

Capacity of (Imperfect) Stegosystems



Andrew Ker

adk@comlab.ox.ac.uk

Oxford University Computing Laboratory

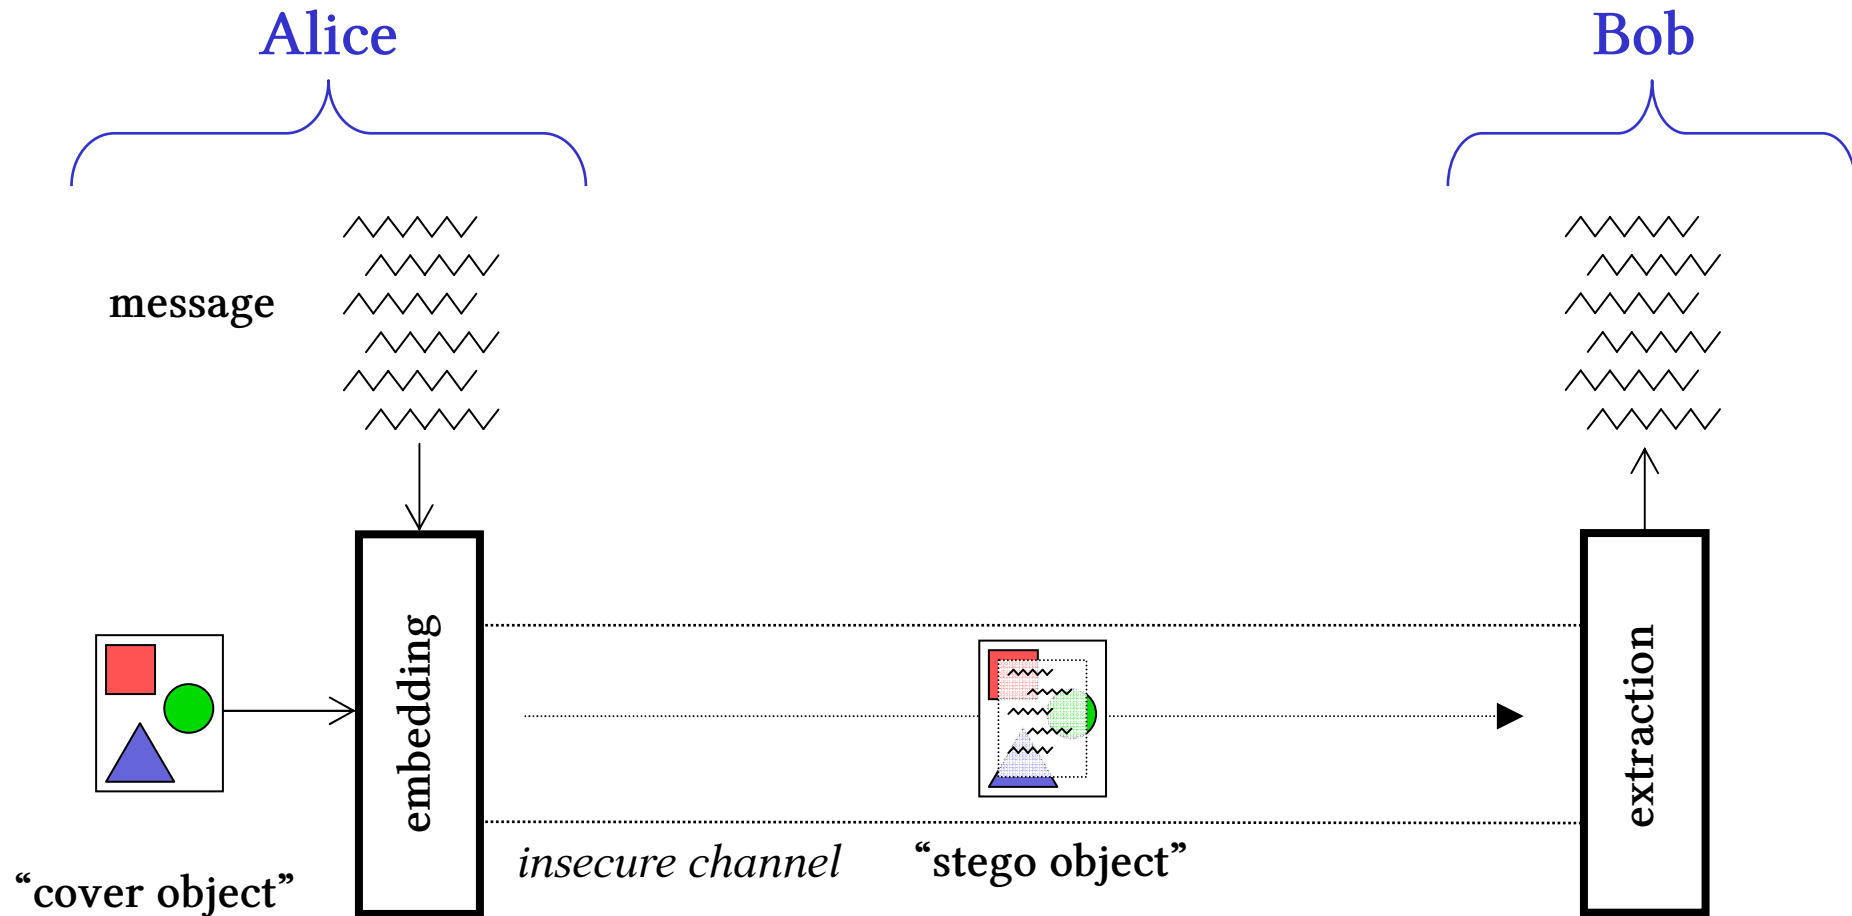
Steganography & Steganalysis seminar, Paris, 25 March 2010

Capacity of (Imperfect) Stegosystems

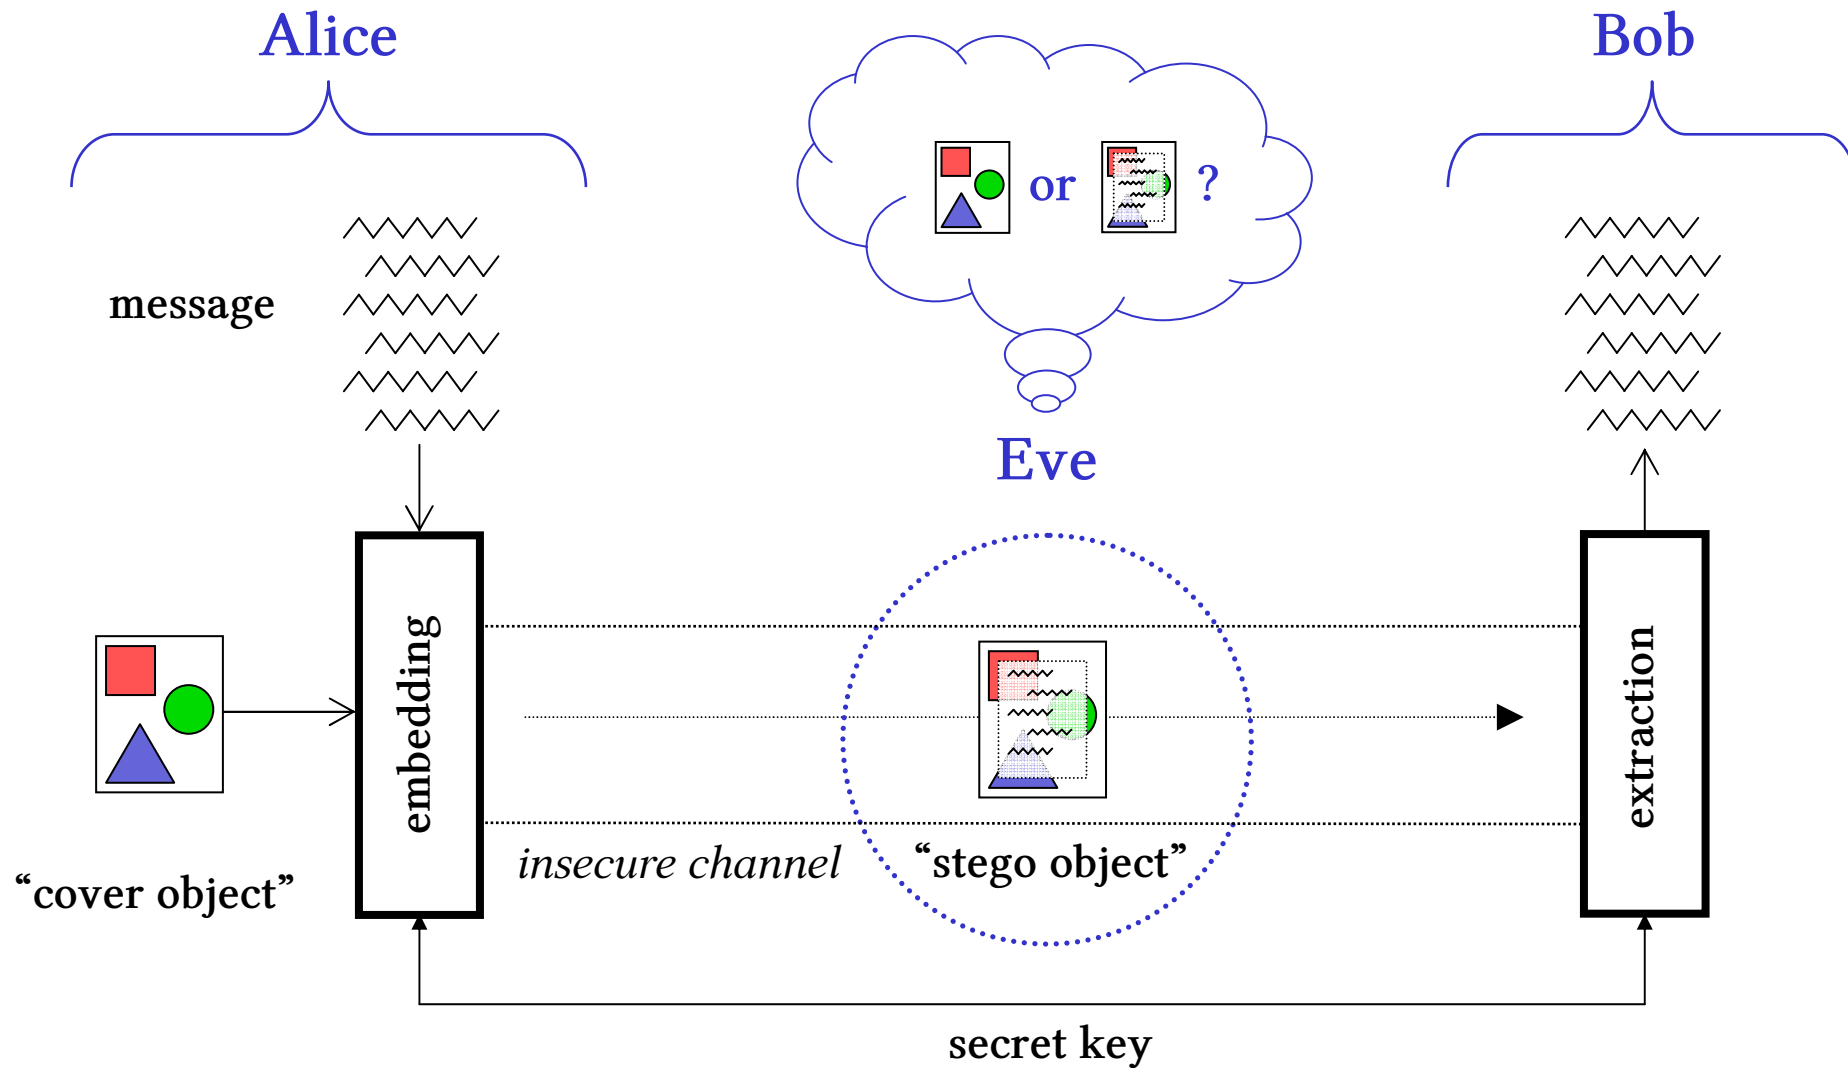
Outline

- Steganography; example
- Perfect and imperfect stegosystems
 - Approaches to capacity
- The Square Root Law
 - Proof of simplified theorem
 - Extensions & further work
- Batch steganography & pooled steganalysis

Steganography





Steganography








cover: 512×512 raw bitmap, 24 bit colour













































Perfect steganography

A stegosystem is **perfect** if the distribution of the stego objects matches exactly the distribution of the covers.

In this case,

- the embedding is undetectable;
 - the **capacity** is linear in the amount of data transmitted.
-
- *We must consider the possibility of dependence between the objects in the stego and cover channels.*
 - *How can we ever know if a real-world stegosystem is perfect?*

Perfect steganography

Perfect stegosystems are theoretically possible:

1. If the embedder knows exactly the distribution of the covers, they can match it.

- Various methods exist, but may require a lot of work at the embedder.
- Is it realistic to know the exact distribution of the covers?

Perfect steganography

Perfect stegosystems are theoretically possible:

1. If the embedder knows exactly the distribution of the covers, they can match it.

- Various methods exist, but may require a lot of work at the embedder.
- Is it realistic to know the exact distribution of the covers?

2. If the embedder has unlimited access to covers, they can sample until they find a sequence which match their message.

- This is the “rejection sampler”, which requires a lot of work.
- The payload would have to be very small.

Capacity of imperfect steganography

Not “how much can you hide undetectably”,
rather “how much is reliably detectable”?

Fix:

- cover source,
- embedding method,
- limit on “risk” (minimum detector error).

What is the largest payload which can safely be embedded?

Could use it to compare: embedding methods,
covers,
detectors, ...

But difficult to answer – must reason about EVERY detector.

Information theoretic bounds

If X has density function f , and Y has density function g , then the *Kullback-Leibler divergence* from X to Y is

$$D_{\text{KL}}(X, Y) = - \int f(x) \log \frac{g(x)}{f(x)} dx$$

Information Processing Theorem: $D_{\text{KL}}(h(X), h(Y)) \leq D_{\text{KL}}(X, Y)$

Therefore, if trying to classify an observation as X or Y , the error rates α and β must satisfy

$$\alpha \log \frac{\alpha}{1 - \beta} + (1 - \alpha) \log \frac{1 - \alpha}{\beta} \leq D_{\text{KL}}(X, Y).$$

Information theoretic bounds

$$\alpha \log \frac{\alpha}{1 - \beta} + (1 - \alpha) \log \frac{1 - \alpha}{\beta} \leq D_{\text{KL}}(X, Y)$$

Moulin *et al.* – proposes statistical models for cover and stego media to compute $D_{\text{KL}}(\text{cover objects}, \text{stego objects})$ in terms of payload size.

Conclusion: If the embedding rate is fixed then the probability of false negative (missed detection) tends to zero exponentially fast.

Security can be measured by the “error exponent”.

Information theoretic bounds

$$\alpha \log \frac{\alpha}{1 - \beta} + (1 - \alpha) \log \frac{1 - \alpha}{\beta} \leq D_{\text{KL}}(X, Y)$$

Moulin *et al.* – proposes statistical models for cover and stego media to compute D_{KL} (cover objects, stego objects) in terms of payload size.

Conclusion: If the embedding rate is fixed then the probability of false negative (missed detection) tends to zero exponentially fast.

Security can be measured by the “error exponent”.

Problem: Statistical models for covers are (very) inaccurate.

> distinction between **artificial** covers – mathematical objects
and **empirical** covers – realisations of reality.

Information theoretic bounds

$$\alpha \log \frac{\alpha}{1 - \beta} + (1 - \alpha) \log \frac{1 - \alpha}{\beta} \leq D_{\text{KL}}(X, Y)$$

Moulin *et al.* – proposes statistical models for cover and stego media to compute $D_{\text{KL}}(\text{cover objects}, \text{stego objects})$ in terms of payload size.

Conclusion: If the embedding rate is fixed then the probability of false negative (missed detection) tends to zero exponentially fast.

Security can be measured by the “error exponent”.

Problem: Statistical models for covers are (very) inaccurate.

Problem: If fixed-rate embedding leads to certain detection, why do it?

The Square Root Law

“The amount of information you can hide securely is asymptotically proportional to the square root of the space you have to hide it in.”

- Suggested sublinear capacity after empirical study in 2004.
- Conjectured square root relationship in 2005.
- First proved a square root law in 2006.
- New and improved square root laws in 2008, 2009, 2010, ...

SRL Theorem

- Assuming
- Covers consist of n independent random bits (pixels) (X_1, \dots, X_n) .
 - Payload, size m , affects pixels independently with probability m/n .
 - Unaltered pixels are 1 with probability p , pixels used for payload are 1 with probability q .
 - $p \neq 0, 1, \quad p \neq q$.
1. If $m/\sqrt{n} \rightarrow \infty$ as $n \rightarrow \infty$ then, for sufficiently large n , an arbitrarily accurate detector exists.
 2. If $m/\sqrt{n} \rightarrow 0$ as $n \rightarrow \infty$ then, for sufficiently large n , every detector must have an arbitrarily high inaccuracy.

Proof

1. If $m/\sqrt{n} \rightarrow \infty$ then an arbitrarily accurate detector exists.

W.l.o.g. $p < q$. The detector is

payload detected if $Y = \#\{X_i = 1\}$ is greater than a critical threshold $Y^ = np + c\sqrt{n}$.*

If no payload, $Y \sim \text{Bi}(n, p)$, and

$$P(Y > Y^*) = P(Y - E[Y] > c\sqrt{n}) \leq \exp(-2c^2)$$

which can be made arbitrarily small.

↑
{ Hoeffding's inequality }

If payload, $Y \sim \text{Bi}(n, \frac{m}{n}q + (1 - \frac{m}{n})p)$,

$$P(Y \leq Y^*) = P(Y - E[Y] \leq c\sqrt{n} - (q - p)m) = O(\exp(-\frac{m^2}{n}))$$

which tends to zero as $n \rightarrow \infty$.

Proof

2. If $m/\sqrt{n} \rightarrow 0$ then detectors must have an arbitrarily high inaccuracy.

Without payload, “1” has probability p ; with payload, $\frac{m}{n}q + (1 - \frac{m}{n})p$.

By the information processing theorem, *any* detector has false positive rate α and false negative rate β satisfying

$$\begin{aligned}\alpha \log \frac{\alpha}{1-\beta} + (1-\alpha) \log \frac{1-\alpha}{\beta} &\leq D_{\text{KL}}(\mathbf{X} \mid \text{no embedding}, \mathbf{X} \mid \text{embedding}) \\ &= nD_{\text{KL}}(X_1 \mid \text{no embedding}, X_1 \mid \text{embedding}) \\ &= np \log \frac{\frac{m}{n}q + (1 - \frac{m}{n})p}{p} + n(1-p) \log \frac{1 - \frac{m}{n}q - (1 - \frac{m}{n})p}{1-p}\end{aligned}$$

$$\begin{aligned}\{ \log(1+z) \geq z - z^2 \} &\longrightarrow \leq n \left(\frac{m}{n} \right)^2 \left(\frac{(q-p)^2}{p} + \frac{(q-p)^2}{1-p} \right) \\ &\longrightarrow 0.\end{aligned}$$

Therefore $\alpha \rightarrow 1 - \beta$.

The Square Root Law

“The amount of information you can hide securely is asymptotically proportional to the square root of the space you have to hide it in.”

- Information hiding is unlike cryptography or communication theory.
- Everything based on embedding “rate” should be reconsidered...

Extensions

1. More realistic cover models.

SRL Theorem

- Assuming
- Covers consist of n independent random bits (pixels) (X_1, \dots, X_n) .
 - Payload, size m , affects pixels independently with probability m/n .
 - Unaltered pixels are 1 with probability p , pixels used for payload are 1 with probability q .
 - $p \neq 0, 1, \quad p \neq q$.
1. If $m/\sqrt{n} \rightarrow \infty$ as $n \rightarrow \infty$ then, for sufficiently large n , an arbitrarily accurate detector exists.
 2. If $m/\sqrt{n} \rightarrow 0$ as $n \rightarrow \infty$ then, for sufficiently large n , every detector must have an arbitrarily high inaccuracy.

SRL Theorem

Assuming

- Covers consist of n independent random bits (pixels) (X_1, \dots, X_n) .

Unrealistic (artificial) cover model.

- Immediate extension to covers of independent random pixels of finitely many colours.
- With difficult analysis, can be extended further...

1. If $m/\sqrt{n} \rightarrow \infty$ as $n \rightarrow \infty$ then, for sufficiently large n , an arbitrarily accurate detector exists.
2. If $m/\sqrt{n} \rightarrow 0$ as $n \rightarrow \infty$ then, for sufficiently large n , every detector must have an arbitrarily high inaccuracy.

SRL Theorem

- Assuming
- Covers consist of n realisations from a Markov chain (X_1, \dots, X_n) .
 - Payload, size m , affects pixels independently with probability m/n .
 - (The Markov chain is nontrivial.)
 - (The stego object has a different distribution from the covers.)
1. If $m/\sqrt{n} \rightarrow \infty$ as $n \rightarrow \infty$ then, for sufficiently large n , an arbitrarily accurate detector exists.
 2. If $m/\sqrt{n} \rightarrow 0$ as $n \rightarrow \infty$ then, for sufficiently large n , every detector must have an arbitrarily high inaccuracy.

SRL Theorem

Assuming • Covers consist of n realisations from a Markov chain (X_1, \dots, X_n) .

Conjecture: holds for all Markov random fields with “exponential forgetting” property.

1. If $m/\sqrt{n} \rightarrow \infty$ as $n \rightarrow \infty$ then, for sufficiently large n , an arbitrarily accurate detector exists.
2. If $m/\sqrt{n} \rightarrow 0$ as $n \rightarrow \infty$ then, for sufficiently large n , every detector must have an arbitrarily high inaccuracy.


Extensions

1. More realistic cover models.
2. Consider the “secret key” size.

SRL Theorem

Assuming

- Covers consist of n independent random bits (pixels) (X_1, \dots, X_n) .

- Payload, size m , affects pixels independently with probability m/n .
- 

Slightly unrealistic embedding model:

- “Uniform” embedding is not “independent” embedding.
- If the payload is of a certain size, embedding in location i means that embedding in location j is marginally less likely.

So even in the i.i.d. cover model, the stego images should not consist of i.i.d. pixels.

SRL Theorem

- Assuming
- Covers consist of n independent random bits (pixels) (X_1, \dots, X_n) .
 - Payload affects m locations chosen uniformly from all possible embedding paths.
 - Unaltered pixels are 1 with probability p , pixels used for payload are 1 with probability q .
 - $p \neq 0, 1, \quad p \neq q$.
1. If $m/\sqrt{n} \rightarrow \infty$ as $n \rightarrow \infty$ then, for sufficiently large n , an arbitrarily accurate detector exists.
 2. If $m/\sqrt{n} \rightarrow 0$ as $n \rightarrow \infty$ then, for sufficiently large n , every detector must have an arbitrarily high inaccuracy.

SRL Theorem

- Assuming
- Covers consist of n independent random bits (pixels) (X_1, \dots, X_n) .
 - Payload affects m locations chosen uniformly from all possible embedding paths.

Still unrealistic!

There are $\frac{n!}{(n-m+1)!}$ possible embedding paths (choose m ordered locations from n , without replacement).

If $m \sim \sqrt{n}$, sender and recipient need to share $\log \frac{n!}{(n-\sqrt{n})!} \sim m \log m$ bits of information to locate the payload, i.e.

they need a secret key larger than the payload transmitted!

SRL Theorem

- Assuming
- Covers consist of n independent random bits (pixels) (X_1, \dots, X_n) .
 - Payload affects m locations chosen from a set of 2^k possible embedding paths (i.e. a secret key of k bits).
 - Unaltered pixels are 1 with probability p , pixels used for payload are 1 with probability q .
 - $p \neq 0, 1, \quad p \neq q$.

Then

if $k/m \rightarrow 0$ and $m \rightarrow \infty$ as $n \rightarrow \infty$ then, for sufficiently large n , an arbitrarily accurate detector exists.

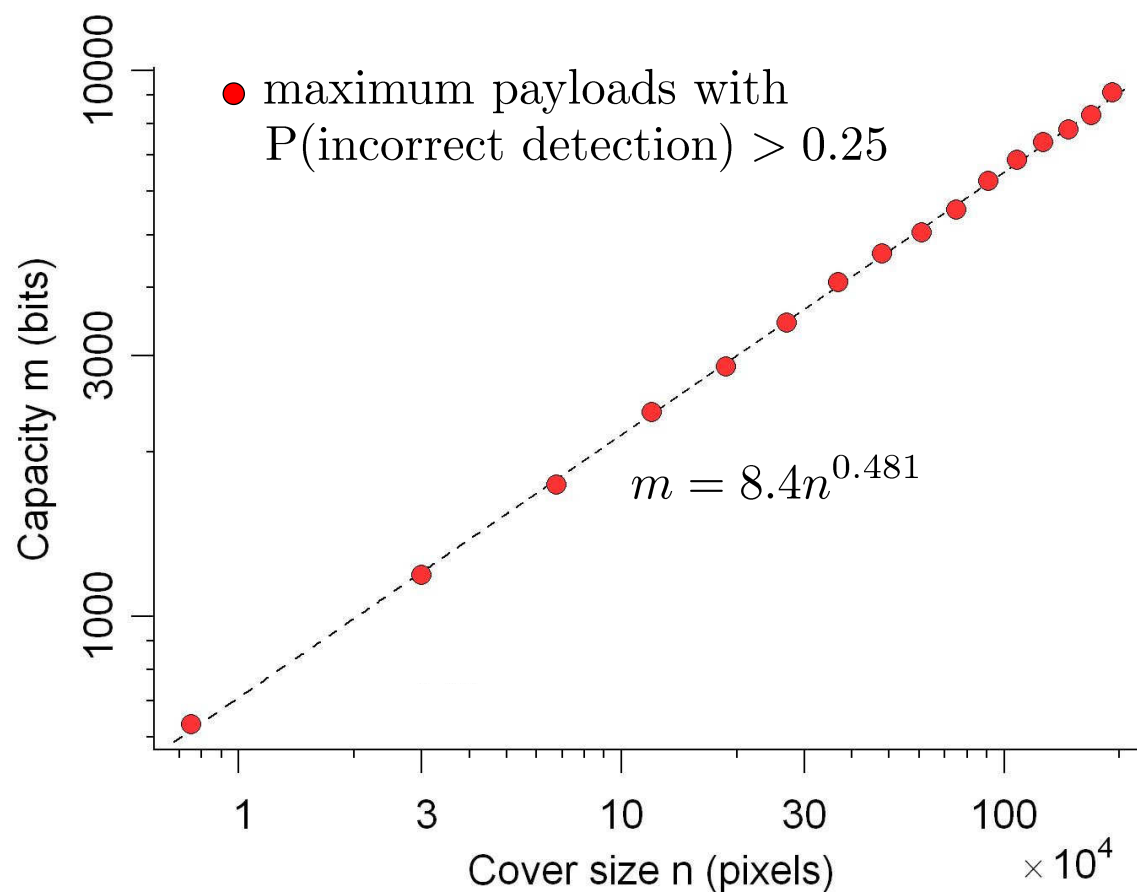
“ k must be at least linear in m if a square root law is to hold”

Extensions

1. More realistic cover models.
2. Consider the “secret key” size.
3. Empirical studies.

Empirical studies

Square root law observed in performance of contemporary detectors.



A. Ker, T. Pevný, J. Kodovský, and J. Fridrich. *The Square Root Law of Steganographic Capacity*.
In Proc. 10th Multimedia and Security Workshop, ACM, 2008

Extensions

1. More realistic cover models.
2. Consider the “secret key” size.
3. Empirical studies.
4. Estimation of multiplicative constant *root rate*.

SRL Theorem

- Assuming
- Covers consist of n independent random bits (pixels) (X_1, \dots, X_n) .
 - Payload, size m , affects pixels independently with probability m/n .
 - Unaltered pixels are 1 with probability p , pixels used for payload are 1 with probability q .
 - $p \neq 0, 1, \quad p \neq q$.
1. If $m/\sqrt{n} \rightarrow \infty$ as $n \rightarrow \infty$ then, for sufficiently large n , an arbitrarily accurate detector exists.
 2. If $m/\sqrt{n} \rightarrow 0$ as $n \rightarrow \infty$ then, for sufficiently large n , every detector must have an arbitrarily high inaccuracy.
 3. If $m \sim r\sqrt{n}$ as $n \rightarrow \infty$..?

SRL Theorem

- Assuming
- Covers consist of n independent random bits (pixels) (X_1, \dots, X_n) .
 - Payload, size m , affects pixels independently with probability m/n .
 - Unaltered pixels are 1 with probability p , pixels used for payload are 1 with probability q .

The root rate r limits the asymptotic performance of any detector. It is related to Fishers Information I , by

$$r \propto 1/\sqrt{I}.$$

Lower bounds on false positive & negative rates determine upper bounds on r , which is the true “capacity”.

3. If $m \sim r\sqrt{n}$ as $n \rightarrow \infty$..?



Extensions

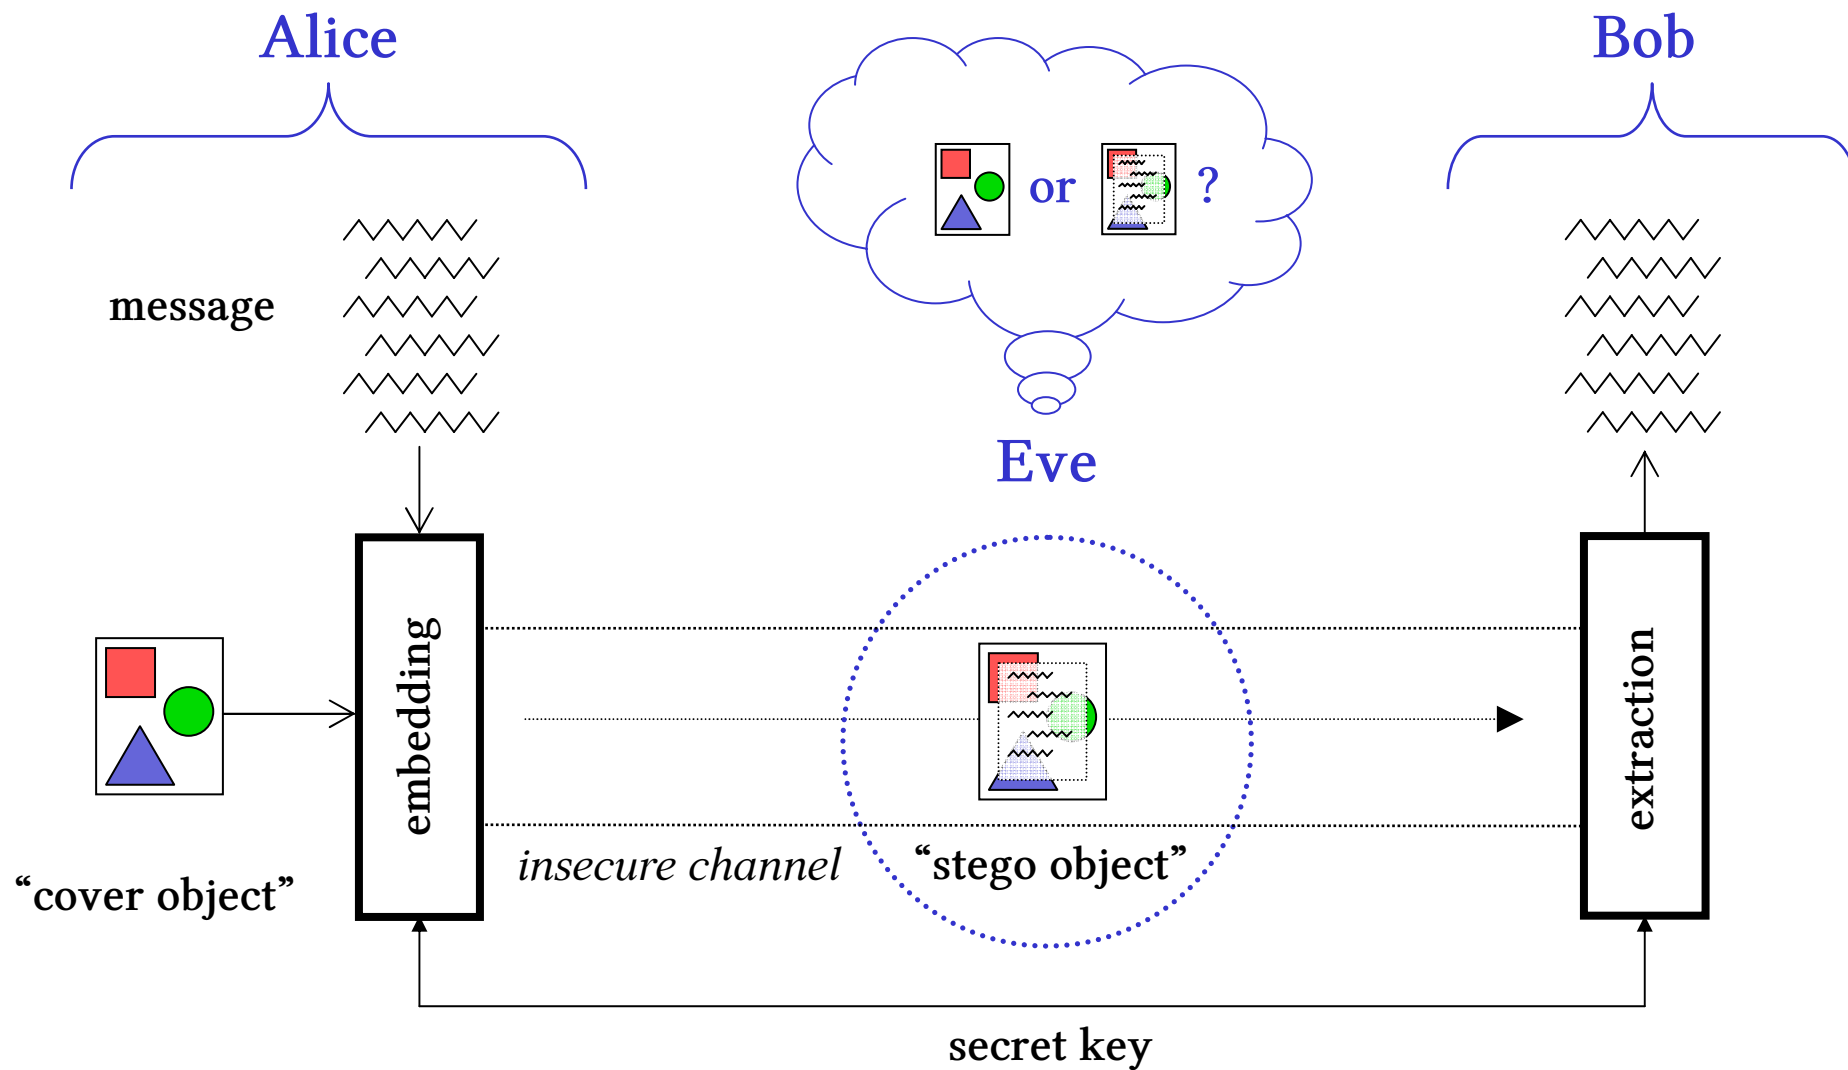
1. More realistic cover models.
2. Consider the “secret key” size.
3. Empirical studies.
4. Estimation of multiplicative constant *root rate*.

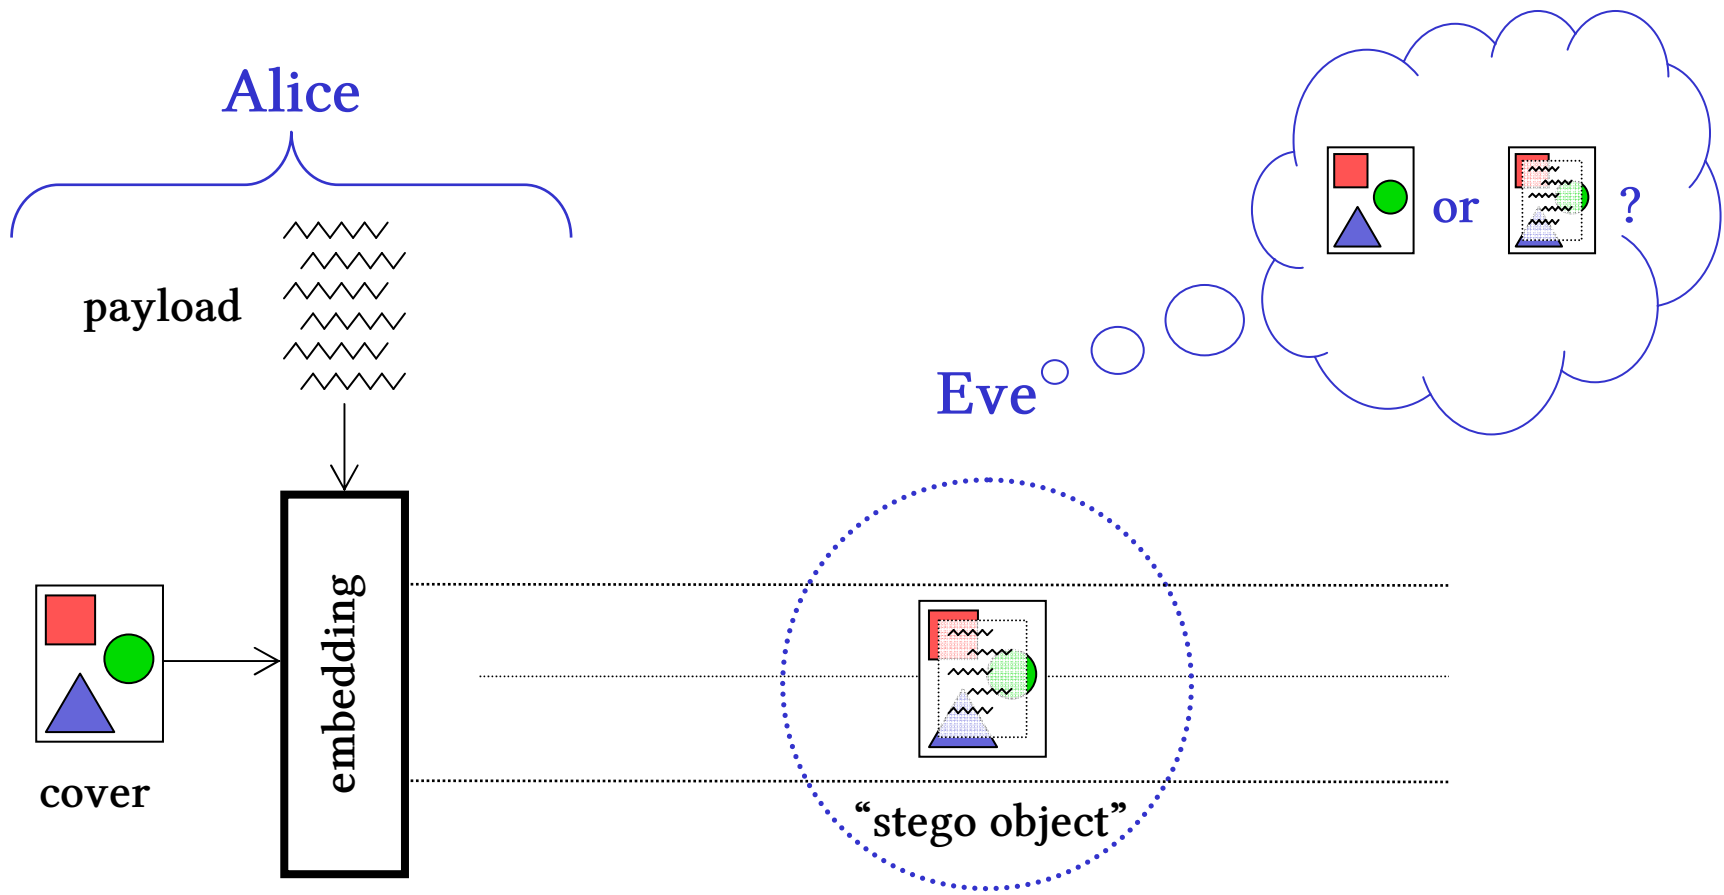
*We can even estimate Fisher Information, hence root rate, empirically.
But the computational challenges are considerable.*

Extensions

1. More realistic cover models.
2. Consider the “secret key” size.
3. Empirical studies.
4. Estimation of multiplicative constant *root rate*.
5. Detectors with an imperfect cover (or stego) model.
6. Embedders with learning behaviour.

Batch steganography






Alice

payload
size M



N covers



$embed\ p_1$



$embed\ p_2$



$embed\ p_3$

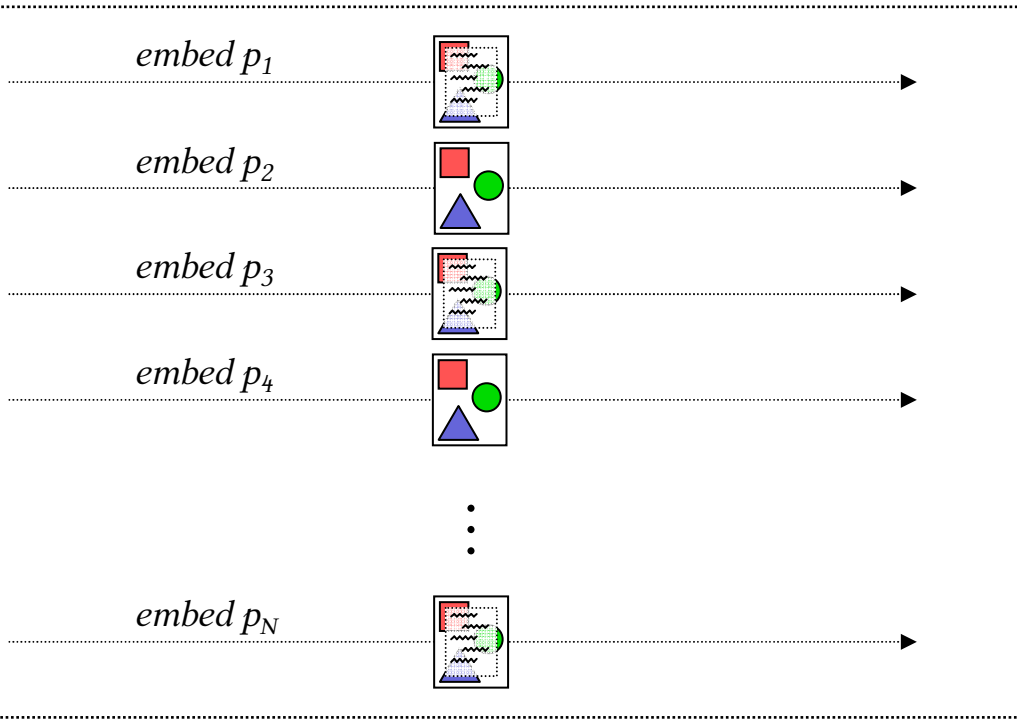


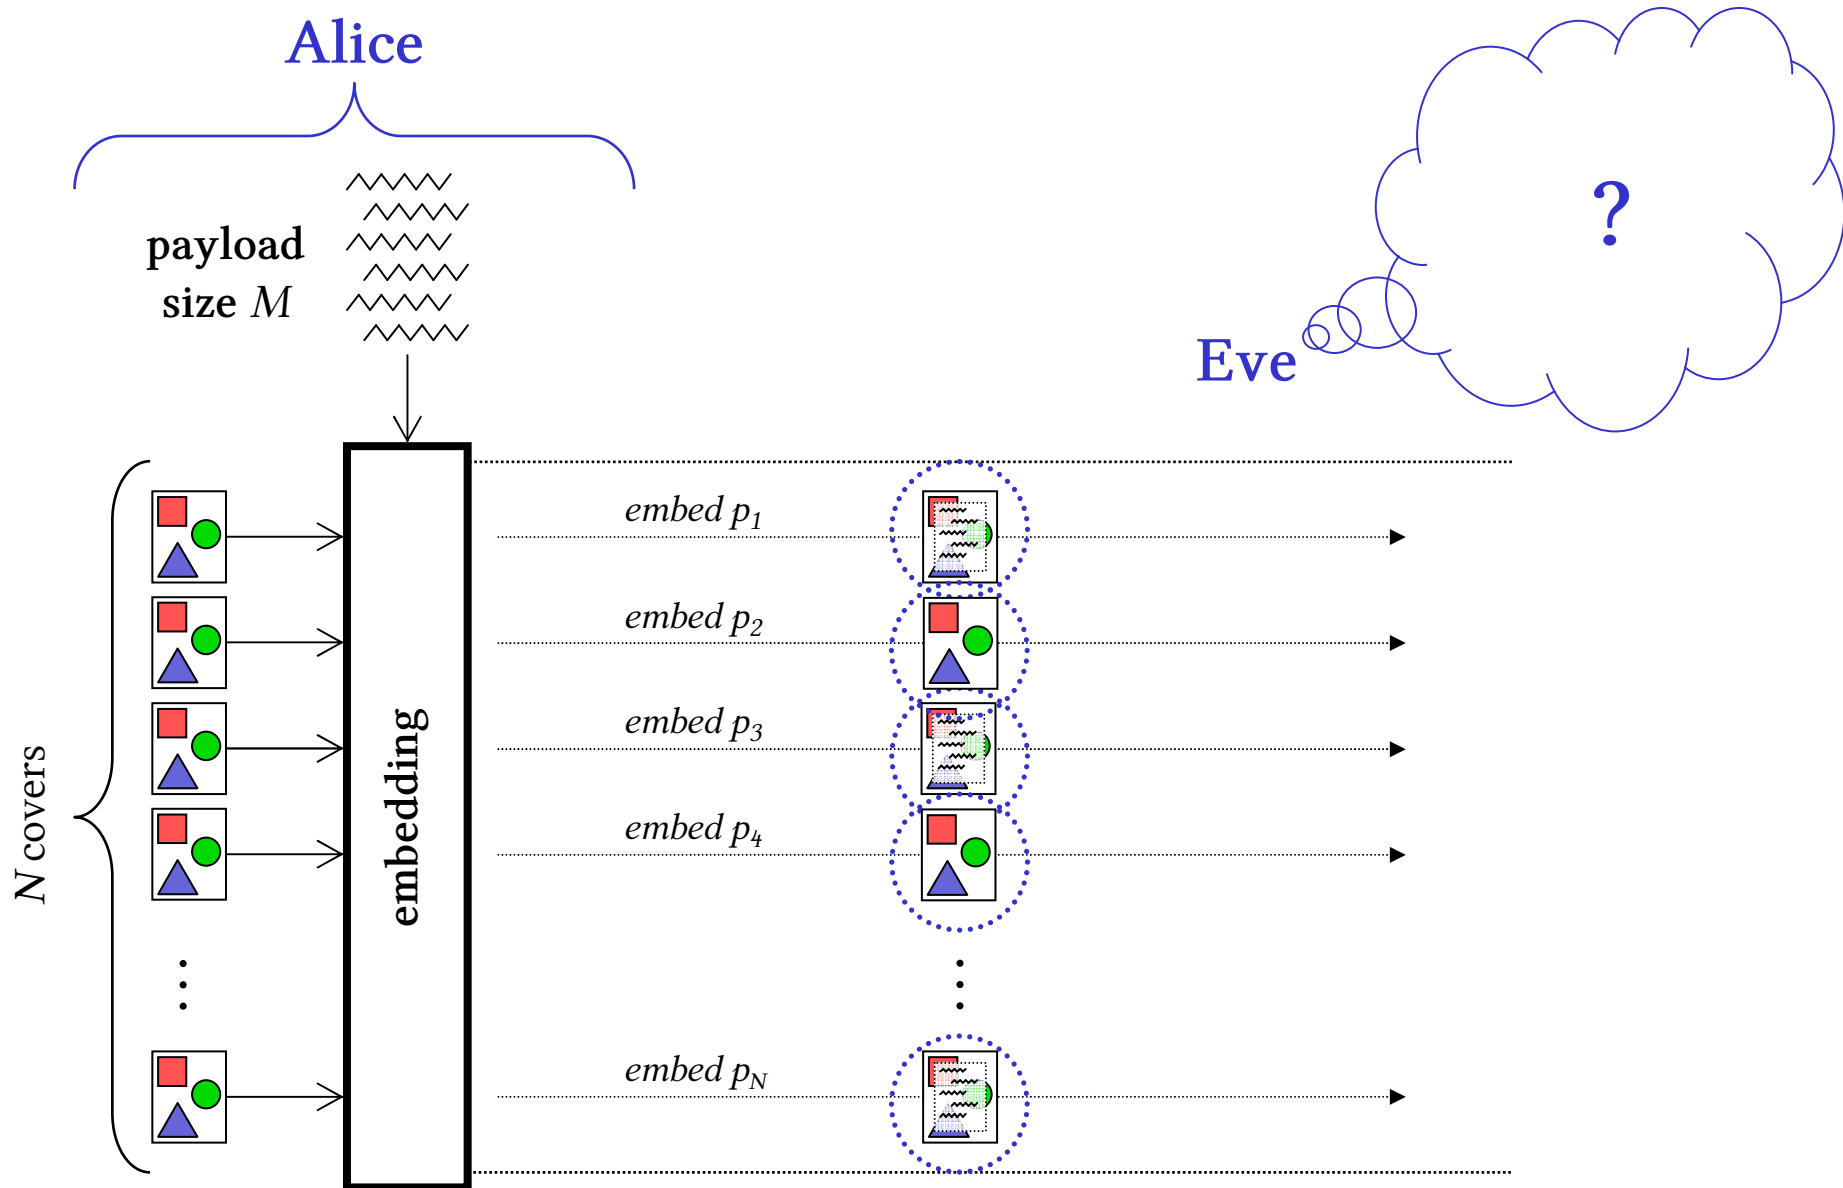
$embed\ p_4$

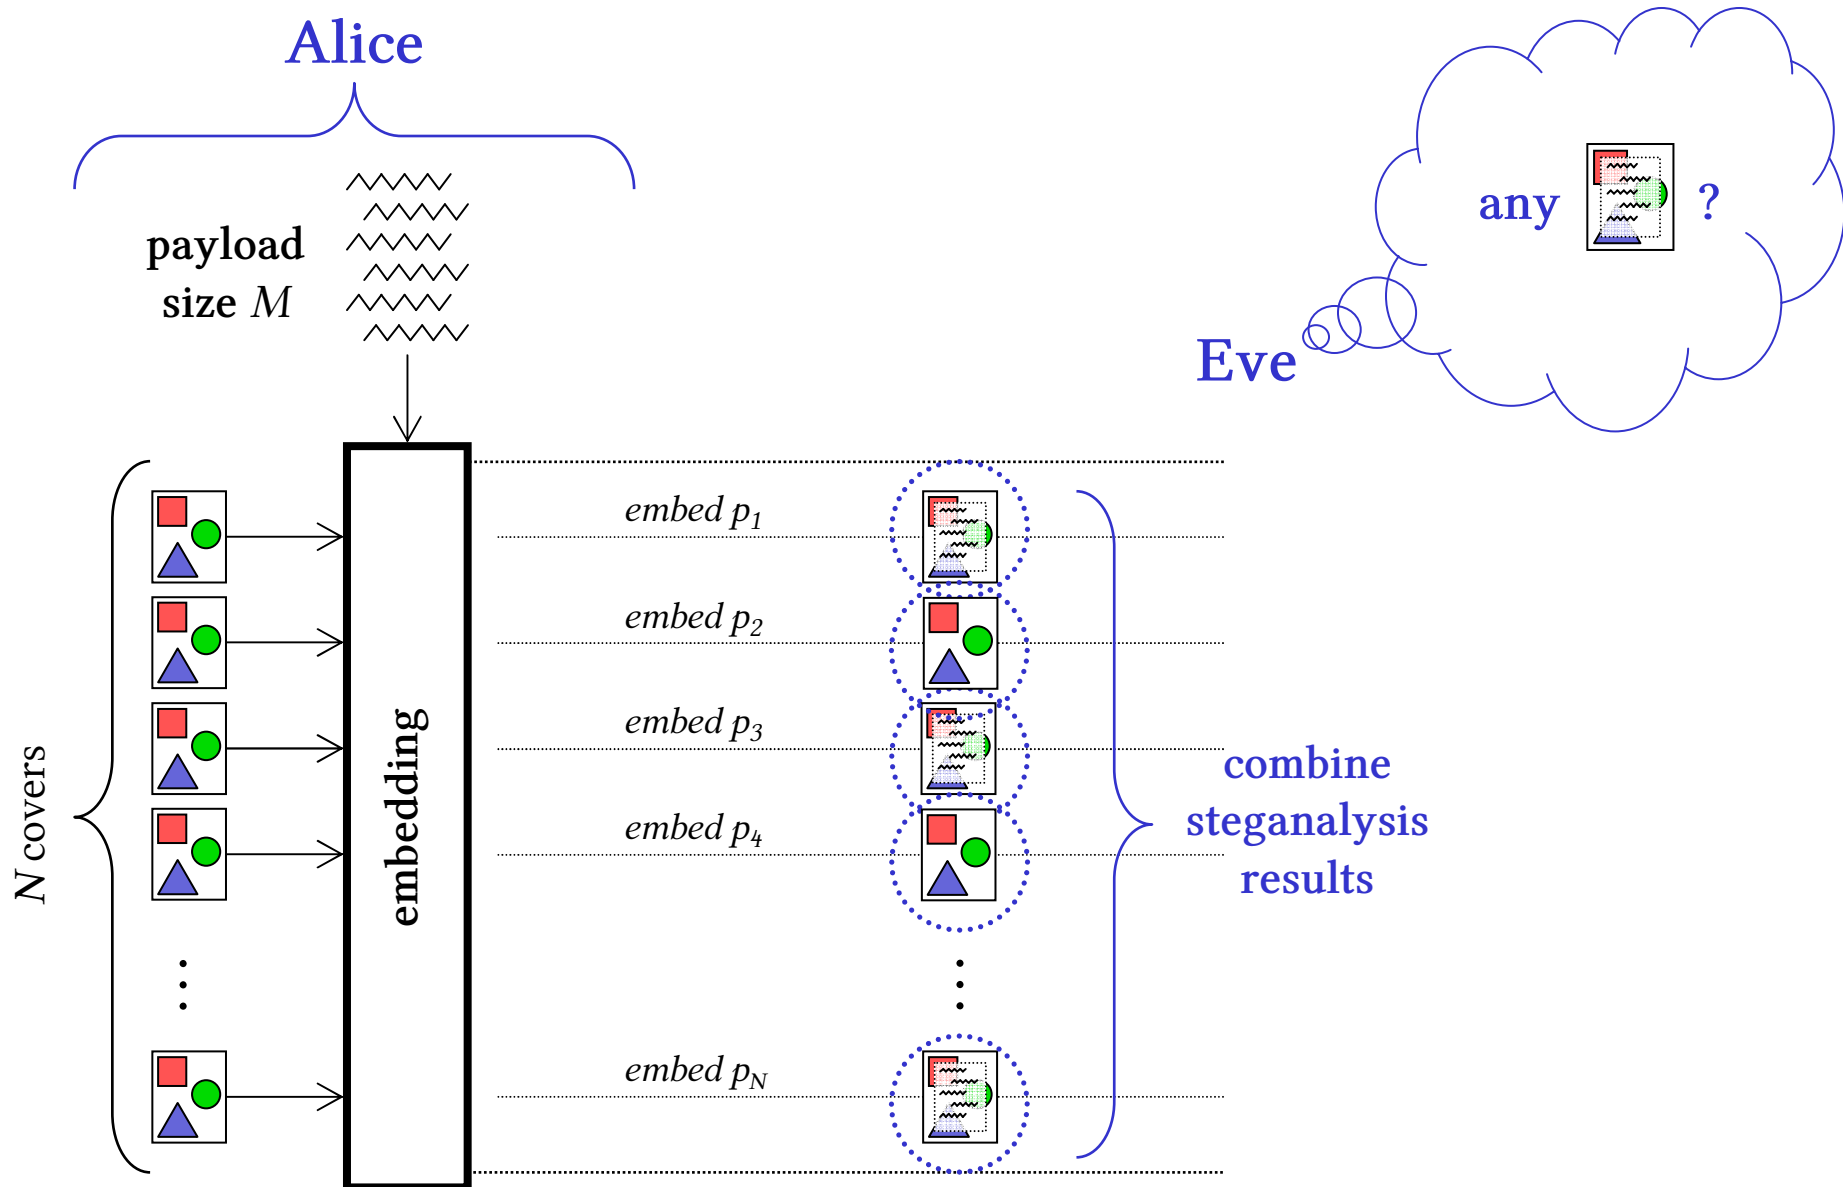


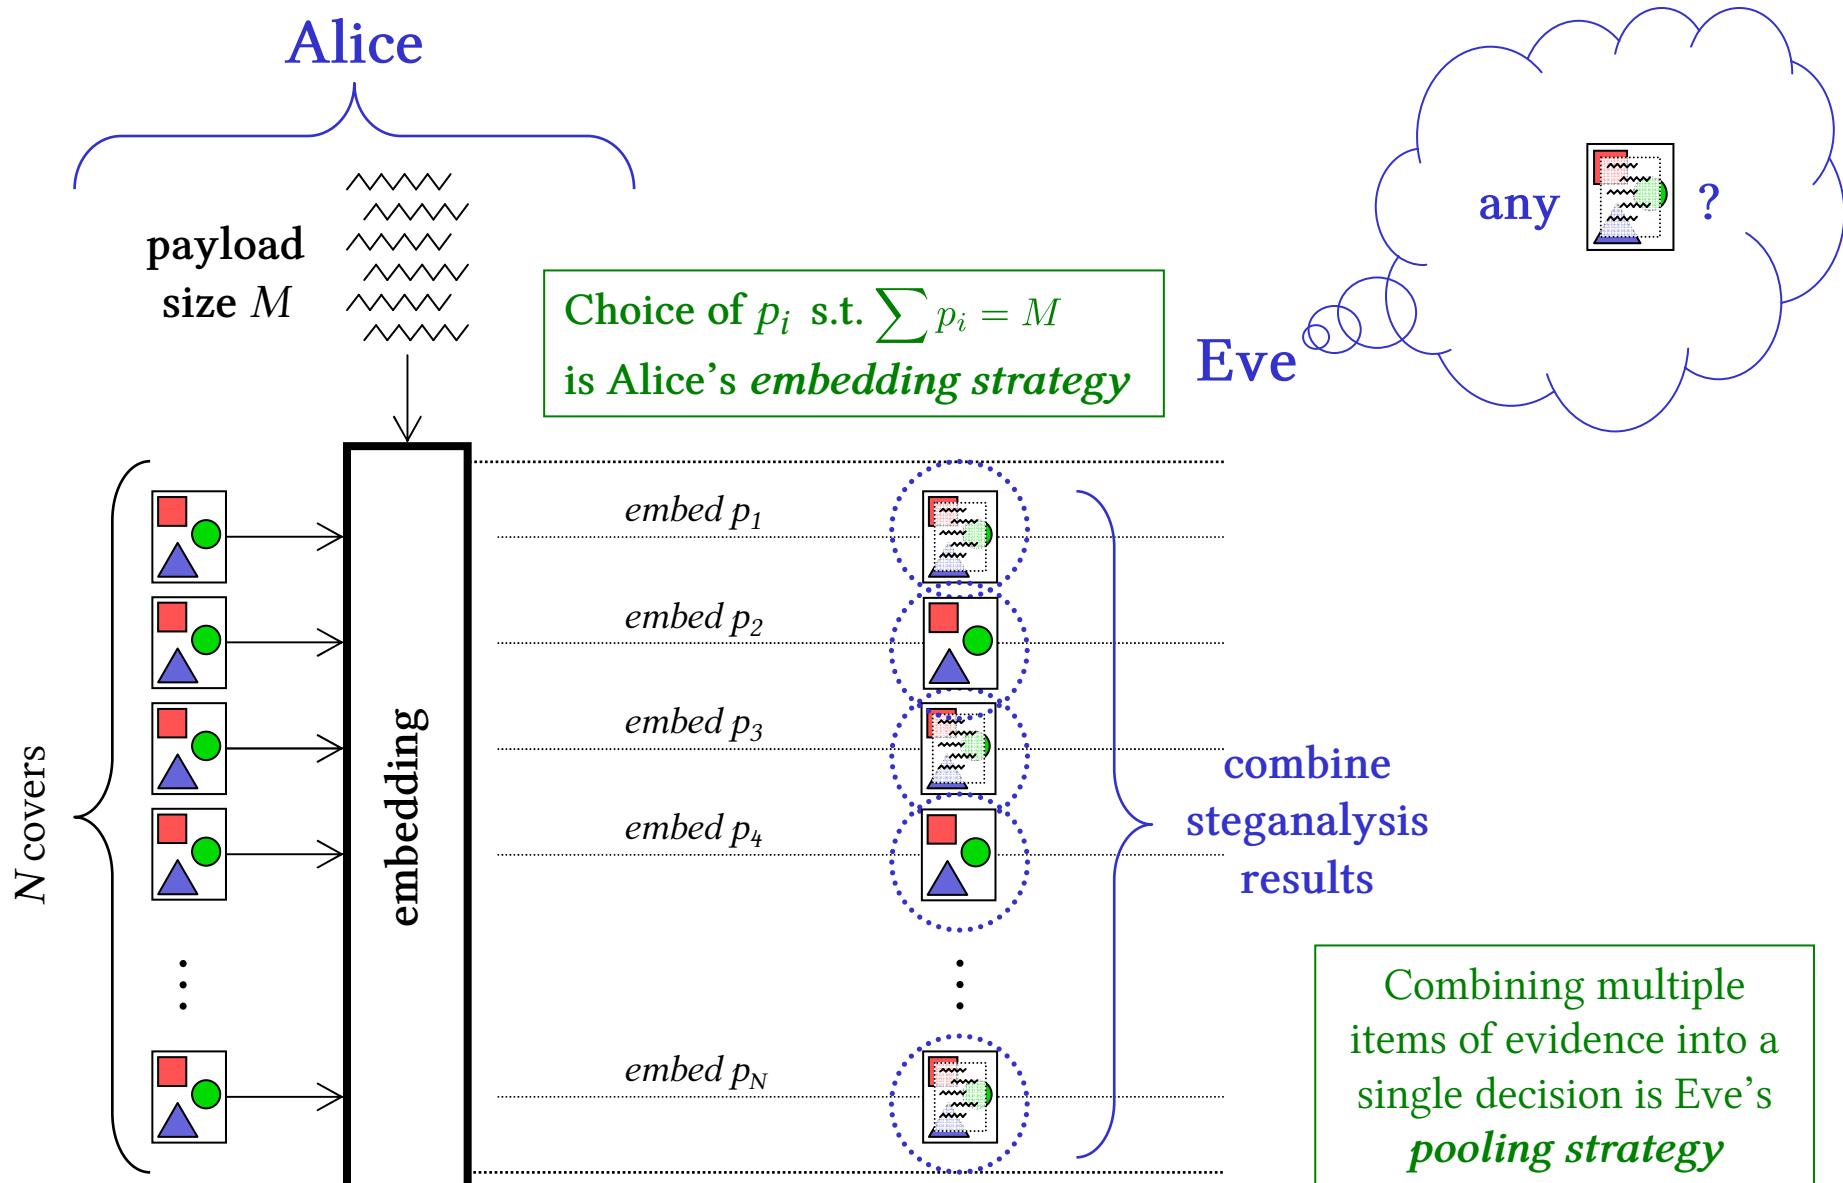
...

$embed\ p_N$









Batch steganography

What we know:

- *Under mild conditions, capacity follows a square root law:
 $M = O(\sqrt{N})$.*
- *Batch steganography & pooled steganalysis lead to a natural game-theoretic problem, with no “pure” solution.*
- *Only a few special cases are solved.*

What we don't yet know:

- *General results about the game.*
- *The best embedding strategy.*
- *Any good pooling strategies (a very practical problem).*

Conclusions

- Imperfect steganography is **not** like communication in noisy channels: capacity is not linear.
 - *We should be wary of embedding “rates”.*
 - *This may have practical implications.*
- There is interesting work yet to do, extending the square root law.
 - *The challenges are mainly in statistics and analysis.*
- Fisher Information should be a focus.
 - *Takes us towards a genuine capacity estimate.*
- Batch steganography & pooled steganalysis deserve more attention.

End