# Fourth-Order Structural Steganalysis and Analysis of Cover Assumptions

Andrew D. Ker

Oxford University Computing Laboratory, Parks Road, Oxford OX1 3QD, England

## ABSTRACT

We extend our previous work on *structural steganalysis* of LSB replacement in digital images, building detectors which analyse the effect of LSB operations on pixel groups as large as four. Some of the method previously applied to triplets of pixels carries over straightforwardly. However we discover new complexities in the specification of a cover image model, a key component of the detector. There are many reasonable *symmetry* assumptions which we can make about parity and structure in natural images, only some of which provide detection of steganography, and the challenge is to identify the symmetries a) completely, and b) concisely.

We give a list of possible symmetries and then reduce them to a complete, non-redundant, and approximately independent set. Some experimental results suggest that all useful symmetries are thus described. A weighting is proposed and its approximate variance stabilisation verified empirically. Finally, we apply symmetries to create a novel *quadruples* detector for LSB replacement steganography. Experimental results show some improvement, in most cases, over other detectors. However the gain in performance is moderate compared with the increased complexity in the detection algorithm, and we suggest that, without new insight, further extension of structural steganalysis may provide diminishing returns.

**Keywords:** Steganography, structural steganalysis, digital image forensics

## 1. INTRODUCTION: LSB STEGANOGRAPHY

It is a popular fallacy that replacement of Least Significant Bits (LSBs) of pixel values provides an undetectable way to hide a secret message in a bitmap image. Whilst it provides a high-capacity and visually imperceptible method of embedding, there is now substantial literature[1–5] to show that replacement of LSBs introduces statistically-detectable changes in the structure of an image. Nonetheless, perhaps due to the extreme simplicity of the embedding algorithm* it remains a popular method for the less-discerning steganographer. In this paper we extend previous work to further refine our capability for detection of LSB replacement: the focus is on reliable detection of ever-smaller hidden messages.

We take "LSB replacement" to mean the commonly-accepted embedding model whereby the cover image is traversed in a pseudorandom order (generated by a secret key shared between sender and recipient) and the cover pixel values altered by replacing least significant bits by the secret message bit stream. Colour images have their colour components treated separately. The pseudorandom order is important both to maintain the secrecy of the message and to spread the stego noise around the cover image, in the case when the hidden message is less than the maximum possible of one secret bit per cover byte.

The most sensitive detectors for LSB replacement steganography make use of the *structural* property of the embedding algorithm – namely that even cover pixels can never be decremented and odd pixels never incremented when least significant bits are replaced – although sometimes the reliance on this property is implicit. Early structural detectors included the methods of *Pairs*[2] and *Sample Pairs*.[3] With one exception (the slightly less sensitive detector *RS*,[1] which has little theoretical explanation) such detectors use the structural property as it relates to pairs of pixels. In recent work[5] we gave a novel framework in which the structural property can be expressed precisely for any size group of pixels, and included as a case study a detector known as *Triples Analysis* which exploits the structural property in triplets of pixels. Performance of the Triples detector is somewhat better than that of the other detectors.

Further author information: E-mail: adk@comlab.ox.ac.uk, Telephone: +44 1865 283530

*In Ref. 4 is a two-line programme for LSB replacement embedding.

In this paper we extend the study to pixel groups of size four and thus create a new "fourth-order" detector which we naturally call *Quadruples Analysis*. Part of the Triples method extends immediately, but there are many more complications in specifying the cover model for quadruplets. The main contribution of this paper is a close analysis of the cover model, which points the way to the extension to arbitrary groups of pixels (the cover model for arbitrary groups was the missing ingredient, in Ref. 5, for a totally general structural detector). Some experimental results here show that Quadruples Analysis outperforms previously known detectors when applied to images with short hidden messages (it complements the Triples detector by performing best when the Triples detector performs worse), but the gain is moderate and the performance is worse for larger messages. This suggests that further extensions may offer limited reward.

To conclude this introduction, we must note that these structural detectors only work because of the structural property of LSB replacement, and that different embedding methods will not be detected. In particular, the small modification known as LSB matching (in which the hidden message is still the LSBs of the stego stream, but the cover is modified nondeterministically to achieve this) is immune to all of the above detectors, and to that presented here.

## 2. FRAMEWORK FOR STRUCTURAL STEGANALYSIS

The detection framework of Ref. 5 is as follows. We define a macroscopic property of images which depends on the length of hidden data $p$, a vector $\mathbf{S}(p)$. We determine how $\mathbf{S}(p)$ depends on $p$ and $\mathbf{S}(0)$ then invert the process: given an image we hypothesise a value for $p$ and compute what this would imply for $\mathbf{S}(0)$. Given a model for (macroscopic properties of) cover images we find the value of $p$ which leads to a value of $\mathbf{S}(0)$ closest to the model: this is the estimator for the true value of $p$.

The first part of this framework we address in this section and the case of quadruplets of pixels is a straightforward application of Ref. 5. The cover image model is considered in depth in Sect. 3 and the question of what we mean by "closest to the model" is addressed in Sect. 4, where we also give experimental evaluation of performance.

According to the framework of Ref. 5, the feature vectors $\mathbf{S}$ are made up of the cardinalities of *trace subsets*:

$$\begin{aligned} \mathcal{E}_{x_1,\ldots,x_{n-1}} &= \{(s_1,\ldots,s_n) \in \mathcal{T} \mid x_1 \text{ even, and } s_{i+1} - s_i = x_i \text{ for each } i\} \\ \mathcal{O}_{x_1,\ldots,x_{n-1}} &= \{(s_1,\ldots,s_n) \in \mathcal{T} \mid x_1 \text{ odd, and } s_{i+1} - s_i = x_i \text{ for each } i\} \end{aligned}$$

where $n$ is the "order" of the detector (2 for Couples, 3 for Triples, 4 for Quadruples) and $\mathcal{T}$ is some parent set of values taken from an image, typically tuples of pixel values from all adjacent groups of $n$ pixels. Assuming byte values for pixel intensities, these definitions make sense for indices in the range $-255$ to 255.

We suppose that a fixed cover image has $N$ pixels (counting colour components separately) and that a hidden message of length $2pN$ is embedded by LSB replacement ($0 \le p \le 0.5$). The sizes of the trace subsets $\mathcal{E}_{\vec{x}}$ (respectively $\mathcal{O}_{\vec{x}}$) are counted prior to embedding as $e_{\vec{x}}$ (respectively $o_{\vec{x}}$) and after embedding as $E'_{\vec{x}}$ (respectively $O'_{\vec{x}}$). As explained in Ref. 5, the trace subsets are predictably affected by LSB replacement (which flips each pixel's LSB independently at random with probability $p$, so long as one assumes that the locations for the hidden message are selected randomly, for example by pseudorandom permutation of the cover). The structural framework connects the random variables $E'_{\vec{x}}$ and $O'_{\vec{x}}$ to the cover properties $e_{\vec{x}}$ and $o_{\vec{x}}$ via $p$. In a detector, the quantities $E'_{\vec{x}}$ and $O'_{\vec{x}}$ are observable simply by counting trace subsets in an image; then we can use assumptions about cover images to provide information about $p$.

We write $\mathbf{S}(p)$ for a vector of some trace subsets (precisely which depends on $n$) after embedding. The connection between $\mathbf{S}(p)$ and $\mathbf{S}(0)$ is computed in terms of the matrices

$$T_1(p) = \begin{pmatrix} 1-p & p \\ p & 1-p \end{pmatrix}$$

and $T_n(p) = T_{n-1}(p) \otimes T_1(p)$. It is easy to compute $T_n^{-1}(p)$ so as to invert the process ($T_n^{-1}$ is again an $n$-fold Kronecker product – the reader is directed to Ref. 5 for the formula).

Considering just pairs of pixels, which leads to a detector we have known as *Couples Analysis*[†], Ref. 5 gives the equation

$$\begin{pmatrix} E[E'_{2m}] \\ E[O'_{2m-1}] \\ E[E'_{2m+1}] \\ E[O'_{2m}] \end{pmatrix} = T_2(p) \begin{pmatrix} e_{2m} \\ o_{2m-1} \\ e_{2m+1} \\ o_{2m} \end{pmatrix} \text{ and therefore } \begin{pmatrix} e_{2m} \\ o_{2m-1} \\ e_{2m+1} \\ o_{2m} \end{pmatrix} \approx T_2^{-1}(p) \begin{pmatrix} E'_{2m} \\ O'_{2m-1} \\ E'_{2m+1} \\ O'_{2m} \end{pmatrix}$$

(using the Law of Large Numbers to replace the expectation by the observed random variable). This expresses properties of the cover image in terms of the stego image and $p$. Note that such an equation holds for any $m$. A similar equation is given in Ref. 5 for triplets of pixels, where an 8-by-8 matrix is used to connect vectors of 8 trace subsets. In the case of quadruplets of pixels, it is a straightforward application of Ref. 5 to deduce that, for each $l$, $m$, and $n$,

$$\begin{pmatrix} e_{2l,2m,2n} \\ o_{2l-1,2m,2n} \\ e_{2l+1,2m-1,2n} \\ o_{2l,2m-1,2n} \\ e_{2l,2m+1,2n-1} \\ o_{2l-1,2m+1,2n-1} \\ e_{2l+1,2m,2n-1} \\ o_{2l,2m,2n-1} \\ e_{2l,2m,2n+1} \\ o_{2l-1,2m,2n+1} \\ e_{2l+1,2m-1,2n+1} \\ o_{2l,2m-1,2n+1} \\ e_{2l,2m+1,2n} \\ o_{2l-1,2m+1,2n} \\ e_{2l+1,2m,2n} \\ o_{2l,2m,2n} \end{pmatrix} \approx T_4^{-1}(p) \begin{pmatrix} E'_{2l,2m,2n} \\ O'_{2l-1,2m,2n} \\ E'_{2l+1,2m-1,2n} \\ O'_{2l,2m-1,2n} \\ E'_{2l,2m+1,2n-1} \\ O'_{2l-1,2m+1,2n-1} \\ E'_{2l+1,2m,2n-1} \\ O'_{2l,2m,2n-1} \\ E'_{2l,2m,2n+1} \\ O'_{2l-1,2m,2n+1} \\ E'_{2l+1,2m-1,2n+1} \\ O'_{2l,2m-1,2n+1} \\ E'_{2l,2m+1,2n} \\ O'_{2l-1,2m+1,2n} \\ E'_{2l+1,2m,2n} \\ O'_{2l,2m,2n} \end{pmatrix} \tag{1}$$

In effect, and up to approximation, we have *undone* the effect of a known amount of steganography as far as measuring the size of trace subsets is concerned. Now if we can find an equation involving the trace subsets in the cover image, for example if we believe that $e_{1,0,0} \approx o_{1,0,0}$ (and we will see that this is approximately true for cover images), then we can select the correct values of $l, m, n$ and the correct component of Eq. (1) to give an equation for $p$. The application, so far, of the general framework and calculations of Ref. 5 has been straightforward. But finding *all* the properties of cover images helpful for steganalysis is a little more difficult, as we see in the next section.

## 3. COVER IMAGE SYMMETRIES

The connection between trace subsets in cover and stego images, as described in the previous section, is without value unless we have some properties of the cover image with which to form an equation for $p$. We call an approximate equation of the form

$$e_{a,b,c} \approx e_{p,q,r} \text{ or } e_{a,b,c} \approx o_{p,q,r}$$

a *symmetry* if we believe that the approximate equality holds in all natural cover images. (We do not state explicitly how closely the equality must hold – this is a qualitative judgment, and the ultimate decision is made based on whether the steganography detectors it produces are at all accurate.)

In the case of Couples, there is a symmetry $e_m \approx o_m$ for each $m$. Approximate equality is verified empirically by examining some cover images, and it can also be explained because we do not expect any correlation between the parity of an individual pixel and the difference between that pixel and its neighbour. Matters are slightly more complicated for Triples (see Ref. 5 for brief details) where we see not only $e_{m,n} \approx o_{m,n}$ but also additional symmetries such as $e_{m,n} \approx e_{n,m}$.

[†]In Ref. 5 it is demonstrated that Couples Analysis is practically equivalent to the standard method Sample Pairs Analysis[3] (SPA) – the only distinction being that SPA conflates the trace subset $e_m$ with either $e_{-m}$ or $o_{-m}$ (depending on the parity of $m$) and has a less simple justification.

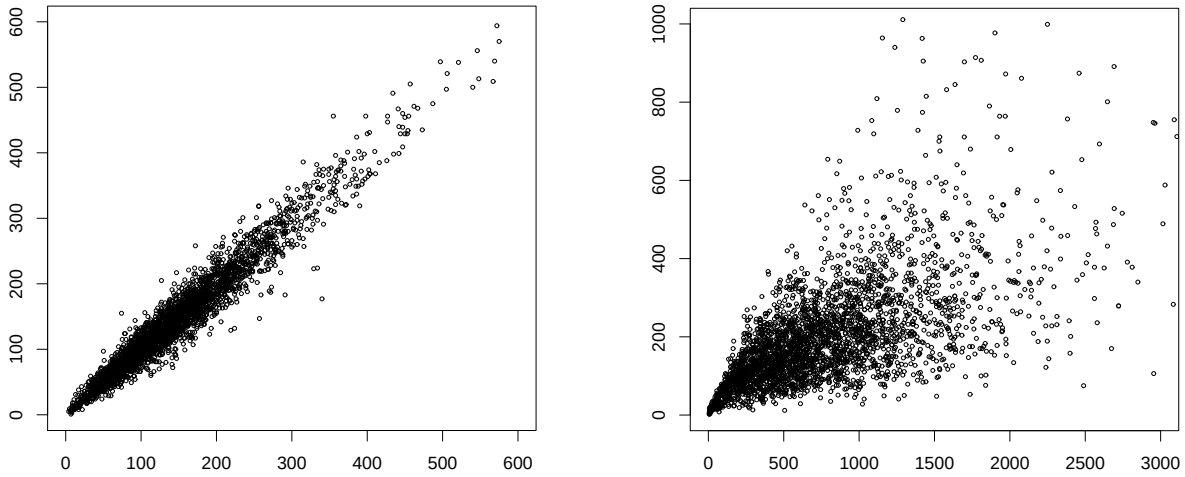


**Figure 1.** Left, scatterplot of $e_{0,1,2}$ (*x-axis*) against $o_{0,1,2}$ (*y-axis*) for 3000 cover images, showing strong correlation. Right, scatterplot of $e_{1,0,0}$ against $o_{1,1,1}$, showing much weaker correlation.

### 3.1. A search for all symmetries

There certainly are similar symmetries in the quadruplet trace subsets. For example, Fig. 1 shows that the trace subsets $\mathcal{E}_{0,1,2}$ and $\mathcal{O}_{0,1,2}$ are usually very close in size – this particular scatterplot came from computing the sizes of those two trace subsets in one particular set of 3000 natural images (described below). On the other hand, the figure also displays two trace subsets which are not closely correlated.

Rather than try to imagine all possible symmetries for the quadruplet trace subsets – which would risk missing some – we performed a systematic search using a set of 3000 never-compressed grayscale bitmap images. These images were downloaded from `http://photogallery.nrcs.usda.gov`; originally very high resolution colour images apparently scanned from film, for all our testing we reduced them in size to approximately $640 \times 450$ pixels. Except where otherwise indicated, we converted the images to grayscale prior to use. Hopefully, this set of images is representative of grayscale natural images as a whole and exposes all pairs of trace subsets which are generally close, in covers.

The size of each trace subset, $e_{a,b,c}$ and $o_{a,b,c}$ for $a, b, c \in \{-4, \ldots, 4\}$, was computed for each cover image. We then measured how "similar" each trace subset was to each other. The measure of similarity needs some care, because straightforward correlation is not suitable. For example, the trace subsets $e_{1,0,0}$ and $e_{1,1,0}$ are fairly close in magnitude because they both count tuples of pixels with mostly equal values, so both tend to be large in images with areas of almost flat colour, and both small in noisy images. This "symmetry" is not due to the fine structure of trace subsets, but rather it is due to a confounding external parameter and it vanishes when one considers only images with similar local variance. Therefore we must find a way to normalise for the sizes of the trace subsets involved. In view of the discussion in Subsection 3.3, below, we measured the average *stablised* deviation, defined to be the average of $(p - q)/\sqrt{p + q}$ (as $p$ and $q$ range over the paired elements of the two trace subsets). The choice of exactly which measure of deviation to use is not so important, as long as some account is taken of magnitude; we noted that other normalised similarity measures (e.g. average relative deviation $\sum |p - q|/(p + q)$) gave very similar results.

The stablised deviation was computed for every pair of trace subsets (considering all trace subsets with $a, b, c \in \{-4, \ldots, 4\}$ meant making over 1 million such comparisons) and the closest trace subsets listed. In Fig. 2 we show the closest few trace subsets to $\mathcal{E}_{1,0,0}$, ranked by their average stabilised deviation, and again for the closest few trace subsets to $\mathcal{E}_{0,1,2}$. In the former case there is a clear, small, list of trace subsets whose sizes are always much more closely-related to $\mathcal{E}_{1,0,0}$ than the others. The case of $\mathcal{E}_{0,1,2}$ is more complex, with quite a large number of trace subsets close to it (the list continues for about 50 trace subsets with deviation less than 0.5, an then a rapid increase with most other trace subsets deviating by at least 1.5).

| Deviation from $\mathcal{E}_{1,0,0}$ | | | |
|---|---|---|---|
| $\mathcal{E}_{1,0,0}$ | 0.0000 | $\mathcal{E}_{0,1,-1}$ | 2.4946 |
| $\mathcal{O}_{0,0,-1}$ | 0.0671 | $\mathcal{O}_{1,-1,0}$ | 2.5743 |
| $\mathcal{E}_{0,0,1}$ | 0.0988 | $\mathcal{O}_{0,1,-1}$ | 2.5953 |
| $\mathcal{E}_{-1,0,0}$ | 0.1726 | $\mathcal{O}_{0,-1,1}$ | 2.6323 |
| $\mathcal{O}_{-1,0,0}$ | 0.1778 | $\mathcal{O}_{-1,1,0}$ | 2.6414 |
| $\mathcal{O}_{0,0,1}$ | 0.2648 | $\mathcal{E}_{1,-1,0}$ | 2.7279 |
| $\mathcal{O}_{0,-1,0}$ | 0.3253 | $\mathcal{O}_{-1,0,1}$ | 2.7674 |
| $\mathcal{E}_{0,1,0}$ | 0.4210 | $\mathcal{E}_{0,-1,1}$ | 2.9676 |
| $\mathcal{E}_{0,0,-1}$ | 0.4403 | | ⋮ |
| $\mathcal{O}_{1,0,0}$ | 0.5275 | | |
| $\mathcal{E}_{0,-1,0}$ | 0.6426 | | |
| $\mathcal{O}_{0,1,0}$ | 0.7611 | | |

| Deviation from $\mathcal{E}_{0,1,2}$ | | | |
|---|---|---|---|
| $\mathcal{E}_{0,1,2}$ | 0.0000 | $\mathcal{O}_{2,-1,-1}$ | 0.2207 |
| $\mathcal{E}_{-2,-1,0}$ | 0.0207 | $\mathcal{E}_{1,1,-2}$ | 0.2280 |
| $\mathcal{O}_{0,1,2}$ | 0.0486 | $\mathcal{O}_{1,1,-2}$ | 0.2352 |
| $\mathcal{O}_{-1,1,2}$ | 0.0667 | $\mathcal{E}_{-1,2,0}$ | 0.2530 |
| $\mathcal{O}_{-2,-1,0}$ | 0.0697 | $\mathcal{E}_{1,-2,0}$ | 0.2657 |
| $\mathcal{O}_{-2,1,1}$ | 0.0803 | $\mathcal{O}_{1,-2,0}$ | 0.2782 |
| $\mathcal{E}_{-1,-1,2}$ | 0.0945 | $\mathcal{E}_{2,-1,0}$ | 0.2858 |
| $\mathcal{E}_{-2,1,1}$ | 0.1007 | $\mathcal{E}_{0,-2,1}$ | 0.2913 |
| $\mathcal{E}_{0,2,-1}$ | 0.1843 | $\mathcal{O}_{2,-1,0}$ | 0.3054 |
| $\mathcal{O}_{0,2,-1}$ | 0.1924 | $\mathcal{O}_{0,-2,1}$ | 0.3285 |
| $\mathcal{O}_{-1,2,0}$ | 0.2148 | | ⋮ |
| $\mathcal{E}_{2,-1,-1}$ | 0.2160 | | |

**Figure 2.** Left, the closest trace subsets to $\mathcal{E}_{1,0,0}$, ranked by average stablised deviation. Right, closest trace subsets to $\mathcal{E}_{0,1,2}$.

Producing all such lists, we then searched for patterns, trying to identify laws which explain all the close pairs of trace subsets. The conclusion is that the cover image symmetries for quadruplets all derive from:

1. Parity Symmetry: $e_{a,b,c} \approx o_{a,b,c}$.
2. Inversion Symmetry: $e_{a,b,c} \approx o_{-a,-b,-c}$.
3. Permutative Symmetry: for any permutation $\pi$, $e_{a,b,c} \approx e_{\pi(a,b,c)}$.

Such symmetries explain almost all of the close trace subsets we observed. Note that symmetries can be composed to show additional approximate equalities between trace subsets. For example, returning to Fig. 2, we see the low deviation between the sizes of $\mathcal{E}_{1,0,0}$ and $\mathcal{E}_{0,-1,0}$ can be explained by the following chain of reasoning: $e_{1,0,0}$ should be close to $o_{1,0,0}$ because of parity symmetry; this in turn should be close to $e_{-1,0,0}$ by inversion symmetry; this in turn should be close to $e_{0,-1,0}$ by permutative symmetry.

Furthermore, all these symmetries are quite plausible: they would follow from invariance of cover image properties under addition of one to each pixel, inversion of whole image, and local permutation of image, respectively. Because of the confluence of experimental evidence and plausible reasoning about natural images, we will accept this list of symmetries.

We should note, however, that we have not managed to explain quite all the observed close trace subsets. Some can be see in the right-hand table of Fig. 2: for example the closeness of $e_{0,1,2}$ and $o_{-1,1,2}$ is not explained by any of the listed symmetries, even in combination. However such "additional" symmetries seem always to be explained by the fact that two quadruplets which differ only by a single value at a single pixel will occur similarly often. A brief consideration of the trace subsets $\mathcal{E}_{0,1,2}$ and $\mathcal{O}_{-1,1,2}$ shows that the difference is only in an increment of the first pixel. We will discount such symmetries for now (although not without reservations) because they are not genuinely fourth-order symmetries. For the same reasons, and with the same reservations, we will restrict our attention to permutative symmetries only when $\pi$ is a cyclic permutation (i.e. when $\pi$ does not fix any element) because other permutative symmetries do not involve the whole quadruple. Although we will not use these symmetries in the detector, we must bear them in mind for what follows.

Now the symmetries of parity, inversion, and permutation tell us that lots of trace subsets should be equal to lots of others, in cover images; we must involve as many as possible in the Quadruples detector if we aim to maximise detection power. However we must ask three additional questions: 1) do all these cover assumptions distinguish innocent cover images from stego images, 2) are all cover assumptions equally subject to error, and 3) is violation of different cover assumptions independent? These questions are addressed in the following three subsections.

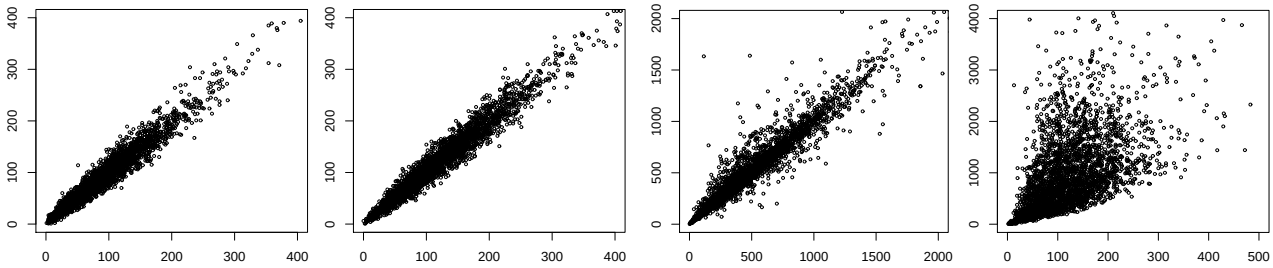

**Figure 3.** From left to right: scatterplots of $e_{0,-2,2}$ against $o_{0,-2,2}$ in covers; the sizes of the same trace subsets after maximal embedding; scatterplots of $e_{0,-1,1}$ against $o_{0,-1,1}$; the sizes of the same trace subsets after maximal embedding. The symmetry $e_{0,-2,2} \approx o_{0,-2,2}$ does not discriminate cover images from stego images whereas the symmetry $e_{0,-1,1} \approx o_{0,-1,1}$ does.

### 3.2. Identification of discriminating symmetries

Not all the symmetries are useful for detection of LSB steganography. This is true in the cases of Couples and Triples: as noted in Ref. 5 the symmetry $e_m \approx o_m$ was discarded for even $m$ in the Couples detector, because if true for covers it is also true for stego images when $m$ is even. Similar calculations caused us to discard parity symmetry $e_{m,n} \approx o_{m,n}$ in the Triples detector in Ref. 5 when *both* $m$ *and* $n$ are even.

The same problem occurs in the case of Quadruples. Consider Fig. 3: the left two scatterplots plot the trace subsets $e_{0,-2,2}$ against $o_{0,-2,2}$ in cover images and then maximally-embedded stego images. The cover symmetry is not broken by steganography. But the other two scatterplots in that figure show that the symmetry between $e_{0,-1,1}$ and $o_{0,-1,1}$ in covers *is* broken in stego images.

We must determine which of all the cover symmetries listed in Subsect. 3.1 are broken by LSB steganography. To answer the question, we consider what happens to each trace subset if *all* LSBs are flipped – this is at least easier than considering the case of a general amount of embedding. But it is still awkward to compute as it depends on the parity of the indices and there is no neat formula. Elementary calculations give the results shown in Tab. 1. Note that LSB flipping is self-inverse, so the table can be read either as indicating the resulting trace subset after LSB flipping, given the original, or as diagnosing the original trace subset given its flipped version.

| $a$ | $b$ | $c$ | $\mathcal{E}_{a,b,c}$ becomes | $\mathcal{O}_{a,b,c}$ becomes |
|---|---|---|---|---|
| even | even | even | $\mathcal{O}_{a,b,c}$ | $\mathcal{E}_{a,b,c}$ |
| odd | even | even | $\mathcal{O}_{a-2,b,c}$ | $\mathcal{E}_{a+2,b,c}$ |
| even | odd | even | $\mathcal{O}_{a,b-2,c}$ | $\mathcal{E}_{a,b+2,c}$ |
| odd | odd | even | $\mathcal{O}_{a-2,b+2,c}$ | $\mathcal{E}_{a+2,b-2,c}$ |
| even | even | odd | $\mathcal{O}_{a,b,c-2}$ | $\mathcal{E}_{a,b,c+2}$ |
| odd | even | odd | $\mathcal{O}_{a-2,b,c+2}$ | $\mathcal{E}_{a+2,b,c-2}$ |
| even | odd | odd | $\mathcal{O}_{a,b-2,c+2}$ | $\mathcal{E}_{a,b+2,c-2}$ |
| odd | odd | odd | $\mathcal{O}_{a-2,b+2,c-2}$ | $\mathcal{E}_{a+2,b-2,c+2}$ |

**Table 1.** The effect of flipping all LSBs, on each trace subset.

Now we take each cover image assumption in turn, e.g. $e_{1,1,-2} \approx o_{1,1,-2}$ (parity symmetry). After flipping all LSBs, we read from the table that $\mathcal{E}_{1,1,-2}$ will have cardinality $o_{-1,3,-2}$ and that $\mathcal{O}_{1,1,-2}$ will have cardinality $e_{3,-1,-2}$. But we can see that $e_{3,-1,-2}$ is assumed equal to $e_{-1,3,-2}$ (permutative symmetry) which in turn is equal to $o_{-1,3,-2}$ (parity symmetry). Therefore this cover image assumption is *not* useful for steganalysis because it is not violated by LSB flipping.

We can carry out a similar calculation for all cover symmetries identified in Subsect. 3.1. The results of such analysis is stated in the following 3 lemmas.

LEMMA 3.1. *Parity Symmetry discriminates between cover and stego images as long as either one or three of* $a$*,* $b$*,* $c$ *are odd, or if two of* $a$*,* $b$*,* $c$ *are odd and nonequal.*

*Proof.* There are eight cases, depending on the parity of $a$, $b$, and $c$. We illustrate only the case $a$ odd, $b$ odd, $c$ even. Parity symmetry states that $e_{a,b,c} \approx o_{a,b,c}$ in covers. Consider the effect of flipping all LSBs in a cover. Then the size of trace subset $\mathcal{E}_{a,b,c}$ becomes (according to Tab. 1) $o_{a-2,b+2,c}$ and the size of $\mathcal{O}_{a,b,c}$ becomes $e_{a+2,b-2,c}$.

Now we ask for what values of $a$, $b$, and $c$ is it possible that there is no discrimination, i.e. whether the symmetry still holds after LSB flipping. We must examine every possible combination of parity, inversion, and permutation symmetry to see whether they can, between them, deduce $o_{a-2,b+2,c} \approx e_{a+2,b-2,c}$. Because parity and inversion symmetries commute with any other symmetry, and the permutative symmetries form a group, it is only necessary to consider the 7 combinations which involve each type of symmetry zero or one time each.

Parity symmetry alone will explain the equation only if $a-2 = a+2$, $b+2 = b-2$, and $c = c$: impossible. Inversion symmetry alone will explain the equation if $a-2 = -(a+2)$, $b+2 = -(b-2)$, and $c = -c$. This implies that $a = 0$ but we assumed that $a$ was odd: impossible. We can rule out the possibility that the equation is explained by permutation symmetry alone, or by a combination of parity and inversion symmetry, or by all three, because such combinations can only relate trace subsets $e_{\ldots}$ to $e_{\ldots}$ and $o_{\ldots}$ to $o_{\ldots}$.

If the relation is to be explained by a combination of permutation and inversion symmetry, we must have that $(a-2, b+2, c)$ is a permutation of $(-(a+2), -(b-2), -c)$. Because of the parity of $a$, $b$, $c$, the only option is $a-2 = -(b-2)$, $b+2 = -(a+2)$, $c = -c$. These equations are inconsistent. Finally, if the relation is to be explained by a combination of permutation and parity symmetry, we must have $(a-2, b+2, c)$ is a permutation of $(a+2, b-2, c)$. Again because of the parity of $a$, $b$, and $c$, the only possibility is that $a+2 = b+2$, $a-2 = b-2$, $c = c$ which implies that $a = b$. In such a situation, we *can* deduce the symmetry $o_{a-2,b+2,c} \approx e_{a+2,b-2,c}$ and the original case of parity symmetry holds even after LSB flipping.

We have shown that, as long as $a$ and $b$ are odd and $c$ is even, only when $a = b$ does parity symmetry fail to discriminate between cover and stego images. The same calculations must then be repeated for each other parity combination of $a$, $b$ and $c$ to deduce the stated result. □

LEMMA 3.2. *Inversion Symmetry never discriminates between cover and stego images.*

*Proof.* It is easy to see, from Tab. 1, that if $\mathcal{E}_{a,b,c}$ turns into $\mathcal{O}_{p,q,r}$ when all LSBs are flipped, then $\mathcal{O}_{-a,-b,-c}$ turns into $\mathcal{E}_{-p,-q,-r}$. Therefore each case of inversion symmetry still holds after LSB flipping. □

LEMMA 3.3. *The case of Permutative Symmetry involving the cyclic permutation $e_{a,b,c} \approx e_{b,c,a}$ discriminates between cover and stego images as long as $a$ is odd and either: $b \neq a$ with $b$ odd and $c$ even, or $c \neq a$ with $c$ odd and $b$ even, or $b \neq c$ with both $b$ and $c$ odd. The other cyclic permutation discriminates in the same cases with $a$ and $c$ swapped.*

*Proof.* As for Lemma 3.1. □

Lemmas 3.1 and 3.3 give precise conditions under which we can make use of parity and permutative symmetries for the purposes of steganalysis. In view of Lemma 3.2 we will not make use of any instances of inversion symmetry.

### 3.3. Variance stabilisation

We must take care how we interpret deviations from the cover symmetries. Take, for example, the parity symmetry $e_{0,1,2} \approx o_{0,1,2}$. If the two quantities involved are large, i.e. lots of groups of pixels have successive differences $0, 1, 2$, then we might expect larger natural deviations between the two quantities $e_{0,1,2}$ and $o_{0,1,2}$. See Fig. 4, left, which clearly shows that the deviations from the symmetry vary in a way highly dependent on the size of the trace subsets involved. We must account for this in any analysis of deviations: we need a variance-stablising function.

We propose the following idea of a "platonic natural image" which explains deviation from the cover symmetries' exact equalities. Suppose a symmetry equates $a$, the size of a trace subset $\mathcal{A}$, with $b$, the size of a trace subset $\mathcal{B}$. In a natural image, we imagine that the total number of quadruples in these trace subsets $a+b$ is fixed, but that each is allocated to $\mathcal{A}$ or $\mathcal{B}$ independently, at random.

LEMMA 3.4. *Under this assumption, $a-b$ is approximately normally distributed with mean 0 and variance $a+b$. Therefore $(a-b)/\sqrt{a+b}$ has variance independent of the magnitude of $a$ and $b$.*

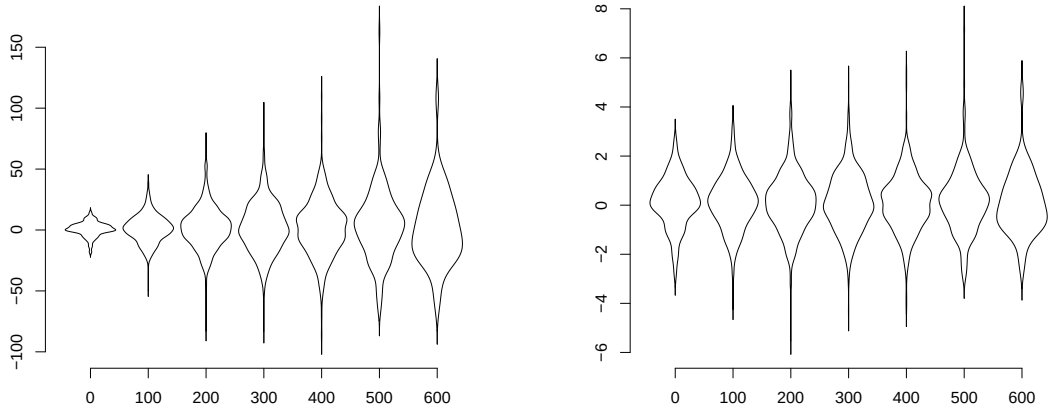


**Figure 4.** Left, densities ("violinplot") of $e_{0,1,2} - o_{0,1,2}$ across 3000 cover images, grouped by magnitude of $e_{0,1,2} + o_{0,1,2}$ rounded to the nearest 100 (*x-axis*). Clearly the variation depends heavily on the magnitude. Right, $(e_{0,1,2} - o_{0,1,2})/\sqrt{e_{0,1,2} + o_{0,1,2}}$, indicating that variance has been approximately stabilised in this case.

*Proof.* We have that $a$ is binomially distributed with parameters $a + b$ and $1/2$. The normal approximation holds under very weak assumptions ($a + b > 10$ is usually considered sufficient, and the sizes of our trace subsets should be orders of magnitude higher than this) which says that $a$ is normal with mean $(a + b)/2$ and variance $(a + b)/4$. Now $b = a + b - a$ and we assumed $a + b$ was fixed, so that $b$ is determined by $a$. Then $a - b = a - (a + b - a) = 2a - (a + b)$, with the second term constant, and so this has a normal distribution with mean 0 and variance $a + b$. □

We call $(e_{a,b,c} - o_{a,b,c})/\sqrt{(e_{a,b,c} + o_{a,b,c})}$ the *stabilised* deviation from the parity symmetry assumption, and analogously for the other symmetries. Note that this coincides with the component of a $\chi^2$ test for equal distribution between all $e_{a,b,c}$ and all $o_{a,b,c}$. Fig. 4 shows density plots for raw and stabilised deviations, illustrating that the scaled deviation has a more stable distribution, as the size of the trace subsets varies, than the raw deviation.

We add here that an idea of a "platonic natural image" deserves much more scrutiny, and the suggested process is hardly a plausible description of the way images are created. But this issue is beyond the scope of this work. For now we accept the conclusion of how to compute a stabilised deviation, without necessarily accepting the assumption which suggested it.

Finally, we note that we have stabilised the variance with respect to the magnitude of the trace subsets involved – this is *not* to say that the variation in such quantities is necessarily stable as $a, b, c$ vary, but further investigation shows that this also to be approximately true, with some exceptions for zero values of $a$, $b$ and $c$.

### 3.4. Independence of symmetries

Finally, we consider independence of different cover image assumptions. If we observe two cover symmetries whose deviation, in natural images, is highly correlated then we should exclude one of the two.

Figure 5 shows the correlation coefficient of stabilised deviations from one particular parity symmetry (respectively permutative symmetry), against a large number of other parity symmetries (respectively permutative symmetries), again computed using the set of 3000 grayscale bitmaps. We observe that the correlations are quite small (of course the outlier is the exact correlation between the symmetry and itself). Statistical tests of independence (details omitted) show that about half of the correlations between parity symmetries *are* significantly higher than zero, but we can see that none are very highly correlated. On the right of Fig 5 we see that there is more dependence between different instances of permutative symmetry than parity symmetry, but still most instances are only weakly correlated.

We performed a systematic computation of correlation between stabilised deviations from every pair of discriminating symmetries, subject to the limitation that all indices were in the range $\{-4, \ldots, 4\}$. There are approximately 1000 such symmetries, leading to approximately 1 million tests. Although many symmetries have

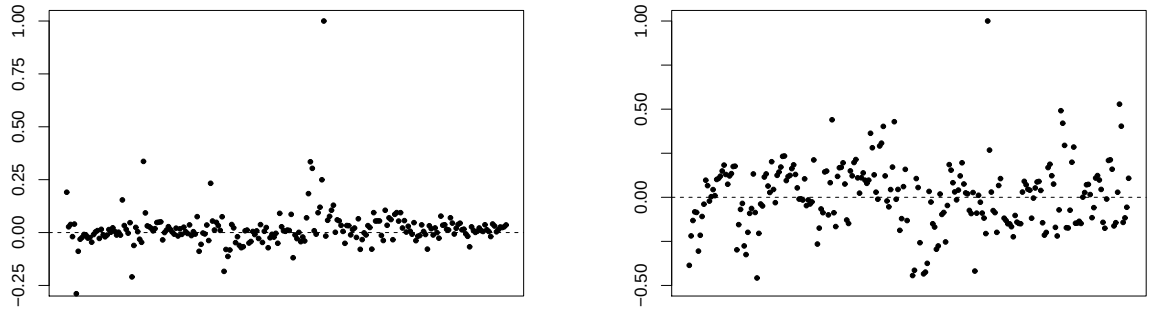


**Figure 5.** Left, plot of correlation coefficient of stabilised deviations from $e_{0,1,2} \approx o_{0,1,2}$ and $e_{a,b,c} \approx o_{a,b,c}$ for all $a, b, c \in \{-3, \ldots, 3\}$. Right, plot of correlation coefficient of stabilised deviations from $e_{1,0,0} \approx e_{0,0,1}$ and $e_{a,b,c} \approx e_{b,c,a}$ for all $a, b, c \in \{-3, \ldots, 3\}$. In both cases the outlier is the pair of identical symmetries.

significantly greater than zero correlation, almost none have strong correlation above about 0.5. The exception is that correlation between permutative symmetries $e_{a,b,c} \approx e_{\pi_1(a,b,c)}$ and $e_{a,b,c} \approx e_{\pi_2(a,b,c)}$, for different $\pi_1$ and $\pi_2$ (but the same $a, b, c$) are sometimes (although, curiously, not always) very highly correlated. We conclude that we may make use of all discriminating cover symmetries, but should restrict our attention to a single permutation in the case of permutative symmetry.

## 4. A QUADRUPLES STEGANOGRAPHY DETECTOR

With the discriminating cover assumptions identified, we come to the final stage of the detector. The detector is in fact an estimator for $p$: such detectors are sometimes called *quantitative steganalysis*. First, we must decide how to count the trace subsets themselves; although there are many options it does not seem to be an important question. Then we must turn the cover symmetries into approximate equations for $p$; there are some significant choices to be made here. Finally, we measure the performance of the resulting estimators.

### 4.1. Counting trace subsets

The first decision we must make is how to define the trace subsets themselves. Recall that each trace subset is a subset of some parent set $\mathcal{T}$, a set of (in this case) quadruplets of spatially-close pixel values extracted from a particular image under consideration. In the case of Couples we generally use all pairs of adjacent pixels; in the case of Triples, in Ref. 5, we elected to use all horizontal rows of 3 pixels. We have a number of sensible options when the group size is as high as 4:

(i) Use horizontal or vertical blocks of 4-by-1 pixels.

(ii) Use 2-by-2 blocks, considering the block in the order $\begin{array}{|c|c|}\hline 1 & 2 \\ \hline 3 & 4 \\ \hline \end{array}$ (or its transpose).

(iii) Use 2-by-2 blocks, considering the block in the order $\begin{array}{|c|c|}\hline 1 & 2 \\ \hline 4 & 3 \\ \hline \end{array}$ (or its transpose).

Each of these options can be modified to force the groups of pixels used to be disjoint, or to allow overlaps (whereby each pixel appears in 4 different trace subsets – strictly speaking this violates an independence assumption needed for the structural framework of Ref. 5, but the effect is small enough to be negligible).

In fact, this question appears to be a non-issue. We repeated a number of the experiments which follow with each of the above options, and observed practically no difference in overall detector performance (although there were particular cases when one method or the other appeared a bit better). Whichever we use, we are considering the evidence of all the pixels in the image and it is reasonable to believe that it does not make much difference what order we use them in, and whether we use the same pixels repeatedly. Without clear evidence in favour of any one method we chose option (iii), allowing overlapping blocks.

### 4.2. Description of detector

A symmetry such as $e_{a,b,c} \approx o_{a,b,c}$ gives rise to an approximate equation for $p$ as follows. First, select the appropriate values for $l, m, n$ so that Eq. (1) applies to the trace subsets involved (it may be different $l, m, n$ for $e_{a,b,c}$ than for $o_{a,b,c}$) and take the component required. Counting the quantities $E'_{2l,2m,2n}, \ldots, O'_{2l,2m,2n}$ in the image under consideration, and assuming both that the approximate equality given by the cover symmetry is exact, and that Eq. (1) is exact, gives an equation for $p$.

We omit the equation itself – the calculations are rather tedious and an analogue of those already given in Ref. 5. We note that it is very convenient to make the substitution $q = 1/(1-2p)$, for then $T_4^{-1}(q)$ is in the form of the 4-fold Kronecker tensor product of the matrix

$$\frac{1}{2}\begin{pmatrix} 1+q & 1-q \\ 1-q & 1+q \end{pmatrix}$$

with itself. This leads to a polynomial equation in $q$, easier to solve than the rational polynomial equation for $p$ achieved by direct application of Eq. (1).

The procedure is not quite as simple as suggested in the previous paragraph, because the equation for $q$ is a polynomial of degree 4. As such it might have no roots, or multiple roots. A similar, but simpler, situation is seen in Sample Pairs/Couples Analysis, where it can be shown that the lower root of two is almost always correct, and that the hidden message length is usually close to 1 when there are no roots. In the Quadruples case we have no such simple solutions. We propose that the best way to solve the multiple-root problem is first to compute a prior estimate for $p$ using some other steganalysis method (we used the Triples estimator) and then to select whichever root of the quartic equation for $q$ leads to $p$ closest to the prior estimate. In practice this seems to work quite well, although it is a bit unwieldy. The no-root problem seems insoluble, and all we can do is fail to give an estimate in that case. Thankfully, we will see below that we have multiple equations and that in practice not all of them will fail to have roots.

So each cover symmetry gives a separate estimator for the unknown $p$ (unless it fails through lack of roots). But any individual symmetry only involves a very small part of the image, because it only uses two out of a large number of trace subsets. Even limiting to small values of $a, b, c$ (for example to those with an absolute value less than 4) leads to hundreds of discriminating examples of parity symmetry alone, and therefore hundreds of equations which estimate $p$. If we want to make use of all the evidence in the image, we must somehow combine all these estimators.

We propose three possible options, each of which is inspired by other steganalysis methods. Suppose that each cover symmetry is expressed as a deviation $\epsilon_i(p) \approx 0$. Then either:

(i) Solve $\sum \epsilon_i(p) = 0$ for $p$. This amounts to summing all the quartic polynomials in $q$, described above, and is analogous to the approach used by Dumitrescu in Ref. 3.

(ii) Find $p$ to minimise the squared error $\sum \epsilon_i(p)^2$. Closest in spirit to the principle "find the value of $p$ to match most closely the cover image model" this is the approach used for the Least Squares Method[6] and Triples.[5] In this case we have to solve a polynomial of degree 7 in $q$.

(iii) Solve each equation $\epsilon_i(p) = 0$ individually, producing separate estimators $p_i$, and then combine these estimators by taking their mean or median. This approach is similar to ideas suggested in Ref. 4 for improving steganalysis performance.

In view of the discussion of Subsect. 3.3 we might prefer to use the *stabilised* deviation for $\epsilon_i(p)$. However this leads to a practical problem: the quantity $(e_{a,b,c} - o_{a,b,c})/\sqrt{e_{a,b,c} + o_{a,b,c}}$ (to take a case of parity symmetry, for example) is quite a complex function of $p$. Certainly it is not a rational polynomial, and there seems to be no substitution to turn it into an equation which we can solve with any simplicity. In fact, the stabilised deviation is a rather pathological function of $p$ with many extreme minima. However there is an alternative method of stabilisation: instead of dividing by the square root of the trace subset sizes themselves, we divide by the square root of $E'_{2l,2m,2n} + \cdots + O'_{2l,2m,2n}$, where the sum is the 16 elements of the vector displayed in Eq. (1) and $l, m, n$ are such that $E'_{a,b,c}$ appears in the sum (when $O'_{a,b,c}$ appears in a different sum, we must add both sets of 16 elements together). This quantity has the very useful property of being invariant under LSB operations (see

Ref. 5 for an explanation), while still being a good proxy for the size of $e_{a,b,c}$ and $o_{a,b,c}$, so dividing by it does not increase the complexity of the deviation, as a function of $p$.

To decide which of these procedures to use, we performed some initial experiments. It rapidly became clear that (i) and (ii) produce quite poor detectors. Although there is some improvement when the stabilised deviations are used (which amounts to a weighting on the individual quartic equations for $q$), the detectors remain less accurate than others in the literature, in particular than the Triples Analysis detector of Ref. 5. The reasons for this bear further study, but it appears to be due to dramatic violations of a few of the cover symmetries exerting a large influence (even after weighting) on the overall result. Perhaps the larger number of symmetries involved in Quadruples Analysis means that there are more opportunities for outliers. For now, we discarded these detectors.

Option (iii), however, produces a good detector. Given the heavy-tailed nature of steganalysis estimation errors[7] we might expect that taking the median of the individual estimators would be more robust to individual outliers in the symmetry assumptions, but in fact as long as we discard implausible estimates of $p$ (e.g. those outside the range $[-0.1, 1]$ or more than 0.5 from the prior estimate by Triples), it seems that both the mean and median produce good results and we test both in the following subsection.

Some further experiments were performed, to select which cover symmetries we should use. We settled on use of all discriminating examples of parity symmetry $e_{a,b,c} \approx o_{a,b,c}$, and all discriminating examples of the cycle permutative symmetry $e_{a,b,c} \approx e_{b,c,a}$, for $a, b, c$ in the range $\{-3, \ldots, 3\}$. There are 400 such symmetries, and so we are taking the mean or median of up to 400 individual estimators (there will be less than 400 in practice, because some give rise to equations with no roots). We restricted our detector to symmetries with small indices for the following reason: trace subsets with larger indices are rather rare, and our use of the Law of Large Numbers to reach Eq. (1) is only justified if the sizes of the subsets are large; rarer trace subsets give rise to estimators so wildly unpredictable that they add nothing to the final outcome.

### 4.3. Experimental results

We have compared the accuracy of Quadruples Analysis, as defined above (both mean and median versions), with leading LSB replacement detectors from the literature: the method of RS,[1] Couples Analysis (practically equivalent to Sample Pairs Analysis[3]), the Least Squares Method variation on Couples[6] (the published version of which contains some bugs, which we corrected), and Triples Analysis.[5] Using the distributed steganalysis project outlined in Ref. 8 we tested these detectors against the set of 3000 cover images used in the previous section. Tests were repeated before and after these images were converted to grayscale, so as to see any performance differences in colour covers[‡]. So as to separate characteristics of uncompressed and previously JPEG compressed covers (where, in previous work,[4, 5] we have noted a substantial difference in steganalysis performance) we also repeated the tests with the cover images subject to JPEG compression and decompression prior to embedding. For JPEG compression we selected "quality factor" 90, representing mild compression; we also experimented with other quality factors, but the results produced were not different enough to warrant inclusion in this paper.

Note that in this section we will redefine $p$ to mean the proportionate hidden message length, i.e. the hidden message length is now $pN$ rather than $2pN$ as before. This change of notation is unfortunate, but corresponds to the usual conventions of the literature.

We begin by showing some simple histograms of the estimators, in Fig. 6. These compare the distributions of the Triples and Quadruples estimators (both forms of Quadruples, using either the mean or median of individual estimators). It is immediately apparent that the distributions of the Quadruples estimators are rather different from those of Triples Analysis: it appears that they are asymmetric, with a much lighter right tail than left tail, and overall a lower dispersion (more accuracy). They also seem to have some negative bias. In the first of the displayed histograms, for example, the median of the Triples estimators is -0.00147, but the median of

[‡]The LSB replacement process we tested treated each colour channel separately and embedded in all channels equally; the detector counted the trace subsets of each channel separately and then totalled each trace subset across the 3 channels prior to making the detector computations. There is scope for further research into optimization of both embedding and detection stages in the case of colour images.

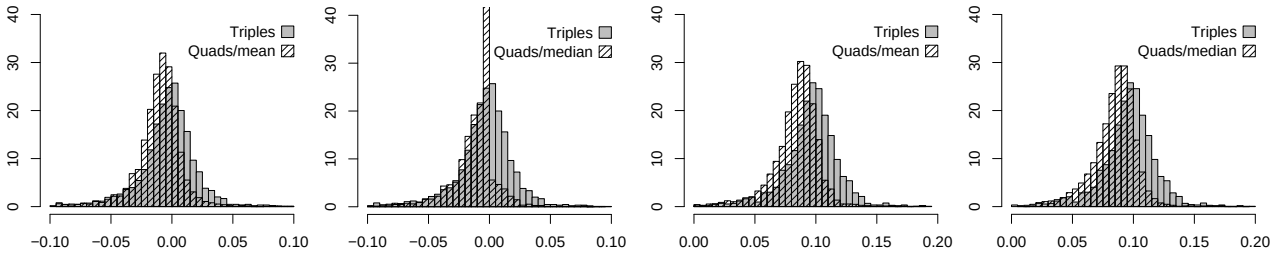


**Figure 6.** Histograms of Triples and Quadruples estimators, computed from 3000 never-compressed grayscale bitmaps. From left to right: the Triples estimator compared with the Quadruples mean and median estimators (with no embedded data); the same when data is embedded at rate 0.1.

the Quadruples estimators is -0.00868. We will return to the negative bias shortly, but for now we focus on the reduced dispersion of the Quadruples estimators.

When $p = 0$, the Triples estimator has an interquartile range of 0.0229, whereas the Quadruples/mean and Quadruples/median estimators have interquartile ranges of 0.0179 and 0.0149 respectively. In the case when $p = 0.1$, the value for Triples is 0.0231, and for the two Quadruples estimators 0.0196 and 0.0215. If we choose to measure dispersion in terms of standard deviation (which might not be wise, as we have no evidence that the distributions even have finite variance – and if not, then the sample standard deviation will not converge!) then we also see a reduction (of roughly 20%) when moving from Triples to Quadruples. Either way, the Quadruples estimator is more accurate than the Triples estimator.

In fact, the Quadruples estimator should have an additional improvement not shown by these numbers, because of the right tail of the distribution appearing to be light. It is the right tail which is so important for discriminating between cover and stego images, because over-estimates correspond to false positive errors. In this paper we do not make further study of the detectors' discrimination ability (both for reasons of space and because we believe that the Quadruples method should be optimised further, first).

We return to the apparent negative bias in the Quadruples detectors. At present we have no explanation for this phenomenon and expect to perform further study – hopefully to correct it. For now, we try to measure it. In our experiments the bias was more apparent in colour images than grayscale images, so we focus on the former by using our set of 3000 images prior to their conversion to grayscale. Each of these 3000 colour bitmaps had random messages embedded by LSB replacement, with messages of length $0, 0.05, 0.1, \ldots, 0.5$ of the maximum. In Fig. 7, left, we plot the observed bias of Couples, Triples, and Quadruples estimators. In view of possible long tails, we plot the median observed detection value as opposed to mean, although in fact this would make very little difference to the outcome displayed. On the right of Fig. 7 we plot the interquartile range of the estimators.

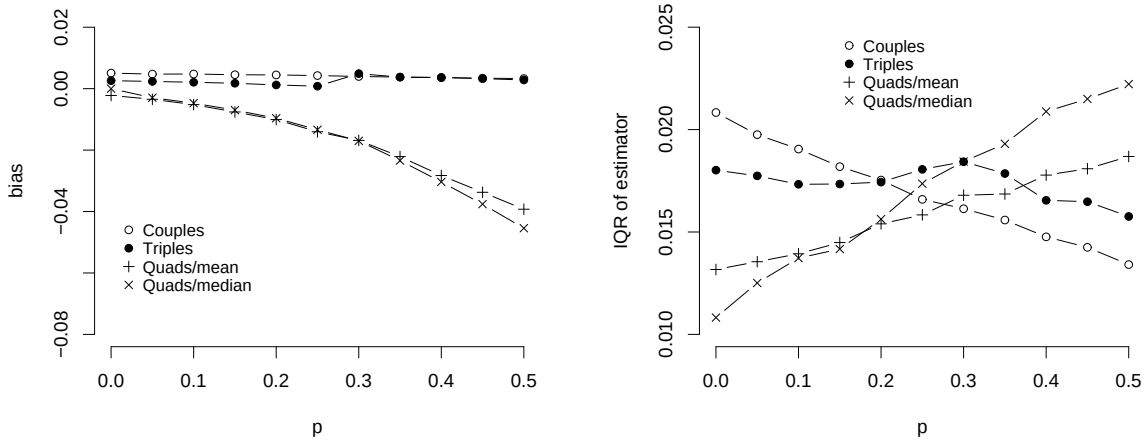


**Figure 7.** Left, bias of detectors, observed in colour never-compressed covers. Right, interquartile range.

**Table 2.** Interquartile range (standard deviation) of estimators, $\times 10^2$, when the true value of $p$ is zero.

| | Never-compressed covers | | JPEG (q.f. 90) covers | |
|---|---|---|---|---|
| | Grayscale | Colour | Grayscale | Colour |
| RS | 2.31 (3.25) | 2.15 (2.67) | 1.78 (2.02) | 8.22 (10.94) |
| Couples | 1.96 (3.29) | 2.08 (2.56) | 1.44 (1.73) | 7.10 (8.56) |
| Couples/LSM | 2.23 (3.59) | 2.41 (2.96) | 1.46 (1.86) | 4.45 (3.90) |
| Triples | 2.29 (3.45) | 1.80 (2.36) | 1.09 (1.90) | 1.40 (2.08) |
| Quadruples/mean | 1.79 (2.73) | 1.32 (1.56) | 1.42 (1.69) | 2.17 (2.03) |
| Quadruples/median | 1.49 (2.56) | 1.08 (1.45) | 1.40 (1.62) | 1.72 (2.03) |

The negative bias of the Quadruples estimators increases with $p$, and apparently the relationship is nonlinear. However the absolute bias remains small at least for values of $p$ less than about 0.2. Very close to $p = 0$ the negative bias is only about as large as the positive bias found in the Couples and Triples detectors. Figure 7, right, is particularly interesting. It shows clearly that the Quadruples detectors are superior for small values of $p$ (less than about 0.2) and inferior for higher values of $p$, more so for the median Quadruples estimator than the mean. The progression from a second-order structural detector (Couples) through a third-order detector (Triples) to our new fourth-order detectors is also apparent.

Taking these two charts together, we must conclude that the Quadruples estimators are going to be rather better for discrimination (between the case $p = 0$ and $p = p_1$) than the estimation for which they were designed. Bias is irrelevant to the discrimination problem. Also, the Quadruples detectors will be superior to the methods of Couples and Triples as long as $p_1$ is fairly small. Of course, the case of $p_1$ small is the most interesting one, when detection is difficult.

In order to test the performance of the Quadruples detectors against a wider range of other algorithms, and over different types of cover images, we now restrict our attention to the case of small $p$, as we did in Ref. 5. Since the distribution of the estimators is at least continuous in $p$, the error distribution when $p = 0$ serves as a good measure here. Bearing in mind the long-tailed distribution of steganalysis estimators[7] we compute accuracy of an estimator both by its standard deviation and the interquartile range, the latter for robustness to outliers. These values are displayed in Tab. 2.

The table tells us something that we did not previously know about the Triples detector: its performance on grayscale never-compressed images is disappointing, and a little inferior to the standard Sample Pairs method (in Ref. 5 we only tested on colour images, and therefore missed this feature). The good news is that in this case the Quadruples detectors provide a good improvement whereas, in the case of colour JPEG covers, the Quadruples detectors cannot quite match the extraordinarily good performance, relative to the others, achieved by the Triples method.

It is interesting to compare magnitudes of interquartile range and standard deviation shown in Tab. 2. The latter is sensitive to the weight of heavy tails, whereas the former disregards the distant tails altogether. Even in cases when the Quadruples estimator has a higher interquartile range than the Triples estimator, it has a lower standard deviation. These results suggest that the improved reliability of the Quadruples estimator takes the form of a reduction in the number of egregious errors. This can be confirmed by inspection of the histograms of the estimators, as in Fig. 6.

Of course, comparing detectors in the case $p = 0$ is favourable to the Quadruples detector (and, to a much lesser extent, the Triples detector) and we must note that, as $p$ increases, the relative performance changes in favour of the traditional second-order detectors. But for small $p$, in particular for $p$ less than about 0.2 (for colour images) or 0.1 (for grayscale images), it remains the superior of all the tested detectors. This suggests that a fruitful steganalysis procedure might be first to test an image using a standard method such as Couples, proceeding to make a more accurate Quadruples estimate only if the prior estimator indicates that the true value of $p$ is small.

By the standards of our previous work[4, 5, 8] the Quadruples method has received relatively little testing here. The aim was only to indicate the value of the fourth-order structural method. We expect to produce improved fourth-order structural detectors, for which more in-depth testing would be appropriate.

## 5. CONCLUSIONS

We have extended the method of Triples Analysis to Quadruples – the main complexity was in determining the cover image assumptions, and selection of those which discriminate between cover and stego images. Based on the general framework of Ref. 5 and quadruplet cover symmetries described here, we have implemented a detector. It is on balance a superior detector of short hidden messages, but the advantage over the Triples method is not enormous, suggesting that optimal structural steganalysis may be found for group sizes close to 4. Also, the new detector is substantially more complex than previous generations of structural detectors.

However we do not claim that our use of the fourth-order structural properties is in any way optimal. The main aim of this paper has been to present a fourth-order detector, with focus on finding the cover assumptions, rather than optimising performance. We have not investigated whether the decision to discard cover symmetries which were not "genuine" fourth-order was detrimental. And there remains the possibility that quadruple-based methods might provide a further substantial improvement because we have used fairly naive methods at the final stage when the cover assumptions were combined with equations given by the structural properties. Potentially we could make much better use of the experimental evidence about cover symmetries: speculatively, we propose that it might be possible to create a learning machine which is trained on cover images and forms its own conclusion about the weight to be given to various cover assumptions, the extent to which they discriminate cover images from stego images, and dependencies amongst them.

Finally, if we were to proceed still further to fifth-order structural steganalysis we would run into particular problems with the analogue of Eq. (1), because the use of the Law of Large Numbers to equate the realization of the trace subsets $E'_{\vec{x}}$ and $O'_{\vec{x}}$ with their expectations will be less dependable: the higher the order of the analysis, the more fragmented the trace subsets and the smaller the numbers involved.

## ACKNOWLEDGMENTS

The author is a Royal Society University Research Fellow.

## REFERENCES

1. J. Fridrich, M. Goljan, and R. Du, "Reliable detection of LSB steganography in color and grayscale images," *Proc. ACM Workshop on Multimedia and Security*, pp. 27–30, 2001.
2. J. Fridrich, M. Goljan, and D. Soukal, "Higher-order statistical steganalysis of palette images," in *Security and Watermarking of Multimedia Contents V*, E. J. Delp III and P. W. Wong, eds., *Proc. SPIE* **5020**, pp. 178–190, 2003.
3. S. Dumitrescu, X. Wu, and Z. Wang, "Detection of LSB steganography via sample pair analysis," in *Proc. 5th Information Hiding Workshop*, *Springer LNCS* **2578**, pp. 355–372, 2002.
4. A. Ker, "Improved detection of LSB steganography in grayscale images," in *Proc. 6th Information Hiding Workshop*, *Springer LNCS* **3200**, pp. 97–115, 2004.
5. A. Ker, "A general framework for the structural steganalysis of LSB replacement," in *Proc. 7th Information Hiding Workshop*, *Springer LNCS* **3727**, pp. 296–311, 2005.
6. P. Lu, X. Luo, Q. Tang, and L. Shen, "An improved sample pairs method for detection of LSB embedding," in *Proc. 6th Information Hiding Workshop*, *Springer LNCS* **3200**, pp. 116–127, 2004.
7. R. Böhme, "Assessment of steganalytic methods using multiple regression models," in *Proc. 7th Information Hiding Workshop*, *Springer LNCS* **3727**, pp. 278–295, 2005.
8. A. Ker, "Quantitive evaluation of Pairs and RS steganalysis," in *Security, Steganography, and Watermarking of Multimedia Contents VI*, E. J. Delp III and P. W. Wong, eds., *Proc. SPIE* **5306**, pp. 83–97, 2004.