

Empirical Validation of the Square Root Law of Steganography



Catherine Vlasov
University College
University of Oxford

Honour School of Computer Science
Computer Science Project (Part C)

May 2019

Abstract

According to the Square Root Law, the detectability of a hidden message stays roughly constant if its length is proportional to the square root of the size of the space in which it hides. This is both a collection of theorems about probabilistic models of the medium and an empirical observation. However, the most recent investigation of the phenomenon in practice was done in 2008. Since then, a number of new techniques for hiding (Syndrome-Trellis codes, adaptive embedding) and detection (rich feature models) have been developed.

This project parallels the experiments done in 2008. It provides fresh results by running computation-intensive experiments on two large sets of images, which show that the Square Root Law continues to hold with the current best-known algorithms.

Contents

1	Introduction	1
1.1	Outline	2
2	Background	3
2.1	Steganography	3
2.2	JPEG Compression	5
2.2.1	Decompression	7
2.3	JPEG Steganography	7
2.3.1	Progress in the Last Decade	8
2.3.2	Distortion Function Minimization	8
2.3.3	J-UNIWARD	11
2.4	JPEG Steganalysis	11
2.4.1	Progress in the Last Decade	11
2.4.2	JRM	13
2.4.3	Low-Complexity Linear Classifier	14
2.5	Square Root Law	14
2.5.1	Batch Steganographic Capacity	14
2.5.2	Relation to Capacity of Individual Cover Objects	15
2.5.3	Perfect Versus Imperfect Steganography	15
2.5.4	Caveats	15
2.6	Project Motivation, Goals, and Hypothesis	17
2.6.1	Motivation	17
2.6.2	Goals	17
2.6.3	Hypothesis	18
3	Experiment Design	19
3.1	Images	19
3.1.1	Image Sets	19
3.1.2	Preprocessing	20
3.1.3	Cropping	21
3.2	Embedding	22
3.3	Costs	22
3.4	Features	23

3.5	Classifier	23
3.6	Pipeline	23
3.7	Payload Sizes	25
3.8	Experiments	26
4	Results & Analysis	27
4.1	Experiment Result Visualizations	27
4.1.1	ROC Curves	27
4.1.2	Detectability Versus Cover Size	28
4.2	Results From a Decade Ago	29
4.3	Actor3	29
4.3.1	ROC Curves	29
4.3.2	Detectability Versus Cover Size	32
4.4	BOSSbase	33
4.4.1	ROC Curves	33
4.4.2	Detectability Versus Cover Size	34
5	Conclusion	37
5.1	Further Work	37
5.2	Critical Evaluation	38
	Appendix A	47
A.1	Image Sizes	47
A.2	Proportionality Constants	48
A.3	Payload Sizes	48
A.3.1	Actor3	48
A.3.2	BOSSbase	50

Chapter 1

Introduction

On July 5, 2018, a General Electric (GE) employee emailed himself an image of a sunset [3]. On April 23, 2019, he was arraigned by U.S. authorities [29]. Why? The answer was hidden in plain sight.

Information can be shared over public channels in plaintext or encrypted and in both cases, all messages are visible to eavesdroppers. They can tell which people are communicating, how often they do so, and how much information they are sharing. However, a third option is gaining popularity: hiding one message inside another to make the hidden message invisible to eavesdroppers. This is called steganography: the art of hiding information.

The sunset image that the employee emailed to his personal address while at work looked innocent to the naked eye – and to GE’s intrusion detection system – but it concealed 40 Excel and MATLAB files [3]. Banner advertisements [28], memes [40], and even the 16×16 pixel “favicons” shown beside website names in browser tabs [32] have been used for steganography – in many cases to spread malware. It is therefore important for the security community to understand how steganography works and how to detect it in the real world, where theoretical assumptions do not hold.

Intuitively, it should be easier to identify that an image has a hidden payload if the payload is large than if it is small. This idea is formalized by the “square root law”, which quantifies the relationship between payload size and the accuracy of a binary classifier whose task is to identify images with a hidden payload. Informally, the law states that detectability remains constant if the payload size is proportional to the *square root* of the number of pixels in an image.

The law was conjectured in 2006 [18] then empirically validated in 2008 using state-of-the-art (at the time) embedding and detection algorithms [24]. However, the results in [24] are now – more than a decade later – out of

date given how far the field has progressed. The aim of this project is to empirically validate the law using the current best available tools and therefore provide fresh results. By running computation-intensive experiments on two large sets of images, we find that the square root law still appears to hold.

As discussed in [24], this result is significant because it puts into question what unit of measurement is most meaningful when describing the capacity of an image (the amount of information that can be hidden inside it). Instead of “bits per pixel”, should it be “bits per square root pixel”? The square root law also has practical implications for users of steganography algorithms. When deciding how much payload to embed, should they consider the maximum number of bits that can “physically” fit in an image or some other measurement?

In addition to mirroring the experiments in [24], we explore an idea mentioned in the paper’s conclusion. It is suggested that steganographic capacity might be more precisely modelled as being proportional to $\sqrt{N} \log N$ (where N is the number of pixels in an image) due to the behaviour of modern embedding techniques. We find that this does appear to be the case, which is a new result.

1.1 Outline

Chapter 2 provides background on the field of steganography, including necessary terminology and progress made in the field since 2008, when [24] was published. I will then give technical descriptions of the algorithms I chose, a precise mathematical formulation of the square root law, and an overview of important caveats related to the law’s practical application. Finally, I will discuss the project’s motivations, goals, and hypotheses.

Chapter 3 covers the design, implementation, and running of the experiments. I outline the experiment requirements and then explain each component of my experiment pipeline. The chapter provides sufficient detail to replicate my results, given access to the image sets and relevant code.

Chapter 4 presents and analyzes the experiment results, which align with the expected behaviour.

Finally, Chapter 5 summarizes my findings, describes interesting areas of further research, and discusses how the project goals were achieved.

Chapter 2

Background

To see why this project was undertaken, we need to understand its context: the current state of steganography as well as its state ten years ago.

I will start by introducing key steganography terminology and concepts and describe how JPEG compression works. Then I will discuss some of today’s best-known embedding and detection algorithms and compare them to the best techniques when [24] was written. Finally, I will give a precise formulation of the square root law and lay out this project’s motivation, goals, and hypotheses.

2.1 Steganography

The word “steganography” comes from the combination of the Greek words “steganos” (“covered”) and “graphia” (“writing”) [11]. The first known use of steganography [11] is documented in Herodotus’ influential *The Histories* [15], written in 440 BC.

Simmons’ famous prisoners’ problem [35] is typically used to introduce the field of steganography. Two Prisoners are separated but are allowed to communicate via a medium monitored by a Warden. They want to discuss an escape plan, but can only do so by establishing a **subliminal channel** that they use to communicate by hiding secret information in innocent-looking messages. Their goal is to prevent the Warden from detecting that they are communicating secretly (as opposed to preventing the Warden from reading their hidden messages).

So far, we have informally been referring to two concepts: hiding (also called embedding) and detection. In the context of the prisoners’ problem, the Prisoners are Embedders and the Warden is a Detector. These concepts are in fact part of two separate but related fields: **steganography** and

steganalysis. The former is the study of *hiding* information whereas the latter is the study of *detecting* hidden information. An object that has not been modified is referred to as a **cover object** and an object with an embedded message is called a **stego object**.

The input to an **embedding algorithm** is a cover object, a secret message (called the **payload**), and a secret key shared by the sender and receiver. The output is a stego object. An embedding algorithm without a corresponding method for reading the secret message is useless to both the sender and receiver, so every embedding algorithm has a matching **extraction algorithm**. Its input is a stego object and a secret key and its output is the hidden payload. The embedding and extraction algorithms, together with the sources of cover objects (such as a camera or microphone), secret keys, and payloads as well as the communication channel (such as an email service), make up a **stegosystem** [20].

It is worth noting that in steganalysis, it is typically assumed that the Warden knows which embedding method the Prisoners are using. This is a conservative assumption in line with Shannon’s paraphrasing of Kerckhoffs’ principle: “the enemy knows the system” [34] (in this case the stegosystem).

In a modern context, the objects used for steganography are commonly digital media such as images, audio files, and videos [23]. In this project we will study steganography in JPEG images, which makes use of properties of the JPEG compression algorithm (described in Section 2.2) to hide information.

The aim of steganography is to make it impossible for an eavesdropper to decide whether or not an object contains a hidden payload. A stegosystem that fulfills this goal is called **secure**. In addition, an aim of steganography researchers is to design embedding and extraction algorithms that maximize a stegosystem’s **capacity**: the maximum payload size. However, these two aims are in direct competition. An unsatisfying (and useless) way of fulfilling the first aim is to always choose a payload of size zero. This way, the cover and stego objects are identical and therefore an eavesdropper cannot distinguish between them. However, this does not seem to *maximize* the payload size in any sense of the word, therefore failing to meet the second aim. On the other hand, we can fulfil the second aim by defining capacity as the number of bits B that used to store the cover object (say a JPEG file) on a computer. However, hiding a B -bit payload is very likely to make the stego object look like a garbled sequence of bits. Therefore, a detector will likely identify such objects as stego objects with near perfect accuracy, failing to meet the first aim.

As we can see, it is difficult to define the concept of capacity precisely.

Despite its lack of a formal definition, capacity is often measured in practice as an **embedding rate**: payload size relative to cover size (such as bits per cover pixel). This project does not attempt to define capacity, but rather shines a light on what its most appropriate unit of measurement might be in the context of JPEG steganography.

2.2 JPEG Compression

I will now briefly describe how JPEG compression works, based on the description in [20]. A complete, detailed account can be found in [30].

What makes JPEG compression space-efficient is **quantization**: for any $x \in \mathbb{R}$ and constant **quantization factor** q , we have $x \approx q[x/q]$ (where $[\cdot]$ is the nearest integer function). Instead of storing x , we can store $[x/q]$ (which is small for large enough q) and approximately recover x by multiplying $[x/q]$ by q .

Computers typically store colour images using the RGB format: an array of pixels, where each pixel is three bytes representing the red, green, and blue colour intensities, respectively. JPEG compression of an RGB image involves five steps.

Step 1: Colour space transformation Each pixel is separated into one luminance (brightness) component and two chrominance (colour) components via a linear transformation. The human visual system is more sensitive to changes in brightness than in colour [11], so JPEG reduces the resolution of the chrominance components for storage efficiency. The remaining steps are described only for luminance, but the chrominance components undergo a similar process.

Step 2: Division into blocks The $W \times H$ image is split into $\lceil W/8 \rceil \times \lceil H/8 \rceil$ blocks of 8×8 luminance values. Zero padding is used if either dimension is not a multiple of 8.

Step 3: Discrete cosine transform (DCT) This is the heart of the algorithm and is similar to the discrete Fourier transform. At the expense of being lossy, the DCT transforms the image from the domain of pixels into another domain that is easier to compress due to its sparsity. Specifically, the **spatial domain** (pixels) is transformed to a **frequency domain** (also called the **DCT domain** or **JPEG domain**). Each 8×8 block \mathbf{B} of pixels undergoes the DCT to produce an 8×8 block \mathbf{C} of frequencies (referred to as **coefficients**).

It is easiest to understand the transformation by considering its inverse. The

goal is to choose coefficients $C[u, v]$ for $0 \leq u, v \leq 7$ such that:

$$\mathbf{B} = \sum_{u=0}^7 \sum_{v=0}^7 C[u, v] \mathbf{A}_{u,v}$$

for **DCT modes** (u, v) . The transformation is effectively an orthonormal change of basis where the blocks $\mathbf{A}_{u,v}$ make up the basis and the coefficients \mathbf{C} define the linear combination of these basis blocks. Coefficient $C[0, 0]$ is called the **DC coefficient** and the remaining coefficients are called **AC coefficients**.

The $\mathbf{A}_{u,v}$ are defined as follows for $0 \leq i, j \leq 7$, where $c_0 = 1$ and $c_i = \sqrt{2}$ for all $i > 0$:

$$\mathbf{A}_{u,v}[i, j] = \frac{c_u c_v}{8} \cos\left(\frac{\pi}{8}\left(i + \frac{1}{2}\right)u\right) \cos\left(\frac{\pi}{8}\left(j + \frac{1}{2}\right)v\right)$$

Figure 2.1 shows a few examples of these blocks. The result of this step is that the high- and low-frequency parts of each luminance block \mathbf{B} are separated.




Figure 2.1: Examples of DCT basis blocks (taken from [20]).

Step 4: Quantization Most of the information loss is due to this step because quantization (dividing by a number then rounding to the nearest integer) is a many-to-one operation. Within each block \mathbf{C} , each coefficient is quantized separately, with larger coefficients getting larger quantization factors than smaller coefficients. A matrix of quantization factors \mathbf{Q}_{qf} can be determined according to the **quality factor** qf (an integer between 1 and 100, inclusive) chosen for the compression. Thus, instead of storing $C[i, j]$, we store the following, where $[\cdot]$ once again denotes the nearest integer function:

$$\left[\frac{C[i, j]}{\mathbf{Q}_{qf}[i, j]} \right]$$

Low quality factors result in larger quantization factors, which cause more coefficients to be quantized to zero. Therefore the compression is more efficient at the expense of more information being lost.

Step 5: Encoding This step is not relevant for this project and a full description can be found in [30]. It involves losslessly encoding the matrix of quantization factors as well as the quantized coefficients (stored in a zig-zag ordering in order to increase the chances of runs of zeros for better compressibility).

2.2.1 Decompression

Decompression is possible because of the quantization factors: the quantization step can be approximately reversed by multiplying each quantized coefficient by the corresponding element of the quantization factor matrix and performing the inverse DCT. Since information is lost during the quantization step, the decompressed image will almost certainly differ from the original image.

2.3 JPEG Steganography

Now that we understand how JPEG compression works, we can explore how to use JPEG images as a medium for steganography. In the remainder of this report, we will use m to denote payload size (in bits) and N to denote cover size (in pixels¹), unless otherwise noted.

One of the simplest embedding algorithms, called **least significant bit replacement (LSBR)**, visits the pixels in an image in a pseudorandom order determined by the secret key and uses the least significant bit (LSB) of each pixel to carry one bit of the payload². This works for raw, uncompressed images (though it is highly detectable [20]), but not for JPEGs due to the information loss during the quantization step. However, we can still use this idea, except *after* information is lost.

Instead of overwriting the LSB of each pixel, we can modify the LSB of each quantized luminance³ DCT coefficient by adding or subtracting 1. This is the central idea behind most JPEG embedding algorithms.

Modifying coefficients that are equal to zero tends to have a visually perceptible effect on the decompressed image, particularly in smooth areas (like

¹As we learned in Section 2.2, each 8×8 block of pixels produces an 8×8 block of DCT coefficients. Thus, N denotes both the number of pixels and the number of DCT coefficients in an image.

²This assumes we use a grayscale image where each pixel is represented as a single byte. In the case of an RGB image, we can use each of the three bytes that represent an RGB value to store one bit of payload.

³The luminance component is almost always the one used for embedding because changes to the chrominance components can be visually perceptible [20]. From now on, I will therefore no longer specify that coefficients are luminance coefficients in the context of embedding. Similarly, I will omit the term “quantized” and assume it to be implied because coefficients are only modified in their quantized form.

the sky). Thus, changing zero-valued coefficients is typically avoided and so the capacity of JPEG images is often measured in terms of bits per non-zero DCT coefficient. Changes to the (0,0) DCT mode – which is featured in Figure 2.1 – are also highly detectable because they cause blocks of 8×8 pixels to appear brighter or dimmer. Therefore, this is also generally avoided [20].

2.3.1 Progress in the Last Decade

One of the best-known JPEG embedding methods at the time [24] was written is called **F5**. Each bit of the payload is embedded in a non-zero coefficient by modifying the coefficient’s LSB – if this results in the coefficient becoming equal to zero, then the same bit is embedded again in the next non-zero coefficient [39]. Two consequences of this repeated embedding is that the capacity of images is relatively low and coefficients get pulled towards zero. The latter effect is referred to as “shrinkage” and it means that the number of zero-valued coefficients increases, which increases detectability. In **no shrinkage F5** (referred to as **nsF5**), the effect of pulling coefficients towards zero is eliminated using wet paper codes [12] (a topic beyond the scope of this report).

nsF5 was chosen to represent state-of-the-art embedding in [24], but **adaptive embedding** algorithms have since become the new standard [22]. Their goal is to identify embedding changes (such as modifications to the LSB of a coefficient) that are the least detectable [20]. This is expressed formally in Section 2.3.2.

2.3.2 Distortion Function Minimization

We will now formalize the task of embedding while minimizing a distortion function. This is partially based on the formulation in [9], which I have adapted to describe embedding in JPEGs. I will skip over some details in the interest of conciseness, but a full account can be found in [9].

Let:

- $\mathcal{I} = \mathbb{Z}$ be the set of all possible quantized DCT coefficient values
- $\mathcal{X} = \mathcal{I}^N$ be the set of all possible cover images
- $\mathbf{x} = (x_1, \dots, x_N) \in \mathcal{X}$ be a cover image, expressed as a sequence of N quantized DCT coefficients

Treating cover image \mathbf{x} as a constant, let:

- $\mathcal{I}_i \subseteq \mathcal{I}$ (for $1 \leq i \leq N$) be the set of possible values for the i^{th} quantized DCT coefficient of a stego image produced from \mathbf{x}

- $\mathcal{Y} = \mathcal{I}_1 \times \dots \times \mathcal{I}_N$ be the set of all stego images that can be produced from \mathbf{x}
- $\mathbf{y} = (y_1, \dots, y_N) \in \mathcal{Y}$ be a stego image produced from \mathbf{x}

Each \mathcal{I}_i satisfies $x_i \in \mathcal{I}_i$ because it is possible for a coefficient to remain unchanged. We say embedding is **binary** if $|\mathcal{I}_i| = 2$ for all i or **ternary** if $|\mathcal{I}_i| = 3$ for all i . Changing coefficients by ± 1 , for instance, yields $\mathcal{I}_i = \{x_i - 1, x_i, x_i + 1\}$ for all i . In this project, however, we are interested in binary embedding algorithms for simplicity⁴.

Let $\mathcal{D} : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$ be a distortion function that expresses the total cost of modifying the coefficients that differ in \mathbf{x} and \mathbf{y} . We are interested in additive distortion functions, which can be written as:

$$\mathcal{D}(\mathbf{x}, \mathbf{y}) = \sum_{i=1}^N \rho_i(\mathbf{x}, y_i)$$

where ρ_i is a function that gives a bounded cost for replacing cover pixel x_i with y_i ⁵. To simplify notation, let $c_i = \rho_i(\mathbf{x}, y_i)$. If $y_i \notin \mathcal{I}_i$, we say $\rho_i(\mathbf{x}, y_i) = \infty$. Notice that ρ_i depends on the entire cover image \mathbf{x} , meaning that it can make use of relationships between coefficients. However, it does *not* depend on any other ρ_j (where $i \neq j$), meaning that embedding changes are independent.

Here, we will diverge slightly from the formulation given in [9] and express formulas in terms of individual coefficients rather than the cover as a whole. This makes the presentation a bit simpler.

Our embedding algorithm computes probabilities π_1, \dots, π_N such that π_i is the probability that coefficient x_i is modified. In the context of binary embedding, this means that coefficient x_i is untouched with probability $1 - \pi_i$ and coefficient x_i is either incremented or decremented by 1 with probability π_i . The choice between incrementing or decrementing is made uniformly at random. This is in contrast with ternary embedding where there are three probabilities: one for making no change, one for adding 1, and one for subtracting 1.

Finally, we introduce the binary entropy function from information theory⁶:

$$h(x) = -x \log_2 x - (1 - x) \log_2 (1 - x)$$

⁴It is worth nothing that in general, ternary embedding is better because it can embed the same payload by making fewer changes and therefore achieve a lower detectability [8].

⁵Ideally ρ_i would be determined by the detectability of an embedding change (i.e. modification of a coefficient) but in practice it is set heuristically because there is no way to measure the detectability of a single change.

⁶ [22] explains why entropy is relevant.

Putting everything together, we can express the problem of embedding while minimizing a distortion function as two optimization problems, which are dual to each other:

1. **Payload-limited sender (PLS):** Embed a payload of a fixed length of m bits while minimizing the average distortion:

$$\begin{aligned} & \underset{\pi_1, \dots, \pi_N}{\text{Minimize}} && \sum_{i=1}^N c_i \pi_i \\ & \text{subject to} && \sum_{i=1}^N h(\pi_i) \geq m \end{aligned}$$

2. **Distortion-limited sender (DLS):** Given a fixed average distortion C_ϵ , embed a payload of maximum length:

$$\begin{aligned} & \underset{\pi_1, \dots, \pi_N}{\text{Maximize}} && \sum_{i=1}^N h(\pi_i) \\ & \text{subject to} && \sum_{i=1}^N c_i \pi_i \leq C_\epsilon \end{aligned}$$

From a practical perspective, PLS is a more natural formulation because someone wanting to send a message tends to have a particular message in mind. Since PLS and DLS are dual to each other, the optimal solution for one is also optimal for the other and this can be computed⁷ to be:

$$\pi_i = \frac{\exp(-\lambda c_i)}{\exp(-\lambda c_i) + \exp(-\lambda c_i^*)} = \frac{1}{\exp(-\lambda c_i) + 1} \quad (2.1)$$

where λ is fixed and c_i^* ($= 0$) is the cost of *not* changing coefficient x_i .

We can assume $c_i^* = 0$ because PLS and DLS are invariant to adding a constant to both c_i and c_i^* .

This optimal solution occurs when $\sum_{i=1}^N h(\pi_i) = m$ because $\sum_{i=1}^N h(\pi_i)$ is a monotonically increasing function. We can therefore pick a value of λ such that the equality holds for any value of m we want and compute the probabilities π_i with (2.1). λ can be chosen as follows: exponential search to get an upper (and lower) bound (i.e. increase λ starting from zero until $\sum_{i=1}^N h(\pi_i) < m$), then binary search until, say, $\sum_{i=1}^N h(\pi_i) \in [m, m + 1)$.

The key takeaway here is that we can *simulate* the embedding of an m -bit payload by changing each pixel with probability π_i . This is precisely what is done by J-UNIWARD, the embedding algorithm used in this project (described in Section 2.3.3).

⁷The full derivation can be found in [9]

2.3.3 J-UNIWARD

UNIWARD (“universal wavelet relative distortion”) [17] is a distortion function that forms the basis of several embedding methods. The term “universal” refers to the fact that UNIWARD can be used on any domain, such as the spatial domain or JPEG domain.

The UNIWARD distortion function is the sum of relative changes to wavelet coefficients with respect to the cover image. The wavelet coefficients come from three directional wavelet filter banks, which effectively measure the horizontal, vertical, and diagonal smoothness of an image. A more technical description of filter banks and wavelets can be found in [17]. As is standard practice, an additive approximation of the function is used for the embedding.

J-UNIWARD is one of the embedding algorithms based on the UNIWARD distortion function and it operates in the JPEG domain. It has been shown [17] to significantly outperform nsF5 and another former state-of-the-art algorithm called UED (uniform embedding distortion) and an implementation is freely available online for research purposes [36], thus making it a good choice for this project.

Figure 2.2 gives two examples of stego images produced by J-UNIWARD.

2.4 JPEG Steganalysis

Steganalysis is a supervised binary classification problem. Given a set of training data (consisting of a feature vector and a “cover” or “stego” label for each of a large number of images), a **detector** learns how to distinguish cover images from stego images⁸. To achieve the best results, features used in steganalysis should be sensitive to changes caused by embedding while being indifferent to the content of images.

2.4.1 Progress in the Last Decade

Advances in steganalysis over the last decade were mainly due to the development of richer feature sets.

A decade ago, features derived from DCT coefficients enabled classifiers to achieve the best performance in JPEG steganalysis. This is because features computed directly from the embedding domain (like RGB values for raw colour images and DCT coefficients for JPEG images) were considered to be the most sensitive to changes caused by embedding [13].

⁸In practice, cover and stego images are represented in the feature space. However, it is possible in theory to give the quantized DCT coefficients directly to a classifier instead of features.



(a)



(b)



(c)



(d)



(e)

Figure 2.2: Examples of stego images produced by J-UNIWARD embedding. 2.2(a) is the 320×240 pixel cover image. 2.2(b) and 2.2(c) are stego images with small and large payloads, respectively. 2.2(d) and 2.2(e) show (with enhanced contrast) the result of subtracting 2.2(b) from the cover image and subtracting 2.2(c) from the cover image, respectively. As we can see, the small payload results in very few embedding changes, which are imperceptible even at high contrast in 2.2(d). On the other hand, the large payload causes many embedding changes, which are concentrated in the parts of the image that are *not* smooth (such as the street and parking lot) and therefore less detectable.

However, feature spaces tended to be quite small: the features used in [24] had only 274 dimensions. These so-called “merged” features combined:

1. DCT features – such as histograms of DCT coefficient values, histograms of coefficients of specific DCT modes, and dependencies between coefficients within blocks
2. Markov features – these are derived from a Markov process that models differences between the absolute values of neighbouring DCT coefficients

These features were used to build several binary classifiers based on soft-margin support vector machines (SVMs) with Gaussian kernels [31] to identify stego images produced by different embedding algorithms. Since this technique was considered state-of-the-art in 2008, [24] used the classifier trained to distinguish between cover images and images produced by nsF5.

Today, a typical approach to steganalysis consists of defining an image model and a machine learning model trained to differentiate between cover and stego images (represented in the image model) [25]. Image models are based on the embedding domain, but tend to be much more complex than the feature sets a decade ago. In the context of JPEG steganalysis, they capture properties of individual DCT coefficients (such as frequencies of coefficient values) as well as relationships between coefficients (like how often two values occur beside each other) in the same block *and* in different blocks. This results in much richer – and therefore larger – feature sets that help classifiers achieve better results. The JPEG Rich Model (JRM) is one such model, which I use in this project and I describe in Section 2.4.2.

In 2012, a well-known machine learning tool was put to use in steganalysis as a simpler and more scalable alternative to SVMs. Ensemble classifiers, implemented as random forests, came to light as a new tool with promising results – they had better detection accuracy by combining the results of multiple individual classifiers [26]. However, a few years later, it was shown that a properly regularized (and trained) linear classifier could perform just as well as an ensemble classifier with a much lower computational complexity [5]. For this reason, I chose to use one such linear classifier for steganalysis in my experiments, as discussed in Section 2.4.3.

2.4.2 JRM

The Cartesian-calibrated JPEG Rich Model (CC-JRM) [25] is a state-of-the-art image model that produces a 22 510-dimensional feature space, irrespective of image size. The term “rich” refers to the model’s complexity: it consists of a large number of submodels that are combined to produce the

features. CC-JRM uses some interesting ideas, such as using the differences between the absolute values of coefficients.

2.4.3 Low-Complexity Linear Classifier

For my experiments, I used an implementation of a low-complexity linear classifier that is available online [27]. The classifier takes a set of pairs of cover image features and corresponding stego image features as input and splits them 50/50 into a training set and a testing set, keeping cover-stego pairs together. Then, it selects a regularization parameter with 5-fold cross-validation on the training set and uses it to regularize the classifier. Finally, the classifier runs on the testing set and outputs the probability of error.

2.5 Square Root Law

2.5.1 Batch Steganographic Capacity

The so-called “square root law” (SRL) was first conjectured in 2006 in [18] in the context of **batch steganography**, where a payload is spread over a collection of objects. A year later, it was formalized and proven [19]. Finally, the law was empirically validated in 2008 both in the spatial domain and in the JPEG domain [24].

Theorem 1 states the formalized SRL from [19], as paraphrased in [24]:

Theorem 1 (Batch Steganographic Capacity). *If a steganographer embeds a total payload of m bits into N uniform cover objects, then:*

1. *If $M/\sqrt{N} \rightarrow \infty$ as $N \rightarrow \infty$ then there is a pooled detector which, for sufficiently large N , comes arbitrarily close to perfect detection.*
2. *If $M/\sqrt{N} \rightarrow 0$ as $N \rightarrow 0$ then, by spreading the payload equally between N covers, the performance of any pooled detector must become arbitrarily close to random for sufficiently large N .*

where a **pooled detector** is one that performs steganalysis on each of a sequence of N objects *separately* and then uses the results to decide whether the entire batch of objects contains a payload. Note that m is not fixed – we can think of it as a function of N .

Although not directly relevant to this project, it is interesting to briefly mention two other versions of the SRL: it has been shown to hold if the cover source can be modelled as a stationary Markov chain [10] or as a Markov Random Field (including inhomogeneous Markov chains and Ising models) [21]. These results are significant because they prove that the law holds under more realistic assumptions about the cover source. Theorem 1,

on the other hand assumes that the covers are uniform, which is not realistic. This and other assumptions are discussed in Section 2.5.4.

As we can see, the SRL is in fact a *collection* of theorems. Throughout this paper, however, I only refer to the variant given in Theorem 1.

2.5.2 Relation to Capacity of Individual Cover Objects

How does Theorem 1 – a result about the steganographic capacity of *batches* of object – apply to the capacity of *individual* cover objects? We can answer this by giving another interpretation of a “batch” of objects.

Let us model a single cover object \mathcal{O} as a batch of objects. In particular, if we can split \mathcal{O} into a sequence of small regions, then we can treat these regions as a batch and apply the theorem directly. In the context of JPEG steganography, the division of a cover image into many small regions comes naturally using our knowledge of JPEG compression: simply treat each 8×8 block of coefficients as a separate cover object. However, this transformation has a catch: the theorem requires the cover objects to be uniform. This is discussed in Section 2.5.4.

2.5.3 Perfect Versus Imperfect Steganography

Before discussing caveats of the SRL, I will first clarify the difference between perfect and imperfect steganography and explain why this is relevant. A stegosystem is called **perfectly secure** if stego objects come from the same probability distribution as cover objects, meaning there is no statistical way of reliably distinguishing between the two [38].

It is well-known that the steganographic capacity of perfectly secure stegosystems is *linearly* proportional to the cover size [10]. Perfect security, however, is only possible if the Embedder has full information about their cover object source – an unlikely scenario in practice. It is therefore arguably more important to study the steganographic capacity of *imperfect* stegosystems, which is the case for the SRL as stated in Section 2.5.1.

2.5.4 Caveats

This SRL is a powerful result, but it can only be directly applied in fairly artificial circumstances. In particular, it makes the following assumptions:

- The proof of Theorem 1 relies on *both* the Embedder and the Detector knowing the number of cover objects in each batch.
 - In general, this is unrealistic. But, interestingly, this might be slightly more realistic for JPEGs because given an image, the Detector can compute the number of 8×8 coefficient blocks and

therefore the number of covers potentially used for batch embedding. This is assuming, of course, that the Detector somehow knows that separate JPEGs are not batched together and that the Embedder is not treating images as individual cover objects.

- The cover objects must be *uniform*.
 - In the case of JPEG steganography, this means adjacent blocks of coefficients must be independent. This is not completely implausible when using different cover objects, but is unlikely when dividing one image into many pieces.
 - However, we can consider this a non-issue if we know that JPEG steganalysis methods only measure properties of coefficients *within* each block and treat blocks as independent. A decade ago, this was indeed the case for most JPEG steganalysis, as mentioned in [24]. Detectors could be treated as pooling detectors that combine their independent analysis of each coefficient block and the batch capacity law is therefore obeyed.
 - Unfortunately, this uniformity assumption can no longer be explained away. For instance, the image model I am using to generate image features (JRM) considers both intra-block *and* inter-block coefficient relationships [25]. We must therefore accept that this assumption does not hold when using state-of-the-art techniques.
- The SRL holds, *all other things being equal*.
 - As with any experiment on real-world data, it is impossible to guarantee that all factors other than those being tested are constant.
 - For example, in these experiments we obtain images of different sizes by cropping evenly from the top/bottom and left/right edges of images, therefore keeping the pixels in the middle of the image. It is known [2] that factors such as local variance and cover saturation affect detection accuracy. Therefore, cropping towards the centre of an image might not be representative of the original image, so we should not compare the results for images of different sizes in our experiments.

There is another caveat related to empirically validating the law: we need access to both a perfect embedder⁹ and an optimal detector. These do not

⁹“Perfect” is used in the sense that the algorithm chooses how and where to make embedding changes so as to minimize statistical detectability, rather than making changes that are highly detectable (such as only modifying coefficients corresponding to pixels depicting a clear blue sky where the colour is smooth and changes are very visible).

exist in practice, so we instead have to use the best detection algorithms currently available. This is an approximation that, as we will see in Chapter 4, does not have a major effect on the manifestation of the SRL in practice.

2.6 Project Motivation, Goals, and Hypothesis

Given all the necessary background information on the SRL, state-of-the-art embedding and detection techniques, and what has changed in the last decade, we are now in a position to justify the need for this research, define concrete objectives, and give a hypothesis about the results.

2.6.1 Motivation

The last empirical study of the SRL took place over a decade ago [24]. In the meantime, both steganography and steganalysis techniques have greatly developed. There is therefore a need to provide fresh results regarding the existence of a SRL in practice that reflect the current state of the field.

2.6.2 Goals

As the project title suggest, the aim of this project is to empirically validate the SRL of steganographic capacity. Specifically, the aim is to mimic the experiments conducted in [24], but using state-of-the-art embedding and detection algorithms. [24] studied both spatial domain and JPEG domain steganography, but this project will only study the latter¹⁰.

We have the following concrete goals:

- Find two sets of images that:
 - Are free and legal to use
 - Contain a large number of images (“large” $\approx 10\,000$)
 - Consist of JPEG images all at the same JPEG quality factor (or raw images so that we can compress them ourselves)
- Choose one state-of-the-art algorithm for embedding and one for detection, ideally with implementations freely available online.
- Design and implement experiments to test detection accuracy on multiple image sizes and four payload rates: $O(1)$, $O(\sqrt{N})$, $O(\sqrt{N} \log N)$ ¹¹,

¹⁰Studying the spatial domain would amount to repeating all the experiments using different embedding and detection algorithms and thus would not be very interesting.

¹¹This payload rate was not studied in [24], but was mentioned in the paper’s conclusion as potentially being a more accurate measurement of steganographic capacity for modern embedding techniques. I therefore decided to include it.

and $O(N)$. The experiments should be easy to run and understand, robust to server failure, reasonably fast (ideally less than a day from start to finish), and configurable (with respect to image sizes and payload sizes).

- Run the experiments on both sets of images.
- Plot the results in a way that is easy to compare to the results in [24].

2.6.3 Hypothesis

As discussed in Section 2.5.4, we cannot empirically validate the SRL as stated in Theorem 1 because perfect embedders and detectors do not exist. We must instead study an approximation of the law using the state-of-the-art algorithms.

However, as steganography and steganalysis techniques develop, we can reasonably expect their performance to improve. Therefore, I expect the following for each of the payload rates:

- $O(1)$ – detector accuracy decreases as image size increases
- $O(\sqrt{N})$ and $O(\sqrt{N} \log N)$ – detector accuracy is (approximately) constant for all image sizes
- $O(N)$ – detector accuracy increases as image size increases

Chapter 3

Experiment Design

In order to study the square root law, we need a large number of images, tools for steganography and steganalysis, and a pipeline to process the images. This chapter describes the steps taken to select and preprocess appropriate image sets, build the required tools, assemble a pipeline, and run a large number of computation-intensive experiments. These steps are described one by one and Section 3.6 summarizes how they fit together.

The code used to run experiments is largely my own. It uses, or is based on, existing implementations of steganography and steganalysis algorithms. My supervisor provided me with the image sets and image metadata as well as a faster, optimised version of one of the algorithm implementations.

3.1 Images

3.1.1 Image Sets

First, the Yahoo Flickr Creative Commons 100 Million Dataset (YFCC100M) [37] was a natural choice as it contains nearly 100 million images from many users and is legal to use due to its Creative Commons license. My supervisor provided me with 13 349 images taken by one user with one camera. From now on, I will refer to this image set as “Actor3”.

Second, BOSSbase v1.01 [6] is a set of 10 000 raw grayscale images. It originally consisted of 9 074 images taken by seven different cameras and was put together in 2010 [1]. Since then, 926 images were added to bring the total to 10 000. It is the most popular image set for developing and testing steganographic algorithms, but its most common version is made up of 512×512 cropped versions of the images, which are far too small for this project. Thus, my supervisor provided me with PGM¹ versions of the

¹PGM (portable graymap format) is a file format for grayscale images.

images, which he produced from the raw, full-size images.

Actor3 was my primary image set, so I ran all experiments on it first and only then used BOSSbase as a secondary means of validation.

Why repeat the experiments on two different sets of images? It is *not* the case that because the SRL appears to hold for one image set, it will necessarily hold for another. This is because with any image set, there is a chance that all the images have some property that causes embedders and detectors to behave in a certain way. For instance, if a high proportion of user A’s pictures are zoomed in on grains of sand whereas user B mostly takes pictures of the clear blue sky, detectability will almost certainly be higher in all of user B’s images, regardless of the payload size². This is particularly relevant for Actor3 because all the images were taken by the same user with the same camera (and likely the same camera settings). Hence, it is particularly important to validate our Actor3 experiment findings on another image set. BOSSbase is a good choice because it contains images taken with seven different cameras, so if we observe similar behaviour again then it is likely be a general trend rather than a trend specific to Actor3.

3.1.2 Preprocessing

Actor3 images came in a variety of sizes, the largest and most common size being 3072×2304 pixels. I wrote a script to select all images with these dimensions (using image metadata supplied by my supervisor), convert them to grayscale, and rotate the portrait ones to landscape³. This produced a set of 9 539 images.

The BOSSbase images had many different shapes and sizes. I wrote a script that converts the PGM files to JPEG⁴, reads the image metadata (provided by my supervisor), selects all images that are large enough to be cropped to 2560×1920 , and rotates the portrait ones to landscape. An important difference was that I had to crop at least 32 pixels from each edge of each image. This is because the PGM images were produced directly from the raw images and therefore many of them have a thin black frame on the edges. Figure 3.1 shows an example. My script produced 10 000 images in the same format as the Actor3 images (except for the image dimensions).

²This is an extreme example, but serves to illustrate the point that image sets can have properties that alter the behaviour of embedders and therefore detectors. This is not limited to the subject of the images, which is in fact one of the least important properties of a set of images in practice. An example of a more important property is the ISO setting (which specifies sensitivity to light) of the camera when the images are taken.

³I performed these image transformations using a free command-line utility called `jpegtran`.

⁴This was done using a free command-line utility called `cjpeg`. I used quality factor 80 for the compression, to be consistent with the Actor3 images.



Figure 3.1: Grayscale PGM version of one of the raw BOSSbase images. A thin black frame can be seen on some of the image edges, particularly the top edge.

3.1.3 Cropping

In order to study the relationship between the payload size and the cover size, I had to produce images of several different sizes.

I did this by taking central crops of the images in each image set⁵. A **central crop** retains the pixels in the middle of an image and crops pixels evenly from the top/bottom and left/right edges. I decided to use central crops (rather than random crops, for instance) since the main subject of an image is usually in, or near, the center. This meant that cropping was likely to preserve the general properties of the images and made for a more meaningful comparison of the classifier’s performance on the large images relative to the smaller cropped images.

For each image set, I chose a smallest image size and some intermediate sizes such that the number of pixels in consecutive sizes grows linearly and the 4:3 ratio of the largest image sizes is preserved:

- For Actor3, I chose ten image sizes in total, with a smallest image size of 320×240 .
- For BOSSbase, I chose eight image sizes in total, with a smallest image size of 672×896 . Since it was my secondary image set, I wanted to have fewer image sizes in order to speed up the experiments.

The computed image sizes can be found in Appendix A.1.

I treated the smallest image size for each image set as a means of quickly testing my experiment pipeline and for this reason the corresponding data points in my graphs in Chapter 4 are not connected to the other points.

⁵Once again, I used `jpegtran` for this.

3.2 Embedding

A C++ implementation of J-UNIWARD available online [36] formed the basis of the implementation I used in my experiments, subject to a few important modifications.

In this implementation, the payload size must be specified in terms of bits per non-zero coefficient. However, we are interested in specifying the payload as a specific number of bits. This ensures that the same amount of information is embedded in each cover image, regardless of the coefficients. This has practical implications: users of an embedding algorithm will generally be interested in the detectability of a specific payload (such as “Meet me at 2pm”), without having to worry about the number of non-zero coefficients. So, I modified the implementation so the user can specify the payload size in bits.

Another important observation is that the implementation simulates ternary embedding whereas we are interested in binary embedding. Modifying J-UNIWARD to simulate binary instead of ternary embedding was a non-trivial change that involved updating the Payload-Limited Sender (PLS) implementation.

3.3 Costs

J-UNIWARD simulates embedding by computing a cost c_i for changing each coefficient x_i and then using these costs to compute the probability π_i of changing x_i , according to the PLS formulation described in Section 2.3.2.

In the original ternary implementation, there were costs for changing coefficient x_i by $+1$ or -1 or not changing it at all (call these c_i^{+1} , c_i^{-1} , and c_i^* , respectively). A characteristic of the resulting costs is $c_i^{+1} = c_i^{-1}$. However, since I use binary embedding, I only need the cost of making *any* change – whether it is $+1$ or -1 does not matter. This is simply c_i^{+1} (or equivalently c_i^{-1}).

By default, the J-UNIWARD implementation computes costs from scratch for each image. This is very inefficient since I need to simulate the embedding of many different payload lengths. A much better approach is to let J-UNIWARD compute and save the costs for each image once and read them whenever that image is used for embedding. My supervisor noticed this inefficiency prior to my project and had on hand a faster, optimised version of J-UNIWARD that saves the computed costs to a file and skips the embedding. He provided me with a copy of this code, which I used to compute the costs of all my images. I will refer to this version as J-

UNIWARD-COSTS. Consequently, I had to modify the J-UNIWARD code once more to make it skip the cost computation and instead read the costs from a specified file.

3.4 Features

I used a freely available [7] MATLAB implementation of JRM to generate image features. However, I had to write a script to convert the features (stored in one ASCII file per image) to a MATLAB matrix (stored as a .MAT file) with the fields expected by the classifier.

A consequence of JRM’s design is that the sum of all features of an image is a constant. Thus, I occasionally calculated the sum of the features as a sanity check.

3.5 Classifier

As mentioned in Section 2.4.3, I used a low-complexity linear classifier. A MATLAB implementation is available online [27] and I wrote a MATLAB script to read the cover and stego image features from two .MAT files and run the classifier.

3.6 Pipeline

Figure 3.2 shows how everything comes together.

The first step in the pipeline is the image preprocessing and cropping. This is done once per image set, producing eight (BOSSbase) or ten (Actor3) copies of the images – one for each size. Then, each set of images of each size flows through the rest of the pipeline, independently of the others, once for each payload size.

To understand the rest of the pipeline, consider a concrete example: the 320×240 Actor3 images and a 12 345-bit payload. The final step in the pipeline is feeding the features of all 320×240 cover images and the features of their corresponding stego images to the classifier. These can be computed in parallel. Computing the cover image features is simple: run JRM. For the stego images, we need to run J-UNIWARD-COSTS to compute the costs of changing the coefficients, then J-UNIWARD to simulate the embedding of a 12 345-bit payload in all cover images, and finally JRM to compute the features of the resulting stego images.

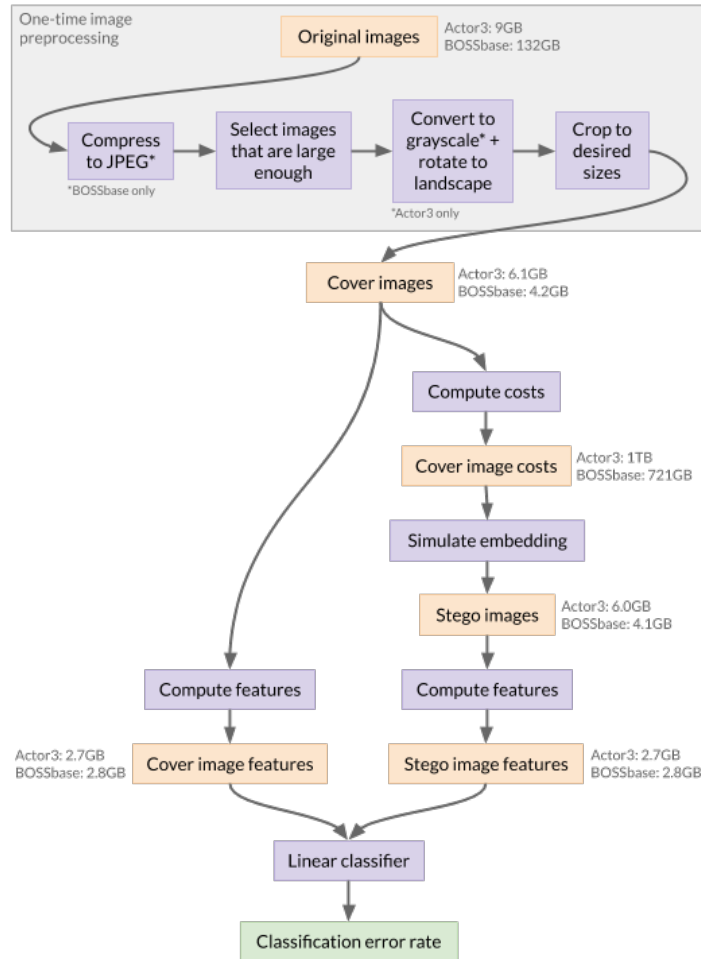


Figure 3.2: Experiment design flowchart. The purple boxes are actions and the orange boxes are inputs and outputs to/from the actions. The “Original images” box is annotated with the total size of all Actor3 and BOSSbase images, respectively, before any preprocessing. All other orange boxes are annotated with the total size of the relevant files for the largest image size and some payload size for each image set. The green “Classification error rate” box is the number output by the linear classifier.

3.7 Payload Sizes

We are interested in four payload rates: $O(1)$, $O(\sqrt{N})$, $O(\sqrt{N} \log N)$, and $O(N)$. In order to get meaningful results, we want the classifier’s probability of error for all experiments to be spread out between 0% and 50%. If the classifier has a probability of error of 0% (or 50%) for multiple experiments, then it is difficult to draw any conclusions. Thus, I somewhat arbitrarily chose to target a probability of error around 20%.

To determine the concrete payload sizes, I first found a combination of image size $W \times H$ and payload size m for each image set such that the classifier’s probability of error (when trained and tested on cover images of size $W \times H$ and stego images of size $W \times H$ with an m -bit payload) is around 20%. I found the following:

- Actor3: 36 223 bits of payload in 1056×792 images had a minimum probability of error of 0.2249
- BOSSbase: 41 287 bits of payload in 1792×1344 images had a minimum probability of error of 0.2028

For each image set, I chose a “middle” image size: 2048×1536 for Actor3 and 1792×1344 for BOSSbase. Then, I estimated the payload size needed to produce an error rate of around 20% in images of these two sizes:

- For Actor3, I had to compute the following, under the assumption that the SRL holds:

$$36\,223 \cdot \sqrt{\frac{2048 \cdot 1536}{1056 \cdot 792}} \approx 70\,252$$

- For BOSSbase, the middle image size is 1792×1344 so no further calculations were necessary: 41 287 is the desired number of bits.

Next, we need four proportionality constants r_1, r_2, r_3, r_4 for each image set such that we embed the following payloads for each image size:

- $O(1)$: $m = r_1 \cdot 1$
- $O(\sqrt{N})$: $m = r_2 \cdot \sqrt{N}$
- $O(\sqrt{N} \log N)$: $m = r_3 \cdot \log N \sqrt{N}$
- $O(N)$: $m = r_4 \cdot N$

Let N_{mid} be the number of bits in the middle image size of an image set and let m_{mid}^* be the number of bits that should be embedded in images of that size to produce a probability of error around 20%. For simplicity, I wanted proportionality constants such that the same number of bits is embedded in the middle image size for all four payload rates.

Hence, I solved the following equations for r_1, \dots, r_4 , given N_{mid} and m_{mid}^* :

$$\begin{aligned} m_{mid}^* &= r_1 \\ m_{mid}^* &= r_2 \cdot \sqrt{N_{mid}} \\ m_{mid}^* &= r_3 \cdot \sqrt{N_{mid}} \cdot \log N_{mid} \\ m_{mid}^* &= r_4 \cdot N_{mid} \end{aligned}$$

For Actor3, $N_{mid} = 2048 \cdot 1536$ and $m_{mid}^* = 70\,252$ and for BOSSbase, $N_{mid} = 1792 \cdot 2344$ and $m_{mid}^* = 41\,287$. The resulting constants are in Table 3.1.

	Actor3	BOSSbase
r_1	70252	41287
r_2	39.610	26.604
r_3	1.8351	1.2549
r_4	0.022333	0.017143

Table 3.1: Proportionality constants chosen for Actor3 and BOSSbase.

Finally, we can compute the payload sizes by multiplying the constants by 1, \sqrt{N} , $\sqrt{N} \log N$, and N for each image size N . The results of these calculations are in Appendix A.3.

In order to more closely mirror the experiments from [24] (as well as provide a more confident validation of the SRL), I computed two additional sets of proportionality constants. They were simply 30% smaller and larger, respectively, than the constants above. I will refer to these as the “small”, “middle”, and “large” constants, respectively. These constants are in A.2 and their corresponding payload sizes are in Appendix A.3.

3.8 Experiments

Everything is now in place to run the experiments. This involves running through the one-time preprocessing steps highlighted in Figure 3.2, once for Actor3 and once for BOSSbase, and then executing the rest of the pipeline once for each image size and payload size. Finally, the results need to be plotted, as described in Chapter 4.

Chapter 4

Results & Analysis

This chapter presents the experiment results and discusses how they align with our expectations. Overall, the experiments required 10 months of computing time and generated 10 TB of data in the form of images, features, and costs.

4.1 Experiment Result Visualizations

Before presenting any results, I will explain the visualization methods I chose.

4.1.1 ROC Curves

First, I will show **receiver operating characteristic (ROC) curves**. ROC curves plot the true positive rate ($1 - f_n$), on the y-axis, against the false positive rate (f_p), on the x-axis, and are a standard tool for evaluating the diagnostic ability of a binary classifier.

The points making up an ROC curve can be used to compute the **minimum probability of error (P_E)**: $P_E = \min \frac{1}{2}(f_n + f_p)$ over all points $(f_p, 1 - f_n)$. Figure 4.1 shows an example of an ROC curve with the minimum P_E (referred to from now on as **MinPE**) point highlighted.

A straight line from the bottom-left corner to the top-right corner corresponds to the performance of a random classifier. On the other hand, the closer a curve is to the top-left corner, the closer the classifier is to perfection. Thus, if we identify the MinPE point on an ROC curve, we expect it to be one of the closest points to the top-left corner of the graph. This makes sense because, in our case, the top-left corner corresponds to $f_p = 0$ (no cover images are identified as stego images) and $1 - f_n = 0 \implies f_n = 1$ (all stego images are identified as stego images).




Figure 4.1: Sample ROC curve. This is the curve generated when the classifier is run on all 2048×1536 Actor3 cover images and corresponding stego images with 70252 bits of payload. The MinPE point is highlighted by a black dot.

4.1.2 Detectability Versus Cover Size

Second, I will plot detectability against cover size for each payload rate (I will call these **detectability graphs**).

How should detectability be measured? [24] considered three metrics (one being $1 - \text{MinPE}$) and they all exhibited similar results. Here we will only consider $1 - \text{MinPE}$, which can be informally interpreted as the classifier’s accuracy.

These graphs have a similar format to those in Section 4.2 for ease of comparison. However, there are a few differences worth noting:

- I express cover size as the number of pixels (or equivalently coefficients), whereas [24] expresses cover size as the number of *non-zero* (AC) DCT coefficients. Both units are correct in their respective contexts because they represent the number of places where embedding changes can be made: the nsF5 embedding algorithm used in [24] ignores zero-valued coefficients, while J-UNIWARD uses all of them.
- I run experiments with four payload rates whereas [24] only studies three payload rates. As I mentioned in Section 2.6.2, this is because the $O(\sqrt{N} \log N)$ payload rate is mentioned in [24] as a potentially better measure of steganographic capacity for modern embedding algorithms that use adaptive embedding.

4.2 Results From a Decade Ago

We first remind ourselves of what we hope our results will look like. Figure 4.2 shows the 2008 results for the detection accuracy of a support vector machine-based classifier for nsF5 embedding in JPEG images.




Figure 4.2: JPEG domain results from [24]. The graphs show the results for $O(1)$ payloads (left), $O(\sqrt{N})$ payloads (middle), and $O(N)$ payloads (right), where N is the number of non-zero coefficients. Each graph has three lines, corresponding to the use of three proportionality constants.


4.3 Actor3

4.3.1 ROC Curves


Figure 4.3 presents ROC curves for each payload rate. As explained in Section 4.1.1, we can assess how well the classifier performs by looking at how close a curve is to the top-left corner of the graph.

$O(1)$ payload The classifier achieves perfect detection with the smallest image size since its MinPE has coordinates $(0, 1)$, as we see in Figure 4.3(a). As image size increases, detector accuracy decreases and approaches that of a random classifier. This is expected since the number of potential embedding locations (coefficients) to choose from increases with cover size, so the embedding algorithm will have more low-cost (as measured by the distortion function) choices. Thus, it can achieve a lower overall distortion in large images than in small images.


$O(\sqrt{N})$ and $O(\sqrt{N} \log N)$ payloads The curves and MinPE points are very close together in both Figures 4.3(b) and 4.3(c). The smallest image size is a notable exception in both and we will treat it as an outlier. The rest of the curves, however, show that detectability is approximately constant for both payload rates. The curves in Figure 4.3(c) are visibly closer together than those in Figure 4.3(b), suggesting that steganographic capacity is more accurately measured as being of order $\sqrt{N} \cdot \log N$ than of order \sqrt{N} .




(a) $O(1)$ payload



(b) $O(\sqrt{N})$ payload




(c) $O(\sqrt{N} \log N)$ payload




(d) $O(N)$ payload


Figure 4.3: Actor3 ROC curves. These results are for the payload sizes generated by the middle proportionality constants. The legend indicates the width of each image size and the MinPE point of each curve is highlighted with a black dot. The gray dotted line represents the performance of a random classifier and is included as a reference point.




(a) $O(1)$ payload



(b) $O(\sqrt{N})$ payload



(c) $O(\sqrt{N} \log N)$ payload



(d) $O(N)$ payload

Figure 4.4: BOSSbase ROC curves. These results are for the payload sizes generated by the middle proportionality constants. The legend indicates the width of each image size and the MinPE point of each curve is highlighted with a black dot. The gray dotted line represents the performance of a random classifier and is included as a reference point.

$O(N)$ payload Figure 4.3(d) shows that as image size increases, detector accuracy *increases* – the opposite of the trend in Figure 4.3(a). This behaviour is expected: if we think of each coefficient modification as sending a nearly imperceptible signal to the classifier, then the aggregation of increasingly many signals (as image size, and therefore payload size, grows) becomes increasingly perceptible. Another interesting observation is that the curves in Figure 4.3(d) are closer to the line depicting a random classifier’s accuracy than the curves in Figure 4.3(a). This is not a result of the payload rate, but rather it is a result of the proportionality constant choice – if the constant was larger, the curves would all be closer to the top-left corner. See Figure 4.5 for a concrete example.

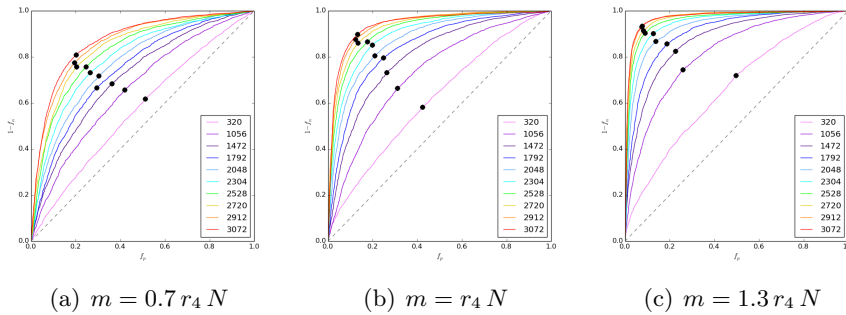


Figure 4.5: Actor3 ROC curves for $O(N)$ payloads for all proportionality constants. The legend indicates the width of each image size and the MinPE point of each curve is highlighted with a black dot. The gray dotted line represents the performance of a random classifier and is included as a reference point.

4.3.2 Detectability Versus Cover Size

In Figure 4.6 we see $1 - \text{MinPE}$ plotted against cover size, where the MinPE points are precisely those highlighted in Figure 4.5. It is more straightforward to assess changes in detectability with these graphs than with ROC curves because we can compare the detectability metric (MinPE) directly on the y-axis. The trends discussed in Section 4.3.1 are also visible here.

$O(1)$ payload Figure 4.6(a) shows that detectability decreases as image size increases.

$O(\sqrt{N})$ and $O(\sqrt{N} \log N)$ payloads Figures 4.6(b) and 4.6(c) show that detectability is approximately constant for $O(\sqrt{N})$ and $O(\sqrt{N} \log N)$ payloads, respectively. The fact that the ROC curves were closer together for $O(\sqrt{N} \log N)$ payloads than for $O(\sqrt{N})$ payloads manifests itself here

as well: the lines in Figure 4.6(c) are flatter than those in Figure 4.6(b) (ignoring the smallest image size, which we also disregarded when assessing the ROC curves). This may not be obvious to the naked eye, so I computed the slope of the best fit line¹ for these two payload rates, using the middle constant and ignoring the smallest image size. The slope for Figure 4.6(c) ($-6.425e5$) is closer to zero than the slope for Figure 4.6(b) ($-1.285e5$), indicating that the former is flatter.

$O(N)$ payload Figure 4.6(d) shows that detectability increases as image size increases.

4.4 BOSSbase

4.4.1 ROC Curves

Figure 4.4 shows the ROC curves for BOSSbase. There is an evident similarity between these curves and the curves for Actor3 in Figure 4.3 for all payload rates. Therefore, all the analysis in Section 4.3.1 applies here as well, though with a few slight differences:

- The smallest BOSSbase images do not stand out as much as the smallest Actor3 images did. This difference is not surprising: the smallest BOSSbase images have around 10 times as many pixels as the Actor3 images but the number of payload bits as a proportion of the number of pixels is 14% for Actor3’s smallest images but only 3% for BOSSbase’s smallest images (for a $O(\sqrt{N})$ payload, as an example). Thus, a much larger proportion of the pixels (coefficients, to be precise) need to be used for embedding in Actor3’s smallest images, resulting in higher detectability. This could have been avoided by choosing the same image sizes for both image sets.
- The ROC curves for the $O(1)$ payloads and $O(N)$ payloads are much more spread out for Actor3 in Figures 4.3(a) and 4.3(d) than for BOSSbase in Figures 4.4(a) and 4.4(d). This can again be attributed to the fact that different image sizes were chosen for the two image sets. The difference in the number of pixels between consecutive image sizes is larger for Actor3 (800 000 pixels) than for BOSSbase (700 000 pixels), so consecutive image sizes can be expected to behave more similarly for BOSSbase than for Actor3.

¹I used simple linear regression with ordinary least squares to fit a line to the points in the detectability graph.

4.4.2 Detectability Versus Cover Size

Given the similarity of the BOSSbase and Actor3 ROC curves, the detectability graphs for BOSSbase are unsurprisingly similar to those for Actor3, so the Actor3 analysis in Section 4.3.2 also applies here. It appears that the SRL is more closely obeyed for $O(\sqrt{N} \log N)$ payloads than for $O(\sqrt{N})$ payloads. As with Actor3, I computed the slope of the best fit line for both payload rates and arrived at the same result: the slope for Figure 4.7(c) ($-7.116e5$) is closer to zero than the slope for Figure 4.7(b) ($-1.0932e4$).

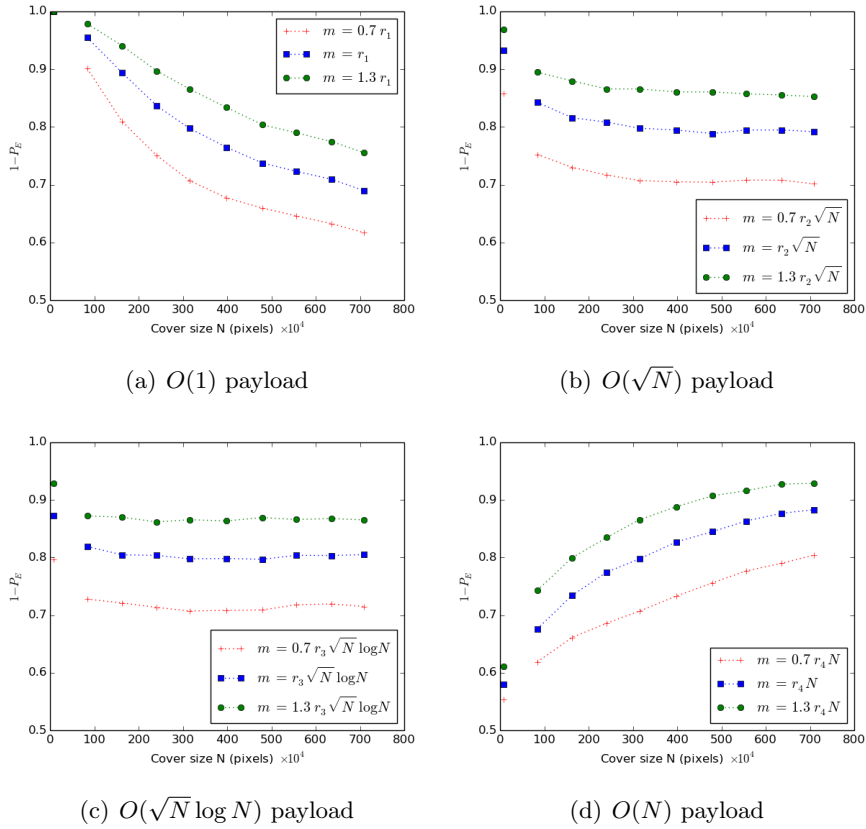


Figure 4.6: Actor3 detectability graphs. These graphs plot $1 - \text{MinPE}$ (on the y-axis) against the cover size (on the x-axis) for all ten image sizes. The legend indicates which line corresponds to which proportionality constant. The points corresponding to the smallest image size are not connected to the rest since we treated them as a way to test the pipeline rather than to get meaningful results, as described in Section 3.1.3.

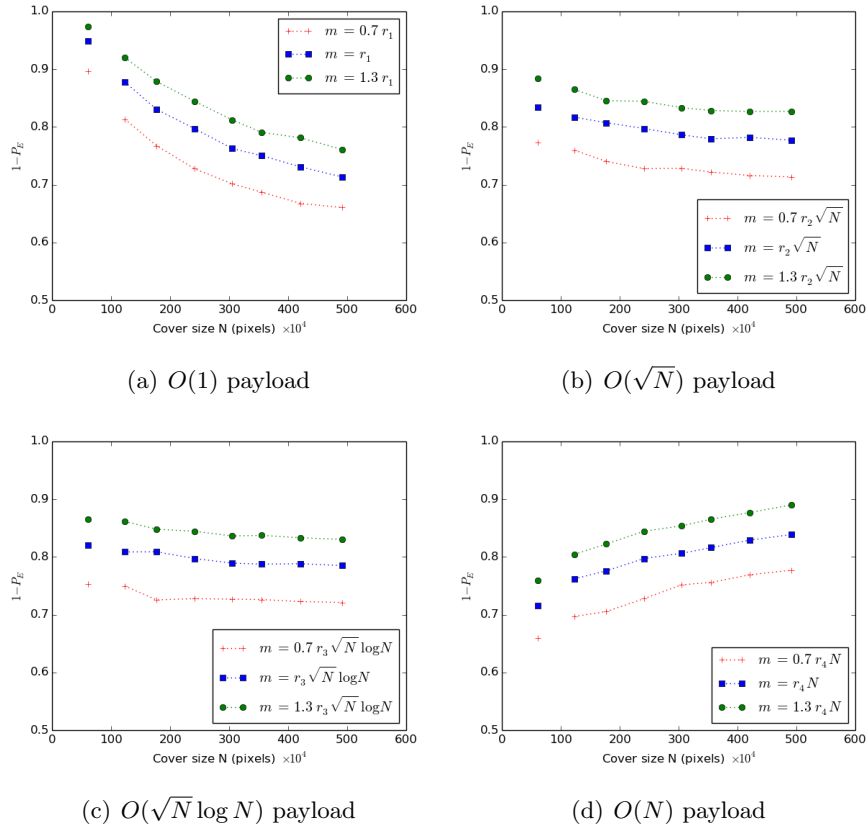


Figure 4.7: BOSSbase detectability graphs. These graphs plot $1 - \text{MinPE}$ (on the y-axis) against the cover size (on the x-axis) for all eight image sizes. The legend indicates which line corresponds to which proportionality constant. The points corresponding to the smallest image size are not connected to the rest since we treated them as a way to test the pipeline rather than to get meaningful results, as described in Section 3.1.3.

Chapter 5

Conclusion

The square root law was first shown to hold in practice in 2008 [24] in both the spatial domain (raw images) and the JPEG domain using state-of-the-art (at the time) steganography and steganalysis techniques. We have shown that the law continues to be observed empirically in the JPEG domain – in fact, we observed that steganographic payload capacity is of order $\sqrt{N} \log N$ rather than \sqrt{N} , which reflects the use of state-of-the-art adaptive embedding algorithms.

Our experimental results on two large image sets align with the results from [24] and confirm our hypotheses from Section 2.6.3:

- For $O(1)$ payloads, detectability decreases as image size increases.
- For $O(\sqrt{N})$ or $O(\sqrt{N} \log N)$ payloads, detectability is approximately constant as image size increases.
- For $O(N)$ payloads, detectability increases as image size increases.

This empirical validation of the square root law has practical implications. Firstly, it indicates to steganographers that secret channels dry out: the more information they want to send, the slower they need to send it in order to avoid detection. It also puts into question the units of measurement used for steganographic capacity: perhaps “bits per square root pixel” is more appropriate than “bits per pixel”.

5.1 Further Work

The embedding and detection tools that we tested are currently the state-of-the-art, but another empirical study will be needed once a new suite of tools are developed, tested, and put to use. In particular, there has been research into the use of deep learning for steganalysis.

For instance, [4] describes a convolutional neural network (CNN) for JPEG steganalysis that is tested on two embedding algorithms (one of which is J-UNIWARD) and has promising results. However, it is important to note that a CNN is not a novel approach and works very much like the current state-of-the-art: it gives a (learned) set of features to a linear classifier (in the output layer). There is therefore no reason to believe that the square root law will manifest itself differently. Arguably, CNNs are not even a good steganalysis tool:

- A CNN does not work for multiple image sizes because the input vector has a fixed size. Fixes like zero-padding affect the network's behaviour.
- Only small image sizes can be supported due to memory constraints. For example, a 512×512 image has to be fed into a network whose input vector has over a quarter of a million dimensions. This makes steganalysis infeasible on large, commonly used image sizes.

Nevertheless, it will be interesting to see if there are breakthroughs in steganography and steganalysis that affect the manifestation of the square root law in practice.

5.2 Critical Evaluation

From a research point of view, this was an interesting project that taught me a lot about the theory of steganography and steganalysis as well as the practice of experiment design and result analysis. There are a few additional topics I could have addressed given more time:

- The classifier was run only once for each combination of image set, image size, and payload size. This means that each MinPE value reported in Chapter 4 is for one training/testing split of the images. If this project's results are published, the classifier should be run multiple times (as a reference point, [24] did this 100 times).
- To more confidently validate the square root law in practice, I could have tested: more than two image sets, more embedding methods (such as UED [14] and MiPOD [33]), more detection methods (such as those described in Section 5.1), and more feature sets (such as DCTR [16]). However, it is unlikely this would have brought any new insight, so I considered it out of scope.
- I did not include error bars in my detectability graphs, unlike in [24]. This was beyond the scope of my project and would have required running the classifier multiple times for each point and performing statistical calculations.
- The classifier split my cover and stego images 50/50 into training and

testing sets and I did not consider the effect of the number of images used for training on detectability. It would have been interesting to study this and plot training set size against accuracy to see how many training images the classifier needs until its detectability becomes (approximately) constant for a given combination of image set, image size, and payload size.

- I would have also liked to study how the classifier performs when trained on images of many different sizes. This would be relevant for steganalysis in practice because it is infeasible to have a trained classifier for every possible image size. Having a classifier that works well for multiple image sizes would be practical.

From a software engineering point of view, this was a fun and challenging project that involved:

- Studying code in several languages from a few different sources and understanding it to a point that I could modify it to meet my needs
- Learning about Bash and Python scripting in order to automate my experiment pipeline
- Learning how to parallelize computations in Python for improved experiment efficiency
- Learning (the hard way) how to write code that can deal with server failure, running out of memory, etc.

Acknowledgements

I would like to thank my supervisor Dr. Andrew Ker for his support not only with this project, but also with my studies over the last four years. I have had an incredible time at Oxford and this is largely due to his wonderful tutorials and project meetings, extra-curricular advice and encouragement, and general thought-provoking discussions. The last year has been both fun and challenging and I am grateful to Dr. Ker for his time and energy.

Bibliography

- [1] P. Bas, T. Filler, and T. Pevný. “Break Our Steganographic System”: The Ins and Outs of Organizing BOSS. pages 59–70, May 2011.
- [2] R. Böhme and A. D. Ker. A Two-Factor Error Model for Quantitative Steganalysis. 6072, February 2006.
- [3] D. Bradbury. GE Engineer Charged for Novel Data Theft, April 2019. <https://www.infosecurity-magazine.com/infosec/ge-engineer-charged-data-theft-1/>, accessed May 2019.
- [4] M. Chen, V. Sedighi, M. Boroumand, and J. Fridrich. JPEG-Phase-Aware Convolutional Neural Network for Steganalysis of JPEG Images. In *Proceedings of the 5th ACM Workshop on Information Hiding and Multimedia Security*, pages 75–84. ACM, June 2017.
- [5] R. Cograñne, V. Sedighi, J. Fridrich, and T. Pevný. Is Ensemble Classifier Needed for Steganalysis in High-Dimensional Feature Spaces? In *2015 IEEE International Workshop on Information Forensics and Security*, pages 1–6, 2015.
- [6] Downloads. <http://dde.binghamton.edu/download/>, accessed May 2019.
- [7] Feature Extractors for Steganalysis. http://dde.binghamton.edu/download/feature_extractors/, accessed May 2019.
- [8] T. Filler and J. Fridrich. Minimizing Additive Distortion Functions with Non-Binary Embedding Operation in Steganography. In *2010 IEEE International Workshop on Information Forensics and Security*, pages 1–6. IEEE, 2010.
- [9] T. Filler, J. Judas, and J. Fridrich. Minimizing Additive Distortion in Steganography Using Syndrome-Trellis Codes. *IEEE Transactions on Information Forensics and Security*, 6(3):920–935, 2011.
- [10] T. Filler, A. D. Ker, and J. Fridrich. The Square Root Law of Steganographic Capacity for Markov Covers. In *Media Forensics and Security*,

volume 7254, page 725408. International Society for Optics and Photonics, February 2009.

- [11] J. Fridrich. *Steganography in Digital Media : Principles, Algorithms, and Applications*. Cambridge University Press, 2009.
- [12] J. Fridrich, T. Pevný, and J. Kodovský. Statistically Undetectable JPEG Steganography: Dead Ends, Challenges, and Opportunities. In *Proceedings of the 9th Workshop on Multimedia & Security*, pages 3–14. ACM, 2007.
- [13] M. Goljan, J. Fridrich, and T. Holotyak. New Blind Steganalysis and Its Implications. *Proceedings of the SPIE*, 6072, February 2006.
- [14] L. Guo, J. Ni, and Y. Q. Shi. Uniform Embedding for Efficient JPEG Steganography. *IEEE Transactions on Information Forensics and Security*, 9:814–825, May 2014.
- [15] Herodotus. *The Histories*. Penguin Classics, 2013. Translated by Tom Holland.
- [16] V. Holub and J. Fridrich. Low-Complexity Features for JPEG Steganalysis Using Undecimated DCT. *IEEE Transactions on Information Forensics and Security*, 10:219–228, February 2015.
- [17] V. Holub, J. Fridrich, and T. Denemark. Universal Distortion Function for Steganography in an Arbitrary Domain. *EURASIP Journal on Information Security*, 2014(1):1, 2014.
- [18] A. D. Ker. Batch Steganography and Pooled Steganalysis. In *International Workshop on Information Hiding*, volume 4437, pages 265–281. Springer, July 2006.
- [19] A. D. Ker. A Capacity Result for Batch Steganography. *IEEE Signal Processing Letters*, 14:525–528, September 2007.
- [20] A. D. Ker. Information Hiding Notes (Advanced Security Course), 2016. <https://www.cs.ox.ac.uk/teaching/materials15-16/advsec/>.
- [21] A. D. Ker. The Square Root Law of Steganography: Bringing Theory Closer to Practice. In *Proceedings of the 5th ACM Workshop on Information Hiding and Multimedia Security*, pages 33–44. ACM, June 2017.
- [22] A. D. Ker. On the Relationship Between Embedding Costs and Steganographic Capacity. In *Proceedings of the 6th ACM Workshop on Information Hiding and Multimedia Security*, pages 115–120. ACM, June 2018.

- [23] A. D. Ker, P. Bas, R. Böhme, R. Cogramme, S. Craver, T. Filler, J. Fridrich, and T. Pevný. Moving Steganography and Steganalysis from the Laboratory into the Real World. In *Proceedings of the First ACM Workshop on Information Hiding and Multimedia Security*, pages 45–58. ACM, June 2013.
- [24] A. D. Ker, T. Pevný, J. Kodovský, and J. Fridrich. The Square Root Law of Steganographic Capacity. In *Proceedings of the 10th ACM Workshop on Multimedia and Security*, pages 107–116. ACM, 2008.
- [25] J. Kodovský and J. Fridrich. Steganalysis of JPEG Images Using Rich Models. In *Media Watermarking, Security, and Forensics 2012*, volume 8303, page 83030A. International Society for Optics and Photonics, 2012.
- [26] J. Kodovsky, J. Fridrich, and V. Holub. Ensemble Classifiers for Steganalysis of Digital Media. *IEEE Transactions on Information Forensics and Security*, 7:432–444, April 2012.
- [27] Low-complexity Linear Classifier. <http://dde.binghamton.edu/download/LCLSMR/>, accessed May 2019.
- [28] L. Mathews. Malware Hidden In Banner Ads Served Up To Millions, December 2016. <https://www.forbes.com/sites/leemathews/2016/12/08/malware-hidden-in-banner-ads-served-up-to-millions/>, accessed May 2019.
- [29] E. Nakashima. U.S. Charges American Engineer, Chinese Businessman with Stealing GE’s Trade Secrets, April 2019. https://www.washingtonpost.com/world/national-security/us-charges-american-engineer-chinese-businessman-with-stealing-ges-trade-secrets/2019/04/23/cb32c78a-65f5-11e9-82ba-fcfeff232e8f_story.html, accessed May 2019.
- [30] W. B. Pennebaker and J. L. Mitchell. *JPEG: Still Image Data Compression Standard*. Chapman & Hall Digital Multimedia Standards Series. Springer US, 1992.
- [31] T. Pevný and J. Fridrich. Merging Markov and DCT features for Multi-Class JPEG Steganalysis. 6505, February 2007.
- [32] D. Sancho. Steganography and Malware: Why and How, May 2015. <https://blog.trendmicro.com/trendlabs-security-intelligence/steganography-and-malware-why-and-how/>, accessed May 2019.
- [33] V. Sedighi, R. Cogramme, and J. Fridrich. Content-Adaptive Steganography by Minimizing Statistical Detectability. *IEEE Transactions on Information Forensics and Security*, 11:1–1, January 2015.

- [34] C. E. Shannon. Communication Theory of Secrecy Systems. *Bell System Technical Journal*, 28(4):656–715, 1949.
- [35] G. J. Simmons. The Prisoners’ Problem and the Subliminal Channel. In D. Chaum, editor, *Advances in Cryptology, CRYPTO ’83*, pages 51–67. Plenum Press, August 1983.
- [36] Steganographic Algorithms. http://dde.binghamton.edu/download/stego_algorithms/, accessed May 2019.
- [37] B. Thomee, B. Elizalde, D. Shamma, K. Ni, G. Friedland, D. Poland, D. Borth, and L. J. Li. YFCC100M: The New Data in Multimedia Research. *Communications of the ACM*, 59:64–73, January 2016.
- [38] Y. Wang and P. Moulin. Perfectly Secure Steganography: Capacity, Error Exponents, and Code Constructions. *IEEE Transactions on Information Theory*, 54(6):2706–2722, 2008.
- [39] A. Westfeld. F5 - A Steganographic Algorithm. *Lecture Notes in Computer Science*, 2137:289–302, October 2001.
- [40] A. Zahravi. Cybercriminals Use Malicious Memes that Communicate with Malware, December 2018. <https://blog.trendmicro.com/trendlabs-security-intelligence/cybercriminals-use-malicious-memes-that-communicate-with-malware/>, accessed May 2019.

Appendix A

A.1 Image Sizes

This section gives the image sizes for Actor3 and BOSSbase, as described in Section 3.1.3.

Width	Height	Total pixels
320	240	76 800
1056	792	836 352
1472	1104	1 625 088
1792	1344	2 408 448
2048	1536	3 145 728
2304	1728	3 981 312
2528	1896	4 793 088
2720	2040	5 548 800
2912	2184	6 359 808
3072	2304	7 077 888

Table A.1: Actor3 image sizes.

Width	Height	Total pixels
896	672	602 112
1280	960	1 228 800
1536	1152	1 769 472
1792	1344	2 408 448
2016	1512	3 048 192
2172	1632	3 544 704
2368	1776	4 205 568
2560	1920	4 915 200

Table A.2: BOSSbase image sizes.

A.2 Proportionality Constants

This section gives the three sets of proportionality constants described in Section 3.7 for both image sets.

	Actor3	BOSSbase
0.7 r_1	49176	28900
r_1	70252	41287
1.3 r_1	91327	53673
0.7 r_2	27.727	18.623
r_2	39.610	26.604
1.3 r_2	51.492	34.585
0.7 r_3	1.2845	0.87844
r_3	1.8351	1.2549
1.3 r_3	2.3856	1.6314
0.7 r_4	0.015633	0.012000
r_4	0.022333	0.017143
1.3 r_4	0.029032	0.022285

Table A.3: Actor3 and BOSSbase proportionality constants.

A.3 Payload Sizes

This section gives the payload sizes (in bits) for Actor3 and BOSSbase for all three sets of their respective proportionality constants, as described in Section 3.7.

A.3.1 Actor3

Width	Height	Pixels	$O(1)$	$O(\sqrt{N})$	$O(\sqrt{N} \log N)$	$O(N)$
320	240	76 800	49 176	7 683	5 777	1 200
1056	792	836 352	49 176	25 356	23 111	13 074
1472	1104	1 625 088	49 176	35 345	33 785	25 404
1792	1344	2 408 448	49 176	43 029	42 261	37 650
2048	1536	3 145 728	49 176	49 176	49 176	49 176
2304	1728	3 981 312	49 176	55 323	56 194	62 239
2528	1896	4 793 088	49 176	60 702	62 410	74 929
2720	2040	5 548 800	49 176	65 312	67 790	86 743
2912	2184	6 359 808	49 176	69 922	73 212	99 421
3072	2304	7 077 888	49 176	73 764	77 762	110 647

Table A.4: Actor3 payload sizes for proportionality constants $0.7r_1, 0.7r_2, 0.7r_3, 0.7r_4$ (given in Table A.3), with the total number of pixels as a reference.

Width	Height	Pixels	$O(1)$	$O(\sqrt{N})$	$O(\sqrt{N} \log N)$	$O(N)$
320	240	76 800	70 252	10 976	8 253	1 715
1056	792	836 352	70 252	36 223	33 016	18 677
1472	1104	1 625 088	70 252	50 493	48 264	36 292
1792	1344	2 408 448	70 252	61 470	60 373	53 786
2048	1536	3 145 728	70 252	70 252	70 252	70 252
2304	1728	3 981 312	70 252	79 033	80 278	88 912
2528	1896	4 793 088	70 252	86 717	89 158	107 042
2720	2040	5 548 800	70 252	93 303	96 843	123 919
2912	2184	6 359 808	70 252	99 889	104 589	142 030
3072	2304	7 077 888	70 252	105 378	111 089	158 067

Table A.5: Actor3 payload sizes for proportionality constants r_1, r_2, r_3, r_4 (given in Table A.3), with the total number of pixels as a reference.

Width	Height	Pixels	$O(1)$	$O(\sqrt{N})$	$O(\sqrt{N} \log N)$	$O(N)$
320	240	76 800	91 327	14 269	10 728	2 229
1056	792	836 352	91 327	47 090	42 921	24 281
1472	1104	1 625 088	91 327	65 641	62 744	47 180
1792	1344	2 408 448	91 327	79 911	78 485	69 922
2048	1536	3 145 728	91 327	91 327	91 327	91 327
2304	1728	3 981 312	91 327	102 743	104 361	115 586
2528	1896	4 793 088	91 327	112 732	115 905	139 154
2720	2040	5 548 800	91 327	121 294	125 895	161 094
2912	2184	6 359 808	91 327	129 856	135 966	184 640
3072	2304	7 077 888	91 327	136 991	144 416	205 487

Table A.6: Actor3 payload sizes for proportionality constants $1.3r_1, 1.3r_2, 1.3r_3, 1.3r_4$ (given in Table A.3), with the total number of pixels as a reference.

A.3.2 BOSSbase

Width	Height	Pixels	$O(1)$	$O(\sqrt{N})$	$O(\sqrt{N} \log N)$	$O(N)$
896	672	602 112	28 900	14 450	13 087	7 225
1280	960	1 228 800	28 900	20 643	19 698	14 745
1536	1152	1 769 472	28 900	24 772	24 252	21 233
1792	1344	2 408 448	28 900	28 900	28 900	28 900
2016	1512	3 048 192	28 900	32 513	33 034	36 577
2176	1632	3 551 232	28 900	35 093	36 021	42 614
2368	1776	4 205 568	28 900	38 190	39 639	50 465
2560	1920	4 915 200	28 900	41 287	43 291	58 981

Table A.7: BOSSbase payload sizes for proportionality constants $0.7r_1, 0.7r_2, 0.7r_3, 0.7r_4$ (given in Table A.3), with the total number of pixels as a reference.

Width	Height	Pixels	$O(1)$	$O(\sqrt{N})$	$O(\sqrt{N} \log N)$	$O(N)$
896	672	602 112	41 287	20 643	18 695	10 321
1280	960	1 228 800	41 287	29 490	28 140	21 064
1536	1152	1 769 472	41 287	35 388	34 646	30 333
1792	1344	2 408 448	41 287	41 287	41 287	41 287
2016	1512	3 048 192	41 287	46 447	47 192	52 253
2176	1632	3 551 232	41 287	50 134	51 459	60 877
2368	1776	4 205 568	41 287	54 557	56 627	72 094
2560	1920	4 915 200	41 287	58 981	61 844	84 259

Table A.8: BOSSbase payload sizes for proportionality constants r_1, r_2, r_3, r_4 (given in Table A.3), with the total number of pixels as a reference.

Width	Height	Pixels	$O(1)$	$O(\sqrt{N})$	$O(\sqrt{N} \log N)$	$O(N)$
896	672	602 112	53 673	26 836	24 304	13 418
1280	960	1 228 800	53 673	38 337	36 582	27 384
1536	1152	1 769 472	53 673	46 005	45 040	39 433
1792	1344	2 408 448	53 673	53 673	53 673	53 673
2016	1512	3 048 192	53 673	60 382	61 350	67 930
2176	1632	3 551 232	53 673	65 174	66 896	79 140
2368	1776	4 205 568	53 673	70 925	73 615	93 722
2560	1920	4 915 200	53 673	76 675	80 398	109 536

Table A.9: BOSSbase payload sizes for proportionality constants $1.3r_1, 1.3r_2, 1.3r_3, 1.3r_4$ (given in Table A.3), with the total number of pixels as a reference.