

Improved Steganographic Embedding Using Feature Restoration

Undergraduate dissertation of Ventsislav K. Chonev

Supervised by Andrew D. Ker

Chapter 1

Introduction

This project is concerned with feature restoration, a problem within the fields of steganography and steganalysis. This chapter introduces both fields briefly, then proceeds to define the problem and outline the structure of the work we have done to solve it.

1.1 Steganography

Cryptography is a science concerned with hiding information. In cryptography, one devises schemes to encrypt messages in such a way that no one apart from the intended recipient may decrypt them. However, a general weakness of cryptography is that, while its methods preserve the secrecy of the exchanged information, they do little to hide the *presence* of communication. Even if an unbreakable cryptographic cipher is used, a passive intruder would still see an unreadable message and become aware of secret communication between the two parties, despite being unable to read it.

The related field of steganography goes further by attempting to conceal the presence of communication altogether. This is done by hiding the message inside a cover - for example, Johannes Trithemius's work *Steganographia* (published in 1606) is ostensibly about magic and the occult but is in fact only a covertext for a treatise of cryptography and steganography.

In the context of modern Computer Science, steganography is often expressed as the *prisoners' problem*. Alice and Bob are in prison, living in separate cells. They each have access to a computer, and may use a communication channel existing between the two computers. This channel is the only way they can exchange information. Alice and Bob each have access to a variety of file formats - text, pictures, video, sound, etc. - and may exchange such files over the communication channel. They are also assumed to share a secret key which they agreed on before being imprisoned. Finally, each of them has access to a random number generator (RNG). The warden Wendy monitors all communication along the channel, and if she becomes suspicious, she will close it for good. It is Alice and Bob's goal to devise an escape plan without arousing Wendy's suspicion. A practical way to do so is to embed messages into cover files using the shared key for encryption and decryption. Steganography is the field concerned with devising such embedding schemes.

1.2 Least Significant Bit Matching

In this project, we will focus exclusively on one steganographic scheme for embedding messages into greyscale images¹, called Least Significant Bit (LSB) matching.

Encoding a message into a cover image is done as follows:

1. Seed the RNG with the steganographic key.
2. Use the RNG to generate a permutation L of the pixels of the cover image.
3. Convert the plaintext into a bitstream B in some standard way (e.g., by concatenating the binary representations of the ASCII codes of the characters). Assume $|B| \leq |L|$.
4. Simultaneously traverse B and L , embedding one bit of B into each pixel in L . To embed a bit b into a pixel p , examine $LSB(p)$ - the parity of p . If $LSB(p) = b$, leave the pixel as it is. Otherwise: if $p = 0$, increment p by 1, if $p = 255$, decrement p by 1, and if $0 < p < 255$, randomly choose to increment or decrement p by 1 with equal probability.

The hidden information B is called the *payload*, and the ratio $\frac{|B|}{|L|}$ is called the *payload percentage*. For example, a payload of 7373 bytes = 58984 bits embedded into a 256×256 image is a 90%-payload. The pixels which hold the bit string B in their LSBs are called *payload pixels*.

Decoding a message from a stego-image is done similarly:

1. Seed the RNG with the steganographic key.
2. Use the RNG to generate a permutation L of the pixels of the stego-image. If the same key was used for the encoding, this permutation will be the same as the one used to create the stego-image.
3. Traverse L . At each pixel p in L , calculate $b = LSB(p)$ and append b at the end of a bit-string B .
4. Convert B into plaintext. If the message length is known, take the prefix of B of that length.

1.3 Steganalysis

Steganalysis is the complement field of steganography; it is concerned with detecting the presence of a hidden payload in a cover file. In the context of the prisoners' problem, steganalysis is Wendy's job: she wants to detect any covert communication between Alice and Bob, without accusing them

¹A grayscale image is just a matrix of 8-bit values. Throughout this project, we used the pgm file format, which is precisely such a matrix, preceded by a short header.

falsely. A common assumption in steganalysis, like in cryptanalysis, is Kerckhoffs's principle: Wendy knows the exact steganographic algorithm used by Alice and Bob, but not their shared key.

Most of modern steganalysis is based on *statistical features*. The idea of feature-based steganalysis is introduced in [4]; since then various feature sets and techniques have been proposed ([5]-[10]). In all cases, the central idea is that covers have a high degree of coherence, which makes their content predictable, unlike the embedded stego signal, which is essentially additive noise. Therefore, feature-based steganalysis performs the opposite to denoising: it strips away the content of the cover to leave only the noise, and then measures chosen statistical characteristics (features) of it. A good feature set should be sensitive to embedding noise, but insensitive to the image content.

Once a feature set is chosen and an extractor is built for it, machine learning techniques are used to train a classifier to distinguish between the features of stego-covers and innocent covers. Classifiers considered in literature include the Fisher Linear Discriminator ([2]), Support Vector Machines and Neural Networks ([3]).

1.4 Feature Restoration

Feature restoration is a technique for reducing the suspiciousness of a stego-image, thereby decreasing the likelihood of detection. When a payload is embedded into a cover image, its features are likely to become highly abnormal and indicative of hidden information. Feature restoration attempts to selectively modify the non-payload pixels of the stego-image in order to bring the features back to normal, and avoid detection by a steganalyzer. This technique is only sensible for high payloads (i.e., above 50%), because efficient techniques for concealing low payloads are known and well-studied. The concept of feature restoration was suggested in [2], but no work has been done to investigate algorithms or to benchmark them against real steganalyzers.

1.5 This Project

This project attempts to fill the void in the literature by devising some algorithms for feature restoration and comparing their performance. Throughout, we will focus solely on the steganographic scheme LSB-matching applied to grayscale images because it is both one of the simplest and least detectable steganographic schemes known. Chapter 2 presents the best known features for detecting LSB-matching, called WAM. Chapter 3 investigates an important property of WAM. Chapter 4 defines a distance metric to measure the suspiciousness of feature vectors. Chapter 5 investigates algorithms for feature restoration.

It should be pointed out that feature restoration is a very difficult problem to solve optimally. Formalising it produces a Binary Integer Quadratic Problem, which is known to be an NP-hard optimisation problem - see Chapter 5 for details. Therefore, our investigation of algorithms for feature restoration will focus on heuristics and approximations rather than ambitiously attempting to solve the problem optimally.

Nearly all of the programming for this project was done in C++. We implemented LSB-matching (see Appendix A for the well-commented code), a feature extractor for WAM (see Appendix B after reading Chapter 2), and the feature restoration algorithms that we devised (see Appendix C after reading Chapter 5). We also wrote much miscellaneous code, such as a reader for *pgm* images, to facilitate the work of these main modules. Finally, everything was glued together using bash scripts to run experiments and Matlab to convert their results into charts.

Chapter 2

Wavelet Absolute Moments

The focus of this project is a set of image features, called Wavelet Absolute Moments (WAM). Their calculation is a multi-stage process involving a wavelet transform. We begin by providing some minimalistic background in wavelets and wavelet transforms. Then we proceed to define the WAM features, and to describe the building of a feature extractor, pointing out some important efficiency considerations.

2.1 Motivation for Wavelet Theory

Wavelet Theory is a field which evolved from Fourier Theory. At the heart of Fourier Theory is the Fourier Transform, which is used to represent a given function in a convenient form. The function $f(x)$ is expressed as a (possibly infinite) sum of sine and cosine waves of varying frequencies and amplitudes. The Fourier series of a periodic $f(x)$ with period L is given by:

$$f(x) = \frac{1}{2}a_0 + \sum_{n=1}^{\infty} a_n \cos\left(\frac{\pi nx}{L}\right) + \sum_{n=1}^{\infty} b_n \sin\left(\frac{\pi nx}{L}\right)$$

$$a_n = \frac{1}{L} \int_{-L}^L f(x) \cos\left(\frac{\pi nx}{L}\right) dx$$

$$b_n = \frac{1}{L} \int_{-L}^L f(x) \sin\left(\frac{\pi nx}{L}\right) dx$$

The Fourier Transform of any $f(x)$ is obtained by extending the series to complex coefficients, replacing the discrete summation with integration, and letting $L \rightarrow \infty$. The transform $F(v)$ gives the amplitudes of the sine and cosine waves of frequency v appearing in the summation of f . Thus, there are two ways to look at a function: its *spatial domain* representation $f(x)$, which allows one to evaluate the function at a specific point, and its *frequency domain* representation $F(v)$, which allows one to examine its frequency content.

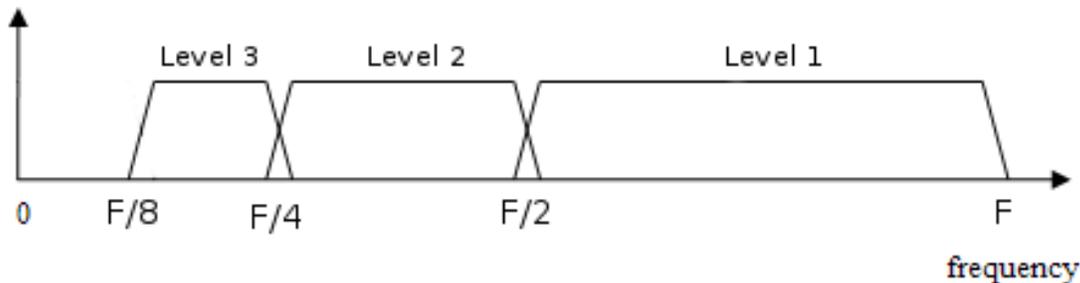
A drawback of Fourier analysis is that sine and cosine waves are global. Hence, it is difficult to examine an interesting region of a function in isolation because all the waves contribute to it. Instead, one might like to zoom in on the interesting section by knowing which summands are relevant to it. This has prompted the emergence of Wavelet Theory.

2.2 The Discrete Wavelet Transform

Wavelet Theory provides a transformation similar to the Fourier Transform, but using a basis of functions localised in space - wavelets. This transformation will allow us to speak of a *wavelet domain*, just like the Fourier Transform produced a frequency domain. The basis consists of dilated, compressed and translated versions of a *mother wavelet*, which is just a specially chosen function satisfying certain requirements that formalise the notion of being a localised wave. In this context, translation and scaling are only allowed in the x direction.

Allowing unconstrained translation and scaling of the mother wavelet would produce an uncountably infinite basis, with a high degree of redundancy. In order to make the wavelet basis countable, we allow only translations by a fixed step $a = 1$ and dilation and compression by a fixed factor $b = 2$. As we are only interested in wavelets that have some overlap with the analysed signal, we can obtain an upper and a lower bound on the translation parameter.

To bound the scaling parameter, consider the Fourier spectrum of the daughter wavelets and the signal. The qualifying requirements for a mother wavelet imply that its spectrum is a band. Fourier Theory shows that compressing such a function by a factor of 2 doubles each frequency present in the spectrum, effectively doubling the spectrum's width and shifting it up. Dilating the mother wavelet has the opposite effect. If we choose the mother wavelet carefully, the spectra of wavelets at consecutive scales will touch each other, or overlap slightly:



Typically, we are interested in bandlimited signals $f(x)$. Such signals have a maximum frequency, so compressing the mother wavelet multiple times will eventually produce wavelets whose frequency spectrum does not intersect the spectrum of $f(x)$, providing a lower bound on the scaling parameter s . However, an upper bound does not exist. Adding an extra scale level to the wavelet basis halves the width of the spectrum left to cover down to 0, so a finite basis of daughter wavelets cannot be enough to represent $f(x)$. The solution is to add an extra function (*father wavelet*) to the basis. Its spectrum is low-pass, so it is a placeholder in the basis for an infinite number of daughter wavelets. Now the basis is finite.

The Discrete Wavelet Transform (DWT) is a mechanism to obtain the wavelet domain representation of a discrete signal for a given wavelet basis. The algorithm is inductive: if the signal is expressed as a weighted sum of daughter wavelets up to a certain scale s and the father wavelet for

that scale, an inductive step may be made to include the next scale $s + 1$ into the basis. This incorporates more dilated wavelets into the basis and reveals finer detail about the signal. The formal details may be found in [12].

In practice, however, the DWT is computed using signal filtering techniques. A Quadrature Mirror Filter (QMF) is designed, specific to the chosen wavelet basis. A QMF is a pair of filters (g, h) , such that g is low-pass and admits only frequencies up to some threshold, whereas h is high-pass and admits only frequencies above that threshold. At each stage, the signal is fed into both filters using a discrete convolution. The operation produces low-frequency output and high-frequency output:

$$x = \text{inputSignal}$$

$$\text{filterLength} = \text{length}[g] = \text{length}[h]$$

$$yLow[i] = \sum_{j=0}^{\text{filterLength}} g[\text{filterLength} - j - 1] * x[i + j]$$

$$yHigh[i] = \sum_{j=0}^{\text{filterLength}} h[\text{filterLength} - j - 1] * x[i + j]$$

The high-pass output $yHigh$ yields the wavelet coefficients for that scale. The low-pass output $yLow$ corresponds to the father wavelet coefficients, and is a coarse approximation of the signal. Then $yLow$ is *downsampled* - its every other entry is discarded, halving its length. The downsampled $yLow$ is given as input to the next stage. This is Mallat's algorithm for computing the DWT.

The algorithm generalises to 2D. Given a 2D signal x and a QMF (g, h) , we can filter each row of x using one of (g, h) , and then each column of the result using one of (g, h) . Thus, 2D filtering may be performed in four ways:

- Using g both for row filtering and column filtering gives the output matrix LL.
- Using g for rows and h for columns gives LH.
- Using h for rows and g for columns gives HL.
- Using h for rows and columns gives HH.

All four output matrices are downsampled by 2, discarding every other row and column. LL is a coarse approximation of the input signal. After downsampling, it is used as input to the next level. LH is typically referred to as the horizontal output band **H**, HL - as the vertical band **V**, and HH - as the diagonal band **D**.

2.3 Definition of WAM Features

Now that we have some knowledge of what a wavelet domain is, we are in a position to define WAM. Given a 2D signal S , its WAM features are defined as follows (following the exposition in [2, 3]):

1. Compute a one-level 2D DWT of S using the Daubechies QMF of length 8, *without* down-sampling. Discard the low-frequency output LL. Denote the three output bands H , V , and D , respectively.
2. Let σ_0^2 denote the noise variance of the signal; for LSB-matching we have $\sigma_0^2 = 0.5$. Estimate the local variance of each element in H , V , and D , using windows of sizes 3, 5, 7 and 9:

$$\sigma_H^2(x, y) = \max(0, \min(h_3, h_5, h_7, h_9) - \sigma_0^2) = \max(0, \min(h_3, h_5, h_7, h_9) - 0.5)$$

where $h_i = \frac{1}{|I|} \sum_{(a,b) \in I} (H(a,b))^2$ is the average squared wavelet coefficient in the $i \times i$ neighbourhood of (x, y) in H . Similarly for $\sigma_V^2(x, y)$ and $\sigma_D^2(x, y)$.

3. Apply a Wiener filter to each of H , V , and D to obtain the denoised wavelet coefficients, then subtract them from the coefficients to leave only the noise residuals:

$$R_H(x, y) = H(x, y) - H_{den}(x, y) = H(x, y) - \frac{\sigma_H^2(x, y)}{\sigma_0^2 + \sigma_H^2(x, y)} H(x, y) = \frac{0.5}{0.5 + \sigma_H^2(x, y)} H(x, y)$$

and similarly for $R_V(x, y)$ and $R_D(x, y)$.

4. Compute the average noise residual in each band: $\overline{R_H}$, $\overline{R_V}$ and $\overline{R_D}$.
5. Finally, compute the first nine central absolute moments of the residuals in each band:

$$M_H^k = \frac{1}{|I|} \sum_{(x,y) \in I} |R_H(x, y) - \overline{R_H}|^k$$

where I is the index set of H , and k takes values 1...9. Similarly for M_V^k and M_D^k .

6. The WAM features of S are the 27-dimensional vector

$$[M_H^1, \dots, M_H^9, M_V^1, \dots, M_V^9, M_D^1, \dots, M_D^9]^T$$

It is worth clarifying the steps in the pipeline. Step 1 transforms the image into the wavelet domain. This is done because it has been shown that features taken from the wavelet domain are more sensitive than ones taken from the spatial domain. Steps 2 and 3 strip away the image content, in order to leave only the noise: the wavelet coefficients are denoised using a quasi-Wiener filter¹, then the denoised coefficients are subtracted from the coefficients themselves, leaving only the noise content. The Wiener filter requires an estimation of local variance, which necessitates step 2. Finally, absolute central moments of the noise residuals are known to provide a good characterisation of the noise of the image ([3]), so they are used as features.

¹The formula given in step 3 is exactly that of the Wiener filter, except we are applying it in the wavelet domain, rather than the frequency domain, hence the 'quasi'.

2.4 WAM Feature Extractor

For the purposes of this project, we built a feature extractor for WAM from the ground up. This included implementing directly the 2D DWT, the estimation of local variance, and the calculation of noise residuals and moments. This tool will be at the heart of our feature restoration algorithms. With it, we can gauge how the features change when some pixels are perturbed in a particular way. Using such queries, we will selectively modify specific non-payload pixels of the stego image in order to reduce the distance² to a target innocent feature vector.

Because we will be relying heavily on the WAM oracle in feature restoration, it is crucial to ensure its efficiency. A very significant part of the effort spent on this project was aimed at building the feature extractor and optimising it. This section gives a brief overview of the challenges involved in creating this tool and how we overcame them. The full code with detailed comments appears in the appendix.

Firstly, we implemented the 2D DWT; this is straightforward from the definition given above. The variance estimation performed in step 2 of the WAM algorithm is more challenging. For each output band $w \in \{LH, HL, HH\}$, for each element $w(x, y)$, we are required to compute the averages of the squared 3×3 , 5×5 , 7×7 and 9×9 regions around (x, y) in w . Directly traversing each of these regions for every (x, y) is wasteful, and would imply visiting each position 81 times. Instead, we used the transformation: $w \rightarrow v$, where $v(x, y) = \sum_{i \leq x, j \leq y} (w(i, j))^2$. This transformation may be

computed much more efficiently from w :

```

PrefixedSquaresForm(matrix w of size N by M)
{
  for i in [0..N) do
    for j in [0..M) do
      v(i, j) = square(w(i, j));

  for i in [0..N) do
    for j in [1..M) do
      v(i, j) = v(i, j-1) + v(i, j);

  for j in [0..M) do
    for i in [1..N) do
      v(i, j) = v(i-1, j) + v(i, j);
  return v;
}

```

The benefit of this form is that it allows us to compute the sum of the squared values in any rectangular region of w using at most 3 additions, making the efficient calculation of local variance trivial:

$$\text{sumOfSquaresInRectangle}(0, 0, X_1, Y_1) = v(X_1, Y_1)$$

$$\text{sumOfSquaresInRectangle}(0, Y_0, X_1, Y_1) = v(X_1, Y_1) - v(X_1, Y_0 - 1)$$

$$\text{sumOfSquaresInRectangle}(X_0, 0, X_1, Y_1) = v(X_1, Y_1) - v(X_0 - 1, Y_1)$$

²For a notion of distance to be defined later.

$$\text{sumOfSquaresInRectangle}(X_0, Y_0, X_1, Y_1) = v(X_1, Y_1) - v(X_0 - 1, Y_1) - v(X_1, Y_0 - 1) + v(X_0 - 1, Y_0 - 1)$$

Calculating the residuals and the moments is straightforward from the definition. We should point out that the moments are the most expensive part of the pipeline.

At this point, we have an oracle capable of computing the WAM features of any given image. For a WAM calculation on a **fresh** image, this pipeline is very efficient. However, most of the time we will be interested in the effect on the features of perturbing some pixels of an image whose features we have already calculated. The best we can do with the oracle outlined above is to perform the perturbation, calculate the WAM features of the resulting image from scratch, and then undo the perturbation. This wastes a lot of computation power, especially considering the local nature of most of the steps involved in the calculation. The following improvements are aimed at lazily recalculating the features of a perturbed image based on cached results of the WAM pipeline prior to the perturbation.

Firstly, consider how the DWT coefficients change when a **single pixel** (X, Y) is perturbed. As 2D filtering is local, most of the wavelet coefficients will remain the same. In fact, changes will be present only in the 8×8 neighbourhood $[X - 7 \dots X] \times [Y - 7 \dots Y]$. Therefore, if we cache the results LL, LH, HL, HH, then we can redo the filtering only on the 8 rows and 8 columns in question.

A similar results holds for local variance and the residuals. As the variance around a DWT coefficient depends only on the 9×9 neighbourhood around it, changing a single pixel (X, Y) of the image affects only the variances in the region $[X - 11 \dots X + 4] \times [Y - 11 \dots Y + 4]$. The recalculation region for residuals is the same as for variances. If we cache the computed variances, perturbing a single pixel would require us to recompute a very small number of them. Unfortunately, there is no good way to update the prefixed squares structure outlined above, so instead we recalculate the dirty variances directly from the definition - by traversing the 9×9 neighbourhood around each changed DWT coefficient - and then update the residuals in the same region.

Unfortunately, the 27 moments need to be recomputed fully, because they are global. Nevertheless, performing the above lazy steps has significantly speeded up the WAM oracle. If we would like to know how the features change when **one** particular pixel is perturbed, we can perform a lazy query. It does the perturbation, updates the local quantities lazily as above, calculates the moments to obtain the feature vector of the perturbed image, and then performs the opposite perturbation in order to return to the original.

Often we will be interested in the effect on the features of **multiple** perturbations, rather than single-pixel ones. To query the effect of such a group change efficiently, we will perform lazy updates on the local quantities around each pixel involved, then calculate the moments only *once* at the end, then undo the changes. A final point on improving the feature extractor is that if too many pixels are involved in a perturbation, it might be cheaper to calculate *all* the variances using the prefixed squares form as above, instead of traversing the 9×9 region around each perturbed DWT coefficient. The exact number of pixels at the threshold will depend on the size of the image; our tests show that for 256×256 images it is about 3.

For testing purposes, we compared output with M. Goljan's original WAM code. In doing so, we uncovered two bugs in the original code. Firstly, in the local variance estimation of $3 \times 3, \dots, 9 \times 9$ regions around the point of interest, a fixed denominator of 9, 25, 49 or 81 is used, even near the

edges, where the elements contributing to the summation are fewer. Secondly, the filtering routine underlying the DWT performs a convolution with the periodicity extension of the signal being filtered, rather than the signal itself. Both of these errors are relevant only near the edges (if at all), so we reimplemented them in order to be able to compare output with the original implementation of WAM. We believe they are inconsequential to the problem of feature restoration.

Chapter 3

Linearity of the WAM Features

This chapter investigates the linearity of the WAM features.

3.1 Definition of Linearity

Let f be a vector containing certain computable features of an image A . Consider two single-pixel changes on A : C_1 and C_2 . Applying C_1 on its own produces an image with feature vector v_1 . Likewise, C_2 on its own produces v_2 . Performing both changes yields v_3 . Let $\delta_1 = v_1 - f$, $\delta_2 = v_2 - f$, $\delta_3 = v_3 - f$. The feature set is said to be linear if $\delta_1 + \delta_2 = \delta_3$ for all pairs of changes C_1, C_2 on all images A .

3.2 Motivation

The aims of this chapter are to determine whether the WAM features are linear in the sense defined above, and, if they are not, to find out under what conditions two changes exhibit some degree of linearity. The motivation is that it would generally be helpful to have an intuitive grasp of how changing an image affects its WAM features. This intuition will be useful in designing algorithms for feature restoration.

From the definition of the WAM features, we can expect that they are not linear. Some nonlinearity is introduced by the calculation of local variances, and much more by the calculation of the nine moments for each subband. Therefore, it is reasonable to expect that $\delta_1 + \delta_2 = \delta_3$ will almost never hold. Moreover, due to the local nature of the variance estimation, we expect nonlinearities to be more evident for changes of pixels which are close to each other.

3.3 Experiments and Conclusions

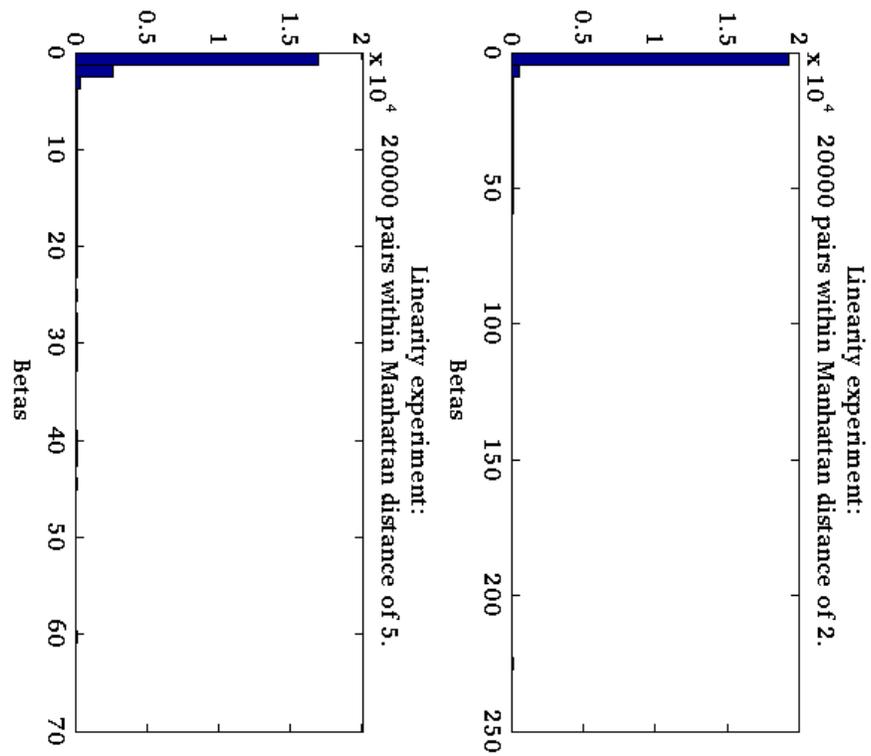
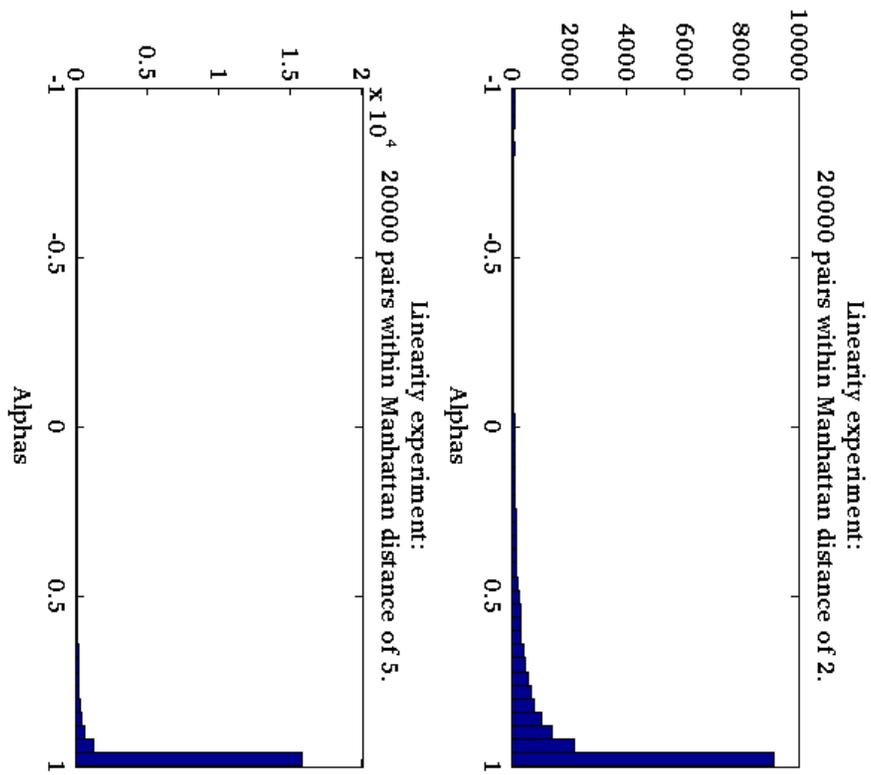
To measure linearity, we introduce two metrics. Firstly, $\alpha = \frac{\delta_3 \cdot (\delta_1 + \delta_2)}{\|\delta_3\| \|\delta_1 + \delta_2\|}$ is the cosine of the angle between $\delta_1 + \delta_2$ and δ_3 . Secondly, $\beta = \frac{\|\delta_1 + \delta_2\|}{\|\delta_3\|}$ is the ratio of their sizes. For linear features, both metrics must be 1 for any pair of single-pixel changes.

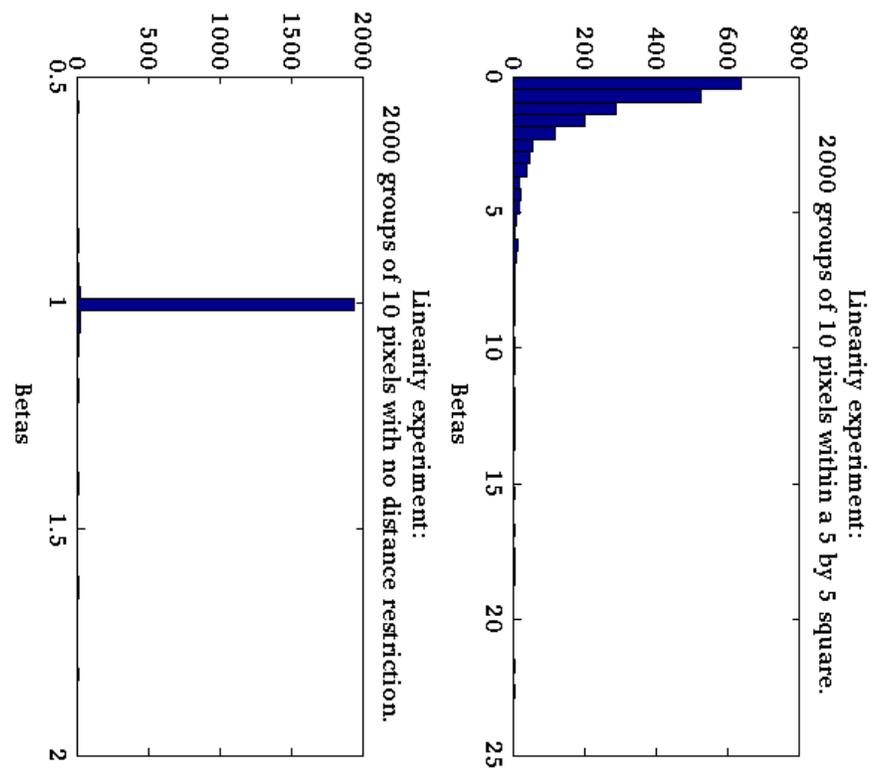
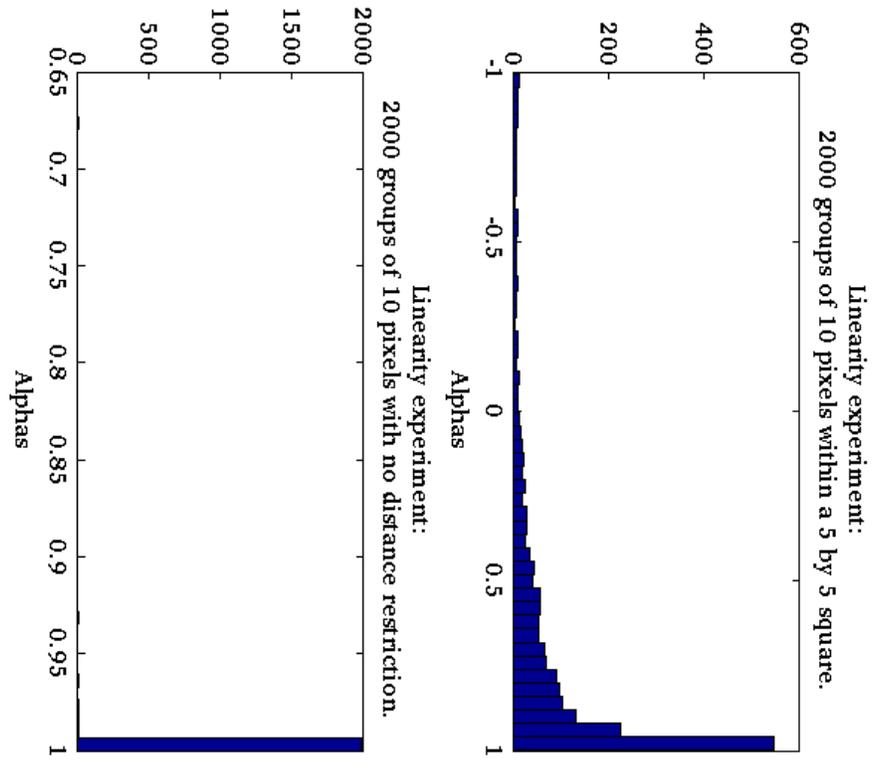
To investigate the linearity of the WAM features, we ran the following experiments. On a fixed image, randomly choose a pair of pixels with Manhattan distance¹ within Y . Randomly choose to perturb their values by 1 or -1 . Query the WAM oracle to obtain δ_1 , δ_2 , and δ_3 - the effect of performing the first change, the second change or both changes, respectively. Then calculate the two metrics and record them. Repeat X times. The experiment was performed with $X = 20000$, $Y = 2, 5, 15, 30, \infty$. Below we show histograms of α and β for $X = 20000$, $Y = 2, 5$.

To test linearity on arbitrarily large groups of changes, we also ran a generalisation of the experiment: Pick X groups of Z changes, contained within a $(a \times b)$ neighbourhood of pixels. For each group, query the oracle to obtain $\sum_{i=1}^Z \delta_i$ (the sum of the individual effects) and δ_{Z+1} (the effect of all the changes together), and calculate the metrics. Histograms for $(X, Z, a, b) = (2000, 10, 5, 5)$ and $(2000, 10, \infty, \infty)$ are shown below.

Contrary to our expectations, pairs of pixel changes which exhibit linearity are very common, particularly when the pixels affected are far apart. The nonlinear nature of the WAM features manifests itself only for pairs of changes which are *very* close together. The result holds for larger groups of changes as well. Knowledge of this pseudo-linearity will be a helpful tool ahead.

¹The Manhattan distance between two points (x_1, y_1) and (x_2, y_2) is $|x_1 - x_2| + |y_1 - y_2|$.





Chapter 4

Feature Distance Metrics

In feature restoration the goal is to make a stego-image less suspicious by changing its non-payload pixels so that the distance between the image's feature vector and a particular target feature vector is minimised. Therefore, a *notion of distance* is needed. This chapter shows why Euclidean distance is a poor choice for feature restoration and defines a distance measure more suited to the task.

4.1 Shortcomings of Euclidean Distance

The main disadvantage of Euclidean distance is its assumption that vector components are uncorrelated. When this assumption fails, the existing correlations are not taken into account by the norm. This could lead to difficulty if we are trying to gauge whether a particular feature vector may pass for innocent, based on an estimation of the cluster of all natural covers. To appreciate this, consider the figure below:



The example is a simplified one as it is in two dimensions, but it serves to illustrate the point. The cluster shown in the figure represents the features of all innocent-looking images. When a payload is embedded into image A, the resulting stego-image is likely to be well outside the cluster. B and C are two candidate restored images. B is inside the cluster and looks less suspicious than C. However, their Euclidean distances to the original A are the same, so they are equally good solutions under Euclidean distance. The WAM features are indeed correlated, because they are 3 sets of 9 successive moments. This indicates the need to use a non-Euclidean notion of distance on WAM feature vectors.

4.2 The Mahalanobis Norm

4.2.1 Notation

The expectation of a random variable x_i is $E(x_i)$. The covariance of two random variables x_i, x_j with $E(x_i) = \mu_i$ and $E(x_j) = \mu_j$ is $Cov(x_i, x_j) = E((x_i - \mu_i)(x_j - \mu_j))$. Extend this notation to random vectors. For a multivariate, random vector $x = (x_1, x_2, \dots, x_n)^T$ with $E(x_i) = \mu_i$, the expected value is $E(x) = (\mu_1, \mu_2, \dots, \mu_n)^T = \mu$, and the covariance matrix is $S = E((x - \mu)(x - \mu)^T)$, so that $S_{ij} = E((x_i - \mu_i)(x_j - \mu_j)) = Cov(x_i, x_j)$.

Note that the covariance matrix of a real-valued random vector is symmetric, therefore, by the finite-dimensional spectral theorem, it is always diagonalizable by an orthogonal matrix Q : $S = Q\Lambda Q^T$ where $Q^T Q = Q Q^T = I$ and Λ is a diagonal matrix containing the eigenvalues of S . Recall also that any invertible covariance matrix is positive definite, which makes its eigenvalues strictly positive¹. This implies the existence of $\Lambda^{-\frac{1}{2}}$, obtained by replacing every diagonal entry of Λ by the root of its reciprocal.

4.2.2 Whitening

Let the covariance matrix of the random vector x be S , diagonalized as $Q\Lambda Q^T$, and let the mean of x be μ . Define the linear transformation $w(x) = \Lambda^{-\frac{1}{2}} Q^T (x - \mu)$. We can make two observations about w .

Firstly, $E(w)$ is the zero vector. This is immediate from linearity of expectation:

$$E(w) = E(\Lambda^{-\frac{1}{2}} Q^T (x - \mu)) = \Lambda^{-\frac{1}{2}} Q^T (\mu - \mu) = 0$$

Secondly, the covariance matrix of w is the identity matrix I :

$$\begin{aligned} E((w - E(w))(w - E(w))^T) &= E(ww^T) = E(\Lambda^{-\frac{1}{2}} Q^T S Q \Lambda^{-\frac{1}{2}}) = \\ &= E(\Lambda^{-\frac{1}{2}} Q^T Q \Lambda Q^T Q \Lambda^{-\frac{1}{2}}) = E(\Lambda^{-\frac{1}{2}} I \Lambda \Lambda^{-\frac{1}{2}}) = I \end{aligned}$$

¹In general, covariance matrices are positive semi-definite, which allows 0 to be an eigenvalue, making the matrix singular. This is the case exactly when one of the random variables is a linear combination of the others, which would mean the random vector was not entirely random to begin with. In this discussion, we ignore such pathological cases.

Now, consider applying the transformation w to the whole image of the random vector x . The first observation says that the resulting cluster will be centred around the origin. The second says that any two distinct components of the transformed vector will have zero covariance, so there will be **no correlation** among the vector components. The transformation squashes the cluster into a multi-dimensional sphere around the origin. This transformation is called *whitening*, and the resulting vectors are called *white*.

4.2.3 Mahalanobis's Distance

For a whitened the cluster, the issues discussed above are no longer present. A natural idea in defining a norm that accounts for correlations is to whiten the cluster using w , and then use Euclidean distance on the result. The distance between the centre of the cluster μ and a vector x is:

$$\begin{aligned} \|w(x) - w(\mu)\|_E &= \|w(x)\|_E = \|\Lambda^{-\frac{1}{2}}Q^T(x - \mu)\|_E = \sqrt{((x - \mu)^T Q(\Lambda^{-\frac{1}{2}})^T)(\Lambda^{-\frac{1}{2}}Q^T(x - \mu))} = \\ &= \sqrt{(x - \mu)^T Q\Lambda^{-1}Q^T(x - \mu)} = \sqrt{(x - \mu)^T S^{-1}(x - \mu)} \end{aligned}$$

More generally, a measure of distance between two arbitrary points x and y is:

$$\begin{aligned} \|w(x) - w(y)\|_E &= \sqrt{w(x)^T w(x) - 2w(x)^T w(y) + w(y)^T w(y)} = \\ &= \sqrt{(x - \mu)^T S^{-1}(x - \mu) - 2(x - \mu)^T S^{-1}(y - \mu) + (y - \mu)^T S^{-1}(y - \mu)} \\ &= \sqrt{x^T S^{-1}x + y^T S^{-1}y - 2x^T S^{-1}y} = \sqrt{(x - y)^T S^{-1}(x - y)} \end{aligned}$$

Thus, the Mahalanobis norm of x with mean μ and covariance matrix S is defined by $\|x\|_M = \sqrt{(x - \mu)^T S^{-1}(x - \mu)}$, and the Mahalanobis distance of x and y is $Dist_M(x, y) = \sqrt{(x - y)^T S^{-1}(x - y)}$. The Mahalanobis distance is a measure of dissimilarity. A small distance means x and y are equally likely to belong to the cluster. Conversely, a large distance means one is well within the cluster, and the other is far from it. This notion of distance was originally introduced in [11].

In order to use the Mahalanobis distance, we need two estimations. Firstly, the feature vector μ - the average of the WAM vectors of all natural cover images. Here, by 'natural' we mean ones that carry no steganographic payload. Secondly, the 27 by 27 covariance matrix S of the WAM vector, along with its inverse. Computing these exactly is impossible, as it would mean calculating the WAM features of all natural images.

For the purposes of estimating μ and S , we used 17500 natural images, IMG_1 to IMG_{17500} , each with dimensions 256×256 , cropped from 500 images with dimensions 1500×2000 . The original 3Mpx images were taken with a Minolta DIMAGE A1 camera, and have never been subjected to compression. The estimation of the mean vector was:

$$\mu = \frac{1}{17500} \sum_{i=1}^{17500} WAM(IMG_i)$$

We obtained the unbiased estimate of the covariance matrix:

$$\text{Cov}(x_i, x_j) = \frac{1}{17499} \sum_{t=1}^{17500} (\text{WAM}(IMG_t)_i - \mu_i) (\text{WAM}(IMG_t)_j - \mu_j)$$

Then we inverted the covariance matrix using Matlab. The choice was motivated by the fact that Matlab's implementation of the Gauss-Jordan algorithm is known to have a good degree of numerical stability, hopefully reducing precision-related problems to a minimum.

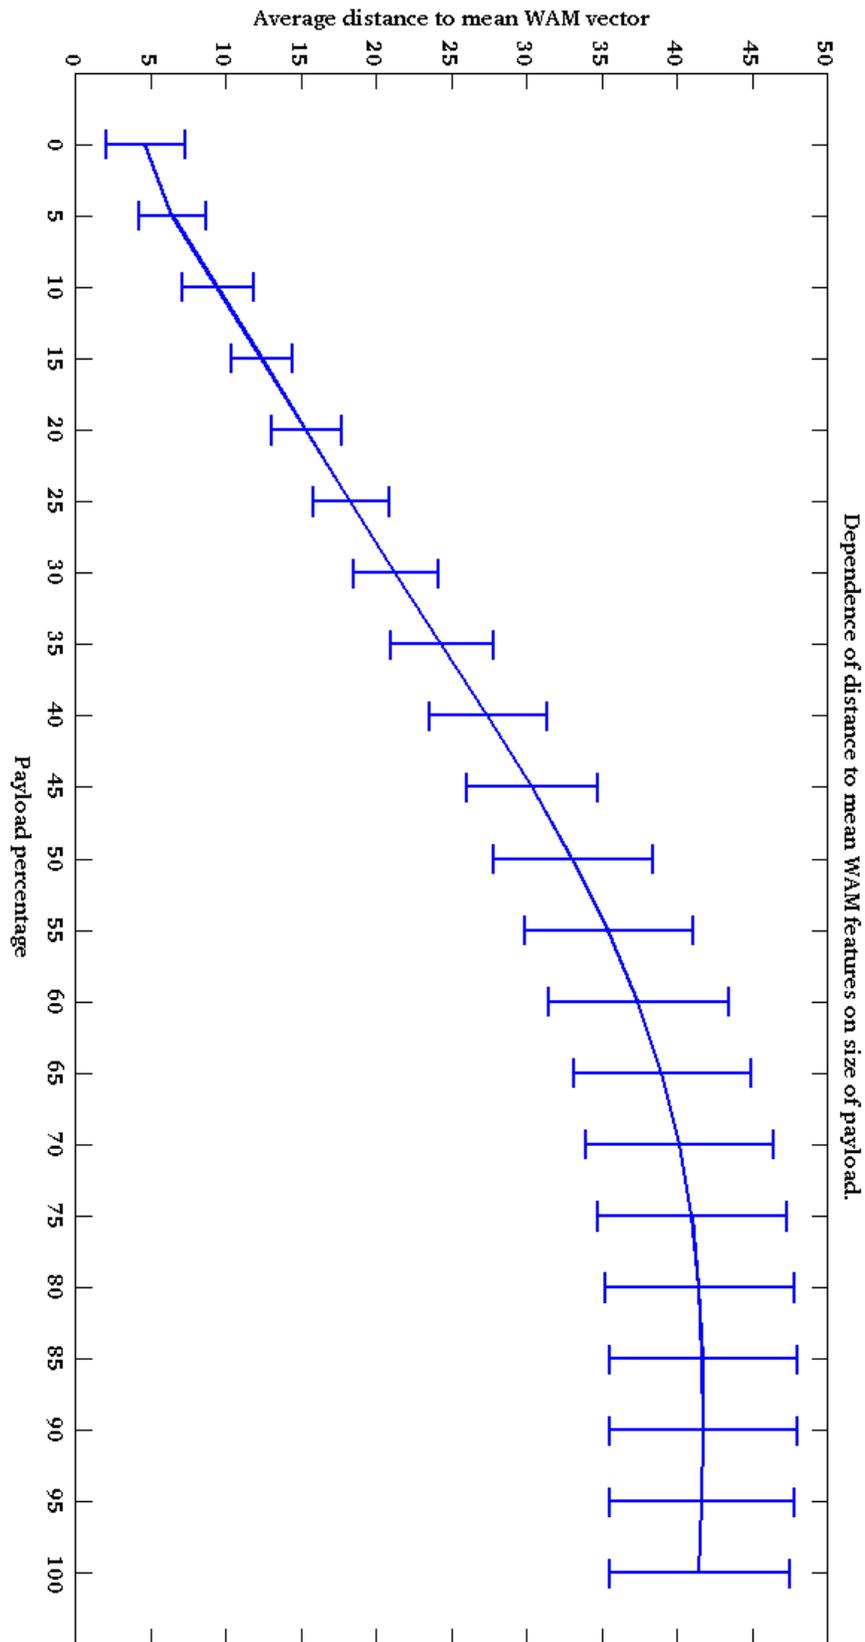
4.3 Contextualising the Distance

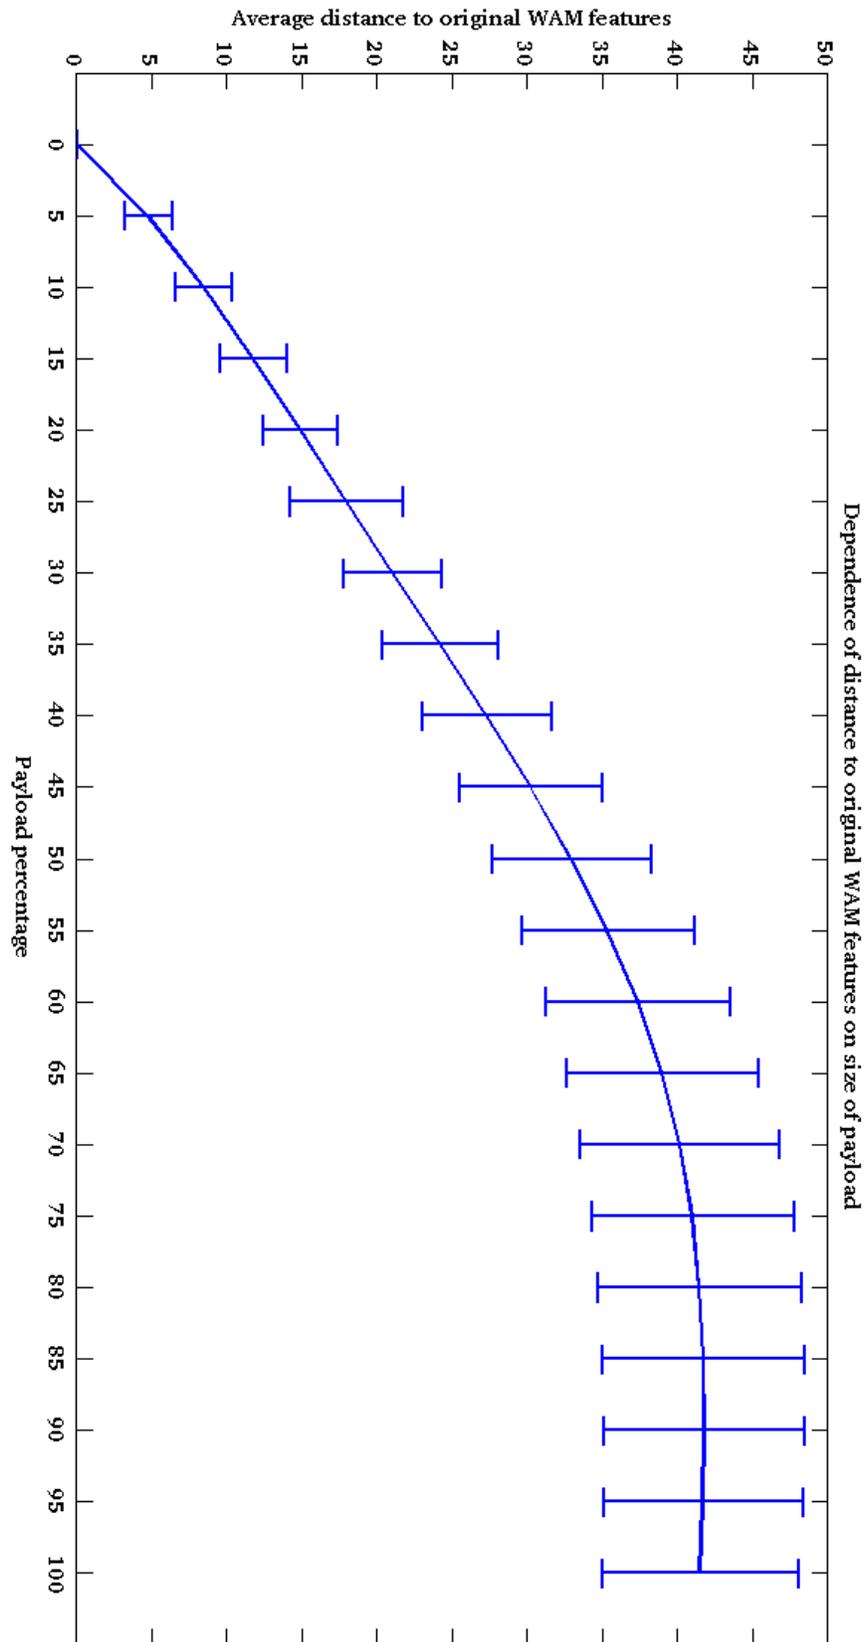
Let g and f be the feature vectors of the original image and the stego-image, respectively. We may choose from two distance measures to minimise in feature restoration:

- $\text{Dist}(g, f)_M$, the Mahalanobis distance to the **original** image's WAM vector
- $\text{Dist}(\mu, f)_M$, the Mahalanobis distance to the estimated **mean** WAM vector

Since the WAM features are sensitive to noise, and a larger payload introduces more noise, it is to be expected that the distance measures depend on the payload percentage. This section provides context for the numbers, and gives an intuition of how embedding various payloads changes the distance measures.

We estimated the average value of the distances for a number of fixed payloads. For each payload percentage, we embedded a message of the appropriate length into the covers IMG_1 to IMG_{17500} , calculated their WAM features, and computed the distances. Finally, the average and the standard deviation of the distances at that payload were computed. The results appear in the two figures below. They look almost exactly the same, so at this point we make the decision that we will only use $\text{Dist}(\mu, f)_M$ to measure the performance of feature restoration algorithms. With this choice of target, we will actually be making the features plausible, instead of restoring them.





Chapter 5

Feature Restoration

Now we are in a position to attack the problem of feature restoration. We have chosen a steganographic scheme for encoding and decoding messages into and from greyscale images (LSB-matching) and a particular set of features for steganalysis (WAM), and have defined a notion of an image's suspiciousness (the Mahalanobis distance between its WAM vector and the estimated mean WAM vector μ). Feature restoration is the problem of deliberately altering the non-payload pixels of a stego-image in order to reduce its suspiciousness. At our disposal, we have two black-box tools: an efficient oracle able to calculate the features of any image and foresee the effect of any perturbation, and a distance calculator able to measure the suspiciousness of any feature vector. We are also aware that the features are pseudo-linear.

In this chapter, we present a number of algorithms we have devised for feature restoration. We have implemented each of these algorithms and benchmarked its performance. Each section will describe an algorithm, along with the ideas motivating it, and give results from our benchmarking. We conclude the chapter with a comparison and a pointer to further work that should be done on the topic.

To benchmark each algorithm A, we used a set of 100 images of size 256×256 . For each payload size of interest, we embedded a payload of that size into each of the 100 images, obtaining a stego set for that payload. Then we ran A on each of the 100 stego images, giving it an allowance of 50000 WAM oracle queries per image. We kept track of the distance to μ at each query count, and averaged this distance over the 100 stego images, obtaining a curve of the average distance to μ at each query count. The decision to measure distance reduction versus queries performed instead of clock time was based on the idea that query usage is a performance measure independent of implementation details and hardware¹. The payloads we considered in this way were 50% (a low payload), 90% (a high payload but nonetheless allowing some degree of freedom), and 99% (an extremely high and constraining payload).

Some of our algorithms have a parameter space. In these cases, we use smaller-scale benchmarking in order to explore the parameter space and converge on a good set of parameters. This uses only

¹Just to give the reader an intuition of how queries used translate to time: 1000 queries take about 11 seconds when each query refers to a single-pixel change, and about 30 seconds when each query refers to 25 changes. This is on an Intel(R) Core(TM) i5 CPU running at 2.27 GHz with 4 GB of RAM.

20 stego images, at only one payload percentage (90%). Then the algorithm with the chosen parameters is benchmarked fully.

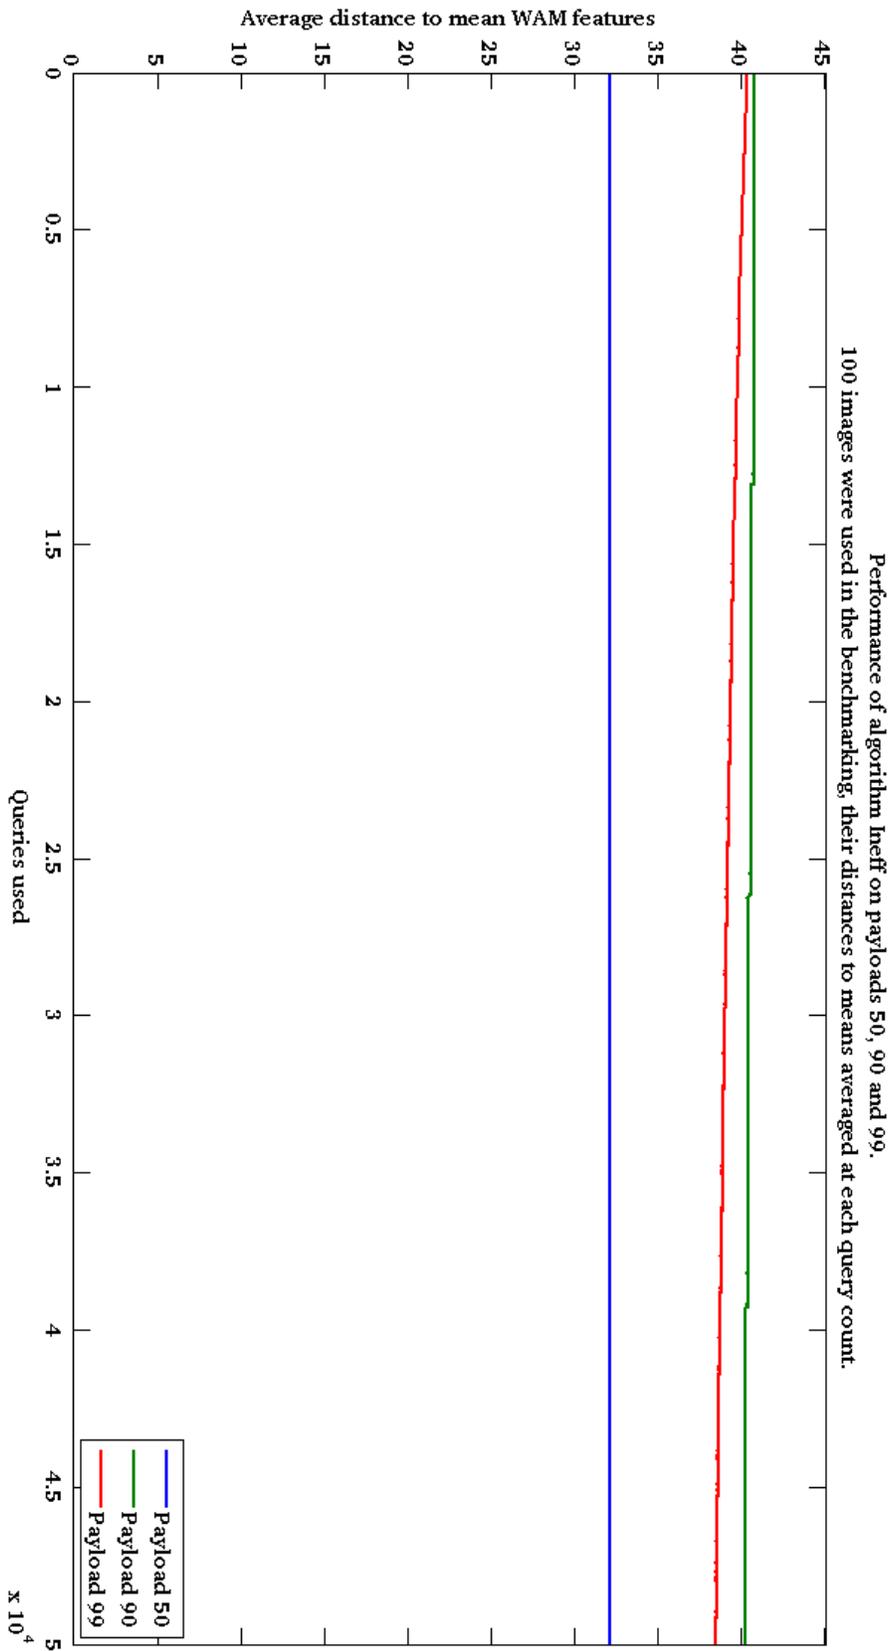
A final point concerns the maximal value by which a restoration algorithm may change each non-payload pixel. It is clear that such a maximum must be set and adhered to in order to prevent distortions of the image which are visible to the naked eye. However, setting such a limit is not easy because the question of which distortions are visible is subjective. The matter is further complicated by the fact that cameras often have imperfections which result in artifacts appearing in natural images. We have fixed this maximum change cap at 10. This limit, though perhaps somewhat generous, was chosen for two reasons. Firstly, 10 appears to be the threshold where distortions of the output image are invisible to the casual observer, but barely noticeable upon very close inspection (zooming in). Secondly, we were interested in finding out the extent to which the nature of WAM allows the distance to μ to be reduced; therefore, we elected not to constrain the feature restoration algorithms too much with subjective notions of how noticeable a distortion is and whether it may be plausibly blamed on the camera. If one is indeed concerned about masking one's attempt at feature restoration against a human observer, one may use a smaller limit in the interest of remaining undetected. The limit was implemented and adhered to in our algorithms and in their benchmarking.

5.1 Greedy and Inefficient Algorithm

We begin with the only algorithm for feature restoration appearing in literature. It was suggested in [2], more to introduce the concept of feature restoration than as a serious attempt at solving the problem.

The algorithm performs a number of iterations. On each iteration, it considers each non-payload pixel in turn, and queries the WAM oracle twice - once for the effect of perturbing the pixel by 1, and once for -1. Throughout the iteration, the algorithm keeps track of the perturbation which brings the feature vector closest to μ . At the end of the iteration, the algorithm realises the best change found. Iterations are performed for as long as there are queries available. If an iteration finds no distance-reducing change, the algorithm terminates.

Unfortunately, this algorithm is very inefficient. At each iteration it uses a very large number of queries - $2(1-p)NM$, where p is the payload percentage and (N, M) are the dimensions of the image - but only performs one actual change. The smaller the payload, the less feasible this algorithm becomes, as each iteration requires a larger number of oracle queries. The care taken in selecting the best available change at each iteration would be justified if this change entailed a very significant distance reduction compared to runners-up. However, our experiments showed that at each iteration, a very large number of the changes considered entail a distance reduction virtually as good as that of the best change, so the queries spent on them are essentially wasted. The performance graphs of the algorithm are flat:

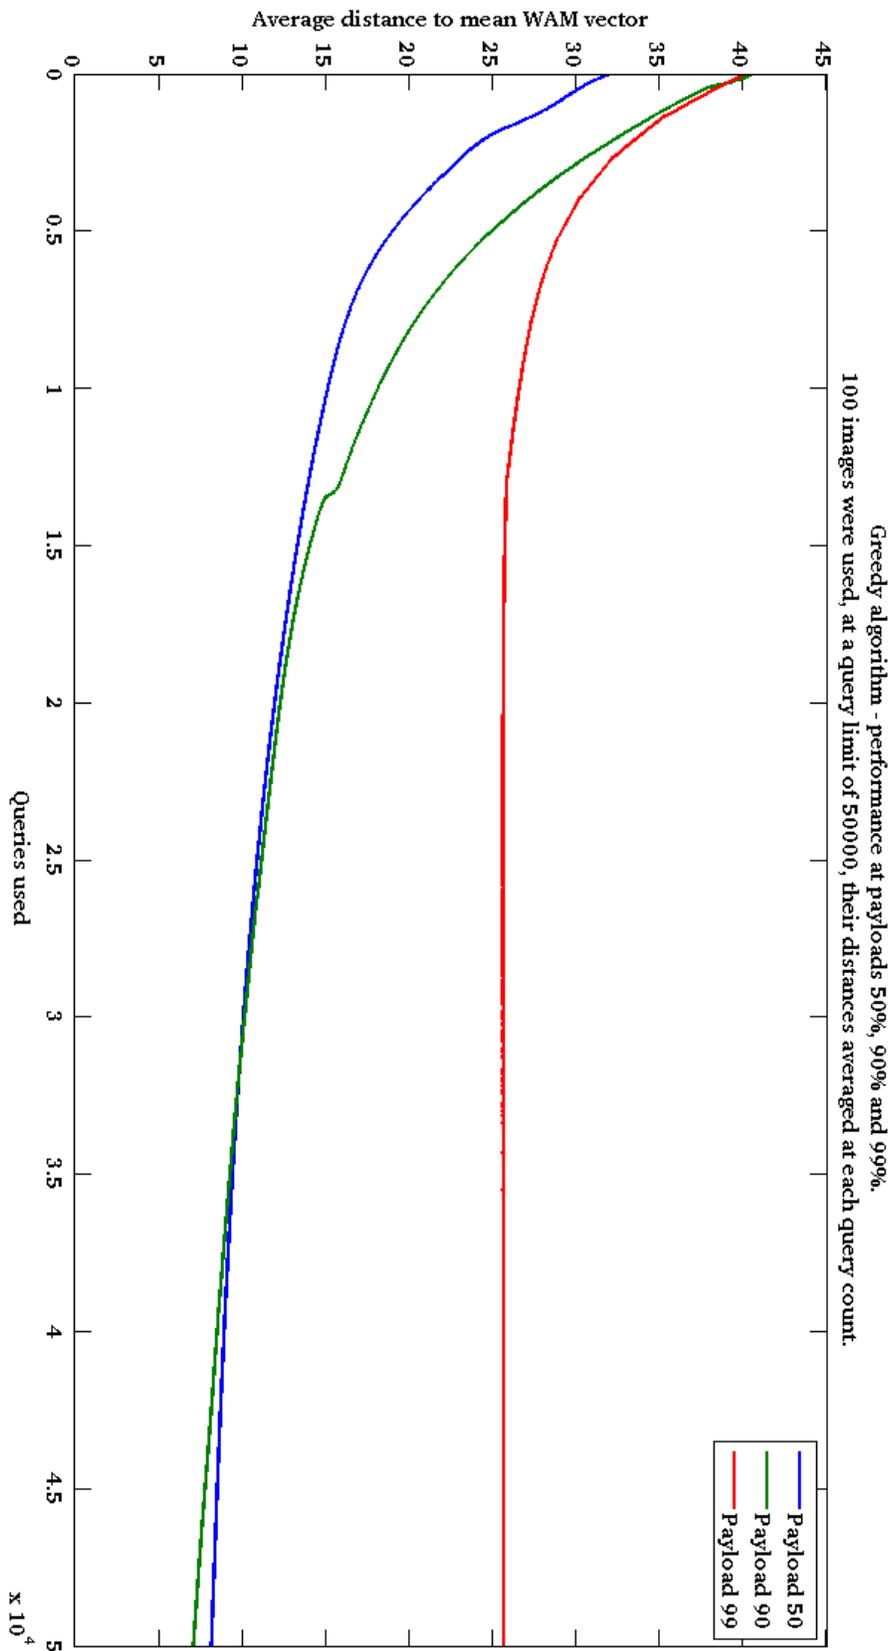


5.2 Greedy Algorithm

While the algorithm in the previous section is infeasible except for very large payloads, it offers a good starting point for refinement. As we mentioned earlier, the inefficient algorithm discards many reducing changes in its search for the best one, thereby wasting far too many queries. Many of these discarded changes in fact entail a large distance reduction. Moreover, having seen the pseudo-linearity of the WAM features, we can expect that if a change C_1 is distance-reducing *before* change C_2 is performed, then C_1 is likely to still be distance-reducing *after* C_2 is performed, especially if the pixels involved in C_1 and C_2 are far apart. These observations suggest a much more efficient greedy algorithm.

Traverse the non-payload pixels of the image in order. For each pixel (i, j) , query the WAM oracle for the result of perturbing (i, j) by 1 and -1. If neither change is reducing, move on. If one of the two changes is reducing, perform it and move on. If both are reducing, perform the one which entails a greater distance reduction and move on. If all non-payload pixels are traversed in this way, wrap around and traverse them again. Repeat for as long as oracle queries are available.

The performance charts for this greedy algorithm are found below. An interesting point to note in the chart is that the curve corresponding to our experiments at 90% payload eventually overtakes the curve for 50%. We attribute this to the difference in the number of non-payload pixels. At 50%, a single traversal of all the pixels takes much more queries than at 90%. Consequently, at 90%, each individual pixel is visited by the algorithm much more frequently. The better performance of the algorithm at 90% payload suggests that there are often pixels which should receive concentrated attention and be perturbed multiple times for maximum distance reduction. This insight leads to the algorithm in the next section.



5.3 Variance Sort Algorithm

A possible weakness of the Greedy algorithm is that it spends an equal proportion of its query allowance on each non-payload pixel, when it might be better to traverse the pixels in a carefully chosen order and, when visiting a pixel, attempt to do as much work on it as possible, instead of simply perturbing it by +1 or -1 and moving on. This sets two questions: firstly, how to decide how much work is needed at each pixel, and secondly, what order to traverse them in?

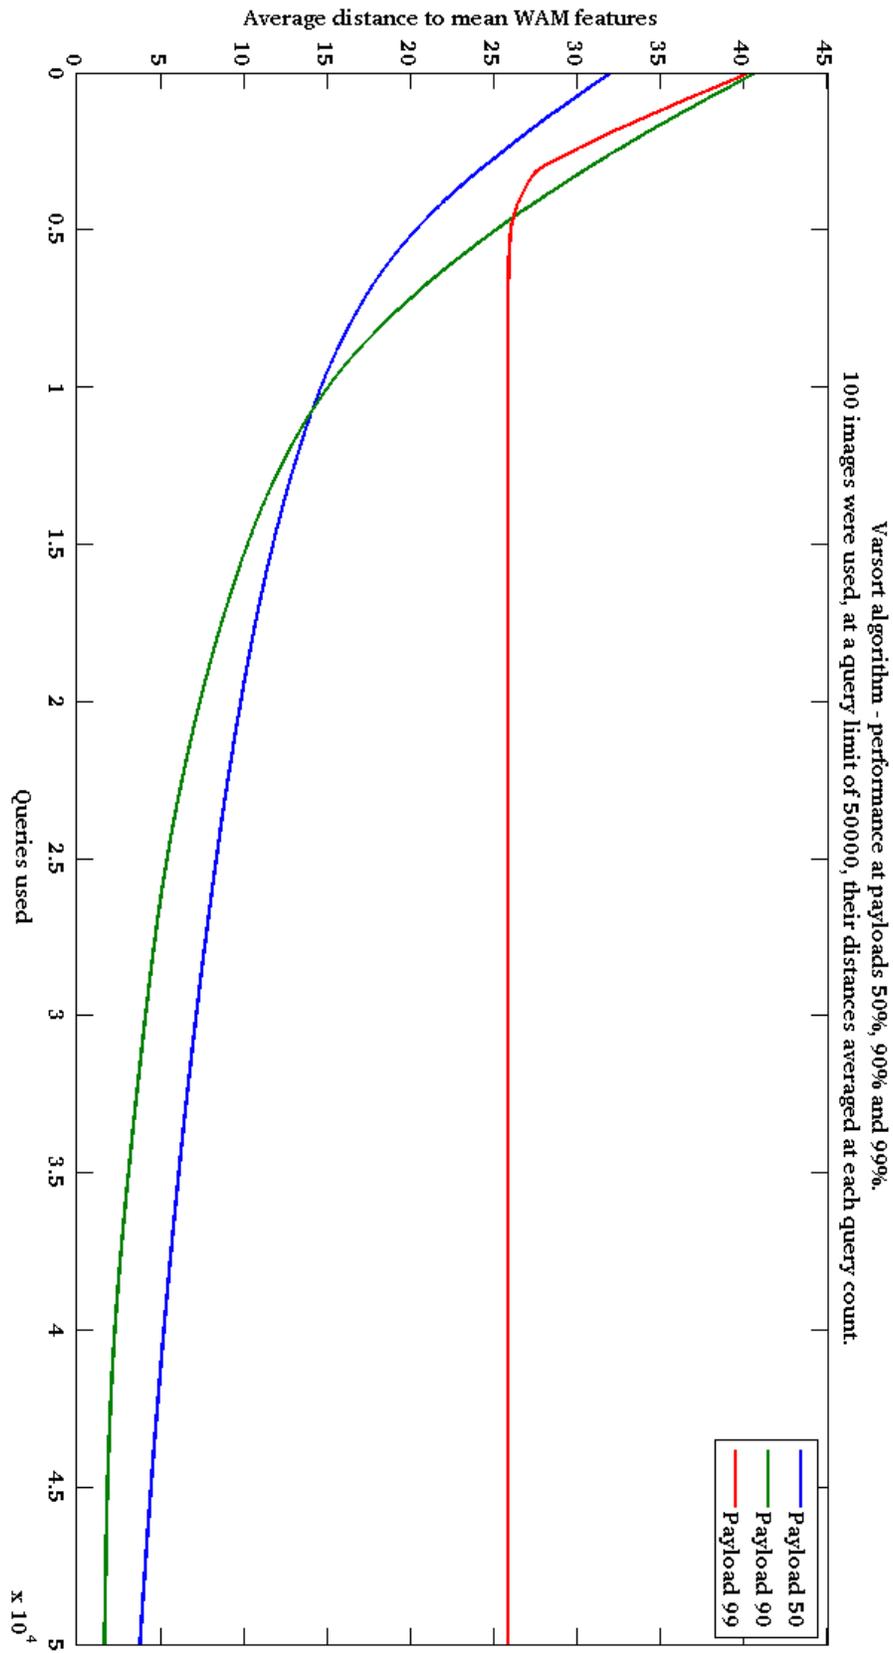
A good answer to the first question is to begin by querying the oracle for the effects of perturbing the pixel (i, j) by +1 and -1 in order to determine a *direction of change*. If both changes increase the distance to μ , ignore this pixel. If one change increases the distance, but the other decreases it, choose the latter. If both changes decrease the distance, choose the one which entails a greater distance reduction. After a direction $d = \pm 1$ is chosen, repeatedly perturb the pixel by d , querying the oracle for the effect of the change (i, j, d) at each step, for as long as this change reduces the distance to μ .²

To answer the second question, we remind ourselves that the WAM features are a measure of noise. Noisy regions of the image contribute most to the distance to μ , so we should focus on them first. Therefore, we decided to estimate the variance in the *spatial* domain at each pixel, and to consider pixels in the way outlined above in order of decreasing variance. In this way, we hope to systematically restore the noisiest bits of the stego-image first. The algorithm is:

1. Estimate the variance in the spatial domain around each non-payload pixel (i, j) . To do this, consider the 5×5 square of pixel values with (i, j) at its centre; let these be x_1 to x_n (most of the time n will be 25, near the edges it will be less). Their mean is $\bar{x} = \frac{1}{n} \sum_{k=1}^n x_k$, and their variance is the unbiased estimator $\widehat{Var}(x) = \frac{1}{n-1} \sum_{k=1}^n (x_k - \bar{x})^2$. Do this for all non-payload pixels, and sort them in order of decreasing variance to obtain the sorted list of pixels L .
2. Traverse L . For each pixel (i, j) in L , perform step 3. If you reach the end of L , wrap around to its beginning. If you run out of oracle queries, terminate.
3. At pixel (i, j) , query the WAM oracle twice - with $(i, j, 1)$ and $(i, j, -1)$. If both changes increase the distance to μ , move on. Otherwise, select $v = \pm 1$ such that (i, j, v) entails a greater reduction of the distance to μ . Then repeatedly perform change (i, j, v) and query the oracle with (i, j, v) ; if the oracle predicts an increase in the distance then stop and move on to another pixel. Also, (i, j) is allowed to differ by at most 10 from its value at the beginning of the restoration algorithm. If this limit is reached, move on to another pixel.

The performance chart for the algorithm appears below. The Varsort algorithm is a definite improvement over the Greedy one. At payloads 50% and 90%, it completely conceals the payload within the query allowance, whereas at payload 99% it brings the distance down to the same level as the Greedy algorithm, but in fewer queries.

²And, of course, for as long as the cumulative change at pixel (i, j) is within the limit of 10 that we set at the beginning.



5.4 Genetic Algorithm

The next algorithm attempts to combine randomness as a source of candidate perturbations with a deliberate process of selecting the most beneficial ones. It was inspired by genetic algorithms.

Recall that a *change* or *singleton change* is just a triple (x, y, v) where x and y refer to the position of a pixel, and v is the value by which it is perturbed. Recall also that a “group change” is a set of changes. On a given image, an *individual* is defined as a pair $(gc, distRed)$ where gc is a group change and $distRed$ is a real number - the amount by which the distance between the WAM features of the image and μ would be reduced if gc were performed. A *population* is defined to be a set of individuals.

Our genetic algorithm will maintain a population P which changes over time. The following operators are defined on it:

- *Inward Migration.* This is the mechanism by which the population grows. A singleton change is randomly chosen, with its value component equal to 1 or -1. Then the WAM oracle is queried with the change to obtain the distance reduction it entails. If this reduction is negative (that is, the change *increases* the distance to μ), the change is discarded and another attempt is made. When a singleton distance-reducing change is found, it is wrapped with its distance reduction label to form an individual, and is inserted into P .
- *Merge.* This is a mechanism to hopefully obtain good distance-reducing changes from existing ones in the population and to eliminate bad ones. All the individuals in P are sorted in decreasing order of their distance reduction labels. The worst few individuals are discarded; the exact number or proportion may be specified as an argument to the operation. Then the remaining individuals are paired up: the best with the 2nd best, the 3rd with the 4th, and so on. In case of an odd number of individuals, the worst is discarded. Then the two individuals in each pair are *merged*: the group change of the result is the union of the two group changes³, and the distance label of the result is obtained by querying the WAM oracle.
- *Outward Migration.* This is the mechanism which alters the image by applying individuals to it. The individual with the greatest distance reduction label in P is found and removed from P . If its distance reduction label is positive, its group change is applied to the image. If not, it is discarded. It is also discarded if performing the group change would make one or more pixels differ from their original values by more than our limit of 10.

Note that distance reduction labels may become outdated. An individual’s distance reduction label is correct when the individual is introduced into P with an inward migration (because the migration explicitly queries the WAM oracle), but afterwards the image may be modified by outward migrations, making distance reduction labels inaccurate. However, the pseudo-linearity of WAM implies that the inaccuracy is likely to be small. Thus, a distance reduction label in P should be seen as very untrustworthy only if a large number of outward migrations have been performed since the individual was introduced into the population.

³If the two group changes being unioned both refer to some pixel (x, y) with perturbation values u and v , respectively, then the result of the union refers to (x, y) with value $u+v$.

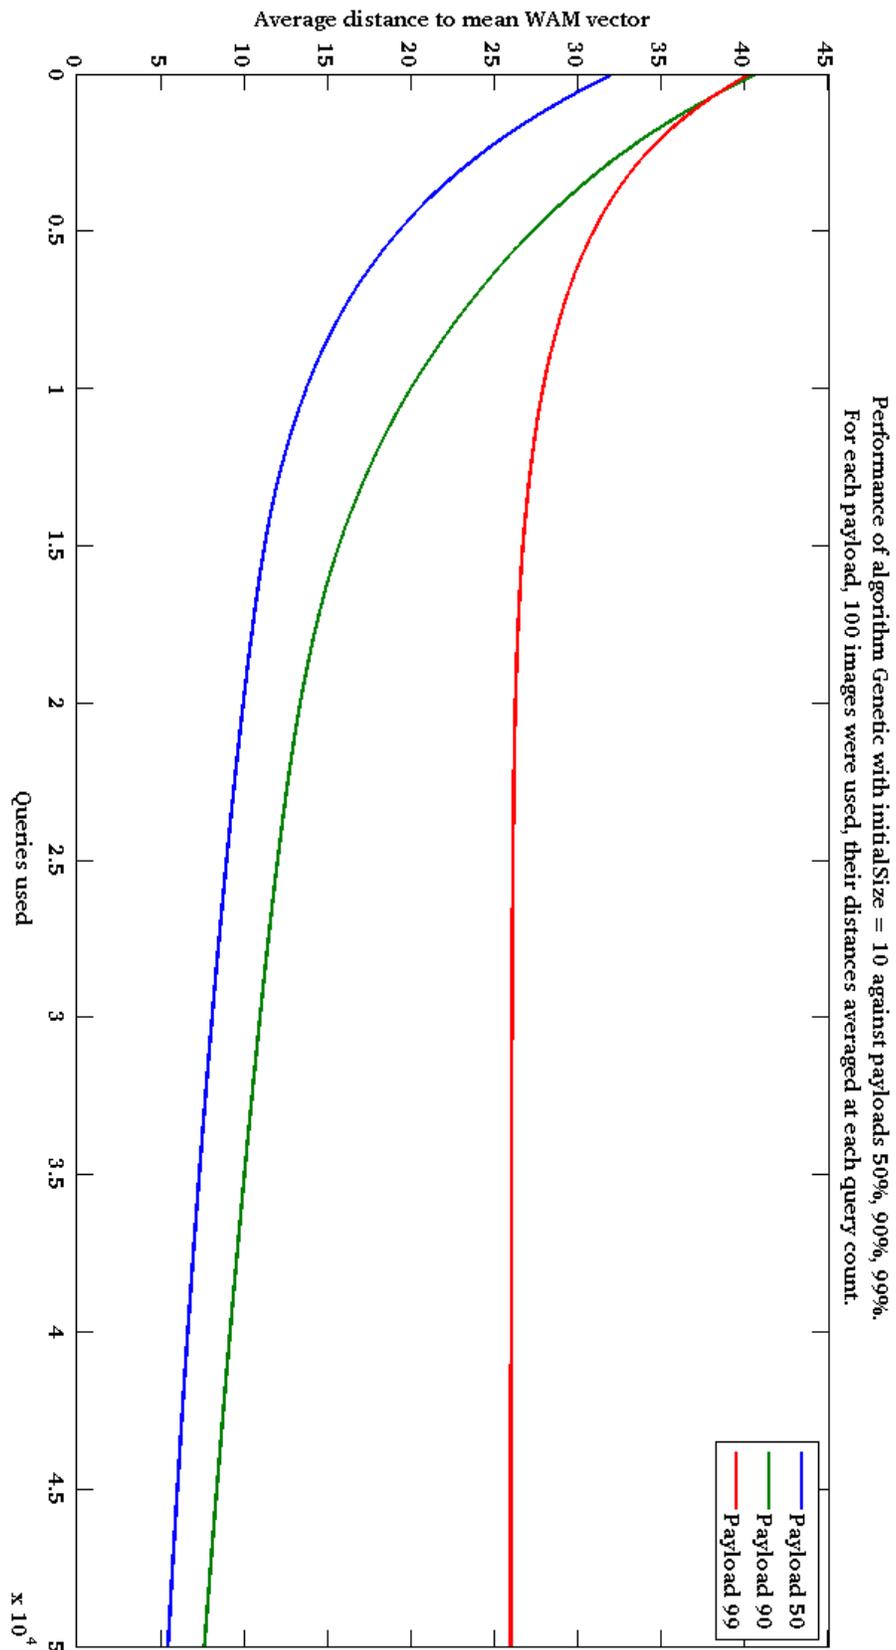
The mechanism of merging serves three purposes. Firstly, it eliminates the worst elements of the population, in the spirit of genetic algorithms. Secondly, it attempts to obtain better candidate group changes by combining good ones appearing in the population, as is common in genetic algorithms. More subtly, however, it refreshes the distance-reduction labels of the whole population: recall that the label of the union of two individuals is obtained by querying the WAM oracle directly. Thus, all the labels appearing in P are accurate immediately after a Merge, so applying this operation occasionally should serve to increase their reliability.

Having defined the operations above, we only need to decide how to combine them in order to obtain a full algorithm. Our algorithm has a parameter *initialSize*. Throughout, P will vary in size from *initialSize* to $2.1 * \textit{initialSize}$. Our algorithm is:

1. Initialise P by performing an Inward Migration *initialSize* times.
2. Iterate for as long as oracle queries are available:
3. On each iteration, perform two Inward Migrations and one Outward Migration. Thus, on each iteration, the size of P grows by 1.
4. If this size reaches $2.1 * \textit{initialSize}$, perform a Merge which begins by discarding $0.1 * \textit{initialSize}$ individuals. Then the merging process halves the size of the population, bringing it down to *initialSize*.
5. When the query allowance is exhausted, perform Outward Migrations until P becomes empty, then terminate.

In our implementation, the population is a priority queue, with higher priority given to individuals with greater distance reduction, to speed up the operations defined above. As a minor implementation issue, we point out that care was needed to ensure that the algorithm does not use more queries than it is allowed to, considering that both Merge and Inward Migration use more than one query at a time. Trivial modifications were made to allow these routines to stop partway through and avoid exceeding the query limit.

To obtain a good value for *initialSize*, we performed some partial benchmarking with a number of possible values (see Appendix D). Our results showed that a small *initialSize* is preferable. Therefore, we proceeded to fully benchmark the algorithm with *initialSize* = 10. The results are in the chart below. The algorithm achieves a good distance reduction overall, but is outclassed by the Greedy and Varsort algorithms.



5.5 Random Algorithm

Now that we have seen an algorithm which partly uses randomness, it is interesting to explore one which is purely random. It is a very simple concept, but the results from it yielded two important insights into feature restoration.

The algorithm has two parameters, *lots* and *changesPerLot*, and works as follows:

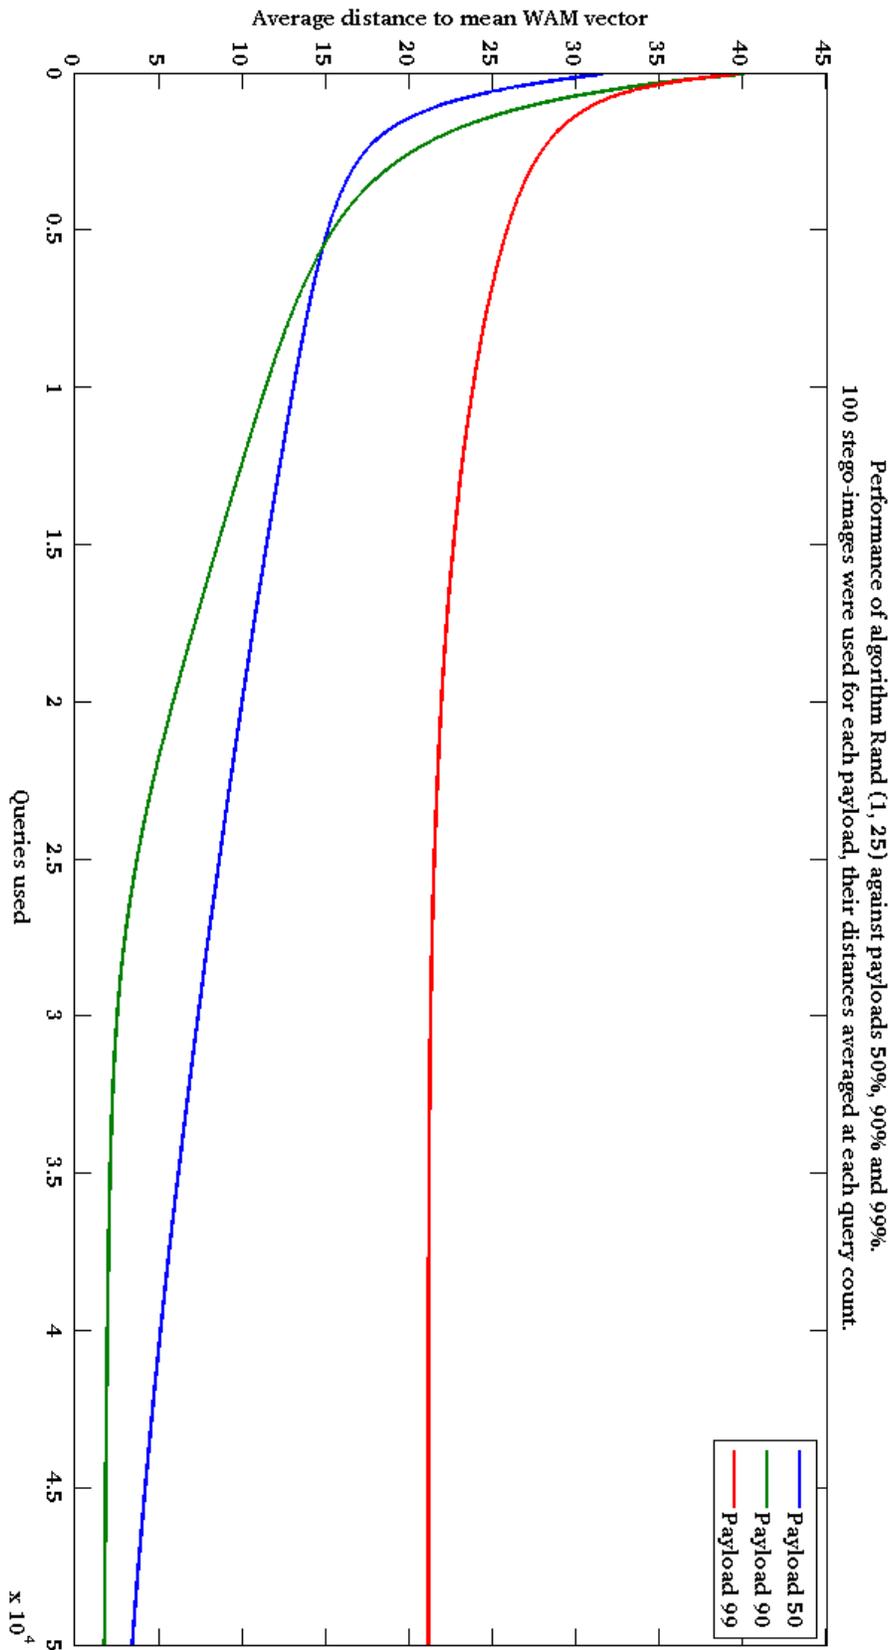
1. Iterate for as long as there are queries available:
2. On each iteration, randomly choose *lots* group changes. Each group change must contain *changesPerLot* singleton changes. Each singleton change must have value equal to -1, 0 or 1; at least one singleton change must have a non-zero value. (Equivalently, each group change contains **upto** *changesPerLot* singleton changes, each with value ± 1 .) Query the WAM oracle for the effect of each of the *lots* group changes. Perform the one which reduces the distance to μ most, discard the rest. If all the group changes increase the distance, perform none.
3. Whilst iterating, keep track of how many consecutive void iterations (ones which perform no change) have been done. If this number reaches 100, reduce *changesPerLot* by 25%, to a minimum of 1.

Increasing the parameter *lots* makes the algorithm less greedy and more cautious - at each iteration, more work is done and more queries are spent, in the hope that this will yield a better distance-reducing change. Increasing the parameter *changesPerLot* has the effect of making the algorithm conserve its queries by including more singleton changes in each group change and querying the oracle only once for their cumulative effect.

The motivation for step 3 is that if the algorithm is consistently unable to find *large* distance-reducing group changes, then perhaps there are very few of them. Perturbing many pixels of the image creates noise more often than removing it, so it is probably worth switching to a finer distance reduction tool, namely, smaller group changes, hence the adaptive behaviour in the step.

To obtain a good set of parameters, we partially benchmarked the algorithm with $(lots, changesPerLot) \in [1..4] \times [1..4]$ (see Appendix E). The general shape of the 16 curves is the same. They reveal two insights. Firstly, for a fixed parameter *lots*, better distance reduction is given by increasing *changesPerLot*. That is, it is beneficial to conserve queries by attempting to do more work on each perturbation. Secondly, for a fixed parameter *changesPerLot*, better distance reduction is given by reducing *lots*. Thus, spending many queries on each iteration in order to find a very good distance-reducing change is wasteful. The message is that greediness is better than caution.

Having seen how the parameters affect the distance reduction, we chose $(lots, changesPerLot) = (1, 25)$ as a suitable contender against the remaining algorithms. The results of its full benchmarking appear below. They look quite promising and improve on the Varsort algorithm.



5.6 Quadratic Programming Algorithm

The final approach we present for the problem of feature restoration is the least heuristic one. The idea is to use Quadratic Programming.

A Quadratic Problem (QP) is the problem of assigning values to a vector x of unknowns so that $x^T H x + f^T x$ is minimised, and x satisfies $lb \leq x \leq ub$, $Ax \leq b$, $A_1 x = b_1$, where H , A , A_1 are known matrices and f , lb , ub , b , b_1 are known vectors. An Integer Quadratic Problem (IQP) has the same form, with the additional restriction that x must be an integer vector. Finally, a Binary Integer Quadratic Problem (BIQP) is an IQP where the constraint $lb \leq x \leq ub$ is replaced by the restriction that x must be a binary vector.

We begin by showing how to express feature restoration as a BIQP. Throughout, we assume linearity of the features - the effect of performing many changes is the sum of the effects of performing each change on its own. Let the WAM features of the stego-image are $stego$. Suppose we have chosen a set of single-pixel changes C_1 to C_n , the set of all admissible single-pixel changes on the image. Change C_i alters the feature vector by δ_i . Let also x be a binary vector of length n , so that x_i specifies whether change C_i is performed or not. Then the features of the restored image are:

$$restored = stego + \sum_{i=1}^n x_i \delta_i = stego + Vx, \text{ where}$$

$V = [\delta_1 | \dots | \delta_n]$ is the $27 \times n$ matrix with columns δ_1 to δ_n . If S is the 27×27 covariance matrix of the WAM features, the Mahalanobis distance between the mean WAM vector μ and $restored$ is:

$$\begin{aligned} \|restored\|_M^2 &= \|stego + Vx\|_M^2 = \\ & (stego + Vx - \mu)^T S^{-1} (stego + Vx - \mu) = \\ & (stego - \mu)^T S^{-1} (stego - \mu) + (Vx)^T S^{-1} (Vx) - 2(stego - \mu)^T S^{-1} (Vx) = \\ & x^T H x + f^T x + c, \text{ where} \end{aligned}$$

$$H = V^T S^{-1} V, \quad f = 2V^T S^{-1} (stego - \mu), \quad \text{and} \quad c = (stego - \mu)^T S^{-1} (stego - \mu)$$

In feature restoration, the goal is to choose the values of x , subject to $x_i \in \{0, 1\}$, so that $\|restored\|_M$ is minimised. Additionally, we wish to specify that some changes should not be performed together. For example, if a pixel (x, y) may be perturbed by -2, -1, 1 or 2, our set of changes will include $C_i = (x, y, -2)$, $C_j = (x, y, -1)$, $C_k = (x, y, 1)$, $C_l = (x, y, 2)$, but at most one of these changes may be performed. This is easily specified with linear constraints: $x_i + x_j + x_k + x_l \leq 1$. The resulting problem is a BIQP.

It must be stressed that BIQP is very difficult to solve, completely infeasible for large-scale instances with lots of variables and linear constraints. Expressing feature restoration as a BIQP serves to give us a sense of its difficulty, and motivates focusing on heuristics and approximations, instead of ambitiously attempting to solve the problem optimally. We also reiterate that the problem is cast as a BIQP *assuming full linearity*. Without this simplifying assumption, it is even harder to solve optimally.

Nonetheless, we are now aware that reducing the distance to the mean feature vector is, in effect, minimising a quadratic objective. While it is infeasible to solve a BIQP obtained as above, a natural idea is to repeatedly obtain smaller instances of QPs, using fewer pixels and fewer restrictions. At each step, the feature restoration algorithm will produce a QP and give it to a solver, then recover the solution, and repeat. The difficulty lies in deciding how many of the restrictions of BIQP to relax. If we create too difficult quadratic programming problems, the solver will not be able to feasibly calculate solutions to them. On the other hand, if we relax too many constraints, we could be losing too much optimality.

The first constraint that we dispose of is integrality of x_i . A number of reasons motivate this decision. Firstly, although there are plenty of freely available Integer Programming packages which minimise a *linear* objective, it is very difficult to find ones capable of minimising a *quadratic* objective. In fact, we only found one package in the public domain which advertised this ability. However, we were disappointed to discover that it did not work correctly. On the other hand, quadratic programming *without* integer restrictions has been well-studied, and there are many, reasonably efficient QP solvers available. In particular, Matlab offers a routine *quadprog* which solves precisely QPs as defined at the start of this section. Therefore, we decided to relax the integrality constraints on the x_i , and replace them with $0 \leq x_i \leq 1$. When our feature restoration recovers the output from the QP solver, each x_i will be rounded.

We also decided to omit the linear constraints which specify which changes must not be performed together. This decision was arrived at purely by experimenting with Matlab's *quadprog* routine: adding linear constraints, even small ones such as $x_i + x_j \leq 1$, drastically increases the time the routine takes to terminate. Therefore, to build each QP, we will select a set of n pixels, then obtain $2n$ changes from them: $C_{2i} = (x_i, y_i, 1)$, $C_{2i+1} = (x_i, y_i, -1)$, but we will **not** add constraints $x_{2i} + x_{2i+1} \leq 1$. Thus, our algorithm is:

1. Repeat steps 2-8 for as long as queries are available:
2. Choose a value of n . (See below.)
3. Select n non-payload pixels, and the set of $2n$ changes C_1 to C_{2n} - one for perturbing each pixel by 1, and one for -1.
4. Use the WAM oracle to calculate δ_1 to δ_{2n} , and set them as the columns of a matrix V .
5. Obtain the current wam features of the image: *features*.
6. Create the QP: minimise $x^T H x + f^T x$, where $H = V^T S^{-1} V$ and $f = 2V^T S^{-1}(\text{features} - \mu)$.
7. Give the QP to a quadratic programming solver, wait for it to finish, and recover its output x .
8. For each $x_i \geq 0.5$, perform change C_i .

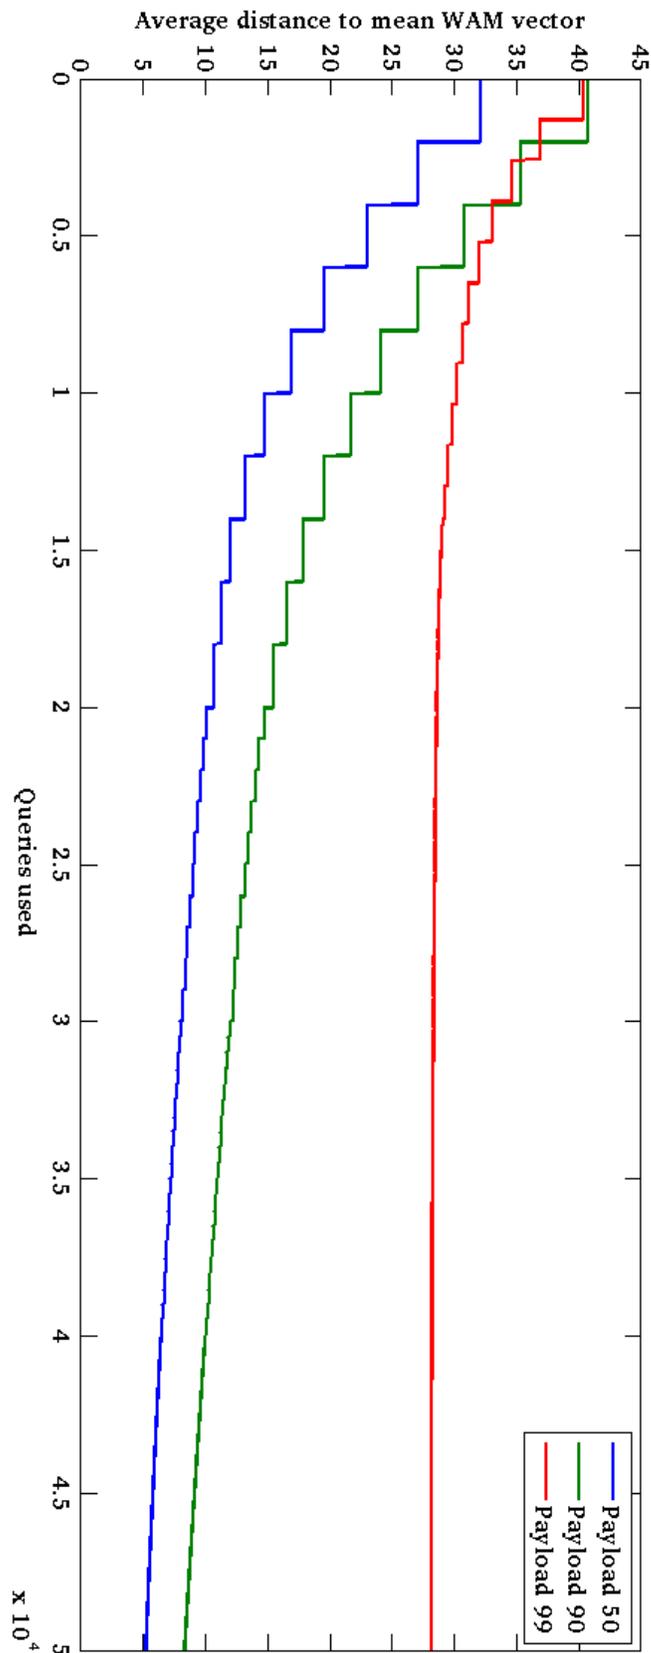
The only remaining question is how to specify n : the number of pixels used to create each QP. We experimented with some values. Past 1000, the solver needs too much time to solve each instance. Moreover, for a fixed value of n , we found that the QPs at the start of the algorithm are solved very quickly, whereas later they begin to take a disproportionately large amount of time. Therefore, we

used $n = 1000$ during the first 40% of the queries, then $n = 500$ for the next 20%, $n = 250$ for the next 20%, and $n = 25$ for the final 20% of the queries. Of course, if fewer pixels are available for the restoration, as is the case during the first iterations of the algorithm against payloads of 99%, we just use all of them.

We needed to be careful about measuring the performance of this algorithm, because it includes a time-expensive component (namely, the QP solving) which could potentially dominate the cost of querying the oracle. With the values we have chosen, the algorithm's timescale is similar to that of the Random algorithm with parameters (1, 25). Therefore, it is fair to benchmark the QP algorithm using queries as a unit of time, especially considering that for larger image sizes the cost of queries will dominate the cost of running the QP solver.

A programming challenge was to link the feature restoration code with Matlab's quadratic problem routine. To do so, we wrote a Matlab script which repeatedly checks for the presence of a file, signalling that it should read some input, solve a QP and output. After outputting the solution to the QP, it prints another flag file to signal the feature restoration to read Matlab's output and continue its work. Thus, we effectively used the presence of flag files as a semaphore in order to automate the feature restoration algorithm.

The chart for this algorithm appears below.



Performance of algorithm QP against payloads of 50%, 90% and 99%.
100 images were used, their distances averaged at each query count.

5.7 Comparison and Conclusion

In order to compare our feature restoration algorithms, we overlaid their performance charts for each payload percentage. This produced the **three final comparison charts** - one for 50% payload, one for 90% and one for 99%, all shown below.

At 50% payload, the two best are the Random algorithm with parameters (1, 25), and the Variance Sort algorithm. In 50000 queries, they reduce the distance to μ to a point corresponding to no payload at all. However, the Random algorithm's chart is steeper at the start, so it should be preferred if the query allowance is smaller - say, 10000 to 15000.

At 90% payload, the situation is identical. The Random and Variance Sort algorithms outclass the rest, achieving an apparent payload of 0%, with the Random algorithm doing so in fewer queries.


At 99% payload, the Random algorithm performs significantly better than the rest, achieving a greater distance reduction both in the short term and long term. In 50000 queries, it reaches an apparent payload of about 30%. This result is definitely more than we expected to achieve at the start of this project. The Variance Sort, Greedy and Genetic algorithms all achieve an apparent payload of about 40%. Surprisingly, the QP algorithm does not perform particularly well.


It must be stressed that the Random algorithm's good performance is due to the amount of work it does per query used. For each oracle query, it perturbs up to 25 pixels, whereas the other algorithms typically use queries on single-pixel changes. Initially, we suspected this behaviour might turn out to be very suboptimal. The results have proven us wrong; the project has discovered that the keys to efficient feature restoration are:


1. *Greediness*. When you spot a beneficial change, perform it, instead of being cautious and looking for even better ones.
2. *Conserving queries*. Attempt to maximise the ratio of changes performed per query used.

Also, our results show that smaller payloads like 50% are more difficult to feature restore than payloads of around 90%. Whilst perhaps counter-intuitive, this is because at 90% the search space of available perturbations is much smaller.

We believe this project was a successful investigation because it brought these principles to light. We implemented an efficient WAM oracle, devised a number of algorithms for feature restoration and obtained encouraging results from them, along with the above insights into the problem, setting the basis for further work.







5.8 Further Work

In order to fit this investigation into the timescale of a 3rd-year project, we had to exclude a number of important questions from its scope. Further work should focus on:

1. Devising feature restoration algorithms which conserve their queries similarly to the Random algorithm with a large second parameter.
2. Testing feature restoration against real steganalyzers to see if it truly makes the stego images less detectable.
3. Experimenting with various payload percentages in order to determine the best proportion of pixels to reserve for feature restoration. We suspect the optimal payload size is close to 90%.
4. Experimenting with various image sizes (this project only worked with 256×256), and image sets with different characteristics (source camera, use of compression, etc.).
5. Using feature restoration on different distance metrics and feature sets: we focused solely on WAM and Mahalanobis's distance, but our feature restoration algorithms treat them as black boxes, making them applicable to other sets of features and measures of distance.

Fortunately, we will be addressing these questions over the summer of 2010 under an EPSRC project titled "Large-Scale Benchmarking of Machine Learning Steganalysis", and hope to answer some of them by September.

Bibliography

- [1] Kodovsky, J., and Fridrich, J., “On completeness of feature spaces in blind steganalysis,” in *[Proceedings of the 10th ACM Multimedia & Security Workshop]* (2008)
- [2] Goljan, M., Fridrich, J., and Holotyak, T., “New blind steganalysis and its implications,” in *[Security, Steganography and Watermarking of Multimedia Contents VIII]*, *Proc. SPIE* **6072**, 0101–0113 (2006).
- [3] Ker, A., and Lubenko, I., “Feature reduction and payload location with WAM steganalysis,” in *[Electronic Imaging, Media Forensics and Security XI]*, *Proc. SPIE* **6072**, 0A01–0A13 (2009).
- [4] Avcıbaşı, I., Memon, N., and Sankur, B., “Steganalysis using image quality metrics,” in *[Electronic Imaging, Security and Watermarking of Multimedia Contents II]*, *Proc. SPIE* **4314**, 523–531, (2001).
- [5] Farid, H., and Lyu, S.,: “Detecting hidden messages using higher-order statistics and support vector machines,” in *[5th International Workshop on Information Hiding]*, *LNCS* **2578**, 340–354, (2002).
- [6] Lyu, S. and Farid, H., “Steganalysis using color wavelet statistics and one-class support vector machines,” in *[Security, Steganography, and Watermarking of Multimedia Contents VI]*, *Proc. SPIE* **5306**, 35–45 (2004).
- [7] Harmsen, J., and Pearlman, W., “Steganalysis of additive noise modelable information hiding,” in *[Proc. SPIE Electronic Imaging, Security, Steganography, and Watermarking of Multimedia Contents V]*, 131–142 (2003).
- [8] Ker, A., “Steganalysis of LSB matching in grayscale images,” in *[IEEE Signal Processing Letters]* **12**(6), 441–444 (2005).
- [9] Xuan, G., Shi, Y., Gao, J., Zou, D., Yang, C., Zhang, Z., Chai, P., Chen, C., and Chen, W., “Steganalysis based on multiple features formed by statistical moments of wavelet characteristic functions,” in *[Proc. 7th Information Hiding Workshop]*, *Springer LNCS* **3727**, 262–277 (2005).
- [10] Holotyak, T., Fridrich, J., and Voloshynovskiy, S., “Blind statistical steganalysis of additive steganography using wavelet higher order statistics,” in *[9th IFIP TC-6 TC-11 Conference on Communications and Multimedia Security]*, *Springer LNCS* **3677**, 273–274 (2005).

- [11] Mahalanobis, P., "On the generalised distance in statistics," in [*Proc. National Institute of Sciences of India*], Vol. **2**(1) (1936).
- [12] Mallat, S., "A theory for multiresolution signal decomposition: the wavelet representation," in [*IEEE Transactions on Pattern Analysis and Machine Intelligence Vol 2*] (1989).

Appendix A

Code for LSB-matching

```
1 // LSBEncoder.cpp
2 // A class which takes care of encoding.
3 #include "../perms.cpp"
4 #include "../bitconversion.cpp"
5 class LSBEncoder
6 {
7 public:
8 LSBEncoder() { }
9 // Public routine. It embeds a given message into an image specified by its name using a
10 // given key.
11 // The parameters also specify the name of the stego image and the name of the file to be
12 // used as a
13 // log of the encoding.
14 void messageIntoPgm(string message, string imageName, string newName, string logFile, int
15 key, matrixInt &usedPixels, matrixInt &changesMade)
16 {
17 matrixDouble image;
18 pgmToMatrix(imageName, image); // Read the image.
19 messageIntoMatrix(message, image, key, usedPixels, changesMade); // Put a message into
20 // it.
21 matrixToPgm(image, newName); // Write the image back.
22 // The log contains information about which pixels contain and the payload, and how
23 // much
24 // each pixel has been changed: -1 or +1 for payload pixels, 0 for non-payload.
25 FILE *fout = fopen(logFile.c_str(), "w");
26 FECHO(fout, "%s\n", imageName.c_str());
27 FECHO(fout, "%s\n", newName.c_str());
28 writeMatrix(fout, usedPixels);
29 writeMatrix(fout, changesMade);
30 fclose(fout);
31 }
32 private:
33 // Encoding procedure. Converts the message into a bit string, initialises the RNG with
34 // the given key
35 // and uses Knuth shuffle to generate a random permutation of the pixels of the image.
36 // Then it
37 // traverses the pixels in the order specified by the permutation, and embeds each bit of
38 // the message
39 // into a pixel as specified by the algorithm for LSB matching.
40 void messageIntoMatrix(string message, matrixDouble &image, int key, matrixInt &
41 usedPixels, matrixInt &changesMade)
42 {
43 int N, M;
44 int i, j;
```

```

36     vecBool bitMessage;
37     vecInt perm;
38     bitMessage = messageToBits(message); // Obtain a bitstring.
39     N = image.size();
40     M = image[0].size();
41     // Prepare matrices.
42     usedPixels.clear();
43     changesMade.clear();
44     usedPixels.resize(N);
45     changesMade.resize(N);
46     for(i = 0; i < N; i++) { usedPixels[i].resize(M); changesMade[i].resize(M); }
47     for(i = 0; i < N; i++)
48         for(j = 0; j < M; j++)
49             {
50                 usedPixels[i][j] = 0;
51                 changesMade[i][j] = 0;
52             }
53     srand(key); // Initialise the RNG with key.
54     randomPermutation(N*M, perm); // Get a random permutation.
55     // Do the encoding.
56     for(i = 0; i < bitMessage.size(); i++)
57         {
58             int pixel = perm[i];
59             int pixelX = pixel / M;
60             int pixelY = pixel % M;
61             int pixelValue = (int) image[pixelX][pixelY];
62             int bit = bitMessage[i];
63             int change = 0;
64             usedPixels[pixelX][pixelY] = 1;
65             if(LSB(pixelValue) == bit) continue;
66             else
67                 {
68                     if(pixelValue == 255) change = -1;
69                     else if(pixelValue == 0) change = 1;
70                     else if(rand() % 2 == 0) change = 1;
71                     else change = -1;
72                 }
73             image[pixelX][pixelY] = pixelValue + change;
74             changesMade[pixelX][pixelY] = change;
75         }
76     };
77 };
78 // LSBDecoder.cpp
79 // A class which performs decoding.
80 #include "../perms.cpp"
81 #include "../bitconversion.cpp"
82 class LSBDecoder
83 {
84 public:
85     LSBDecoder() { }
86     // A public routine with two versions. Given the name of a pgm file and a key, extract
87     // the message
88     // embedded into it using LSB matching. Optionally, a message length may be provided, if
89     // not all
90     // pixels hold payload.
91     string messageFromPgm(string fileName, int key, int expectedBitLength)
92     {
93         matrixDouble image;
94         pgmToMatrix(fileName, image);
95         return messageFromMatrix(image, key, expectedBitLength);
96     }
97     string messageFromPgm(string fileName, int key)
98     {
99         matrixDouble image;
100        pgmToMatrix(fileName, image);
101        return messageFromMatrix(image, key);

```

```

100     }
101     private:
102     // Decoding procedure. Initialises the RNG with key, generates a permutation of the
103     // pixels,
104     // traverses them in the order specified by the permutation, and collects the LSBs of the
105     // traversed pixels. Then it lumps them into groups of 8, thus getting a byte message.
106     // expectedLength is there to tell us when to stop the traversal. Obviously there will be
107     // pixels into which no information has been embedded, so there's no sense in looking
108     // at their LSBs. If this parameter is missing, then keep going until all pixels are
109     // traversed.
110     string messageFromMatrix(matrixDouble &image, int key, int expectedBitLength)
111     {
112         int N, M;
113         int i, j;
114         vecBool bitString;
115         vecInt perm;
116         N = image.size();
117         M = image[0].size();
118         if(expectedBitLength > N*M) expectedBitLength = N*M;
119         srand(key);
120         randomPermutation(N*M, perm);
121         for(i = 0; i < expectedBitLength; i++)
122         {
123             int pixel = perm[i];
124             int pixelX = pixel / M;
125             int pixelY = pixel % M;
126             int pixelValue = (int) image[pixelX][pixelY];
127             bitString.push_back(LSB(pixelValue));
128         }
129         return bitsToMessage(bitString);
130     }
131     string messageFromMatrix(matrixDouble &image, int key)
132     { return messageFromMatrix(image, key, image.size()*image[0].size()); }
133 };
134 // perms.cpp
135 // An invocation of an STL routine for obtaining a random permutation.
136 void randomPermutation(int N, vecInt &ans)
137 {
138     int i, j;
139     ans.clear();
140     ans.resize(N);
141     for(i = 0; i < N; i++) ans[i] = i;
142     random_shuffle(ans.begin(), ans.end());
143 }
144 // bitconversion.cpp
145 // Two routines to convert between bit strings and byte strings.
146 // Take a bit-string and convert it to a regular string by grouping bits into
147 // groups of 8. <bits> is assumed to have a length which is divisible by 8.
148 string bitsToMessage(vecBool bits)
149 {
150     int letters = bits.size() / 8;
151     int i, j;
152     string ans = "";
153     for(i = 0; i < letters; i++)
154     {
155         int c = 0;
156         for(j = 0; j < 8; j++) c += bits[8*i+j]*(1 << (7-j));
157         c -= 128;
158         ans.append(1, (char)c);
159     }
160     return ans;
161 }
162 // Take a string and break each character into its bit representation.
163 vecBool messageToBits(string message)
164 {
165     int N = message.size();
166     vecBool ans(N);
167     for(int i = 0; i < N; i++)
168     {
169         int c = message[i];
170         for(int j = 0; j < 8; j++)
171             ans[8*i+j] = (c >> (7-j)) & 1;
172     }
173 }

```

```
164 int i, j;
165 for(i = 0; i < message.length(); i++)
166 {
167     int c = message[i];
168     c += 128;
169     vecBool tmp;
170     for(j = 0; j < 8; j++)
171     {
172         tmp.push_back(c % 2);
173         c /= 2;
174     }
175     // tmp is the reversed bit-representation of message[i]. Append it in reverse to ans.
176     for(j = 7; j >= 0; j--) ans.push_back(tmp[j]);
177 }
178 return ans;
179 }
```

Appendix B

Code for WAM

```
1 // DWTCalculation.cpp
2 // A class offering routines for the 2D DWT.
3
4 #define maxDim 2048
5 #define maxFilterLength 16
6
7 class DWTCalculation
8 {
9 /*
10 extractRow(MX, i) means x becomes row i of MX.
11 extractColumn(MX, j) means x becomes column j of MX.
12 putRow(MX, i) means row i of MX becomes x.
13 putColumn(MX, j) means column j of MX becomes x.
14 */
15 #define extractRow(MX, i) { for(j = 0; j < M; j++) x[j] = (MX)[(i)][j]; xLength = M; }
16 #define extractColumn(MX, j) { for(i = 0; i < N; i++) x[i] = (MX)[i][(j)]; xLength = N; }
17 #define putRow(v, MX, i) { for(j = 0; j < M; j++) (MX)[(i)][j] = (v)[j]; }
18 #define putColumn(v, MX, j) { for(i = 0; i < N; i++) (MX)[i][(j)] = (v)[i]; }
19
20 // These arrays are used for convenience, as temporary storage of the input and output
21 // of the filtering routine. x is the input to doFiltering(). yLow and yHigh are its
22 // output. lowFilter and highFilter are the two filters.
23 double yLow[maxDim], yHigh[maxDim];
24 double lowFilter[maxFilterLength], highFilter[maxFilterLength];
25 double x[maxDim];
26 int xLength;
27 int filterLength;
28
29 public:
30 DWTCalculation()
31 {
32 // Set the two filters to be the Daubechies 8-tap QMF.
33 int i;
34 vecDouble filter;
35 filter.resize(8);
36 filter[0] = 0.230377813309;
37 filter[1] = 0.714846570553;
38 filter[2] = 0.630880767930;
39 filter[3] = -0.027983769417;
40 filter[4] = -0.187034811719;
41 filter[5] = 0.030841381836;
42 filter[6] = 0.032883011667;
43 filter[7] = -0.010597401785;
44 filterLength = filter.size();
```

```

45   for(i = 0; i < filterLength; i++) lowFilter[i] = filter[filterLength-i-1];
46   for(i = 0; i < filterLength; i++) highFilter[i] = filter[i];
47   for(i = 0; i < filterLength; i += 2) highFilter[i] = -highFilter[i];
48   }
49
50   // DWT procedure. Its input is the image. The filter is fixed to be 8-tap Daubechies.
51   // The matrices L and H store the intermediate results of the DWT calculation – the
52   // results of row-wise filtering. These are then filtered column-wise to
53   // produce the four output matrices LL, LH, HL and HH. The procedure's signature is
54   // STL-oriented, in sync with the rest of my code. However, the inner workings favour
55   // arrays over vectors. This is intentional and aims to improve efficiency, as array
56   // access is somewhat faster than vector access.
57   public:
58   void oneLevelOf2DDWT(matrixDouble &image, matrixDouble &L, matrixDouble &H, matrixDouble
      &LL, matrixDouble &LH, matrixDouble &HL, matrixDouble &HH, bool downsample)
59   {
60     register int i, j;
61
62     int N, M; // Dimensions of the image.
63     N = image.size();
64     M = image[0].size();
65     // Specify dimensions of LL, LH, HL, HH to be NxM. Later we'll downsample,
66     // reducing each dimension by half.
67     resizeMatrix(HH, N, M)
68     resizeMatrix(HL, N, M)
69     resizeMatrix(LH, N, M)
70     resizeMatrix(LL, N, M)
71     resizeMatrix(L, N, M)
72     resizeMatrix(H, N, M)
73     // Filter each row. Store the low-frequency output in LL, and the high-frequency
      output in HH.
74     for(i = 0; i < N; i++)
75     {
76       extractRow(image, i)
77       doFiltering();
78       putRow(yLow, L, i)
79       putRow(yHigh, H, i)
80     }
81     for(j = 0; j < M; j++)
82     {
83       extractColumn(L, j)
84       doFiltering();
85       putColumn(yLow, LL, j)
86       putColumn(yHigh, LH, j)
87       extractColumn(H, j)
88       doFiltering();
89       putColumn(yLow, HL, j)
90       putColumn(yHigh, HH, j)
91     }
92     if(downsample)
93     {
94       // Now downsample rows by discarding every other.
95       N /= 2;
96       for(i = 0; i < N; i++)
97       {
98         LL[i] = LL[i+i];
99         LH[i] = LH[i+i];
100        HL[i] = HL[i+i];
101        HH[i] = HH[i+i];
102      }
103      // Now downsample columns by discarding every other.
104      M /= 2;
105      for(i = 0; i < N; i++)
106        for(j = 0; j < M; j++)
107          {
108            LL[i][j] = LL[i][j+j];

```

```

109         LH[i][j] = LH[i][j+j];
110         HL[i][j] = HL[i][j+j];
111         HH[i][j] = HH[i][j+j];
112     }
113     // Now set dimensions of LL, LH, HL and HH to be (M/2)x(N/2).
114     resizeMatrix(HH, N, M)
115     resizeMatrix(HL, N, M)
116     resizeMatrix(LH, N, M)
117     resizeMatrix(LL, N, M)
118 }
119 }
120
121 // A lazy recalculation of the DWT. If the image is changed at (X, Y), and L, H, LL,
122 // LH, HL, HH are the DWT results prior to the change, do the least amount of work
123 // to update the results to reflect the changed image.
124 public:
125 void recalculateDWT(
126     matrixDouble &image,
127     matrixDouble &L,
128     matrixDouble &H,
129     matrixDouble &LL,
130     matrixDouble &LH,
131     matrixDouble &HL,
132     matrixDouble &HH,
133     int X, int Y)
134 {
135     int c, i, j;
136     int N, M;
137     N = image.size();
138     M = image[0].size();
139     extractRow(image, X)
140     doFiltering();
141     putRow(yLow, L, X)
142     putRow(yHigh, H, X)
143     for(j = Y - 7; j <= Y; j++)
144     {
145         if(j >= 0) c = j; else c = j + M;
146         extractColumn(L, c)
147         doFiltering();
148         putColumn(yLow, LL, c)
149         putColumn(yHigh, LH, c)
150         extractColumn(H, c)
151         doFiltering();
152         putColumn(yLow, HL, c)
153         putColumn(yHigh, HH, c)
154     }
155 }
156 // The filtering routine.
157 private:
158 void doFiltering()
159 {
160     register int i, j, k;
161     // Convolve.
162     double sum0, sum1;
163     for(i = 0; i < xLength; i++)
164     {
165         sum0 = sum1 = 0;
166         for(j = 0; j < filterLength; j++)
167         {
168             k = i + j; if(k >= xLength) k -= xLength; // Note the periodicity extension.
169             sum0 += lowFilter[filterLength - j - 1]*x[k];
170             sum1 += highFilter[filterLength - j - 1]*x[k];
171         }
172         yLow[i] = sum0;
173         yHigh[i] = sum1;
174     }

```

```

175     }
176   };
177   #undef putRow
178   #undef putColumn
179   #undef extractRow
180   #undef extractColumn
181   #undef maxDim
182   #undef maxFilterLength
183
184
185   // WAMCalculation.cpp.
186   // A stateless class capturing the calculation of the WAM features.
187
188   class WAMCalculation
189   {
190   private:
191     matrixDouble prefixedSquares;
192
193   public:
194     WAMCalculation() { }
195
196     vecDouble calculateWAM(
197       matrixDouble &image,
198       matrixDouble &L,
199       matrixDouble &H,
200       matrixDouble &LL,
201       matrixDouble &LH,
202       matrixDouble &HL,
203       matrixDouble &HH,
204       matrixDouble &variancesH,
205       matrixDouble &variancesV,
206       matrixDouble &variancesD,
207       matrixDouble &residualsH,
208       matrixDouble &residualsV,
209       matrixDouble &residualsD,
210       vecDouble &momentsH,
211       vecDouble &momentsV,
212       vecDouble &momentsD
213     )
214     {
215       vecDouble features;
216
217       // Dimensions of the image.
218       int N = image.size();
219       int M = image[0].size();
220
221       // Calculate the DWT.
222       DWTCalculation DWT = DWTCalculation();
223       DWT.oneLevelOf2DDWT(image, L, H, LL, LH, HL, HH, false);
224
225       // Calculate the 9 moments from H.
226       computeVariances(LH, variancesH);
227       computeResiduals(LH, variancesH, residualsH);
228       computeMoments(residualsH, momentsH);
229
230       // Calculate the 9 moments from V.
231       computeVariances(HL, variancesV);
232       computeResiduals(HL, variancesV, residualsV);
233       computeMoments(residualsV, momentsV);
234
235       // Calculate the 9 moments from D.
236       computeVariances(HH, variancesD);
237       computeResiduals(HH, variancesD, residualsD);
238       computeMoments(residualsD, momentsD);
239
240       // Assemble the 27-d feature vector.

```

```

241 features.insert(features.end(), momentsH.begin(), momentsH.end());
242 features.insert(features.end(), momentsV.begin(), momentsV.end());
243 features.insert(features.end(), momentsD.begin(), momentsD.end());
244
245 // Return it.
246 return features;
247 }
248
249 // The calculation of local variances. Given a matrix of dwt coefficients,
250 // estimate the local variance of the submatrix (fromX, fromY) - (toX, toY).
251 void computeVariances(matrixDouble &dwt, matrixDouble &var, int fromX, int toX, int fromY
    , int toY)
252 {
253     int N = dwt.size();
254     int M = dwt[0].size();
255     int i, j, r, c;
256     for(i = fromX; i < toX; i++)
257         for(j = fromY; j < toY; j++)
258             {
259                 r = (i + N) % N;
260                 c = (j + M) % M;
261                 var[r][c] = computeOneVariance(r, c);
262             }
263 }
264 // Routine to compute the variances of the whole dwt matrix.
265 void computeVariances(matrixDouble &dwt, matrixDouble &var)
266 {
267     int N = dwt.size();
268     int M = dwt[0].size();
269     resizeMatrix(var, N, M);
270     transform(dwt);
271     computeVariances(dwt, var, 0, N, 0, M);
272 }
273
274 // Each coefficient's local variance is estimated separately, based on four
275 // windows centred around it - of sizes 3, 5, 7 and 9.
276 double computeOneVariance(int x, int y)
277 {
278     double v[6] = {0, 0, 0, 0, 0};
279     int d;
280     int X, Y, F, G;
281     int N = prefixedSquares.size();
282     int M = prefixedSquares[0].size();
283     v[1] = v[2] = v[3] = v[4] = 0;
284
285
286     for(d = 1; d <= 4; d++) // For each window size 2*d+1:
287     {
288         // The window is (X, Y) - (F, G).
289         X = x - d;
290         Y = y - d;
291         F = x + d;
292         G = y + d;
293
294         // Crop window at edges.
295         if(X < 0) X = 0;
296         if(Y < 0) Y = 0;
297         if(F >= N) F = N-1;
298         if(G >= M) G = M-1;
299
300         // Use prefixedSquares to get the numerator.
301         if(!X && !Y) v[d] = prefixedSquares[F][G];
302         else if(!X) v[d] = prefixedSquares[F][G] - prefixedSquares[F][Y-1];
303         else if(!Y) v[d] = prefixedSquares[F][G] - prefixedSquares[X-1][G];
304         else v[d] = prefixedSquares[F][G] - prefixedSquares[F][Y-1] - prefixedSquares[X-1][G]
            + prefixedSquares[X-1][Y-1];

```

```

305
306     #ifdef WAM_fix_denominator
307         // Use a fixed denominator. We used this.
308         v[d] /= (2*d + 1)*(2*d + 1);
309     #else
310         // Use a non-fixed denominator.
311         v[d] /= (F-X+1)*(G-Y+1);
312     #endif
313     }
314
315     // Find the least v[d].
316     double best = v[1];
317     for(d = 2; d <= 4; d++) best = MIN(best, v[d]);
318
319     // Subtract a half, clamp to 0.
320     best -= 0.5;
321     if(best < 0) best = 0;
322
323     // Return.
324     return best;
325 }
326
327 // The calculation of all residuals, given the dwt and variance matrices. For a
328 // coefficient x with local variance estimate v, the residual is 0.5*x / (0.5+v).
329 void computeResiduals(matrixDouble &dwt, matrixDouble &var, matrixDouble &res)
330 {
331     int N = dwt.size();
332     int M = dwt[0].size();
333     resizeMatrix(res, N, M);
334     computeResiduals(dwt, var, res, 0, N, 0, M);
335 }
336
337 // The calculation of only some residuals.
338 // Those appearing in the window (fromX, fromY) - (toX, toY).
339 void computeResiduals(matrixDouble &dwt, matrixDouble &var, matrixDouble &res, int fromX,
340     int toX, int fromY, int toY)
341 {
342     int i, j, r, c;
343     int N = dwt.size();
344     int M = dwt[0].size();
345     for(i = fromX; i < toX; i++)
346         for(j = fromY; j < toY; j++)
347             {
348                 r = (i + N) % N;
349                 c = (j + M) % M;
350                 res[r][c] = 0.5*dwt[r][c] / (0.5 + var[r][c]);
351             }
352 }
353 // The calculation of moments, given the residuals. The m-th moment is defined by
354 // (sum(i, j) of |residuals[i][j] - average|^m) / (M*N), where average is the
355 // mean value of the matrix residuals. We're interested in the first nine moments.
356 void computeMoments(matrixDouble &res, vecDouble &moments)
357 {
358     int N = res.size();
359     int M = res[0].size();
360     int i, j, m;
361     double average = 0;
362
363     // Initialise.
364     moments.resize(9);
365     for(m = 0; m < 9; m++) moments[m] = 0;
366
367     // Find average residual.
368     for(i = 0; i < N; i++)
369         for(j = 0; j < M; j++)

```

```

370     average += res[i][j];
371     average /= (N*M);
372
373     // Do the moments.
374     for(i = 0; i < N; i++)
375         for(j = 0; j < M; j++)
376             {
377                 double R = res[i][j] - average;
378                 R = ABS(R);
379
380                 double v = R;
381                 for(m = 0; m < 9; m++)
382                     {
383                         moments[m] += v;
384                         v *= R;
385                     }
386             }
387     for(m = 0; m < 9; m++) moments[m] /= (M*N);
388 }
389
390 // Routine which computes the prefixed squares form of mx and
391 // puts it into prefixedSquares. Useful for variances.
392 void transform(matrixDouble &mx)
393     {
394         int i, j;
395         int N = mx.size();
396         int M = mx[0].size();
397         resizeMatrix(prefixedSquares, N, M);
398
399         for(i = 0; i < N; i++)
400             for(j = 0; j < M; j++)
401                 prefixedSquares[i][j] = mx[i][j]*mx[i][j];
402
403         for(i = 0; i < N; i++)
404             for(j = 1; j < M; j++)
405                 prefixedSquares[i][j] += prefixedSquares[i][j-1];
406
407         for(j = 0; j < M; j++)
408             for(i = 1; i < N; i++)
409                 prefixedSquares[i][j] += prefixedSquares[i-1][j];
410     }
411 };
412
413
414 // WAMUpdater.cpp
415 // This is the most important class of the WAM pipeline. It holds an image, and
416 // all the information from its WAM calculation. It provides an interface to modify
417 // the image whilst keeping the WAM information in sync with it. It also offers
418 // an interface for hypothetical queries: "I don't want to change this image
419 // in this way yet, but if I did, what would the effect be?"
420
421 class WAMUpdater
422 {
423
424 private:
425
426     matrixDouble image; // The image.
427     matrixDouble L, H, LL, LH, HL, HH; // The results from the DWT calculation.
428     matrixDouble variancesH, variancesV, variancesD; // The variance matrices.
429     matrixDouble residualsH, residualsV, residualsD; // The residuals matrices.
430     vecDouble momentsH, momentsV, momentsD; // The three sets of moments.
431     int queriesUsed; // A counter of many queries have been answered so far.
432
433     // References to the two above classes.
434     WAMCalculation W;
435     DWTCalculation DWT;

```

```

436
437 public:
438 // Obvious access methods.
439 int getN() { return image.size(); }
440 int getM() { return image[0].size(); }
441 int getQueriesUsed() { return queriesUsed; }
442 int getPixelValueAt(int x, int y) { return (int) image[x][y]; }
443
444 // To get the WAM vector, concatenate the 3 sets of moments.
445 vecDouble getWAM()
446 {
447     vecDouble features;
448     features.insert(features.end(), momentsH.begin(), momentsH.end());
449     features.insert(features.end(), momentsV.begin(), momentsV.end());
450     features.insert(features.end(), momentsD.begin(), momentsD.end());
451     return features;
452 }
453
454 // Null constructor.
455 WAMUpdater() { }
456
457 /*
458 Initialisation. Given a file containing an image,
459 do a first-time WAM calculation to initialise all
460 the info.
461 */
462 WAMUpdater(string imageFile)
463 {
464     // Read in.
465     pgmToMatrix(imageFile, image);
466     int N, M;
467     N = image.size();
468     M = image[0].size();
469
470     // Initialise query counter.
471     queriesUsed = 0;
472
473     // Set up sizes.
474     resizeMatrix(variancesH, N, M)
475     resizeMatrix(variancesV, N, M)
476     resizeMatrix(variancesD, N, M)
477     resizeMatrix(residualsH, N, M)
478     resizeMatrix(residualsV, N, M)
479     resizeMatrix(residualsD, N, M)
480     momentsH.resize(9);
481     momentsV.resize(9);
482     momentsD.resize(9);
483
484     // Calculate everything.
485     W.calculateWAM(
486     image,
487     L, H, LL, LH, HL, HH,
488     variancesH,
489     variancesV,
490     variancesD,
491     residualsH,
492     residualsV,
493     residualsD,
494     momentsH,
495     momentsV,
496     momentsD);
497 }
498
499 // Check method:
500 // Can I perturb pixel (x, y) by v?
501 bool canMakeOneChange(int x, int y, int v)

```

```

502     {
503         return (x >= 0 &&
504             x < image.size() &&
505             y >= 0 &&
506             y < image[0].size() &&
507             (int)image[x][y] + v >= 0 &&
508             (int)image[x][y] + v <= 255);
509     }
510 bool canMakeOneChange(Change c)
511 { return canMakeOneChange(c.x, c.y, c.v); }
512
513 /*
514  Check method for group changes.
515  We're relying on no duplicate pixels
516  appearing in the group change. In the rest
517  of our code, we arrange this to be so.
518  */
519 bool canMakeGroupChange(vector<Change> v)
520     {
521         for(int i = 0; i < v.size(); i++) if(!canMakeOneChange(v[i])) return false;
522         return true;
523     }
524
525 // Query routine.
526 vecDouble QueryOneChange(Change C)
527     {
528         // Count the query.
529         queriesUsed++;
530         // Make the change.
531         makeOneChange(C);
532         // Remember the features.
533         vecDouble features = getWAM();
534         // Undo the change.
535         undoOneChange(C);
536         // Return.
537         return features;
538     }
539
540 // Same but for group queries.
541 vecDouble QueryGroupChange(vector<Change> &v)
542     {
543         queriesUsed++;
544         makeGroupChange(v);
545         vecDouble features = getWAM();
546         undoGroupChange(v);
547         return features;
548     }
549
550 /*
551  The interface to change the image.
552  Given a singleton change, cast it
553  to a group change and apply it.
554  */
555 void makeOneChange(Change C)
556     {
557         vector<Change> v;
558         v.push_back(C);
559         makeGroupChange(v);
560     }
561
562 // Routine to apply a group change.
563 void makeGroupChange(vector<Change> &v)
564     {
565         int i;
566         // Firstly change the image.
567         for(i = 0; i < v.size(); i++)

```

```

568     {
569     image[v[i].x][v[i].y] += v[i].v;
570     image[v[i].x][v[i].y] = MIN(image[v[i].x][v[i].y], 255);
571     image[v[i].x][v[i].y] = MAX(image[v[i].x][v[i].y], 0);
572     // Recalculate the DWT stuff around each changed pixel.
573     if(v[i].v != 0)
574         recalculateDWT(v[i].x, v[i].y);
575     }
576     // Recalculate the variances and residuals around the changes v.
577     recalculateVariances(v);
578     recalculateResiduals(v);
579     // Recalculate the moments.
580     recalculateMoments();
581     }
582
583     // Mechanism for undoing changes.
584     // Yes, I'm aware that it's
585     // code duplication.
586     void undoOneChange(Change C)
587     {
588     vector<Change> v;
589     v.push_back(C);
590     undoGroupChange(v);
591     }
592
593     void undoGroupChange(vector<Change> &v)
594     {
595     int i;
596     for(i = 0; i < v.size(); i++)
597     {
598     image[v[i].x][v[i].y] -= v[i].v;
599     image[v[i].x][v[i].y] = MIN(image[v[i].x][v[i].y], 255);
600     image[v[i].x][v[i].y] = MAX(image[v[i].x][v[i].y], 0);
601     recalculateDWT(v[i].x, v[i].y);
602     }
603     recalculateVariances(v);
604     recalculateResiduals(v);
605     recalculateMoments();
606     }
607
608     private:
609     // To recalculate the DWT stuff around (x, y),
610     // just invoke the relevant method in the DWT class.
611     void recalculateDWT(int x, int y)
612     {
613     DWT.recalculateDWT(image, L, H, LL, LH, HL, HH, x, y);
614     }
615
616     // Direct recalculation of moments.
617     void recalculateMoments()
618     {
619     W.computeMoments(residualsH, momentsH);
620     W.computeMoments(residualsV, momentsV);
621     W.computeMoments(residualsD, momentsD);
622     }
623
624     // Recalculation of variances around group change v.
625     void recalculateVariances(vector<Change> &v)
626     {
627     int i, x, y;
628     // If big group, recalculate variances using prefixedSquares.
629     if(v.size() > 2) recalculateVariancesByTransform(v);
630     // If small, just recalculate the dirty variances by walking over the 9x9 region around
        them.
631     else recalculateVariancesFromDefinition(v);
632     }

```



```

699 // Divide by a fixed denominator.
700 v[1] /= 9;
701 v[2] /= 25;
702 v[3] /= 49;
703 v[4] /= 81;
704 #else
705 // Or a variable one. We used fixed.
706 for(d = 1; d <= 4; d++) v[d] /= den[d];
707 #endif
708 // Find the smallest.
709 double best, best1, best2;
710 best1 = MIN(v[1], v[2]);
711 best2 = MIN(v[3], v[4]);
712 // Subtract a half and clamp to zero.
713 best = MIN(best1, best2) - 0.5;
714 if(best < 0) best = 0;
715 // Return.
716 return best;
717 }
718 #undef OK
719
720 // If there are too many variances to update,
721 void recalculateVariancesByTransform(vector<Change> v)
722 {
723     int i, x, y;
724     // Calculate the prefixed squares form.
725     W.transform(LH);
726     // Then update the dirty variances.
727     for(i = 0; i < v.size(); i++)
728     {
729         if(v[i].v == 0) continue;
730         x = v[i].x; y = v[i].y;
731         W.computeVariances(LH, variancesH, x - 11, x + 5, y - 11, y + 5);
732     }
733     // Same.
734     W.transform(HL);
735     for(i = 0; i < v.size(); i++)
736     {
737         if(v[i].v == 0) continue;
738         x = v[i].x; y = v[i].y;
739         W.computeVariances(HL, variancesV, x - 11, x + 5, y - 11, y + 5);
740     }
741     // Same.
742     W.transform(HH);
743     for(i = 0; i < v.size(); i++)
744     {
745         if(v[i].v == 0) continue;
746         x = v[i].x; y = v[i].y;
747         W.computeVariances(HH, variancesD, x - 11, x + 5, y - 11, y + 5);
748     }
749 }
750
751 // To recalculate the residuals around a group change v,
752 void recalculateResiduals(vector<Change> &v)
753 {
754     int i, x, y;
755     for(i = 0; i < v.size(); i++)
756     {
757         if(v[i].v == 0) continue;
758         // Take each x, y in v.
759         x = v[i].x;
760         y = v[i].y;
761         // And redo the residuals in [x-11..x+4][y-11..y+4].
762         W.computeResiduals(LH, variancesH, residualsH, x - 11, x + 5, y - 11, y + 5);
763         W.computeResiduals(HL, variancesV, residualsV, x - 11, x + 5, y - 11, y + 5);
764         W.computeResiduals(HH, variancesD, residualsD, x - 11, x + 5, y - 11, y + 5);

```

```
765     }
766   }
767
768   // Output methods.
769   public:
770   void outputImage(string fileName)
771     { matrixToPgm(image, fileName); }
772
773   void outputWAM(string fileName)
774     {
775     vecDouble features = getWAM();
776     writeVector(fileName, features);
777     }
778
779   void outputAll(string fileName)
780     {
781     writeMatrix(fileName+".dwt.H", LH);
782     writeMatrix(fileName+".dwt.V", HL);
783     writeMatrix(fileName+".dwt.D", HH);
784
785     writeMatrix(fileName+".var.H", variancesH);
786     writeMatrix(fileName+".var.V", variancesV);
787     writeMatrix(fileName+".var.D", variancesD);
788
789     writeMatrix(fileName+".res.H", residualsH);
790     writeMatrix(fileName+".res.V", residualsV);
791     writeMatrix(fileName+".res.D", residualsD);
792
793     writeVector(fileName+".mom.H", momentsH);
794     writeVector(fileName+".mom.V", momentsV);
795     writeVector(fileName+".mom.D", momentsD);
796     }
797   };
```

Appendix C

Code for Feature Restoration

```
1 // Restorer.cpp
2 // A superclass for our restorers. It promises the apply() routine, which is implemented
   by subclasses.
3
4 class Restorer
5 {
6     protected:
7         WAMUpdater W;
8         vecDouble goal;
9         matrixInt usedPixels;
10        matrixInt changesMade;
11
12
13        string originalFile;
14        string stegoFile;
15
16        // A list of all the non-payload pixels available for restoration.
17        vector<Pixel> availablePixels;
18
19    public:
20
21        virtual void apply() = 0;
22
23        Restorer(string encodingLogFile)
24        {
25            char x[128];
26            int i, j;
27            FILE *fin = fopen(encodingLogFile.c_str(), "rb");
28
29            fscanf(fin, "%s", x);
30            originalFile = string(x);
31
32            fscanf(fin, "%s", x);
33            stegoFile = string(x);
34
35
36            loadMatrix(fin, usedPixels);
37            loadMatrix(fin, changesMade);
38
39            W = WAMUpdater(stegoFile);
40            fclose(fin);
41
42            prepareMahalanobis();
43            setGoal("means");
```

```

44
45 int N, M; N = W.getN(); M = W.getM();
46 // Find all the non-payload pixels, and put them into a vector. This will allow quick
    // generation of valid group changes.
47
48 for(i = 0; i < N; i++)
49     for(j = 0; j < M; j++)
50         if(!usedPixels[i][j])
51             availablePixels.push_back(newPixel(i, j));
52 availablePixels.resize(availablePixels.size());
53 }
54
55
56 /*
57 Routine to set the goal. The choice is between aiming for the original features, and
    // aiming for the mean features. We are going to use the mean features.
58 */
59 void setGoal(string instruction)
60 {
61     if(instruction == "means")
62         { goal = means; }
63     else if(instruction == "original")
64         {
65             WAMUpdater F = WAMUpdater(originalFile);
66             goal = F.getWAM();
67         }
68     else
69         { exit(0); }
70 }
71
72 // RANGE: [a..b)
73 int randomNumberInRange(int a, int b)
74 { return a + (rand() % (b-a)); }
75
76 protected:
77 // Returns a combination of k elements, drawn from [0..m), where m is the number of
    // available pixels.
78
79 vecInt randomSample(int k)
80 {
81     int i, ind;
82     bool hasDups;
83     int m = availablePixels.size();
84
85     vecInt combination;
86     combination.resize(k);
87
88     do
89     {
90         for(i = 0; i < k; i++)
91             {
92                 ind = randomNumberInRange(0, m);
93                 combination[i] = ind;
94             }
95         sort(combination.begin(), combination.end());
96         hasDups = false;
97         for(i = 1; i < k && !hasDups; i++)
98             if(combination[i] == combination[i-1]) hasDups = true;
99         } while(hasDups);
100
101
102 return combination;
103 }
104
105
106 vector<Change> randomNodupGroupChange(int size, int vmin, int vmax)

```

```

107 {
108   int i, ind, v, x, y;
109   bool hasDups;
110   vector<Change> groupChange;
111   vecInt indices;
112
113   groupChange.resize(size);
114   indices.resize(size);
115
116   vmax += 1;
117
118   do
119   {
120     indices = randomSample(size);
121     for(i = 0; i < size; i++)
122     {
123       ind = indices[i];
124       x = availablePixels[ind].x;
125       y = availablePixels[ind].y;
126       do
127       { v = randomNumberInRange(vmin, vmax); }
128       while(ABS(changesMade[x][y] + v) > PIXEL_CHANGE_CAP);
129
130       groupChange[i] = newChange(x, y, v);
131     }
132   } while (!W.canMakeGroupChange(groupChange));
133
134   return groupChange;
135 }
136
137 };
138
139 // RestorerIneff.cpp
140 // The restorer for the inefficient algorithm.
141
142 class RestorerIneff : public Restorer
143 {
144 private:
145   int iterations;
146
147 public:
148   RestorerIneff(string encodingLog, int ITERATIONS) : Restorer(encodingLog)
149   {
150     iterations = ITERATIONS;
151   }
152
153
154   void apply()
155   {
156     int N, M;
157     int it;
158     int i, j, v;
159     Change C, bestChange;
160     double dist, bestdist;
161
162     N = W.getN();
163     M = W.getM();
164
165     vecDouble cur = W.getWAM();
166     bestdist = Distance(goal, cur);
167     for(it = 1; it <= iterations; it++)
168     {
169       ECHO("%d_%.121f\n", W.getQueriesUsed(), bestdist);
170
171       bestChange = newChange(-1, -1, -1);
172       for(i = 0; i < N; i++)

```

```

173     for(j = 0; j < M; j++)
174         if(!usedPixels[i][j])
175             for(v = -1; v <= 1; v += 2) // v in {-1, 1}
176                 {
177                     C = newChange(i, j, v);
178
179                     if(W.canMakeOneChange(C) && ABS(changesMade[i][j] + v) <= PIXEL_CHANGE_CAP)
180                         {
181                             vecDouble w = W.QueryOneChange(C);
182                             dist = Distance(w, goal);
183                             if(dist < bestdist)
184                                 { bestChange = C; bestdist = dist; }
185                         }
186                 }
187     if(bestChange.x == -1) break;
188     else
189     {
190         W.makeOneChange(bestChange);
191         changesMade[bestChange.x][bestChange.y] += bestChange.v;
192     }
193
194     cur = W.getWAM();
195     bestdist = Distance(goal, cur);
196 }
197 ECHO("%d_%.12lf\n", W.getQueriesUsed(), bestdist);
198 W.outputImage("restored.pgm");
199 }
200 };
201
202
203 // RestorerGreedy.cpp
204 // The restorer class for the Greedy algorithm.
205
206 class RestorerGreedy : public Restorer
207 {
208 private:
209     int queryLimit;
210
211 public:
212     RestorerGreedy(string encodingLog, int QUERIES) : Restorer(encodingLog)
213     {
214         queryLimit = QUERIES;
215     }
216
217 void apply()
218 {
219     int N, M; // Dimensions of the image.
220     int m; // Number of pixels available for restoration.
221
222     // Temporary variables used below.
223     int x, y, pixelValue, v;
224     double hypDistPlus, hypDistMinus, bestHypDist, curDist;
225     vecDouble hypWAMplus, hypWAMminus, bestHypWAM, curWAM;
226     int i, j;
227
228     // Dimensions.
229     N = W.getN();
230     M = W.getM();
231
232     // WAM features of the image and their distance to the goal features.
233     curWAM = W.getWAM();
234     curDist = Distance(curWAM, goal);
235
236     // Create a list of the pixels available for restoration.
237     vector<Pixel> availablePixels;
238     for(i = 0; i < N; i++)

```

```

239     for(j = 0; j < M; j++)
240         if(!usedPixels[i][j])
241             availablePixels.push_back(newPixel(i, j));
242     m = availablePixels.size();
243     availablePixels.resize(m);
244
245     // Iterate while there are queries available.
246     while(W.getQueriesUsed() < queryLimit)
247         for(i = 0; i < m && W.getQueriesUsed() < queryLimit; i++)
248             {
249                 ECHO("%.12lf_%d\n", curDist, W.getQueriesUsed());
250                 // For each non-payload pixel (x, y),
251                 x = availablePixels[i].x;
252                 y = availablePixels[i].y;
253                 pixelValue = W.getPixelValueAt(x, y);
254
255                 // Calculate the WAM of the image obtained by perturbing the pixel by 1.
256                 // Avoid possibility of incrementing pixel outside 8-bit range.
257                 if(pixelValue != 255 && ABS(changesMade[x][y] + 1) <= PIXEL_CHANGE_CAP)
258                     {
259                         hypWAMplus = W.QueryOneChange(newChange(x, y, 1));
260                         hypDistPlus = Distance(hypWAMplus, goal);
261                     }
262                 else hypDistPlus = 1000000;
263
264                 // Same for -1.
265                 // Calculate the hypothetical distances of the two possibilities.
266                 if(pixelValue != 0 && ABS(changesMade[x][y] - 1) <= PIXEL_CHANGE_CAP)
267                     {
268                         hypWAMminus = W.QueryOneChange(newChange(x, y, -1));
269                         hypDistMinus = Distance(hypWAMminus, goal);
270                     }
271                 else hypDistMinus = 1000000;
272
273                 // Determine which of these two changes leads to a greater distance reduction.
274                 if(hypDistPlus < hypDistMinus)
275                     {
276                         bestHypDist = hypDistPlus;
277                         bestHypWAM = hypWAMplus;
278                         v = 1;
279                     }
280                 else
281                     {
282                         bestHypDist = hypDistMinus;
283                         bestHypWAM = hypWAMminus;
284                         v = -1;
285                     }
286
287                 // Finally, check if this distance reduction is positive. If so, perform the change
288                 if(bestHypDist < curDist)
289                     {
290                         W.makeOneChange(newChange(x, y, v));
291                         changesMade[x][y] += v;
292                         curDist = bestHypDist;
293                         curWAM = bestHypWAM;
294                     }
295             }
296     ECHO("%.12lf_%d\n", curDist, W.getQueriesUsed());
297     W.outputImage("restored.pgm");
298 }
299 };
300
301 // RestorerGenetic.cpp
302 // The restorer for our Genetic algorithm
303

```

```

304 struct PopulationIndividual
305 {
306     GroupChange vc;
307     double distanceReduction;
308 };
309
310 bool operator <(PopulationIndividual a, PopulationIndividual b)
311 {
312     if(a.distanceReduction < b.distanceReduction) return true;
313     else return false;
314 }
315
316 PopulationIndividual newIndividual(GroupChange gc, double distRed)
317 {
318     PopulationIndividual ans;
319     ans.vc = gc;
320     ans.distanceReduction = distRed;
321     return ans;
322 }
323
324 typedef PopulationIndividual Individual;
325 typedef priority_queue<PopulationIndividual> Population;
326
327 class RestorerGenetic : public Restorer
328 {
329 private:
330     Population P; // The population for the genetic algorithm.
331     int initialSize; // Parameters of the genetic algorithm.
332     int queryLimit; // What it says on the tin.
333
334 public:
335     // Constructor.
336     RestorerGenetic(string encodingLogFile, int initSize, int queries) : Restorer(
337         encodingLogFile)
338     {
339         // Set state space of the genetic algorithm restorer.
340         initialSize = initSize;
341         queryLimit = queries;
342     }
343
344 void apply()
345 {
346     initialiseGeneticAlgorithm();
347     iterate();
348     dwindle();
349     W.outputImage("restored.pgm");
350 }
351
352 private:
353 void initialiseGeneticAlgorithm()
354 {
355     int i, j;
356     int N, M;
357     N = W.getN(); M = W.getM();
358
359     vecDouble curWAM;
360     double curDist;
361     curWAM = W.getWAM();
362     curDist = Distance(curWAM, goal);
363     ECHO("%.12lf_%d\n", curDist, W.getQueriesUsed());
364
365     for(i = 0; i < initialSize; i++)
366         migrateIn();
367 }
368
369 void iterate()

```

```

369 {
370     vecDouble curWAM;
371     double curDist;
372
373     curWAM = W.getWAM();
374     curDist = Distance(curWAM, goal);
375     ECHO("%.12lf_%d\n", curDist, W.getQueriesUsed());
376
377     while(1)
378     {
379         if(W.getQueriesUsed() >= queryLimit) break;
380         migrateIn();
381         if(W.getQueriesUsed() >= queryLimit) break;
382         migrateIn();
383         if(W.getQueriesUsed() >= queryLimit) break;
384
385         migrateOut();
386         curWAM = W.getWAM();
387         curDist = Distance(curWAM, goal);
388         ECHO("%.12lf_%d\n", curDist, W.getQueriesUsed());
389
390         if(P.size() == 2*initialSize + initialSize / 10) merge();
391     }
392 }
393
394 void migrateIn()
395 {
396     vector<Change> groupChange;
397     vecDouble hypWAM, curWAM;
398     double curDist, hypDist, distRed;
399
400     curWAM = W.getWAM();
401     curDist = Distance(curWAM, goal);
402     while(1)
403     {
404         if(W.getQueriesUsed() >= queryLimit) return;
405         do
406             { groupChange = randomNodupGroupChange(1, -1, 1); }
407         while(groupChange[0].v == 0);
408         hypWAM = W.QueryGroupChange(groupChange);
409         hypDist = Distance(hypWAM, goal);
410         distRed = curDist - hypDist;
411         if(distRed > 0.000001) break;
412     }
413     P.push(newIndividual(groupChange, distRed));
414 }
415
416 void migrateOut()
417 {
418     int i;
419     int x, y, v;
420     Individual best;
421
422     best = P.top(); P.pop(); // Extract the top element from P.
423
424     // Check that applying it does not violate the pixel change cap.
425     for(i = 0; i < best.vc.size(); i++)
426     {
427         x = best.vc[i].x;
428         y = best.vc[i].y;
429         v = best.vc[i].v;
430         // If it does, break early. If not, finish normally.
431         if(ABS(changesMade[x][y] + v) > PIXEL_CHANGE_CAP) break;
432     }
433     // If haven't finished normally, then applying best violates pixel change cap. Discard
         it.

```

```

434     if (i != best.vc.size()) return;
435
436     // Check that applying best will not take a pixel outside [0..255]. Also, that the
437     // reduction label is positive.
438     if (W.canMakeGroupChange(best.vc) && best.distanceReduction > 0)
439     {
440         // If so, make the change.
441         W.makeGroupChange(best.vc);
442         // Record that pixels have been touched.
443         for (i = 0; i < best.vc.size(); i++)
444             changesMade[best.vc[i].x][best.vc[i].y] += best.vc[i].v;
445     }
446 }
447
448 void merge()
449 {
450     Individual first, second;
451     Individual merged;
452     vecDouble hypWAM, curWAM;
453     double curDist, hypDist, distRed;
454     int i;
455
456     vector<Individual> mergeResults;
457     for (i = 0; i < initialSize; i++)
458     {
459         first = P.top(); P.pop();
460         second = P.top(); P.pop();
461         if (W.getQueriesUsed() < queryLimit)
462         {
463             merged.vc = mergeTwoGroupChanges(first.vc, second.vc);
464
465             hypWAM = W.QueryGroupChange(merged.vc);
466             curWAM = W.getWAM();
467             curDist = Distance(curWAM, goal);
468             hypDist = Distance(hypWAM, goal);
469             merged.distanceReduction = curDist - hypDist;
470             mergeResults.push_back(merged);
471         }
472         else
473         {
474             mergeResults.push_back(first);
475             mergeResults.push_back(second);
476         }
477     }
478
479     while (!P.empty()) P.pop();
480
481     for (i = 0; i < mergeResults.size(); i++)
482         P.push(mergeResults[i]);
483 }
484
485 void dwindle()
486 {
487     while (!P.empty())
488         migrateOut();
489
490     vecDouble curWAM = W.getWAM();
491     double curDist = Distance(curWAM, goal);
492     ECHO("%.12lf_%d\n", curDist, W.getQueriesUsed());
493 }
494
495 };
496
497
498 // RestorerRand.cpp

```

```

499 // The restorer for the Random algorithm.
500 class RestorerRand : public Restorer
501 {
502 private:
503 int LOTS, CHANGES;
504 int queryLimit;
505
506 public:
507 RestorerRand(string encodingLog, int lots, int changes, int queries) : Restorer(
    encodingLog)
508 {
509     LOTS = lots;
510     CHANGES = changes;
511     queryLimit = queries;
512 }
513
514 void apply()
515 {
516     int i, j;
517     int iterationsWithoutReduction = 0;
518     int N = W.getN();
519     int M = W.getM();
520
521     double hypDist, bestDist;
522     vector<Change> groupChange, bestGroupChange;
523     vecDouble curWAM, hypWAM;
524
525     while(W.getQueriesUsed() < queryLimit)
526     {
527         curWAM = W.getWAM();
528         bestDist = Distance(goal, curWAM);
529         bestGroupChange.clear();
530
531         ECHO("%.12lf_%.12lf\n", bestDist, W.getQueriesUsed());
532
533         if(iterationsWithoutReduction == 100)
534         {
535             iterationsWithoutReduction = 0;
536             CHANGES = MAX(1, (CHANGES*3)/4);
537         }
538
539         for(i = 0; i < LOTS; i++) // Randomly choose LOTS groups of pixels to change.
540         {
541             // Each group has CHANGES pixels in it. Each is perturbed by 1, 0 or -1.
542             // Effectively, the group has upto CHANGES perturbations, each of value 1 or -1.
543             groupChange = randomNodupGroupChange(CHANGES, -1, 1);
544
545             // For large values of CHANGES it is virtually impossible for the group change to
546             // zeroes. However, for a small value of CHANGES, it is worth checking. If so,
547             // reselect.
548             if(CHANGES <= 10)
549             {
550                 for(j = 0; j < groupChange.size(); j++)
551                     if(groupChange[j].v != 0)
552                         break;
553                 if(j == groupChange.size())
554                     { i--; continue; }
555             }
556
557             // Change is not void. Query the oracle for its effects.
558             // If better than best, record it.
559             hypWAM = W.QueryGroupChange(groupChange);
560             hypDist = Distance(goal, hypWAM);
561             if(hypDist < bestDist)

```

```

562         bestDist = hypDist;
563         bestGroupChange = groupChange;
564     }
565
566     }
567     if(bestGroupChange.empty())
568     { iterationsWithoutReduction++; }
569     else
570     {
571         iterationsWithoutReduction = 0;
572         W.makeGroupChange(bestGroupChange);
573         for(i = 0; i < bestGroupChange.size(); i++)
574             changesMade[bestGroupChange[i].x][bestGroupChange[i].y] += bestGroupChange[i].v;
575     }
576 }
577 curWAM = W.getWAM();
578 bestDist = Distance(goal, curWAM);
579 ECHO("%.12lf_%.d\n", bestDist, W.getQueriesUsed());
580 W.outputImage("restored.pgm");
581 }
582 };
583
584 // RestorerQP.cpp
585 // The restorer for the quadratic programming algorithm.
586 class RestorerQP : public Restorer
587 {
588 private:
589     int queryLimit;
590     int COUNT;
591
592     vector<Change> changes;
593
594     matrixDouble A;
595     matrixDouble V;
596     vecDouble b;
597
598 public:
599     RestorerQP(string encodingLog, int queries) : Restorer(encodingLog)
600     {
601         queryLimit = queries;
602         COUNT = 2000;
603     }
604
605     void apply()
606     {
607         clock();
608         srand(time(0));
609
610         vecDouble curWAM = W.getWAM();
611         double curDist = Distance(curWAM, goal);
612         ECHO("%.12lf_%.d\n", curDist, W.getQueriesUsed());
613
614         while(W.getQueriesUsed() < queryLimit)
615         {
616             if(W.getQueriesUsed() >= 20000) COUNT = 1000;
617             if(W.getQueriesUsed() >= 30000) COUNT = 500;
618             if(W.getQueriesUsed() >= 40000) COUNT = 50;
619
620             pickChanges();
621             calculateModel();
622             writeMatlabInput();
623             signalMatlab();
624             idleUntilMatlabIsReady();
625             performChanges();
626         }
627         W.outputImage("restored.pgm");

```

```

628     }
629
630 void pickChanges()
631 {
632     int i, x, y, v;
633
634     random_shuffle(availablePixels.begin(), availablePixels.end());
635     changes.clear();
636     for(i = 0; i < availablePixels.size() && changes.size() < COUNT && changes.size() + W.
        getQueriesUsed() < queryLimit; i++)
637     {
638         x = availablePixels[i].x;
639         y = availablePixels[i].y;
640         if(ABS(changesMade[x][y]) == PIXEL_CHANGE_CAP) continue;
641         Change c = newChange(x, y, -1);
642         changes.push_back(c);
643         c = newChange(x, y, 1);
644         changes.push_back(c);
645     }
646     sort(changes.begin(), changes.end());
647 }
648
649 void calculateModel()
650 {
651     int i, j;
652     vecDouble f, v, delta;
653     int n = changes.size();
654
655     A.clear();
656     V.clear();
657     b.clear();
658     f = W.getWAM();
659     V.resize(27);
660     for(i = 0; i < 27; i++) V[i].resize(n);
661     for(i = 0; i < n; i++)
662     {
663         Change c = changes[i];
664         v = W.QueryOneChange(c);
665         delta = v - f;
666         for(j = 0; j < 27; j++)
667             V[j][i] = delta[j];
668     }
669     // A = V'*covInv*V;
670     matrixDouble Vprimed = transpose(V);
671     matrixDouble a = Vprimed*covInv;
672     A = a*V;
673     // b = 2*V'*covInv*(f - means);
674     vecDouble c = f - means;
675     for(i = 0; i < 27; i++) c[i] *= 2;
676     b = a*c;
677 }
678
679 void writeMatlabInput()
680 {
681     int i, j;
682     int n = changes.size();
683     // Describe the quadratic objective. Note that Matlab's
684     // input is 1/2x'Hx, so I need to double my A before putting it in.
685     FILE *fout = fopen("H.dat", "wb");
686     for(i = 0; i < n; i++)
687     {
688         for(j = 0; j < n; j++)
689             FECHO(fout, "%.12lf_", 2*A[i][j]);
690         FECHO(fout, "\n");
691     }
692     fclose(fout);

```

```

693
694 // Describe the linear objective. Matlab's f is my b.
695 fout = fopen("f.dat", "wb");
696 for(i = 0; i < n; i++)
697     FECHO(fout, "%.12lf_", b[i]);
698     FECHO(fout, "\n");
699 fclose(fout);
700 // Linear constraint matrix.
701 // A constraint for each i, i + 1. We did not use these! The Matlab script ignores them
702
703 fout = fopen("A.dat", "wb");
704 for(i = 0; i < n; i += 2)
705     {
706         for(j = 0; j < n; j++)
707             if(j == i || j == i + 1)
708                 FECHO(fout, "1_");
709             else
710                 FECHO(fout, "0_");
711             FECHO(fout, "\n");
712         }
713     fclose(fout);
714 // Right-hand side of the linear constraints.
715 fout = fopen("b.dat", "wb");
716 for(i = 0; i < n; i += 2)
717     FECHO(fout, "1_");
718     FECHO(fout, "\n");
719     fclose(fout);
720 }
721
722 void signalMatlab()
723 {
724     FILE *fout;
725     fout = fopen("signalToMatlab", "wb");
726     FECHO(fout, "112233\n");
727     fclose(fout);
728 }
729
730 void idleUntilMatlabIsReady()
731 {
732     int x;
733     FILE *fin;
734     while(1)
735     {
736         fin = fopen("signalFromMatlab", "rb");
737         if(fin != NULL)
738         {
739             fscanf(fin, "%d", &x);
740             if(x == 998877) break;
741         }
742         fclose(fin);
743         remove("signalFromMatlab");
744     }
745 }
746
747 void performChanges()
748 {
749     int n = changes.size();
750     int i;
751     vector<Change> gc;
752     double x;
753     FILE *fin = fopen("matlabOutput", "rb");
754     for(i = 0; i < n; i++)
755     {
756         fscanf(fin, "%lf", &x);
757         if(x >= 0.5) { gc.push_back(changes[i]); }
758     }
759 }

```

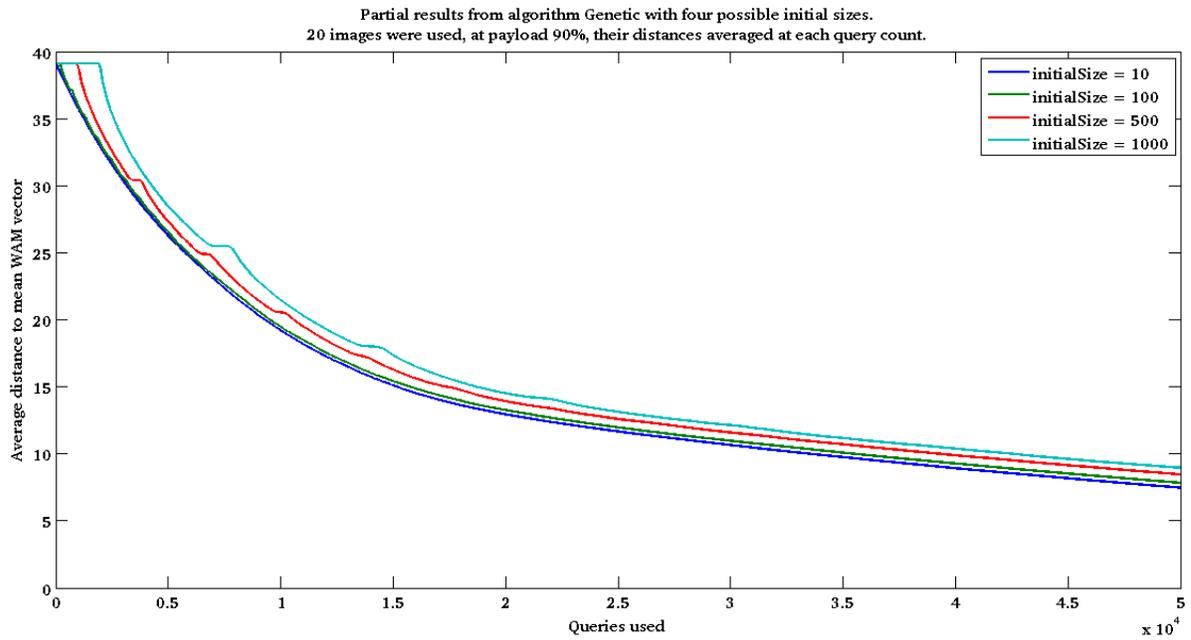
```

758     fclose(fin);
759     W.makeGroupChange(gc);
760     vecDouble curWAM = W.getWAM();
761     double curDist = Distance(curWAM, goal);
762     for(i = 0; i < gc.size(); i++)
763     {
764         int pixelx = gc[i].x;
765         int pixely = gc[i].y;
766         int changev = gc[i].v;
767         changesMade[pixelx][pixely] += changev;
768     }
769     ECHO("%.121f_%.121f\n", curDist, W.getQueriesUsed());
770 }
771 };
772
773 % The Matlab script we used to automate the QP algorithm.
774 while true
775     clear;
776     fid = fopen('signalToMatlab', 'r');
777     if(fid == -1) continue; end;
778     x = fscanf(fid, '%d', 1);
779     fclose(fid);
780     if x ~= 112233 continue; end;
781     load -ascii H.dat;
782     load -ascii f.dat;
783     n = 2000;
784     lb = ones(n, 1) - ones(n, 1);
785     ub = ones(n, 1);
786     tic
787     xs = quadprog(H, f, [], [], [], [], lb, ub);
788     toc
789     save -ascii -double 'matlabOutput' xs;
790     delete('signalToMatlab');
791     delete('H.dat');
792     delete('f.dat');
793     fid = fopen('signalFromMatlab', 'w');
794     fprintf(fid, '998877');
795     fclose(fid);
796 end

```

Appendix D

Partial Benchmarking of Genetic Algorithm



Appendix E

Partial Benchmarking of Random Algorithm

