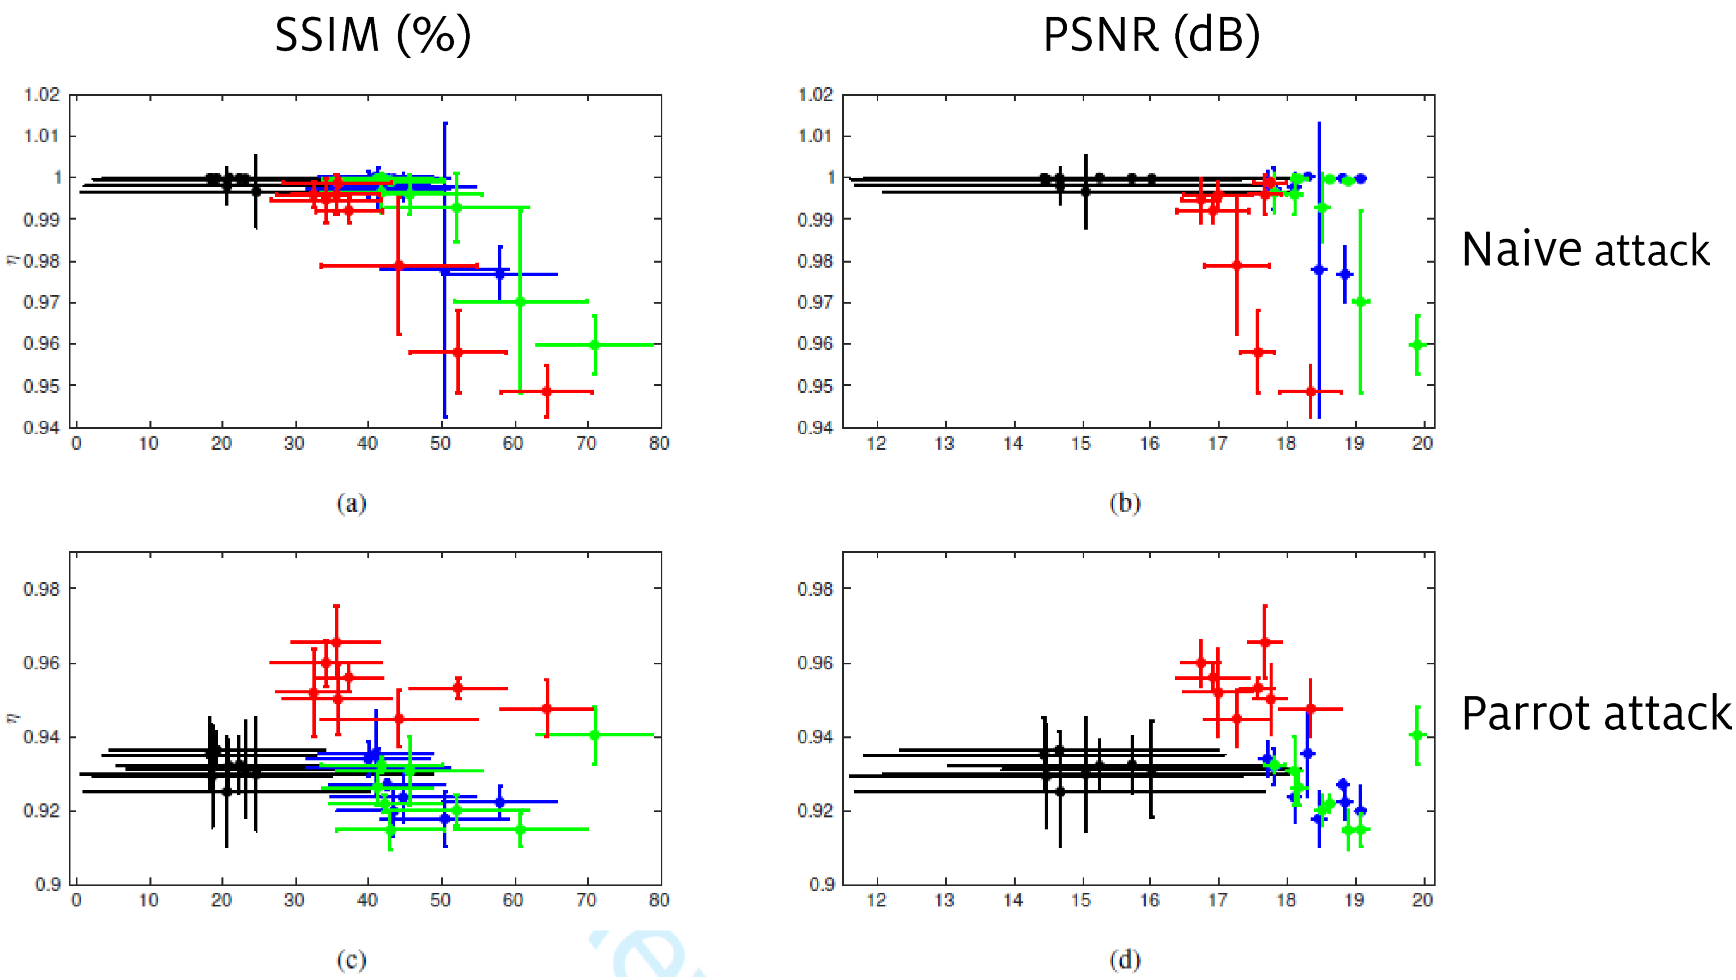

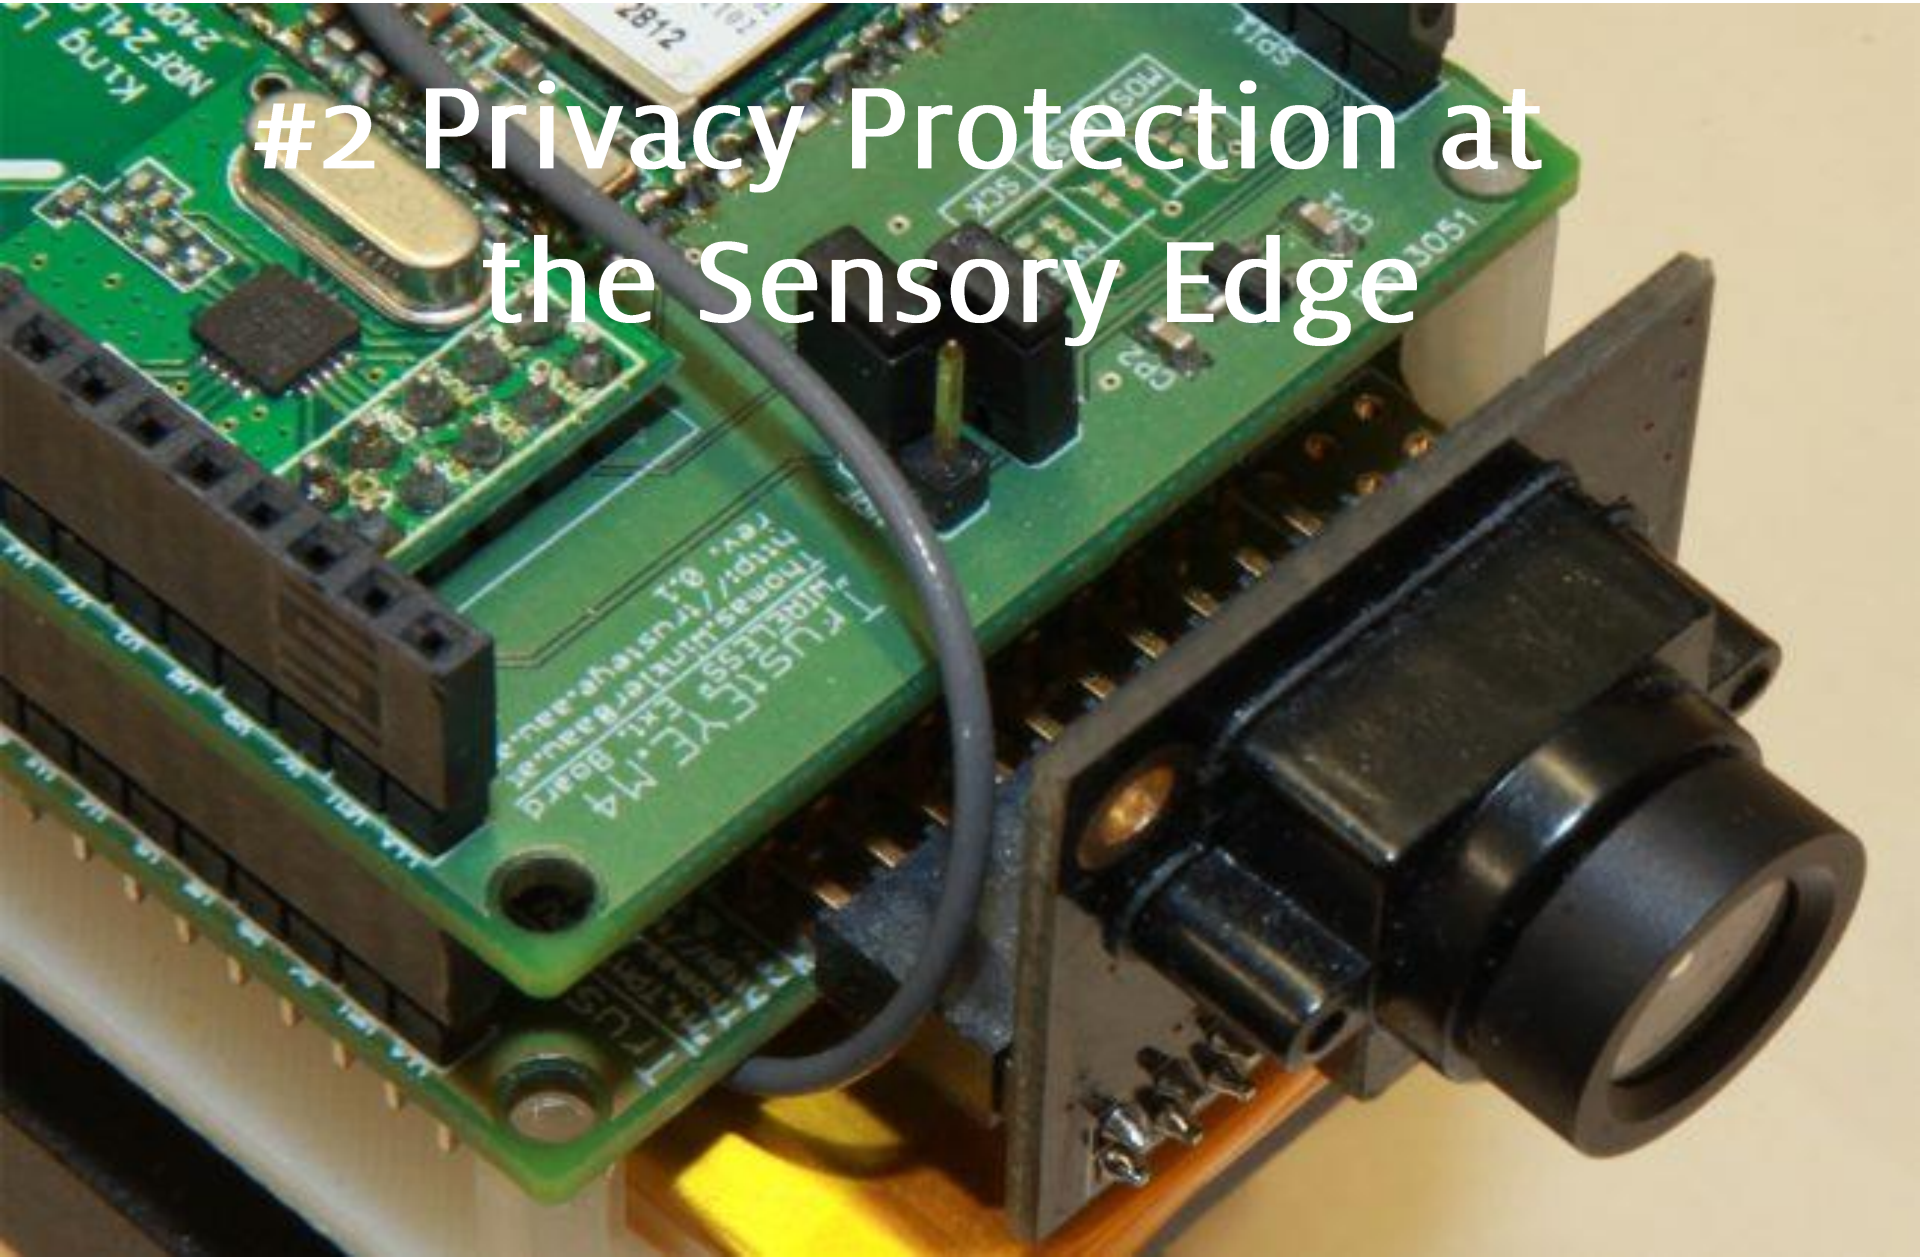

#2 Privacy Protection at the Sensory Edge

Onboard Protection on Camera

- Most cameras have no onboard protection, rarely software protection
- TrustCAM with TPM-based security features
  - Trusted boot
  - Integrity/authenticity by TPM-protected RSA keys
  - Freshness/timestamping for outgoing images
  - Multi-level encryption as privacy protection
  - Authentic user feedback

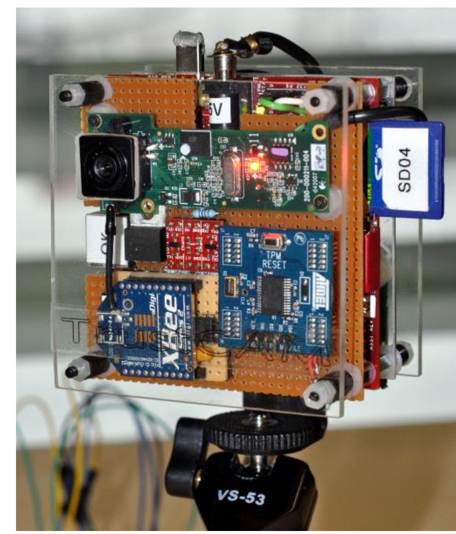

- Successful feasibility study, but security functionality was highly intertwined with application code

[Winkler, Rinner. Securing embedded smart cameras with trusted computing. EURASIP Journal on Wireless Communications and Networking, 2011]

Secure and Privacy-aware Camera

- Vision: TrustEYE - security and privacy protection as a feature of the image sensor instead of the camera

- Benefits:
  - Sensor delivers protected and pre-filtered data
  - Strong separation btw. trusted and untrusted domains
  - Camera software does no longer have to be trustworthy
  - Security can not be bypassed by application developers
  - TrustEYE is anchor for secure inter-camera collaboration

[Winkler, Erdelyi, Rinner. TrustEYE.M4: Protecting the Sensor - not the Camera. In Proc. AVSS 2014]

http://trusteye.aau.at/

TrustEYE Architecture

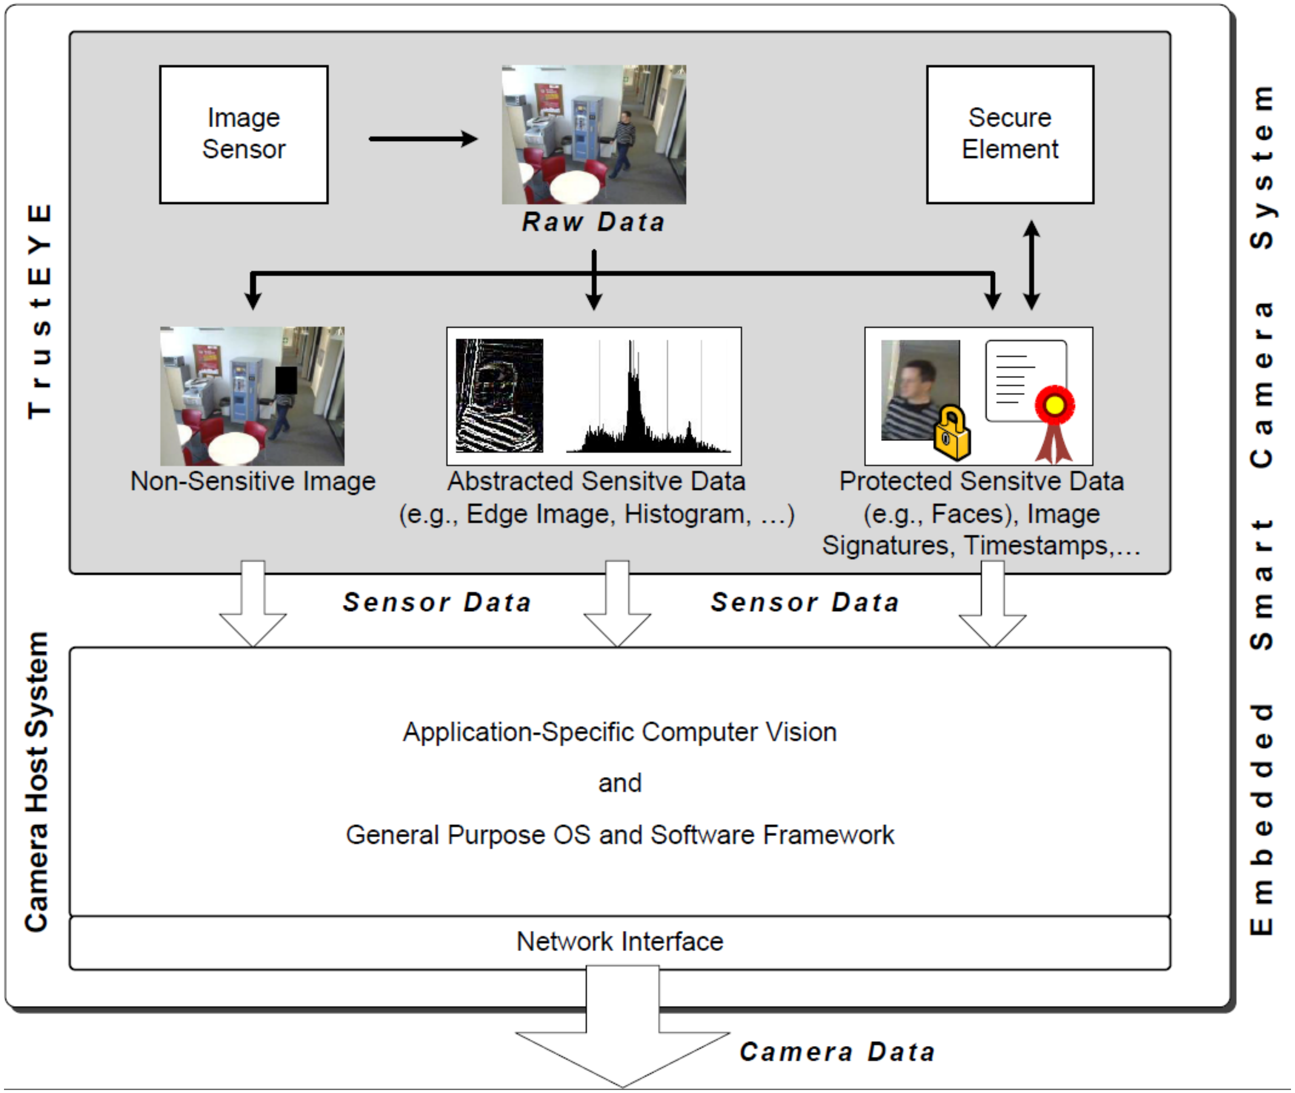

TrustEYE Platform

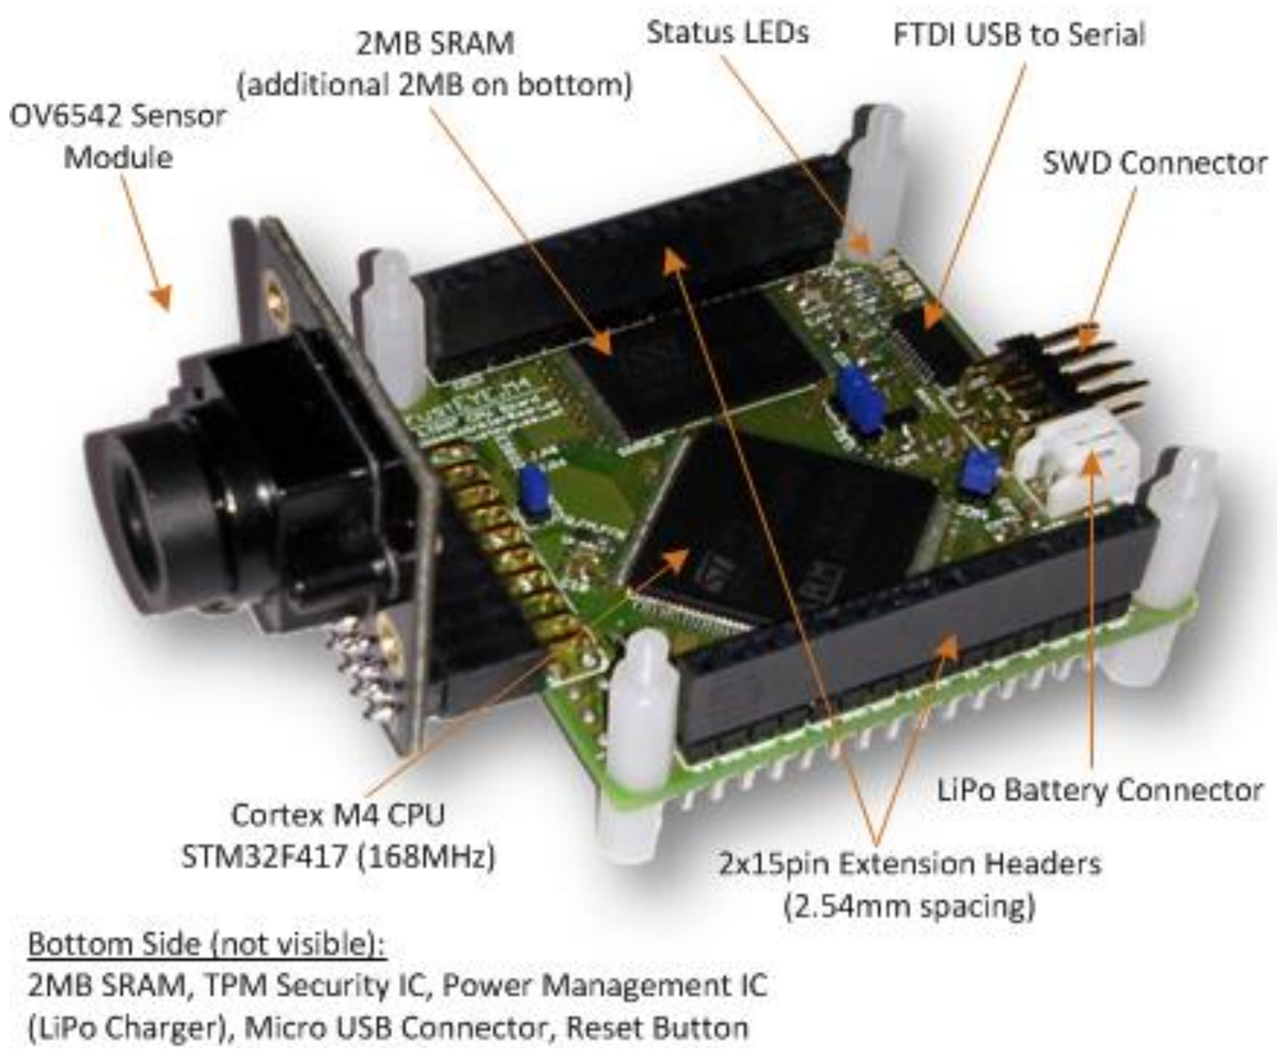

Cartooning Privacy Filter

- Abstract parts or entire image by blurring and color filtering
- Cartooning pipeline

ROI-based cartooning

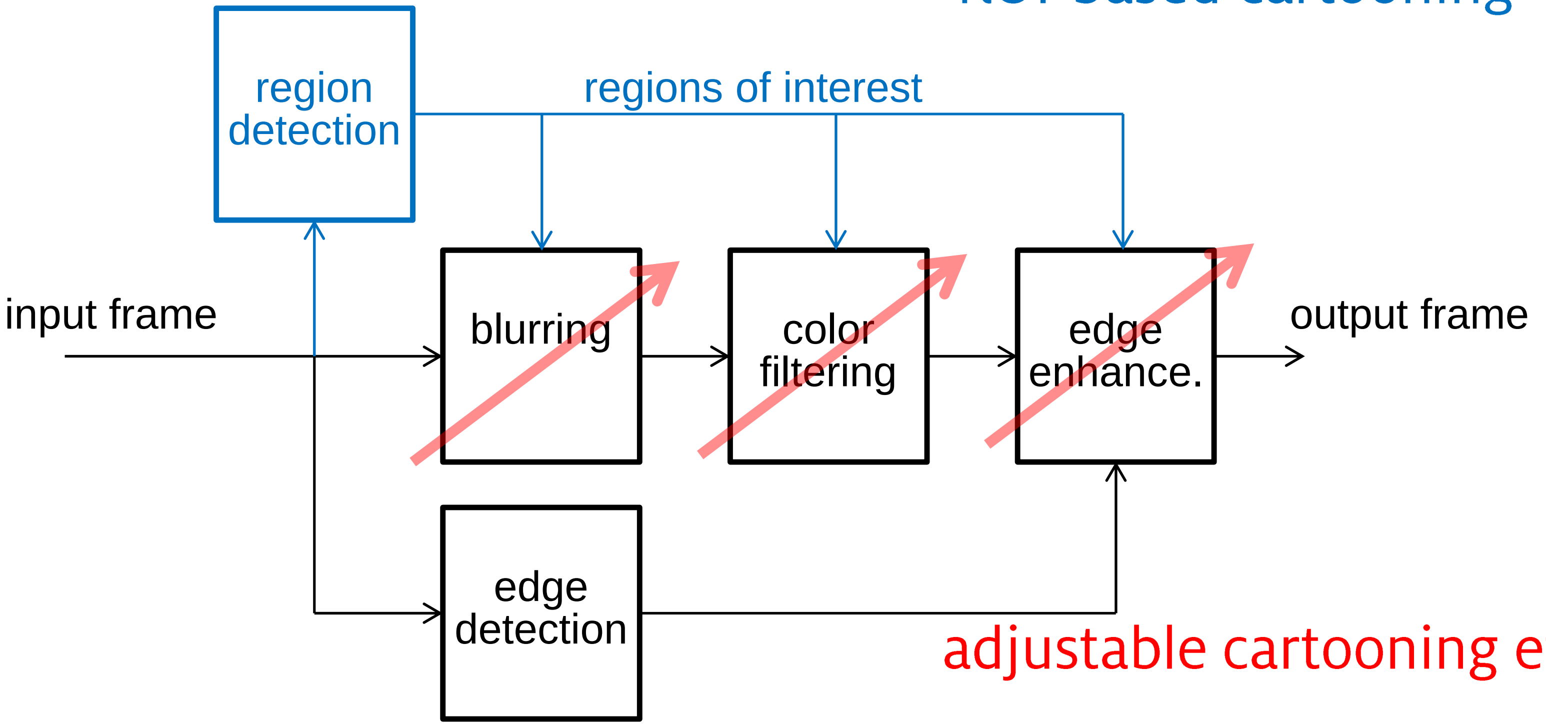

adjustable cartooning effect

[Erdelyi et al. Adaptive Cartooning for Privacy Protection in Camera Networks. In Proc. AVSS 2014.]

Adaptive Cartooning Filter


original


cartooning (small)


cartooning (std)


cartooning (strong)

[© Mediaeval Dataset]

TrustEYE Demo

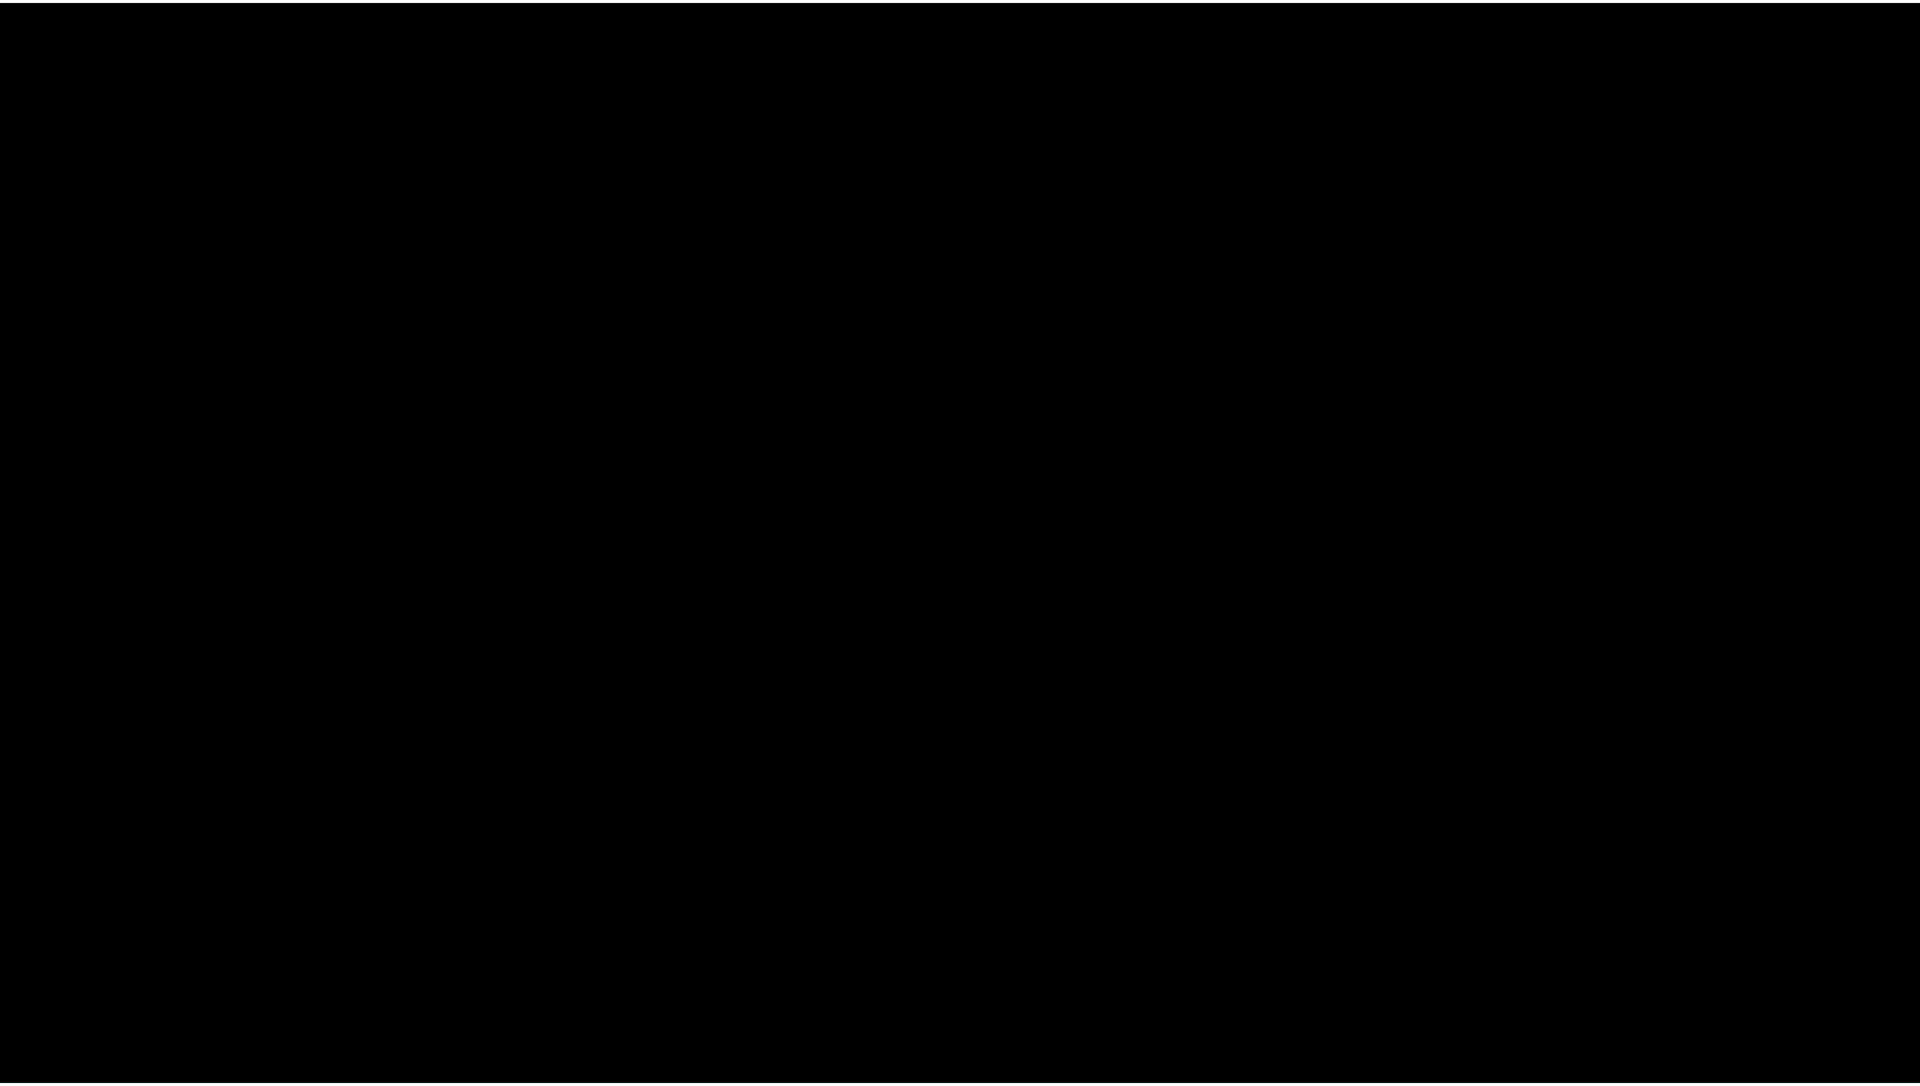

Trustworthy Sensing

- Exploit intrinsic hardware properties as key storage and avoid dedicated security chip
- Physically Unclonable Functions (PUFs) extracts fingerprints
  - Secure key generation & storage
  - Attestation of sensed data
  - Verified boot of sensor controller
  - Little system overhead

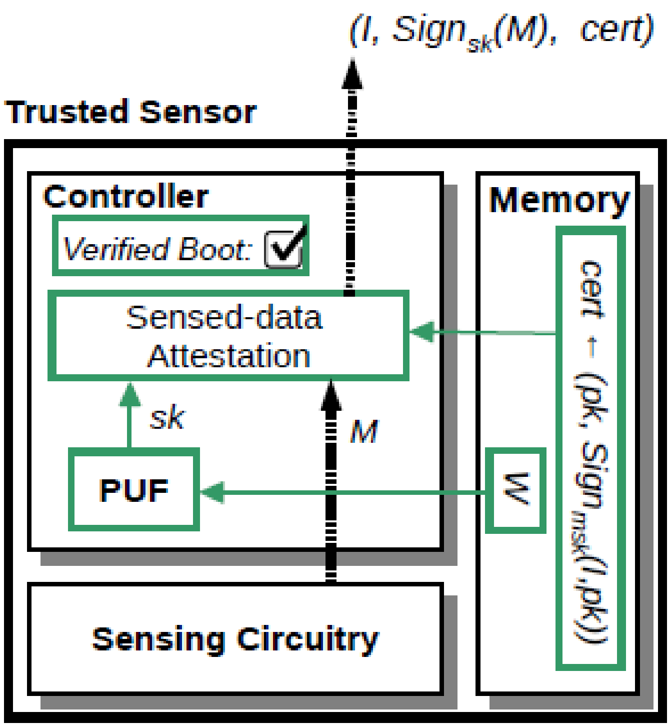

[Haider, Hoeberl, Rinner. Trusted Sensors for Participatory Sensing and IoT Applications based on Physically Unclonable Functions. In Proc. IoTPTS 2016]

Prototype SoC Implementation

- Xilinx Zynq 7010 (FGPA & dual Cortex ARM9 cores)
  - Ring-oscillator PUF with error correction to generate 128 bit keys
  - BLS signature scheme for data attestation
- Security overhead
  - 230 Bytes storage
  - 2210 logic cells
  - 6 ms for attestation

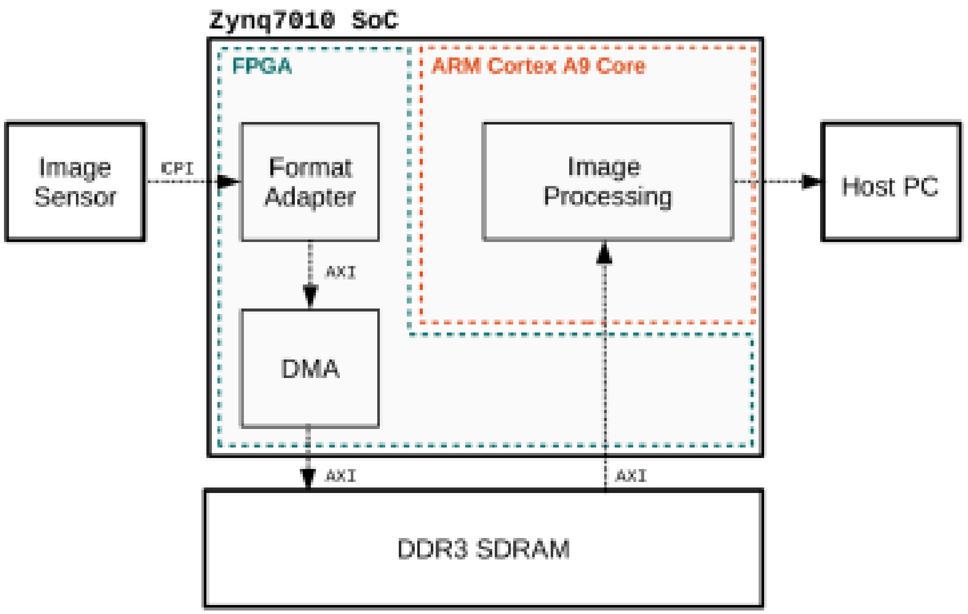

Conclusion

- Privacy protection important for commercial and private aerial imaging
- No single best protection method available. Tradeoff between protection, utility and resource usage
- Mostly image distortion used for protection, some can adapt the filter strength to scene
- Increase privacy awareness

Acknowledgements

Pervasive Computing group

Institute of Networked and Embedded Systems

http://nes.aau.at

http://bernhardrinner.com

Funding support

- KWF/FWF “Trustworthy Sensing and Cooperation in Visual Sensor Networks”
- FFG “Progressing towards Secure, Cooperating Smart Cameras”