Steganalysis of Overlapping Images

Jimmy Whitaker

JimmyMWhitaker@gmail.com

Andrew Ker

adk@cs.ox.ac.uk

UNIVERSITY OF OXFORD

DEPARTMENT OF COMPUTER SCIENCE

SPIE/IS&T Electronic Imaging, San Francisco, 11 February 2015

Real-world images

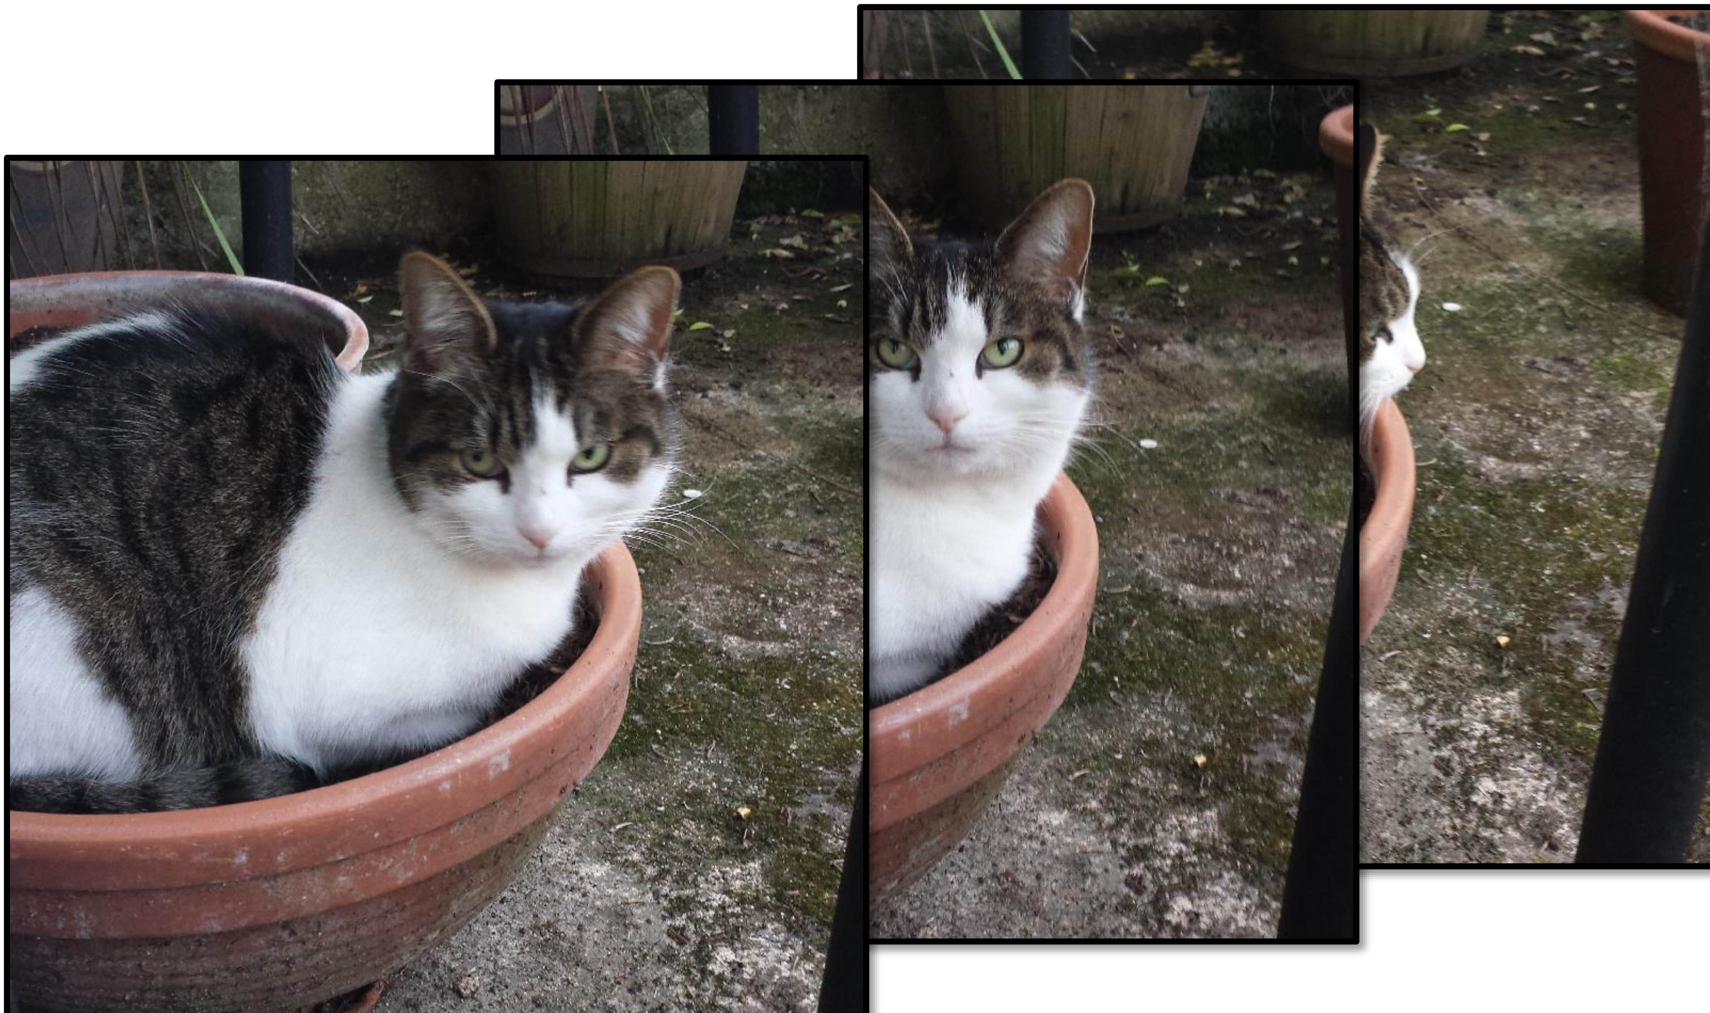

Real-world images

Real-world images

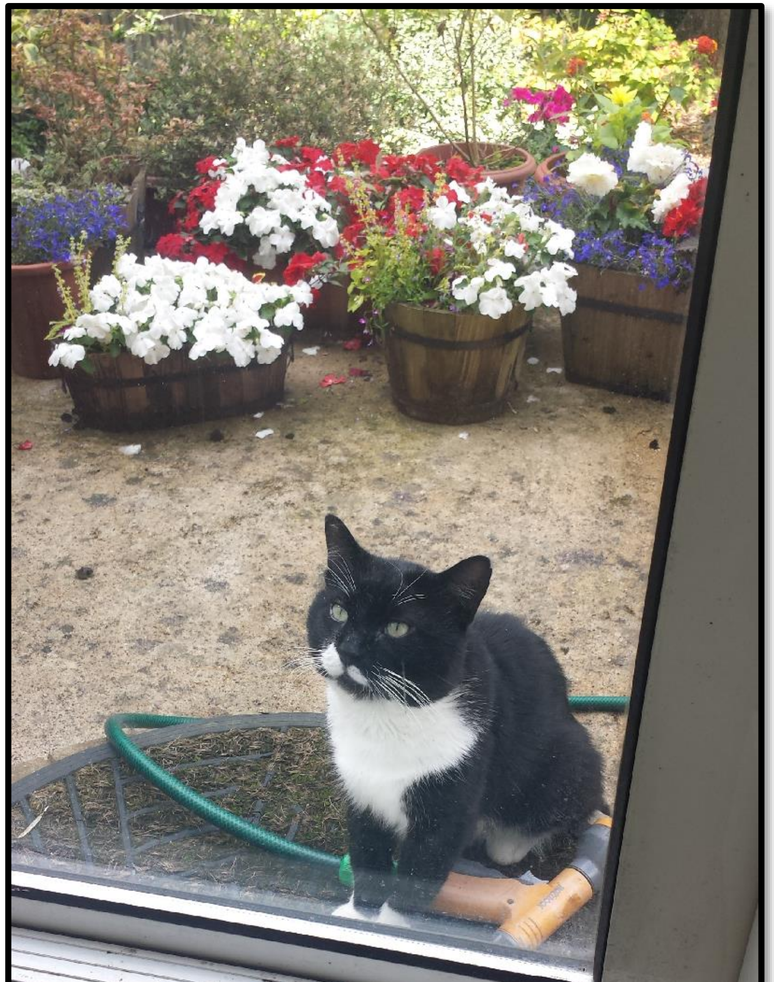

Real-world images

- Are very likely to include a cat.
- Probably contain multiple captures of similar scenes: overlapping images.

Steganalysis

Fundamental difficulty: stego noise is an extremely small signal.

- ***Filtering***

Apply noise reduction filters, keeping only the residual noise.

Use many diverse filters.

- ***Calibration***

Process a stego image to learn about the cover.

- JPEG decompress-crop-recompress [Fridrich et al., 2002]
- Spatial-domain calibration (unsuccessful) [Ker, 2005]
- Contrast parts of an image likely to contain payload with other parts. [Denemark et al., 2014; Carnein et al., 2014]

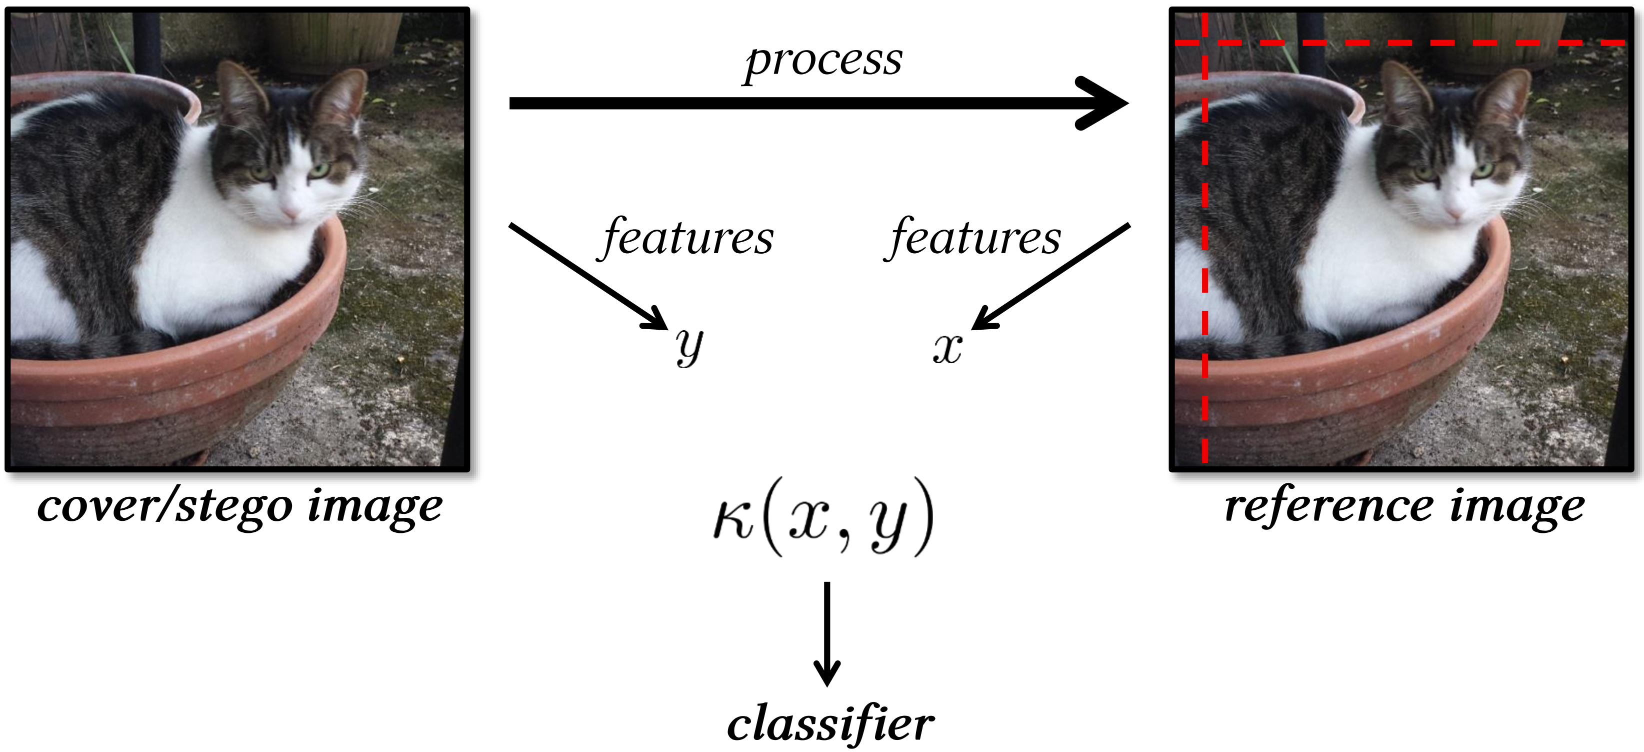

- *Calibration*

Process a stego image to learn about the cover.

- JPEG decompress-crop-recompress [Fridrich et al., 2002]

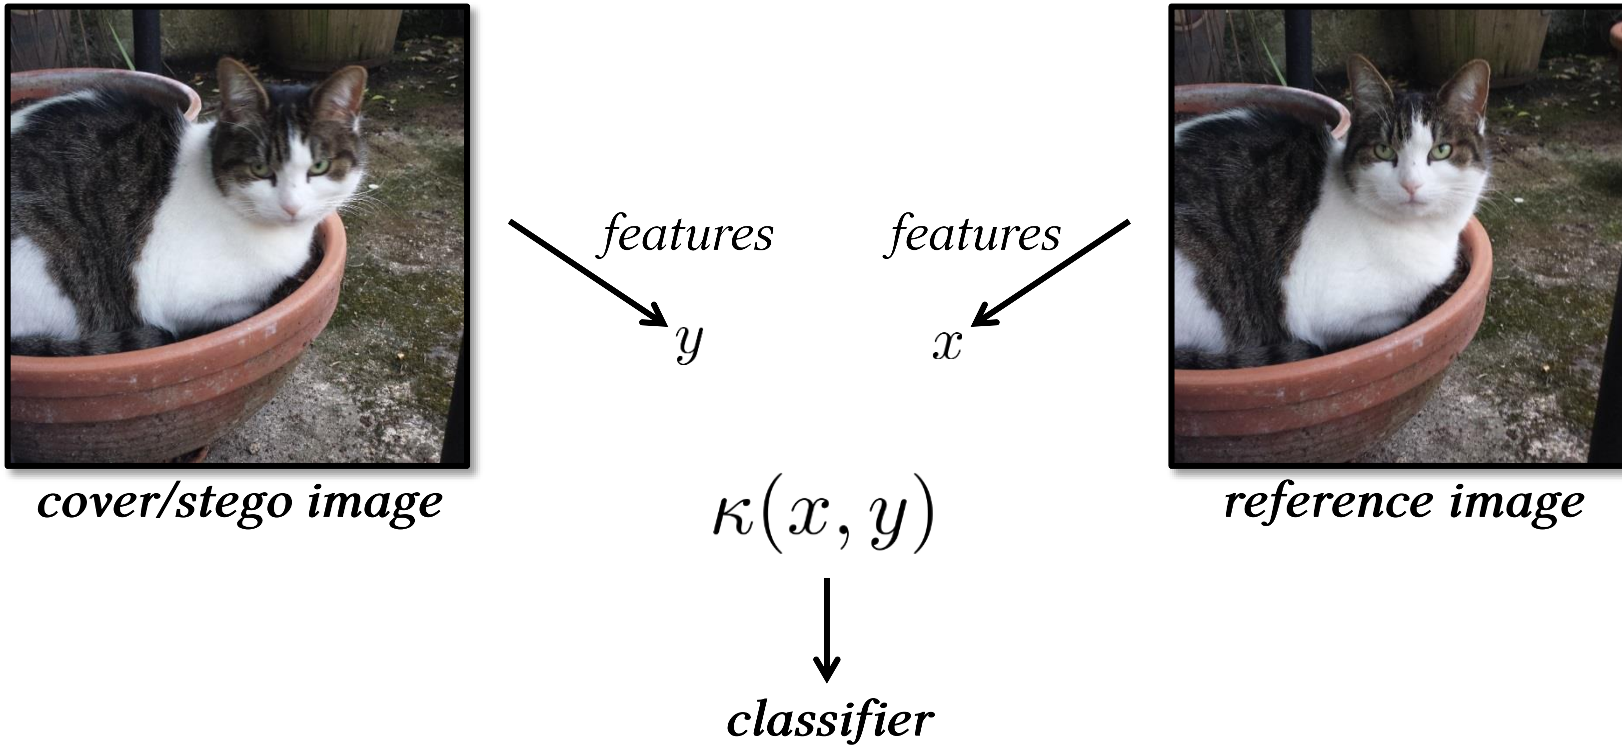

- *Calibration*

Process a stego image to learn about the cover.

- JPEG decompress-crop-recompress [Fridrich et al., 2002]

Investigation

In laboratory conditions,

- analyst has access to the cover source
- stego method & payload size known

given two images
with overlapping content,

- identical camera settings
- one is known to be cover

can one be used to calibrate the other?

Study limited to uncompressed images.

Overlapping image dataset

All taken with Canon G16.

Overlapping image dataset

All camera settings fixed for each scene.

Overlapping image dataset

100% overlap

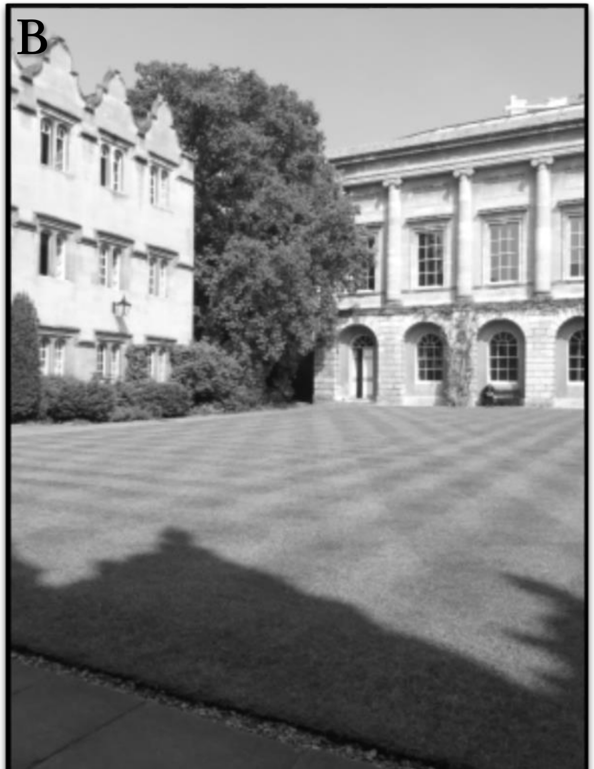

Overlapping image dataset

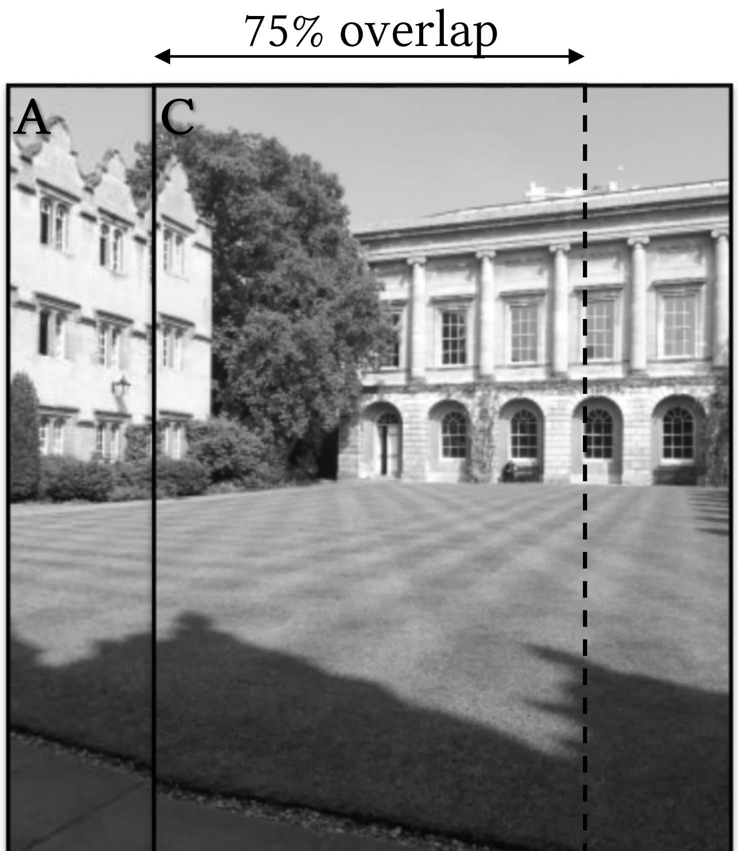

Overlapping image dataset

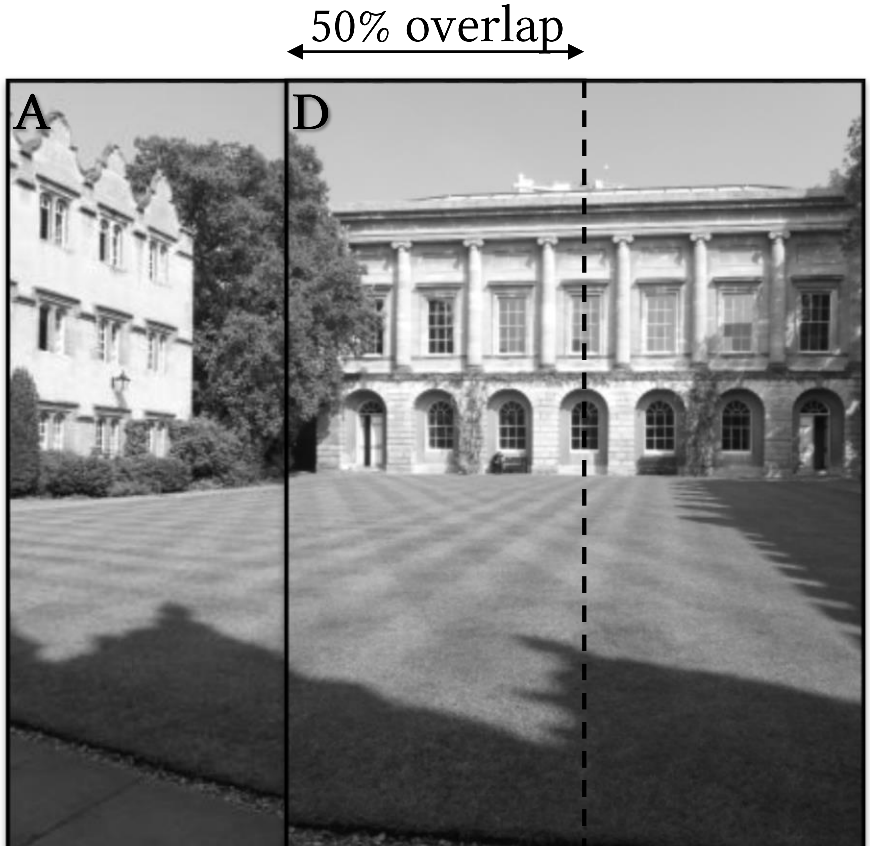

Overlapping image dataset

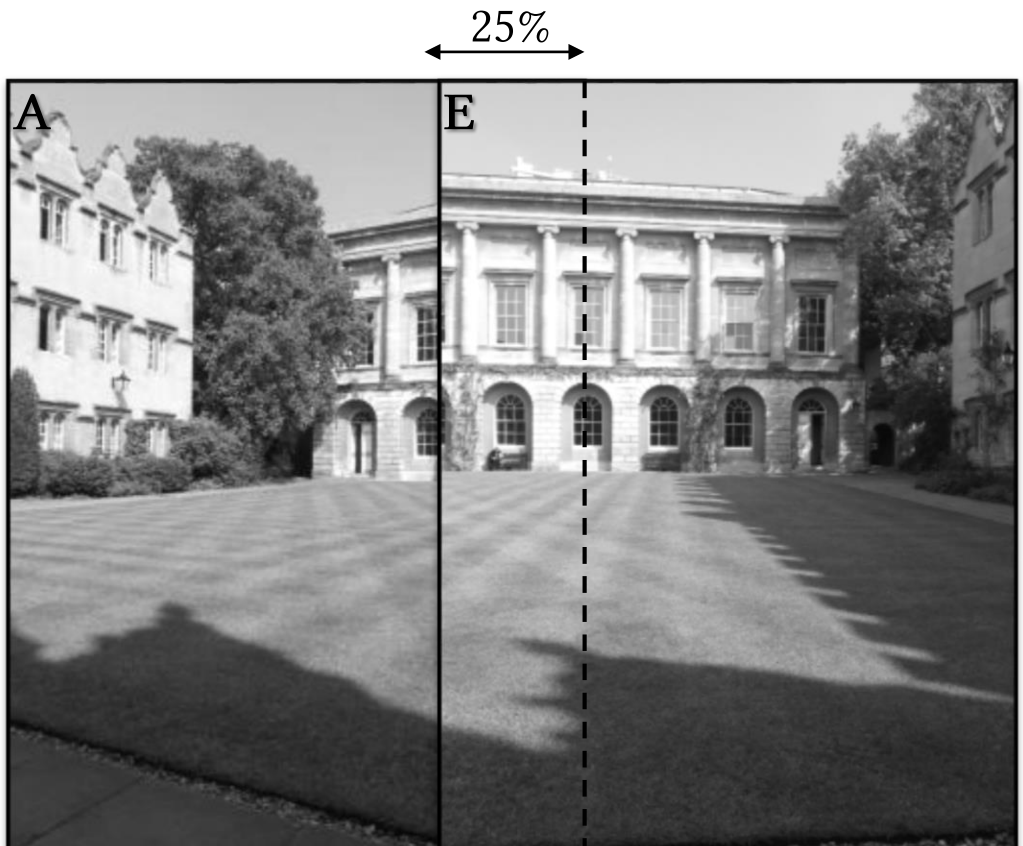

Overlapping image dataset

Overlapping image dataset

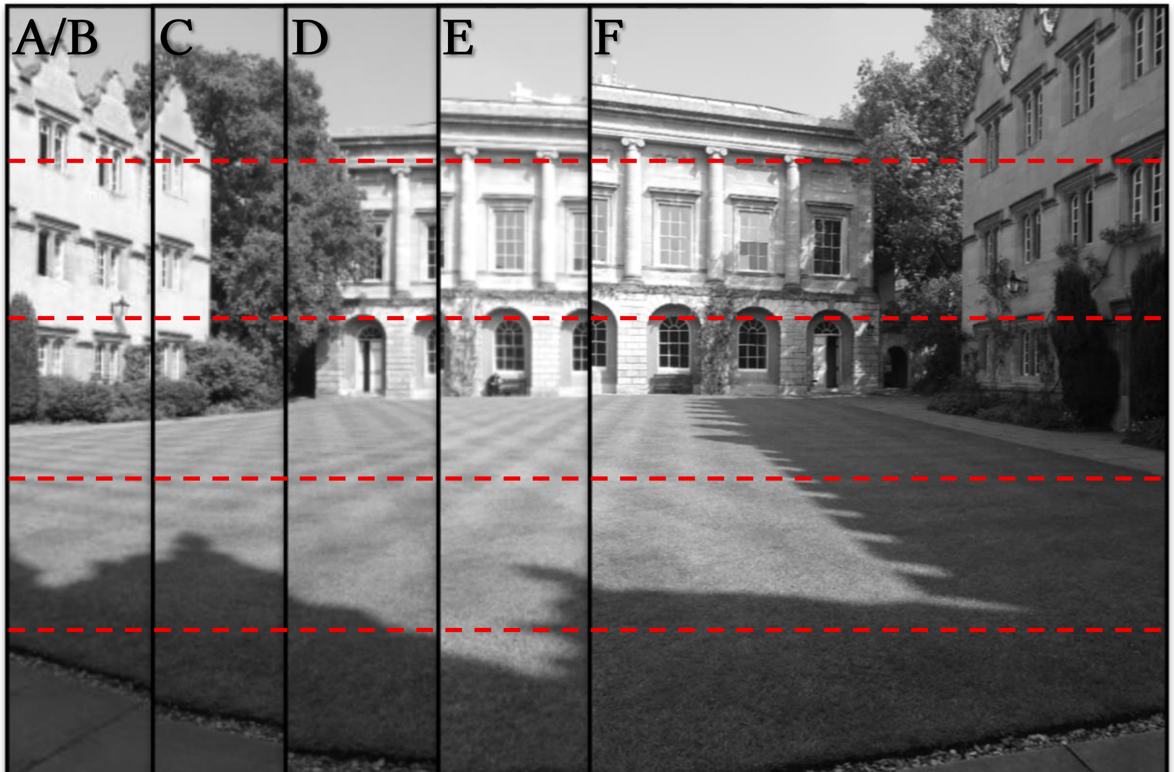

5×500 images @ 3000×800 (2.4Mpix) in each set.

Captured RAW, converted to grayscale using camera software.

Experiments

Embedding

- HUGO @ 0.05/0.1 bpp
- LSBM @ 0.01/0.02 bpp

Features

SPAM	Laplacian filter, residual co-occurrences	[2009]
SRM	Diverse filters, residual co-occurrences	[2012]
PSRM	Diverse filters, random convolutions, histograms	[2013]

Experiments

Embedding

- HUGO @ 0.05/0.1 bpp
- LSBM @ 0.01/0.02 bpp

Features

▪ SPAM	Laplacian filter, residual co-occurrences	**686-dim**
▪ SRM	Diverse filters, residual co-occurrences	**12753-dim**
▪ PSRM	Diverse filters, random convolutions, histograms	**8070-dim**

Experiments

Calibration

- $\kappa(x,y) = y$ *- no calibration (baseline)*
- $\kappa(x,y) = y - x$ *- classical calibration*
- $\kappa(x,y) = x \cdot y$ *- cartesian calibration*
- $\kappa(x,y) = x \cdot (y - x)$
- $\kappa(x,y) = x \cdot y \cdot (y - x)$

… some based on normalized difference are in the paper or Jimmy's dissertation.

Experiments

Calibration

- $\kappa(x, y) = y$
- $\kappa(x, y) = y - x$
- $\kappa(x, y) = x \cdot y$
- $\kappa(x, y) = x \cdot (y - x)$
- $\kappa(x, y) = x \cdot y \cdot (y - x)$

Classifier

Kodovský's ensemble of FLDs.

- Chose best base learner subdimension $d_{sub} \in \{50, 100, 150, 200, 250, 300, 400, 500, 600, 800, 1000, 1500, 2000, 3000\}$.
- 5-fold cross-validation optimizing OOB error, measuring mean testing error.

Cropping

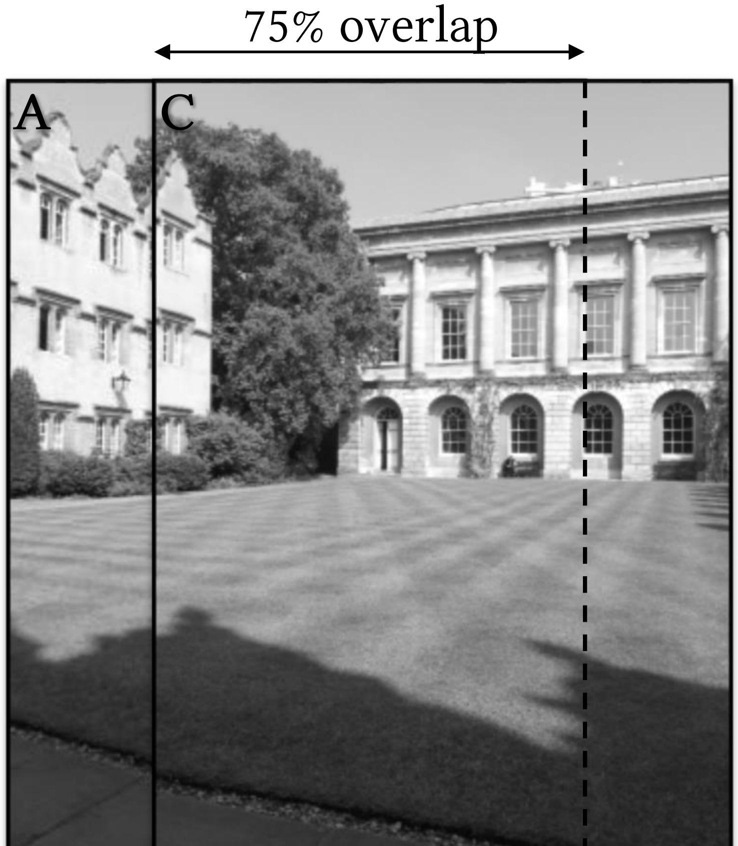

Cropping

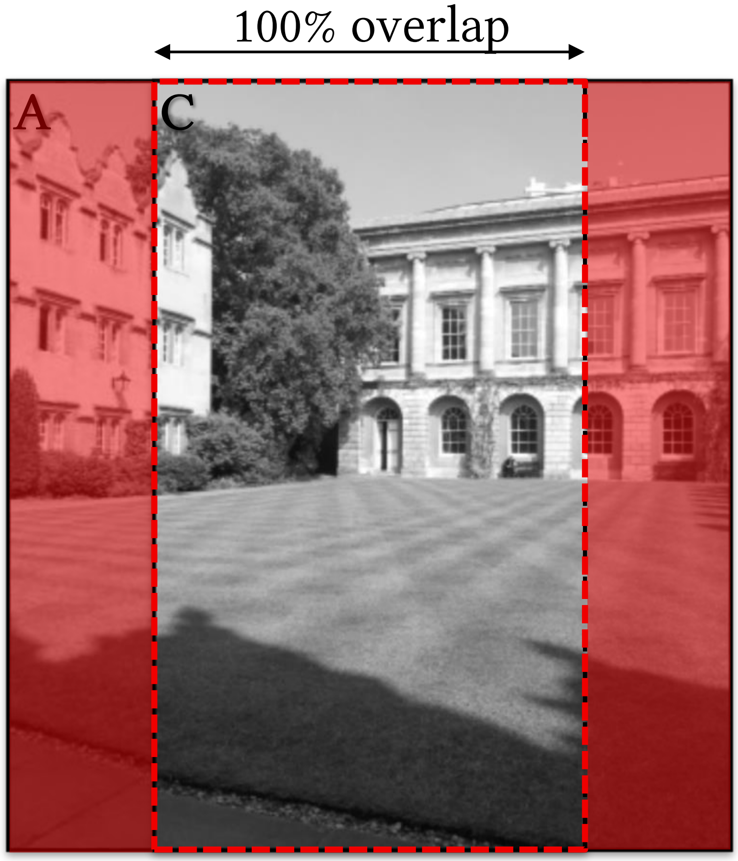

Results

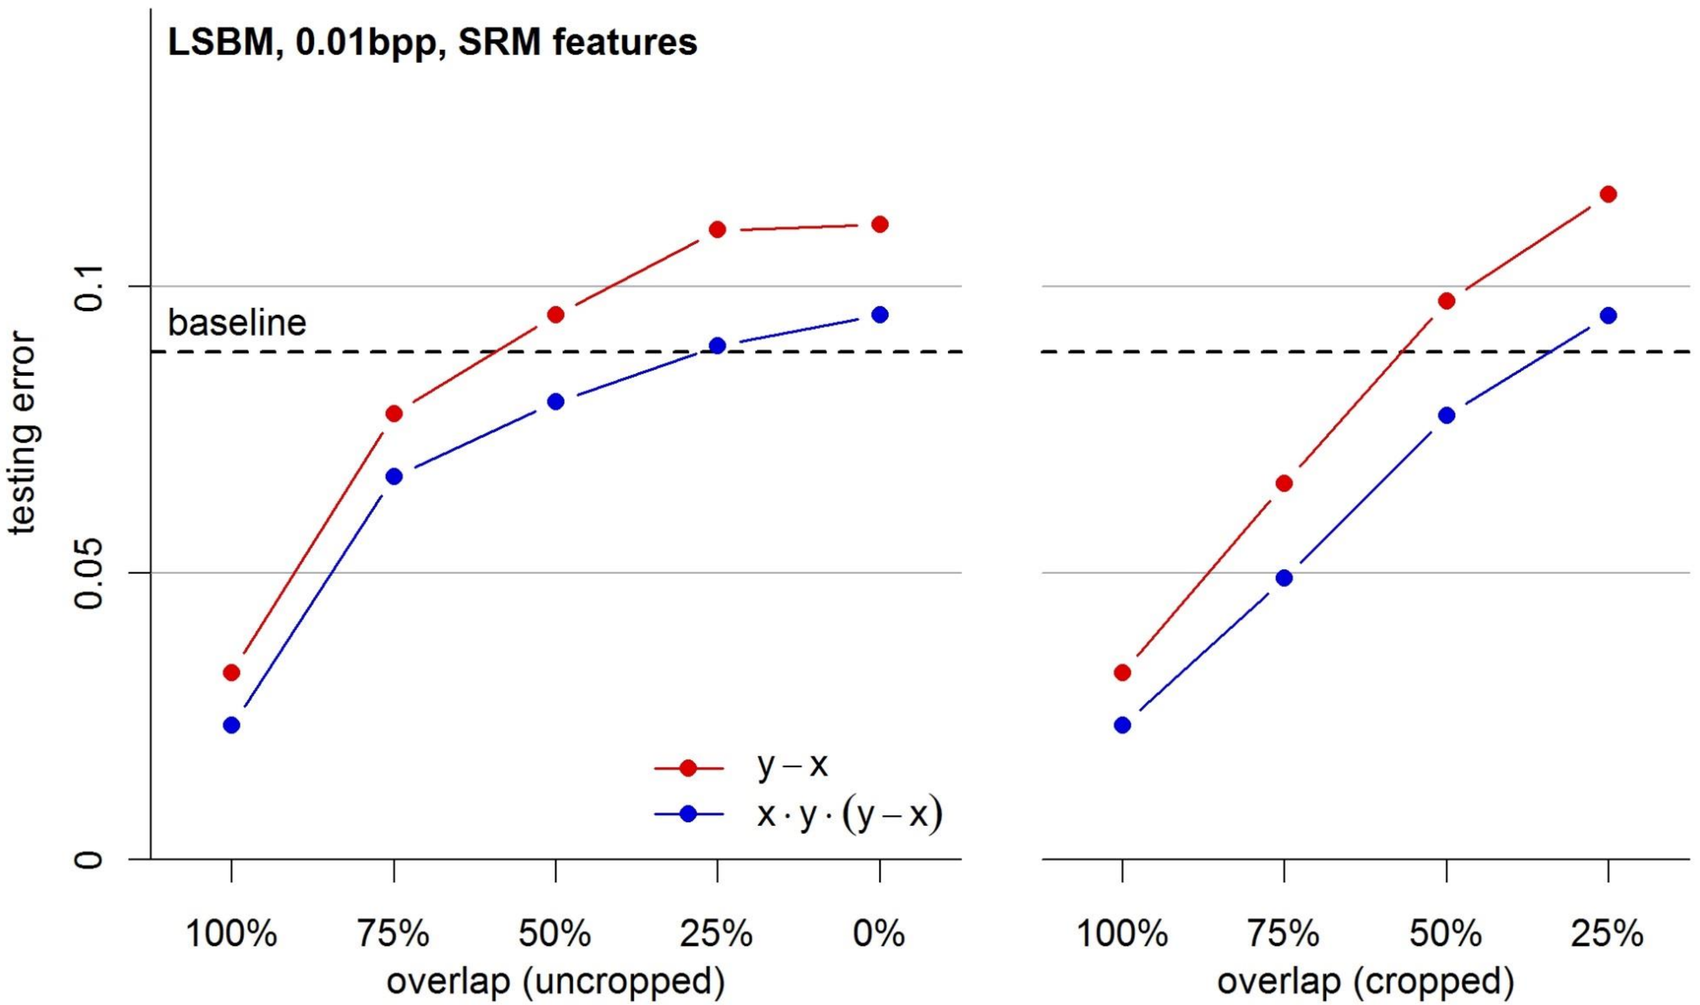

LSBM, 0.01bpp, SRM features
testing error
0.1
0.05
0
baseline
y – x
x · y · (y – x)
100%
75%
50%
25%
0%
overlap (uncropped)
100%
75%
50%
25%
overlap (cropped)

Results

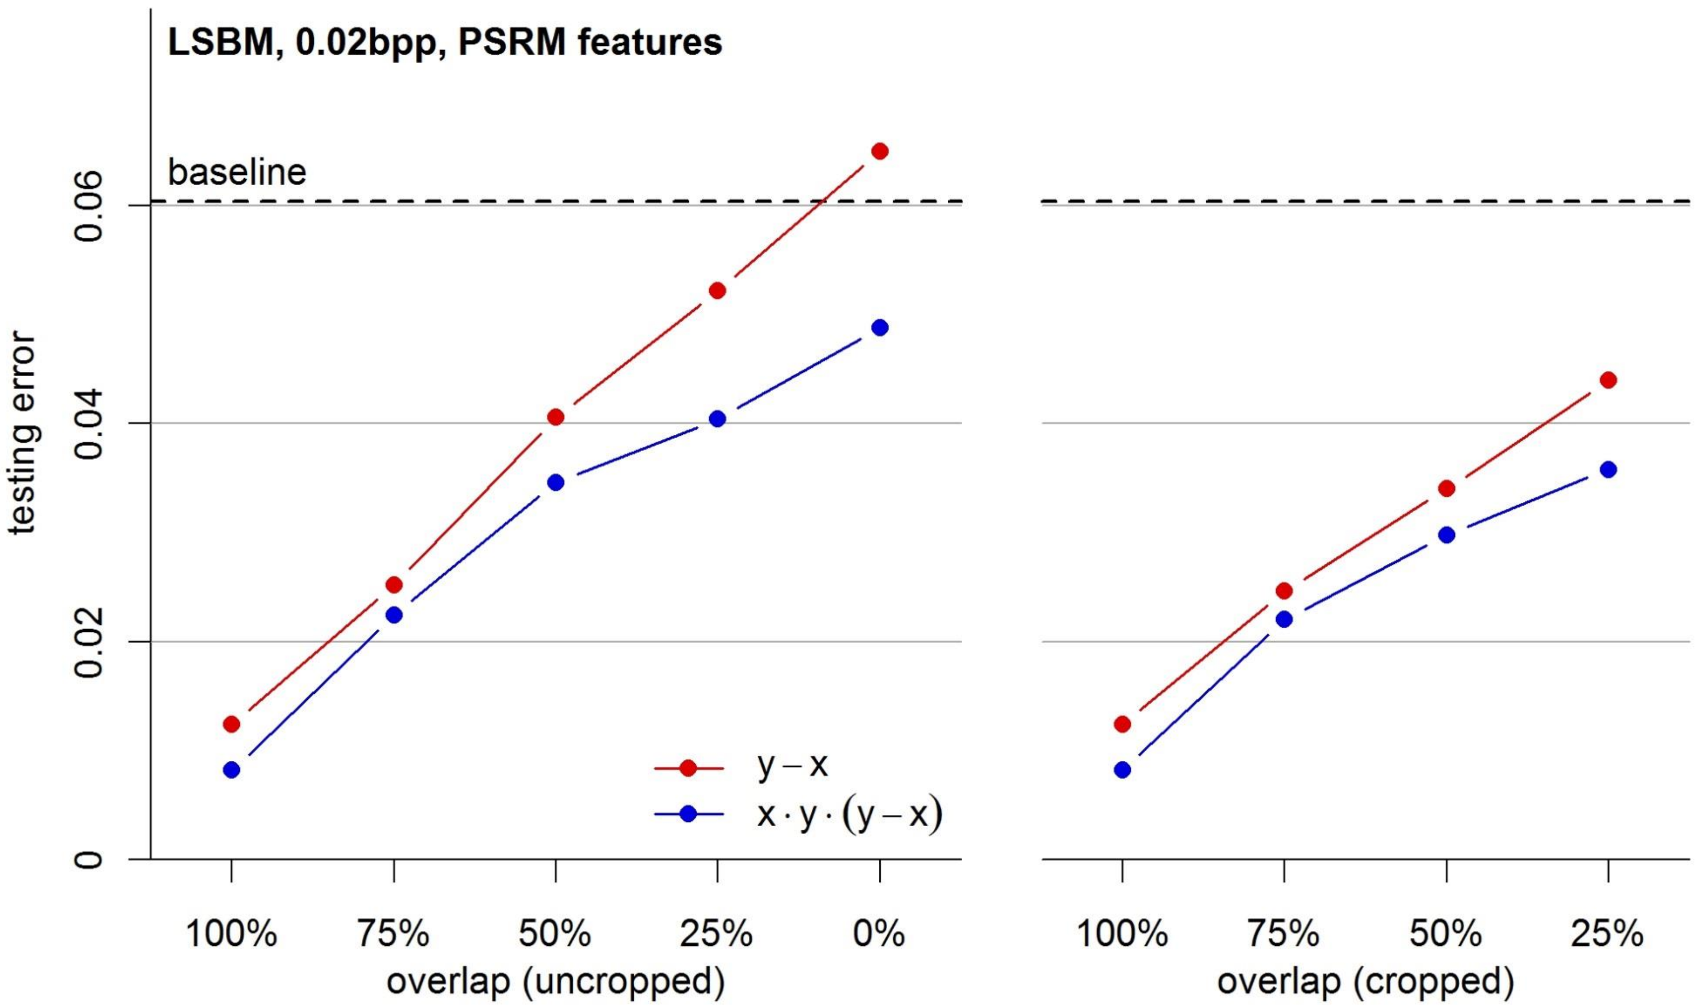

LSBM, 0.02bpp, PSRM features
baseline
testing error
0
0.02
0.04
0.06
y − x
x · y · (y − x)
100%
75%
50%
25%
0%
overlap (uncropped)
100%
75%
50%
25%
overlap (cropped)

Results

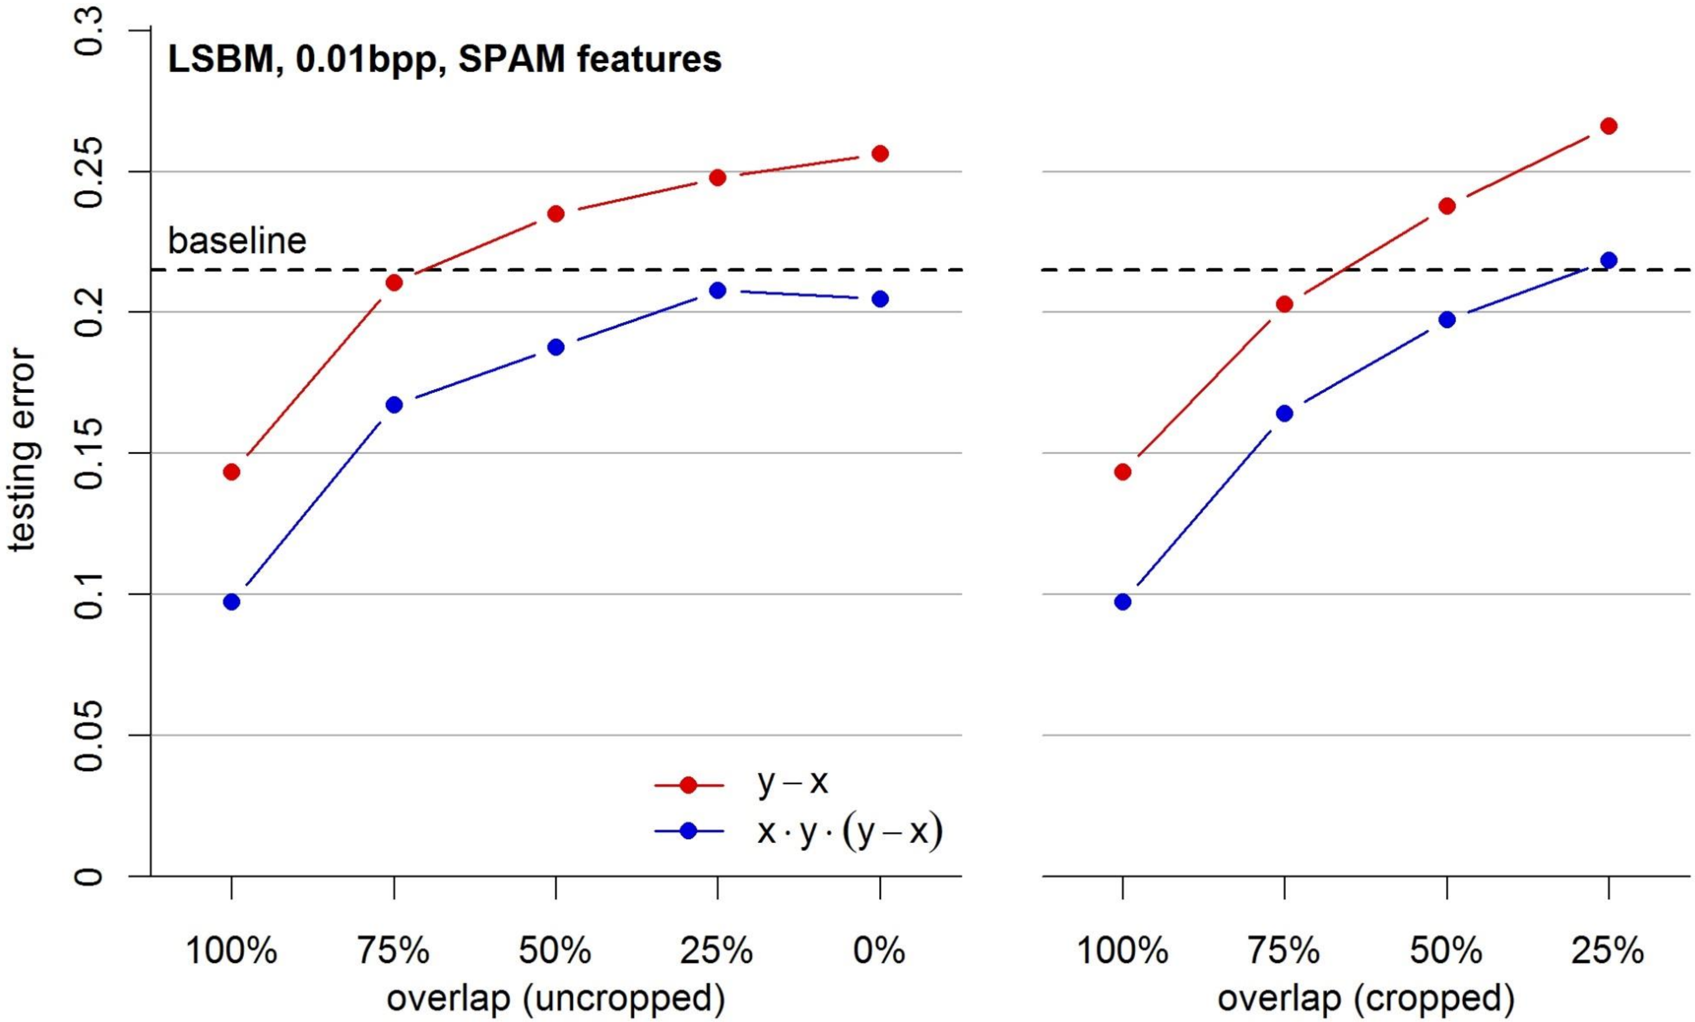

LSBM, 0.01bpp, SPAM features
baseline
testing error
0.3
0.25
0.2
0.15
0.1
0.05
0
y – x
x · y · (y – x)
100%
75%
50%
25%
0%
overlap (uncropped)
100%
75%
50%
25%
overlap (cropped)

Results


HUGO, 0.05bpp, SRM features
testing error
0
0.05
0.1
0.15
0.2
0.25
0.3
baseline
y − x
x · y · (y − x)
100%
75%
50%
25%
0%
overlap (uncropped)
100%
75%
50%
25%
overlap (cropped)

Robustness

Mismatched payload

Seems quite robust.

Mismatched reference

Robust if we use $\kappa(x, y) = y - x$ and a ‘double-sided’ classifier.

Mismatched amount of overlap

Not very robust: scope for further work.

Distance

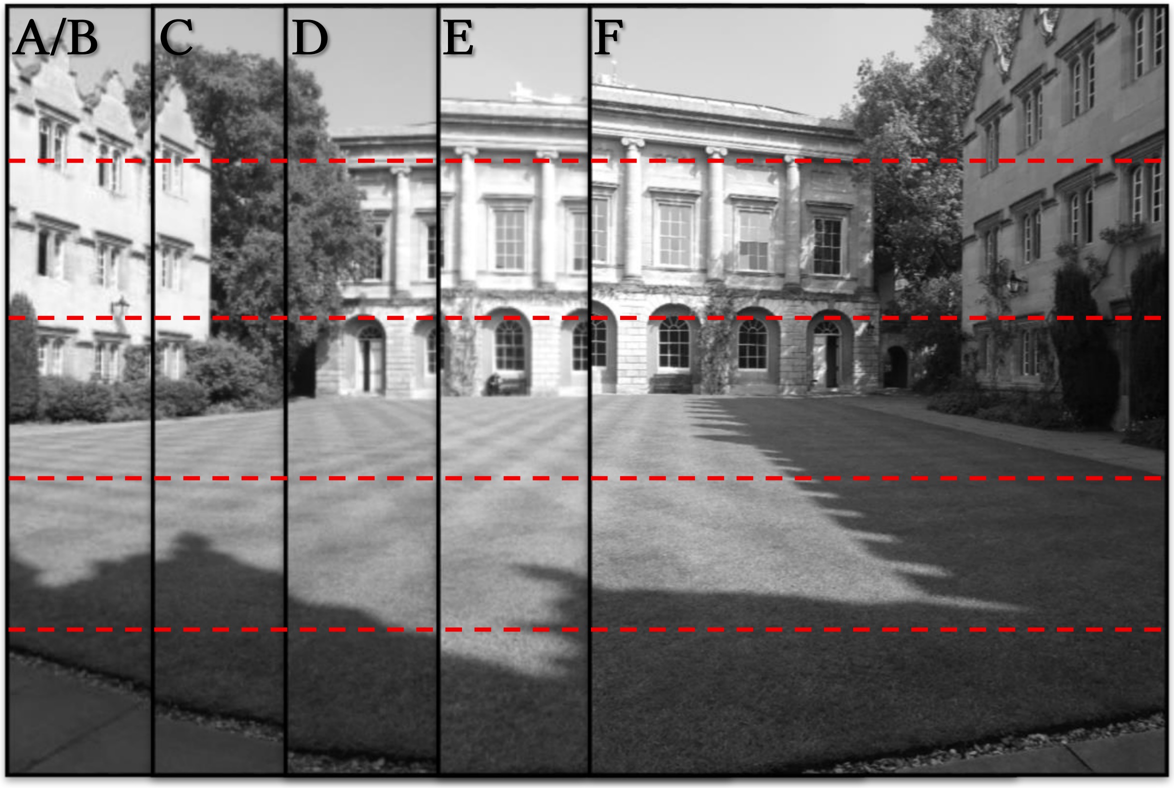

How ‘far apart’ are these images, and how far is a stego object?

Distance

Whitened (Mahalanobis-like) distance

- Apply PCA to pooled cover & stego features.
- Keep all numerically-significant components.
- Normalize each dimension, measure Euclidean distance.

HUGO 0.05 bpp SRM features	**mean distance to stego image**	**mean distance to cover, with overlap**				
		100%	75%	50%	25%	none
Whitened distance:	0.034	0.063	0.281	0.445	0.564	0.650

Scaled so that mean distance between different covers is 1.

Distance

Projected distance

- Train numerically-stabilized FLD on all cover & stego features.
- Project features onto separating vector.

HUGO 0.05 bpp SRM features	mean distance to stego image	mean distance to cover, with overlap				
		100%	75%	50%	25%	none
Whitened distance:	0.034	0.063	0.281	0.445	0.564	0.650
Projected distance:	4.076	1.507	1.594	1.682	1.705	1.694

Scaled so that mean distance between different covers is 1.

Illustration

Illustration

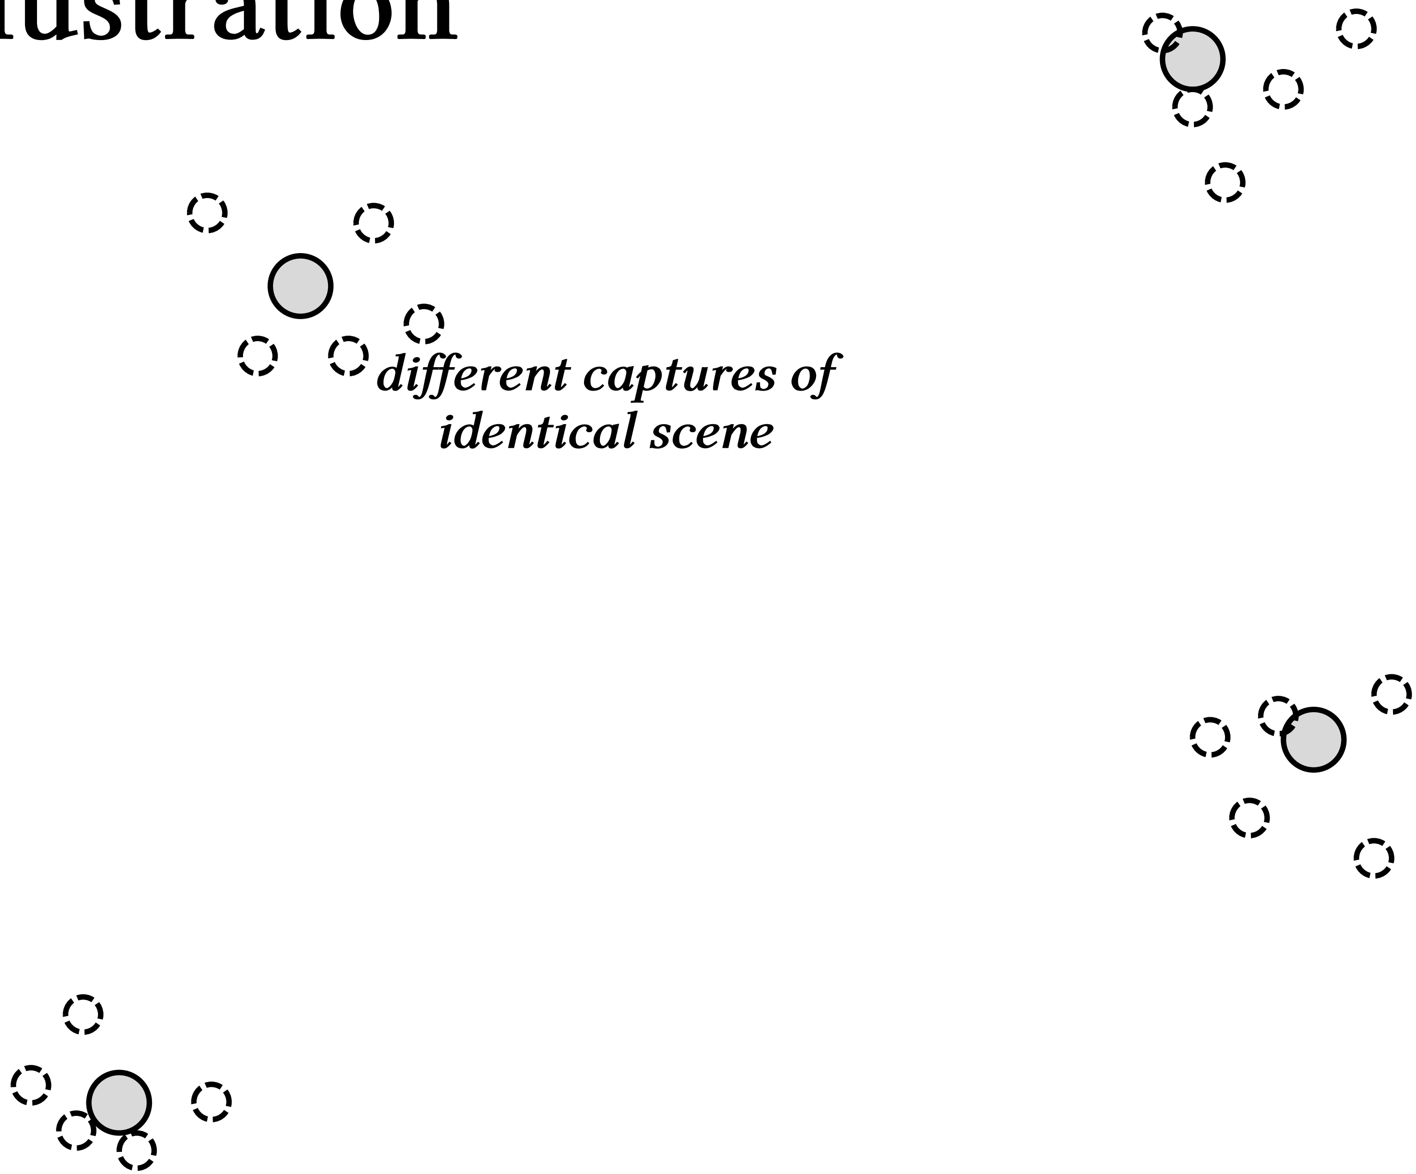

Illustration

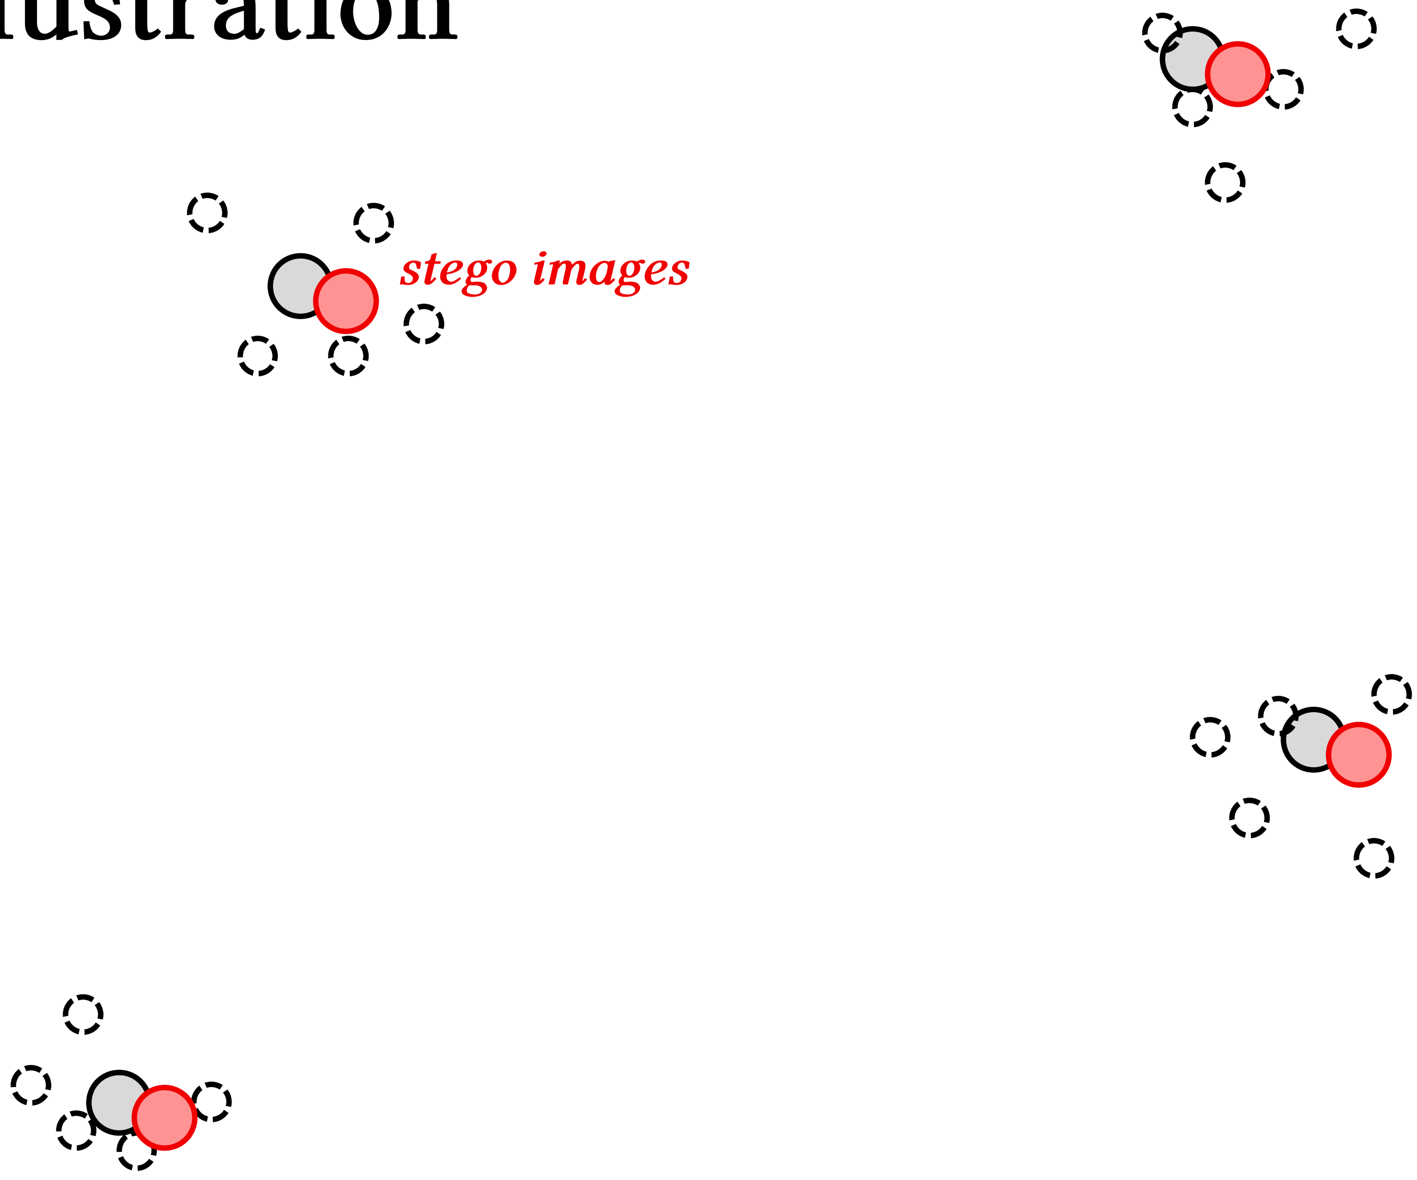

Conclusions

- Images overlapping by 75% or more make classification better.

  Seems good detectors benefit more than bad ones.
  Should be a regressor for difference in payload?

- Turning it into a forensic tool:

  Automatically identifying overlap ✓
  Checking camera settings ✓
  Developing training data ?

- Limitations:

  Controlled conditions.
  Stable camera.
  Only considered uncompressed images.

Conclusions

- Images overlapping by 75% or more make classification better.

  Seems good detectors benefit more than bad ones.
  Should be a regressor for difference in payload?

- Turning it into a forensic tool:

  Automatically identifying overlap ✓
  Checking camera settings ✓
  Developing training data ?

Pilot study on JPEG images
(*q.f. 80, nsF5 @ 0.02 bpnc, JRM features*)

Uncalibrated error	5.6%
Calibrated by decompress-crop-recompress	4.9%
Calibrated by 100% overlapping image	4.7%