Steganalysis in resized images

Jan Kodovský, Jessica Fridrich

ICASSP 2013

BINGHAMTON
UNIVERSITY
STATE UNIVERSITY OF NEW YORK

Outline

1. Steganography – basic concepts
2. Why we study steganalysis in resized images
3. Eye-opening experiment on BOSSbase
4. Experiments on a controlled data set
5. Theoretical analysis (**new scaling law**)
6. Conclusion

Steganography

- Goal: Hide message in a cover object so that its presence cannot be established
  - cover (**x**) + message (**m**) = stego object (**y**)
- $\mathbf{x} \sim P_{\mathbf{x}}$, $\mathbf{y} \sim P_{\mathbf{y}}$, $P_{\mathbf{x}} = P_{\mathbf{y}} \Rightarrow$ **perfect security**
- When **x** are digital media, steganography is **imperfect**, $P_{\mathbf{x}} \neq P_{\mathbf{y}}$
  - secure payload $\propto \sqrt{n}$, n is cover size (the square root law **SRL**)
  - maximize payload for a given level of statistical detectability (value of the KL divergence $D_{\mathrm{KL}}(P_{\mathbf{x}}||P_{\mathbf{y}}) > 0$)
  - steganographic security evaluated empirically for a **given image source**

Why we study resizing

1. Imagery attached to e-mails is usually resized.
2. Image-sharing portas, e.g, Flickr and Picassa, offer several resized versions of images.
3. Steganalysis and steganography is benchmarked on standard databases, which are usually resized (e.g., BOSSbase, BOWS2).
4. Resizing changes statistical properties of pixels, which has a **strong** effect on steganalysis $\Rightarrow$ SRL no longer holds.

Eye-opening experiment

BOSSbase (10,000 images of size 512×512)

- Introduced in 2010 – BOSS competition
- Standard cover source for benchmarking today

Image processing pipeline

- Demosaicking from RAW images (7 different cameras)
- Converting to 8-bit grayscale
- Resizing to the smaller side 512 pixels – `convert` (ImageMagick)
- Central-cropping to 512×512

Steganalysis of HUGO (Pevný, 2010)

- Detector: binary ensemble classifier (Kodovský, 2012)
- Feature vector: 12,753 dim. Spatial Rich Model (Fridrich, 2012)
- Evaluation: $P_{\mathrm{E}} = \min \frac{1}{2}(P_{\mathrm{FA}} + P_{\mathrm{MD}})$ on the testing set

Eye-opening experiment

BOSSbase (10,000 images of size 512×512)

- In
- S

Image

- D
- C
- R
- C

Stega

- D
- F
- Evaluation: $P_{\mathrm{E}} = \min \frac{1}{2}(P_{\mathrm{FA}} + P_{\mathrm{MD}})$ on the testing set

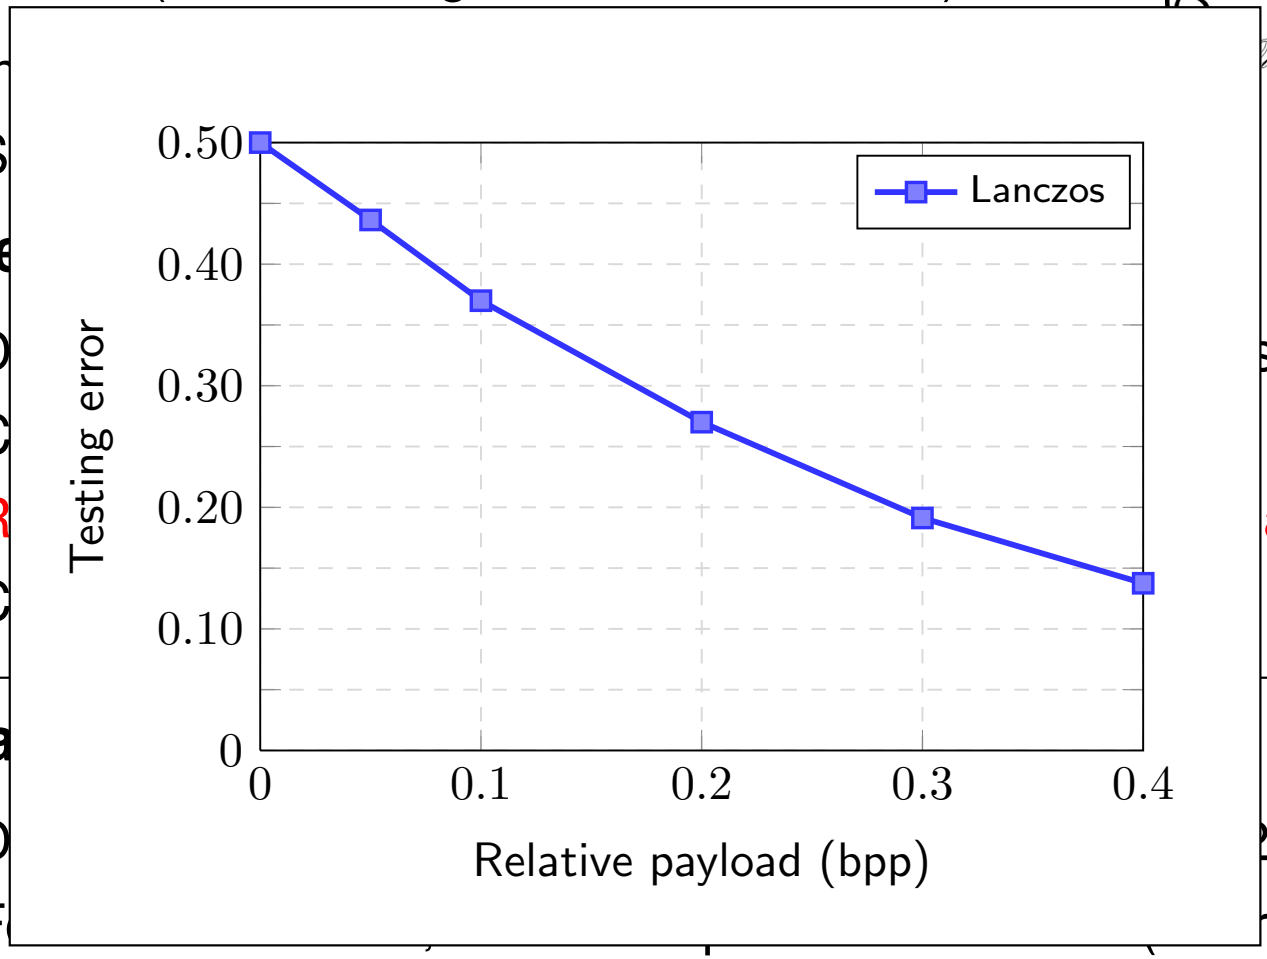

Eye-opening experiment

BOSSbase (10,000 images of size 512×512)

- In
- S

Image

- D
- C
- R
- C

Stega

- D
- F
- Evaluation: $P_{\mathrm{E}} = \min \frac{1}{2}(P_{\mathrm{FA}} + P_{\mathrm{MD}})$ on the testing set

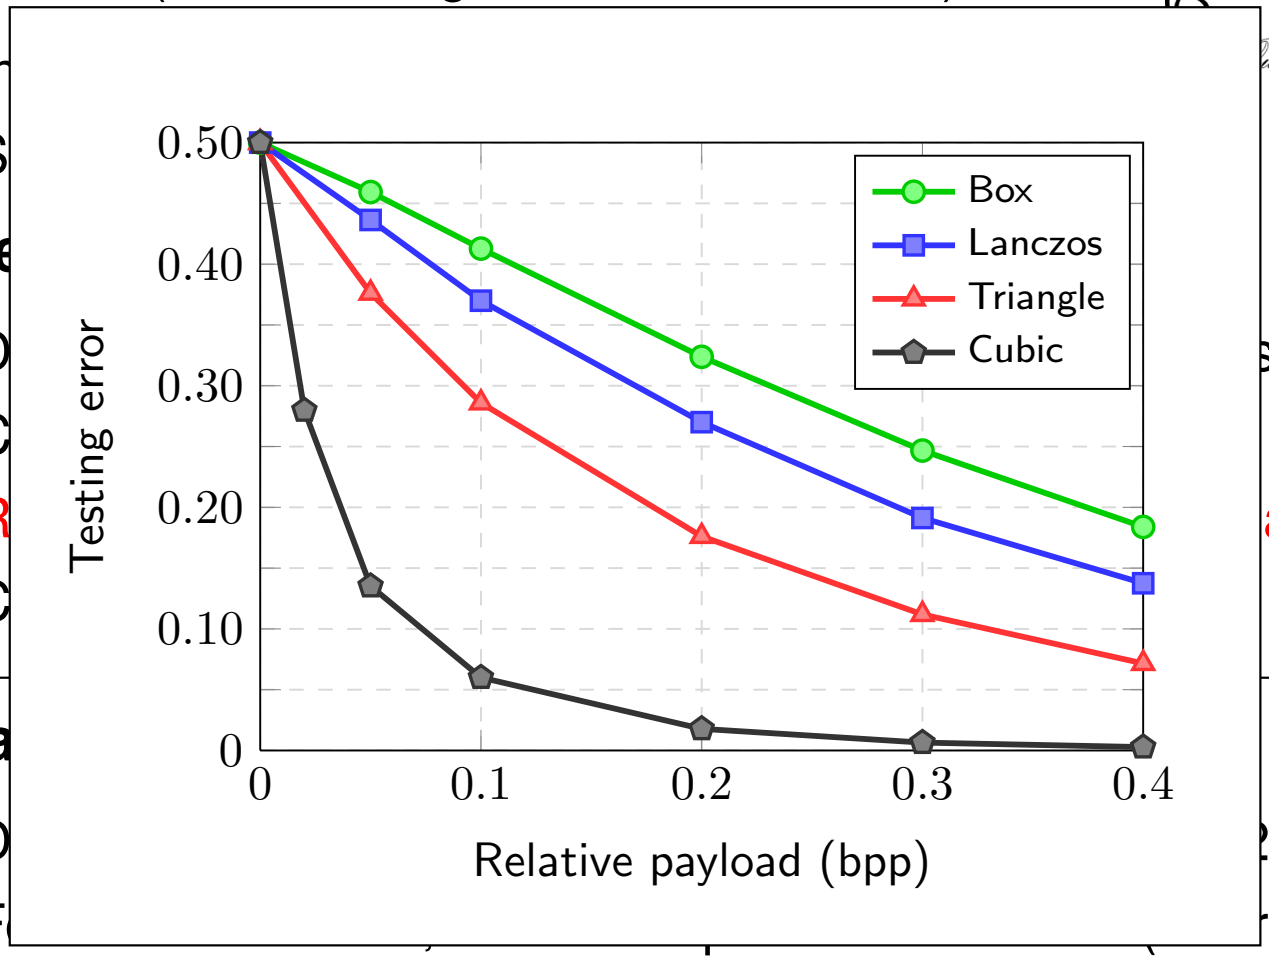

Formalization

Image registered by camera

discrete sampling function

$$\mathbf{X}(x,y) = Q\left(C_\Theta(x,y) \cdot f(x,y)\right)$$

scalar quantizer

2D scene (reality)

Resized image

interpolation kernel

$$\mathbf{X}^{(k)}(x,y) = Q\left(C_{\Theta^{(k)}}(x,y) \cdot (\mathbf{X} * \varphi)(x,y)\right)$$

- $(\mathbf{X} * \varphi)(x,y)$ serves as an approximation of reality
- Kernel function φ satisfies $\int_{\mathbb{R}^2} \varphi(x,y)\mathrm{d}x\mathrm{d}y = 1$
- $\Theta, \Theta^{(k)}$ ... parameters of the sampling function (rectangular grid of spatial locations)

Controlled testing environment

Image source

- 1,000 images from a single camera model
  - Canon EOS 400D (RAW images available at BOSS' website)

Simplification

- HUGO $\rightarrow$ LSB Matching, fixed change rate β
- SRM $\rightarrow$ 4D cooc. of quantized $\mathbf{R} = \mathbf{X} - \mathbf{K} * \mathbf{X}$ (169), $\mathbf{K} = \begin{pmatrix} -0.25 & 0.5 & -0.25 \\ 0.5 & 0 & 0.5 \\ -0.25 & 0.5 & -0.25 \end{pmatrix}$
- ImageMagick's `convert` $\rightarrow$ Matlab's `imresize` (different kernels)

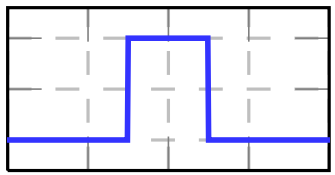

Box

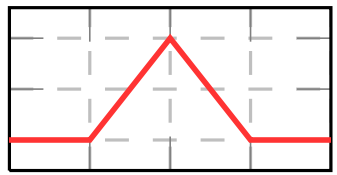

Triangle

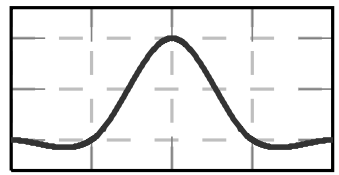

Cubic

- After resizing, images are always central-cropped to 512×512
  - Eliminates effects of the SRL

Choice of the kernel

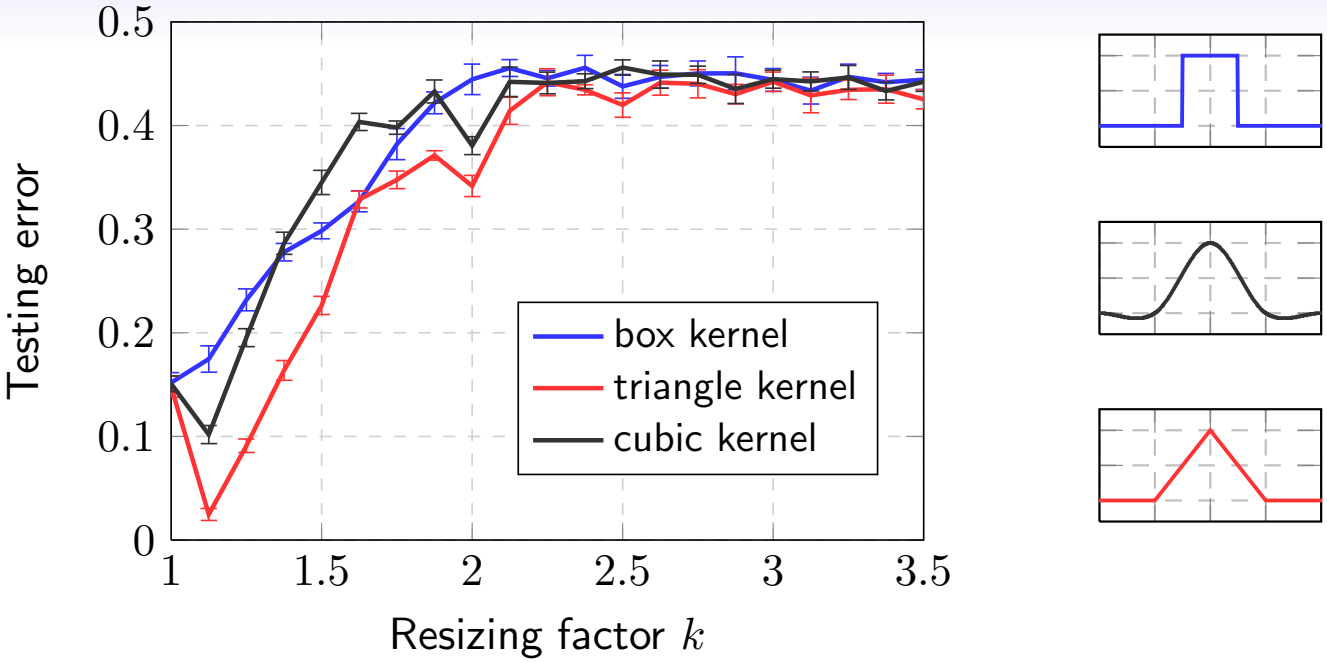

- Resizing factor k . . . downsampling to $1/k$ of the original size
- Differences due to different combinations of pixels during interpolation
- In general, detection error grows with higher k (downsampling w/o anti-aliasing weakens pixel dependencies)
- The error is not necessarily monotonous in k

Choice of the kernel

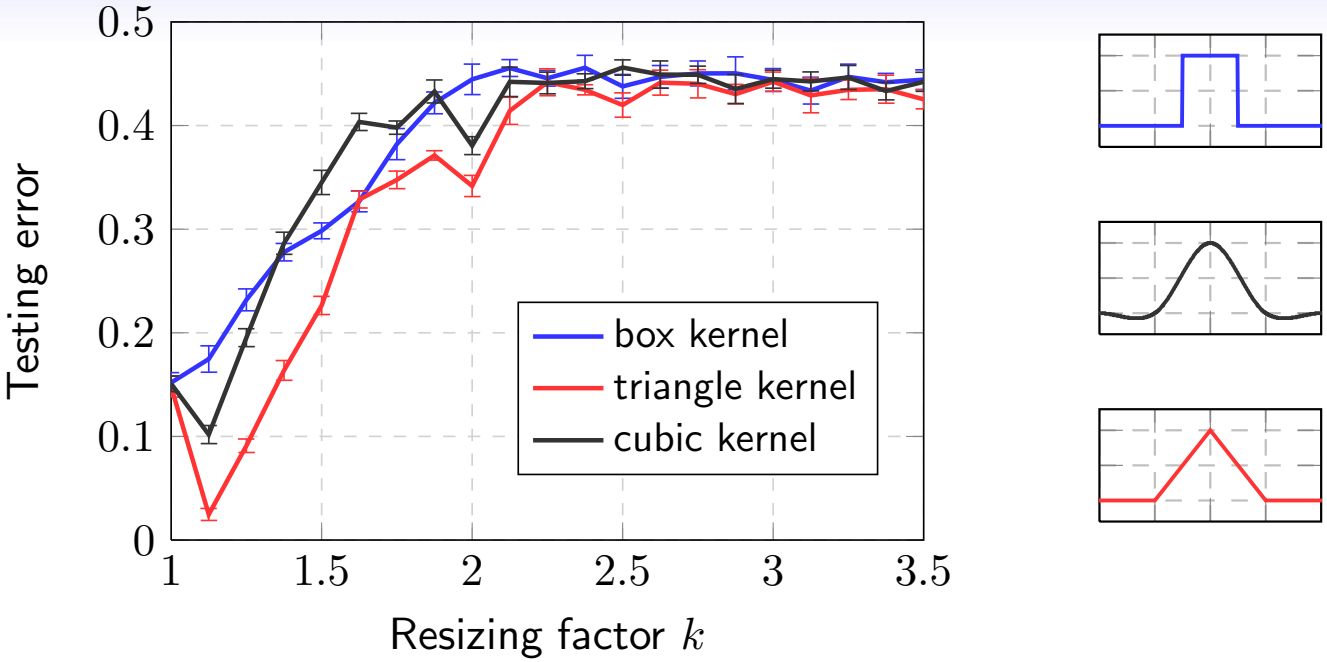

Two important factors influencing the performance

- Distance between pixels at resolution *k*
- Position of the first pixel in the resized image

Choice of the kernel

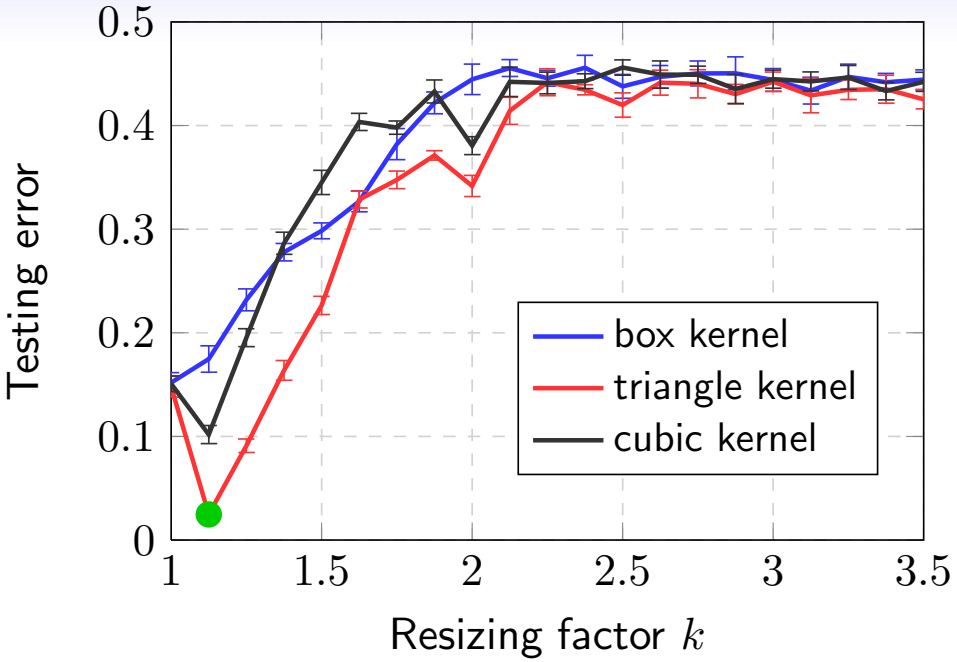

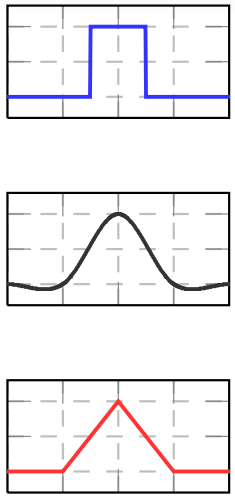

Triangle kernel, $k \approx 1.1$

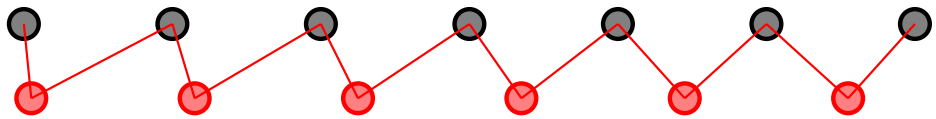

original image
resized image

- $k > 1 \Rightarrow$ Original pixels contribute to *two* pixels of the resized image
  $\Rightarrow$ increased strength of dependencies $\Rightarrow$ easier steganalysis

Choice of the kernel

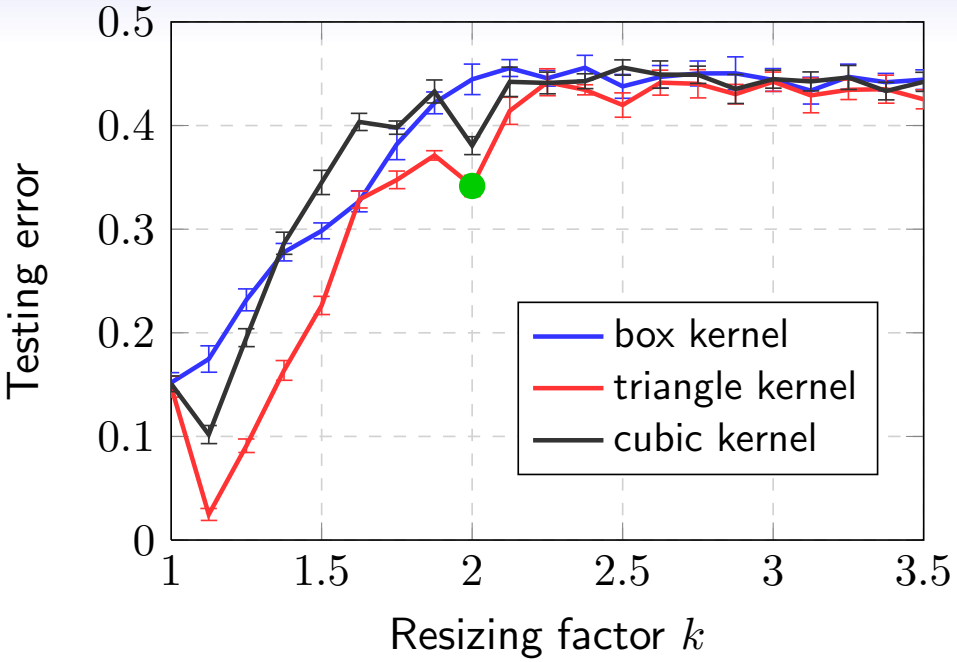

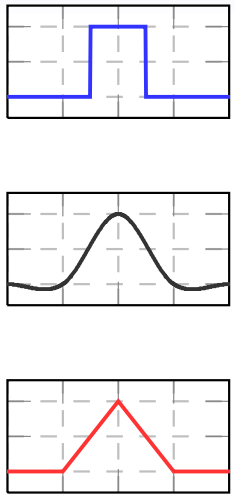

Triangle kernel, $k = 2$ (downsampling by 50%)

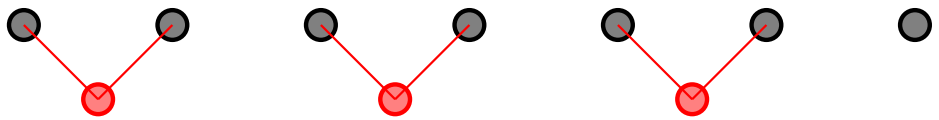

original image
resized image

- Perfect synchronization $\Rightarrow$ averaging $\Rightarrow$ increases local correlations
- Position of the first pixel is important

Choice of the kernel

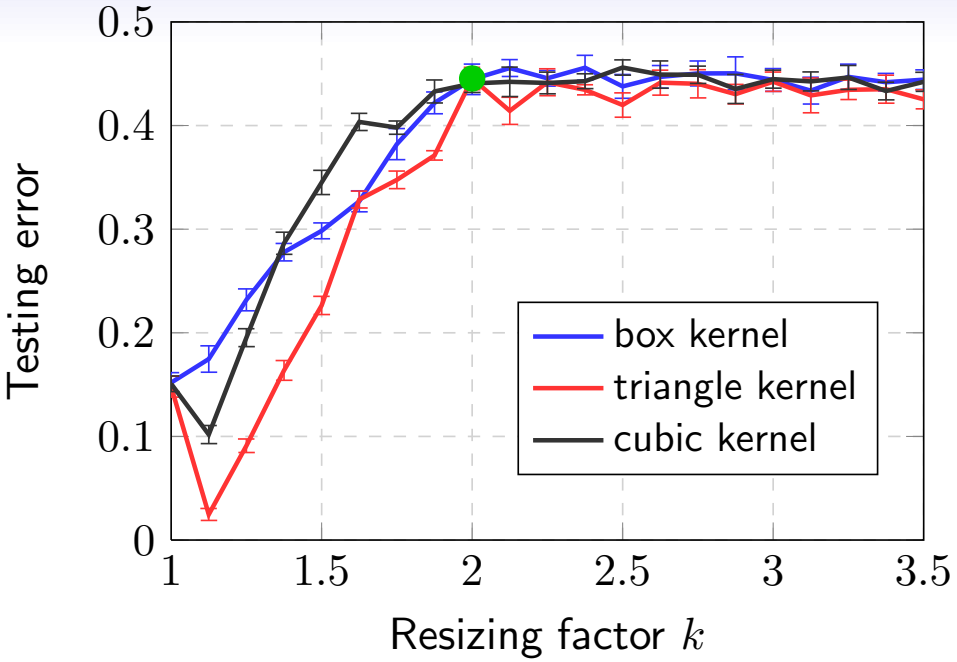

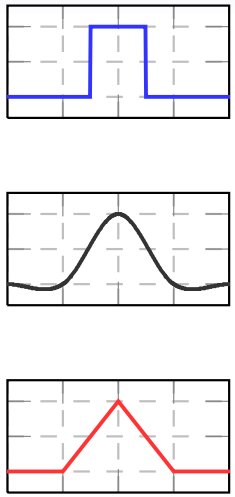

Triangle kernel, $k = 2$ (downsampling by 50%)

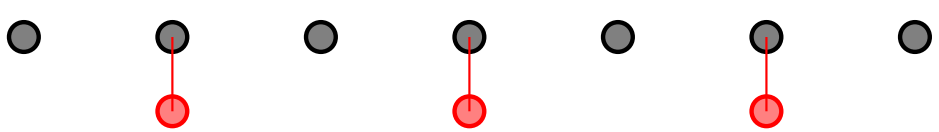

- If the pixels were aligned $\Rightarrow$ subsampling $\Rightarrow$ weaker dependencies
- All three kernels would become identical (zeros at integers)

Theoretical scaling law

Question: How does secure payload scale w.r.t. resolution?

Assumptions:

- Box kernel: nearest neighbor interpolation (subsampling)
- Image model: pixel rows $=$ first-order Markov chain
  - TPM $\mathbf{A} = (a_{ij})$, $a_{ij} = \frac{1}{Z_i} \exp\left\{-\left(\frac{|i-j|}{\tau}\right)^{\gamma}\right\}$
  - Parameters τ and γ estimated from 500 images

- TPM at resolution k is $\mathbf{A}^k$ (generalized matrix power)
- $D_{\mathrm{KL}}(k;\beta) \approx \frac{1}{2} n \beta^2 I(k)$, $I(k) =$ steganographic FI rate
  - Filler (IH, 2009) – closed-form expression for $I(k)$ for any mutually-independent embedding operation
- For constant statistical detectability over k:

$$D_{\mathrm{KL}}(1,\beta) = D_{\mathrm{KL}}(k;\alpha(k)\beta) \quad \Rightarrow \quad \alpha(k) = \sqrt{I(1)/I(k)}$$

Scaling factor for Canon 400D

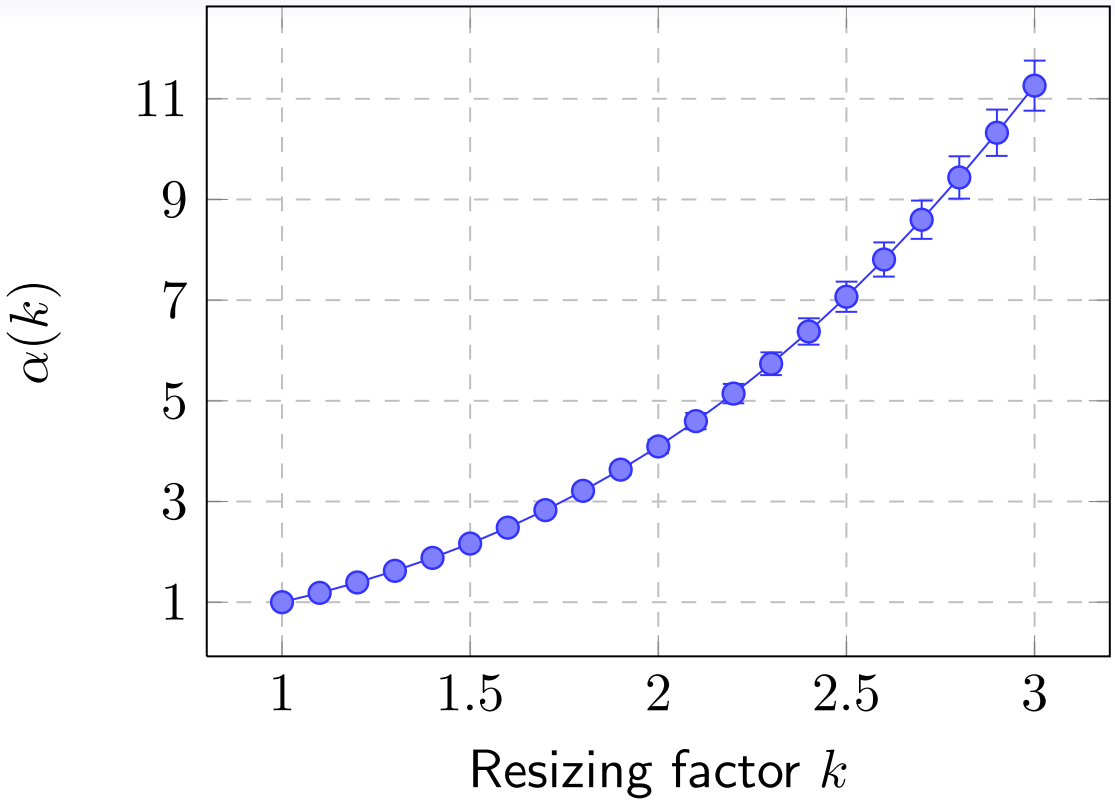

$\alpha(k)$ is the scaling factor by which we need to modify the change rate β at resolution k in order to keep the same level of statistical detectability as with change rate β at full resolution ($k = 1$)

Experimental verification of the scaling law

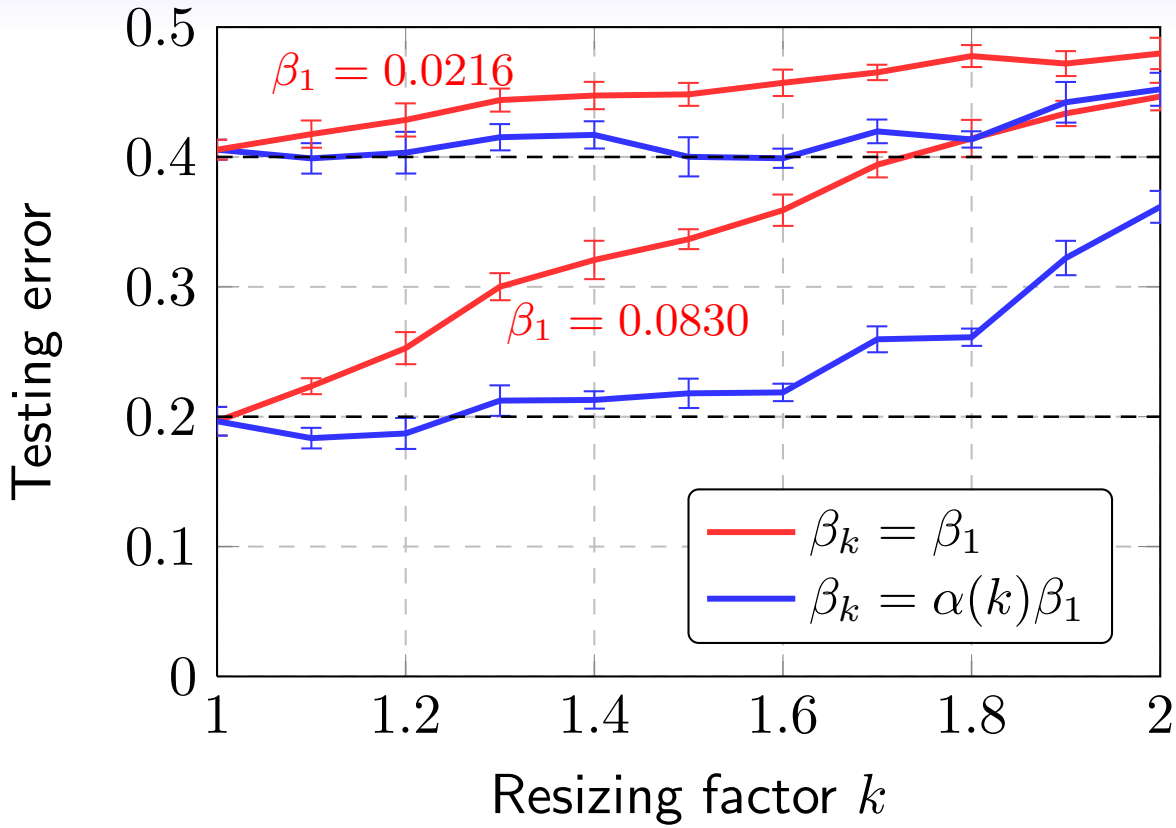

- $\beta_k = \beta_1$ . . . constant change rate
- $\beta_k = \alpha(k)\beta_1$ . . . change rate adjusted w.r.t. derived scaling law

Experimental verification of the scaling law

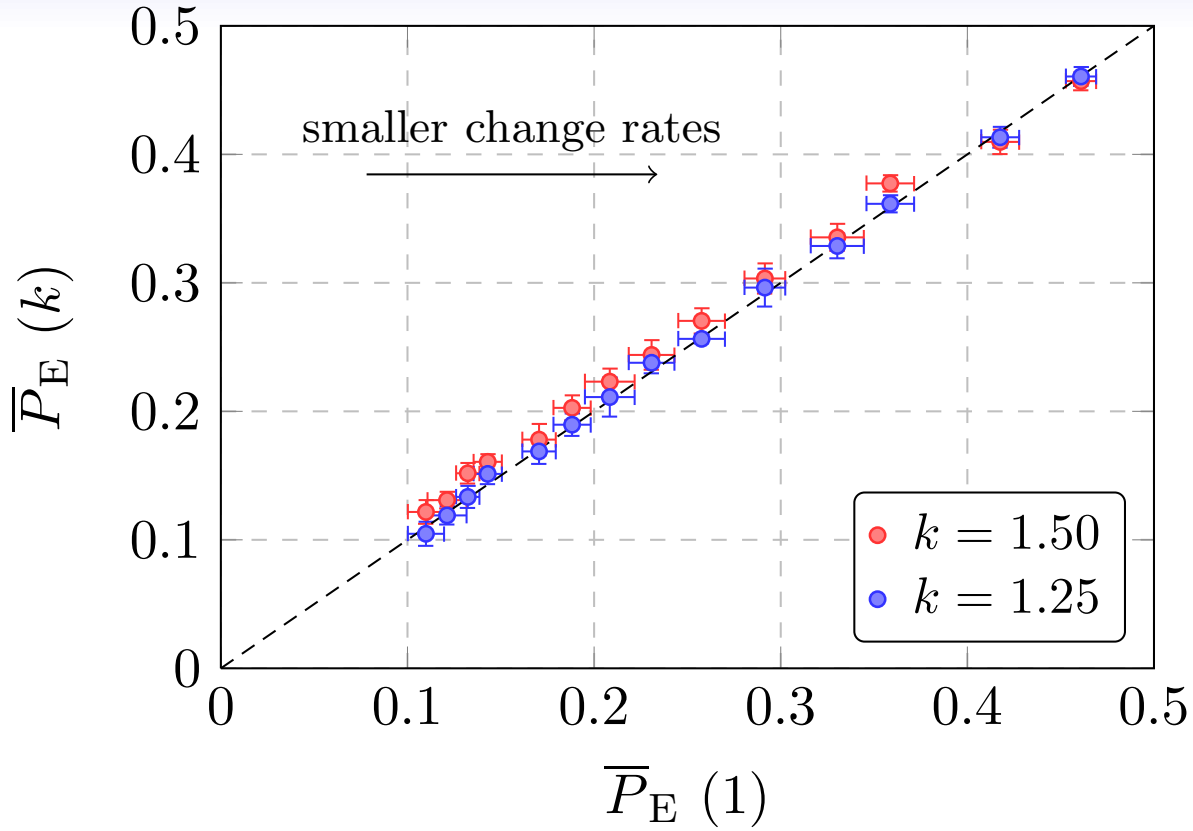

- **x-axis**: $\overline{P}_{\mathrm{E}}(1)$ for full-resolution images, change rate β ($k = 1$)
- **y-axis**: $\overline{P}_{\mathrm{E}}(k)$ for resized images, change rate $\alpha(k)\beta$

Summary and future effort

- Resized images are ubiquitous (also used for benchmarking stego)
- Resizing changes statistical properties of pixels
- Resizing factor, interpolation kernel, and other parameters have a **profound** effect on detectability of embedding changes
- Derived secure payload scaling under resizing for
  - nearest neighbor interpolation
  - Markov chain model of pixel rows
  - Mutually-independent embedding operation
- Experimentally verified to hold for resizing factor ($1 \leq k \lesssim 1.7$)

Journal version (under review)

- Extended experimental section & more camera models
- Effects of anti-aliasing and grid-alignment on security