Batch Steganography and Pooled Steganalysis

Andrew Ker

adk@comlab.ox.ac.uk

Royal Society University Research Fellow
Oxford University Computing Laboratory

8th Information Hiding Workshop
11 July 2006

“The Prisoners’ Problem”

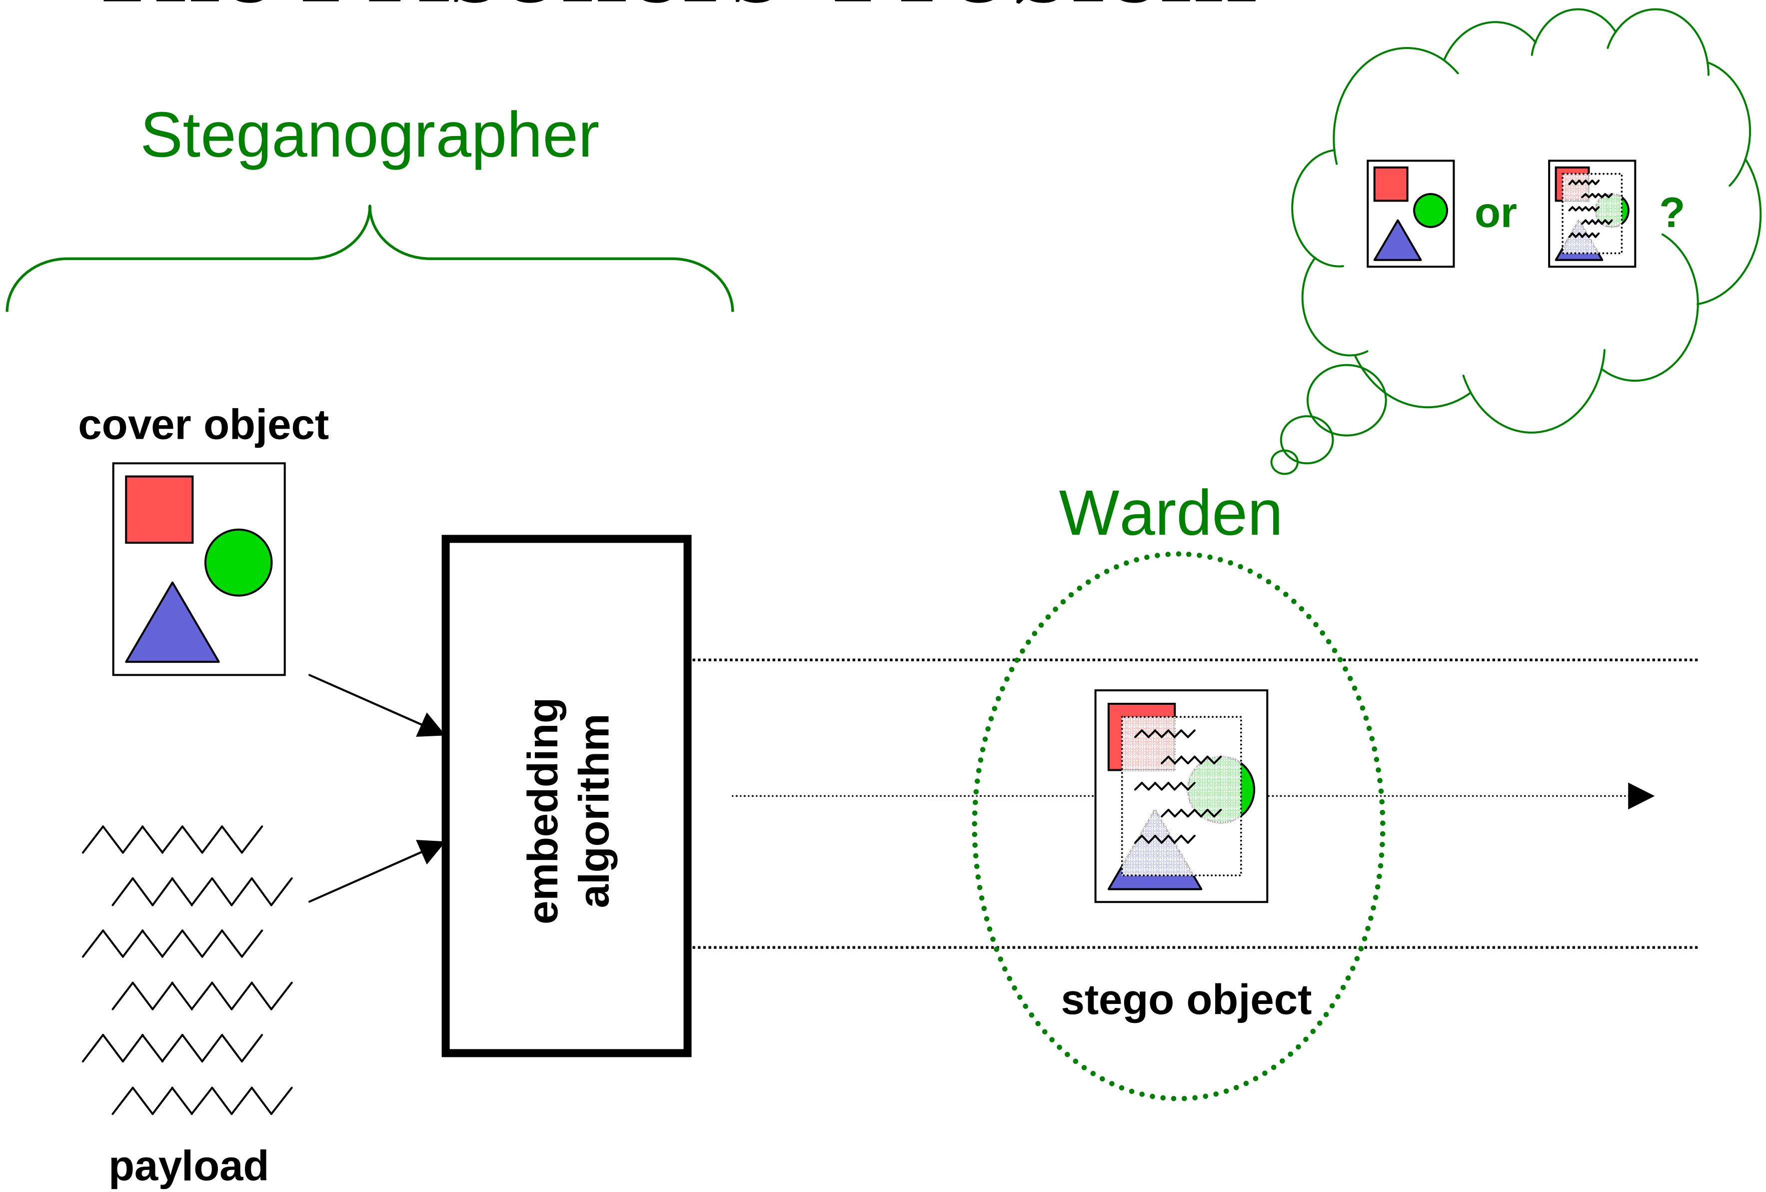

...more realistic?

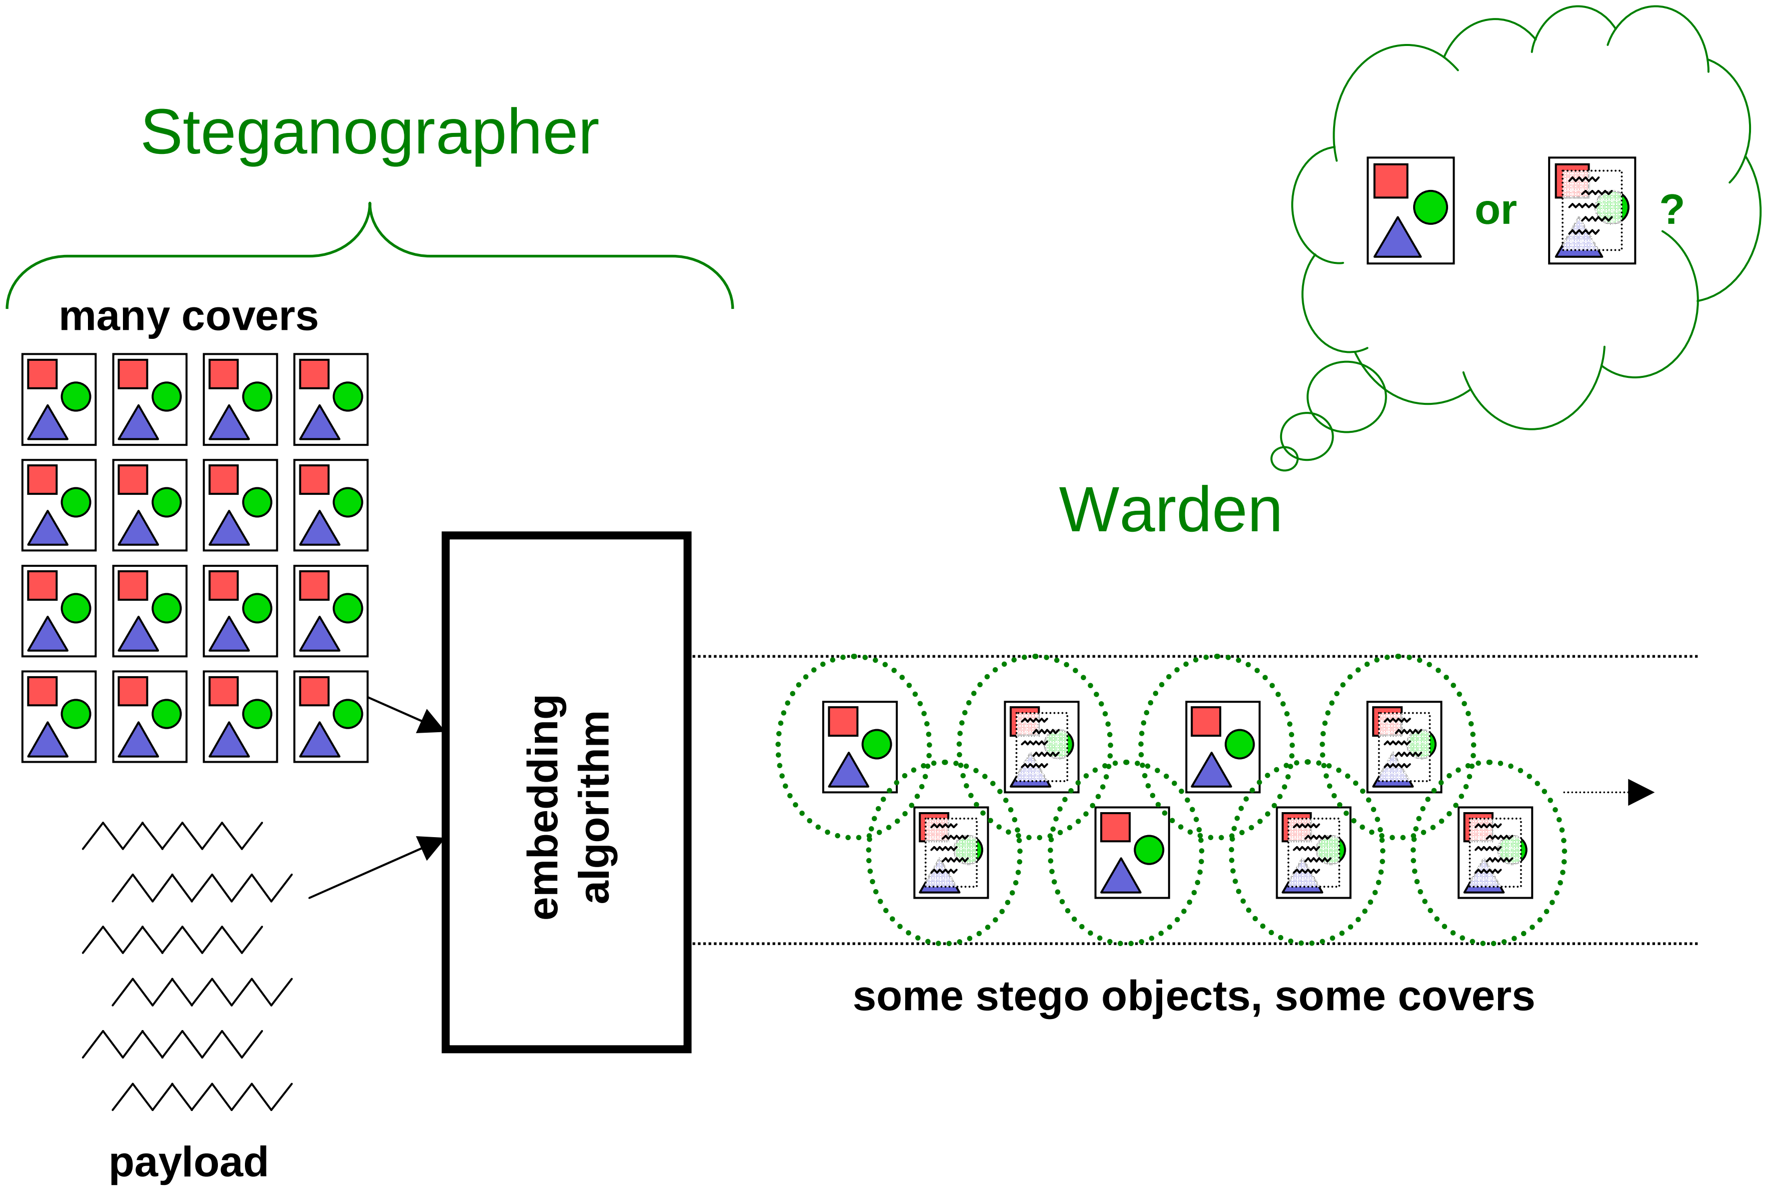

...more realistic?

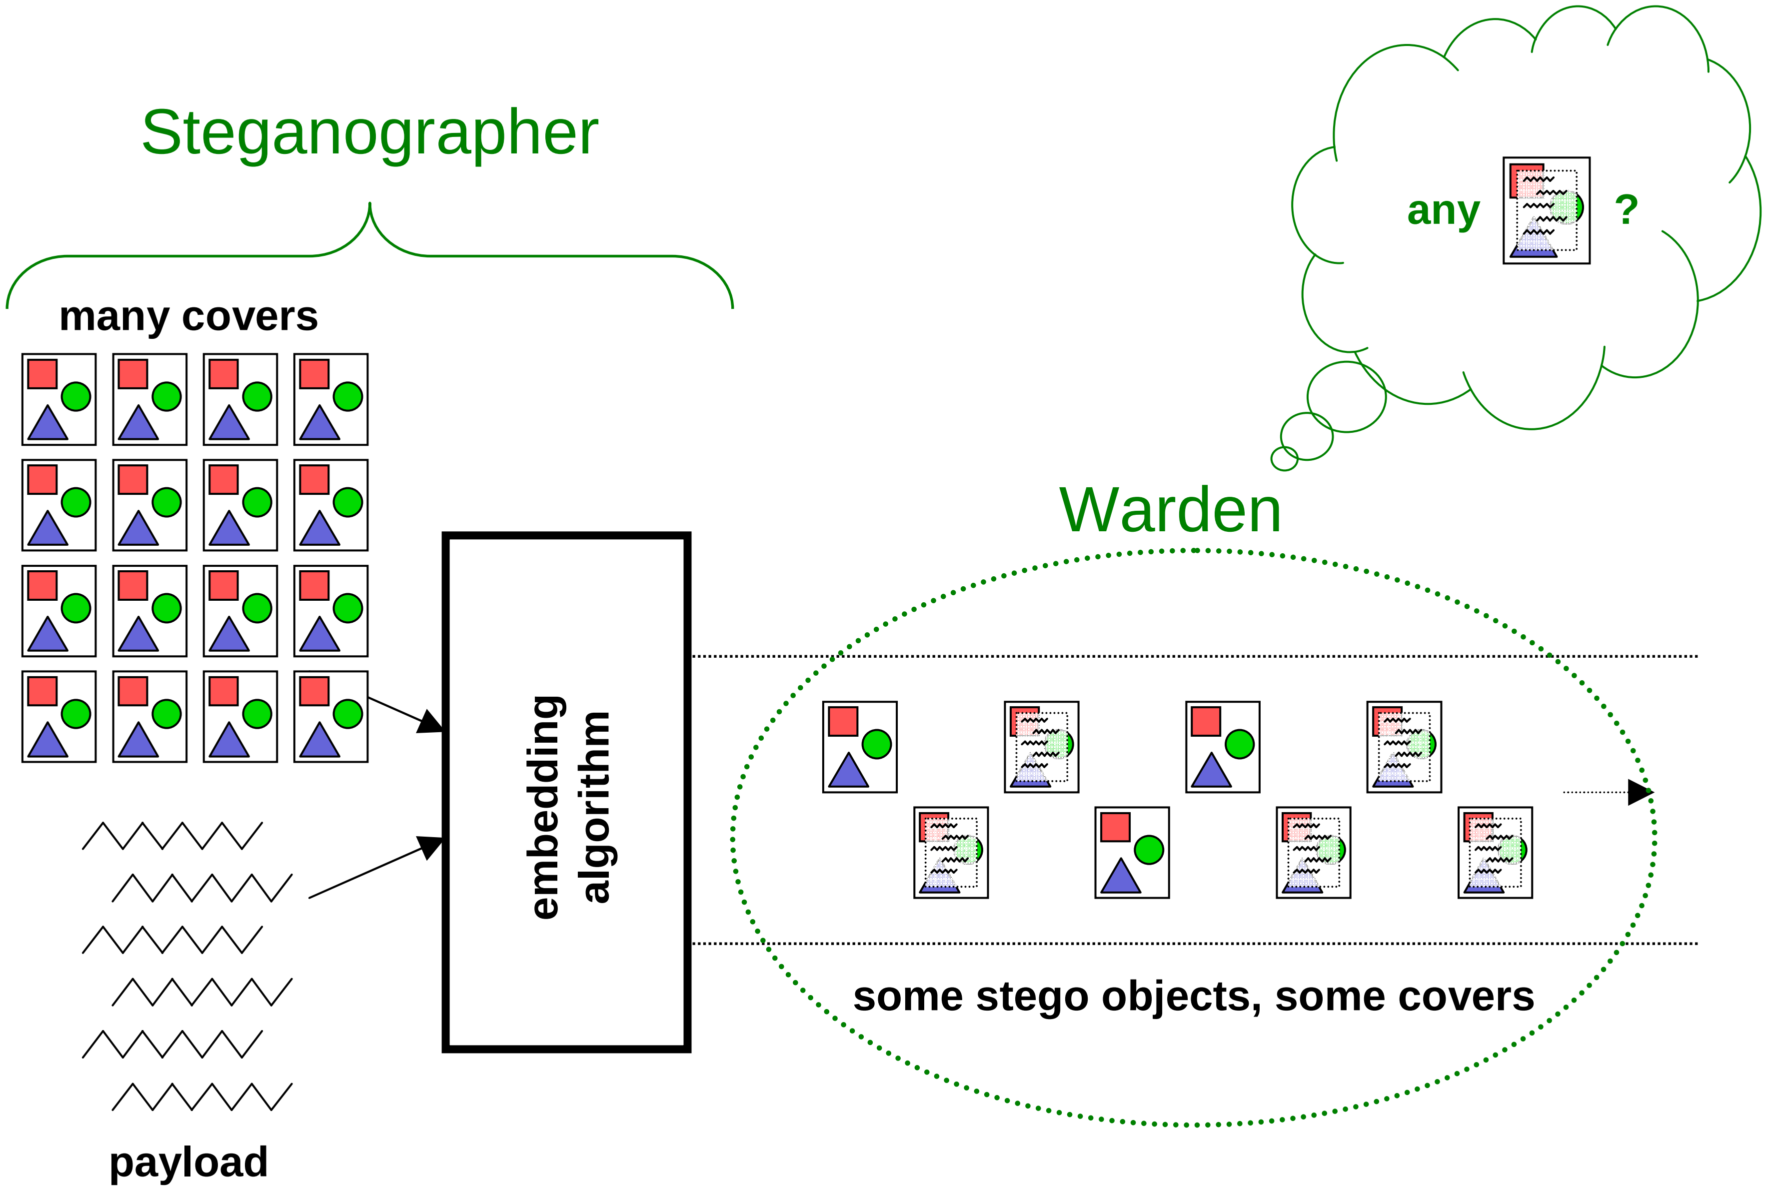

Batch Steganography

The Steganographer:

- has N covers each with same capacity C,
- wants to embed a payload of BNC,

  *$B<1$ is the proportional **bandwidth***

- embeds Cp in each of Nr covers, leaving the other $N(1-r)$ alone.

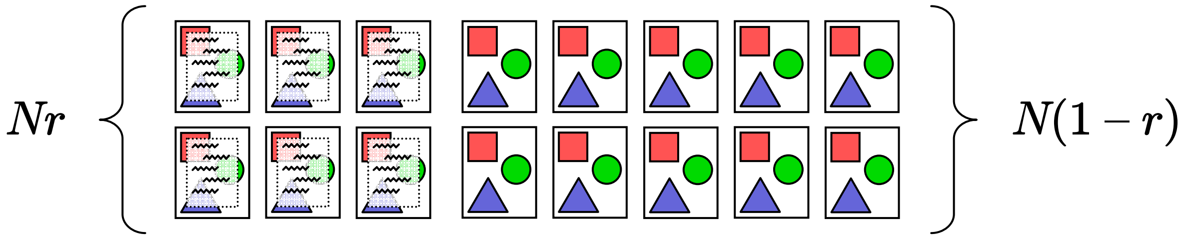

*p is the **proportion of capacity** used when a cover is embedded in*
*r is the **rate** at which covers are used*

constraints: $rp = B \qquad p \leq 1 \qquad r \leq 1$

Pooled Steganalysis

The Warden:

- has a quantitative steganalysis method which estimates the proportionate payload in each cover: $X_1, X_2, \ldots, X_N$

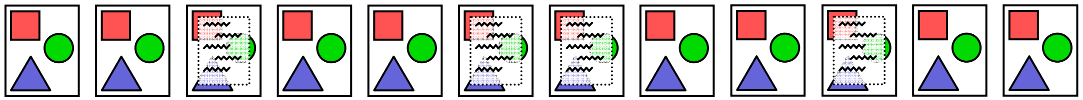

$X_1 \quad X_2 \quad X_3 \quad \ldots \quad X_N$

- wants to pool this evidence to answer the hypothesis test

$$\begin{aligned} H_0 &: r = 0 \\ H_1 &: p, r > 0 \end{aligned}$$

- for now, does not aim to estimate B, r, p or separate individual stego objects from covers.

Assumptions

- N fixed
- The Shift Hypothesis:

  If proportion of capacity p is embedded in cover i,

  $$X_i = p + \epsilon_i$$

  where the error ϵ_i is independent of p

  Will write ψ for error pdf
  Ψ for error cdf

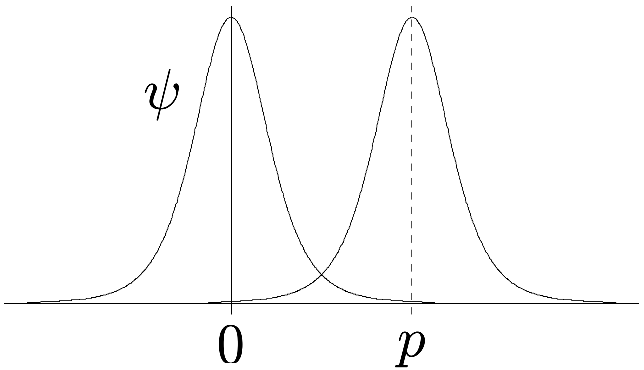

- Assumptions about the shape of ψ:

  "Bell shaped"
  Symmetric about 0
  Unimodal
  Suitably smooth

  But we do **not** assume finite variance

Outline

- Three pooling strategies:

  I: ***Count positive observations***

  II: ***Average observation***

  III: ***Generalised likelihood ratio test***

  for $H_0 : r = 0$
  $H_1 : p, r > 0$

- For each, consider
  - False positive rate @ 50% false negatives,
  - Steganographer's best embedding counterstrategy,
  - How performance depends on B and N.

- Results of some simulation experiments
- Conclusions

I: Count Positive Observations

- Pooled statistic: $\sharp P = |\{ X_i : X_i > 0\}|$

This is just the sign test for whether the median of observed dist is greater than 0

- Null distribution: $H_0 : \sharp P \sim \mathrm{Bi}(N, \frac{1}{2}) \approx \mathrm{N}(\frac{N}{2}, \frac{N}{4})$

- Stego distribution: $H_1 : \sharp P \sim \mathrm{Bi}(N(1-r), \frac{1}{2}) + \mathrm{Bi}(Nr, \Psi(p))$

$$\mathrm{median}(\sharp P) \approx \tfrac{1}{2}N + Nr(\Psi(p) - \tfrac{1}{2})$$

- Median p-value: $\Phi\left(-2BN^{\frac{1}{2}}(\frac{\Psi(p)-\frac{1}{2}}{p})\right)$

An increasing function of p; steganographer should take p=1 r=B

II: Average Observation

- Pooled statistic: $\bar{X} = \frac{1}{N} \sum X_i$
- Null distribution: $H_0 : \bar{X} \stackrel{\cdot}{\sim} \mathrm{N}(0, \sigma^2/N)$
- Stego distribution: $H_1 : \mathrm{median}(\bar{X}) \approx rp = B$
- Median p-value: $\Phi(-\frac{1}{\sigma} B N^{\frac{1}{2}})$

Independent of choice of p

III: Likelihood Ratio

- Pooled statistic: $$\ell = \log \frac{L(X_1, \ldots, X_N; \hat{r}, \hat{p})}{L(X_1, \ldots, X_N; r=0, p=0)}$$ ★

Likelihood function based on *mixture* pdf $f(x) = (1-r)\psi(x) + r\psi(x-p)$

- Null distribution: $\ell \mathrel{\dot\sim} \lambda \chi_d^2$ ★

Theorem [see Appendix]
Under some assumptions... (omitted here)

In the limit as $N \to \infty$, for small B, $\mathrm{E}[\ell]$ is maximized when p=1, r=B, and then ★

$$E[\ell] \sim \frac{NB^2}{2} \int \frac{\psi'(x)^2}{\psi(x)} + \psi''(x)\,\mathrm{d}x$$

- Median (mean) p-value: maximized when p=1, r=B
  function of NB^2

Strategies Summarised

Pooling strategy	*Best steg. strategy*	*False +ve rate at 50% false −ve*	*Total capacity $\propto BN \propto$*
Count positive observations	$p = 1$ $r = B$	decreasing function of $BN^{\frac{1}{2}}$	$N^{\frac{1}{2}}$
Average observation	any	decreasing function of $BN^{\frac{1}{2}}$	$N^{\frac{1}{2}}$ ★
Generalised Likelihood Ratio Test (ψ known)	$p = 1$ $r = B$ (for small B)	decreasing function of B^2N	$N^{\frac{1}{2}}$

Experimental Results

- Covers: A set of 14000 grayscale images
- Steganography: LSB Replacement
- Steganalysis: "Sample Pairs" [Dumitrescu, IHW 2002]
- N=10, 100, 1000

For a random batch of size N, compute $\sharp P, \bar{X}, \ell$

5000 ***samples with no steganography, to fit null distributions***

500 ***samples each with a range of*** p, r ***such that*** $rp=B=0.01$

Measure false positive rate @ 50% false negatives

Experimental Results: $B = 0.01$

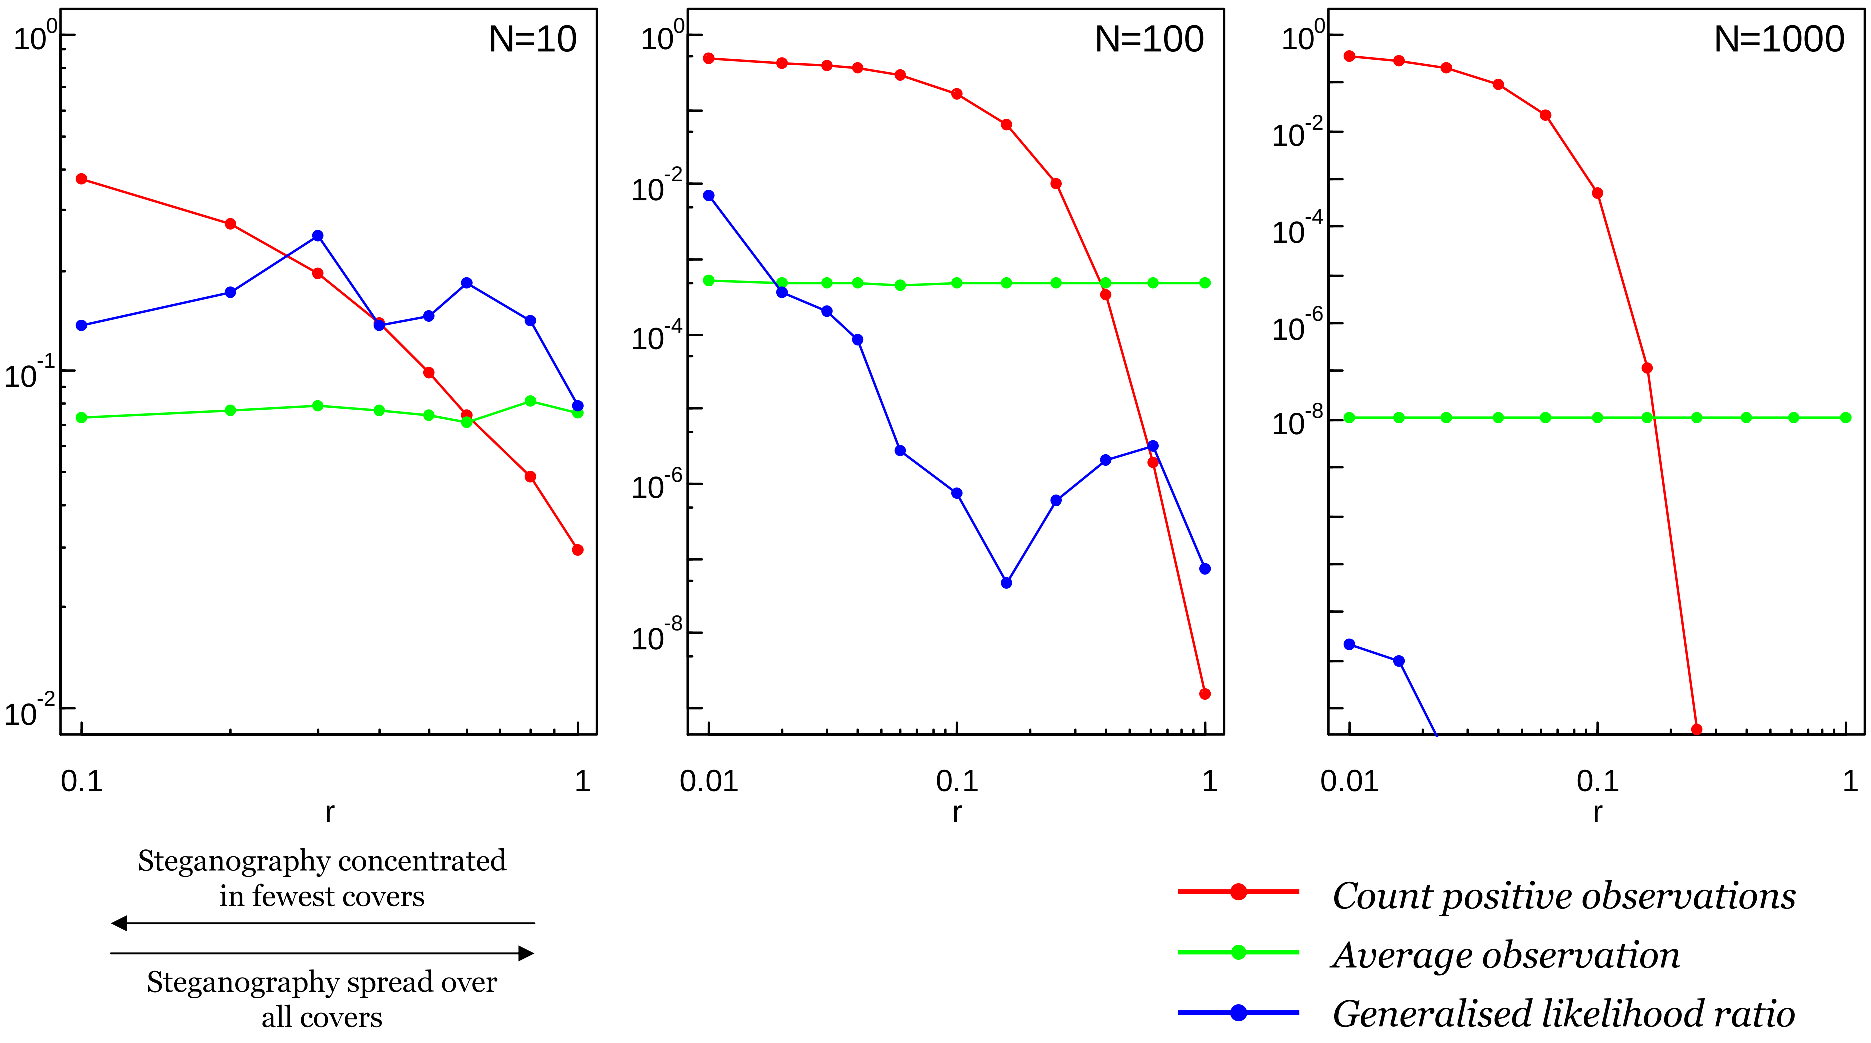

Not in this talk

- ★Technical statistical difficulties.
- Empirical investigation of relationship between B and N.
- A critical problem: bias in the quantitative steganalysis method.

Further Work

- Other strategies for Warden

  e.g. "count observations greater than some threshold t"

- Try to relax some of the assumptions

  Uniformity of covers/embedding
  Shift hypothesis

Conclusions

- Batch steganography and pooled steganalysis are interesting and relevant problems.

  Complicated by the plethora of possible pooling strategies for the Warden.
  Mathematical analysis can be intractable.

- Common theme: B should shrink as N grows, for fixed risk.

  Conjecture: Steganographic capacity is proportional to the ***square root*** *of the total cover size.*

- Common theme: Steganographer should concentrate the steganography.

  Not true for all pooling strategies!
  Nonetheless, seems to be true for all "sensible" pooling strategies...
  Lessons for adaptive embedding?

The End

adk@comlab.ox.ac.uk