The Square Root Law in Stegosystems with Imperfect Information

Andrew Ker

adk@comlab.ox.ac.uk

Royal Society University Research Fellow
Oxford University Computing Laboratory

12th Information Hiding Conference
Calgary, Canada, 29 June 2010

Perfect and imperfect embedding

Perfect embedding preserves **all** statistics of the cover source.

- It is undetectable.
- It has a linear capacity law.

It can be accomplished in two ways:

1. The ‘rejection sampler’.
   - *Unrealistic to achieve nontrivial capacity.*
2. Match distribution of cover source.
   - *Böhme argues that perfect knowledge of a real ‘empirical’ cover source is impossible.*

We contend that **all practical steganography is imperfect**.

- Capacity follows a ‘Square Root Law’.

Classic square root law

Cover consists of 'pixels', which may be changed into 'stego pixels'.

- **Cover pixels: i.i.d. bits, 1 with probability p,**
- **Stego pixels: i.i.d. bits, 1 with probability q,**
- **Embedding: overwrite each pixel, independently, with probability γ,**
- **p known to the detector, $p \neq 0, 1$, $p \neq q$.**

As cover size $n \to \infty$,

1. If $\gamma^2 n \to \infty$ then an asymptotically perfect detector exists.

2. If $\gamma^2 n \to 0$ then we have asymptotically perfect security.

The critical rate is $\gamma = O(1/\sqrt{n})$

Usually, payload size $M \propto n\gamma$: $M = O(\sqrt{n})$

Classic square root law

Cover consists of 'pixels', which may be changed into 'stego pixels'.

- Cover pixels: i.i.d. **with pdf $p(x)$,**
- Stego pixels: i.i.d. **with pdf $q(x)$,**
- Embedding: overwrite each pixel, independently, with probability γ,
- **$p(x)$ known to the detector, $\forall x.p(x) \neq 0, 1$, $\exists y.p(y) \neq q(y)$.**

As cover size $n \to \infty$,

1. If $\gamma^2 n \to \infty$ then an asymptotically perfect detector exists.

2. If $\gamma^2 n \to 0$ then we have asymptotically perfect security.

The critical rate is $\gamma = O(1/\sqrt{n})$

Usually, payload size $M \propto n\gamma$: $M = O(\sqrt{n})$

Classic square root law

Cover consists of ‘pixels’, which may be changed into ‘stego pixels’.

- Cover pixels: realisations of a Markov chain,
- Stego pixels: random function of cover pixels,
- Embedding: change each pixel, independently, with probability γ,
- Cover source known to the detector, nontrivial, not preserved by stego.

As cover size $n \to \infty$,

1. If $\gamma^2 n \to \infty$ then an asymptotically perfect detector exists.

2. If $\gamma^2 n \to 0$ then we have asymptotically perfect security.

The critical rate is $\gamma = O(1/\sqrt{n})$

Usually, payload size $M \propto n\gamma$: $M = O(\sqrt{n})$

Classic square root law

Cover consists of 'pixels', which may be changed into 'stego pixels'.

- Cover pixels: i.i.d. bits, 1 with probability p,
- Stego pixels: i.i.d. bits, 1 with probability q,
- Embedding: use randomly selected fixed number γn,
- p known to the detector, $p \neq 0, 1$, $p \neq q$.

As cover size $n \to \infty$,

1. If $\gamma^2 n \to \infty$ then an asymptotically perfect detector exists.

2. If $\gamma^2 n \to 0$ then we have asymptotically perfect security.

The critical rate is $\gamma = O(1/\sqrt{n})$

Usually, payload size $M \propto n\gamma$: $M = O(\sqrt{n})$

Classic square root law

- Cover pixels: i.i.d. bits, 1 with probability p,
- Stego pixels: i.i.d. bits, 1 with probability q,
- Embedding: overwrite each pixel, independently, with probability γ,
- p known to the detector, $p \neq 0, 1$, $p \neq q$.

imperfect embedding

covers not deterministic

perfect information

The Square Root Law in Stegosystems with Imperfect Information

Outline

- Imperfect steganography
- Square root laws
- Imperfect information
  - Enforcing ignorance
  - Modified square root law
- Embedding with learning

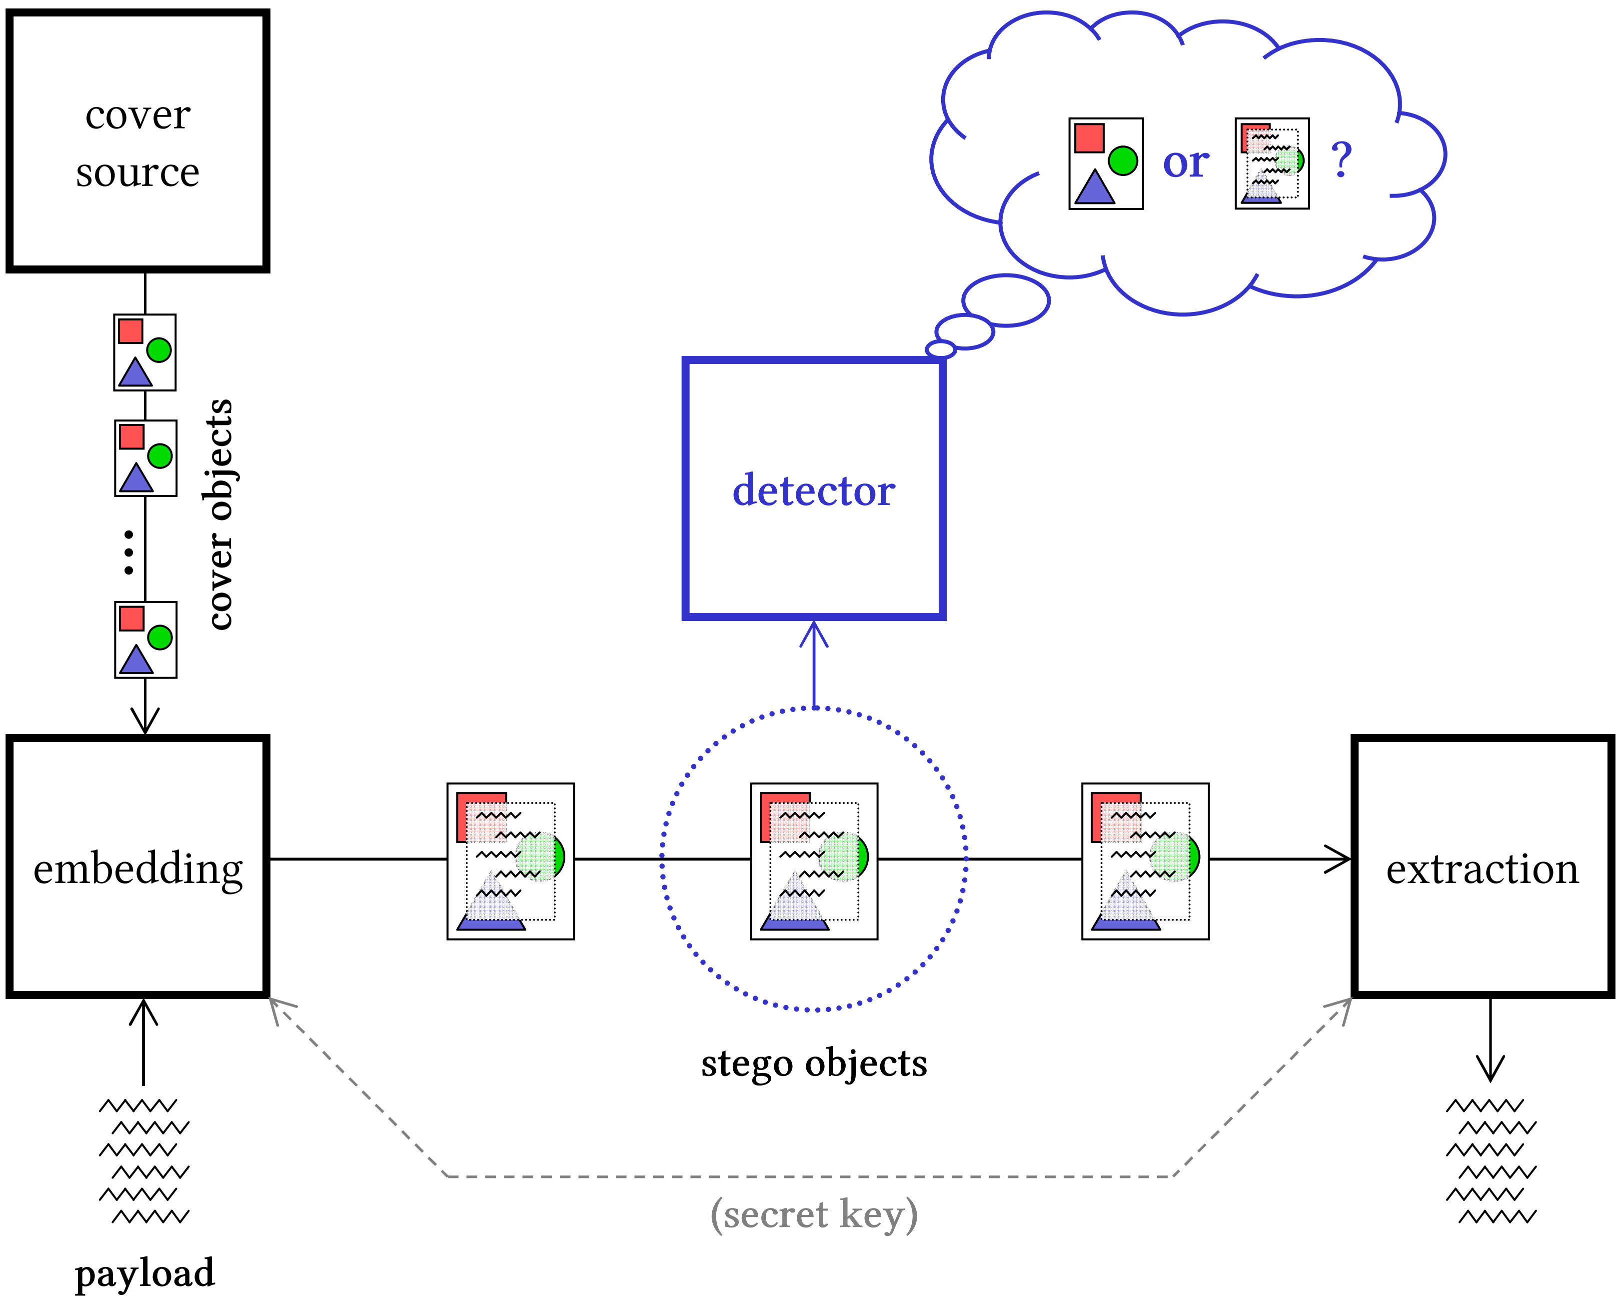

cover source
cover objects
or
?
detector
embedding
stego objects
extraction
(secret key)
payload

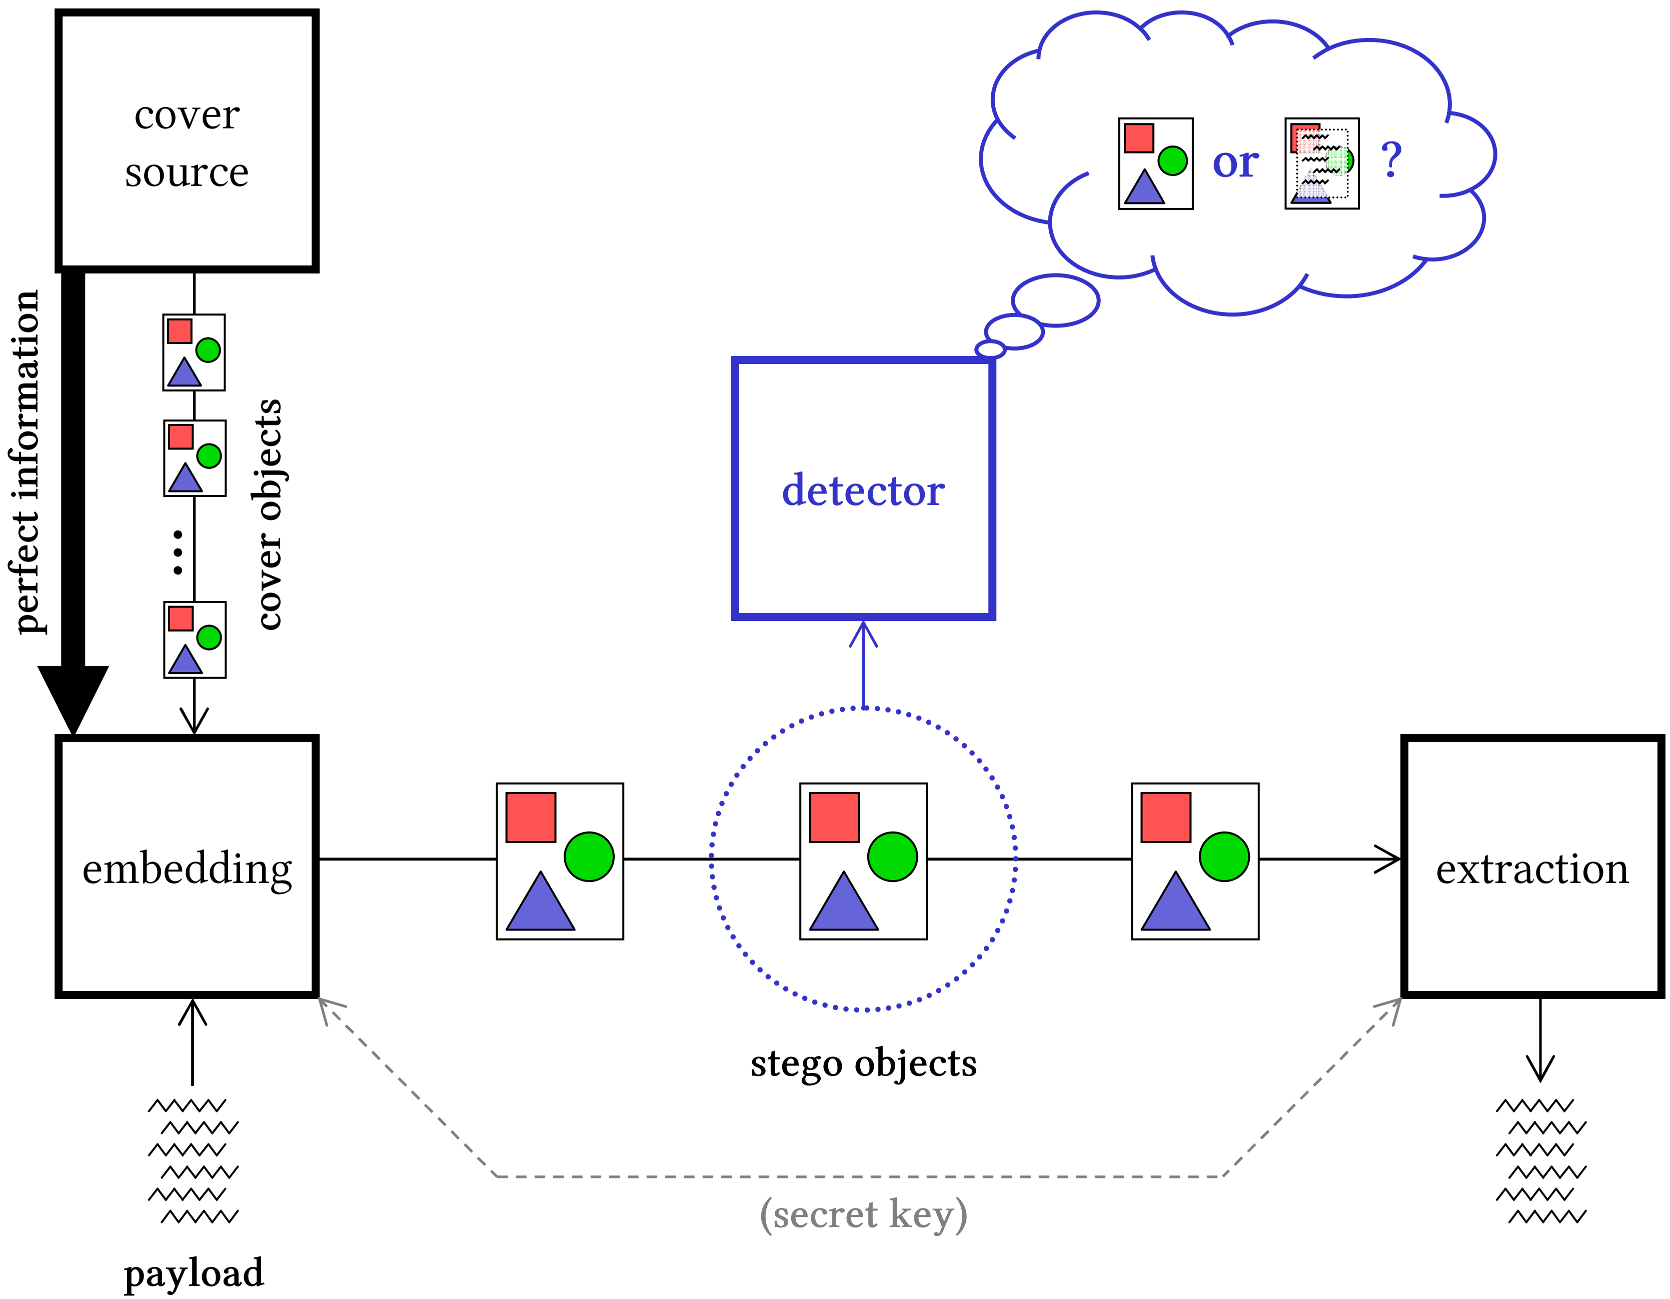
cover
source
perfect information
cover objects
embedding
detector
or
?
stego objects
extraction
(secret key)
payload

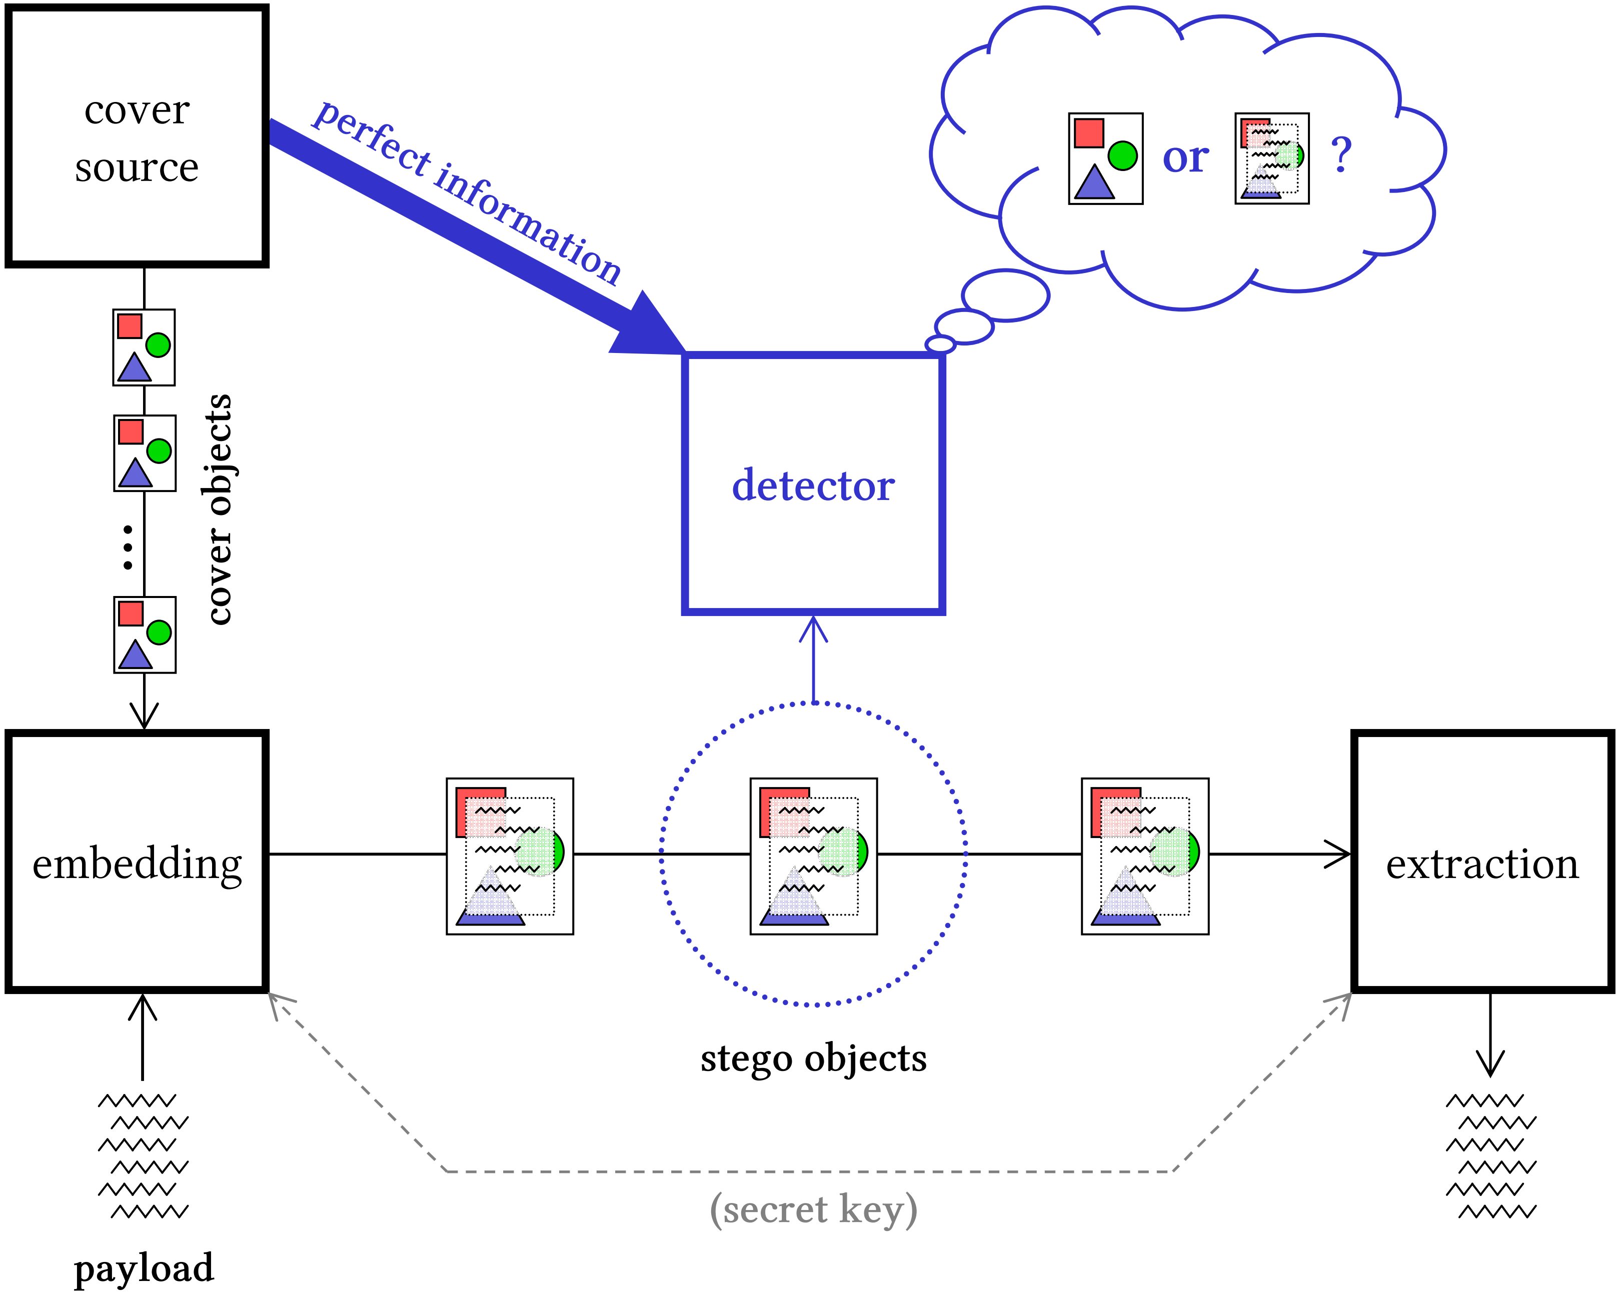
cover
source
perfect information
or
?
cover objects
detector
embedding
extraction
stego objects
payload
(secret key)

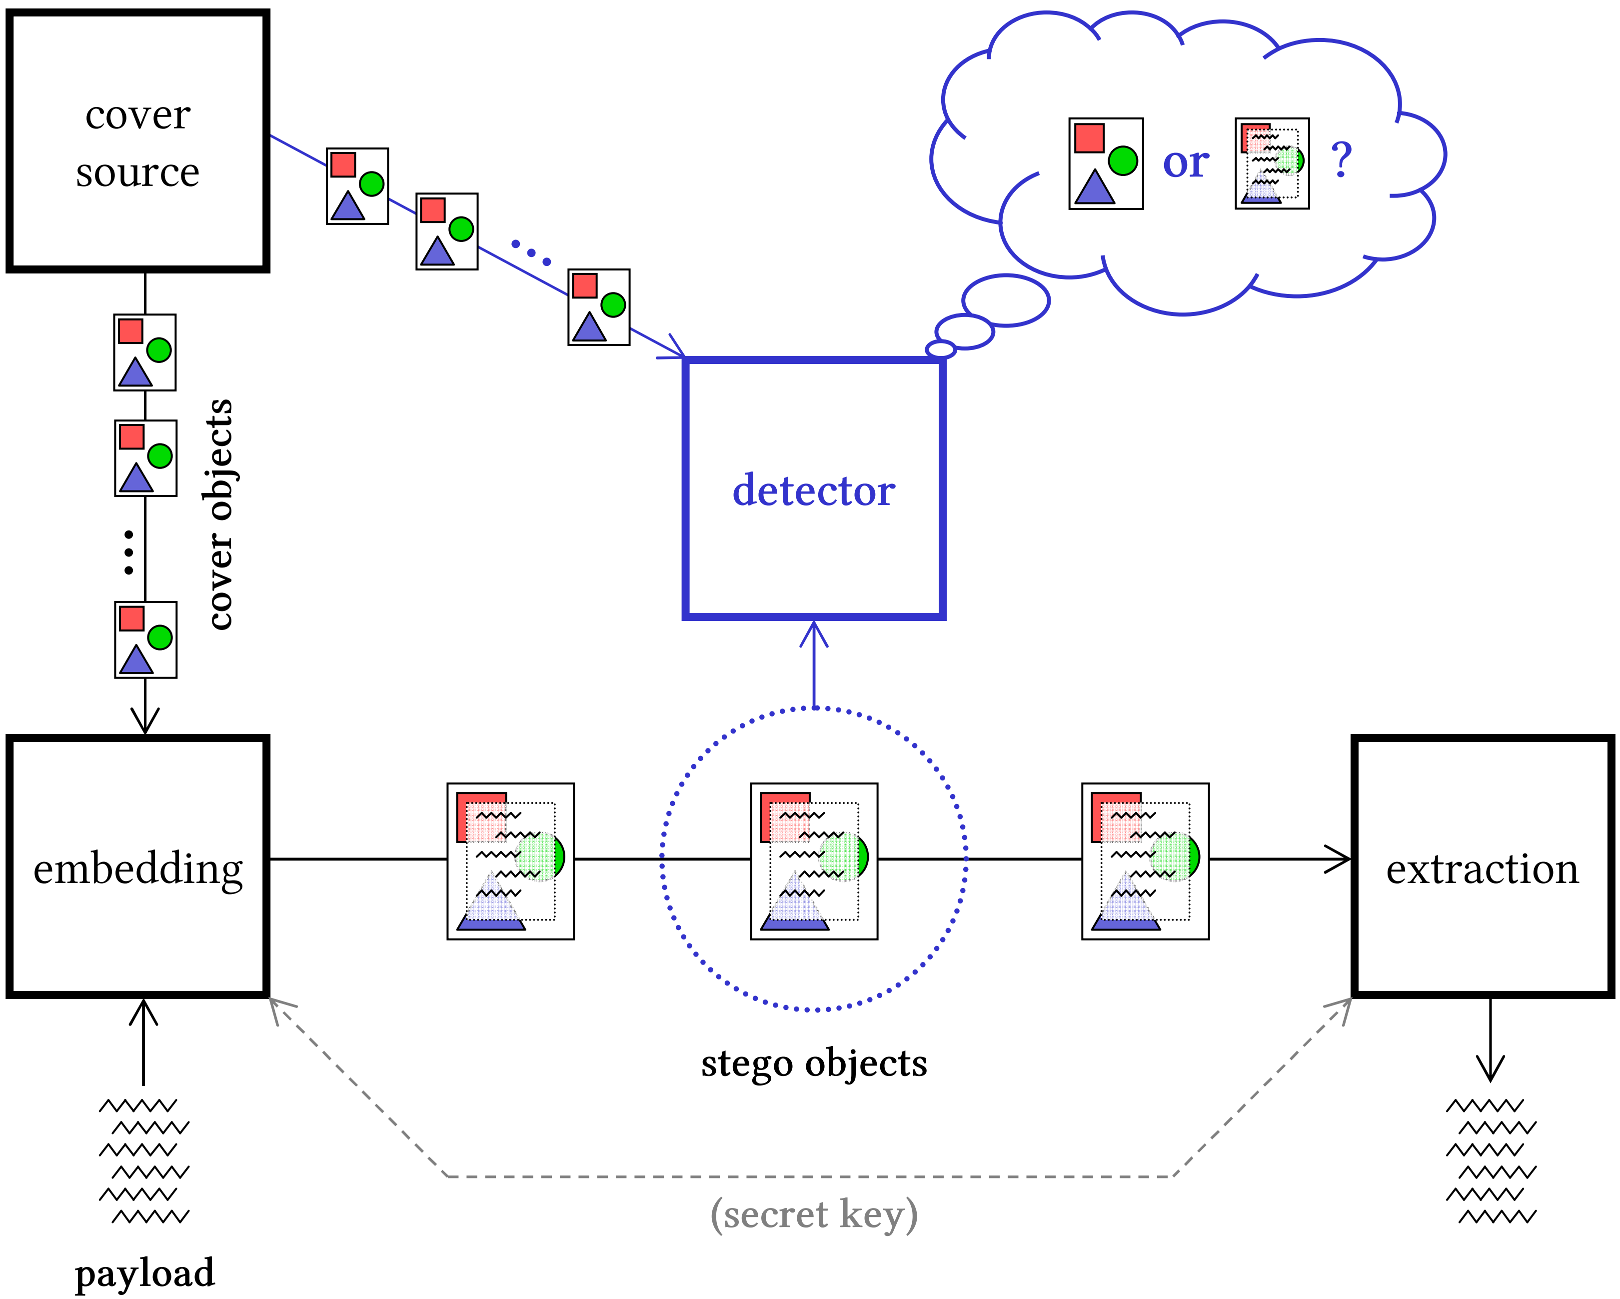

cover source
cover objects
detector
or
?
embedding
stego objects
extraction
payload
(secret key)

Imperfect information

Assume that the detector has access to a cover oracle, from which they can estimate characteristics of the cover source.

Questions:

- Are finitely many oracle accesses sufficient to restrict the embedder to a square root law? **(*No*)**
- Are exponentially many oracle accesses required? **(*No*)**

Imperfect information SRL

- Cover pixels: i.i.d. bits, 1 with probability p,
- Stego pixels: i.i.d. bits, 1 with probability q,
- Embedding: overwrite each pixel, independently, with probability γ,
- Detector has no prior knowledge of p, $p \neq 0, 1$, $p \neq q$.
- Detector has m bits from a cover oracle, also i.i.d., 1 with probability p.

As cover size $n \to \infty$,

1. If . . . then an asymptotically perfect detector exists.

2. If . . . then we have asymptotically perfect security.

Imperfect information SRL

- Cover pixels: i.i.d. bits, 1 with probability p,
- Stego pixels: i.i.d. bits, 1 with probability q,
- Embedding: overwrite each pixel, independently, with probability γ,
- Detector has no prior knowledge of p, $p \neq 0, 1$, $p \neq q$.
- Detector has m bits from a cover oracle, also i.i.d., 1 with probability p.

Detector sees:

m cover oracle bits	$(X_1, \dots, X_m)$	$X_i \sim \mathrm{Ber}(p)$
n suspect bits	$(Y_1, \dots, Y_n)$	$Y_i \sim \mathrm{Ber}(p + \gamma(q - p))$

and wants to know whether $\gamma > 0$.

Enforcing ignorance

Detector sees:

m cover oracle bits	$(X_1, \ldots, X_m)$	$X_i \sim \mathrm{Ber}(p)$
n suspect bits	$(Y_1, \ldots, Y_n)$	$Y_i \sim \mathrm{Ber}(p + \gamma(q - p))$

and wants to know whether $\gamma > 0$.

Asymptotic security is usually proved by showing that

$$\mathrm{D}_{\mathrm{KL}}(\text{cover objects} \,\|\, \text{stego objects}) \to 0$$

as $n \to \infty$.

Fails: cannot take account of a lack of knowledge by the detector.

Enforcing ignorance

Detector sees:

~~m cover oracle bits~~	~~$(X_1, \ldots, X_m)$~~	~~$X_i \sim \mathrm{Ber}(p)$~~
n suspect bits	$(Y_1, \ldots, Y_n)$	$Y_i \sim \mathrm{Ber}(p + \gamma(q - p))$

and wants to know whether $\gamma > 0$.

Asymptotic security is usually proved by showing that

$$\mathrm{D_{KL}}(\text{cover objects} \,\|\, \text{stego objects}) \to 0$$

as $n \to \infty$.

Fails: cannot take account of a lack of knowledge by the detector.

Even if $m = 0$, the KLD is positive.

Enforcing ignorance

Detector sees:

m cover oracle bits	$(X_1, \ldots, X_m)$	$X_i \sim \mathrm{Ber}(p)$
n suspect bits	$(Y_1, \ldots, Y_n)$	$Y_i \sim \mathrm{Ber}(p + \gamma(q - p))$

and wants to know whether $\gamma > 0$.

Try imposing a uniform prior on p?

Fails: If p were random we could repeat the experiment to test $p + \gamma(q - p)$ for uniformity.

Distribution of $p + \gamma(q - p)$ if $\gamma = 0$: 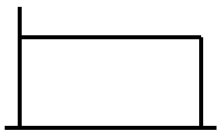 $\gamma > 0$: 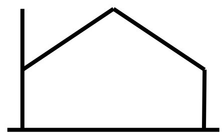

Enforcing ignorance

Detector sees:

~~m cover oracle bits~~ $\cancel{(X_1, \ldots, X_m)}$ $\cancel{X_i \sim \mathrm{Ber}(p)}$

n suspect bits $(Y_1, \ldots, Y_n)$ $Y_i \sim \mathrm{Ber}(p + \gamma(q - p))$

and wants to know whether $\gamma > 0$.

Try imposing a uniform prior on p?

Fails: If p were random we could repeat the experiment to test $p + \gamma(q - p)$ for uniformity.

Distribution of $p + \gamma(q - p)$ if $\gamma = 0$: 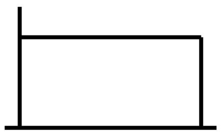 $\gamma > 0$: 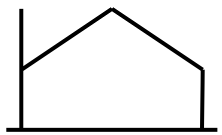

Even if $m = 0$, the KLD is positive.

Enforcing ignorance

Detector sees:

m cover oracle bits	$(X_1, \ldots, X_m)$	$X_i \sim \mathrm{Ber}(p)$
n suspect bits	$(Y_1, \ldots, Y_n)$	$Y_i \sim \mathrm{Ber}(p + \gamma(q - p))$

and wants to know whether $\gamma > 0$.

Impose unbiasedness:

A detector is unbiased if, no matter what p,

$$\mathbf{Pr(true\ {+ve}) \geq Pr(false\ {+ve}).}$$

The statistics literature tells us that the most powerful (optimal) unbiased test for Bernoulli probabilities depends only on $\sum Y_i \mid \left(\sum X_i + \sum Y_i\right)$.

Imperfect information SRL

- Cover pixels: i.i.d. bits, 1 with probability p,
- Stego pixels: i.i.d. bits, 1 with probability q,
- Embedding: overwrite each pixel, independently, with probability γ,
- Detector unbiased for p, $p \neq 0, 1$, $p \neq q$,
- Detector has m bits from a cover oracle, also i.i.d., 1 with probability p.

As cover size $n \to \infty$,

1. If $\gamma^2 \dfrac{n\,m}{n+m} \to \infty$ then an asymptotically perfect detector exists.
2. If $\gamma^2 \dfrac{n\,m}{n+m} \to 0$ then we have asymptotically perfect security.

The critical rate is $\gamma = O\big(1/\sqrt{1/m + 1/n}\big)$

Interpretation

The critical rate is $\gamma = O\big(1/\sqrt{1/m + 1/n}\big)$

If m is finite (does not grow with n) then the critical rate is $\gamma = O\big(1\big)$:

- *finite information at the detector leads to linear capacity.*

If m is at least linear in n, then the critical rate is $\gamma = O\big(1/\sqrt{n}\big)$:

- *linearly many oracle accesses suffice to restrict the embedder to a square root law.*

If m is sublinear in n, then the critical rate is intermediate.

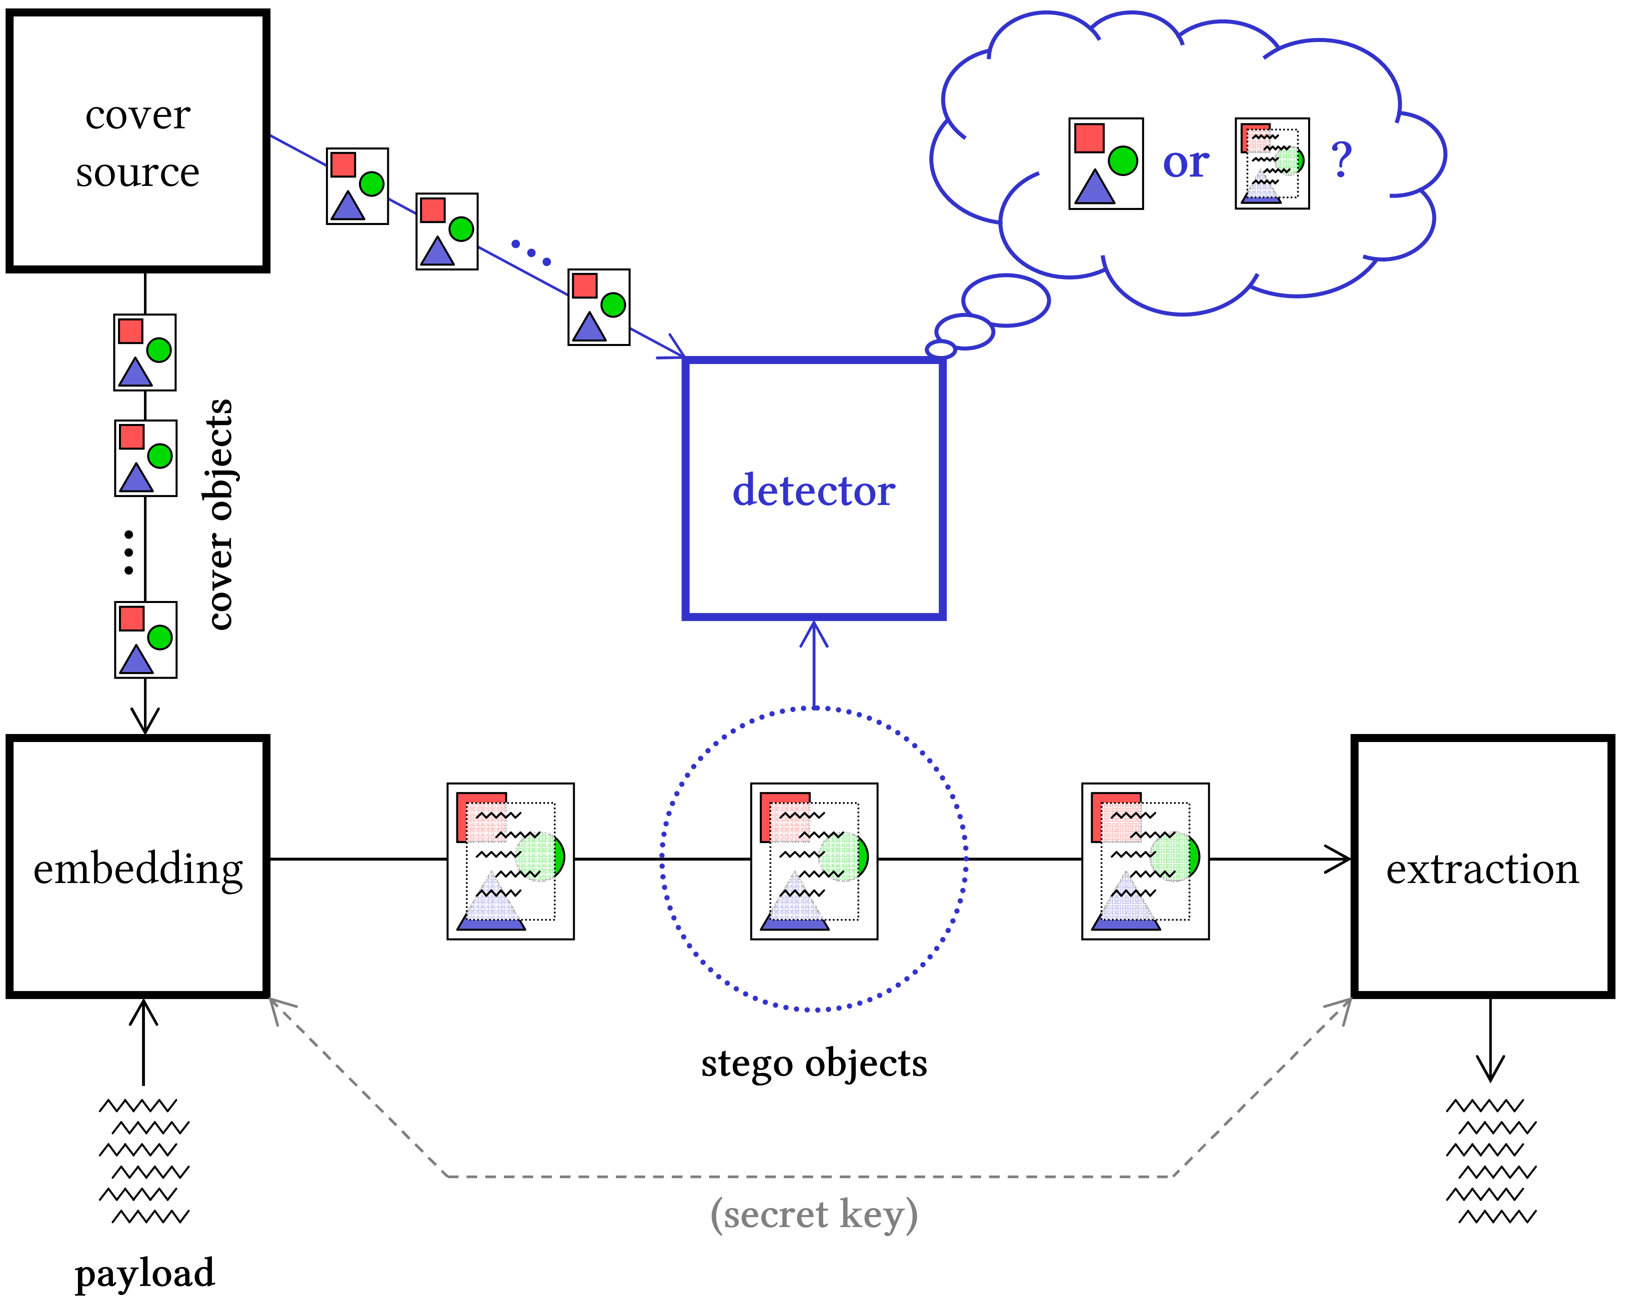

cover source
cover objects
detector
or
?
embedding
stego objects
extraction
payload
(secret key)

Conclusions

- Reasoning about imperfect information is difficult.
  - *KL divergence alone is not sufficient.*
  - *Statistical concepts of unbiasedness and invariance may be useful.*

- The square root law still holds in the imperfect information case...

  ... as long as the detector has linearly many cover oracle accesses.

- 'Embedding with learning' needs more theoretical scrutiny.
  - *We may be heading back towards a linear capacity law.*

- Consider the epistemology of steganography.
  - *Assuming perfect knowledge of the cover source is unrealistic.*
  - *Kerckhoffs' Principle should not be used blindly.*
  - *There may be many variants of the 'steganography problem'.*