Automated Image Forgery Detection through Classification of JPEG Ghosts

Dipl.-Inf. **Christian Riess** (joint work with Fabian Zach and Elli Angelopoulou)
August 30th, 2012
Pattern Recognition Lab (CS 5)
University of Erlangen-Nuremberg

FRIEDRICH-ALEXANDER
UNIVERSITÄT
ERLANGEN-NÜRNBERG

TECHNISCHE FAKULTÄT

Image forensics

- Is a picture authentic?
- Has a picture been taken with a particular camera?
- Emerging application in information security, signal processing and computer vision

Image forensics

- Is a picture authentic?
- Has a picture been taken with a particular camera?
- Emerging application in information security, signal processing and computer vision

Image forensics

- Is a picture authentic?
- Has a picture been taken with a particular camera?
- Emerging application in information security, signal processing and computer vision

Image forensics

- Is a picture authentic?
- Has a picture been taken with a particular camera?
- Emerging application in information security, signal processing and computer vision

Approaches to the detection of digital forgeries

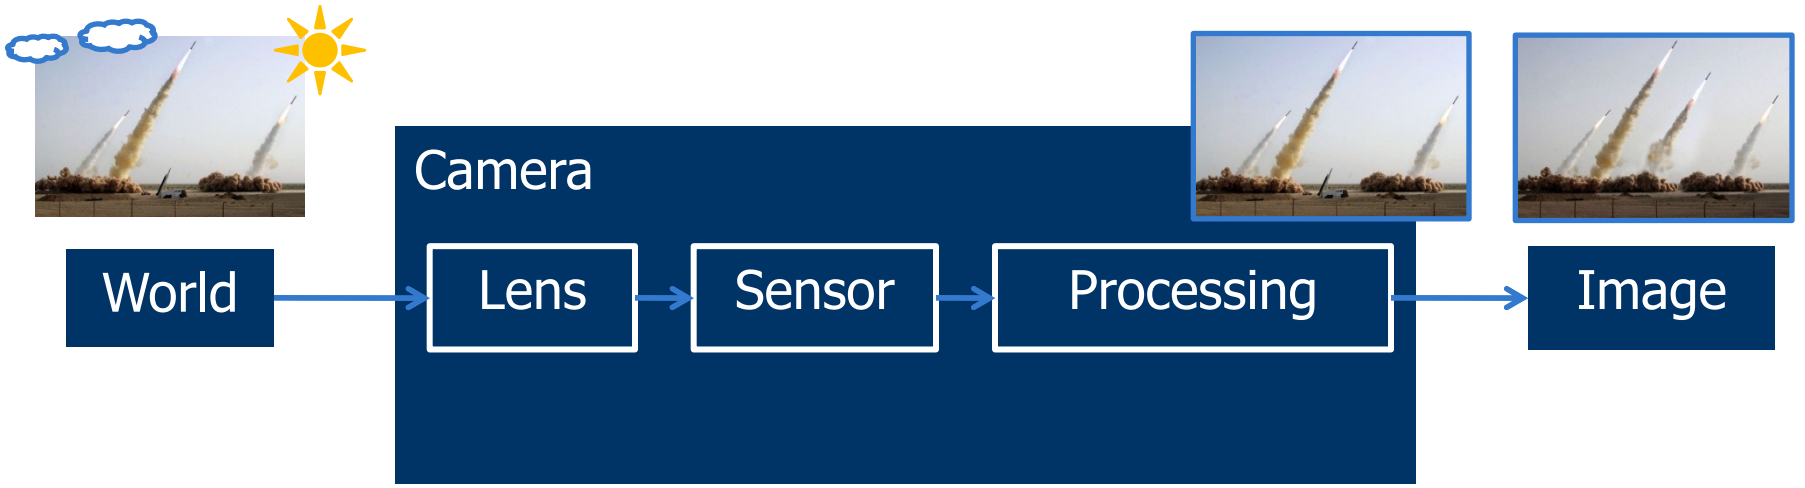

Approaches to the detection of digital forgeries

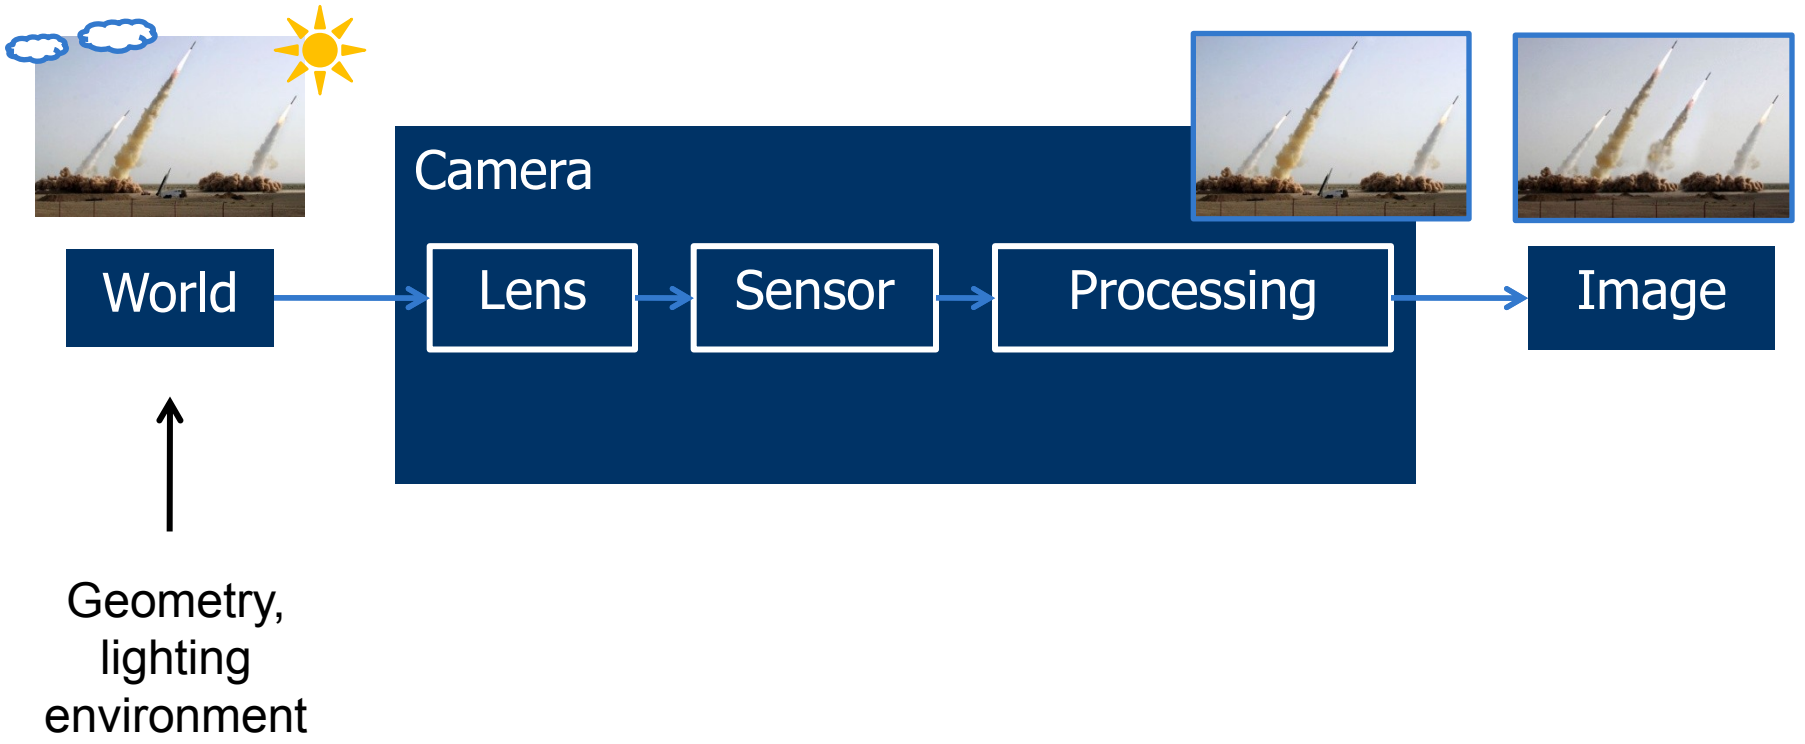

Approaches to the detection of digital forgeries

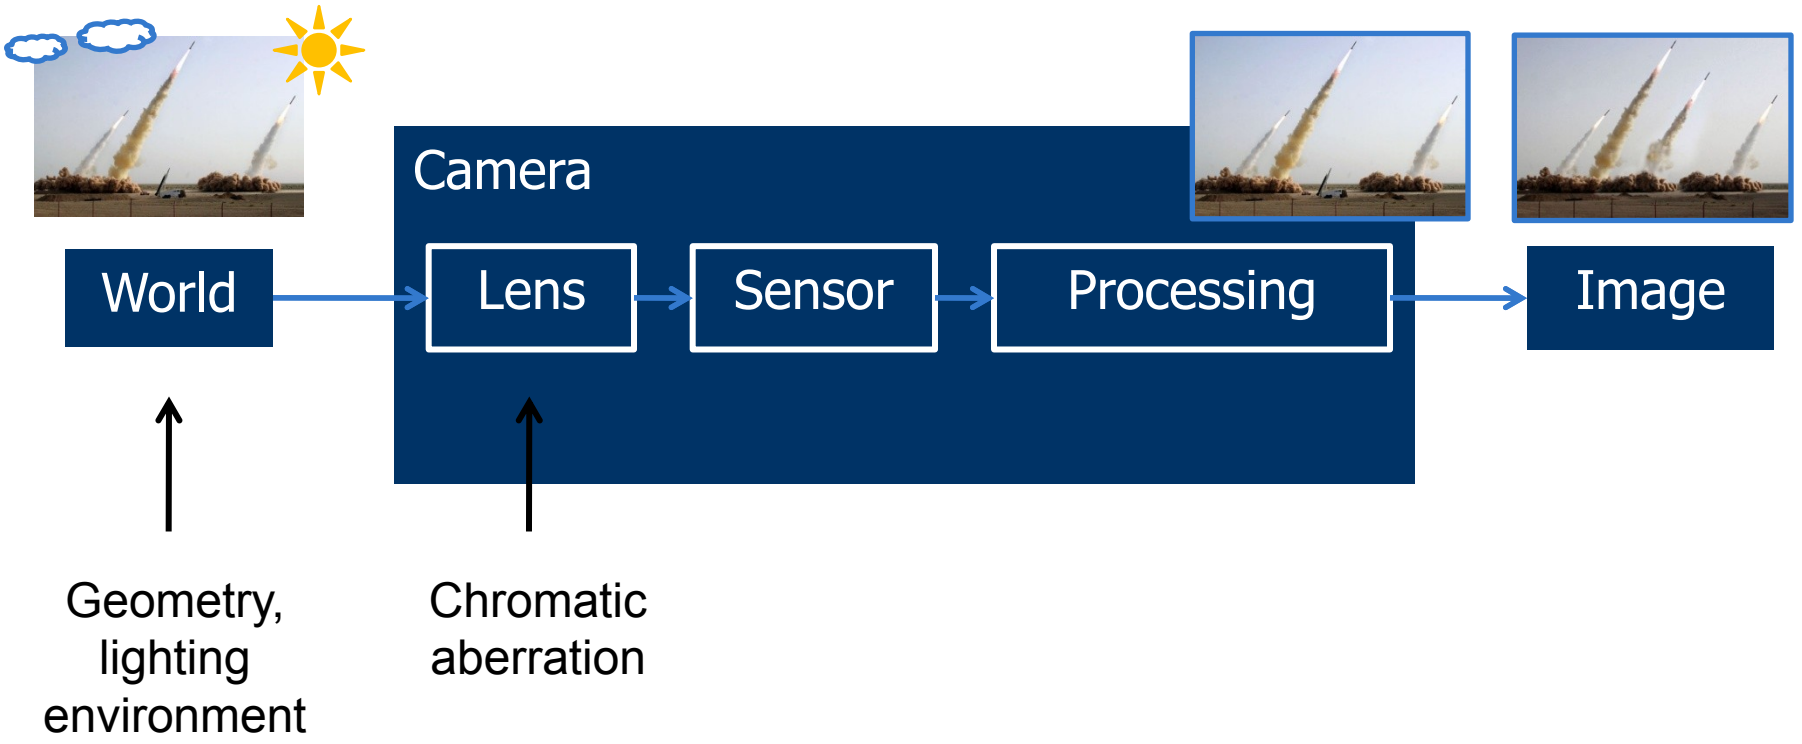

Approaches to the detection of digital forgeries

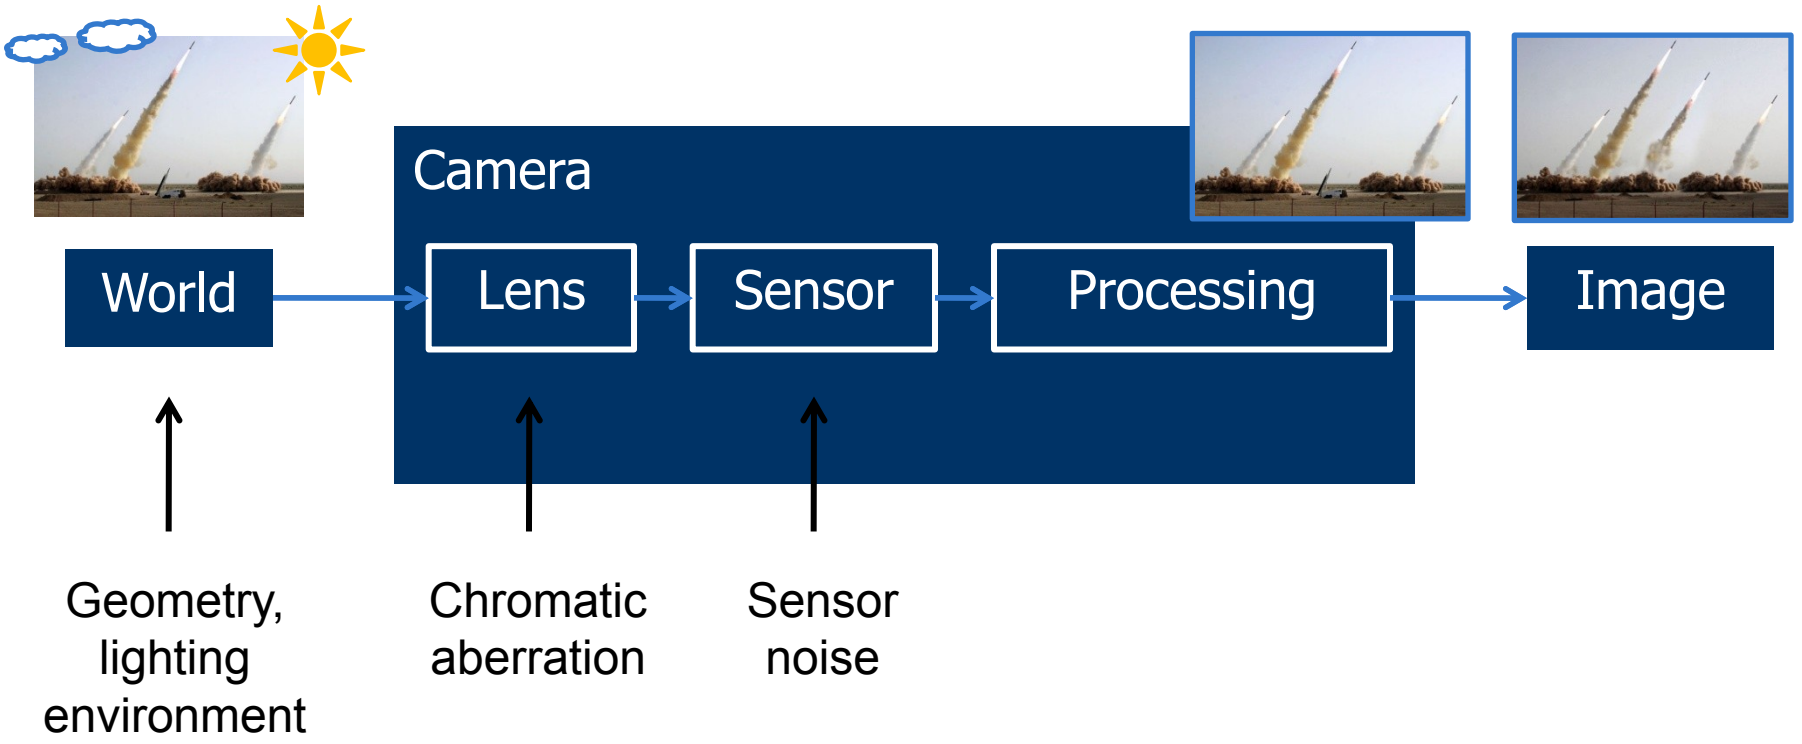

Approaches to the detection of digital forgeries

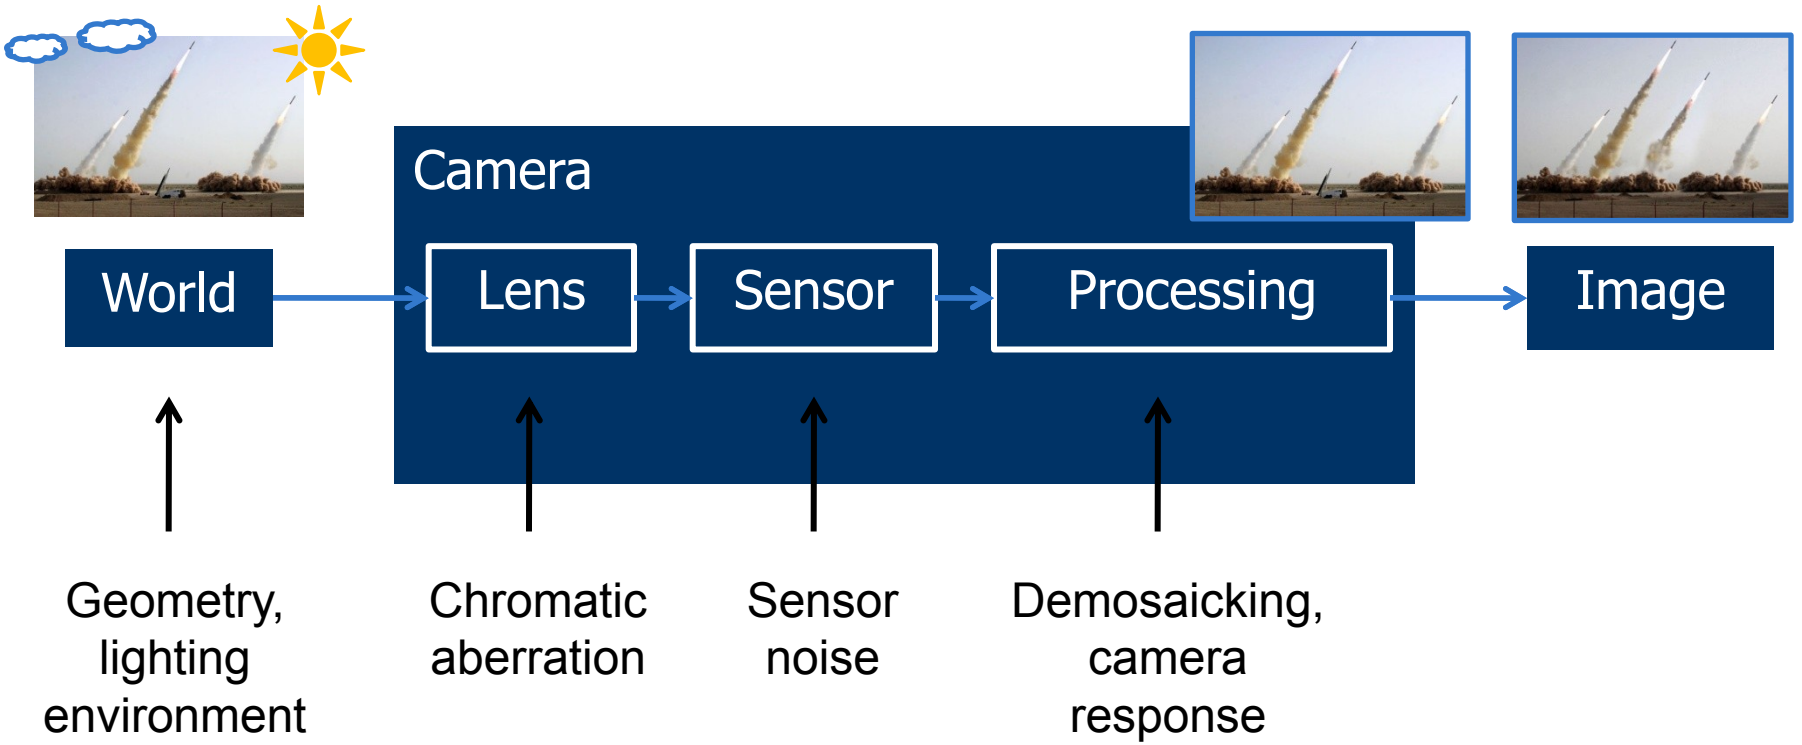

Approaches to the detection of digital forgeries

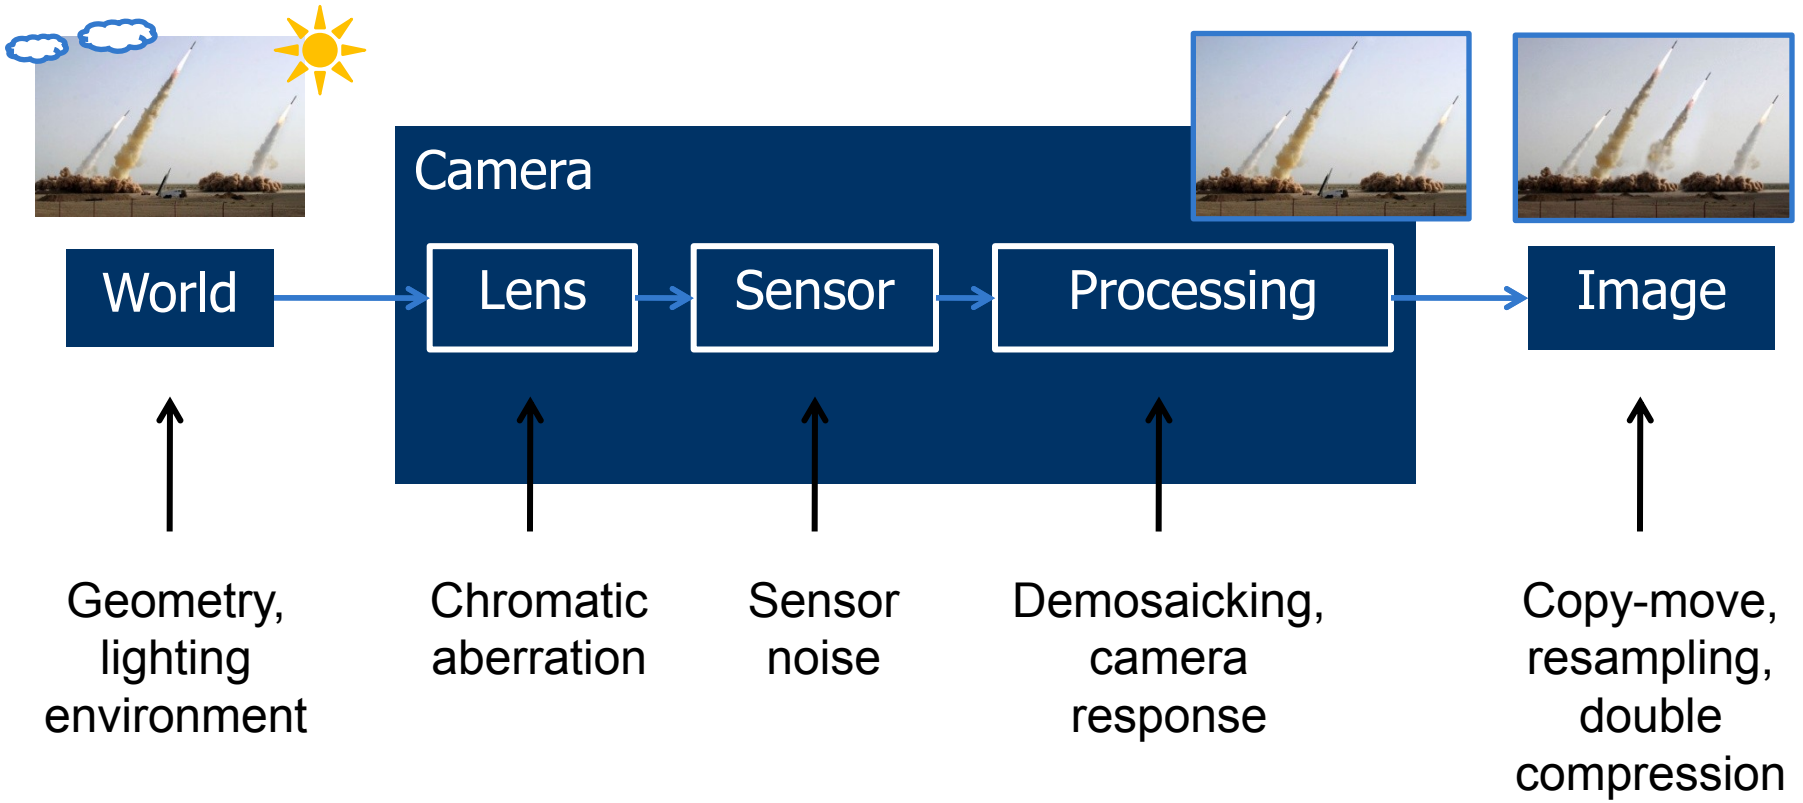

Approaches to the detection of digital forgeries

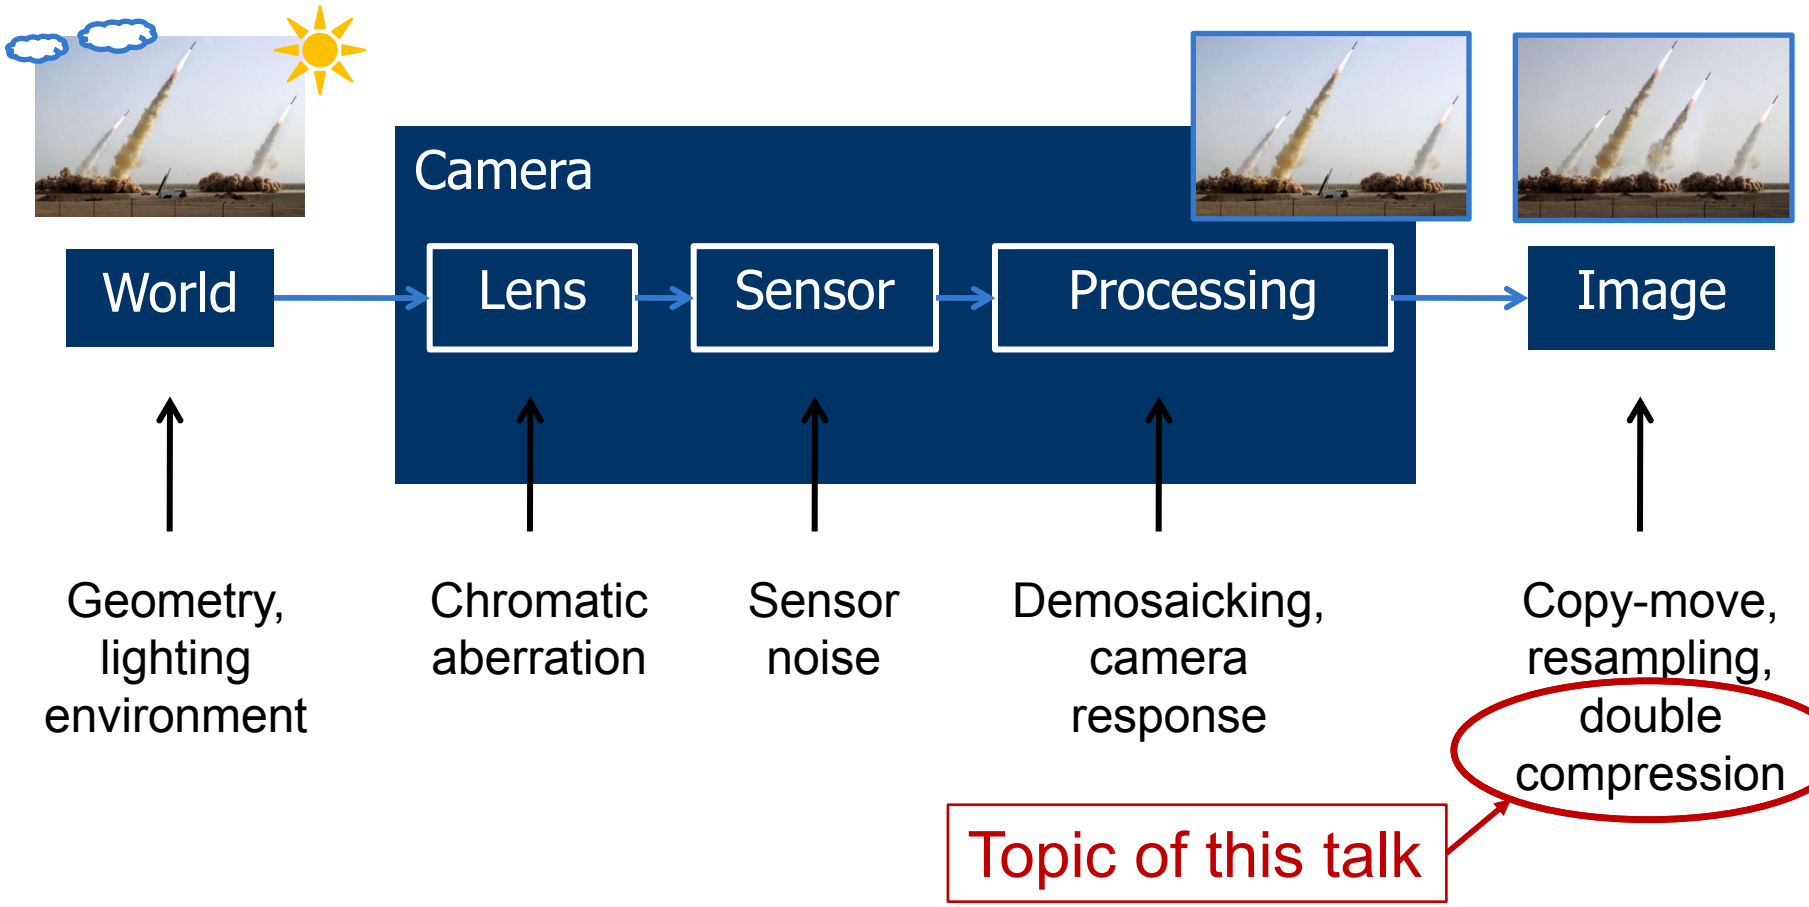

JPEG compression

- JPEG compression is block-based and lossy

Quality 100

Quality 20

- JPEG block grid: 8 by 8 pixels
- If recompressed, the new grid can be aligned/misaligned to previous grid

JPEG compression

- JPEG compression is block-based and lossy

Quality 100

Quality 20

- JPEG block grid: 8 by 8 pixels
- If recompressed, the new grid can be aligned/misaligned to previous grid

- This work: Classification of single- vs double-compressed JPEG blocks

JPEG artifacts as manipulation cues

- JPEG is very popular → “bad men” use JPEG, too
- If an image is recompressed, the statistics of JPEG artifacts change
- Forensic scenario:
  - First JPEG compression: in camera
  - Second JPEG compression: e.g. in a postprocessing tool
  - If only **part of** an image is double-compressed? Maybe a manipulation!

JPEG artifacts as manipulation cues

- JPEG is very popular → “bad men” use JPEG, too
- If an image is recompressed, the statistics of JPEG artifacts change
- Forensic scenario:
  - First JPEG compression: in camera
  - Second JPEG compression: e.g. in a postprocessing tool
  - If only **part of** an image is double-compressed? Maybe a manipulation!

Forensic question: “Are the JPEG artifacts consistent?”

Manipulation scenario: image splicing

Background, JPEG compressed

Foreground, JPEG compressed

Forgery, JPEG compressed

Manipulation scenario: image splicing

Background, JPEG compressed

Foreground, JPEG compressed

Forgery, JPEG compressed

Foreground is adjusted to background

(e.g., by cutting, cropping, painting, smearing)

→ this destroys original (primary) JPEG artifacts

Background: **double-compressed**

Cat (foreground): **single-compressed** **(due to editing)**

Detecting forgeries from JPEG inconsistencies

- Common approach:
  - Inspect coefficients of the discrete cosine transform
    (e.g. Lukas et al. 2003, He et al. 2005, Ye et al. 2007, Huang et al. 2010)

  - Tricky: case-by-case analysis, depending on compression parameters
    - e.g. primary (first) compression == secondary compression,
    - or secondary quantization factors multiples of primary factors

Detecting forgeries from JPEG inconsistencies

- Common approach:
  - Inspect coefficients of the discrete cosine transform (e.g. Lukas et al. 2003, He et al. 2005, Ye et al. 2007, Huang et al. 2010)
  - Tricky: case-by-case analysis, depending on compression parameters
    - e.g. primary (first) compression == secondary compression,
    - or secondary quantization factors multiples of primary factors

JPEG Ghosts [1] exploit a different cue

- Applicable if primary compression < secondary compression
- Simple to explain and implement
- But: manual browsing of 100's of images required ☹

[1] Hany Farid, "Exposing Digital Forgeries from JPEG Ghosts" in IEEE Transactions on Information Forensics and Security, vol. 1, no. 4, 2009, pp. 154-160.

The JPEG Ghost observation

- Read JPEG compression parameters from image header
- Recompress an image I with lower qualities q_1, q_2,…
- Look at difference images [redacted], [redacted] ,…
- …and the “ghost” appears

The JPEG Ghost observation

- Read JPEG compression parameters from image header
- Recompress an image I with lower qualities q_1, q_2,…
- Look at difference images [redacted], [redacted] ,…
- …and the “ghost” appears

The JPEG Ghost observation

- Read JPEG compression parameters from image header
- Recompress an image I with lower qualities q_1, q_2,…
- Look at difference images [redacted], [redacted] ,…
- …and the “ghost” appears

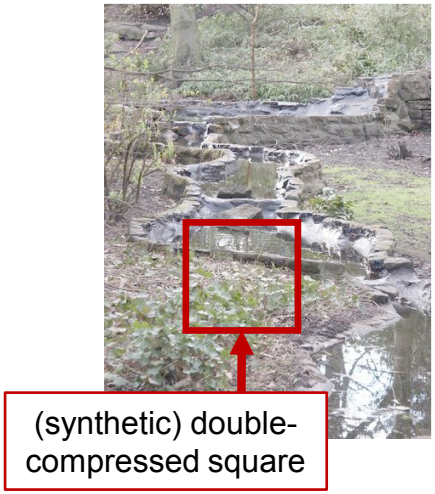

The JPEG Ghost observation

- Read JPEG compression parameters from image header
- Recompress an image I with lower qualities q_1, q_2,…
- Look at difference images [redacted], [redacted] ,…
- …and the “ghost” appears

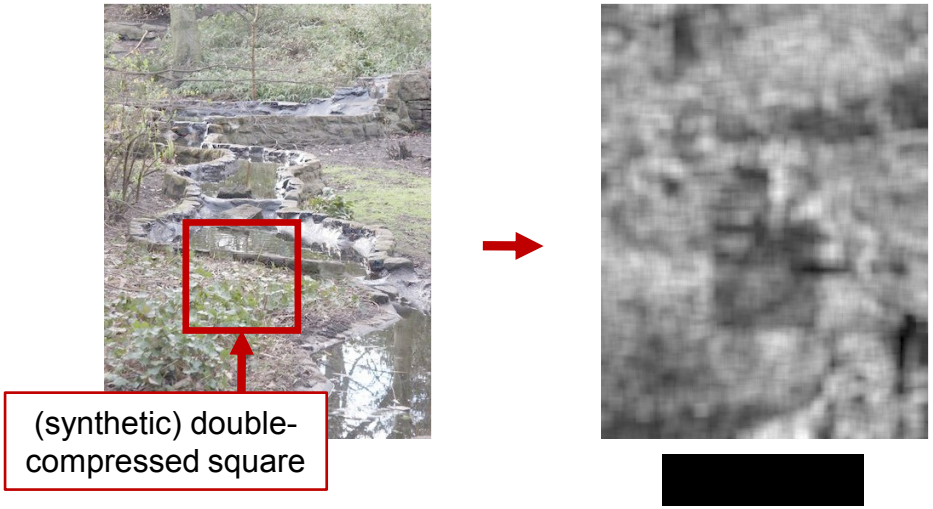

The JPEG Ghost observation

- Read JPEG compression parameters from image header
- Recompress an image I with lower qualities q_1, q_2,…
- Look at difference images [redacted], [redacted],…
- …and the “ghost” appears

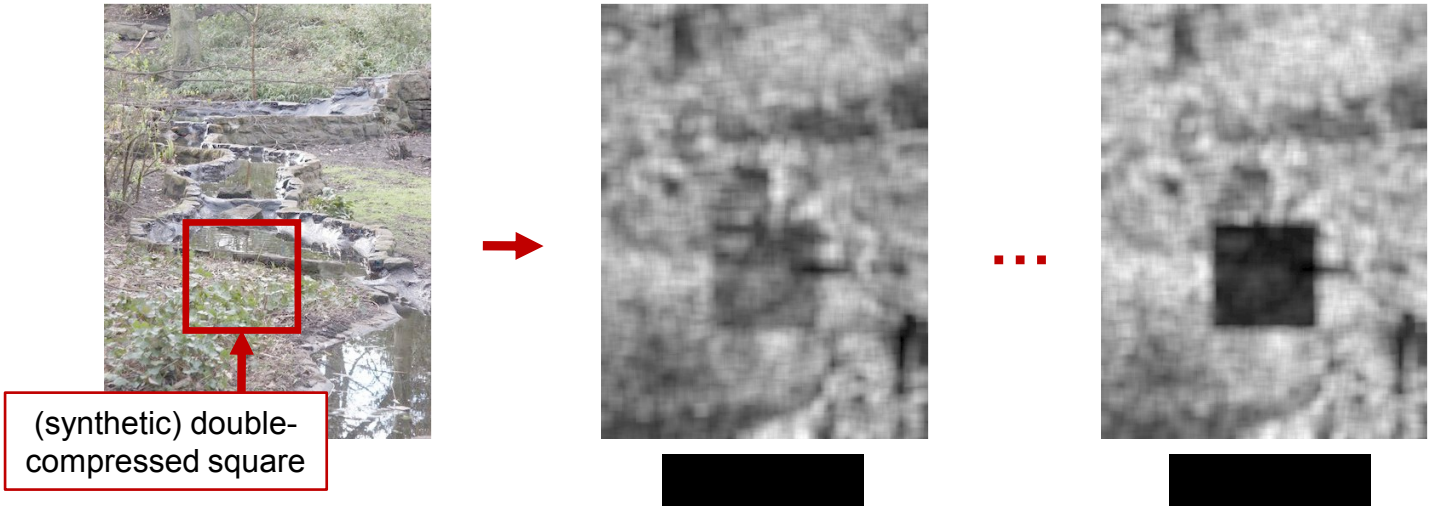

Finding Ghosts requires patience

Ghost hardly visible if

Finding Ghosts requires patience

Ghost hardly visible if

- block grid is not aligned

Finding Ghosts requires patience

Ghost hardly visible if

- block grid is not aligned **[try all 64 alignments!]**

Finding Ghosts requires patience

Ghost hardly visible if

- block grid is not aligned **[try all 64 alignments!]**
- first and second compression parameters of similar magnitude

Finding Ghosts requires patience

Ghost hardly visible if

- block grid is not aligned **[try all 64 alignments!]**
- first and second compression parameters of similar magnitude **[look sharper!]**

Finding Ghosts requires patience

Ghost hardly visible if

- block grid is not aligned **[try all 64 alignments!]**
- first and second compression parameters of similar magnitude **[look sharper!]**

→ up to 500 low-contrast b/w images.

Finding Ghosts requires patience

Ghost hardly visible if

- block grid is not aligned **[try all 64 alignments!]**
- first and second compression parameters of similar magnitude **[look sharper!]**

→ up to 500 low-contrast b/w images.

Contribution of this work:

Complete Automation of the JPEG Ghost Scheme

Extraction of Ghost data

- For each JPEG block
  - track [redacted], [redacted],… over recompression steps q_1, q_2,…
  - extract features,
  - classify the block.

Extraction of Ghost data

- For each JPEG block
  - track [redacted], [redacted],… over recompression steps q_1, q_2,…
  - extract features,
  - classify the block.

Extraction of Ghost data

- For each JPEG block
  - track [redacted], [redacted],… over recompression steps q_1, q_2,…
  - extract features,
  - classify the block.

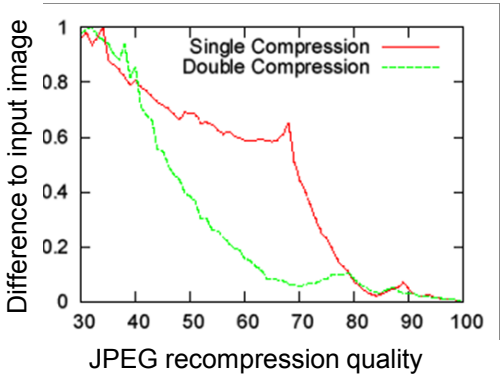

Extraction of Ghost data

- For each JPEG block
  - track [redacted], [redacted],… over recompression steps q_1, q_2,…
  - extract features,
  - classify the block.

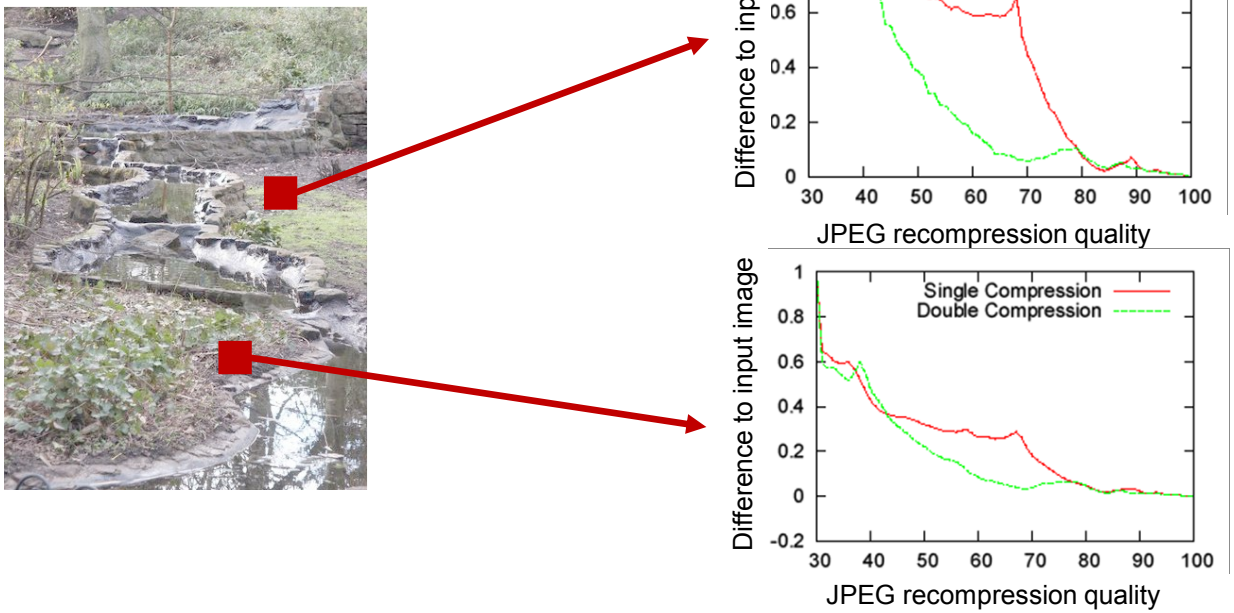

Feature extraction

- Weighted mean
- Median
- Slope of regression line
- Y-axis intercept of regression line
- Weighted sum of points below 0.5
- SSD between diagonal and curve

Single compression

Double compression

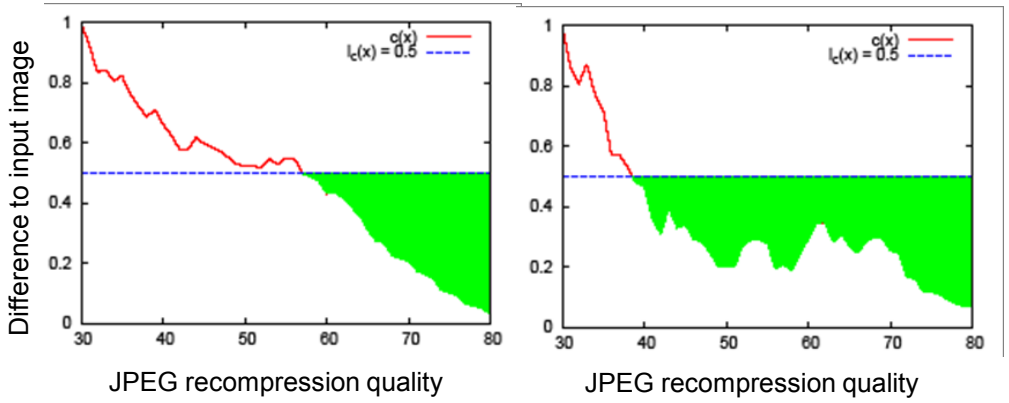

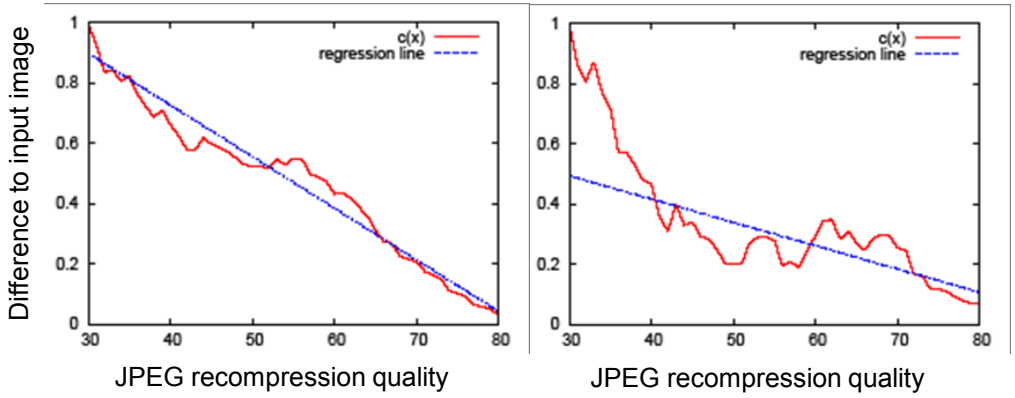

Feature histograms: double vs. single-compressed

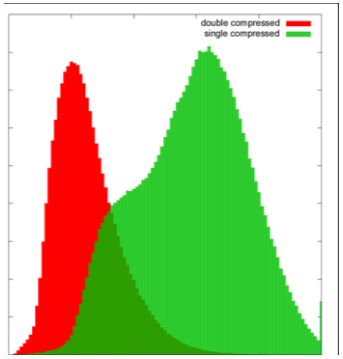

Median

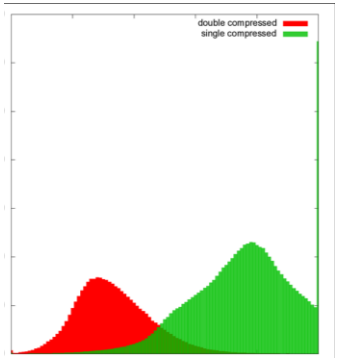

Slope of regression line

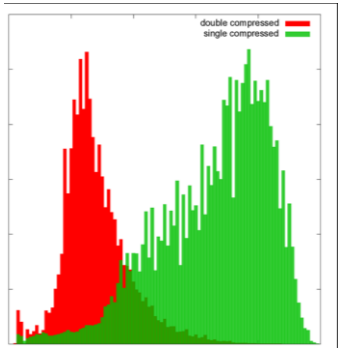

% of points below 0.5

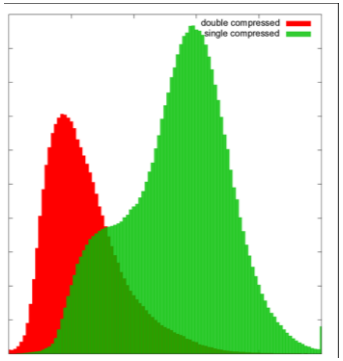

Weighted mean

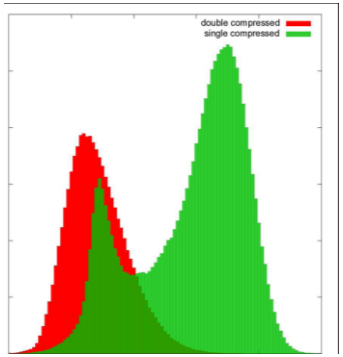

Y-intercept of regression line

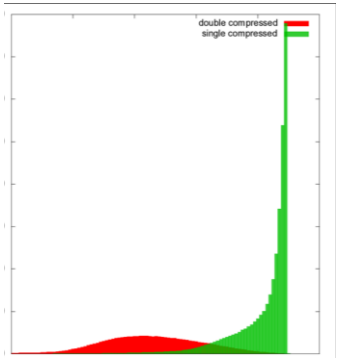

SSD below diagonal

Classification and evaluation

- Classifiers:
  - Naïve Bayes
  - Multilayer Perceptron
  - AdaBoost
  - Random Forests
- Ghost Embeddings in UCID dataset [2]
  - 1338 images, 512x384 pixels
  - Three compression variants:
    a) Purely single-compressed
    b) 192x192 pixels double-compressed, remainder single-compressed
    c) The opposite of b)
- Specificity/Sensitivity on 8x8, 16x16, 32x32 and 64x64 windows

[2] G. Schaefer and M. Stich, “UCID – An Uncompressed Colour Image Database,” in SPIE Storage and Retrieval Methods and Applications for Multimedia, Jan. 2004, pp. 472-480.

Quantitative results (aligned grid, per image)

- Metric: Specificity / Sensitivity

$$\text{Spec.} = \frac{\text{TN}}{\text{TN+FN}} \qquad \text{Sens.} = \frac{\text{TP}}{\text{TP+FN}}$$

- Results by quality difference δ between first and second compression
- Best performance on 8x8 pixels (→ 1 JPEG block):
  - δ = 5: Specificity 0.82, Sensitivity 0.86
  - δ = 20: Specificity 0.997, Sensitivity 0.93

- Comparison to 8x8 pixel block method by Lin et al.:
  - δ = 5: Specificity 0.58, Sensitivity 0.64
  - δ = 20: Specificity 0.70, Sensitivity 0.60

Qualitative results (1)

Qualitative results (2)

Summary

- **Full automation of the JPEG Ghost scheme** for distinguishing single- and double-JPEG compression

- Features:
  - Recompress image block with various (lower) quality levels
  - Differences to input image block serve as basis for feature extraction
  - 6 “straightforward-to-compute” features such as
    - median of differences
    - y-axis intercept of regression line

- **Competitive detection rates** with the simplicity of the Ghost scheme

  e.g. AdaBoost:

  specificity = 0.82, sensitivity = 0.86 at quality difference=5, 8x8 pixels