

Lightweight Detection of Additive Watermarking in the DWT-Domain

Roland Kwitt

Peter Meerwald

Andreas Uhl

Technical Report 2010-04

May 2010

Department of Computer Sciences

Jakob-Haringer-Straße 2
5020 Salzburg
Austria
www.cosy.sbg.ac.at

Technical Report Series

Lightweight Detection of Additive Watermarking in the DWT-Domain

Roland Kwitt, Peter Meerwald, and Andreas Uhl

Abstract

This article aims at lightweight, blind detection of additive spread-spectrum watermarks in the DWT domain. We focus on two host signal noise models and two types of hypothesis tests for watermark detection. As a crucial point of our work we take a closer look at the computational requirements of the detectors. This involves the computation of the detection response, parameter estimation and threshold selection. We show that by switching to approximate host signal parameter estimates or even fixed parameter settings we achieve a remarkable improvement in runtime performance without sacrificing detection performance. Our experimental results on a large number of images confirm the assumption that there is not necessarily a trade-off between computation time and detection performance.

Index Terms

Watermarking, Wavelet, Statistical Signal Detection, Parameter Estimation

I. INTRODUCTION

WATERMARKING has been proposed as a technology to ensure copyright protection by embedding an imperceptible, yet detectable signal in digital multimedia content such as images or video. For blind watermarking, i.e. when detection is performed without reference to the unwatermarked host signal, the host interferes with the watermark signal.

Many detection approaches for additive watermarks embedded in Discrete Cosine Transformation (DCT) or Discrete Wavelet Transformation (DWT) coefficients have been proposed in literature so far [1]–[4]. The perceptual characteristics and distributions of transform domain coefficients have been extensively studied for image compression [5] and these results can be applied to watermarking in order to permit watermark embedding in significant signal components through modeling of human perception. For blind watermarking, the host transform coefficients are considered as noise from the viewpoint of signal detection. If we assume Gaussian noise, it is known that the optimal detector is the straightforward linear-correlation (LC) detector [6]. Unfortunately, DCT and DWT coefficients do not obey a Gaussian law in general, which renders the LC detector suboptimal in these situations and modeling the host signal becomes crucial for detection performance. A first approach, exploiting the fact that DCT or DWT coefficients do not follow a Gaussian law is proposed in [1]. The authors derive an optimal detector for an additive bipolar watermark sequence in DCT transform coefficients following a Generalized Gaussian Distribution (GGD). In [4], it is shown that the low- to mid-frequency DCT coefficients excluding the DC coefficient can also be modeled by the family of symmetric alpha-stable distributions ($S\alpha S$) and a detector is derived for Cauchy distributed DCT coefficients by following the same scheme as is presented in [1]. However, both approaches are based on the strong assumption that the watermark embedding power is known to the detector. In [2], a new watermark detector based on the Rao hypothesis test [7] is proposed for watermark detection in Generalized Gaussian distributed noise. The detector is asymptotically optimal (e.g. for large data records) and does not depend on knowledge about the embedding power any more. In [8] the same scheme is employed to derive a watermark detector in Cauchy distributed noise.

Since detection runtime requirements are important to certain applications we are not only concerned about the detection performance of the watermark detectors but also about their computational behavior. With the objective of lightweight watermark detection in mind, we compare several state-of-the-art detectors from a computational viewpoint. This includes the computation of the detection response, parameter estimation as well as threshold selection. We extend our previous results [9] with large-scale experiments and focus on the issue of host signal parameter estimation which is often neglected in this research area but crucial w.r.t. detection. We show that a considerable runtime improvement can be achieved by switching to approximate or even fixed parameter settings without sacrificing detection performance.

The remainder of this article is structured as follows: In Section II we review two statistical models for DWT coefficients and introduce the detection problem from a hypothesis testing viewpoint. Parameter estimation issues are discussed in Section III. The impact of fast, approximate parameter estimation or even fixed settings on the detection performance is discussed and evaluated in Section IV. We further provide extensive experimental results over a large database of images in Section V and a computational analysis including runtime measurements in Section VI. Finally, Section VII concludes the paper with a discussion of applications and an outlook on further research.

II. STATISTICAL MODELS AND DETECTION PROBLEM

First, we introduce some notation and define the watermark embedding rule. For a J -scale pyramidal DWT we obtain three detail subbands per decomposition level $j \leq J$, denoted by \mathbf{H}_j (horizontal detail subband), \mathbf{V}_j (vertical detail subband) and \mathbf{D}_j (diagonal detail subband). The detail subbands are given in matrix notation. The number of transform coefficients of a detail subband on level j is given as N_j (without loss of generality we assume square images). When it is not necessary to speak of a specific subband we omit the subscript j and the coefficients are given as $x[1], \dots, x[N]$ (vector notation \mathbf{x}). This vector arises by simple concatenation of the row vectors of the appropriate subband coefficient matrix. By adhering to this convention, the elements of the bipolar watermark sequence used for marking an arbitrary subband are denoted by $w[t], 1 \leq t \leq N$ with $w[t] \in \{+1, -1\}$. For the rest of the paper, small boldface letters denote vectors, big boldface vectors denote matrices. The rule for additive embedding of the watermark sequence in the detail subband transform coefficients is given by

$$y[t] = x[t] + \alpha w[t], \quad 1 \leq t \leq N \quad (1)$$

where $\alpha \in \mathbb{R}$ denotes the watermark embedding power and $y[t]$ denotes a watermarked transform coefficient. To derive a hypothesis test, we assume that the transform coefficients $x[t]$ represent an i.i.d. random sample drawn from some underlying probability density function (PDF). In the context of watermarking the coefficients are referred to as host signal noise. The concrete noise model is the key element to derive a watermark detector.

A. Statistical Models for Host Signal Noise

It is commonly accepted that the marginal distributions of the detail subband coefficients of natural images are highly non-Gaussian but can be well modeled by the GGD (see [5], [10]). Employing the parametrization of [11], the PDF of the GGD with scale parameter $a > 0$ and shape parameter $c > 0$ is

$$p(x|a, c) = \frac{c}{2a\Gamma(1/c)} \exp\left(-\left|\frac{x}{a}\right|^c\right) \quad (2)$$

with $-\infty < x < \infty$. In contrast to the Gaussian distribution (which arises as a special case of the GGD for $c = 2$), the GGD is a leptokurtic distribution which allows heavy-tails. A second model which has recently appeared in literature is the two-parameter Cauchy distribution, which is a member of the family of $S\alpha S$ distributions [12]. This model has already been successfully employed for blind DCT-domain watermarking utilizing low- to mid-frequency DCT coefficients [4] and DWT detail subband coefficients [8]. The PDF of the Cauchy distribution with location parameter $-\infty < \delta < \infty$ and shape parameter $\gamma > 0$ is given by [13]

$$p(x|\gamma, \delta) = \frac{1}{\pi} \frac{\gamma}{\gamma^2 + (x - \delta)^2}, \quad (3)$$

with $-\infty < x < \infty$. The Cauchy distribution with $\delta = 0$ (which is symmetric around zero) will be abbreviated by $p(x|\gamma) := p(x|\gamma, 0)$. In contrast to the Gaussian distribution, the tails of the Cauchy distribution decay at a rate slower than exponential, hence we observe heavy-tails in the PDF.

B. Hypothesis Tests for Watermark Detection

In general, we can follow several ways to derive a watermark detector. First, we restate our watermark detection problem as a hypothesis testing problem:

$$\begin{aligned} \mathcal{H}_0 : y[t] &= x[t] & t &= 1, \dots, N \\ \mathcal{H}_1 : y[t] &= x[t] + \alpha w[t] & t &= 1, \dots, N \end{aligned} \quad (4)$$

Here, \mathcal{H}_0 denotes the *null hypothesis* (no or other watermark) and \mathcal{H}_1 denotes the *alternative hypothesis* (watermarked). In case of the Neyman-Pearson (NP) approach to signal detection, it is well-known that the detector which maximizes the probability of detection to a given probability of false-alarm (i.e. deciding \mathcal{H}_1 though \mathcal{H}_0 is true) is the (log) Likelihood-Ratio Test (LRT), given by

$$L(x) := \log\left(\frac{p(x; \mathcal{H}_1)}{p(x; \mathcal{H}_0)}\right) > \log(\tau) =: T, \quad (5)$$

where $p(x|\mathcal{H}_0)$ denotes the PDF under \mathcal{H}_0 , $p(x|\mathcal{H}_1)$ denotes the PDFs under \mathcal{H}_1 and T denotes the detection threshold. Conditioned on the host signal noise models from Section II-A, we obtain the LRT statistics for the Generalized Gaussian model (LRT-GG) [1]

$$\rho(\mathbf{y}) = \frac{1}{a^c} \sum_{t=1}^N (|y[t]|^c - |y[t] - \alpha w[t]|^c) \quad (6)$$

and the Cauchy model (LRT-C) [4]

$$\rho(\mathbf{y}) = \sum_{t=1}^N \log\left(\frac{\gamma^2 + y[t]^2}{\gamma^2 + (y[t] - \alpha w[t])^2}\right). \quad (7)$$

The assumption leading to these detectors is the complete knowledge of the PDFs under both hypothesis. In a realistic case however, the PDFs are incompletely specified (missing knowledge of a, c or γ) and the watermark embedding power α is unknown at the detection stage. This is the general setup for composite hypothesis testing. The unknown noise parameters (denoted by θ) are termed the nuisance parameters since they are not directly related to the detection problem, but affect the shape of the PDFs under \mathcal{H}_0 and \mathcal{H}_1 . On this basis, we formulate a two-sided parameter test:

$$\begin{aligned}\mathcal{H}_0 &: \alpha = 0, \theta \\ \mathcal{H}_1 &: \alpha \neq 0, \theta\end{aligned}\quad (8)$$

In practice, it is assumed that the watermark embedding process does not alter the noise characteristics significantly which allows to estimate θ and α from the watermarked image. The classic approach now is to derive a so called *Generalized Likelihood-Ratio Test (GLRT)*, where the unknown parameters are replaced by their Maximum-Likelihood estimates (MLE). However, ML estimation of α under the alternative hypothesis is quite difficult in case of non-Gaussian noise [14]. To remedy this problem, we can switch to the *Rao* hypothesis test which has the same asymptotic performance as the GLRT but only requires ML estimates under \mathcal{H}_0 . The derivation of a Rao test for the hypothesis formulation of Eq. (8) was first discussed by Kay [15] in case of non-Gaussian noise. Conditioned on the particular host signal noise models we obtain the Rao test statistics for the Generalized Gaussian model (Rao-GG) [2]

$$\rho(\mathbf{y}) = \frac{\left(\sum_{t=1}^N \text{sgn}(y[t])w[t]|y[t]|^c\right)^2}{\sum_{t=1}^N |y[t]|^{2c}} \quad (9)$$

and the Cauchy model (Rao-C) [8]

$$\rho(\mathbf{y}) = \left[\sum_{t=1}^N \frac{y[t]w[t]}{\gamma^2 + y[t]^2} \right]^2 \frac{8\gamma^2}{N}. \quad (10)$$

We remind that in case we assume a Gaussian host signal model, the corresponding LRT, which is a NP detector, is the linear-correlator (LC) with the test statistic given by

$$\rho(\mathbf{y}) = \frac{1}{N} \sum_{t=1}^N y[t]w[t]. \quad (11)$$

Note that for the computation of the Rao and LC detection statistics we do not require knowledge of the embedding strength α at all.

C. Threshold Determination

Determining a reasonable detection threshold T differs depending on the type of hypothesis test and host signal noise model. Further, the computation of the distribution parameters of the detection responses under \mathcal{H}_0 has a considerable impact on the overall runtime performance as we will see in Section VI. We briefly recapitulate the threshold selection process for the detectors of Section II-B to get a notion of the required computational steps. According to [2], [4] the detection statistics Eqs. (6), (7) and (11) follow a Gaussian distribution under both \mathcal{H}_0 and \mathcal{H}_1 . Given that $\mu_{\mathcal{H}_1}$ denotes the mean of the detection statistic under \mathcal{H}_1 and $\mu_{\mathcal{H}_0}$ denotes the mean under the null hypothesis, it can easily be shown that $\mu_{\mathcal{H}_0} = -\mu_{\mathcal{H}_1}$. For the LC detector we have $\mu_{\mathcal{H}_0} = \bar{y}\bar{w}$, $\sigma_{\mathcal{H}_0}^2 = \sigma_y^2 \bar{w}^2 / N$ where \bar{y} and σ_y^2 denote the sample mean and variance of the received signal \mathbf{y} and \bar{w} denotes the sample mean of the watermark (which is ≈ 0 since $w[t] \in \{+1, -1\}$, hence $\mu_{\mathcal{H}_0} \approx 0$ and $\sigma_{\mathcal{H}_0}^2 = 1/N\sigma_y^2$). Given that $y_1[t] := y[t] - \alpha$ and $y_2[t] := y[t] + \alpha$, the location parameter $\mu_{\mathcal{H}_0}$ in case of the LRT-C detector is computed through [4]

$$\mu_{\mathcal{H}_0} = \sum_{t=1}^N \log \left(\frac{\gamma^2 + y[t]^2}{[(\gamma^2 + y_1[t]^2)(\gamma^2 + y_2[t]^2)]^{1/2}} \right) \quad (12)$$

and $\sigma_{\mathcal{H}_0}^2$ is given by

$$\sigma_{\mathcal{H}_0}^2 = \frac{1}{4} \sum_{t=1}^N \left(\frac{\gamma^2 + y_1[t]^2}{\gamma^2 + y_2[t]^2} \right)^2. \quad (13)$$

The LRT-GG detection distribution parameters [1] are

$$\mu_{\mathcal{H}_0} = \sum_{t=1}^N a^c |y[t]|^c - \frac{1}{2} \sum_{t=1}^N c^a (|y_1[t]|^c - |y_2[t]|^c) \quad (14)$$

and

$$\sigma_{\mathcal{H}_0}^2 = \frac{1}{4} \sum_{t=1}^N a^{2c} (|y_2[t]|^c - |y_1[t]|^c)^2. \quad (15)$$

For all LRT detectors, the relationship of the variance under \mathcal{H}_0 and \mathcal{H}_1 is $\sigma_{\mathcal{H}_0}^2 = \sigma_{\mathcal{H}_1}^2$. The probability of false alarm (denoted by P_f) can now be formulated as $P_f = \mathbb{P}(\rho(\mathbf{y}) > T | \mathcal{H}_0) = 1/2 \operatorname{erfc}((T - \mu_{\mathcal{H}_0})/\sqrt{2}\sigma)$ which allows to set T according to the NP criterion $T = \sqrt{2}\sigma \operatorname{erfc}^{-1}(2P_f) + \mu_{\mathcal{H}_0}$. Considering the Rao hypothesis tests, the situation is somewhat different. As already noted, the Rao test exhibits the same behavior as the GLRT for large sample sizes (i.e. asymptotically). According to [6], test statistics Eq. (9) and Eq. (10) follow a χ_1^2 distribution with one degree of freedom under \mathcal{H}_0 . Exploiting the relation $Q_{\chi_1^2}(x) = 2Q(\sqrt{x})$ where $Q(\cdot)$ denotes the Q-function of the Gaussian distribution and $Q_{\chi_1^2}(\cdot)$ denotes the Q-function of the χ_1^2 distribution, we can write $P_f = \mathbb{P}(\rho(\mathbf{y}) > T | \mathcal{H}_0) = Q_{\chi_1^2}(T) = \operatorname{erfc}((T/2)^{1/2})$ and set the desired threshold for a given P_f as $T = 2(\operatorname{erfc}^{-1}(P_f))^2$. Under the alternative hypothesis both Rao test statistics follow a non-central Chi-Square distribution $\chi_{1,\lambda}^2$ with one degree of freedom and *non-centrality* parameter λ , see [2] and [8], respectively. In Section IV we will see how to estimate λ from the experimental detection responses to analyze the performance.

III. PARAMETER ESTIMATION

We discuss the parameter estimation issues related to the Generalized Gaussian and Cauchy host signal noise model using Maximum-Likelihood and approximative solutions.

A. Generalized Gaussian Distribution

GGD parameters estimation has been extensively covered in literature, we only provide a brief overview of the main results. Maximum Likelihood estimation is studied in the work of Varanasi et al. [16] including both, joint parameter estimation, and situations where one parameter is already known. Do and Vetterli provide a Newton-Raphson algorithm to find the root of the resulting transcendental equation in [17] which involves computation of the Digamma and Trigamma function. We refer to their algorithm for any computational discussion. The starting value for the Newton-Raphson iteration is usually obtained using the moment estimate of β , presented by Birney et al. [5] and Mallat [10]. However, even moment matching is challenging since it requires a numerical solution to the function inversion problem $F^{-1}(m_1^2/m_2)$ with

$$F(c) := \frac{\Gamma^2(2/c)}{\Gamma(1/c)\Gamma(3/c)} \quad (16)$$

and

$$m_1 = \frac{1}{N} \sum_{t=1}^N |x[t]| \quad \text{and} \quad m_2 = \frac{1}{N} \sum_{t=1}^N x[t]^2. \quad (17)$$

This problem can either be solved using the combination of a lookup-table and some sort of interpolation method (see [17]), or by employing the approximation of Krupinski [18] where the author proposes to define an invertible approximation $R(c) = \exp(k + lc^m)$ to $F(c)$ and solves a non-linear curve fitting problem for certain ranges of c . In this work, we do not split the range of c and obtain $k = -0.2667$, $l = -0.4172$ and $m = -1.1585$ as the corresponding coefficients (using the MATLAB Curve Fitting Toolbox). Moment matching then reduces to the simple function evaluation $R^{-1}(c)$ which has a closed-form expression. It is obvious that from a computational viewpoint we favor the invertible approximation to $F(c)$. However, the question arises whether the detection process actually benefits from an accurate estimate of c , or if it is possible to set c to a fixed value. This idea is mentioned by Hernandez et al. [1], however it is not further discussed there. In Section V we will see that fixed settings actually lead to quite good detection results on a large set of images. As a last point we note that there exists another recently proposed method to estimate the GGD parameters [19] which is based on exploiting a convex shape equation. However, this method again requires to solve a numerical root finding problem, which is why we omit any further discussion here.

B. Cauchy distribution

Regarding the estimation of the Cauchy distribution parameter γ , we start with the ML estimation approach. Given that $x[1], \dots, x[N]$ denote realizations of N i.i.d. random variables following a Cauchy distribution with $\delta = 0$, the ML estimate of γ is given as the root of (see [13])

$$\frac{1}{N} \sum_{t=1}^N \frac{2}{1 + (x[t]/\gamma)^2} - 1 =: h(\gamma) \quad (18)$$

which has to be solved numerically. For Newton-Raphson root-finding, we require $h'(\gamma)$. The first derivative of $h(\gamma)$ is given as

$$h'(\gamma) := \frac{\partial h}{\partial \gamma} = \frac{4\gamma}{N} \sum_{t=1}^N \frac{x[t]^2}{(\gamma^2 + x[t]^2)^2} \quad (19)$$

which leads to $\hat{\gamma}_{n+1} = \hat{\gamma}_n - h(\hat{\gamma}_n)/h'(\hat{\gamma}_n)$ as the update step in the n -th ($n > 1$) iteration. As starting value $\hat{\gamma}_1$ we use an estimation based on the method of sample quantiles x_p [13],

$$\hat{\gamma}_1 = 0.5(x_p - x_{1-p}) \tan(\pi(1-p)), \quad (20)$$

Algorithm 1 Calculate detection responses $\rho_{i,j,k}$

Require: Images $\mathcal{I}_1, \dots, \mathcal{I}_L$; Parameters $\theta_1, \dots, \theta_B$

for all \mathcal{I} in $\{\mathcal{I}_1, \dots, \mathcal{I}_L\}$, $1 \leq i \leq L$ **do**

for $1 \leq j \leq W$ **do**

$\mathcal{W} \leftarrow \text{generateWatermark}()$

$\mathcal{E} \leftarrow \text{embed}(\mathcal{I}, \mathcal{W})$

for all θ in $\{\theta_1, \dots, \theta_B\}$, $1 \leq k \leq B$ **do**

$\rho_{i,j,k} \leftarrow \text{detect}(\mathcal{E}, \mathcal{W}, \theta)$

end for

end for

end for

with $0.5 < p < 1$. A computationally fast algorithm for approximate biased quantile estimation is presented by Zhang and Wang [20], since quantile estimation usually requires to sort the data first. In this work, we use another fast algorithm originally presented by Trihrintzis and Nikias [21] in the context of S α S parameter estimation. Given that we assume $\delta = 0$ and $\alpha = 1$ (here α denotes the S α S parameter) the proposed estimator for γ is given by¹

$$\hat{\gamma} = \left[\frac{\sqrt{3}}{2N} \sum_{t=1}^N |x[t]|^{1/3} \right]^3 \quad (21)$$

which can be computed with linear complexity. Last, we note that moment estimation is not possible in case of the Cauchy distribution, since the moments do not exist. As with the approximate GGD shape parameter c , it is of course possible to use Eq. (21) or Eq. (20) alone to compute the detection statistics for the detectors conditioned on the Cauchy distribution. In our experiments we use Eq. (21) and further evaluate the possibility of setting γ to a fixed value.

IV. PARAMETER ESTIMATION AND DETECTION PERFORMANCE

In this section we study the impact of the different parameter estimation approaches and fixed parameter settings on the detection performance by plotting the probability of missing the watermark as a function of the Cauchy γ and GGD c shape parameter, an idea first presented in [1]. We decompose the images using a two-scale DWT with biorthogonal CDF 9/7 filters and select the \mathbf{H}_2 subband for watermark embedding. For each image we choose α such that a DWR of 16 dB is obtained. The resulting PSNR values are listed in Table I for three different DWR settings. To obtain an empirical estimate of the probability of miss we first fix the desired probability of false-alarm. We will use $P_f = 10^{-6}$ throughout this work. Next, we calculate the detection statistic parameters under \mathcal{H}_0 and determine a detection threshold according to Section II-C. We then estimate the detection statistic parameters $\lambda, \mu_{\mathcal{H}_1}$ and $\sigma_{\mathcal{H}_1}^2$ from the detection responses under \mathcal{H}_1 when embedding and detecting $W = 1000$ random watermarks. The algorithm to obtain the detection responses for each image and shape parameter is listed in Algorithm 1. We have set $\theta_1, \dots, \theta_B$ representative for the range of GGD and Cauchy shape parameters.

In case of the LRT detectors, we can simply use the sample mean and sample variance to obtain estimates for $\mu_{\mathcal{H}_1, i, k}$ and $\sigma_{\mathcal{H}_1, i, k}^2$ for each image and shape parameter. The empirical probability of miss is then calculated from $\hat{P}_{m, i, k} = 1 - Q((T_{i, k} - \hat{\mu}_{\mathcal{H}_1, i, k}) / \hat{\sigma}_{i, k})$. In case of the Rao detectors, we exploit the fact that a non-central Chi-Square random variable with one degree of freedom and non-centrality parameter λ is equivalent to the square of a Gaussian random variable with mean $\sqrt{\lambda}$ and unit variance. Hence, we can estimate $\lambda_{i, k}$ from the responses of the i -th image using shape parameter θ_k as follows:

$$\hat{\lambda}_{i, k} = \left(\frac{1}{W} \sum_{j=1}^W \sqrt{\rho_{i, j, k}} \right)^2 \quad (22)$$

The empirical probability of missing the watermark is given by $\hat{P}_{m, i, k} = 1 - Q(\sqrt{T_{i, k}} - \sqrt{\hat{\lambda}_{i, k}}) + Q(\sqrt{T_{i, k}} + \sqrt{\hat{\lambda}_{i, k}})$. Fig. 2 shows a comparison of the LRT-C and Rao-C detector performance as a function of the γ parameter for the six test images of Fig. 1 for $P_f = 10^{-6}$. The parameter γ varies from 0.1 to 45 with a step size of 0.1. The performance of the detectors w.r.t. to the MLE and its fast approximation (i.e. quantile estimation) are marked by a *circle* 'o' and *diamond* '◇', respectively. We observe that, except for *Bridge* and *Peppers*, the MLE and approximate γ are almost visually indistinguishable. In general, the approximation is very close to the ML estimate for all images and the difference has a negligible impact on watermark detection performance. Further, it is evident that neither the MLE nor the approximation leads to the minimum probability of miss. Nevertheless, the parameter estimates are close enough to achieve good results. Fig. 3 shows the same type of plot for the LRT-GG and Rao-GG detector. The GGD shape parameter c varies from 0.02 to 4 with a step size of 0.01. Again, the

¹As proposed in [21], we set the fractional moment p to $1/3$ which leads to this simplified expression.



Fig. 1. Grayscale test images (from top left to bottom right: *Lena*, *Barbara*, *Elaine*, *Boat*, *Bridge*, *Peppers*)

TABLE I
DOCUMENT-TO-WATERMARK RATIO (DWR) IN DB AND RESULTING PSNR (DB) FOR THE TEST IMAGES

DWR (dB)	PSNR (dB)					
	<i>Lena</i>	<i>Elaine</i>	<i>Barbara</i>	<i>Boat</i>	<i>Bridge</i>	<i>Peppers</i>
16	47.40	50.66	47.70	46.50	46.83	47.30
20	50.86	52.96	51.03	50.19	50.45	50.80
23	52.38	57.52	52.69	51.77	51.95	52.30

fast estimate (i.e. invertible approximation to $F(c)$) is reasonably close to the ML estimate. In case of the LRT-GG detector, peak performance is missed by only a small margin. However, in case of the Rao-GG detector, the estimates of c are far off the parameter value achieving the minimum probability of miss.

Based on these experiments, we draw the first conclusion that fast, approximate estimation of the shape parameter for both noise models can actually replace the ML estimate without significantly sacrificing detection performance. A second conclusion is that the parameter estimation approaches are in many cases not equivalent to searching the parameter value minimizing the probability of miss (or equivalently maximizing the probability of detection). This has already been noted earlier for the LRT-GG detector by Hernandez et al. [1]. For the Rao-GG detector the estimates do not achieve good detection performance at all for many images. In Nikolaidis' work [2], the detection results for the Rao-GG detector vary widely across the test images and the detector performs even worse than the LC detector. We presume that the reason for the bad results has its roots in the behavior we observe in our plots here.

Regarding the use of fixed parameter values, an image independent GGD shape parameter $c = 0.8$ has already been proposed for DWT coefficients [1] to save the estimation effort. Our comparative plots in Fig. 3 allow to immediately read off the resulting performance and compare with the performance of the LC detector (i.e. $c = 2$). In case of the Rao-GG detector, the shape of the probability of miss curve indicates a parameter setting of $c \approx 1$ as a reasonable choice. For the Cauchy host signal model, no fixed γ parameter value has been proposed to the best of our knowledge. A good candidate for an image independent parameter seems to be $\gamma = 8$ for the DWT details subbands. In the next section we present large-scale experimental results which confirm our observations.

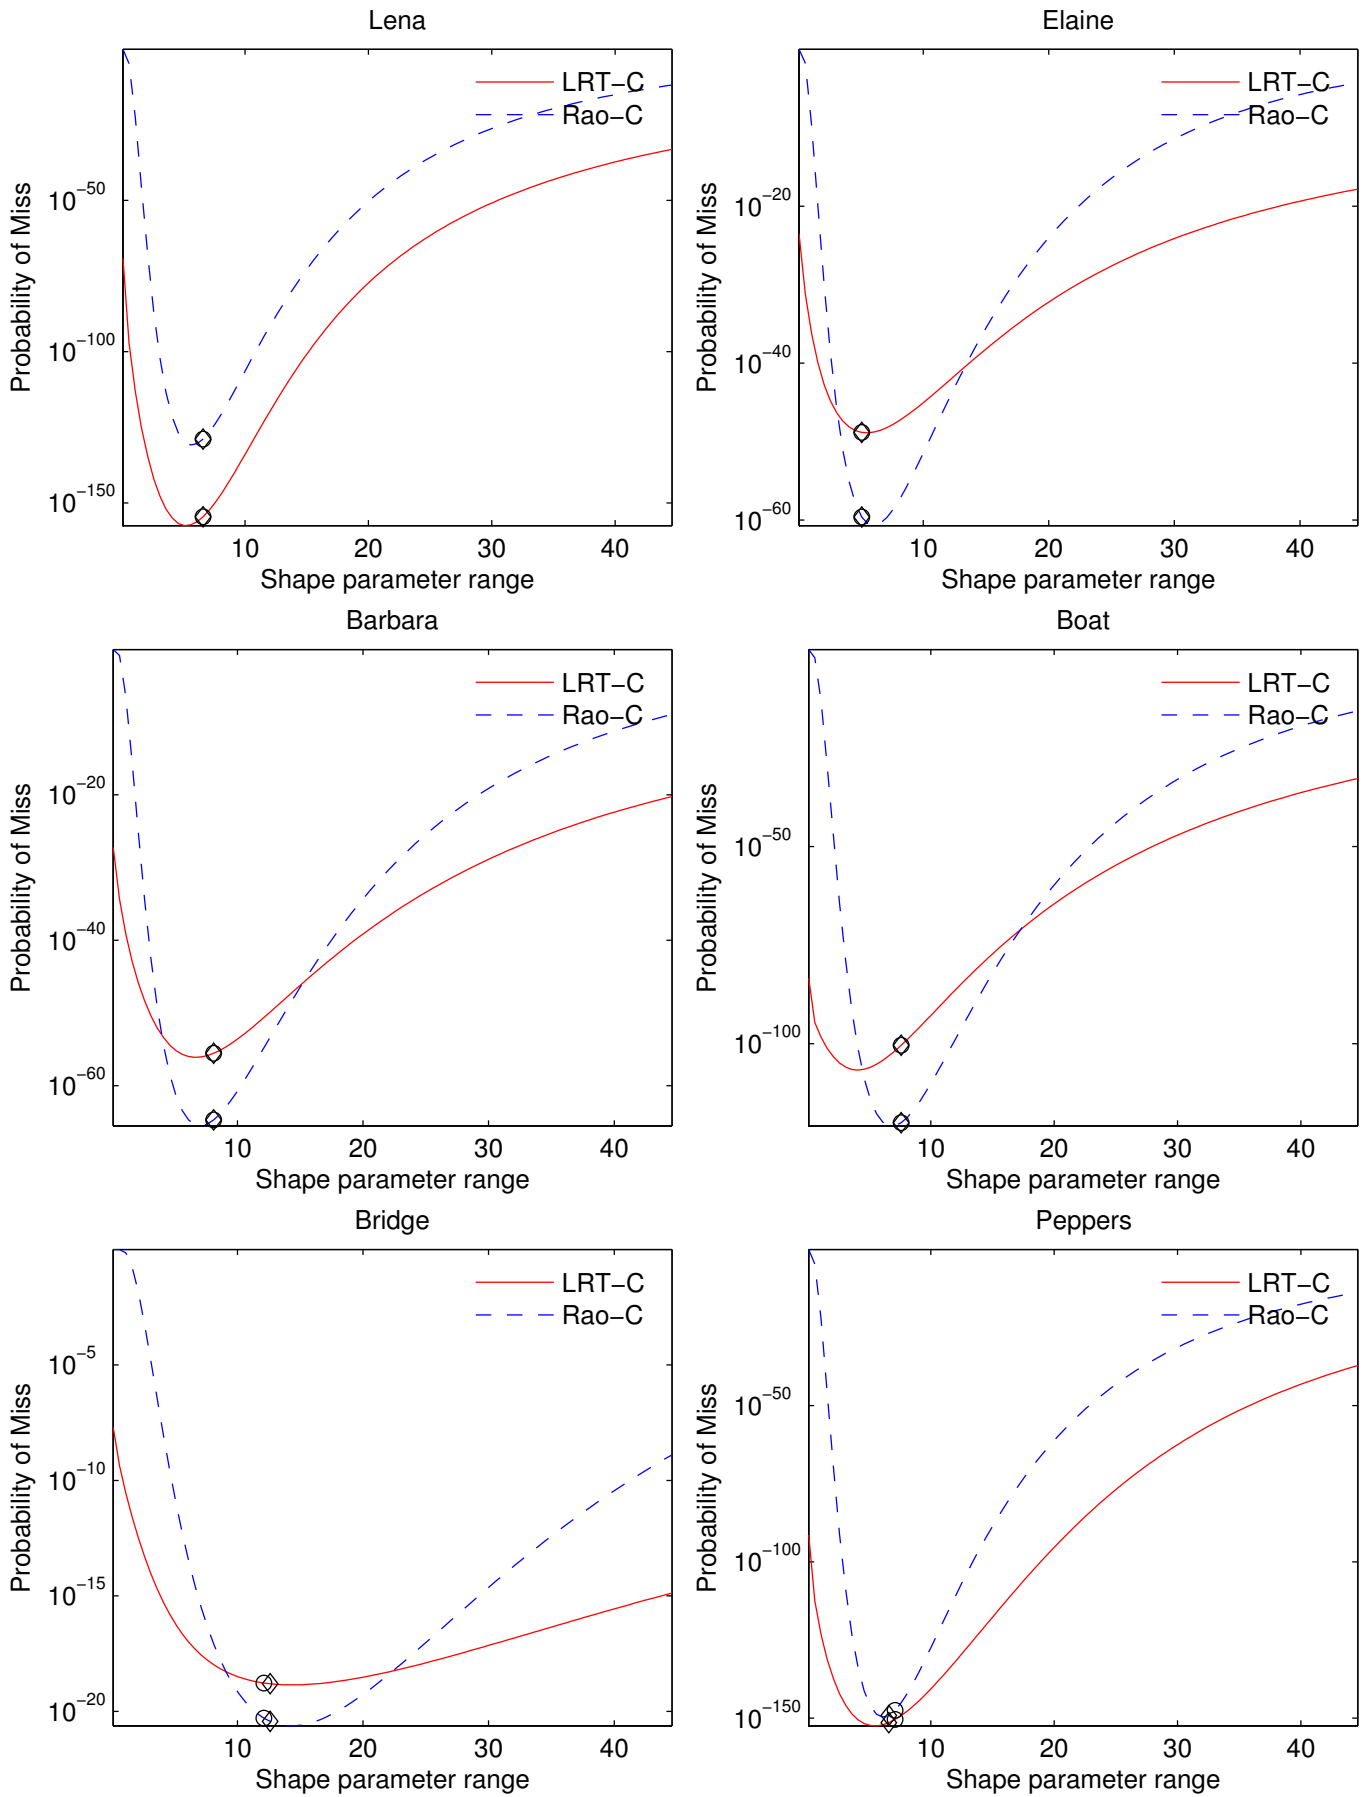


Fig. 2. Probability of miss comparison of the LRT-C and Rao-C detector as a function of γ at DWR 16 dB and $P_f = 10^{-6}$.

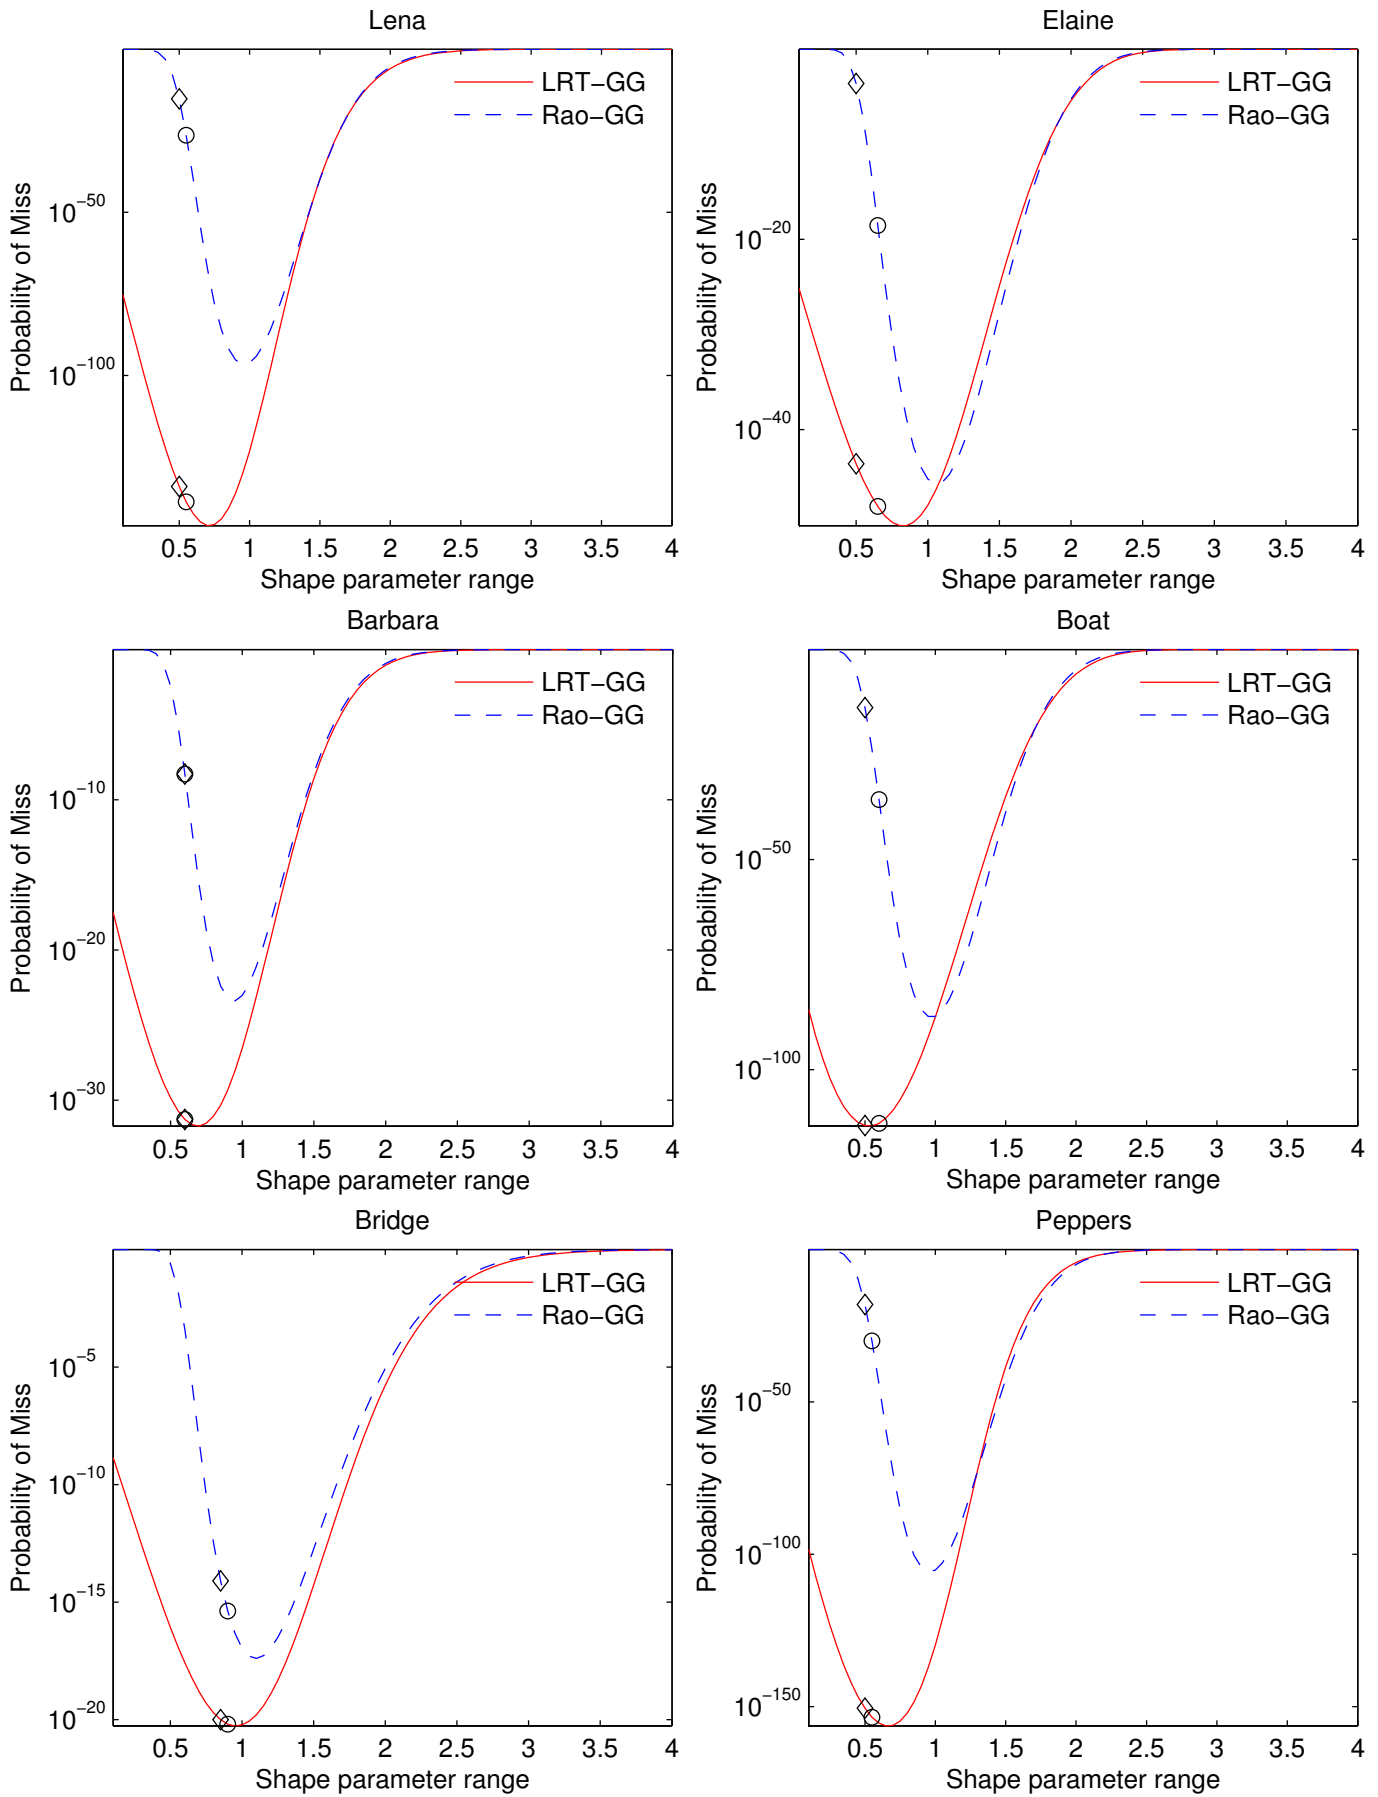


Fig. 3. Probability of miss comparison of the LRT-GG and Rao-GG detector as a function of γ at DWR 16 dB and $P_f = 10^{-6}$.

Algorithm 2 Calculate optimum $\hat{P}_{m,i}^*$

Require: $\hat{P}_{m,i,k}$, $1 \leq i \leq L$, $1 \leq k \leq B$
for $1 \leq i \leq L$ **do**
 $s_i \leftarrow \arg \min_k \hat{P}_{m,i,k}$
 $\hat{P}_{m,i}^* \leftarrow \hat{P}_{m,i,s_i}$, $\theta_i^* \leftarrow \theta_{i,s_i}$
end for

Algorithm 3 Calculate sub-optimal $\hat{P}_{m,i}^s$

Require: Empirical $\hat{P}_{m,i,k}$, Target parameters t_1, \dots, t_L ,
for $1 \leq i \leq L$ **do**
 $r_i \leftarrow \arg \min_{k \in \{1, \dots, B\}} |t_i - \theta_k|$
 $\hat{P}_{m,i}^s \leftarrow \hat{P}_{m,i,r_i}$
end for

V. LARGE-SCALE TESTS

In this section we evaluate the detection performance of the detectors as a function of the parameter setting over a large set of test images. We choose $L = 1000$ images from the UCID [22] color image database for our experiments. The images were first converted to grayscale and scaled down to 50% of the original size (i.e. 768×512). To compare the performance of the detectors we choose the probability of miss as the evaluation criterion. The procedure to obtain the empirical probability of miss for each detector and shape parameter was explained in the last section.

In order to get an impression of the impact of the shape parameters c and γ , we first determine the optimal probability of miss $\hat{P}_{m,i}^*$ and the corresponding optimal parameter value θ_i^* for each image according to Algorithm 2. Next, since our aim is to compare the optimal parameter setting to the MLE, approximate and fixed settings, we calculate the corresponding (sub-optimal) probability of miss $\hat{P}_{m,i}^s$ according to Algorithm 3. What we basically do is to search the parameter θ_i , $1 \leq i \leq B$ that is closest (in l_1 sense) to either the ML estimate or the approximate parameter estimate of each image, denoted by t_i .

To obtain comparative plots, we choose the following approach: first, we sort the optimal probability of miss values in ascending order to obtain $\hat{P}_{m,v_1}^* \leq \dots \leq \hat{P}_{m,v_L}^*$ where v_1, \dots, v_L denotes the index set of sorted probabilities. Second, we arrange the sub-optimal probabilities using this index set to obtain $\hat{P}_{m,v_1}^s, \dots, \hat{P}_{m,v_L}^s$. Third, we plot both vectors of probabilities (optimal and sub-optimal) on a semi-logarithmic scale. Figs. 4–7 show the corresponding plots for the different detectors and a selection of shape parameter choices. Note that we have limited the axis scaling to show only the last 600 images since probabilities below 10^{-100} are less interesting in practice.

For the LRT-GG detector we observe that the approximate estimate of c as well as the fixed parameter setting of $c = 0.8$ do not lead to great deviations from the optimal curve. Especially in the interesting top-right hand area of the plots the probability of miss is almost at the optimum. A fixed value of $c = 1.2$, which is between a Laplace and a Gaussian distribution, shows the expected worst performance. These results are in accordance with what we have seen in Fig. 3. In case of the Rao-GG detector (see Fig. 5) both MLE and the approximation lead to bad performance. The best results are obtained from the fixed settings $c = 0.8$ and $c = 1$. The good results for $c = 1$ confirm our presumption of the last section (see Fig. 3) where we expected $c \approx 1$ to be a good choice. Regarding the LRT-C detector, shown in Fig. 6, there is no noteworthy difference between the fixed settings of $\gamma = 3$, $\gamma = 8$ and the MLE as well as the approximation. However, since we favor good performance especially for P_m values close to 1, we would reject $\gamma = 3$ due to slight deviations from the optimum in the top-right hand area of the plot. Last, the Rao-C detector shows good performance for both MLE and the fast approximation. While the fixed setting of $\gamma = 8$ deviates only slightly from the optimum, the setting $\gamma = 3$ sacrifices performance and turns out to be impractical.

We can conclude that for all four detectors, we can find a fixed parameter setting which leads to good detector performance over the whole set of natural images and only slightly deviates from the optimum. This fact is remarkable since it allows to completely eliminate the computationally expensive estimation step of the detection process.

The image received by the watermark detector might have been subjected to image processing operations. In Fig. 8 we test JPEG compression (with compression factor $Q = 70$) applied on the watermarked image. We observe the performance of the LRT-GG and Rao-GG detector in the left and right column, resp. As expected, the probability of missing the watermark increases for both detectors due to the added noise. The fixed parameter settings still perform well and the approximate parameter estimates closely match the MLE approach in both cases. In Fig. 9 we turn to the detectors based on the Cauchy host signal model and observe that $\gamma = 8$ is still a reasonable fixed host signal parameter setting, even under moderate JPEG compression. Comparing the LRT and Rao detectors, however, we notice that the performance of the Rao-C and Rao-GG detector increase relative to the LRT methods and the Cauchy shape parameter estimates are closer to their optimal value than the GGD parameter estimates. We conjecture that the Cauchy model might be more robust and a better fit for the wavelet coefficients of a JPEG-compressed image.




Fig. 4. LRT-GG probability of miss comparison over all images for various choices of c at $P_f = 10^{-6}$ and DWR 16 dB.

VI. COMPUTATIONAL ANALYSIS

In order to quantify the term lightweight detection, we take a closer look at the computational effort which is required to employ the various detectors. This includes a discussion of the number of required arithmetic operations to calculate the detection statistics, parameter estimation issues and the determination of detection thresholds. By arithmetic operations we understand the number of additions & subtractions ($+$, $-$), multiplications & divisions (\times , \div), logarithms & exponentiations (\log, x^r) with floating point numbers as well as absolute & signum ($|\cdot|$, sgn) operations.

In Table II we provide the number of operations as a function of the input vector length N (constant terms are omitted). From these numbers it is obvious that the LC detector is by far the simplest in terms of arithmetic operations, since it involves only summations and multiplications of floating point numbers. Only the watermarked coefficients and the watermark sequence itself are involved. However, the Rao-C detector is only slightly more expensive, since Eq. (10) merely involves exponentiations with integer exponents (can be formulated as multiplications), additions as well as multiplications which can be very efficiently performed with few CPU cycles. In contrast to that, the LRT-C, LRT-GG and Rao-GG detector require computations of the logarithm or exponentiations with floating point numbers, which are both expensive in terms of CPU cycles.

Regarding parameter estimation effort, we have seen that fixed parameter settings and approximate solutions lead to quite good results. However, for the sake of completeness we take a look at the computational effort required to compute the MLE. The LC detector defined in Eq. (11) is the simplest one, since it requires no parameter estimation at all, followed by the Rao-C and LRT-C detector, which both require an estimate of the shape parameter γ of the Cauchy distribution. The Rao-GG and LRT-GG detectors are most affected by the estimation procedure, since the Newton-Raphson update steps in case of the GGD are more expensive than the update steps in case of the Cauchy distribution. We also note that there is almost no difference

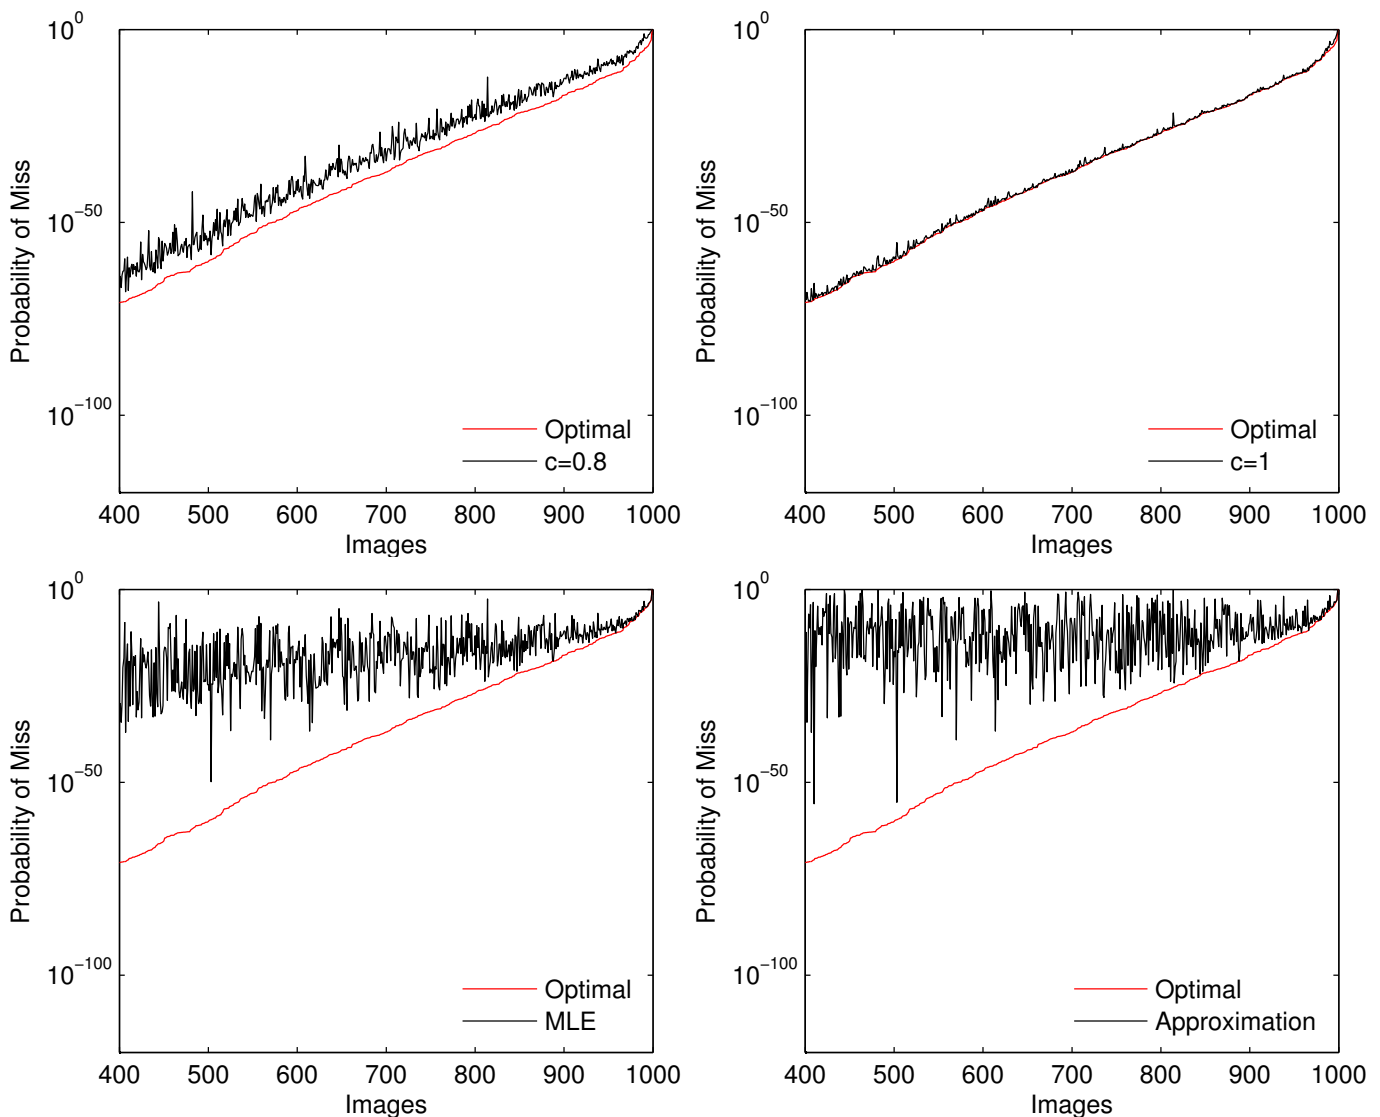


Fig. 5. Rao-GG probability of miss comparison over all images for various choices of c at $P_f = 10^{-6}$ and DWR 16 dB.

TABLE II
NUMBER OF ARITHMETIC OPERATIONS

Detector	Operations			
	$+, -$	\times, \div	x^r, \log	$ \cdot , \text{sgn}$
LC, Eq. (11)	N	N		
LRT-GG [1], Eq. (6)	$3N$	N	$2N$	$2N$
LRT-C [4], Eq. (7)	$4N$	$4N$	N	
Rao-GG [2], Eq. (9)	$2N$	$3N$	N	$2N$
Rao-C [8], Eq. (10)	$2N$	$3N$		

in the number of required iterations to reach our predefined convergence bound that the next estimate in the iteration is in a range of 10^{-6} of the last estimate.

Finally, we cover the effort for the determination of detection thresholds. In case of the LC, LRT-GG and LRT-C detector, we have to compute the mean and variance of the normally distributed detection statistic under the null-hypothesis \mathcal{H}_0 to determine a suitable threshold (see Section II-C). Since detection statistic parameters depend on the noise-model parameters, threshold determination is signal dependent and has to be performed for each image. In contrast to that, the Rao detectors do not require to compute detection statistic parameters at all, since the detection statistic under \mathcal{H}_0 follows a Chi-Square distribution with one degree of freedom. This perfectly fits into the lightweight idea since we can precompute a detection threshold for a desired probability of false-alarm and use this threshold for all images.

To quantify the computational differences in practice, we perform a comparative runtime measurement for all five detectors,

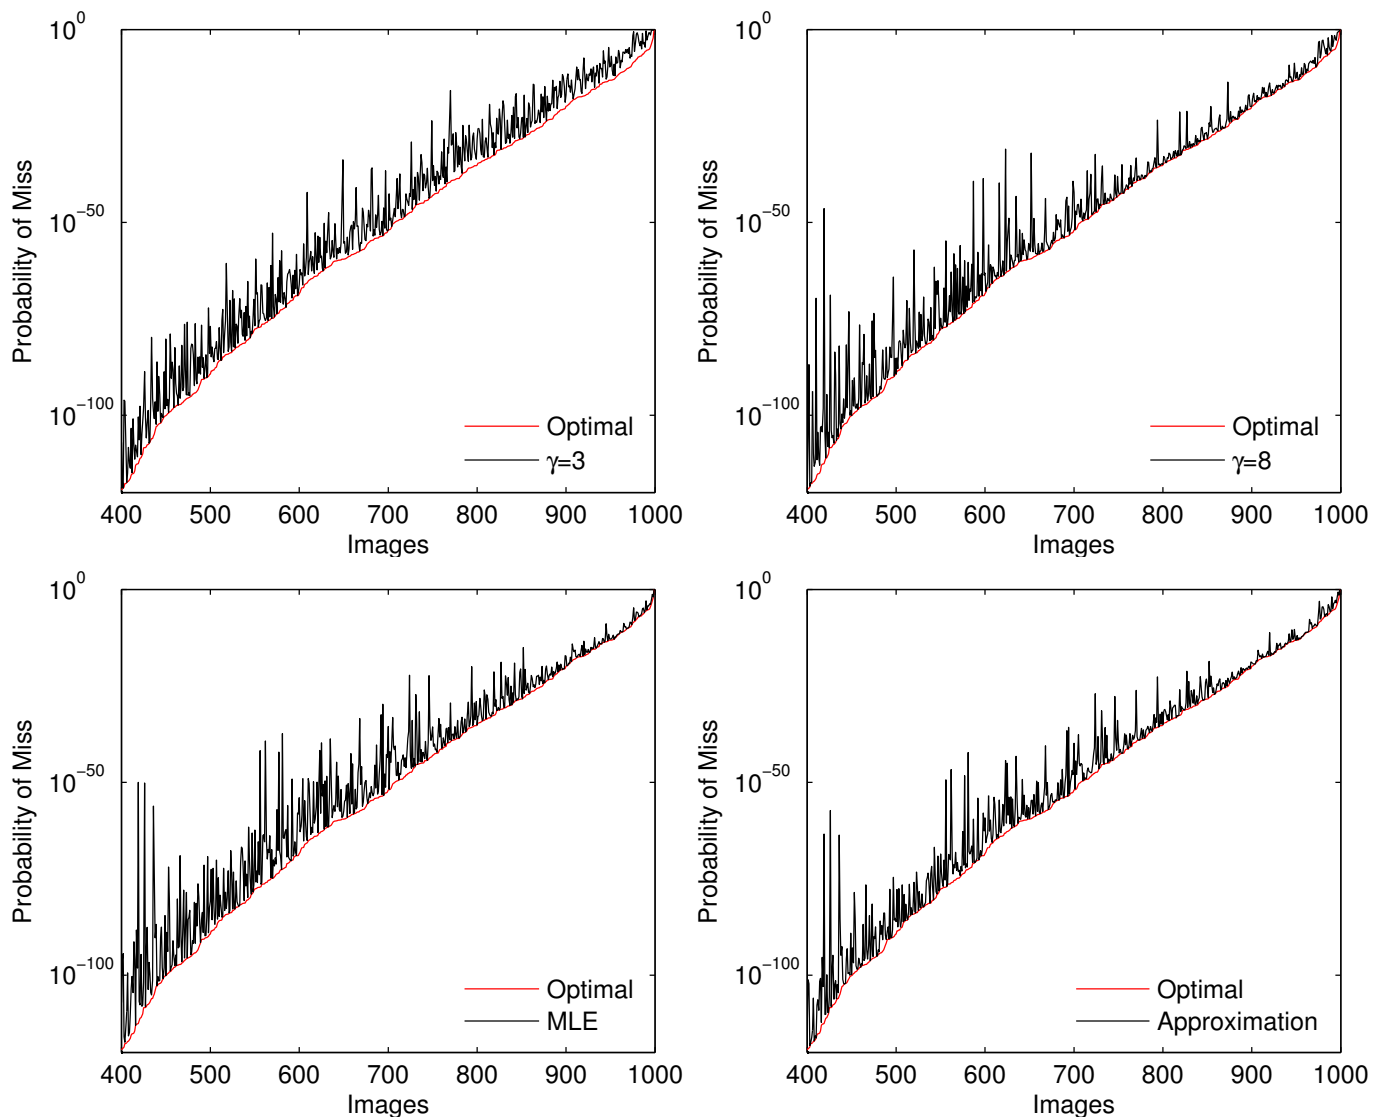


Fig. 6. LRT-C probability of miss comparison over all images for various choices of γ at $P_f = 10^{-6}$ and DWR 16 dB.

implemented in the MATLAB programming environment². The runtime is measured on a Intel Core2 Duo 2.66 GHz system with 2 GB of memory running MATLAB 7.6. We emphasize that the focus is more on relative runtime differences than on absolute values. We choose a real-world signal (\mathbf{V}_1 subband of 512×512 *Lena*) and average the runtimes over 100 runs. The plots in Figs. 10 and 11 illustrate the three main runtime contributions, namely host signal model parameter estimation, threshold selection and computation of the detection statistic. In Fig. 10 we use ML estimation in the estimation step, whereas in Fig. 11 we use the fast approximations. Of course, in case of fixed parameter settings, we can ignore the contribution of the estimation step. As we can see, for the LRT-C, LRT-GG and LC detector the threshold selection dominates the computation of the actual detection responses. Further, the runtime differences in the detection response calculation are in good agreement with the analytical results of Table II. As expected, the estimation process in the ML case has the greatest runtime impact in all cases. Comparing the fast parameter estimation methods, we observe that the GG shape estimation is more efficient than the fast Cauchy estimation approach, since we only need to insert the sample moments m_1 and m_2 in Eq. (17). In contrast to that, the Cauchy model parameter estimation in Eq. (21) requires exponentiation with a floating point number. For the Rao-GG and LRT-GG detector, the computation of the detection response is computationally more expensive than the estimation procedure. Overall, the Rao-C detector provides the best compromise between detection performance and computational effort.

VII. CONCLUSION

In this article, we have taken a closer look at four state-of-the-art detectors in the field of additive spread-spectrum watermarking in the DWT domain. We reviewed the a-priori requirements and assumptions which are taken as a basis for the

²Source code is available at <http://www.wavelab.at/sources>.




Fig. 7. Rao-C probability of miss comparison over all images for various choices of γ at $P_f = 10^{-6}$ and DWR 16 dB.

derivation of each detector. We then discussed parameter estimation issues in a Maximum-Likelihood framework as well as some approximate solutions. Our first experimental results showed that neither ML estimation nor the approximations lead to the optimal detection performance over a set of test images. More extensive experiments on a large number of images revealed that it is actually possible to find fixed parameter settings for each host signal noise model which allow competitive or even better detection performance than the approaches using estimates. This is in accordance with the idea of lightweight detection since the parameter estimation process poses a computational bottleneck. Regarding the computational effort to compute the detection statistics itself as well as detection thresholds, we favor the Rao-C detector due to a small number of arithmetic operations and one-time threshold determination.

We believe that a lightweight detection approach is a good match when the resources of the embedding side are constrained and thus allow only a simple embedding strategy, such as the scenario described by Nelson et al. [23] where the watermark is added by CMOS sensor hardware. Further, fast detectors might be useful when trying to combat watermark desynchronization (e.g. due to geometric transformation attacks) with exhaustive search over the transform space [24].

Perceptual shaping of the watermark or multiplicative embedding [25] complicate the formulation of the detection statistics. Liu et al. [26] propose to transform the cover signal into a perceptually uniform domain where simple additive embedding can be employed and derive a locally-optimum detector for the GGD host signal model. Fast parameter estimation is directly applicable to Liu's framework, but we have not derived alternative detectors based on the Cauchy model or Rao test. However, we caution that the effort to compute the perceptual model or the respective transform may outweigh the actual detection effort.

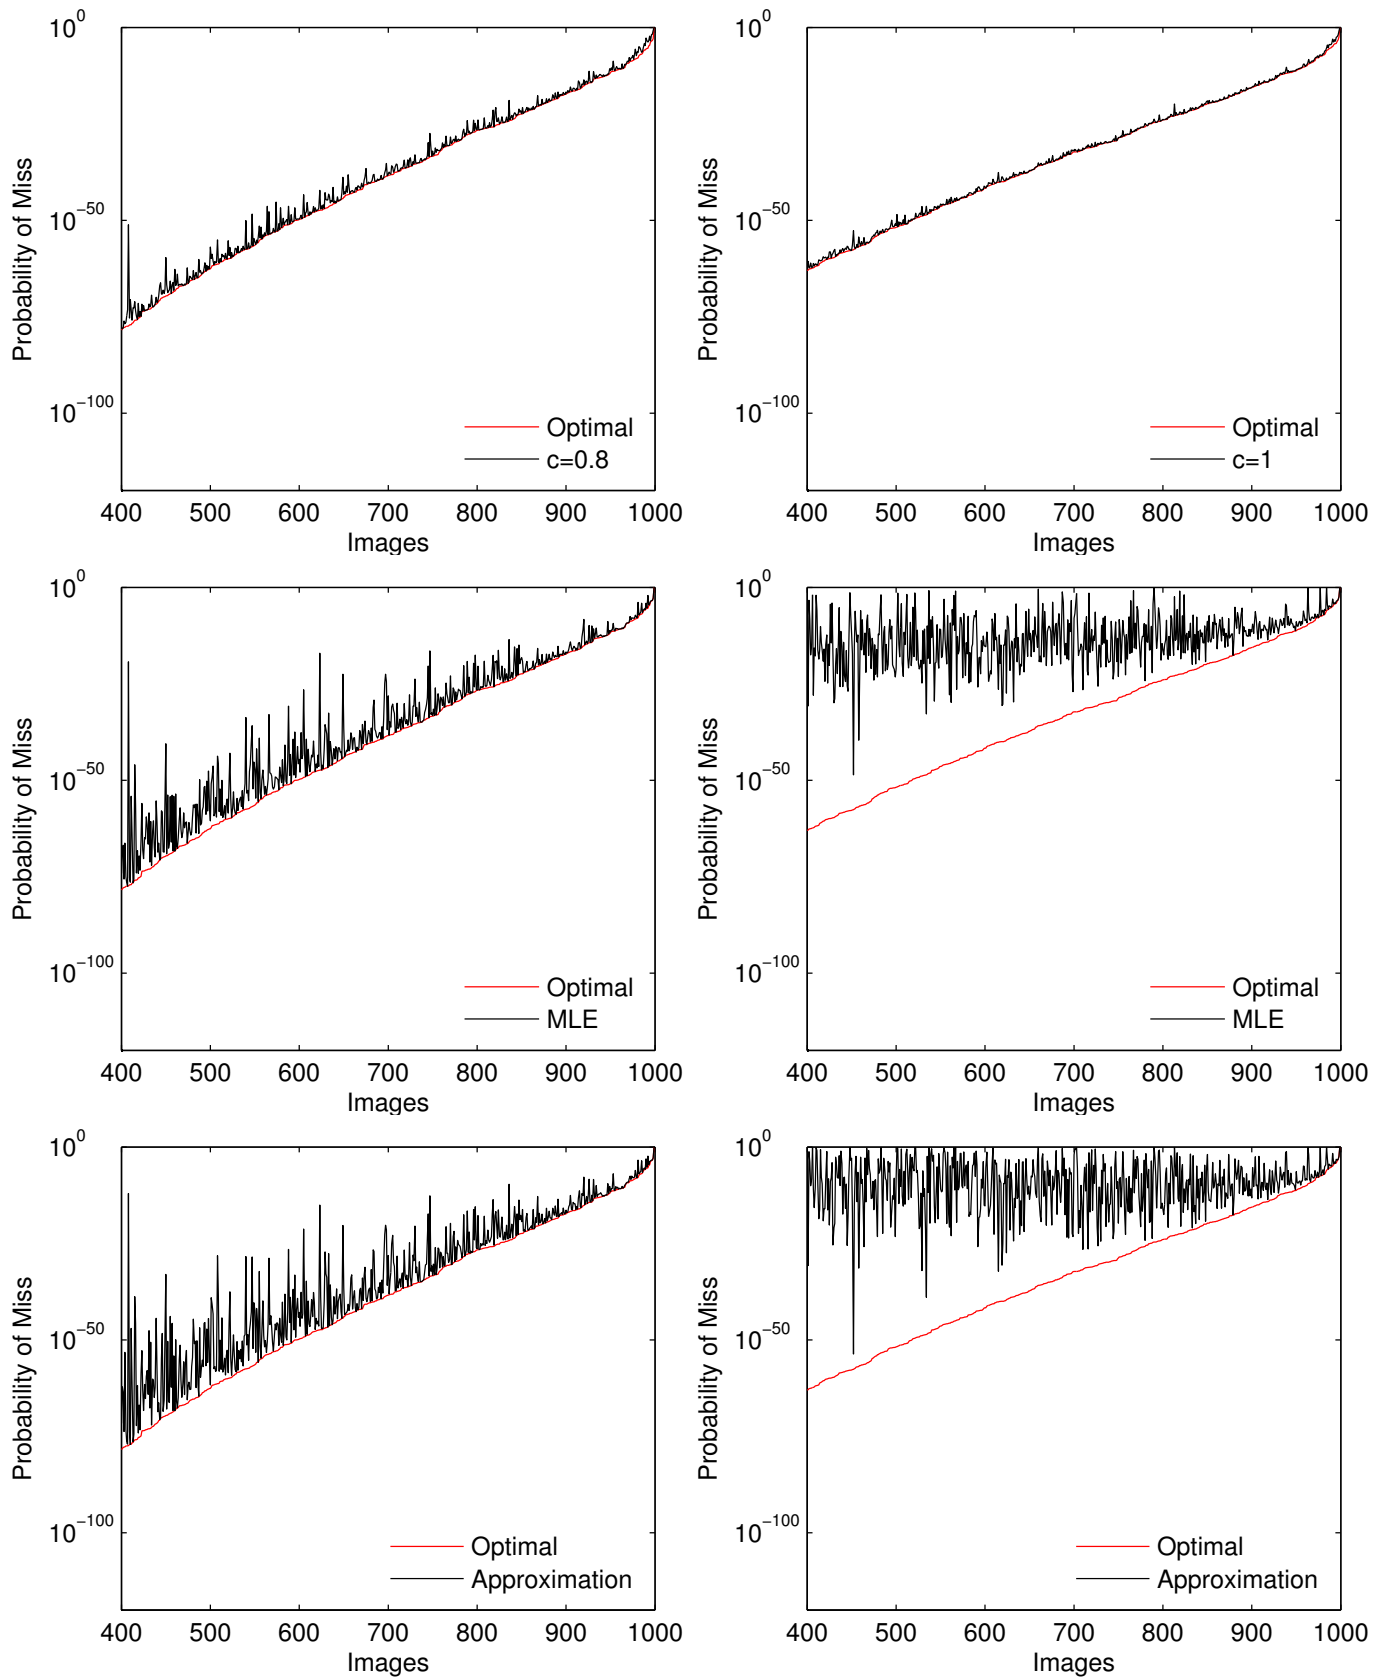


Fig. 8. Probability of miss after JPEG compression ($Q = 70$) for the LRT-GG (left) and Rao-GG (right) detector at $P_f = 10^{-6}$ and DWR 16 dB.

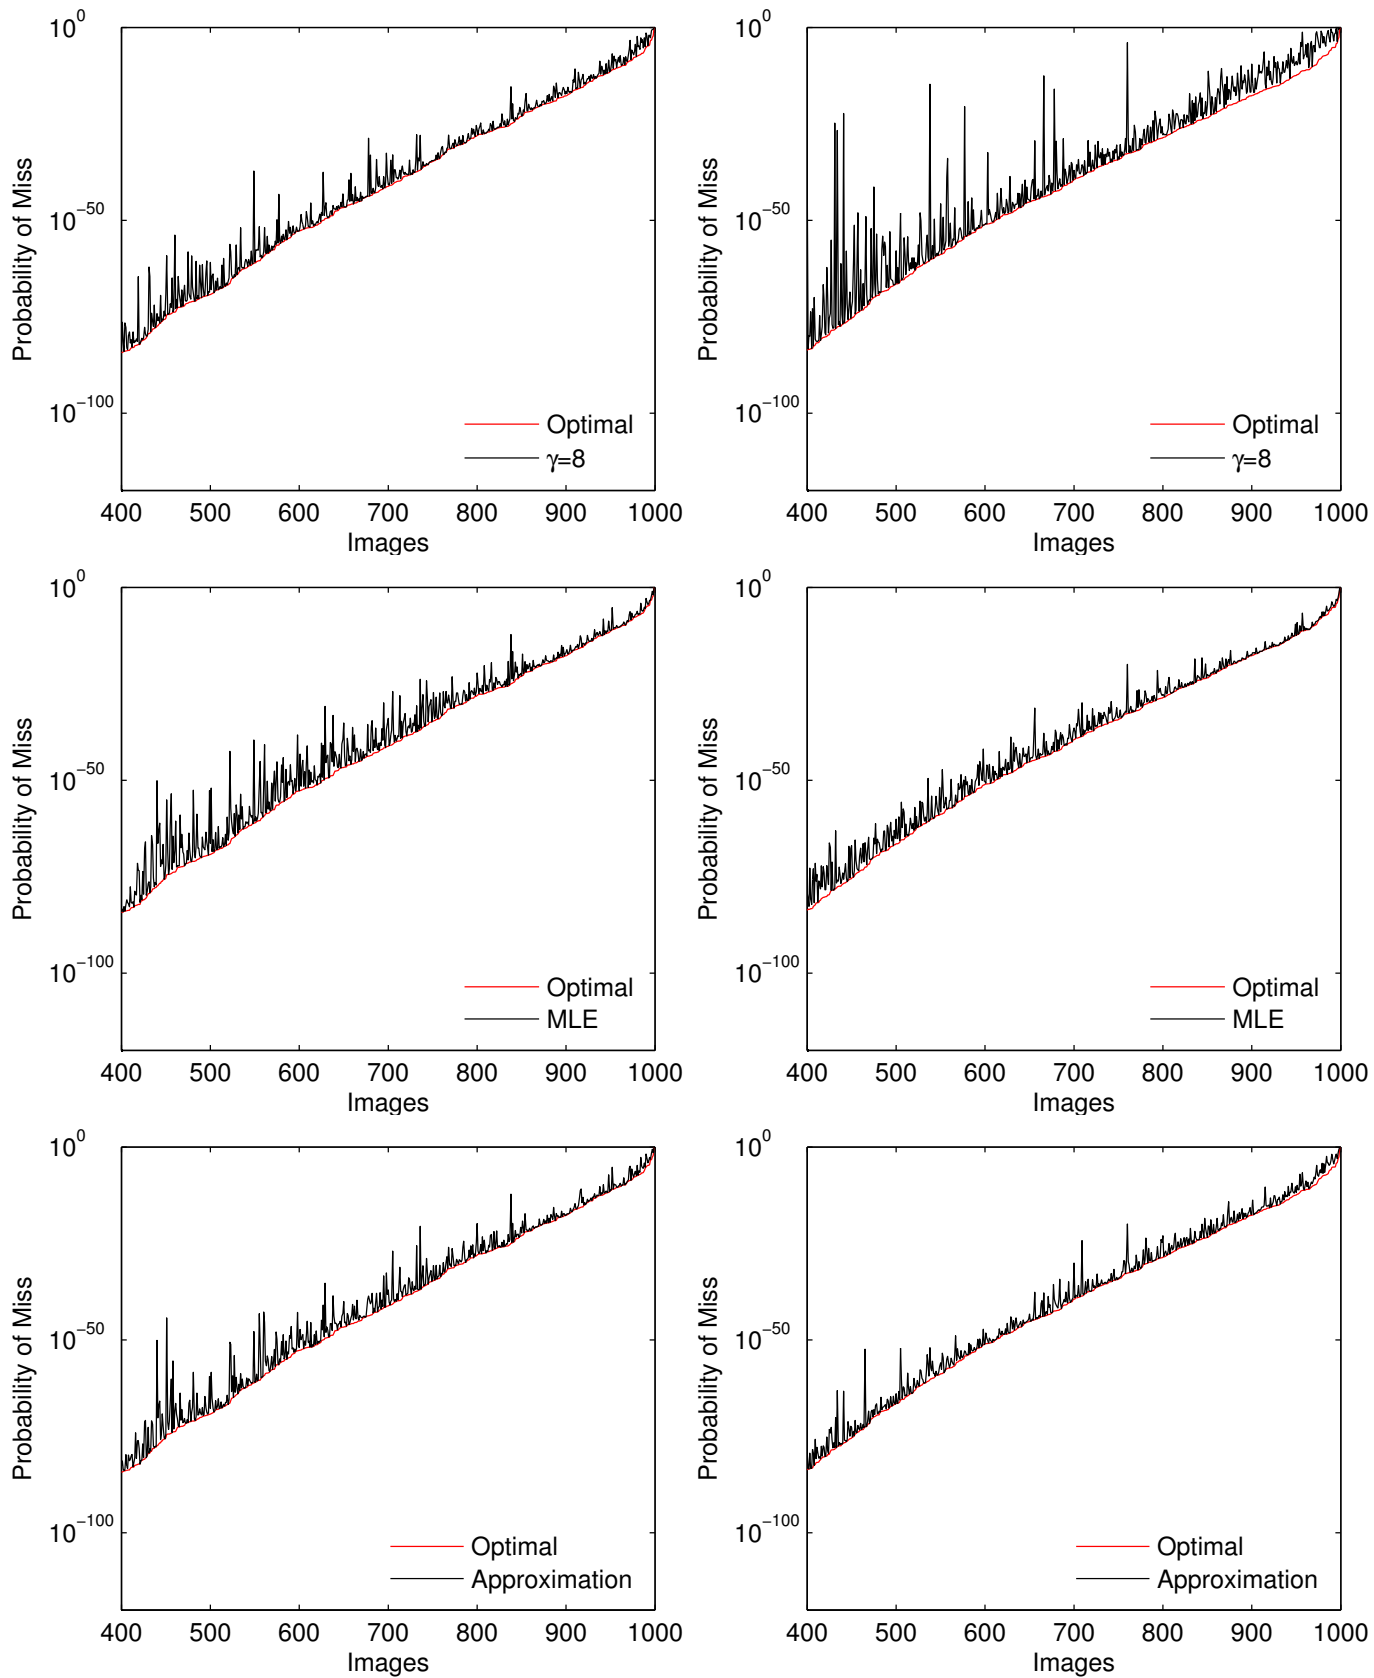


Fig. 9. Probability of miss after JPEG compression ($Q = 70$) for the LRT-C (left) and Rao-C (right) detector at $P_f = 10^{-6}$ and DWR 16 dB.




Fig. 10. Runtime measurement (using MLE) for a signal of length $N = 67600$ (V_1 subband of the 512×512 version of *Lena*).




Fig. 11. Runtime measurement (using fast estimation) for a signal of length $N = 67600$ (V_1 subband of the 512×512 version of *Lena*).

REFERENCES

- [1] J. R. Hernández, M. Amado, and F. Pérez-González, "DCT-domain watermarking techniques for still images: Detector performance analysis and a new structure," *IEEE Transactions on Image Processing*, vol. 9, no. 1, pp. 55–68, Jan. 2000.
- [2] A. Nikolaidis and I. Pitas, "Asymptotically optimal detection for additive watermarking in the DCT and DWT domains," *IEEE Transactions on Image Processing*, vol. 12, no. 5, pp. 563–571, May 2003.
- [3] Q. 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, Sep. 2001.
- [4] A. Briassouli, P. Tsakalides, and A. Stouraitis, "Hidden Messages in Heavy-Tails: DCT-Domain Watermark Detection Using Alpha-Stable Models," *IEEE Transactions on Multimedia*, vol. 7, no. 4, pp. 700–715, Aug. 2005.
- [5] K. A. Birney and T. R. Fischer, "On the modeling of DCT and subband image data for compression," *IEEE Transactions on Image Processing*, vol. 4, no. 2, pp. 186–193, Feb. 1995.
- [6] S. M. Kay, *Fundamentals of Statistical Signal Processing: Detection Theory*, A. Oppenheim, Ed. Prentice-Hall, 1998, vol. 2.
- [7] C. R. Rao, *Linear Statistical Inference and Its Applications*, ser. Probability and Mathematical Statistics. Wiley, 1973.
- [8] R. Kwitt, P. Meerwald, and A. Uhl, "A lightweight Rao-Cauchy detector for additive watermarking in the DWT-domain," in *Proceedings of the ACM Multimedia and Security Workshop, MMSEC '08*. Oxford, UK: ACM, Sep. 2008, pp. 33–41.
- [9] —, "Efficient detection of additive watermarking in the DWT-domain," in *Proceedings of the 17th European Signal Processing Conference, EUSIPCO '09*. Glasgow, UK: EURASIP, Aug. 2009, pp. 2072–2076.
- [10] S. Mallat, "A Theory for Multiresolution Signal Decomposition: The Wavelet Representation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 11, no. 7, pp. 674–693, Jul. 1989.
- [11] S. Nadarajah, "A generalized normal distribution," *Journal of Applied Statistics*, vol. 32, pp. 685–694, Sep. 2005.
- [12] C. Nikias and M. Shao, *Signal Processing with Alpha-Stable Distributions and Applications*. Wiley-Interscience, 1995.
- [13] K. Krishnamoorthy, *Handbook of Statistical Distributions with Applications*. Chapman & Hall, 2006.
- [14] S. M. Kay, *Fundamentals of Statistical Signal Processing: Estimation Theory*. Prentice-Hall, 1993, vol. 1.
- [15] —, "Asymptotically optimal detection in incompletely characterized non-gaussian noise," *IEEE Transactions on Acoustics, Speech and Signal Processing*, vol. 37, no. 5, pp. 627–633, May 1989.
- [16] M. Varanasi and B. Aazhang, "Parametric Generalized Gaussian density estimation," *Journal of the Acoustical Society of America*, vol. 86, no. 4, pp. 1404–1415, Oct. 1989.
- [17] M. Do and M. Vetterli, "Wavelet-based texture retrieval using Generalized Gaussian density and Kullback-Leibler distance," *IEEE Transactions on Image Processing*, vol. 2, no. 2, pp. 146–158, Feb. 2002.
- [18] R. Krupinski and J. Purczynski, "Approximated fast estimator for the shape parameter of Generalized Gaussian distribution," *Signal Processing*, vol. 86, no. 2, pp. 205–211, Feb. 2006.
- [19] K.-S. Song, "A globally convergent and consistent method for estimating the shape parameter of a Generalized Gaussian distribution," *IEEE Transaction on Information Theory*, vol. 52, no. 2, pp. 510–527, Feb. 2006.
- [20] Q. Zhang and W. Wang, "An efficient algorithm for approximate biased quantile computation in data streams," in *Proceedings of the 16th ACM Conference on Information and Knowledge Management, CIKM '07*. Lisboa, Portugal: ACM, Nov. 2007, pp. 1023–1026.

- [21] G. A. Trihrintzis and C. L. Nikias, "Fast estimation of the parameters of alpha-stable impulsive interference," *IEEE Transactions on Signal Processing*, vol. 44, no. 6, pp. 1492–1503, Jun. 1996.
- [22] G. Schaefer and M. Stich, "UCID - an uncompressed colour image database," in *Proceedings of SPIE, Storage and Retrieval Methods and Applications for Multimedia*, vol. 5307. San Jose, CA, USA: SPIE, Jan. 2004, pp. 472–480.
- [23] G. R. Nelson, G. A. Julien, and O. Yadid-Pecht, "CMOS image sensor with watermarking capabilities," in *Proceedings of the IEEE International Symposium on Circuits and Systems, ISCAS '05*, vol. 5. IEEE, May 2005, pp. 5326–5329.
- [24] M. Barni, "Effectiveness of exhaustive search and template matching against watermark desynchronization," *IEEE Signal Processing Letters*, vol. 12, no. 2, pp. 158–161, Feb. 2005.
- [25] Q. Cheng and T. Huang, "Robust optimum detection of transform domain multiplicative watermarks," *IEEE Transactions on Signal Processing*, vol. 51, no. 4, pp. 906–924, Apr. 2003.
- [26] W. Liu, L. Dong, and W. Zeng, "Optimum detection for spread-spectrum watermarking that employs self-masking," *IEEE Transactions on Information Forensics and Security*, vol. 2, no. 4, pp. 645–654, Dec. 2007.