OPTIMAL DETECTION OF MULTIPLICATIVE WATERMARKS USING IMAGE HASHES

BY

JILLIAN LEIGH CANNONS

B.S., University of Manitoba, 2000

THESIS

Submitted in partial fulfillment of the requirements
for the degree of Master of Science in Electrical Engineering
in the Graduate College of the
University of Illinois at Urbana-Champaign, 2002

Urbana, Illinois

ABSTRACT

Digital image watermarking is a method by which a message can be hidden in an image for the purpose of information protection. In this thesis, the host data into which the message are embedded is chosen as the highest-magnitude coefficients from a subset of the full-frame image transform, and the watermark is inserted through multiplicative embedding. To complete the watermarking system, a detector that is capable of accurately determining the presence or absence of a watermark in a given image must be developed.

This thesis focuses on the formulation of statistically optimal watermark detectors that incorporate the use of side information in the form of an image hash, creating a joint hashing/watermarking system. The classical Weibull and power exponential distributions are considered as probabilistic models of the host data. The notched power exponential distribution is also proposed to better represent the data selected for watermarking. For each distribution, detectors are derived using the techniques of likelihood ratio testing, maximum likelihood estimation, and locally optimal detection. Chernoff bounds on performance are found for the likelihood ratio test detectors, and Chernoff distance is used to quantify the difficulty of the detection problem, including that of an eavesdropper's problem.

Through analysis and experimentation using synthesized data and real-world test images, the use of the notched power exponential distribution and likelihood ratio testing is observed to produce a vastly superior detector, as measured by the probabilities of false alarm and detection. Although the detector derivations do not include models of attacks, the effects of additive white Gaussian noise and JPEG compression are explored, and the notched power exponential likelihood ratio test detector again yields the highest-quality performance. The exemplary results obtained clearly demonstrate that the combination of statistical modeling and detection theory permit the development of promising detectors, thereby facilitating the use of joint digital image hashing/watermarking in practical applications.

ACKNOWLEDGMENTS

I would like to express my sincere thanks to Professor Pierre Moulin, my adviser, who has been supportive, patient, and perceptive throughout my pursuit of a Master's degree. Also, thank you to my family for their far-reaching motivation and encouragement in all my academic endeavors. Finally, to Ryan Szypowski, who has provided focus and insight at times of difficulty in my research, and whose friendship and love instill in me the belief that I can achieve any goals to which I aspire.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

LIST OF SYMBOLS

$\mathbf{A}$	transform coefficients of the original image
$\mathbf{B}$	transform coefficients of the watermarked image
$\tilde{\mathbf{B}}$	transform coefficients of the attacked watermarked image
$D(t_0^*)$	Chernoff distance
E_{P_D}	error exponent associated with P_D
E_{P_F}	error exponent associated with P_F
H_0	hypothesis zero
H_1	hypothesis one
L	size of the candidate set
M	total number of transform coefficients/image pixels
N	number of transform coefficients selected for watermarking
P_D	probability of detection
P_F	probability of false alarm
P_M	probability of miss
$P(z)$	cumulative density function
$\mathbf{Z}$	additive white Gaussian noise
$\mathbf{a}$	pixels of the original image
$\mathbf{b}$	pixels of the watermarked image
$\tilde{\mathbf{b}}$	pixels of the attacked watermarked image
$\mathbf{h}$	hash of the image corresponding to the candidate region
i, j, k	vector indices
$\mathbf{m}$	message to be embedded
p_i	probability distribution under hypothesis i

$\mathbf{r}$	full message corresponding to the candidate region
s	unspecified embedding strength
s^*	specified embedding strength
$\hat{s}$	estimate of embedding strength
t_i	parameter of the cumulant generating function under hypothesis i
t_i^*	maximizing value of t_i
$\mathbf{x}$	selected transform coefficients of the original image
$\mathbf{y}$	selected transform coefficients of the watermarked image
$\tilde{\mathbf{y}}$	selected transform coefficients of the attacked image
α	probability distribution parameter
β	probability distribution parameter
δ	threshold parameter of the notched probability distribution
δ_i	watermarked version of the threshold parameter
$\tilde{\delta}$	maximum of δ and δ_i
γ	threshold against which the log-likelihood ratio is compared
κ	fraction of probability mass discarded from the power exponential distribution in forming a notched probability distribution
λ	probability of choosing H_0 in Regions 1 and 4 using the simple detector
μ_i	cumulant generating function under hypothesis i
ν	fraction of the image contained in the candidate set
ρ	maximum allowable false alarm probability for Neyman-Pearson testing
τ	threshold against which the maximum likelihood estimate or the modified likelihood ratio is compared
$\mathcal{C}$	set of image transform coefficients which are candidates for watermarking
$\mathcal{D}$	decision rule
$\mathcal{E}_2$	event that a transform coefficient lies in Region 2 for the notched power exponential distribution
$\mathcal{E}_3$	event that a transform coefficient lies in Region 3 for the notched power exponential distributions
$\mathcal{I}^+$	set of indices for which $m_i = 1$
$\mathcal{I}^-$	set of indices for which $m_i = -1$

$\mathcal{G}$	function used to simplify μ_i expression
$\mathcal{N}(0, \sigma^2)$	zero-mean Gaussian distribution with variance σ^2
$\mathcal{S}$	acceptable range of $\hat{s}$ values

LIST OF ACRONYMS

AWGN	additive white Gaussian noise
BHT	binary hypothesis test(ing)
DCT	discrete cosine transform
DFT	discrete Fourier transform
DWT	discrete wavelet transform
iid	independent and identically distributed
JPEG	Joint Photographic Experts Group
LOD	locally optimal detection
LRT	likelihood ratio test(ing)
MLE	maximum likelihood estimation
MSE	mean squared error
NPE	notched power exponential (distribution)
PE	power exponential (distribution)
ROC	receiver operating characteristic
SD	simple detector (for notched distributions)
WB	Weibull (distribution)

CHAPTER 1

INTRODUCTION

1.1 Purpose

The purpose of this thesis is to develop statistically optimal watermark detectors that combine the use of hashes into the watermarking process, creating a joint hashing/watermarking technique. Theoretical bounds are also found on the behavior of these detectors. The degree of difficulty experienced by an eavesdropper attempting to detect the watermark is also considered. Monte Carlo simulations are then performed using both synthetic and real-world data. Finally, various attacks on the watermarked object are introduced to study the robustness of the derived detectors.

1.2 Problem

The desire and ability to hide information in a commonly used medium have been present in society for thousands of years. Throughout the generations, the techniques used to accomplish these covert goals have varied, and, with the current prevalence of digital multimedia data, these methods continue to evolve. Data hiding has a vast array of applications, including the newly developing subfield of digital watermarking. Watermarking is a type of information hiding where the purpose is data protection. Typically, a watermark is employed to protect the rights of a creator of a multimedia object. For example, a watermark can be introduced as a means of identifying the owner of a digital image, or it could contain a unique serial number for each individual who has been given access rights to a digital movie. Regardless of the specific instance of watermarking, the unifying feature is the utilization of data hiding techniques to benefit the owner of a multimedia object. In general, the insertion of a watermark should be imperceptible in the host object (invisible watermarking); however, the less common visible watermarking is preferred in certain situations. Also, since the reliability of a watermark rests on security, a property of a watermarking method is that it should be robust against attacks. A trade-off exists between the two properties of invisibility and robustness, thereby providing a challenge in developing an ideal watermarking scheme.

The process of watermarking can be divided into three distinctly different portions, namely, data modeling, watermark embedding, and watermark detection. Since much previous watermarking research has focused on the embedding process, this thesis mainly considers the development of detection techniques, specifically for multiplicative embedding. However, the derivations found herein are easily applied to other embedding processes. To produce more robust detectors, side information is provided to the detector in the form of a hash of the original object, creating a joint hashing/watermarking system. A hash function is a function that takes in a set of data (in this case the original multimedia object) and a cryptographic key, and produces a simplified representation of that data. This resulting hash is quite small and, thus, is easily provided to the watermark detector. The particular hash function selected for use in this thesis leads to an expansion of the methods currently employed in data modeling. Although joint hashing/watermarking systems have been examined in other works [1], they have not been considered from a statistical standpoint. By formulating the watermarking process in the context of a communications system, the theories of signal detection and estimation can be employed. Furthermore, performance bounds can be calculated to determine the worst-case behavior of the detectors. Similarly, measures can be developed to illustrate the relative difficulty of different detection problems, including the problem seen by an eavesdropper attempting to detect a watermark. Using these insights from the field of communications, this thesis develops and analyzes statistically optimal means of detecting a watermark that has been multiplicatively embedded into a digital image.

1.3 Thesis Organization

This thesis is divided into eight chapters. Chapter 1 presents the problem addressed by the thesis and describes the need for further solutions. Chapter 2 formulates the joint hashing/watermarking system in a mathematical framework so that optimal detectors can be developed in Chapter 3. Chapter 4 then derives theoretical bounds on the performance of the studied detectors. In Chapter 5, experiments are discussed and performed using synthetically generated data to determine the abilities of the detection techniques. In Chapters 6 and 7, the developed detectors are applied to real-world test images in both the absence and presence of an attack, respectively. Finally, Chapter 8 provides conclusions based on the development and analysis presented in the thesis, and offers recommendations for the future directions of the watermarking problem.

CHAPTER 2

PROBLEM FORMULATION

In this chapter the watermarking problem is presented mathematically. First, the watermarking system, including the embedding technique, is discussed. This scheme is then specialized for use with digital images using a joint hashing/watermarking setup. Then, three different methods that may be used to model the host data are developed. Finally, for the purpose of implementation, techniques allowing synthetic host data to be generated are presented.

2.1 Basic Watermarking Problem

The basic watermarking problem can be described pictorially by the diagram in Figure 2.1. First, there is a vector of host data, $\mathbf{x}$, where the $\{x_i\}$ are assumed to be independent, into which a specific, known watermark vector, $\mathbf{m}$, is to be embedded. The result of this embedding operation, $\mathbf{y}$, may then be passed through an attack channel, creating $\tilde{\mathbf{y}}$, before being presented to the detector. Receiving this signal, the detector must determine whether or not the specific message is present. To examine this system in more detail, first consider the watermark embedding process. A variety of methods is available to insert a watermark into a set of host data; this thesis, however, is concerned with multiplicative watermark embedding [2]. In this setup, each element of the watermarked data is generated according to the formula

$$y_i = x_i(1 + sm_i)$$

where s is the strength of the watermark embedding. The sole restriction placed on s is that is should be a real number within the interval $[0, 1)$. This limitation is applied because the watermark is to be visually undetectable in the host data. For simplicity, each element

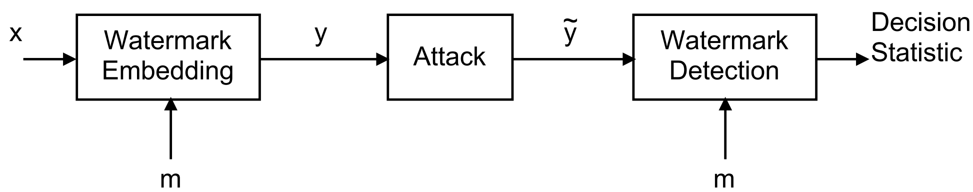

Figure 2.1 Basic watermarking process.

of the watermark message to be embedded is generated uniformly and independently from a binary distribution with values $\{-1, 1\}$; however, the use of other modulation techniques and of error correcting codes is also possible. By examining this embedding formula, it is easily seen that the amount of error introduced into an element is proportional to the size of the element itself.

With the watermark now embedded into the host data, the watermarked object is released into the public domain. Thus, the detector has no knowledge of the manipulations performed on the object, which could include simple image processing operations or even attempts to remove the watermark itself. Hence, an attack block is introduced into the diagram, producing the vector $\tilde{\mathbf{y}}$.

Finally, the corrupted version of the watermarked data is presented to a detector, which is responsible for determining whether or not the specific watermark is present in the sequence. Such a detector may yield a number of possible types of output. First, the detector may simply give a yes or no answer regarding the presence of the specific watermark. Alternately, the detector could present a statistic indicating how certain it is that the given watermark is present, which may then be compared against a threshold to determine a yes or no answer. Finally, the detector can give an estimate of the actual embedding strength parameter s, which again can be compared against a threshold to determine a binary answer.

2.2 Image Watermarking Problem

With the basic watermarking process formalized, the system is now specialized for use in a particular application. First, the source of the host data must be considered. Depending on the goals of the watermark embedder, a multitude of sources of host data is possible. This thesis deals with the problem of watermarking a digital grayscale image using a joint hashing/watermarking technique. The specialized setup is described pictorially in Figure 2.2. In this context, a common source of host data is image transform coefficients, as opposed to the actual pixel intensities themselves. This choice often results in watermarks that are more robust to attacks. Transforms such as the *discrete Fourier transform* (DFT), the *discrete cosine transform* (DCT), and the *discrete wavelet transform* (DWT) are all viable candidates for use in watermarking. The DCT is selected for use in this thesis for simplicity, and the full-frame transform is employed to increase robustness because attacks in this domain will be reflected throughout the entire spatial image. An additional advantage of full-frame DCT watermarking is that it is inherently robust against image resizing attacks.

To further increase the security of the watermarking system, a subset of the transform coefficients is defined, and forms a set of candidate points for embedding. This candidate set is of size L and is described by its indices into the full-frame coefficients, $\mathcal{C} = \{c_1, c_2, \ldots, c_L\}$, where the c_i are coefficient indices. Note that $\mathcal{C}$ will often be used to refer to the transform coefficients themselves, as well as their index set. The parameter ν denotes the fraction of coefficients included in the candidate set, $\nu = \frac{L}{M}$, where M is the total number of transform coefficients. This set acts as a secret key, which is image independent, and must be known by both the embedder and the detector.

The watermark will then be embedded in the transform coefficients of $\mathcal{C}$ having the highest magnitudes. More specifically, the host data set, $\mathbf{x}$, contains only those coefficients with magnitudes greater than a fixed parameter, δ. These coefficients are dominant in the image and, thus, are expected to be quite well preserved in the presence of a moderate attack. Moreover, these high-magnitude coefficients are most robust against attacks (small-magnitude coefficients may simply be discarded by an attacker) and, thus, are most effective in terms of detection performance. It should be noted that since the same candidate set and coefficient threshold are used for various images, the length of the host vector, denoted by $N \leq L$, will be image dependent.

The method by which the watermark is constructed depends partially on the specific image watermarking application. One technique is for the encoder to be given a binary sequence, $\mathbf{r}$ with $r_i \in \{-1, 1\}$, of the same size (L) as the candidate set. Each element of this longer sequence corresponds to a coefficient in $\mathcal{C}$. Those elements of $\mathbf{r}$ coinciding with the

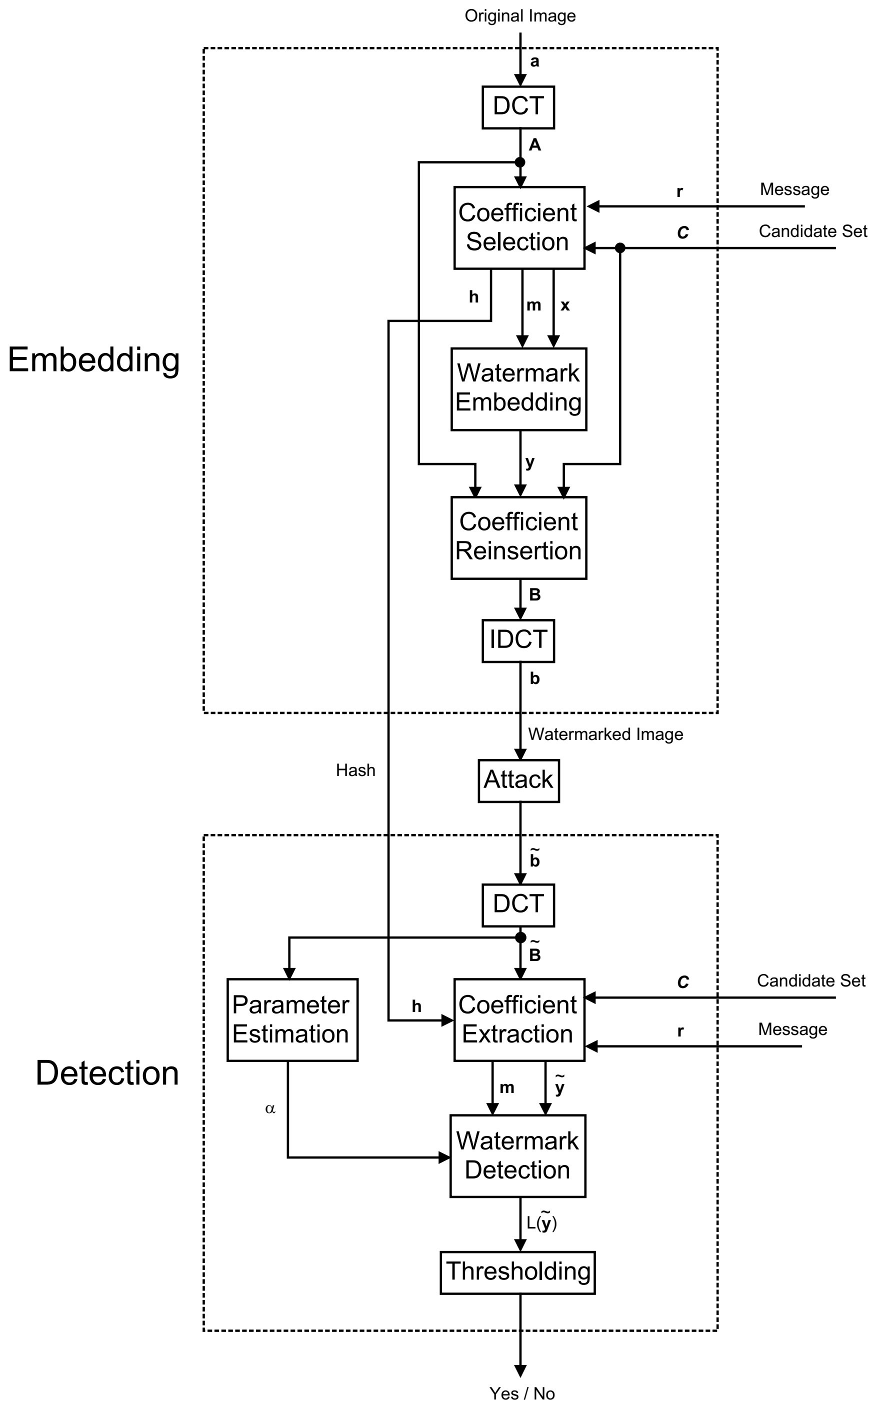

Figure 2.2 Image watermarking process.

coefficients of $\mathbf{x}$ are denoted by $\mathbf{m}$. Thus, $\mathbf{m}$ is of length N and forms the actual watermark message that will be embedded in the image.

Using the above method, the precise value of $\mathbf{m}$ is dependent on the original image. If more control is desired over the sequence actually embedded, the encoder could be altered to accept a shorter sequence (of length less than N) that, with padding, will be embedded into the host vector, thereby ensuring the presence of the specified sequence. The encoder would then generate and output the full message, $\mathbf{r}$, which includes the vector $\mathbf{m}$ at appropriate locations, and random -1 and 1 padding for the remainder of the sequence. Although this method guarantees that a particular sequence will be embedded, the full message, $\mathbf{r}$, will be dependent on the original image.

With the development of the image watermarking encoder complete, the amount of distortion introduced by the multiplicative embedding is now considered. To quantify this amount, the *mean squared error* (MSE) distortion measure is defined as

$$D(\mathbf{u}, \mathbf{v}) = \frac{1}{M} \sum_{k=1}^{M} |u_k - v_k|^2. \tag{2.1}$$

Although this measure is known to be not particularly well suited to quantifying the change in an image as perceived by the human visual system, it is widely used because it is conceptually and computationally simple. When the measure is applied to this specific multiplicative setup, the MSE is given by

$$\begin{aligned} D_1 = D_1(\mathbf{a}, \mathbf{b}) &= \frac{1}{M} \sum_{k=1}^{M} |a_k - b_k|^2 \\ &= \frac{1}{M} \sum_{i=1}^{M} |A_i - B_i|^2 \\ &= \frac{1}{M} \sum_{i=1}^{N} |x_i - y_i|^2 \\ &= \frac{1}{M} \sum_{i=1}^{N} |x_i - x_i(1 + s^* m_i)|^2 \\ &= \frac{1}{M} \sum_{i=1}^{N} (x_i s^*)^2 \end{aligned}$$

where

$\mathbf{a}$ = the pixels of the original image,

$\mathbf{b}$ = the pixels of the watermarked image,

$\mathbf{A}$ = the transform coefficients of the original image, and

$\mathbf{B}$ = the transform coefficients of the watermarked image.

Note that Parseval's theorem is used to express the distortion in the transform domain.

The focus of the problem specification is now turned to the watermark detector. For each image, the detector will be presented with a vector of side information that is similar to a hash of the image. This vector, $\mathbf{h}$, contains elements that correspond to each coefficient in the candidate set. Each h_i will be either 0 or 1, with 1 indicating a coefficient magnitude greater than δ, and 0 indicating a magnitude less than or equal to δ. Equivalently, each h_i can be viewed as a 1-bit quantized version of the corresponding original image coefficient.

With the specification of the candidate set, the hash vector, the full message, and the possibly watermarked image, the detector is able to determine which coefficients belong to the possibly watermarked vector, $\tilde{\mathbf{y}}$, and the corresponding watermark, $\mathbf{m}$, for which to search. Note that since additional information is provided to the detector (as opposed to simply the message), this scheme is not considered to be a blind watermarking system. However, since the original image is not required at the detector, it is also not a non-blind system. Rather, this image watermarking setup is best described as a joint hashing/watermarking system.

2.3 Host Data Modeling

Since this thesis deals with the optimal detection of watermarks, it relies heavily on the statistical properties of the data involved. Thus, a realistic probabilistic representation of the host data must be determined. In the case of image transform coefficients, the *Weibull distribution* (WB) or the *power exponential distribution* (PE) are often utilized to model the resulting distributions. However, this thesis advocates the use of the *notched power exponential distribution* (NPE) to reflect the fact that only large coefficients are selected for watermarking, as specified by the hash. These three distributions are now defined, and their properties are discussed.

2.3.1 Weibull distribution

The Weibull distribution [3] is given by

$$p(x) = \begin{cases} \frac{\beta}{\alpha}\left(\frac{x}{\alpha}\right)^{\beta-1} \exp\left\{-\left(\frac{x}{\alpha}\right)^{\beta}\right\} & \text{if } x > 0 \\ 0 & \text{else} \end{cases}$$

for $x \in \mathbb{R}$, $\alpha > 0$, and $\beta > 0$. A plot of this distribution is given in Figure 2.3. The mean of the Weibull distribution is given by

$$\mathrm{E}[X] = \alpha\Gamma\left(1 + \frac{1}{\beta}\right)$$

and the variance is given by

$$\begin{aligned} \mathrm{Var}[X] &= \alpha^2\Gamma\left(1 + \frac{2}{\beta}\right) - (\mathrm{E}[X])^2 \\ &= \alpha^2\left(\Gamma\left(1 + \frac{2}{\beta}\right) - \left[\Gamma\left(1 + \frac{1}{\beta}\right)\right]^2\right) \end{aligned}$$

where $\Gamma(z) = \int_0^\infty e^{-t}t^{z-1}dt with z > 0$, is the Gamma function. By varying the two parameters of the distribution, α and β, the mean, variance, and general shape of the curve can be refined.

The Weibull distribution has been used to model the magnitude of the discrete Fourier transform and discrete cosine transform coefficients of an actual image [3], [4]. However, the parameters of the distribution may not be known. From the above expressions for mean and variance, it is possible to estimate α and β from sample data by computing the sample mean and variance, where

$$\text{Sample } \mathrm{E}[x] = \frac{1}{n}\sum_{i=1}^{n} x_i$$

and

$$\text{Sample } \mathrm{Var}[x] = \frac{1}{n-1}\sum_{i=1}^{n}\left(x_i - \bar{x}\right)^2.$$

From these values, a system of two equations can be solved to yield estimates of distribution parameters.

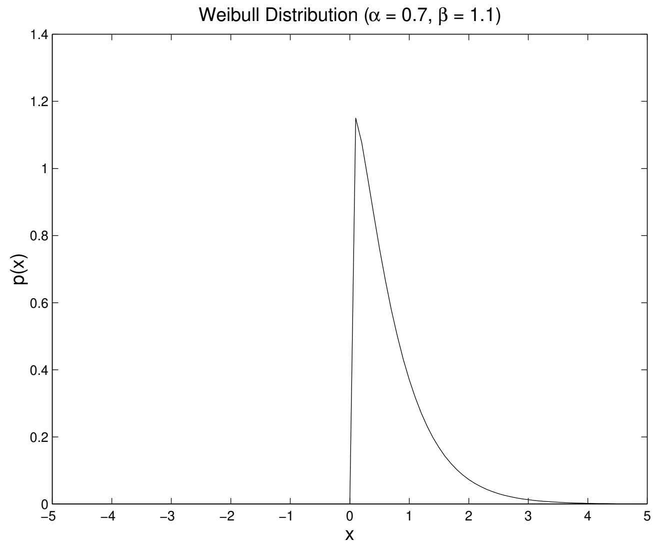

Figure 2.3 Weibull distribution with parameters $\alpha = 0.7$ and $\beta = 1.1$.

2.3.2 Power exponential distribution

The power exponential distribution (also called the generalized error distribution or the generalized Gaussian distribution) is commonly used to model the distribution of an image's discrete cosine transform coefficients [5]. The power exponential distribution contains two parameters, α and β, with α relating to the variance and β relating to the heaviness of the distribution tails. The distribution itself is given by

$$p(x) = C \exp\left\{ - \left| \frac{x}{\alpha} \right|^{\beta} \right\}$$

where

$$C = \frac{\beta}{2\alpha\Gamma\left(\frac{1}{\beta}\right)}$$

for $x \in \mathbb{R}$, $\alpha > 0$, and $\beta > 0$. Figure 2.4 contains a plot of the distribution. The second moment (variance) and fourth moment of the power exponential distribution [6] are given by

$$\mathrm{Var}[x] = M_2 = \frac{\alpha^2 \Gamma\left(\frac{3}{\beta}\right)}{\Gamma\left(\frac{1}{\beta}\right)}$$

and

$$M_4 = \frac{\alpha^4 \Gamma\left(\frac{5}{\beta}\right)}{\Gamma\left(\frac{1}{\beta}\right)}.$$

Note that when $\beta = 1$, which is a common value when the distribution is utilized to model image transform coefficients, the power exponential distribution reduces to the Laplacian distribution [5].

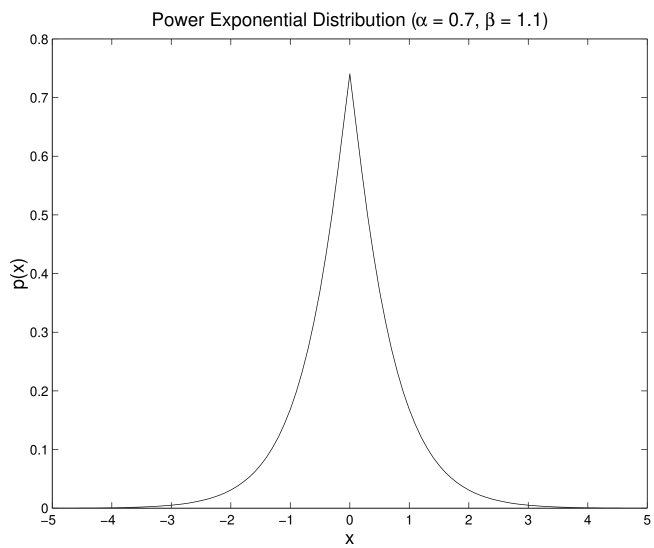

Figure 2.4 Power exponential distribution with parameters $\alpha = 0.7$ and $\beta = 1.1$.

When using the power exponential distribution to model the coefficients used in watermarking, values of α and β must be determined on the basis of the actual image coefficients. By computing the second and fourth moments of the data, the two parameters can be estimated by solving a system of two equations.

2.3.3 Notched power exponential distribution

In the multiplicative watermarking scheme, the image transform coefficients can be modeled using the power exponential distribution. However, the high-magnitude coefficients selected for watermarking reside in the tail portions of the power exponential distribution and occur with lower probabilities. Therefore, the actual distribution of the host data is a notched power exponential distribution. In addition to the original power exponential distribution parameters, α and β, a new positive parameter, δ, is introduced. Here, $\{x : |x| > \delta\}$ is the range of values of the coefficients that are watermarked, where $\delta > 0$. For convenience, the derived parameter $\kappa \in (0, 1)$ is also introduced to represent the cumulative probability under the power exponential distribution of the coefficients that are not selected. These two parameters are shown pictorially with respect to the power exponential distribution in Figure 2.5. The shaded areas represent the portions contained within the notched power exponential distribution.

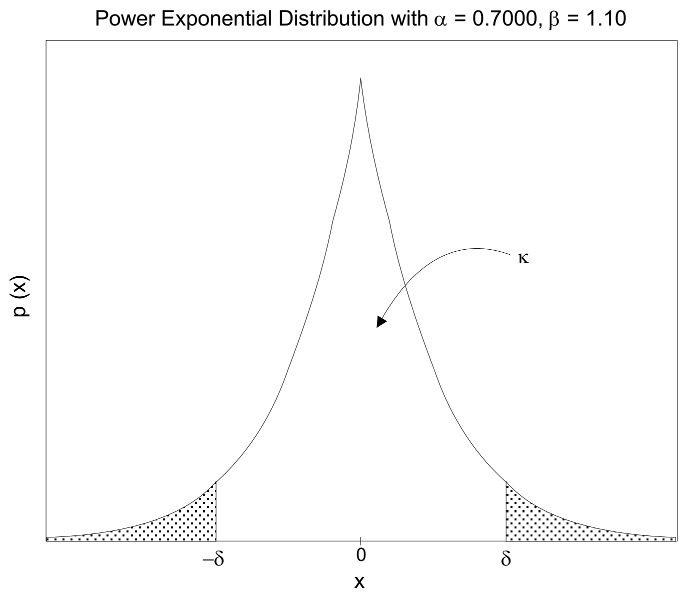

Figure 2.5 Power exponential distribution with notched distribution parameters δ and κ.

Due to the symmetry of the distribution, the parameters δ and κ are related according to the equation

$$\begin{aligned}
\frac{1-\kappa}{2} &= \int_{\delta}^{\infty} p(x)dx \\
&= \int_{\delta}^{\infty} C \exp\left\{-\left(\frac{x}{\alpha}\right)^{\beta}\right\} dx \\
&= \frac{\beta}{2\alpha\Gamma\left(\frac{1}{\beta}\right)} \int_{\delta}^{\infty} \exp\left\{-\left(\frac{x}{\alpha}\right)^{\beta}\right\} dx.
\end{aligned}$$

Performing the substitution $w = \left(\frac{x}{\alpha}\right)^{\beta}$, this equation becomes

$$\begin{aligned}
\frac{1-\kappa}{2} &= \int_{\left(\frac{\delta}{\alpha}\right)^{\beta}}^{\infty} \frac{\beta}{2\alpha\Gamma\left(\frac{1}{\beta}\right)} \exp\{-w\} \frac{\alpha}{\beta} w^{-\left(1-\frac{1}{\beta}\right)} dw \\
&= \frac{1}{2\Gamma\left(\frac{1}{\beta}\right)} \int_{\left(\frac{\delta}{\alpha}\right)^{\beta}}^{\infty} \exp\{-w\} \frac{\alpha}{\beta} w^{-\left(1-\frac{1}{\beta}\right)} dw \\
1-\kappa &= \frac{1}{\Gamma\left(\frac{1}{\beta}\right)} \Gamma\left(\frac{1}{\beta}, \left(\frac{\delta}{\alpha}\right)^{\beta}\right)
\end{aligned}$$

$$\kappa \;=\; 1-\frac{1}{\Gamma\left(\frac{1}{\beta}\right)}\Gamma\left(\frac{1}{\beta},\left(\frac{\delta}{\alpha}\right)^{\beta}\right) \tag{2.2}$$

where $\Gamma(a,z)=\int_z^{\infty}\exp\{-t\}t^{a-1}dt$ with $a>0$, is the incomplete Gamma function.

Thus, the notched power exponential can be expressed as

$$p(x)=\begin{cases}\frac{C}{1-\kappa}\exp\left\{-\left|\frac{x}{\alpha}\right|^{\beta}\right\} & \text{if } |x|>\delta\\ 0 & \text{else}\end{cases}$$

where

$$C=\frac{\beta}{2\alpha\Gamma\left(\frac{1}{\beta}\right)} \tag{2.3}$$

for $x\in\mathbb{R}$, $\alpha>0$, $\beta>0$, $\delta>0$, and $0<\kappa<1$.

In order to use the notched power exponential distribution to model selected image transform coefficients, the distribution parameters must be computed from the data. The α and β parameters are given by the same expressions as in the power exponential case, using all of the transform coefficients. The value of the notch threshold, δ, is identical to the fixed system parameter used to determine the host vector in the image watermarking scheme. Finally, κ can be computed from the other three parameters by using (2.2).

2.4 Host Data Generation

Watermark detection and host data modeling are two separate problems and should be treated as such. Thus, to accurately evaluate a watermark detector under a specified statistical model for the host data, synthetic data are generated. For each of the three distributions considered, the cumulative density function, $P(z)$, is found. Then, random numbers $\{v_i\}$ between $[0,1]$ can be generated according to a uniform distribution. Finally, for each v_i generated, the cumulative density function is inverted to find the data point, $z_i=P^{-1}(v_i)$. Using this process, it is possible to generate data according to the required distributions.

2.4.1 Weibull distribution

The cumulative density function for the Weibull distribution is given by

$$\begin{aligned} P(z) &= \int_0^z \frac{\beta}{\alpha}\left(\frac{x}{\alpha}\right)^{\beta-1} \exp\left\{-\left(\frac{x}{\alpha}\right)^{\beta}\right\} dx \\ &= 1-\int_z^\infty \frac{\beta}{\alpha}\left(\frac{x}{\alpha}\right)^{\beta-1} \exp\left\{-\left(\frac{x}{\alpha}\right)^{\beta}\right\} dx. \end{aligned}$$

Performing a change of variables, $w = \left(\frac{x}{\alpha}\right)^{\beta}$, yields

$$\begin{aligned} P(z) &= 1-\int_{\left(\frac{z}{\alpha}\right)^{\beta}}^\infty \frac{\beta}{\alpha}\left(\frac{x}{\alpha}\right)^{\beta-1} \exp\left\{-w\right\} \left(\frac{\alpha}{x}\right)^{\beta-1} \frac{\alpha}{\beta} dw \\ &= 1-\int_{\left(\frac{z}{\alpha}\right)^{\beta}}^\infty \exp\left\{-w\right\} dw \\ &= 1-\exp\left\{-\left(\frac{z}{\alpha}\right)^{\beta}\right\}. \end{aligned}$$

This function is now inverted to solve for the data point, z:

$$\begin{aligned} \exp\left\{-\left(\frac{z}{\alpha}\right)^{\beta}\right\} &= 1-P(z) \\ -\left(\frac{z}{\alpha}\right)^{\beta} &= \ln(1-P(z)) \\ \left(\frac{z}{\alpha}\right)^{\beta} &= -\ln(1-P(z)) \\ \frac{z}{\alpha} &= (-\ln(1-P(z)))^{\frac{1}{\beta}} \\ z &= \alpha(-\ln(1-P(z)))^{\frac{1}{\beta}}. \end{aligned}$$

2.4.2 Power exponential distribution

For the power exponential distribution, the cumulative density function is given by

$$\begin{aligned} P(z) &= \int_{-\infty}^z C\exp\left\{-\left|\frac{x}{\alpha}\right|^{\beta}\right\} dx \\ &= 1-\int_z^\infty C\exp\left\{-\left|\frac{x}{\alpha}\right|^{\beta}\right\} dx. \end{aligned}$$

Considering the case when $z \geq 0$ yields

$$P(z) = 1-\int_z^\infty C\exp\left\{-\left(\frac{x}{\alpha}\right)^{\beta}\right\} dx.$$

Using a change of variables, $w = \left(\frac{x}{\alpha}\right)^\beta$, gives

$$\begin{aligned}
P(z) &= 1 - C \int_{\left(\frac{z}{\alpha}\right)^\beta}^{\infty} \exp\{-w\} \frac{\alpha}{\beta} w^{-\left(1-\frac{1}{\beta}\right)} dw \\
&= 1 - \frac{C\alpha}{\beta} \int_{\left(\frac{z}{\alpha}\right)^\beta}^{\infty} \exp\{-w\} w^{\left(\frac{1}{\beta}-1\right)} dw \\
&= 1 - \frac{C\alpha}{\beta} \Gamma\left(\frac{1}{\beta}, \left(\frac{z}{\beta}\right)^\beta\right) \\
&= 1 - \frac{1}{2\Gamma\left(\frac{1}{\beta}\right)} \Gamma\left(\frac{1}{\beta}, \left(\frac{z}{\beta}\right)^\beta\right).
\end{aligned}$$

Similarly, for the case when $z < 0$,

$$\begin{aligned}
P(z) &= \int_{-\infty}^{z} C \exp\left\{-\left(\frac{-x}{\alpha}\right)^\beta\right\} dx \\
&= \frac{1}{2\Gamma\left(\frac{1}{\beta}\right)} \Gamma\left(\frac{1}{\beta}, \left(\frac{-z}{\alpha}\right)^\beta\right).
\end{aligned}$$

Finally, combining the two cases for z yields

$$P(z) = \begin{cases} 1 - \frac{1}{2\Gamma\left(\frac{1}{\beta}\right)} \Gamma\left(\frac{1}{\beta}, \left(\frac{z}{\alpha}\right)^\beta\right) & \text{if } z \geq 0 \\ \frac{1}{2\Gamma\left(\frac{1}{\beta}\right)} \Gamma\left(\frac{1}{\beta}, \left(\frac{-z}{\alpha}\right)^\beta\right) & \text{if } z < 0. \end{cases}$$

Due to the difficulty in inverting the incomplete Gamma function, this expression is not solved explicitly for the data point, z. To determine the required value, numerical techniques are employed.

2.4.3 Notched power exponential distribution

For the purpose of data generation, the notched power exponential distribution is quite similar to the power exponential distribution. The cumulative density function is

$$P(z) = \int_{-\infty}^{z} p(x) dx.$$

Consider the case when $z > \delta$. Then

$$\begin{aligned}
P(z) &= 1 - \int_{z}^{\infty} \frac{C}{1-\kappa} \exp\left\{-\left(\frac{x}{\alpha}\right)^\beta\right\} dx \\
&= 1 - \frac{1}{2(1-\kappa)\Gamma\left(\frac{1}{\beta}\right)} \Gamma\left(\frac{1}{\beta}, \left(\frac{z}{\beta}\right)^\beta\right)
\end{aligned}$$

where the expression for the integral is identical to that present for the case of the power exponential distribution. Now, consider when $-\delta \leq z \leq \delta$. Here,

$$P(z) = \frac{1}{2}$$

because the notched power exponential is symmetric, and has value 0 between $-\delta$ and δ. Finally, when $z < -\delta$,

$$\begin{aligned} P(z) &= \int_{-\infty}^{z} \frac{C}{1-\kappa} \exp\left\{ -\left(\frac{-x}{\alpha}\right)^{\beta} \right\} dx \\ &= \frac{1}{2(1-\kappa)\Gamma\left(\frac{1}{\beta}\right)} \Gamma\left(\frac{1}{\beta}, \left(\frac{-z}{\alpha}\right)^{\beta}\right) \end{aligned}$$

where again the integral was determined in the power exponential case. Therefore, combining these results yields

$$P(z) = \begin{cases} 1 - \frac{1}{2(1-\kappa)\Gamma\left(\frac{1}{\beta}\right)} \Gamma\left(\frac{1}{\beta}, \left(\frac{z}{\alpha}\right)^{\beta}\right) & \text{if } z > \delta \\ \frac{1}{2} & \text{if } -\delta \leq z \leq \delta \\ \frac{1}{2(1-\kappa)\Gamma\left(\frac{1}{\beta}\right)} \Gamma\left(\frac{1}{\beta}, \left(\frac{-z}{\alpha}\right)^{\beta}\right) & \text{if } z < -\delta. \end{cases}$$

Once again, an explicit formula for z is not found because of the complexity involved in inverting the incomplete Gamma function. However, the expression can be solved numerically to yield the value for the generated data point.

CHAPTER 3

WATERMARK DETECTION

This chapter develops three main techniques for detecting a watermark within a set of data. These detectors are based on *likelihood ratio testing* (LRT), *maximum likelihood estimation* (MLE), and *locally optimal detection* (LOD). All three are derived from statistical detection theory, and each one is formulated for the three distributions, namely, Weibull, power exponential, and notched power exponential. A fourth detector is also constructed for use with notched distributions as a simplification of the likelihood ratio test detector.

3.1 Simple Likelihood Ratio Test

The first detection technique covered is based upon *binary hypothesis testing* (BHT) using a likelihood ratio test [7]. In this setup, it is assumed that the embedding strength, s, is a known, non-random parameter, say s^*. Then, the detection problem is formulated as a choice between two hypotheses, hypothesis 0 (H_0) and hypothesis 1 (H_1). H_0 states that the data do not contain any watermark, while H_1 states that the data contain a specific watermark. Each hypothesis has associated with it a distribution for the data, $p_0(\mathbf{y})$ and $p_1(\mathbf{y})$. To determine which hypothesis to select, a decision rule is developed that maps each possible data vector to one hypothesis or the other. From statistical decision theory, it is known that the optimal decision rule is given by a likelihood ratio test, with the ratio being the distribution under H_1 to the distribution under H_0,

$$L(\mathbf{y}) = \frac{p_1(\mathbf{y})}{p_0(\mathbf{y})}.$$

For the sake of simplicity, the log of this ratio, the log-likelihood ratio, $\ln L(\mathbf{y})$, is often considered.

To perform the actual likelihood ratio test, this log-likelihood ratio is compared against a threshold value, γ. If the ratio is greater than the threshold, then H_1 is selected; while if it is less than the threshold H_0 is chosen.

The remaining question in the development of a likelihood ratio test-based detector is how to best choose the threshold value. The selected approach is based on Neyman-Pearson hypothesis testing [7]. The probability of false alarm, P_F, is defined to be the probability with which the detector falsely detects the watermark in a set of data. If a specific value of this probability, ρ, is selected as the maximum permissible P_F, then the corresponding threshold, γ, can be determined such that $P_F = \rho$. In the case of the log-likelihood statistic, the false alarm probability is given by

$$\begin{aligned} P_F &= P\left[\text{choosing } H_1 \mid H_0 \text{ is true}\right] \\ &= P_0\left[\ln L(\mathbf{y}) > \gamma\right] \\ &= \int_{\gamma}^{\infty} p_{0,L}(\ell)d\ell \end{aligned}$$

where $p_{0,L}(\ell)$ is the distribution of $\ln L(\mathbf{y})$ under H_0. Thus, if this distribution is computed, the desired threshold value can be found by setting the resulting P_F to ρ and solving for γ.

In summary, the likelihood ratio test detector is formed by calculating a log-likelihood ratio based on two hypotheses to yield a decision statistic. This statistic is then compared against a threshold to determine a yes or no answer as to the presence of the specific watermark in the data vector.

3.1.1 Weibull distribution

The first distribution considered for modeling the coefficients is the Weibull distribution [3]. Here the two hypotheses used in the binary hypothesis test are

$$\begin{aligned} H_0 \ &: \ s = 0 \\ & \quad\ y_i = x_i, \quad 1 \le i \le N \end{aligned}$$

$$\begin{aligned} H_1 \quad &: \quad s = s^* \\ & \quad y_i = x_i \left(1 + s^* m_i\right), \quad 1 \le i \le N \end{aligned}$$

where N is the number of (high-magnitude) coefficients that are watermarked. The distribution under H_0 is given by

$$\begin{aligned} p_0(y_i) &= p_{x_i}(y_i) \\ &= \frac{\beta}{\alpha}\left(\frac{y_i}{\alpha}\right)^{\beta-1} \exp\left\{-\left(\frac{y_i}{\alpha}\right)^{\beta}\right\} \end{aligned}$$

$$\begin{aligned} p_0(\mathbf{y}) &= \prod_{i=1}^{N} p_0(y_i) \\ &= \prod_{i=1}^{N} \frac{\beta}{\alpha}\left(\frac{y_i}{\alpha}\right)^{\beta-1} \exp\left\{-\left(\frac{y_i}{\alpha}\right)^{\beta}\right\} \end{aligned}$$

where $p_{x_i}(x_i)$ is the Weibull distribution of the host data. The product forming the distribution of the vector follows from the assumption that all of the y_i are *independent and identically distributed* (iid). Similarly, the distribution under H_1 can be written as

$$\begin{aligned} p_1(y_i) &= \frac{1}{1+s^* m_i} p_{x_i}\left(\frac{y_i}{1+s^* m_i}\right) \\ &= \frac{\beta}{\alpha(1+s^* m_i)}\left(\frac{y_i}{\alpha(1+s^* m_i)}\right)^{\beta-1} \exp\left\{-\left(\frac{y_i}{\alpha(1+s^* m_i)}\right)^{\beta}\right\} \end{aligned}$$

$$\begin{aligned} p_1(\mathbf{y}) &= \prod_{i=1}^{N} p_1(y_i) \\ &= \prod_{i=1}^{N} \frac{\beta}{\alpha(1+s^* m_i)}\left(\frac{y_i}{\alpha(1+s^* m_i)}\right)^{\beta-1} \exp\left\{-\left(\frac{y_i}{\alpha(1+s^* m_i)}\right)^{\beta}\right\}. \end{aligned}$$

With the distributions of the data under each hypothesis found, it is now possible to evaluate the log-likelihood ratio:

$$\begin{aligned} \ln L(\mathbf{y}) &= \ln \frac{p_1(\mathbf{y})}{p_0(\mathbf{y})} \\ &= \ln \frac{\prod_{i=1}^{N} \frac{\beta}{\alpha(1+s^* m_i)}\left(\frac{y_i}{\alpha(1+s^* m_i)}\right)^{\beta-1} \exp\left\{-\left(\frac{y_i}{\alpha(1+s^* m_i)}\right)^{\beta}\right\}}{\prod_{i=1}^{N} \frac{\beta}{\alpha}\left(\frac{y_i}{\alpha}\right)^{\beta-1} \exp\left\{-\left(\frac{y_i}{\alpha}\right)^{\beta}\right\}} \end{aligned}$$

$$
\begin{aligned}
&= \ln\left(\prod_{i=1}^{N} \frac{1}{(1+s^* m_i)^{\beta}} \exp\left\{-\left(\frac{y_i}{\alpha(1+s^* m_i)}\right)^{\beta} + \left(\frac{y_i}{\alpha}\right)^{\beta}\right\}\right) \\
&= \sum_{i=1}^{N} -\beta \ln(1+s^* m_i) + \sum_{i=1}^{N}\left[-\left(\frac{y_i}{\alpha(1+s^* m_i)}\right)^{\beta} + \left(\frac{y_i}{\alpha}\right)^{\beta}\right] \\
&= \sum_{i=1}^{N}\left[-\beta \ln(1+s^* m_i) - \left(\frac{y_i}{\alpha(1+s^* m_i)}\right)^{\beta} + \left(\frac{y_i}{\alpha}\right)^{\beta}\right]. \qquad (3.1)
\end{aligned}
$$

To perform the actual likelihood ratio test, the log-likelihood ratio must be compared against a threshold value, γ. It should be noted that for a fixed watermark, $\mathbf{m}$, the first term in the summation in the ratio is a constant and can thus be incorporated into the threshold to form a new threshold, if desired. At this point, a Neyman-Pearson approach can be applied to find an appropriate value for the threshold, from which a binary decision can be made regarding the presence of the watermark in the data.

3.1.2 Power exponential distribution

The likelihood ratio test detector is now developed for the case when the power exponential distribution is used to model the host coefficients [8]. In this case, the two hypotheses are given by

$$
\begin{aligned}
H_0 \ &: \ s = 0 \\
&\quad\ y_i = x_i, \quad 1 \le i \le N
\end{aligned}
$$

$$
\begin{aligned}
H_1 \ &: \ s = s^* \\
&\quad\ y_i = x_i\left(1 + s^* m_i\right), \quad 1 \le i \le N
\end{aligned}
$$

where N is the number of (high-magnitude) coefficients that are watermarked. The distribution under H_0 is

$$
p_0(y_i) = C \exp\left\{-\left|\frac{y_i}{\alpha}\right|^{\beta}\right\}
$$

$$
p_0(\mathbf{y}) = \prod_{i=1}^{N} C \exp\left\{-\left|\frac{y_i}{\alpha}\right|^{\beta}\right\}
$$

while the distribution under H_1 is

$$\begin{aligned} p_1(y_i) &= \frac{1}{1+s^*m_i} p_{x_i}\left(\frac{y_i}{1+s^*m_i}\right) \\ &= \frac{C}{1+s^*m_i} \exp\left\{-\left|\frac{y_i}{\alpha\,(1+s^*m_i)}\right|^\beta\right\} \end{aligned}$$

$$p_1(\mathbf{y}) = \prod_{i=1}^{N} \frac{C}{1+s^*m_i} \exp\left\{-\left|\frac{y_i}{\alpha\,(1+s^*m_i)}\right|^\beta\right\}.$$

Thus, the log-likelihood ratio is given by

$$\begin{aligned} \ln L(\mathbf{y}) &= \ln\frac{p_1(\mathbf{y})}{p_0(\mathbf{y})} \\ &= \ln\frac{\prod_{i=1}^{N} \frac{C}{1+s^*m_i} \exp\left\{-\left|\frac{y_i}{\alpha\,(1+s^*m_i)}\right|^\beta\right\}}{\prod_{i=1}^{N} C\exp\left\{-\left|\frac{y_i}{\alpha}\right|^\beta\right\}} \\ &= \ln\left(\prod_{i=1}^{N} \frac{1}{1+s^*m_i} \exp\left\{-\left|\frac{y_i}{\alpha\,(1+s^*m_i)}\right|^\beta + \left|\frac{y_i}{\alpha}\right|^\beta\right\}\right) \\ &= \sum_{i=1}^{N} \ln\left(\frac{1}{1+s^*m_i} \exp\left\{-\left|\frac{y_i}{\alpha\,(1+s^*m_i)}\right|^\beta + \left|\frac{y_i}{\alpha}\right|^\beta\right\}\right) \\ &= \sum_{i=1}^{N} \left(\ln\frac{1}{1+s^*m_i} - \left|\frac{y_i}{\alpha\,(1+s^*m_i)}\right|^\beta + \left|\frac{y_i}{\alpha}\right|^\beta\right). \end{aligned} \tag{3.2}$$

As in the case of the Weibull distribution, a threshold value can be compared against the decision statistic to find a yes or no detector output.

It is interesting to note that this statistic is nearly identical to the simple likelihood ratio test statistic (3.1) for the Weibull distribution. The differences are the presence of the absolute values and a scaling factor of β on the constant term. This similarity is expected because the two distributions have similar forms, with the power exponential possessing absolute value terms.

3.1.3 Notched power exponential distribution

Using likelihood ratio testing, a watermark detector can be found for the case when the marked coefficients are modeled by the notched power exponential distribution. The two hypotheses are stated as

$$\begin{aligned} H_0 \quad &: \quad s = 0 \\ & \quad\;\; y_i = x_i, \quad 1 \le i \le N \end{aligned}$$

$$\begin{aligned} H_1 \quad &: \quad s = s^* \\ & \quad\;\; y_i = x_i \left(1 + s^* m_i\right), \quad 1 \le i \le N \end{aligned}$$

where N is the number of (high-magnitude) coefficients that are watermarked. The distribution under H_0 is

$$p_0(y_i) = \begin{cases} \frac{C}{1-\kappa} \exp\left\{ -\left|\frac{y_i}{\alpha}\right|^\beta \right\} & \text{if } |y_i| > \delta \\ 0 & \text{else} \end{cases}$$

$$p_0(\mathbf{y}) = \begin{cases} \displaystyle\prod_{i=1}^{N} \frac{C}{1-\kappa} \exp\left\{ -\left|\frac{y_i}{\alpha}\right|^\beta \right\} & \text{if } |y_i| > \delta \ \forall i \\ 0 & \text{else.} \end{cases}$$

The distribution under H_1 is

$$\begin{aligned} p_1(y_i) &= \frac{1}{1+s^* m_i} p_{x_i} \left(\frac{y_i}{1+s^* m_i} \right) \\ &= \begin{cases} \frac{C}{(1-\kappa)(1+s^* m_i)} \exp\left\{ -\left|\frac{y_i}{\alpha(1+s^* m_i)}\right|^\beta \right\} & \text{if } |y_i| > \delta_i \\ 0 & \text{else} \end{cases} \end{aligned}$$

$$p_1(\mathbf{y}) = \begin{cases} \displaystyle\prod_{i=1}^{N} \frac{C}{(1-\kappa)(1+s^* m_i)} \exp\left\{ -\left|\frac{y_i}{\alpha\left(1+s^* m_i\right)}\right|^\beta \right\} & \text{if } |y_i| > \delta_i \ \forall i \\ 0 & \text{else} \end{cases}$$

where C is defined in (2.3), and $\delta_i = \delta(1 + s^* m_i)$.

The likelihood ratio is given by

$$L(\mathbf{y}) = \frac{p_1(\mathbf{y})}{p_0(\mathbf{y})}.$$

Due to the conditions in the definition of the two distributions, the likelihood ratio is also defined conditionally as follows,

$$L(\mathbf{y}) = \begin{cases} \displaystyle\prod_{i=1}^{N} \frac{1}{1+s^* m_i} \exp\left\{ -\left|\frac{y_i}{\alpha\,(1+s^* m_i)}\right|^{\beta} + \left|\frac{y_i}{\alpha}\right|^{\beta} \right\} & \text{if } (|y_i| > \delta_i \ \forall i) \text{ and } (|y_i| > \delta \ \forall i) \\ \infty & \text{if } (|y_i| > \delta_i \ \forall i) \text{ and} \\ & \quad (\exists i : |y_i| \leq \delta \text{ and } m_i = -1) \\ 0 & \text{if } (\exists i : |y_i| \leq \delta_i \text{ and } m_i = 1) \text{ and} \\ & \quad (|y_i| > \delta \ \forall i) \\ \displaystyle\prod_{i=1}^{N} \frac{1}{1+s^* m_i} \exp\left\{ -\left|\frac{y_i}{\alpha\,(1+s^* m_i)}\right|^{\beta} + \left|\frac{y_i}{\alpha}\right|^{\beta} \right\} & \text{else.} \end{cases}$$

Thus the log-likelihood ratio is

$$\ln L(\mathbf{y}) = \begin{cases} \displaystyle\sum_{i=1}^{N} \left(\ln \frac{1}{1+s^* m_i} - \left|\frac{y_i}{\alpha\,(1+s^* m_i)}\right|^{\beta} + \left|\frac{y_i}{\alpha}\right|^{\beta} \right) & \text{if } (|y_i| > \delta_i \ \forall i) \text{ and } (|y_i| > \delta \ \forall i) \\ \infty & \text{if } (|y_i| > \delta_i \ \forall i) \text{ and} \\ & \quad (\exists i : |y_i| \leq \delta \text{ and } m_i = -1) \\ -\infty & \text{if } (\exists i : |y_i| \leq \delta_i \text{ and } m_i = 1) \text{ and} \\ & \quad (|y_i| > \delta \ \forall i) \\ \displaystyle\sum_{i=1}^{N} \left(\ln \frac{1}{1+s^* m_i} - \left|\frac{y_i}{\alpha\,(1+s^* m_i)}\right|^{\beta} + \left|\frac{y_i}{\alpha}\right|^{\beta} \right) & \text{else.} \end{cases} \tag{3.3}$$

These results are quite intuitive. In the first case, when the y_i are in the appropriate ranges, the statistic is the same as in the case of the power exponential. When one of the distribution conditions is met but the other is not, one of the hypotheses is impossible, which is clearly indicated by the presence of infinities in the statistic. Finally, when the range conditions fail in both hypotheses (a zero-probability event), the regular statistic is again calculated.

Figure 3.1 illustrates the four regions present in the decision statistic. Note that Regions 1 and 4 can occur for either $m_i = -1$ or $m_i = 1$, while Region 2 only occurs when $m_i = -1$, and Region 3 only occurs when $m_i = 1$.

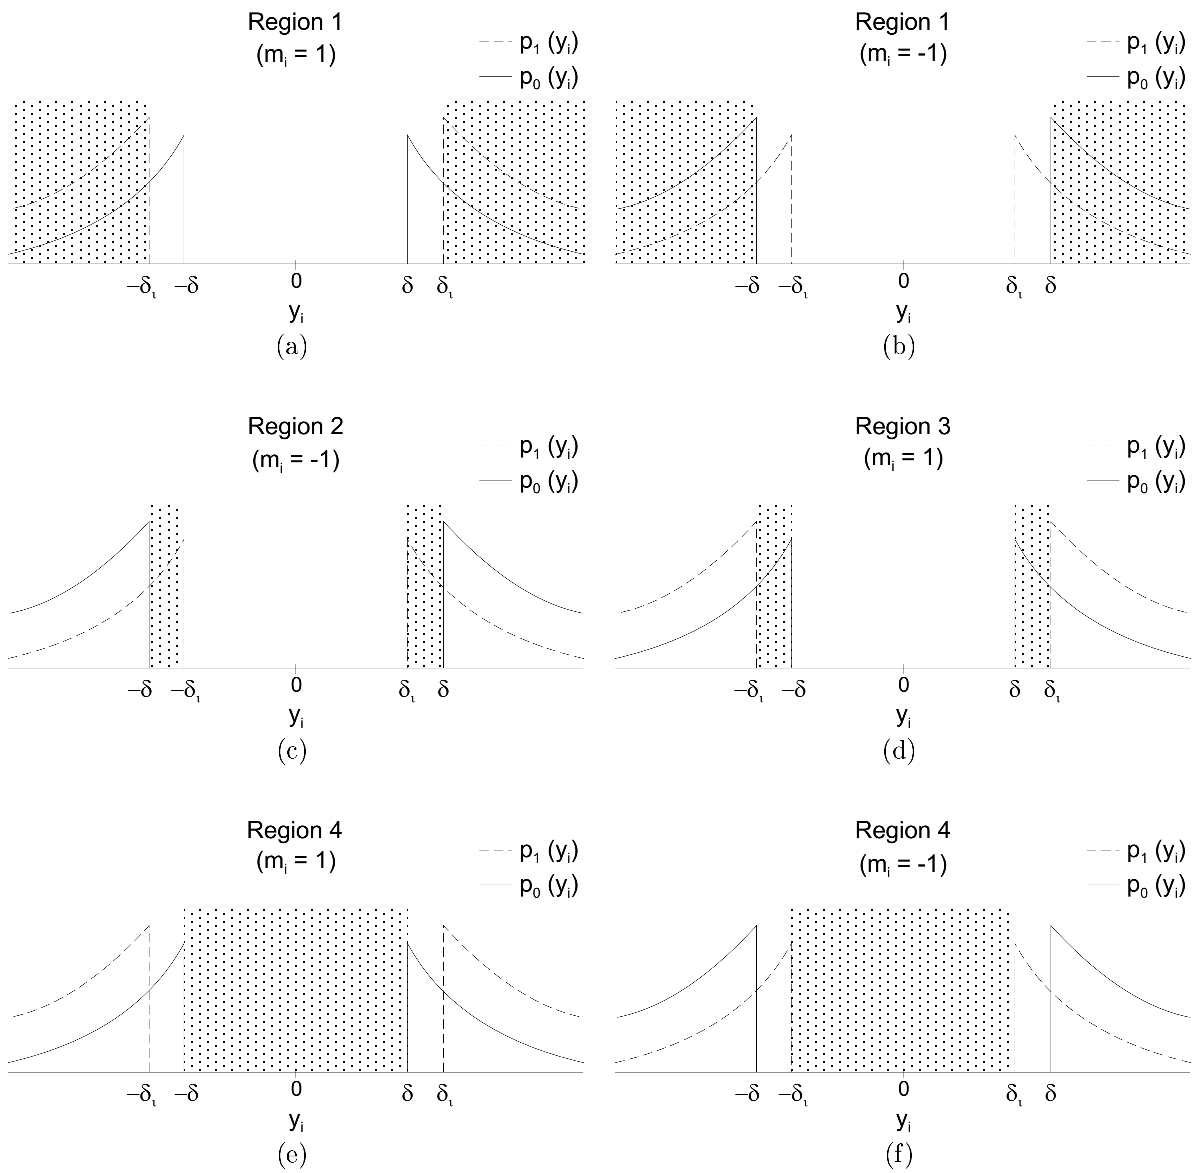

Figure 3.1 Four regions utilized in the likelihood ratio test statistic for the notched power exponential distribution.

3.2 Simple Detector for Notched Distributions

The likelihood ratio test detector developed for the notched power exponential distribution reveals a key benefit resulting from the use of a notched distribution: it is sometimes possible to decisively eliminate one of the two hypotheses simply by comparing the received data samples to the notch thresholds. Using this idea, a simplified version of the likelihood ratio test detector can be developed for any notched distribution. Consider the two hypotheses presented to the LRT detector,

$$\begin{aligned} H_0 \quad &: \quad s = 0 \\ & \quad\ \ y_i = x_i, \quad 1 \leq i \leq N \end{aligned}$$

$$\begin{aligned} H_1 \quad &: \quad s = s^* \\ & \quad\ \ y_i = x_i \left(1 + s^* m_i\right), \quad 1 \leq i \leq N \end{aligned}$$

where N is the number of (high-magnitude) coefficients that are watermarked. As seen in the previous development, a decision as to the presence of the watermark can be made with certainty if a single data sample, y_i, is found to lie either in Region 2 or 3, as shown in Figure 3.1 (c) and (d). Now, define $\mathcal{E}_2$ and $\mathcal{E}_3$ to be the events that either of these cases occurs. More formally,

$$\mathcal{E}_2 = \{\mathbf{y} : \delta_i < |y_i| < \delta \text{ and } m_i = -1 \text{ for some } i\}$$

and

$$\mathcal{E}_3 = \{\mathbf{y} : \delta < |y_i| < \delta_i \text{ and } m_i = 1 \text{ for some } i\}.$$

Now, the probability that $\mathcal{E}_3$ occurs under H_0 is found,

$$\begin{aligned} P_0[\mathcal{E}_3] &= P_0[\exists i : \delta < |y_i| < \delta_i] \\ &= 1 - \prod_{i=1}^{\frac{N}{2}} P_0[|y_i| > \delta(1 + s^*)] \\ &= 1 - \left(2P_0(-\delta(1 + s^*))\right)^{\frac{N}{2}} \end{aligned}$$

where $P_0(z) = \int_{-\infty}^{z} p_0(x)dx$ is the cumulative density function of $p_0(x)$. Note that in the product, only $\frac{N}{2}$ terms are included because $\mathcal{E}_3$ only deals with $m_i = 1$. Now, since for a notched distribution, $P_0(-\delta) = \frac{1}{2}$, and since $-\delta(1+s^*) < -\delta$, it is clear that $0 < P_0(\delta(1+s^*)) < \frac{1}{2}$. Thus, $0 < 2P_0(-\delta(1+s^*)) < 1$, implying that $(2P_0(-\delta(1+s^*)))^{\frac{N}{2}}$ approaches 0 as N approaches infinity. Consequently, $P_0[\mathcal{E}_3]$ approaches 1 as N approaches infinity. In other words, if no watermark is present, the probability of a data point falling in Region 3 goes to 1 as N goes to infinity.

Similarly, the probability that $\mathcal{E}_2$ occurs under H_1 is found as

$$\begin{aligned} P_1[\mathcal{E}_2] &= P_1[\exists i : \delta_i < |y_i| < \delta] \\ &= 1 - \prod_{i=1}^{\frac{N}{2}} P_1[|y_i| > \delta] \\ &= 1 - (2P_1(-\delta))^{\frac{N}{2}} \end{aligned}$$

where $P_1(z) = \int_{-\infty}^{z} p_1(x)dx$ is the cumulative density function of $p_1(x)$. Once again, only $\frac{N}{2}$ terms have been included in the product because $\mathcal{E}_2$ can only occur when $m_i = -1$. By an argument similar to that given above, $(2P_1(-\delta))^{\frac{N}{2}}$ approaches 0 as N approaches infinity. Hence, $P_1[\mathcal{E}_2]$ approaches 1 as N approaches infinity. In other words, if a watermark is present, the probability of a data point lying in Region 2 goes to 1 as N goes to infinity.

With this knowledge of the asymptotic behavior of $\mathbf{y}$, a *simple detector* (SD) can be designed. Using the same four regions defined for the LRT detector, the simple detector is given by

$$\mathcal{D}(\mathbf{y}) = \begin{cases} \text{choose } H_0 \text{ with probability } \lambda, & \text{if } (|y_i| > \delta_i \ \forall i) \text{ and } (|y_i| > \delta \ \forall i) \\ \qquad H_1 \text{ with probability } 1-\lambda & \\ \text{choose } H_1 & \text{if } (|y_i| > \delta_i \ \forall i) \text{ and} \\ & \quad (\exists i : |y_i| \le \delta \text{ and } m_i = -1) \\ \text{choose } H_0 & \text{if } (\exists i : |y_i| \le \delta_i \text{ and } m_i = 1) \\ & \quad \text{and } (|y_i| > \delta \ \forall i) \\ \text{choose } H_0 \text{ with probability } \lambda, & \text{else} \\ \qquad H_1 \text{ with probability } 1-\lambda & \end{cases} \tag{3.4}$$

where λ is a parameter that permits preference to be given to a specific hypothesis in the case of Regions 1 and 4. From the above derivation, it is known that, asymptotically, the middle two cases in the decision statistic will occur with certainty. Hence, this simple detector is equivalent to the optimal LRT detector for notched distributions, as N goes to infinity. Finally, it should be noted that a definite yes or no decision is produced by the detector in each of the possible cases.

3.3 Maximum Likelihood Estimation

The previous detector employed binary hypothesis testing using a likelihood ratio to determine, for a known embedding strength, s, the likelihood that a specific watermark is present in a signal. Another method of watermark detection is to estimate the embedding strength given the signal data [4]. This estimate will indicate the strength of presence of a specific watermark in a set of data. Then, the estimate is used as the decision statistic in a binary hypothesis test to determine the presence of the watermark. However, it should be noted that this use of hypothesis testing is suboptimal because a likelihood ratio test is the optimal binary hypothesis test.

The assumption of an unknown embedding strength that is made in the derivation of this detector is often motivated as a simple means of combating an attack on the watermarked data. Although the actual embedding strength is known, an attack can be modeled, in a simplistic sense, as a change in this parameter. Then, the strength is no longer known at the detector and, thus, must be estimated. This approach is not an optimal means of incorporating an attack and, hence, results in a suboptimal detector. However, this detector has been fairly widely studied and facilitates interesting comparisons with the likelihood ratio test detector.

To develop the estimation detector, the technique utilized to perform the estimation given the signal data is first defined. As the name suggests, maximum likelihood estimation is a means of estimating the most likely value for a parameter on the basis of knowledge of the distribution of the data given this parameter. Thus, the detector should estimate a value for s knowing the distribution of $\mathbf{y}$ given s, $p_s(\mathbf{y})$. The maximum likelihood estimate is given by

$$\hat{s} = \underset{s}{\operatorname{argmax}}\, p_s(\mathbf{y}).$$

To proceed, the function $p_s(\mathbf{y})$ is maximized over s to find the MLE. The resultant maximizing value for s will be the maximum likelihood estimate, $\hat{s}$.

Now, a binary hypothesis test is employed to provide an answer stating whether or not the watermark is present. The decision statistic, $\hat{s}$, is compared against a threshold, τ, and the watermark is declared to be present if $\hat{s} > \tau$. Otherwise, the detector states that no watermark is present. To select τ, the distribution of the output statistic (MLE) can be determined and then, for example, a Neyman-Pearson constraint on the probability of false alarm could be used to solve for τ. In any case, τ should be between 0 and 1 because it represents a cutoff in the embedding strength, which is known to be between 0 and 1.

In summary, the maximum likelihood estimation-based detector computes the MLE for s on the basis of the data vector, yielding a measure of the strength of presence of the specific watermark. This output is then compared against a threshold to determine a yes or no answer for the presence of the watermark.

3.3.1 Weibull distribution

The first distribution considered for the maximum likelihood estimation detector is the Weibull distribution [4]. From the likelihood ratio test section, the required distribution of $\mathbf{y}$ given s is found to be

$$\begin{aligned} p_s(\mathbf{y}) &= \prod_{i=1}^{N} \frac{\beta}{\alpha(1+sm_i)} \left(\frac{y_i}{\alpha(1+sm_i)} \right)^{\beta-1} \exp\left\{ -\left(\frac{y_i}{\alpha(1+sm_i)} \right)^{\beta} \right\} \\ &= \frac{\beta^N}{\alpha^{\beta N}} \prod_{i=1}^{N} \frac{1}{(1+sm_i)} \left(\frac{y_i}{1+sm_i} \right)^{\beta-1} \exp\left\{ -\left(\frac{y_i}{\alpha(1+sm_i)} \right)^{\beta} \right\}. \end{aligned}$$

Now, the natural logarithm of this distribution is found to simplify calculations,

$$\begin{aligned} \ln p_s(\mathbf{y}) &= N\ln\beta - \beta N \ln\alpha + \sum_{i=1}^{N} \left[\ln\left(\frac{y_i^{\beta-1}}{(1+sm_i)^{\beta}} \right) - \left(\frac{y_i}{\alpha(1+sm_i)} \right)^{\beta} \right] \\ &= N\ln\beta - \beta N \ln\alpha + \sum_{i=1}^{N} \left[(\beta-1)\ln y_i - \beta\ln(1+sm_i) - \left(\frac{y_i}{\alpha(1+sm_i)} \right)^{\beta} \right]. \end{aligned}$$

To maximize this expression, its derivative is set to 0 and then solved for $\hat{s}$, the maximizing value of s; hence,

$$0 = \left. \frac{\partial \ln p_s(\mathbf{y})}{\partial s} \right|_{s=\hat{s}}$$

$$= \sum_{i=1}^{N}\left[\frac{-\beta m_i}{1+\hat{s}m_i} + \frac{y_i^\beta}{\alpha^\beta}\left(\frac{\beta m_i}{(1+\hat{s}m_i)^{\beta+1}}\right)\right].$$

Since s is small, the approximation of a first-order Taylor series is used, yielding

$$\begin{aligned}
0 &= \sum_{i=1}^{N}\left[\beta m_i(1-\hat{s}m_i) + \frac{\beta m_i y_i^\beta}{\alpha^\beta}(1-(\beta+1)\hat{s}m_i)\right] \\
&= \sum_{i=1}^{N}\left[-\beta m_i + \beta\hat{s}m_i^2 + \frac{\beta m_i y_i^\beta}{\alpha^\beta} - (\beta+1)\left(\frac{\beta\hat{s}m_i^2 y_i^2}{\alpha^\beta}\right)\right] \\
\sum_{i=1}^{N}\left[\beta m_i - \frac{\beta m_i y_i^\beta}{\alpha^\beta}\right] &= \sum_{i=1}^{N}\left[\left(\beta m_i^2 - (\beta+1)\left(\frac{\beta m_i^2 y_i^2}{\alpha^\beta}\right)\right)\hat{s}\right] \\
\hat{s} &= \frac{\sum_{i=1}^{N}\left[\beta m_i - \frac{\beta m_i y_i^\beta}{\alpha^\beta}\right]}{\sum_{i=1}^{N}\left[\beta m_i^2 - (\beta+1)\left(\frac{\beta m_i^2 y_i^2}{\alpha^\beta}\right)\right]} \\
&= \frac{\sum_{i=1}^{N} y_i^\beta m_i}{(\beta+1)\sum_{i=1}^{N} y_i^\beta - N\alpha^\beta}
\end{aligned}$$

where it was noted that $m_i \in \{-1, 1\}$ so $m_i^2 = 1$, and that $\sum_{i=1}^{N} m_i = 0$ because m_i is a zero-mean sequence. Thus, the MLE for s is given by

$$\hat{s} = \frac{\sum_{i=1}^{N} y_i^\beta m_i}{(\beta+1)\sum_{i=1}^{N} y_i^\beta - N\alpha^\beta}. \tag{3.5}$$

Hence, an expression for the estimated embedding strength of the given watermark has been found. The detector can now compare this value against an appropriate threshold, for example, one based on a Neyman-Pearson constraint, to decide whether or not the specific watermark is present in the data.

3.3.2 Power exponential distribution

As expected, the derivation of the maximum likelihood estimation-based detector for the case of power exponentially distributed data is quite similar to that for Weibull distributed data. Again, the maximum likelihood estimate of s is given by

$$\hat{s} = \underset{s}{\operatorname{argmax}}\, p_s(\mathbf{y})$$

where $p_s(\mathbf{y})$ is the distribution of the data, $\mathbf{y}$, given the strength parameter, s. From the likelihood ratio test section for the power exponential distribution, this conditional distribution is

$$p_s(\mathbf{y}) = \prod_{i=1}^{N} \frac{C}{1+sm_i} \exp\left\{ - \left| \frac{y_i}{\alpha\,(1+sm_i)} \right|^{\beta} \right\}.$$

This function is now maximized over s to find the MLE, noting again that it is equivalent to maximize $\ln p_s(\mathbf{y})$ because $\ln(x)$ is an increasing function:

$$\begin{aligned}
\ln p_s(\mathbf{y}) &= N \ln C - \sum_{i=1}^{N} \ln(1+sm_i) - \sum_{i=1}^{N} \left| \frac{y_i}{\alpha(1+sm_i)} \right|^{\beta} \\
&= N \ln C - \sum_{i=1}^{N} \ln(1+sm_i) - \sum_{i=1}^{N} \left| \frac{y_i}{\alpha} \right|^{\beta} \frac{1}{(1+sm_i)^{\beta}}.
\end{aligned}$$

The derivative of this expression is now set to 0 and solved for $\hat{s}$, the maximizing value of s:

$$\begin{aligned}
0 &= \left. \frac{\partial \ln p_s(\mathbf{y})}{\partial s} \right|_{s=\hat{s}} \\
&= -\sum_{i=1}^{N} \frac{m_i}{1+\hat{s}m_i} + \sum_{i=1}^{N} \left| \frac{y_i}{\alpha} \right|^{\beta} \frac{\beta m_i}{(1+\hat{s}m_i)^{\beta+1}}.
\end{aligned}$$

Since s is small, the approximation of a first-order Taylor series is again used, yielding

$$\begin{aligned}
0 &= \sum_{i=1}^{N} \left[-m_i(1-\hat{s}m_i) + \beta m_i \left| \frac{y_i}{\alpha} \right|^{\beta} (1-(\beta+1)\hat{s}m_i) \right] \\
&= \sum_{i=1}^{N} \left[-m_i + \hat{s}m_i^2 + \beta m_i \left| \frac{y_i}{\alpha} \right|^{\beta} - \beta m_i \left| \frac{y_i}{\alpha} \right|^{\beta} (\beta+1)\hat{s}m_i \right] \\
&= \sum_{i=1}^{N} \left[\hat{s} + \beta m_i \left| \frac{y_i}{\alpha} \right|^{\beta} - \beta \left| \frac{y_i}{\alpha} \right|^{\beta} (\beta+1)\hat{s} \right] \\
&= \sum_{i=1}^{N} \left[\beta m_i \left| \frac{y_i}{\alpha} \right|^{\beta} + \left(1 - \beta \left| \frac{y_i}{\alpha} \right|^{\beta} (\beta+1) \right) \hat{s} \right] \\
\sum_{i=1}^{N} \beta m_i \left| \frac{y_i}{\alpha} \right|^{\beta} &= \sum_{i=1}^{N} \left(\beta \left| \frac{y_i}{\alpha} \right|^{\beta} (\beta+1) - 1 \right) \hat{s}
\end{aligned}$$

$$
\begin{aligned}
\hat{s} &= \frac{\displaystyle\sum_{i=1}^{N} \beta m_i \left|\frac{y_i}{\alpha}\right|^{\beta}}{\displaystyle\sum_{i=1}^{N} \left(\beta \left|\frac{y_i}{\alpha}\right|^{\beta} (\beta+1) - 1\right)} \\
&= \frac{\displaystyle\sum_{i=1}^{N} |y_i|^{\beta} m_i}{\displaystyle(\beta+1)\sum_{i=1}^{N} |y_i|^{\beta} - \frac{N\alpha^{\beta}}{\beta}}
\end{aligned}
$$

where it is noted that $m_i \in \{-1, 1\}$ so $m_i^2 = 1$, and that $\displaystyle\sum_{i=1}^{N} m_i = 0$ because m_i is a zero-mean sequence. Thus, the MLE for s is given by

$$
\hat{s} = \frac{\displaystyle\sum_{i=1}^{N} |y_i|^{\beta} m_i}{\displaystyle(\beta+1)\sum_{i=1}^{N} |y_i|^{\beta} - \frac{N\alpha^{\beta}}{\beta}}. \tag{3.6}
$$

As in the case of the Weibull distribution, Neyman-Pearson constraints can be utilized to provide a binary detection response based on the estimate.

Similar to the likelihood ratio test, the MLE statistic for the power exponential distribution bears resemblance to that for the Weibull distribution, Equation (3.5). The major difference is the presence of absolute values in the power exponential case, reflecting its inclusion of both positive and negative valued data.

3.3.3 Notched power exponential distribution

The maximum likelihood estimate is now determined for the case of the notched power exponential distribution. From the likelihood ratio testing section, the distribution of the output data $\mathbf{y}$ given s for this distribution is

$$
p_s(\mathbf{y}) = \begin{cases} \displaystyle\prod_{i=1}^{N} \frac{C}{(1-\kappa)(1+sm_i)} \exp\left\{-\left|\frac{y_i}{\alpha\,(1+sm_i)}\right|^{\beta}\right\} & \text{if } |y_i| > \delta_i \ \forall i \\ 0 & \text{else }. \end{cases}
$$

To find the MLE, this distribution should be maximized over the parameter s. However, since the distribution is defined over two ranges where the boundary is a function of s (because $\delta_i = \delta(1 + sm_i)$), the maximization problem becomes more complicated.

To maximize $p_s(\mathbf{y})$, the maximizing value of s, $\hat{s}$, must ensure that $|y_i| > \delta_i\ \forall i$. This requirement will present a restriction on the range of possible values for $\hat{s}$. The data $\mathbf{y}$ is split into two groups depending on the value of m_i. Let $\mathcal{I}^+$ denote the set of indices for which $m_i = 1$ and let $\mathcal{I}^-$ denote the set of indices for which $m_i = -1$. Then,

$$\begin{aligned}
|y_i| &> \delta_i = \delta(1+\hat{s}) \quad \forall i \in \mathcal{I}^+ \\
\hat{s} &< \frac{|y_i|}{\delta} - 1 \quad \forall i \in \mathcal{I}^+ \\
\hat{s} &< \frac{\min\limits_{i\in\mathcal{I}^+} |y_i|}{\delta} - 1
\end{aligned}$$

which provides an upper bound on $\hat{s}$. Similarly,

$$\begin{aligned}
|y_i| &> \delta_i = \delta(1-\hat{s}) \quad \forall i \in \mathcal{I}^- \\
\hat{s} &> 1 - \frac{|y_i|}{\delta} \quad \forall i \in \mathcal{I}^- \\
\hat{s} &> 1 - \frac{\min\limits_{i\in\mathcal{I}^-} |y_i|}{\delta}
\end{aligned}$$

which provides a lower bound on $\hat{s}$. Combining these restrictions with the known range of embedding strengths yields an overall bound of

$$\begin{cases}
1 - \frac{\min\limits_{i\in\mathcal{I}^-}|y_i|}{\delta} < \hat{s} < \min\left(1, \frac{\min\limits_{i\in\mathcal{I}^+}|y_i|}{\delta} - 1\right) & \text{if } 1 - \frac{\min\limits_{i\in\mathcal{I}^-}|y_i|}{\delta} > 0 \\
0 \le \hat{s} < \min\left(1, \frac{\min\limits_{i\in\mathcal{I}^+}|y_i|}{\delta} - 1\right) & \text{else.}
\end{cases}$$

Let this range of acceptable $\hat{s}$ values be denoted by $\mathcal{S}$. Therefore, the MLE is given by

$$\hat{s} = \operatorname*{argmax}_{s\in\mathcal{S}} \prod_{i=1}^{N} \frac{C}{(1-\kappa)(1+sm_i)} \exp\left\{ - \left| \frac{y_i}{\alpha\,(1+sm_i)} \right|^{\beta} \right\}. \tag{3.7}$$

This expression can be solved numerically to yield the estimate of the embedding strength. As before, the resulting estimate may now be compared against a threshold in the range [0, 1) to yield a binary output from the detector.

3.4 Locally Optimal Detector

The final detector considered makes use of the knowledge that s takes on small values. Thus, a locally most powerful test is considered, as described by Poor [7]. This test finds a decision rule for the case when the data distributions under the two hypotheses are known to be quite close together. To apply this type of test to the watermark detection problem, a composite binary hypothesis test is formulated. A composite binary hypothesis test is one in which at least one of the hypotheses includes a range of values for the variable in question. In this case, hypothesis 0 states the embedding strength of the watermark in the data is s_0, which here will be taken to be 0, meaning no watermark is present. Hypothesis 1 states that the watermark is present with some non-zero strength.

Then, the locally most powerful test is given by

$$\underset{\mathcal{D}}{\operatorname{argmax}}\, P_D'(\mathcal{D}, s_0 = 0)$$

subject to the constraint that $P_F(\mathcal{D}, s_0) \leq \rho$, where P_D is the probability of detection (the probability of correctly detecting the watermark when it is present), P_F is the probability of false alarm, and $\mathcal{D}$ is a decision rule. The locally optimal test can be formulated as a modified likelihood ratio test in which $p_1(\mathbf{y})$ in the standard ratio is replaced by $\left.\frac{\partial p_s(\mathbf{y})}{\partial s}\right|_{s=s_0}$.

This modified likelihood ratio will yield a decision statistic. As seen previously, a comparison threshold can be found by considering the distribution of the statistic, and then applying a ρ-level Neyman-Pearson test, for example, to determine an appropriate threshold, τ. Using this threshold, the detector can form a binary answer for the presence of the watermark.

In summary, the detector based upon a locally optimal test is applied by computing a decision statistic, based on a modified likelihood ratio, which indicates the strength of presence of the specific watermark in the data. The statistic can then be compared against a threshold to yield a binary output.

3.4.1 Weibull distribution

The locally optimal detector is first developed for the situation when the host data is modeled by the Weibull distribution. Here, the two hypotheses are given by

$$\begin{aligned} H_0 \quad &: \quad s = 0 = s_0 \\ & \quad\;\; y_i = x_i, \quad 1 \leq i \leq N \end{aligned}$$

$$
\begin{aligned}
H_1 \quad &: \quad s > 0 \\
& \qquad y_i = x_i \left(1 + s m_i\right), \quad 1 \leq i \leq N.
\end{aligned}
$$

The locally most powerful test can now be formulated as a modified likelihood ratio test. Thus, the partial derivative $\left.\frac{\partial p_s(\mathbf{y})}{\partial s}\right|_{s=s_0}$ is required:

$$
\begin{aligned}
p_s(\mathbf{y}) &= \frac{\beta^N}{\alpha^{\beta N}} \prod_{i=1}^{N} \frac{y_i^{\beta-1}}{(1+sm_i)^{\beta}} \exp\left\{-\left(\frac{y_i}{\alpha(1+sm_i)}\right)^{\beta}\right\} \\
\frac{\partial p_s(\mathbf{y})}{\partial s} &= \frac{\beta^N}{\alpha^{\beta N}} \sum_{i=1}^{N} \left(\left[\frac{-\beta y_i^{\beta} m_i}{(1+sm_i)^{\beta}} \exp\left\{-\left(\frac{y_i}{\alpha(1+sm_i)}\right)^{\beta}\right\}\right.\right. \\
&\qquad \left. + \frac{y_i^{\beta} m_i}{\alpha^{\beta}(1+sm_i)^{\beta+1}} \exp\left\{-\left(\frac{y_i}{\alpha(1+sm_i)}\right)^{\beta}\right\} \frac{y_i^{\beta-1}}{(1+sm_i)^{\beta}}\right] \\
&\qquad \left. \times \prod_{j=1, j\neq i}^{N} \frac{y_j^{\beta-1}}{(1+sm_j)} \exp\left\{-\left(\frac{y_j}{\alpha(1+sm_j)}\right)^{\beta}\right\}\right) \\
\left.\frac{\partial p_s(\mathbf{y})}{\partial s}\right|_{s=0} &= \frac{\beta^N}{\alpha^{\beta N}} \sum_{i=1}^{N} \left(\left[-\beta y_i^{\beta} m_i \exp\left\{-\left(\frac{y_i}{\alpha}\right)^{\beta}\right\} + \frac{y_i^{\beta} m_i}{\alpha^{\beta}} \exp\left\{-\left(\frac{y_i}{\alpha}\right)^{\beta}\right\} y_i^{\beta-1}\right]\right. \\
&\qquad \left. \times \prod_{j=1, j\neq i}^{N} y_j^{\beta-1} \exp\left\{-\left(\frac{y_j}{\alpha}\right)^{\beta}\right\}\right).
\end{aligned}
$$

Now, recall that

$$
p_0(\mathbf{y}) = \frac{\beta^N}{\alpha^{\beta N}} \prod_{i=1}^{N} \left(\frac{y_i}{\alpha}\right)^{\beta} \exp\left\{-\left(\frac{y_i}{\alpha}\right)^{\beta}\right\}.
$$

Hence, the modified likelihood ratio is given by

$$
\begin{aligned}
L(\mathbf{y}) &= \frac{\left.\frac{\partial p_s(\mathbf{y})}{\partial s}\right|_{s=0}}{p_0(\mathbf{y})} \\
&= \sum_{i=1}^{N} \frac{\left[-\beta y_i^{\beta-1} m_i \exp\left\{-\left(\frac{y_i}{\alpha}\right)^{\beta}\right\} + \frac{y_i^{\beta} m_i}{\alpha^{\beta}} \exp\left\{-\left(\frac{y_i}{\alpha}\right)^{\beta}\right\} y_i^{\beta-1}\right]}{\left(\frac{y_i}{\alpha}\right)^{\beta-1} \exp\left\{-\left(\frac{y_i}{\alpha}\right)^{\beta}\right\}} \\
&= \sum_{i=1}^{N} (-\beta m_i \alpha^{\beta-1} + y_i^{\beta} m_i \alpha^{-1}) \\
&= \sum_{i=1}^{N} y_i^{\beta} m_i \alpha^{-1} \qquad (3.8)
\end{aligned}
$$

where, as in the development of the MLE, the fact that $\mathbf{m}$ is a zero-mean sequence has again been used. A reasonable threshold against which to compare the output statistic may now be determined in order to yield a binary detector response.

It is interesting to note the relation between this statistic and the commonly used correlation detector [2], whose output is given by $\sum_{i=1}^{N} y_i m_i$. Since the statistic given in Equation (3.8) has locally optimal properties, it is expected that the detector based upon it will have a lower probability of error than the standard correlation detector. This improvement occurs because, for the locally optimal case, the separation of the distribution of the statistic for when the watermark is present and when it is not present, is greater than in the basic correlator case.

3.4.2 Power exponential distribution

In the case of the locally optimal detector for power exponentially distributed data [8], the composite hypothesis test is formulated as follows:

$$
\begin{aligned}
H_0 \quad &: \quad s = 0 = s_0 \\
&\quad\;\; y_i = x_i, \quad 1 \leq i \leq N \\
\\
H_1 \quad &: \quad s > 0 \\
&\quad\;\; y_i = x_i \left(1 + s m_i\right), \quad 1 \leq i \leq N.
\end{aligned}
$$

Now, the partial derivative required in the modified likelihood ratio is given by

$$
\begin{aligned}
p_s(\mathbf{y}) &= \prod_{i=1}^{N} \frac{C}{1 + s m_i} \exp\left\{ -\frac{|y_i|^{\beta}}{\alpha^{\beta}(1 + s m_i)^{\beta}} \right\} \\
\frac{\partial p_s(\mathbf{y})}{\partial s} &= p_s(\mathbf{y}) \frac{\partial \ln p_s(\mathbf{y})}{\partial s} \\
&= \left(\prod_{j=1}^{N} \frac{C}{1 + s m_j} \exp\left\{ - \left| \frac{y_j}{\alpha(1 + s m_j)} \right|^{\beta} \right\} \right) \\
&\quad \times \left(- \sum_{i=1}^{N} \frac{m_i}{1 + s m_i} + \sum_{i=1}^{N} \left| \frac{y_i}{\alpha} \right|^{\beta} \frac{\beta m_i}{(1 + s m_i)^{\beta+1}} \right).
\end{aligned}
$$

Recalling

$$p_0(\mathbf{y}) = \prod_{i=1}^{N} C \exp\left\{ - \left| \frac{y_i}{\alpha} \right|^{\beta} \right\},$$

the modified likelihood ratio is given by

$$\begin{aligned} L(\mathbf{y}) &= \frac{\left.\frac{\partial p_s(\mathbf{y})}{\partial s}\right|_{s=s_0}}{p_0(\mathbf{y})} \\ &= \frac{\prod_{i=1}^{N} C \exp\left\{ - \left| \frac{y_i}{\alpha} \right|^{\beta} \right\} \sum_{i=1}^{N} \left| \frac{y_i}{\alpha} \right|^{\beta} \beta m_i}{\prod_{i=1}^{N} C \exp\left\{ - \left| \frac{y_i}{\alpha} \right|^{\beta} \right\}} \\ &= \frac{\beta}{\alpha^{\beta}} \sum_{i=1}^{N} |y_i|^{\beta} m_i. \end{aligned} \tag{3.9}$$

Once again, a threshold can now be utilized so that the detector provides a yes or no response.

Again, a strong parallel is seen between this statistic and that for the Weibull distribution under locally optimal detection, which is given in Equation (3.8). The only differences between the two are the presence of the absolute value around the y_i for the case of the power exponential distribution and a scaling factor.

3.4.3 Notched power exponential distribution

The final distribution considered for the locally optimal detector is the notched power exponential distribution. To begin, the composite hypothesis test is formulated as follows,

$$\begin{aligned} H_0 &: s = 0 = s_0 \\ & \quad y_i = x_i, \quad 1 \le i \le N \\ H_1 &: s > 0 \\ & \quad y_i = x_i\,(1 + s m_i)\,, \quad 1 \le i \le N. \end{aligned}$$

Here, the distribution of the output conditioned on the embedding strength is given by

$$p_s(\mathbf{y}) = \begin{cases} \prod_{i=1}^{N} \frac{C}{(1-\kappa)(1+s m_i)} \exp\left\{ - \left| \frac{y_i}{\alpha\,(1+s m_i)} \right|^{\beta} \right\} & \text{if } |y_i| > \delta_i \ \forall i \\ 0 & \text{else.} \end{cases}$$

To be able to perform the required derivative with respect to s, the condition $|y_i| > \delta_i$ should be written explicitly in terms of s. As in the case of maximum likelihood estimation, the vector $\mathbf{y}$ is divided into two groups. Let $\mathcal{I}^+$ denote the set of indices for which $m_i = 1$, and let $\mathcal{I}^-$ denote the set of indices for which $m_i = -1$. Then,

$$\begin{aligned} |y_i| &> \delta(1+s) \ \ \forall i \in \mathcal{I}^+ \\ \frac{|y_i|}{\delta} &> 1+s \ \ \forall i \in \mathcal{I}^+ \\ s &< \frac{|y_i|}{\delta} - 1 \ \ \forall i \in \mathcal{I}^+ \\ s &< \frac{\min\limits_{i\in\mathcal{I}^+} |y_i|}{\delta} - 1. \end{aligned}$$

Similarly,

$$\begin{aligned} |y_i| &> \delta(1-s) \ \ \forall i \in \mathcal{I}^- \\ \frac{|y_i|}{\delta} &> 1-s \ \ \forall i \in \mathcal{I}^- \\ s &> 1 - \frac{|y_i|}{\delta} \ \ \forall i \in \mathcal{I}^- \\ s &> 1 - \frac{\min\limits_{i\in\mathcal{I}^-} |y_i|}{\delta}. \end{aligned}$$

Combining these bounds and the known bound on s, $[0,\,1)$, provides the overall condition

$$\begin{cases} 1 - \frac{\min\limits_{i\in\mathcal{I}^-} |y_i|}{\delta} < \hat{s} < \min\left(1, \frac{\min\limits_{i\in\mathcal{I}^+} |y_i|}{\delta} - 1\right) & \text{if } 1 - \frac{\min\limits_{i\in\mathcal{I}^-} |y_i|}{\delta} > 0 \\ 0 \le s < \min\left(1, \frac{\min\limits_{i\in\mathcal{I}^+} |y_i|}{\delta} - 1\right) & \text{else.} \end{cases}$$

For simplicity, let this range on s be denoted by $\mathcal{S}$. Note that if $\|\mathcal{I}^-\| \gg 1$, then the first case occurs with high probability. The conditional distribution can now be rewritten with the ranges shown in terms of s:

$$p_s(\mathbf{y}) = \begin{cases} \prod\limits_{i=1}^{N} \frac{C}{(1-\kappa)(1+sm_i)} \exp\left\{ -\left| \frac{y_i}{\alpha\,(1+sm_i)} \right|^{\beta} \right\} & \text{if } s \in \mathcal{S} \\ 0 & \text{else.} \end{cases}$$

When the derivative of this function is used in the modified likelihood ratio test, it is evaluated at the point $s = s_0 = 0$. If 0 is not included in $\mathcal{S}$, then the numerator of the ratio is equal to 0. However, if 0 is included in $\mathcal{S}$, the derivative is non-zero. Thus, the conditions must be found for which the lower bound on $\mathcal{S}$ includes 0:

$$\begin{aligned} 1 - \frac{\min_{i \in \mathcal{I}^-} |y_i|}{\delta} &\leq 0 \\ \min_{i \in \mathcal{I}^-} |y_i| &\geq \delta. \end{aligned}$$

Only when this condition is met is the slope of $p_s(\mathbf{y})$ not equal to 0 at $s = 0$. More precisely, the derivative is given by

$$\begin{aligned} \left.\frac{\partial p_s(\mathbf{y})}{\partial s}\right|_{s=s_0} &= p_s(\mathbf{y}) \left.\frac{\partial \ln p_s(\mathbf{y})}{\partial s}\right|_{s=s_0} \\ &= \begin{cases} \left(\prod_{i=1}^{N} \frac{C}{(1-\kappa)} \exp\left\{ -\left|\frac{y_i}{\alpha}\right|^\beta \right\} \right) \times \left(\sum_{i=1}^{N} \left|\frac{y_i}{\alpha}\right|^\beta \beta m_i \right) & \text{if } \min_{i \in \mathcal{I}^-} (|y_i|) \geq \delta \\ 0 & \text{else} \end{cases} \end{aligned}$$

where the fact that $\mathbf{m}$ is a zero-mean sequence is employed. Now, recall that

$$p_0(\mathbf{y}) = \begin{cases} \prod_{i=1}^{N} \frac{C}{1-\kappa} \exp\left\{ -\left|\frac{y_i}{\alpha}\right|^\beta \right\} & \text{if } |y_i| > \delta \ \forall i \\ 0 & \text{else.} \end{cases}$$

Thus, the modified likelihood ratio is given by

$$\begin{aligned} L(\mathbf{y}) &= \frac{\left.\frac{\partial p_s(\mathbf{y})}{\partial s}\right|_{s=s_0}}{p_0(\mathbf{y})} \\ &= \begin{cases} \frac{\beta}{\alpha^\beta} \sum_{i=1}^{N} |y_i|^\beta m_i & \text{if } \min_{i \in \mathcal{I}^-} (|y_i|) \geq \delta \text{ and } (|y_i| > \delta \ \forall i) \\ \infty & \text{if } \min_{i \in \mathcal{I}^-} (|y_i|) \geq \delta \text{ and } (\exists i : |y_i| \leq \delta) \\ 0 & \text{if } \min_{i \in \mathcal{I}^-} (|y_i|) < \delta \text{ and } (|y_i| > \delta \ \forall i) \\ \infty & \text{if } \min_{i \in \mathcal{I}^-} (|y_i|) < \delta \text{ and } (\exists i : |y_i| \leq \delta) \end{cases} \\ &= \begin{cases} \frac{\beta}{\alpha^\beta} \sum_{i=1}^{N} |y_i|^\beta m_i & \text{if } \min_{i \in \mathcal{I}^-} (|y_i|) \geq \delta \text{ and } (|y_i| > \delta \ \forall i) \\ \infty & \text{if } \min_{i \in \mathcal{I}^-} (|y_i|) < \delta \text{ and } (\exists i : |y_i| \leq \delta) \end{cases} \end{aligned} \tag{3.10}$$

where the final reduction can be performed because the inner two cases can never occur. Once again, the result of the likelihood ratio can be compared against a threshold to produce a binary answer as to the presence of the watermark in the given data.

Although this modified likelihood ratio test does not take on the same form as the ratio for the LRT of the same distribution, conditions still exist for which the statistic will attain infinite values. In these instances, the presence of the watermark is clear and is reflected as such in the statistic. Thus, it is expected that the performance of the detector will be enhanced by the presence of the possible infinities.

CHAPTER 4

CHERNOFF BOUNDS

Simple binary hypothesis testing using a likelihood ratio test forms a conceptually simple statistical detector. However, the computation of such decision statistics can often be quite complex. Thus, bounds on the performance indicators P_D and P_F are desired. The approach taken here is to consider Chernoff bounds [7], which are large-deviation bounds. The bound provided for P_F is an upper bound, while the bound provided for P_D is a lower bound. Hence, the bound on a *receiver operating characteristic* (ROC) curve (P_D vs. P_F) is also a lower bound. Therefore, Chernoff bounds provide an indication of the worst-case performance of a detector using binary hypothesis testing with a likelihood ratio test. In this chapter, Chernoff bounds are formulated for the general case of multiplicative watermarking, and then specialized to the likelihood ratio test detectors based on the Weibull, power exponential, and notched power exponential distributions.

4.1 Chernoff Bounds for Multiplicative Watermarking

As noted in the development of the various likelihood ratio test detectors, when multiplicative watermarking is employed to insert a watermark into a set of host data, the distribution of each element of the output under H_1 is a scaled version of the corresponding distribution under H_0:

$$p_1(y_i) = \frac{1}{1 + s^* m_i} p_0 \left(\frac{y_i}{1 + s^* m_i} \right).$$

This property allows Chernoff bounds to be constructed in a general sense, and later specialized to individual distributions.

To begin, the cumulant generating function of $\ln L(\mathbf{y})$ under H_0, $\mu_0(t_0)$, for $t_0 > 0$, is defined as [7]

$$\mu_0(t_0) = \ln \int p_0^{1-t_0}(\mathbf{y}) p_1^{t_0}(\mathbf{y}) d\mathbf{y}.$$

Using the above scaling property, the cumulant generating function can be rewritten entirely in terms of p_0:

$$\begin{aligned}
\mu_0(t_0) &= \ln \int \prod_{i=1}^{N} p_0^{1-t_0}(y_i) \left(\frac{1}{1+s^* m_i}\right)^{t_0} p_0^{t_0}\left(\frac{y_i}{1+s^* m_i}\right) d\mathbf{y} \\
&= \ln \prod_{i=1}^{N} \int p_0^{1-t_0}(y_i) \left(\frac{1}{1+s^* m_i}\right)^{t_0} p_0^{t_0}\left(\frac{y_i}{1+s^* m_i}\right) dy_i \\
&= \sum_{i=1}^{N} \ln \mathcal{G}\left(t_0, p_0, \frac{1}{1+s^* m_i}\right)
\end{aligned}$$

where the function $\mathcal{G}(t, f, z) = \int f^{1-t}(y) z^t f^t(yz) dy$ is introduced to simplify notation. The cumulant generating function under H_1 is similarly given for $t_1 < 0$ by

$$\begin{aligned}
\mu_1(t_1) &= \mu_0(t_1 + 1) \\
&= \sum_{i=1}^{N} \ln \mathcal{G}\left(t_1 + 1, p_0, \frac{1}{1+s^* m_i}\right).
\end{aligned}$$

With the cumulant generating functions defined, a bound on the probability of false alarm can be written as [7]

$$P_F \leq \exp\{-t_0 \gamma + \mu_0(t_0)\}$$

where γ is the decision threshold. Then, the Chernoff bound is defined as the bound resulting from choosing a value, t_0^*, which maximizes $t_0\gamma - \mu_0(t_0)$:

$$P_F \leq \exp\{-t_0^* \gamma + \mu_0(t_0^*)\}.$$

Similarly, a bound on P_M, and hence P_D, is given by

$$\begin{aligned}
P_M &\leq \exp\{-t_1 \gamma + \mu_1(t_1)\} \\
1 - P_D &\leq \exp\{-t_1 \gamma + \mu_1(t_1)\} \\
P_D &\geq 1 - \exp\{-t_1 \gamma + \mu_1(t_1)\}.
\end{aligned}$$

Finally, the Chernoff bound on P_D is found by maximizing $t_1\gamma - \mu_1(t_1)$ over t_1, yielding

$$P_D \geq 1 - \exp\{-t_1^* \gamma + \mu_1(t_1^*)\}.$$

These Chernoff bounds developed for multiplicative watermarking may now be specialized for the three modeling distributions simply by substituting the desired p_0 into the function $\mathcal{G}$ defining the cumulant generating functions.

In addition to examining detector performance using Chernoff bounds on the P_D and P_F curves, another related measure is often utilized. The Chernoff distance is defined as $D(t_0^*) = -\mu_0(t_0^*)$, and measures the degree of separation between the distributions of the log-likelihood statistic under the two hypotheses. Thus, a higher Chernoff distance corresponds to a stronger detector.

4.2 Weibull Distribution

The first Chernoff bound specialized for the Weibull distribution is one on the probability of false alarm, P_F. Consider the cumulant generating function of $\ln L(\mathbf{y})$ under H_0, $\mu_0(t_0)$, for $t_0 > 0$:

$$\begin{aligned}
\mu_0(t_0) &= \ln \prod_{i=1}^{N} \mathcal{G}\left(t_0, p_0, \frac{1}{1+s^* m_i}\right) \\
&= \ln \prod_{i=1}^{N} \int \left(\frac{\beta}{\alpha}\right)^{1-t_0} \left(\frac{y_i}{\alpha}\right)^{(\beta-1)(1-t_0)} \exp\left\{-(1-t_0)\left(\frac{y_i}{\alpha}\right)^{\beta}\right\} \\
&\quad \times \left(\frac{\beta}{\alpha(1+s^* m_i)}\right)^{t_0} \left(\frac{y_i}{\alpha(1+s^* m_i)}\right)^{(\beta-1)t_0} \exp\left\{-t_0\left(\frac{y_i}{\alpha(1+s^* m_i)}\right)^{\beta}\right\} dy_i \\
&= \ln \prod_{i=1}^{N} \int \left(\frac{\beta}{\alpha}\right)\left(\frac{1}{1+s^* m_i}\right)^{\beta t_0} \left(\frac{y_i}{\alpha}\right)^{\beta-1} \\
&\quad \times \exp\left\{-(1-t_0)\left(\frac{y_i}{\alpha}\right)^{\beta} - t_0\left(\frac{y_i}{\alpha(1+s^* m_i)}\right)^{\beta}\right\} dy_i \\
&= \ln \prod_{i=1}^{N} \int \left(\frac{\beta}{\alpha}\right)\left(\frac{1}{1+s^* m_i}\right)^{\beta t_0} \left(\frac{y_i}{\alpha}\right)^{\beta-1} \\
&\quad \times \exp\left\{\left(\frac{y_i}{\alpha}\right)^{\beta}\left[-1+t_0-t_0\left(\frac{1}{1+s^* m_i}\right)^{\beta}\right]\right\} dy_i \\
&= \ln \prod_{i=1}^{N} \int_0^{\infty} \left(\frac{\beta}{\alpha}\right)\left(\frac{1}{1+s^* m_i}\right)^{\beta t_0} \left(\frac{y_i}{\alpha}\right)^{\beta-1} \\
&\quad \times \exp\left\{\left(\frac{y_i}{\alpha}\right)^{\beta}\left[-1+t_0-t_0\left(\frac{1}{1+s^* m_i}\right)^{\beta}\right]\right\} dy_i \qquad (4.1)
\end{aligned}$$

$$
\begin{aligned}
&= \ln \prod_{i=1}^{N} \left(\frac{\beta}{\alpha}\right) \left(\frac{1}{1+s^*m_i}\right)^{\beta t_0} \left(\frac{1}{\left[-1+t_0-t_0\left(\frac{1}{1+s^*m_i}\right)^{\beta}\right]}\right) \left(\frac{\alpha}{\beta}\right) \\
&\quad \times \left[\exp\left\{\left(\frac{y_i}{\alpha}\right)^{\beta}\left[-1+t_0-t_0\left(\frac{1}{1+s^*m_i}\right)^{\beta}\right]\right\}\right]\Bigg|_{y_i=0}^{\infty} \\
&= \ln \prod_{i=1}^{N} \left(\frac{1}{1+s^*m_i}\right)^{\beta t_0} \left(\frac{-(1+s^*m_i)^{\beta}}{-(1+s^*m_i)^{\beta}+t_0(1+s^*m_i)^{\beta}-t_0}\right) \\
&= \ln \prod_{i=1}^{N} \left(\frac{1}{1+s^*m_i}\right)^{\beta t_0-\beta} \left(\frac{-1}{-(1+s^*m_i)^{\beta}+t_0(1+s^*m_i)^{\beta}-t_0}\right) \\
&= \sum_{i=1}^{N} \ln\left((1+s^*m_i)^{\beta-\beta t_0}\left(\frac{-1}{-(1+s^*m_i)^{\beta}+t_0(1+s^*m_i)^{\beta}-t_0}\right)\right) \\
&= \sum_{i=1}^{N} \left\{(\beta-\beta t_0)\ln(1+s^*m_i) - \ln\left[(1+s^*m_i)^{\beta}-t_0(1+s^*m_i)^{\beta}+t_0\right]\right\} \\
&= -\sum_{i=1}^{N} \left\{(\beta t_0-\beta)\ln(1+s^*m_i) + \ln\left[(1+s^*m_i)^{\beta}-t_0(1+s^*m_i)^{\beta}+t_0\right]\right\}.
\end{aligned}
$$

Note that for the integral to converge in (4.1), it is required that

$$
\begin{aligned}
-1+t_0-\frac{t_0}{(1+s^*m_i)^{\beta}} &< 0 \\
t_0\left(1-\frac{1}{(1+s^*m_i)^{\beta}}\right) &< 1 \\
t_0\left(\frac{(1+s^*m_i)^{\beta}-1}{(1+s^*m_i)^{\beta}}\right) &< 1 \\
t_0 \underset{m_i=-1}{\overset{m_i=1}{\lessgtr}} &\frac{(1+s^*m_i)^{\beta}}{(1+s^*m_i)^{\beta}-1}.
\end{aligned}
$$

But when $m_i = -1$, $\frac{(1+s^*m_i)^{\beta}}{(1+s^*m_i)^{\beta}-1} < 0$, so the bound $t_0 > 0$ suffices. However, the upper bound holds for all i for which $m_i = 1$. But, since only a single value of t_0 is selected, if $m_i = 1$ for any i, then the upper bound must be included. Thus, the overall range restriction is given by

$$
\begin{cases}
0 < t_0 < \frac{(1+s^*)^{\beta}}{(1+s^*)^{\beta}-1} & \text{if } \exists i : m_i = 1 \\
0 < t_0 & \text{else.}
\end{cases}
$$

Using the equation for $\mu_0(t_0)$ and the above bound, a bound on P_F is given by

$$
P_F \leq \exp\{-t_0\gamma + \mu_0(t_0)\}.
$$

By maximizing $t_0\gamma - \mu_0(t_0)$ numerically over the allowable range of t_0, the tightest bound on P_F can be achieved. The maximizing value of t_0 will be denoted by t_0^*. Then, the Chernoff bound on P_F is given by

$$P_F \leq \exp\left\{-t_0^*\gamma + \mu_0(t_0^*)\right\}.$$

Similarly, a lower bound on the probability of miss (not detecting a watermark when it is present), P_M, and, hence, an upper bound on the probability of detect, $P_D = 1 - P_M$, can be found. Consider the cumulant generating function of $\ln L(\mathbf{y})$ under H_1, $\mu_1(t_1)$, for $t_1 < 0$:

$$\begin{aligned}
\mu_1(t_1) &= \mu_0(t_1+1) \\
&= -\sum_{i=1}^{N}\{(\beta(t_1+1)-\beta)\ln(1+s^*m_i) \\
&\qquad + \ln\left[(1+s^*m_i)^\beta - (t_1+1)(1+s^*m_i)^\beta + (t_1+1)\right]\} \\
&= -\sum_{i=1}^{N}\left\{\beta t_1 \ln(1+s^*m_i) + \ln\left[(1+s^*m_i)^\beta - (t_1+1)(1+s^*m_i)^\beta + (t_1+1)\right]\right\}.
\end{aligned}$$

Here the restriction on t_1 for convergence is given by

$$t_1 \underset{m_i=-1}{\overset{m_i=1}{\lessgtr}} \frac{(1+s^*m_i)^\beta}{(1+s^*m_i)^\beta - 1} - 1.$$

But when $m_i = 1$, $\frac{(1+s^*m_i)^\beta}{(1+s^*m_i)^\beta-1} - 1 > 0$, so the bound $t_1 < 0$ suffices. The lower bound holds for all i for which $m_i = -1$. But, since only one value of t_1 is chosen, the lower bound is included if $m_i = -1$ for any i. Thus the overall bounds on t_1 are given by

$$\begin{cases} \frac{(1-s^*)^\beta}{(1-s^*)^\beta-1} - 1 < t_1 < 0 & \text{if } \exists i : m_i = -1 \\ t_1 < 0 & \text{else.} \end{cases}$$

Then, a bound is given for the above restrictions by

$$P_D \;\geq\; 1 - \exp\left\{-t_1\gamma + \mu_1(t_1)\right\}.$$

In order to achieve the tightest bound, $t_1\gamma - \mu_1(t_1)$ is numerically maximized over t_1 to find the maximizing value t_1^*. Hence, the Chernoff bound is given by

$$P_D \geq 1 - \exp\left\{-t_1^*\gamma + \mu_1(t_1^*)\right\}.$$

4.3 Power Exponential Distribution

Chernoff bounds are now developed for the case when the coefficients are modeled using the power exponential distribution. First, an upper bound on the probability of false alarm, P_F, is found. To begin, consider the cumulant generating function of $\ln L(\mathbf{y})$ under H_0, $\mu_0(t_0)$, for $t_0 > 0$:

$$
\begin{aligned}
\mu_0(t_0) &= \ln \prod_{i=1}^{N} \mathcal{G}\left(t_0, p_0, \frac{1}{1+s^* m_i}\right) \\
&= \ln \prod_{i=1}^{N} \int \left(C^{1-t_0} \exp\left\{ -(1-t_0) \left| \frac{y_i}{\alpha} \right|^{\beta} \right\} \right. \\
&\qquad \left. \times \left(\frac{C}{1+s^* m_i} \right)^{t_0} \exp\left\{ -t_0 \left| \frac{y_i}{\alpha\,(1+s^* m_i)} \right|^{\beta} \right\} \right) dy_i \\
&= \ln \prod_{i=1}^{N} \int \frac{C}{(1+s^* m_i)^{t_0}} \exp\left\{ - \left| \frac{y_i}{\alpha} \right|^{\beta} + t_0 \left| \frac{y_i}{\alpha} \right|^{\beta} - t_0 \left| \frac{y_i}{\alpha\,(1+s^* m_i)} \right|^{\beta} \right\} dy_i \\
&= \ln \left[C^N \prod_{i=1}^{N} \frac{2\Gamma\left(\frac{1}{\beta}\right)}{(1+s^* m_i)^{t_0} \left(\frac{1}{\alpha^\beta} - \frac{t_0}{\alpha^\beta} + \frac{t_0}{\alpha^\beta (1+s^* m_i)^\beta} \right)^{\frac{1}{\beta}} \beta} \right] \\
&= \ln \left[\left(\frac{\beta}{2\alpha\Gamma\left(\frac{1}{\beta}\right)} \right)^N \prod_{i=1}^{N} \frac{2\Gamma\left(\frac{1}{\beta}\right)}{(1+s^* m_i)^{t_0} \left(\frac{1}{\alpha^\beta} - \frac{t_0}{\alpha^\beta} + \frac{t_0}{\alpha^\beta (1+s^* m_i)^\beta} \right)^{\frac{1}{\beta}} \beta} \right] \\
&= \ln \left[\prod_{i=1}^{N} \frac{1}{(1+s^* m_i)^{t_0} \left(1 - t_0 + \frac{t_0}{(1+s^* m_i)^\beta} \right)^{\frac{1}{\beta}}} \right] \\
&= \ln \left[\prod_{i=1}^{N} \frac{1}{\left((1+s^* m_i)^\beta - t_0\,(1+s^* m_i)^\beta + t_0 \right)^{\frac{1}{\beta}} (1+s^* m_i)^{t_0 - 1}} \right] \\
&= \sum_{i=1}^{N} \ln \left[\frac{1}{\left((1+s^* m_i)^\beta - t_0\,(1+s^* m_i)^\beta + t_0 \right)^{\frac{1}{\beta}} (1+s^* m_i)^{t_0 - 1}} \right] \\
&= -\sum_{i=1}^{N} \ln \left[\left((1+s^* m_i)^\beta - t_0\,(1+s^* m_i)^\beta + t_0 \right)^{\frac{1}{\beta}} (1+s^* m_i)^{t_0 - 1} \right] \\
&= -\sum_{i=1}^{N} \left[\frac{1}{\beta} \ln \left((1+s^* m_i)^\beta - t_0\,(1+s^* m_i)^\beta + t_0 \right) + (t_0 - 1) \ln (1+s^* m_i) \right].
\end{aligned}
$$

Note that for the integral in the above expression to converge, the same restriction on t_0 as present in the case of the Weibull distribution holds, namely,

$$-1 + t_0 - \frac{t_0}{(1 + s^* m_i)^\beta} < 0.$$

This expression leads to the overall restrictions of

$$\begin{cases} 0 < t_0 < \frac{(1+s^*)^\beta}{(1+s^*)^\beta - 1} & \text{if } \exists i : m_i = 1 \\ 0 < t_0 & \text{else.} \end{cases}$$

Using the equation for $\mu_0(t_0)$ and the above bound, a bound on P_F is given by

$$P_F \leq \exp\left\{-t_0\gamma + \mu_0(t_0)\right\}.$$

To achieve the tightest bound, $t_0\gamma - \mu_0(t_0)$ is maximized with respect to t_0. By performing this maximization numerically to find the maximizing value, t_0^*, the Chernoff bound on P_F can be written as

$$P_F \leq \exp\left\{-t_0^*\gamma + \mu_0\left(t_0^*\right)\right\}.$$

Now, using a similar construction, a lower bound on the probability of miss, P_M, and, hence, an upper bound on the probability of detect, $P_D = 1 - P_M$, is found. To begin, consider the cumulant generating function of $\ln L(\mathbf{y})$ under H_1, $\mu_1(t_1)$, for $t_1 < 0$:

$$\begin{aligned} \mu_1(t_1) &= \mu_0(t_1 + 1) \\ &= -\sum_{i=1}^{N} \left[\frac{1}{\beta} \ln\Big((1 + s^* m_i)^\beta - (t_1 + 1)(1 + s^* m_i)^\beta + (t_1 + 1)\Big) + t_1 \ln(1 + s^* m_i)\right]. \end{aligned}$$

Here the restriction on t_1 for convergence is also identical to that in the Weibull distribution case, namely,

$$t_1 \underset{m_i=-1}{\overset{m_i=1}{\lessgtr}} \frac{(1 + s^* m_i)^\beta}{(1 + s^* m_i)^\beta - 1} - 1.$$

Hence, the overall bounds on t_1 are given by

$$\begin{cases} \frac{(1-s^*)^\beta}{(1-s^*)^\beta - 1} - 1 < t_1 < 0 & \text{if } \exists i : m_i = -1 \\ t_1 < 0 & \text{else.} \end{cases}$$

Thus, a bound is given for the above restrictions by

$$\begin{aligned}
P_M &\leq \exp\{-t_1\gamma + \mu_1(t_1)\} \\
1 - P_D &\leq \exp\{-t_1\gamma + \mu_1(t_1)\} \\
P_D &\geq 1 - \exp\{-t_1\gamma + \mu_1(t_1)\}.
\end{aligned}$$

Again, wishing to obtain the tightest bound, $t_1\gamma - \mu_1(t_1)$ is numerically maximized with respect to t_1 to find the maximizing value, t_1^*. Then, the Chernoff bound on P_D is

$$P_D \geq 1 - \exp\{-t_1^*\gamma + \mu_1(t_1^*)\}.$$

4.4 Notched Power Exponential Distribution

Finally, the Chernoff bounds on the probability of false alarm and the probability of detect are now found for the case of the notched power exponential distribution. To begin, an upper bound on the probability of false alarm, P_F, is found. Consider the cumulant generating function of $\ln L(\mathbf{y})$ under H_0, $\mu_0(t_0)$, for $t_0 > 0$, where $\tilde{\delta}_i = \max(\delta_i, \delta)$:

$$\begin{aligned}
\mu_0(t_0) &= \ln \prod_{i=1}^{N} \mathcal{G}\left(t_0, p_0, \frac{1}{1+s^*m_i}\right) \\
&= \ln\left[2^N \prod_{i=1}^{N} \left(\frac{C}{1-\kappa}\right)^{1-t_0} \left(\frac{C}{(1-\kappa)(1+s^*m_i)}\right)^{t_0}\right. \\
&\qquad \left.\times \int_{\tilde{\delta}_i}^{\infty} \exp\left\{\left(\frac{-1}{\alpha^\beta} + \frac{t_0}{\alpha^\beta} - \frac{t_0}{\alpha^\beta(1+s^*m_i)^\beta}\right) y_i^\beta\right\} dy_i\right] \\
&= \ln\left[2^N \prod_{i=1}^{N} \left(\frac{C}{1-\kappa}\right)^{1-t_0} \left(\frac{C}{(1-\kappa)(1+s^*m_i)}\right)^{t_0}\right. \\
&\qquad \left.\times \frac{\Gamma\left(\frac{1}{\beta}, \left(\frac{1}{\alpha^\beta} - \frac{t_0}{\alpha^\beta} + \frac{t_0}{\alpha^\beta(1+s^*m_i)^\beta}\right)\tilde{\delta}_i^\beta\right)}{\beta\left(\frac{1}{\alpha^\beta} - \frac{t_0}{\alpha^\beta} + \frac{t_0}{\alpha^\beta(1+s^*m_i)^\beta}\right)^{\frac{1}{\beta}}}\right] \\
&= \ln\left[2^N \prod_{i=1}^{N} \left(\frac{C}{1-\kappa}\right) \left(\frac{1}{1+s^*m_i}\right)^{t_0} \frac{\Gamma\left(\frac{1}{\beta}, \left(\frac{1}{\alpha^\beta} - \frac{t_0}{\alpha^\beta} + \frac{t_0}{\alpha^\beta(1+s^*m_i)^\beta}\right)\tilde{\delta}_i^\beta\right)}{\beta\left(\frac{1}{\alpha^\beta} - \frac{t_0}{\alpha^\beta} + \frac{t_0}{\alpha^\beta(1+s^*m_i)^\beta}\right)^{\frac{1}{\beta}}}\right] \\
&= \ln\left[2^N \prod_{i=1}^{N} \left(\frac{\beta}{2\alpha\Gamma\left(\frac{1}{\beta}\right)(1-\kappa)}\right) \left(\frac{1}{1+s^*m_i}\right)^{t_0} \frac{\Gamma\left(\frac{1}{\beta}, \left(\frac{1}{\alpha^\beta} - \frac{t_0}{\alpha^\beta} + \frac{t_0}{\alpha^\beta(1+s^*m_i)^\beta}\right)\tilde{\delta}_i^\beta\right)}{\beta\left(\frac{1}{\alpha^\beta} - \frac{t_0}{\alpha^\beta} + \frac{t_0}{\alpha^\beta(1+s^*m_i)^\beta}\right)^{\frac{1}{\beta}}}\right]
\end{aligned}$$

$$
\begin{aligned}
&= \ln\left[\prod_{i=1}^{N}\left(\frac{1}{\Gamma\left(\frac{1}{\beta}\right)(1-\kappa)}\right)\left(\frac{1}{1+s^*m_i}\right)^{t_0}\frac{\Gamma\left(\frac{1}{\beta},\left(\frac{1}{\alpha^\beta}-\frac{t_0}{\alpha^\beta}+\frac{t_0}{\alpha^\beta(1+s^*m_i)^\beta}\right)\tilde{\delta}_i^\beta\right)}{\left(1-t_0+\frac{t_0}{(1+s^*m_i)^\beta}\right)^{\frac{1}{\beta}}}\right] \\
&= \sum_{i=1}^{N}\left[-\ln\left(\Gamma\left(\frac{1}{\beta}\right)\right)-\ln(1-\kappa)+\ln\left(\left(\frac{1}{1+s^*m_i}\right)^{t_0}\right)\right. \\
&\qquad +\ln\left(\Gamma\left(\frac{1}{\beta},\left(\frac{1}{\alpha^\beta}-\frac{t_0}{\alpha^\beta}+\frac{t_0}{\alpha^\beta(1+s^*m_i)^\beta}\right)\tilde{\delta}_i^\beta\right)\right) \\
&\qquad \left.-\ln\left(\left(1-t_0+\frac{t_0}{(1+s^*m_i)^\beta}\right)^{\frac{1}{\beta}}\right)\right] \\
&= \sum_{i=1}^{N}\left[-\ln\left(\Gamma\left(\frac{1}{\beta}\right)\right)-\ln(1-\kappa)+t_0\ln\left(\frac{1}{1+s^*m_i}\right)\right. \\
&\qquad +\ln\left(\Gamma\left(\frac{1}{\beta},\left(\frac{1}{\alpha^\beta}-\frac{t_0}{\alpha^\beta}+\frac{t_0}{\alpha^\beta(1+s^*m_i)^\beta}\right)\tilde{\delta}_i^\beta\right)\right) \\
&\qquad \left.-\frac{1}{\beta}\ln\left(1-t_0+\frac{t_0}{(1+s^*m_i)^\beta}\right)\right]
\end{aligned}
\tag{4.2}
$$

where $\Gamma(a,z) = \int_z^\infty \exp\{-t\}t^{a-1}dt$ with $a > 0$, is the incomplete Gamma function. As seen in the case of the power exponential distribution, t_0 must be bounded such that the integral converges. The resulting bound is the same as that found in the case of the power exponential, namely,

$$
\begin{cases} 0 < t_0 < \frac{(1+s^*)^\beta}{(1+s^*)^\beta-1} & \text{if } \exists i : m_i = 1 \\ 0 < t_0 & \text{else.} \end{cases}
$$

Using this equation for $\mu_0(t_0)$, a bound on P_F is given by

$$P_F \leq \exp\{-t_0\gamma + \mu_0(t_0)\}.$$

To achieve the tightest bound, $t_0\gamma - \mu_0(t_0)$ is maximized with respect to t_0. This maximization is performed numerically to find the maximizing value, t_0^*. Then, the Chernoff bound on P_F can be written as

$$P_F \leq \exp\{-t_0^*\gamma + \mu_0(t_0^*)\}.$$

Now, using a similar construction, a lower bound on the probability of miss, P_M, and hence an upper bound on the probability of detect, $P_D = 1 - P_M$, is found. To begin, consider the cumulant generating function of $\ln L(\mathbf{y})$ under H_1, $\mu_1(t_1)$ for $t_1 < 0$:

$$\begin{aligned}
\mu_1(t_1) &= \mu_0(t_1+1) \\
&= \sum_{i=1}^{N} \Bigg[-\ln\Gamma\left(\frac{1}{\beta}\right) - \ln(1-\kappa) + (t_1+1)\ln\left(\frac{1}{1+s^* m_i}\right) \\
&\qquad + \ln\Gamma\left(\frac{1}{\beta}, \left(\frac{1}{\alpha^\beta} - \frac{t_1+1}{\alpha^\beta} + \frac{t_1+1}{\alpha^\beta(1+s^* m_i)^\beta}\right)\tilde{\delta}_i^\beta\right) \\
&\qquad - \frac{1}{\beta}\ln\left(-t_1 + \frac{t_1+1}{(1+s^* m_i)^\beta}\right)\Bigg].
\end{aligned}$$

Again, restrictions on t_1 for convergence are present and are the same as in the case of the power exponential,

$$\begin{cases} \frac{(1-s^*)^\beta}{(1-s^*)^\beta - 1} - 1 < t_1 < 0 & \text{if } \exists i : m_i = -1 \\ t_1 < 0 & \text{else.} \end{cases}$$

Thus, a bound is given by

$$\begin{aligned}
P_M &\leq \exp\{-t_1\gamma + \mu_1(t_1)\} \\
1 - P_D &\leq \exp\{-t_1\gamma + \mu_1(t_1)\} \\
P_D &\geq 1 - \exp\{-t_1\gamma + \mu_1(t_1)\}.
\end{aligned}$$

Again, wishing to obtain the tightest bound, $t_1\gamma - \mu_1(t_1)$ is maximized with respect to t_1. By solving this equation numerically for the maximizing value, t_1^*, the Chernoff bound on P_D can be written as

$$P_D \geq 1 - \exp\{-t_1^*\gamma + \mu_1(t_1^*)\}.$$

CHAPTER 5

THEORETICAL ANALYSIS

This chapter describes the analyzes and experiments performed in order to test the watermark detectors using synthetic data, and presents their results. First, the Chernoff bounds on the likelihood ratio test detectors are presented to study the behavior of the detectors over a range of watermark lengths. The performance of the simplified detector is then derived and examined. Next, comments are made on the degree of difficulty experienced by an eavesdropper attempting to detect the watermark. Finally, for a specific watermark length, the performance of the detectors is evaluated through Monte Carlo simulation using data synthesized according to the corresponding distributions. By comparing the Monte Carlo simulation results for the likelihood ratio test detector with the Chernoff bounds, the tightness of the bounds can be assessed.

5.1 Evaluation of Chernoff Bounds

In order to demonstrate the behavior of the derived Chernoff bounds, the bounds are calculated for each of the three distributions, averaging over 100 messages, for a range of message sizes. The averaging is performed because of the dependence of the bound on the message, $\mathbf{m}$. This dependence diminishes as the length of the message increases because then the number of instances of $m_i = 1$ and of $m_i = -1$ will each become closer to $\frac{N}{2}$. However, small message sizes are also included in the experiments, so multiple simulations are utilized and averaged. The computed bounds may be used to evaluate detector performance without the need for the large Monte Carlo simulations that are required for very high detection probabilities and low false alarm probabilities.

Figures 5.1 and 5.2 contain the Chernoff bound results for the Weibull and power exponential distributions, respectively. Figure 5.3 contains the bounds for the notched power exponential distribution on a logarithmic scale. Some general comments can be made regarding the plots for all of the distributions. In the case of P_D vs. Threshold and P_F vs. Threshold, the graphs demonstrate that as the number of watermark bits is increased, the transition from high probabilities to low probabilities is spread out across a larger range of thresholds. For the P_D vs. P_F graph, the curves approach the top left-hand corner as the number of bits increases. This result is expected because the number of bits used for the watermark is increased; hence, more data are present, making the watermark easier to detect. Thus, higher detection probabilities are achieved for the same false alarm probabilities. The latter plot demonstrates the bound on the receiver operating characteristic curve. For the Weibull and power exponential distributions, many of the sequence lengths result in bounds lying below the line $P_D = P_F$. These bounds are trivial bounds because better performance could be achieved simply by flipping a coin (yielding $(P_D, P_F) = (0.5, 0.5)$), by always choosing H_0 ($(P_D, P_F) = (0, 0)$), or by always choosing H_1 ($(P_D, P_F) = (1, 1)$). The logarithmic plots for the notched power exponential ROC curve are utilized because, unlike the Weibull and power exponential bounds, the detection and false alarm probabilities become extremely close to 1 and 0, respectively, as the number of bits is increased. Thus, the log scale allows these precise values to be examined more explicitly. However, to plot P_D more clearly in a log fashion, the affine transform $P_D - 1 = -P_M$ is used. Finally, the curves for 128, 256, and 512 watermark bits are not present on the P_D or ROC graphs for the notched power exponential case because their detection probabilities were identically equal to 1 for the number of simulations performed.

The Chernoff distances are computed for each of the three coefficient modeling distributions, for the specific threshold of zero, and for watermark length of 32. The results are given in Table 5.1 for embedding strengths of 0.05 and 0.10. Similar to the ROC curves, the Chernoff distances demonstrate clearly higher performance by the notched power exponential distribution over both the Weibull and power exponential distributions.

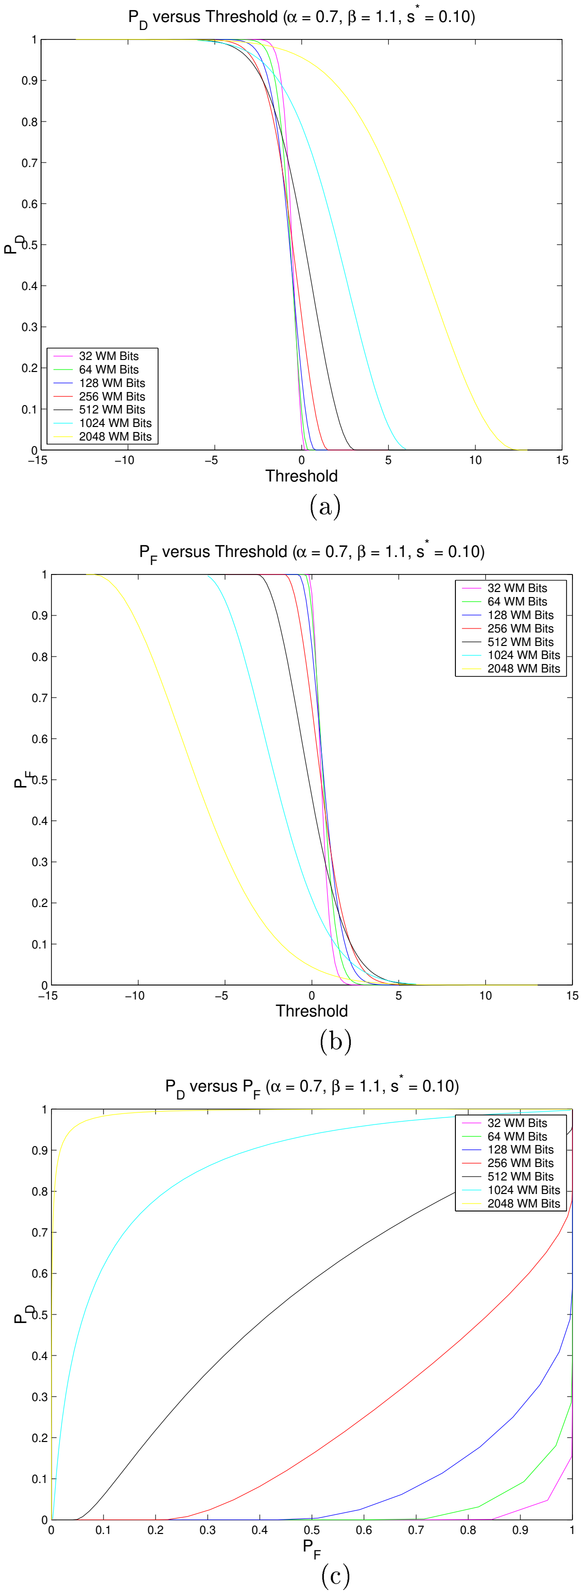

Figure 5.1 Chernoff bounds on the WB LRT detector ($s^* = 0.10$).

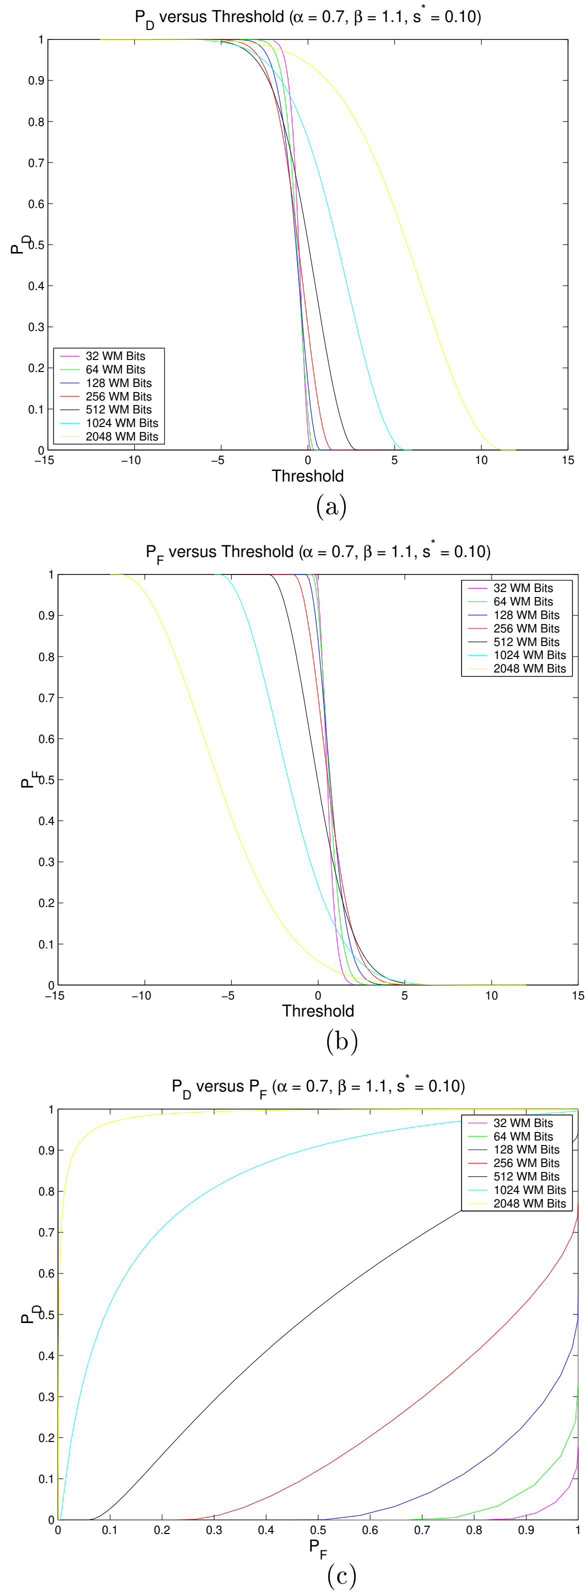

Figure 5.2 Chernoff bounds on the PE LRT detector ($s^* = 0.10$).

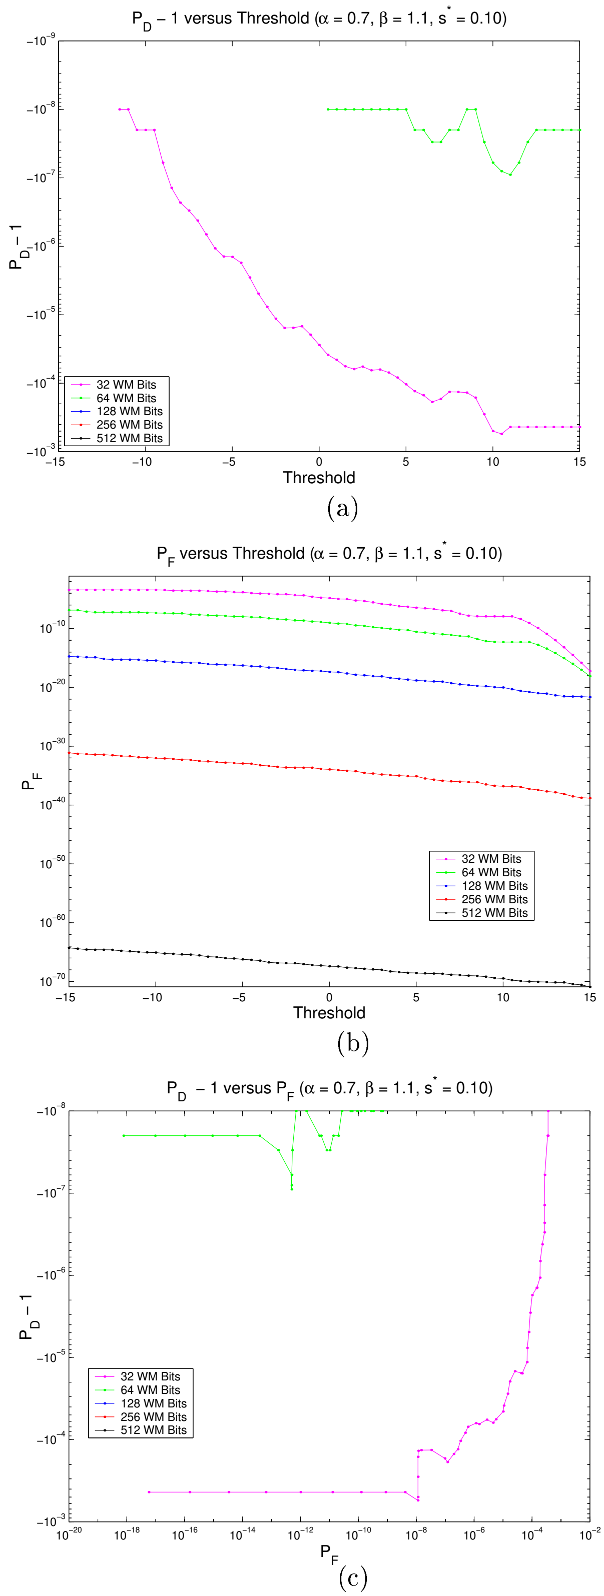

Figure 5.3 Chernoff bounds on the NPE LRT detector (log) ($s^* = 0.10$).

Table 5.1 Chernoff distances for the three modeling distributions using a threshold of zero.

Distribution	Chernoff distance	
	$s^* = 0.05$	$s^* = 0.10$
Weibull	0.0120	0.0497
Power exponential	0.0111	0.0452
Notched power exponential	5.0135	31.6715

5.2 Performance of the Simple Detector for Notched Distributions

In the previous section, bounds were presented on the likelihood ratio test detector based on the notched power exponential distribution. The performance of the related simple detector (3.4) is now analyzed. First, exact expressions for the detection and false alarm probabilities are derived. Then, error exponents are computed, indicating the rate of convergence of these probabilities as message length increases.

To begin, P_D and P_F are considered for the SD for any notched distribution, and will be later specialized to the notched power exponential. Using this detector, false alarms will occur with probability $1 - \lambda$ when the data fail to lie in Region 3 under H_0. Mathematically,

$$\begin{aligned} P_F &= (1-\lambda)\left(1 - P_0[\mathcal{E}_3]\right) \\ &= (1-\lambda)\left(1 - 1 + \left(2P_0(-\delta(1+s^*))\right)^{\frac{N}{2}}\right) \\ &= (1-\lambda)\left(2P_0(-\delta(1+s^*))\right)^{\frac{N}{2}}. \end{aligned} \tag{5.1}$$

Similarly, misses will occur with probability λ when the data fail to lie in Region 2 under H_1. Thus,

$$\begin{aligned} P_M &= \lambda\left(1 - P_1[\mathcal{E}_2]\right) \\ &= \lambda\left(1 - 1 + \left(2P_1(-\delta)\right)^{\frac{N}{2}}\right) \\ &= \lambda\left(2P_1(-\delta)\right)^{\frac{N}{2}} \end{aligned}$$

$$\begin{aligned} P_D &= 1 - P_M \\ &= 1 - \lambda\left(2P_1(-\delta)\right)^{\frac{N}{2}}. \end{aligned} \tag{5.2}$$

For the case of the notched power exponential distribution,

$$
\begin{aligned}
P_F &= (1-\lambda)\left[2\int_{-\infty}^{-\delta(1+s^*)} \frac{C}{1-\kappa}\exp\left\{-\left|\frac{y}{\alpha}\right|^{\beta}\right\} dy\right]^{\frac{N}{2}} \\
&= (1-\lambda)\left[2\frac{C}{1-\kappa}\int_{\delta(1+s^*)}^{\infty} \exp\left\{-\left(\frac{y}{\alpha}\right)^{\beta}\right\} dy\right]^{\frac{N}{2}}.
\end{aligned}
$$

Performing the substitution $w = \left(\frac{y}{\alpha}\right)^{\beta}$ yields

$$
\begin{aligned}
P_F &= (1-\lambda)\left[\frac{2C}{1-\kappa}\int_{\left(\frac{\delta(1+s^*)}{\alpha}\right)^{\beta}}^{\infty} \exp\{-w\}\frac{\alpha}{\beta}w^{\frac{1}{\beta}-1}dw\right]^{\frac{N}{2}} \\
&= (1-\lambda)\left[\frac{2C\alpha}{(1-\kappa)\beta}\Gamma\left(\frac{1}{\beta},\left(\frac{\delta(1+s^*)}{\alpha}\right)^{\beta}\right)\right]^{\frac{N}{2}} \\
&= (1-\lambda)\left[\left(\frac{1}{(1-\kappa)\Gamma\left(\frac{1}{\beta}\right)}\right)\Gamma\left(\frac{1}{\beta},\left(\frac{\delta(1+s^*)}{\alpha}\right)^{\beta}\right)\right]^{\frac{N}{2}}.
\end{aligned}
$$

Similarly,

$$
\begin{aligned}
P_D &= 1-\lambda\left[2\int_{-\infty}^{-\delta} \frac{C}{(1-\kappa)(1-s^*)}\exp\left\{-\left|\frac{y}{\alpha(1-s^*)}\right|^{\beta}\right\} dy\right]^{\frac{N}{2}} \\
&= 1-\lambda\left[2\frac{C}{(1-\kappa)(1-s^*)}\int_{\delta}^{\infty} \exp\left\{-\left(\frac{y}{\alpha(1-s^*)}\right)^{\beta}\right\} dy\right]^{\frac{N}{2}}.
\end{aligned}
$$

Performing the substitution $w = \left(\frac{y}{\alpha(1-s^*)}\right)^{\beta}$ yields

$$
\begin{aligned}
P_D &= 1-\lambda\left[\frac{2C}{(1-\kappa)(1-s^*)}\int_{\left(\frac{\delta}{\alpha(1-s^*)}\right)^{\beta}}^{\infty} \exp\{-w\}\frac{\alpha(1-s^*)}{\beta}w^{\frac{1}{\beta}-1}dw\right]^{\frac{N}{2}} \\
&= 1-\lambda\left[\frac{2C\alpha(1-s^*)}{(1-\kappa)(1-s^*)\beta}\Gamma\left(\frac{1}{\beta},\left(\frac{\delta}{\alpha(1-s^*)}\right)^{\beta}\right)\right]^{\frac{N}{2}} \\
&= 1-\lambda\left[\left(\frac{1}{(1-\kappa)\Gamma\left(\frac{1}{\beta}\right)}\right)\Gamma\left(\frac{1}{\beta},\left(\frac{\delta}{\alpha(1-s^*)}\right)^{\beta}\right)\right]^{\frac{N}{2}}.
\end{aligned}
$$

These expressions facilitate the computation of detection and false alarm probabilities for the simplified detector. However, the resulting equations are rather complex and produce values for P_D and P_F that are extremely close to 1 and 0, respectively, which makes comparisons difficult. Thus, error exponents for these probabilities are now found.

To give meaning to the use of error exponents, it is first noted that the expressions for P_F and P_D (Equations (5.1) and (5.2)) can be rewritten as

$$P_F = (1-\lambda)\exp\left\{\frac{N}{2}\ln\left(2P_0(-\delta(1+s^*))\right)\right\}$$

and

$$P_D = 1-\lambda\exp\left\{\frac{N}{2}\ln\left(2P_1(-\delta)\right)\right\}.$$

From these equations, it is clear that the error probabilities P_F and $P_M = 1 - P_D$ behave as decaying exponentials. Error exponents provide an indication of the rate of convergence as N approaches ∞, with a higher value corresponding to a faster convergence. The error exponent indicating the rate of convergence of P_F to 0 is given by

$$\begin{aligned}
E_{P_F} &= -\lim_{N\to\infty}\frac{1}{N}\ln P_F \\
&= -\lim_{N\to\infty}\frac{1}{N}\ln\left[(1-\lambda)\exp\left\{\frac{N}{2}\ln\left(2P_0(-\delta(1+s^*))\right)\right\}\right] \\
&= -\lim_{N\to\infty}\left[\frac{1}{N}\ln(1-\lambda)+\frac{1}{2}\ln\left(2P_0(-\delta(1+s^*))\right)\right] \\
&= -\frac{1}{2}\ln\left(2P_0(-\delta(1+s^*))\right)
\end{aligned}$$

while the error exponent corresponding to the rate of convergence of P_D to 1 is given by

$$\begin{aligned}
E_{P_D} &= -\lim_{N\to\infty}\frac{1}{N}\ln(1-P_D) \\
&= -\lim_{N\to\infty}\frac{1}{N}\ln\left[\lambda\exp\left\{\frac{N}{2}\ln\left(2P_1(-\delta)\right)\right\}\right] \\
&= -\lim_{N\to\infty}\left[\frac{1}{N}\ln\lambda+\frac{1}{2}\ln\left(2P_1(-\delta)\right)\right] \\
&= -\frac{1}{2}\ln\left(2P_1(-\delta)\right).
\end{aligned}$$

It is interesting to note that the error exponents do not depend on the randomization probability, λ.

To specialize these expressions to the notched power exponential distribution, the appropriate cumulative density functions are inserted, yielding

$$E_{P_F} = -\frac{1}{2}\ln\left[\left(\frac{1}{(1-\kappa)\Gamma\left(\frac{1}{\beta}\right)}\right)\Gamma\left(\frac{1}{\beta},\left(\frac{\delta(1+s^*)}{\alpha}\right)^{\beta}\right)\right]$$

and

$$E_{P_D} = -\frac{1}{2}\ln\left[\left(\frac{1}{(1-\kappa)\Gamma\left(\frac{1}{\beta}\right)}\right)\Gamma\left(\frac{1}{\beta},\left(\frac{\delta}{\alpha(1-s^*)}\right)^{\beta}\right)\right].$$

For a given set of parameters, the expressions for the error exponents can be evaluated numerically. Using the parameter values $\alpha = 0.1700$, $\beta = 1.1$, $\delta = 0.7760$, and $s^* = 0.10$, the resulting error exponents are

$$E_{P_F} = 0.2978$$

and

$$E_{P_D} = 0.3310.$$

Note that the error exponents for P_F and P_D are of similar values, indicating that P_F approaches 0 at approximately the same rate that P_D approaches 1. The simplified detector can be compared against the LRT notched power exponential detector, whose error exponent (for P_F) is related to the Chernoff distance. Under the assumption of a message with an equal number of elements having -1 and $+1$ values, the Chernoff distance (derived in Section 4.4) can be written in terms of two constants, K^+ and K^-, which are equal to the summand in (4.2) with $m_i = 1$ and $m_i = -1$ substituted, respectively,

$$D(t_0^*) = \frac{N}{2}\left(K^+ + K^-\right).$$

Then, since the Chernoff bound using t_0^* is tight in the exponent, the error exponent for P_F is given by

$$\begin{aligned} E_{P_F} &= -\lim_{N\to\infty}\frac{1}{N}\ln P_F \\ &= -\lim_{N\to\infty}\frac{1}{N}\ln\exp\left\{-t_0^*\gamma - D(t_0^*)\right\} \\ &= -\lim_{N\to\infty}\frac{1}{N}\left(-t_0^*\gamma - D(t_0^*)\right) \end{aligned}$$

$$
\begin{aligned}
&= \lim_{N\to\infty} \frac{D(t_0^*)}{N} \\
&= \frac{1}{2}\left(K^+ + K^-\right).
\end{aligned}
$$

Thus, the error exponent is found by calculating the Chernoff distance and dividing by the length of the message used in its computation. This distance is given for the NPE using $N = 32$ in Table 5.1, yielding an error exponent of $E_{P_F} = 0.9897$. The NPE value is higher than that for the SD, thereby quantifying the performance decrease resulting from the use of the simplified detector. However, the corresponding error exponents for the LRT detectors for the Weibull and power exponential distributions (0.0016 and 0.0014, respectively) are lower than that for the SD, demonstrating a strong performance by the simple detector, despite its simplicity.

5.3 An Eavesdropper's Detection Problem

Previously, Chernoff distances were employed to evaluate the difficulty of the detection problem for various modeling distributions. Another problem worthy of study is the relative difficulty of the image watermarking detection problem as seen by the detector compared to that seen by an eavesdropper. The goal of the eavesdropper is to detect whether a message is present in an image without knowledge of the candidate region, $\mathcal{C}$.

When an eavesdropper obtains an image, two hypothesis are possible: H_0, which states the message is not present; and H_1, which states that the message is present. Under H_0, the distribution of the output, p_0, is simply given by a power exponential. However, since the eavesdropper does not know the candidate set, the distribution under H_1, p_1, is unknown. Thus, the eavesdropper would assume a mixture distribution, $\tilde{p}_1$. To formulate this distribution, a number of addition probability distributions are first defined:

$$
\begin{aligned}
p_0 &: \text{unmarked power exponential} \\
p_0(y) &= C \exp\left\{ -\left|\frac{y}{\alpha}\right|^{\beta} \right\} \\
p_2 &: \text{unmarked bandpass power exponential} \\
p_2(y) &= \begin{cases} \frac{C}{1-\kappa} \exp\left\{ -\left|\frac{y}{\alpha}\right|^{\beta} \right\} & \text{if } |y| \le \delta \\ 0 & \text{else} \end{cases}
\end{aligned}
$$

$$
\begin{aligned}
p_3 &: \text{ unmarked notched power exponential} \\
p_3(y) &= \begin{cases} \frac{C}{1-\kappa} \exp\left\{ -\left|\frac{y}{\alpha}\right|^\beta \right\} & \text{if } |y| > \delta \\ 0 & \text{else} \end{cases} \\
p_4^+ &: \text{ positively marked notched power exponential} \\
p_4^+(y) &= \frac{1}{1+s^*} p_3 \left(\frac{y}{1+s^*} \right) \\
p_4^- &: \text{ negatively marked notched power exponential} \\
p_4^-(y) &= \frac{1}{1-s^*} p_3 \left(\frac{y}{1-s^*} \right) \\
\tilde{p}_4 &: \text{ mixture watermarked notched power exponential} \\
\tilde{p}_4(y) &= \frac{1}{2} \left[p_4^+ + p_4^- \right]
\end{aligned}
$$

Note that in defining the mixture distributions, the assumption that the message is distributed as an iid binary sequence is employed. If a different distribution were to be present (for example if error correcting codes were utilized), then the mixtures need be adjusted accordingly. Using these distributions, the distribution of the output data as seen by the eavesdropper under H_1 is given by

$$\tilde{p}_1 = \nu(1-\kappa)\tilde{p}_4 + \nu\kappa p_2 + (1-\nu)p_0.$$

Assuming that the data are iid, the Chernoff distance between the distributions under H_0 and H_1 as seen by the eavesdropper is given by

$$D(t_0^*, p_0, \tilde{p}_1) = MD(t_0^*, p_0, \nu(1-\kappa)\tilde{p}_4 + \nu\kappa p_2 + (1-\nu)p_0). \tag{5.3}$$

The situation is different for the actual detector, who has knowledge of the candidate set. Thus, for each coefficient, the detector knows which pair of distributions to consider for the two hypotheses. Hence, the Chernoff distance as seen by the detector is given by

$$
\begin{aligned}
D(t_0^*, p_0, p_1) &= M\nu(1-\kappa)\left(\frac{1}{2}\right) D(t_0^*, p_0, p_4^+) + M\nu(1-\kappa)\left(\frac{1}{2}\right) D(t_0^*, p_0, p_4^-) \\
&\quad + M\nu\kappa D(t_0^*, p_o, p_2) + M(1-\nu)D(t_0^*, p_0, p_0) \\
&= M\nu(1-\kappa)\left(\frac{1}{2}\right) D(t_0^*, p_0, p_4^+) + M\nu(1-\kappa)\left(\frac{1}{2}\right) D(t_0^*, p_0, p_4^-) \\
&\quad + M\nu\kappa D(t_0^*, p_o, p_2).
\end{aligned} \tag{5.4}
$$

A comparison may now be made between the Chernoff distances, (5.3) and (5.4), seen by the eavesdropper and the detector, to provide insight into the relative difficulty of the detection problems. With the above formulations and by the concavity of Chernoff distance, it is clear that

$$D(t_0^*, p_0, p_1) > D(t_0^*, p_0, \tilde{p}_1).$$

Thus, as is to be expected, the Chernoff distance seen by the detector is larger than that seen by an eavesdropper. Hence, the detection problem is more difficult for the eavesdropper than for the detector.

To quantify this effect, the Chernoff distances are evaluated over the range of possible ν values using the parameters $\alpha = 0.17$, $\beta = 1.1$, $\delta = 0.7760$, and $s^* = 0.10$. Note that when $\nu = 0$, the candidate set has size zero; while when $\nu = 1$, the candidate set consists of the entire M coefficients. The resulting Chernoff distance curves are given in Figure 5.4. These curves illustrate that the detector observes a significantly larger Chernoff distance than the eavesdropper and, hence, will be better able to detect the message.

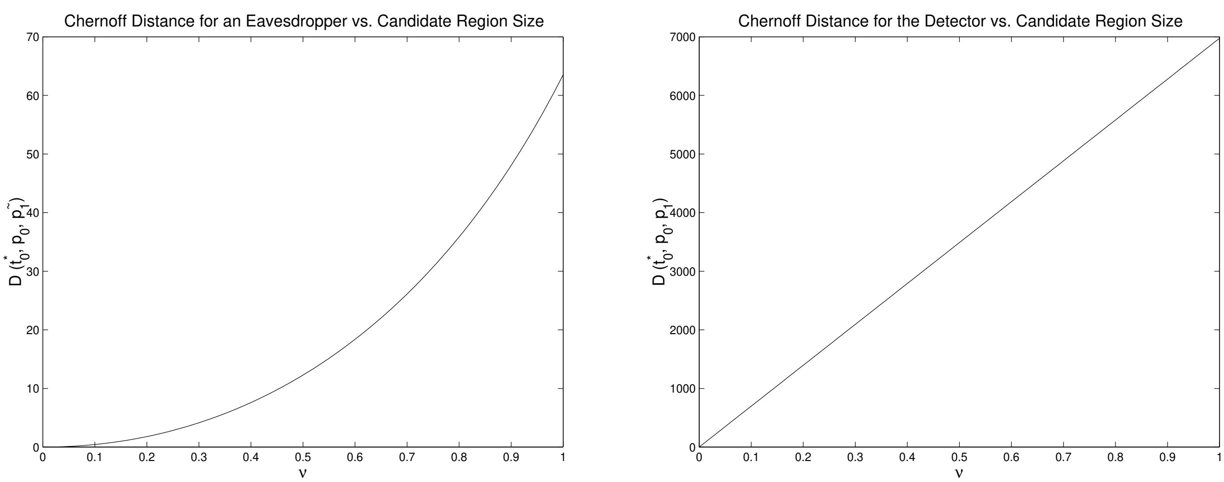

(a) (b)

Figure 5.4 Chernoff distances as seen by (a) an eavesdropper and (b) the detector as a function of ν.

5.4 Monte Carlo Simulations Using Synthetic Data

To begin testing the Weibull, power exponential, and notched power exponential detectors, data are generated as previously discussed according to each of the three distributions. More specifically, $512^2 = 262\ 144$ coefficients are generated for each distribution to model image transform coefficients. The Weibull data utilize parameters of $\alpha = 0.4437$ and $\beta = 1.1$, while the power and notched power exponential use the parameter values of $\alpha = 0.1700$, $\beta = 1.1$, $\delta = 0.7760$, and $\kappa = 0.9961$ (approximately 0.4% of the coefficients are watermarked), as shown in Figure 5.5. The values selected for the distribution parameters reflect the estimated parameters for the Lena image, as will be presented in the next chapter.

As discussed previously, multiplicative watermarking is used to insert the watermark. Experiments are performed using the two embedding strengths of 0.05 and 0.10. These values are selected to be quite small, so that the distortion introduced into the original coefficients is minimal. Using these watermarked coefficients, the performance of the watermark detectors for a given distribution can be determined. A Monte Carlo simulation for each distribution is performed to analyze each of the three main detectors. The simulations are done over a range of thresholds, with either 10 000 or 500 000 runs for each threshold (depending on the required precision), where each run contains a new message. In each run, 32 coefficients are watermarked, and different sets of coefficients are used for the different runs. A large number of runs is utilized so that data can be gathered for an average message and coefficient set, and so as to achieve a higher degree of precision in the probability values. From these data, P_D and P_F curves can be plotted. For the simplified detector, a simulation utilizing 500 000 runs is performed for the case where the detector is specialized to the notched power exponential distribution.

The choice of 32 watermark coefficients is motivated by considering the problem of embedding a watermark within a small portion of an image. A typical image size is 512×512, hence, 32 is only a small fraction of the available coefficients. However, many image coding techniques allow an image to be manipulated in smaller blocks. For example, image transforms can be taken block-wise, and different levels of details may be stored for each block. Thus, it is possible, and often desirable, to embed a watermark in only a reduced region of the image. Furthermore, the small message length results in extremely minimal distortion introduced by the watermarking process. Hence, the choice of 32 coefficients is reasonable.

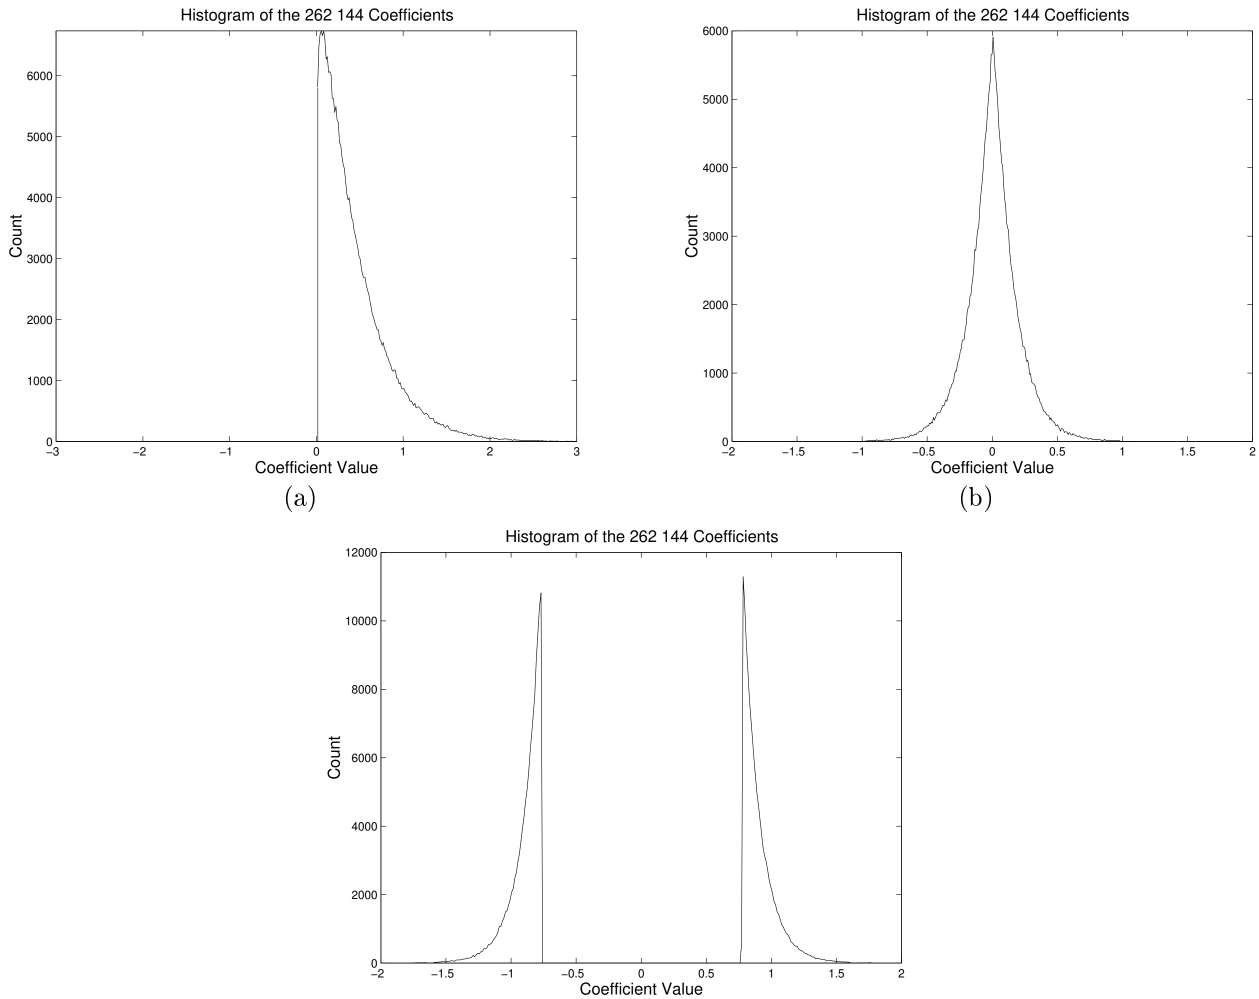

(a) (b) (c)

Figure 5.5 Generated coefficients for the (a) Weibull distribution with $\alpha = 0.4437$ and $\beta = 1.1$, (b) power exponential distribution with $\alpha = 0.1700$ and $\beta = 1.1$, and (c) notched power exponential distribution with $\alpha = 0.1700$, $\beta = 1.1$, $\delta = 0.7760$, and $\kappa = 0.9961$.

5.4.1 Simple likelihood ratio test

The first detector examined is that based on likelihood ratio testing. For each distribution, three figures will be presented. The first will show the distribution of the test statistic under each hypothesis for two embedding strengths, $s^* = 0.05$ and $s^* = 0.10$. Due to the difficulty in computing these distributions analytically, the curves are found using Monte Carlo simulations. The separation between the distributions under H_0 and H_1 determines detector performance.

The second and third figures will present estimates of P_D and P_F based on Monte Carlo simulations, with the corresponding Chernoff bounds overlaid. One figure will contain these results for an embedding strength of $s^* = 0.05$, while the other is for the case when $s^* = 0.10$. In each of these figures, the P_D vs. P_F curves are the best performance indicators. Ideally, these curves will be present in the top left of the graphs, demonstrating that high detection probabilities result for low false alarm probabilities. Finally, since the Monte Carlo results are plotted with the Chernoff bounds, it is observed that for the P_D vs. Threshold and P_D vs. P_F plots, the Monte Carlo curves are necessarily above the Chernoff bounds, while for the P_F vs. Threshold, the experimental curve is below the Chernoff bound.

5.4.1.1 Weibull distribution results

To evaluate the likelihood ratio test based-detector, the host coefficients are first modeled using a Weibull distribution. Figure 5.6 shows the distribution of the test statistic, defined by Equation (3.1), under the two hypotheses. From this figure, it is clear that the H_0 and H_1 distributions are more separated when a higher embedding strength is employed.

Now, the performance of the detector is considered in terms of P_D and P_F curves. Figures 5.7 and 5.8 each contain three plots showing the output of the Monte Carlo simulation compared to the Chernoff bounds for embedding strengths of 0.05 and 0.10, respectively. First, it is noted that the bounds are much tighter (closer to the simulation curve) for extremely high and low probabilities of detection and false alarm. For example, when $s^* = 0.10$ for the simulation point $P_D = 0.8$, the Chernoff bound is only approximately $P_D = 0.2$, while for $P_D = 0.99$, the bound is closer to 0.97. It is in these regions of extreme probabilities where Chernoff bounds most closely bound the performance. However, as

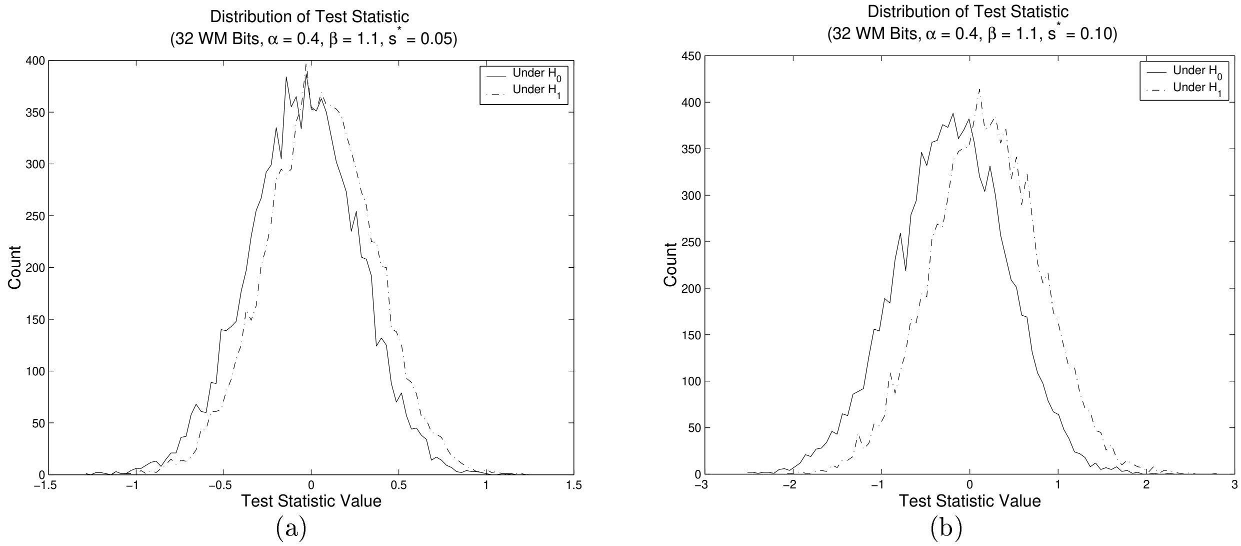

Figure 5.6 Distribution of the test statistic for synthetic data using the WB distribution and LRT with an embedding strength of (a) 0.05 and (b) 0.10.

seen in the previous chapter, the Chernoff bounds on the ROC curves do not convey any information because they are below the trivial bound of $P_D = P_F$. Also, the curves depict a small improvement gained by increasing the embedding strength. For an embedding strength of 0.05, to obtain a detection probability of approximately 65%, the false alarm probability is approximately 50%. Conversely, for an embedding strength of 0.10, this same detection probability is achieved with a lower false alarm probability, just under 40%. Finally, these figures demonstrate that the likelihood ratio test detector using the Weibull distribution does not perform overly well, and is likely not suitable for practical applications.

5.4.1.2 Power exponential distribution results

The likelihood ratio test-based detector is now evaluated when the power exponential distribution is used to model the host coefficients. Figure 5.9 shows the distribution of the test statistic, defined by Equation (3.2), under the two hypotheses. These results are similar to those in the case of the Weibull distribution; the increased embedding strength causes the two distributions to become further apart. Thus, an increase in performance is seen for higher embedding strengths.

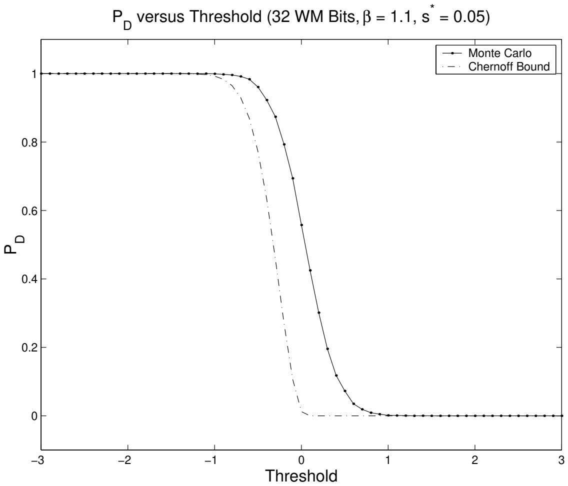

(a)

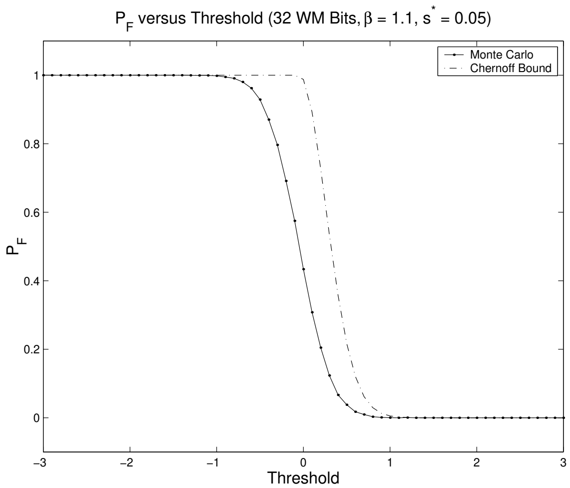

(b)

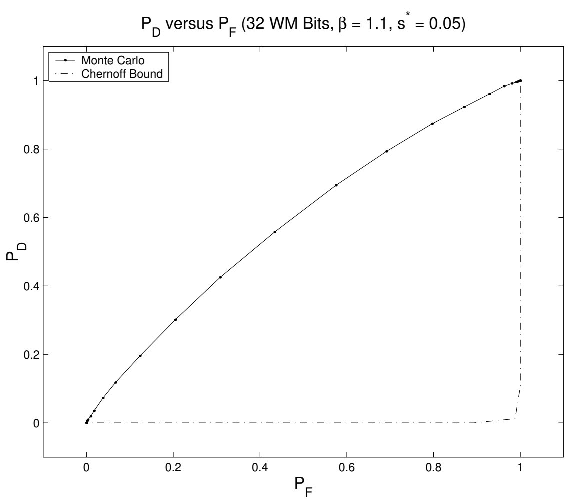

(c)

Figure 5.7 P_D and P_F curves for the WB LRT for synthetic data ($s^* = 0.05$).

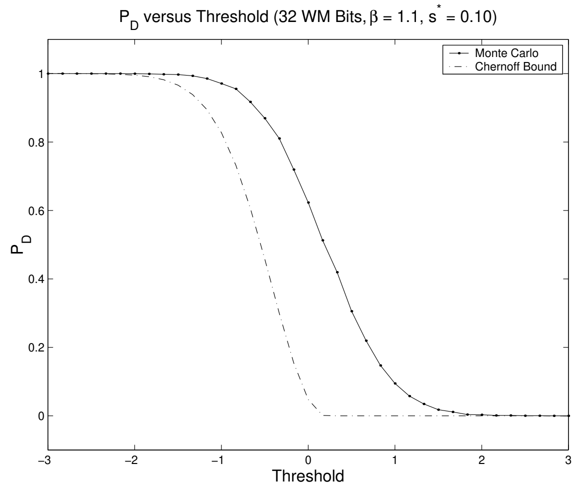

(a)

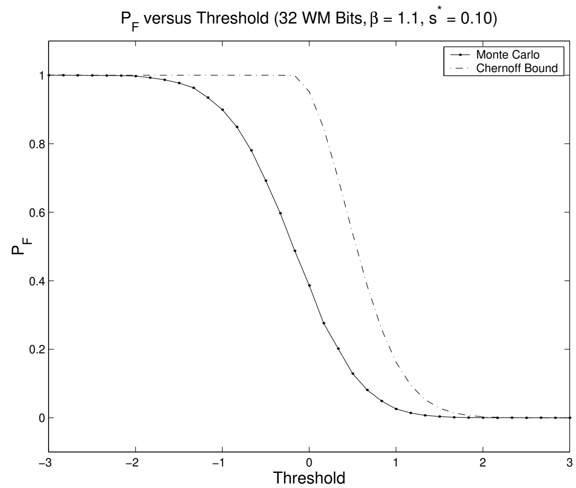

(b)

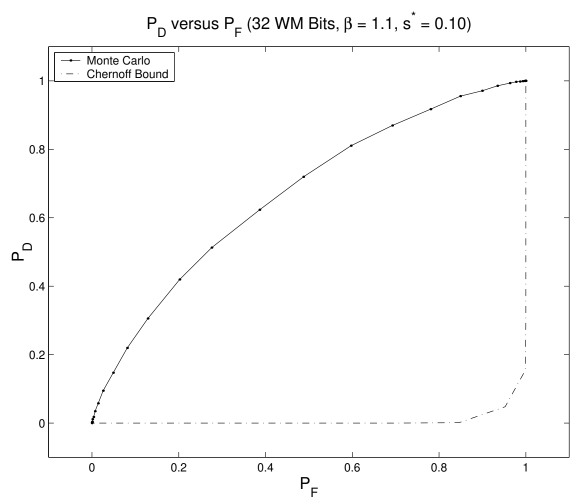

(c)

Figure 5.8 P_D and P_F curves for the WB LRT for synthetic data ($s^* = 0.10$).

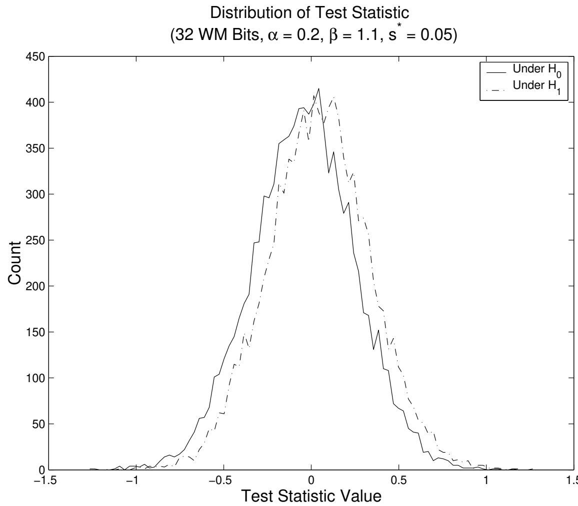

(a)

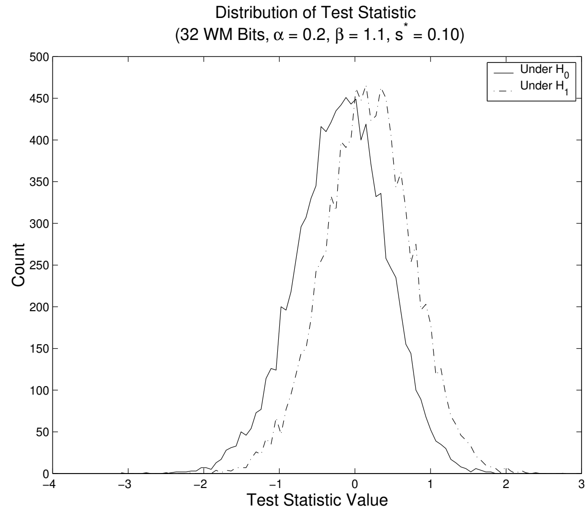

(b)

Figure 5.9 Distribution of the test statistic for synthetic data using the PE distribution and LRT with an embedding strength of (a) 0.05 and (b) 0.10.

Now, the P_D and P_F curves for the Monte Carlo simulations and their Chernoff bounds are considered. Figures 5.10 and 5.11 each contain the three probability plots, the first figure for an embedding strength of 0.05, and the second for 0.10. It can immediately be seen that the power exponential results are almost identical to the Weibull results. In actual fact, the Weibull results are just slightly better. Here, a detection probability of just under 65% is achieved at the cost of approximately a 50% false alarm probability for $s^* = 0.05$. It is again noted that the curves do indeed obey the Chernoff bounds (which are uninformative in the ROC case), and that the bounds become tighter for probabilities near 0 and 1. Also, an increase in performance is observed as the embedding strength is raised. However, it is concluded that the high false alarm probabilities render this detector likely infeasible in a practical sense.

5.4.1.3 Notched power exponential results

The final likelihood ratio test-based detector studied is that developed for the case when the host data are modeled using the newly defined notched power exponential distribution. To begin, the distribution of the decision statistic, defined by Equation (3.3), under each of the two hypotheses is considered. Figure 5.12 gives these distributions for embedding strengths of

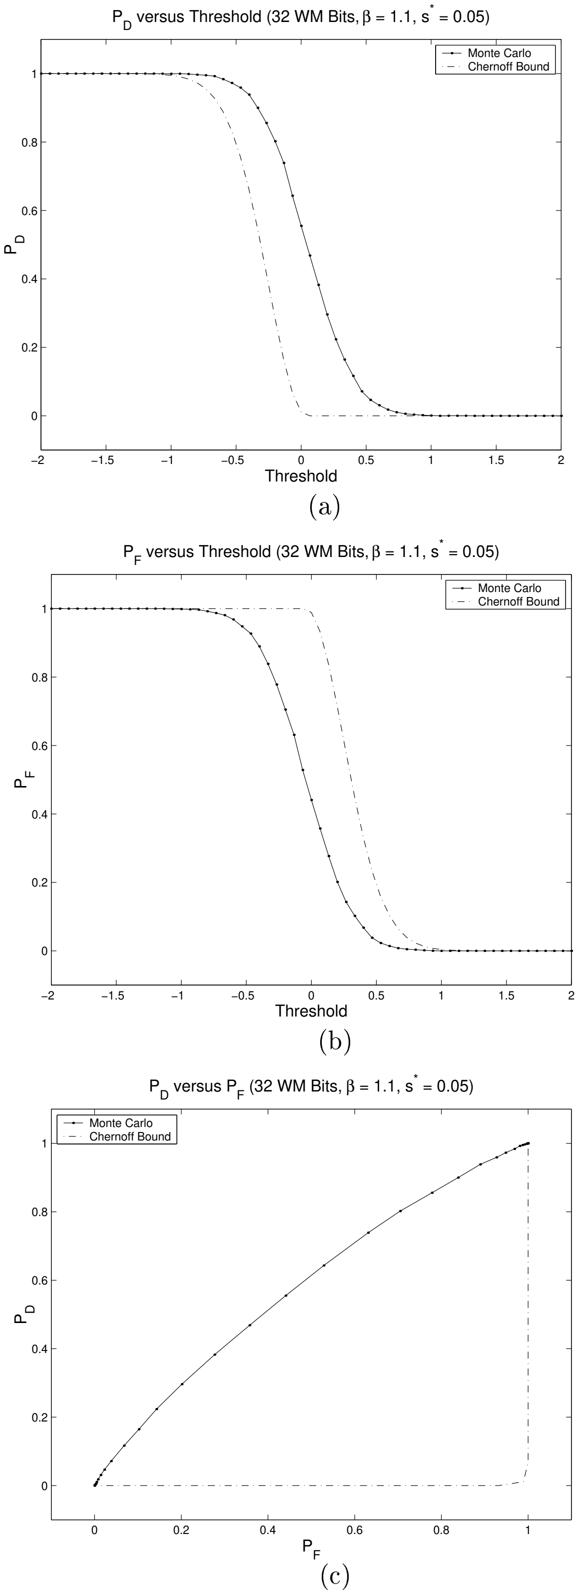

Figure 5.10 P_D and P_F curves for the PE LRT for synthetic data ($s^* = 0.05$).

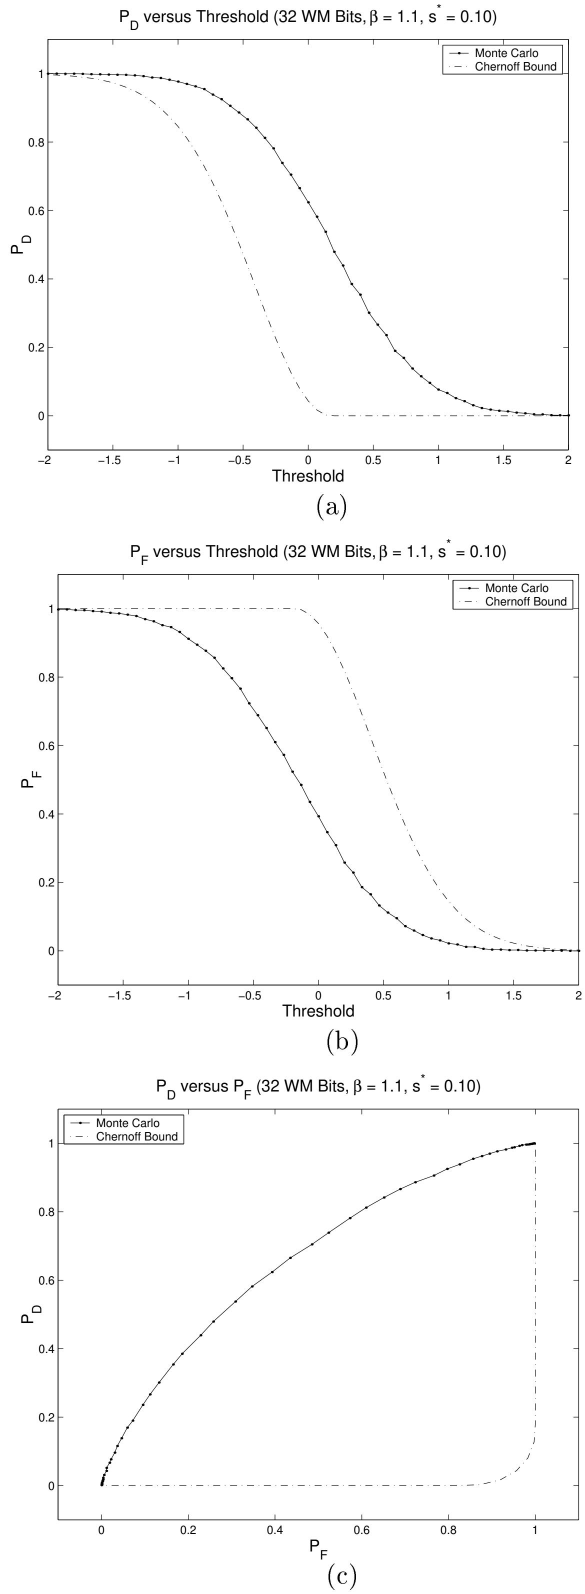

Figure 5.11 P_D and P_F curves for the PE LRT for synthetic data ($s^* = 0.10$).

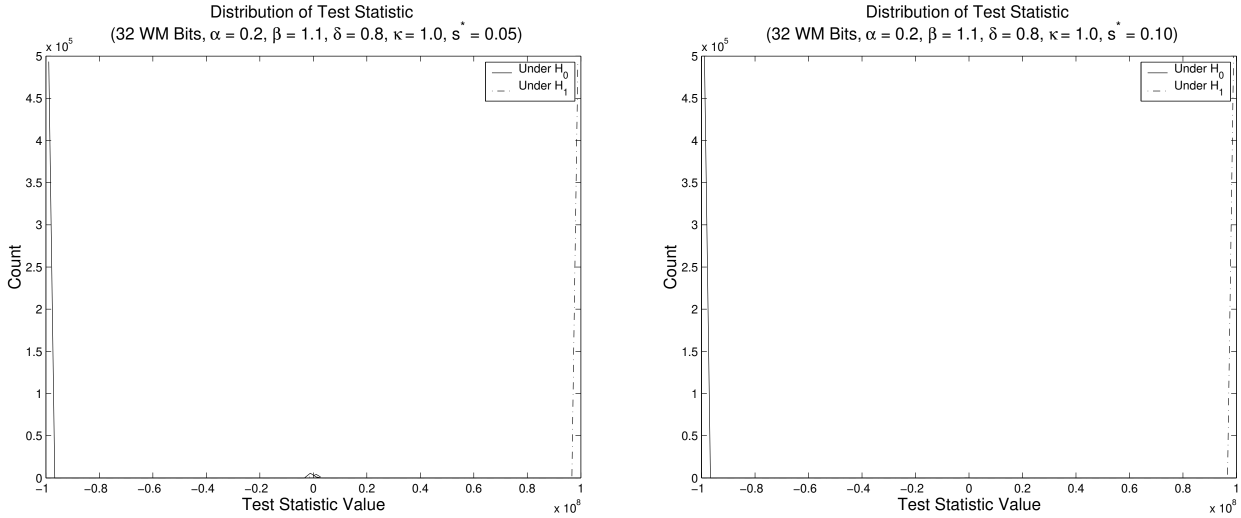

(a) (b)

Figure 5.12 Distribution of the test statistic using for synthetic data the NPE distribution and LRT with an embedding strength of (a) 0.05 and (b) 0.10.

0.05 and 0.10. Since it is difficult to represent the infinities that can be present in the decision statistic for this distribution, the implementation of the detector uses $\pm 1 \cdot 10^8$ instead. In performing this simplification, it is assumed that the noninfinite statistics will be much less than this value, which appears to be quite valid for the simulations in question. From the graphs, it is observed that for both embedding strengths, most of each distribution's mass lies at $+\infty$ or $-\infty$. The congregation of the mass at these infinity points indicates that the majority of times the decision statistic is evaluated, a definitive response is provided for the presence or absence of the watermark (because one hypothesis is not possible). Furthermore, these infinities result when at least one data point lies below the notch threshold, δ_i or the distribution threshold, δ, yielding $-\infty$ and ∞, respectively. For both hypotheses, the curves possess much smaller amounts of mass around the origin. These sections represent the noninfinite statistics in the Monte Carlo simulations, and the plots demonstrate that they are quite diminished as the embedding strength increases. The decrease in mass of the noninfinite statistics with increased embedding strength is caused by the larger strength perturbing the data points greater distances. Thus, it is possible for more points to be moved below the thresholds, making the infinite statistics more probable. From these plots it is clear that the separation between the distributions under H_0 and H_1 is profound.

With the favorable results shown in the distribution of the test statistic, the Monte Carlo simulations and Chernoff bounds for the notched power exponential detector are now considered. Figures 5.13 and 5.14 each contain the three plots used to evaluate detector performance. First, it should be noted that, in the graphs, P_D only ranges from 0.99 to 1 and P_F only ranges from 0 to 0.01 because values outside these ranges do not occur in the bounds or simulation. In order to more accurately represent these precise probabilities, 500 000 runs are utilized in generating the P_D and P_F curves. As expected, the experimental probabilities satisfy the Chernoff bounds. Also, contrary to the Weibull and power exponential distributions, the notched power exponential distribution displays a significant increase in performance as the embedding strength is increased; the range of P_D shrinks to approximately [0.9995, 1], while that of P_F decreases to approximately [0, 0.0005]. The impact of the reduced and shrinking ranges is reflected in the ROC curves, which are now much closer to the top left corner, especially when $s^* = 0.10$. For example, even for the small embedding strength, a detection probability just over 99.8% incurs a false alarm probability of approximately 0.3%; while for an embedding strength of $s^* = 0.10$, a 99.98% detection is achieved with a 0% false alarm probability (to the degree of accuracy provided by the simulation runs). This is an astonishing improvement in performance over the detectors where the host coefficients are modeled using the Weibull or power exponential distributions. The major contributing factor to this improvement is the presence of the δ and δ_i thresholds against which the y_i are compared. If even one y_i value is below one threshold but not the other, one hypothesis can be eliminated immediately. From the plots, it is seen that, even with as few as 32 watermark bits, this situation occurs quite frequently. Thus, employing the notched power exponential distribution to model the data yields a system that could quite easily be used in practical situations.

5.4.2 Simple detector for notched distributions

The simple detector derived from the likelihood ratio test detector is now evaluated when the notched power exponential distribution is employed. The behavior of the detector is examined in terms of the number of times a hypothesis is selected with certainty versus the number of times one is selected at random using the decision rule (3.4). Table 5.2 contains

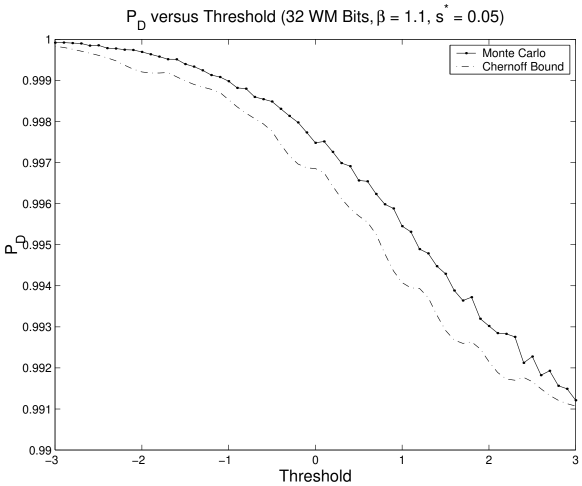

(a)

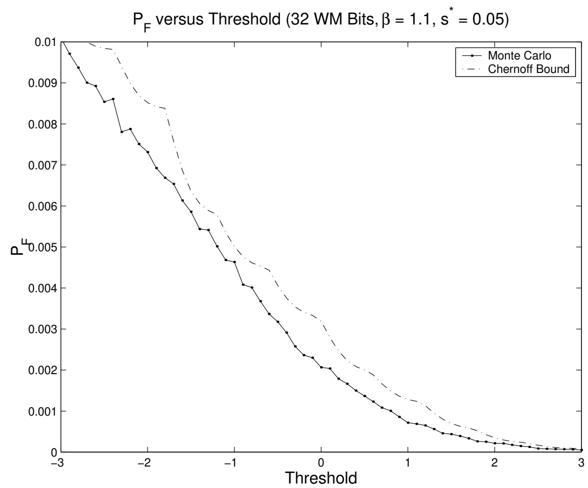

(b)

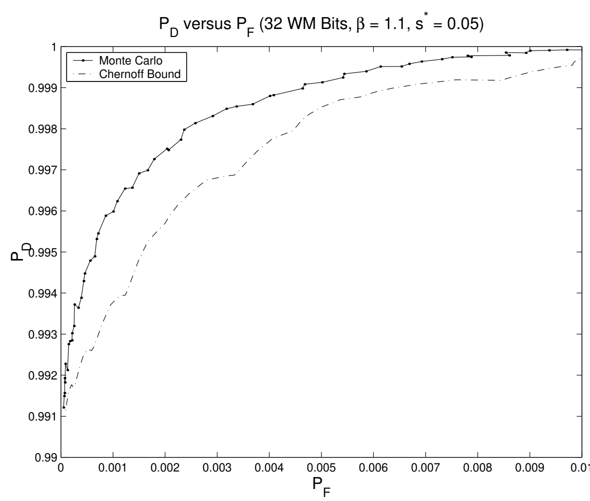

(c)

Figure 5.13 P_D and P_F curves for the NPE LRT for synthetic data ($s^* = 0.05$).

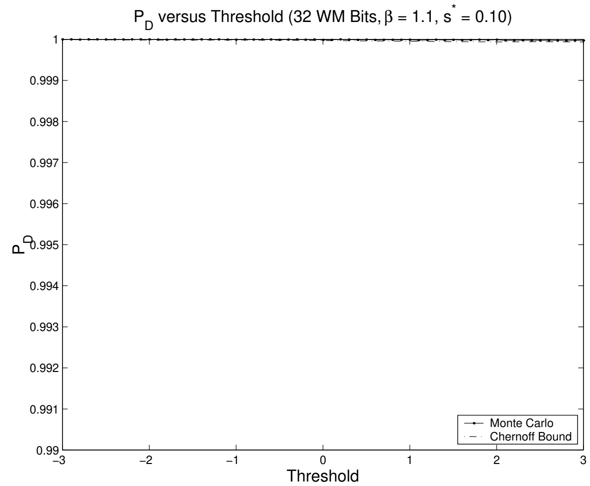

(a)

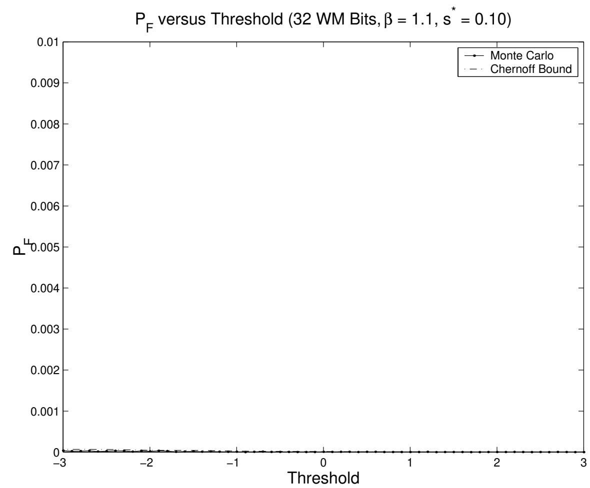

(b)

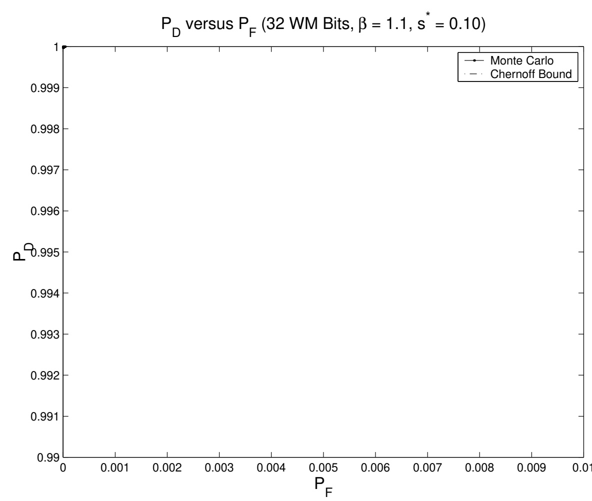

(c)

Figure 5.14 P_D and P_F curves for the NPE LRT for synthetic data ($s^* = 0.10$).

Table 5.2 Simple detector output for synthetically generated data.

Situation	Count	
	$s^* = 0.05$	$s^* = 0.10$
Certain H_0 under H_0	493 529	499 853
Certain H_1 under H_0	0	0
Random selection under H_0	6471	147
Certain H_0 under H_1	0	0
Certain H_1 under H_1	494 862	499 932
Random selection under H_1	5138	68

these counts for the watermark embedding strengths of 0.05 and 0.10 using a value of $\lambda = \frac{1}{2}$. For each embedding strength, 500 000 runs are performed first using data that have not been watermarked (denoted "under H_0"), and then for watermarked data (denoted "under H_1"). From these data, it is clear that a decision is made with certainty for the majority of the trials, with the percentage increasing with the embedding strength. Since the number of times a decision is made randomly is small, the performance of the detector is expected to be quite comparable to that of the likelihood ratio test detector for the NPE distribution.

The comparison between the simple detector and the likelihood ratio test detector is made more explicit by considering the detection and false alarm probabilities over a range of λ values. The ROC curve produced by the simplified detector is overlaid with that resulting from the LRT detector, as shown in Figure 5.15 for both embedding strengths. In both subfigures, the axes have been reduced to reflect these precise probabilities. These plots illustrate the lower performance attained by the simple detector by the presence of its ROC curve to the lower right of the LRT ROC curve. However, the simple detector still performs extremely well, demonstrating the benefit of the use of notched distributions.

5.4.3 Maximum likelihood estimation

The next detector considered is that based upon the technique of maximum likelihood estimation. Because of the poor performance of the Weibull and power exponential distributions for the likelihood ratio test detector, the MLE detectors based on these distributions are rel-

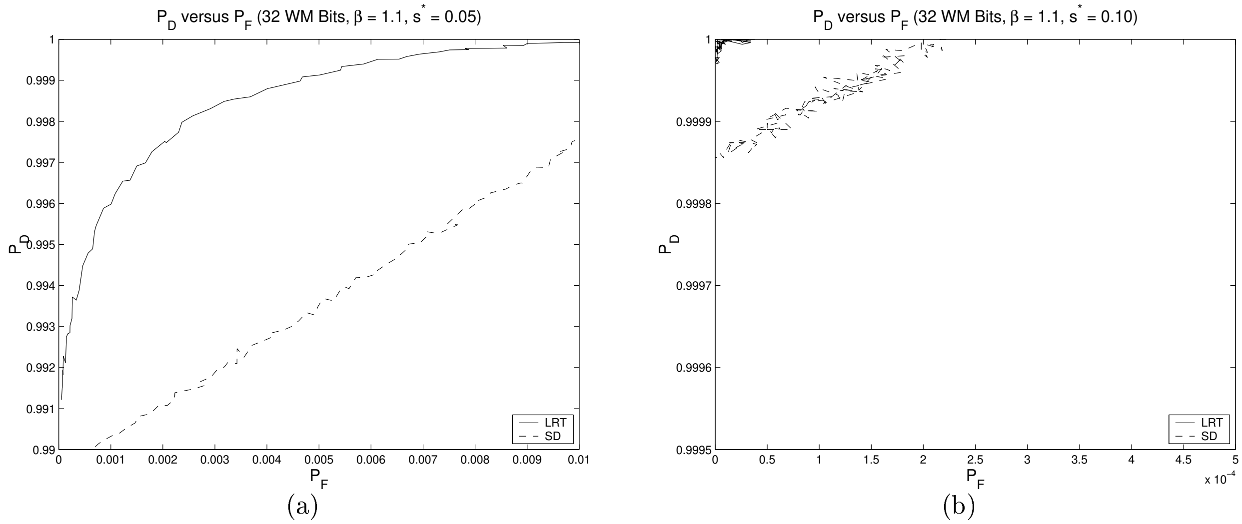

Figure 5.15 ROC curves for the LRT and SD for the NPE distribution with synthetic data using an embedding strength of (a) 0.05 and (b) 0.10.

egated to Appendix A. For the notched power exponential distribution, the three analytical plot types used for the previous detector are again employed to examine detector performance.

The distributions of the test statistics (estimates), defined by Equation (3.7), found through Monte Carlo simulations are presented in Figure 5.16. Unlike the corresponding figures found in the appendix for the Weibull and power exponential distributions, the separation between the curves under H_0 and H_1 is quite significant for the notched distribution. Under hypothesis 1, a clear peak is visible at the true value of the embedding strength. The degree of separation, which is due to this peak, becomes more pronounced as the embedding strength is increased.

To facilitate further examination, curves of the P_D and P_F statistics are given in Figures 5.17 and 5.18. The first notable observation is the rapid decrease in detection probability as the threshold is raised above the actual embedding strength. This result is quite intuitive because it is illogical to threshold a strength estimate using a value significantly above the true strength. Next, by comparing the two figures, an increase in performance is seen for higher embedding strengths, as is expected. For the low embedding strength, a detection probability of approximately 98% can be obtained at the cost of a 5% false alarm

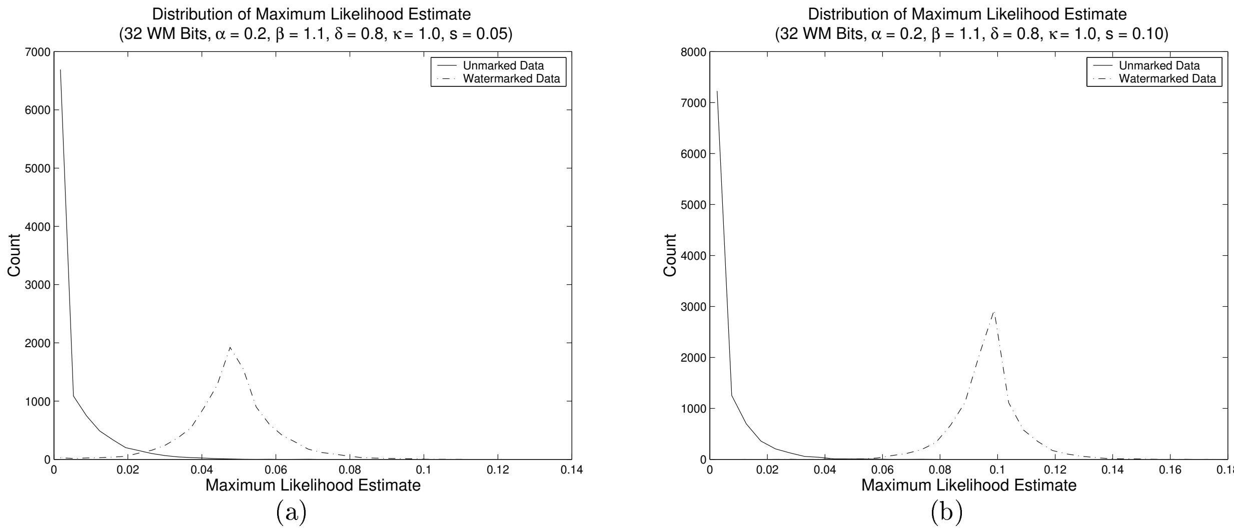

Figure 5.16 Distribution of the test statistic using the NPE distribution and MLE with an embedding strength of (a) 0.05 and (b) 0.10.

probability. While these statistics possibly still render this detector impractical in real-world situations, they do represent a significant improvement over the Weibull and power exponential maximum likelihood estimators. However, the notched power exponential MLE detector performs substantially worse than the likelihood ratio test detector for the same distribution, whose ROC curves are presented in Figures 5.13 and 5.14.

5.4.4 Locally optimal detection

The final detector examined is the locally optimal detector. Once again, only the notched power exponential distribution is considered in this section, while the Weibull and power exponential detectors are analyzed in Appendix A. As before, three figures are presented for the notched power exponential distribution detector in order to analyze the experimental results.

The distributions of the test statistic, defined by Equation (3.10), as found through Monte Carlo simulations, are given in Figure 5.19. These plots demonstrate once again the large degree of separation that the notched power exponential detectors achieve. Nearly all of the distribution mass under H_1 is located at the ∞ point (represented here by $1 \cdot 10^8$), while, under H_0, all the mass is clustered near the origin. An extremely small portion of the

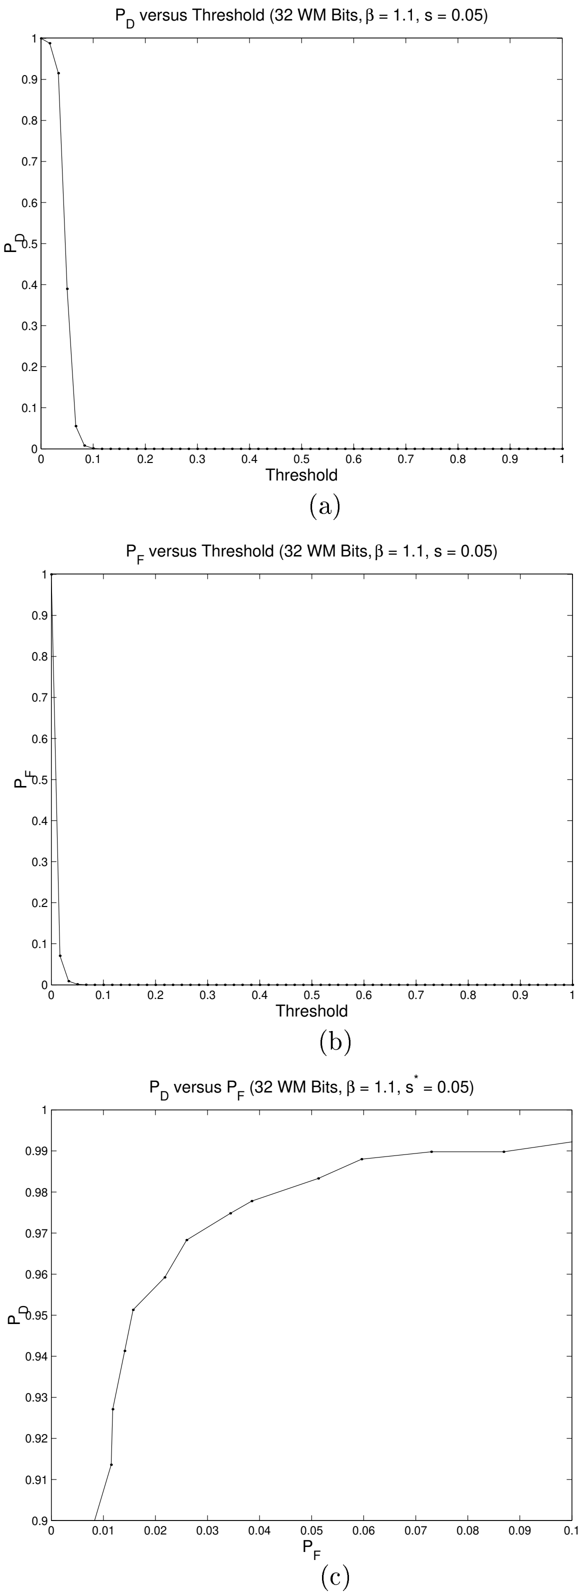

Figure 5.17 P_D and P_F curves for the NPE MLE for synthetic data ($s^* = 0.05$).

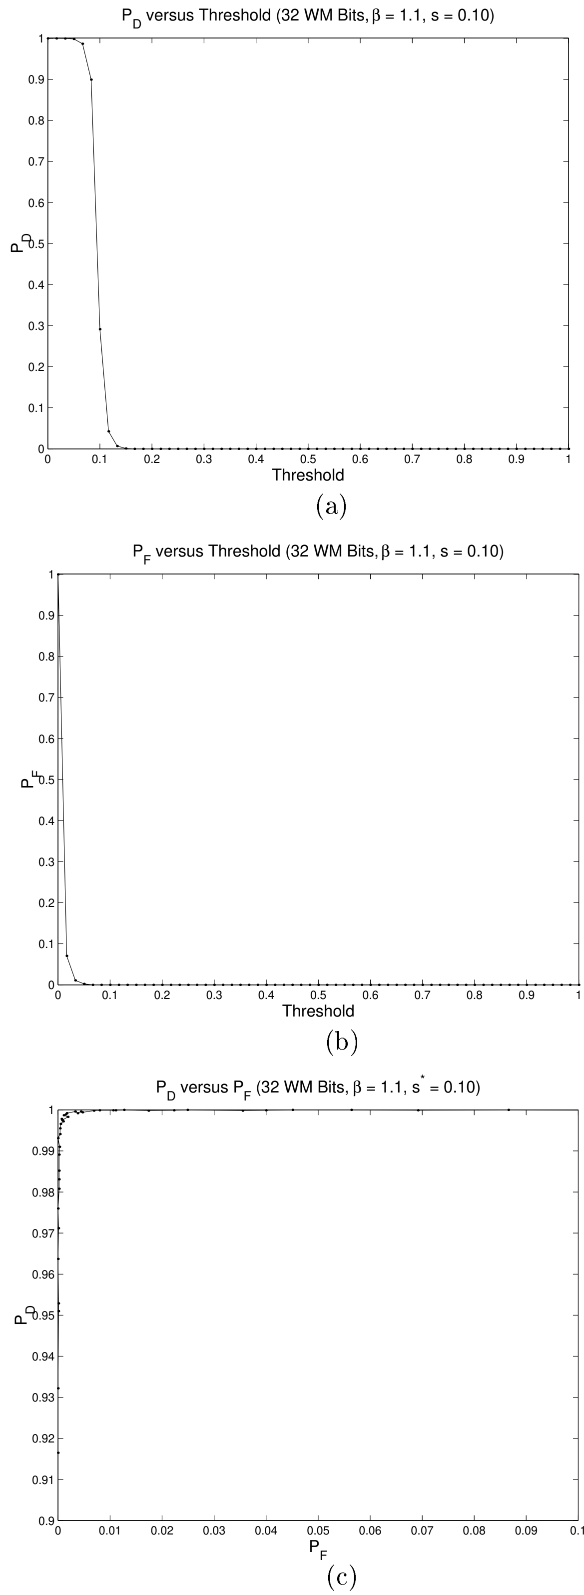

Figure 5.18 P_D and P_F curves for the NPE MLE for synthetic data ($s^* = 0.10$).

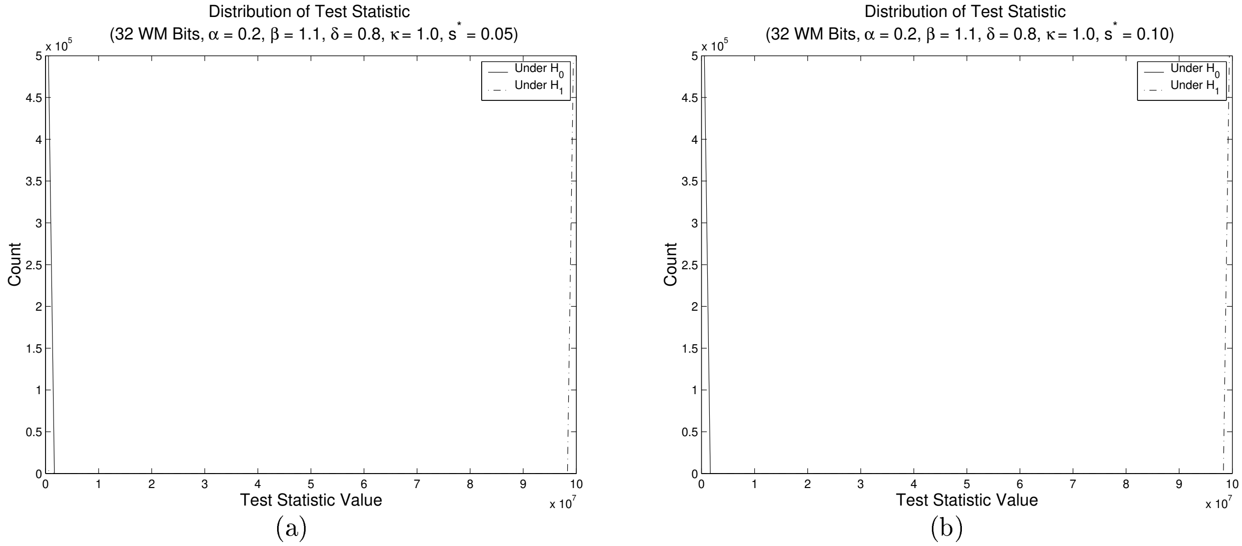

Figure 5.19 Distribution of the test statistic using the NPE distribution and LOD with an embedding strength of (a) 0.05 and (b) 0.10.

H_1 curve is also located near the origin, indicating those trials for which noninfinite decision statistics are found. When the two sub-figures are compared, the small portion of the H_1 curve becomes imperceptible as the embedding strength is increased.

Figures 5.20 and 5.21 contain the P_D and P_F curves resulting from the Monte Carlo simulations. Note that the axes have been limited so that P_D is shown in the range [0.99, 1]. The values taken on by the detection probability are all contained in this range for the simulations. As before, 500 000 runs for each threshold are utilized to achieve a higher degree of accuracy for these small probabilities. These plots clearly indicate that the use of the notched power exponential to model the coefficients has again produced results superior to the other two distributions, especially when the embedding strength is 0.10. For the low embedding strength, this detector achieves approximately a 99.9% probability of detection for approximately a 50% probability of false alarm. When the stronger embedding is considered, the detector responds with approximately a 100% detection probability for a 50% false alarm probability. These values reflect the separation seen in the distributions of the test statistic. This separation is caused by the ease with which H_0 can often be ruled out simply by considering the data points. These results are just slight worse than those for the corresponding likelihood ratio test detector, Figures 5.13 and 5.14. As a result, the use of

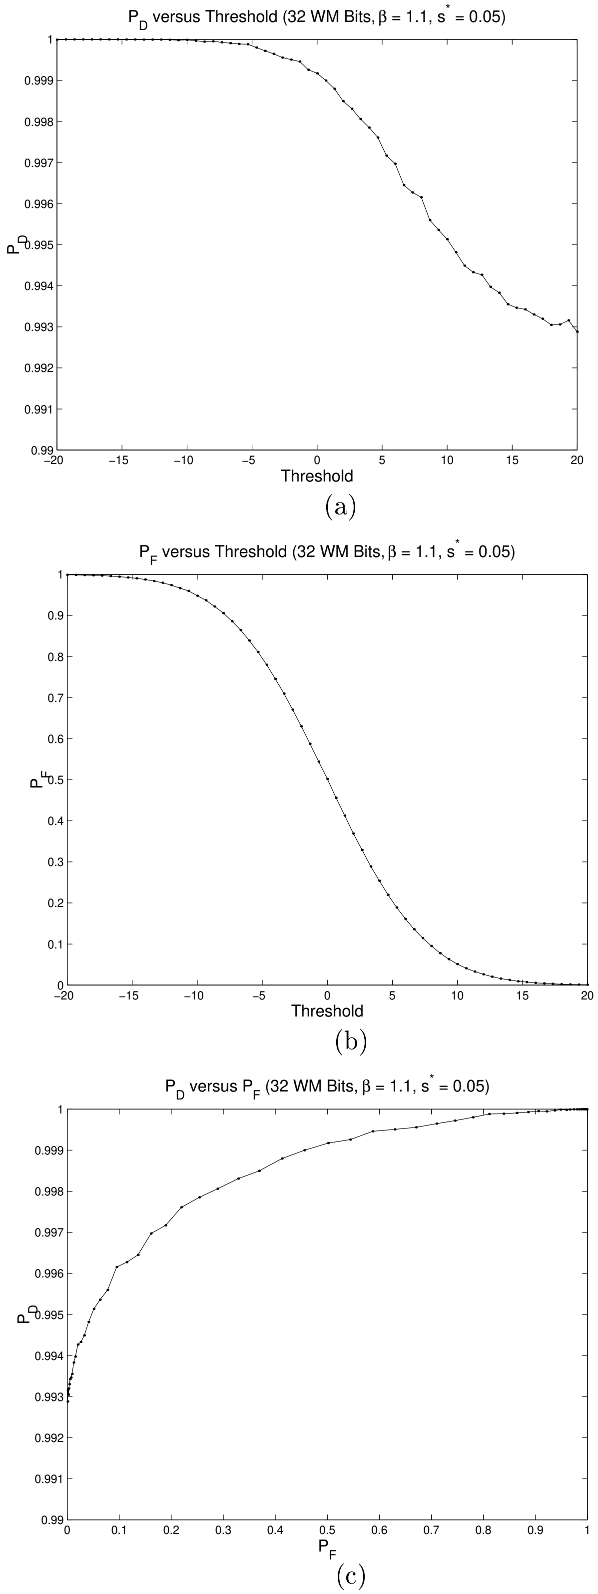

Figure 5.20 P_D and P_F curves for the NPE LOD for synthetic data ($s^* = 0.05$).

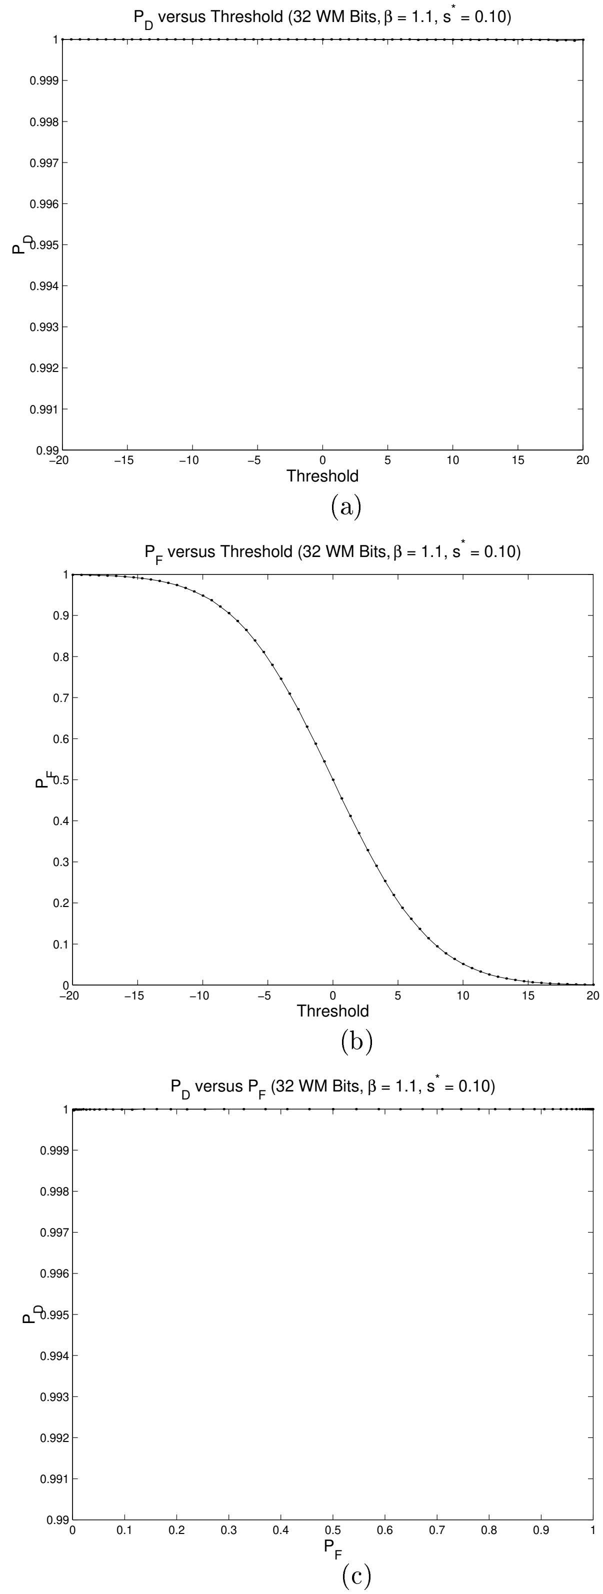

Figure 5.21 P_D and P_F curves for the NPE LOD for synthetic data ($s^* = 0.10$).

the notched power exponential with locally optimal detection would not likely be used in real-world situations.

5.5 Summary of Simulation Results

This chapter has presented the results of simulations conducted on the various detector-distribution pairs that are developed in this thesis. Each of the detectors was evaluated using data generated synthetically according to the appropriate distribution. The data shown in the previous figures and those in the appendix that examine the detectors in terms of their probability of detection and false alarm are now summarized in Table 5.3. The probabilities are shown in percentage form, rounded to the nearest whole number. Data points have been selected from the curves in an attempt to facilitate comparison between the detectors and distributions. More specifically, points have been chosen to demonstrate the value of P_F that must be tolerated in order to achieve a P_D of approximately 98%. For the detectors for which such a detection probability is not possible, the point corresponding to the highest possible P_D value is included.

Table 5.3 Selected results for synthetically generated data.

Distribution	Detector	$s^* = 0.05$		$s^* = 0.10$	
		$P_D(\%)$	$P_F(\%)$	$P_D(\%)$	$P_F(\%)$
Weibull	LRT	98	96	98	93
Weibull	MLE	58	50	67	49
Weibull	LOD	98	95	98	92
Power exponential	LRT	98	97	98	93
Power exponential	MLE	58	50	66	49
Power exponential	LOD	98	96	98	93
Notched power exponential	LRT	99	0	100	0
Notched power exponential	MLE	98	0	98	0
Notched power exponential	LOD	99	0	100	0
Notched power exponential	SD	99	0	100	0

A number of general conclusions can be drawn from the data contained in the table. The main observation to be made is the difference in performance between the detectors using the newly developed notched power exponential distribution and those based on the Weibull or power exponential distributions. The addition of the notch thresholds (δ and δ_i) in the NPE distribution provides additional information for detecting the watermark. By comparing the data points to these thresholds, it is often possible to completely rule out a hypothesis. Thus, the results agree with the intuitive conclusion that the notched power exponential provides a superior means of modeling the coefficients selected for watermarking, and thus yields superior results. To further illustrate the high quality of the notched power exponential distribution, Figure 5.22 contains plots of the ROC curves resulting from the LRT detector for each of the three modeling distributions. These plots are depicted on a log scale using $P_D - 1$ to illustrate the small probabilities under the NPE distribution. As a result, the notched power exponential distribution data appear rather jagged for the higher embedding strength because of the finite number of simulations performed. Also, the Weibull and power exponential curves lie so close together that it is difficult to discriminate between them at this scale.

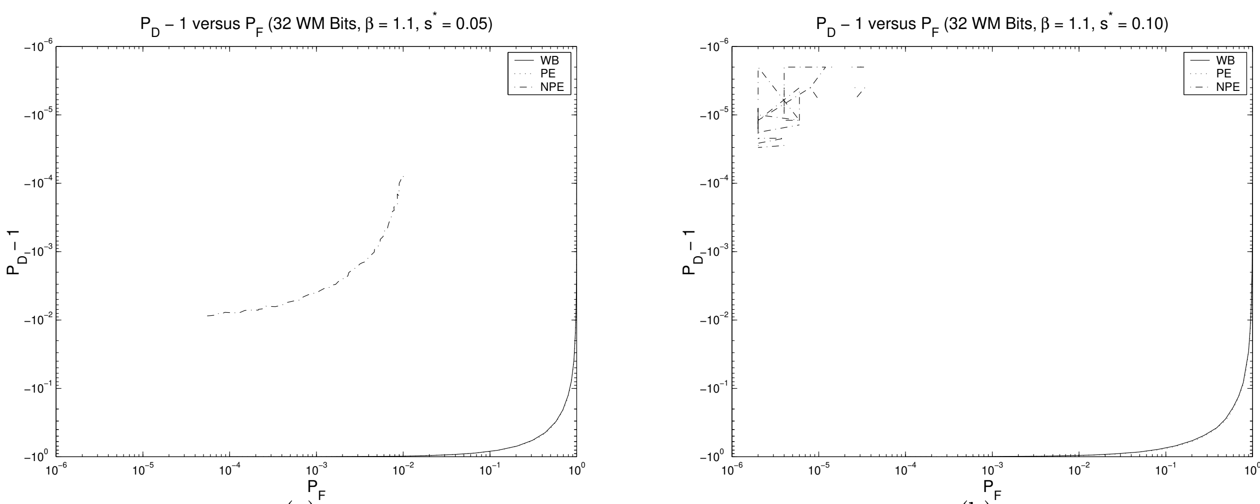

(a) (b)

Figure 5.22 ROC curves for the LRT using the WB, PE, and NPE for embedding strengths of (a) 0.05 and (b) 0.10.

It is also possible to compare solely the detectors developed in this thesis for a particular distribution. First, it is noted that the maximum likelihood estimation detectors are often not capable of producing high detection probabilities, regardless of the false alarm probability. This result is due to the fact that the detector is actually estimating the watermark embedding strength. Thus, reasonable threshold values are restricted to the range [0, 1). Therefore, the highest detection probability is that found using a threshold of 0, which need not be close to 100%. The other generalization that can be made is that the likelihood ratio test detector tends to perform better than the other three detectors. Also, for the NPE distribution, the simple detector performs just slightly worse than its optimal LRT counterpart. To highlight the performance difference between the detectors, Figure 5.23 contains plots of the resulting ROC curves for the notched power exponential distribution, for embedding strengths of 0.05 and 0.10. Note that a log scale has been utilized and, hence, $P_D - 1$ is plotted instead of P_D. The curves appear somewhat jagged in these plots, in particular for the higher embedding strength because of the finite number of Monte Carlo simulations performed. However, these graphs clearly demonstrate the superior performance of the LRT detector and the poor performance of the MLE detector.

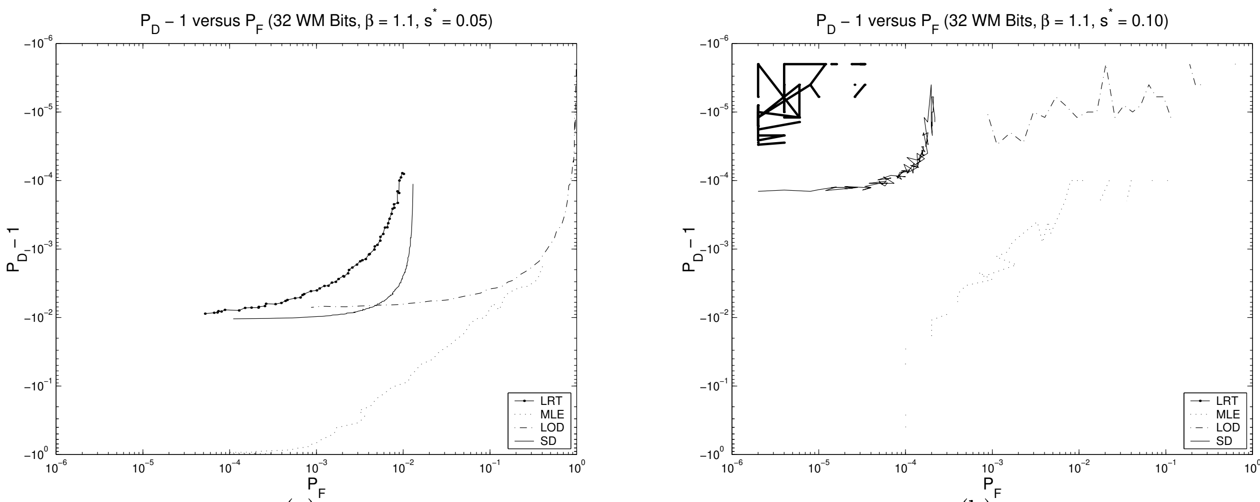

(a) (b)

Figure 5.23 ROC curves for the LRT, MLE, LOD, and SD using the NPE for embedding strengths of (a) 0.05 and (b) 0.10.

Combining these observations yields the overall conclusion that the likelihood ratio test detector derived for coefficients modeled by the notched power exponential distribution is far better suited for watermark detection than the other detector-distribution pairs considered in this thesis.

CHAPTER 6

IMAGE WATERMARKING EXPERIMENTS

The previous chapter examined the detectors based on likelihood ratio testing, maximum likelihood estimation, and locally optimal detection using data synthesized to follow the appropriate distribution. Now, these detectors are applied to standard real-world test images to ascertain how well they perform on coefficients that do not necessary follow the distributions upon which they are based.

6.1 Implementation Details

The image watermarking experiments performed utilize the procedure described in Section 2.2. As in the case of the synthetic data, 262 144 coefficients are used, but here the coefficients are found by performing the discrete cosine transform of a 512×512 test image. For simplicity, the fractional size of the candidate set, ν, is taken to be 1, so that the entire set of transform coefficients can be included in the watermarking process.

A fixed threshold of $\delta = 0.7760$ is employed, resulting in a value of $N = 32$ for the test images considered. The other distribution parameters, α and β, must be estimated because the true distribution of the coefficients is unknown. As described in Section 2.3, these parameters can be found by considering the sample moments of the data. However, during experimentation, it was found that the detectors yield better performance if the value of the β parameters is selected to be around 1, rather than calculated from the data. Thus, for experimentation, a fixed value of $\beta = 1.1$ is utilized to improve performance without overly simplifying the decision statistics. Using this selected β, α can then be estimated from the sample moments.

With the details of the problem setup formalized, the performance of the detectors on actual image data can now be analyzed. Because of the overwhelmingly superior performance of the notched power exponential distribution with synthetic data, this chapter focuses only on its use for modeling the image coefficients. However, the Weibull and power exponential distributions are considered in Appendix B. For each detector type, a Monte Carlo simulation is performed to analyze the performance in terms of detection and false alarm probabilities. The simulations are performed over a range of thresholds, with either 10 000 or 40 000 runs for each threshold (depending on the required precision), where a new message is generated for each trial.

6.2 Watermarking Lena DCT Coefficients

To commence the study of the detectors using actual image data, coefficients are drawn from the DCT of the Lena test image, which is shown in Figure 6.1. Histograms of the 262 144 DCT coefficients and those selected for watermarking are displayed in Figure 6.2. The graph depicted for the 32 coefficients is rather blocky due to the limited amount of data being plotted.

The estimated values for the distribution parameters for the Lena coefficients are calculated using a fixed value of $\beta = 1.1$, and are shown in Table 6.1 for all three of the possible modeling distributions.

Table 6.1 Estimated distribution parameters for the Lena DCT coefficients ($\beta = 1.1$).

Distribution	Parameters		
	α	δ	κ
Weibull	0.4437	-	-
Power exponential	0.1700	-	-
Notched power exponential	0.1700	0.7760	0.9961

Figure 6.1 The Lena image.

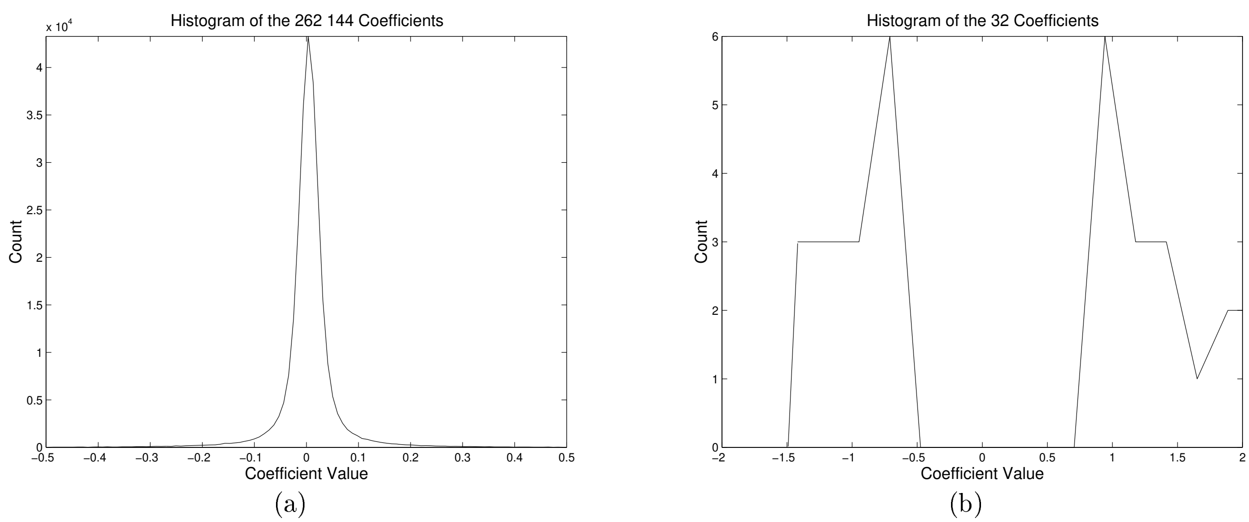

Figure 6.2 (a) DCT coefficients of the Lena image and (b) those coefficients selected for watermarking.

6.2.1 Simple likelihood ratio test

The setup used to present the likelihood ratio test results for the Lena data is identical to that for synthetic data: three figures are presented demonstrating the distribution of the test statistic, and the detection and false alarm probabilities for embedding strengths of $s^* = 0.05$ and $s^* = 0.10$.

To begin, the distributions of the decision statistic for the notched power exponential, as found through Monte Carlo simulations, are considered. Figure 6.3 gives these distributions for embedding strengths of 0.05 and 0.10. For the ease of viewing, the infinities present in the decision statistic have again been represented by the values $\pm 1 \cdot 10^8$. The graphs demonstrate that most of the mass of the distribution lies at these infinite endpoints, with a small amount near the origin. The small sections represent the non-infinite values of the statistic and diminish as the embedding strength is increased. From these figures, it is clear that there is an extreme degree of separation between the distributions under H_0 and H_1.

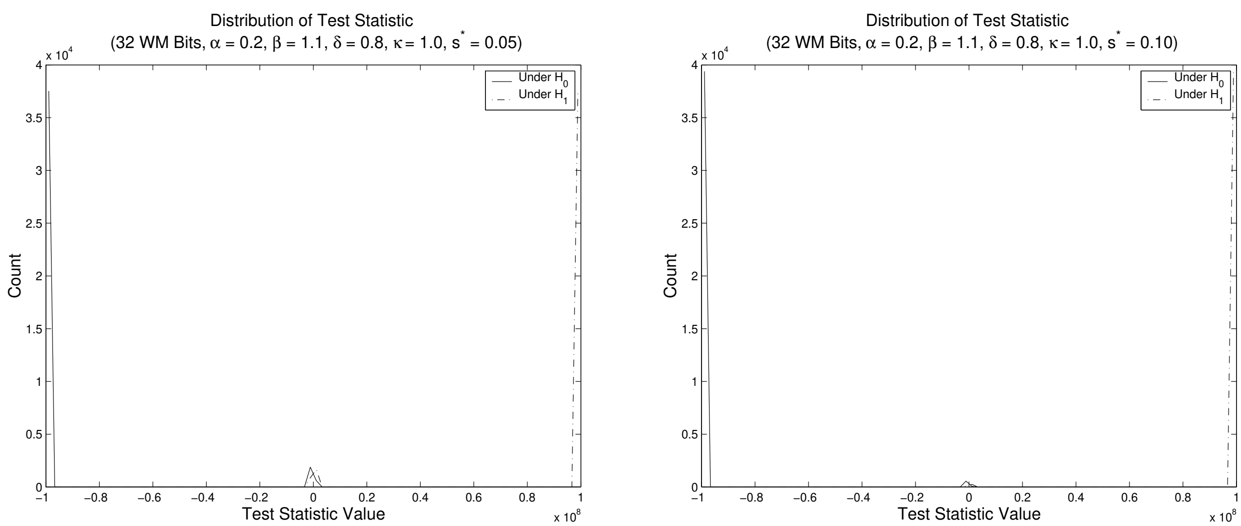

(a) (b)

Figure 6.3 Distribution of the test statistic for the Lena image using the NPE distribution and LRT with an embedding strength of (a) 0.05 and (b) 0.10.

At this point, the P_D and P_F curves are examined for the Monte Carlo simulations on the Lena data. Figures 6.4 and 6.5 each contain the three plots of the detection and false alarm probabilities. It should be noted that the curves are graphed on a reduced scale because of the small probability values. P_D is only plotted over the range of 0.9 to 1, while P_F is

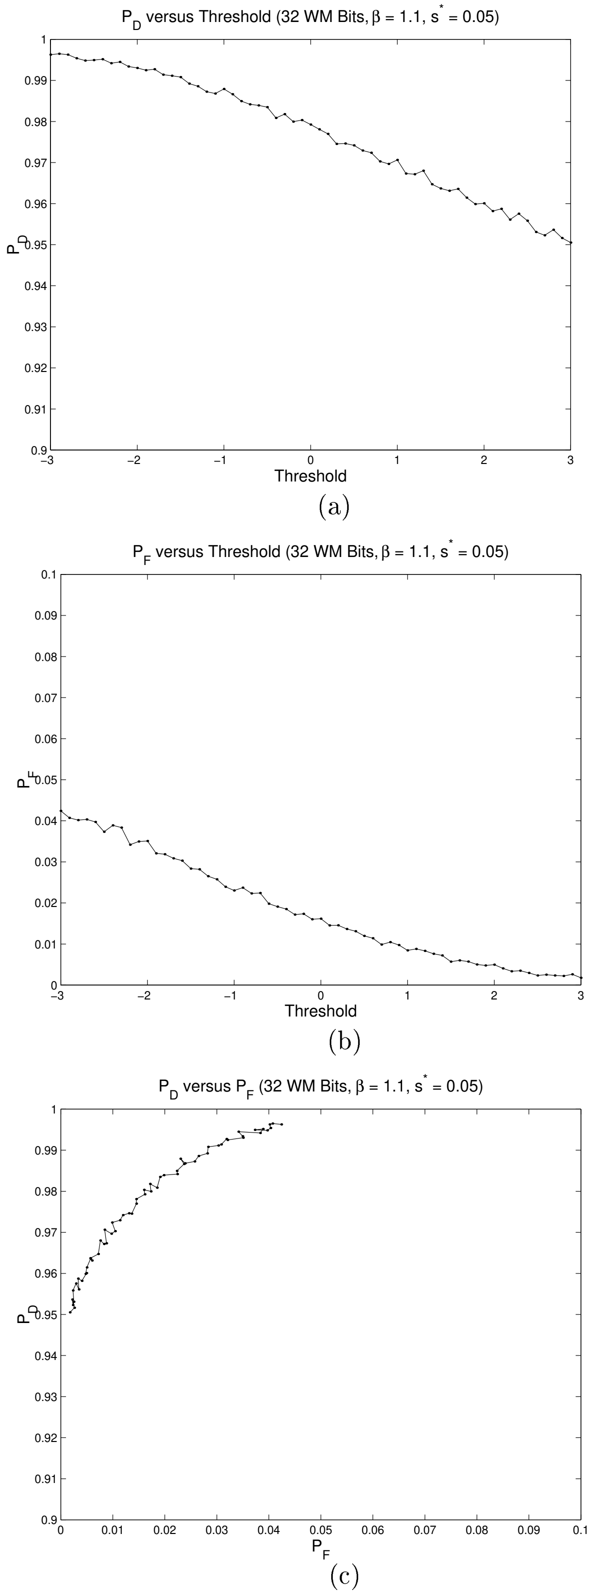

Figure 6.4 P_D and P_F curves for the NPE LRT for Lena data ($s^* = 0.05$).

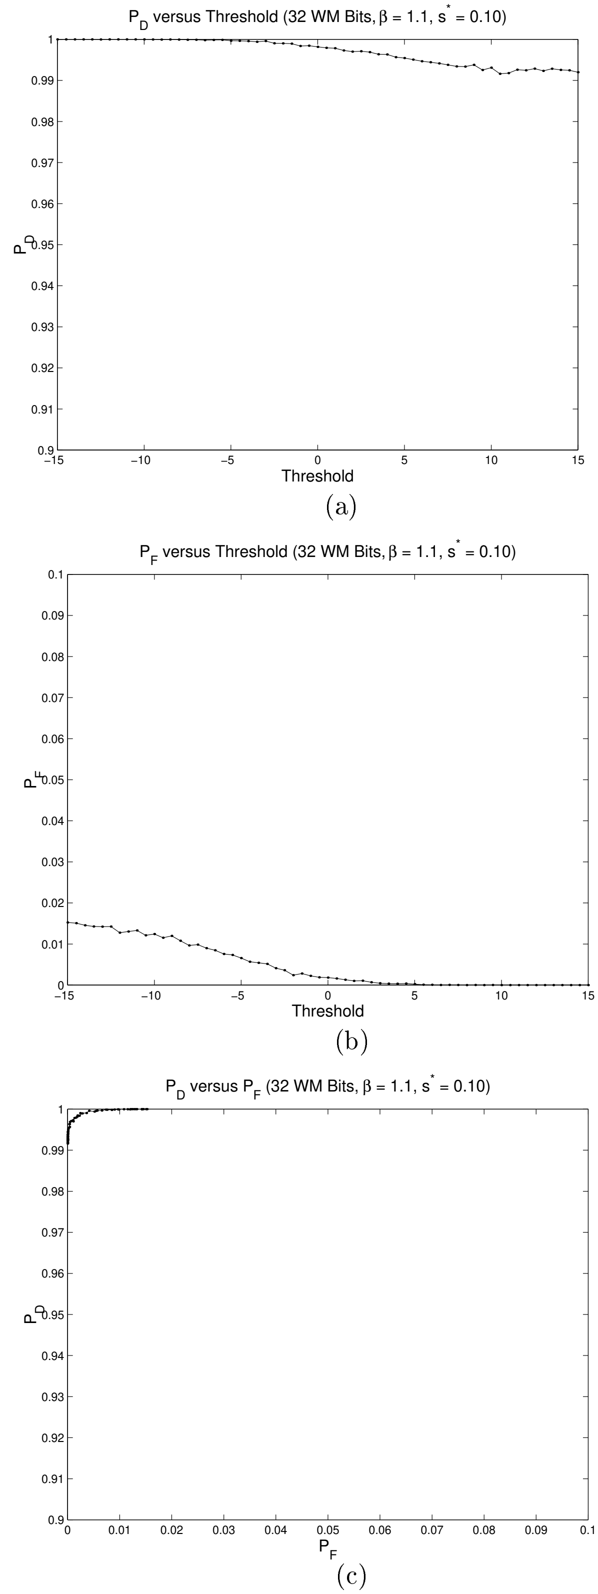

Figure 6.5 P_D and P_F curves for the NPE LRT for Lena data ($s^* = 0.10$).

shown between 0 and 0.1. Values outside these ranges do not occur in the simulation results. Since the resultant probabilities are quite small, the curves appear slightly jagged because of the finite number of simulations performed. By comparing the plots in the two figures, it is seen that, unlike with the Weibull and power exponential distributions considered in the appendix (Figures B.2, B.3, B.5, and B.6), a significant increase in performance is achieved by increasing the embedding strength. However, even at the low embedding strength, the detector performs extremely well; a detection probability of approximately 98% is achievable with a false alarm probability of only 2%. These favorable values are reflected in the ROC curves, which are situated near the top left corner of a standard range plot. The plots clearly indicate that using the notched power exponential distribution to model the coefficients of the Lena image results in a powerful detector. The considerable performance gain with this distribution over the Weibull and power exponential is attributed to the ability to eliminate a hypothesis if a data point is found to lie below one of the notch thresholds, δ or δ_i. It is also expected, in an intuitive sense, that the notched distribution will achieve better results than either the Weibull or power exponential distributions because only the largest magnitude coefficients are selected for watermarking. Thus, if the entire distribution of the DCT coefficients is well modeled using a power exponential distribution, the notched distribution should be the natural choice to model the few selected for watermarking. Consequently, the results demonstrate that the likelihood ratio test detector based on the notched power exponential distribution can quite conceivably be used in practical situations.

It is also interesting to compare these ROC curves with those presented when the LRT detector is used with the notched power exponential on synthetic data (Figures 5.13 and 5.14). In the case of the synthetic data with $s^* = 0.05$, for P_D values lying in the range [0.99, 1], P_F values only occur between 0 and 0.01. These data represent a decreased performance when the detector is utilized with the Lena coefficients. A similar drop is noticed when $s^* = 0.10$. These drops in performance indicate that the Lena DCT coefficients do not follow a notched power exponential distribution exactly. However, the use of this distribution to model that of the coefficients does indeed produce a detector of high quality.

6.2.2 Simple detector for notched distributions

The simplified detector specialized to the notched power exponential distribution is now considered for use in detecting a watermark in the Lena image. As in the case for synthetic data, a table is employed to demonstrate the behavior of the detector for $\lambda = \frac{1}{2}$. Table 6.2 contains counts of the decisions produced by the rule (3.4) under each of the two hypotheses. Although not as frequently as for synthetic data, a decision is made with certainty for the majority of the 40 000 simulation trials. Thus, strong performance is expected from the detector.

Table 6.2 Simple detector output for data from the Lena DCT coefficients.

Situation	Count	
	$s^* = 0.05$	$s^* = 0.10$
Certain H_0 under H_0	37 518	39 370
Certain H_1 under H_0	0	0
Random selection under H_0	2480	630
Certain H_0 under H_1	0	0
Certain H_1 under H_1	37 496	39 653
Random selection under H_1	2504	347

To further the comparison between the simplified detector and the likelihood ratio test detector, plots of the ROC curves for the LRT detector are shown in Figure 6.6 with those for the simple detector overlaid. Note that the axes have been truncated in both subfigures to more clearly illustrate the precise probability values. In both subfigures, the simple detector curve lies to the lower right of the LRT ROC curve, indicating the decrease in performance. However, the simplified detector still provides strong results for detecting a watermark embedded in a test image.

6.2.3 Maximum likelihood estimation

Maximum likelihood estimation of the embedding strength, s, with the notched power exponential distribution is now considered for detecting a watermark that has been embedded

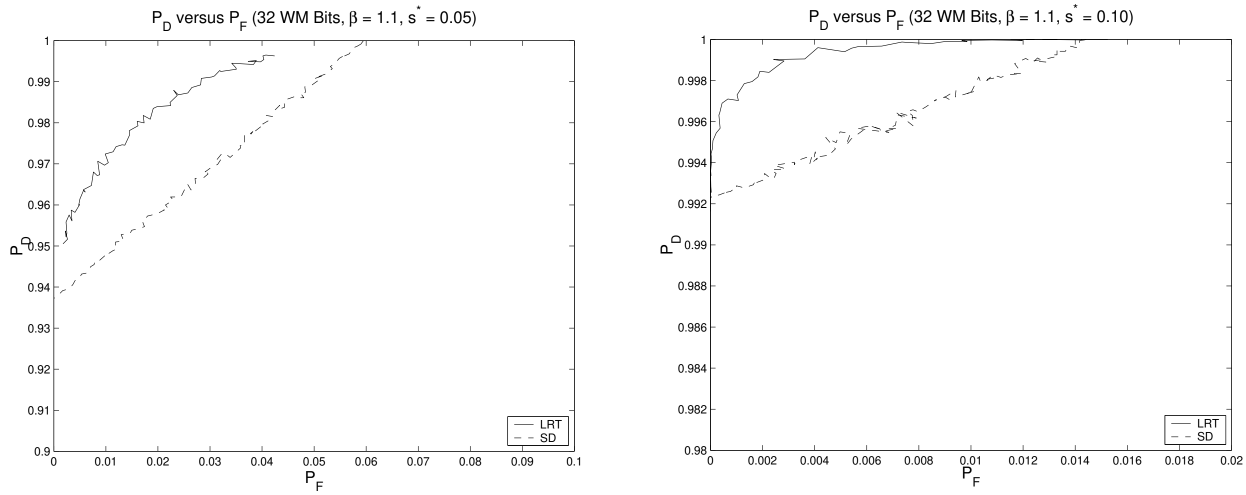

(a) (b)

Figure 6.6 ROC curves for the LRT and SD for the NPE distribution with Lena data using an embedding strength of (a) 0.05 and (b) 0.10.

in the DCT coefficients of the Lena image. As with the other detectors, three figures are presented for performance analysis.

The distributions of the test statistics resulting from 10 000 Monte Carlo simulations are provided in Figure 6.7. Similar to the synthetic data case, the competing distributions are well separated. For the case of H_1, a clear peak is observed in the estimate around the true strength value. The location of this peak is more pronounced and further separated from the H_0 distribution when a higher embedded strength is employed.

The performance of the detector is now examined in the context of P_D and P_F curves. These plots are shown in Figures 6.8 and 6.9 for embedding strengths of 0.05 and 0.10, respectively. Once again, a sharp decrease in the probability of detection is observed as the threshold is moved above the true embedding strength, as expected. By comparing the ROC curves in the two figures, it is apparent that the detector performs better when a higher embedding strength is utilized. However, even when the true strength is 0.05, the detector achieves a 98% detection probability for a false alarm probability of approximately 50%. Although these results are not as strong as the corresponding likelihood ratio test detector (Figures 6.4 and 6.5), they do represent a significant improvement over the MLE detectors for Weibull and power exponential distributions (Figures B.8, B.9, B.11, and B.12) .

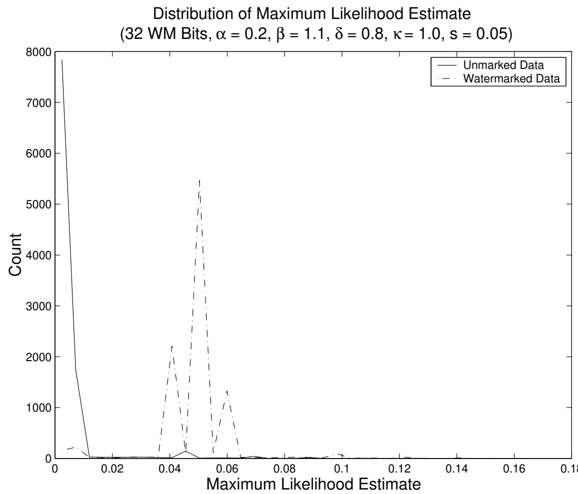

(a)

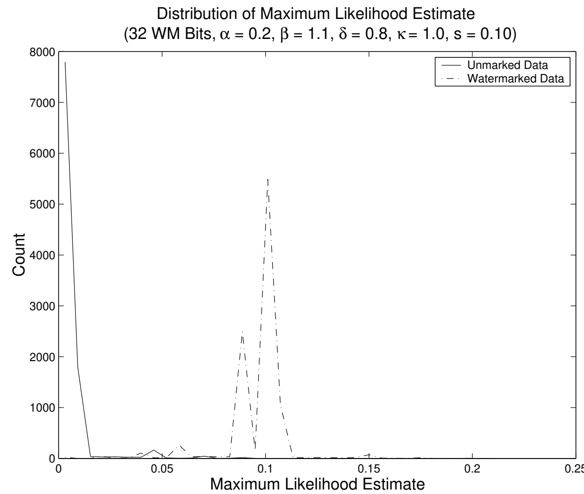

(b)

Figure 6.7 Distribution of the test statistic for the Lena image using the NPE distribution and MLE with an embedding strength of (a) 0.05 and (b) 0.10.

However, because of the stronger performance of the LRT detector, the maximum likelihood estimation detector using the notched power exponential distribution does not likely yield adequate robustness for use in practical applications.

The ROC curves resulting from the application of the MLE detector on actual watermarked image coefficients are now compared with those for the synthetic data case, Figures 5.17 and 5.18. As expected, the ROC curve is more favorable for the synthetic data than the Lena coefficients. However, the use of the notched power exponential distribution to model the image transform coefficients still produces a reasonably accurate detector.

6.2.4 Locally optimal detection

The final detector examined for use in watermark detection for the Lena image coefficients is that based on locally optimal detection. As before, three figures are presented to analyze the performance of the locally optimal detector resulting from the use of the notched power exponential distribution.

The distributions of the test statistic under each of the two hypotheses are given for embedding strengths of 0.05 and 0.10 in Figure 6.10. From this figure, it is clear that most of the distribution mass under H_1 lies at the endpoint ∞, which is represented by $1 \cdot 10^8$ for

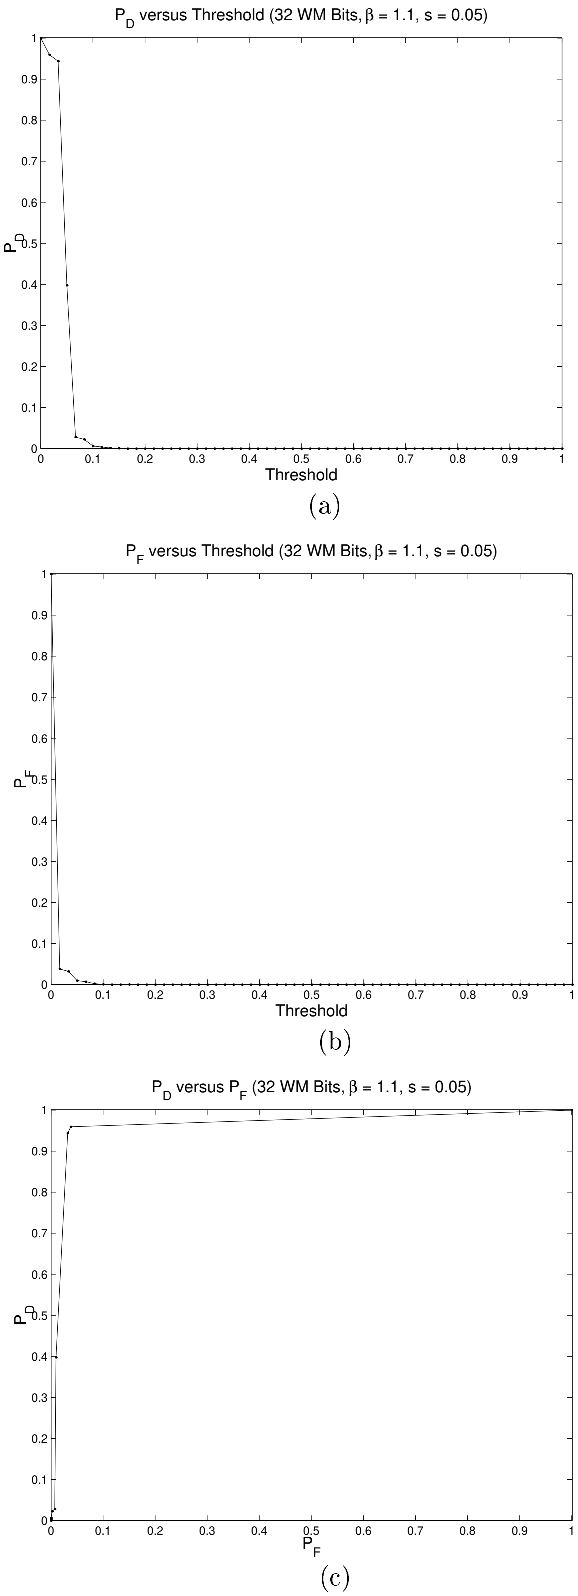

Figure 6.8 P_D and P_F curves for the NPE MLE for Lena data ($s^* = 0.05$).

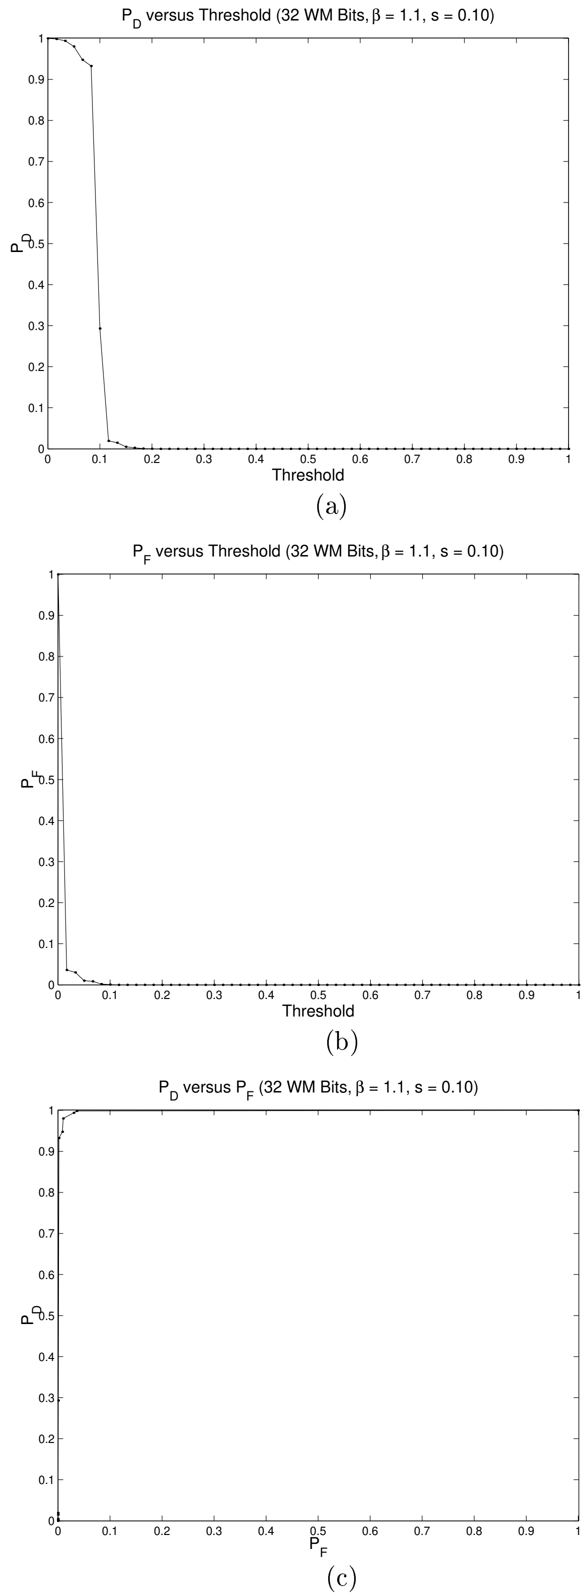

Figure 6.9 P_D and P_F curves for the NPE MLE for Lena data ($s^* = 0.10$).

the purpose of viewing, while under H_0, all of the mass is clustered near 0. A small portion of the H_1 mass is located near the origin, indicating that non-infinite decision statistics do indeed occur under this hypothesis. The probability of their occurrence, however, decreases as the embedding strength is increased. On the basis of the given plots, strong performance is expected from the detector because a large degree of separation is present between the distributions under the two hypotheses.

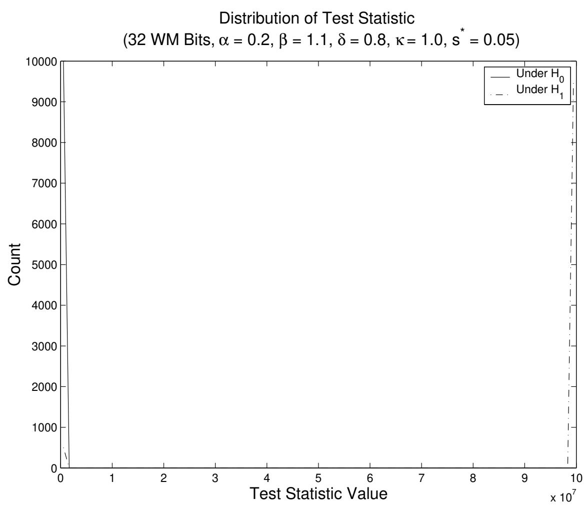

(a)

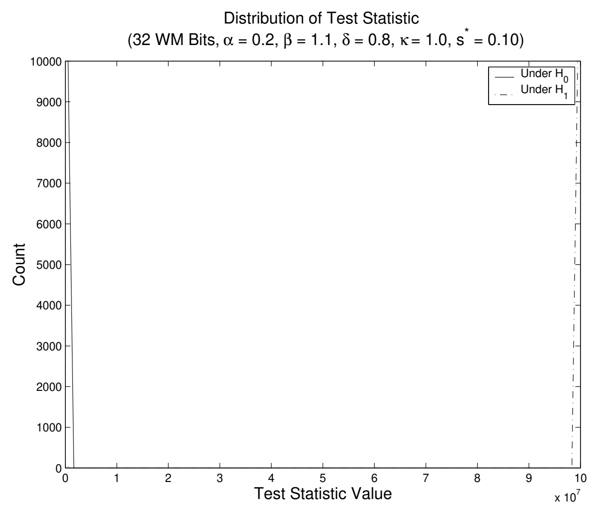

(b)

Figure 6.10 Distribution of the test statistic for the Lena image using the NPE distribution and LOD with an embedding strength of (a) 0.05 and (b) 0.10.

The detector is now examined in terms of the P_D and P_F curves generated from Monte Carlo simulations on the Lena data. Figures 6.11 and 6.12 each contain three plots of the detection and false alarm probabilities. Since the values taken on by the detection probability in the simulation are in a limited range, the axes are restricted to [0.9, 1]. Because of this increase in scale, the curves appear more jagged, which is an artifact of the finite number of simulations performed and the use of a different random message for each trial. When the two groups of plots are compared, it is clear that an improvement in performance results from increasing the embedding strength. For the case when $s^* = 0.05$, the detector performs moderately, with a detection probability of approximately 98% yielding a false alarm probability just over 50%. These statistics are significantly worse than those for the corresponding likelihood ratio test detector; however, they represent a large gain

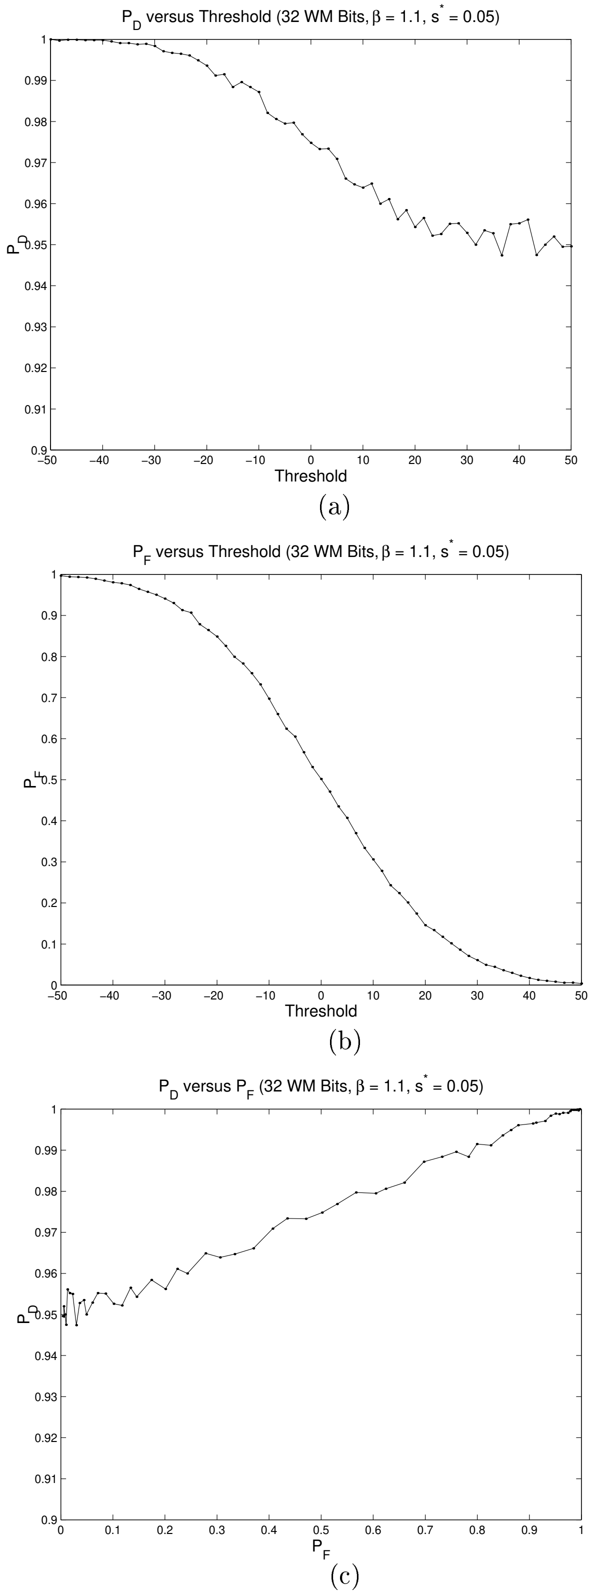

Figure 6.11 P_D and P_F curves for the NPE LOD for Lena data ($s^* = 0.05$).

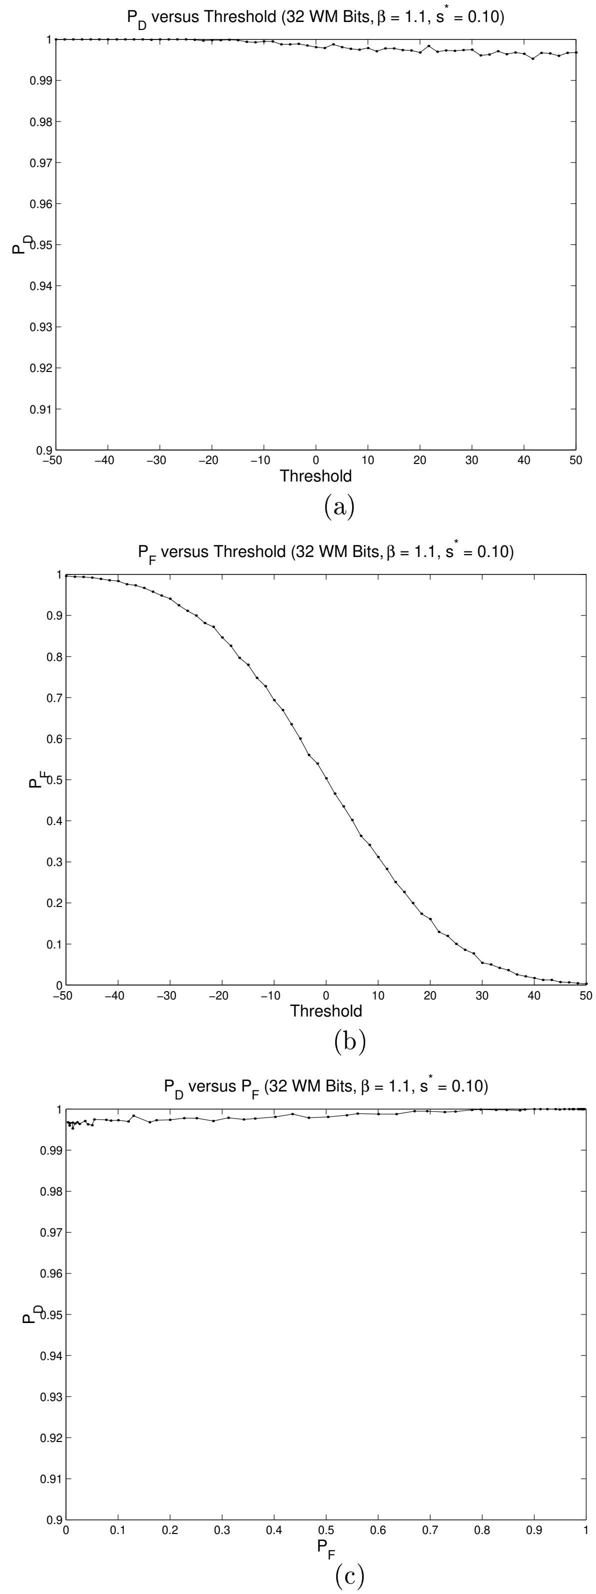

Figure 6.12 P_D and P_F curves for the NPE LOD for Lena data ($s^* = 0.10$).

over the locally optimal detectors based on the Weibull and power exponential distributions (Figures B.14, B.15, B.17, and B.18). This result supports the intuitive belief that the notched distribution should offer better performance because it more closely models the coefficients chosen for watermarking. In a mathematical sense, the improvement is caused by the ability to eliminate a hypothesis when a data element is found below one of the notch thresholds. However, given the strong performance of the LRT detector, the locally optimal detector based on the notched power exponential distribution does not yield results that would support its use in a real-world application.

It is interesting to compare the data resulting from using the NPE locally optimal detector on coefficients from the Lena image with those gathered when synthetic data is utilized (Figures 5.20 and 5.21). As observed for both the likelihood ratio test and maximum likelihood estimation detectors, a decrease in performance is present for the Lena image data. This result indicates that the DCT coefficients of the Lena image do not follow the notched power exponential distribution exactly.

6.3 Watermarking Peppers DCT Coefficients

In order to further test the developed detectors on actual image data, the experiments performed for the Lena image are repeated using the Peppers image given in Figure 6.13. The notched power exponential distribution detectors are considered in this chapter, while the Weibull and power exponential distribution analyzes are presented in Appendix B. Histograms of the DCT coefficients of this image and the 32 selected for watermarking are shown in Figure 6.14.

The values of the distribution parameters for all three modeling distributions are estimated by using the moments of the sample data and assuming, as before, a value of $\beta = 1.1$. The results are shown in Table 6.3.

Monte Carlo simulations are now performed for each detector to evaluate the performance in terms of P_D and P_F. Each simulation considers a range of threshold values, with 40 000 runs for each threshold, where a new message is created for each run.

Figure 6.13 The Peppers image.

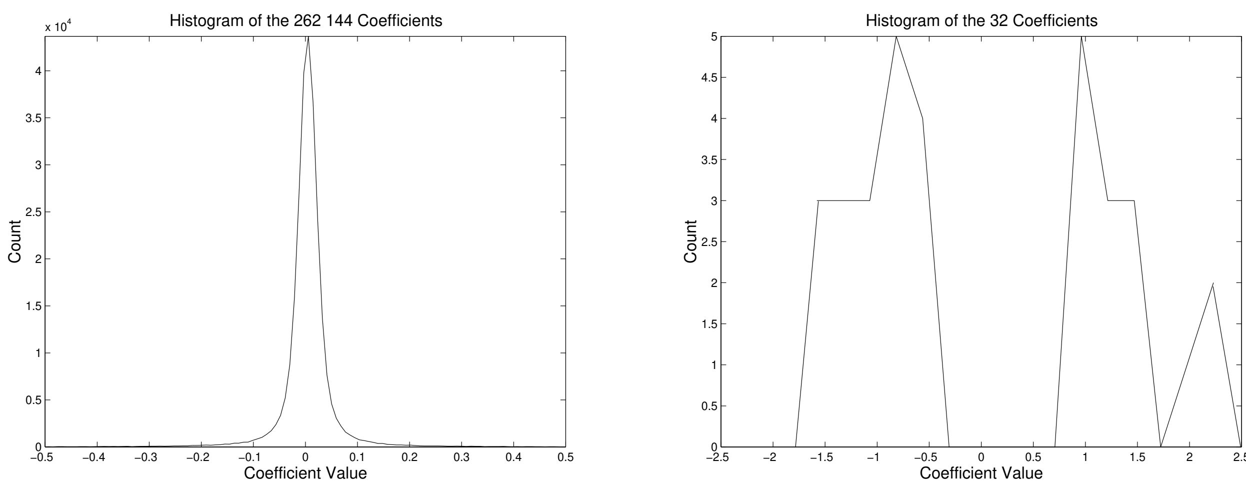

(a) (b)

Figure 6.14 (a) DCT coefficients of the Peppers image and (b) those coefficients selected for watermarking.

Table 6.3 Estimated distribution parameters for the Peppers DCT coefficients ($\beta = 1.1$).

Distribution	Parameters		
	α	δ	κ
Weibull	0.4593	-	-
Power exponential	0.1734	-	-
Notched power exponential	0.1734	0.7758	0.9956

6.3.1 Simple likelihood ratio test

The first detector examined for the Peppers image is that based on likelihood ratio testing with the notched power exponential distribution. The results are again presented in the forms of the distributions of the test statistic, and the P_D and P_F curves from the simulations.

First, Figure 6.15 presents the distributions of the test statistics under each hypothesis. Once again the two distributions do indeed contain most of their mass at infinite values (represented here by $\pm 1 \cdot 10^8$), with only small amounts of mass near the origin. These small peaks are reduced to near invisibility on the depicted scale for the higher embedding strength.

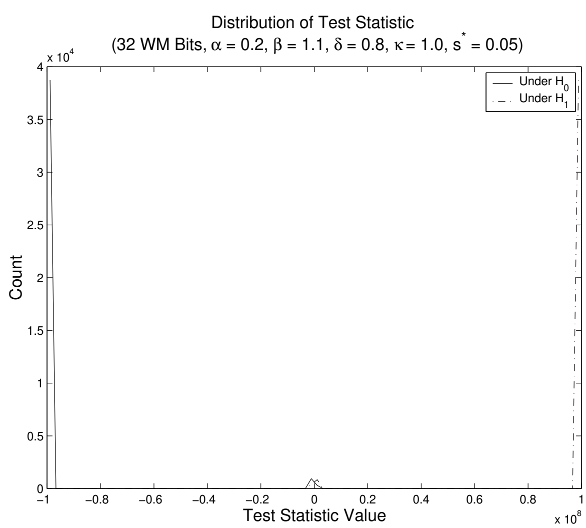

(a)

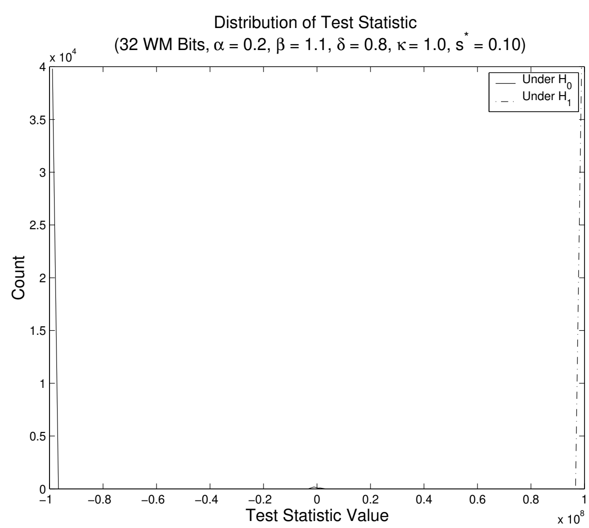

(b)

Figure 6.15 Distribution of the test statistic for the Peppers image using the NPE distribution and LRT with an embedding strength of (a) 0.05 and (b) 0.10.

The detection and false alarm probability results for the Monte Carlo simulations on the Peppers data are given in Figures 6.16 and 6.17. Note that the plots are shown on reduced scales because a limited range of probabilities occurs during the simulations. Thus, the curves still remain somewhat jagged because of the randomness introduced by the different message in each run and the finite number of runs performed. When the graphs in the two figures are compared, a significant increase in performance is noticed for the higher embedding strength. However, even when $s^* = 0.05$, a detection probability of approximately 99.25% is achieved for a false alarm probability of approximately 1%. These rates increase to an astonishing 99.99% detection with a $6 \cdot 10^{-6}$% probability of false alarm when $s^* = 0.10$. Thus, it is clear that the use of the likelihood ratio test detector with the notched power exponential distribution offers great promise for use in real-world applications.

As in the case of the Lena image, a drop in performance is observed when the detector is applied to the Peppers image coefficients as opposed to synthetic data (Figures 5.13 and 5.14). However, the size of this decrease is very similar to that for the Lena image, suggesting that the performance is not highly dependent on the particular real-world image that is watermarked.

6.3.2 Simple detector for notched distributions

The next detector examined for the Peppers image is the simple detector using the notched power exponential distribution. Table 6.4 contains statistics on the output of the detector under each of the two hypotheses using the derived decision rule (3.4) with $\lambda = \frac{1}{2}$. The values contained are similar to those for the Lena image, with slightly more certain decisions. Thus, slightly stronger performance than that for the Lena image is expected.

The simple detector is now compared against the likelihood ratio test detector. Figure 6.18 shows the ROC curve for the LRT detector with that for the simple detector overlaid. Once again the axes have been reduced for each subfigure to highlight the small range of probabilities. For both embedding strengths, the ROC curves for the SD lie to the lower right of those for the LRT detector, demonstrating that a slight drop in performance results from simplifying the detector. However, the derived detector still performs well when applied to the Peppers image.

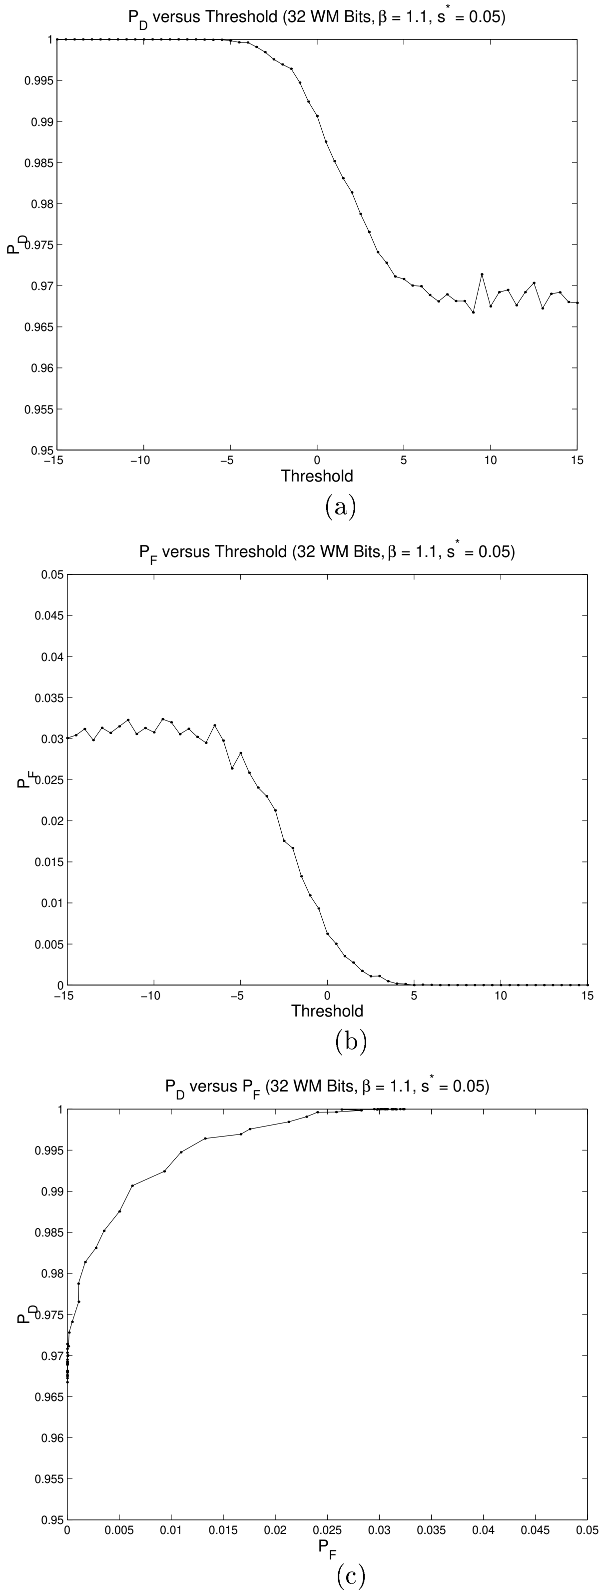

Figure 6.16 P_D and P_F curves for the NPE LRT for Peppers data ($s^* = 0.05$).

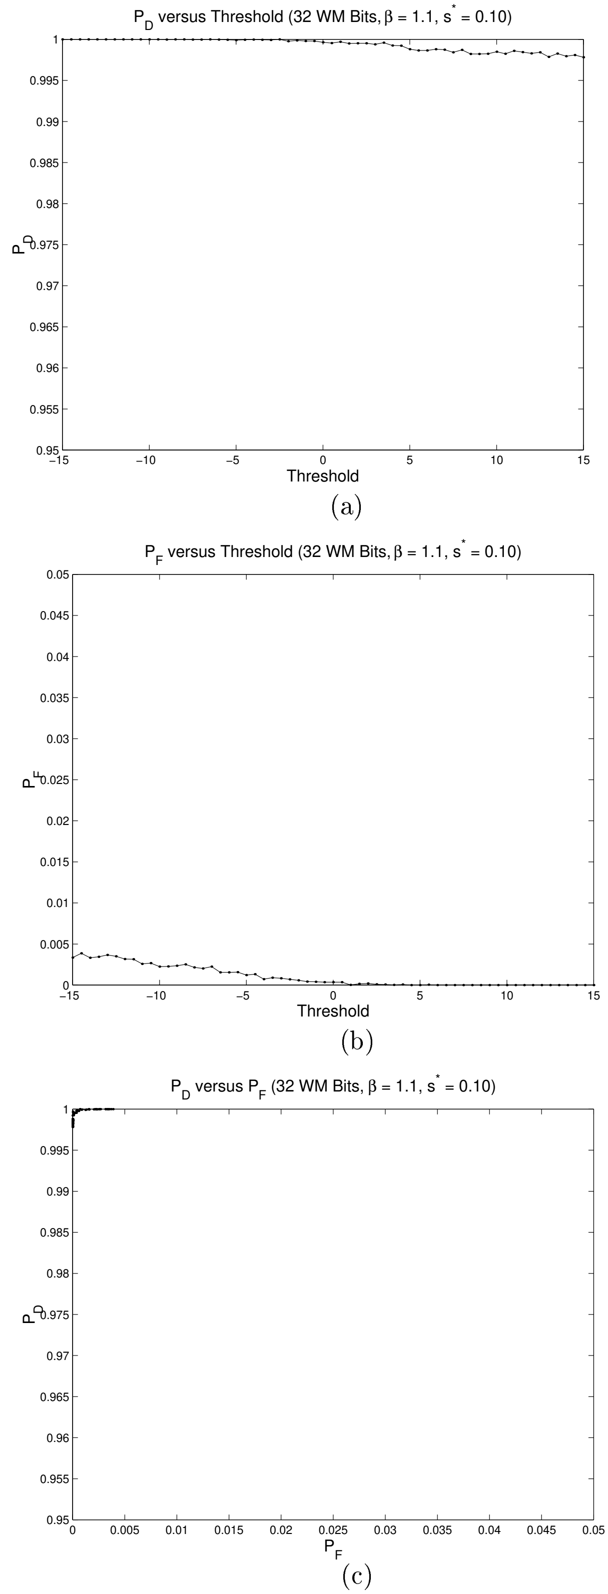

Figure 6.17 P_D and P_F curves for the NPE LRT for Peppers data ($s^* = 0.10$).

Table 6.4 Simple detector output for data from the Peppers DCT coefficients.

Situation	Count	
	$s^* = 0.05$	$s^* = 0.10$
Certain H_0 under H_0	38 762	39 846
Certain H_1 under H_0	0	0
Random selection under H_0	1238	154
Certain H_0 under H_1	0	0
Certain H_1 under H_1	38 736	39 906
Random selection under H_1	1264	94

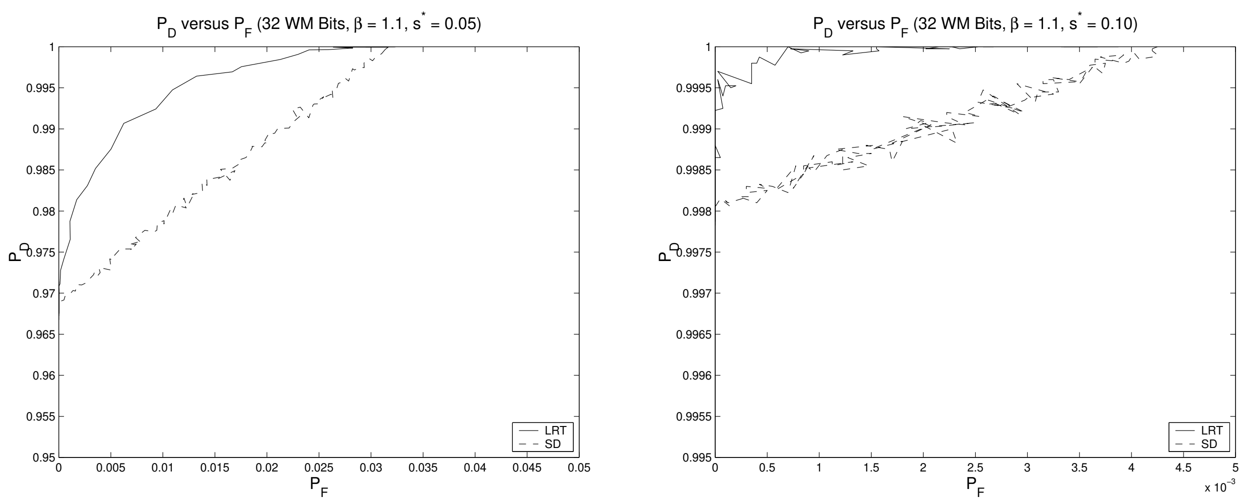

(a) (b)

Figure 6.18 ROC curves for the LRT and SD for the NPE distribution with Peppers data using an embedding strength of (a) 0.05 and (b) 0.10.

6.3.3 Maximum likelihood estimation

The detector based upon maximum likelihood estimation and the notched power exponential distribution is evaluated on its ability to detect a watermark embedded into the DCT coefficients of the Peppers image. Once again, three figures are presented to demonstrate the results.

Figure 6.19 contains plots of the distributions of the estimates for two different embedding strengths, found through 10 000 Monte Carlo simulations. As seen for the Lena image, there is a clear separation between the distributions under H_0 and H_1, with H_1 exhibiting strong peaks around the true values of the embedding strength, particularly for the higher embedding strength.

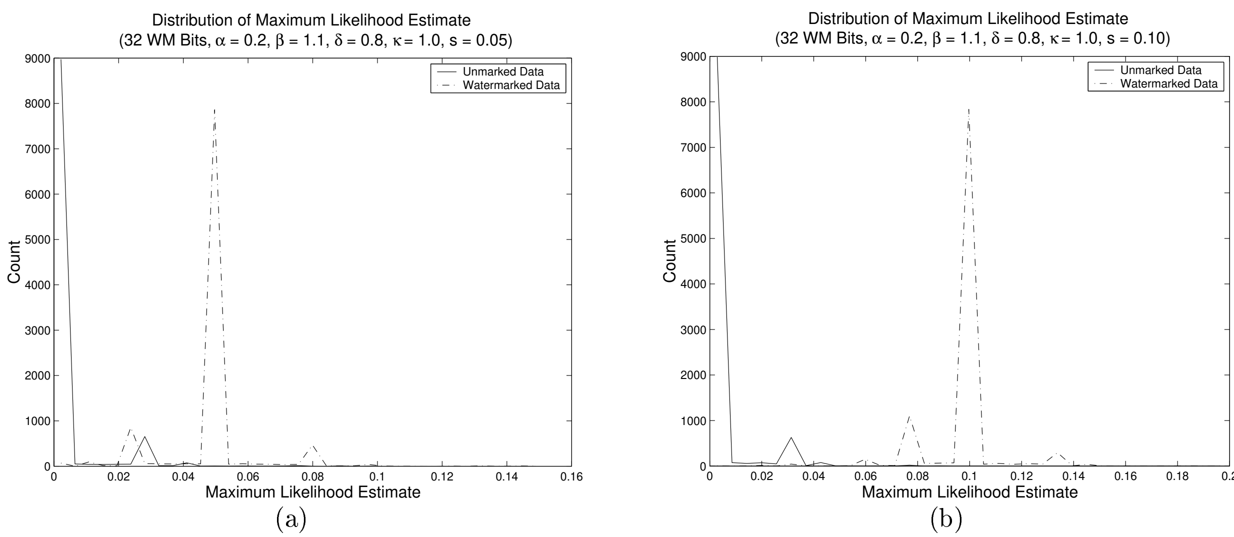

Figure 6.19 Distribution of the test statistic for the Peppers image using the NPE distribution and MLE with an embedding strength of (a) 0.05 and (b) 0.10.

In order to further study the detector, curves for the performance indicators P_D and P_F are given in Figures 6.20 and 6.21. The typical expected result for MLE is found, where the detection probability falls off dramatically as the threshold is increased above the actual embedding strength. Through a comparison of the two figures, it is evident that increasing the embedding strength contributes to the detector's performance, as expected. For the lower embedding strength, a detection probability of approximately 98% is achieved for a false alarm probability of 50%. Although the use of the notched power exponential distribution

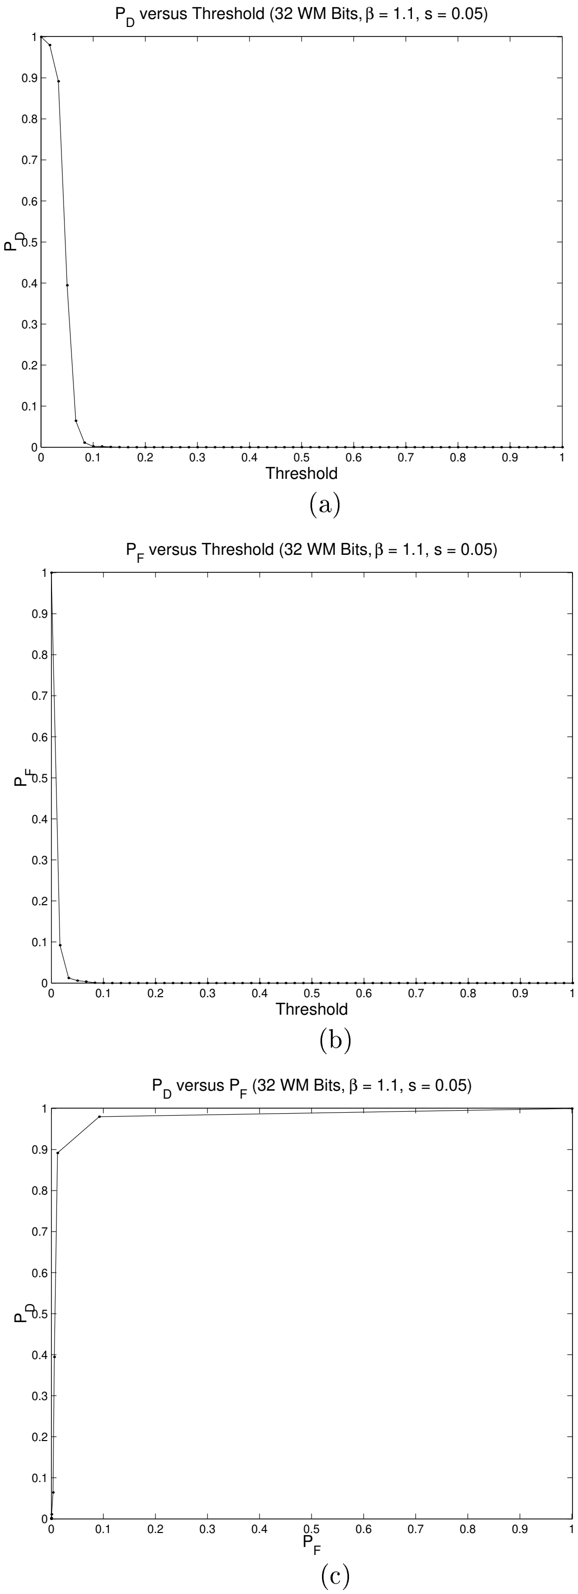

Figure 6.20 P_D and P_F curves for the NPE MLE for Peppers data ($s^* = 0.05$).

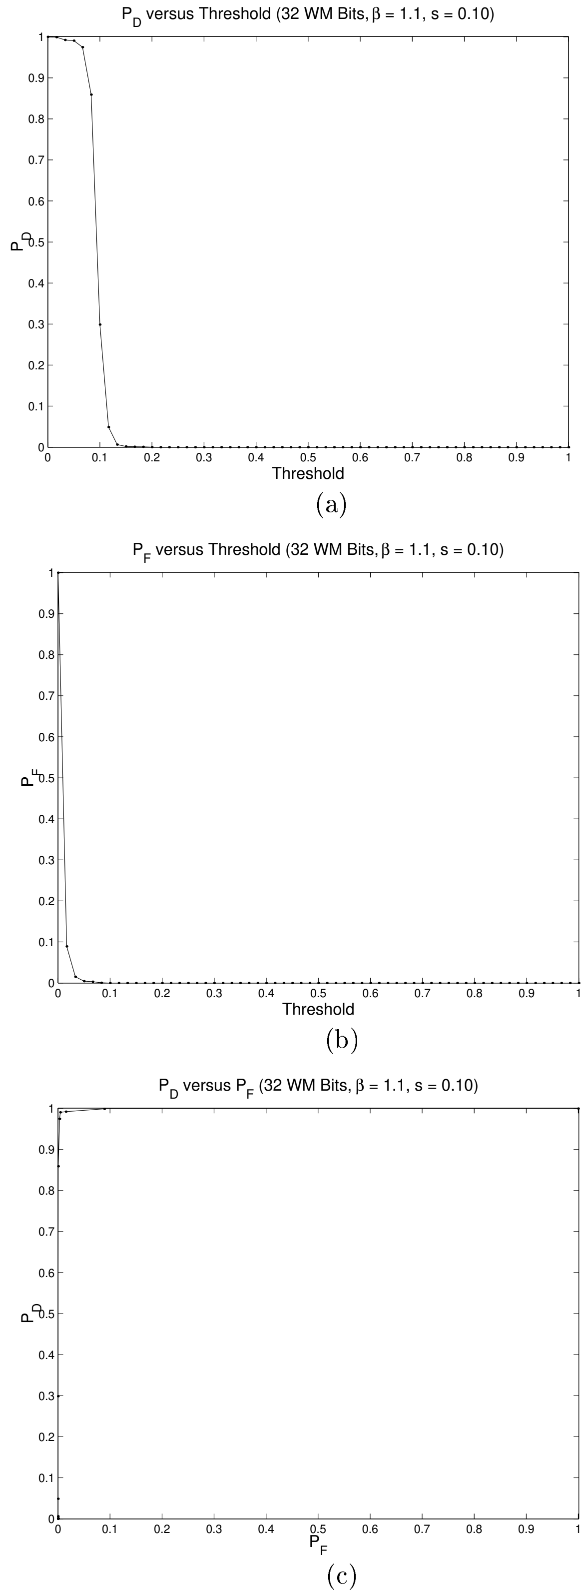

Figure 6.21 P_D and P_F curves for the NPE MLE for Peppers data ($s^* = 0.10$).

improves the maximum likelihood estimation detector substantially, the resulting detector does not perform as well as the corresponding likelihood ratio test detector. Hence, it is doubtful that this detector would be employed in real-world situations.

By comparing these results to those for synthetic data, Figures 5.17 and 5.18, a drop in performance is observed. However, the magnitude of this decrease is comparable to that observed for the Lena image, again suggesting the watermarking detectors are not highly sensitive to the particular selection of the real-world image.

6.3.4 Locally optimal detection

To complete the study of the notched power exponential detectors when a watermark is inserted into the Peppers image, locally optimal detection is considered. As before, the experimental results are presented in the form of three figures.

The distributions of the test statistics under each hypothesis are shown in Figure 6.22. In these plots, the infinite value that is possible for the statistic is represented by $1 \cdot 10^8$. The graphs demonstrate that all of the mass of the distribution under H_0 is clustered near 0. On the other hand, under H_1, most of the mass is located at infinity, particularly for the high embedding strength, where the non-infinite lobe near 0 is all but invisible. Thus, a strong performance is expected from the detector.

P_D and P_F curves are now presented to complete the examination of the detector in question. Figures 6.23 and 6.24 contain the necessary graphs for embedding strengths of 0.05 and 0.10, respectively. The axes in the figures have been truncated such that P_D is shown between [0.95, 1]. Values outside this area do not occur in the Monte Carlo simulations performed. As a consequence of the increased scale, the curves appear rather jagged because a finite number of simulations is performed, each using a different random message. By comparing the curves for the two embedding strengths, a significant improvement is observed for the higher case. However, even for $s^* = 0.05$, a detection probability of approximately 98% can be achieved with a false alarm probability of 10%. These figures improve to approximately 99.95% and 0.44% for $s^* = 0.10$. Thus, the detector offers a commanding improvement over those for the Weibull and power exponential distributions given in the appendix, but is slightly worse than the corresponding likelihood ratio test detector. As a

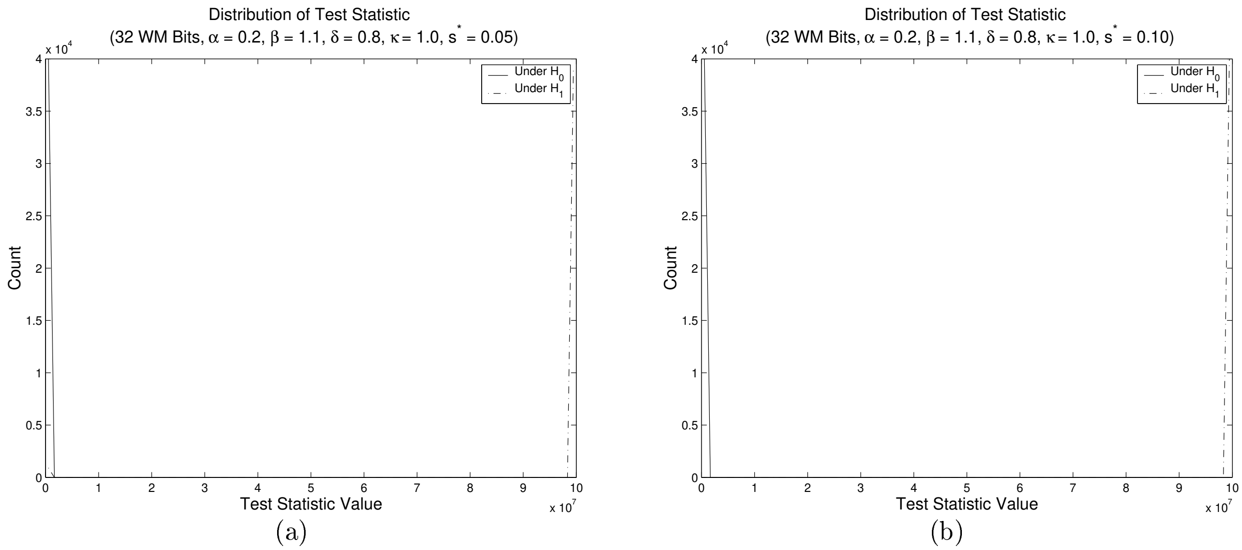

Figure 6.22 Distribution of the test statistic for the Peppers image using the NPE distribution and LOD with an embedding strength of (a) 0.05 and (b) 0.10.

result, its utilization in practical applications is unlikely.

It is also interesting to compare the results from the Pepper image data to those for synthetic data. As expected, the detector does not perform as well on the image data as it does on the synthetic coefficients. However, the degree of the drop is quite similar to that observed for the Lena image, which again supports the belief that the detectors are not significantly impacted by the real-world image selected for watermarking.

6.4 Summary of Results

This chapter examined the four types of watermarking detectors based on the notched power exponential distribution discussed in this thesis for use on real-world test images. The detection and false alarm probability results from these experiments, and those in the appendix for the Weibull and power exponential distributions, are condensed into Tables 6.5 and 6.6 for the Lena and Peppers images, respectively. The probabilities are shown in percentage form and are rounded to the nearest whole number. As with the summary for synthetic data, the data points are selected to demonstrate the false alarm probability that must be incurred in order to achieve a detection probability of approximately 98%. When this probability can not be achieved, the highest P_D value is chosen.

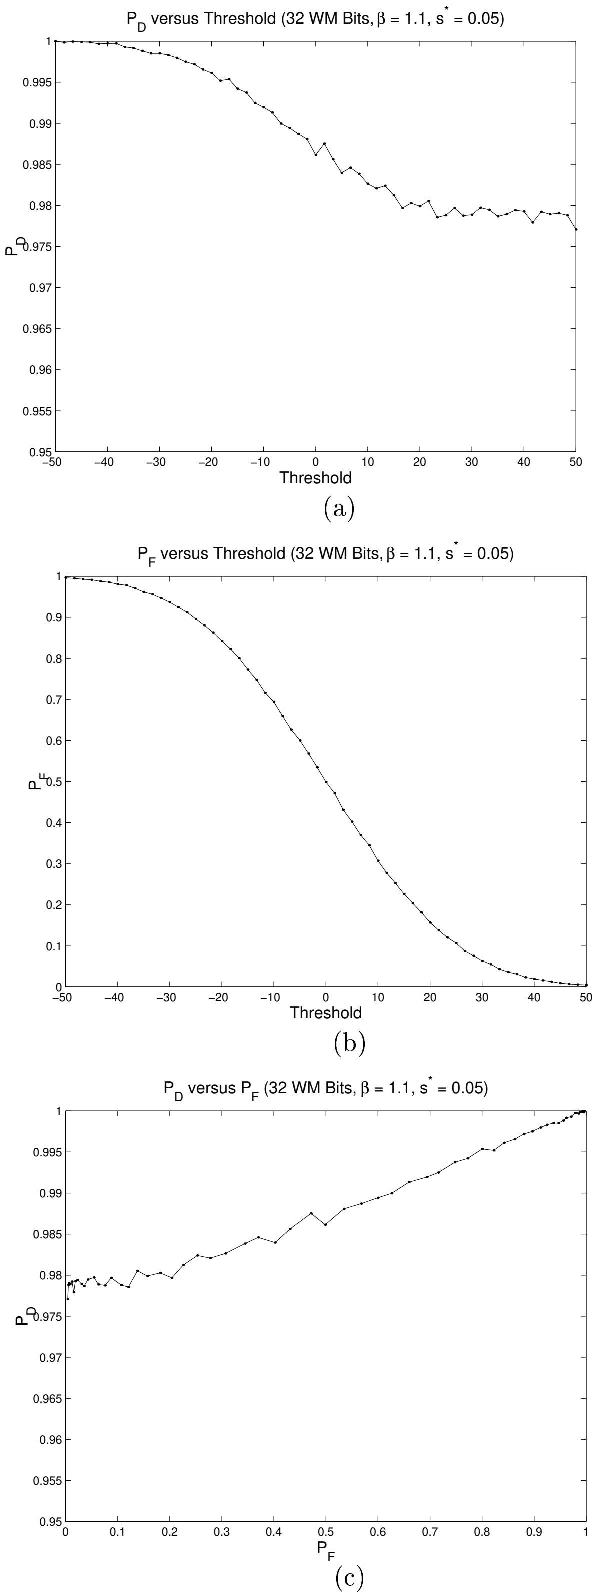

Figure 6.23 P_D and P_F curves for the NPE LOD for Peppers data ($s^* = 0.05$).

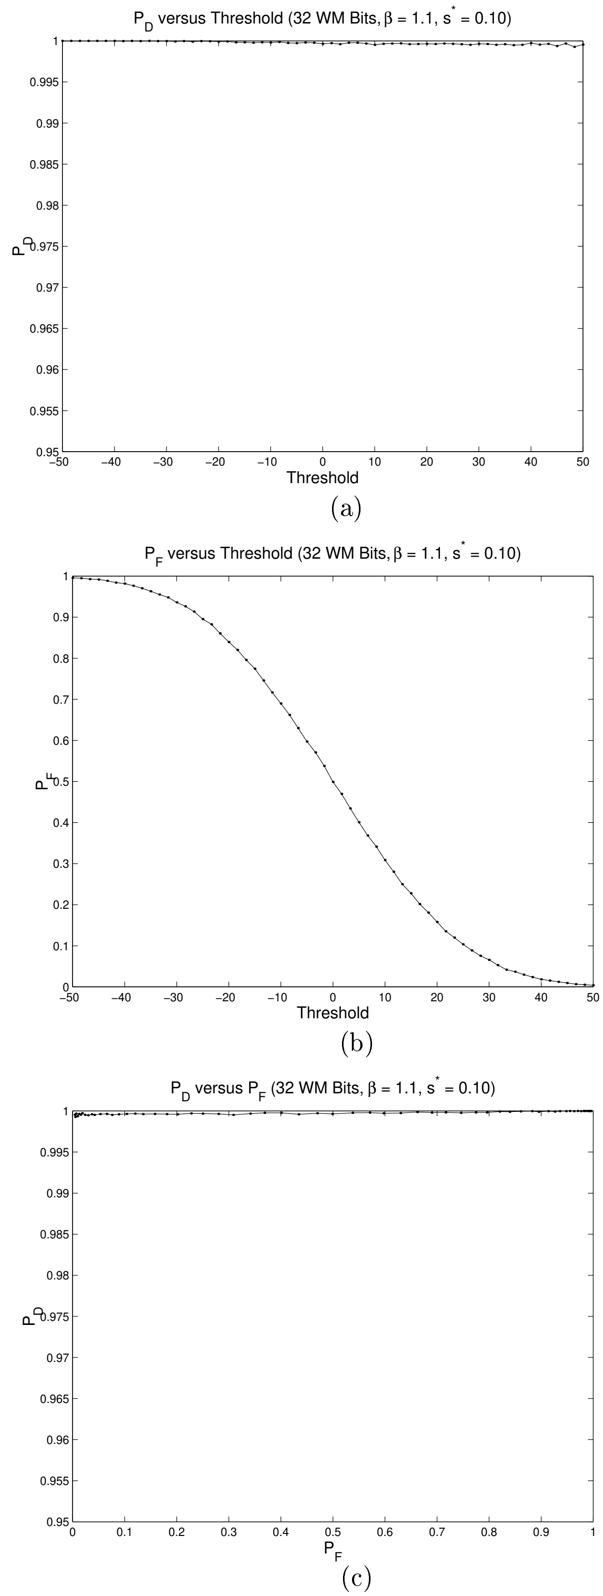

Figure 6.24 P_D and P_F curves for the NPE LOD for Peppers data ($s^* = 0.10$).

Table 6.5 Selected results for data from the Lena DCT coefficients.

Distribution	Detector	$s^* = 0.05$		$s^* = 0.10$	
		$P_D(\%)$	$P_F(\%)$	$P_D(\%)$	$P_F(\%)$
Weibull	LRT	98	95	98	88
Weibull	MLE	62	51	72	50
Weibull	LOD	98	88	98	64
Power exponential	LRT	98	95	98	91
Power exponential	MLE	62	51	72	50
Power exponential	LOD	97	84	98	63
Notched power exponential	LRT	98	2	99	0
Notched power exponential	MLE	96	4	98	1
Notched power exponential	LOD	98	61	100	1
Notched power exponential	SD	97	3	100	1

Table 6.6 Selected results for data from the Peppers DCT coefficients.

Distribution	Detector	$s^* = 0.05$		$s^* = 0.10$	
		$P_D(\%)$	$P_F(\%)$	$P_D(\%)$	$P_F(\%)$
Weibull	LRT	98	95	98	87
Weibull	MLE	61	49	71	50
Weibull	LOD	98	88	98	63
Power exponential	LRT	98	95	98	91
Power exponential	MLE	61	50	72	50
Power exponential	LOD	98	88	98	62
Notched power exponential	LRT	98	0	100	0
Notched power exponential	MLE	98	1	98	0
Notched power exponential	LOD	98	1	100	0
Notched power exponential	SD	99	2	100	0

The conclusions that can be drawn from these tables in terms of the best performing detector and distribution are the same as those stated for the synthetic data in Section 5.5. Namely, the likelihood ratio test detector based on the notched power exponential distribution is clearly superior to the other detector-distribution pairs considered.

Comparisons can also be made between the experimental results gathered from the test images and those found for the synthetic data, Table 5.3. In general, a decrease in performance is observed when the detectors are applied to actual image DCT coefficients. This result is expected because, here, the distributions are simply being used to model the distribution of the transform coefficients, which are not necessarily distributed according to the studied distributions. However, the performance drop is observed for both the Lena and Peppers test images, and is of approximately the same amount: generally a 1% - 2% increase in P_F for the same P_D for the notched power exponential distribution. Thus, the detectors appear to be fairly robust to the image selected for watermarking.

CHAPTER 7

IMAGE WATERMARKING EXPERIMENTS WITH ATTACKS

The previous chapter examined the performance of the developed detectors for use in detecting a watermark embedded into a real-world test image. The results demonstrated that the likelihood ratio test detector coupled with the notched power exponential distribution yields the most promising means of watermark detection. This chapter again considers the detection of a watermark in a natural image; however, an attack is now included in the system. Because of its strong performance, only the LRT detector is evaluated, and the notched power exponential distribution is compared with the power exponential distribution. It is important to note that the development of these detectors did not include a model of an attack; hence, they are no longer optimal. However, the behavior of the detectors in the presence of an attack remains worthy of study. This chapter discusses the types of attacks introduced and then presents the results when they are applied to a watermarked version of the Lena test image.

7.1 Attack Types

Two attack methods are considered for application on an image in an attempt to reduce a detector's ability to accurately determine the presence of a watermark. These techniques include *additive white Gaussian noise* (AWGN) and *Joint Photographic Experts Group* (JPEG) compression. This section describes the implementation of each attack, as well as the resulting amount of distortion introduced. In both cases, the attacker does not know which transform coefficients are watermarked; hence, the attacks are applied to the entire image. In the formulation of the distortion, the mean squared error measure is used to determine the difference between the watermarked image and the attacked watermark image.

7.1.1 Additive white Gaussian noise

The first type of attack considered is the addition of white Gaussian noise to the watermarked image. A matrix, having the same size as the image, of noise is generated randomly according to a zero-mean Gaussian distribution, $\mathbf{Z} \sim \mathcal{N}(0, \sigma^2)$, with variance σ^2. The matrix is then added to the watermarked image, or, equivalently, the transform coefficients of the watermarked image, to produce an attacked image (or attacked coefficients).

In order to derive the amount of distortion introduced by this attack, some additional notation must first be defined:

Let $\tilde{\mathbf{b}}$ = the pixels of the attacked watermarked image.

Let $\tilde{\mathbf{B}}$ = the transform coefficients of the attacked watermarked image = $\mathbf{B} + \mathbf{Z}$.

Then, the distortion (as defined in Equation (2.1)) introduced by the AWGN is given by

$$\begin{aligned} D_2 = D_2(\mathbf{b}, \tilde{\mathbf{b}}) &= \frac{1}{M}\sum_{k=1}^{M} |b_k - \tilde{b}_k|^2 \\ &= \frac{1}{M}\sum_{i=1}^{M} |B_i - \tilde{B}_i|^2 \\ &= \frac{1}{M}\sum_{i=1}^{M} |B_i - B_i - Z_i|^2 \\ &= \frac{1}{M}\sum_{i=1}^{M} |Z_i|^2 \end{aligned}$$

where Parseval's theorem is employed to change the equation to the transform domain. The expected distortion is given by

$$\begin{aligned} E[D_2] &= E\left[\frac{1}{M}\sum_{i=1}^{M} |Z_i|^2\right] \\ &= \frac{1}{M}\sum_{i=1}^{M} \sigma^2 \\ &= \sigma^2. \end{aligned}$$

Thus, the expected strength of the AWGN attack is controlled directly by altering the variance of the noise introduced.

7.1.2 JPEG compression

The next type of attack considered is that of JPEG compression. JPEG compression is a means of reducing the number of bytes required to represent an image by removing some of the information contained in the image's transform. More specifically, the pixels of an image are considered in 8×8 blocks. The DCT of each block is taken, and the transform coefficients are quantized to reduce the number of bits required to represent the data values. The image is then further compressed without loss of information using entropy encoding. The degree of quantization employed is determined through the use of a quality factor in the range [0, 100], where a higher number produces a better-quality image.

The distortion introduced by the JPEG compression is given by

$$\begin{aligned} D_2 = D_2(\mathbf{b}, \tilde{\mathbf{b}}) &= \frac{1}{M}\sum_{k=1}^{M} |b_k - \tilde{b}_k|^2 \\ &= \frac{1}{M}\sum_{i=1}^{M} |B_i - \tilde{B}_i|^2 \\ &= \frac{1}{M}\sum_{i=1}^{M} |B_i - Q_i(B_i)|^2 \\ &\leq \frac{1}{M}\sum_{i=1}^{M} \left|\frac{q_i}{2}\right|^2 \end{aligned}$$

where Q_i is the quantizer applied to coefficient i (as defined by the JPEG standard using the specified quality factor), and q_i is the step size corresponding to the applied quantizer. Because of the complexity of this expression, it is difficult to directly relate the quality factor to the corresponding distortion. Hence, to create a suitable attack, a quality value is selected and the resulting distortion is computed. Through experimentation, the quality can be chosen such that the desired amount of distortion is introduced.

7.2 Simulation Results

To evaluate the performance of the likelihood ratio test detectors based on the notched power exponential and power exponential distributions, Monte Carlo simulations are run for each of the noise types previously discussed. The number of simulations performed is 40 000 for the case of AWGN and 10 000 for JPEG compression. The latter reduction is due to the large amount of computation required in implementing the attack. To illustrate detector performance, plots of the test statistic distributions, P_D vs. Threshold, P_F vs. Threshold, and P_D vs. P_F are included, with curves for the NPE and the PE displayed concurrently. Thus, the higher performance expected from the notched distribution should be clearly visible.

7.2.1 Lena image with additive white Gaussian noise

To examine the LRT detector in the presence of an AWGN attack, the Lena image is watermarked multiplicatively, as before. The distortion introduced through watermarking has a value of $D_1 = 0.000\ 193$. Two sets of Monte Carlo simulations are performed for each distribution, the first where $D_2 = D_1$ and the second where $D_2 = 10D_1$. The appropriate vectors of white Gaussian noise are created and added to the image to produce the attacked images shown in Figure 7.1. Because of the small amount of distortion introduced in the watermarking process, the attacker is limited to a correspondingly little amount of noise when $D_2 = D_1$. This restriction renders the distortion just barely visible in the more smooth regions of the image. However, for the larger noise magnitude, the decreased quality is quite visible throughout the entire image.

The distributions of the test statistic for the power exponential distribution, as found using 40 000 Monte Carlo simulations, are given in Figure 7.2. Little change in the distribution is present as the attack strength is increased. However, these distributions are slightly less separated than those for the Lena image when no attack is present, as seen in Figure B.4. Thus, it appears that the performance of the LRT detector based on the power exponential distribution is not significantly hindered by the AWGN attack.

(a) (b)

Figure 7.1 The Lena image with AWGN with (a) $D_2 = D_1$ and (b) $D_2 = 10D_1$.

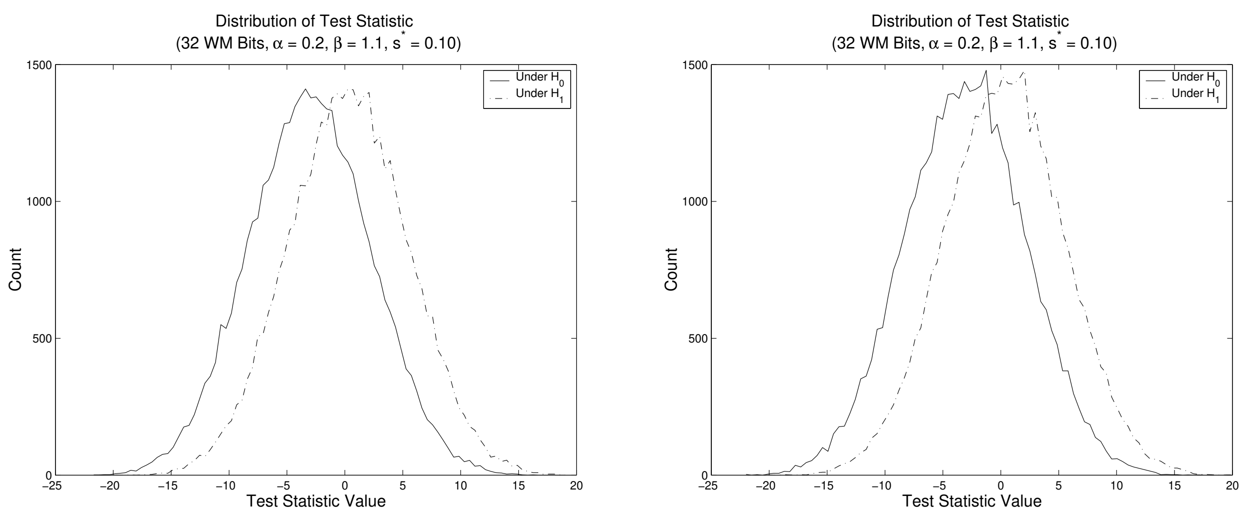

(a) (b)

Figure 7.2 Distribution of the test statistic using the PE distribution and LRT under an AWGN attack with (a) $D_2 = D_1$ and (b) $D_2 = 10D_1$.

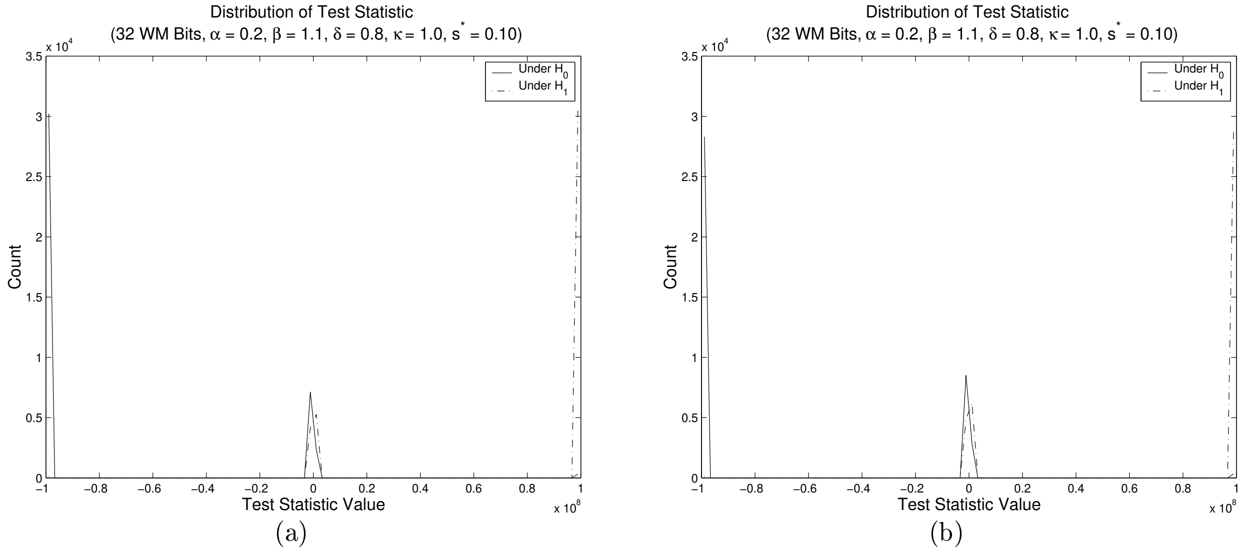

Figure 7.3 Distribution of the test statistic using the NPE distribution and LRT under an AWGN attack with (a) $D_2 = D_1$ and (b) $D_2 = 10D_1$.

Using another 40 000 Monte Carlo simulations, the distributions of the test statistic for the notched power exponential are found for each distortion case and are depicted in Figure 7.3. Although, in both cases, most of the distribution mass is located at the infinity points, the center lobes are much larger than those seen when no attack is present, Figure 6.3. This increase in area is more pronounced for the higher-variance attack. The movement of mass to the inner region is caused by the introduction of noise shifting more data points to lie within Regions 1 and 4, thereby demonstrating the sensitivity of the detector to points around the threshold. Also, the addition of noise causes a small number of false instances of Regions 2 and 3, as shown by the mass located at ∞ under H_0 and at $-\infty$ under H_1. Overall, these curves demonstrate that the LRT detector based on the notched power exponential distribution is moderately affected by the presence of noise, causing an expected decrease in performance.

To more clearly illustrate the performance of the detectors, curves of the detection and false alarm probabilities are considered. Figures 7.4 and 7.5 depict these probabilities for the two attacker distortion strengths, with both the power exponential and notched power exponential results shown. The curves for the power exponential are virtually identical to those for the unattacked case, given in Figure B.6. On the other hand, a substantial decrease

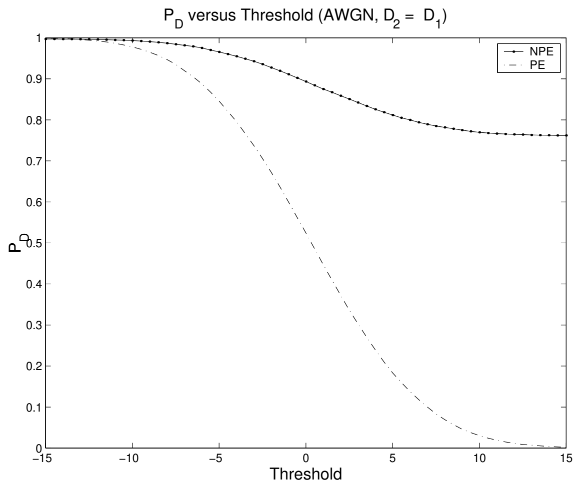

(a)

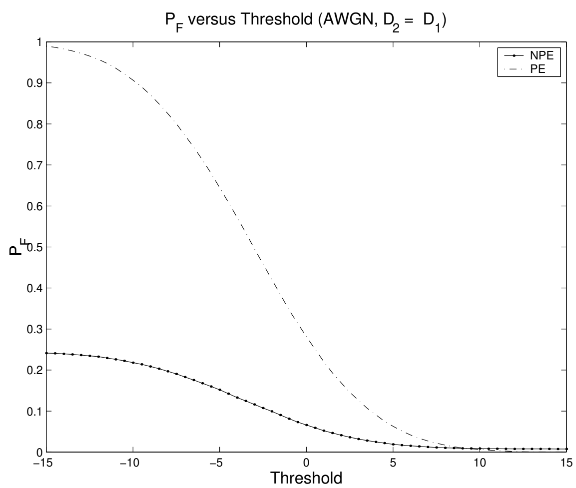

(b)

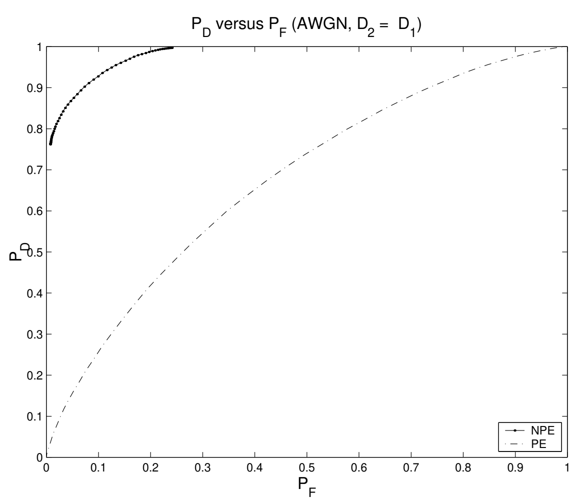

(c)

Figure 7.4 P_D and P_F curves for the NPE and PE LRT (AWGN, $D_2 = D_1$).

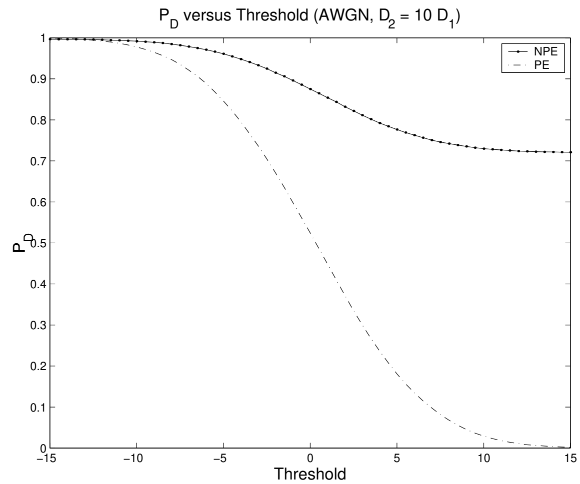

(a)

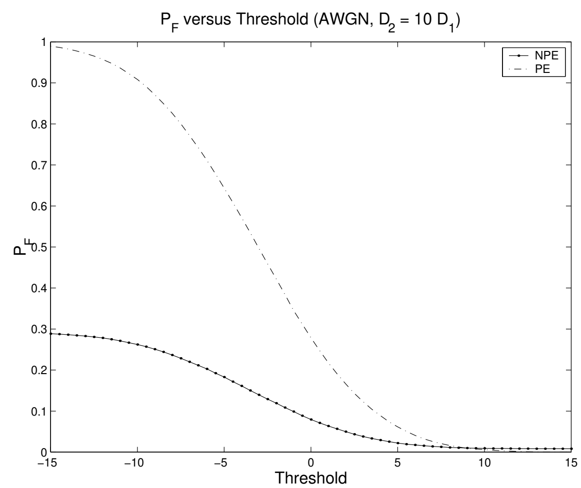

(b)

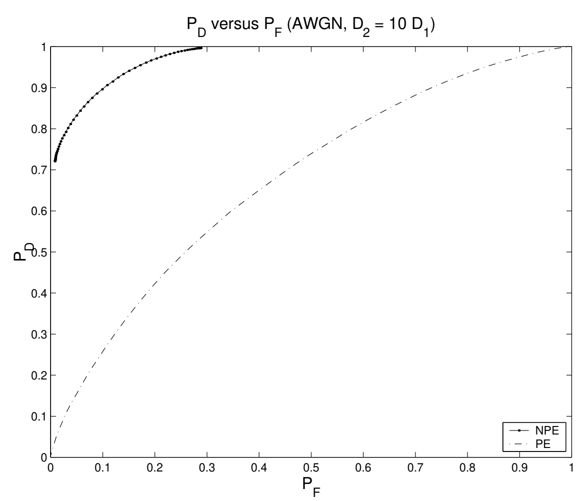

(c)

Figure 7.5 P_D and P_F curves for the NPE and PE LRT (AWGN, $D_2 = 10D_1$).

in performance is present for the NPE case when compared with the data resulting from no attack, Figure 6.5. This decrease is amplified in the higher-noise case, and is consistent with the change in the statistic distribution previously observed. However, the performance of the likelihood ratio test detector using the notched power exponential distribution remains notably superior to that resulting from the use of power exponential distribution.

7.2.2 Lena image with JPEG compression

The likelihood ratio test detector is now examined in the case where an attacker uses JPEG compression to manipulate the watermarked image. Once again, Monte Carlo simulations are performed for compression qualities resulting in mean squared errors of $D_2 = D_1$ and $D_2 = 10D_1$. The quality factors required to achieve these distortion levels are 81 and 6, respectively, and the two compressed images are presented in Figure 7.6. The small value of D_1 causes the visible compression artifacts present in the first image to be extremely minimal. However, when the attacker is permitted to introduce 10 times the amount of distortion as the embedding process, typical blocking artifacts caused by JPEG compression are quite prevalent.

Because of the large amount of computation required for JPEG compression, only 10 000 Monte Carlo simulations are performed to evaluate each detector. The resulting distributions of the test statistic under both compression strengths are given for the power exponential in Figure 7.7. Similar to the attack using Gaussian noise, these curves are nearly identical to those found when no attack is present, Figure B.4. Thus, little performance degradation is expected for the LRT power exponential detector.

Considering the distribution of the statistic when the notched power exponential is employed, Figure 7.8, somewhat different results from the AWGN attack are observed. For the lower attack strength, the statistic distribution appears quite similar to the unattacked case (Figure 6.3); most of the mass lies at the infinity points with only a small amount near the 0 point. However, the slight decrease in mass at the endpoints observed when $D_2 = D_1$ is significantly amplified when $D_2 = 10D_1$. Thus, the LRT NPE detector possesses a robustness to lower rates of JPEG compression, but suffers considerably when severe compression is applied.

(a) (b)

Figure 7.6 The Lena image with JPEG compression with (a) $D_2 = D_1$ and (b) $D_2 = 10D_1$.

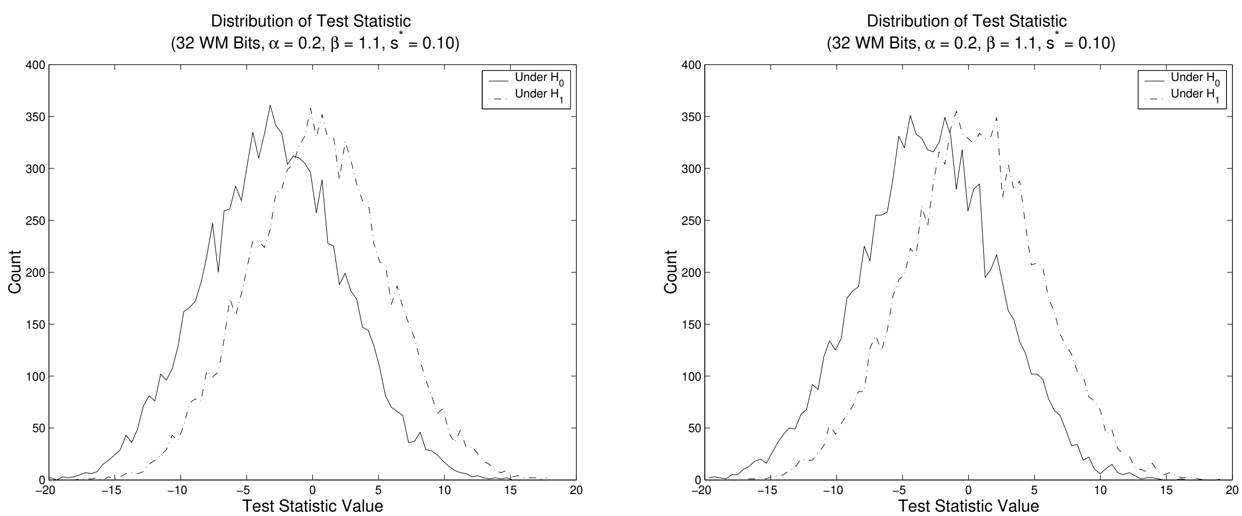

(a) (b)

Figure 7.7 Distribution of the test statistic using the PE distribution and LRT under a JPEG compression attack with (a) $D_2 = D_1$ and (b) $D_2 = 10D_1$.

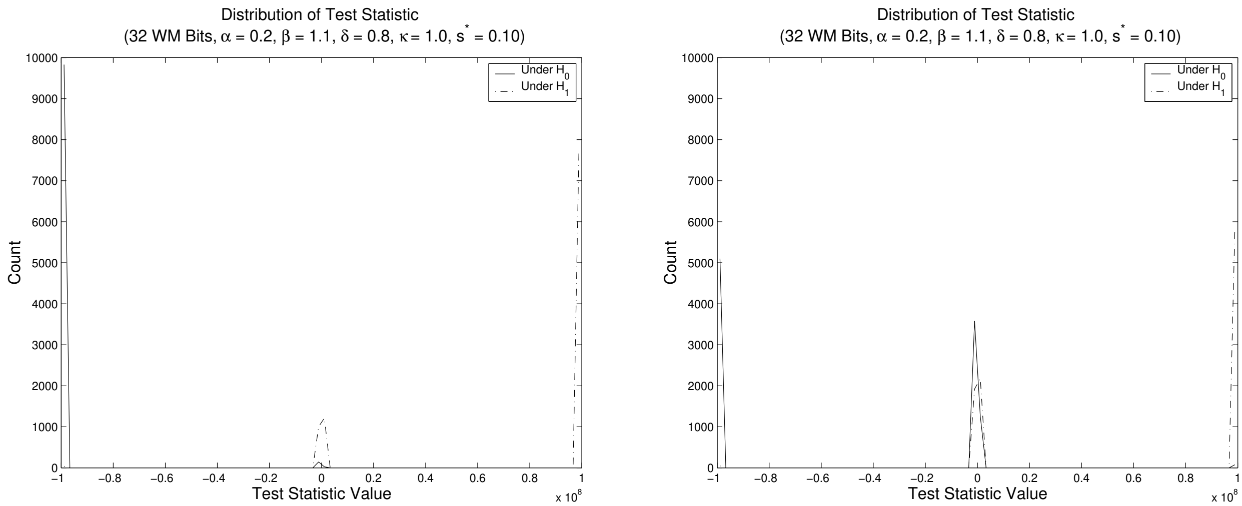

(a) (b)

Figure 7.8 Distribution of the test statistic using the NPE distribution and LRT under a JPEG compression attack with (a) $D_2 = D_1$ and (b) $D_2 = 10D_1$.

The performance of the likelihood ratio test detectors is now gauged using the indicators of detection and false alarm probabilities. Figures 7.9 and 7.10 present the resulting curves for both the PE and NPE for the two attack strengths. As expected, the curves for the power exponential are nearly identical to those for the no-attack case. Conversely, those for the notched power exponential display a decrease in performance. This drop is only slight for the low-attack case, but is more pronounced for the larger attack. However, these figures again demonstrate the clear superiority of the notched power exponential distribution over the power exponential distribution for modeling the selected coefficients.

7.3 Summary of Results

This chapter has presented the results of experiments employed to evaluate the performance of the likelihood ratio test detector for the Lena test image when an attack is present. Tables 7.1 and 7.2 contain selected data points from the P_D and P_F curves generated through Monte Carlo simulations for AWGN and JPEG attacks, respectively. Points have been selected to illustrate the behavior of the detectors for a detection probability of approximately 98%. The probabilities are shown in percentage form and are rounded to the nearest whole number.

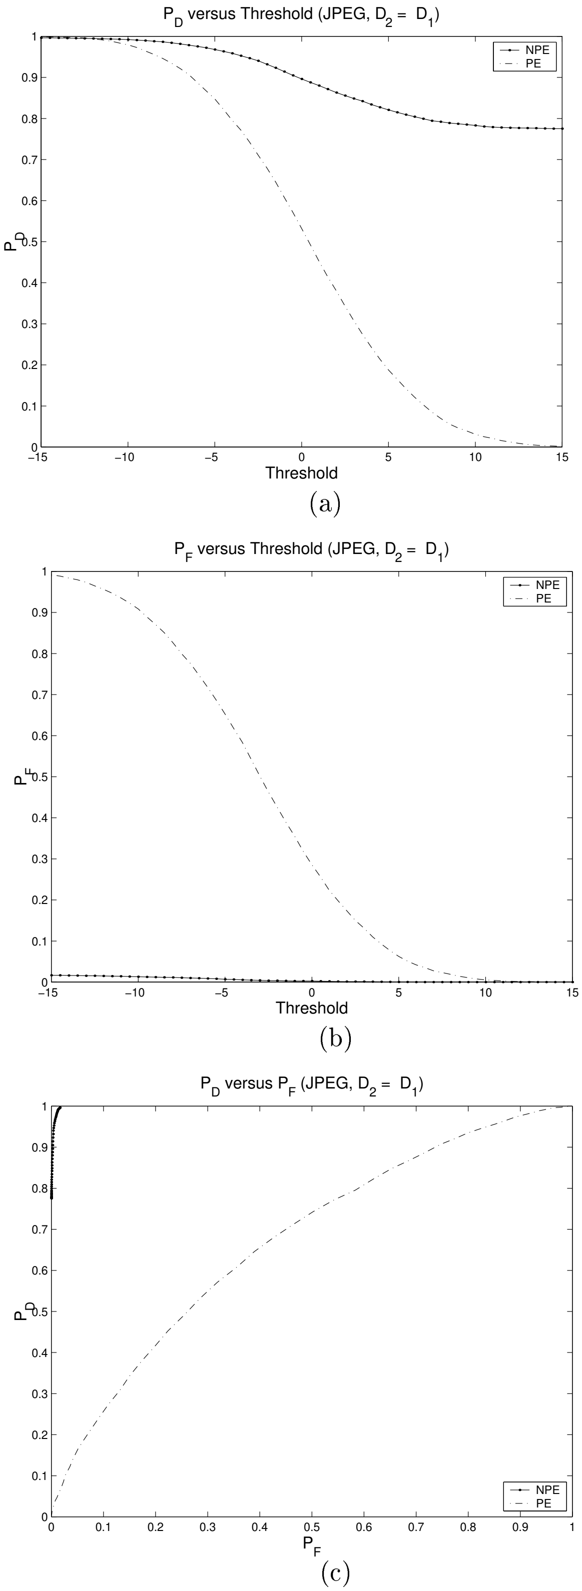

Figure 7.9 P_D and P_F curves for the NPE and PE LRT (JPEG, $D_2 = D_1$).

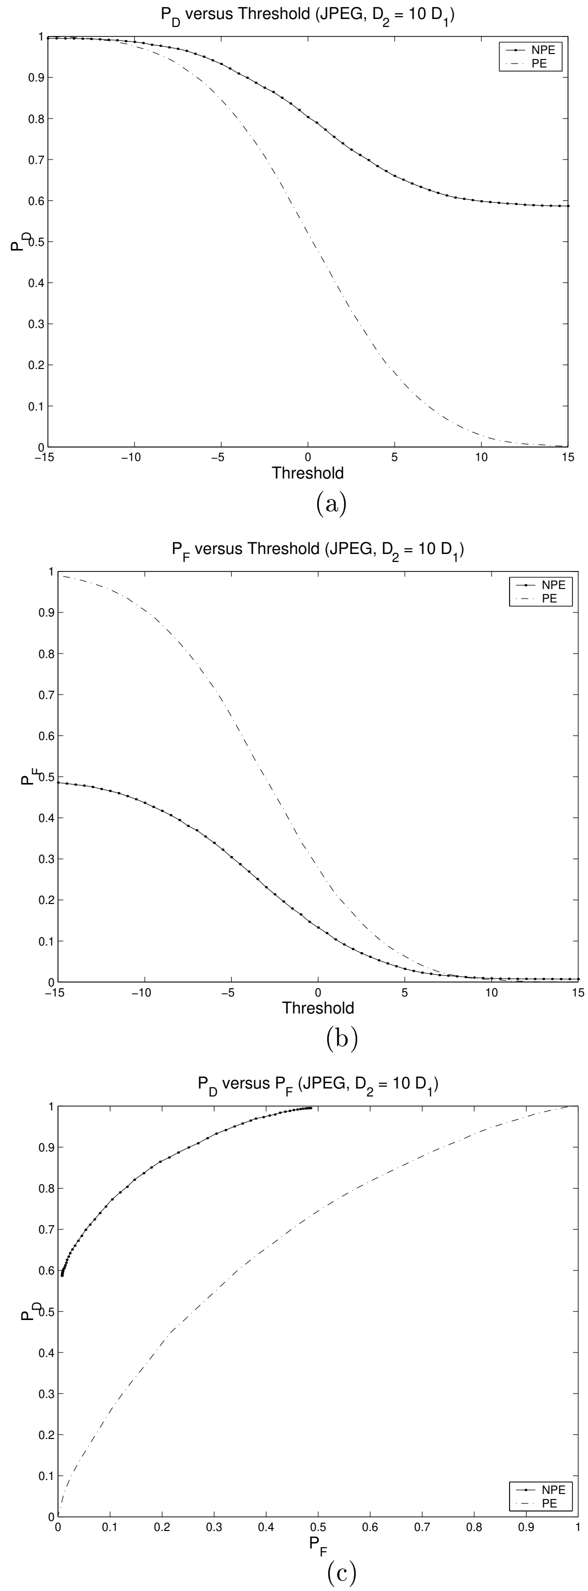

Figure 7.10 P_D and P_F curves for the NPE and PE LRT (JPEG, $D_2 = 10D_1$).

Table 7.1 Selected results for data from the Lena DCT coefficients with an AWGN attack.

Distribution	No Attack		$D_2 = D_1$		$D_2 = 10D_1$	
	$P_D(\%)$	$P_F(\%)$	$P_D(\%)$	$P_F(\%)$	$P_D(\%)$	$P_F(\%)$
Power exponential	98	91	98	91	98	91
Notched power exponential	100	1	98	18	98	22

Table 7.2 Selected results for data from the Lena DCT coefficients with a JPEG compression attack.

Distribution	No Attack		$D_2 = D_1$		$D_2 = 10D_1$	
	$P_D(\%)$	$P_F(\%)$	$P_D(\%)$	$P_F(\%)$	$P_D(\%)$	$P_F(\%)$
Power exponential	98	91	98	91	98	92
Notched power exponential	100	1	98	1	98	42

The values given in these tables reinforce the observations made previously regarding the performance of the LRT detectors based on the two considered distributions. While the use of the power exponential distribution results in very little performance loss when an attack is applied, the outstanding behavior of the notched power exponential distribution translates to a stronger detector, despite significant drops in performance with the inclusion of an attack. However, it should again be noted that neither of these two detectors were derived with attacks included in the models; hence, they are not optimal. Thus, it is expected that the performance loss experienced by the notched power exponential detector will be reduced if the attack were to be included in the formulation of the detector.

CHAPTER 8

CONCLUSIONS

8.1 Summary

This thesis deals with the problem of detecting a watermark that has been multiplicatively embedded into the highest-magnitude transform coefficients of a digital image. In the setup considered, the detector has access to side information about the original image in the form of an image hash (a 1-bit quantized version of a secret subset of the original image DCT coefficients), creating a joint hashing/watermarking system. The inclusion of the side information permits the development of detectors that offer extremely high performance, even for small messages. Various probability distributions, including the Weibull distribution, the power exponential distribution, and the notched power exponential distribution, are examined for use in modeling the statistical distribution of the coefficients selected for watermarking. Through the use of signal detection and estimation techniques prevalent in the field of communications, three detectors are developed and then further specialized to the considered probability distributions. The detectors are based on likelihood ratio testing, maximum likelihood estimation, and locally optimal detection. A fourth detector is also constructed as a simplified version of the likelihood ratio test detector for the case of notched distributions. The precise image watermarking setup employed is semi-blind because some side information is required at the detector; however, the original image need not be available.

Chernoff bounds are formulated for the likelihood ratio test detectors to provide insight into their worst-case performance. Evaluation of the bounds reveals tremendous increases in accuracy when the notched power exponential is utilized for modeling, as opposed to the two more traditional approaches. These insights are supported by the study of the corresponding Chernoff distances, which provide an indication of the relative difficulty of the detection tasks. The Chernoff distance as seen by an eavesdropper attempting to detect the message is also considered, quantifying the increase in difficulty for this detection problem.

Monte Carlo simulations are employed to analyze the behavior of the detectors using both data generated synthetically to follow the appropriate modeling distribution, and data gathered from real-world test images. The performance is measured in terms of the probability of falsely detecting a watermark, and the probability of correctly detecting a watermark. In each of the two simulation settings, the notched power exponential distribution displays vastly superior performance over the Weibull and power exponential distributions, offering decreases in false alarm probabilities between 90% to 50%, depending on the detector type, for the same detection probability. Furthermore, the likelihood ratio test detector provides a much higher degree of performance than those based upon maximum likelihood estimation and locally optimal detection. In the case of the notched power exponential distribution, the simplified version of the likelihood ratio test detector still produces remarkably strong results, with only a small departure from the optimal LRT. Combining these results, it is clear that the use of a likelihood ratio test and the notched power exponential distribution forms a detector that is unrivaled in the watermark detection simulations.

Although the development of the detectors assumes no attacks are made on the watermarked images, further experiments are performed to observe the repercussions of additive white Gaussian noise and JPEG compression on the output of the LRT detectors for the power exponential and notched power exponential distributions. Trials are performed using attack distortions up to 10 times that caused by the embedding process itself, in the mean squared error sense. The false alarm and detection probabilities for the power exponential are only slightly affected by the introduction of an attack, while the notched power exponential detector is moderately hindered. However, the use of the power exponential distribution results in a detector that, in the presence of an attack, is still considerably inferior to that based on the notched power exponential distribution. Moreover, the inclusion of an attack model in the detector derivation will only strengthen these results, with larger improvements expected for the notched power exponential case, where the degradation due to the attacks is more severe.

The analyses and experiments performed in this thesis demonstrate that the use of statistical modeling, and signal detection and estimation theory provides a structured framework in which optimal watermark detectors can be developed and assessed. The exemplary results attained clearly reveal the promise of joint digital image hashing/watermarking as a viable means of information protection for real-world applications.

8.2 Recommendations

The field of digital watermarking is still in its initial stages of development, providing a vast array of topics worthy of study. Several possible recommendations for future directions of research based upon the particular techniques developed in this thesis follow:

- The simulations performed only consider the use of the entire set of transform coefficients as candidates for embedding ($\nu = 1$). Further experiments could be conducted to examine the effects on the detector performance of reducing the size of this set to heighten robustness.

- Multiplicative embedding is currently used to insert the watermark into the highest-magnitude coefficients because it is believed these locations will better withstand attacks. Characteristics of the sensitivity of the human visual system to variations in transform coefficients could be incorporated into the embedding process to better disguise the watermark, while simultaneously increasing security.

- The detectors derived in this thesis do not incorporate models of any of the possible attacks that could be performed on the watermarked image. By compensating for such processes, the watermark detectors would become more suitable for use in practical situations.

- Digital grayscale images represent only a single source of multimedia objects to which watermarking techniques may be applied. Although the detection approaches described in this thesis are independent of the host data, further considerations must be made to frame the basic watermarking system for use with additional objects such as color images, audio sequences, and video data.

8.3 Contributions

The contributions made by the research and development of this thesis may be summarized as follows:

- The construction of a joint image hashing/watermarking system based on multiplicative embedding into high-magnitude image transform coefficients.
- The extension of the power exponential distribution to the notched power exponential distribution to better model the coefficients selected for watermarking.
- The development of detectors based on likelihood ratio testing, maximum likelihood estimation, and locally optimal detection using the notched power exponential distribution.
- The derivation of Chernoff bounds on the performance of the likelihood ratio test detector for the notched power exponential distribution.
- The formulation of Chernoff distances for the notched power exponential likelihood ratio test detector, and for the detection problem as seen by an eavesdropper.
- The simulation of the considered detectors, highlighting the substantial performance gains resulting from the use of notched distributions.

REFERENCES

[1] G. Voyatzis and I. Pitas, "The use of watermarks in the protection of digital multimedia products," *Proceedings of the IEEE*, vol. 87, no. 7, pp. 1197–1207, 1999.

[2] I. Cox, J. Kilian, T. Leighton, and T. Shamoon, "Secure spread spectrum watermarking for multimedia," *IEEE Transactions on Image Processing*, vol. 6, no. 12, pp. 1673–1687, 1997.

[3] A. De Rosa, M. Barni, F. Bartolini, V. Cappellini, and A. Piva, "Optimum decoding of non-additive full frame DFT watermarks," in *Proceedings of the Third International Information Hiding Workshop*, 1999, pp. 167–179.

[4] J. Oostveen, T. Kalker, and J. P. Linnartz, "Optimal detection of multiplicative watermarks," in *Proceedings of the European Signal Processing Conference*, vol. 5, 2000, pp. 2973–2976.

[5] F. Müller, "Distribution shape of two-dimensional DCT coefficients of natural images," *Electronics Letters*, vol. 29, no. 22, pp. 1935–1936, 1993.

[6] S. Choi, A. Cichocki, and S.-I. Amari, "Flexible independent component analysis," *Journal of VLSI Signal Processing*, vol. 26, no. 1/2, pp. 25–38, 2000.

[7] H. Poor, *An Introduction to Signal Detection and Estimation.* New York: Springer–Verlag, 2nd ed., 1994.

[8] Q. C. Cheng and T. S. Huang, "An additive approach to transform-domain information hiding and optimum detection structure," *IEEE Transactions on Multimedia*, vol. 3, no. 3, pp. 273–284, 2001.

APPENDIX A

ADDITIONAL THEORETICAL ANALYSIS

This appendix presents analyzes of the maximum likelihood estimation detector and locally optimal detector for the Weibull and power exponential distributions when synthetic data are utilized. The setup of the experiments performed is identical to that described in Chapter 5.

A.1 Maximum Likelihood Estimation

The detectors based upon maximum likelihood estimation are now considered for the cases when the Weibull and power exponential distributions are used to model the distribution of image transform coefficients. As before, three plots are employed to examine the performance of the detectors.

A.1.1 Weibull distribution results

The MLE detector based on the Weibull distribution is now investigated. The distributions of the test statistic, defined by Equation (3.5), under the two hypotheses are shown in Figure A.1. These plots demonstrate that the distributions of the estimates under each hypothesis nearly completely overlap. However, the higher embedding strength does result in slightly more separation, as expected. Under both hypotheses, a large number of the estimates are of value 0, thus, a peak exists at this point. A sharp knee results around 0.05 because of the large peak and the finite number of bins utilized when plotting the histograms of the distributions.

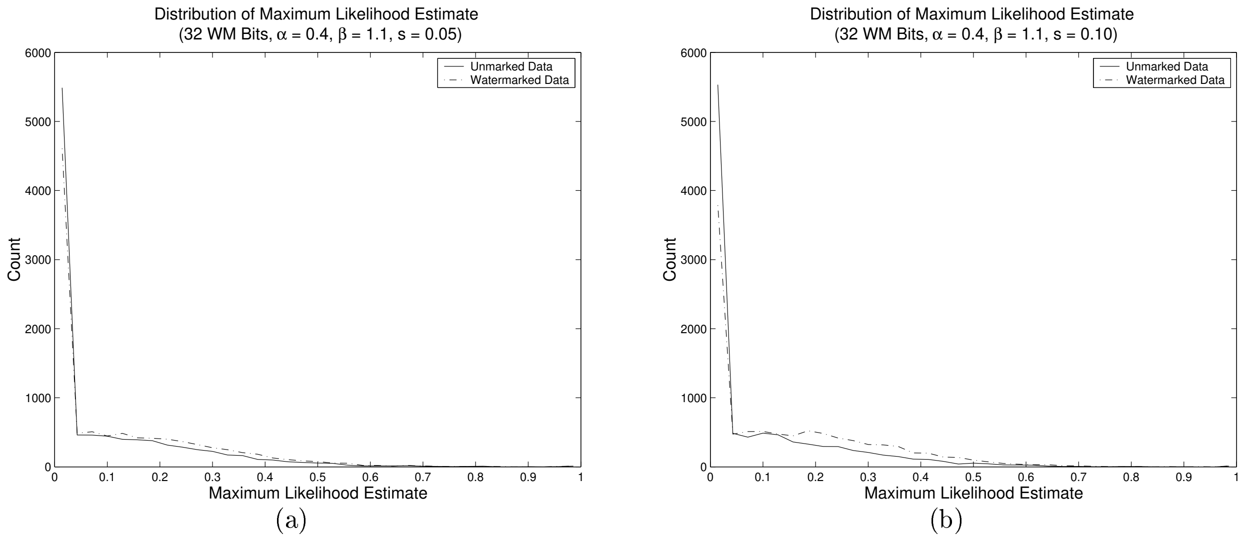

Figure A.1 Distribution of the test statistic using the WB distribution and MLE with an embedding strength of (a) 0.05 and (b) 0.10.

The next two figures, A.2 and A.3, contain the results of the Monte Carlo simulations using the MLE detector. First, it should be noted that the threshold against which the estimate is compared is in the range [0, 1] because the strength must be in this range. As a result, full ranges of probabilities are not necessarily present in the ROC curves. Since the figures contain estimates for actual embedding strengths of 0.05 and 0.10, it is expected that the detection probability will be low when the threshold is above these values. By comparing Figure A.2 and Figure A.3, the expected increase in performance is observed as the embedding strength is increased. However, neither strength presents a robust detector. With approximately a 50% false alarm probably for approximately a 60% detection probability for $s^* = 0.05$, the MLE detector performs slightly worse than the corresponding binary hypothesis testing detector, whose results are presented in Figure 5.7 (and Figure 5.8 for the $s^* = 0.10$ case). Hence, it is doubtful that the MLE detector based on the Weibull distribution will suffice in practical applications.

A.1.2 Power exponential distribution results

Next, the maximum likelihood estimator that is based on modeling the host coefficients using the power exponential distribution is evaluated. The distributions of the test statistic, defined

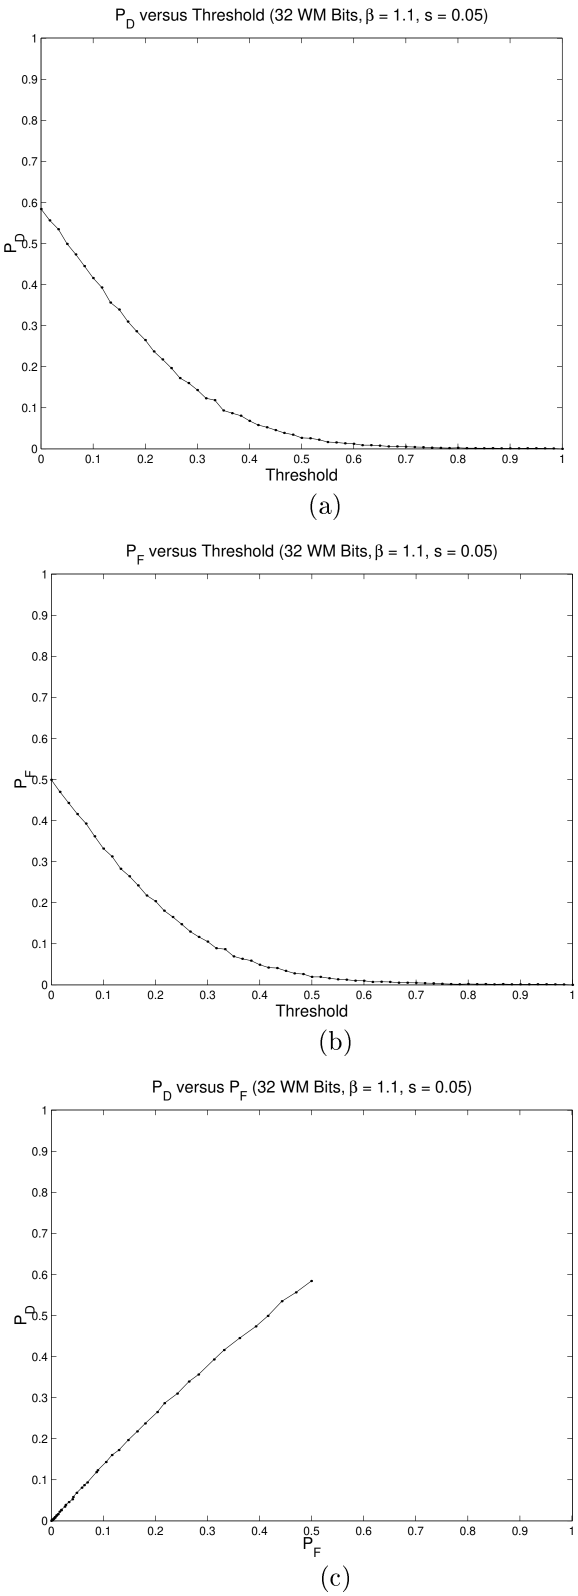

Figure A.2 P_D and P_F curves for the WB MLE for synthetic data ($s^* = 0.05$).

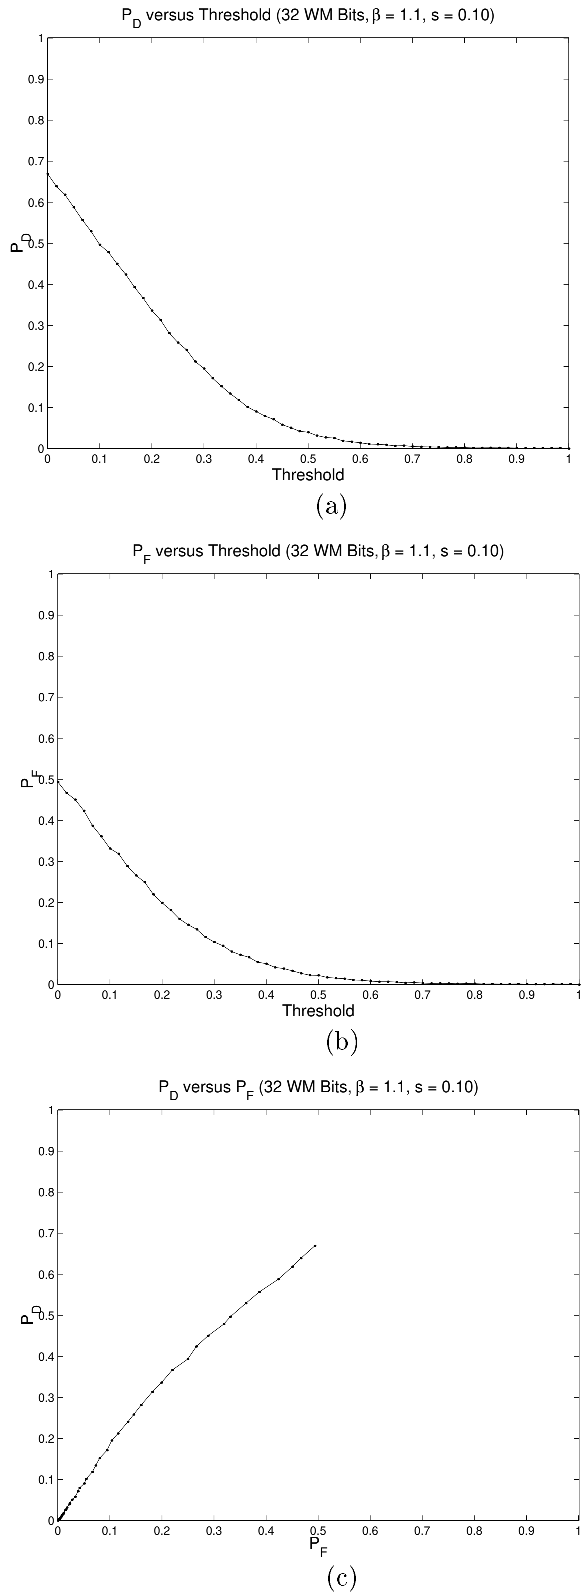

Figure A.3 P_D and P_F curves for the WB MLE for synthetic data ($s^* = 0.10$).

by Equation (3.6), as found using Monte Carlo simulation are given in Figure A.4. Here the results are quite similar to those for the MLE detector using the Weibull distribution: the distributions are greatly overlapped and a slight improvement is gained by increasing the embedding strength.

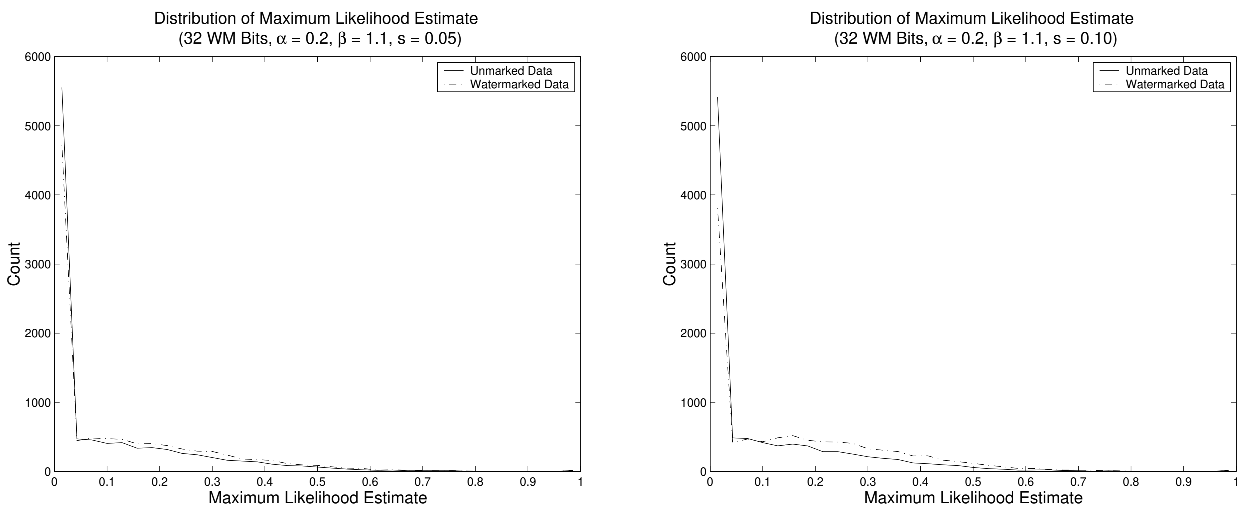

(a) (b)

Figure A.4 Distribution of the test statistic using the PE distribution and MLE with an embedding strength of (a) 0.05 and (b) 0.10.

The performance indicators of P_D and P_F are shown in the next two figures, Figure A.5 and Figure A.6. The curves depicted are nearly identical to those for the case of the Weibull distribution. An increase in performance as the embedding strength is increased is observed; however, the detector performs rather poorly overall. Once again, for a detection probability of approximately 60%, a false alarm probability of approximately 50% must be tolerated. These results are slightly worse than those for the likelihood ratio test detector for the same distribution, whose ROC curves are given in Figures 5.10 and 5.11. Thus, it appears that the MLE detector using the power exponential distribution is not a viable solution.

A.2 Locally Optimal Detection

This section examines the watermark detectors based on locally optimal detection for the Weibull and power exponential distributions. Once again, three figures are utilized to analyze the experimental results.

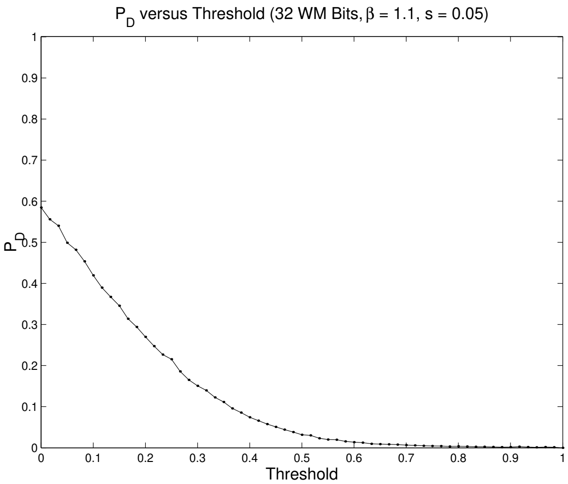

(a)

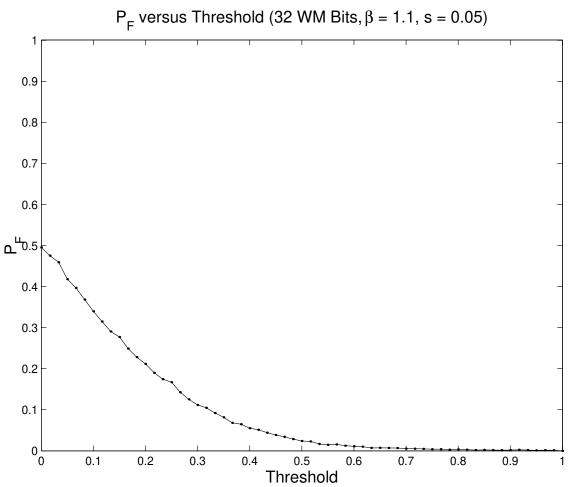

(b)

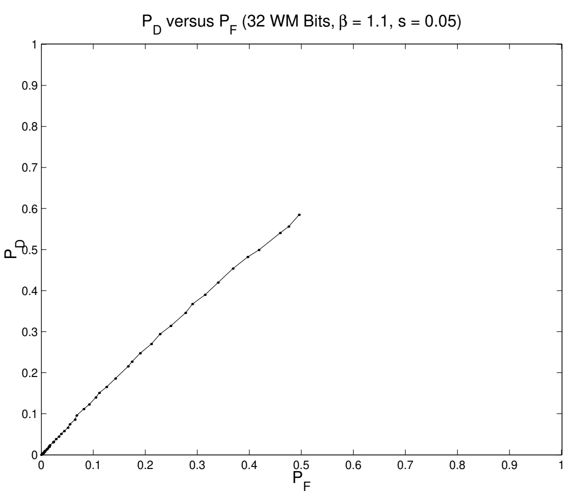

(c)

Figure A.5 P_D and P_F curves for the PE MLE for synthetic data ($s^* = 0.05$).

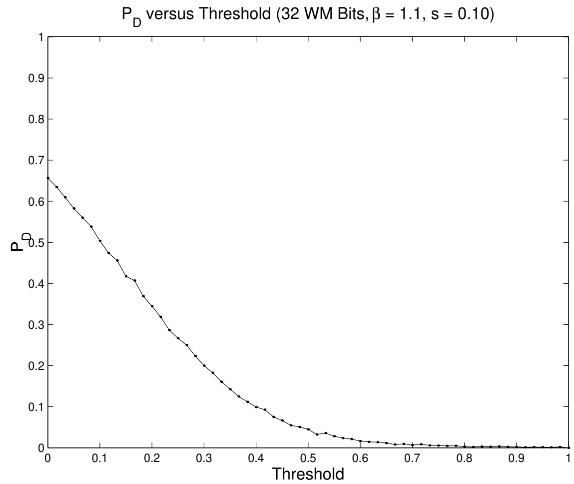

(a)

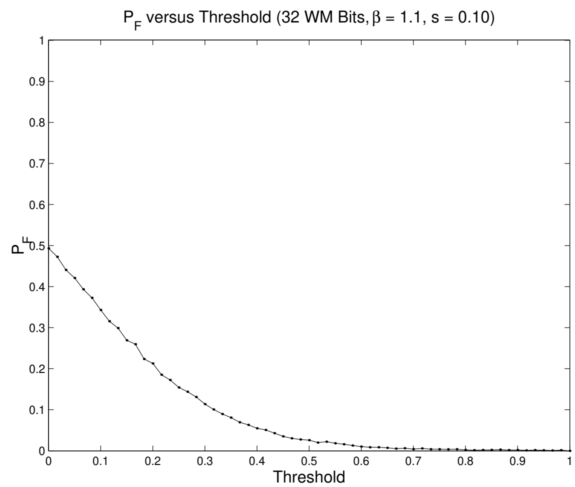

(b)

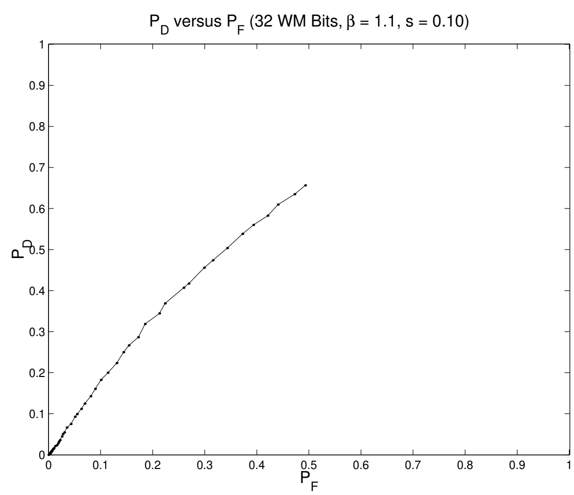

(c)

Figure A.6 P_D and P_F curves for the PE MLE for synthetic data ($s^* = 0.10$).

A.2.1 Weibull distribution results

To further evaluate the locally optimal detector, the detector created for the case when the host coefficients are modeled using the Weibull distribution is considered. The distributions of the test statistic, defined by Equation (3.8), under the two hypotheses are given in Figure A.7. This figure shows the two distributions overlap quite significantly, with slightly more separation present for the higher embedding strength.

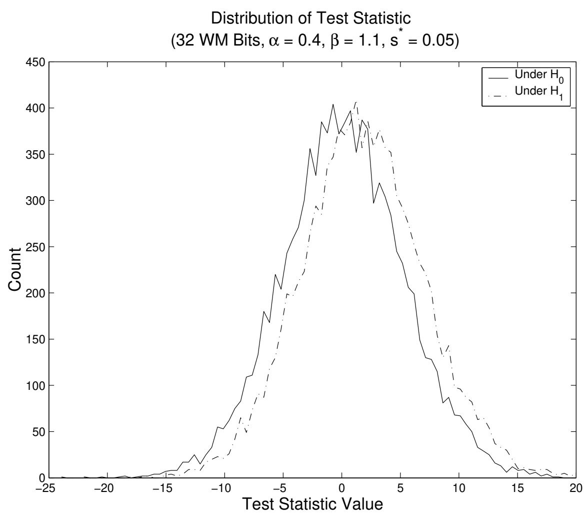

(a)

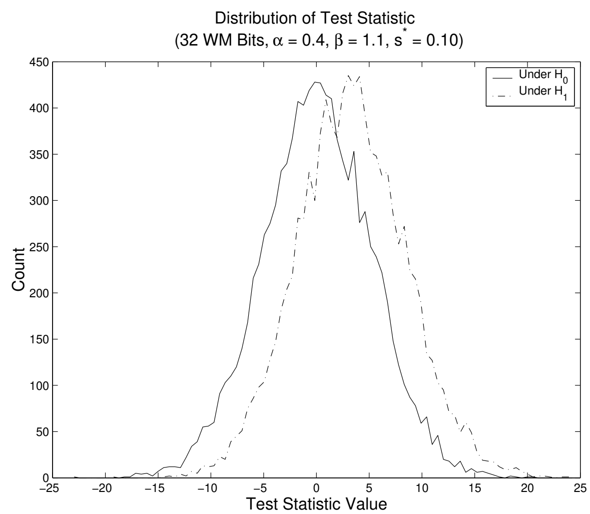

(b)

Figure A.7 Distribution of the test statistic using the WB distribution and LOD with an embedding strength of (a) 0.05 and (b) 0.10.

The Weibull distribution locally optimal detector is now examined using the P_D and P_F statistics. Figures A.8 and A.9 contain the results of the Monte Carlo simulations using this detector. The ROC plots illustrate a slight performance gain when the embedding strength is increased to 0.10 from 0.05, which is to be expected. However, in general, the results are not promising. In order to achieve approximately a 64% detection probability for the lower embedding strength, close to a 50% false alarm probability results. These statistics lie just below those for the corresponding binary hypothesis testing detector, and just above the maximum likelihood detector. Therefore, it appears that the locally optimal detector constructed using the Weibull distribution does not offer sufficient performance to be employed in practice.

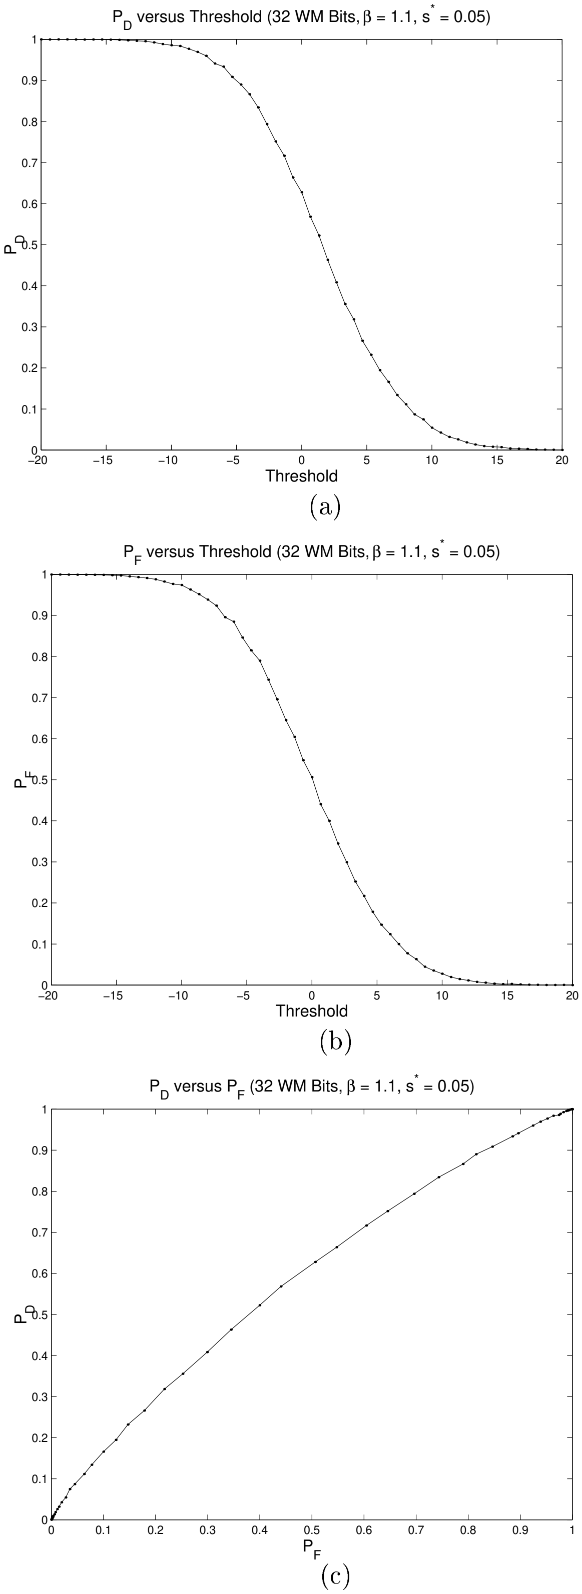

Figure A.8 P_D and P_F curves for the WB LOD for synthetic data ($s^* = 0.05$).

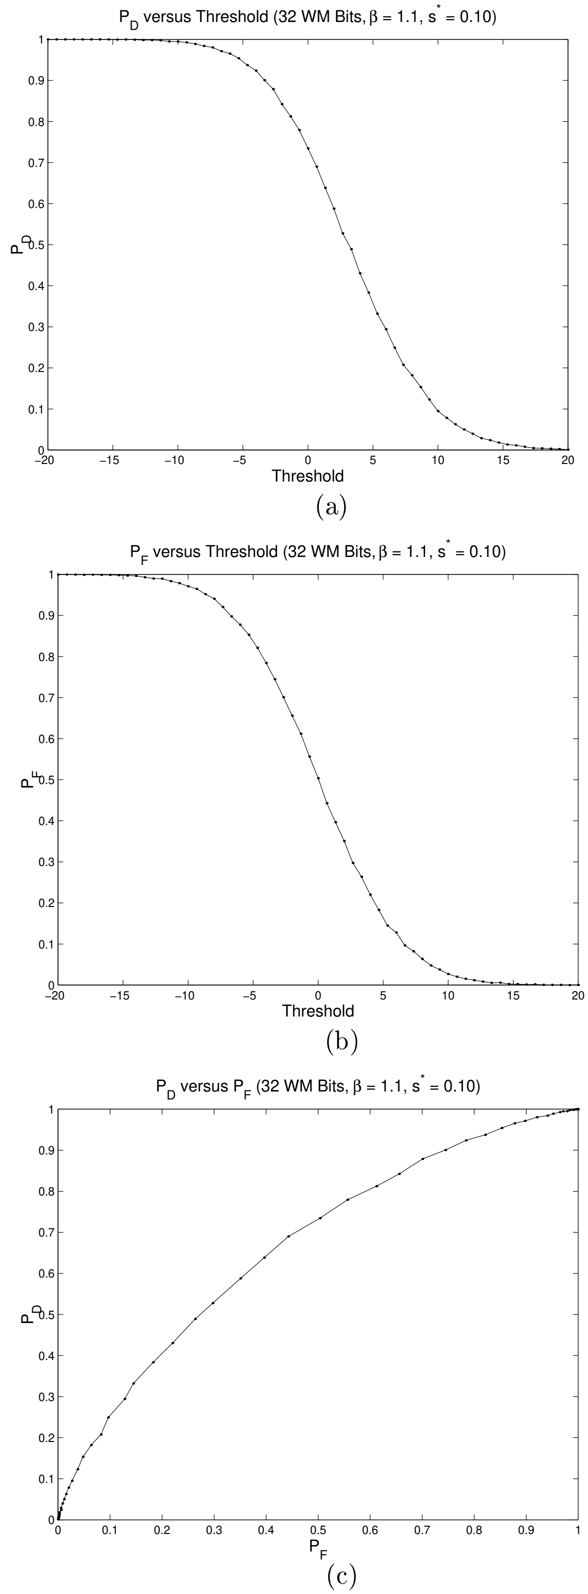

Figure A.9 P_D and P_F curves for the WB LOD for synthetic data ($s^* = 0.10$).

A.2.2 Power exponential distribution results

The next distribution considered for the locally optimal detector is the power exponential distribution. The distributions of the test statistics, defined by Equation (3.9), resulting from Monte Carlo simulations on this detector are displayed in Figure A.10. The results in this situation are quite similar to those of the LOD using the Weibull distribution: the statistic distributions display much overlap with a slight improvement in separation for the higher embedding strength.

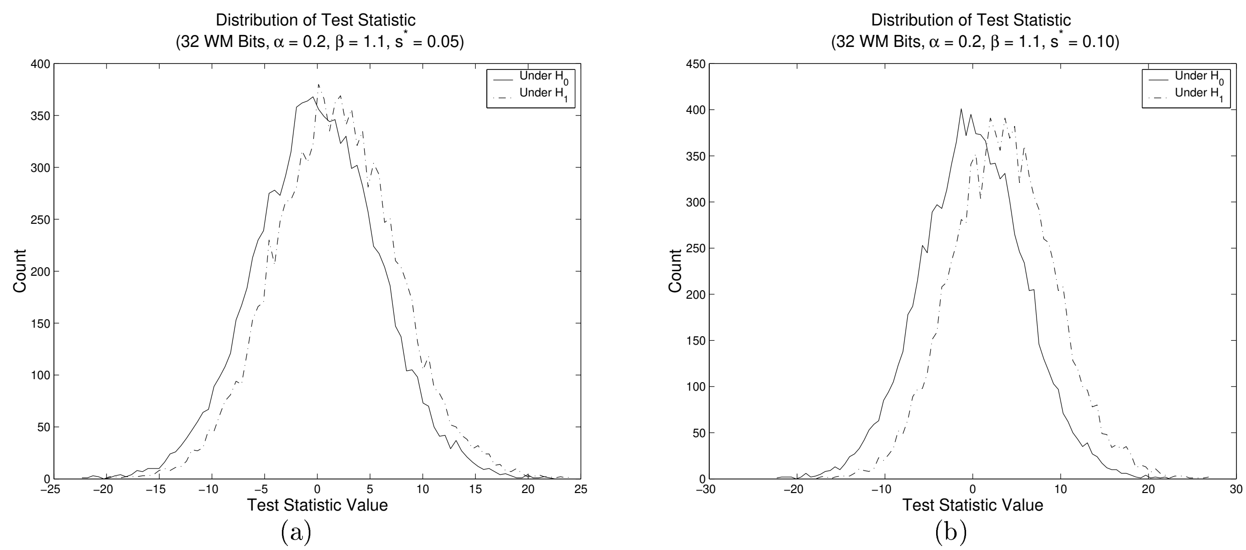

Figure A.10 Distribution of the test statistic using the PE distribution and LOD with an embedding strength of (a) 0.05 and (b) 0.10.

Figures A.11 and A.12 demonstrate the performance of the locally optimal detector in terms of P_D and P_F. These results are quite similar to those seen previously. The ROC curve moves slightly more to the top left corner of the plot as the embedding strength is increased; but overall the detector does not perform well. In order to achieve a detection probability of approximately 65% with the lower embedding strength, a false alarm probability of approximately 50% results. These statistics are more favorable than those for the maximum likelihood estimator of the same distribution, and comparable to those of the corresponding likelihood ratio test detector. Consequently, it is unlikely that the locally optimal detector using the power exponential distribution to model the coefficients will provide an adequate watermark detector in practice.

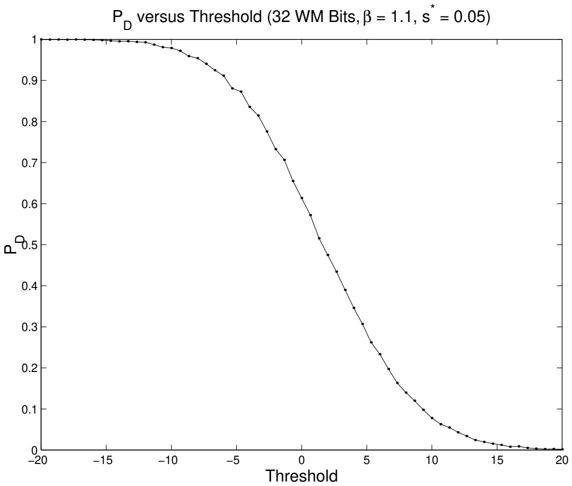

(a)

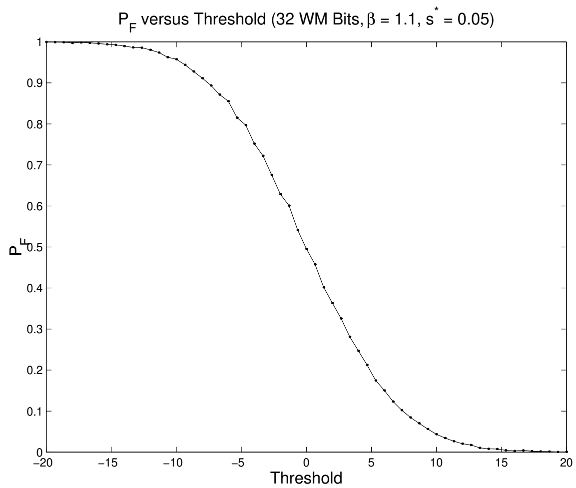

(b)

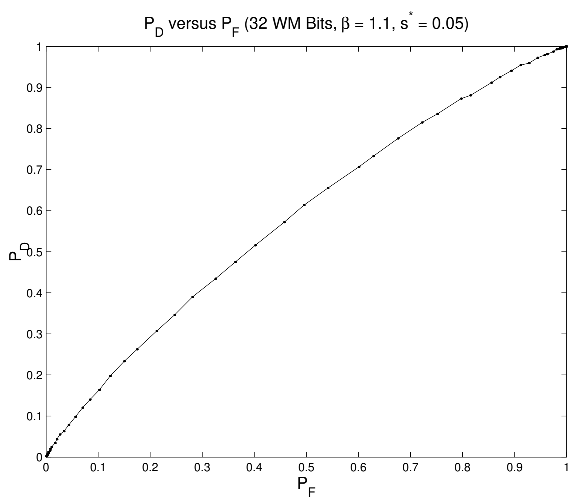

(c)

Figure A.11 P_D and P_F curves for the PE LOD for synthetic data ($s^* = 0.05$).

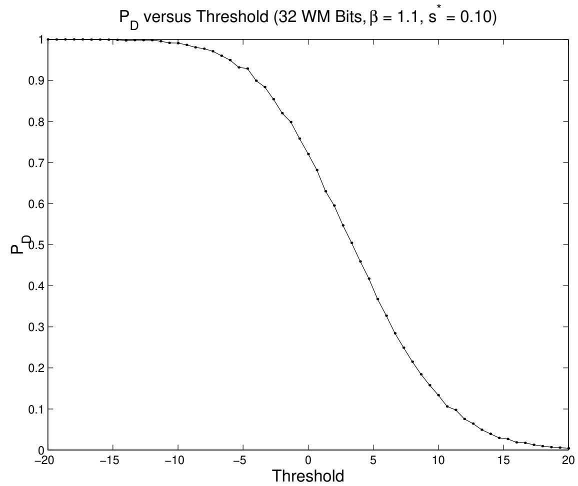

(a)

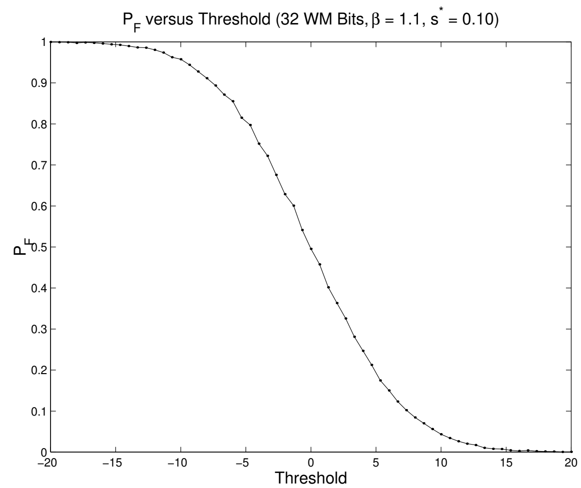

(b)

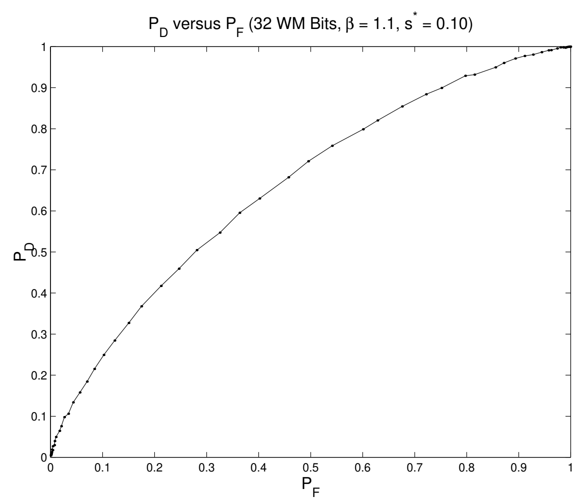

(c)

Figure A.12 P_D and P_F curves for the PE LOD for synthetic data ($s^* = 0.10$).

APPENDIX B

ADDITIONAL IMAGE WATERMARKING EXPERIMENTS

The analyzes of the watermark detectors based on the Weibull and power exponential distributions for use with the Lena and Peppers images are presented in this appendix. The setup of the experiments performed is nearly identical to that described for the notched power exponential distribution in Chapter 6. The main difference is that only 10 000 trials are performed for each Monte Carlo simulation because the resulting probabilities for these distributions are do not require the same fine degree of precision as those for the notched power exponential. Also, note that since the Weibull distribution requires non-negative data, the absolute values of the DCT coefficients are utilized for watermarking.

B.1 Watermarking Lena DCT Coefficients

In this section, the Lena image, as shown in Figure 6.1, is utilized as the source of coefficients into which the watermark is to be inserted.

B.1.1 Simple likelihood ratio test

First, the detectors based upon likelihood ratio testing are examined for the Lena image through the use of the three typical figures.

B.1.1.1 Weibull distribution results

To continue the analysis of the likelihood ratio test detector on the Lena image, the Weibull distribution is utilized to model the selected coefficients. The distribution of the test statistic under each hypothesis is given in Figure B.1. As expected, this figure demonstrates that the separation between the two distributions increases as the embedding strength increases.

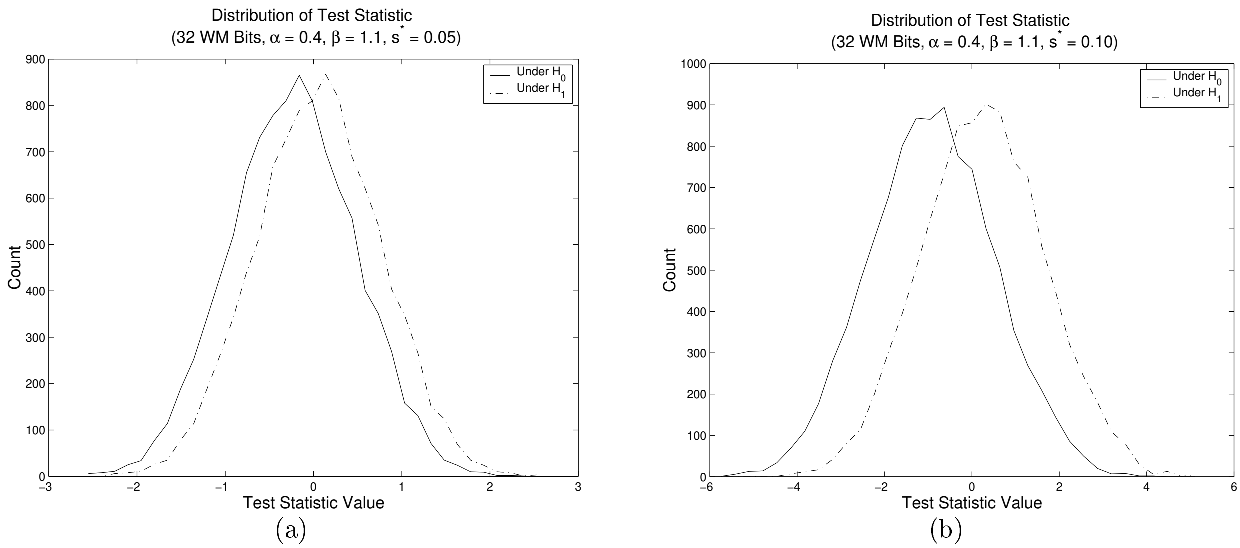

Figure B.1 Distribution of the test statistic for the Lena image using the WB distribution and LRT with an embedding strength of (a) 0.05 and (b) 0.10.

At this point the performance of the detector is considered using the detection and false alarms probabilities. Figures B.2 and B.3 each contain three plots demonstrating the outcome of the Monte Carlo simulations for embedding strengths of 0.05 and 0.10, respectively. From these graphs, it is seen that the detector performs better as the embedding strength is increased. However, for an embedding strength of 0.05, a false alarm probability of approximately 50% must be tolerated to achieve a detection probability just under 70%. Thus, using the Weibull distribution to model the chosen Lena coefficients appears to not result in a practical watermark detector.

B.1.1.2 Power exponential distribution results

The next distribution evaluated for use in a likelihood ratio test detector is the power exponential distribution. Figure B.4 shows the distribution of the test statistic under each of the two hypotheses. These results are similar to those in the case of the Weibull distribution: the increased embedding strength causes the two distributions to become further apart. It should be noted that the jaggedness in the second plot is simply due to an increase in the number of bins used to compute the histogram.

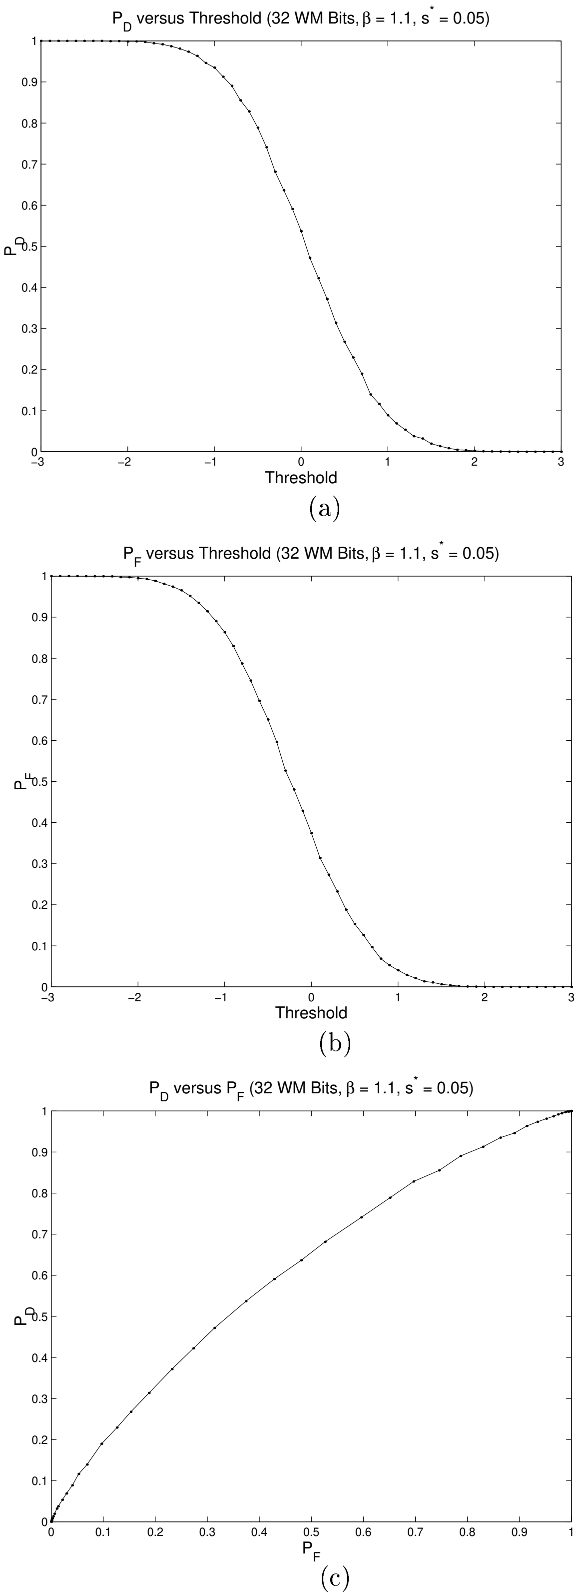

Figure B.2 P_D and P_F curves for the WB LRT for Lena data ($s^* = 0.05$).

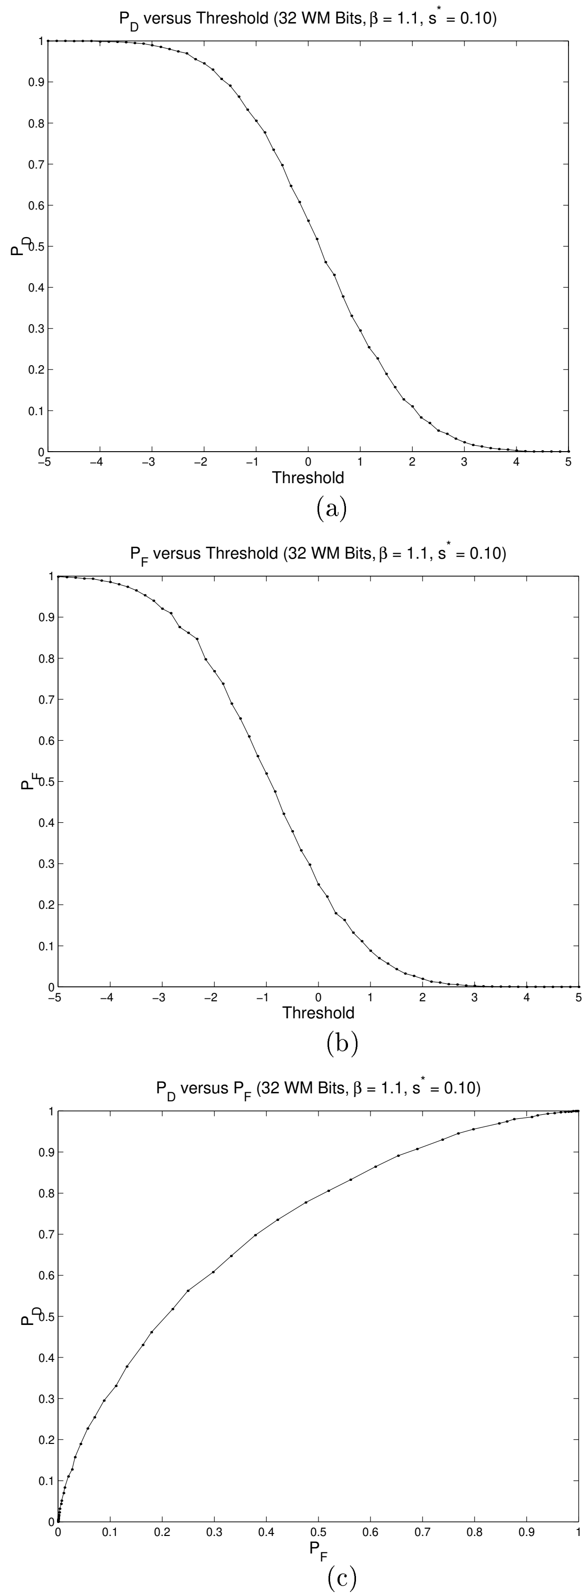

Figure B.3 P_D and P_F curves for the WB LRT for Lena data ($s^* = 0.10$).

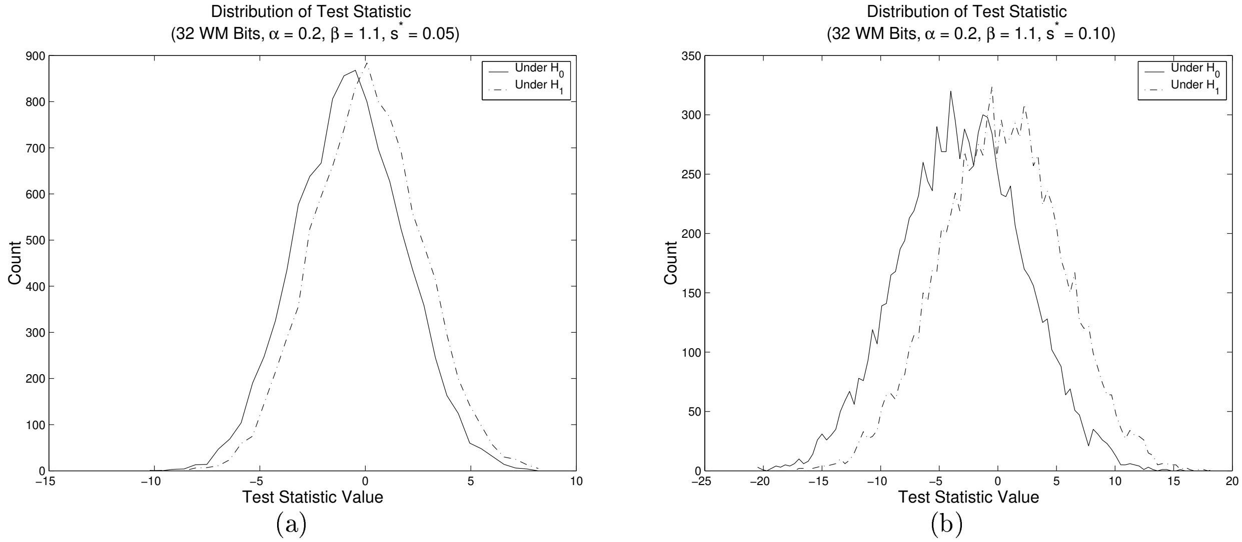

Figure B.4 Distribution of the test statistic for the Lena image using the PE distribution and LRT with an embedding strength of (a) 0.05 and (b) 0.10.

The P_D and P_F curves for the Monte Carlo simulations are now considered. Figures B.5 and B.6 present the three performance plots: the first figure for an embedding strength of 0.05, and the second for a strength of 0.10. Once again, it is noticed that the ROC moves slightly closer to the top left corner when the embedding strength is increased, thus indicating that the watermark is easier to detect. Unfortunately, to achieve a detection probability of approximately 64%, a false alarm probability of approximately 50% is incurred. This performance is slightly worse than that demonstrated for the binary hypothesis testing detector using the Weibull distribution to model the Lena image coefficients. Thus, the detector based on the power exponential distribution is likely also unfit for real-world use.

B.1.2 Maximum likelihood estimation

The detector based on maximum likelihood estimation is now considered for the Weibull and power exponential distributions. Once again, figures of the distribution of the test statistic and P_D and P_F curves are employed to analyze detector performance.

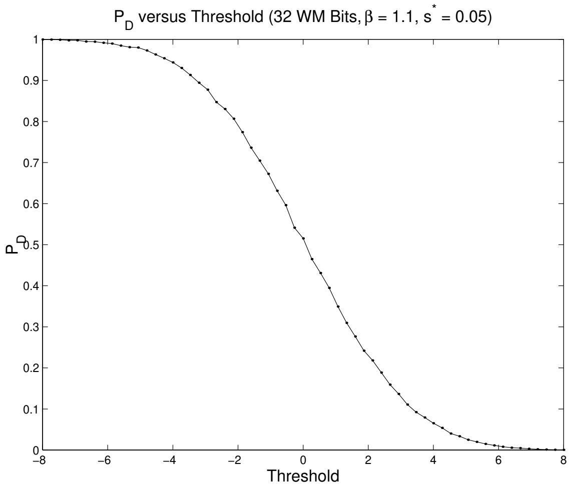

(a)

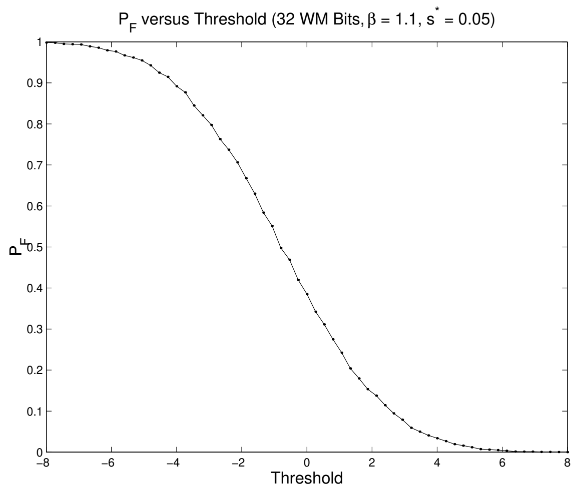

(b)

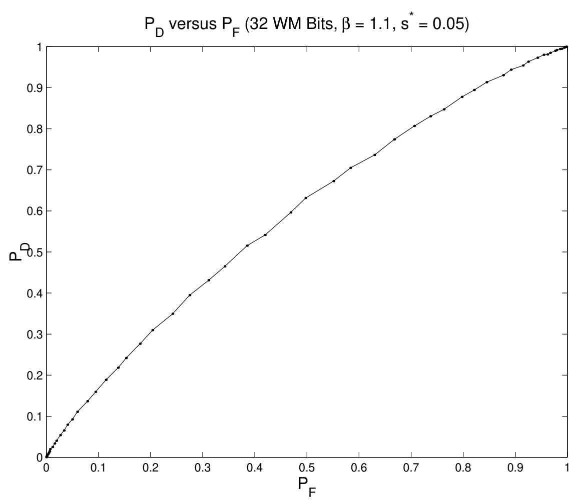

(c)

Figure B.5 P_D and P_F curves for the PE LRT for Lena data ($s^* = 0.05$).

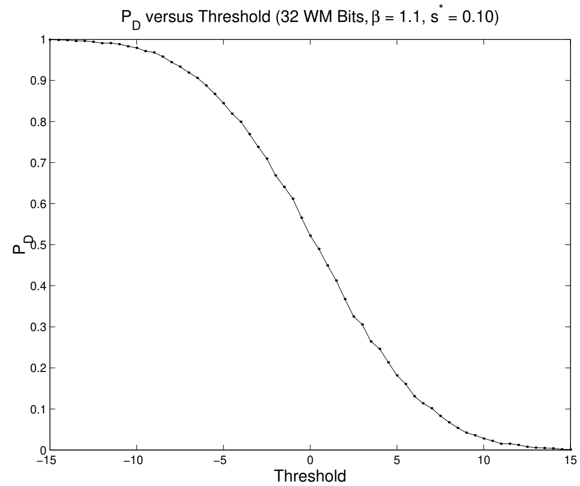

(a)

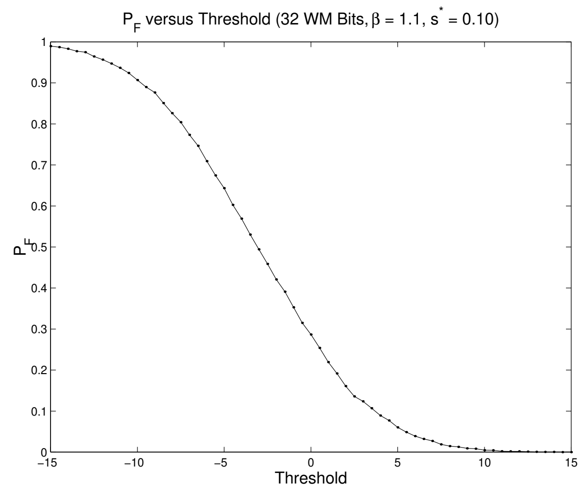

(b)

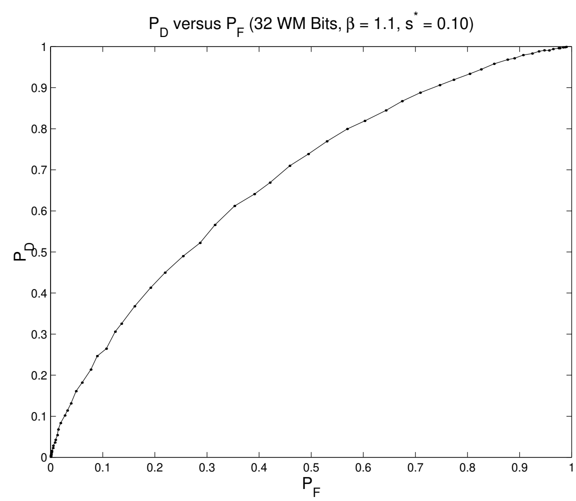

(c)

Figure B.6 P_D and P_F curves for the PE LRT for Lena data ($s^* = 0.10$).

B.1.2.1 Weibull distribution results

To continue the evaluation of the maximum likelihood estimation detectors, the case when the Weibull distribution is used to model the DCT coefficients is considered. The distribution of the estimate under each of the two hypotheses is given in Figure B.7. For both $s = 0.05$ and $s = 0.10$, the two distributions are nearly completely overlapping, but do contain larger values near the actual embedding strengths. The two distributions also become slightly more separated as the watermark embedding strength is increased, which is to be expected.

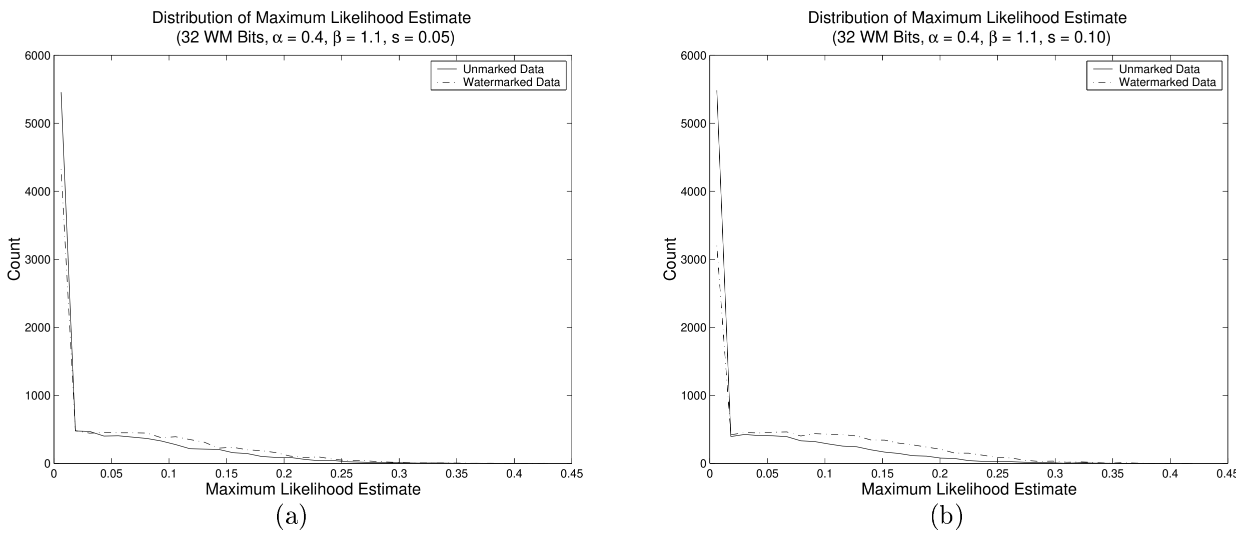

Figure B.7 Distribution of the test statistic for the Lena image using the WB distribution and MLE with an embedding strength of (a) 0.05 and (b) 0.10.

Figures B.8 and B.9 display the result of the Monte Carlo simulation in terms of detection and false alarm probabilities. As previously with maximum likelihood estimation, the range of threshold is limited to [0, 1) because the embedding strength must be inside this range. Consequently, P_D and P_F do not necessarily vary over the full range of 0 to 1. It is observed that, in the figures, when the threshold is raised above the actual embedding strength, the detection probability falls off significantly because the threshold is greater than the value being estimated. When the two figures are compared, an increase in performance is observed for the higher embedding strength. However, the detector performance is quite lacking. To achieve a detection probability of approximately 62% when $s = 0.05$, approximately a 50% chance of producing a false alarm is incurred, which is slightly worse than the

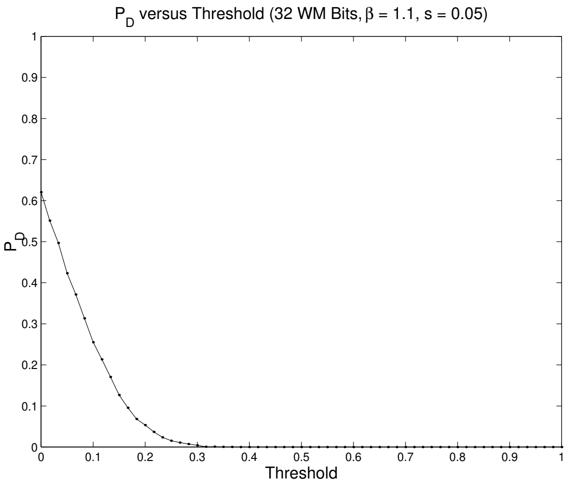

(a)

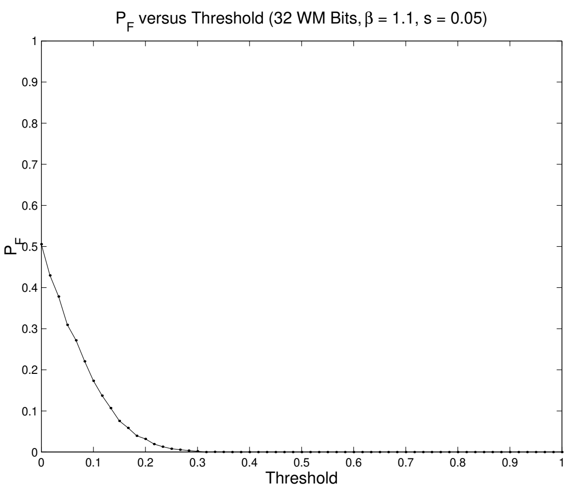

(b)

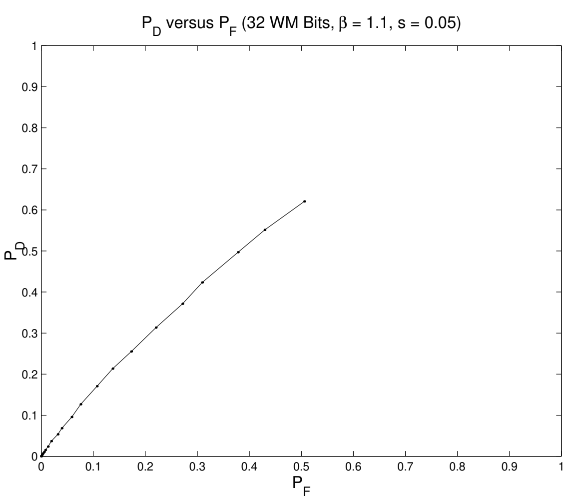

(c)

Figure B.8 P_D and P_F curves for the WB MLE for Lena data ($s^* = 0.05$).

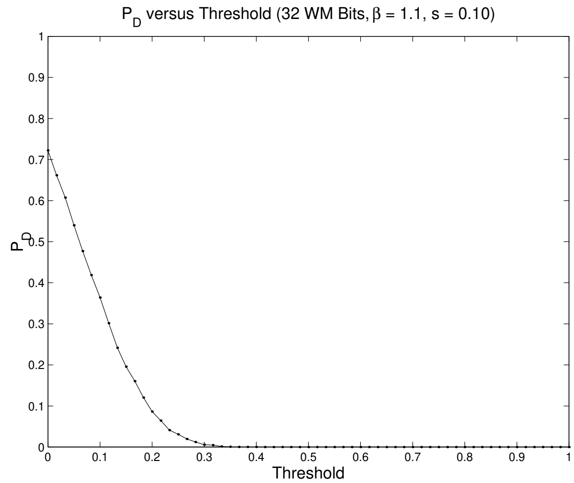

(a)

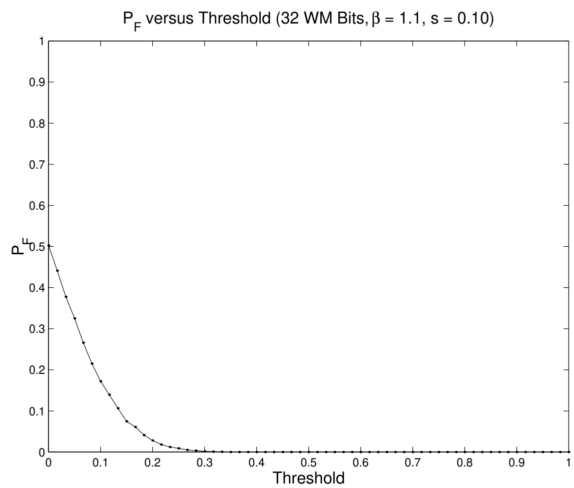

(b)

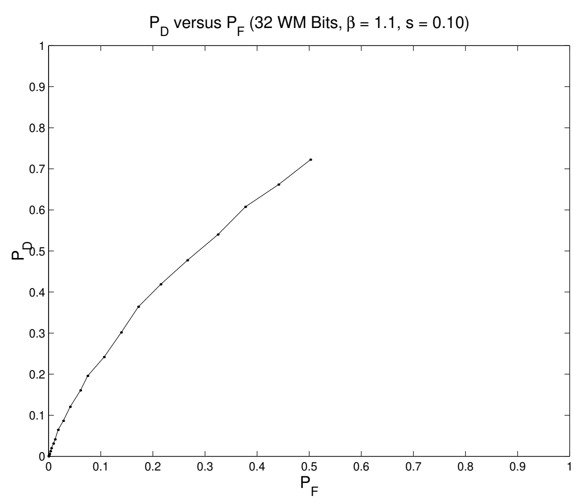

(c)

Figure B.9 P_D and P_F curves for the WB MLE for Lena data ($s^* = 0.10$).

corresponding likelihood ratio test detector. These statistics indicate that the use of maximum likelihood estimation and the Weibull distribution to model the Lena image coefficients likely will not provide a practical watermark detector. This conclusion is expected given the poor performance of this detector on synthetic data.

B.1.2.2 Power exponential distribution results

The maximum likelihood estimator used to detect a watermark's presence is now evaluated for the Lena image data when the coefficients are modeled using the power exponential distribution. To begin, Figure B.10 illustrates the Monte Carlo simulation results for the distribution of the estimate under each of the two hypotheses. As in the case of the Weibull distribution, the distributions are quite overlapping and display values in the areas around the true embedding strengths. It is also observed that the distributions become slightly more separated as the embedding strength is increased.

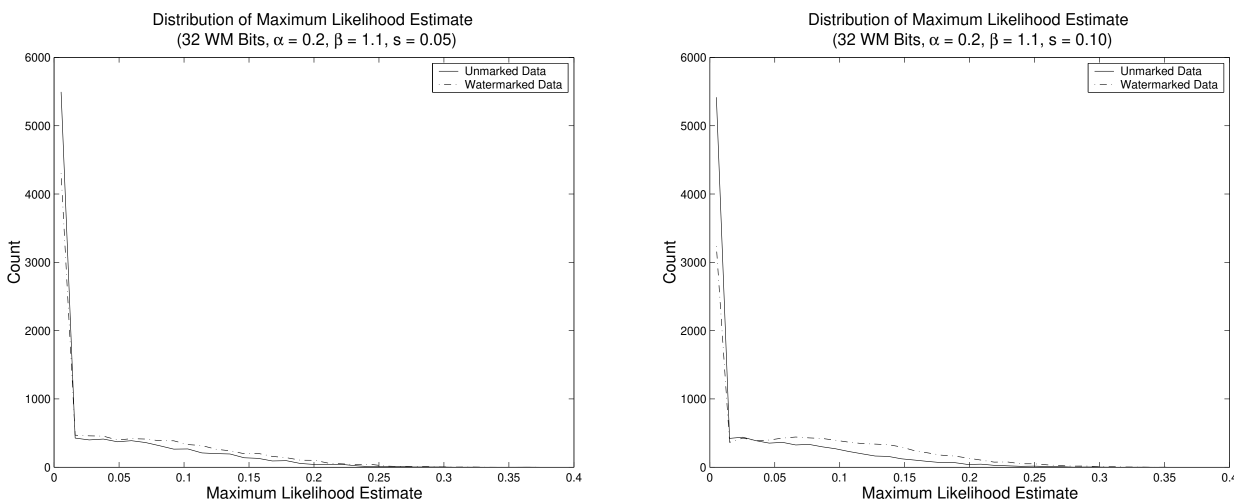

(a) (b)

Figure B.10 Distribution of the test statistic for the Lena image using the PE distribution and MLE with an embedding strength of (a) 0.05 and (b) 0.10.

Figures B.11 and B.12 contain the performance indicator plots for embedding strengths of 0.05 and 0.10, respectively. These plots are quite similar to those found for the Weibull distribution. The detection probability drops off significantly as the threshold is increased above the actual embedding strength, as expected. Also observable is an increase in performance with the higher embedding strength. However, the detector performs rather poorly

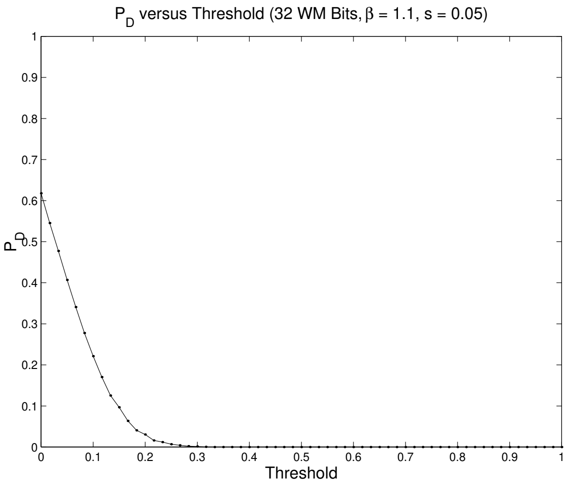

(a)

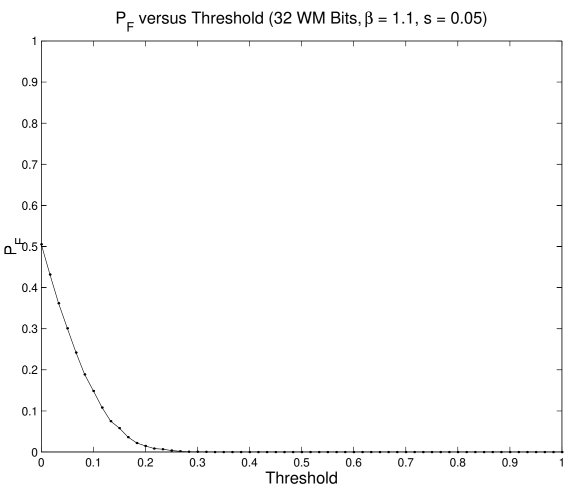

(b)

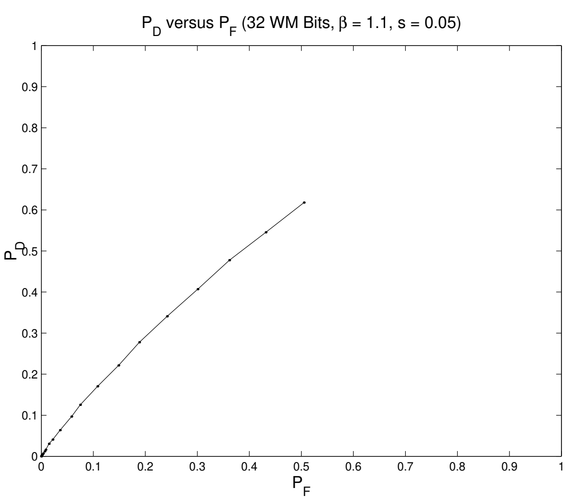

(c)

Figure B.11 P_D and P_F curves for the PE MLE for Lena data ($s^* = 0.05$).

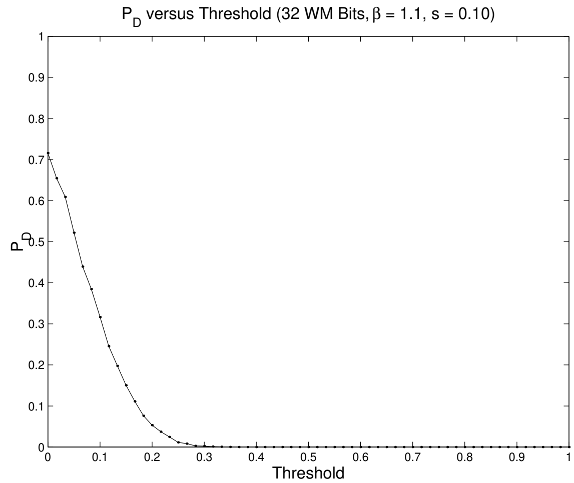

(a)

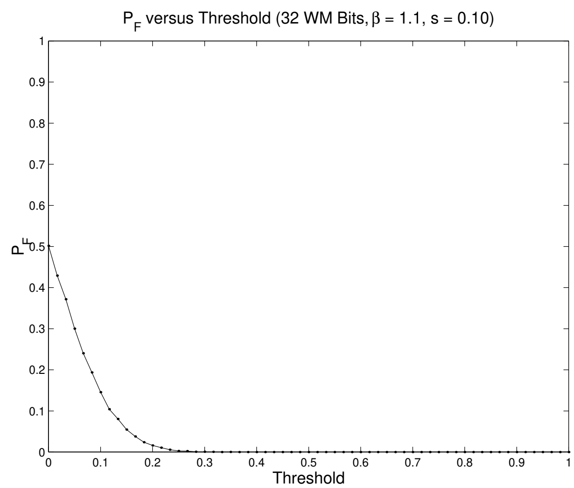

(b)

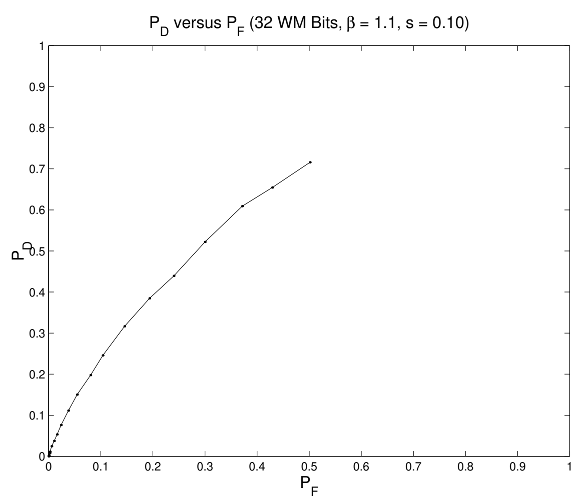

(c)

Figure B.12 P_D and P_F curves for the PE MLE for Lena data ($s^* = 0.10$).

overall. For the case when $s = 0.05$, a detection probability of approximately 62% can only be achieved with a false alarm probability of approximately 50%. These statistics are slightly worse than the corresponding likelihood ratio test results. Thus, it is doubtful that using maximum likelihood estimation with the power exponential distribution will yield a detector that is suitable for real-world applications.

B.1.3 Locally optimal detection

The analysis of the locally optimal detector on the Lena image data is now studied for the cases when the Weibull and power exponential distributions are employed to model the image transform coefficients. Again, three figures are utilized for each distribution to examine the detector performance.

B.1.3.1 Weibull distribution results

The first additional distribution evaluated for use with locally optimal detection is the Weibull distribution. Figure B.13 provides the distribution of the test statistic under each of the two hypotheses as found through Monte Carlo simulations. Considering the plot for the low embedding strength, the two distributions are still fairly overlapping; however, for the larger embedding strength, a moderate degree of separation is indeed present.

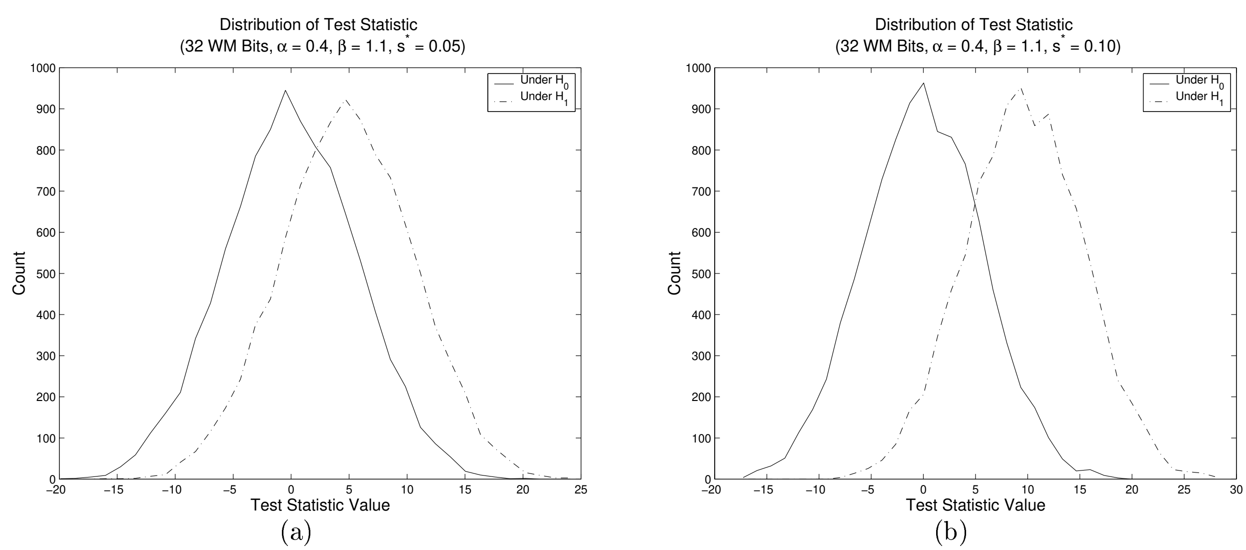

Figure B.13 Distribution of the test statistic for the Lena image using the WB distribution and LOD with an embedding strength of (a) 0.05 and (b) 0.10.

The P_D and P_F curves for the two embedding strengths, found through Monte Carlo simulations, are given in Figures B.14 and B.15. By comparing ROC curves in the two figures, a clear improvement in performance is seen for the higher embedding strength, as expected. However, the detector performs only moderately well in general. For the case when $s^* = 0.05$, a detection probability of approximately 80% is achieved with a false alarm probability of approximately 50%. These results offer an improvement over both the corresponding likelihood ratio test and maximum likelihood estimator detectors; however, it is doubtful that the detector will yield strong enough performance for use in real-world applications.

B.1.3.2 Power exponential distribution results

The locally optimal detector formulated based on the power exponential distribution is now evaluated using the Lena DCT coefficients. First, the distributions of the test statistics under H_0 and H_1 are given in Figure B.16. These graphs are quite similar to those for the Weibull distribution locally optimal detector. The separation between the two distributions increases to a moderate amount as the embedding strength is increased.

Figures B.17 and B.18 demonstrate the performance of the detector in terms of detection and false alarm probabilities. Considering the ROC curves, a clear increase in performance is observed as the embedding strength is raised. Unfortunately, the general performance is somewhat lacking. With the low embedding strength, approximately a 50% false alarm probability must be tolerated to achieve a detection probability just over 80%. These statistics represent an improvement over the corresponding binary hypothesis testing and maximum likelihood estimation detectors. However the results indicate that modeling the selected DCT coefficients using the power exponential distribution for locally optimal detection likely does not provide a detector that is suitable for practical situations.

B.2 Watermarking Peppers DCT Coefficients

This section uses the Peppers test image, shown in Figure 6.13 to further evaluate detector performance for the Weibull and power exponential distributions.

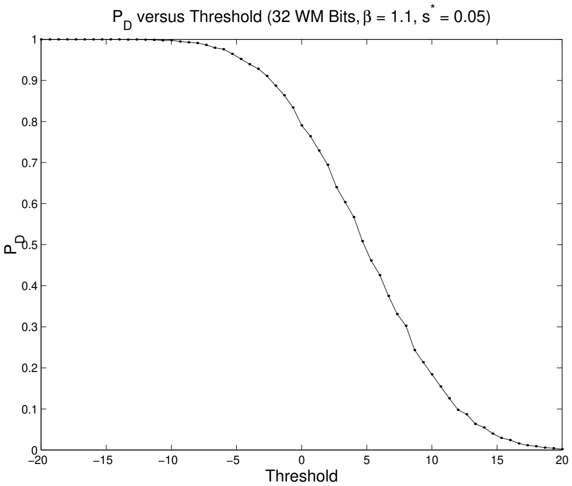

(a)

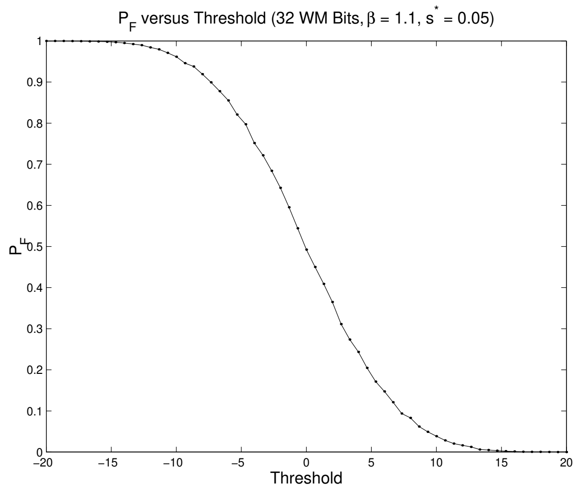

(b)

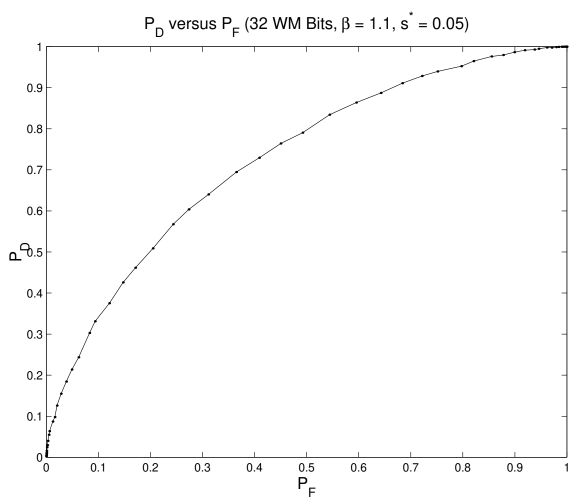

(c)

Figure B.14 P_D and P_F curves for the WB LOD for Lena data ($s^* = 0.05$).

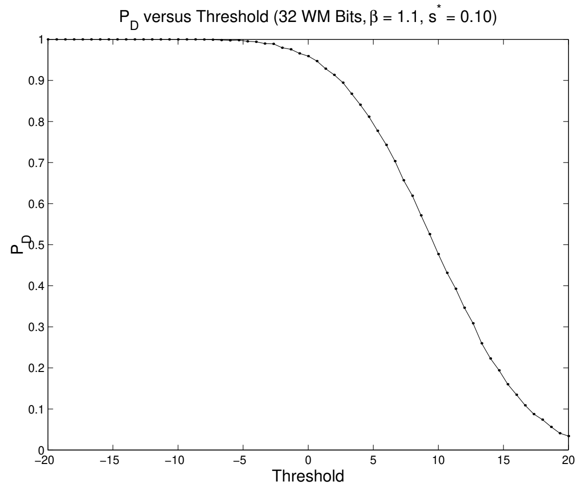

(a)

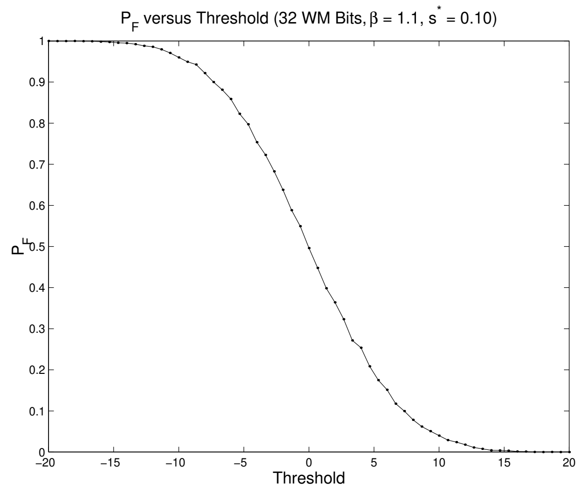

(b)

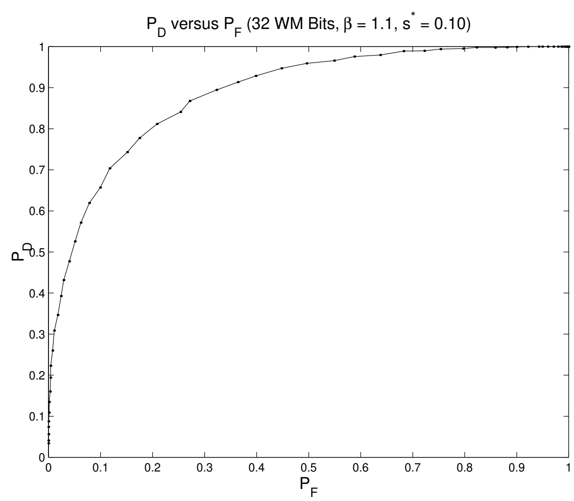

(c)

Figure B.15 P_D and P_F curves for the WB LOD for Lena data ($s^* = 0.10$).

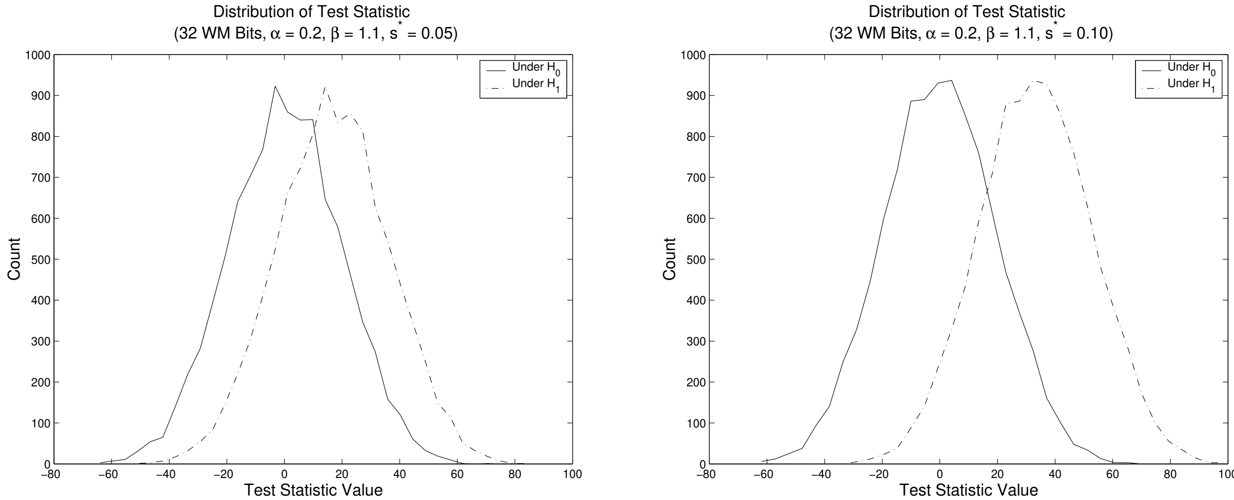

(a) (b)

Figure B.16 Distribution of the test statistic for the Lena image using the PE distribution and LOD with an embedding strength of (a) 0.05 and (b) 0.10.

B.2.1 Simple likelihood ratio test

The first experiments for the Weibull and power exponential distributions for the Peppers image are those using the binary hypothesis testing detectors. Once again, the experimental results are presented in the three standard figures.

B.2.1.1 Weibull distribution results

To begin, the coefficients from the Peppers image used for watermarking are to be modeled using the Weibull distribution. The resulting distributions of the test statistics are depicted in Figure B.19. The separation between the two distributions is quite small for the low-embedding case and only slightly better for the stronger embedding strength.

Figures B.20 and B.21 contain graphs that demonstrate the detector's performance in terms of detection and false alarm probabilities. A comparison of the two figures reveals that an increase in performance occurs when the embedding strength is made larger, as to be expected. However, for the lower embedding strength, a detection probability of approximately 68% is only realized with a false alarm probability of approximately 50%. Thus, it is quite unlikely that this detector will be effective in a practical situation.

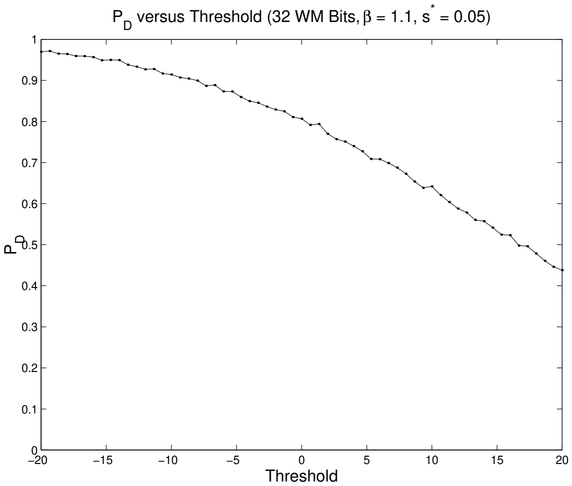

(a)

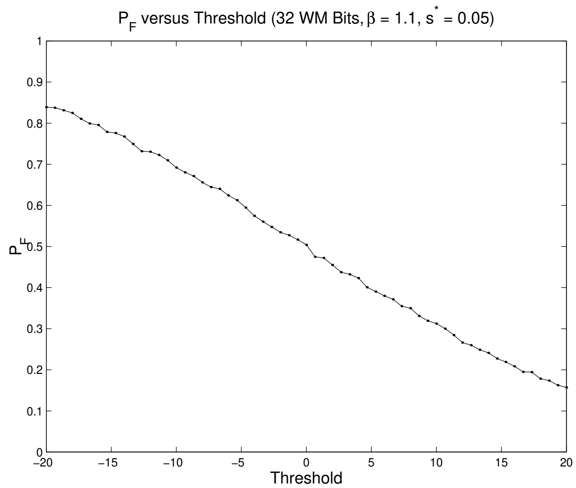

(b)

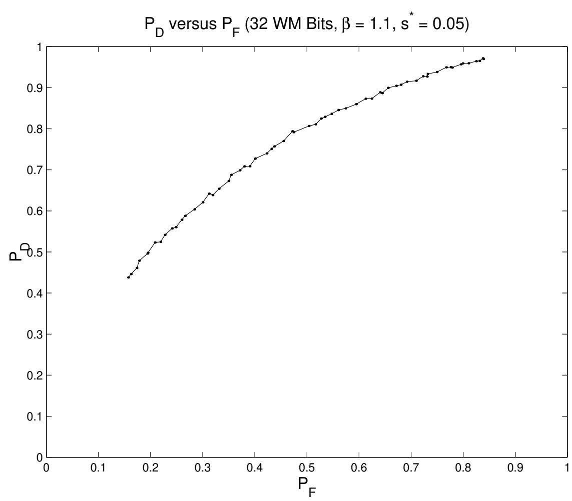

(c)

Figure B.17 P_D and P_F curves for the PE LOD for Lena data ($s^* = 0.05$).

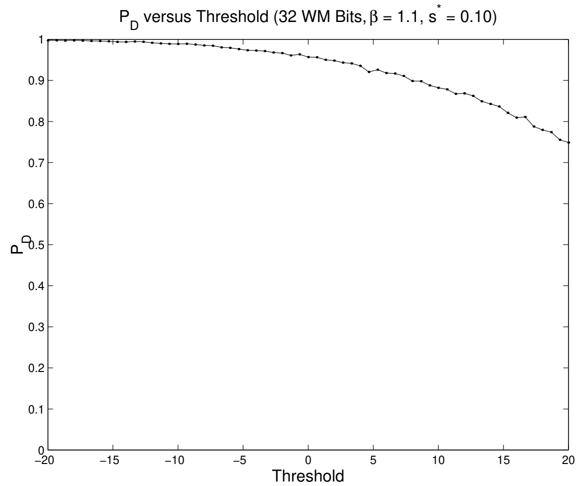

(a)

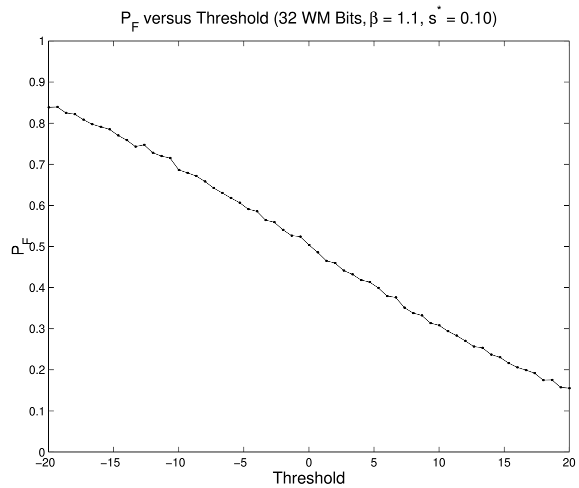

(b)

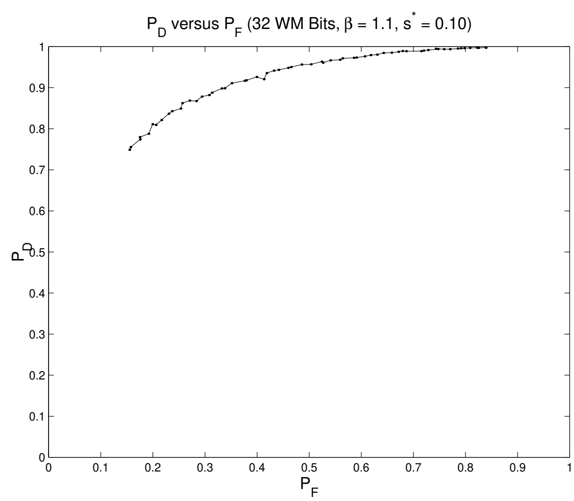

(c)

Figure B.18 P_D and P_F curves for the PE LOD for Lena data ($s^* = 0.10$).

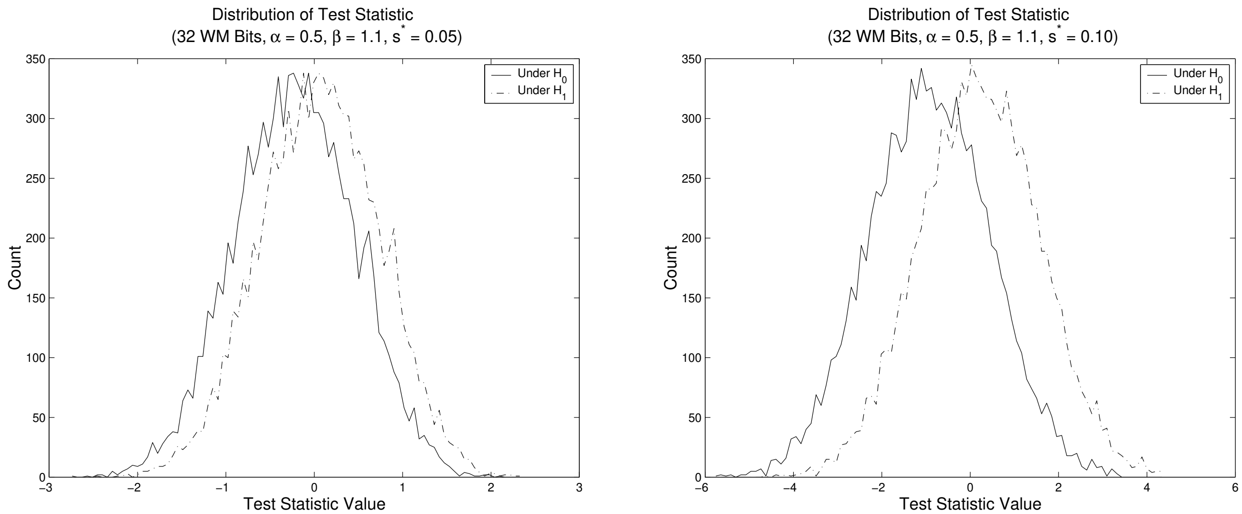

(a) (b)

Figure B.19 Distribution of the test statistic for the Peppers image using the WB distribution and LRT with an embedding strength of (a) 0.05 and (b) 0.10.

B.2.1.2 Power exponential distribution results

The binary hypothesis detector developed for the power exponential distribution is now examined for the DCT coefficients gathered from the Peppers image. The distributions of the test statistic under H_0 and H_1 are given in Figure B.22. As before, the statistics are not well separated, with only a slight increase when the embedding is strengthened.

Figures B.23 and B.24 contain curves of the probabilities P_D and P_F. Although a slight gain is observed as the embedding strength is increased, the detector does not offer high performance. In order to obtain just over a 60% detection probability, a 50% false alarm probability is incurred. These statistics likely render the detector-distribution pair unusable in a practical situation.

B.2.2 Maximum likelihood estimation

Further testing is now performed using the Weibull and power exponential distributions to model the Peppers image coefficients for use with the maximum likelihood estimation detector. As before, the results are illustrated in the typical thee figure formats.

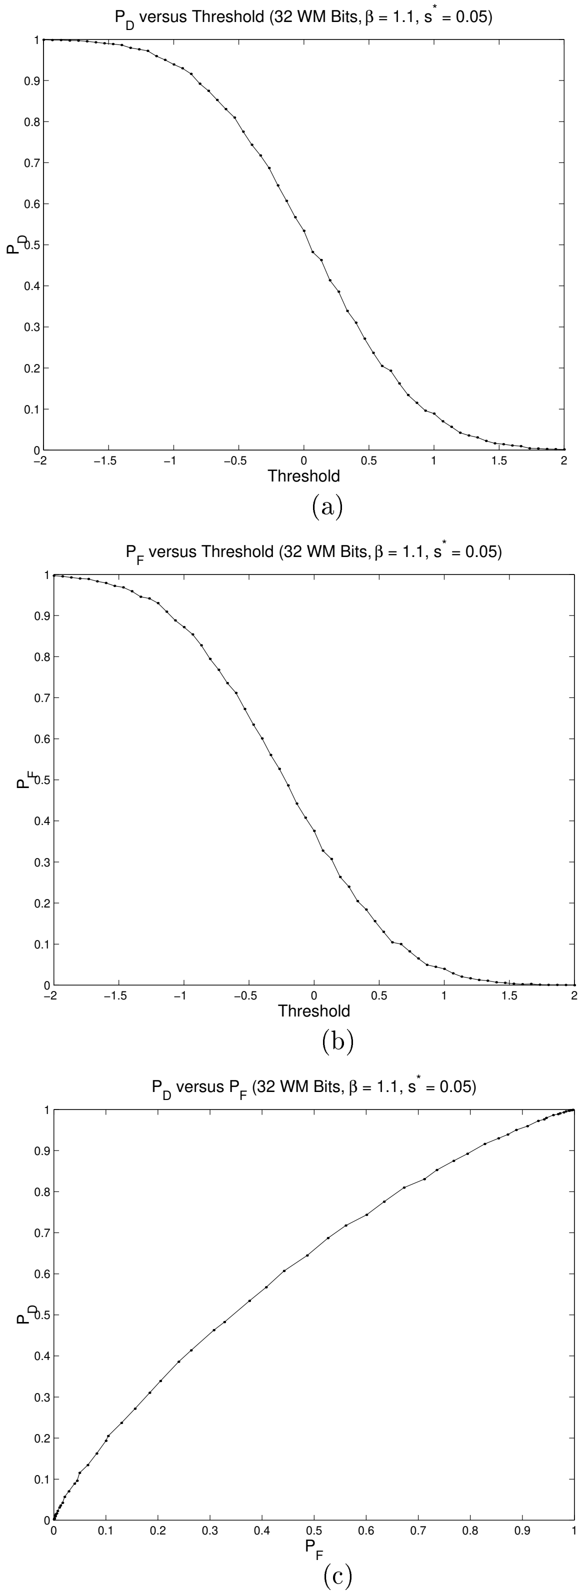

Figure B.20 P_D and P_F curves for the WB LRT for Peppers data ($s^* = 0.05$).

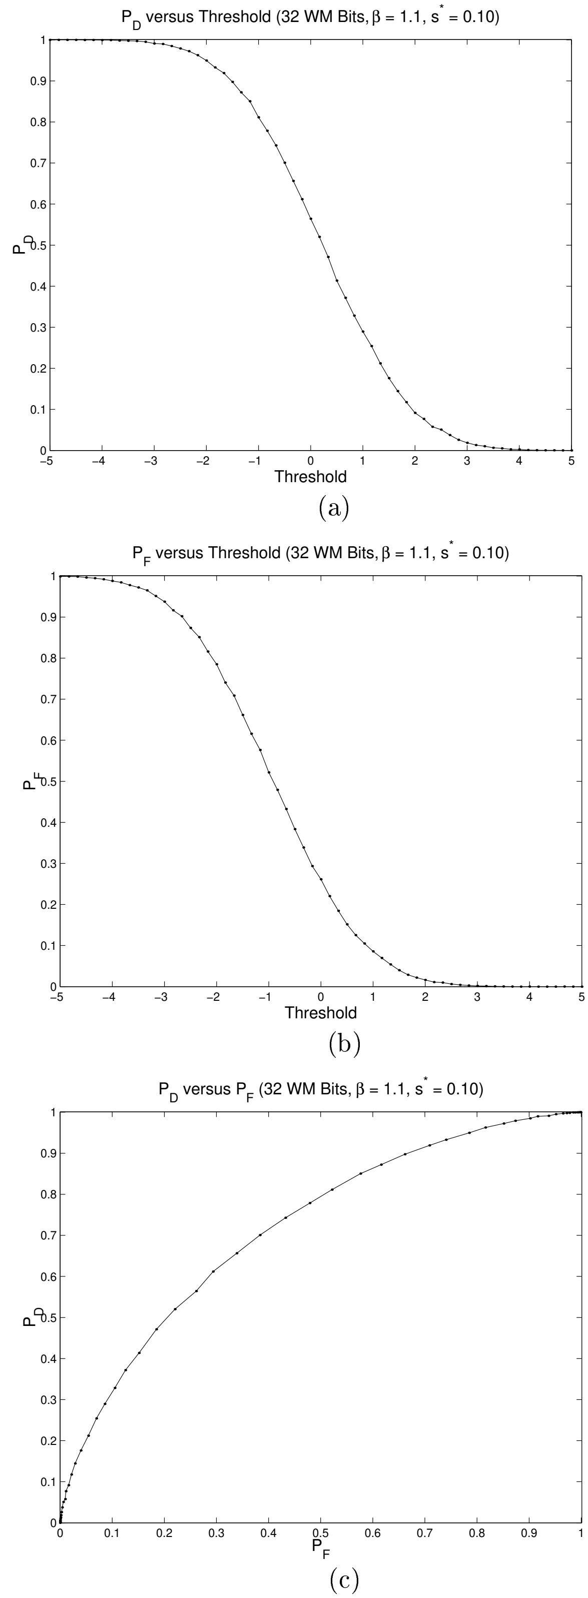

Figure B.21 P_D and P_F curves for the WB LRT for Peppers data ($s^* = 0.10$).

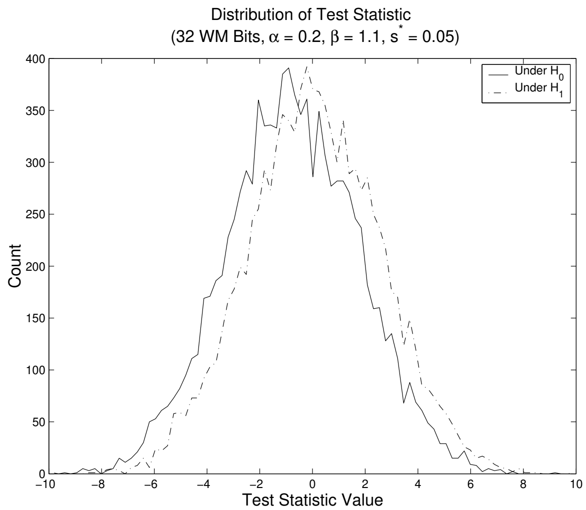

(a)

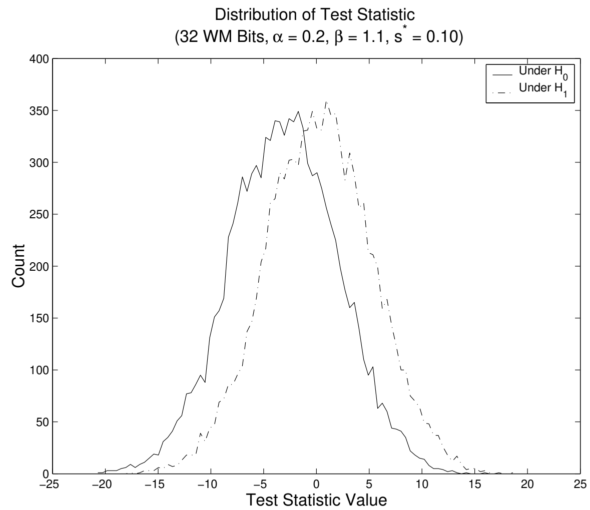

(b)

Figure B.22 Distribution of the test statistic for the Peppers image using the PE distribution and LRT with an embedding strength of (a) 0.05 and (b) 0.10.

B.2.2.1 Weibull distribution results

The first additional distribution investigated for use with MLE is the Weibull distribution. Figure B.25 presents a plot of the distributions of the test statistics for two embedding strengths. The distributions under H_0 and H_1 overlap quite severely, with only a slight increase in separation for the higher embedding strength. However, the estimates for the embedding strength are indeed gathered around the actual value of the strength.

The performance indicators of P_D and P_F are shown in Figures B.26 and B.27. Since the decision statistic takes the form of an estimate, the range of thresholds is limited to [0, 1). Thus, the maximum achievable detection probability is not necessarily equal to 1. The graphs illustrate that, if the threshold is increased much above the actual strength value, the detection probability drops off significantly, as expected. Also, the detector displays better performance for the higher embedding strength. However, the detector is overall quite poor. To achieve a P_D of approximately 62% for the low embedding strength, a P_F of approximately 50% results. Thus, this detector-distribution pair is not a desirable detection solution.

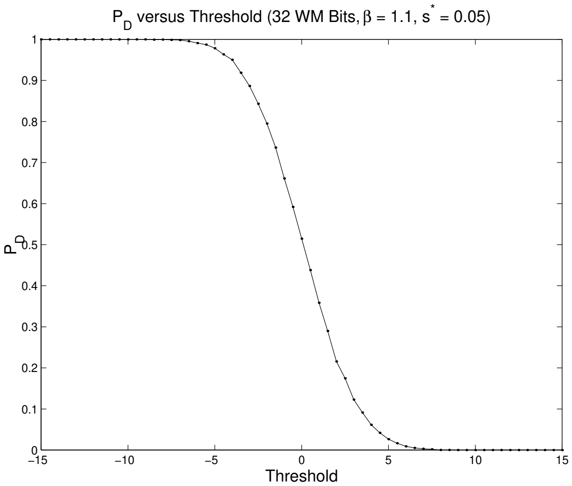

(a)

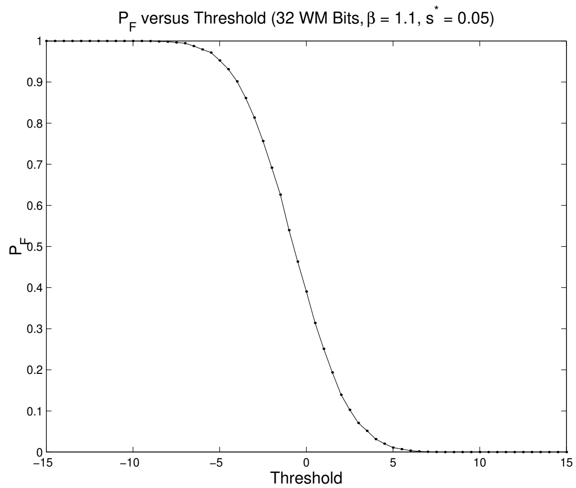

(b)

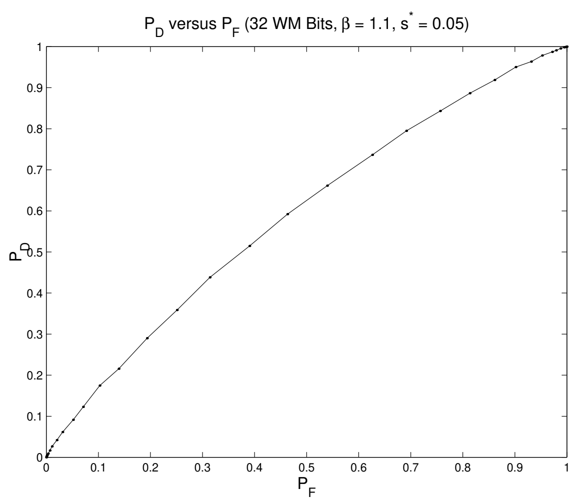

(c)

Figure B.23 P_D and P_F curves for the PE LRT for Peppers data ($s^* = 0.05$).

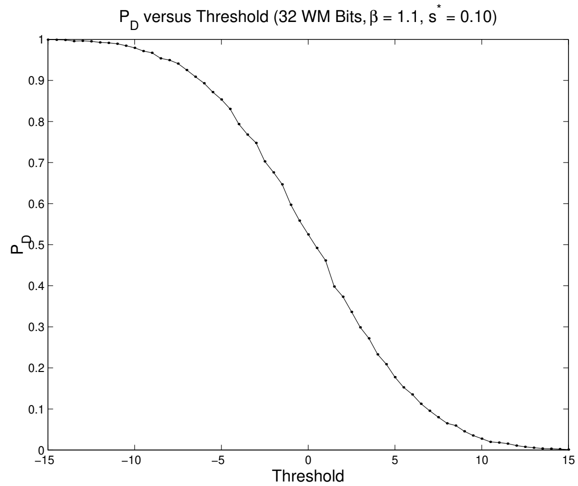

(a)

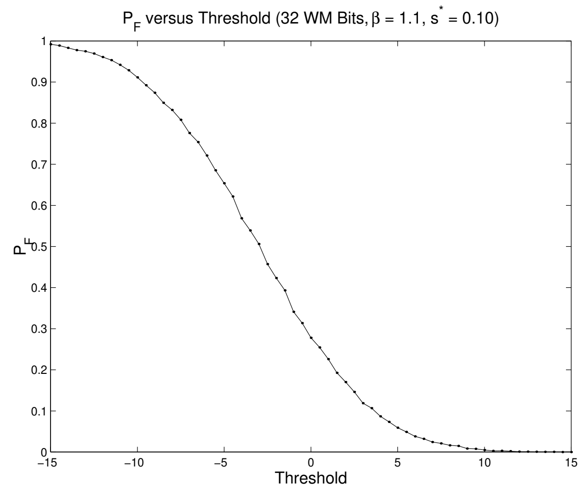

(b)

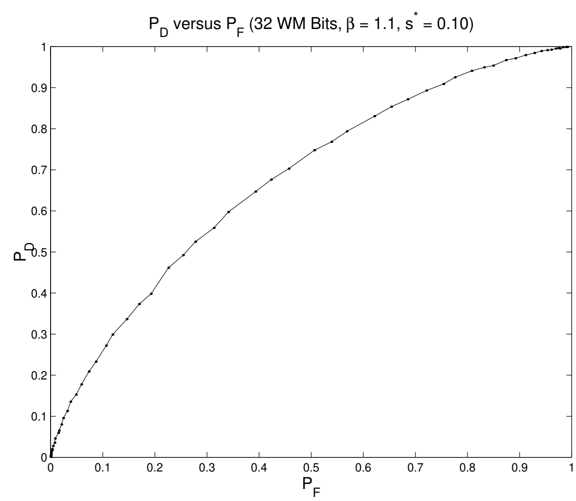

(c)

Figure B.24 P_D and P_F curves for the PE LRT for Peppers data ($s^* = 0.10$).

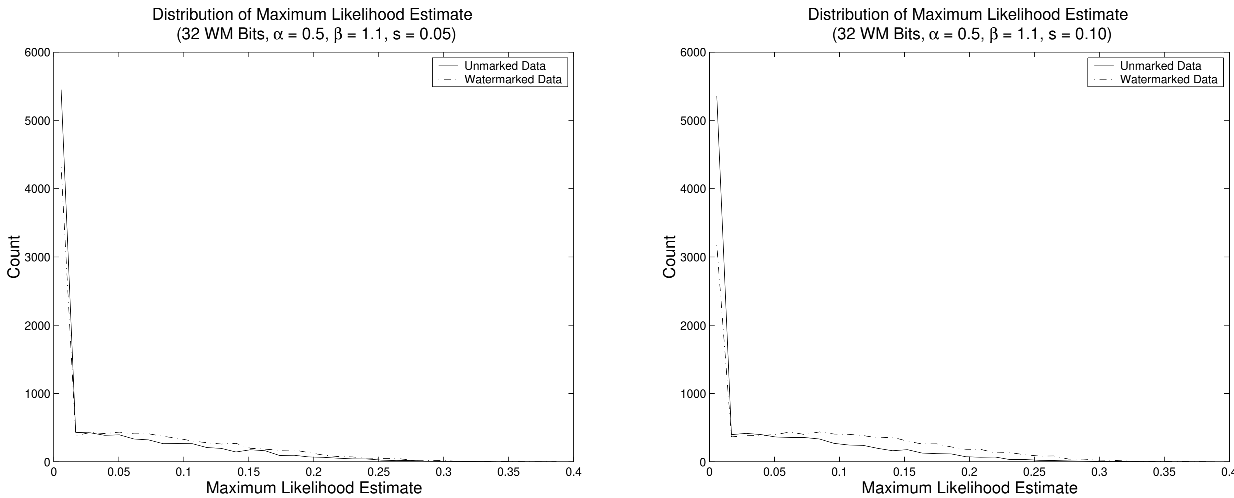

(a) (b)

Figure B.25 Distribution of the test statistic for the Peppers image using the WB distribution and MLE with an embedding strength of (a) 0.05 and (b) 0.10.

B.2.2.2 Power exponential distribution results

Maximum likelihood estimation is now examined for the case when the watermarked Peppers DCT coefficients are modeled using the power exponential distribution. The distributions of the test statistic under the two hypotheses are shown in Figure B.28. These results are quite similar to those for the Weibull distribution MLE: the statistic curves overlap severely with only a small increase in separation as the embedding strength is increased. Although, the estimates do somewhat accurately reflect the true embedding strength.

The next two figures, B.29 and B.30, demonstrate the performance of the detector in terms of detection and false alarm probabilities. These curves exhibit the behavior seen previously for maximum likelihood estimation. The detection probability achieves a maximum of only approximately 60% (at the expense of just under a 50% probability of false alarm) and drops off as the threshold is increased above the actual embedding strength. A slight improvement is present as the embedding strength is increased, but overall the detector does not offer adequate performance to be employed in a realistic setting.

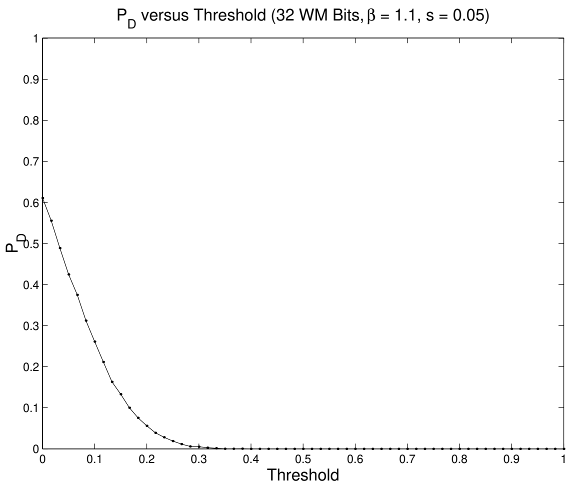

(a)

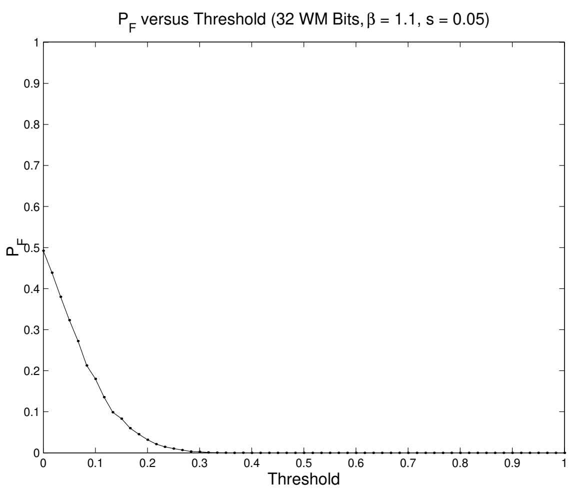

(b)

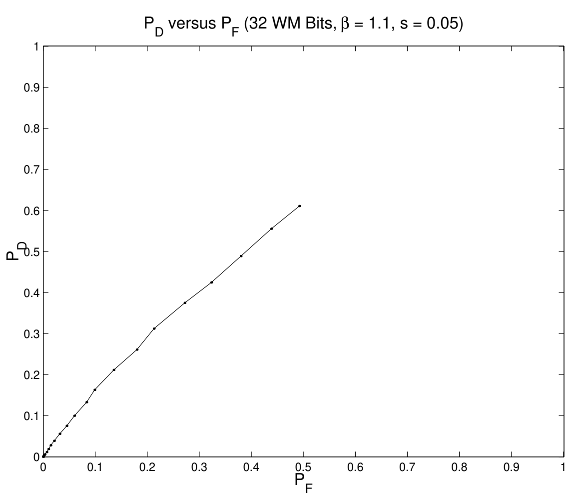

(c)

Figure B.26 P_D and P_F curves for the WB MLE for Peppers data ($s^* = 0.05$).

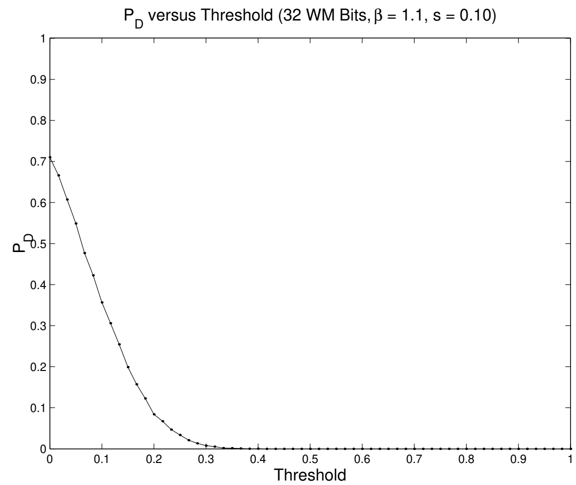

(a)

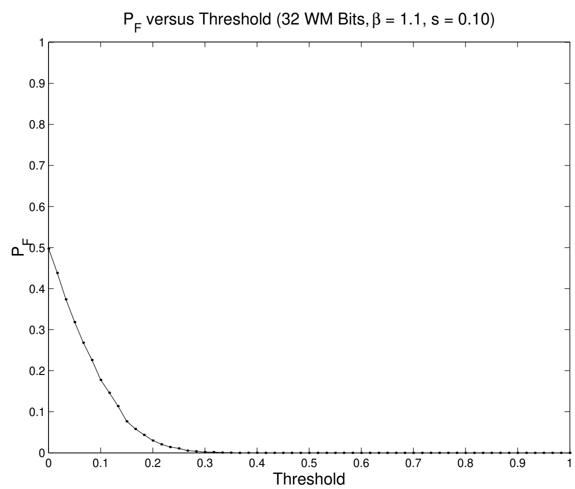

(b)

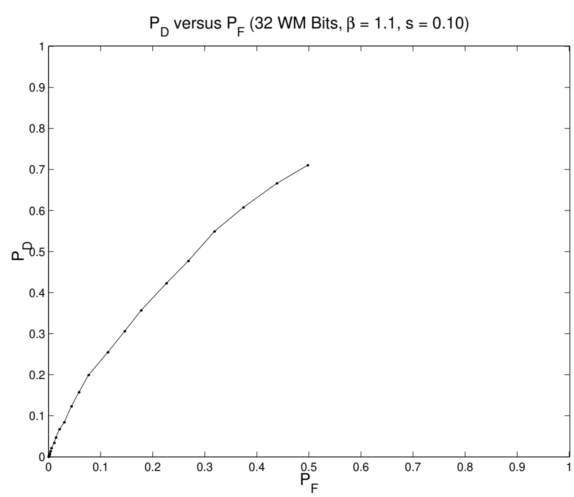

(c)

Figure B.27 P_D and P_F curves for the WB MLE for Peppers data ($s^* = 0.10$).

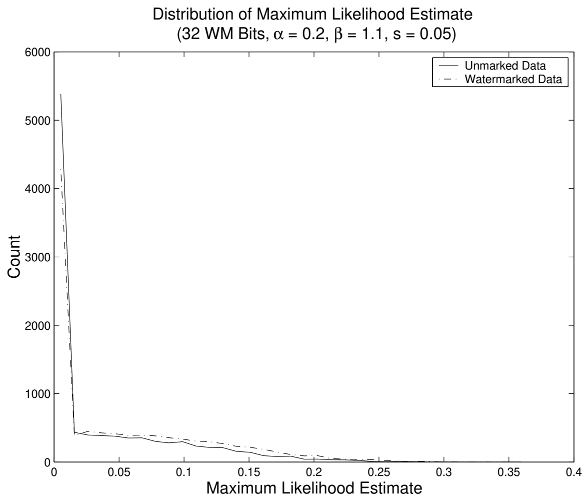

(a)

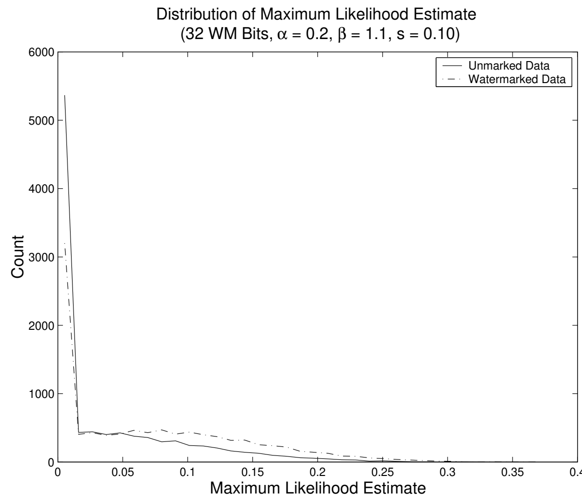

(b)

Figure B.28 Distribution of the test statistic for the Peppers image using the PE distribution and MLE with an embedding strength of (a) 0.05 and (b) 0.10.

B.2.3 Locally optimal detection

The final detector considered for use with the Pepper image and the Weibull and power exponential distributions is the locally optimal detector. Once again, three figures are utilized to illustrate the results.

B.2.3.1 Weibull distribution results

To begin, the locally optimal detector derived though modeling the selected coefficients using the Weibull distribution is considered. Figure B.31 provides the test statistic distributions under H_0 and H_1. For the lower embedding strengths, the two distributions are fairly overlapped: however, when the strength is raised to 0.10, a relatively moderate degree of separation is achieved.

Figures B.32 and B.33 demonstrate the performance of the detector with respect to detection and false alarm probabilities. A noticeable increase in performance is present as the embedding strength is increased from 0.05 to 0.10. However, for the lower strength, a detection probability of approximately 80% can only be attained at the cost of a false alarm probability of approximately 50%. As a result, the detector will not likely be acceptable in practical applications.

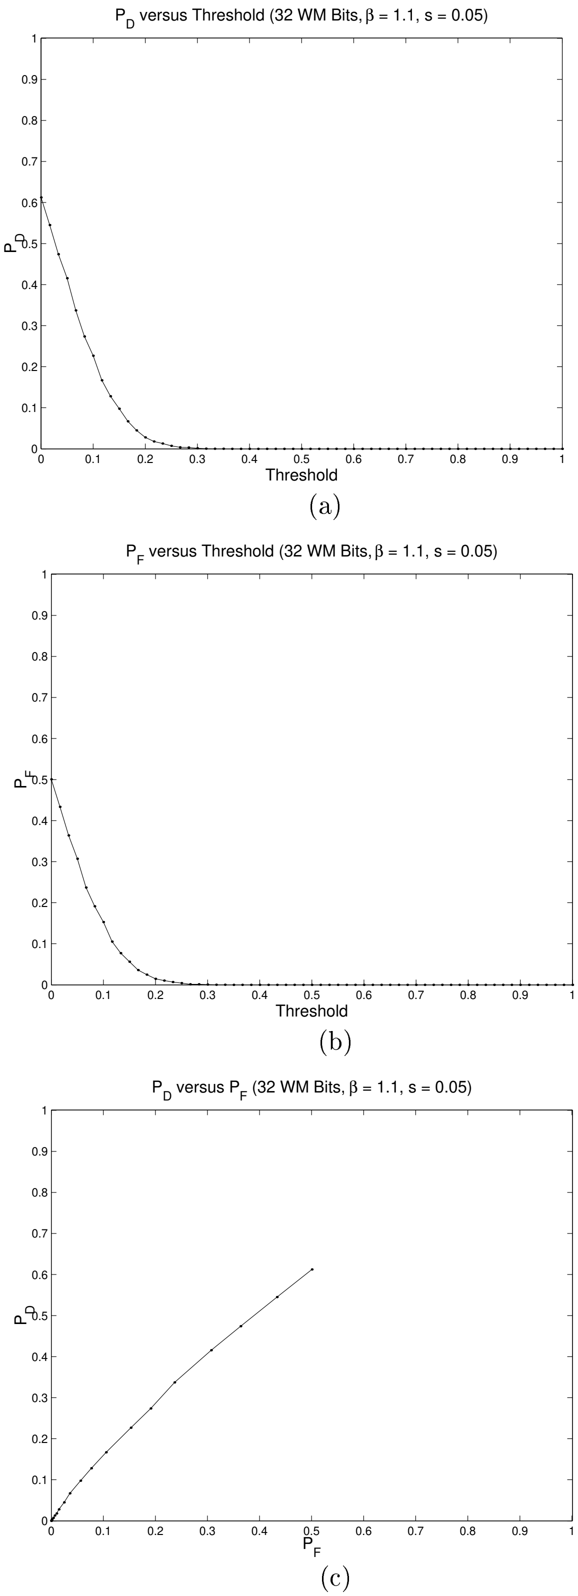

Figure B.29 P_D and P_F curves for the PE MLE for Peppers data ($s^* = 0.05$).

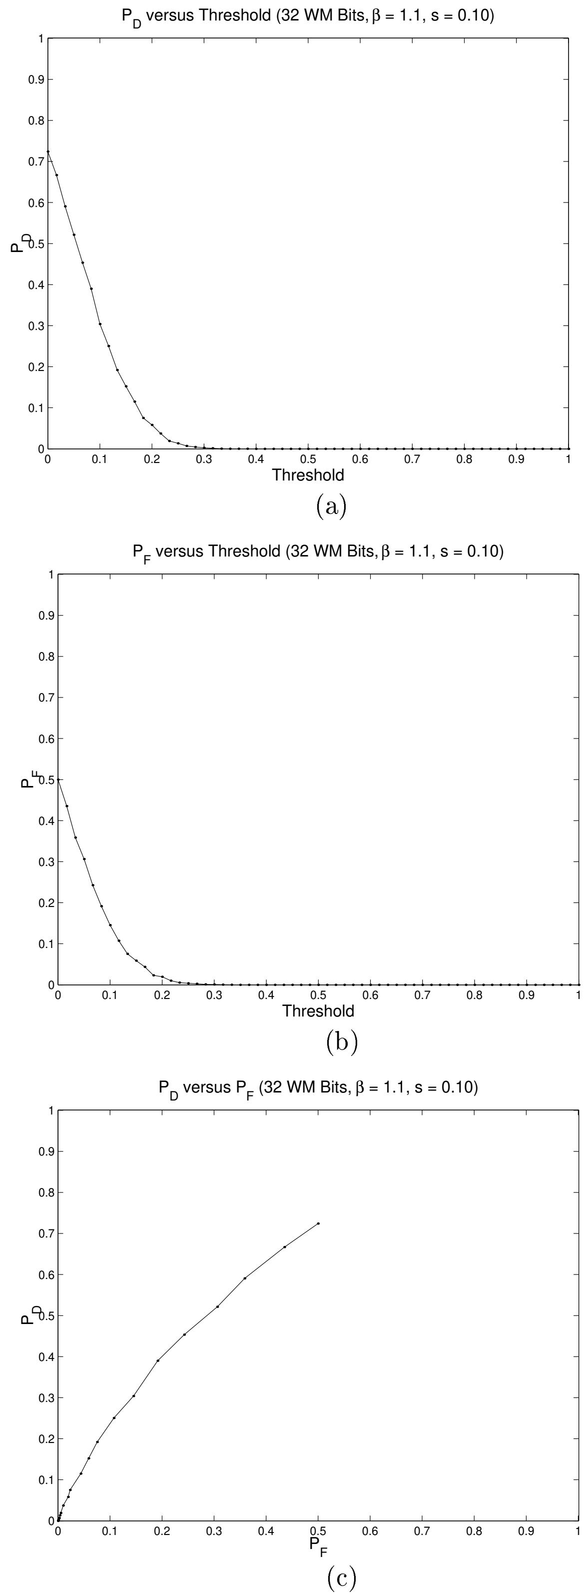

Figure B.30 P_D and P_F curves for the PE MLE for Peppers data ($s^* = 0.10$).

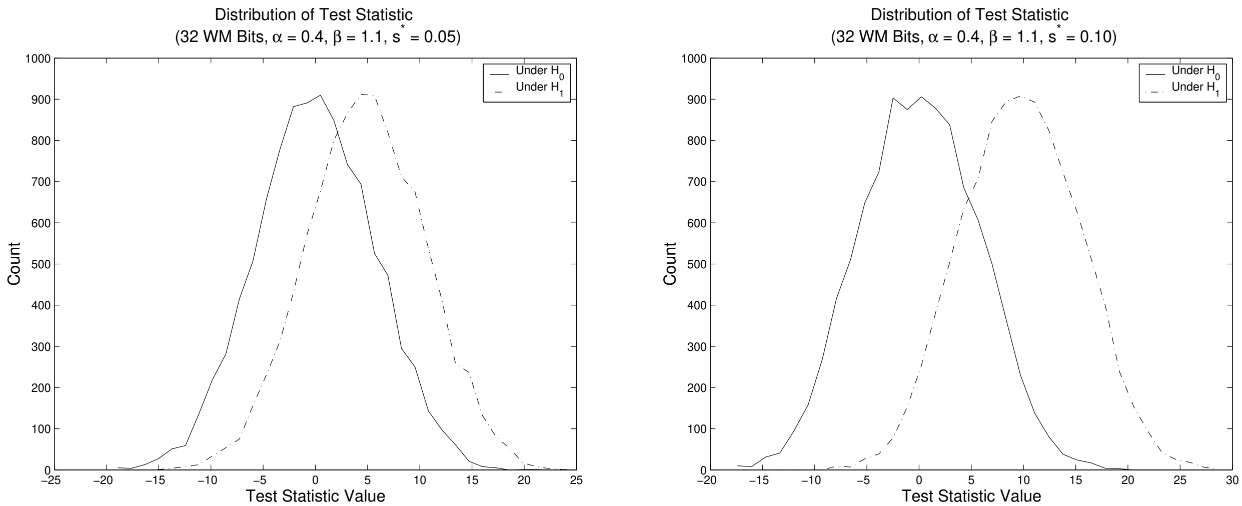

(a) (b)

Figure B.31 Distribution of the test statistic for the Peppers image using the WB distribution and LOD with an embedding strength of (a) 0.05 and (b) 0.10.

B.2.3.2 Power exponential distribution results

The locally optimal detector is now examined for the case when the selected DCT coefficients from the Peppers image are modeled using the power exponential distribution. First, the distributions of the test statistics under the two hypotheses are illustrated in Figure B.34. Similar to the locally optimal detector for the Weibull distribution, the separation between the two distributions of the test statistic increases to a relatively moderate amount as the embedding strength is increased.

The detector's behavior is now considered in the context of the detection and false alarm probabilities. These curves are presented in Figure B.35 and B.36 for the two embedding strengths. When the figures are compared, it is clear that the stronger presence of the watermark aids in its detection. However, the overall performance is still fairly low. For the low embedding, a false alarm probability of approximately 50% is required to yield a detection probability of approximately 80%. Thus, the power exponential locally optimal detector is likely not suitable for most applications.

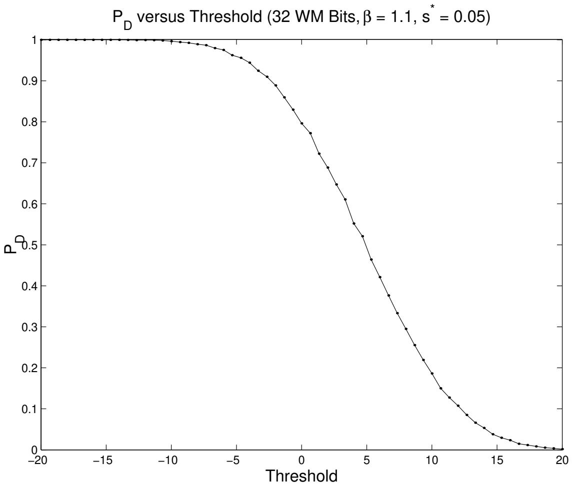

(a)

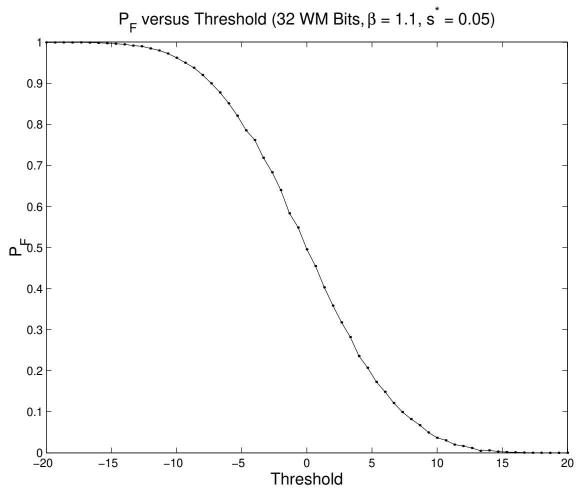

(b)

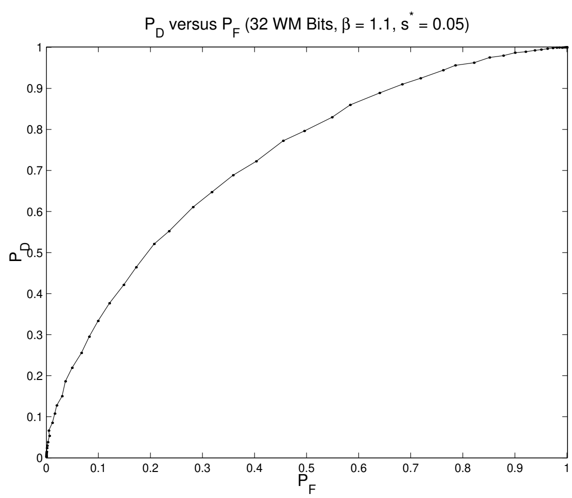

(c)

Figure B.32 P_D and P_F curves for the WB LOD for Peppers data ($s^* = 0.05$).

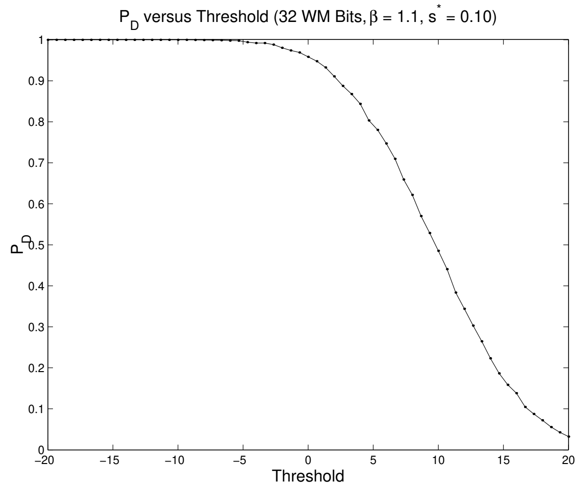

(a)

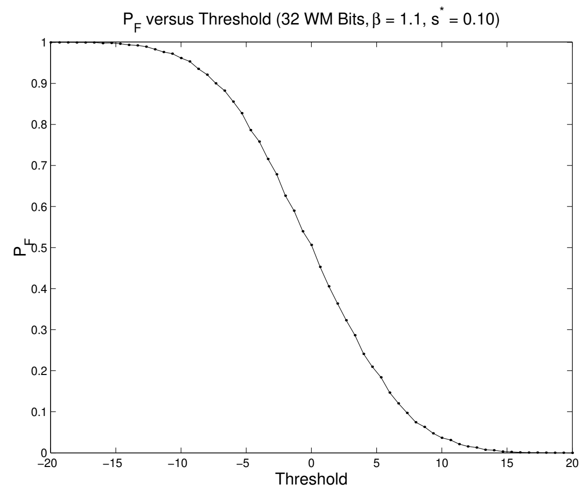

(b)

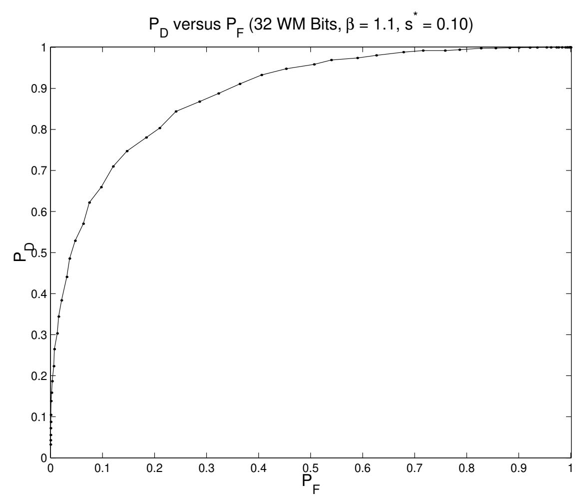

(c)

Figure B.33 P_D and P_F curves for the WB LOD for Peppers data ($s^* = 0.10$).

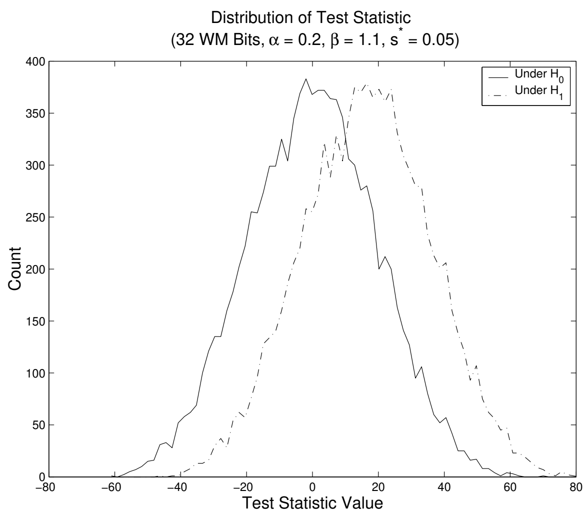

(a)

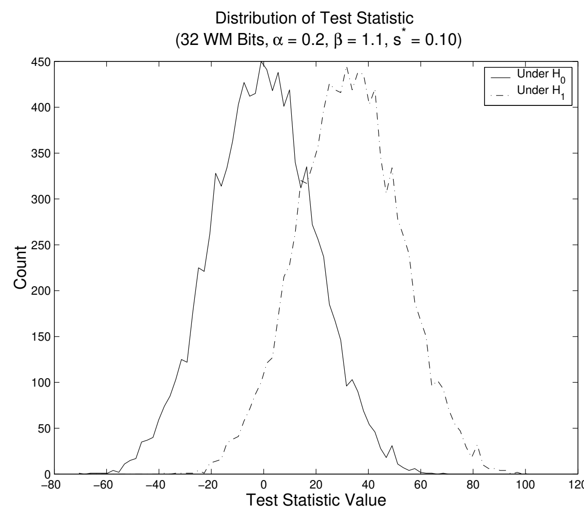

(b)

Figure B.34 Distribution of the test statistic for the Peppers image using the PE distribution and LOD with an embedding strength of (a) 0.05 and (b) 0.10.

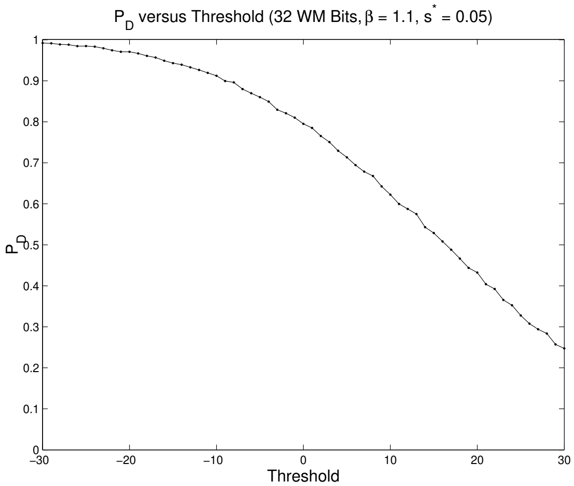

(a)

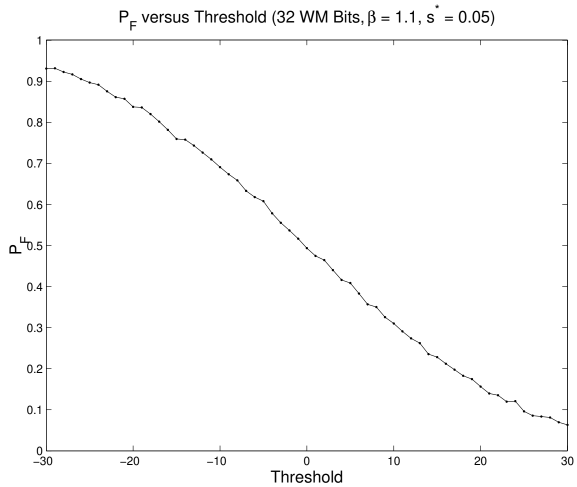

(b)

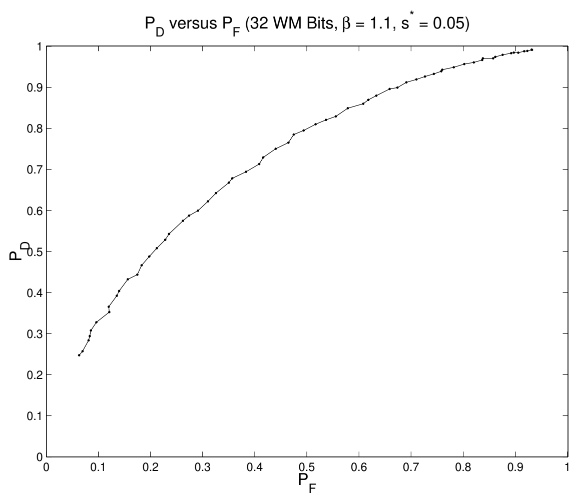

(c)

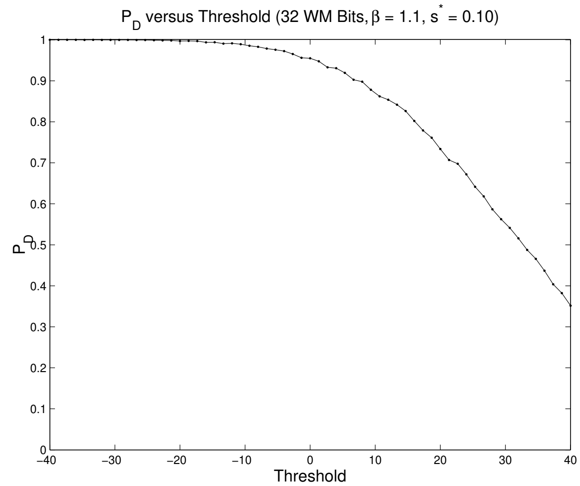

(a)

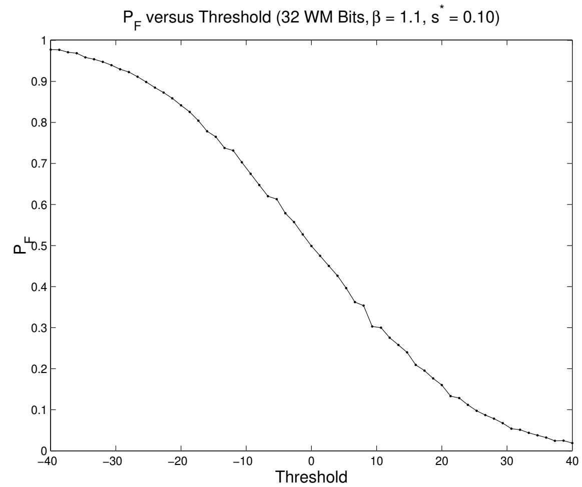

(b)

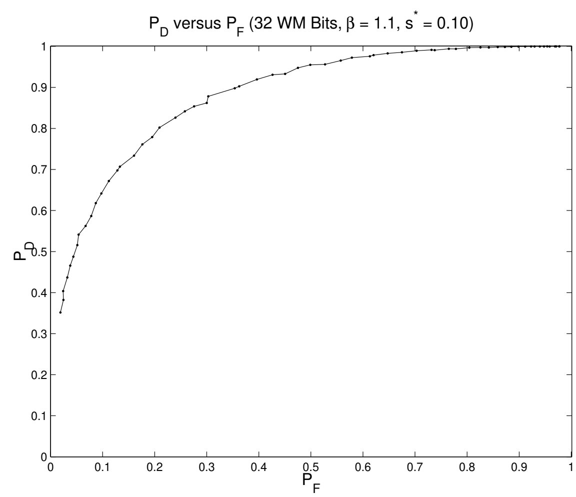

(c)

Figure B.35 P_D and P_F curves for the PE LOD for Peppers data ($s^* = 0.05$).

Figure B.36 P_D and P_F curves for the PE LOD for Peppers data ($s^* = 0.10$).