

Environmentally Adaptive Noise Estimation for Active Sonar

**A thesis submitted in partial fulfilment
of the requirement for the degree of Doctor of Philosophy**

Robert Bareš

**Cardiff University
School of Computer Science**

22nd February 2012

Declaration

This work has not previously been accepted in substance for any degree and is not concurrently submitted in candidature for any degree.

Signed (candidate) Date

Statement 1

This thesis is being submitted in partial fulfilment of the requirements for the degree of PhD.

Signed (candidate) Date

Statement 2

This thesis is the result of my own independent work/investigations, except where otherwise stated. Other sources are acknowledged by explicit references.

Signed (candidate) Date

Statement 3

I hereby give consent for my thesis, if accepted, to be available for photocopying and for inter-library loan, and for the title and summary to be made available to outside organisations.

Signed (candidate) Date

Abstract

Noise is frequently encountered when processing data from the natural environment, and is of particular concern for remote-sensing applications where the accuracy of data gathered is limited by the noise present. Rather than merely accepting that sonar noise results in unavoidable error in active sonar systems, this research explores various methodologies to reduce the detrimental effect of noise.

Our approach is to analyse the statistics of sonar noise in trial data, collected by a long-range active sonar system in a shallow water environment, and apply this knowledge to target detection. Our detectors are evaluated against simulated targets in simulated noise, simulated targets embedded in noise-only trial data, and trial data containing real targets.

First, we demonstrate that the Weibull and K-distributions offer good models of sonar noise in a cluttered environment, and that the K-distribution achieves the greatest accuracy in the tail of the distribution. We demonstrate the limitations of the Kolmogorov-Smirnov goodness-of-fit test in the context of detection by thresholding, and investigate the upper-tail Anderson-Darling test for goodness-of-fit analysis. The upper-tail Anderson-Darling test is shown to be more suitable in the context of detection by thresholding, as it is sensitive to the far-right tail of the distribution, which is of particular interest for detection at low false alarm rates. We have also produced tables of critical values for K-distributed data evaluated by the upper-tail Anderson-Darling test.

Having established suitable models for sonar noise, we develop a number of detection statistics. These are based on the box-car detector, and the generalized likelihood ratio test with a Rician target model. Our performance analysis shows that both types of detector benefit from the use of the noise model provided by the K-distribution. We also demonstrate that for weak signals, our GLRT detectors are able to achieve greater probability of detection than the box-car detectors. The GLRT detectors are also easily extended to use more than one sample in a single test, an approach that we show to

increase probability of detection when processing simulated targets.

A fundamental difficulty in estimating model parameters is the small sample size. Many of the pings in our trial data overlap, covering the same region of the sea. It is therefore possible to make use of samples from multiple pings of a region, increasing the sample size. For static targets, the GLRT detector is easily extended to multi-ping processing, but this is not as easy for moving targets. We derive a new method of combining noise estimates over multiple pings. This calculation can be applied to either static or moving targets, and is also shown to be useful for generating clutter maps. We then perform a brief performance analysis on trial data containing real targets, where we show that in order to perform well, the GLRT detector requires a more accurate model of the target than the Rician distribution is able to provide. Despite this, we show that both GLRT and box-car detectors, when using the K-distribution as a noise model, can achieve a small improvement in the probability of detection by combining estimates of the noise parameters over multiple pings.

Table of Contents

Abstract	i
List of Symbols and Operators	ix
Acknowledgements	xi
1 Introduction	1
1.1 Historical Context	1
1.2 Research Hypothesis	7
1.3 Objectives	7
1.4 Thesis Outline	8
1.5 Contributions	11
2 Background	15
2.1 Basic Principles	16
2.2 Propagation and Noise	18
2.2.1 Noise	19
2.2.2 Propagation Loss	20
2.2.3 The Speed of Sound in Water	22
2.2.4 Multipath	23
2.3 Sonar Signal Processing Chain	24
2.3.1 Active Sonar	25
2.3.2 Passive Sonar	29

2.4	Hydrophone Arrays	31
2.5	Beamforming	33
2.6	Matched Filters	38
2.6.1	Matched Filter in the Time Domain	38
2.6.2	Matched Filter in the Frequency Domain	39
2.6.3	Matched Filter Output	41
2.7	Normalization	44
2.8	Thresholding	48
2.8.1	Hypothesis Testing	48
2.8.2	Performance Analysis	54
2.9	Target Tracking	55
2.9.1	Clustering	56
2.9.2	Tracking	57
3	Trial Data	61
3.1	Introduction	61
3.2	Data Format	63
3.2.1	The Initial Blast	65
3.2.2	Reverberation and Clutter	66
3.3	Non-Acoustic Data	67
3.3.1	Coordinate Space	71
3.3.2	Spatial Translation	72
4	Noise Distributions	79
4.1	Noise Distributions	80
4.2	Parameter Estimation	84
4.2.1	Estimation Theory	84
4.2.2	The $z \log(z)$ Estimator	89
4.2.3	The Method of Moments Estimator	90

4.2.4	The Bayesian Approaches	92
4.2.5	Large Values of ν	94
4.3	Basic Random Number Generation	96
4.3.1	The Inversion Method	97
4.3.2	The Acceptance-Rejection Method	99
4.3.3	The Modular Method	102
4.4	Generating Sonar Noise	104
4.5	Generating Autocorrelated Sonar Noise	106
5	Statistical Analysis of Trial Data	111
5.1	Analysis of Parameter Estimates	112
5.2	Goodness-of-Fit Testing	119
5.3	Analysis of Trial Data	123
5.3.1	Goodness-of-Fit	123
5.3.2	Non-Rayleigh Environments	125
5.4	Discussion	129
6	Target Distributions	131
6.1	Embedding Models	132
6.1.1	The Rice Distribution	132
6.1.2	The Homodyned K-Distribution	133
6.1.3	The Generalized K-Distribution	135
6.1.4	The Nakagami Distribution	137
6.1.5	The Generalized Gamma Distribution	139
6.2	Swerling Models	140
6.2.1	The Swerling 0 Model	141
6.2.2	The Swerling I Model	142
6.2.3	The Swerling III Model	142
6.3	Generating Data	143

6.3.1	The Rice Distribution	143
6.3.2	The Generalized K-Distribution	144
6.3.3	The Homodyned K-Distribution	145
7	Target Detection	149
7.1	Experimental Design	150
7.1.1	Number of Detection Tests	151
7.1.2	Spatial Distribution of Detection Tests	152
7.1.3	Target Simulation	153
7.2	The Box-car Detector	156
7.3	Likelihood Ratio Test	158
7.3.1	Generalized Likelihood Ratio Test	160
7.3.2	Extended Detection	160
7.4	Comparison of GLR statistics	162
7.4.1	Experimental Scheme	163
7.5	Theoretical ROC curves	170
7.5.1	Analysis of Results	172
7.6	Application to Real Sonar Noise	175
7.6.1	Box-car Detectors	176
7.6.2	GLRT Detectors	177
7.7	Spatial Variation of Performance	179
7.8	Application to Real Targets	184
8	Multi-ping Processing	191
8.1	Combining Parameter Estimates	193
8.2	Performance of Combined Estimates	199
8.3	Multi-ping Composite Images	203
8.3.1	Coordinate Space	203
8.3.2	Registration	204

8.3.3	Visualisation	206
8.3.4	Temporal Statistics	214
8.4	Multi-ping Detection	219
8.5	Application to Real Sonar Noise	221
8.6	Application to Real Targets	224
9	Conclusion	229
9.1	Contributions	231
9.2	Limitations and Future Work	233
A	Simulated ROC Curves	237
B	ROC Curves for Real Noise	255
C	Spatial Performance	271
	References	281

List of Symbols and Operators

r	radius from a central point, range from the sonar system	21
$(f * g)$	convolution of f with g	39
$\mathcal{F}\{\cdot\}$	Fourier transform	39
\otimes	point-wise multiplication	39
X	complex signal from a matched filter, may include a target signal	41
$ \cdot $	absolute value, euclidean norm	41
ϕ	phase	41
A_X	amplitude of received signal, may include a target signal	41
b	bearing	42
S	complex target signal without noise	45
$x \sim F(\theta)$	x is distributed according to the distribution F	45
H_0	null hypothesis	49
H_1	alternative hypothesis	49
T	threshold	49
\tilde{X}	normalized amplitude	49
$P(\cdot)$	probability of an event occurring	53
$F(\cdot)$	cumulative distribution function	53
α	false alarm rate	53
$\text{erf}(\cdot)$	Gauss error function	53
$\{\varphi, \lambda\}$	latitude and longitude	74
R	radius of the earth	74

$\lfloor \cdot \rfloor$	floor operator, rounding down	77
λ	K-distribution scale parameter	81
ν	K-distribution shape parameter	81
$\Gamma(\cdot)$	gamma function	81
$K_\nu(\cdot)$	modified Bessel function of the second kind	81
p	probability of an event occurring	85
θ	parameter for some abstract function	85
$L(\cdot)$	likelihood of an event occurring	86
m_k	the k 'th non-central sample moment	86
$\hat{\theta}$	estimate of a parameter θ	87
$E(\cdot)$	expected value	88
z	data in the intensity domain	90
$\psi(\cdot)$	digamma function	91
$\gamma(\cdot, \cdot)$	lower incomplete gamma function	101
$\rho(\cdot)$	autocorrelation function	106
$\sup \cdot $	supremum operator	119
$I \cdot $	indicator function	119
D_n	Kolmogorov-Smirnov statistic	119
ω^2	Cramér-von-Mises statistic	120
$\Phi(\cdot)$	weighting function	121
A_n^2	upper-tail Anderson-Darling statistic	121
I_0	modified Bessel function of the first kind	132
A_Y	noise amplitude	133
Y	complex noise signal from a matched filter	145
N_D	number of detection tests needed for a chosen false alarm rate	152
$\lceil \cdot \rceil$	ceiling operator, rounding up	152
N_P	number of pings in a dataset	152
Θ	parameter space	159

Acknowledgements

I would like to express my thanks to my supervisor, Dafydd Evans. This PhD was an incredible opportunity for me, and I deeply appreciate his relentless encouragement. Dafydd's enthusiasm for research has been an inspiration. His sound advice and critical eye were instrumental in this work. I would also like to thank my industrial supervisors, Stephen Long and Roger Benton. Throughout my thesis Stephen and Roger provided a great deal of technical guidance, and many helpful suggestions. Without them, this project would not have been possible.

Further, I want to express my appreciation for the enjoyable atmosphere in our research laboratory. This was due in no small part to my student colleagues, who with many a friendly conversation provided a fun and supportive work environment. I am especially grateful to Ahmed Alazzawi, Dashan Lu, Ian Cooper, Martin Chorley, Rich Coombs, Alysia Skilton, Diego Pizzocaro, Matthew Williams and Jonathan Quinn in Cardiff, and Trevor Wood at Oxford.

Post-graduate education can be a particularly stressful time, and so it is only right that I thank my closest friends for their support. Thanks are due to Serkan Ahmet, Mengle Zhang, Leon Gregory, Steve Turner and Franziska Kohlt.

Finally I wish to thank my parents, Lynda and Petr, who have always encouraged me to pursue my dreams. I am lucky in knowing that their unassuming love and support is just a phone-call away. I am forever grateful for everything they have done for me.

The work undertaken in this thesis was funded by the Engineering and Physical Sciences Research Council (EPSRC) and Thales Underwater Systems Ltd, Cheadle Heath, UK. The project was arranged and supported by the Industrial Mathematics Knowledge Transfer Network. Access to trial data was arranged by Roger Benton and Les Hart of Thales Underwater Systems.

Chapter 1

Introduction

“

Above all, the spur of necessity under war conditions made men's brains act with greater vigour, and science responded to the demands.

— WINSTON CHURCHILL, 1935

1.1 Historical Context

Over 50 million years have passed since *microchiropterans*, or microbats, first evolved the ability to detect and hunt their prey through the use of sound¹. This process, known as *acoustic location* or *echo-ranging*, is also used by members of the *cetacea* order of whales and dolphins, and three genera of mouse-like mammals [53, 65]. In this context,

¹An account of the evolution of microbats can be found in the work of Altringham et al. [14]

the use of acoustics has provided an evolutionary advantage, either by offsetting the relative weakness of existing senses, such as bats' generally poor eyesight, or as a result of adaptation to the environment, as with dolphins and other cetacea.

Captain Linwood Howeth has assembled an extensive historical account of the use of underwater acoustics [57], an outline of which is presented here. As stated by Captain Howeth himself, the first realization that sound carried through water is lost in antiquity. Although Sri Lankan tribes are known to have communicated with their boats at sea by hitting a clay pot under the water, it was not until 1490 that acoustic location was rediscovered by Leonardo da Vinci. Leonardo found that, by placing one end of a tube into the sea and the other end to his ear, he was able to hear nearby ships.

In the 1850s, the American Lighthouse Board recognised the problem of signalling ships to warn them of shallow waters. Despite the use of Fresnel lenses² to concentrate beams of light, and timed bells for use during fog, these signalling systems proved erratic. In the 1890s, Lucien Blake developed an alternative system, using an underwater bell and a corresponding microphone at the receiving vessel. Blake later took his experience of underwater signalling to the Submarine Signalling Company, and was their chief engineer in 1906 and 1907 [83].

The sinking of the Titanic on the 14th of April 1912 strengthened the market for navigational aids. In the same month, the Submarine Signalling Company employed Reginald Fessenden to improve on Blake's earlier design. In Blake's system, the underwater microphone (*hydrophone*) was sensitive to other sounds in the sea, which made it difficult to detect the sound of a warning bell. Fessenden set about designing a more discerning microphone for use underwater, and was awarded five patents for his work between 1914 and 1920. This device, known as the *Fessenden Oscillator*, is

²Lenses for use in lighthouses should ideally have a short focal length, equal to the distance between the light source and the lens itself, and a large aperture - the area of the lens through which the light will travel. Conventional lenses must be large in order to achieve these two properties, but the staggered surface of a Fresnel lens achieves this in a lightweight design.

also capable of transmitting sound, and was one of the first systems to be used for sonar echo-ranging as it is known today.

Following the outbreak of World War I, with German U-boats breaching terms of the Hague conventions of 1899 and 1907, there was an internationally recognised need for improved methods of *anti-submarine warfare*. This need spurred further research in acoustic location, with notable contributions from the Royal Navy Anti-Submarine Division and the Americans' newly founded Subcommittee on Submarine Detection by Sound. As the application of acoustic location had shifted from navigation to the detection of enemy targets, technology had to adapt in order to meet new requirements. Existing systems were only required to operate over short ranges, to detect obstacles in close proximity. In contrast, for military applications, the early detection of an enemy unit beyond firing-range would provide a strategic advantage.

Despite the short detection range offered by the Fessenden Oscillator at the outbreak of WWI, the device was fitted to 10 Royal Navy submarines in 1915. Unfortunately the system did not perform well. While British forces concentrated research efforts on a system known as ASDIC, the Americans continued to refine the Fessenden Oscillator. While further prototypes were developed throughout the war, the use of acoustic location achieved only a limited impact on the threat posed by German U-boats. A far greater impact was achieved by the use of destroyers to escort merchant vessels in convoy.

When military co-operation between Britain and America resumed during World War II, it became apparent that the two countries had achieved similar advances in echo-ranging during the inter-war years. As detection ranges had improved from just a few hundred feet in 1927 to around five miles by 1929, echo-ranging equipment was ready to play a key role in the naval battles of World War II. Large numbers of ships and patrol boats were fitted with the new echo-ranging equipment, which led in part to the allies' success in the battles for the Atlantic and Pacific oceans. The advances in echo-ranging technology also had an influence on naval tactics. For example, although

some vessels were only equipped with torpedo detectors, torpedo countermeasures were in use during the Battle of the Atlantic. Acoustic torpedoes, which home in on a target by following its acoustic emissions, could be thrown off course by a decoy. This was achieved by releasing a sound transmitter behind the vessel, which reduced naval losses. Furthermore, buoys equipped with echo-ranging equipment and a radio transmitter, known as *sonobuoys*, were successfully dropped from aircraft in order to survey the underwater environment. If a target was detected, this information would be relayed back to the aircraft, allowing focused deployment of available resources.

Advances in the understanding of underwater acoustics offered further advantages. For example, during the Battle of the Pacific, the U.S. Navy had obtained charts depicting regions which contained large shoals of snapping shrimp. Given the noisy nature of these shrimp, which snap their claws for both hunting and communication, the natural environment could be used to mask the sound of the American submarines' machinery. The impact of using echo-location in shallow waters was also recognised, with the development of an additional feature known as *reverberation gain control*. This behaves in a similar fashion to a volume control, and is capable of reducing the effect of loud echoes received from rocky outcrops and nearby sea floor surfaces - which, while useful for navigational purposes, are of no interest to target detection.

In 1946 the American system became known as *sonar*, which was initially³ an acronym for **s**ound, **n**avigation and **r**anging. Despite the predominantly military-oriented history of sonar, the technology has resulted in a broad range of civilian applications. Of particular note is the *fathometer*, developed during the inter-war period by the Submarine Signalling Company for measuring sea depth, known as *echo-sounding*. Echo-sounding equipment is commonly applied in oceanographic surveys to establish the *bathymetry* (topography) of the sea floor, or sub-floor profiles indicating geological characteristics such as the presence of oil deposits. This same technology is also used

³Despite its etymology as an acronym, through popular use sonar has become a noun, entering the Oxford English Dictionary in 1986

in commercial *fish-finders*, where experienced users can identify changes in reported depth corresponding to the movement of fish beneath the sonar equipment.

Since 1994 there has been increasing focus on the environmental impact of high power sonar on marine species [27], some of which use the same frequencies for communication. This interest was stimulated by a publication from the American Committee on Low-Frequency Sound and Marine Mammals, which cited a lack of data as the main obstacle to evaluating the effects of sonar on marine life [98]. Beyond the potential for hearing loss, the use of military sonar has since been connected to a number of mass strandings of whales following exercises conducted by NATO and the U.S. Navy [46]. These occurrences have been exacerbated by the recent development of low-frequency systems, which operate over long ranges and at high power. Although the processes involved are still not well understood, initial predictions suggested that particularly sensitive marine mammals, such as beaked whales, may be distressed by sonar signals and in some cases become beached while attempting to flee [78]. As stated by the Director of the Marine Mammal Protection Project at the American Natural Resources Defense Council,

There is no longer a serious scientific debate about the connection between sound and marine mammal mortality. A range of experts, from the International Whaling Commission's ("IWC") Scientific Committee to the U.S. Navy's own commissioned scientists, have agreed that the evidence linking mass strandings to mid-frequency sonar is "convincing" and "overwhelming."

— Joel Reynolds [112, p762]

Consequently, methods have been proposed for mitigating the impact of military sonar on marine life, for example to gradually increase the power of a sonar system to allow

animals to leave the area [78], regional monitoring for marine life prior to commencing exercises, and ensuring that operations occur at a sufficient distance from marine sanctuaries [13]. Although military and civil policy-makers have different priorities, balancing military effectiveness and environmental concerns, the recommendations of many research groups have resulted in legislation.

A crude means of improving sonar range and performance is to continue increasing the power of the system. However, it is generally agreed that this is unsustainable. Beyond environmental concerns, from an engineering standpoint, as power increases there is an increased risk of performance degradation resulting from *cavitation*, where bubbles form on the sonar equipment due to excessive changes in pressure [136, p6]. Furthermore, increasing power will invariably increase the amplitude of echoes from the local environment, which is already a limiting factor in current systems. Evidently, an alternative approach to improving the performance of sonar systems would be beneficial, which leads to the overall purpose of this research project, and the research hypothesis underpinning the work.

A means of environmental mitigation proposed by Levine et al. [78], is to more carefully design the transmitted signal. In the report, they propose that changes to the signal transmitted by the sonar equipment may reduce or even eliminate damage to marine life. Whether in terms of the frequency range, power of the transmissions, or the duration of transmissions, this approach also offers further scope for performance improvements. If current sonar performance can be matched by a more advanced system at a lower power, then we have essentially “hit two birds with one stone”. On the one hand the maximum power of the system could be reduced, reducing the likelihood of harming marine life and ensuring compliance with environmental legislation. On the other hand, should greater sonar performance be required, increasing the power to current levels would provide substantial gains.

1.2 Research Hypothesis

In this research we aim to improve sonar performance through advances in signal processing. If a sonar system is designed to detect a target object, signals received from any other source are of no interest, and are considered *noise*. Considering the cunning tactics exploited by submarines, such as hiding amongst shrimp, and the difficulties associated with sonar in shallow water, there is an obvious interest in the effects of the natural environment on sonar performance. Consequently, this research is focused on mitigating the effects of noise in an environmentally adaptive manner.

Our hypothesis is that, by analysing the statistical properties of sonar noise and applying the results of this analysis to the problem of target detection, better detection performance can be achieved. Our hope is that this will lead to more sensitive systems and improve operation in difficult environments.

It should be noted that the analytic framework developed in this thesis is not restricted for use only in sonar. The same principles apply to many other forms of remote sensing and signal processing applications. While the most obvious of these is *radar*, with which sonar has a great deal in common, sources of environmental noise exist in many other systems. Other potential applications for our work include medical imaging systems [56], *global navigation satellite systems* [86], and the *non-destructive evaluation* of structural integrity [76].

1.3 Objectives

While research in this field is often of theoretical interest, it is still desirable to achieve practical results. Thales Underwater Systems Ltd. have kindly provided real-world data,

collected from a long-range sonar system during shallow-water trials. The broad aim of this work is to attempt to achieve greater performance than current sonar systems; we define the following objectives:

1. Identify statistical models for sonar noise in shallow water environments
2. Evaluate the accuracy of these models using trial data
3. Apply these models of sonar noise to current methods of target detection
4. Develop and evaluate alternative methods of detection which may be better suited to these noise models
5. Develop and evaluate generic ways of combining measurements, to make more effective use of available data

This forms a natural flow from theoretical interests to practical implementations. Comparing these implementations in terms of a quantitative performance analysis then allowed us to critically evaluate the success of this work.

1.4 Thesis Outline

In this chapter we outline the motivation for this project, followed by our research hypothesis and a list of our research objectives. We now present an outline of the remainder of the thesis, followed by a summary of our original contributions.

Chapter 2 provides background material for this research. This includes a description of the basic principles of active sonar and the associated signal processing needed for target detection. Some of the ideas discussed are not used explicitly in later chapters,

but are included in order to firmly establish our work within the field of sonar signal processing.

Chapter 3 describes the trial data provided by Thales Underwater Systems. The chapter begins with a description of how the data was collected, and the format in which the data was recorded. The contents of the trial data are then briefly discussed, followed by the methods needed to process the data in terms of GPS coordinates.

Chapter 4 begins by introducing a number of statistical models for sonar noise. We then describe generic algorithms for parameter estimation, and identify appropriate estimators for the various noise models. This is followed by an overview of random number generation, and the processes by which synthetic sonar noise can be generated.

Chapter 5 contains a statistical analysis of the sonar noise in our trial data, based on the models identified in Chapter 4. This includes an analysis of the previously defined parameter estimators for the K-distribution, which are compared in terms of variance and bias, to identify the most appropriate estimator for use in this thesis. In order to evaluate how accurately each model matches the trial data provided by Thales Underwater Systems, we describe and apply goodness-of-fit testing. The chapter is then concluded by evaluating our findings, and identifying the most accurate models of the noise in our trial data, which are of critical importance to subsequent chapters.

Chapter 6 describes statistical models for targets in a sonar environment. Initially, we examine models for sonar noise which can incorporate an additional signal. We then examine models for target signals which can be embedded in these noise models. Lastly, as with Chapter 4, techniques are presented which enable us to generate synthetic sonar targets according to these models.

Chapter 7 contains the definition and evaluation of various detection schemes, applied to both synthetic and real data. The chapter begins by describing the techniques used

in our performance analysis, including the number of detection tests to evaluate, and practical considerations to avoid bias. The standard box-car detector from Chapter 2 is reintroduced, and extended to incorporate any noise model of our choosing. We then outline the theory of *Generalized Likelihood Ratio Tests* (GLRTs), and design various GLRT detectors. These GLRT detectors are compared using synthetic data from Chapters 4 and 6, which enables us to select those with the most promising performance. We then apply a selection of box-car and GLRT detectors to simulated noise environments with simulated targets, and analyse their performance. This analysis is then extended to their application in real sonar noise, where they are applied to simulated targets embedded in the trial data. Lastly, these detectors are applied to a new dataset with real targets, to which Thales Underwater Systems provided limited access in a secure environment.

Chapter 8 describes methods for processing data from more than one ping. This begins with the derivation of a novel method of parameter estimation, whereby noise estimates calculated from more than one ping of a region can be combined. This method is compared against an existing technique, and its advantages demonstrated. The method is then applied to the generation of a clutter map, which visualises the sources of clutter in our trial data. We then define a framework for the detection of sonar targets over multiple pings, applicable to both box-car and GLRT detectors, and capable of processing static or moving targets. This detection framework is applied to simulated targets embedded in our trial data, and briefly evaluated against real targets in real sonar noise.

Chapter 9 provides the conclusion to this thesis. We recap our contributions and examine the limitations of this research, including suggestions for future work.

1.5 Contributions

This thesis includes the following contributions:

1. We demonstrate that the Kolmogorov-Smirnov test, commonly used in sonar signal processing literature, is a poor choice when evaluating noise distributions in the context of standard detection algorithms. Standard detection algorithms require a good fit to the tail of the noise distribution in order to achieve reliable results. The upper-tail Anderson-Darling test is proposed as a suitable alternative.
2. We apply goodness-of-fit tests to long-range sonar trial data provided by Thales Underwater Systems. We show that the K-distribution provides a good fit to this trial data, and compute tables of critical values for the upper-tail Anderson-Darling test applied to K-distributed data.
3. We extend the standard (Gaussian) box-car detection algorithm to arbitrary noise distributions. This facilitates the application of an environmentally adaptive threshold, which varies with the local noise statistics for any chosen noise distribution. Using this algorithm, we compare the performance of various noise distributions, and find that the K-distribution is the only model of sonar noise which achieves increased probability of detection over the standard (Gaussian) box-car algorithm.
4. We define and evaluate detectors for targets in sonar data based on the *generalized likelihood ratio test*, making use of appropriate distributions for sonar noise, and the Rician distribution for targets. These detectors are also applied within a framework for multi-ping processing. We show that our GLRT detectors can provide a significant increase in detection performance, providing that an accurate model of the target is available.

5. We propose and evaluate a new method for combining estimates of the K-distribution shape parameter. This method has low computational cost, and unlike the method proposed by Abraham and Lyons [4], is not restricted to a particular parameter estimation method. We also demonstrate its ability to achieve reduced variance in the parameter estimate, compared to both direct estimation and the method proposed by Abraham and Lyons [4].
6. We describe a method for the generation of high-resolution clutter maps from sonar trial data. This method is based on previous work by Prior [106], and Abraham and Prior [6]. Unlike previous efforts, our method is not based on heavily sub-sampled data, and includes an iterative algorithm, that processes blocks of data for consecutive regions of the sea. We also demonstrate that if a scatterer map is based on K-distribution shape parameters, and if data is available over more than one ping, it is advantageous to combine these parameter estimates using our new method. This provides greater detail of small-scale scatterers in the clutter map, compared to clutter maps based on median sampling or the arithmetic mean.

This research has resulted in the following publications:

Refereed Journal Papers

- R. Bareš et al., “Noise estimation in long-range matched-filter envelope sonar data,” *Oceanic Engineering, IEEE Journal of*, vol. 35, no. 2, [230–235, Apr. 2010, ISSN: 0364-9059. DOI: 10.1109/JOE.2009.2036947
- R. Bareš, “On combining estimates of the K-distribution shape parameter,” To be submitted to *IEEE Signal Processing Letters*, 2012

Conference Papers

- R. Bareš et al., “Environmentally adaptive noise estimation for active sonar,” *UDT Europe*, Jun. 2009

Chapter 2

Background

“

I must confess that my imagination refuses to see any sort of submarine doing anything but suffocating its crew and floundering at sea.

— H. G. WELLS (1866–1946)

In this chapter we provide background material for this research. This includes a description of the basic principles of active sonar and the associated signal processing needed for target detection. Some of the ideas discussed are not used explicitly in later chapters, but are included in order to firmly establish our work within the field of sonar signal processing.

2.1 Basic Principles

The two main applications of sonar are defined in its initial use as an acronym. Sound for navigation and ranging (SONAR) is concerned with the application of *underwater acoustics* for the purpose of sensing the underwater environment. To achieve this a sonar system must include an acoustic sensor. While microphones are used for acoustic sensing above water, a *hydrophone* must be used underwater. Hydrophones operate in the same fashion to microphones, converting acoustic energy in the form of pressure waves into electrical energy in the form of a voltage. The main difference in the requirements of a hydrophone is that its *impedance* must be matched to that of water rather than air. Consequently while most common microphones operate by *electromagnetic induction*, where an induction coil vibrates in a magnetic field causing a voltage in the coil, hydrophones use a different approach. *Piezoelectric transducers* are a suitable alternative, producing a change in voltage as a result of a change in pressure. Useful materials with piezoelectric properties include certain families of crystal, such as quartz, and ceramics such as barium-titanate. Hydrophones are not restricted to piezoelectric materials; for example, *ferromagnets* produce a change in their magnetic field when pressure is applied, a phenomenon known as *magnetostriction* [12].

Electromagnetic induction can also be used to produce sound; the hardware of microphones and loudspeakers are broadly the same. Similarly, piezoelectric and magnetostrictive materials can also be used to produce sound waves, and when applied to underwater acoustics, this is referred to as a *projector*.

There are two basic types of sonar system, *active* and *passive*. *Active sonar* systems transmit a pulse of acoustic energy into the water, then listen for echoes reaching the hydrophone from any reflecting surfaces. The transmitted pulse is known as a *ping*. This technique is able to operate over long ranges, but broadcasts the location of the sonar system itself, which is often undesirable in military applications. In contrast,

passive sonar systems do not require projectors; a passive system listens for acoustic signals which are naturally radiated from objects of interest. The source of the radiated signals may be the hum of a submarine's engine, the propeller of a ship, or mammals communicating in the ocean. Passive systems are both covert, and provide additional information about the detected object. The two different types of sonar system are illustrated in Figure 2.1.




Figure 2.1: Basic principles of sonar

In the case of active sonar, the distance to a reflecting surface is easily estimated from the time it takes for an echo to reach a hydrophone. Although the speed of sound in seawater relies on a number of factors (to be discussed in Section 2.2), it is approximately equal to 1500 meters per second under normal operating conditions. The time between the transmission of a ping and its echo reaching the hydrophone is proportional to twice the distance to the reflecting surface.

Since passive systems have no control over the emission of acoustic energy, this technique does not apply. For passive sonar, the sonar system and towing vessel, known collectively as the *sonar platform*, must perform a manoeuvre in order to estimate the distance to an acoustic source. There are various techniques for achieving this; a simple method is to record the angle (*bearing*) to the acoustic source, travel to a new position, and then record the new angle to the source. Under the assumption that the source has not moved between these two recordings, basic triangulation can be used to calculate

the location of the object of interest.

In both active and passive sonar, distance alone will rarely provide a sufficient understanding of the underwater environment. The methods required to calculate the bearing to the reflecting surface (or acoustic source) are complicated (see Section 2.5). If a sonar system establishes that a target is nearby, identifying the target's location, defined by its range and bearing, is a priority for any sonar operator.

2.2 Propagation and Noise

Sonar uses acoustic pressure waves to sense the underwater environment. In this section, we introduce physical factors that affect the range and motion of these pressure waves, also known as their *propagation*. Sonar signals account for only a small proportion of the acoustic energy in the sea, and there is a substantial level of ambient acoustic energy from the environment. Furthermore, we also consider the echo of a transmitted sonar signal to be noise if it has not been reflected by a target. Consequently, the majority of the acoustic energy recorded by the hydrophones is *noise*. The factors contributing to this noise will be discussed in greater detail below.

Perhaps the simplest model of the operating environment could be specified by:

1. The environment is silent, other than the transmitted ping,
2. A ping travels at constant speed at all times, and
3. A ping travels in a straight line to and from the target.

Environmental and systematic effects challenge these assumptions, resulting in noisy sonar data which causes difficulties during processing.

2.2.1 Noise

The underwater environment is not silent, and contains a large number of acoustic sources. In sonar theory, noise can be broken into two categories, *sea-noise* and *self-noise*. As the name suggests sea-noise is comprised of acoustic sources which are external to the sonar platform and associated vessels, but are of no interest to the sonar operator. Examples of these sources may include rain falling on the sea surface, marine life and shipping. In contrast, self-noise is generated by the sonar equipment, the vessel to which it is attached, and the flow of water at the hydrophones. It should be noted that self-noise varies with the speed at which the sonar platform travels through the water and the operating frequency of the hydrophones. The main source of self-noise is caused by the propulsion mechanism used by the attached vessel. As a consequence, self-noise has maximum impact at the stern of the vessel. Together sea and self-noise form the *background noise* in which sonar signals are transmitted and received, complicating the detection process.

In active sonar, reflections of the ping from surfaces which are of no interest to the sonar operator are also considered to be noise. These additional sources are broken into two categories, *reverberation* and *clutter*. Echoes from the sea floor are known as reverberation. In shallow water environments, as the sea floor is close to the sonar system, reverberation quickly becomes a limiting factor for detection. However, unwanted reflections are not only received from the sea floor. Further unwanted reflections are received from objects such as rocky outcrops or objects in the water volume, known as *scatterers*. These reflected signals are referred to as clutter. While reverberation is predominantly a problem at short ranges, and tends to vary slowly with range and bearing, levels of clutter can vary quickly and can also occur at greater distances. The noise sources affecting active sonar are summarised in Figure 2.2.

A further concern for active sonar is the amount of acoustic energy reflected by a




Figure 2.2: Summary of noise sources in active sonar

surface. This is known as the *reflection index* of the surface, and depends on the physical composition of the material. For example, a man-made vehicle consisting of a metal shell with a hollow interior reflects a large fraction of any incoming acoustic energy, and therefore has a high reflection index. In contrast, sand or mud would absorb a large amount of the acoustic energy and have a correspondingly low reflection index. However, dense structures such as rock can also have a high reflection index, which leads to difficulty in distinguishing between natural and man-made objects in the ocean.

2.2.2 Propagation Loss

There is often very little, if any, prior knowledge of the environment in which a sonar system operates. Because of this, the most common form of active ping transmission is *omni-directional*, where the acoustic energy is radiated isotropically, which allows an active sonar system to survey the environment in every direction. Consequently, the acoustic energy transmitted from the projector is immediately spread outwards in a sphere, resulting in a low intensity wave in any single direction. This is known as *spherical spreading*, and is one of the factors which causes echoes reaching the hydrophone to be far quieter than the originally transmitted signal. The surface area A

of a sphere with radius r is given by

$$A = 4\pi r^2, \quad (2.1)$$

so the decrease in acoustic energy in a particular direction is proportional to $1/r^2$, quickly leading to large losses as range increases. This effect is less pronounced in shallow waters, where the sea boundaries are relatively close to the transmitter. In shallow water the sonar environment would be better modelled as a cylinder, with acoustic energy spreading laterally as shown in Figure 2.3. In this case, the spreading loss is related to the surface area A of a cylinder with radius r and height h , which is:

$$A = 2\pi rh. \quad (2.2)$$

As the decrease in acoustic energy is proportional to a coefficient of $1/r$, the power




Figure 2.3: Comparison between spherical spreading in an open environment, and cylindrical spreading in a shallow environment

loss due to spreading has a reduced impact in shallow water compared with deep water environments. Despite this, the acoustic energy still decreases significantly as range increases.

Though projectors exist that are not omnidirectional, they too are prone to some form of spreading loss. Some of the acoustic energy is also absorbed by the water in the form of heat, and reflected away by scatterers in the sea volume. For both active and passive

sonar, these factors are still a concern at the target, where spreading and absorption losses occur as the acoustic energy travels back to the hydrophone. Spreading and absorption losses are collectively known as *propagation loss*. It is also important to note that propagation loss is a two-way process for active sonar, due to the transmission and reception stages of the detection process. This also has an impact on reverberation received by active sonar systems. As echoes received from the sea floor are subject to propagation loss, they decrease slowly as the range increases.

2.2.3 The Speed of Sound in Water

As noted in the previous section, the speed of sound in water is not constant. Rather, the speed of sound in sea water depends on three main factors: salinity, temperature, and pressure. These factors depend on the local region, the season and the current depth of the acoustic energy. This variability is best portrayed as a *sound speed profile* (SSP), which defines the speed of sound in terms of depth for a particular geographic location and time of year. An example of a SSP can be seen in Figure 2.4. As an acoustic wave travels deeper into the water, changes in the speed at which sound travels can result in refraction, causing the acoustic wave to bend upwards or downwards. This can cause effects such as *dead zones* where no acoustic waves are able to travel, *convergence zones* where a large number of acoustic paths meet in the same location, and *wave guides* within which the sound wave bounces back and forth between the boundaries of thermal layers. All of these effects are most evident when operating in deep water. As can be seen in Figure 2.4, there is only a slight change in speed over the first 200 meters of depth. Though these effects have little impact on this research, the interested reader can find further information in the book by A. D. Waite [136].




Figure 2.4: Example of a sound speed profile

2.2.4 Multipath

In shallow water the variability of the speed of sound is sufficiently small that we may approximate it by a constant value. However, this environment gives rise to another problem. Due to the shallow depth, there is greater likelihood of echoes reflecting off multiple surfaces, such as the sea surface, the sea floor, and other objects in the environment. These are known as *multipath* echoes and cause two main problems. Firstly, as an acoustic wave comes into contact with a surface, not all of the acoustic energy is reflected so absorption losses occur. Secondly, since the transmitted pulse and echo may not be travelling along a direct path between the sonar platform and reflecting surface, the time taken for the echo to reach the hydrophone becomes a less accurate estimate of distance. This can lead to a number of echoes arriving from a single surface at different times, with each echo having travelled along a different route, as shown in Figure 2.5. In general, the multipath echoes will take longer to arrive, and will therefore be reported as being located at a greater distance from the platform. In some rare circumstances, this may not be the case. For example, if the acoustic wave travels through rock in the sea floor, the denser medium leads to an increased speed of sound, so a multipath echo may actually reach the hydrophone *before* an echo travelling by the direct route.




Figure 2.5: Example of the path taken by a direct (red) and multipath (blue) echo received by a hydrophone

2.3 Sonar Signal Processing Chain

As stated in Chapter 1, this work is focused on the detection of targets by an active sonar system. The various stages involved in this process, from the reception of sonar echoes at the hydrophones to the output display seen by a sonar operator, make up the *sonar signal processing chain*. An example of a signal processing chain for active sonar systems is shown in Figure 2.6, where three main stages are defined.

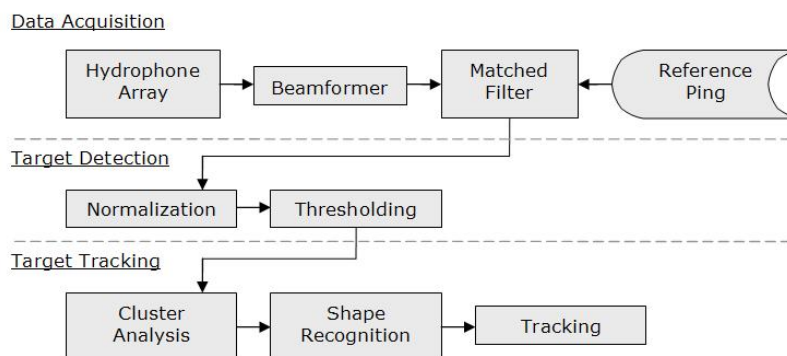


Figure 2.6: Summary of Sonar Processing Chain

During the *data acquisition* phase, data is recorded and processed into a suitable format, as used throughout the rest of this project. This data is then passed to the *detection phase* which is our main area of interest, and potential detections are then passed to the *target tracking* phase which combines detections over time before the results are presented

to an operator. From the beginning to the end of this chain, the signal processing algorithms aim to provide a reliable and effective means of detecting targets within the operating range of the sonar system.

2.3.1 Active Sonar

In active sonar we must begin by transmitting an acoustic signal into the water. This transmission is known as a *ping*, and the waveform (amplitude over time) of this ping depends on a number of factors.

As transmitted signals are subject to propagation loss, the detection range of a sonar system is related to the amount of acoustic energy transferred into the environment. This energy is dependent on the ping duration, and so increasing the length of a transmitted ping may increase the range of a sonar system.

The main frequency and *bandwidth* of the ping are important, as the effective range of a sonar system also depends on these values. Bandwidth denotes the range of the frequencies in the signal, from the maximum to the minimum. Therefore the main frequency of the system is not, on its own, a complete description of the frequency content of the ping. We will now briefly consider how absorption losses depend on the frequency of a sonar system. One model for absorption losses in sea water, described by Ainslie and McColm [11], considers three main factors. In fresh water, the main factor influencing absorption losses is the conversion of acoustic energy into heat due to the viscosity of the liquid. However, the salts dissolved in sea water also contribute to the absorption of acoustic energy. This process, known as *ionic relaxation*, is primarily caused by the salts of Magnesium Sulphate and Boric Acid. Given a depth z (km), a temperature T ($^{\circ}C$), salinity S (ppt), and frequency f (kHz), an estimate of the

absorption loss η per kilometre is given by

$$\begin{aligned}
 \eta = & \underbrace{0.106 \frac{f_1 f^2}{f^2 + f_1^2} e^{(pH-8)/0.56}}_{\text{Boric Acid Contribution}} \\
 & + \underbrace{0.52 \left(1 + \frac{T}{43}\right) \left(\frac{S}{35}\right) \frac{f_2 f^2}{f^2 + f_2^2} e^{-z/6}}_{\text{Magnesium Sulphate Contribution}} \\
 & + \underbrace{0.00049 f^2 e^{-(T/27+z/17)}}_{\text{Pure Water Viscosity}} \quad (2.3)
 \end{aligned}$$

Where

$$f_1 = 0.78 (S/35)^{1/2} e^{T/26} \quad (2.4)$$

$$f_2 = 42 e^{T/17}. \quad (2.5)$$

We remark that the operating frequency of the sonar system (f) is a factor in each of the three parts of the expression for η . Even in the limiting case of pure water, absorption losses are proportional to f^2 , thus rapidly increasing with any increase in frequency.

The consequences of (2.3) are encountered in the next chapter, where we observe that the maximum range of the system decreases as the operating frequencies increase. Furthermore, depending on its material composition, a surface will reflect a different percentage of acoustic energy at each frequency. The outcome of these two effects are that low frequency systems suffer from far greater levels of reverberation and clutter, though they are capable of detection at a much longer range. Conversely high frequency systems have high levels of attenuation, reducing the impact of reverberation and clutter, but due to the decreased wavelength, high frequency sonar is also capable of greater angular resolution when applied to short range detection.

Two main types of ping exist, *continuous wave* (CW) pings which provide a constant tone, and *frequency modulated* (FM) pings which involve a gradual sweep across a

range of frequencies. When using a CW ping the frequency-content of the ping does not vary with time, and the *range resolution* of the sonar system will be $cT/2$, where c denotes the speed of sound in water and T denotes the duration of the transmitted ping. FM pings have the advantage that their bandwidth can be used to compress the received signal. With FM pings, the frequency-content of the ping does vary with time, which results in *pulse compression* while correlating the received signal with the transmitted signal, as the auto-correlation function is short compared to the duration of the transmission. The resolution of a sonar system employing FM pings is $c/2B$, where B denotes bandwidth. Therefore, unlike CW pings, with FM pings the duration of the transmission can be increased to transfer more energy into the environment while maintaining the sonar systems range resolution.

Although the frequencies of a transmitted ping are known, we cannot know whether the surface from which the echo returns is stationary. If the surface is moving, then the frequency of the echo will be subject to *Doppler shift*, and will no longer match the original waveform. In the case of a hyperbolic FM ping, where the frequency of the ping increases hyperbolically over time, this Doppler shift results in a shift along the hyperbolic curve of the original frequencies, and so a large portion of the received echo will still match the transmitted signal, as shown in Figure 2.7. This makes doppler-shifted hyperbolic FM pings easy to process, but the shift in the frequency spectrum can result in a shift in the estimated range of the received signal, and as Doppler information is lost we cannot easily calculate the speed of a target. The need for greater Doppler information is met by CW pings. CW pings do not vary in frequency, and therefore have the same *frequency spectrum* for the duration of the transmission. The frequency spectrum of a square CW pulse is shown in Figure 2.8. If the echo of a CW ping is subject to Doppler shift, the transmitted and received signals will not be well matched. Therefore, in order to match the transmitted and received signals, a bank of Doppler-shifted copies of the transmitted signal is often used. By comparing received signals to this bank of reference signals, the received signal can be tested for echoes of the

transmitted ping while simultaneously estimating the extent of the Doppler-shift.

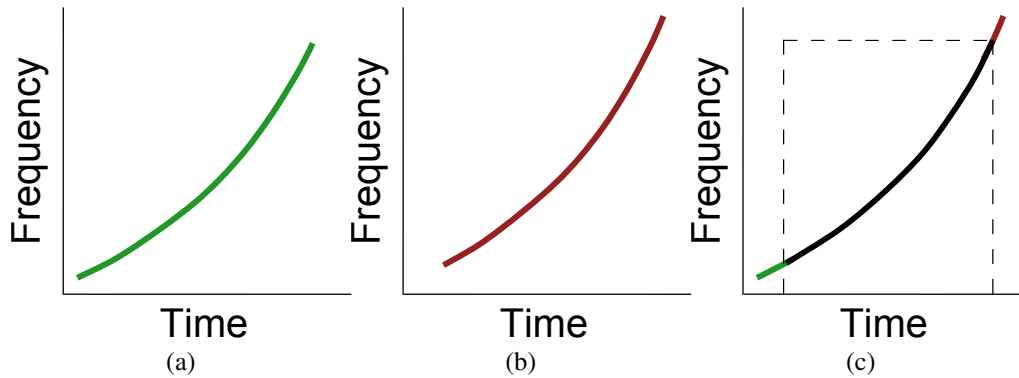


Figure 2.7: Example spectrogram of (a) transmitted FM ping, (b) received echo from an FM ping with doppler shift, (c) large overlap between transmitted and received signals




Figure 2.8: Example frequency spectrum of a CW pulse

After a ping has been transmitted, the hydrophones are switched on and begin receiving echoes. These signals are first passed through amplifiers on their way to the beam-forming and subsequent matched-filtering stages of the processing chain. Here the raw acoustic waveforms (recorded by the hydrophones) are separated into a set of echo amplitudes and the bearings from which they arrived. In this form the data is processed by detection algorithms, which attempt to discriminate between artefacts due to sonar noise and any potential targets. Having estimated the range of each potential target (from the time of arrival and the estimated speed of sound), these potential targets are passed through automated target tracking algorithms before being displayed for the sonar operator.

2.3.2 Passive Sonar

Although this project is concerned with active sonar systems, we provide a brief account of the differences between active and passive systems and provide an example of passive sonar data.

Passive sonar aims to detect sounds which are emitted by the target itself. For military applications, such sounds might include the internal machinery of a vessel, or the sound of its propeller blades. In other applications the acoustic signals could include marine wildlife of particular interest to fishing or conservation operations. These signals are generally not intended to reach the sonar hydrophones. Due to a lack of control over the frequency and timing of passive signals, detection in passive sonar is a challenging task.

The signal processing chain for a passive system shares some of the equipment used by active systems. Data from multiple hydrophones are passed to a beamformer in order to establish the bearing from which the sound was emitted. In active sonar, the transmitted signal is much louder than most sources of background noise, and so natural acoustic signals are relatively weak compared to echoes from the ping. In contrast with active sonar systems, passive targets emit far weaker signals, and so natural acoustic signals have a large impact on passive data. A useful method of analysis is to calculate the frequency spectrum of each beam. A *spectrogram* is obtained by plotting the amplitude of each frequency against time, which is then compared to a large bank of spectrograms for known acoustic sources. An example of a spectrogram can be seen in Figure 2.9, where strong vertical lines represent high amplitude frequencies. The process of target detection is then reduced to a process of pattern-matching, which not only identifies whether the source is a valid target, but in many cases provides a more exact identification. Unfortunately individual spectrograms can not help us localise the target.




Figure 2.9: Simulated spectrogram of passive sonar data. Courtesy of T. Wood [139]

As passive systems have no control over transmission time, it is not possible to estimate distance using only the time of arrival alone. Consequently, the location of a possible target has to be determined by estimating the bearing to the target, moving the sonar system to a new location, and estimating the new bearing to the target. With knowledge of where the operating vessel has travelled, the range can then be calculated by applying basic triangulation. It is also possible to estimate bearing by switching between spectrograms on different beams, looking for the beam with highest amplitude. Normalizing and thresholding the data is a simple method for identifying any particularly strong signals which should be investigated. By plotting the location of these signals in bearing and time, it is possible to monitor the angular motion of an acoustic source relative to the sonar system.

In passive sonar, it is also necessary to consider noise from the environment. Rather than reverberation, a significant concern is *transient* acoustic signals from shipping activity, marine mammals, and natural geological phenomena. All of these sources of

noise must be accounted for within the signal processing chain.

2.4 Hydrophone Arrays

A hydrophone array consists of two or more hydrophones arranged in a fixed pattern. Additional hydrophones can be used to increase the sensitivity of a sonar system, by combining the received signal by each hydrophone. In the same way that using multiple loudspeakers can provide directivity in terms of stereo or surround-sound, multiple hydrophones can be used to determine the direction from which an acoustic wave has travelled. This process is known as *beamforming*, described in Section 2.5.

In the design of a hydrophone array, the first issue is to decide the number of hydrophones to include, which is one of the main factors determining the sensitivity and angular resolution of the system. A further consideration is the pattern (or shape) formed by the hydrophones. A variety of different array shapes exist; the most common of these are depicted in Figure 2.10. We will now discuss each of them in greater detail.

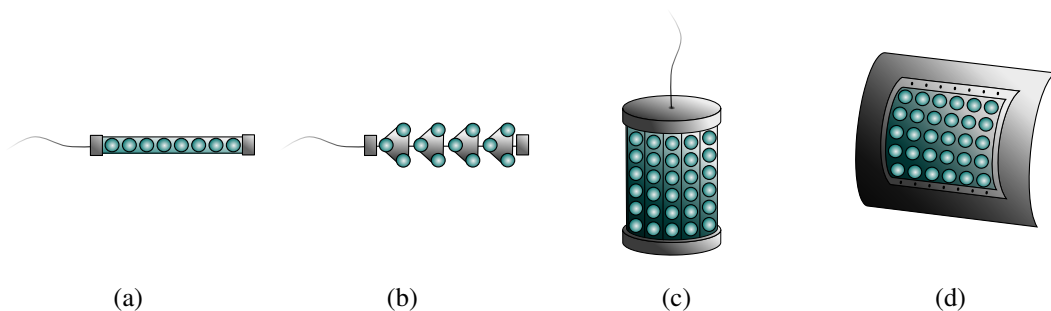


Figure 2.10: Examples of hydrophone arrays: (a) towed, (b) cardioid, (c) cylindrical, (d) hull-mounted

Towed Arrays

Towed arrays are usually a linear array of hydrophones and are deployed behind the towing vessel. A towed array is suitable where the number of hydrophones and their spacing make it prohibitive to attach the array to a vessel. Additional hardware can be used to manipulate the operating depth of a towed array, which would not be possible if the array were mounted to the hull of a surface ship. De-coupling the sensitive hydrophones from the towing vessel also has the advantage of reducing the effect of self-noise. However, as this array shape is one-dimensional it is not possible to determine from which side the acoustic signal originated. This is known as *left-right ambiguity*. Traditionally this ambiguity is resolved manually by performing a manoeuvre with the towing vessel and observing how the sonar output changes in response. Modern alternatives include the use of *cardioid* or *triplet* arrays, where each of the original hydrophones is replaced by three hydrophones in a triangular composition. The addition of hydrophones on the left and right hand side allows the system to resolve the left-right ambiguity. More information can be found in the 2006 paper from Baldacci et al. [16].

Two-dimensional Arrays

In contrast to linear arrays, the two-dimensional nature of cylindrical and hull-mounted arrays provide the sonar system with two axes of angular resolution, which allows for a three-dimensional exploration of the underwater environment. Cylindrical sonar arrays are particularly well suited for use as a *dipping sonar*. Their 360° coverage and decoupled design allow them to be winched into the water from a helicopter. This is typically of use in anti-submarine warfare (ASW) operations. An example of this application is the AN/AQS-13 dipping sonar used by the Sikorsky SH-3D Sea King ASW helicopter, shown in Figure 2.11.

Though decoupling the array from the vessel reduces self-noise, it introduces some problems. For navigation purposes, it is necessary to detect the topography of the ocean in the vessels immediate vicinity. In this case deploying a sonar array away from the ship would be of reduced benefit, and a hull-mounted sonar system would be most appropriate.



Figure 2.11: An AQS-13 dipping sonar being lowered from an SH-3D Sea King helicopter from Helicopter Anti-submarine Squadron 2 (HS-2). Courtesy of the U.S. Department of Defense [103]

For a technical overview of sonar array design the interested reader should examine the work of A. D. Waite [136].

2.5 Beamforming

In this section we provide a description of beam patterns and the methods required to determine the angle from which an acoustic signal has arrived. This technique is known

as *beamforming*.

The beam patterns presented here are essentially a map of a hydrophone array's sensitivity to acoustic energy arriving from each bearing. In the example shown as a *flower plot* in Figure 2.12, the *main lobes* (to which the array is most sensitive), are at 0° and 180° . The array is also sensitive at other bearings, which are known as *side-lobes*.




Figure 2.12: Basic beam pattern. Courtesy of A. D. Waite [136]

By manipulating the phase or timing of the signal received at each hydrophone, the beam pattern of an array can be *trained*, in order to concentrate the array's sensitivity in a chosen direction. This process is applied by a *beamformer*, which executes time or frequency domain methods to simultaneously produce a set of beams, with each beam having high sensitivity in a particular direction. An example of such a calculation is shown in Figure 2.13. For a linear array, to train the main lobe of the beam pattern to some angle θ , we must consider the time taken for a wave arriving from this angle to reach each of the hydrophones in turn. If the spacing between hydrophones along the array is d , the additional distance travelled to reach each successive hydrophone (d^*) is calculated by basic trigonometry as:

$$d^* = d \sin \theta \quad (2.6)$$

If d^* is measured in metres, and estimating the speed of sound in water as 1,500 metres per second, the additional time τ (in seconds) it takes for the signal to reach each hydrophone can be estimated as:

$$\hat{\tau} = \frac{d \sin \theta}{1500}. \quad (2.7)$$

Let each hydrophone be numbered from $i = 1 \dots n$, with angles measured such that $i = 1$ is 0 degrees from the centre of the array and $i = n$ corresponds to 180 degrees. To train the main lobe of the array to receive a signal from angle θ , the signal received from each hydrophone should be processed with a respective time delay of $\hat{\tau}(i - 1)$ seconds.

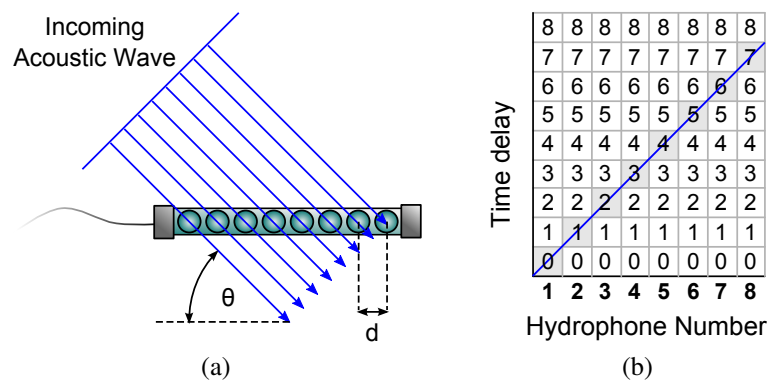


Figure 2.13: Example of time-domain beamforming (a) Scenario with an incoming acoustic wave at 45° from an array (b) Time delay (as a multiple of τ) applied to the signal from each hydrophone, in order to train the main lobe of the arrays' beam pattern in the direction of the incoming wave

Beamforming is applied after data is collected from the hydrophones. Thus it is possible to perform this step in parallel, to generate the output from all beams at once. This is useful, as a large number of beams are required to accurately distinguish the bearing of an incoming echo. Although it appears that an increase in the number of beams must increase the angular resolution of the sonar system, there is a limiting factor in this regard. For example, the main lobe in Figure 2.12 is rather wide. Improvements in resolution are dependent on *beamwidth*, which eventually leads to diminishing returns

as successive beams increasingly overlap each other (see Figure 2.14). The greater the overlap between adjacent beams, the greater the correlation between the received signal along each beam.




Figure 2.14: Example of diminishing returns when adding further beams (a) Original beams (b) Further beams included, visibly overlapping the original beams

If a beam is trained in a chosen direction, the main lobes of the beam pattern will align with this direction of interest. Unfortunately, due to the presence of side-lobes, data recorded by the array will not be perfectly isolated from acoustic energy at other bearings. Furthermore, when the array is trained in a new direction, these side-lobes will change with the bearing of the main lobe, as shown in Figure 2.15. Care must be taken to minimize the side lobes, so that one can be sure that data returned by the beamformer is indeed from the intended bearing. With this in mind, the beamformer can attempt to reduce these side lobes by *shading* the received data [136]. Applying a shading function has the effect of reducing side-lobes while simultaneously increasing the width of the main lobe, and decreasing the array sensitivity. Shading must be applied carefully to ensure that the desired level of angular discrimination between beams is maintained, as well as ensuring adequate sensitivity for the desired detection performance. A number of different shading functions exist, including the Binomial and Dolph-Chebyshev functions [44].

The same principles can be used to intentionally introduce a *null* or *blind zone* in the beam pattern, resulting in directions of low sensitivity. In military applications, this may be useful for masking the effects of a sonar *jammer* - a countermeasure designed




Figure 2.15: Demonstration of beam patterns for a line array trained to 4 different bearings. Courtesy of A. D. Waite [136]

to emit a large amount of acoustic energy, intended to overwhelm nearby sonar systems and thus prevent them from operating effectively. By training low-sensitivity regions of a beam towards the source of a jammer, it is possible to mitigate some of its effects.

2.6 Matched Filters

Once the beamformer has provided acoustic data for a chosen bearing, we must consider the influence of noise on this data (see Section 2.2). If we are to process echoes with a view to detecting, localising or characterising some reflecting surface, it would be advantageous to reduce the influence of noise.

In active sonar, the exact form of the transmitted ping is known, and can be used to detect the presence of echoes of itself. Let us assume that we have received a vector of data X from the beamformer, for a transmitted signal P . In this case the matched filter is the best linear approach [132] to maximising the *signal-to-noise* ratio of the received signal. In other words, the matched filter will increase the relative amplitude of echoes of the transmitted signal, or equivalently, decrease the relative amplitude of noise sources.

2.6.1 Matched Filter in the Time Domain

Matched filters make use of *convolution* to calculate the correlation between a transmitted signal and a received signal. This calculates the similarity between the two signals as one is shifted over the other. If the two signals are similar at a particular location t , then the output $M[t]$ of the matched filter at this point will be large. Similarly if they differ greatly then the output $M[t]$ is small. An example of this process can be seen in Figure 2.16. Though the echo signal is barely visible in Figure 2.16c, the matched filter

output reduces the impact of the noise while successfully accentuating the signal, as shown in Figure 2.16d.

Given a received signal X and transmitted signal P , the matched filter is defined as the convolution of the received signal with a conjugated time-reversed version of the transmitted signal. This can be expressed as follows:

$$M[t] = \sum_{i=-\infty}^{\infty} H[t - i]X[i], \quad (2.8)$$

where $H[x]$ is the time-reversed complex-conjugate of P . As can be seen, the transmitted ping P slides over the received signal as the convolution is calculated. Thus in temporal regions where the received signal is a close match to the transmitted signal, the output will be higher than regions where the signals differ greatly. This has the effect of estimating the relative amplitude of any echoes, minimizing the influence of noise sources. However, it should be noted that strong noise sources may still cause large values, and so the presence of a peak in the matched-filter output is not sufficient to assume an echo of our transmitted signal.

2.6.2 Matched Filter in the Frequency Domain

The matched filter can also be applied in the frequency domain, using the convolution theorem. Let us denote the convolution of the vector f with the vector g by $(f * g)$, the Fourier transform by $\mathcal{F}\{\cdot\}$ and point-wise multiplication by \otimes . The convolution theorem states that

$$(f * g) = \mathcal{F}^{-1} \{ \mathcal{F}\{f\} \otimes \mathcal{F}\{g\} \}. \quad (2.9)$$

This is an attractive representation as pointwise-multiplication is more computationally efficient than computing a convolution, and this approach is often more efficient despite the overhead of calculating Fourier transforms.

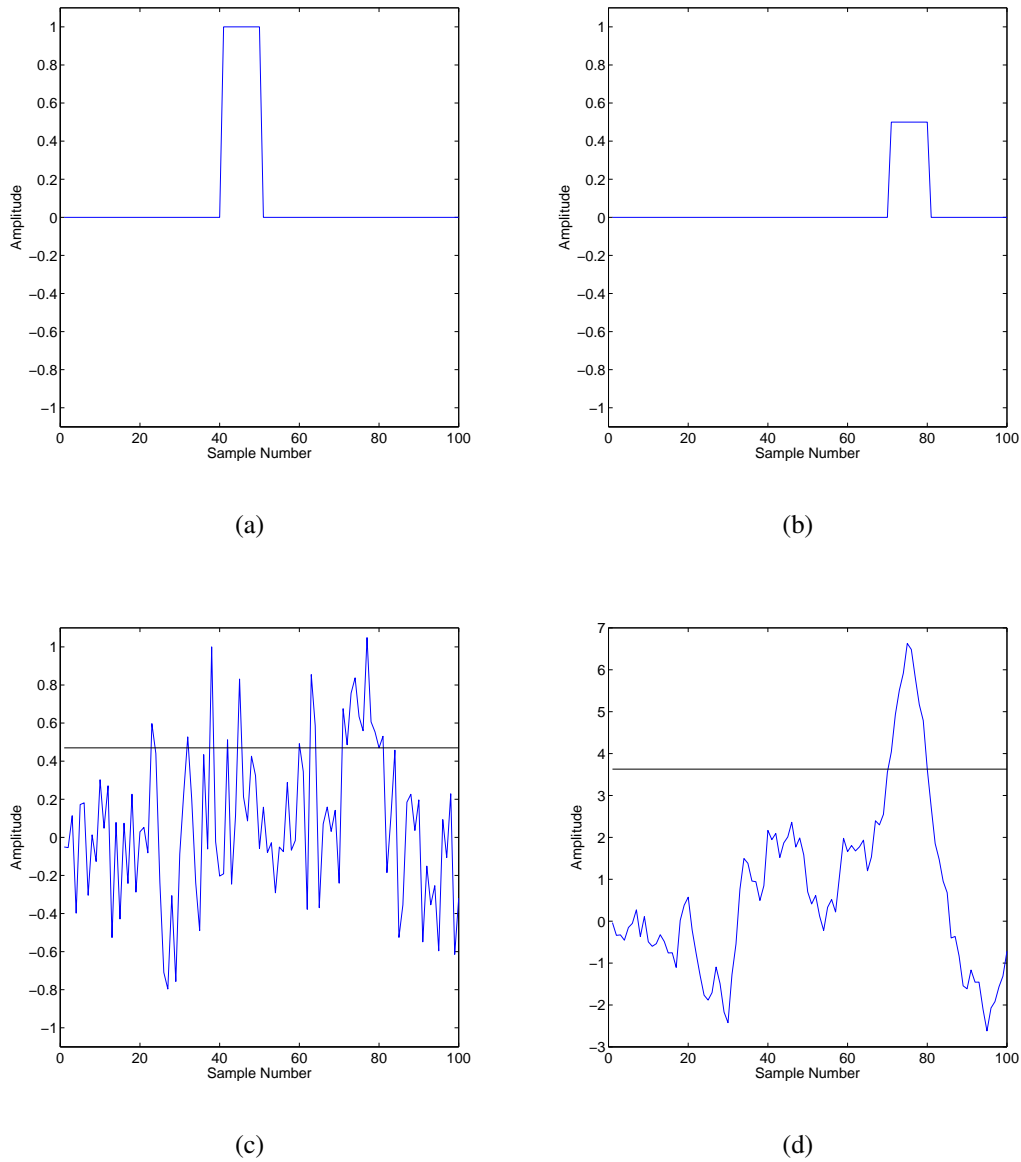


Figure 2.16: Example of a matched filter (a) transmitted signal (b) returning signal echo (c) received signal including noise (d) output of matched filter. Thresholds are set at the required level to detect the signal echo at samples 70 to 80 (see Section 2.8).

2.6.3 Matched Filter Output

When applied to complex data, such as our own, matched-filters return a sequence of complex numbers $X = \Re\{X\} + i\Im\{X\}$, containing the *amplitude* and *phase* of any echoes of the transmitted signal at a particular frequency f ,

$$\phi(f) = \text{atan2} \left(\frac{\Im\{X(f)\}}{\Re\{X(f)\}} \right) \quad (\text{Phase}) \quad (2.10)$$

$$A_X(f) = |X(f)| = \sqrt{\Re\{X(f)\}^2 + \Im\{X(f)\}^2} \quad (\text{Amplitude}) \quad (2.11)$$

where atan2 is the four-quadrant inverse tangent function. The amplitude is proportional to the signal with additive noise, while the rate of change of phase is related to the velocity of the reflecting surface and the sonar platform, due to Doppler shift.

As phase is not meaningful over large bandwidths it must be defined against specific frequencies, as in Equation 2.10. Conversely, amplitude can be calculated over all frequencies, which we then denote A_X . This is advantageous as although reflecting surfaces can exhibit different frequency responses, it is convenient to summarise the amplitude of an echo as a whole. In this research we deal exclusively with amplitude data of the form:

$$A_X = |X| = \sqrt{\Re\{X\}^2 + \Im\{X\}^2} \quad (\text{Overall Amplitude}) \quad (2.12)$$

It should be noted that the matched filters in this section have been simplified to demonstrate the basic concepts. The examples calculate the cross-correlation between the transmitted and received signals. A more sophisticated implementation makes use of a covariance matrix for the noise to take into account any prior knowledge we may have of the noise process.

By applying the matched filter to each beam, we obtain a two-dimensional matrix of

amplitudes. This matrix contains the amplitude of the echo received as a function of bearing and time. Since time can be used to estimate the distance between the sonar system and the reflecting surface, this matrix is often referred to as a *range-bearing* matrix. We refer to the location of each value in a range-bearing matrix as a cell, such that the value at cell (r, b) corresponds to the output of the matched filter at range r and beam number b . These matrices can also be used to visualize sonar data, an example of this is shown in Figure 2.17, with range on the horizontal axis and bearing on the vertical axis. The colour coding applied to this image serves to highlight differences in amplitude, ranging from low amplitudes shown in blue to high amplitudes shown in red. As can be seen in Figure 2.17, the data contains both large-scale features such as reverberation, and small-scale features including clutter. While small regions of high amplitude may be due to reflections from specific scatterers or a sonar target, large-scale features are determined by the sonar environment and the bathymetry of the sea floor (see Section 2.2).

In some cases it is possible to replace the beamformer and matched filter by a single calculation. This is an ongoing field of research, see for example the matched filter array processing method of Renomeron et al. [111] and the beamspace adaptive matched filter method of Yang et al. [141]. While these methods have the potential to improve performance when used alongside the new techniques presented here, such considerations are outside the scope of this thesis.

The output of the matched filter marks the end of the data acquisition phase (see Figure 2.6). At this stage in the processing chain, we have obtained amplitude data for each point in range-beam space, representing the presence and relative amplitude of a sonar echo at a chosen range and bearing relative to the sonar platform. This amplitude data is then passed on to the detection phase, described in the following two sections.

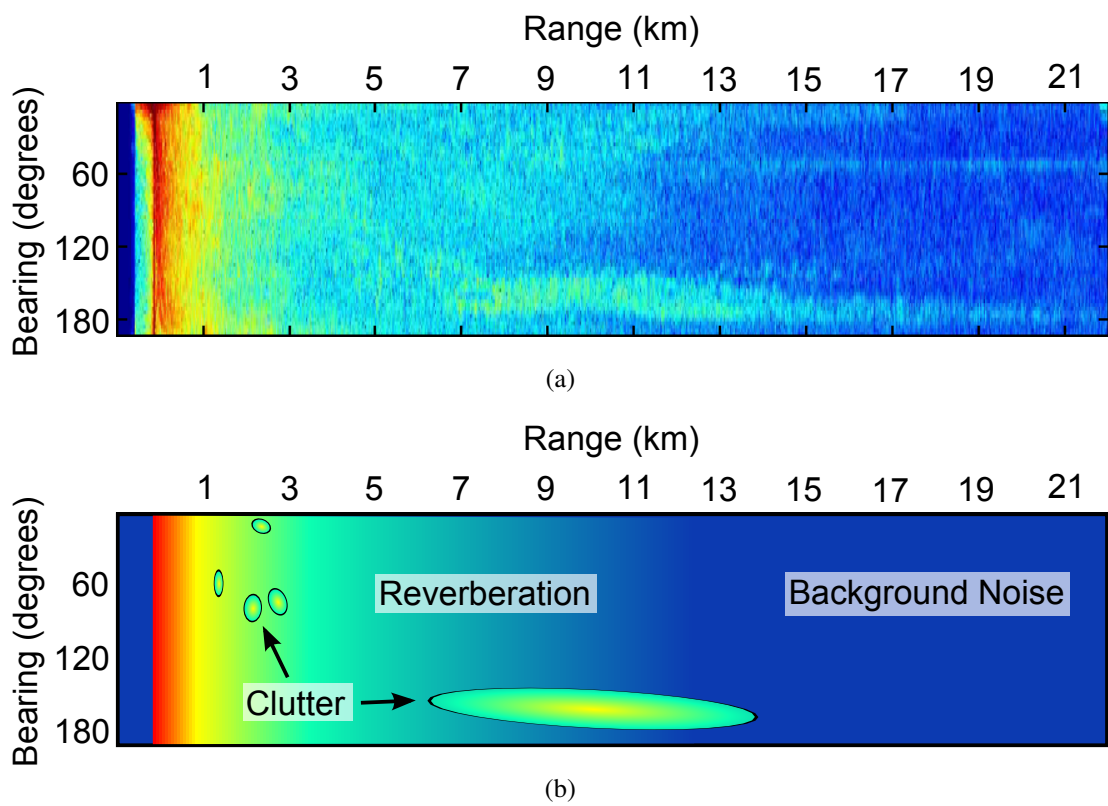


Figure 2.17: Example range-bearing plot, with and without markings to denote regions of noise

2.7 Normalization

Normalization is the first stage of the detection phase, as summarized in Figure 2.6. Detection is described by Kay [69] as the process of determining the existence of a *feature of interest*. In military applications, the feature of interest is likely to be a man-made vehicle, but this could just as easily be shoals of fish or different types of rock on the sea floor. The entire detection process should be tuned towards detecting the feature of interest in order to reject spurious results. Yet the distinction between target and non-target data is rarely clear-cut. A common approach is to find a quantitative metric which returns a large value for features of interest, and a low value otherwise.

In the previous section, we introduced the concept of a range-bearing plot. We also noted that the large-scale features in this data are caused by noise from the environment. To accurately detect sonar targets (which are a form of small-scale feature), it is necessary to remove these large gradual trends from the data, to prevent them from influencing the output of our detection algorithm. This is achieved using a process called *normalization*.

In general terms, normalization aims to reduce the complexity of data by enforcing a chosen statistical property. In photography, normalization is a process which can improve the contrast of an image taken in less than ideal lighting conditions. Equivalently, in many mathematical problems normalization is used to ensure that the solution of an equation is expressed with respect to a standard scale. An example of this is the calculation of the arithmetic mean for the n-tuple $X = \{x_1, \dots, x_n\}$

$$\hat{\mu} = \frac{\sum_{i=1}^n x_i}{n}, \quad (2.13)$$

where the denominator normalizes $\hat{\mu}$ such that it lies on the same scale as the original sample points $\{x_1, \dots, x_n\}$. By tailoring the normalization stage to our needs, it can provide the quantitative metric we need for detection. In the case of sonar signal

processing, early theoretical work in the field made the assumption that the amplitude of sonar noise was distributed according to a Gaussian distribution. Through the use of time-domain matched filters and considering sonar noise as *additive*, the received signal X is of the form:

$$X = S + Y \quad (2.14)$$

where S denotes target signal and Y denotes noise signal. If $S = 0$ and so no target signal is present, we call this *noise-only data*. However, if a target signal is present ($S \neq 0$) we call this *embedded target data*, by virtue of the target signal existing in the presence of noise. To recover the target signal from X , we must first estimate the noise Y with mean μ and standard deviation σ . If these values are known then we can reduce the noise component to a unit Gaussian through normalization.

$$\text{Let } Y \sim N(\mu, \sigma)$$

$$\left[X = S + Y \right] \sim N(S + \mu, \sigma) \quad (2.15)$$

$$\text{Let } \epsilon \sim N(0, 1)$$

$$\left[\frac{X - \mu}{\sigma} = S + \epsilon \right] \sim N(S, 1) \quad (2.16)$$

where $Y \sim N(\mu, \sigma)$ denotes that Y is distributed according to a Gaussian distribution with mean μ and standard deviation σ . By subtracting the mean and dividing by the standard deviation we have obtained a normalized signal with unit Gaussian noise.

It should be noted that while this Gaussian normalisation is justified for real Gaussian values, it is also often applied to the amplitude of complex values, despite the fact that even if these values are distributed according to a complex Gaussian, the amplitude of these values would actually be Rayleigh distributed. Thus, when we apply Gaussian

normalization we are using the equation:

$$\tilde{X} = \frac{A_X - \mu}{\sigma} \quad (2.17)$$

where \tilde{X} denotes the normalized value of A_X , μ is the mean of A_X , and σ is the standard deviation of A_X , whether or not A_X is truly Gaussian.

If the population parameters are not known, then estimates can be used for μ and σ , which we will denote $\hat{\mu}$ and $\hat{\sigma}$. These estimates should be calculated from data which is representative of the local noise environment surrounding the cell under test. However, if we attempt to calculate these estimates from the data which we are testing for a target, it is quite possible that the presence of the target signal S will corrupt the estimate of the noise parameters. In order to avoid samples of target data, we must consider the expected length of our target. In this thesis we assume that our system is attempting to detect a target of 100 metres in length. If we assume that a target is centred at the cell under test, then we can obtain noise-only samples which are not corrupted by target data at 50 metres before and 50 metres after the cell under test. It is assumed that these two regions of 50 metres around the cell under test may contain some target data, and these are referred to as the *guard bands*. If the width of the guard bands is large then there is less likelihood of corrupting the noise estimate with data including a target, however this also implies that the samples used to calculate the noise estimate are taken further away from the cell under test, so it is less likely that these samples are representative of the noise environment at the cell under test. This method of sampling embedded target and noise-only data is known as the *box-car* sampling scheme, by virtue of its graphic appearance as demonstrated in Figure 2.18. Box-car normalizers use this scheme to estimate the parameters of the noise model, moving forward in range and collecting a new set of samples at each cell.

After applying this box-car normalizer, the normalized amplitude data quantifies the number of standard deviations between the original sample value and the noise statistics

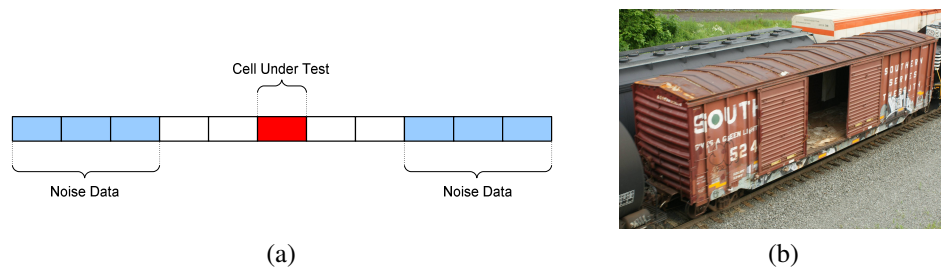


Figure 2.18: Graphical depiction of box-car sampling (a) demonstrating the location of embedded target and noise-only samples (b) compared to the style of freight vehicle by the same name, photograph courtesy of Ryan Kunkle

in its neighbourhood. Thus, a relatively large value suggests a sample which differs greatly from its local environment, such as an echo reflected from a target. Since the large-scale features seen in Figure 2.17 vary slowly between samples, they are effectively removed in the normalized amplitude data. A comparison between the original amplitude data and its normalized counterpart is shown in Figure 2.19. As anticipated, local features are maintained while the large slow-changing features of the sonar environment are greatly reduced.

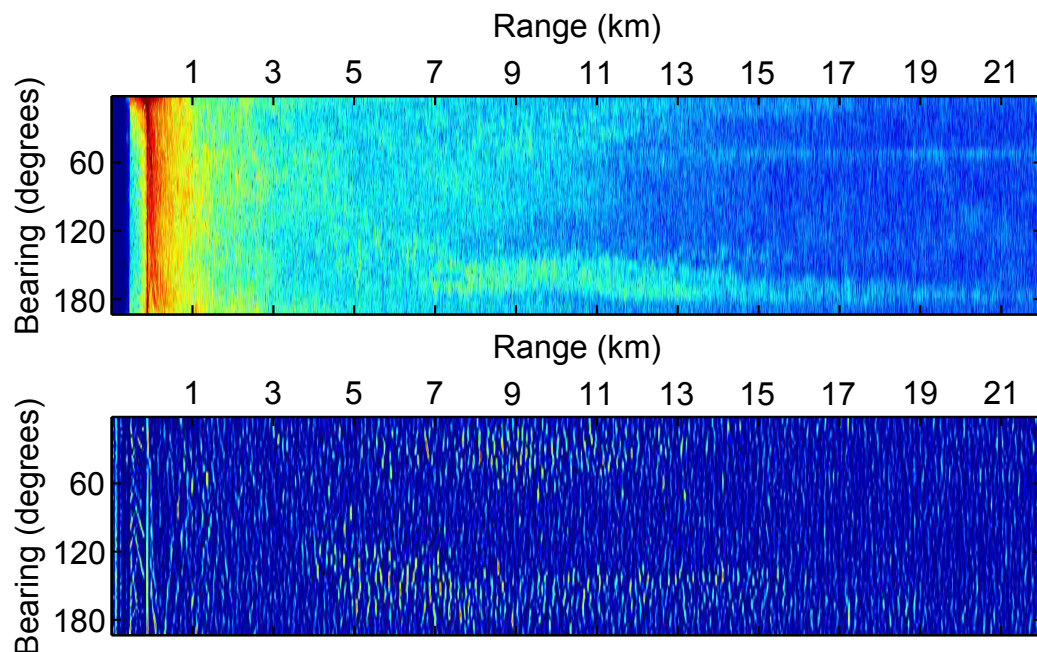


Figure 2.19: Comparison of amplitude data before and after normalization

2.8 Thresholding

The previous two sections have dealt with methods of reducing the noise in the sonar data in order to assist the detection process. By this stage in the signal processing chain, the normalized amplitude data exhibits high values in cells where the received echo is significantly louder than in the surrounding neighbourhood. As described in Section 2.7, a simple target model is an echo from the surface of the target, combined with an additive noise component. As neither the target component nor the noise component can be negative in amplitude, this implies that on average, in an environment with homogeneous noise, data containing a portion of a target is of inherently higher amplitude than noise-only data.

2.8.1 Hypothesis Testing

Let us now consider target detection in terms of *hypothesis testing*, where we wish to test the validity of one or more hypotheses according to available data. Before we can formulate a detection test we must take care in defining our hypotheses. In this scenario, a normalized sample of sonar data may either be from a noise-only environment, or otherwise contains target data. Thus, we wish to determine whether a target is present in a particular range cell (the *cell under test*). This leads to two possibilities, which in standard terminology are known as a *null hypothesis* (H_0) and an *alternative hypothesis* (H_1). Given a cell in our amplitude data, at location (r, b) , our two hypotheses can be defined as

$$H_0 : \text{the cell}(r, b) \text{ does not contain target data, i.e. } S(r, b) = 0, \text{ and} \quad (2.18)$$

$$H_1 : \text{the cell}(r, b) \text{ contains target data, i.e. } S(r, b) \neq 0. \quad (2.19)$$

We must now consider how the available data can be used to decide whether the null hypothesis H_0 can be rejected in favour of the alternative hypothesis H_1 . This is achieved by using available data, including the cell under test, to form a *detection statistic*. A detection statistic should be chosen such that changes in its value signify a departure from the null hypothesis. More precisely, a useful detection statistic will increase in value as the data in cell (r, b) exhibits properties attributable to an embedded target, and thus appears less likely to be noise-only.

To resolve the possible values of the detection statistic into a decision pertaining to the hypotheses, a *threshold* value is used. A detection statistic greater than this threshold suggests that H_0 is incorrect, in which case it is rejected in favour of the alternative hypothesis, and we conclude that the cell under test is likely to contain target data. Conversely, if the detection statistic is less than or equal to the threshold value, the null hypothesis is not rejected, and we conclude that the cell under test is not likely to contain target data. So our test is defined by the condition

$$\text{if } \left(\tilde{X}(r, b) > T \right), \text{ reject } H_0 \quad (2.20)$$

where the test statistic \tilde{X} is chosen from a cell at range r and beam b , and compared to a *threshold*, T .

Thresholds can be applied in a number of ways, see Sezgin and Sankur (2004) [116]. *Global thresholding* applies the same threshold value throughout a dataset. This is effective when embedded target and noise-only data are easily separable, for example when noise is homogeneous over the dataset and the target is very different to its background. Unfortunately this approach often breaks down in the presence of non-homogeneous data. For example, suppose we wish to detect a signal embedded in time-varying noise (see Figure 2.20). As shown in Figure 2.20c, although it is possible to detect the target signal using a global threshold, it is not possible to accurately distinguish between the embedded target and the noise-only data. A more advanced

Table 2.1: Contingency table for a detector with two hypotheses

	H_0 is true	H_1 is true
H_0 is assumed	True positive	False positive
H_1 is assumed	False negative	True negative

approach is that of *adaptive thresholding*, where the threshold is chosen depending on the local characteristics of the data, making it better suited to the problem of sonar detection. This approach is successfully applied in Figure 2.20c, where the target signal is detected above the threshold while consistently rejecting the noise-only data. Normalization converts the amplitude data to a standard form, such that threshold values are universally applicable to the data. The same adaptive threshold applied in Figure 2.20c is reapplied to the normalized amplitude data in Figure 2.20d. As the local data characteristics should now be constant, the adaptive threshold need not change between samples. Although the sample threshold value is now applied globally, this is still an adaptive detection algorithm. This is the approach applied in this thesis, where we first normalize the sonar amplitude data, and then apply hypothesis testing with a single threshold.

The calculation of a suitable threshold value is a delicate matter. If the threshold is set too low then many cells will be classed as a target and included in subsequent processing, but many of these will be artefacts due to noise only. Conversely, setting the threshold too high will reduce artefacts due to noise, but at the expense of also excluding cells which do contain an embedded target.

The behaviour of a detector must be quantified in terms of its ability to discriminate between available hypotheses. There are four possible outcomes to our testing procedure (see Table 2.1).

Given the outcome of the detection procedure, the probability of being correct and the probability of being incorrect are complementary. Therefore, the probabilities of

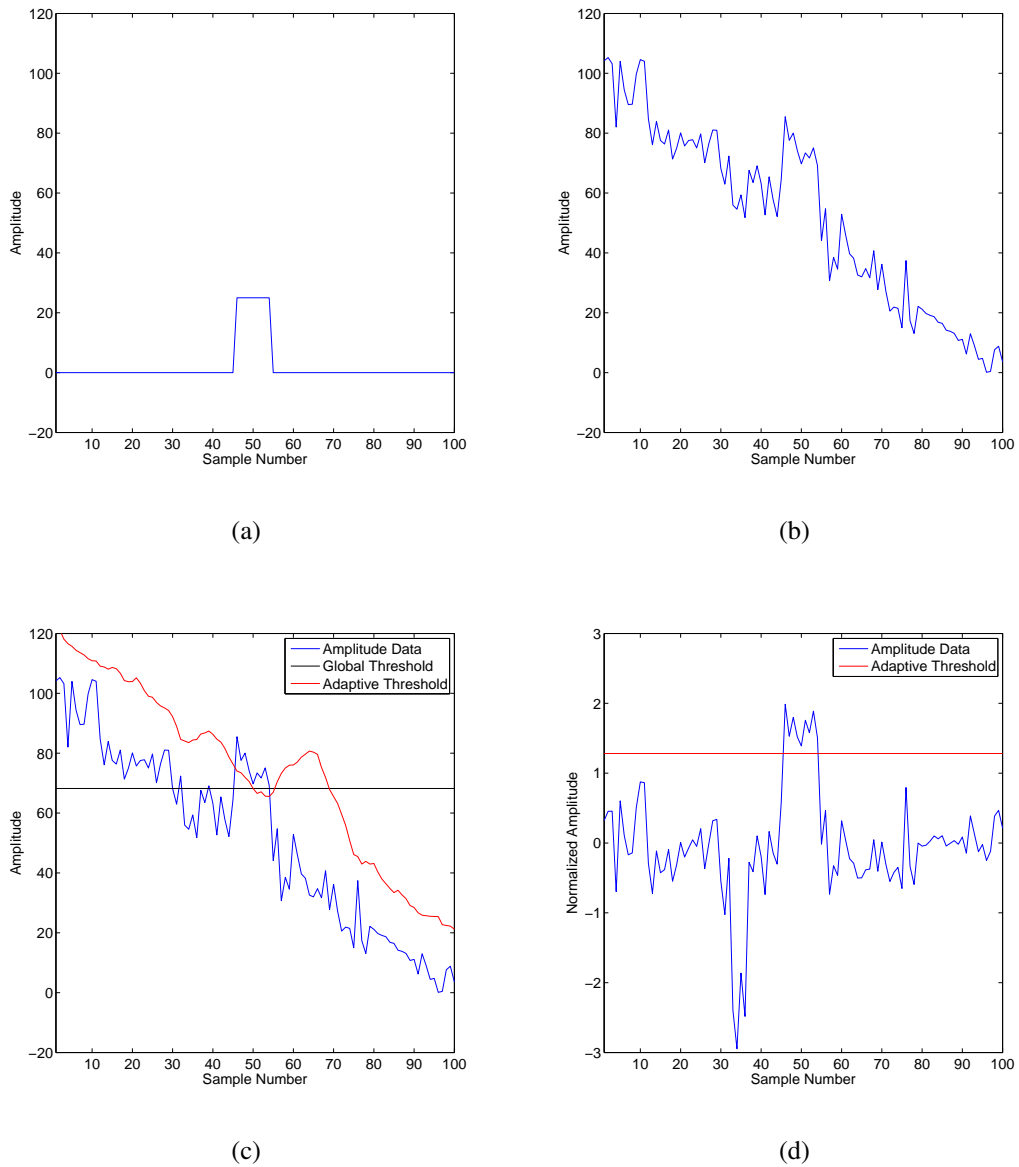


Figure 2.20: Example of global and adaptive thresholding in the presence of time-varying noise (a) target signal (b) target embedded in additive time-varying noise (c) thresholds applied to this data (d) adaptive threshold applied to the data after normalization

Table 2.2: Contingency table of probabilities for a detector with two hypotheses

	H_0 is true	H_1 is true
H_0 is assumed	P_d	P_{fa}
H_1 is assumed	$1 - P_d$	$1 - P_{fa}$

all four states can be fully specified in terms of two probabilities, one for each of the two outputs of the detector. In sonar theory, the standard performance criteria are the following:

- *Probability of Detection* (P_d): the probability with which an embedded target sample is correctly identified as a target, calculated as

$$P_d(T) = p(\tilde{X} > T | H_1) \quad (2.21)$$

- *Probability of False Alarm* (P_{fa}): also known as the *false alarm rate*, this is the probability with which a noise-only sample is incorrectly identified as a target. In hypothesis testing this is formally referred to as a *Type I Error* (False Positive), and calculated as

$$P_{fa}(T) = p(\tilde{X} > T | H_0) \quad (2.22)$$

Thus Table 2.1 can be rewritten as Table 2.2.

The selection of an optimal threshold value is therefore concerned with maximising P_d while simultaneously minimizing P_{fa} . As illustrated in Figure 2.21, both P_d and P_{fa} will vary with a change in threshold. There is an obvious trade-off between the two quantities, but the most common approach is to calculate thresholds for a *constant false alarm rate* (CFAR). As our theory assumes that the normalized data should be distributed according to a standardised distribution, the probability of false alarm α

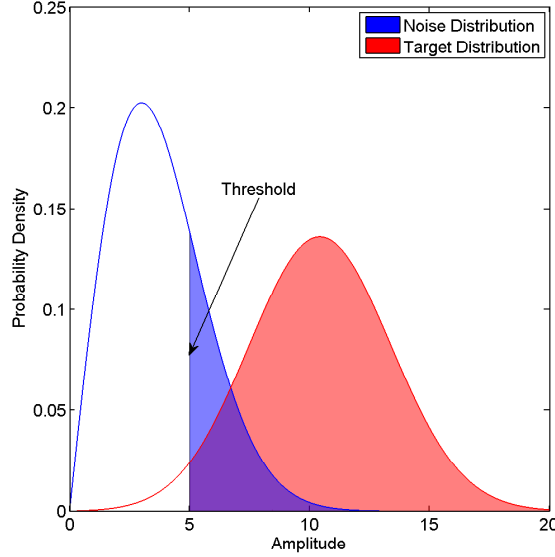


Figure 2.21: Demonstration of thresholding. The shaded areas under the graph correspond to the Probability of Detection (red) and Probability of False alarm (blue) for a threshold at an amplitude of 5

achieved by a threshold T , can be calculated by observing the following relation:

$$\begin{aligned}
 \alpha &= p(\tilde{X} > T | H_0) \\
 &= 1 - p(\tilde{X} \leq T | H_0) \\
 &= 1 - F(T),
 \end{aligned} \tag{2.23}$$

where $F(\cdot)$ is the cumulative distribution function (CDF) of \tilde{X} under the null hypothesis (noise only). For the unit Gaussian distribution, this leads to the following equation in terms of α ,

$$T = \left(\sqrt{2}\right) \operatorname{erf}^{-1}(1 - 2\alpha), \tag{2.24}$$

where $\operatorname{erf}(\cdot)$ denotes the Gauss error function. Thus, providing the data is distributed according to a unit Gaussian distribution, a suitable threshold T can be found for a constant false alarm rate α . However, if the assumption of a Gaussian distribution is

inaccurate then the use of (2.24) will be ineffective at maintaining the desired P_{fa} .

2.8.2 Performance Analysis

As the detection performance of a sonar system depends to a large extent on the selection of an appropriate threshold, we need a threshold-independent method of comparing processing methods. A sensible method of quantitative performance analysis is to generate *Receiver Operating Characteristic* (ROC) curves. ROC curves are a plot of P_d against P_{fa} , and a quick glance can provide insight into the detection performance obtained at various thresholds. The ROC curve for the embedded target and noise-only distributions in Figure 2.21 is shown in Figure 2.22a.




Figure 2.22: Example ROC curves: (a) Linear scale, (b) Base-10 logarithmic scale

In analysing ROC curves, good systems correspond to curves entering the top left corner of the graph. Such curves suggest that, in varying the threshold, it is possible to greatly increase the probability of detection for a relatively small increase in the probability of false alarm. Conversely, ROC curves close to the diagonal indicate poor performance,

with the extreme case at the diagonal itself showing that the detector performs no better than a random guess. If the ROC curve of a detector falls below the diagonal, this indicates that the detection criteria have been poorly defined (one would do better to invert the criteria, reversing this effect).

As sonar systems become increasingly sensitive, there is growing interest in very low false alarm rates. When a ROC curve is plotted on a linear scale, false alarm rates of less than or equal to 10^{-2} appear at the far left of the plot. If the graph were to occupy 100 pixels, equally spaced from a P_{fa} of 0 to 1, there would be no way to evaluate performance lower than 10^{-2} , as all of these positions would occupy a single pixel. Therefore, when evaluating the performance of a system at low false alarm rates, the ROC curve is plotted on a logarithmic scale as shown in Figure 2.22b. With the logarithmic scale provided in this example we are then able to evaluate the performance of a system at false alarm rates even lower than one millionth of a percent (10^{-8}). Though such a low P_{fa} may seem excessive, let us consider a two-dimensional array of amplitude data, containing 64 beams and 185,000 samples per beam. At a supposedly 'low' P_{fa} of 10^{-2} (one false alarm in every hundred cells) this would already result in over 100,000 false alarms. Additionally, in order to achieve only one or two false alarms in an array of this size, a P_{fa} of 10^{-7} would be required. Consequently, the ROC curves presented in this thesis are all provided with P_{fa} on a base-10 logarithmic scale.

2.9 Target Tracking

The previous stages in the sonar signal processing chain are responsible for

1. collecting data from the hydrophones, and
2. attempting to identify and localise cells containing a target.

Given a sufficiently large number of cells, even very low P_{fa} constraint can lead to a large number of false alarms. In many applications it would not be appropriate to treat every detection (including false alarms) as a definite target, and so the information is passed on to a human operator for interpretation. Though much of the received signal has been filtered out, and the detection stage has identified cells of particular interest to assist the task of identifying real targets, operator fatigue remains a serious concern. To ease this problem, *tracking* algorithms can be used to track detections over successive pings. As tracks build up, it is then easier for an operator to determine whether an object is behaving like a target.

The detection stage identifies the range and bearing of cells deemed to be of interest in the data; these are sometimes referred to as *features*. As a target will generally occupy more than one cell of data, it is often the case that many features will correspond to a single target.

2.9.1 Clustering

To reduce the complexity of the problem these should be grouped together using a *clustering* algorithm. One such algorithm is *k-means clustering* [84], which first selects a number of clusters at random locations and assigns each point to the nearest cluster. The centre of each cluster is then recomputed from its assigned points. By iteratively assigning points to the nearest cluster then moving the cluster centre to the centroid of its assigned points, this technique converges towards a natural partitioning of the data.

Each cluster can be characterised by its amplitude statistics and geometry. By testing certain criteria against each cluster, unwanted detections can be filtered out prior to tracking. For example, if we were attempting to detect a submarine, we may apply prior knowledge by excluding clusters beyond a size of say 30 by 200 meters. Such bounds

would still be sufficient to include even the largest of known submarines (the 20 by 175 metre Russian Typhoon class SSBN, which gained notoriety through the Tom Clancy novel *The Hunt for Red October* [36]).

2.9.2 Tracking

The final stage of the sonar processing chain is then initiated by passing the remaining clusters to a tracking algorithm. The typical output from such an algorithm is shown in Figure 2.23. Perhaps the simplest approach is to assume that a single target is present in a single cluster and attempt to track this cluster from one ping to the next in the presence of Gaussian noise. If we consider this from a Bayesian perspective, we might predict the targets position from available data, then improve this estimate with the data provided in subsequent pings. A tractable means of implementation is the Kalman filter, proposed by Rudolph Kalman in 1960 [67].

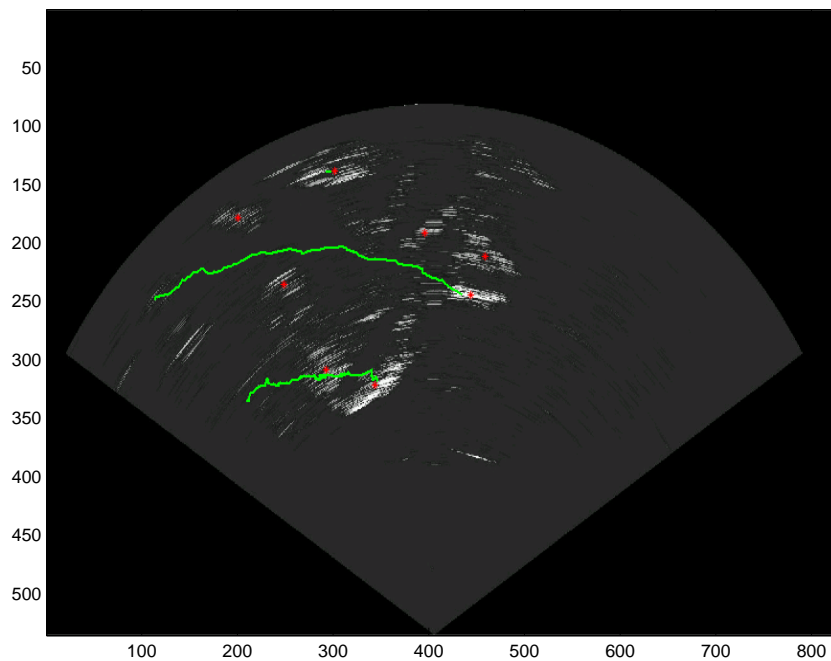


Figure 2.23: Example output from a GM-PHD tracker in cluttered sonar data. Red markers denote clusters, green lines denote tracks. Courtesy of D. Clark [38]

The assumption of a single target is unlikely to be correct, and it is probable that multiple tracks will be needed to differentiate between reverberation, clutter, and actual target clusters. This requires the application of algorithms for *multiple-target tracking* [99]. Such algorithms often attempt to reduce the problem to a single-target case. This is achieved by splitting the data into measurements for each separate track, which can then be processed by single-target Kalman filters, a method known as the *Multiple Hypothesis Tracker* (MHT) [110]. Although the multiple-target Bayesian filter is computationally intractable, its first moment can be approximated. This approximation is known as the *Probability Hypothesis Density* (PHD) filter, introduced by Mahler in 2003 [85]. The PHD filter has two possible implementations, the particle-based *sequential Monte Carlo* (SMC) PHD filter [134], and the *Gaussian-mixture* (GM) PHD filter [135].

Subsequent research has determined that the SMC-PHD filter can provide greater performance than the standard MHT when applied to generic simulated data [101]. Though Clark found in 2006 that the SMC-PHD filter could only demonstrate performance comparable to the MHT in low-noise sonar environments with high SNR targets [37], his later application of the GM-PHD filter was more successful [38]. He reached the conclusion that the GM-PHD filter can perform better than the SMC-PHD filter in high-clutter environments. That the GM-PHD filter is able to outperform MHT in the presence of reverberation and clutter was verified by Wood in 2009 [140], who succeeded in applying the GM-PHD filter to the sonar trial data used in this thesis. Although a great deal of progress has been made with the PHD filter, the problem of multiple-target tracking is still an area of active research.

In this chapter we have provided the necessary background to sonar signal processing required for later chapters. This has included an overview of active sonar, and a description of the statistical signal processing needed to perform target detection. Target tracking has also been discussed, to establish how the output of the target detection stage feeds into the later stages of the sonar signal processing chain. In the following

chapter we will introduce the trial data provided by Thales Underwater Systems and the methods needed to process this data.

Chapter 3

Trial Data

“

Data is a precious thing and will last longer than the systems themselves.

TIM BERNERS-LEE (1955–)

In this chapter we describe the trial data provided by Thales Underwater Systems. The chapter begins with a description of how the data was collected, and the format in which the data was recorded. The contents of the trial data are then briefly discussed, followed by the methods needed to process the data in terms of GPS coordinates.

3.1 Introduction

The sonar trial data used in this research was provided by Thales Underwater Systems. This included two datasets. The primary dataset contained no cooperating targets

and will be discussed in greater detail in this chapter. Unless stated otherwise, when discussing trial data we are referring to this primary dataset, which is used throughout this thesis. A second dataset was made available for use in later chapters (see Section 7.8 and Section 8.6). This second dataset contained real target data, and was therefore only accessible for a limited period of time in a secure environment.

The trial data was obtained by a towed long-range active sonar system, and gathered in a shallow water environment in the south-western approaches of the United Kingdom.

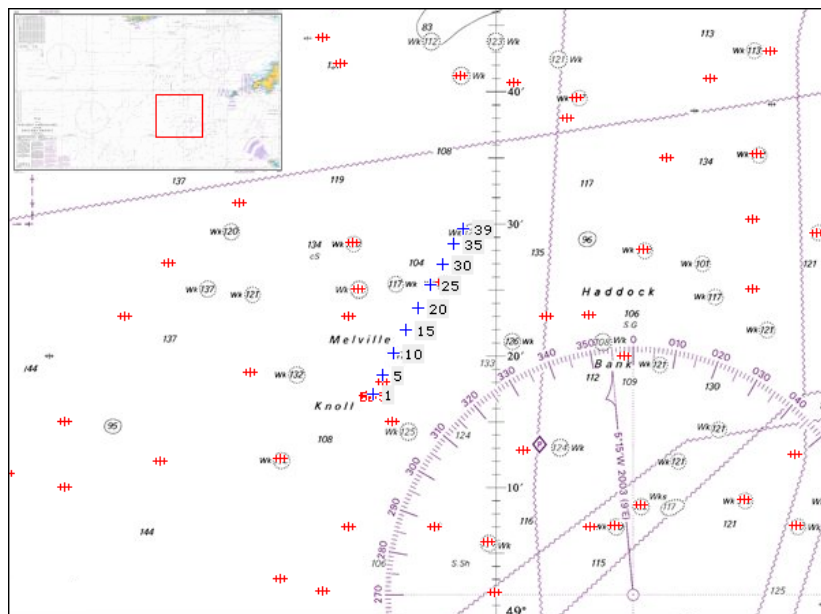


Figure 3.1: Map of the south-western approaches, with markings for the position of the towing vessel at each ping, and shipwrecks in the region

This region ($7.5^{\circ}\text{W} - 8.5^{\circ}\text{W}$, $49^{\circ}\text{N} - 50^{\circ}\text{N}$) is known to contain variable levels of clutter and a large numbers of wrecks, with a typical water depth between 120 and 150 metres. The system consisted of a towed omni-directional projector and a separately towed array, providing a nested three-octave line array. During this trial a total of 39 pings were transmitted into the water with care taken to avoid any overlap between the transmitted and returning signals from consecutive pings. The position of the towing vessel at each ping can be seen in Figure 3.1¹. For each ping, the array was towed in a straight line, the

¹Based on Admiralty Chart 2649. Reproduced by permission of the Controller of Her Majesty's Stationery Office and the UK Hydrographic Office. Additional wreck markings provided by the wrecksite.eu community database.

projector used broadband FM chirp pulses, and the hydrophone array recorded echoes with no resolution of left-right ambiguity. Range-bearing matrices were calculated by the sonar processing chain for each ping, with the data collected following the application of a matched-filter. The processing included time-domain beamforming and matched-filtering in three frequency bands. The time-domain beamformer provided 64 beams from 0 to 180 degrees in azimuth, equally spaced in cosine space (see Figure 3.2). The three frequency bands covered 900Hz to 7.2kHz in frequency. Splitting the data into these three frequency bands has allowed us to compare the performance of the sonar system at different operating frequencies and bandwidths without any increase in the number of transmitted pings. The three frequency bands were spaced along this range with respective bandwidths of 900Hz, 1.6kHz and 3.8kHz [21], which we refer to as the Low, Medium and High frequency bands respectively.




Figure 3.2: Layout of the 64 beams in our trial data, with markings for 0, 90 and 180 degrees in bearing

3.2 Data Format

Each ping was provided as a pair of two-dimensional matrices of single precision floating-point values, containing the real and imaginary components of the matched-filter output. Each matrix is approximately 64 by 185,000 values in size, corresponding to the 64 beams, and a range of 22.5km at a sampling rate of 8.22 cells per metre.

The associated amplitude arrays, described in (2.12), hold values in the range $[0, 7 \times 10^{10}]$. Because of this considerable range, and due to the presence of some extreme outliers, the amplitude data is best visualized on a logarithmic scale as shown in the range-beam plots of Figure 3.3

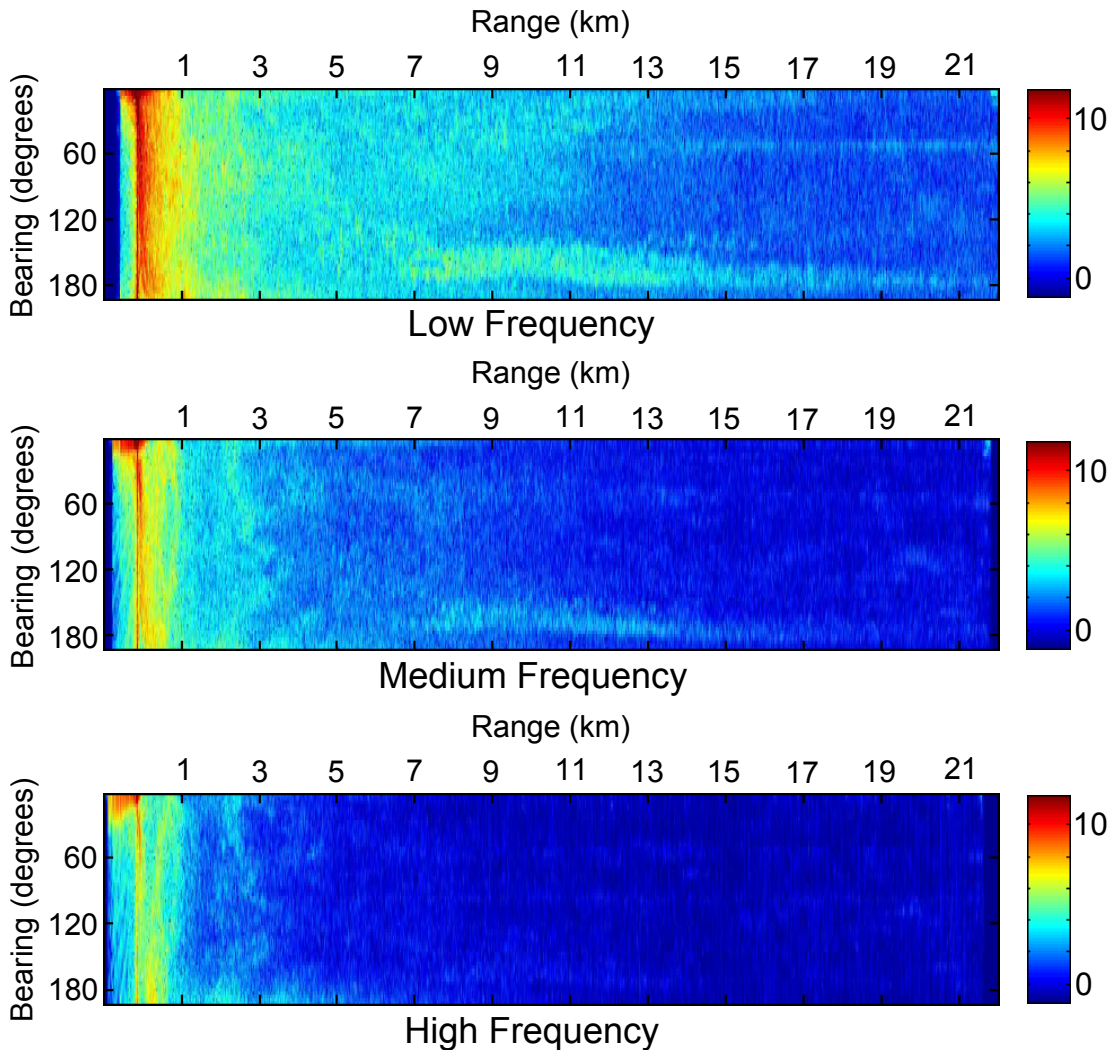


Figure 3.3: Trial data from Low, Medium, and High frequency bands, respectively. All three plots use a logarithmic scale for amplitude

The colour scale on these figures is such that red denotes high amplitude while blue denotes low amplitude. As can be seen on the right hand side of these plots, the maximum range of the sonar system is greatest at low frequency, and decreases as frequency increases. As stated in Section 2.3, low frequency bands suffer from far

greater levels of reverberation and clutter. This can be seen by the large areas of high amplitude in Figure 3.3, which also decrease as frequency increases.

3.2.1 The Initial Blast

In Figure 3.3, a large amplitude is observed in all three frequency bands at 0° bearing. This is caused by the transmitted signal travelling directly along the path between the projector and the hydrophones. The array is towed in a straight line, so the angle between the projector and hydrophone is approximately 0° . Since this route is so short, there is very little spreading. Additionally, the signal has not been reflected by a boundary or object in the sea, so there is a minimal loss in signal amplitude. This phenomenon is known as the *initial blast*, shown in Figure 3.4. The initial blast also extends across some beams adjacent to 0° , as the high amplitude signal overwhelms the beamformer and enters the side-lobes of these adjacent beams. In this project we have made the assumption that the projector, the towed array, and the towing vessel are all positioned at the same location, and so the initial blast is taken to be the position of the sonar system²; only beyond this point is useful data obtained from the hydrophones. It should be noted that the cell number for the initial blast is easily identified by searching for the largest amplitude in the matrix. This correction has been applied to the range axis of all sonar figures in this thesis, in order that the initial blast corresponds to a value of 0km in range. Recall that each matrix holds data for 22.5km. If we remove data prior to the initial blast, we find that the useful sonar data is closer to 21.5km in range.

²Though this assumption is not entirely accurate, no sensor data was available to correct for the distance between the hydrophone array and the projector.




Figure 3.4: Example of the path taken by the initial blast

3.2.2 Reverberation and Clutter

As well as the initial blast, there are also high amplitudes in Figure 3.3 caused by *reverberation* from the sea floor (see Section 2.2). This data was collected in a shallow-water environment, so there is a substantial area of sea floor at short range from the sonar system, and a large proportion of the acoustic energy is reflected directly from the sea floor surrounding the vessel. The rate at which this amplitude decreases with range is an interesting demonstration of the losses caused by spreading and multi-path echoes over short distances. As discussed in Section 2.6, the amplitude of reverberation decreases smoothly with range and therefore leads to large-scale features. However, further small-scale details are visible in these plots such as the horizontal region at a bearing of 160 degrees. This is *clutter* (see Section 2.2), and is slightly different to reverberation in that it is comprised of discrete scatterers with a higher reflection index than the surrounding region. Clutter is often found at longer ranges, where reverberation is no longer a limiting factor, and is therefore no longer masking the presence of other objects. It is exactly these two environmental factors, reverberation and clutter, which we aim to address in pursuit of improved detection in active sonar systems.

3.3 Non-Acoustic Data

So far the data discussed in this chapter has been relative to the location of the sonar platform. However, the sonar platform is known to change location between each ping. To conduct more sophisticated processing of the trial data, it would be advantageous to represent the data with respect to a fixed coordinate system, so that identical locations in the sea can be compared over multiple pings, regardless of the changing position of the sonar platform. This is best achieved with non-acoustic data, such as GPS data, digital compass data, or inertial navigation systems.

The positions of the towing vessel, shown in Figure 3.1, were provided by GPS coordinates. Throughout this project we have considered the hydrophone array and projector to be located at the GPS location recorded at the towing vessel. These coordinates make use of the world geodetic system 1984 (WGS84) as a reference framework, which models the earth's surface as a spheroid. Data represented in this manner can be converted to a two-dimensional plane by use of projection (the same techniques required to display three dimensional scenes on two-dimensional media such as computer monitors). Though we provide a brief introduction to geographic projections, further reading is available in the book by Snyder [122].


There are a number of different methods for geographic projection, each having specific properties and applications. It is therefore important to appreciate the extent to which an object is corrupted by applying an inappropriate projection technique. Various projection methods are available, each preserving specific metric properties such as relative distance, angle, or area. These methods are frequently most accurate near a chosen reference point, and it is difficult to preserve more than one geometric relation across large regions. For example, in Figure 3.5 three lines are drawn with constant bearing. Under the Mercator projection these lines appear to be linear, as this projection preserves angles. However, these lines are distorted by the Plate Carrée projection.

Consequently, projections are often classified both by the type of surface they use for approximating the spheroid, and also by their metric properties.

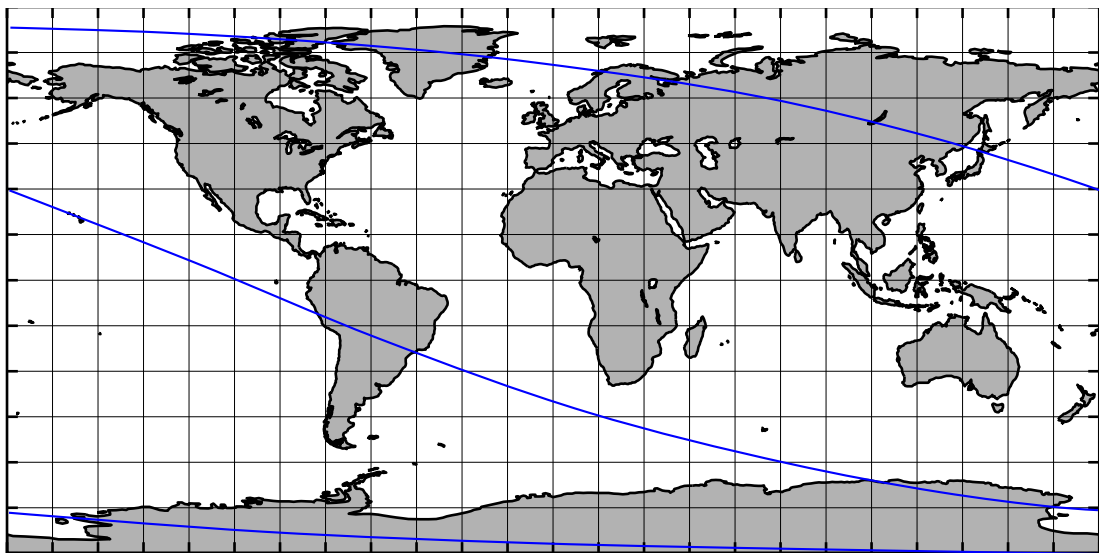
We will now examine three common projections, illustrated in Figure 3.6. The most commonly recognised projection is known as *orthographic* projection, which provides a valid perspective projection of a view of the earth from outer space. Images such as Figure 3.6a provide a close match to data provided by satellites, and is the projection used by Google Earth [51]. This method is known to preserve angles relative to the centre of the projection (the reference point used). Unfortunately, as with most projections, it suffers from singularities at the north and south poles where the meridians meet. In order to counteract this deficit, the coordinate system of the sphere can be rotated prior to projection, having the effect of moving the singularities to the equator, and thus providing a tractable means of mapping the polar regions. The orthographic projection also suffers from distortion near the edges of the image, and thus the reference point must be changed regularly.

The second of these projections, demonstrated in Figure 3.6b, is the Mercator projection. This method approximates the surface of the spheroid by projecting onto the surface of an enclosing cylinder, which is then unwrapped at the poles. Used by Google Maps, the Mercator projection is able to preserve local angles between lines, ideal for the reproduction of shapes required for the depiction of road systems [52]. It is also used by the UK Hydrographic Office for their admiralty charts, as shown in Figure 3.1. It should be noted that this method suffers from inaccurate calculations of area and angle.

Lastly, Figure 3.6c shows the Plate Carrée projection. This simple projection is similar to the Mercator projection, except that all parallels and meridians are equally spaced. This projection preserves distances along meridians, is synonymous to a Cartesian grid, and is thus a simple means of plotting GPS coordinates. Due to its simplicity, this method is only suitable over less than 100 kilometres, where the ellipsoid nature of the earth can be reasonably approximated by a flat surface.



(a)




(b)


Figure 3.5: Example with lines of constant bearing (a) preserved by the Mercator projection, (b) distorted by the Plate Carrée projection. Courtesy of Furuti [47]



(a)



(b)



(c)

Figure 3.6: Various projections of the world map, (a) Orthographic (b) Mercator (c) Plate Carrée

As the data provided in this project covers less than 50 kilometres, we have elected to apply the Plate Carrée projection in Chapter 8 of this thesis. However, if data is provided over larger distances, these distortions will become increasingly problematic and it will be necessary to implement a more accurate (but computationally expensive) projection for the required level of accuracy.

3.3.1 Coordinate Space

Beamforming has the effect of simulating multiple directional microphones, positioned at a range of pre-defined angles. As a result the data is obtained in polar coordinates, centred on the sonar array.

As the sonar system travels in a straight line over the sea floor, some environmental features are visible over multiple pings. Comparing the range-beam plots for these pings, one might expect the positions of these features within the plots to move in a straight line, relative to the sonar array. However, due to the use of polar coordinates, this change in relative position is likely to be curved, as demonstrated in Figure 3.7.




Figure 3.7: Example of how linearity is distorted by polar coordinates (a) a linear feature passing through 18 beams (b) the resulting range-beam plot

While it may seem intuitive that the 64 beams in our trial data should be spaced linearly between 0 and 180 degrees, this is not the case. The angle covered by each of the beams

is a result of the array design and the beamformer itself. In our case they are spaced equally in cosine space. In Figure 3.2 it was observed that the first beam occupied a larger angle than the beam at 90 degrees to the array. Thus we obtain greater angular resolution at 90 degrees to the sonar array (*broadside*) and a far more coarse resolution at the *bow* and *stern* of the array. This can be seen in the trial data in Figure 3.8.

The trial data in this figure has been pre-processed for use in a later chapter, but the effect of converting polar sonar data in cartesian coordinates is evident. As we are unable to resolve left-right ambiguity, the Cartesian representation has been calculated such that all data received in the original ping is taken to have originated from the right-hand side of the array. The high amplitude reverberation seen at the bottom of Figure 3.8a, closest to the sonar platform, is now positioned in a very small region encircling the hydrophone array. Comparing the first few beams of Figure 3.8a with Figure 3.8b, it is also shown that as range increases the data stored in polar coordinates occupies a greater area in Cartesian coordinates. This is because the data has a certain angular resolution, and as range increases, the area occupied by the sector also increases. This does not imply that the transmitted signal has necessarily been reflected by this entire region, but simply suggests that the signal was reflected from somewhere in the region. This effect decreases as beamwidth decreases, as can be seen for data at 90° bearing, where there is far less angular spreading in the Cartesian representation.

3.3.2 Spatial Translation

Trigonometric techniques such as the spherical law of cosines readily provide translation and rotation between GPS coordinates.

Let us make the assumption of a spherical earth, which is reasonable for trials conducted over small areas. In this case the International Union of Geodesy and Geophysics define




Figure 3.8: Figures with colour-coding depicting the same data in two different coordinate spaces (a) polar coordinates, (b) Cartesian coordinates

the mean radius of the earth to be $R \approx 6371$ km. If we then denote latitude as φ and longitude as λ , the distance r in kilometres between the current location $\{\varphi_1, \lambda_1\}$ and a further location $\{\varphi_2, \lambda_2\}$ is given by the spherical law of cosines, as follows:

$$r = R \times \arccos(\sin(\varphi_1) \sin(\varphi_2) + \cos(\varphi_1) \cos(\varphi_2) \cos(\lambda_2 - \lambda_1)) \quad (3.1)$$

In the past it has been noted that this formula is prone to rounding errors, and the *haversine* formula was recommended as an alternative [120]. However, with the introduction of IEEE 754 64-bit floating point arithmetic, these rounding errors have reduced to the extent that they are no longer a serious concern [1]. Applying the spherical law of cosines with 64-bit data types leads to typical rounding errors of less than a tenth of a metre. Given that our trial data contains approximately 8 samples per metre, the spherical law of cosines is satisfactory. If greater accuracy is required, Vincenty's formula can be used instead [133]. This sophisticated model assumes an elliptical earth and is accurate to 0.5mm, even over large distances.

The north-based bearing θ between these points can also be calculated by trigonometry:

$$\theta = \text{atan2} \left(\frac{\sin(\lambda_2 - \lambda_1) \cos(\varphi_1)}{\cos(\varphi_1) \sin(\varphi_2) - \sin(\varphi_1) \cos(\varphi_2) \cos(\lambda_2 - \lambda_1)} \right) \quad (3.2)$$

where atan2 is the four-quadrant inverse tangent function.

As can be seen in Figure 3.9, the ship travelled in an approximately straight line over a distance of 25 km. Over this region of sea floor, a planar approximation to the earth's surface should not therefore lead to significant errors. Given the calculated distance and angle between each ping, we once again have data in polar space. Converting these polar coordinates to Cartesian coordinates provides us with a relative estimate of horizontal and vertical displacement between pings, suitable for use with the Plate Carrée projection we have chosen.

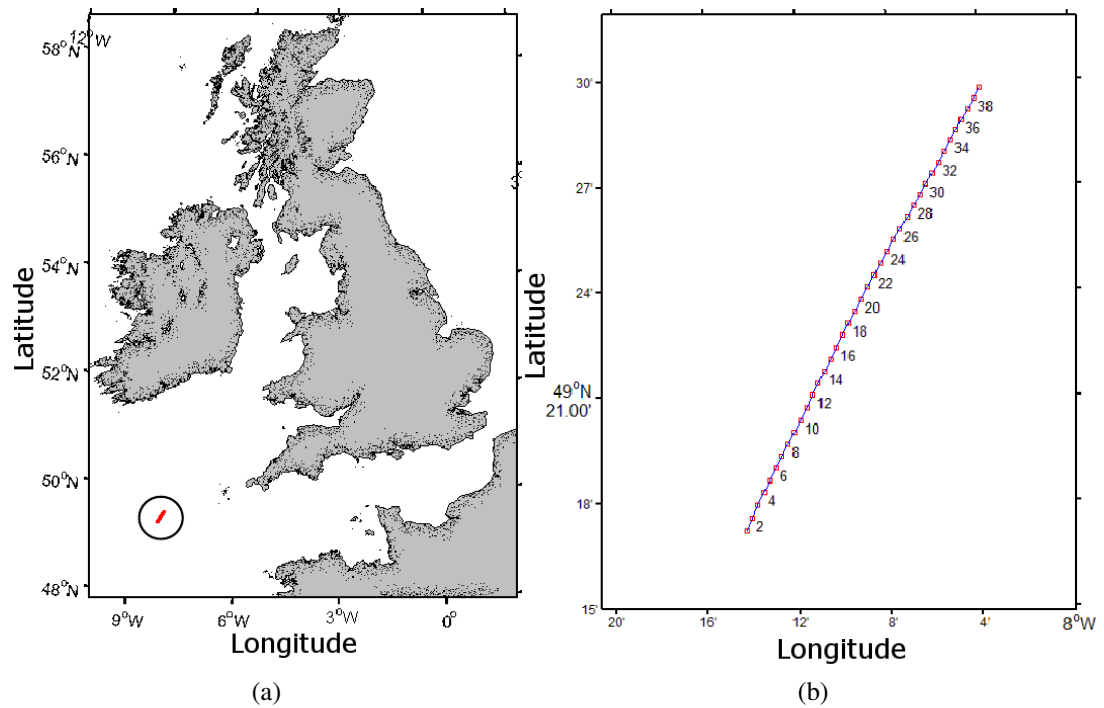


Figure 3.9: Maps depicting the location of the ship at each ping of sonar data (a) relative to the UK, (b) locally with ping numbers displayed

In the absence of non-acoustic data specifying the bearing of the sonar array at each ping (such as a digital compass reading or the output of an inertial navigation system), the bearing of the sonar array can also be estimated from the GPS coordinates. A naïve approach would be to take the north-based bearing between the current ping and the next. However, this assumes that the sonar platform travels in a perfectly straight line between locations. Alternatively, if θ_p is the angle between the current ping and the next ping, then the instantaneous bearing of the sonar array (denoted θ_S) can be estimated by averaging θ_p with the previous north-based bearing θ_{p-1} :

$$\theta_S = \frac{1}{2} (\theta_p + \theta_{p-1}). \quad (3.3)$$

This procedure is illustrated in Figure 3.10.




Figure 3.10: Calculating the north-based bearing of each ping. Red dots denote pings, black lines denote the relative range and bearing between consecutive pings, blue lines denote the instantaneous bearing of each ping estimated by averaging the two local bearings.

Equations (3.1) and (3.3) provide the range r and the north-based bearing θ between the current location of the sonar array and a further location. In order to convert the north-based bearing to a bearing relative to a ping, we must account for two factors: the orientation of the sonar system, and the presence of left-right ambiguity. To account for the orientation of the sonar system we subtract the north-based bearing of the sonar

system θ_S ,

$$\theta' = \theta - \theta_S \quad (3.4)$$

It should be noted that in this equation θ_S is calculated with respect to the previous and current ping. In contrast θ is calculated with respect to the current ping and a chosen location. The value θ' may be negative, and so the modulo remainder should be taken with respect to 360° , as defined by Knuth [75]:

$$\begin{aligned} \theta'_+ &= \theta' - 360 \left\lfloor \frac{\theta'}{360} \right\rfloor \\ &= (\theta - \theta_S) - 360 \left\lfloor \frac{\theta - \theta_S}{360} \right\rfloor \end{aligned} \quad (3.5)$$

where $\left\lfloor \frac{\theta'}{360} \right\rfloor$ denotes the largest value less than or equal to $\frac{\theta'}{360}$. Due to left-right ambiguity we only have data for 180 degrees, rather than duplicating these cells, the bearing can be mapped to this angular range. If the sonar-oriented bearing is greater than 180 degrees then we subtract this value from 360, this is expressed mathematically as:

$$\theta'_{180+} = \begin{cases} \theta'_+ & \text{if } \theta'_+ \leq 180 \\ 360 - \theta'_+ & \text{if } \theta'_+ > 180 \end{cases} \quad (3.6)$$

These calculations convert the original north-based bearing between the current ping and a second location, to a bearing relative to the sonar system, which lies within the angular range covered by our 64 beams. Given the current and previous locations of the sonar system, we are now able to calculate the range and bearing between the sonar system and any arbitrary GPS coordinate. This will prove useful in Chapter 7 and Chapter 8 where we perform target detection.

In this chapter we have described the trial data provided by Thales Underwater Systems. This included details of how the data was collected, its format, and examples of the environmental and systematic features present in the data. This trial data is used extensively throughout this thesis as it contains real sonar noise, collected from a known

location. Consequently the latter half of this chapter has focused on how this trial data can be interpreted in terms of Cartesian coordinates, and the methods needed to calculate locations relative to each ping in the trial data. We will now proceed to examine various distributions which can be used to model sonar noise.

Chapter 4

Noise Distributions

“

All noise is waste. So cultivate quietness in your speech, in your thoughts, in your emotions. Speak habitually low. Wait for attention and then your low words will be charged with dynamite.

— ELBERT HUBBARD (1856–1915)

The first step in achieving improved detection performance is to find an accurate model of target free (noise-only) data. As described in Chapter 2, the properties of the noise distribution have a direct impact on the probability of false alarm, which is related to the probability of detection by the detection threshold. Many distributions exist which can be used to model noise. Some of these have been proposed as suitable models for use in sonar, and we will examine them here.

This chapter begins by introducing a number of statistical models for sonar noise. We then describe generic algorithms for parameter estimation, and identify appropriate

estimators for the various noise models. This is followed by an overview of random number generation, and the processes by which synthetic sonar noise can be generated.

4.1 Noise Distributions

If we assume that the *central limit theorem* leads the real and imaginary components of the matched filter output to have a Gaussian distribution, the amplitude of this complex data then has a Rayleigh distribution. This is commonly used in both radar and sonar [3, 136, 138]. Under this assumption, which is reasonable in cases where scatterers are of approximately equal strength, the model is justified for use in noise-limited environments without clutter or reverberation. The cumulative distribution function (CDF) of the Rayleigh distribution, with scale denoted by σ , is defined as:

$$F_{\text{RAYLEIGH}}(x; \sigma) = 1 - e^{-x^2/2\sigma^2}. \quad (4.1)$$

It is accepted in the literature, and demonstrated by our trial data, that not all sonar data can be modelled by the Rayleigh distribution. This scenario is often caused by reverberation and clutter, which are frequently found in shallow water environments. In the case of radar systems, the K -distribution was developed to account for additional modulations in amplitude which are not captured by the Rayleigh distribution [62]. Specifically, the K -distribution can be derived from a Rayleigh distribution modulated by a Gamma distribution. In 2002, Abraham and Lyons [3] proposed the K -distribution as a model for post-matched-filter sonar data from a finite number of point scatterers. Experimental studies have demonstrated that the K -distribution does indeed improve on the standard Rayleigh model [3, 31]. The K -distribution is characterized by its *scale* and *shape* parameters, denoted by λ and ν respectively, and reduces to the Rayleigh distribution in the limit as $\nu \rightarrow \infty$. With the scale parameter defined as the second

moment of the distribution [22], the cumulative distribution function can be written as

$$F_K(x; \nu, \lambda) = 1 - \frac{2}{\Gamma(\nu)} \left(\frac{\nu}{\lambda}\right)^{\frac{\nu}{2}} x^\nu K_\nu \left(2x \left(\frac{\nu}{\lambda}\right)^{\frac{1}{2}}\right) \quad (4.2)$$

where $\Gamma(\cdot)$ is the gamma function and $K_\nu(\cdot)$ is the modified Bessel function of the second kind. In the limit $\nu \rightarrow \infty$ this is equivalent to the Rayleigh distribution. As the shape parameter ν decreases, the departure from the Rayleigh distribution can be used to model sonar data from a decreased number of scatterers [3]. This is analogous to the modelling of a more cluttered environment, which provides physical justification for the use of the K -distribution to model clutter. The sensitivity of the shape parameter ν is demonstrated by plots of the K -distribution probability density function in Figure 4.1.

Beyond physically justified models of sonar data, distributions are often chosen for the simple reason that they provide a good fit to the data. Sonar clutter is observed to exhibit a heavy-tailed amplitude [3]. Therefore the log-normal and Weibull distributions are often used to model sonar noise, both of which have heavier tails than the Rayleigh distribution.

The log-normal distribution is the exponential of a Gaussian distribution. It is characterised by its scale and shape parameters, μ and σ , and has the following cumulative distribution function

$$F_{\text{LOGNORMAL}}(x; \mu, \sigma) = \frac{1}{2} \left(1 - \operatorname{erf} \left[-\frac{\ln x - \mu}{\sigma\sqrt{2}} \right] \right). \quad (4.3)$$

It should be noted that the parameters μ and σ specify the underlying Gaussian distribution, and should not be confused with the mean and standard deviation of the log-normal distribution.

The Weibull distribution is characterized by its scale and shape parameters, α and β ,

and has the following cumulative distribution function

$$F_{\text{WEIBULL}}(x; \alpha, \beta) = 1 - e^{-\left(\frac{x}{\alpha}\right)^\beta}. \quad (4.4)$$

This is part of the *extreme value* family of distributions, and includes both the Rayleigh and exponential distributions as special cases. This distribution can also provide a variety of shapes similar to those of the K-distribution, as seen in Figure 4.2.

Another motivating factor when choosing an appropriate distribution is its ease of use. The Gaussian distribution is often used for this reason. For example, this is applied in the normalization scheme we defined in (2.17) of Section 2.7. This model is known to be incorrect [5], and was quickly found to offer a poor fit to our trial data. The Gaussian distribution will not be considered in our statistical analysis. The Gaussian distribution is, however, an important baseline due to its widespread use, and will be seen again in Chapter 7.




Figure 4.1: K-distribution probability density function for various values of the shape parameter ν




Figure 4.2: Weibull probability density function for various values of the shape parameter β

4.2 Parameter Estimation

All of the distributions described above have parameters, so the probability density functions of these distributions are not fully defined until we choose values for each of them. If we know the correct parameter values for each distribution, then these can be substituted directly into the equations. However, when modelling data collected in the real-world, this is rarely the case. In the absence of this information, one approach to the problem is to fit the distribution to available data, attempting to select parameter values which accurately model our observations. This is achieved through a process known as *parameter estimation*.

We will now outline the basic theory for two main methods of parameter estimation, Maximum Likelihood Estimation (MLE), and Method of Moments (MME). From the distributions described above, the parameters of the Rayleigh, log-normal and Weibull distributions are easily estimated by MLE. However, the shape parameter of the K-distribution is not.

As noted previously in the literature [3, 109], computing the MLE of the K -distribution shape parameter is computationally intensive, requiring an extensive numerical search and repeated calculation of values of the modified Bessel function $K_\nu(\cdot)$ [66]. This difficulty makes the MLE infeasible for the analysis of large data sets such as the trial data provided by Thales Underwater Systems. To resolve this difficulty, this section will finish with a summary of different techniques for estimating ν .

4.2.1 Estimation Theory

Parameter estimation is an optimisation problem, and there are two main approaches which can be taken. It is either necessary to identify a statistic which quantifies the

desirability of a parameter value, increasing as the estimate becomes more accurate, or to identify a statistic which quantifies the error between the estimate and the true value of the parameter.

Maximum Likelihood Estimation

Maximum likelihood estimates take the first approach, attempting to maximize the *likelihood function* of a distribution, which quantifies how likely it is that a parameter would take a certain value given the observed data. As the name suggests, likelihood is related to probability. It is a common misconception to think of likelihood as a synonym for probability, but there are distinct differences crucial to its correct interpretation.

Let us first consider a *probability density function*, $f(x|\theta)$, where θ represents the distribution parameters. Such functions are understood as a mapping from x to a probability of observing the value x for some fixed parameters θ . In other words:

$$x \mapsto f(x|\theta) \quad (4.5)$$

Observe that

$$f(x|\theta) = \frac{f(\theta|x)f(x)}{f(\theta)}. \quad (4.6)$$

Clearly this reasoning can be reversed such that our calculation is in terms of θ . This is known as a likelihood function, and is denoted $L(\theta|x)$. For a probability density function we think of the parameters as fixed and the data as a (random) variable. In contrast, for a likelihood function, the data is fixed and the parameters are (random) variables. As a result of this, while a probability density function must integrate to 1, the equivalent quantity for a likelihood function need not

$$\int f(x|\theta), dx = 1 \quad \text{but} \quad \int L(\theta|x)d\theta \neq 1 \quad \text{necessarily.} \quad (4.7)$$

A further subtlety is that a likelihood function is actually an equivalence relation, with a non-zero coefficient referred to as the *constant of proportionality*.

In maximum likelihood estimation [68], the estimation technique must calculate the distribution parameters θ for which the data x is most ‘likely’ to be observed. By maximizing $L(\theta|x)$ we find the value(s) of θ which best match the probability density function for the input data x . The actual value of the likelihood function is immaterial, it is the relative magnitude of this value compared to the likelihood for other parameters which is of interest.

In some cases this maximization can be performed analytically. For example, the commonly recognised calculations for mean and variance are maximum likelihood estimates for the Gaussian distributions parameters μ and σ^2 , respectively. However, for many distributions this is not possible and so the likelihood function must be maximized numerically.

Method of Moments Estimation

An alternative approach is to make use of formulas which relate the moments of a distribution to an unknown parameter we wish to calculate. This is known as the *method of moments*, or *moment-matching*. This can either be performed analytically, by equating the unknown parameter to a function of one or more moments of the distribution, or numerically by attempting to minimize the error between the two quantities. In the absence of prior knowledge, the moments of a distribution can be approximated by sample moments

$$\hat{m}_k = \frac{1}{n} \sum_{i=1}^n x_i^k \quad (4.8)$$

where \hat{m}_k is the k 'th sample moment of the observed data $X = \{x_1, \dots, x_n\}$. As a simple example, in the case of the Rayleigh distribution it is known that the scale

parameter σ is related to the mean (or first moment) by:

$$m_1 = \sigma \sqrt{\frac{\pi}{2}}. \quad (4.9)$$

By substituting the sample moment \hat{m}_1 and rearranging we obtain an estimate of σ by method of moments:

$$\hat{\sigma} = \frac{\sum_{i=1}^n x_i}{n\sqrt{\pi/2}}. \quad (4.10)$$

We can also make use of higher moments (though these are often more sensitive to statistical noise). For example the variance of the Rayleigh distribution is related to σ as follows:

$$\text{If } X \sim f_{\text{RAYLEIGH}}(\sigma) \quad (4.11)$$

$$\text{Then } \text{Var}(X) = \frac{4 - \pi}{2} \sigma^2 \quad (4.12)$$

and can be estimated from sample moments as:

$$\text{Var}(X) = \hat{m}_2 - \hat{m}_1^2. \quad (4.13)$$

Substituting (4.9) into (4.13) and simplifying leads to another estimator for σ :

$$\hat{\sigma} = \sqrt{\frac{\sum_{i=1}^n x_i^2}{2n}} \quad (4.14)$$

It can be shown that this estimator is also the MLE for σ . For a number of simple distributions, an analytic formula for the MLE can be found by MME approaches, but this is not always the case. Generally speaking, MME estimators are easy to implement, but if they do not also maximize the likelihood of the distribution they tend to provide lower accuracy than an MLE estimator.

Now that we have identified two common methods for the derivation of parameter

estimators, let us consider how the accuracy of an estimator can be quantified.

Estimator Accuracy

There are two main statistics used in the evaluation of an estimator's accuracy. These values are useful as they can be used to calculate the error bounds of an estimator's output, or for comparison against other candidate methods of estimation. If we consider the output of an estimator to be a random variable, the distribution of values is often seen to be Gaussian, and so the mean and variance are *sufficient statistics*. A sufficient statistic is a statistic which contains all available information regarding a statistical model, such that no other statistic exists which can provide further information. This implies that the mean and variance are often capable of fully specifying the statistical properties of the distribution of the output of an estimator. From these two statistics we can obtain the *bias*, and *variance* of an estimator, for specific parameter values.

An estimator should, on average, provide an estimate which is close to the true value most of the time. This is quantified by *bias*, which is defined as the average error between the true and estimated values:

$$\text{Bias}(\hat{\theta}) = E(\hat{\theta}) - \theta, \quad (4.15)$$

where $\hat{\theta}$ denotes an estimate of the true value θ , and $E(\cdot)$ denotes expectation. If the bias of an estimator is zero, we say that the estimator is unbiased, and so the estimator will, on average, provide the true value of the parameter.

The other main consideration in estimator accuracy is *variance*, as a quantification of spread. In other words, how far individual estimates of a sample will lie from the average output of the estimation process. As the variance of an estimator increases, a greater proportion of its estimates will differ greatly from the average result, therefore

decreasing our confidence in its accuracy.

The concepts of bias and variance will be put to use in Section 5.1, where we will analyse the accuracy of various estimators for the K-distribution shape parameter. However, in the above discussion we have implicitly assumed that the data used to estimate a parameter is from the same distribution as that parameter. In cases where data from one distribution is used to estimate parameters for another, there is no guarantee that an accurate or sensible fit will be obtained. As there may be no true value with which to compare the output of the estimation process, the bias and variance may no longer be suitable measures of performance. In the context of sonar, this difficulty may arise if a noise distribution does not suitably model the relevant noise sources in the environment. The trial data used in this thesis was collected in littoral waters and is therefore dominated by reverberation and background noise. In deep waters, where these sources of noise are less prevalent, other sources of noise may be significant. If a distribution is unable to accommodate these, then the estimation process may lead to an unsatisfactory fit between the distribution and the data. Methods for quantifying goodness-of-fit are discussed in detail in Section 5.2. We will now continue by defining the shape parameter estimators.

4.2.2 The $z \log(z)$ Estimator

This estimator was derived by Blacknell and Tough [22], and is known to provide a more robust result than higher-order method of moment estimates [60]. In this case the estimator arises from the limiting case of the fractional exponent decreasing. While our data is recorded in terms of signal amplitude, this estimator was derived in terms of intensity. The two are easily interchanged, as intensity z is related to amplitude x by

the equation $z \approx x^2$. From Equation (16) of [22], the estimator is defined as:

$$1 + \frac{1}{\nu} = \frac{E(z \log z)}{E(z)} - E(\log z) \quad (4.16)$$

where $E(\cdot)$ denotes expectation. Calculating these values from available samples and inverting the equation leads to the following:

$$\hat{\nu} = \left[\frac{\frac{1}{n} \sum_{i=1}^n z_i \log(z_i)}{\frac{1}{n} \sum_{i=1}^n z_i} - \frac{\sum_{i=1}^n \log(z_i)}{n} - 1 \right]^{-1} \quad (4.17)$$

The authors' analysis showed that this technique provides at least comparable performance to other closed-form estimators, such as the normalized log estimator, V and W estimators proposed by C. J. Oliver [100]. As a result, Oliver's estimators will not be studied further. The main advantages to the $z \log(z)$ estimator are its low computational complexity and easy implementation. However, for large shape values (tending towards Rayleigh distribution, see Section 4.1) the variance of this estimator rapidly increases, becoming a limiting factor. The Cramér-Rao lower bound for the K-distribution, which defines the minimum possible variance of an unbiased estimator, has been approximated by Abraham and Lyons [3] as ν^2/n . It is therefore not surprising that it is difficult to estimate the shape parameter of the K-distribution if it is large.

This method was extended by Hu in 2009 [59], proposing a $z^r \log(z)$ estimator with low variance for values in the range $0 < \nu < 2$, but as this range is already easy to accurately estimate, this new proposal is not particularly useful in our work.

4.2.3 The Method of Moments Estimator

Proposed by Abraham and Lyons [3], this estimator combines an iterative method of moments estimate, based on the first two moments of the distribution, with a starting value provided by approximating the gamma function. Through Newton-Raphson

iteration we aim to solve the following equation:

$$\frac{m_2}{m_1^2} = \frac{4\nu\Gamma^2(\nu)}{\pi\Gamma^2\left(\nu + \frac{1}{2}\right)} \quad (4.18)$$

for ν , where $\Gamma(\cdot)$ is the gamma function, and m_k denotes the k 'th moment of the available data.

To implement this estimator we must first define the prior calculation of the moments, and 3 functions:

$$\beta = \frac{m_2}{m_1^2} = \frac{\frac{1}{n} \sum_{i=1}^n x_i^2}{\left(\frac{1}{n} \sum_{i=1}^n x_i\right)^2} \quad (4.19)$$

$$h(\nu) = \frac{4\nu\Gamma^2(\nu)}{\pi\Gamma^2\left(\nu + \frac{1}{2}\right)} \quad (4.20)$$

$$g(\nu) = [h(\nu) - 1]^{-1} - [\beta - 1]^{-1} \quad (4.21)$$

$$g'(\nu) = \frac{-\nu^{-1}h(\nu)}{[h(\nu) - 1]^2} \left\{ 1 + 2\nu \left[\psi(\nu) - \psi\left(\nu + \frac{1}{2}\right) \right] \right\} \quad (4.22)$$

where $\psi(\cdot)$ denotes the digamma function [10]. The Newton-Raphson iteration is then incrementally calculated as follows:

$$\hat{\nu}_i := \hat{\nu}_{i-1} - \frac{g(\hat{\nu}_{i-1})}{g'(\hat{\nu}_{i-1})} \quad (4.23)$$

until $\hat{\nu}$ has suitably converged or a maximum number of iterations has been reached. The starting point is chosen by applying an approximation to the gamma function and inverting (4.18), resulting in:

$$\hat{\nu}_1 = \frac{1}{4} \left[\log\left(\frac{\pi}{4}\beta\right) \right]^{-1}. \quad (4.24)$$

Alternatively, the seed $\hat{\nu}_1$ can be set to the output of the $z \log(z)$ estimator, which tends to provide better performance at low shape parameter values. Note that if $\hat{\nu}_1 < 0$ then $\psi(\nu)$ is undefined and the process will fail. In our implementation, this special case

was handled by arbitrarily setting $\hat{\nu}_1 = 1$. Equally, the iterative MME does not always converge, and in some cases results in an estimate which is not real and greater than zero. This is most often seen when the value of ν is large, leading to high variance as shown by the Cramér-Rao lower bound [3]. Abraham and Lyons found this scenario occurred most frequently for large shape values and/or low sample sizes, where the following inequality did not hold:

$$\frac{m_2}{m_1^2} > \frac{4}{\pi} \quad (4.25)$$

in these cases an alternative estimation method should be used, or the parameter approximated by a large value of ν .

4.2.4 The Bayesian Approaches

We now describe two methods of Bayesian estimation, recently suggested by Abraham and Lyons as reliable methods of obtaining shape parameter estimates [4]. This overcomes the difficulties experienced in the two estimation algorithms described above when no solution could be found.

Both of these techniques make use of an intermediate statistic, the *Detection Index* (D), which is defined by

$$D = \frac{\mu^2}{\sigma^2} \quad (4.26)$$

where μ is the mean and σ the standard deviation. Bayesian estimation of this value is performed using a posterior distribution specified in terms of a non-informative prior, and a Gamma distribution to approximate the K-distribution. This posterior distribution is found to be a truncated Gamma distribution. By taking the mean of this distribution we obtain the estimate \hat{D} , which can then be transformed back to the K-distribution

shape parameter ν through the following approximation

$$\hat{\nu} \approx \frac{1}{4 \log \left[\frac{(1+\hat{D}^{-1})\pi}{4} \right]} \quad (4.27)$$

From this approximation, we now define two different ways of estimating ν . The difference between the two methods, is in how the Gamma distribution parameters $\{\hat{a}, \hat{b}\}$ are estimated. Code listings for both methods as used in this work can be found in [4, Appendix II]. This method is essentially an extension to Raghavans estimator [107], which also approximated the K-distribution by a Gamma distribution, producing an estimate by inversion.

Analytic Estimator (MMA)

For the *analytic* estimator, the gamma distribution parameters $\{\hat{a}, \hat{b}\}$ are estimated by moment-matching against \hat{D} . This is achieved by using sufficient statistics for D which have been carefully chosen to be invariant to the K-distribution scale parameter. These scale-invariant moment estimates are used in the following equations

$$\hat{a} = \frac{\hat{m}_1^2}{\hat{m}_2 - \hat{m}_1^2} \quad (4.28)$$

$$\hat{b} = \frac{\hat{m}_2 - \hat{m}_1^2}{\hat{m}_1} \quad (4.29)$$

which can then be substituted into (4.27).

Bootstrap Estimator (MMB)

For the *bootstrap* method, an MLE approach is taken. Given that the Newton-Raphson solution to this MLE is known to converge rapidly and we intend on implementing

bootstrapping, this estimator approximates the MLE by the Newton-Raphson iterations starting value, which is accurate for large a :

$$\hat{a} \approx \frac{\hat{D}_a}{2(\hat{D}_a - \hat{D}_g)} \quad (4.30)$$

$$\hat{b} = \frac{\hat{D}_a}{\hat{a}}. \quad (4.31)$$

In this approximation \hat{D}_a and \hat{D}_g denote arithmetic and geometric means respectively. These means are calculated from 301 estimates of \hat{D} evaluated using sample estimates of $\{\hat{\mu}, \hat{\sigma}\}$ in (4.26), where the 301 estimates of \hat{D} are obtained by 300 resamplings of the original data. Once again the estimated Gamma distribution parameters are substituted into (4.27) to obtain the shape parameter estimate.

4.2.5 Large Values of ν

As stated in Section 4.2.3, the Cramér-Rao lower bound for the K-distribution shape parameter suggests that variance will be high for large values of ν . To compound this difficulty, numerical algorithms for computing the Bessel function $K_\nu(\cdot)$, which are needed to compute the K-distribution density functions, become unstable at large values of ν . We avoid these difficulties by simply reverting to a Rayleigh distribution when estimation of ν fails or $\hat{\nu} \geq 30$.

This value 30 was selected by taking a K-distribution with scale $\lambda = 100$, and calculating the shape parameter at which the *sum of absolute error* between the pdf of the K-distribution and the pdf of a moment-matched Rayleigh distribution becomes two percent. The shape value was then rounded down to an integer. This was achieved by noting that the second moment of the Rayleigh distribution is defined as:

$$m_2 = 2\sigma^2, \quad (4.32)$$

so given that λ is the second moment of the K-distribution, a suitable scale value for the Rayleigh distribution would be $\sqrt{\lambda/2}$. The sum of absolute error can then be estimated for each shape parameter as follows:

$$\text{SAE}(\nu, \lambda) = \sum_{x=1}^{1000} \left| f_{\text{K}}(x, \nu, \lambda) - f_{\text{RAYLEIGH}}\left(x, \sqrt{\lambda/2}\right) \right|. \quad (4.33)$$

Calculating this statistic for shape parameters between 1 and 100, we obtain the plot presented in Figure 4.3. From this data it is found that an absolute error of two percent is obtained at approximately $\nu = 30.5$, rounding down to the nearest integer we obtain our boundary of $\nu \geq 30$, at which data is instead modelled by a Rayleigh distribution, which is equivalent to assuming the limit $\nu \rightarrow \infty$.




Figure 4.3: Sum of absolute error between a K-distribution with scale $\lambda = 100$ and a moment-matched Rayleigh distribution, for various values of the shape parameter ν

4.3 Basic Random Number Generation

The ability to develop a controlled environment within which algorithms can be tested and verified provides useful insights into the functionality and accuracy of a system. Such resources are necessary to ensure that algorithms perform as expected and can provide an alternative to real data for controlled public demonstrations. The rest of this chapter will focus on random number generators for synthetic sonar data.

Random numbers are often needed for simulations in the sciences, but truly random sequences are notoriously difficult to obtain [137]. As a result we must use pseudo-random numbers, generated as a series by a computer algorithm (starting from an initial *seed*). Such algorithms are known as pseudo-random number generators (PRNG), and generally produce uniformly distributed pseudo-random values in the interval $[0,1]$.

In the literature, there are two approaches to generating simulated sonar noise. The most complex approach is to simulate a full sonar environment. In [35, 141], the authors used the U.S. Navy's CASS-GRAB model (Comprehensive Acoustic Sonar Simulation with Gaussian Ray Bundles) and the VENUS model (Virtual Underwater Environment Simulation), respectively. Both of these systems take environmental parameters as inputs, such as a sound speed profile, water depth, and sea floor characteristics. A scenario is then designed with user-selected projectors and receivers, operating at specified frequencies. These frameworks allow for the simulation of acoustic wave propagation, using variations of ray-tracing algorithms, in order to estimate the echoes received for each cell.

While these simulation packages are designed with sonar principles & effects in mind, the need to calculate the transfer of acoustic energy through the simulated environment leads to very high computational complexity. For the purposes of this project, it is not necessary to simulate the entire underwater environment. To simulate post-matched-

filter data, it will be sufficient to generate random variables with properties which match those of the trial data.

This section deals with generic methods which transform uniform random numbers to random numbers belonging to another distribution. We assume that a uniform PRNG is already available. Many PRNG algorithms exist, as discussed by Marsaglia [90], such as the linear congruential generator (which provides the default PRNG in many programming languages), multiply-with-carry methods, and the Mersenne Twister [96]. In our simulations the Mersenne Twister (MT19937) has been used as it is the default PRNG in MATLAB and known to reliably provide random number streams with a large repetition period, $2^{19937} - 1$, which is more than sufficient for our purposes.

We will now outline the three main methods used to generate non-uniform random numbers from uniform PRNGs. The interested reader can find further information regarding these methods in the paper by M. Brysbaert or the book written by L. Devroye [26, 41].

4.3.1 The Inversion Method

A simple method for generating non-uniform random numbers is obtained by inverting the cumulative distribution function (CDF) of the desired distribution. This is established from the theorem that if a set of samples X has the cumulative distribution function F , then $F(X)$ is uniformly distributed on the range $[0, 1]$. Therefore, providing this CDF has an inverse, we can transform a uniformly distributed random variable into our distribution of choice. Let us define U to be a set of samples from a standard uniform distribution, and the tilde symbol (\sim) to denote “distributed according to”. The

inversion method can then be formalised as follows:

$$\begin{aligned} X &\sim F \\ F(X) &\sim U \\ \therefore F^{-1}(U) &\sim F. \end{aligned} \tag{4.34}$$

For example, the Weibull distribution with scale parameter α and shape parameter β , is defined by the CDF:

$$F_{\text{WEIBULL}}(x; \alpha, \beta) = 1 - e^{-\left(\frac{x}{\alpha}\right)^\beta}. \tag{4.35}$$

Random numbers can be generated from this distribution by inverting the CDF:

$$\begin{aligned} U &= 1 - e^{-\left(\frac{x}{\alpha}\right)^\beta} \\ X &= \alpha \left(-\log(1 - U)\right)^{\frac{1}{\beta}} \end{aligned} \tag{4.36}$$

With the availability of a function providing uniformly distributed random numbers (which we will name `RAND`), the inversion method is easily transferred into pseudo-code, provided by Algorithm 4.1.

Algorithm 4.1 `RANDW`: Random number generator for `Weibull`(α, β)

`U` \leftarrow `RAND`

`X` \leftarrow $\alpha \left(-\log(1 - U)\right)^{\frac{1}{\beta}}$

return `X` as a random value distributed according to `Weibull`(α, β)

In this manner Weibull distributed values are obtained through a basic transformation (4.36) of a uniform random variable.

This method of random number generation is often chosen due to its simplicity and the speed of its calculations. Unfortunately, in some cases it is not possible to find an analytic expression, or suitable approximation, for the inverse function $F^{-1}(\cdot)$. In such cases we must use the Acceptance-Rejection method.

4.3.2 The Acceptance-Rejection Method

When calculating the inverse of a CDF is computationally intensive, or even impossible, an alternative is to apply the Acceptance-Rejection method. Rather than making use of the CDF of the desired distribution, this method requires a well-defined probability density function (PDF) instead.

We assume that for the desired distribution, with PDF $f(X)$, there exists another distribution $g(X)$ which we are readily able to sample from, and a constant c such that

$$f(x) \leq c \times g(x) \forall x, \quad (4.37)$$

as shown in Figure 4.4.




Figure 4.4: Example of a distribution $g(X)$ and constant c from which a distribution $f(X)$ can be sampled. Courtesy of MathWorks [94]

Ideally the function $g(X)$ will be a close match to the desired distribution $f(X)$ so that the area between $c \times g(X)$ and $f(X)$ is small. The acceptance rejection algorithm then executes as follows:

Evidently, the bigger the difference between the scaled distribution g and the target

Algorithm 4.2 Generic algorithm for the Acceptance-Rejection method

```

repeat
   $X \leftarrow$  a random value from the distribution  $g(X)$ 
   $U \leftarrow$  a random value from the uniform distribution on the closed interval  $[0, 1]$ 
   $T \leftarrow c \times g(X)/f(X)$ 
until  $U \times T \leq 1$ 
return  $X$  as the random value generated according to  $f(X)$ 

```

distribution f , the more time will be spent calculating unsuitable random values. As a result, in the absence of a computationally tractable inversion of the desired distribution, this generic method may be seen as an alternative.

Fortunately a number of authors have taken this principle a step further. By generating random values from known distributions and applying certain restrictions to their output it is possible to form computationally efficient algorithms to generate data for a number of distributions. In some cases, these slightly more esoteric approaches are even more efficient than the inversion method. We will now examine a couple of examples, based on Algorithm 4.2.

Generating Gaussian Distributed Values

The Gaussian distribution, with mean μ and variance σ^2 , is defined by the CDF:

$$F_{\text{GAUSSIAN}}(x; \mu, \sigma) = \frac{1}{2} \left[1 + \operatorname{erf} \left(\frac{x - \mu}{\sqrt{2}\sigma} \right) \right] \quad (4.38)$$

where $\operatorname{erf}(\cdot)$ denotes the Gauss error function. As it is difficult to invert this special function, the Gaussian distribution is a candidate for Acceptance-Rejection sampling. The PDF of the distribution is:

$$f_{\text{GAUSSIAN}}(x; \mu, \sigma) = \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{(x-\mu)^2}{2\sigma^2}} \quad (4.39)$$

A number of Acceptance-Rejection based methods exist for the Gaussian distribution, such as the ratio method [72], polar method [88], and the more complicated ziggurat method [91]. Due to its simplicity and speed, the polar method has been used in this project, and is defined in Algorithm 4.3. This is an optimization of the Box-Muller transform [25] from which samples are selected from within a circle of unit radius.

Algorithm 4.3 RANDN: Random number generator for $\text{Gaussian}(\mu, \sigma)$

repeat

$x \leftarrow (-1 + 2 \times \text{RAND})$

$y \leftarrow (-1 + 2 \times \text{RAND})$

$s \leftarrow (x^2 + y^2)$

until $s < 1$

return either $\left(\mu + \sigma x \sqrt{\frac{-1 \log(s)}{s}} \right)$ or $\left(\mu + \sigma y \sqrt{\frac{-1 \log(s)}{s}} \right)$

Generating Gamma Distributed Values

The Gamma distribution, with shape k and scale ω , is defined by the CDF:

$$F_{\text{GAMMA}}(x; k, \omega) = \frac{\gamma(k, x/\omega)}{\Gamma(k)} \quad (4.40)$$

where $\gamma(\cdot, \cdot)$ denotes the lower incomplete gamma function, and $\Gamma(\cdot)$ the gamma function. In order to apply the inversion method we would have to provide an inverse for $\text{erf}\left(\frac{x-\mu}{\sqrt{2\sigma^2}}\right)$ and $\gamma(k, x/\omega)$. Unfortunately there are no closed-form solutions and so the Acceptance-Rejection method is necessary. An efficient algorithm for generating Gamma distributed random numbers was developed by Marsaglia and Tsang [89]. This is defined in Algorithm 4.4, where RANDN refers to the algorithm shown in 4.3.

Algorithm 4.4 RANDG: Random number generator for Gamma(k, ω)

```

if  $k \geq 1$  then
   $d \leftarrow (k - 1/3)$ 
else
   $d \leftarrow (k + 2/3)$ 
end if
 $c \leftarrow (1/\sqrt{9d})$ 
repeat
  repeat
     $x \leftarrow \text{RANDN}$ 
     $v \leftarrow (1 + cx)^3$ 
  until  $v > 0$ 
   $y \leftarrow \text{RAND}$ 
until  $(y < 1 - 0.0331x^4)$  or  $(\log(y) < \frac{x^2}{2} + d[1 - v + \log(v)])$ 
if  $k \geq 1$  then
  return  $(\omega dv)$ 
else
   $z \leftarrow \text{RAND}$ 
  return  $(\omega dv) z^{1/k}$ 
end if

```

4.3.3 The Modular Method

Many distributions are related to other distributions, for which samples are easier to generate. In some cases a random number generator may be available for a distribution which is a generalized form of the desired distribution, where a constraint on the random number generator's parameters will produce the required output. In other cases a transformation or mixture of distributions is needed.

Consider the Rayleigh distribution, with scale parameter σ . Samples from this distribution can be obtained by inverting its CDF

$$\begin{aligned}
 U &= 1 - e\left(\frac{-x^2}{2\sigma^2}\right) \\
 X &= \sigma\sqrt{-2\log(1 - U)}
 \end{aligned}
 \tag{4.41}$$

However, the Rayleigh distribution is a special case of the Weibull distribution (4.35), with the shape parameter $\beta = 2$ and the scale parameter redefined as $\alpha = \sqrt{2}\sigma$. Directly substituting these values into the Weibull random number generator (4.36) leads to the same result:

$$\begin{aligned} X &= \alpha (-\log(1 - U))^{\frac{1}{\beta}} \\ &= (\sqrt{2}\sigma) (-\log(1 - U))^{\frac{1}{2}} \\ &= \sigma \sqrt{-2 \log(1 - U)}. \end{aligned} \quad (4.42)$$

Another representation of the Rayleigh distribution is as the magnitude (or Euclidean norm) of a complex random variable, containing a real and imaginary component distributed according to the Gaussian distribution. This is calculated as the square root of the sum of each component squared, as demonstrated in Figure 4.5.

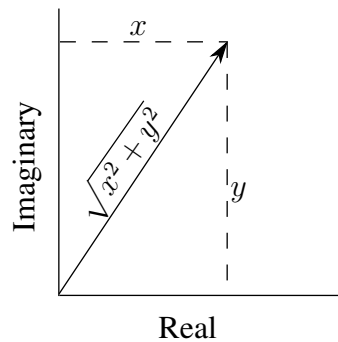


Figure 4.5: Graphical depiction of the magnitude of a complex random variable of the form $x + yi$

Therefore an alternative random number generator for the Rayleigh distribution would

be:

$$\begin{aligned}
 Y_0 &\sim N(0, \sigma^2) \\
 Y_1 &\sim N(0, \sigma^2) \\
 X &= \sqrt{Y_0^2 + Y_1^2} \sim f_{\text{RAYLEIGH}}(\sigma)
 \end{aligned} \tag{4.43}$$

This method allows more complicated random number generators to be defined in a modular fashion, making use of simpler building blocks. Naturally, of the two methods for generating a Rayleigh distribution, the inversion method (4.42) is favourable in terms of efficiency. However, the modular method (4.43) is a better representation of how data is obtained from a sonar system, where the magnitude is taken from the complex matched-filter output. We will demonstrate in Chapter 6 how this advantage provides a simple means of generating targets in a Rayleigh background.

4.4 Generating Sonar Noise

In this section we will demonstrate how to generate samples from the previously defined noise distributions (see Section 4.1). Samples from these random number generators are necessary in order to perform the statistical analysis provided in Chapter 5, and the performance analysis of Chapter 7.

As previously mentioned, the Rayleigh distribution provides a basic model for a noise-limited sonar environment. In strictly noise-limited environments there are no dominant scatterers present, such as a target or reverberation from the sea floor, resulting in the simple Rayleigh distribution. From (4.42) the random number generator for the Rayleigh distribution is:

$$X = \sigma \sqrt{-2 \log(U)}. \tag{4.44}$$

If reverberation and other forms of clutter are to be included in this model, the Rayleigh distribution should be modulated by a Gamma distribution, resulting in a *K-distribution*. This does not make the assumption that all scatterers are of approximately equal strength and can thus be considered a generalisation of the Rayleigh model of sonar noise. The K-distribution, with shape k and scale λ , has the following CDF:

$$F_K(x; \nu, \lambda) = 1 - \frac{2}{\Gamma(\nu)} \left(\frac{\nu}{\lambda}\right)^{\frac{\nu}{2}} x^\nu K_\nu \left(2x \left(\frac{\nu}{\lambda}\right)^{\frac{1}{2}}\right). \quad (4.45)$$

As can be seen the K-distribution CDF includes two special functions, the gamma function $\Gamma(\cdot)$ and the modified Bessel function of the second kind $K_\nu(\cdot)$. This makes random number generation by inversion difficult, as neither of these functions can be inverted analytically. Since the K-distribution is a Rayleigh variable modulated by a Gamma variable, we can apply the modular method [3]. For our parametrisation this requires a rescaling of the Gamma distribution, and leads to the following random number generator:

$$\begin{aligned} Z &\sim f_{\text{RAYLEIGH}}(1) \\ W^2 &\sim f_{\text{GAMMA}}\left(\nu, \frac{\lambda}{2\nu}\right) \\ X &= Z \times W \\ &\sim f_K(\nu, \lambda) \end{aligned} \quad (4.46)$$

Alternatively, this departure from the Rayleigh distribution can be modelled by the similar Weibull distribution. Though the Weibull distribution with shape β and scale α is also a generalization of the Rayleigh distribution, there is no need to perform the modular approach seen above and so we have used the formula previously derived by inversion:

$$X = \alpha (-\log(1 - U))^{\frac{1}{\beta}}. \quad (4.47)$$

Finally, the heavy-tailed nature of sonar reverberation and clutter has also been modelled

by the log-normal distribution. As this name suggests, this distribution is a simple transformation of a Gaussian distribution. In fact, the logarithm of data obtained from a log-normal distribution is Gaussian distributed. Reversing this process, the random number generator we have used for the log-normal distribution with location μ and scale σ , is defined in terms of the exponential function:

$$X = e^{N(\mu, \sigma)}. \quad (4.48)$$

4.5 Generating Autocorrelated Sonar Noise

Though the assumption of independence is made throughout this thesis, now is a convenient stage to consider the consequences of correlation in the sonar data. We first begin by generating autocorrelated noise for the K-distribution, and then consider the repercussions of this correlation with respect to sonar detection.

While other methods exist to generate autocorrelated random variables [123], we will apply the method of *circulant embedding* [43] which was applied to the K distribution by Redding [109]. Let us first start by choosing an autocorrelation function and generating a two-dimensional autocorrelated Gaussian random field. The autocorrelation function we will use here is the exponential function

$$\rho_g(x, y) = e^{-\sqrt{x^2+y^2}/4} \quad (4.49)$$

Given the width and height of the random field required, we can convert this autocorrelation function to a point-spread function (PSF) centred in the middle of a matrix of twice these dimensions. The PSF for (4.49) can be seen in Figure 4.6.

A simple method of generating autocorrelated data would be to perform a convolution,




Figure 4.6: Point-spread function for exponential autocorrelation, centred within an n by m matrix

between this PSF, and the independent, and identically distributed (IID) samples from a Gaussian random field. However, with large samples it is often more efficient to perform convolution in the frequency domain. From the convolution theorem:

$$f * g = \mathcal{F}^{-1} \{ \mathcal{F}\{f\} \otimes \mathcal{F}\{g\} \} \quad (4.50)$$

where $*$ is the convolution operator, \otimes a point-wise multiplication, and \mathcal{F} denotes a fast Fourier transform (FFT).

By applying this operation in the frequency domain, we replace a convolution by two FFTs, a point-wise multiplication, and a further inverse FFT. While the application of the FFTs lead to computational overhead, the benefits of this method have been demonstrated over quite small sample sizes. For example, a comparison by Strum and Kirk [129] found this technique to be faster than direct convolution after only 128 samples. However, a further important observation by Dietrich and Newsam [43] is that the FFT of an uncorrelated Gaussian random field is also Gaussian. Therefore we need only perform one FFT on the PSF and an inverse FFT on the product. Once this convolution has been performed, the top-left quadrant of the result is a Gaussian random field with autocorrelation defined by (4.49). It is important to note that this field is not necessarily distributed according to a standard Gaussian distribution, so it should be rescaled by subtracting its mean and dividing by its standard deviation. Furthermore,

the FFT of the PSD need only be calculated once, and so further realizations can be generated with the computational cost of a single inverse FFT and point-wise multiplication.

By passing these Gaussian distributed values through the CDF of the Gaussian distribution (4.38) we obtain a uniformly distributed autocorrelated random field. To generate non-uniform autocorrelated random fields from this data, this field can be used with the inversion or modular methods described in Section 4.3. Note that this can not be used with rejection-based methods, as the rejection stage would corrupt the autocorrelation. It should also be noted that autocorrelation is only invariant to some linear transformations, and so if a specific autocorrelation is required then this may require correction. For example, in the basic case of transforming correlated Gaussian samples to correlated uniform samples, the autocorrelation function of the Gaussian data must be modified as follows [42]:

$$\rho_g'(x, y) = 2 \sin \left(\frac{\pi}{2} \rho_g(x, y) \right) \quad (4.51)$$

To generate autocorrelated K-distributed values in the amplitude domain, Redding suggests that for synthetic *radar* returns, there tends to be correlation in the gamma modulation but not in the Rayleigh background noise. Given that the gamma random number generator provided in Section 4.3 is a rejection-based method we must take a different approach. For integer shape parameters (ν), it is possible to generate gamma distributed values from realisations of the exponential distribution, or equivalently from n realisations of a uniform distribution as follows:

$$\sum_{i=1}^n -\log(U_i) \sim \text{Gamma}(n, 1). \quad (4.52)$$

As noted by Redding, these exponential distributions can also be generated as the mean of two squared random samples from a Gaussian distribution. Thus, we are able to

generate gamma distributed data with half-integer shape parameters, as follows:

$$\frac{1}{n} \sum_{i=1}^n N(0, 1)_i^2 \sim \text{Gamma}(n/2, 1), \quad (4.53)$$

which is easily rescaled for the desired scale parameter. Thus an autocorrelated K-distributed random field, as defined by (4.42, 4.46, 4.53), and with a half-integer shape parameter ν , can be generated as follows:

$$\left[\frac{-2\lambda \log(U)}{4\nu^2} \sum_{i=1}^{2\nu} \mathcal{N}(0, 1)_i^2 \right]^{\frac{1}{2}} \sim K(\nu, \lambda) \quad (4.54)$$

where $\mathcal{N}(0, 1)$ is an autocorrelated Gaussian random field, as generated above.

Introducing autocorrelation in this manner is a form of averaging, and so it is interesting to note that the autocorrelation in this data leads to increased variance in the samples' parameter estimates. As stated by Redding [109]: “more samples are required in the correlated case than the uncorrelated case to achieve the same variance on the estimators”. This is analogous to considering correlation in the data to be a form of over-sampling. As over-sampling can not provide any new information, introducing correlation at a fixed sample size, with a resulting loss of information, has the same effect as reducing the sample size. This also implies that statistical methods and performance predictions, applied to simulated data without correlation, may suffer from a deterioration in performance when applied to real data, where short-term correlations are likely to exist.

In this chapter various statistical models for sonar noise have been introduced. This was followed by an overview of parameter estimation, and specific estimators for the K-distribution shape parameter. We then describe techniques for the generation of random samples from these noise models, which can be used to simulate a sonar environment.

The methods presented here are vital to the work in subsequent chapters, as they will be used to analyse sonar data, in addition to the derivation, implementation and evaluation of various detection algorithms. In the subsequent chapter we analyse the performance of the K-distribution shape parameter estimates, and perform a statistical analysis of the sonar noise in our trial data.

Chapter 5

Statistical Analysis of Trial Data

“

Do not put your faith in what statistics say until you have carefully considered what they do not say.

— WILLIAM W. WATT (1912–)

Detection algorithms often aim to attain a constant false alarm rate (CFAR), but to accomplish this, accurate threshold values must be found. For a low probability of false alarm, only a low percentage of noise-only samples may exceed the threshold; these occur in the tails of the distribution. Therefore, improved statistical models of sonar noise, particularly in the tails of the distribution, will reduce the number of false alarms, providing greater reliability in difficult environments. To this end, we must choose one or more of the noise distributions defined in Chapter 4. As detection algorithms must be robust to unexpected environments, it may be necessary to switch between multiple models to maintain performance. The selection of a noise distribution must be

an informed decision, and so we perform a statistical analysis on the noise in the trial data provided by Thales Underwater Systems.

This chapter includes an analysis of the previously defined parameter estimators for the K-distribution, which are compared in terms of variance and bias, to identify the most appropriate estimator for use in this thesis. In order to evaluate how accurately each model matches the trial data provided by Thales Underwater Systems, we describe and apply goodness-of-fit testing. The chapter is then concluded by evaluating our findings, and identifying the most accurate models of the noise in our trial data, which are of critical importance to subsequent chapters.

5.1 Analysis of Parameter Estimates

Before we can analyse how closely each noise distribution matches our trial data, we must first choose one of the methods for estimating the shape parameter of the K-distribution. In Section 4.2 we summarised four techniques of moderate computational complexity, which could feasibly be applied to large datasets such as ours. This should also be an informed choice, and so we must perform a quantitative analysis of each estimator's performance.

Traditionally the performance of parameter estimates tend to be evaluated by examining two properties, bias and variance. While the bias quantifies the expected error between the estimate and the true value, the variance quantifies the degree of dispersion from the estimates mean. If we consider a parameter estimate to be a random variable, and assume that the distribution of the parameter estimate is Gaussian, then these two values constitute a sufficient statistic. This is often the case, particularly for MLEs. The symmetry of the Gaussian distribution is also advantageous as extreme values are more

likely to cancel out, such that the sample mean over the estimate is a robust calculation of central tendency.

However, as we must compare two statistics to decide which method of estimation is most suitable, it can be hard to choose the trade-off between bias and variance. For example in scenarios where a Minimum Variance Unbiased estimator (MVUE) exists, it may still be advantageous to apply a biased estimator which achieves lower variance. If we have no reason for choosing low variance or low bias then it is useful to consider a combination of these two statistics. The most common of these is the Mean Squared Error (MSE) which is defined as

$$\begin{aligned} \text{MSE}(\hat{\theta}) &= \frac{1}{n} \sum_{i=1}^n (\hat{\theta}_i - \theta)^2 \\ &= \text{var}(\hat{\theta}) + \text{bias}(\hat{\theta})^2. \end{aligned} \quad (5.1)$$

Unfortunately in cases where the Gaussian assumption does not hold, the sample mean is no longer a robust measure of central tendency and is prone to errors from extreme values or outliers. This problem is known to occur when the distribution of an estimate is skewed, and in these cases the sample median is a more useful measure [28]. If an estimate were to be median-unbiased then this would be equivalent to minimizing its mean absolute error (MAE). This can be seen from the form of the MLE estimate of the Laplace distribution scale parameter, where $x^{(0.5)}$ denotes the median of x :

$$\hat{b} = \frac{1}{n} \sum_{i=1}^n |x_i - x^{(0.5)}|, \quad (5.2)$$

Thus the variance, mean-bias, median-bias, MSE, and MAE have been chosen to evaluate the performance of K-distribution shape parameter estimates.

A Monte-Carlo simulation was performed by generating K-distributed samples of size $n = 1000, 2500, 5000,$ and $10,000$ and shape parameters in the range $[0.05, 30]$. Each

simulation was performed with 10,000 iterations. The shape parameter of each sample was then estimated by the $z \log(z)$, MME, MME- $z \log(z)$, MMB and MMA estimators. From these shape estimates the MSE, MAE, variance, median and mean bias were calculated with respect to the true shape parameter. These can be seen for the sample size $n = 2500$ in Figures 5.1-5.2




Figure 5.1: Analysis of estimates for the K-distribution shape parameter for sample size $n = 2500$ (a) MSE, (b) Percentage mean-bias

On initial inspection of the plots in Figure 5.1 it is evident that as the shape parameter decreases below 5 the various estimators split into two classes of performance. Of these estimators the MMA, MMB and MME all perform poorly for low shape parameters but perform favourably for larger values. From Figure 5.1b this poor performance at low shape parameters is seen to be due to a high positive bias in the estimate. This is unfortunate as this is in the region of the shape parameter where minor deviations lead to a large change in the distribution statistics, as opposed to the minor changes seen at large shape values as the distribution slowly tends towards a Rayleigh. As all three of these methods apply an approximation to the gamma distribution this is the likely cause of the bias, which is not observed in the $z \log(z)$ estimates. As noted by Abraham and

Lyons [4] the approximation is suitably accurate down to a shape parameter of 1, at which point a newton-raphson iteration may be a desirable alternative.

It is also seen that the $z \log(z)$ estimate behaves rather erratically in terms of mean-bias. This behaviour was verified by resimulating these values at increments of 0.1, over 2,000,000 simulations ¹. Following this far more intensive calculation, the mean-bias was still not seen to demonstrate a useful trend, with little reduction in its erratic variation. At first this may seem concerning, but recalling that the mean and variance are sufficient statistics for a Gaussian distributed parameter estimate it became evident that the $z \log(z)$ estimates are often considerably non-Gaussian. This hypothesis was subsequently verified for low sample sizes and at large shape parameters, where the distribution of estimates tended away from a Gaussian and exhibited higher skew or kurtosis. Under these circumstances the MAE and median-bias is of greater interest, and are shown in Figure 5.2.

It is interesting that the $z \log(z)$ provides such low MSE at low shape values. This is seen to be predominantly due to a low bias at low shape values as seen in Figures 5.1b and 5.2b, which is why it was applied as the starting seed to the MME estimator [19]. This mixture of the two detection schemes is referred to in these Figures as the MME- $z \log(z)$. The use of the $z \log(z)$ as the initial seed leads to an overall lower variance estimator, with consistently lower error than the standard MME over all sample sizes tested and performance statistics discussed here. From a more accurate starting value the iterative nature of the MME overcomes the heavy non-gaussianity of the $z \log(z)$, leading to lower variance and lower expected error in terms of mean-bias.

Another important consideration is that of reliability, as addressed by Abraham and Lyons [4]. In our parametrisation of the K-distribution, useful shape values are those

¹As these Monte-Carlo simulations are embarrassingly parallel this was achieved by reimplementing the random number generator and an optimised $z \log(z)$ estimator in C++. This provided a lightweight executable which was deployed on Cardiff Universities CONDOR pool, a high-throughput parallel computing system with over 8000 nodes.




Figure 5.2: Analysis of estimates for the K-distribution shape parameter for sample size $n = 2500$ (a) MAE, (b) Percentage median-bias

which are finite, real, and greater than 0. Shape estimators which do not achieve such a value are considered to have failed. Throughout our simulation these failures were tallied and can be seen in Figure 5.3.

Evidently the MMA and MMB estimators have achieved their stated goal of providing a reliable means of estimating K-distribution shape parameters, with no estimation failures observed. The $MME - z\log(z)$ estimator compares unfavourably with the MMA, MMB and $z\log(z)$ estimators, because, despite improvements in variance and bias there is an associated increase in failure rate.

With a decrease in sample size the expected increase in estimation failures was observed for the MME and $z\log(z)$ estimators, with a corresponding decrease at larger sample sizes. Meanwhile the MMA and MMB estimators remained robust with no observed failures. However, changes in sample sizes did demonstrate a difficulty with the MMA and MMB estimators, in that they are not consistent. An example is provided in Figure 5.4 where the MMB provides a lower variance at a sample size of 1000 than at 2500.




Figure 5.3: Failure rate of the parameter estimators for various shape parameters. Sample size $n = 2500$

Lack of consistency is undesirable as an increase in samples should logically lead to a more accurate result. This observation would be of greater concern had the MMA and MMB estimates not provided the best overall MSE at all sample sizes.

Of all the methods observed here the MMA and MMB are also shown to provide the best performance (in the least-squares sense) at large shape values where estimation is troublesome. While the MMA estimator is the lowest variance and computational complexity the MMB provides slightly lower bias. Unfortunately at low shape values neither of these estimators are able to produce accurate estimates without the further computational burden required to improve the accuracy of the Gamma approximation. At shape parameters lower than 5 it is therefore advisable to attempt estimation using the $z \log(z)$ method. Although the MME- $z \log(z)$ is a reasonable combination of the MME and $z \log(z)$ estimators, in light of the above findings and its high failure-rate this estimator can not be recommended.

However, the MMA and MMB estimators were not yet defined when this chapter's statistical analysis of trial data was conducted, and so the MME- $z \log(z)$ estimator is




Figure 5.4: Variance for the MMA and MMB estimators, demonstrating lack of consistency with respect to sample size

used in the subsequent sections. The $z \log(z)$ seed provides an estimate with lower overall variance, compared to the gamma approximation used in [3], and also provides reduced bias for small values of ν . This method of parameter estimation is far more tractable than the MLE for the K -distribution and is comparable in runtime to the MLE for the Weibull distribution.

Unfortunately, the iterative MME does not always yield a meaningful result. This is most often seen when the value of ν is large, leading to high variance as shown by the Cramér-Rao lower bound [3]. To compound this difficulty, numerical algorithms for computing the Bessel function $K_\nu(\cdot)$, which are needed to compute the K -distribution, become unstable at large values of ν . We avoid these difficulties by simply using a Rayleigh distribution whenever $\nu > 30$, employing the following relation for the Rayleigh scale parameter σ ,

$$\sigma = \sqrt{\lambda/2}. \quad (5.3)$$

5.2 Goodness-of-Fit Testing

To quantify the extent to which real sonar data can be represented by a candidate distribution, a *goodness-of-fit* test must be used. As noted by Chernobai [32], goodness-of-fit statistics can be formulated in terms of either a supremum or quadratic calculation. These statistics tend to be a form of quantitative error between the anticipated, and the observed model. For example, the most commonly known supremum goodness-of-fit test is the Kolmogorov-Smirnov (KS) test.

The KS test statistic [93] is defined to be the maximum absolute difference between the empirical distribution function $F_n(x)$ of the observed data and the cumulative distribution function $F(x)$ being tested,

$$D_n = \sup_x |F_n(x) - F(x)| \text{ where } F_n(x) = \frac{1}{n} \sum_{i=1}^n I(X_i \leq x) \quad (5.4)$$

and $I(X_i \leq x)$ is the indicator function, demonstrated in Figure 5.5. For a given




Figure 5.5: Demonstration of the KS statistic D_n calculated between an empirical distribution and the CDF of the Normal distribution.

significance level, if the computed value of D_n exceeds the critical value under the null hypothesis that the observed data is from distribution F , then the null hypothesis is rejected and we conclude that the data does not come from distribution F . While the KS

test assumes that $F(x)$ is completely specified, we rely on *estimates* of the distribution parameters and must therefore adopt Lilliefors' modification [80] of the KS test when computing critical values.

Differences between the shape of the empirical distribution and the candidate distribution may change the output of a goodness-of-fit statistic. It is only when this difference between the two distributions is greater than that expected by random chance, for a specific level of significance, that the test will reject the null hypothesis. It is therefore possible that a goodness-of-fit test may not reject the hypothesis that data from a different distribution is from the candidate distribution. However, this should only occur if the two distributions are similar, or if they differ in a manner which is not captured by the chosen goodness-of-fit statistic. For this reason it is necessary to carefully select a goodness-of-fit test which is able to evaluate the relevant properties of a distribution for a particular application. Since the majority of sample points are observed near the centre of a distribution, the KS statistic is relatively insensitive to deviations in the tails. In the present context, this is a serious drawback because the main reason for seeking improved noise estimates is to improve target detection, which usually occurs in the upper tail of the observed distribution. To address this problem, we also use the Anderson-Darling goodness-of-fit test [15, 126], which gives more weight to observations in the tails of the distribution.

In order to derive the Anderson-Darling statistic, we must first start with the Cramér-von-Mises (CvM) statistic

$$\omega^2 = \int_{-\infty}^{\infty} [F_n(x) - \hat{F}(x)]^2 d\hat{F}(x) \quad (5.5)$$

this statistic provides equal weighting to all samples, while the Anderson-Darling statistic provides a weighting that increases sensitivity in the tails of a distribution. Let

us apply this weighting to the above formula through a weighting function $\Phi(\hat{F}(x))$

$$\omega^2 = \int_{-\infty}^{\infty} [F_n(x) - \hat{F}(x)]^2 \Phi(\hat{F}(x)) d\hat{F}(x). \quad (5.6)$$

From this equation, the standard CvM statistic can be obtained by setting $\Phi(x) = 1$, meanwhile the Anderson-Darling statistic uses the weighting $\Phi(x) = [x(1-x)]^{-2}$, leading to:

$$\omega^2 = \int_{-\infty}^{\infty} \frac{[F_n(x) - \hat{F}(x)]^2}{[\hat{F}(x)(1 - \hat{F}(x))]^2} d\hat{F}(x). \quad (5.7)$$

As we are interested in only the upper tail of the distribution, we use a weighting of $\Phi(x) = (1-x)^{-2}$. This is equivalent to applying Sinclairs derivation of the *upper-tail* Anderson-Darling test [119], defined as

$$A_n^2 = \frac{n}{2} - 2 \sum_{i=1}^n F(x_i) - \sum_{i=1}^n \left(2 - \frac{2i-1}{n}\right) \log(1 - F(x_i)) \quad (5.8)$$

The weighting applied in the upper-tail Anderson-Darling test can be seen in Figure 5.6.




Figure 5.6: Demonstration of the weighting applied in the upper-tail Anderson-Darling test, providing greatest sensitivity in the tail of the distribution

The distribution of Lilliefors' KS statistic and the upper-tail AD statistic depend on the distribution function under investigation, the number of sample points n , and whether or not the distribution parameters are estimated from the data [80, 126]. Tables of critical values for the distribution of the (two-tailed) AD statistic for a number of different distributions have previously appeared in the literature [34, 125, 127, 128]. To the best of our knowledge, such tables have not been published for the upper-tail AD test applied to the K -distribution. We have therefore performed a Monte Carlo simulation to estimate these critical values over a range of different shape parameter values, sample sizes and significance levels. Since the distribution of A_n^2 is known to be scale-invariant, we define the scale parameter to be $\lambda = 1$ throughout [40]. The procedure is outlined as follows:

1. For a given shape parameter ν generate a sample of n K -distributed random values.
2. Compute estimates of the distribution parameters for each set of sample points.
3. Compute A_n^2 using the sample points and the parameter estimates, binning the value of A_n^2 against the nearest shape parameter under consideration.
4. Repeat steps 1) to 3) to obtain 100,000 realizations of the upper-tail AD test statistic for each shape parameter value.
5. Compute the empirical distribution function of these 100,000 realizations.
6. Estimate the critical values corresponding to the required significance levels.

The results are shown in Tables 5.4 and 5.5 for samples of size $n = 1000$ and $n = 100$ respectively. The critical values are seen to decrease as ν increases, which is to be expected as the K distribution becomes less sensitive to ν . It should be noted that these

tables were calculated using the method of [66], and are only valid if ν is estimated by the iterative MME seeded with the $z \log z$ estimator of [22].

Although Table 5.5 is not used in our analysis of the data, comparing Table 5.4 with Table 5.5 clearly shows that the distribution of A_n^2 depends on the sample size n .

5.3 Analysis of Trial Data

For each ping of data, consecutive sample windows of $n = 1000$ adjacent data points were examined (in range), with a 75% overlap between consecutive windows. In order to examine predominantly reverberant and cluttered data, the data was processed from the direct blast to two-thirds of the sonar systems range. Distribution parameters were estimated in each window, and the associated cumulative distribution functions were then computed using the parameter estimates. The estimated distribution functions were compared against the observed distribution in each window, evaluated using Lilliefors' KS test and the upper-tail AD test. The results of these tests were averaged over all windows and across all pings in each frequency band.

5.3.1 Goodness-of-Fit

Reverberation from the sea floor and clutter from sea floor features often lead to a heavy tailed distribution of amplitude returns. This is related to the frequency of the ping, as higher frequencies are more quickly attenuated by the sea volume, leading to reduced returns and decreased range. Non-Gaussian properties which lead to heavy tailed non-Rayleigh data are therefore expected at low frequency, while predominantly Rayleigh data is anticipated at high frequency. In such scenarios, medium frequency offers a trade-off between range and noise.

For each distribution, Table 5.1 shows the percentage of sample windows for which the null hypothesis was not rejected by Lilliefors' KS test at the $\alpha = 0.05$ significance level. The table suggests that the *K*-distribution, Weibull and Log-normal distributions offer improved performance over the Rayleigh distribution across all frequency bands. The improvement is particularly evident in the low frequency band, where the data exhibits high levels of clutter, but there are also small regions of clutter seen at high frequency which are not fit by the Rayleigh distribution.

Similarly, Table 5.2 shows the percentage of sample windows for which the null hypothesis was not rejected by the upper-tail AD test, where the critical values for the *K*-distribution are computed from Table 5.4 by interpolation. The results show that the *K* and Weibull distributions provide a consistently better fit to the tail of the data than the Rayleigh distribution. However, across all 3 frequency bands the *K*-distribution provided a better fit than the Weibull distribution. Despite the heavy-tailed nature of the Weibull distribution, by applying the two-tailed Anderson-Darling test it was observed that the Weibull distribution provided the best fit to the lower tail of the data, which is of little interest in the present context. In addition, the log-normal distribution was also observed to provide a poor fit to the tail of the data.

In both statistical tests variance was seen in the fit of the distributions. This was frequency dependent, relating to levels of noise in the data. Due to this relationship, variance was seen to increase as frequency decreased. At Medium and High frequency the greatest variance was observed at short range and in regions containing clutter, where the distributions provided their worst fit to the sonar noise. There was far higher variance at Low frequency. This variance was seen to be lowest in poorly fit regions, suggesting a consistent difficulty in modeling reverberation and clutter.

5.3.2 Non-Rayleigh Environments

The relationship between the distribution parameters and environmental factors was investigated, where we found that scale parameters offered little information over observing the raw data. Additionally, the Weibull shape parameter behaved similarly to the log-normal variance and provided no significant information. However, the K -distribution shape parameter (ν) was seen to be sensitive to regions of clutter; small values of ν were observed exclusively in cluttered, non-Rayleigh regions. This is to be expected, because the K -distribution increasingly diverges from the Rayleigh distribution as ν decreases. Estimators for this parameter offer a means of identifying regions where the noise will be under-estimated by the Rayleigh model. This is shown in the range-bearing plots in Fig. 5.7.

If we assume that the Rayleigh model applies in regions of high clutter, we will inevitably underestimate the probability of false alarms P_{fa} . To investigate the effect on false alarm rates when applying the Rayleigh model to non-Rayleigh data, we compute a threshold x_α for the estimated Rayleigh distribution of a sample window at a given false alarm rate α , then compute the true false alarm rate that this threshold would achieve if the sample window contains K or Weibull distributed noise.

Under the assumption that the data has this Rayleigh distribution, the threshold corresponding to the probability of false alarm α is given by

$$x_\alpha = \sigma \sqrt{-2 \log(1 - \alpha)} \quad (5.9)$$

If the data actually has a hypothetical distribution $f(x; \theta)$, the probability of false alarm achieved with the threshold x_α is given by $1 - F(x_\alpha; \theta)$, where F is the CDF of the distribution. These probabilities were calculated for the K and Weibull CDFs in each sample window, using the previously estimated parameters. They were then averaged

over all windows and pings in each frequency band, and compared against five values of false alarm rate α from which the thresholds x_α were calculated. The results are shown in Table 5.3. Note that these values are only of theoretical interest since we make the assumption that the K and Weibull distributions are correctly fit to the data, which is not necessarily the case.




Figure 5.7: Demonstration of the K -distribution shape parameter highlighting regions of non-Rayleigh clutter. (a) Medium frequency matched filter envelope sonar data, (b) Shape parameter with threshold to show $\nu \leq 10$ where white denotes small values, (c) White sample windows where Lilliefors' KS test has rejected the Rayleigh hypothesis at 0.05 significance

Table 5.1: Proportion of non-rejection for Lilliefors' KS test at 0.05 significance in low, medium and high frequency bands (%).

Distribution	LF	MF	HF
K	94.09	99.22	99.38
Weibull	88.44	96.46	96.47
Log-normal	84.41	97.81	97.96
Rayleigh	72.01	88.75	89.71

Table 5.2: Proportion of non-rejection for the upper-tail AD test at 0.05 significance in low, medium and high frequency bands (%).

Distribution	LF	MF	HF
K	40.48	65.66	88.27
Weibull	39.20	60.62	84.72
Rayleigh	31.10	47.93	75.82
Log-normal	0.07	0.03	0.06

Table 5.3: P_{fa} obtained when selecting a threshold by assuming Rayleigh distributed noise

Target P_{fa}	K -distribution			Weibull distribution		
	LF	MF	HF	LF	MF	HF
0.01	0.0149	0.0140	0.0138	0.0128	0.0129	0.0128
0.001	0.0033	0.0031	0.0030	0.0025	0.0031	0.0027
0.0001	0.0012	0.0013	0.0011	0.0010	0.0018	0.0013
0.00001	0.0006	0.0008	0.0006	0.0006	0.0014	0.0010
0.000001	0.0004	0.0005	0.0004	0.0005	0.0013	0.0009

Table 5.4: Upper-tail Anderson-Darling test critical values for the K -distribution, various significance levels, sample size 1000

Shape	Significance level				
	0.25	0.10	0.05	0.025	0.01
0.05	0.383	0.567	0.711	0.855	1.059
0.25	0.309	0.440	0.540	0.640	0.766
0.5	0.289	0.410	0.504	0.600	0.728
0.75	0.268	0.375	0.458	0.538	0.650
1	0.252	0.349	0.424	0.498	0.595
1.5	0.239	0.329	0.396	0.461	0.548
2	0.236	0.324	0.386	0.450	0.531
3	0.230	0.313	0.376	0.438	0.518
4	0.224	0.305	0.368	0.429	0.512
5	0.224	0.306	0.368	0.427	0.511
10	0.220	0.301	0.362	0.424	0.499
15	0.221	0.302	0.363	0.423	0.503
20	0.219	0.297	0.357	0.418	0.496
25	0.218	0.298	0.359	0.419	0.497
30	0.217	0.297	0.357	0.416	0.494

Table 5.5: Upper-tail Anderson-Darling test critical values for the K -distribution, various significance levels, sample size 100

Shape	Significance level				
	0.25	0.10	0.05	0.025	0.01
0.05	0.368	0.541	0.683	0.822	1.016
0.25	0.278	0.386	0.465	0.543	0.644
0.5	0.259	0.357	0.431	0.504	0.603
0.75	0.241	0.329	0.395	0.462	0.548
1	0.235	0.318	0.380	0.444	0.528
1.5	0.227	0.308	0.368	0.431	0.506
2	0.222	0.303	0.363	0.424	0.501
3	0.221	0.301	0.362	0.421	0.504
4	0.222	0.301	0.362	0.422	0.499
5	0.221	0.300	0.360	0.421	0.498
10	0.219	0.300	0.359	0.421	0.498
15	0.216	0.293	0.353	0.413	0.489
20	0.213	0.290	0.347	0.404	0.480
25	0.213	0.291	0.349	0.406	0.486
30	0.215	0.294	0.354	0.413	0.491

5.4 Discussion

Compared to the standard Rayleigh model, we have shown that both the K and Weibull distributions offer improved models of clutter-limited matched filter envelope data. In contrast, the log-normal distribution performed poorly according to the upper-tail AD test, and is therefore a poor model of noise in this dataset.

Using the upper-tail AD test, we have shown that the K and Weibull distributions are better able to model the tail of the observed data than the Rayleigh distribution, which is of great importance in performance prediction. Our analysis of the data has also shown that the K -distribution provides a better model of the tails than that provided by the Weibull distribution. While the Weibull distribution gave a reasonable fit to the noise data we must therefore conclude that, in addition to providing a useful physical model of sonar clutter, the K -distribution is also the most suitable model for accurate selection of detection thresholds. Additionally as a result of the compound nature of the K -distribution it may also be possible to extend the distribution to include further physical properties, in the same manner that the KA -distribution was obtained by modeling further properties of radar returns via the Class A model [138][pp109-122].

The negative effect of using the Rayleigh distribution to model data with non-Rayleigh properties was found to be most significant when setting low false alarm rates, which is often necessary in real-time applications where targets must be quickly identified. We have also shown that an estimate of the K -distribution shape parameter ν provides a simple way of identifying regions of high clutter in the data.

Considering the above findings, the upper-tail AD results show that Lilliefors' KS test alone can not demonstrate the suitability of a distribution when a good fit in the tail is important. This is particularly relevant to the problem of target detection in radar and sonar systems.

While upper-tail AD test critical values for the K -distribution were not derived analytically, the improved accuracy offered by such a derivation would be of considerable interest.

This chapter has provided an analysis of the parameter estimators for the K -distribution, and established that the K -distribution is an accurate model of the sonar noise in our trial data. It has also been shown that the Kolmogorov-Smirnov test is unsuitable for evaluating goodness-of-fit in the tail of the candidate distribution, and that the upper-tail Anderson-Darling test is better suited to this problem. From these findings and the previous chapter we are able to accurately estimate and model the sonar noise in our trial data. These findings will be applied to target detection later in this thesis, but we must first consider what we are trying to detect. In the next chapter we introduce models for target signals and examine the relationship between the target and its surrounding noise environment.

Chapter 6

Target Distributions

“

All models are wrong, but some are useful.

— GEORGE E. P. BOX (1919–)

Detection performance depends on the statistics of both target and noise data, as noted in Chapter 2. The design of a detection scheme must make some assumptions of a targets statistics, in order to separate a target from naturally occurring noise in the environment. Furthermore, the probability of detection cannot be calculated without targets present, and it is therefore necessary to include target data in the analysis of a detectors performance.

This chapter describes statistical models for targets in a sonar environment. Initially, we examine models for sonar noise which can incorporate an additional signal. We then examine models for target signals which can be embedded in these noise models. Since the trial data provided by Thales Underwater Systems only contains sonar noise

(see Chapter 3), it will be necessary to simulate targets in this thesis. Consequently, techniques are presented which enable us to generate synthetic targets according to these models.

6.1 Embedding Models

In this section we will define five distributions which can be used to model a target signal embedded in sonar noise. We begin with three distributions which are explicit extensions of noise distributions we have seen previously in Chapter 4.

6.1.1 The Rice Distribution

The Rice (or Rician) distribution is an extension of the Rayleigh distribution (Section 4.1), and has the following probability density function:

$$f_{\text{RICE}}(x|S, \sigma)_R = \frac{x}{\sigma^2} \exp\left(\frac{-(x^2 + S^2)}{2\sigma^2}\right) I_0\left(\frac{Sx}{\sigma^2}\right) \quad (6.1)$$

where I_0 is the modified Bessel function of the first kind, σ represents the *scale* of the Rayleigh background, and the target signal is represented by S . The impact of the target signal parameter S is demonstrated by plots of the Rician distribution probability density function in Figure 6.1. The Rice distribution is obtained by embedding a signal in Rayleigh distributed data. In addition to the Rayleigh and Nakagami distributions, the Rice distribution is also used in communications theory to model scattered signals that reach a receiver by multiple paths [49]. Estimation is possible using either MME or MLE techniques [108, 118].

The target amplitude is equal to the difference between the second moment of the

Rice distribution and the second moment of the Rayleigh distributed background noise. Therefore, if A_X represents a set of amplitude values containing a target signal embedded in Rayleigh distributed noise, and A_Y represents a set of amplitude values which are Rayleigh distributed with the same scale parameter σ , then we could estimate the target signal S as follows:

$$A_Y \sim f_{\text{RAYLEIGH}}(\sigma) \quad (6.2)$$

$$A_X \sim f_{\text{RICE}}(\nu, \sigma) \quad (6.3)$$

$$\hat{\sigma} = \sqrt{\frac{E(A_Y^2)}{2}} \quad (6.4)$$

$$\hat{S} = \sqrt{E(A_X^2) - 2\sigma^2} \quad (6.5)$$

$$= \sqrt{E(A_X^2) - E(A_Y^2)} \quad (6.6)$$

where $E(\cdot)$ denotes expectation, or the arithmetic mean of the value. The estimate (6.4) of the noise-only amplitude σ is based on the first and second moment of the Rayleigh distribution, while the estimate (6.6) for S is based on the MME estimate of the Rician distribution. When estimating the parameters of the Rician distribution by the MME, \hat{S} is normally calculated with an estimate of σ taken from Rician distributed data. This alternative formulation is included to demonstrate the relationship between the second moment of the Rician and Rayleigh distributions, and is not applied in this work.

6.1.2 The Homodyned K-Distribution

The Homodyned K-distribution, or HK-distribution has the probability density function:

$$f_{\text{HK}}(x|S, \nu, \lambda) = \frac{2x\lambda^\nu}{\Gamma(\nu)} \int_0^\infty y^{\nu-2} \exp(-\lambda y) \exp\left(\frac{-(S^2 + x^2)}{y}\right) I_0\left(\frac{2xS}{y}\right) dy \quad (6.7)$$

where I_0 is the modified Bessel function of the first kind, ν and λ are the shape and scale of the K-distributed noise, and S is the target signal. The impact of the target signal parameter S is demonstrated by plots of the Homodyned K-distribution probability density function in Figure 6.2. This distribution, introduced by Jakeman [61], is an extension of the K-distribution. The term “Homodyned” refers to the mixing of two signals of the same frequency. On this basis, the HK-distribution models a signal embedded in K-distributed noise. Its probability density function is an integral over a modified Bessel function of the first kind, and in general cannot be solved analytically. While it is possible to evaluate the distribution numerically, this lack of an analytical solution still causes great difficulty in obtaining parameter estimates.

Two main approaches exist for estimating parameters of the HK-distribution [55, 56, 58, 92], these are the MME and estimation via level sets. Estimation using level sets is similar to the method behind the Hough transform, where multiple curves intersect to identify a unique solution [124]. In this case, each curve is generated for a particular parameter value, rather than as a transformation of all data points. It should be noted that there is no closed-form derivation of these curves and they must be calculated a priori through simulation [58, 92].

On the other hand the MME can be implemented either by a direct analytic approach or through numerical minimization. However, the sample sizes used in this thesis are too small for the MME approach to reliably obtain valid (real and positive) parameter values analytically, and numerical minimizations often converge on local minima, resulting in estimates providing a poor match to the observed data.

It should be noted that there is some controversy associated with this distribution. Jakeman himself has stated that, from a mathematical perspective, the Homodyned K distribution is an unattractive model [63]. It is also known that authors in the field of ultrasound have chosen alternatives such as the Nakagami and generalized gamma distributions, which we will define shortly [58, 92]. Due to the need for numerical eval-

uation, and our difficulties in obtaining accurate parameter estimates, the Homodyned K-distribution will not be used in this thesis, and is only included to illustrate how a target signal should be embedded in K-distributed sonar noise.

6.1.3 The Generalized K-Distribution

The generalized K-distribution, or GK-distribution has the probability density function:

$$\begin{aligned} f_{\text{GK}}(x|S, \nu, \lambda) &= \frac{2x\lambda^\nu}{\Gamma(\nu)} I_0(2xS) \int_0^\infty y^{\nu-2} \exp(-(\lambda + S^2)y) \exp\left(\frac{-x^2}{y}\right) dy \\ &= \frac{4x^\nu \lambda^\nu}{\Gamma(\nu)(\lambda + S^2)^{(\nu-1)/2}} I_0(2xS) K_{\nu-1}(2x\sqrt{\lambda + S^2}) \end{aligned} \quad (6.8)$$

where I_0 is the modified Bessel function of the first kind, K_ν is the modified Bessel function of the second kind, ν and λ are respectively the shape and scale of the K-distributed noise, and S is the target signal. The impact of the target signal parameter S is demonstrated by plots of the Generalized K-distribution probability density function in Figure 6.3. This distribution, first introduced by Barakat [17], is a generalization of both the Rician and K-distributions [63]. Recall that the K-distribution is a Rayleigh distribution modulated by a gamma distribution (Section 4.1), and that a constant amplitude embedded in a Rayleigh distribution results in a Rician distribution. The GK-distribution combines these two processes, resulting in a Rician distribution modulated by a gamma distribution.

The GK-distribution has been used in radar research [39, 48, 97], where it provides a continuous and physically consistent transition among different scattering scenarios. The distribution parameters have also been used as feature descriptors for the detection of oil spills [48, 97].

In contrast to the HK-distribution, the GK-distribution also modulates the target signal




Figure 6.1: Rician distribution probability density function for various values of the target signal S




Figure 6.2: Homodyned K-distribution probability density function for various values of the target signal S

by the gamma distribution. The amplitude of this distribution can be expressed in closed form as demonstrated in (6.8) above. Distribution parameters can be estimated by MME [17] or MLE [102]. However, as with the MME for the HK-distribution, the sample sizes used in this thesis are too small for analytic approaches to provide accurate parameter estimates. Furthermore the MLE method, which must be evaluated numerically, suffers from long computation time and difficulties with local maxima.

6.1.4 The Nakagami Distribution

The Nakagami distribution has the probability density function:

$$f_{\text{NAKAGAMI}}(x|\mu, \omega) = \frac{2\mu^\mu}{\Gamma(\mu)\omega^\mu} x^{2\mu-1} \exp\left(-\frac{\mu x^2}{\omega}\right) \quad (6.9)$$

where μ is the shape, and ω is the spread of the distribution. The impact of the shape parameter μ is demonstrated by plots of the Nakagami distribution probability density function in Figure 6.3. Due to its simplicity, as compared to (6.7) and (6.8), the Nakagami distribution has been favoured by some in the ultrasound community over the HK and GK-distributions [24]. This distribution includes the Rayleigh distribution as a special case and can also approximate the Rician distribution [117]. Parameters are easily estimated through MLE and MME [24, 108, 117].

Although this distribution is easy to apply, the parametrisation does not include an explicit parameter for the amplitude of the target signal, as provided by the Rician distribution. This makes it difficult to specify an explicit target signal, but the Nakagami distribution may still prove useful if it is capable of accurately fitting target data.




Figure 6.3: Generalized K-distribution probability density function for various values of the target signal S




Figure 6.4: Nakagami distribution probability density function for various values of the shape parameter μ

6.1.5 The Generalized Gamma Distribution

The generalized gamma distribution, or GGamma distribution has the probability density function:

$$f_{\text{GG}}(x|c, \alpha, \beta) = \frac{c}{\Gamma(\alpha)\beta^\alpha} x^{c\alpha-1} \exp\left(-\frac{x^c}{\beta}\right) \quad (6.10)$$

We have reparametrised this distribution such that α, β are the parameters of the underlying gamma distribution, and c is an exponential factor (such that $c = 1$ recovers the underlying gamma distribution). The generalized gamma distribution can be derived through statistical-mechanical methods [79], demonstrating a physical basis for its use in the modelling of random processes. This distribution is capable of reducing to at least seven special cases, including the Rayleigh, Gamma, Nakagami and Weibull distributions [108]. The source of its general applicability is found in the two shape parameters which allow the upper and lower tail to be adjusted independently. As a result the generalized gamma distribution is capable of modelling cluttered (Weibull) and target (approximately Rician) data with great flexibility.

As noted in [73] the generalized gamma distribution can be obtained by raising a gamma distributed variable to a power. This transformation leads to a simple iterative estimation method described in [50], which can use either MME or MLE to estimate the underlying parameters.

It should be noted that, as with the Nakagami distribution, there is no single parameter representing the amplitude of the target signal, but once again this distribution may be useful if it is capable of accurately fitting data containing a target.

6.2 Swerling Models

The so-called Swerling target models are well established in radar literature [121, 131]. These models describe the time-varying nature of the signal received from a target, in the absence of noise. Embedding a Swerling model in sonar noise then allows us to model the amplitude of a target signal received in a sonar environment.

As they were originally defined for radar, the Swerling models are often defined in terms of *square-law* detectors, which apply to data in the intensity domain. Since our data is defined in the amplitude domain, we must take care to use the equivalent Swerling models for linear detectors.

As noted by Richards [115] there is a great deal of confusion in this literature. Those familiar with the field often name distributions in the intensity domain after the equivalent distribution in the amplitude domain. For example, the Swerling I model, represented by the exponential distribution in the intensity domain, is often erroneously described as Rayleigh. This may lead the unwary to believe that a Rayleigh distribution is appropriate for modelling data in the intensity domain, and thus the square root of a Rayleigh distributed variable is required in the amplitude domain. We will state each of the models as defined in the intensity domain for radar [121], followed by an explicit description of their application in the amplitude domain.

The Swerling distributions are derived from the following probability density function:

$$f_{\text{SWERLING}}(x|\mu, m) = \frac{m}{\Gamma(m)\mu} \left(\frac{mx}{\mu}\right)^{m-1} e^{-\frac{mx}{\mu}}, x > 0 \quad (6.11)$$

where μ is the first moment of the data and m depends on the target model used, with $m = \infty$ for a Swerling 0 target, $m = 1$ for a Swerling I target, and $m = 2$ for a Swerling III target. The various shapes provided by the Swerling models are demonstrated in Figure 6.5. It should be noted that the other Swerling models (II and IV) signify that

the target data varies between radar pulses rather than between radar scans. There is no equivalent distinction in sonar and so the Swerling II and Swerling IV models are not discussed here.




Figure 6.5: Comparison of the Swerling 0, Swerling I, and Swerling III target models for a mean value of 2

6.2.1 The Swerling 0 Model

The simplest approach is to model the target amplitude as a constant value. This was first investigated by Marcum [87], but has since been incorporated into the set of Swerling models, where it is commonly known as Swerling 0 or Swerling V. In this case $m \rightarrow \infty$ and so (6.11) reduces to the degenerate distribution:

$$f(x|\mu) = \begin{cases} 1, & \text{if } x = \mu \\ 0, & \text{if } x \neq \mu \end{cases} \quad (6.12)$$

This is the basic case of applying a non-varying target to a noise distribution, as demonstrated with the Rician, HK and GK-distributions in Section 4.4.

6.2.2 The Swerling I Model

The Swerling I model is suited to a target signal from a large number of independent surfaces of approximately equal strength. Interpreted as a Poisson process, where the number of independent events occurring during an interval of time is examined, this leads to an exponential distribution, which is modelled by a Swerling distribution with $m = 1$:

$$f(x|\mu) = \frac{1}{\mu} \exp\left(-\frac{x}{\mu}\right) \quad (6.13)$$

Thus, Swerling I targets are Rayleigh distributed in the amplitude domain. This should not be surprising in view of the absence of a dominant scatterer, which for the Rician embedded model also leads to a Rayleigh distribution (in the absence of a target).

6.2.3 The Swerling III Model

If there is a single dominant scatterer with a large number of independent scatterers of approximately equal strength, we obtain a Swerling distribution with $m = 2$:

$$f(x|\mu) = \frac{4x}{\mu^2} \exp\left(-\frac{2x}{\mu}\right) \quad (6.14)$$

In the amplitude domain, this is approximately equal to the Rician distribution [113]. Interestingly, this approximation makes the assumption that the ratio of the strength of the dominant scatterer to that of the secondary scatterers is $1 + \sqrt{2}$ [114]. Of course, this restriction does not necessarily hold in real sonar data, so a Rician distribution is more useful in the present context.

6.3 Generating Data

We now consider how synthetic target signals, following the models described above, can be embedded into the noise-only trial data described in Chapter 3. Data containing both target and noise signals will be referred to as having an *embedded target* distribution.

6.3.1 The Rice Distribution

We first define the amplitudes of a target signal to be an additional component in the complex data returned by the matched filter, with its phase ϕ determined by the proportion to which the target resides in the real or imaginary component. Therefore the amplitude of a target embedded in noise can be defined in terms of (4.43) as follows:

$$\begin{aligned}
 Y_0 &\sim \text{Gaussian}(S \times \cos(\phi), \sigma^2) \\
 Y_1 &\sim \text{Gaussian}(S \times \sin(\phi), \sigma^2) \\
 X &= \sqrt{Y_0^2 + Y_1^2} \\
 &\sim \text{Rician}(\sigma, S)
 \end{aligned} \tag{6.15}$$

where S is the amplitude of the target signal, and σ is the scale of the Rayleigh distributed background noise.

Thus for a target signal with amplitude S embedded in a noise-limited sonar environment, the embedded target distribution is a Rician distribution. As the phase of the signal is not used in this project, the target signal need only be in either the real or imaginary domain, so by arbitrarily setting $\phi = 0$, we obtain (6.15) in the simplified

form

$$\begin{aligned}
 Y_0 &\sim N(S, \sigma^2) \\
 Y_1 &\sim N(0, \sigma^2) \\
 X &= \sqrt{Y_0^2 + Y_1^2} \\
 &\sim \text{Rician}(\sigma, S).
 \end{aligned} \tag{6.16}$$

Since the Rician distribution is an extension of the Rayleigh distribution to include a target amplitude, this can be directly applied to the modular random number generator seen in (4.46).

6.3.2 The Generalized K-Distribution

By modulating a Rician variable by a Gamma variable we obtain the *generalized K-distribution*,

$$\begin{aligned}
 Z &\sim \text{Rician}(1, S) \\
 W^2 &\sim \text{Gamma}(\nu, \frac{\lambda}{2\nu}) \\
 X &= Z \times W \\
 &\sim GK(\nu, \lambda, S).
 \end{aligned} \tag{6.17}$$

While this may seem a reasonable approach to embedding a target signal S in K -distributed noise, it does not correspond to the way a target is combined with noise in the real world. As seen in (6.17) both the noise and target components are modulated by the gamma distribution. However, target signals are not influenced by clutter or reverberation, and so a gamma modulated target component is inappropriate.

6.3.3 The Homodyned K-Distribution

Rather than modulating a Rician variable by a Gamma variable, in order to obtain a target signal S embedded in K-distributed noise, we should instead modulate the underlying Rayleigh component of a Rician variable, leaving the target signal as-is,

$$\begin{aligned}
 Y_0 &\sim N(0, 1) \\
 Y_1 &\sim N(0, 1) \\
 W^2 &\sim \text{Gamma}\left(\nu, \frac{\lambda}{2\nu}\right) \\
 X &= \sqrt{(S + Y_0 \times W)^2 + (Y_1 \times W)^2} \\
 &\sim HK(\nu, \lambda, S).
 \end{aligned} \tag{6.18}$$

This generates random samples from a *homodyned K-distribution*. The subtle differences between the K-distribution, generalized K-distribution and homodyned K-distribution are demonstrated in Figures 6.6-6.7. By keeping the target signal independent of the (K-distributed) noise we achieve a random sample which is a more accurate representation of sonar targets in the real world.

The form of (6.18) suggests that, given complex sonar noise from a matched filter in a real system, embedded target data can be easily simulated by a simple addition, before calculating the amplitude of the signal. If Y denotes noise signal as a complex valued output from a matched filter, and S denotes target signal without noise, then this can be written as:

$$Y = \Re(Y) + \Im(Y) \tag{6.19}$$

$$A_X = |X| = |S + Y| \tag{6.20}$$

$$= \sqrt{(S + \Re(Y))^2 + \Im(Y)^2}. \tag{6.21}$$

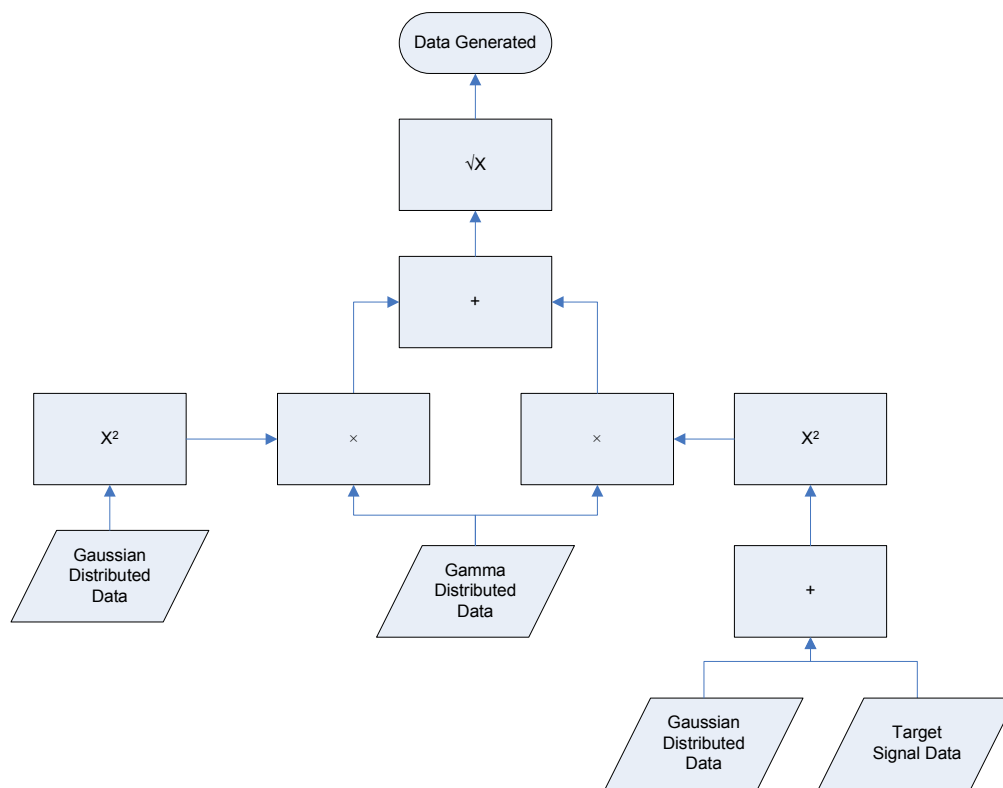


Figure 6.6: Generation of generalized K-distributed values, given 2 independent Gaussian variables, a Gamma variable and a target signal.

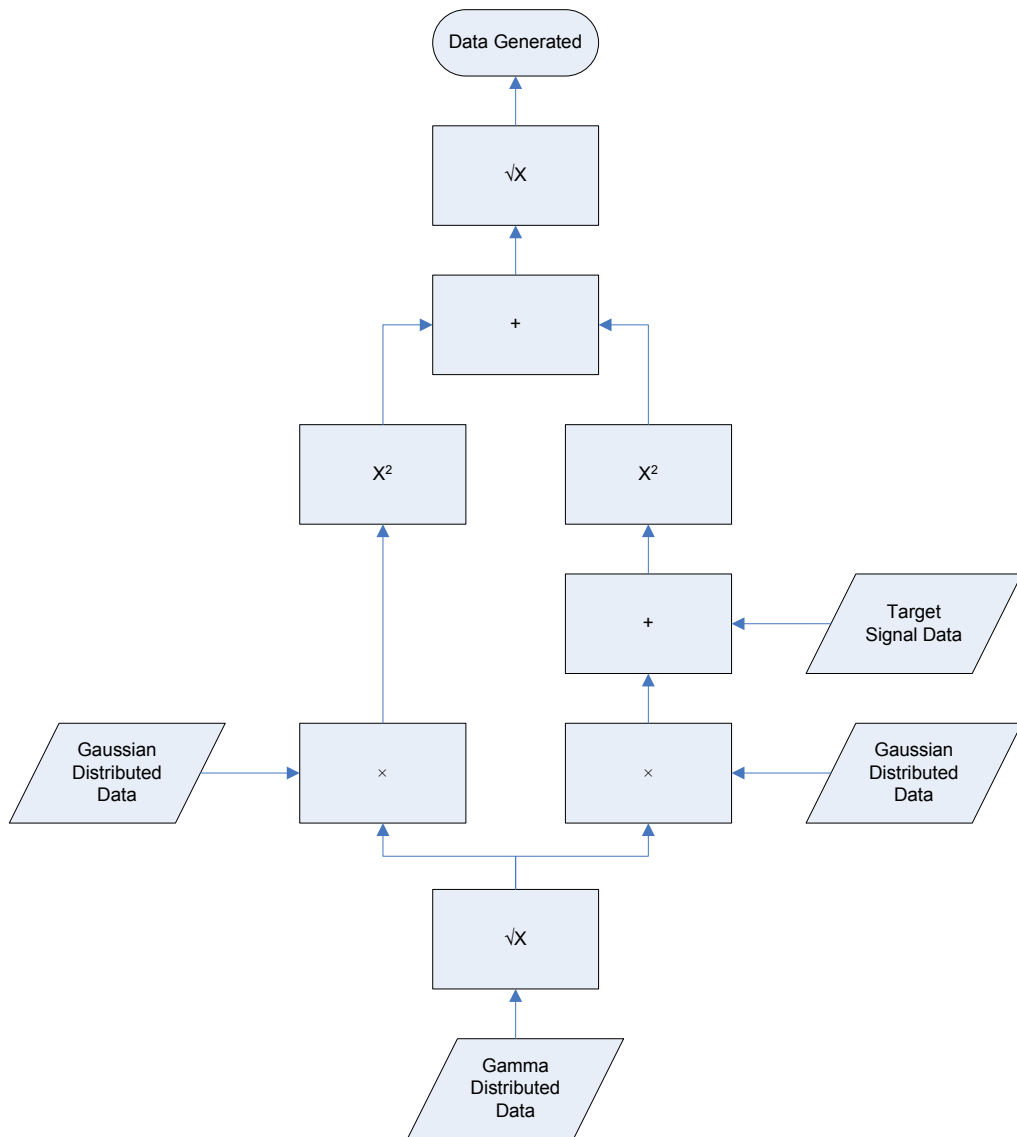


Figure 6.7: Generation of homodyned K-distributed values, given 2 independent Gaussian variables, a Gamma variable and a target signal.

Given that the K-distribution is a generalization of the Rayleigh distribution, this model evidently applies to both noise and clutter-limited environments. This will prove a useful tool in Chapters 7 and 8, where synthetic targets are embedded directly into our trial data.

This chapter has introduced a number of statistical models for a target signal in a noise environment, of which the Homodyned K-distribution is the most accurate representation of a signal in K-distributed noise. We have outlined the Swerling models for target signals in the absence of noise, and provided methods for the generation of Swerling 0 targets in various representations of the noise environment. Between the discussion of targets provided here and the noise discussed in Chapter 4, we have established the two classes of data which a sonar detector should accurately identify. In the next chapter these models are used to define and validate various detection schemes. This will also include a performance analysis based on simulated targets embedded in our noise-only trial data, and a brief application to real targets.

Chapter 7

Target Detection

“

You can't defend. You can't prevent. The only thing you can do is detect and respond.

— BRUCE SCHNEIER (1963–)

When performing detection in sonar signal processing, it is necessary to first model the noise in the data. Data which does not fit the model is generally considered to include a target signal. Data is tested for the presence of a target by first estimating the noise distribution in the region surrounding the data under investigation (assuming that the surrounding region contains no target), then comparing the data under investigation to the estimated noise distribution. Further advances can be made by including a model for the target embedded in sonar noise, and making a comparison between estimates of the noise-only and embedded target distributions. However, in order to identify whether a different detection algorithm achieves greater performance it is necessary to generate a ROC curve, as described in Chapter 2.

In this chapter we define and evaluate various detection schemes, applied to both synthetic and real data. The chapter begins by describing the techniques used in our performance analysis, including the number of detection tests to evaluate, and practical considerations to avoid bias. The standard box-car detector from Chapter 2 is reintroduced, and extended to incorporate any noise model of our choosing. We then outline the theory of *Generalized Likelihood Ratio Tests* (GLRTs), and design various GLRT detectors. These GLRT detectors are compared using synthetic data from Chapters 4 and 6, which enables us to select the GLR statistics with the most promising performance. We then apply a selection of box-car and GLRT detectors to simulated noise environments with simulated targets, and analyse their performance. This analysis is then extended to their application in real sonar noise, where they are applied to simulated targets embedded in the trial data. Lastly, the detectors are applied to a dataset with real targets, to which Thales Underwater Systems provided limited access in a secure environment.

7.1 Experimental Design

As previously demonstrated in Chapter 2, by varying the threshold value T , the probability of detection and probability of false alarm define the *Receiver Operating Characteristic* (ROC) curve (see Figure 2.22). The detection performance of our detection algorithms will be evaluated by inspecting ROC curves. ROC curves can be produced using real data collected during an experiment, or alternatively from synthetic data generated by simulation. Both techniques have advantages and disadvantages. Although experimentation can demonstrate real-world performance, it can be expensive, and it may be difficult to design experiments which take all possible environmental scenarios into consideration. By contrast, simulations can be executed cheaply and allow for complete control over the environment; however, the observed performance is merely

an estimate of detection performance in the real world. Trial data has been provided by Thales Underwater Systems for this study, which contains real sonar noise but does not include any targets (see Chapter 3). First we investigate data that has been entirely simulated, then data in which simulated targets are embedded in real sonar noise, and finally data for real targets in real noise.

In contrast to dedicated sonar equipment, which tends to include bespoke, highly parallelized hardware [33, 77], the equipment available for our analysis was a standard multi-core desktop PC. Due to the limitations in our processing capacity, it is prohibitively time consuming to analyse detection performance by applying a detection test at every cell in the trial data. We must therefore design a suitable experiment which will capture a representative sample of cells, allowing us to approximate the performance of the system as a whole.

7.1.1 Number of Detection Tests

The most sensitive parameter to consider in a sampled performance analysis is the number of detection tests we should apply, which we denote by N_D . A simple rule of thumb is to derive this value from the smallest false alarm rate we wish to represent on our ROC curve. For example, if we wish to investigate false alarm rates of 10^{-6} , we observe, on average, a single false alarm in every million detection tests. The minimum number of detection tests needed to reach this false alarm rate is one million. It also remains beneficial to perform a larger number of tests, as this will reduce any statistical variation in the results, and so N_D is defined as:

$$N_D \geq \left\lceil \frac{1}{\alpha} \right\rceil \quad (7.1)$$

where α is the desired false alarm rate, and $\lceil \frac{1}{\alpha} \rceil$ denotes the smallest integer value greater than or equal to $\frac{1}{\alpha}$.

Our trial data contains multiple pings of sonar noise, and so the detection tests will be evenly split between each of these. Consequently the number of detection tests applied to each ping will actually be $\lceil N_D/N_P \rceil$, where N_P is the number of pings. In this trial data, some cells contain values sampled prior to the initial blast, which is not useful. Additionally, the low-frequency sonar system used was not designed to detect targets at very short range. Consequently no targets will be simulated, or detection tests executed, within the first 500m after the initial blast. Recalling the useful range of our trial data, which is defined in Section 3.2 to be approximately 21.5km, this leaves 21km of data for processing. With a sampling rate of 8.22 cells per metre, this corresponds to a matrix of 64 beams by 172,620 cells in range, leading to a total of 11,047,680 cells in total per ping. As stated in Section 2.3.1, the resolution of sonar systems employing FM pings is related to the bandwidth of the ping by $c/2B$. From Section 3.1, the bandwidth of the three frequency bands included in this trial data are 900 Hz for LF, 1.6kHz for MF and 3.8kHz for HF. Consequently the range resolution of the three frequency bands are approximately 83.3cm for LF, 46.9cm for MF and 19.7cm for HF, or equivalently 1.2 independent cells per meter at LF, 2.1 independent cells per meter at MF, and 5 independent cells per meter at HF.

7.1.2 Spatial Distribution of Detection Tests

If we assume that all values in this truncated dataset are equally important, a representative sample of cells should be evenly-spaced across each ping, this could be achieved by splitting the number of detection tests evenly between beams, and then distributing the detection tests uniformly in range, using either a uniform random variable, or an evenly spaced grid across the data in range-bearing space. This may lead to unwanted

side-effects if the sampling is not representative of the dataset as a whole.

On closer inspection, the assumption that every cell is as important as any other is flawed. Considering the polar nature of sonar data, as demonstrated in Section 3.3, samples at long range cover a larger area than samples at shorter ranges. Consequently, if a sonar target is as likely to be at one GPS coordinate as any other, using an even grid in polar coordinates results in a bias towards short-range cells, leading to a large number of tests taking place at short range, and too few tests at long range. Therefore when applying sampled detection tests to our trial data, a grid of GPS coordinates was overlaid over the area covered by each ping. By simulating targets on an even grid in real world coordinates, spatial bias is avoided.

To locate the $\{r, \theta\}$ cells required for each simulation, the $\{\varphi, \lambda\}$ location must be converted from GPS coordinates back into polar coordinates. Given the GPS coordinates of the sonar system for a particular ping $\{\varphi_P, \lambda_P\}$ and the GPS coordinates for a detection test $\{\varphi_D, \lambda_D\}$, we can calculate the range and bearing coordinates using equations (3.1) and (3.6) previously defined in Chapter 3.

The bearing is then matched to the closest beam in our trial data, providing the polar coordinates needed to apply our detection test. This method of performance analysis by sampled simulation will be applied in the later stages of this chapter, and also forms the basis of our analysis of multi-ping detection in Chapter 8.

7.1.3 Target Simulation

Further considerations for simulated ROC curves include:

- how the target will be represented in the dataset, and

- how this will change from one ping to another.

In Section 5 we established that the K-distribution is the most accurate model of noise in our trial data, and consequently the Homodyned K-Distribution (Section 6.1) is the most suitable target model for our performance analysis. As is shown in (6.21), the compound form of the Homodyned K-Distribution lends itself well to the simulation of synthetic target signals in complex data. Our trial data was provided as matrices of complex values from a matched filter, so a target signal of our choosing can be embedded in the data by simple addition before calculating the amplitude of each complex value (see Section 6.3). With limited access to real target data, and in the absence of a definitive standard, we elected to model our target signals using the Swerling 0 model (see Section 6.2). This model specifies a non-fluctuating target, and so the target signal is constant for all the cells occupied by the target.

From Section 6.3, we know that a target can be simulated by taking the amplitude of a target signal added to the complex noise-only data obtained from the matched filter. From (6.21) we will embed target signals in our noise-only data as follows:

$$A_X = |S + Y| = \sqrt{(S + \Re(Y))^2 + \Im(Y)^2} \quad (7.2)$$

Signal-to-Noise Ratio

To select a suitable value for S we must examine a measurement of relative target strength, known as the *Signal-to-Noise Ratio* (SNR). In its simplest form, the SNR is simply the amplitude of the target data divided by the amplitude of the noise in the local environment (A_X/A_Y). However, as embedded target data contains both a target signal S and a noise signal Y , we reduce the impact of the noise signal by subtracting A_Y from A_X . For historical reasons, values of SNR are often expressed in terms of

a base-10 logarithmic scale, known as decibels (dB). In this thesis we specify SNR as follows¹:

$$\text{SNR} = 10 \log_{10} \left(\frac{A_X - A_Y}{A_Y} \right) \quad (7.3)$$

Substituting (7.2) into (7.3), we obtain:

$$\text{SNR} = 10 \log_{10} \left(\frac{\sqrt{(S + \Re(Y))^2 + \Im(Y)^2} - A_Y}{A_Y} \right) \quad (7.4)$$

By rearranging (7.4) we can define the amplitude of target signal S needed to achieve a chosen SNR in a background of noise with amplitude A_Y :

$$S = \left[(A_Y \cdot 10^{\text{SNR}/10} + A_Y)^2 - \Im(Y)^2 \right]^{\frac{1}{2}} - \Re(Y) \quad (7.5)$$

While easy to use, these equations only work for single values of embedded target and noise-only data, which we know vary with range and bearing. Consequently, these equations must be extended for use over multiple cells in range. This is done by making use of summary statistics for A_X and A_Y , and so (7.3) is extended to multiple cells by:

$$\text{SNR} = 10 \log_{10} \left(\frac{\bar{A}_X - \bar{A}_Y}{\bar{A}_Y} \right) \quad (7.6)$$

where \bar{A}_X denotes the arithmetic mean of A_X and \bar{A}_Y denotes the arithmetic mean of A_Y , over the extent of the target area. In this extended form, it is no longer possible to derive a closed form solution for the target signal S , and we must therefore solve this numerically. This process begins by calculating the value of A_X needed for a desired

¹Although our definition is convenient as we are dealing exclusively in the amplitude domain, SNR is typically defined as a power ratio, of the form $10 \log_{10}(A_X^2/A_Y^2)$, or equivalently $20 \log_{10}(A_X/A_Y)$. This form, commonly used in the standard sonar equations, can be obtained from Equation 7.3 by the transformation $20 \log_{10} [10^{\text{SNR}/10} + 1]$.

SNR, which we will denote \check{A}_X :

$$\check{A}_X = \left(\sum_{i=1}^n A_Y(i) \right) (10^{\text{SNR}/10} + 1) \quad (7.7)$$

The true value of A_X is defined by (7.2), so to calculate S we must numerically solve:

$$\check{A}_X - \sum_{i=1}^n \left[(S + \Re\{Y(i)\})^2 + \Im\{Y(i)\}^2 \right]^{\frac{1}{2}} = 0 \quad (7.8)$$

This two step process provides an accurate calculation of the value of a Swerling 0 target signal S required for a desired target SNR in data where the background noise has amplitude A_Y . As defined in (7.2), it is then just a case of adding S to the complex noise in which we wish to embed our target, and taking the magnitude.

7.2 The Box-car Detector

A simple detection algorithm which is often applied to sonar data is the *box-car detector*. This is a standard method of detection in sonar [136], as previously described in Section 2.7

From the cumulative distribution function, we are able to calculate the probability of noise-only data (which fits this distribution) exceeding a chosen threshold T . From Section 2.8, using the Gaussian distribution as a model for the null hypothesis, with parameters $\theta = \{\mu, \sigma\}$, this is calculated as:

$$\alpha = 1 - F(T|\theta, H_0) \quad (7.9)$$

$$= 1 - F_{\text{GAUSSIAN}}(T|0, 1) \quad (7.10)$$

$$= 1 - \frac{1}{2} \left[1 + \text{erf} \left(\frac{T}{\sqrt{2}} \right) \right] \quad (7.11)$$

where erf denotes the Gauss error function. Additionally, by inverting this equation we can calculate the threshold required for a chosen false alarm rate:

$$T = \left(\sqrt{2}\right) \operatorname{erf}^{-1}(1 - 2\alpha) \quad (7.12)$$

These two equations require that the original noise is Gaussian distributed, and that the data is normalized using the true population parameters.

We have already established that the noise in our trial data does not have a Gaussian distribution. In Section 5.3 it was shown that the K and Weibull distributions are appropriate models of the sonar noise in our trial data. Additionally, by using empirical estimates of μ and σ we introduce some uncertainty into the model. The above method, which we will henceforth refer to as the *standard box-car* detector, can be generalized for use with any distribution. This only requires that we change the normalization procedure shown in (2.17), and calculate thresholds by substituting the CDF of the required noise distribution into (7.9).

In (7.9), if we subtract the cumulative probability of a value from 1, we obtain the probability of false alarm at which this value would result in a candidate detection. Previously in Section 4.3, we introduced the so-called inversion method of generating random numbers, where random numbers can be generated from a distribution by applying the inverse cumulative distribution function to uniformly generated random numbers. Reversing this process, if a random sample X has cumulative distribution function F , then $F(X)$ is a uniformly distributed random variable on the range $[0, 1]$. Consequently, (7.9) has the effect of mapping an amplitude to the range $[0, 1]$ in a manner which corresponds with the false alarm rate for this amplitude. Rather than normalizing the distribution and then calculating the threshold required for a false alarm rate α , we can simply define the detection statistic to be

$$\tilde{X} = 1 - F(A_X|\hat{\theta}, H_0) \quad (7.13)$$

and reject H_0 if $\tilde{X} > \alpha$. This provides a generic method of applying box-car detection for data with an arbitrary distribution, and will be used for Rayleigh, Weibull and K-distributed noise. This method is similar to the CFAR detectors proposed by Jiang [64] and Abaraham [7]. In contrast to the methods proposed by these authors, this extension of the box-car detector does not require the calculation of a threshold value, as the threshold is the desired false alarm rate itself. Additionally, as the higher moments of the noise distribution are represented by the tail of the cumulative distribution function, this methodology includes an accurate representation of the noise model, unlike the CFAR algorithm proposed by López-Estrada [82] which models the noise distribution by its first two moments and a scaling factor. To compare the performance of (7.13) as a box-car detector against the traditional normalization performed by the standard box-car detection algorithm, we will also apply (7.13) for the Gaussian distribution. Therefore, in this thesis, the *standard box-car* refers to (2.17), whilst the *Gaussian box-car* refers to (7.13) with a Gaussian cumulative distribution function.

7.3 Likelihood Ratio Test

In this section we will examine a detection scheme based on *likelihood* (see Section 4.2). This method can provide an improvement over the basic detection scheme discussed in Section 7.2, as it includes a statistical model for the target. As with the box-car detector, detection is achieved by thresholding a test statistic for a chosen probability of false alarm.

In the simple case where all parameters are known, such tests take the form

$$\left[LR(x) = \frac{L(\theta|x, H_1)}{L(\theta_0|x, H_0)} \right] > T, \quad (7.14)$$

where θ denotes the parameters for the target model describing the alternative hypothesis

H_1 , and θ_0 denotes the parameters for the noise-only null hypothesis H_0 . For example, if the noise-only model is that of a Rayleigh distribution, $\theta_0 = \{\sigma\}$, and if the embedded target model is that of a Rician distribution for a Swerling 0 target, then $\theta = \{\nu, \sigma\}$. When the true parameter values are not known, it is advantageous to select a likelihood function $L(\theta|\cdot)$ where the noise-only model can be expressed as a special case of the embedded target model. This nesting of likelihood functions essentially corresponds to a partition of the parameter space Θ , such that Θ_0 denotes the subset of possible parameters for the noise-only scenario. By applying the restriction $\theta_0 \in \Theta_0$ to the denominator of (7.14), we provide the baseline likelihood of the data containing only noise. The numerator, with $\theta \in \Theta$, then covers both noise-only and embedded target scenarios. At this point (7.14) can be rewritten as follows [29]:

$$\left[LR(x) = \frac{\sup\{L(\theta|x) : \theta \in \Theta\}}{\sup\{L(\theta_0|x) : \theta_0 \in \Theta_0\}} \right] > T, \quad (7.15)$$

where $\sup(\cdot)$ is the supremum operator. Since the parameter space of the numerator Θ encompasses that of the denominator Θ_0 , we have $LR(x) \geq 1$. If this partition of the parameter space is not possible, for example if the embedded target and noise-only models are not nested, then it may be the case that $LR(x) < 1$. However, by calculating the distribution of $LR(x)$ for simulated noise-only and simulated embedded target data, we can assess whether a threshold can be used to separate these two classes, and thus to create a viable detector. It is only necessary that $LR(x)$ provide an indication of how far the observed data deviates from the noise-only hypothesis. The threshold value T can then be chosen such that $LR(x) > T$ achieves a false alarm rate of α when the data is known to be noise-only. This results in a *constant false alarm rate* (CFAR) detector.

It should be noted that, depending on the models used, the threshold T may depend on the parameters θ_0 describing H_0 . These values can either be calculated prior to the analysis and stored as a look-up table, or chosen at runtime (as a global threshold) to achieve a desired number of candidate detections over an entire ping.

7.3.1 Generalized Likelihood Ratio Test

The difference between the likelihood function and the probability density function of a distribution depends on the parameters we consider to be fixed, and the parameters we allow to vary. If all function parameters are fixed, the likelihood and probability density function are proportional to one another. As (7.15) maximizes likelihood functions over their respective parameter spaces, Kelly [70] noted that Maximum Likelihood Estimates (MLE) can be directly applied instead (since in such cases they are equal to the supremum), leading to the simple expression:

$$\left[GLR(x) = \frac{p(x|\hat{\theta}, H_1)}{p(x|\hat{\theta}_0, H_0)} \right] > T \quad (7.16)$$

where $\hat{\theta}$ is an estimate of the parameters for the embedded target model, and $\hat{\theta}_0$ is an estimate of the parameters for the noise-only model. This is known as the *Generalized Likelihood Ratio Test* (GLRT). Although this technique technically requires the use of MLE estimators, we consider the use of non-MLE estimators to be an approximation to the true GLR statistic; we will also refer to these as generalized likelihood ratio tests.

7.3.2 Extended Detection

The methods discussed so far apply to single point detection only. In this section, we extend the framework of the GLRT to incorporate detection statistics over multiple samples. We begin with the basic GLR statistic:

$$GLR(x) = \frac{p(x|\hat{\theta}, H_1)}{p(x|\hat{\theta}_0, H_0)}. \quad (7.17)$$

This statistic applies to a single sample point x , taken from the center of the box-car sampling scheme. We will now include further points which may include a target, as

shown by the modified box-car sampling scheme in Figure 7.1.




Figure 7.1: Box-car showing the location of the samples used for the embedded target and noise estimates

The probability of many *independent* events occurring is the product of their probabilities. Making this assumption of independence, and defining an n-tuple of amplitude data $A_X = (x_1, \dots, x_n)$ which covers the region thought to be occupied by the target, (7.17) can be written as the product

$$GLR(A_X) = \frac{\prod_{i=1}^n p(x_i|\hat{\theta}, H_1)}{\prod_{i=1}^n p(x_i|\hat{\theta}_0, H_0)}. \quad (7.18)$$

However, the successive multiplication of small numbers quickly results in very small numbers, which are difficult to represent accurately on a computer due to limited precision. A solution to this difficulty is to apply these calculations in the logarithmic domain, leading to the following equation:

$$\log [GLR(A_X)] = \sum_{i=1}^n \log [p(x_i|\hat{\theta}, H_1)] - \sum_{i=1}^n \log [p(x_i|\hat{\theta}_0, H_0)]. \quad (7.19)$$

It should be noted that the threshold must also be applied in the logarithmic domain, and so the detection test becomes:

$$\log [GLR(A_X)] > \log (T) \quad (7.20)$$

7.4 Comparison of GLR statistics

This section describes the analysis of various embedded target and noise-only models in the context of a GLRT detector. Here we define nine GLR statistics and compare their performance by applying them to embedded target and noise-only data.

As shown in Figure 7.2, the GLR statistic for a data vector $X = (x_1, \dots, x_n)$ can be visualised by plotting the term corresponding to H_0 in (7.19) against the term corresponding to H_1 . We will term these two log-probabilities a *generalized likelihood pair*. Generalized likelihood pairs on the diagonal line correspond to $\log [GLR(A_X)] = 0$, which is the case when the target and noise models have equal probabilities given the observed values A_X . If the target and noise models are nested, this corresponds to absence of a target. In contrast, if the observed data A_X contains a target, its generalized likelihood pair should lie in the upper-left corner of this plot (indicating that $\log [GLR(A_X)] > 0$).




Figure 7.2: Example plot of generalized likelihood pairs, dashed lines show the resulting $\log(\text{GLR})$ values

More importantly, for a GLR statistic to discriminate between embedded target and noise-only data, it is necessary that these two classes lie at different vertical distances from the diagonal line $\left[p(x|\hat{\theta}, H_1) = p(x|\hat{\theta}_0, H_0) \right]$. This vertical distance, correspond-

ing to the resulting $\log(GLR)$ value for a generalized likelihood pair, is shown by the dashed lines in Figure 7.2. If we consider embedded target and noise-only generalized likelihood pairs as two classes of points, occupying two separate clusters, overlap between these clusters is also undesirable as this demonstrates that they can not be easily separated by a threshold, and that applying a threshold will result in false alarms. Ideally, the generalized likelihood pairs for data containing a target would lie in the top-left corner of the graph (which signifies a strong response from the embedded target model and a weak response from the noise-only model), and data containing only noise would result in generalized likelihood pairs either on the diagonal line or the bottom-right corner (which signifies a weak response from the embedded target model and a strong response from the noise-only model).

7.4.1 Experimental Scheme

We will now form our GLR statistics and apply each of them to simulated data. During this process we will also store the individual generalized likelihood pairs so that we can visually evaluate the behaviour of each statistic. From Table 5.2, the three models which best fit our sonar data are the K, Weibull and Rayleigh distributions. These will be applied alongside three embedded target distributions described in Chapter 6, under the assumption of a Swerling 0 target signal. Although the Homodyned and Generalized K distributions offer the most accurate extension of a Swerling 0 target in K-distributed noise, and the Homodyned K distribution is used to simulate the embedded targets, there are considerable implementation issues involved in estimating their parameters, and calculating their probability density functions. As a result of this, we have chosen the Generalized Gamma (GGamma) distribution as an approximation to the Generalized K distribution, the Rician distribution as an accurate model of Swerling 0 target in Rayleigh noise, and the Nakagami distribution as an approximation to the Rician distribution. From these three noise-only models and three embedded target

Table 7.1: Table defining the nine GLR statistics used in this section, including their embedded target and noise-only models

Embedded Target Model	Noise-only Model		
	K	Weibull	Rayleigh
GGamma	GGamma-K	GGamma-Weibull	GGamma-Rayleigh
Rician	Rician-K	Rician-Weibull	Rician-Rayleigh
Nakagami	Nakagami-K	Nakagami-Weibull	Nakagami-Rayleigh

models, nine GLR statistics are defined, as specified in Table 7.1.

Parameter estimates for each of these models are calculated using the methods described in Sections 4.2 and 6.1. Specifically:

- Parameters for the GGamma distribution are estimated iteratively by minimising the chi-square distance between the estimated distribution and the observed data, with Maximum Likelihood Estimation (MLE) of the underlying gamma distribution parameters.
- Parameters for the Nakagami and Rician distributions are estimated by analytic Method of Moments Estimates (MME).
- Parameters for the Rayleigh and Weibull distributions are estimated by MLE.
- Parameters for the K-distribution are again estimated by iterative MME with a $z \log(z)$ seed.

Each embedded target and noise-only model is applied to 1000 simulations of embedded target data and noise-only data, under a selection of scenarios. The equation for each GLR statistic is formed by taking the logarithm of the probability density function of the embedded target model, and subtracting the logarithm of the probability density function of the noise-only model. Each simulation consists of 1000 samples generated for a specific K-distributed noise environment and target SNR, and 1000 further samples

generated from the same noise environment with no target. The embedded target samples are then evaluated in the form specified by the extended detector in equation (7.19), with $\hat{\theta}$ estimated using the embedded target samples and $\hat{\theta}_0$ estimated using the noise-only samples. To form a comparison, the same detector is then used to evaluate the noise-only samples, without an embedded target, with both $\hat{\theta}$ and $\hat{\theta}_0$ estimated using the noise-only data. This allows us to compare the output of each detector when evaluating embedded target data and when evaluating noise-only data, so that we can identify which detectors are most likely to accurately separate these two classes.

The embedded targets are simulated as Swerling 0 targets embedded in K-distributed noise via the Homodyned K model (see Section 6.3). Each detector is tested against embedded targets with SNRs of 1, 2, 3 and 5 dB in K-distributed noise with shape parameter values ν of 0.05, 0.25, 0.5, 1, 5, 10, 15 and 30. Recall that small values of the shape parameter correspond to heavy clutter while a shape parameter of 30 is approximately Rayleigh distributed, which corresponds to a noise-limited environment (see Section 4.1).

Of the nine detectors tested, all were capable of some differentiation between the embedded target and noise-only classes when evaluating high SNR targets in moderately non-Rayleigh clutter (say $\nu \geq 5$). This can be seen in Figure 7.3, which demonstrates the generalized likelihood pairs of the 1000 simulations for all 9 detectors, evaluating a 5 dB embedded target in K-distributed noise with shape parameter $\nu = 5$.

Overall, very little difference was seen between the GLR statistics when processing noise with shape parameters of $\nu \geq 5$. But as SNR decreased, or clutter became more non-Rayleigh ($\nu < 5$), the differences between the nine detectors became evident (see Figure 7.4). When using the GGamma distribution as an embedded target model, there were a large number of vertical outliers, indicating a large variance in the parameter estimates. Here the noise-only and embedded target classes are both very close to the diagonal and thus it is hard to accurately separate them by thresholding these GLR




Figure 7.3: Generalized likelihood pairs for various GLR statistics. Red points were generated from Swerling 0 targets with an SNR of 5dB, embedded in K-distributed noise with shape parameter $\nu = 5$. Blue points were generated from the K-distributed noise only




Figure 7.4: Generalized likelihood pairs for various GLR statistics. Red points were generated from Swerling 0 targets with an SNR of 5dB, embedded in K-distributed noise with shape parameter $\nu = 0.25$. Blue points were generated from the K-distributed noise only

statistics. This is also the case for the Rician-Rayleigh detector, which included a Rician distribution for the embedded target model and a Rayleigh distribution for the noise-only model. In contrast, applying a Rician distribution for the embedded target model and either a Weibull or K-distribution as the noise-only model is shown to provide two clusters which can be easily separated by thresholding the GLR statistic.

As noted in Section 7.3, when the embedded target and noise-only models are not nested, it may not be the case that $LR(x) \geq 1$, but if the distributions of the test statistic for embedded target data and noise-only data are sufficiently distinct, then it is still possible to form a viable detector. This applies to our use of the Rician distribution as an embedded target model, and the Weibull or K-distribution as a noise-only model. As neither of these detectors are nested, it is not true that $GLR(A_X) \geq 1$, demonstrated

by the noise-only cluster falling below the diagonal line. Interestingly, in this case data from the embedded target class results in values close to $\log[GLR(A_X)] = 0$ (equivalent to $GLR(A_X) = 1$), while data from the noise-only class results in values where $\log[GLR(A_X)] < 0$ (corresponding to $GLR(A_X) < 1$). In both cases, the vertical distance from the diagonal is sufficiently large that the variance of the clusters would not lead to any overlap in the $\log[GLR(A_X)]$ statistic, and so these two detectors appear to offer the best performance for this particular noise environment.

Although we have only provided plots here for two scenarios, a full summary of the nine GLR statistics is provided in the two tables in Figure 7.5. These tables include all target SNRs and noise environments, and list the average distance between the $\log[GLR(A_X)]$ statistics for the embedded target and noise-only classes, and their individual variances. For each scenario, the three test statistics with the greatest inter-class distance are highlighted in green, yellow, and orange, in descending order. Of the test statistics described here, the two which provided the greatest discrimination between embedded target and noise-only scenarios were found to be the Rician-K and Rician-Weibull detectors. Despite the increased complexity offered by the GGamma model, this did not lead to an improvement in performance. This is partially explained by high variance in its parameter estimates, leading to a poor fit to the embedded target data. Furthermore, the estimation procedure for the GGamma distribution repeatedly failed to obtain valid estimates for very small shape parameters, resulting in the “NaN” values in Figure 7.5.

Target Model	T-N	var(T)	var(N)	T-N	var(T)	var(N)	T-N	var(T)	var(N)	T-N	var(T)	var(N)
SNR 1												
	0.05			0.25			0.5			1		
GGamma v K	NaN	NaN	0.1029	0.2314	0.0149	0.0348	0.0529	0.0047	0.0439	0.0041	0.0039	0.0262
GGamma v Weibull	NaN	NaN	0.1047	0.1851	0.0094	0.0471	0.0766	0.0045	0.0541	0.0186	0.0052	0.0132
GGamma v Rayleigh	NaN	NaN	0.1380	-1.1678	0.1063	0.0648	-0.3945	0.0504	0.0462	-0.1412	0.0258	0.0205
Rician v K	13.7818	0.6151	0.1414	1.5193	0.1191	0.0578	0.5546	0.0662	0.0315	0.2029	0.0367	0.0166
Rician v Weibull	12.9617	0.6235	0.1974	1.4645	0.1246	0.0680	0.5765	0.0666	0.0331	0.2191	0.0346	0.0140
Rician v Rayleigh	0.0454	0.0403	0.0460	0.1203	0.0238	0.0236	0.1029	0.0176	0.0130	0.0556	0.0108	0.0056
Nakagami v K	6.3115	0.2903	0.0363	0.3434	0.0409	0.0093	-0.0260	0.0157	0.0060	-0.0659	0.0080	0.0193
Nakagami v Weibull	5.5799	0.2856	0.0676	0.3014	0.0418	0.0134	0.0002	0.0159	0.0047	-0.0518	0.0052	0.0164
Nakagami v Rayleigh	-7.3885	0.3160	0.0870	-1.0544	0.0686	0.0325	-0.4743	0.0373	0.0220	-0.2169	0.0242	0.0158
SNR 2												
	0.05			0.25			0.5			1		
GGamma v K	NaN	NaN	0.0940	0.1888	0.0147	0.0462	0.0485	0.0045	0.0326	0.0251	0.0039	0.0094
GGamma v Weibull	NaN	NaN	0.0446	0.2021	0.0099	0.0472	0.0766	0.0045	0.0357	0.0234	0.0051	0.0098
GGamma v Rayleigh	NaN	NaN	0.1432	-1.4026	0.1115	0.0476	-0.4893	0.0519	0.0337	-0.1609	0.0260	0.0125
Rician v K	14.4663	0.6389	0.1208	1.8166	0.1251	0.0376	0.6971	0.0658	0.0165	0.2542	0.0361	0.0032
Rician v Weibull	13.9275	0.6078	0.1842	1.8355	0.1248	0.0442	0.7259	0.0670	0.0100	0.2492	0.0348	0.0024
Rician v Rayleigh	0.1111	0.0403	0.0489	0.2249	0.0232	0.0190	0.1570	0.0172	0.0049	0.0667	0.0112	0.0021
Nakagami v K	6.4379	0.2797	0.0271	0.2804	0.0405	0.0104	-0.0858	0.0159	0.0202	-0.0145	0.0081	0.0297
Nakagami v Weibull	5.8521	0.2948	0.0558	0.2983	0.0422	0.0054	-0.0575	0.0155	0.0171	-0.0172	0.0052	0.0284
Nakagami v Rayleigh	-7.9377	0.3238	0.0841	-1.3073	0.0672	0.0279	-0.6267	0.0380	0.0174	-0.2002	0.0232	0.0261
SNR 5												
	0.05			0.25			0.5			1		
GGamma v K	NaN	NaN	0.0742	0.2932	0.0155	0.0437	0.2742	0.0046	0.0294	0.3157	0.0038	0.0258
GGamma v Weibull	NaN	NaN	0.0499	0.2806	0.0094	0.0444	0.2706	0.0047	0.0277	0.3063	0.0049	0.0255
GGamma v Rayleigh	NaN	NaN	0.0663	-1.4149	0.1036	0.0533	-0.3133	0.0526	0.0280	0.1159	0.0261	0.0256
Rician v K	15.2991	0.6141	0.0593	2.1219	0.1304	0.0392	0.9841	0.0674	0.0296	0.5738	0.0371	0.0258
Rician v Weibull	15.1814	0.6283	0.0929	2.0998	0.1316	0.0367	0.9803	0.0673	0.0314	0.5658	0.0350	0.0261
Rician v Rayleigh	0.4081	0.0415	0.0425	0.4120	0.0230	0.0360	0.3978	0.0176	0.0305	0.3733	0.0112	0.0253
Nakagami v K	6.1822	0.2876	0.0367	0.4572	0.0409	0.0789	0.3027	0.0155	0.0412	0.3392	0.0083	0.0267
Nakagami v Weibull	6.0959	0.2824	0.0102	0.4428	0.0410	0.0865	0.2997	0.0160	0.0378	0.3289	0.0053	0.0265
Nakagami v Rayleigh	-8.7019	0.3021	0.0508	-1.2534	0.0667	0.0793	-0.2781	0.0374	0.0374	0.1408	0.0240	0.0254

Target Model	T-N	var(T)	var(N)	T-N	var(T)	var(N)	T-N	var(T)	var(N)	T-N	var(T)	var(N)
SNR 1												
	5			10			15			30		
GGamma v K	0.0019	0.0028	0.0027	0.0027	0.0024	0.0031	0.0035	0.0025	0.0033	0.0043	0.0024	0.0036
GGamma v Weibull	-0.0006	0.0037	0.0030	0.0009	0.0029	0.0027	0.0016	0.0027	0.0028	0.0026	0.0023	0.0031
GGamma v Rayleigh	-0.0117	0.0066	0.0028	-0.0023	0.0039	0.0026	-0.0001	0.0032	0.0029	0.0019	0.0027	0.0030
Rician v K	0.0180	0.0083	0.0019	0.0077	0.0041	0.0022	0.0059	0.0029	0.0024	0.0051	0.0019	0.0027
Rician v Weibull	0.0150	0.0071	0.0009	0.0056	0.0037	0.0016	0.0040	0.0027	0.0018	0.0034	0.0015	0.0022
Rician v Rayleigh	0.0047	0.0022	0.0007	0.0022	0.0011	0.0014	0.0021	0.0008	0.0016	0.0026	0.0005	0.0022
Nakagami v K	0.0148	0.0124	0.0065	0.0077	0.0068	0.0028	0.0060	0.0049	0.0025	0.0040	0.0031	0.0027
Nakagami v Weibull	0.0120	0.0103	0.0046	0.0057	0.0052	0.0022	0.0039	0.0039	0.0023	0.0021	0.0024	0.0026
Nakagami v Rayleigh	0.0013	0.0082	0.0043	0.0024	0.0049	0.0019	0.0018	0.0035	0.0019	0.0012	0.0022	0.0026
SNR 2												
	5			10			15			30		
GGamma v K	0.0259	0.0026	0.0073	0.0289	0.0024	0.0074	0.0306	0.0026	0.0080	-0.0027	1.0873	0.0075
GGamma v Weibull	0.0201	0.0038	0.0072	0.0244	0.0030	0.0072	0.0266	0.0025	0.0075	0.0284	0.0024	0.0074
GGamma v Rayleigh	0.0080	0.0065	0.0066	0.0197	0.0038	0.0069	0.0234	0.0032	0.0071	0.0253	0.0026	0.0071
Rician v K	0.0420	0.0084	0.0076	0.0338	0.0041	0.0073	0.0331	0.0030	0.0072	0.0327	0.0019	0.0071
Rician v Weibull	0.0358	0.0072	0.0070	0.0295	0.0037	0.0070	0.0294	0.0026	0.0075	0.0290	0.0016	0.0071
Rician v Rayleigh	0.0234	0.0022	0.0068	0.0250	0.0011	0.0070	0.0254	0.0008	0.0069	0.0263	0.0005	0.0066
Nakagami v K	0.0414	0.0127	0.0072	0.0306	0.0068	0.0074	0.0261	0.0048	0.0075	0.0236	0.0029	0.0081
Nakagami v Weibull	0.0359	0.0109	0.0070	0.0260	0.0053	0.0072	0.0234	0.0046	0.0074	0.0198	0.0022	0.0083
Nakagami v Rayleigh	0.0235	0.0099	0.0065	0.0210	0.0047	0.0072	0.0194	0.0038	0.0075	0.0179	0.0022	0.0081
SNR 5												
	5			10			15			30		
GGamma v K	0.3743	0.0027	0.0206	0.3821	0.0028	0.0204	0.3858	0.0026	0.0197	0.3884	0.0024	0.0191
GGamma v Weibull	0.3637	0.0036	0.0213	0.3759	0.0029	0.0206	0.3799	0.0026	0.0208	0.3818	0.0024	0.0196
GGamma v Rayleigh	0.3454	0.0065	0.0208	0.3640	0.0039	0.0200	0.3705	0.0032	0.0196	0.3756	0.0027	0.0196
Rician v K	0.3947	0.0084	0.0209	0.3891	0.0043	0.0195	0.3884	0.0029	0.0192	0.3907	0.0020	0.0188
Rician v Weibull	0.3853	0.0067	0.0208	0.3837	0.0040	0.0198	0.3832	0.0025	0.0198	0.3842	0.0016	0.0190
Rician v Rayleigh	0.3670	0.0021	0.0200	0.3716	0.0012	0.0192	0.3738	0.0008	0.0201	0.3761	0.0005	0.0197
Nakagami v K	0.3828	0.0127	0.0219	0.3810	0.0069	0.0214	0.3811	0.0053	0.0216	0.3811	0.0030	0.0206
Nakagami v Weibull	0.3749	0.0105	0.0225	0.3721	0.0051	0.0219	0.3750	0.0036	0.0209	0.3745	0.0025	0.0207
Nakagami v Rayleigh	0.3578	0.0096	0.0220	0.3636	0.0048	0.0216	0.3660	0.0035	0.0218	0.3681	0.0024	0.0214

Figure 7.5: Comparison of detectors for target (T) and noise (N) data, showing the average difference and variance in the GLR statistic

7.5 Theoretical ROC curves

In this section, we will analyse the theoretical performance of GLRT and box-car detectors. At this stage we will not be simulating a full sonar environment as described in Section 7.1. Rather, we will apply each detector in various homogeneous environments with constant target signal.

Of the GLR statistics tested in Section 7.4, we have identified the Rician-K and Rician-Weibull statistics as providing the greatest discrimination between embedded target and noise-only samples. These two statistics will be used in this section. In order to validate whether these statistics provide greater detection performance as a result of their ability to adapt to local clutter statistics, the Rician-Rayleigh detector will also be included. These three GLRT detectors will be applied alongside the box-car methods introduced in Section 7.2, including the standard box-car detector popularised by early sonar theory.

Both the box-car and GLRT detectors process data sampled in a box-car configuration. However, as the simulated environments will be homogeneous, rather than carefully simulating the environment and sampling representative points, we can directly provide the relevant samples from the random number generator for the HK-distribution (see Section 6.3).

Each simulated detection will be made with the equivalent of a 100 metre target (corresponding to 822 samples) and 100 metres of noise data (corresponding to a further 822 noise samples). Noise samples will be generated from the K-distribution with the shape parameter ranging from heavy clutter ($\nu = 0.05$) to approximately Rayleigh noise-limited environments ($\nu = 30$). Targets will be generated according to a Swerling 0 target in the Homodyned-K embedding model, with target amplitudes chosen to provide SNRs between 1dB and 8dBs. This provides the basis for establishing the

probability of detection, while the probability of false alarm is established by applying the detectors to noise-only samples.

The samples generated for each simulated target are used to calculate a Method of Moments estimate (MME) for the parameter of the Rician distribution. The noise samples are used to estimate parameters for the three respective noise models: the K, Weibull and Rayleigh distributions. As in Section 7.4, parameter estimates for the Weibull and Rayleigh distributions were obtained by MLE estimation, while the K distribution parameters are once again estimated by iterative MME with a $z \log(z)$ seed.

For each box-car detector, the central sample of the target data is chosen to be the cell under test, denoted by x , and evaluated using $1 - F(x|\hat{\theta})$, where $F(x|\hat{\theta})$ denotes the cumulative distribution function of the noise model with estimated parameters $\hat{\theta}$. This forms the test statistic to which thresholds are applied.

For each single-point GLRT detector, the central sample of the target data is taken as the cell under test and its probability is evaluated by the Rician PDF. We then calculate the probability of observing the same sample for each of the noise distributions. The test statistics are then formed by taking the logarithm of these probabilities and subtracting the noise log-probability from the target log-probability.

For the extended GLRT, rather than considering the central sample and calculating a single probability per model, we instead calculate the sum of the log-probabilities for all 822 samples, as demonstrated earlier in (7.19). The summed log-probabilities of the noise model are then subtracted from the summed log-probabilities of the target model to form our test statistic.

To generate ROC curves for low false alarm rates, each target SNR and noise environment was simulated 10,000 times, corresponding to a minimum theoretical false alarm rate of 10^{-5} . Although thresholds are generally selected based on the local

noise statistics, since the noise statistics remain constant across the ROC curves, global thresholds were applied at intervals between the minimum and maximum GLR values obtained.

7.5.1 Analysis of Results

The complete set of 64 ROC curves produced by this experiment is included in Appendix A. The ROC curves for the box-car detectors and single point GLRT detectors can be found in Figures A.1-A.8, while the ROC curves for the extended GLRT detectors are provided in Figures A.9-A.16. We now provide an analysis of these results, and reproduce a subset of these ROC curves to illustrate our discussion.

Since each ROC curve is generated against a homogeneous environment, there is no variation in the noise statistics. Therefore, it should not be surprising that the normalisation applied in the K, Weibull, Gaussian and Rayleigh box-car detectors achieved equivalent probabilities of detection in each noise environment. This is demonstrated in Figure 7.6, where it is difficult to distinguish between these four detectors as they all yield the same curve. In this case, normalizing by each of these noise distributions simply resulted in a different non-linear scaling of the cell under test, with no impact on the resulting curve. However, changes in performance can still be seen by comparing these normalization schemes between noise environments and target SNRs. Overall, box-car detectors are most effective when processing loud targets (with high SNR) in low clutter environments ($\nu > 30$). As target SNR decreases, or noise increases (as seen between Figures 7.6a and 7.6b), the probability of detection also decreases.

Throughout the simulations, the single-point and extended GLRT detectors demonstrated substantially higher probabilities of detection at low false alarm rates ($\alpha \leq 10^{-2}$).

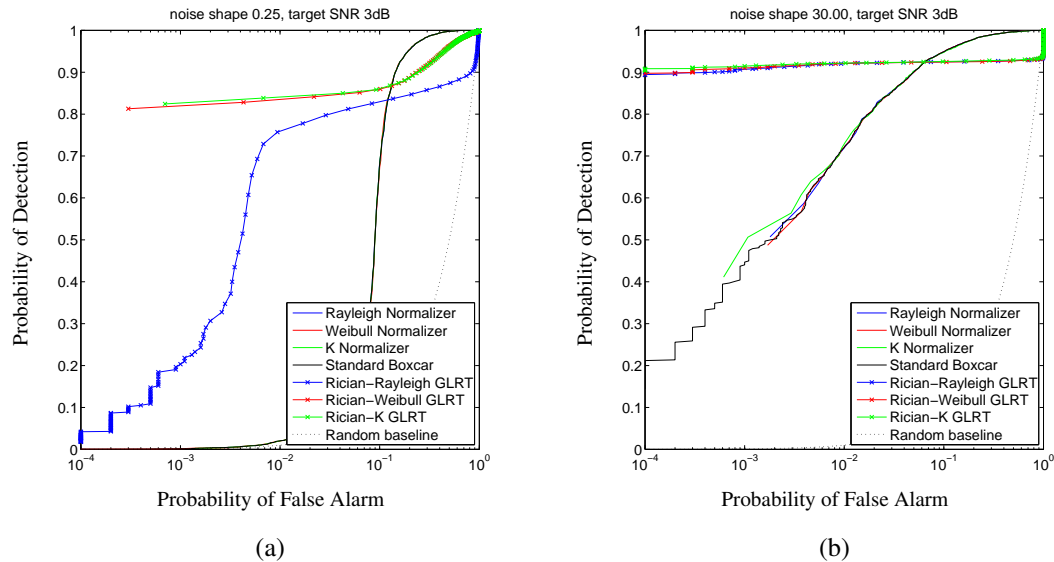


Figure 7.6: Simulated ROC curves for 3dB targets, comparing normalization and single-point GLRT detectors

The GLRT detectors all exhibited a drop in detection at high false alarm rates, leading to the characteristic curve seen in the top-right corner of Figure 7.6b. Although high false alarm rates are not particularly interesting, this suggests that our GLRT detectors are only suitable below a certain level of α , which depends on both the noise environment and target SNR. Yet there is also a plateau in the ROC curve where the false alarm rate can be decreased by orders of magnitude with little impact to detection performance. This property is highly desirable in a detection statistic, and demonstrates the potential benefits of advanced methods of detection. This property is not observed for any of the box-car based detectors.

Comparing the three single-point GLRT detectors, detection performance does depend on the composition of the GLR statistic. Throughout these ROC curves, the Rician-Rayleigh GLRT detector has performed the worst of three. This is to be expected, as it was established in Section 5.3 that the Rayleigh distribution applied in the denominator of this GLR statistic is a poor model of clutter ($\nu < 30$). As clutter decreases (ν

increases), the performance of the Rician-Rayleigh detector rose to that of the Rician-K and Rician-Weibull detectors, but this will be of little use in non-homogeneous sonar noise recorded in the real world, such as our own trial data. In conclusion, we were justified in selecting the Rician-K and Rician-Weibull GLR statistics for our detectors, and their ability to adapt to the local noise environment leads to increased probability of detection.

There was very little difference in performance between the Rician-K and Rician-Weibull GLRT detectors. Unlike the box-car detectors, where embedded target data is assumed to have a higher amplitude than noise data, the GLRT detectors make no such assumption. While the test statistic for box-car detectors are sensitive to changes in the tail of the noise-only distribution, the GLR statistic is sensitive to the tail of the distribution formed by the ratio of the embedded target and noise-only distributions. Although the upper-tail Anderson-Darling test found the K-distribution to be a more accurate model of our trial data in the tail, this does not provide a significant advantage for the GLRT test, where the K and Weibull distributions are equally suitable models of noise-only data.

Applying the GLRT detectors in their extended form provided an even greater performance increase, as shown in Figure 7.7. Here the Rician-Weibull and Rician-K detectors were both capable of almost perfectly discriminating between target and noise data, achieving 100% detection at a false alarm rate of 10^{-4} . The Rician-Rayleigh detector was also able to achieve this level of performance, but still had difficulties in highly cluttered environments (see Figure 7.7a).

In conclusion, the use of GLRT detectors should provide a significant performance advantage over box-car detectors. Although it is anticipated that the performance for real data will be lower than the ROC curves simulated here (due to statistical fluctuations which have not been accounted for), this applies to all of the detectors analysed. As the Rician-K and Rician-Weibull GLRTs performed very similarly, it is believed that




Figure 7.7: Simulated ROC curves for 3dB targets, comparing normalization and extended GLRT detectors

the Rician-Weibull GLRT is preferable, given the ease with which the parameters of the Weibull distribution can be estimated by maximum-likelihood methods. Lastly, the increased sample size applied in the extended GLRT detector was shown to provide a substantial increase in detection performance; this bodes well for the application of the GLRT detector to multi-ping processing (see Chapter 8).

7.6 Application to Real Sonar Noise

In this section we apply the box-car, Rician-Weibull GLRT and Rician-Rayleigh GLRT detectors to real sonar noise with simulated targets. An even grid of GPS coordinates were generated over the region covered by our trial data, with simulations occurring at each point. This process attempts to avoid any bias towards simulating targets at short range from the sonar system, where there are a large number of samples covering a small area. This technique is described in detail in Section 7.1.

Each evaluation of a detectors test statistic used 50 meters of noise samples either side of a 100 metre long, Swerling 0, target. These targets were generated by embedding a constant signal in the raw, complex, data provided by the matched filter, before calculating the amplitude. The target signal was chosen in the manner described in Section 7.1 for SNRs of 1dB, 3dB, 5dB, and 8dB. This allowed us to test each of the detectors against an accurate portrayal of a sonar environment in various frequency bands, rather than relying on homogeneous noise. The full results of this experiment are provided in Appendix B, Figures B.1-B.9.

7.6.1 Box-car Detectors

As previously discussed, applying a noise distribution in a box-car detector is expected to allow us to set a threshold which is adaptive to the noise environment. It was therefore unsurprising that no advantage was seen in the previous simulations when the statistics of the noise environment were constant. Yet in this case, with a wide variety of noise environments ranging from reverberation to noise-limited, we see the advantage of accurately modelling sonar noise. As shown in Figure 7.8, the superior fit of the K-distribution to the right-hand tail of the distribution results in an increase in probability of detection at low false alarm rates. The extent of this improvement depends on target SNR, with smaller differences seen for weak targets. For 5 dB targets in the low frequency band the K-distribution box-car detector is capable of achieving a 50% probability of detection at a false alarm rate of $10^{-4.3}$, which compared to the standard box-car detectors $10^{-3.5}$ false alarm rate, equates to a five-fold reduction in false alarms.

If we compare the results for the other noise distributions, it is clear that normalizing by the Gaussian, Rayleigh or Weibull distributions provides no advantage over the standard box-car detector. In fact, normalizing the Gaussian, Rayleigh or Weibull distributions

by histogram equalisation resulted in detection statistics with lower numerical precision, and so it was not possible to select thresholds for false alarm rates as low as those achieved by the standard box-car detector. To demonstrate this, in Figure 7.8, the Gaussian, Rayleigh and Weibull box-car detectors were incapable of achieving a false alarm rate of 10^{-4} , while the standard and K-distribution box-car detectors both achieved false alarm rates almost an order of magnitude lower. This is due to the fact that the standard box-car detector is not constrained to the range $[0, 1]$, and so it is less likely that a large number of values will equal the maximum value output by this test statistic. In contrast, our other box-car detectors are constrained to this range, but the good fit obtained by the K-distribution in the tail of the data reduces the impact of this limitation.

Considering that the Gaussian distribution is an entirely unsuitable model for the distribution of sonar noise in our trial data, the Rayleigh and Weibull distributions appear to provide nothing more than a slightly different non-linear scaling of the detection statistic, as observed in the previous section for homogeneous noise. It is interesting to note that none of our box-car detectors were capable of achieving detections at 10^{-6} . Although this may be caused in part by numerical precision, there is also an inherent overlap between the distribution of the detection statistic for the target and noise classes. Consequently for target data to be detected at moderate false alarm rates, the target data must have a greater amplitude than the majority of noise samples. This is an intrinsic difficulty with box-car detectors, but as we shall now demonstrate, this can be overcome by GLRT detectors.

7.6.2 GLRT Detectors

As demonstrated in Section 7.5, the form of GLRT detectors applied in this thesis have characteristic ROC curves, achieving a lower probability of detection than box-car detectors at high false alarm rates. This was verified in our trial data as shown in Figure




Figure 7.8: ROC curves for box-car detectors processing LF trial data against Swerling 0 simulated targets

7.9a. Unlike our previous simulations (see Figure 7.6), the probability of detection was not maintained as consistently for decreasing false alarm rates. Despite this, both single and multi-sample GLRT detectors achieved greater probability of detection than the standard box-car detector across all target SNRs and frequency bands. The single-point GLRT's inability to match simulated performance is partly due to the lack of adaptive thresholding based on the noise environment. While the box-car detectors demonstrate that environmentally adaptive thresholding can boost detection performance at low false alarm rates, these findings demonstrate the merits of the GLR statistic on its own. If we compare these results with the K-distribution box-car detector, the single sample GLRT detector only achieved a lower false alarm rate for a 50% probability of detection against targets with SNRs less than 5 dB. However, the multi-sample GLRT achieved consistently better performance, as seen by the high probability of detection in Figure 7.9b. For example, when comparing against the box-car detector in Figure 7.9 for a 50% probability of detection, the single sample Weibull GLRT achieves a two-fold reduction in false alarms, while the multi-sample GLRT achieves a reduction of over two orders

of magnitude.

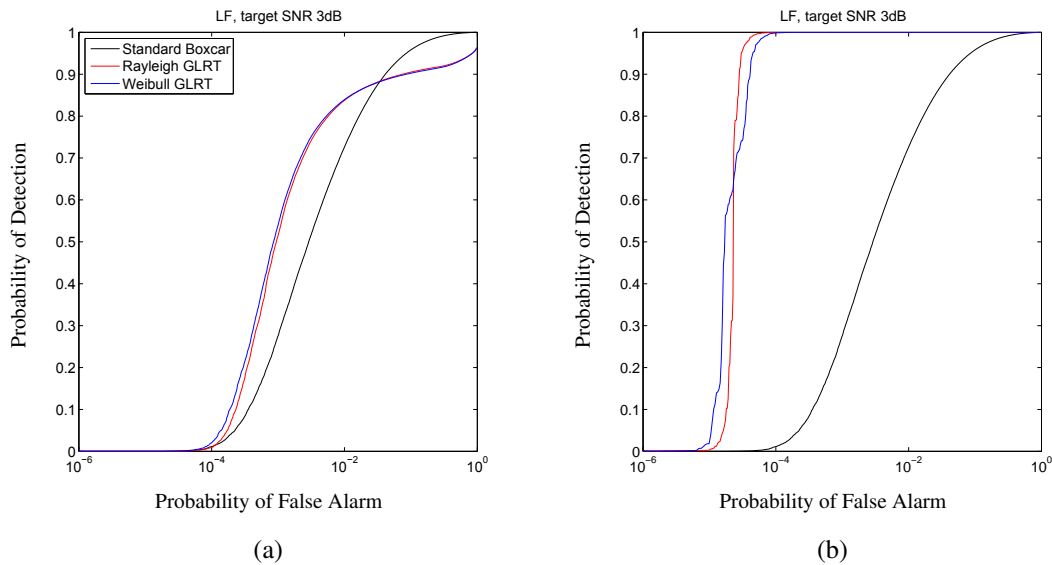


Figure 7.9: ROC curves for a 3dB target in the LF band against Swerling 0 simulated targets. Processed by the (a) single-sample, and (b) multi-sample GLRT detectors

In this section we have demonstrated that environmentally adaptive thresholding provides an increased probability of detection at low false alarm rates, and that this is most significant for strong targets. Meanwhile, the GLRT approach seems better suited to weak targets, and is easily extended to incorporate additional samples of target data for an even higher probability of detection. In the following chapter these detection techniques will be applied over multiple pings.

7.7 Spatial Variation of Performance

As our bias-reduction technique has made use of a non-uniform distribution of targets in range, these simulations are not suitable for the analysis of sonar performance as a function of range and bearing. However, we can cross-reference the noise environments obtained in this trial data with the theoretical performance calculated in Section 7.5.

The median K-distribution shape parameter obtained in the trial data for each range cell is provided in Figure 7.10. Although this figure does not account for the geographic variation of noise sources in the environment, it does summarise the extent to which we could expect to observe non-Rayleigh environments in range and bearing. For example, as anticipated, the Low frequency band appears to be the most heavily affected by clutter, whilst the medium and high frequency bands contain non-Rayleigh noise at ranges which are likely to correspond with reverberation. Additionally, all three frequency bands contain heavily non-Rayleigh noise within the first 500m consistent with reverberation. A more detailed analysis of noise sources in this trial data can be found in Section 8.3.

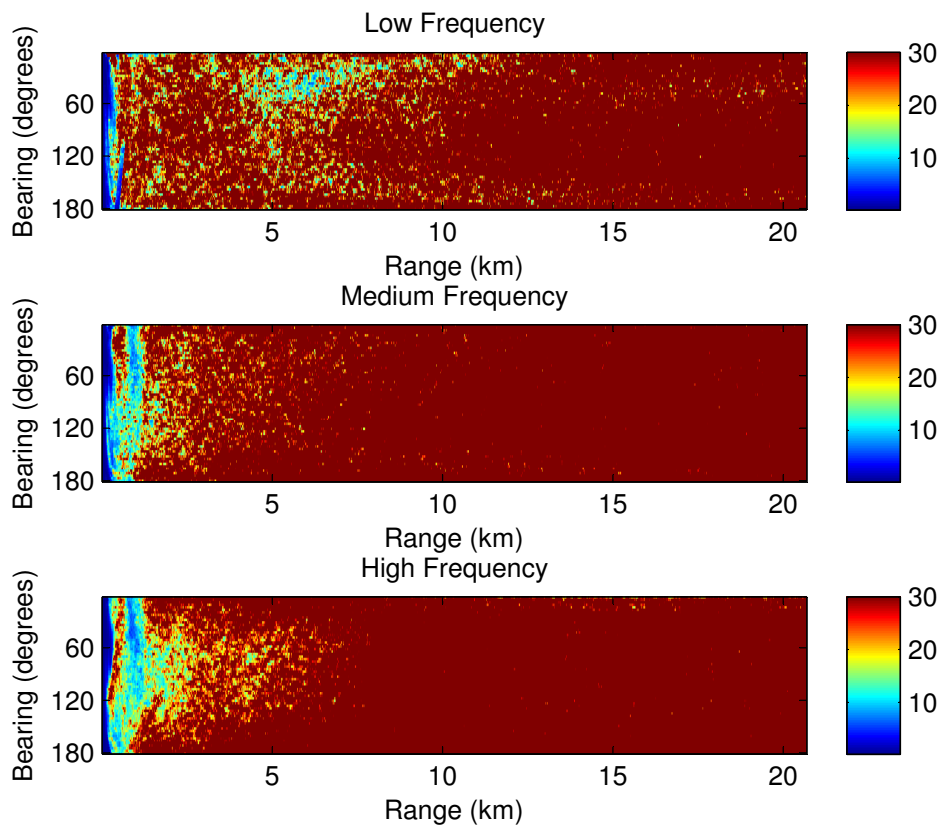


Figure 7.10: Median K-distribution shape parameter (calculated from 39 pings) for each range cell in the Low, Medium, and High frequency bands

These shape parameters have been cross-referenced with the theoretical performance calculations from Section 7.5. The theoretical probability of false alarm in the low frequency band, for a 3dB target at a 50% probability of detection, is provided in Figure 7.11. Equivalent figures for the medium and high frequency bands are available in Appendix C, Figures C.2 and C.3.

As previously demonstrated, the GLRT detectors can theoretically achieve better performance than the box-car detectors. This figure supports the hypothesis that the thresholds applied to the output of the GLRT statistics, which were not previously chosen as a function of ν , may have reduced the effectiveness of this approach to target detection. Consequently, it is likely that if thresholds were selected as a function of ν , that the Weibull GLRT would have achieved a more substantial performance increase over the Rayleigh GLRT in Section 7.6, as the Weibull GLRT is capable of lower probability of false alarm in non-Rayleigh environments.

Although multiple normalization schemes have been applied in this analysis, Figure 7.11 contains only one plot for the box-car detectors. The reason for this is simple, the only difference between these normalisation schemes is how accurately a false alarm rate can be achieved by applying the theoretical threshold for a particular distribution. The consequence of theoretically maintaining a probability of detection of 50%, is that the same probability of false alarm must also be maintained by all box-car detectors, regardless of the prescribed noise distribution. As this is an unlikely scenario, it is useful to consider the accuracy of threshold calculation for a selection of these distributions. Let us make the the assumption that the trial data is perfectly modelled by the K-distribution, with the shape parameters provided in Figure 7.10, and that we wish to apply thresholds for a desired probability of false alarm of 10^{-3} . The theoretical false alarm rates which would be achieved by estimating noise in the Weibull, Rayleigh, and Gaussian box-car detectors are demonstrated in Figure 7.12. Equivalent figures for the medium and high frequency bands are available in Appendix C, Figures C.5 and C.6.

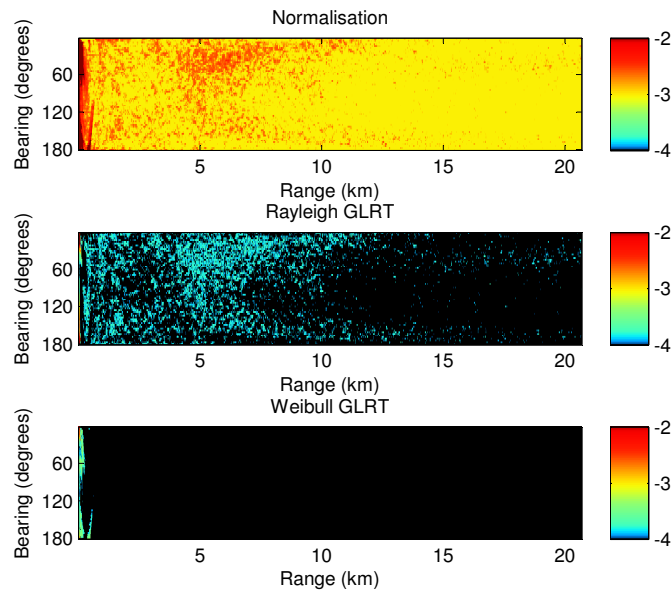


Figure 7.11: Theoretical probability of false alarm (\log_{10}) for normalisation, Rayleigh GLRT, and Weibull GLRT detectors in the Low Frequency band, against a 3dB target at 50% probability of detection

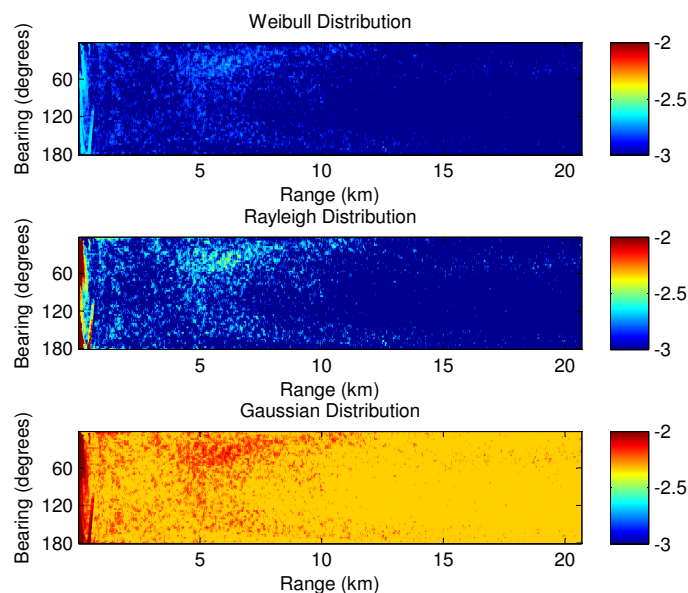


Figure 7.12: Theoretical probability of false alarm (\log_{10}) for Weibull, Rayleigh, and Gaussian normalisation in the Low Frequency band, when selecting a probability of false alarm of 10^{-3} in K-distributed noise

Considering that the Weibull distribution contains the Rayleigh distribution as a special case, it is expected that these two models provide equivalent performance in Rayleigh environments. Additionally, as anticipated, the Weibull distribution also achieves a more accurate probability of false alarm in non-Rayleigh environments.

For convenience the probability of false alarm from each of these box-car detectors is displayed in an alternative form in Figure 7.13. This figure summarises the average probability of false alarm across the 64 beams of trial data, calculated by arithmetic mean. None of these distributions are able to consistently achieve the intended 10^{-3} probability of false alarm. Furthermore, all three distributions have greater probability of false alarm in regions of clutter. Interestingly, the Weibull distribution achieves a probability of false alarm below 10^{-3} at very close range. Recalling that the objective is to accurately achieve the desired false alarm rate, this is indicative of the existence of a shape parameter for which the K and Weibull distributions intersect at the threshold for a 10^{-3} false alarm rate, and should not be considered evidence of the Weibull distributions suitability in such environments. Equivalent figures for the medium and high frequency bands are available in Appendix C, Figures C.8 and C.9.

In this section we have demonstrated that the GLRT detectors would benefit from thresholds selected as a function of ν . We have also demonstrated the relative performance of various box-car detectors, compared to a baseline of K-distributed data. As anticipated, performance tends to be greatest at long range where the noise environment is mostly Rayleigh. However, in regions of clutter and reverberation where non-Rayleigh statistics are present, the use of more complex distributions has a positive impact on the system.




Figure 7.13: Theoretical probability of false alarm (\log_{10}) averaged across all beams, for Weibull, Rayleigh, and Gaussian normalisation in the Low Frequency band, when selecting a probability of false alarm of 10^{-3} in K-distributed noise

7.8 Application to Real Targets

We will now perform a brief performance analysis against real targets in real sonar noise. This dataset is different to the trial data used elsewhere in this thesis, as it included a participating moving target at known locations within each ping. Due to its sensitive nature this data was only available at the offices of Thales Underwater Systems Ltd. This limited access meant that it was not possible to generate the ROC curves in this section to the same resolution as our previous analysis. As with our trial data, this dataset was obtained from a long-range sonar system in a shallow water environment.

As this dataset contained a real target, the experimental method had to be altered to account for this. Rather than selecting points on an even grid in GPS coordinates, we manually select GPS positions which include the target in one or more pings. By applying our detectors to these points we obtain an estimate for the probability of

detection. This is an entirely valid approach, and does not artificially increase the probability of detection as the probability of detection, by definition, must only be calculated from detection tests performed in the presence of an actual target (where H_1 is true). We then select GPS positions on an even grid, but throw out any GPS positions which are within 200 meters of the target in one or more pings (twice the assumed target length). The remaining GPS positions are assumed to contain only noise, and so applying our detectors to these points we obtain an estimate for the probability of false alarm.

During this performance analysis we processed 78 pings of data, with 10,000 detection tests executed per ping. This equates to a theoretical false alarm rate close to 10^{-6} .

Initial tests found that samples from real targets were poorly fit by the models suggested in Chapter 6. This was caused by our assumption that the Swerling target models would provide an appropriate model for the sonar target signal. When embedding a Swerling target in K-distributed noise, all of the Swerling models decrease the skew of the distribution. As the strength of the target signal increases, these distributions either become increasingly Gaussian, Rayleigh, or exhibit a high degree of kurtosis. By contrast, the distribution of the actual sonar target in Rayleigh noise was seen to be similar to a log-normal distribution, with positive skew and a heavier tail.

In an attempt to apply GLRT detectors to this data, we defined the following GLR statistics:

1. Weibull-K: This detector models sonar noise by the K-distribution, to allow for multi-ping processing in the parameter estimate (see Chapter 8). The Weibull distribution is selected as the target model. Although the Weibull distribution is usually applied as a model for sonar noise, the shape of the distribution covers a different range to that of the K-distribution. Furthermore, disparities between

the parameter estimates of the K and Weibull distributions may be useful when estimates of the Weibull distribution parameters are deliberately corrupted by target signal.

2. Lognormal-K: This detector also models sonar noise by the K-distribution. The log-normal distribution is selected as the target model. Although the log-normal distribution was not included in our previous discussion of target models, it has been applied as a model for targets in radar signal processing, with some physical justification [104].

A further property which was not considered in Chapter 6 was the correlation of the target signal. Due to the physical construction of maritime vehicles, some surfaces are likely to produce specular reflections² when acoustic energy is reflected. As the acoustic energy is more focused, less spherical spreading takes place and so the resulting signal amplitude is far larger than would otherwise be the case. Regions of data between specular reflections from the same target, may not include any target signal at all. This may be due to destructive interference, acoustic cladding, or variability in the speed of sound through the water. Consequently, not all samples tested in a region occupied by a target should be expected to produce detections. In order for the GLRT to be successful it need only detect the target at a single point, and so the maximum GLR statistic obtained in the target area was taken for each detection test. This procedure is similar to the maximum detector specified by Abraham [6], which represents a GLRT over the unknown location of the target within this region.


The resulting ROC curves from this analysis are provided in Figures 7.14 and 7.15. In Figure 7.14 we see the performance of our detectors in the low and medium frequency bands, while Figure 7.15 contains ROC curves for the full bandwidth of the sonar

²Specular reflections are characterised by the direct reflection of the incident wave in a single direction, such as light reflecting from a mirror, whereas diffuse reflections are characterised by the incident wave being reflected in a large number of directions, such as light reflecting from a matte surface. This concept applies to acoustics in exactly the same manner.


system. As can be seen throughout these figures, the GLRT and K-distribution box-car detectors fail to achieve greater performance than the standard box-car algorithm.

In conclusion, our attempts to fit alternative models to the sonar target did not yield positive results. Despite our attempts to apply the Weibull or log-normal distributions, in place of the Rician distribution, no benefit was seen in the ROC curves. In fact, there was very little difference in performance between the Weibull-K and lognormal-K detectors.

In this chapter we have defined and evaluate detection schemes based on the standard box-car and GLRT detectors. The standard box-car detector was generalized to an arbitrary noise distribution, and the GLRT detection scheme has been shown to be easily extended to process multiple samples in the region under test. Synthetic data has been used to validate the effectiveness of various GLRT detectors, and it was found that GLRT detectors incorporating a Rician model were the most capable of differentiating between a Swerling 0 target in K-distributed noise, and noise-only environments. Box-car and GLRT detectors were applied to synthetic targets embedded in our noise-only trial data, which verified that GLRT detectors are capable of improved detection performance when compared to box-car based detectors. This analysis also demonstrated that the K-distribution was the only noise model tested which was capable of improving the detection performance of the box-car detector. These detectors were then applied to real target data, where it became clear that the Rician model for embedded target data, and indeed the Swerling models for target signals, did not match the observed data. Despite attempts to apply the Weibull and log-normal distributions as target models, this difficulty was not resolved. Ultimately, when applied to real target data, all of the detectors defined in this chapter achieved comparable results, with no improvement over the standard box-car detector. In the next chapter we will examine ways to make better use of our trial data, as data for the same region exists in more than one ping this will be exploited in order to increase the sample size for our parameter estimates.



(a)



(b)

Figure 7.14: ROC curves generated for real targets at (a) Low frequency (b) Medium Frequency




Figure 7.15: ROC curves generated for real targets in the full sonar bandwidth, without splitting the data into separate frequency bands

During this process we will also examine temporal variation in our noise-only trial data.

Chapter 8

Multi-ping Processing

“

As machines become more and more efficient and perfect, so it will become clear that imperfection is the greatness of man.

ERNST FISCHER (1899–1972)

A fundamental difficulty in estimating model parameters, and in applying the extended GLRT detector, is the small sample size. Many of the pings in our trial data overlap, covering the same region of the sea. It is therefore possible to make use of samples from multiple pings of a region, increasing the sample size. We refer to this as *multi-ping processing*.

In the past, this concept has been applied in Synthetic Aperture Sonar and Synthetic Aperture Radar by taking multiple pings along a chosen track [23, 74]. The data from these pings can then be combined, forming a synthetic array along the length of the

track. This technique provides the same advantages as increasing the size of the array, including an increase in resolution.

In this chapter additional samples are obtained by multi-ping processing after the matched filter, providing temporal statistics and an increased sample size for noise estimation at any given cell. Over short ranges, such an increase in data has been shown to reduce noise in Dual Frequency Identification Sonar systems [71]. Over longer ranges, this leads to an improved understanding of regional clutter statistics, and assists in the classification of noise regimes [106]. Rather than estimating clutter returns from a database of bathymetry and sea floor sediment types, this approach makes use of actual clutter returns from the sonar system itself. The temporal aspect of the data also allows us to investigate the variability of clutter in a region, in addition to classifying sources of clutter in terms of their temporal stability.

The impact of multi-ping processing will naturally be limited by the effects of reverberation, clutter and multipath. If spatial or temporal distance between pings is too great, then scatterers in the sea volume may have moved or may be observed from a different angle. It is therefore likely that the noise distribution (including the view of any correlating sea features) will change over a small number of pings.

This chapter describes methods for processing data from more than one ping. This begins with the derivation of a novel method of parameter estimation, whereby noise estimates calculated from more than one ping of a region can be combined. This method is compared against an existing technique, and its advantages demonstrated. The method is then applied to the generation of a clutter map, which visualises the sources of clutter in our trial data. We then define a framework for the detection of sonar targets over multiple pings, applicable to both box-car and GLRT detectors, and capable of processing static or moving targets. This detection framework is applied to simulated targets embedded in our trial data, and briefly evaluated against real targets in real sonar noise.

8.1 Combining Parameter Estimates

In the case of the K-distribution, the number of samples needed to accurately estimate the shape parameter ν is approximately $\max(200\nu^3, 30\nu^4)$ [105]. Yet, the samples used to estimate this quantity must be taken from a small region, close to the cell under test, in order to justify our assumption that the data comes from a homogeneous and representative noise environment. A solution to this difficulty is to make use of samples from multiple pings of this region.

In this case, it may be desirable to combine the results from previously calculated parameter estimates for each ping, rather than collecting the relevant samples from all of the pings and applying the entire estimation procedure to this data. In this section we will examine methods for combining parameter estimates of the K-distribution shape parameter. This will then be followed by an evaluation of the resulting estimates.

For the MMA and MMB estimators proposed by Abraham and Lyons (see Section 4.2), the authors provide an estimator-specific method for combining estimates. This work serves as our baseline and also demonstrates interest from others in the field. These two estimators were previously shown to depend on the inversion of an intermediary statistic \hat{D} . From equation 4.27 this statistic is fully defined by the parameters of a Gamma distribution with shape a and scale b . If we assume that a number of estimates have been calculated from equally-sized samples of K-distributed data, then these estimates can be combined through the following equations:

$$a = 1 + \sum_{i=1}^n (a_i - 1) \quad (8.1)$$

$$b = \left(\sum_{i=1}^n b_i^{-1} \right)^{-1} \quad (8.2)$$

These two formulae are remarkably simple, but require that the samples be estimated by

the authors' Bayesian method [4]. This method will be henceforth referred to as the *AB* method of combining parameter estimates. We now describe a similar formula acting directly on the K-distribution shape parameter, which will therefore be an estimator-independent method for combining estimates.

Suppose that:

$$X = (x_1, \dots, x_n) \sim K(\nu_x, \lambda_x) \quad (8.3)$$

$$Y = (y_1, \dots, y_n) \sim K(\nu_y, \lambda_y) \quad (8.4)$$

$$\begin{aligned} Z &= (X, Y) \\ &= (x_1, \dots, x_n, y_1, \dots, y_n) \\ &= (z_1, \dots, z_n) \sim K(\nu_z, \lambda_z) \end{aligned} \quad (8.5)$$

In cases where $\lambda_x \approx \lambda_y$ the resulting data may still be approximated by a K distribution, such that (ν_z, λ_z) describe the properties of the tuple $(X, Y) = (x_1, \dots, x_n, y_1, \dots, y_n)$ containing values distributed according to $K(\nu_x, \lambda_x)$ and $K(\nu_y, \lambda_y)$. Since $\hat{\lambda}$ is an estimate of the second moment, and both datasets are of the same length, $\hat{\lambda}_z = \frac{\hat{\lambda}_x + \hat{\lambda}_y}{2}$. However, to estimate the shape parameter of the combined distribution is not as simple.

We explore this relation by simulation with known parameters, estimating the resulting shape parameter. In the case of Figure 8.1a, the two tuples X and Y each consisted of 1,000,000 values. Despite this large number of samples, the variance of the shape estimate is still visible, and was subsequently smoothed by averaging over six realizations (Figure 8.1b). Fitting a smooth surface to these instances yielded:

$$\hat{\nu}_z = \frac{2\hat{\nu}_x\hat{\nu}_y}{\hat{\nu}_x + \hat{\nu}_y} \quad (8.6)$$

This relationship is demonstrated in Figure 8.2a, with the error between (8.6) and the original data demonstrated in Figure 8.2b. As can be seen the greatest errors are obtained




Figure 8.1: Shape value mapping obtained by simulation, with the true shape parameters of the two original tuples along the x and y axis, and the resulting shape parameter estimate along the z axis (a) single realization, (b) smoothed over multiple realizations

at large shape values, where we have already established that the shape estimators are of highest variance (see Section 4.2 and Section 5.1). It was suspected that this increased error is due to errors in the initial shape estimates, rather than a poor mapping of the shape parameter given by (8.6). This was verified by increasing the sample size for each tuple and averaging over increasingly larger numbers of simulations, which consistently decreased the error.




Figure 8.2: Fitted surface, (a) demonstrating agreement with simulated data, (b) error between fitted surface and smoothed simulation

Let us now take a more formal approach to the problem. As the $z \log(z)$ estimate offers

a fairly simple closed-form calculation we will make use of this technique.

Recall from Section 4.2 that the K-distribution scale parameter λ can be estimated for each tuple as follows:

$$\hat{\lambda}_x = \frac{1}{n} \sum_{i=1}^n x_i \quad (8.7)$$

$$\hat{\lambda}_y = \frac{1}{n} \sum_{i=1}^n y_i \quad (8.8)$$

$$\hat{\lambda}_z = \frac{1}{2n} \sum_{i=1}^{2n} z_i \quad (8.9)$$

and that the K-distribution shape parameter ν can be estimated by the $z \log(z)$ estimator as follows:

$$\hat{\nu}_x = \left[\frac{\frac{1}{n} \sum_{i=1}^n x_i \log(x_i)}{\hat{\lambda}_x} - \frac{\sum_{i=1}^n \log(x_i)}{n} - 1 \right]^{-1} \quad (8.10)$$

$$\hat{\nu}_y = \left[\frac{\frac{1}{n} \sum_{i=1}^n y_i \log(y_i)}{\hat{\lambda}_y} - \frac{\sum_{i=1}^n \log(y_i)}{n} - 1 \right]^{-1} \quad (8.11)$$

$$\hat{\nu}_z = \left[\frac{\frac{1}{2n} \sum_{i=1}^{2n} z_i \log(z_i)}{\hat{\lambda}_z} - \frac{\sum_{i=1}^{2n} \log(z_i)}{2n} - 1 \right]^{-1} \quad (8.12)$$

Theorem 8.1.1. *If $\hat{\nu}_x$ is the $z \log(z)$ shape parameter estimate for an n -tuple X , and $\hat{\nu}_y$ is the $z \log(z)$ shape parameter estimate for an n -tuple Y , then if the estimated second moments of X and Y are equal ($\hat{\lambda}_x = \hat{\lambda}_y$), the $z \log(z)$ shape parameter estimate for the $2n$ -tuple $Z = (X, Y)$ is given by*

$$\hat{\nu}_z = \frac{2\hat{\nu}_x\hat{\nu}_y}{\hat{\nu}_x + \hat{\nu}_y}. \quad (8.13)$$

Proof. (8.9) can be rewritten as

$$\hat{\lambda}_z = \frac{1}{2n} \left(\sum_{i=1}^n x_i + \sum_{i=1}^n y_i \right) \quad (8.14)$$

$$= \frac{1}{2} \left(\frac{1}{n} \sum_{i=1}^n x_i + \frac{1}{n} \sum_{i=1}^n y_i \right) \quad (8.15)$$

$$= \frac{\hat{\lambda}_x + \hat{\lambda}_y}{2} \quad (8.16)$$

Given the condition $\hat{\lambda}_x = \hat{\lambda}_y$, this reduces to

$$\hat{\lambda}_z = \frac{2\hat{\lambda}_x}{2} = \frac{2\hat{\lambda}_y}{2} = \hat{\lambda}_x = \hat{\lambda}_y \quad (8.17)$$

Therefore, (8.12) can be rewritten as

$$\hat{\nu}_z = \left[\frac{\frac{1}{2n} (\sum_{i=1}^n x_i \log(x_i) + \sum_{i=1}^n y_i \log(y_i))}{\hat{\lambda}_z} - \frac{\sum_{i=1}^n \log(x_i) + \sum_{i=1}^n \log(y_i)}{2n} - 1 \right]^{-1} \quad (8.18)$$

$$= 2 \left[\frac{\frac{1}{n} (\sum_{i=1}^n x_i \log(x_i) + \sum_{i=1}^n y_i \log(y_i))}{\hat{\lambda}_z} - \frac{\sum_{i=1}^n \log(x_i) + \sum_{i=1}^n \log(y_i)}{n} - 2 \right]^{-1} \quad (8.19)$$

From (8.17) we get:

$$\hat{\nu}_z = 2 \left[\left(\frac{\frac{1}{n} \sum_{i=1}^n x_i \log(x_i)}{\hat{\lambda}_x} + \frac{\frac{1}{n} \sum_{i=1}^n y_i \log(y_i)}{\hat{\lambda}_y} \right) - \left(\frac{\sum_{i=1}^n \log(x_i)}{n} + \frac{\sum_{i=1}^n \log(y_i)}{n} \right) - 2 \right]^{-1} \quad (8.20)$$

$$= 2 \left[\left(\frac{\frac{1}{n} \sum_{i=1}^n x_i \log(x_i)}{\hat{\lambda}_x} - \frac{\sum_{i=1}^n \log(x_i)}{n} - 1 \right) + \left(\frac{\frac{1}{n} \sum_{i=1}^n y_i \log(y_i)}{\hat{\lambda}_y} - \frac{\sum_{i=1}^n \log(y_i)}{n} - 1 \right) \right]^{-1} \quad (8.21)$$

$$= 2 [\hat{\nu}_x^{-1} + \hat{\nu}_y^{-1}]^{-1} \quad (8.22)$$

$$= \frac{2\hat{\nu}_x\hat{\nu}_y}{\hat{\nu}_x + \hat{\nu}_y} \quad (8.23)$$

□

It is interesting to note the form of (8.22). While (8.23) is more computationally efficient, one can view (8.22) as the reciprocal of the mean over $1/\nu$.

Corollary 8.1.2. *The form of (8.23) is a commonly recognised special case of the harmonic mean for two samples. It is trivial to extend the above proof to more than two tuples of equal length. Therefore given m tuples of length n , estimates of the K -distribution shape parameter can be combined by calculating the harmonic mean.*

This is defined as:

$$\hat{\nu} = \frac{m}{\sum_{i=1}^m \hat{\nu}_i^{-1}} \quad (8.24)$$

For this reason, we refer to our method of combining K -distribution shape parameter estimates as the *harmonic-mean* method.

When combining parameters in this way some tolerance must be chosen for $(\hat{\nu}_x, \hat{\nu}_y)$ to ensure radically different noise environments are not merged. Equally it must be

decided how close $(\hat{\lambda}_1, \hat{\lambda}_2)$ must be for practical implementations, such that the regions have a sufficiently similar scale for this calculation to be sensible.

In the context of sonar signal processing, two simple conditions are that parameter estimates should only be combined if they are estimated from samples recorded from a sufficiently similar distance and bearing from the sonar array, and within a short time of each-other. These conditions aim to avoid the combination of parameter estimates when the observed noise environment is likely to have changed systematically, due to factors such as variation in the sound speed profile over time, or the repercussions of increased absorption loss at long range.

8.2 Performance of Combined Estimates

In this section we extend the analysis of single ping parameter estimates (see Section 5.1) to that of combined parameter estimates. We use the same number of simulations and sample sizes as before, but now the samples are split into two or four parts prior to parameter estimation, and then the estimates are recombined by the AB and harmonic-mean methods defined in Section 8.1.

In our previous analysis the MMA and MMB estimators provided the best performance. They are also the only estimators applicable to both the AB and harmonic-mean methods. For this reason the MMA and MMB estimators are used in this section, to compare the performance of the harmonic-mean method to that of the AB method and direct estimation. Simulations applying the harmonic-mean method to the $z \log(z)$ and MME estimators provided results similar to those demonstrated here. Unsurprisingly, from the form of both techniques, neither method of combining estimates causes failures in the estimation procedure, which were shown to be a problem in Section 5.1. The

results of our simulations are summarised in terms of the mean square error, variance and mean-bias in Figures 8.3, 8.4 and 8.5 respectively.

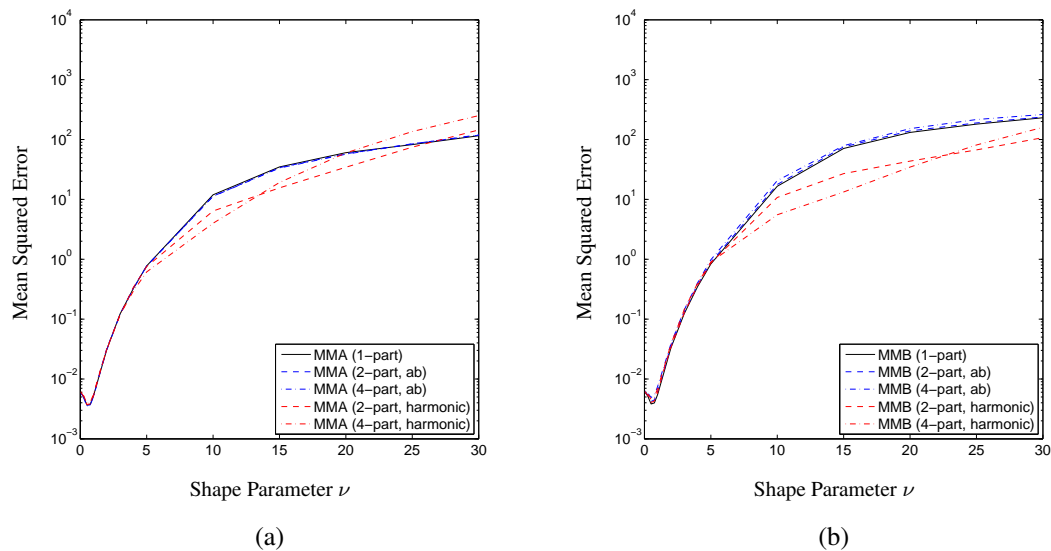


Figure 8.3: MSE for shape parameters in the range $[0.05, 30]$ and a sample size of $n = 2500$, with the sample split into multiple parts and estimates combined by AB and harmonic-mean methods. (a) MMA estimator, (b) MMB estimator

Across all of our simulations the AB method of combining shape parameter estimates achieved performance comparable to direct estimation. This can be seen in terms of the mean square error for a sample size of 2500 in Figure 8.3 and the associated variance in Figure 8.4. The greatest change in the performance of the MMA and MMB estimators was due to a slight increase in bias, shown in Figure 8.5. This resulted in a slight increase in negative-bias for the MMA estimator (underestimating the true value) and a slight increase in positive-bias for the MMB estimator (overestimating the true value). Overall the AB method is impressive, particularly as the greatest impact of bias did not occur at high values of ν , where the K-distribution is difficult to estimate.

In contrast, the harmonic-mean method of combining shape parameter estimates only achieved performance comparable to direct estimation for small values of ν , between 0 and 5. Surprisingly, for larger values of ν this method achieved far lower variance than

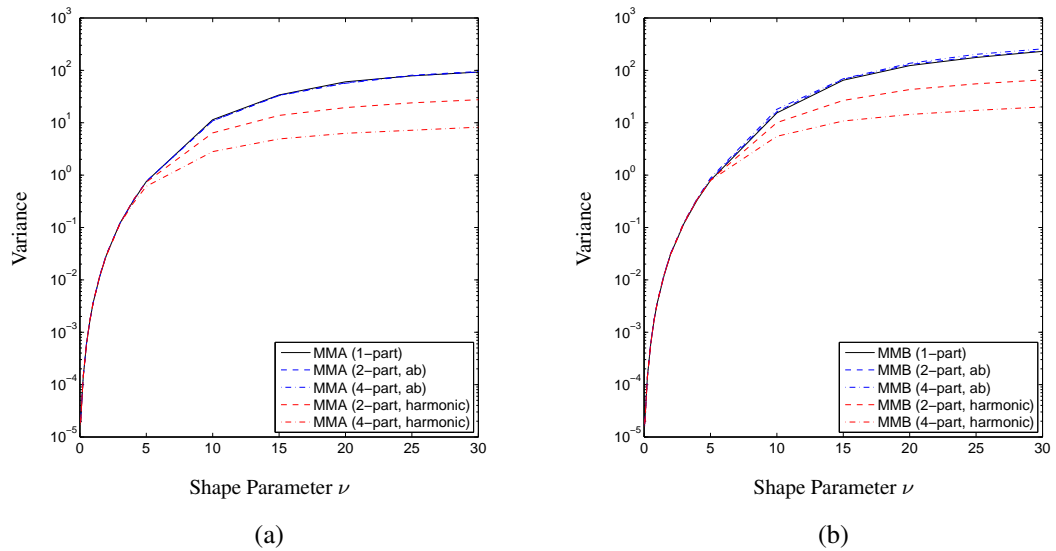


Figure 8.4: Variance for shape parameters in the range $[0.05, 30]$ and a sample size of $n = 2500$, with the sample split into multiple parts and estimates combined by AB and harmonic-mean methods. (a) MMA estimator, (b) MMB estimator

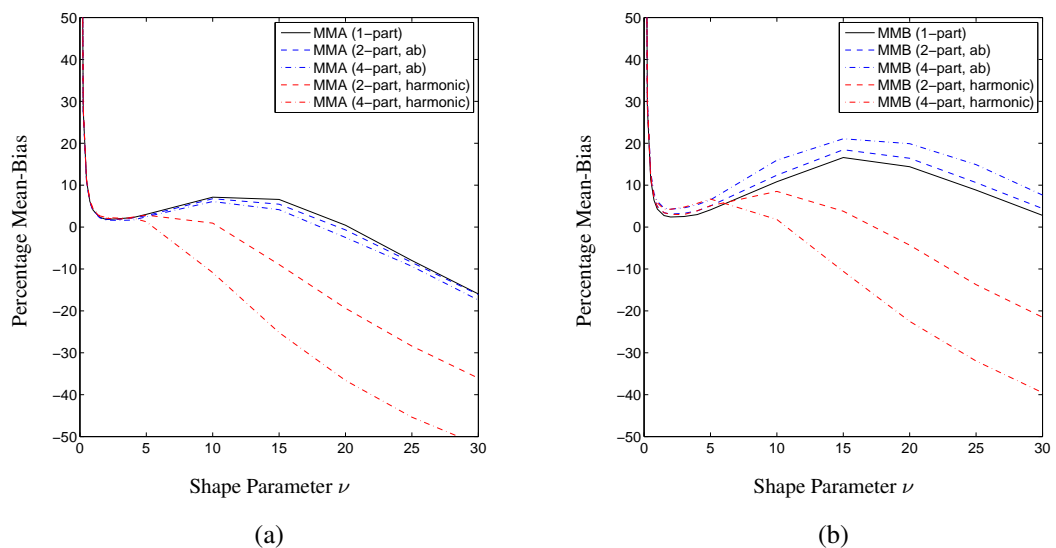


Figure 8.5: Percentage mean-bias for shape parameters in the range $[0.05, 30]$ and a sample size of $n = 2500$, with the sample split into multiple parts and estimates combined by AB and harmonic-mean methods. (a) MMA estimator, (b) MMB estimator

direct estimation, with the lowest variance observed when the data was split the most times prior to estimation. It appears that the mathematical relationship defined by the reciprocal in the harmonic mean captures an inherent property of the K-distribution. It was previously noted by Oliver and Lombardo [81] that the measure for their contrast estimator was proportional to $1/\nu$. Their observation of a linear relationship in the reciprocal of ν is in agreement with our findings. Therefore, when this constraint is applied to the estimation process, the resulting estimate is more accurate. However, as with the AB method, the harmonic-mean method increases the bias of the estimation process, as demonstrated in Figure 8.5. This large negative bias is far worse than that of the AB method, and causes an increase in mean square error for large values of ν . It should be noted that the extent of this bias is dependent on the estimation procedure used, and simulation with the MME and $z \log(z)$ estimators exhibit far lower bias, while still achieving lower variance than direct estimation. Despite this increase in bias, the harmonic-mean method frequently achieved lower mean square error, as shown in Figure 8.3.

The low variance associated with our harmonic-mean method of parameter estimation is particularly advantageous, as it is generally easier to reduce bias than variance. If low bias is more important than low variance for a particular application, the harmonic-mean method can be augmented by a bias-reduction scheme [28]. Such techniques can greatly improve the bias of the harmonic-mean method at low computational cost (e.g. evaluating a polynomial), but in our simulations this also resulted in a large increase in variance, leading to a greater mean square error than direct calculation.

In conclusion, the AB method provides an accurate means of combining partial calculations of shape parameter estimates for the MMA and MMB estimators. The main advantage of the harmonic-mean method is that it can be applied to any estimator of the K-distribution shape parameter, as demonstrated by our positive results here. Although the harmonic-mean method is exact for $z \log(z)$ estimates of samples with the same

scale, there is a trade-off between bias and variance when applied to other estimators of the K-distribution shape parameter.

8.3 Multi-ping Composite Images

In this section we will examine the methods required to achieve a composite visualisation of multiple pings in real-world coordinates. We begin by defining an algorithm to convert our trial data from polar coordinates to cartesian coordinates. This is followed by the registration of each ping in a global coordinate system. We then calculate various summary statistics for each geographic position, generating an individual clutter map from each statistic. These clutter maps are then compared, before briefly analysing the temporal variation between different noise environments.

8.3.1 Coordinate Space

As discussed in Section 3.3, our trial data was obtained in polar coordinates centered on the ship. To visualise the degree of clutter in a region of the sea, this data must be converted into a global spatial representation. We achieve this by creating a matrix which represents a sonar ping in Cartesian coordinates, and then converting the indices for each cell in this matrix into polar coordinates relative to the centre of the ping in question. These polar coordinates are then used to populate the matrix by nearest-neighbour interpolation from our trial data. In MATLAB this can be achieved by applying the code provided by Mathworks Solution ID 1-OYFMA [95]. It should be noted that this code must be modified to make use of nearest-neighbour interpolation, as `griddata` defaults to linear interpolation.

A further consideration is the angular spacing of the beams in our data. In our trial data,

the 64 beams are spaced equally in cosine space; failing to account for this non-linear spacing would result in a warping of the data, which would lead to undesirable artefacts in the multi-ping composite image. This can be achieved by calculating our range of angles linearly between $\cos(0^\circ)$ and $\cos(180^\circ)$, and taking the inverse cosine of this result. Combining this with the aforementioned MATLAB code leads to Algorithm 8.1.

Algorithm 8.1 IM_POL2CART: Convert trial data to cartesian coordinates

```
[nrows, ncols] = size(polar_data);

%prepare matrices containing the polar coordinate data
rho = repmat([1:nrows]', 1, ncols);
theta_values = pi + acos(linspace(cos(0), cos(pi), ncols));
theta = repmat(theta_values, nrows, 1);

%convert the polar coordinates to cartesian
[x,y] = pol2cart(theta, rho);

%use ffgrid to prepare the data
[zz, xx, yy] = ffgrid(x, y, polar_data, 1, 1);

%use griddata to produce the final image
cartesian_data = griddata(x, y, polar_data, xx, yy', 'nearest');
```

This algorithm was used to generate Figure 3.8 in Section 3.3.

8.3.2 Registration

A vital step in multi-ping processing is to calculate the rotation and translation between each ping. This is known as registration, and enables us to reference specific geographical locations (using global coordinates) within the coordinate systems of each ping. Non-acoustic data, such as GPS coordinates for surface ships, can provide functional registration parameters without extensive processing. Additionally, much work has been applied in the estimation of ship location, from hydrodynamic modelling of towed array position [2] to the correction of array shape [54]. In our case, we need

only apply the GPS data because the towing vessel moved in a straight line, but such techniques serve to generalize our methods to other datasets. From the latitude and longitude coordinates provided by GPS, it is possible to calculate the distance and bearing between GPS readings. The distance and bearing between pings should be treated as a polar coordinate, and converted to Cartesian coordinates to calculate the relative change in position (x and y) between pings. Cumulatively summing this change in position between pings produces the translation parameters t_x and t_y , which serves as a simple mapping from WGS84 (see Section 3.3) to Cartesian coordinates. The final step in registration is to suitably rotate each ping of data so that it is oriented correctly given the bearing of the sonar array, defined in Section 3.3 as θ_S .

We will now define the rotation and translation in terms of affine transformations. This is a flexible approach since, depending on implementation requirements, we can apply rotation first and then use translation parameters to process data in the area of interest, or if desired (and subject to memory constraints) apply both the rotation and translation directly.

$$\begin{bmatrix} x' \\ y' \\ 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 & t_x \\ 0 & 1 & t_y \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \cos(\theta) & \sin(\theta) & 0 \\ -\sin(\theta) & \cos(\theta) & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} \quad (8.25)$$

$$\begin{bmatrix} x' \\ y' \\ 1 \end{bmatrix} = \begin{bmatrix} \cos(\theta) & \sin(\theta) & t_x \\ -\sin(\theta) & \cos(\theta) & t_y \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} \quad (8.26)$$

Once the new locations have been calculated, the matrix can be populated by nearest-neighbour interpolation from the original Cartesian data.

Due to memory constraints, we apply rotation and translation separately. The region of sea covered by our trial data was split into blocks of 500 by 500 samples, as shown in Figure 8.6. In MATLAB our Cartesian data was rotated using the function

`imrotate(cartesian_data, theta, 'nearest', 'crop')`. This method allows easy access to data in a small region over multiple pings, without needing the large amount of memory that would be required to hold all of our trial data in memory at once.




Figure 8.6: Grid of 195 blocks of 500x500 samples used to generate each multi-ping composite. The region covered by our trial data is marked in grey

8.3.3 Visualisation

Once registration parameters are found and the relevant blocks are stored for each ping of Cartesian data, it is then possible to continue our analysis. In order to generate a composite image of the sea floor across multiple pings, we must decide how to combine overlapping regions of data.

While generating Figure 5.7, where we compared values of the K-distribution shape parameter to regions of non-Rayleigh data, it was found that shape parameters satisfying

$0 < \nu \leq 10$ were a useful indication of non-Rayleigh environments. This figure was generated by truncating the maximum value of ν at 10, applying the following piecewise function:

$$\text{truncate}(\nu) = \begin{cases} 10 & \text{if } \nu > 10 \\ \nu & \text{if } \nu \leq 10 \end{cases} \quad (8.27)$$

A useful visualisation of non-Rayleigh environments (clutter) may therefore be obtained by the arithmetic mean of (8.27) over multiple pings. Alternatively we might directly combine shape parameter values, using the arithmetic mean, median, or harmonic mean over these pings. So we have four methods of combining overlapping regions of data:

- arithmetic mean of ν truncated at $\nu > 10$
- arithmetic mean of ν
- median value of ν
- harmonic mean of ν

As measures of central tendency, the arithmetic mean, median and harmonic statistics have greater theoretical foundation than the arithmetic mean of a truncated value. Additionally, in regions with homogeneous scale between pings (where λ remains constant), we established in Section 8.1 that the harmonic mean is a reasonable approximation of the shape parameter estimate over multiple pings. These four methods are compared in Figure 8.7 and Figure 8.8, whether they have been applied to our low and medium frequency trial data.

Of these four methods, the arithmetic mean in Figures 8.7 provided composites which demonstrate the large-scale trends in the sonar environment. These trends are also visible in all other composites, but there is limited detail and low contrast. Applying the median greatly improved the contrast between clutter-limited and Rayleigh environments, but

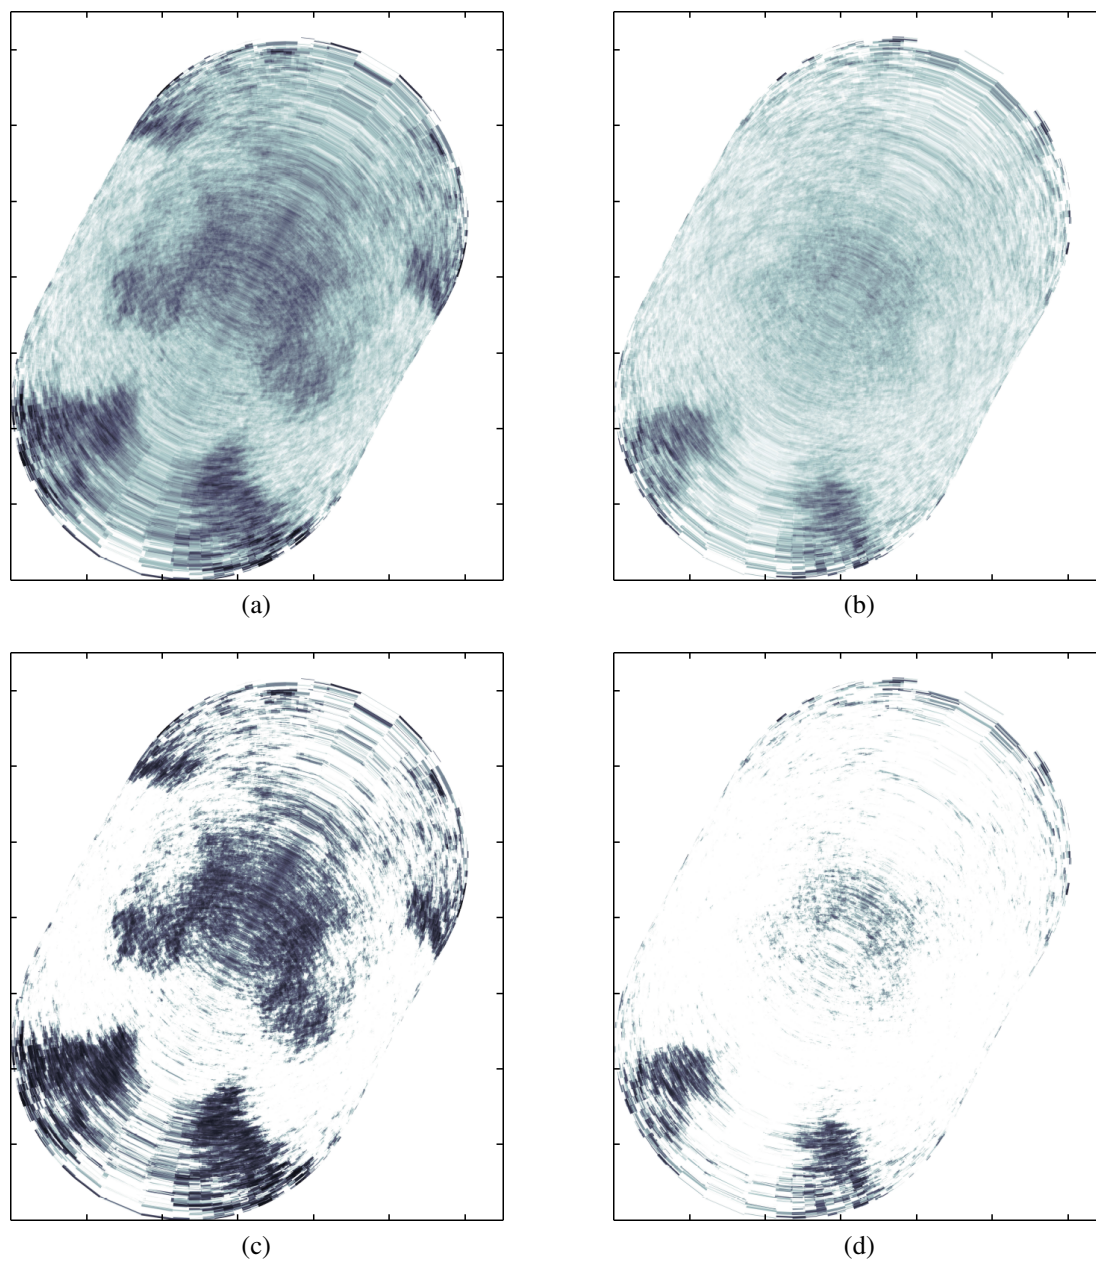


Figure 8.7: K -distribution shape parameter composites (a) low frequency mean (b) medium frequency mean (c) low frequency median (d) medium frequency median




Figure 8.8: K -distribution shape parameter composites (a) low frequency mean of truncated values (b) medium frequency mean of truncated values (c) low frequency harmonic mean (d) medium frequency harmonic mean

as can be seen at medium frequency in Figure 8.7d this method still provides limited detail. This is not surprising, as the median statistic does not merge the shape parameter values over multiple pings, and so some information must be lost when selecting a single value by rank alone.

In a sense, truncating the maximum value of ν to 10 is merely a method of enhancing the contrast between clutter-limited and Rayleigh environments. This can be seen in Figure 8.8. Yet this statistic also provides a greater level of detail than the median. These two advantages are even more evident for the harmonic mean, where a number of small-scale regions of clutter are clearly visible (such as half way down the left hand side of Figure 8.8d). We will now focus on the medium frequency composite generated by the harmonic mean of the shape parameter estimates.

Analysis of Composite Images for Visualisation

Despite the potential localisation error caused by towing the array, registration by GPS coordinates has achieved positive results. The data for each ping has been converted to Cartesian coordinates and superimposed on top of each other, with overlapping regions combined for display. The towing vessel is shown as a sequence of red crosses in Figure 8.9, travelling towards the top right of the figure.

The following findings are observed:

- High levels of clutter are visible at the bottom left of Figure 8.9.
- Further peaks off the centre of the track are of particular interest, as some of these correlate with shipwreck data for the region. In total the wrecks of four merchant ships have been localised, though the identity of two of these is unknown.




Figure 8.9: Composite achieved by converting the K -distribution shape parameter for multiple medium-frequency pings to Cartesian coordinates and performing registration. In order to visualize the dataset, values in overlapping regions were combined by harmonic mean. Here the ship track, from bottom left to top right, is shown by the crosses at the centre of each ping.

- Large arcs can be seen in front of the sonar array which do not relate to any sea features. These are predominantly found in the first few beams and are caused by echoes from the ship itself (best seen at the top right).
- Large values of the shape parameter were obtained directly under the ship track, as expected. While these values are large, this does not denote any particular clutter properties in the region and is merely an artefact of the sonar system's proximity to this area of sea floor. It is known that the ship travelled directly over three additional merchant shipwrecks. Our inability to identify these wrecks from the available data demonstrates the masking effects of reverberation at short range, and corroborates our assertion that this sonar system is not designed for detection at short ranges.

The erroneous inclusion of large shape parameter estimates caused by echoes from the ship, and reverberation under the ship track, could be easily rectified by applying a mask to the ping data before forming the composite image. The mask should remove all data within a chosen *dead range*, and any beams of data which are corrupted by the towing vessel or shipping noise. An example of such a mask is provided in Figure 8.10.

The composite plot of Figure 8.9 demonstrates that meaningful noise statistics can be collected over time from multiple locations. A similar technique was previously employed by Prior [106] to survey the sources of clutter on the Malta Plateau. This was performed from over 140 pings, in a manner similar to our own method for generating a composite sonar map. The main difference in our work has been the application of the K-distribution shape parameter as our chosen noise statistic, while Prior defined statistics in terms of the relative change and standard deviation of raw amplitudes [106]. In subsequent work, Abraham and Prior applied this methodology to a new dataset [9]. This made use of The K-distribution shape parameter, and considered the inverse problem to obtain an estimate of the *independent-scatter density*. However,




Figure 8.10: Figures depicting the masking of sonar data for visualisation: (a) original data (b) mask (c) resulting data

these clutter maps were based on the median value of the shape parameter over multiple pings, resulting in a loss of data which is likely to bias estimates obtained from certain bearings.

In conclusion, the choice of the K-distribution shape parameter as a noise statistic is advantageous as it is a proven indicator of sonar clutter, and a natural by-product of our suggested box-car detector (see Section 7.6). Combining estimates of the K-distribution shape parameter, by the harmonic mean, has been shown to provide a viable estimate of local noise statistics, if data is collected from multiple pings. When applied to clutter maps, this methodology demonstrates good contrast between clutter-limited and Rayleigh environments whilst preserving small-scale features.

8.3.4 Temporal Statistics

While composite images are useful for visualization, the process results in a loss of temporal information. Spatial analysis is common in most fields, but the temporal variation of sonar noise is also of interest. For example Chantler et al. [30] made use of temporal statistics to classify targets in a short-range sonar system.

We will now briefly examine the temporal variation of three noise sources identified in our trial data: reverberation, clutter, and a shipwreck. Reverberation and clutter are defined in Chapter 2, and identified in the trial data in Chapter 3. The shipwreck analysed in this section is an unidentified cargo ship known to the UK Hydrographic Office, at a depth of 130 metres.

In order to select regions of reverberation and clutter it was useful to calculate K-distribution shape parameter estimates for each ping and stack these estimates vertically, forming a three-dimensional representation of the sonar environment with time along the third axis, rather than the two-dimensional composites generated earlier in this

section. An example of this three-dimensional representation is shown in Figure 8.11, where the data has been masked to reduce the effects of short-range reverberation. Low shape parameter estimates can be seen at the bottom of this figure, highlighting two regions of dense clutter. It should be noted that due to left-right ambiguity, these two regions in our figure are likely to represent a single region in real life. Additionally, low shape parameter estimates near the centre of the figure highlight some of the shipwrecks in the area.

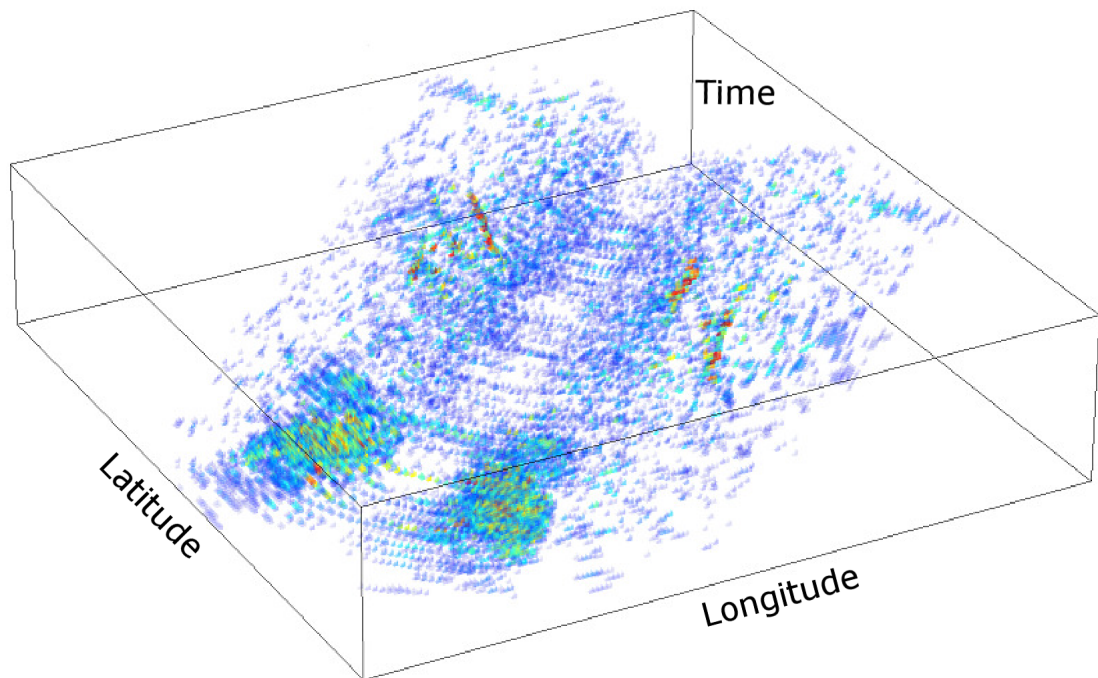


Figure 8.11: Three-dimensional representation of K-distribution shape parameter estimates, stacked such that the z-axis represents time. Blue and white denote large shape parameter estimates (noise limited), while green and red denote small shape parameter estimates (clutter limited).

The selected regions for each noise environment are identified in Figure 8.12. At each of these locations, shape parameter estimates for the K-distribution were collected over multiple pings. Histograms were then generated for each ping to demonstrate the distribution of shape parameter estimates less than 20, corresponding to a difficult to process, non-Rayleigh environment. A small selection of these histograms are shown for the low frequency band in Figure 8.13, and for the medium frequency band in Figure

8.14.

Interesting effects are visible for environments containing reverberation and clutter. At low frequencies cluttered environments tended to contain a large number of shape estimates less than 3. However, at medium frequency there was far less consistency in the shape parameter estimates, as can be seen by the increased range and temporal variation of the peak in Figure 8.14a. This also demonstrates how clutter depends on frequency, as clutter in the medium frequency band (Figure 8.14a) produces a histogram similar to reverberation in the low frequency band (Figure 8.13b). The way in which clutter depends on frequency is related to the makeup of the scattering surface, and so different materials are expected to respond differently. This is true of the shipwreck data, where the temporal characteristics are similar to that of clutter at low frequency, but there is far less variation between frequency bands (Figure 8.13c and 8.14c).

The temporal characteristics of target-like data are of particular interest in this chapter. Figures 8.13c and 8.14c suggest that the temporal statistics of a static target may be sufficiently stable that the GLRT detector can be applied over multiple pings (in the same manner that it was extended over multiple samples).




Figure 8.12: Manually categorised data. Black boxes contain background noise, clutter and reverberation. The red box contains ship wreck data




Figure 8.13: Demonstrative selection of histograms for K-distribution shape estimates calculated at low frequency in (a) clutter (b) reverberation and (c) ship wreck data




Figure 8.14: Demonstrative selection of histograms for K-distribution shape estimates calculated at medium frequency in (a) clutter (b) reverberation and (c) ship wreck data

8.4 Multi-ping Detection

In this section we will describe methods of detection over more than one ping of data. This is an extension of the algorithms described in Chapter 7, and includes a discussion of the additional assumptions made. In the following sections we then offer a performance analysis against simulated targets in real sonar data, followed by a brief performance analysis against real targets in real sonar data.

We will begin by examining the extended GLRT statistic (7.18) defined in Section 7.3.

$$GLR(A_X) = \frac{\prod_{i=1}^n p(x_i|\hat{\theta}, H_1)}{\prod_{i=1}^n p(x_i|\hat{\theta}_0, H_0)}. \quad (8.28)$$

Equivalently, in the logarithmic domain this is defined as

$$\log [GLR(A_X)] = \sum_{i=1}^n \log [p(x_i|\hat{\theta}, H_1)] - \sum_{i=1}^n \log [p(x_i|\hat{\theta}_0, H_0)]. \quad (8.29)$$

As can be seen, in the logarithmic domain, the product of independent probabilities becomes a sum of log-probabilities. Thus if we assume that we have already calculated a set of log-GLR values for the same location and sufficiently similar time-period, we can combine these probabilities in the same fashion. Summing the log-GLRs, as shown below:

$$\log [GLR_m(A_X)] = \sum_{k=1}^m \left\{ \sum_{i=1}^n \log [p(x_i^{(k)}|\hat{\theta}_j^{(k)}, H_1)] - \sum_{i=1}^n \log [p(x_i^{(k)}|\hat{\theta}_0^{(k)}, H_0)] \right\} \quad (8.30)$$

where GLR_m denotes a multi-ping GLRT statistic over m pings. If we recall the form of the box-car method of selecting the target and noise samples used to estimate the model parameters, it becomes evident that this approach assumes a static target. If the target has moved into the location under test in the current ping, but was not present at this location previously, then we would be combining the low probability of a target

associated with previous pings, with the high probability of a target in the current ping. This is likely to result in a false negative, so (8.30) is an inappropriate detection scheme for moving targets.

Rather than using both target and noise models over more than one ping, multi-ping detection can also be achieved by only considering noise samples over multiple pings, combining the shape parameter estimates for use in either the box-car detectors or the noise component of the GLRT models. This should result in a more accurate estimate of the K-distributed background noise, which was shown to increase the performance of single-ping detectors in Section 7.6.

An important concern is the area covered by the noise samples used in this estimate. If these samples are taken from different angles (as will typically be the case) then the area of the sea floor used for the noise estimate is greatly increased, as demonstrated in Figure 8.15. The larger the area covered by the noise samples, the less likely it will be that the noise distribution is homogeneous across this region. This is undesirable as a single noise estimate is not appropriate for inhomogeneous regions, and may result in corruption of the noise estimate.

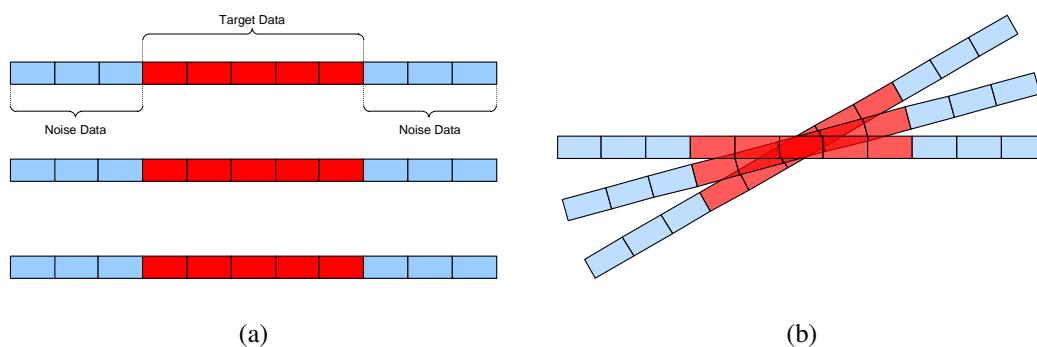


Figure 8.15: Demonstration of multi-ping processing for a static target, (a) box-car configuration for sampling the target and noise data from 3 pings (b) demonstration of possible area covered by the three box-cars

If the assumption of a static target does not hold, we can take a different approach to

estimating noise. If we assume that the target has moved a distance at least equal to its own length between consecutive pings, then it is possible to make use of data from the noise environment which was previously subsumed in the embedded target data. This should lead to a far more accurate estimate of the noise environment, leading to more accurate detection. This also overcomes the issue raised above, as the additional noise samples are taken from an area closer to the target and thus cover a smaller region of the sea, as shown in Figure 8.16.




Figure 8.16: Demonstration of multi-ping processing for a moving target, (a) box-car configuration for sampling the target and noise data from 3 pings (b) demonstration of possible area covered by the three box-cars

With these two methods we are now able to apply multi-ping processing to both static and dynamic targets. The performance of these methods will be evaluated in the following sections.

8.5 Application to Real Sonar Noise

In this section we assume a static target, and are therefore able to apply the multi-ping GLR_m statistic proposed in (8.30). The process is applied to both the single-point and multi-point GLR_m statistics. These simulations are a straight forward extension of the method described in Section 7.1 and implemented in Section 7.6. For synthetic targets

and synthetic noise, multi-ping processing is identical to the extended GLRT defined in Section 7.3, with additional samples obtained over time (from multiple pings) rather than spatially (in range). As we have already performed an equivalent performance analysis in Section 7.5, we skip this step and progress to the performance analysis for synthetic targets embedded in trial data.

Our experimental design demands that simulations take place on a regular grid in GPS coordinates, and is applied to all pings covering these points. These GLR_m statistics have already been computed, and it is trivial to sum these together, calculating the multi-ping GLR_m statistics in the form (8.30). Following this, global thresholds were applied to the GLR_m statistics, from the minimum to maximum values of these statistics over the entire performance analysis. The ROC curves generated by this procedure are provided in Appendix B, Figures B.10 to B.15.

The additional samples provided by the single-point multi-ping GLRT lead to a significant performance increase over the single-point GLRT applied to a single ping. The results shown in Figure 8.17a are typical of this experiment. For the Rician-Weibull GLRT this occurred at all target strengths and frequency bands. Surprisingly, the Rician-Rayleigh GLRT had greater difficulty with loud targets, as shown in Figure 8.17b. We suspect that this is due to the increased *kurtosis* (a sharp peak in the probability density function), and decreased *skew* (the extent to which the right tail is longer than the left tail) obtained when a strong Swerling 0 target is embedded in K distributed noise. This is demonstrated in Figure 6.2 of Section 6.1. By contrast weak target signals result in a distribution with greater similarity to a Rayleigh distribution. The lack of a smooth curve for the Rician-Weibull GLRT in this Figure is also interesting, as this suggests a multi-modal distribution for the GLR_m statistic. For the single-point multi-ping GLRT this effect was most evident for loud targets. Although this is not investigated further in this thesis, it is precisely under these circumstances that a noise-dependent threshold may prove advantageous. The GLRT detectors applied here are inherently adaptive to

local noise characteristics, but as noted in Section 7.3, the distribution of the resulting GLR statistic may still depend on local noise statistics.

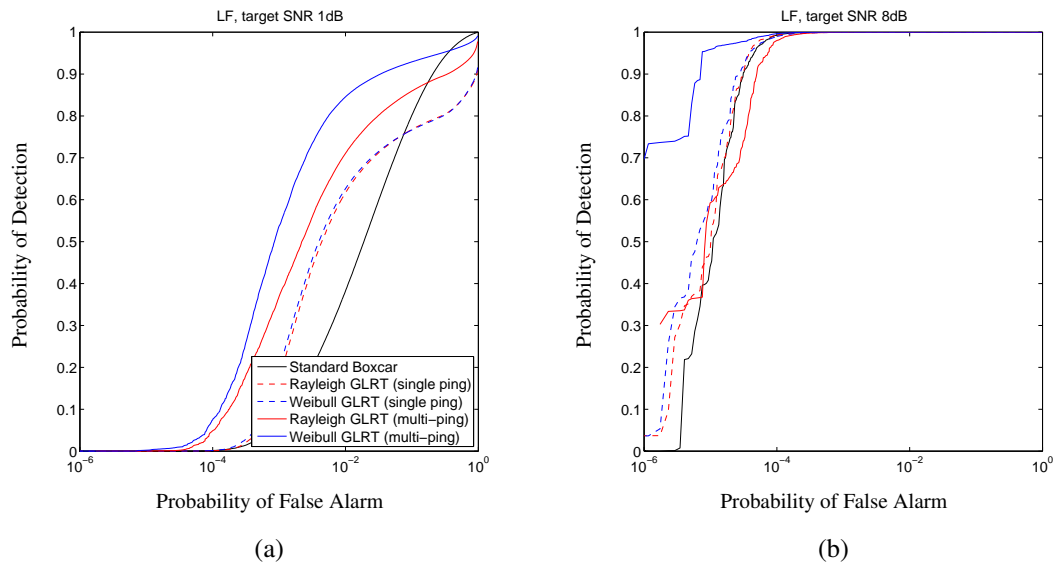


Figure 8.17: ROC curves for simulated detections in the low frequency band against Swerling type 0 simulated targets. Processed by the multi-ping single-sample GLRT detectors (a) 1dB target (b) 8dB target

When multi-ping processing was applied to the extended GLRT detector, perfect discrimination (100% detection) was quickly achieved at the false alarm rates tested. However, the previously identified multi-modal structure of the GLR_m statistic was clearly present in scenarios where perfect discrimination was not achieved, as shown in Figure 8.18. Although the Rician-Rayleigh distribution falls below the detection performance of its single-ping counterpart, this figure demonstrates that higher detection performance can still be achieved at lower false alarm rates, even in scenarios where the noise model is a poor fit to the sonar data.

It should be noted that we have still not considered the accuracy of the target model. The consequences of this will be demonstrated in the following section, as we briefly apply multi-ping processing to real sonar targets.




Figure 8.18: ROC curves for simulated detections in the low frequency band against 1dB Swerling type 0 simulated targets. Processed by the multi-ping multi-sample GLRT detectors

8.6 Application to Real Targets

We now extend the analysis performed in Section 7.8 to multi-ping processing. As the target was moving from ping to ping it was not possible to apply multi-ping processing in the manner described by (8.30). Consequently, we applied multi-ping processing to the noise estimate by combining estimates from multiple pings, using the sampling procedure for a moving target as shown in Figure 8.16. In an attempt to avoid inhomogeneity in the noise environment, noise estimates were only combined if they were measured at a similar distance from the sonar platform, and within a short time of each other. This was achieved by only combining parameter estimates with estimates from the previous and subsequent ping of data covering the region. If $\hat{\nu}^{(S)}$ denotes a shape parameter estimate calculated from samples in the region thought to contain the target, and $\hat{\nu}^{(N)}$ denotes a shape parameter estimate calculated from samples


in the region outside the guard bands, this can be written as:

$$\hat{\nu}_i = 3 \left[\frac{1}{\hat{\nu}_{i-1}^{(S)}} + \frac{1}{\hat{\nu}_i^{(N)}} + \frac{1}{\hat{\nu}_{i+1}^{(S)}} \right]^{-1}. \quad (8.31)$$


The resulting ROC curves from this analysis are provided in Figures 8.19 and 8.20. In Figure 8.19 we see the performance of our detectors in the low and medium frequency bands, while Figure 8.20 contains ROC curves for the full bandwidth of the sonar system. As can be seen, combining K-distribution shape parameter estimates does effect the resulting ROC curve, and can increase probability of detection. At low false alarm rates in low frequency data, the multi-ping Weibull-K and multi-ping lognormal-K GLRT detectors achieved a lower probability of detection than the standard box-car detector. This was not expected, and is likely to be a side-effect of the poor fit of our target models. By contrast, the multi-ping K-distribution box-car detector is not affected by a poor representation of the target model, and therefore did not exhibit a lower probability of detection than the standard box-car.

In conclusion, the application of multi-ping noise estimation to the K-distribution box-car detector consistently achieved equivalent or better probability of detection. This performance analysis demonstrates the advantage of multi-ping noise estimation.

In this chapter we derived a new method of combining K-distribution shape parameter estimates, which is shown to achieve lower variance, but increased bias, when compared to direct estimation and the *AB* method proposed by Abraham and Lyons [4]. We then demonstrated that this method can be applied to the generation of clutter maps with positive results. Returning to target detection, a framework was defined for the detection of sonar targets over multiple pings, which is applicable to both box-car and GLRT detectors, and capable of processing static or moving targets. This detection framework was applied to both simulated targets embedded in our trial data, and real targets in real sonar noise. Although we were unable to increase the probability of



(a)



(b)

Figure 8.19: Multi-ping ROC curves generated for real targets at (a) Low frequency (b) Medium Frequency




Figure 8.20: Multi-ping ROC curves generated for real targets in the full sonar bandwidth, without splitting the data into separate frequency bands

detecting real targets by using a GLRT detector, applying multi-ping noise estimation to the K-distribution noise model does provide a consistent improvement in detection performance for the the box-car detector. The next chapter provides the conclusion to this thesis, including a summary of our contributions, the limitations of our research, and potential avenues for future work.

Chapter 9

Conclusion

“

All things are difficult before they are easy.

— THOMAS FULLER (1608–1661)

Noise is frequently encountered when processing data from the natural environment. This is a particular concern for remote-sensing applications, where the goal is to collect accurate data about an entity or an event, by observing it from a distance. The accuracy of data gathered under these circumstances is limited by the noise present in the environment. In remote sensing techniques, such as active sonar where targets may be masked by noise, the accurate modelling of noise can lead to improved detection algorithms. Rather than merely accepting that sonar noise results in unavoidable error in active sonar systems, it is possible to mitigate the detrimental effect of noise.

It was hypothesised that by analysing the statistical properties of sonar noise, and applying the results of this analysis to the problem of target detection, better detection performance can be achieved.

In this thesis we have evaluated the performance of established models of sonar noise against trial data collected by a long-range active sonar system. This led to the identification of the K and Weibull distributions as the most accurate models of noise in our trial data. We then proceeded to apply these distributions to two different methods of detection. The box-car detector, which is standard in the field of active sonar signal processing, was applied in its generic form. As an alternative to this detector, we developed Generalized Likelihood Ratio Test (GLRT) detectors which also include a model for the statistics of the target. Each of these were initially evaluated within entirely synthetic environments. We then made use of trial data provided by Thales Underwater Systems Ltd., applying each detector to simulated targets embedded in real sonar noise, and then to trial data containing real targets.

Until now, multi-ping processing has been reserved to synthetic aperture processing (which takes place prior to the matched filter), decision based rules for cluster analysis (which takes place after the detection process), and the generation of clutter maps. In this thesis we performed multi-ping processing at both the noise-estimation and detection stages of the sonar processing chain. This included a framework for multi-ping processing with GLRT detectors, which allows for the detection of static or moving targets. As noise statistics are independent of target velocity, both of these methods made use of noise estimation over multiple pings. To this end, we developed a novel computationally efficient method of combining noise estimates over multiple pings. This technique was applied to the multi-ping GLRT detector for moving targets, evaluated against real target data. We also applied this method to the mapping of clutter in our trial data.

Simulation and performance analysis have played vital roles throughout this thesis, from the validation of algorithms, to the estimation of critical values and the analysis of detection performance. These simulations have established that improvements in detection performance can indeed be achieved through improved modelling of sonar

noise statistics. Furthermore, under the assumption of a Swerling 0 target signal, we found that the generalized box-car algorithm is best suited to the detection of strong target signals, while the GLRT detectors achieved the greatest performance increase for weak target signals.

In the remainder of this concluding chapter, we first discuss our contributions to statistical signal processing in active sonar. Following this we examine the limitations of our research and suggest areas for future work.

9.1 Contributions

This thesis includes the following contributions:

- Current sonar literature has made frequent use of the Kolmogorov-Smirnov test when evaluating the goodness-of-fit between noise distributions and sonar trial data. The Kolmogorov-Smirnov test has greatest sensitivity to regions of the candidate distribution where a large number of samples exist. However, if data is to be processed by normalization and thresholding (as is commonly the case), the far tail of the distribution is of greatest interest. Therefore, the Kolmogorov-Smirnov test is not well suited to this problem. We proposed the upper-tail Anderson-Darling test as a suitable alternative. The upper-tail Anderson-Darling test has greatest sensitivity in the far tail of the distribution, and is therefore a better match to the problem at hand (see Chapter 5). The benefit of the Anderson-Darling test was demonstrated in our special issue paper published in the IEEE Journal of Oceanic Engineering [19]. Our recommendation has since been applied by others in the field of sonar signal processing [8].

- Of the noise distributions tested, we verified that the K-distribution provided the best fit to the trial data collected by a long-range sonar system in a shallow water environment. To the best of our knowledge, no tables of critical values existed for K-distribution samples. We tabulated critical values for the K-distribution, with parameters estimated by the MME estimator with a $z \log(z)$ seed, and evaluated by the upper-tail Anderson Darling test (see Section 5.2). These tables were published in our special issue paper for the IEEE Journal of Oceanic Engineering [19].
- We applied a generic box-car algorithm for the detection of targets, applicable to any arbitrary noise distribution (see Section 7.2). This facilitated the application of an environmentally adaptive threshold, which varied with the local noise statistics for any chosen noise distribution. Using this algorithm we simulated ROC curves for Swerling 0 targets, embedded in our trial data. We then compared the performance of the K-distribution, Weibull distribution, Rayleigh distribution and Gaussian distribution as models for sonar noise in the box-car detection scheme. Our analysis revealed that despite the Weibull distribution offering a reasonable fit to cluttered non-Rayleigh data, there was no observable performance increase. In contrast, the K-distribution resulted in an increased probability of detection at low false alarm rates (see Section 7.6). This suggests that there is a critical level of accuracy that a noise model must reach for its implementation to increase probability of detection.
- We defined and evaluated GLRT detectors for the detection of sonar targets, making use of established distributions for sonar noise, and the Rician distribution for sonar targets. We showed that these GLRT detectors can provide far greater detection performance, but in practice they require an accurate model of the sonar target, which is not provided by the Rician distribution (see Chapter 7). If the target model is inaccurate, performance is comparable to the standard box-car detector. Although performance gains can still be achieved by applying multi-ping processing to the noise parameter estimates, this is no better than applying multi-ping processing to the K-distribution box-car detector (see Chapter 8). The inaccurate modelling of sonar

target statistics is a limiting factor to the utility of these detectors.

- We derived and evaluated a new method for combining estimates of the K-distribution shape parameter (see Section 8.1). This method has low computational cost, and unlike the method proposed by Abraham and Lyons [4], is not restricted to any particular parameter estimator. The technique is exact for data estimated by the $z \log(z)$ estimator if the estimated scale parameters are equivalent. This method was also applied to the MMA and MMB estimators (see Section 8.2), and are shown to achieve at-least comparable and often reduced variance compared to both direct estimation and estimates combined by Abraham and Lyons' AB method. However, our method can lead to an increase in bias. This increase in bias can be mitigated by applying bias-reduction techniques, but these can cause an increase in variance. In conclusion, when combining estimates of the K-distribution shape parameter, in cases where our method is not exact, there is a trade-off between bias and variance.
- We described a method for the generation of high-resolution clutter maps from sonar trial data, based on previous work by Prior [106], and Abraham and Prior [9]. This includes an iterative algorithm with low memory requirements, which processes chunks of data for consecutive regions of the sea in turn. We demonstrate that when a clutter map is based on K-distribution shape parameters, if data is available over multiple pings, it is advantageous to combine these parameter estimates by the harmonic method. This provides greater detail of small-scale features, compared to clutter maps based on median sampling or the arithmetic mean.

9.2 Limitations and Future Work

We will now examine the limitations of our research, many of which lead to suggestions for future work.

- In Section 5.2 we provided tables of estimated critical values for the K-distribution evaluated by the upper-tail Anderson-Darling test. These values were calculated by Monte Carlo simulation. Whilst these tables are sufficient to form the comparisons made in our work, it is acknowledged that their accuracy was limited by the Monte Carlo methodology. An analytic derivation of these critical values would be of considerable interest, though this may be difficult as it would still be necessary to consider the properties of the parameter estimator.
- When implementing the generic box-car detector, a serious draw-back is numerical precision. This was encountered in our simulations where, despite increases in probability of detection when using the K-distribution as a noise model, the K-distribution box-car detector failed to achieve the lowest false alarm rates reached by the standard box-car detector. This is likely to be due to rounding errors, particularly within the iterative calculations required to evaluate the Bessel K function when calculating the K-distribution CDF. We have used double-precision floating point, with an available range of approximately $[-10^{308}, 10^{308}]$. Of this range, our test statistic occupies values between $[0, 1]$, with low false alarm rates exclusively occupying values in the range $[0, 10^{-1}]$. Rounding errors and losses in precision may be avoided by extending the range of our test statistic, although this will reintroduce the need to compute threshold values for a chosen false alarm rate.
- Although the GLRT detectors in this thesis are inherently adaptive to local noise characteristics, thresholds in this thesis have been set globally to the detection statistics. It was noted in Section 7.3 that the threshold of the GLR statistic may still depend on the parameters of the noise model. This effect is believed to have caused the multi-modal distribution of the GLR statistics processed in Section 8.5, and may have reduced the effectiveness of the multi-ping GLRT detectors. It is hypothesised that an adaptive threshold, selected as a function of local noise parameters, would result in a further increase in detection performance, and remove the stepped profile of the multi-ping ROC curves seen in this thesis.

- In Section 8.3 we generated a clutter map using the harmonic mean to combine shape parameters. It would be interesting to compare estimates combined by the harmonic mean against in-situ measurements from sea trials. This would enable a more objective comparison of the harmonic mean and the methods applied by Abraham and Prior [6, 106]. It is expected that this will demonstrate increased accuracy in the noise estimate, and may help define guidelines for when estimates should and should not be combined.
- Throughout this thesis we made the assumption of independent, identically distributed samples, for both the noise and target distributions. This limitation became evident in Section 7.8 where we encountered the correlation present in real target data. This may have affected our goodness-of-fit testing, where we may have under-estimated the fit of the distributions by comparing correlated samples against critical values for independent samples. This may have also caused over-estimation of the theoretical detection performance for our extended and multi-ping GLRT detectors, as we simulated targets without correlation, which would be equivalent to a larger number of correlated samples. However, this does not negate our findings. Uncorrelated targets were applied consistently throughout our simulations, and so our comparison of various noise models and detection schemes remains valid. A better understanding of correlation in sonar noise and target signals would provide an additional input to signal processing algorithms, which is expected to improve categorisation, and therefore detection in sonar data.
- Throughout this thesis we assumed that noise samples selected by our algorithms are from a single distribution. Due to the presence of abrupt environmental features, the sonar environment can include discontinuities. Fortunately our trial data was collected in a region with no sharp changes in topography. It is acknowledged that gradual changes in the noise environment are likely to exist, but any serious difficulties arising from this would have been identified during our goodness-of-fit analysis in Chapter 5. However, to generalize our methods it would be interesting to

investigate the modelling of change-points in long-range sonar data, similar to the work undertaken by Fearnhead [45]. These change-points could be incorporated into detectors to mitigate the effects of sharp boundaries. Alternatively, if the local noise environment varies slowly, a mixture model or fuzzy representation may be more appropriate, such as the model proposed by Sun [130].

- In Chapters 7 and 8, we assumed that sonar targets could be modelled by a Swerling target embedded in sonar noise. While undertaking the work in Section 7.8 it became apparent that the Swerling models did not match the available data. This may have been an issue specific to this particular target, or may indicate that the Swerling models are inappropriate for targets recorded by long-range sonar system. More work is needed in the analysis of sonar targets and the development of the models already available. For example, current estimators for the Homodyned and Generalized K distributions are limited to the Swerling 0 case, and frequently exhibit high variance and high failure rates.
- Lastly, the goodness-of-fit tests in Chapter 5 and the performance analysis in Chapters 7 and 8 made use of trial data from a single dataset. It should therefore be noted that some of our findings may not apply to other long-range sonar datasets, particularly if a dataset is obtained from a different sonar system and location. Such differences were outside the scope of this project, but continued research in this field will help identify the findings which can be generalized, and those which require greater care.

Of these suggestions, it is felt that improvements in the modelling of target statistics is of greatest importance to future work in this area. Although improvements in sonar signal processing can be achieved by modelling sonar noise, a better understanding of the target signals probability density and autocorrelation would be of substantial benefit.

Appendix A

Simulated ROC Curves

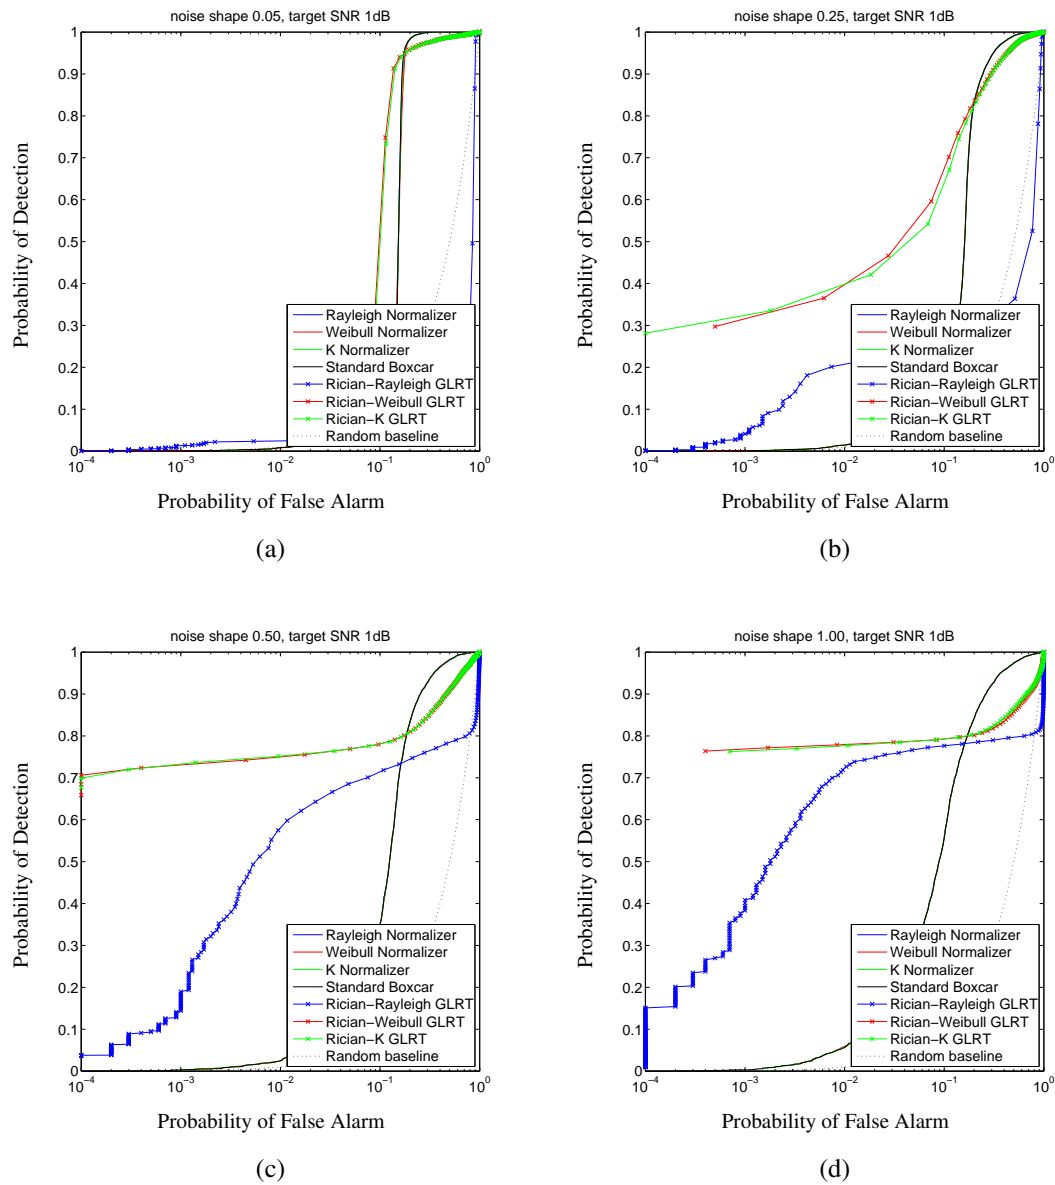


Figure A.1: Simulated ROC curves for 1dB targets and shape parameters 0.05 to 1.00, comparing normalization and single-point GLRT detectors

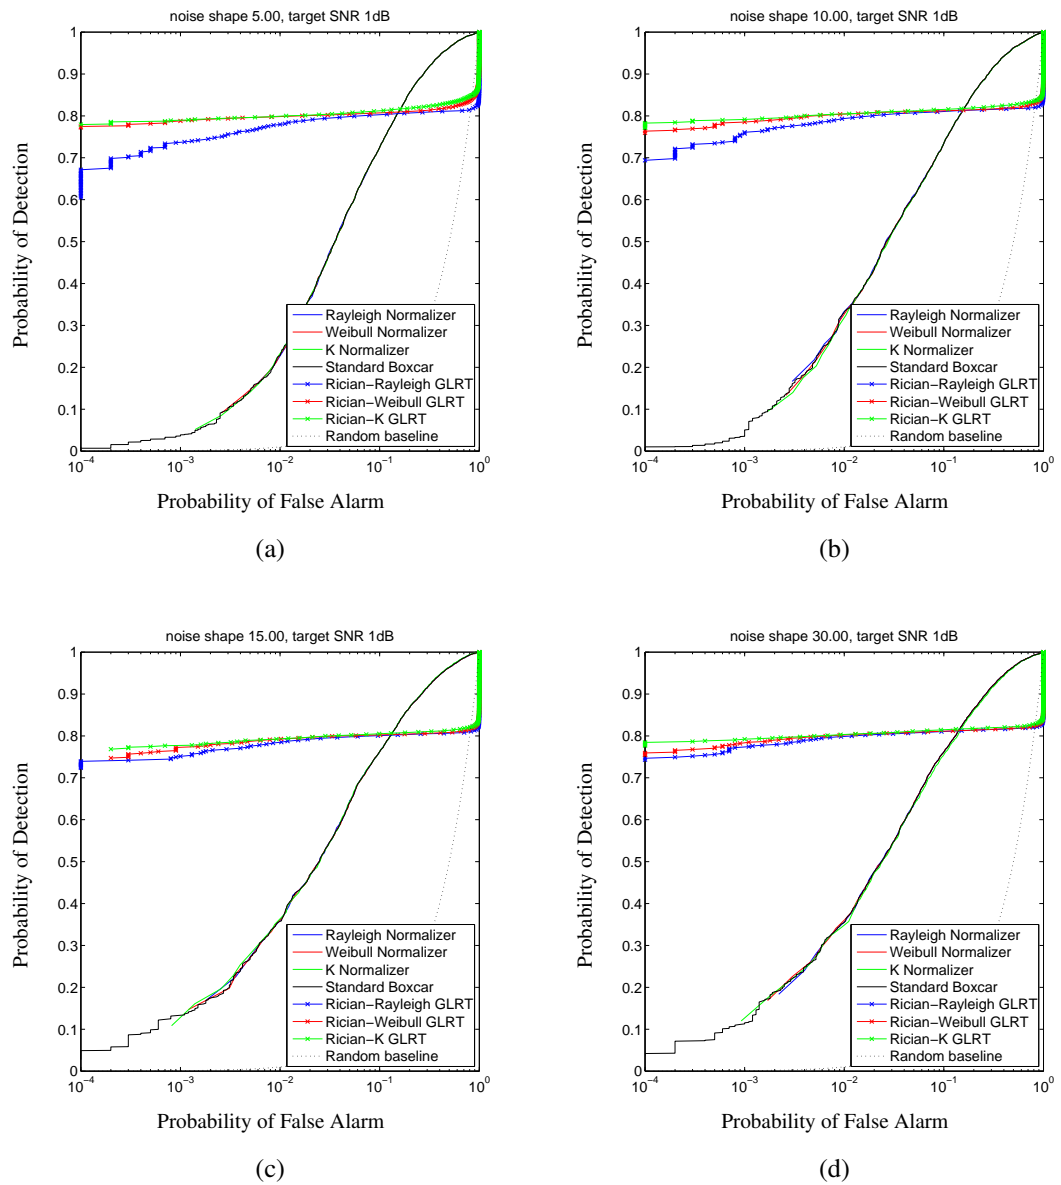


Figure A.2: Simulated ROC curves for 1dB targets and shape parameters 5 to 30, comparing normalization and single-point GLRT detectors

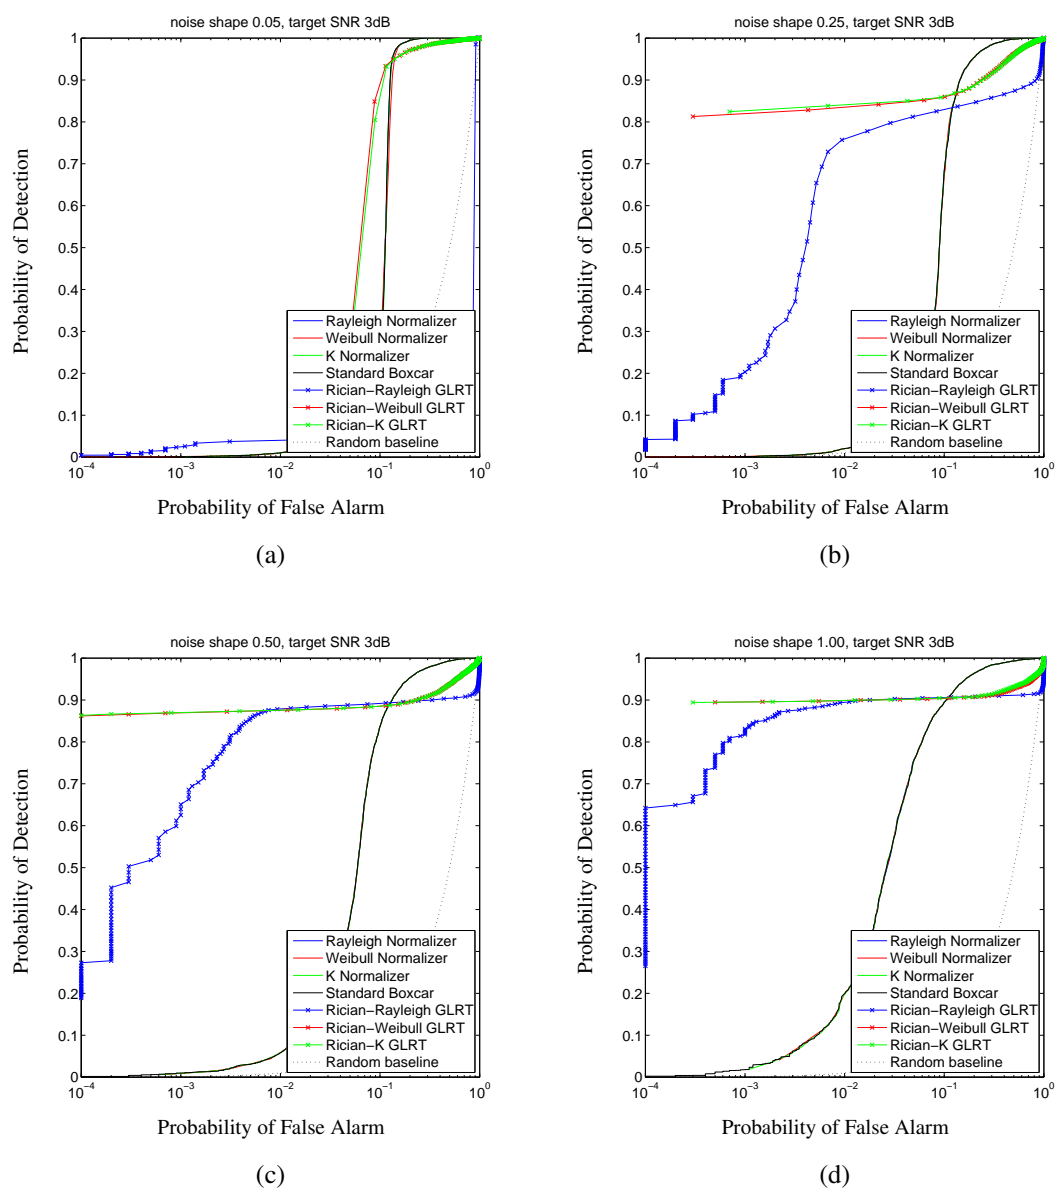


Figure A.3: Simulated ROC curves for 3dB targets and shape parameters 0.05 to 1.00, comparing normalization and single-point GLRT detectors

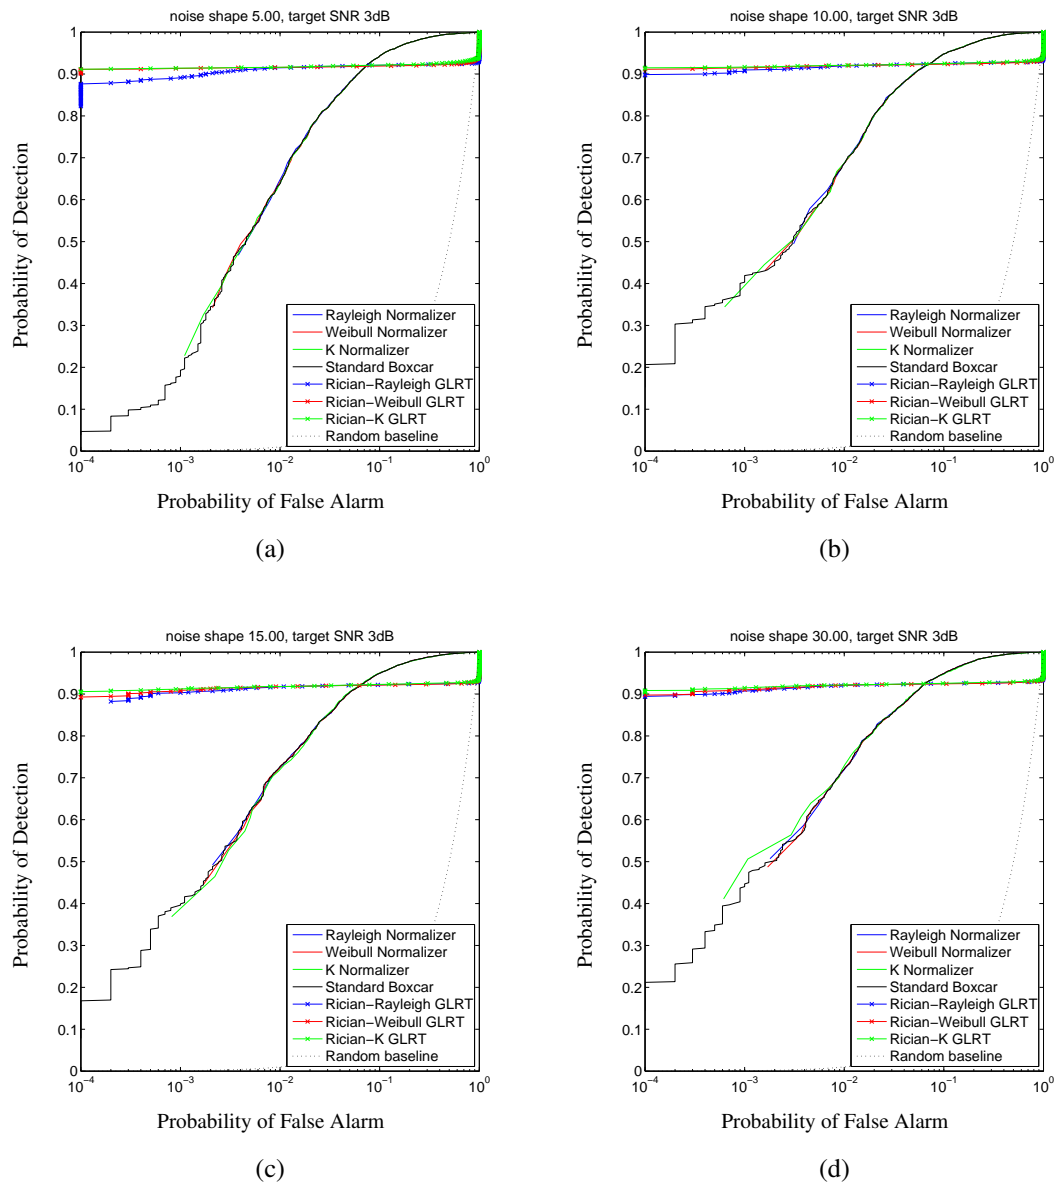


Figure A.4: Simulated ROC curves for 3dB targets and shape parameters 5 to 30, comparing normalization and single-point GLRT detectors

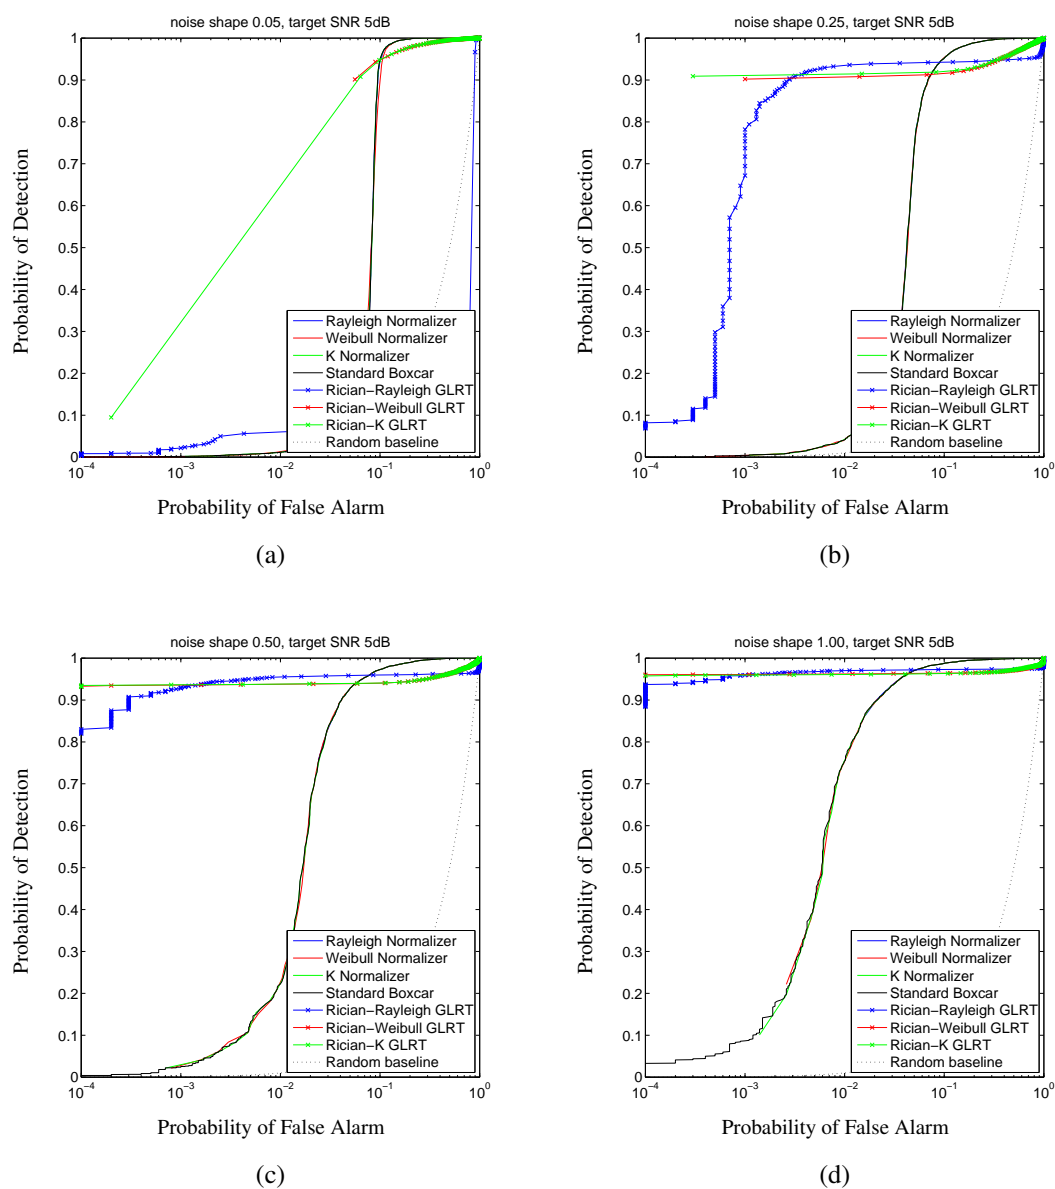


Figure A.5: Simulated ROC curves for 5dB targets and shape parameters 0.05 to 1.00, comparing normalization and single-point GLRT detectors

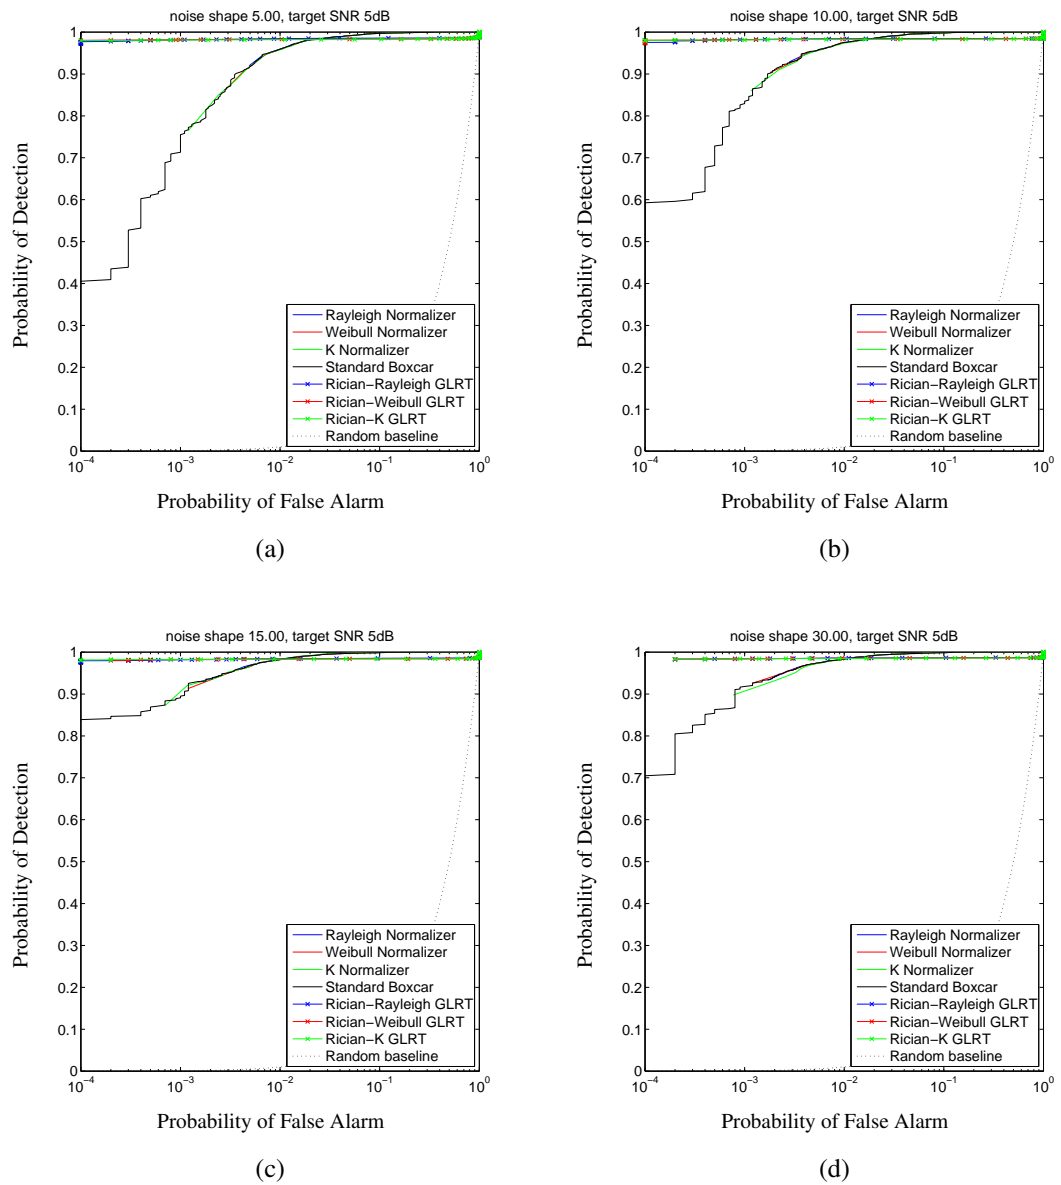


Figure A.6: Simulated ROC curves for 5dB targets and shape parameters 5 to 30, comparing normalization and single-point GLRT detectors

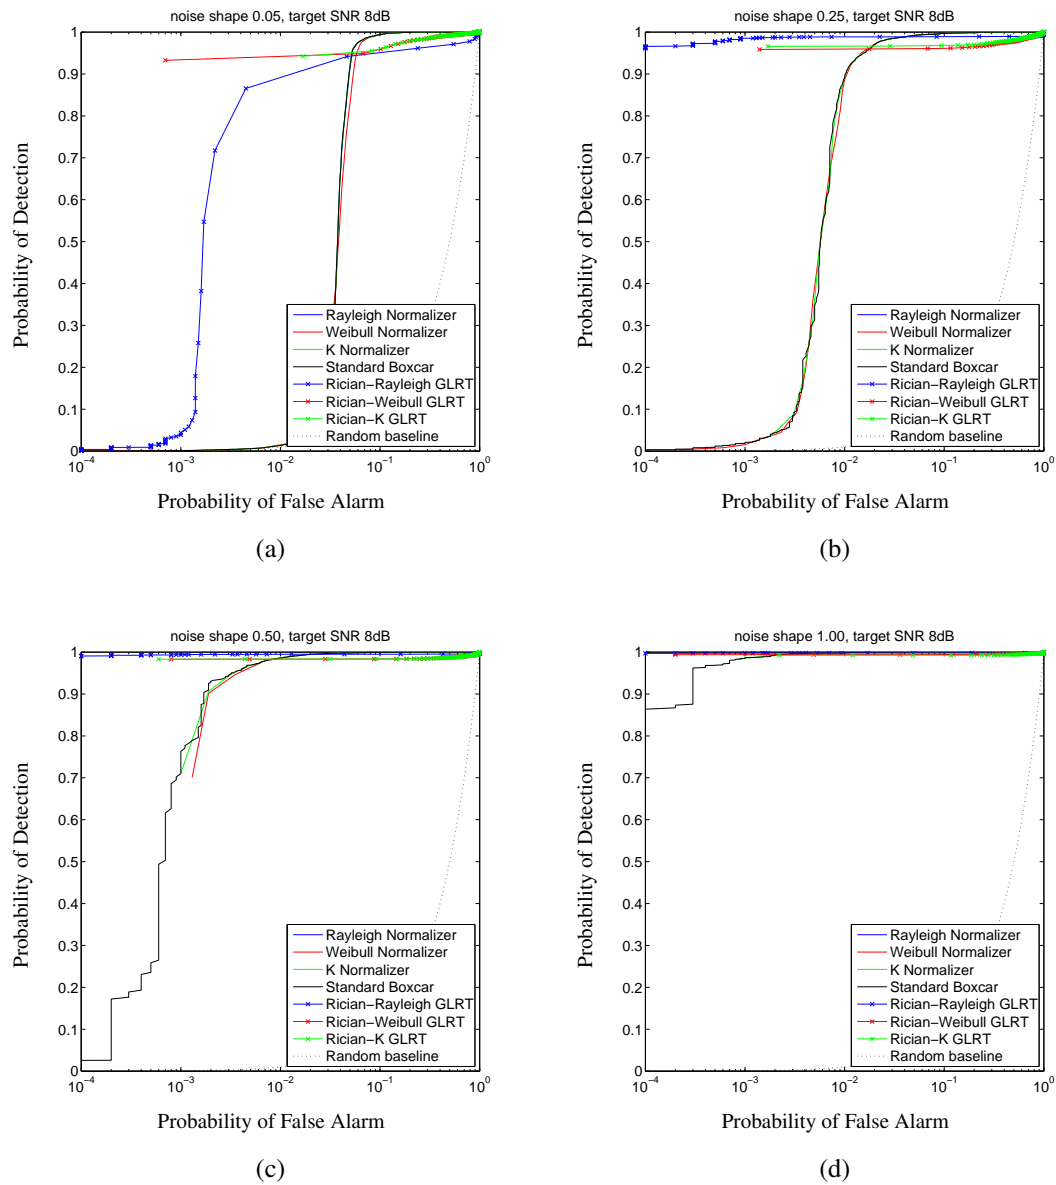


Figure A.7: Simulated ROC curves for 8dB targets and shape parameters 0.05 to 1.00, comparing normalization and single-point GLRT detectors

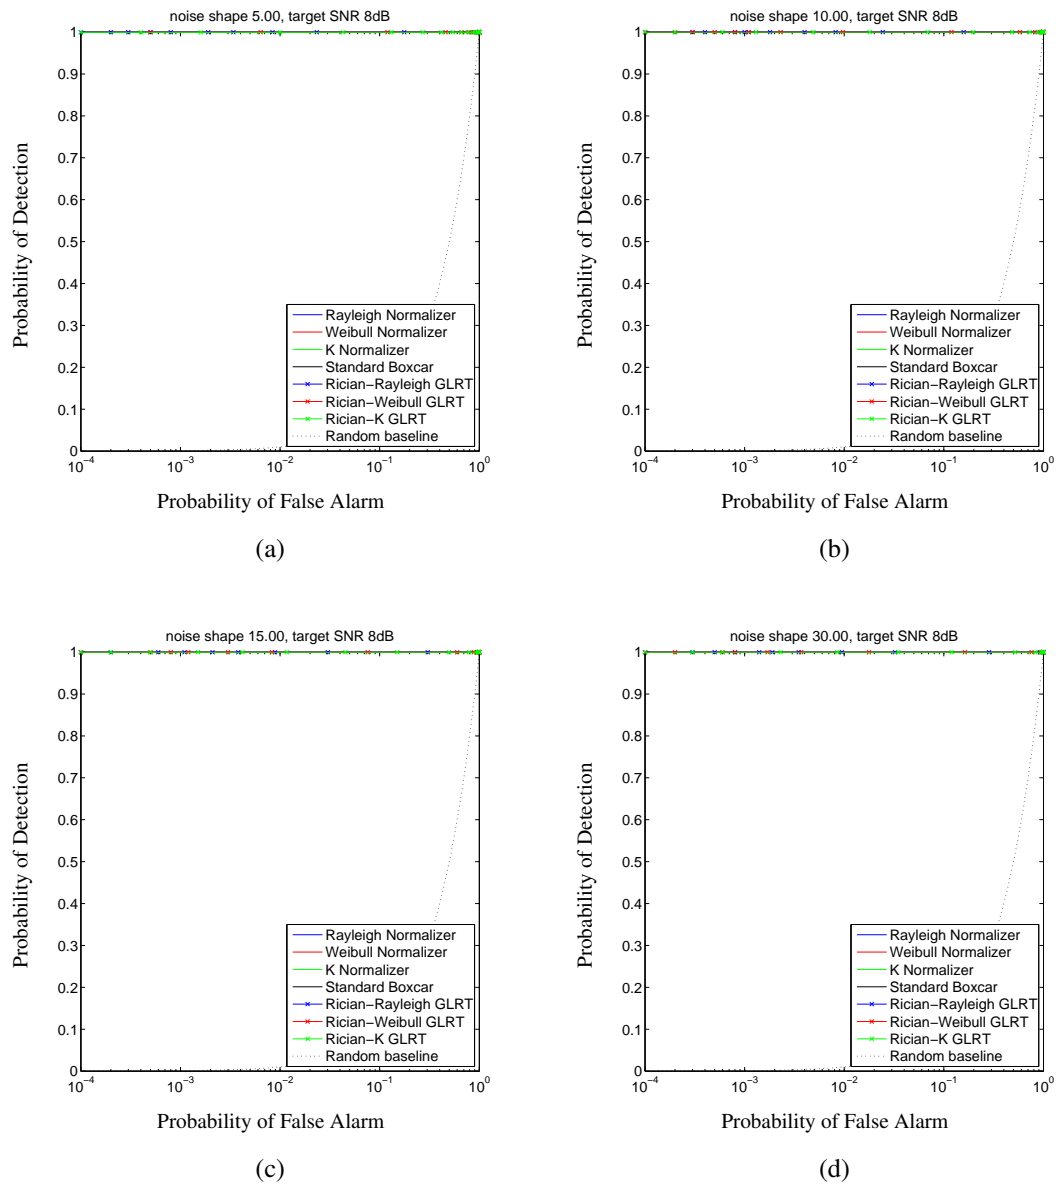


Figure A.8: Simulated ROC curves for 8dB targets and shape parameters 5 to 30, comparing normalization and single-point GLRT detectors

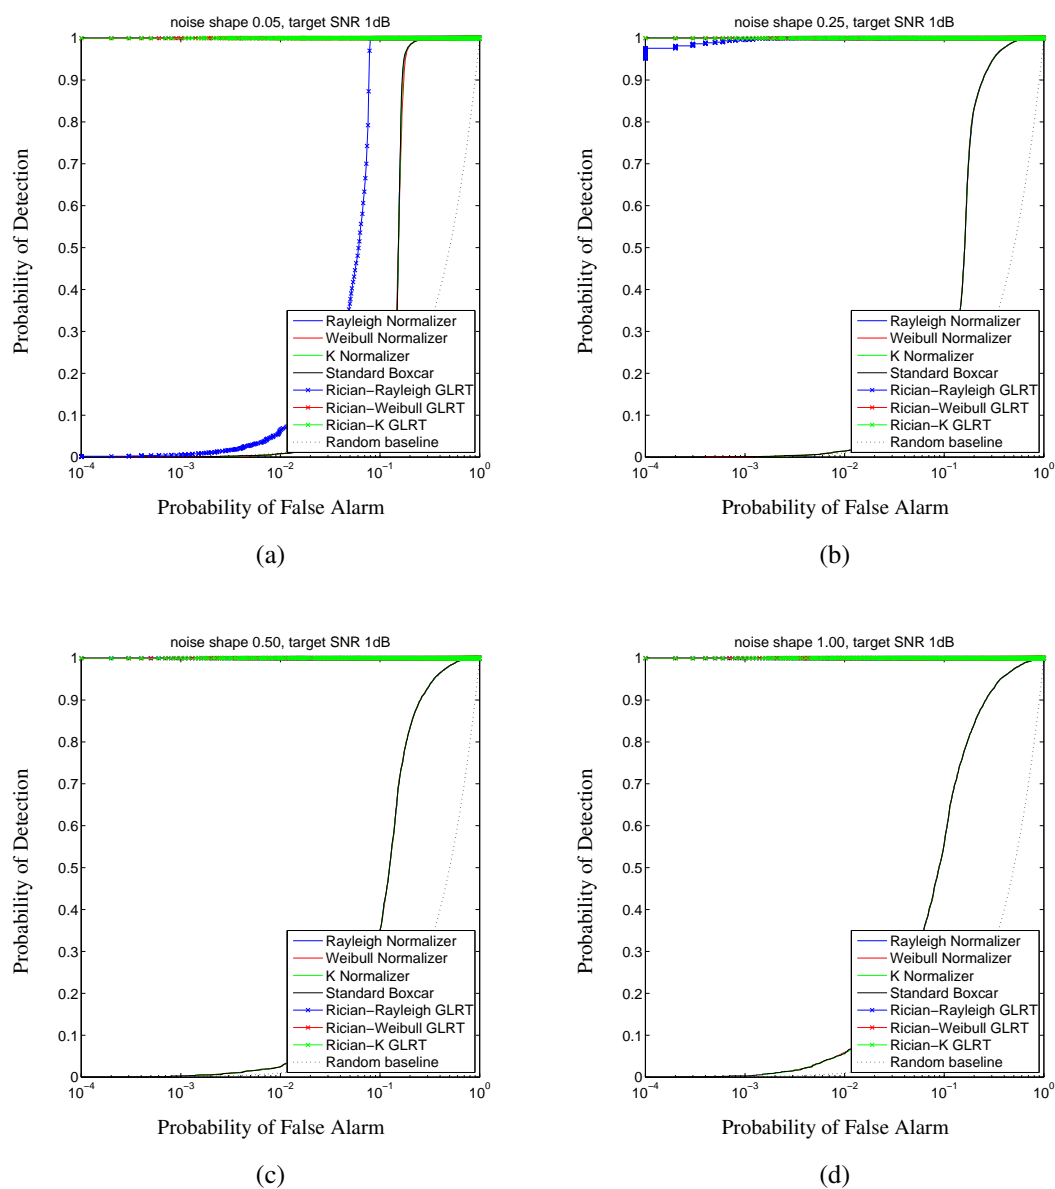


Figure A.9: Simulated ROC curves for 1dB targets and shape parameters 0.05 to 1.00, comparing normalization and combined-probability GLRT detectors

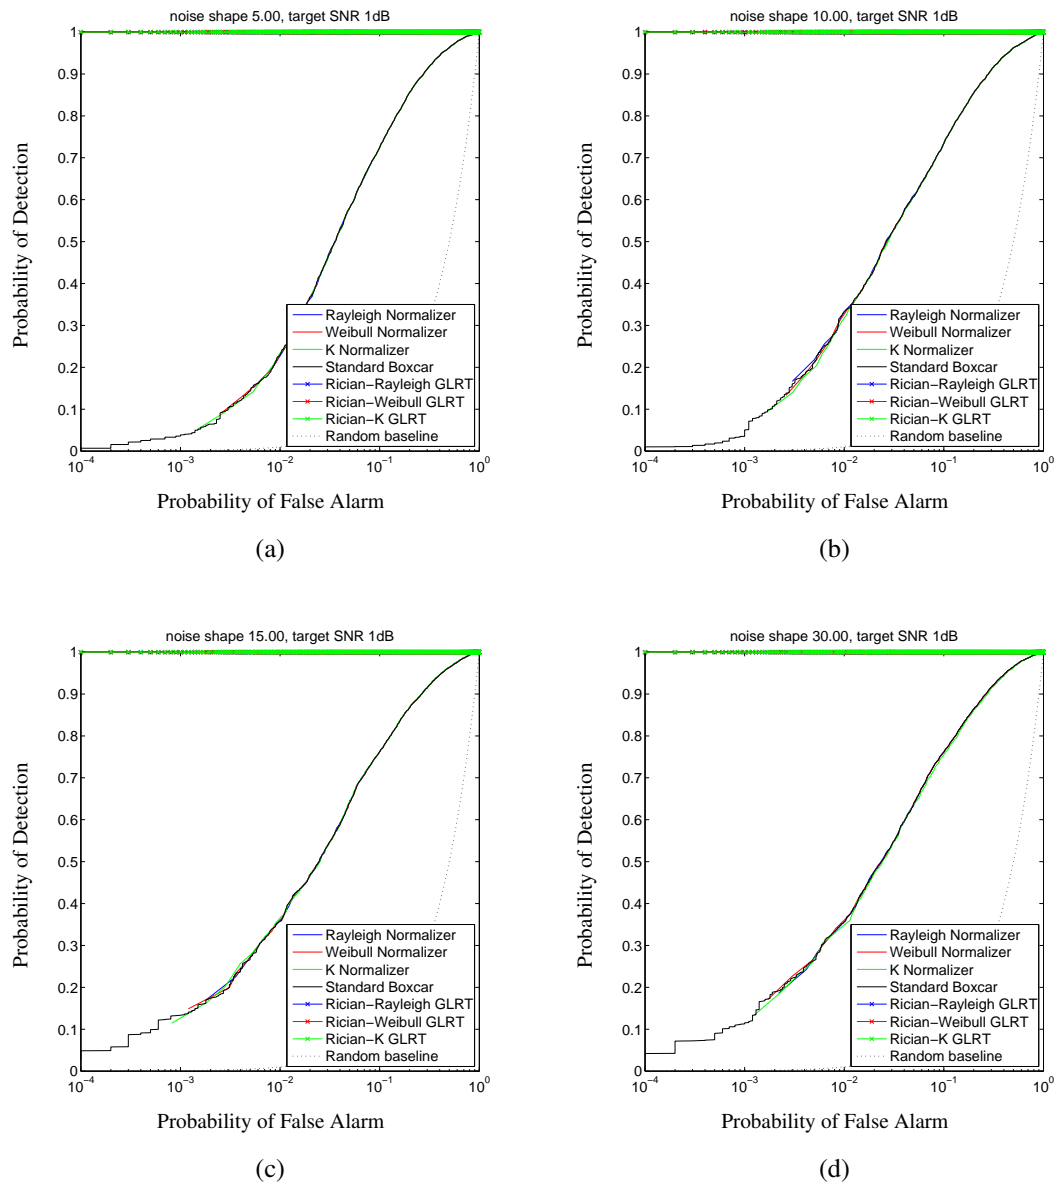


Figure A.10: Simulated ROC curves for 1dB targets and shape parameters 5 to 30, comparing normalization and combined-probability GLRT detectors

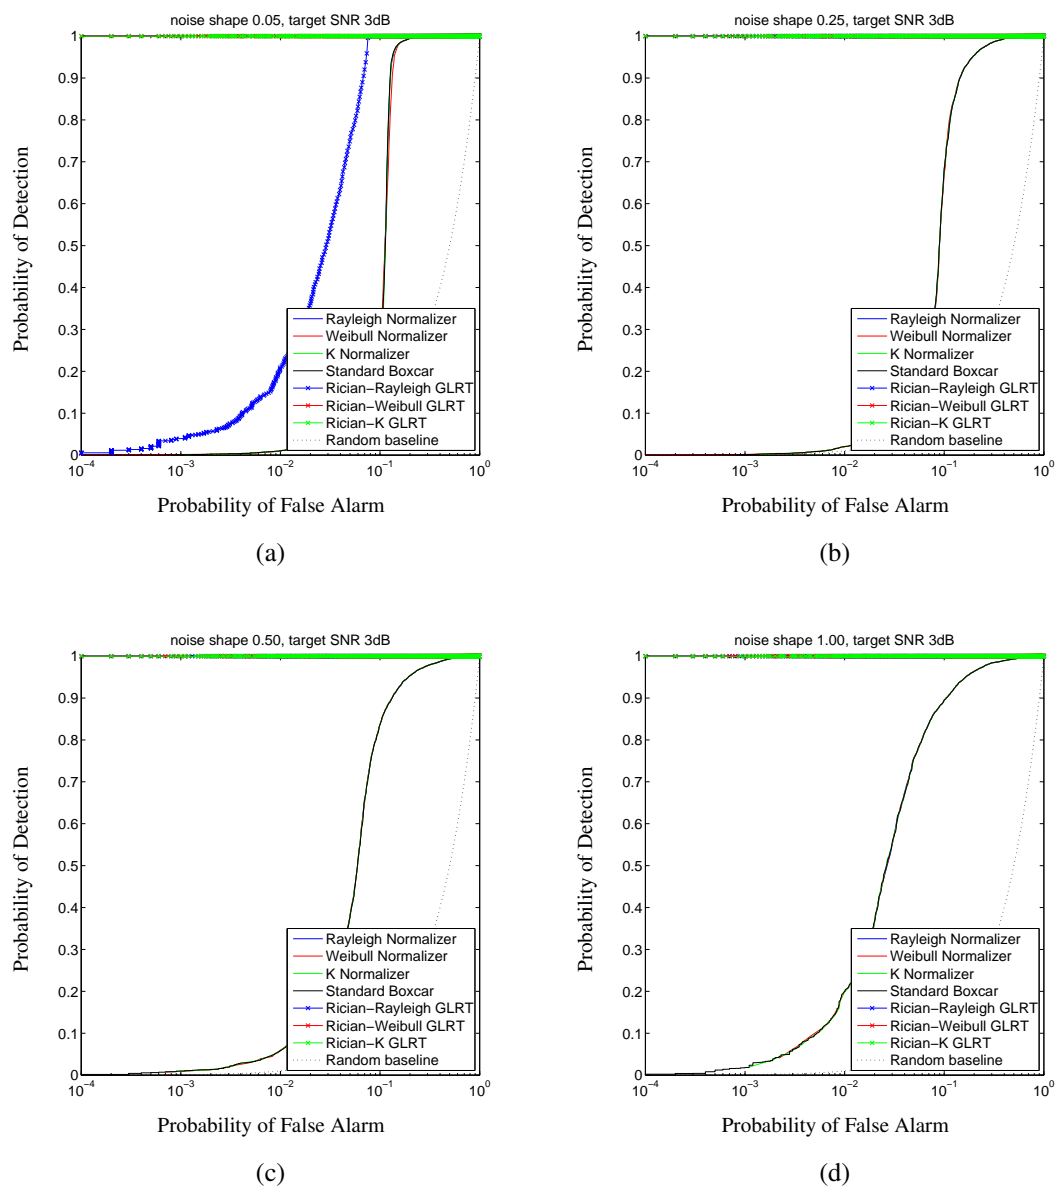


Figure A.11: Simulated ROC curves for 3dB targets and shape parameters 0.05 to 1.00, comparing normalization and combined-probability GLRT detectors

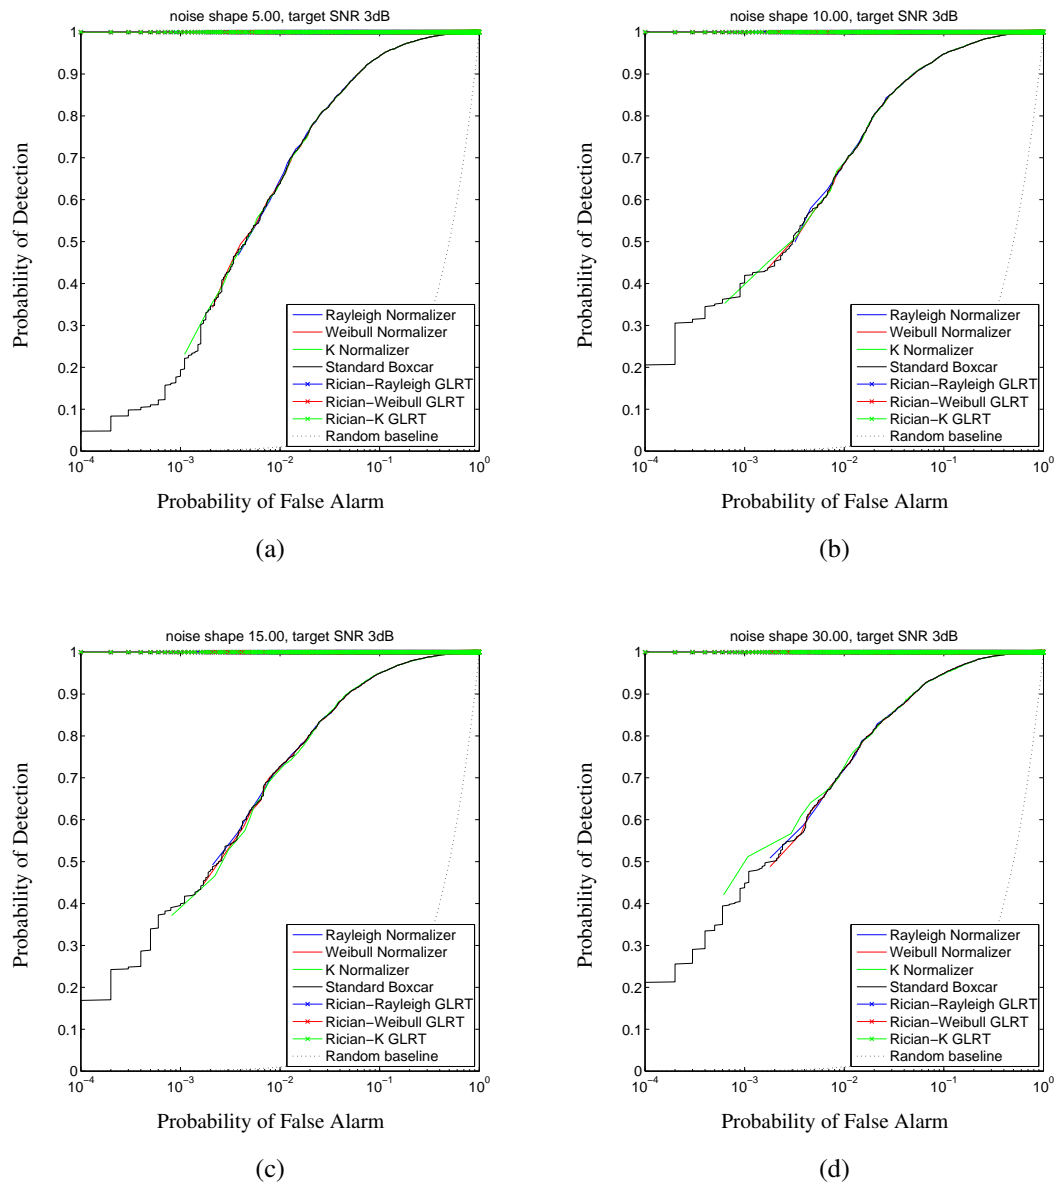


Figure A.12: Simulated ROC curves for 3dB targets and shape parameters 5 to 30, comparing normalization and combined-probability GLRT detectors

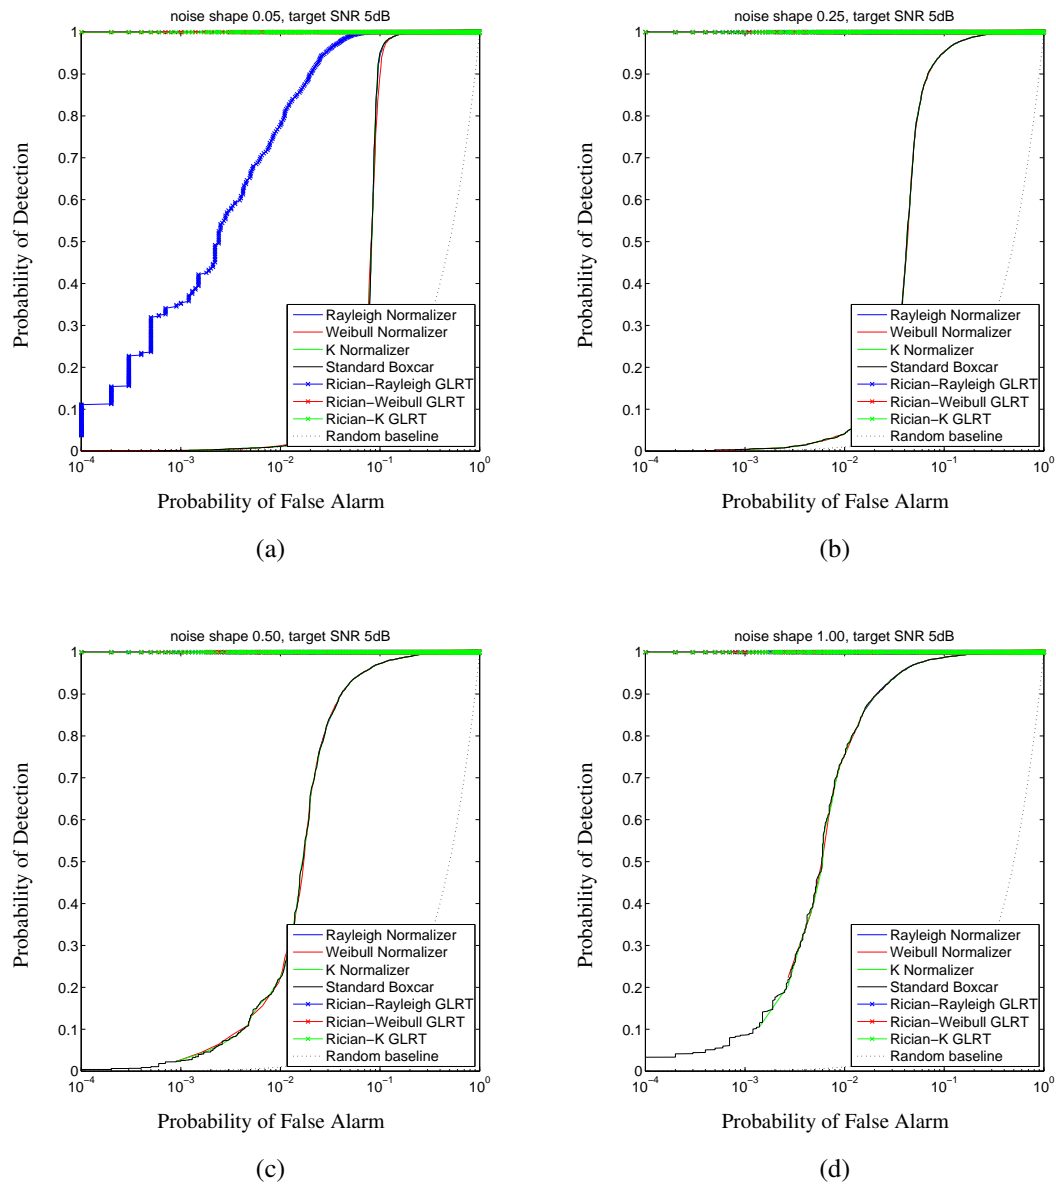


Figure A.13: Simulated ROC curves for 5dB targets and shape parameters 0.05 to 1.00, comparing normalization and combined-probability GLRT detectors

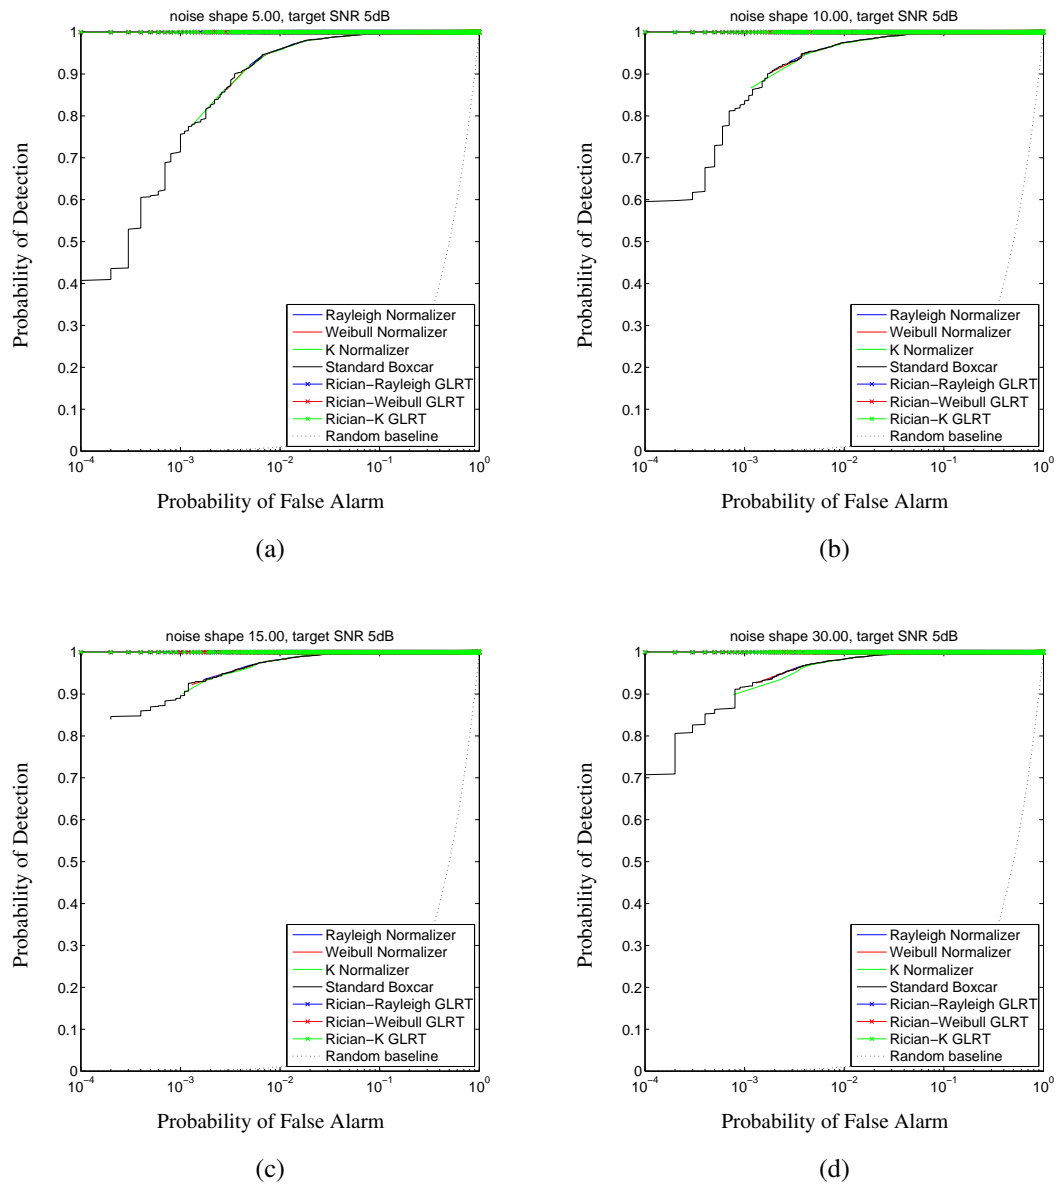


Figure A.14: Simulated ROC curves for 5dB targets and shape parameters 5 to 30, comparing normalization and combined-probability GLRT detectors

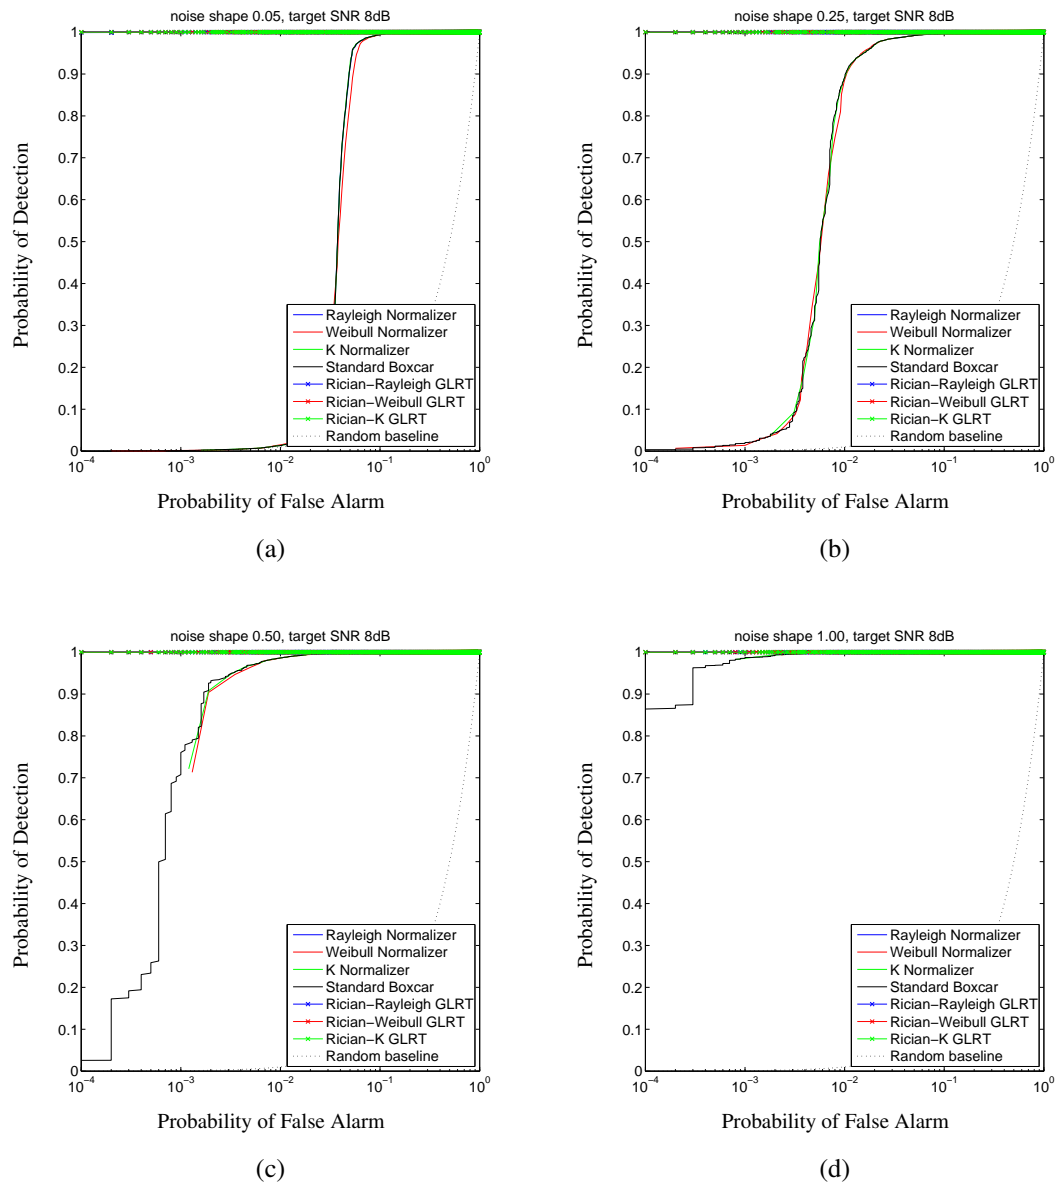


Figure A.15: Simulated ROC curves for 8dB targets and shape parameters 0.05 to 1.00, comparing normalization and combined-probability GLRT detectors

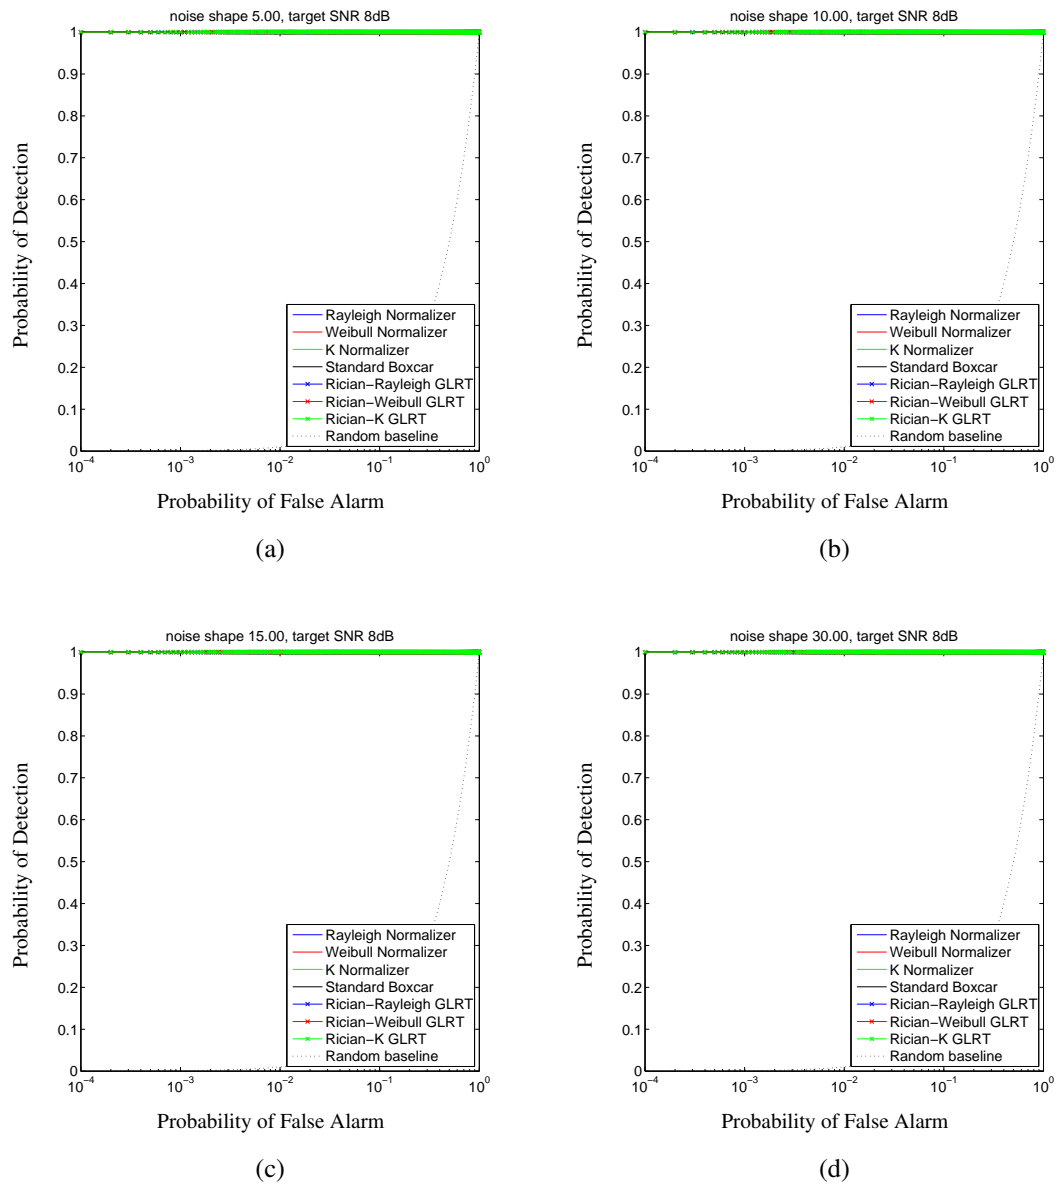


Figure A.16: Simulated ROC curves for 8dB targets and shape parameters 5 to 30, comparing normalization and combined-probability GLRT detectors

Appendix B

ROC Curves for Real Noise

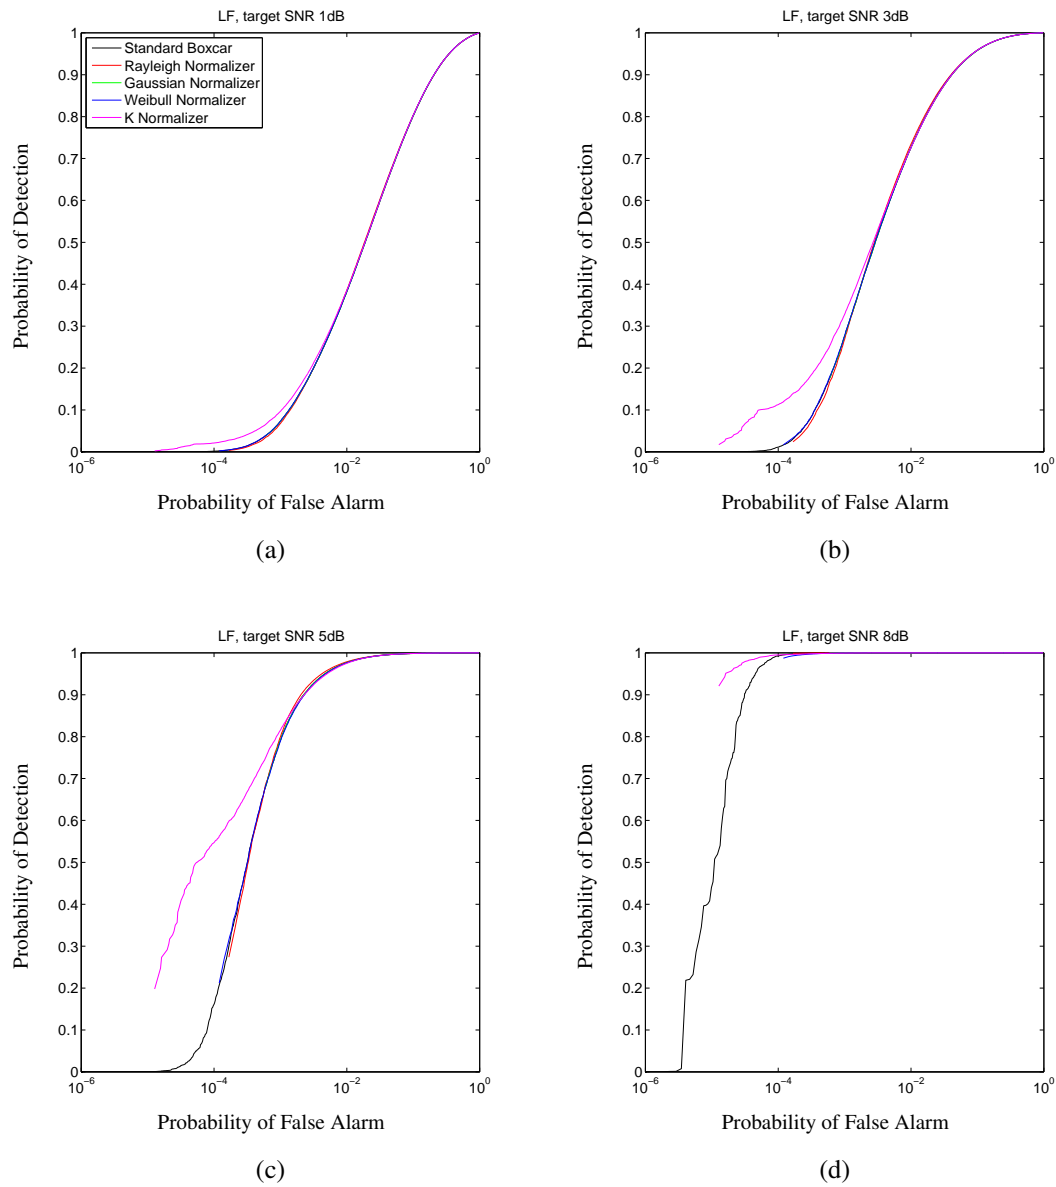


Figure B.1: ROC curves for various target strengths in the low frequency band against Swerling type 0 simulated targets. Processed by the standard boxcar and normalisation detectors

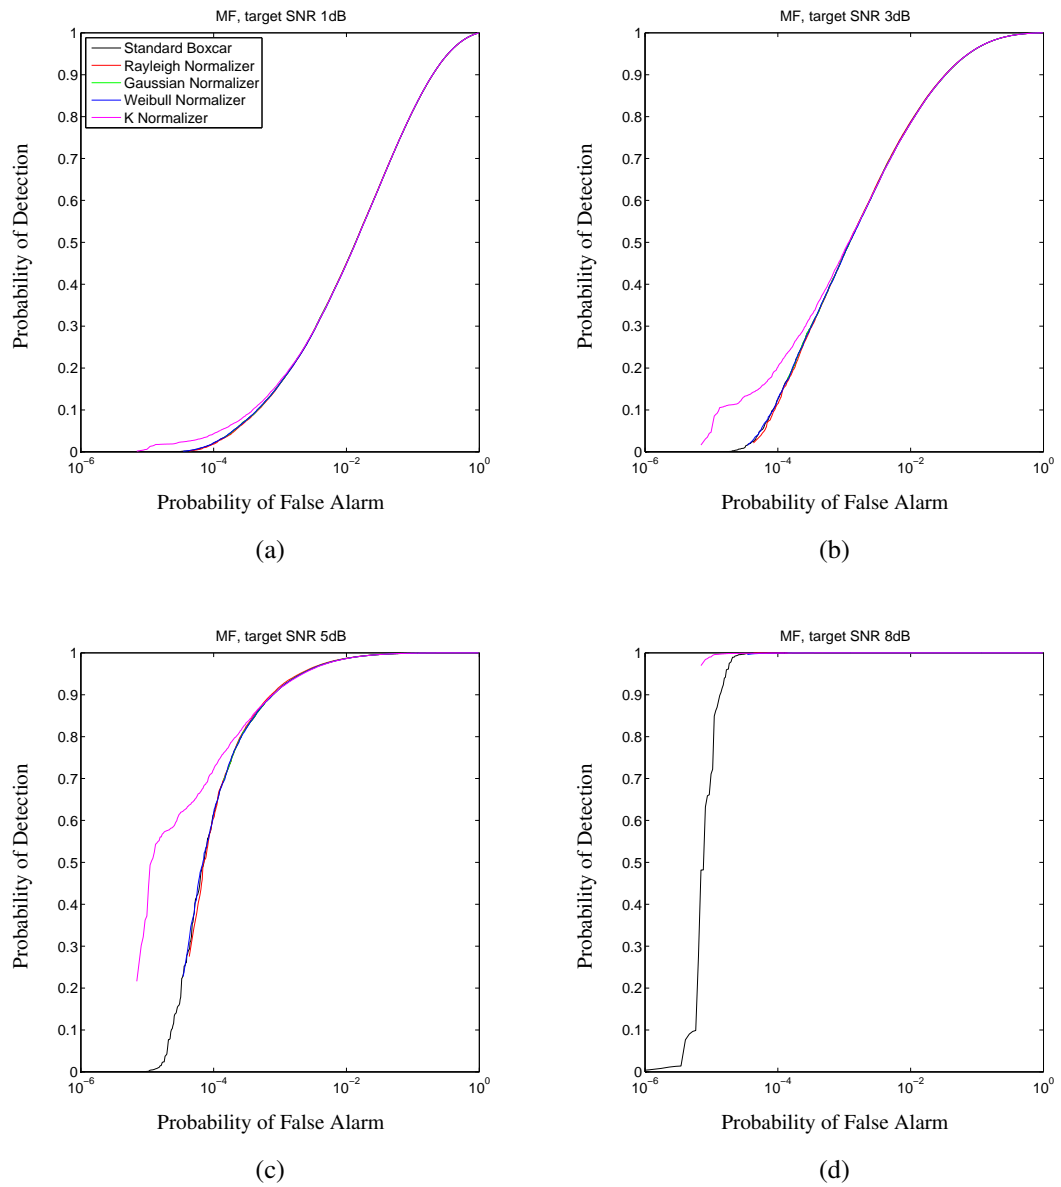


Figure B.2: ROC curves for various target strengths in the medium frequency band against Swerling type 0 simulated targets. Processed by the standard boxcar and normalisation detectors

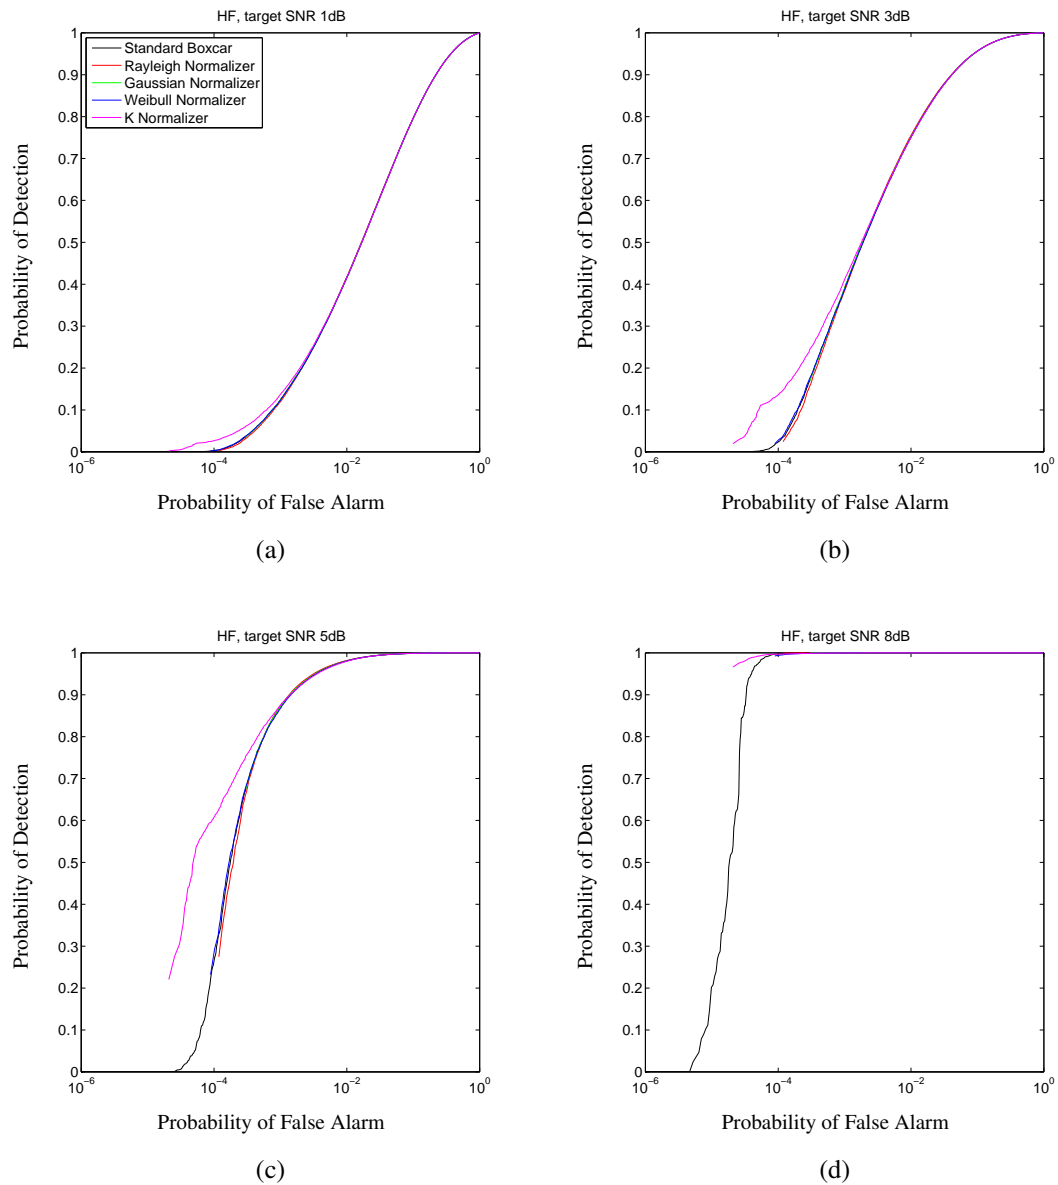


Figure B.3: ROC curves for various target strengths in the high frequency band against Swerling type 0 simulated targets. Processed by the standard boxcar and normalisation detectors

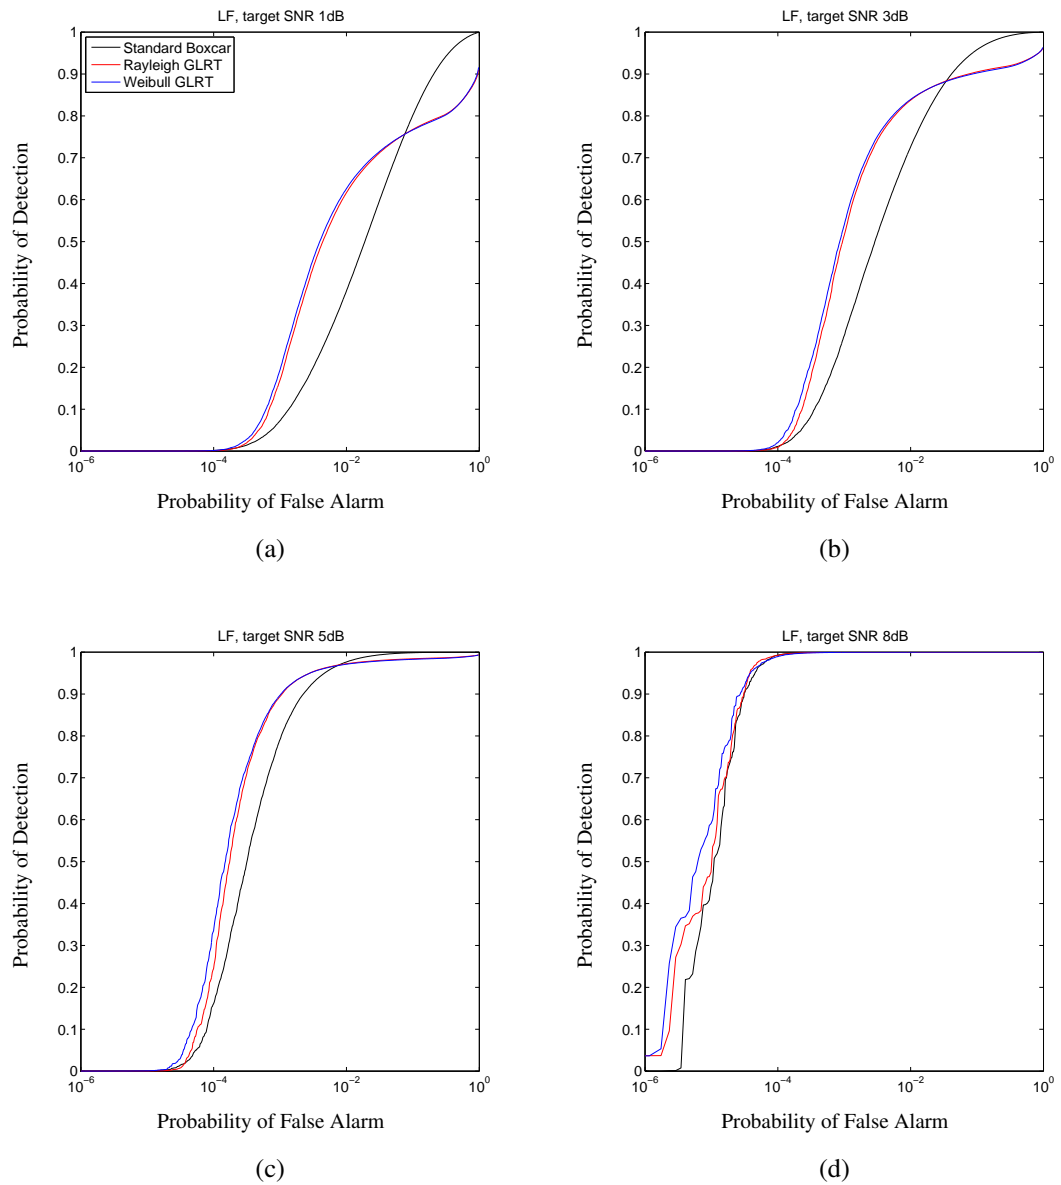


Figure B.4: ROC curves for various target strengths in the low frequency band against Swerling type 0 simulated targets. Processed by the single-sample GLRT detectors

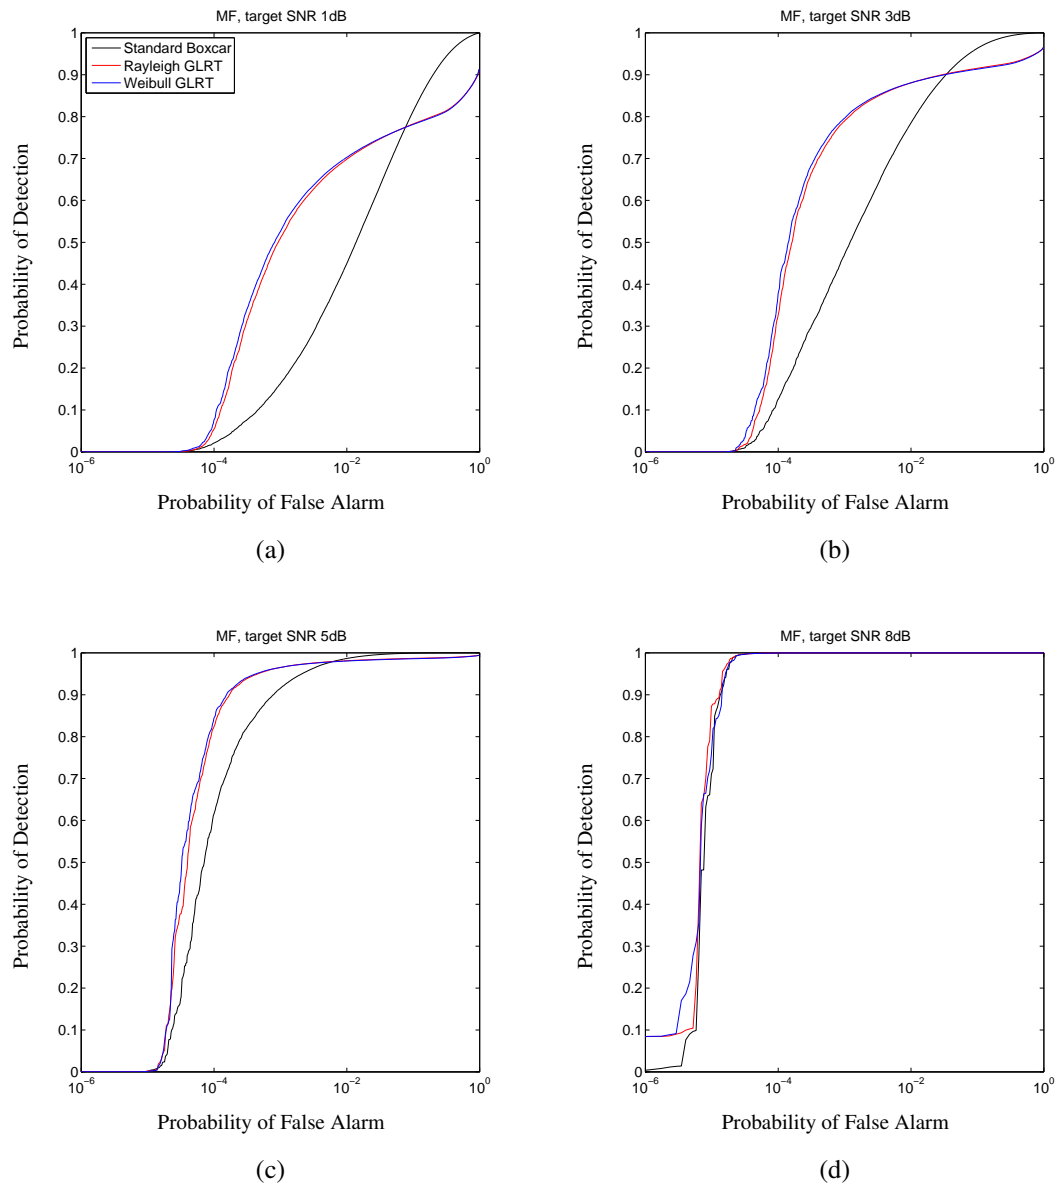


Figure B.5: ROC curves for various target strengths in the medium frequency band against Swerling type 0 simulated targets. Processed by the single-sample GLRT detectors

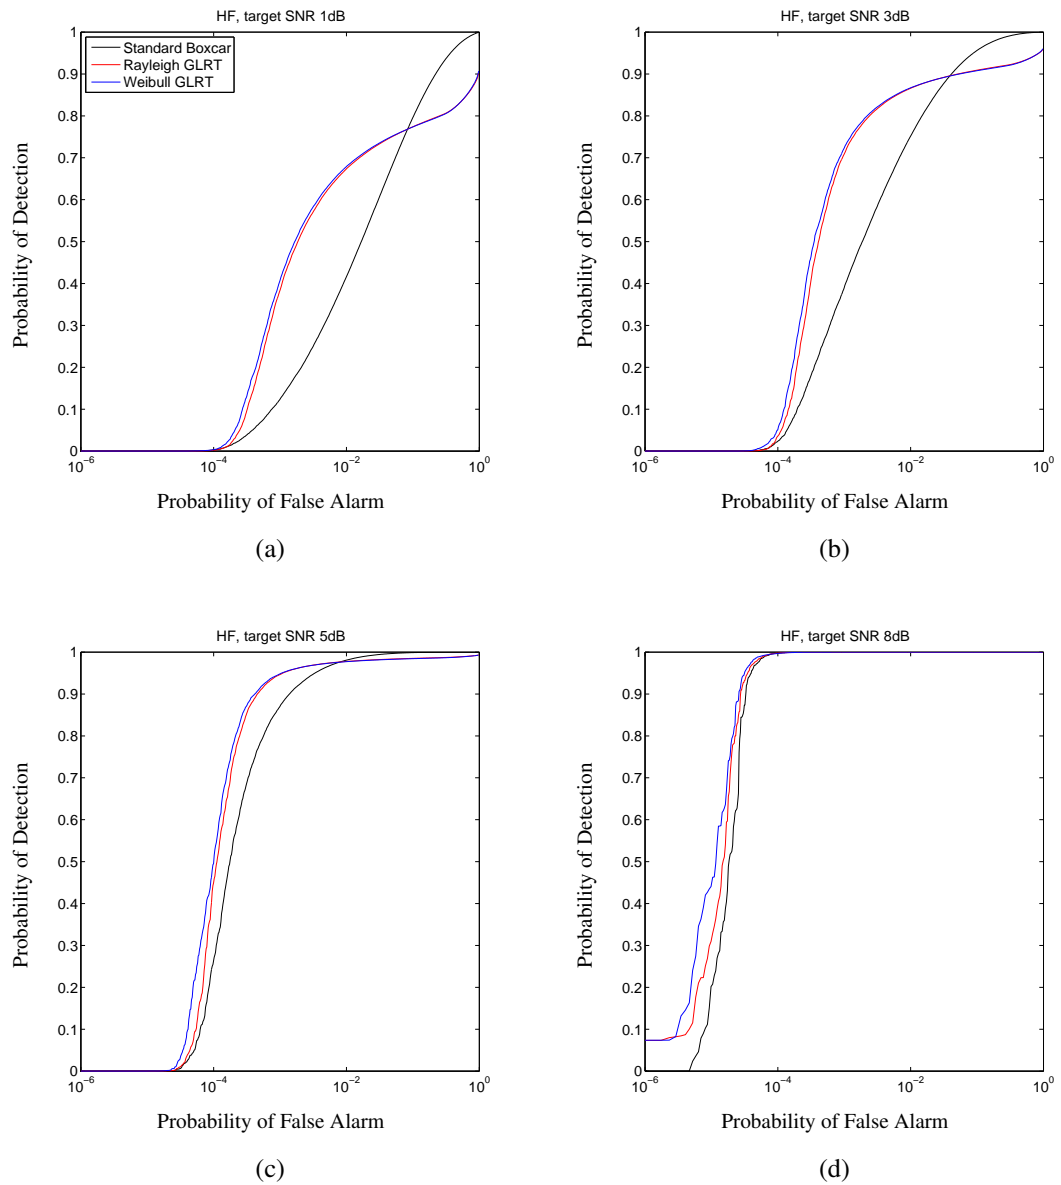


Figure B.6: ROC curves for various target strengths in the high frequency band against Swerling type 0 simulated targets. Processed by the single-sample GLRT detectors

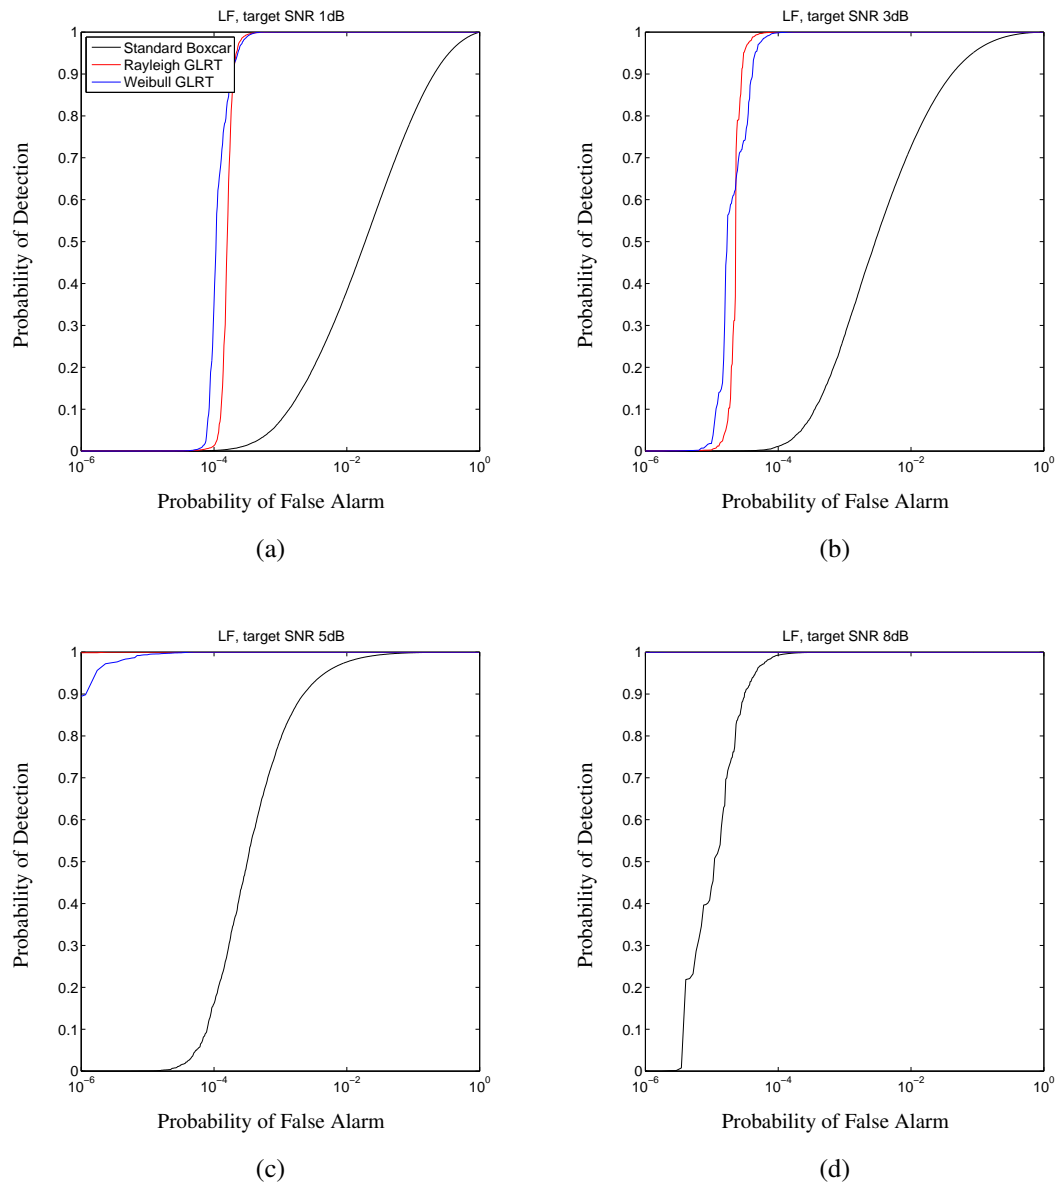


Figure B.7: ROC curves for various target strengths in the low frequency band against Swerling type 0 simulated targets. Processed by the multi-sample GLRT detectors

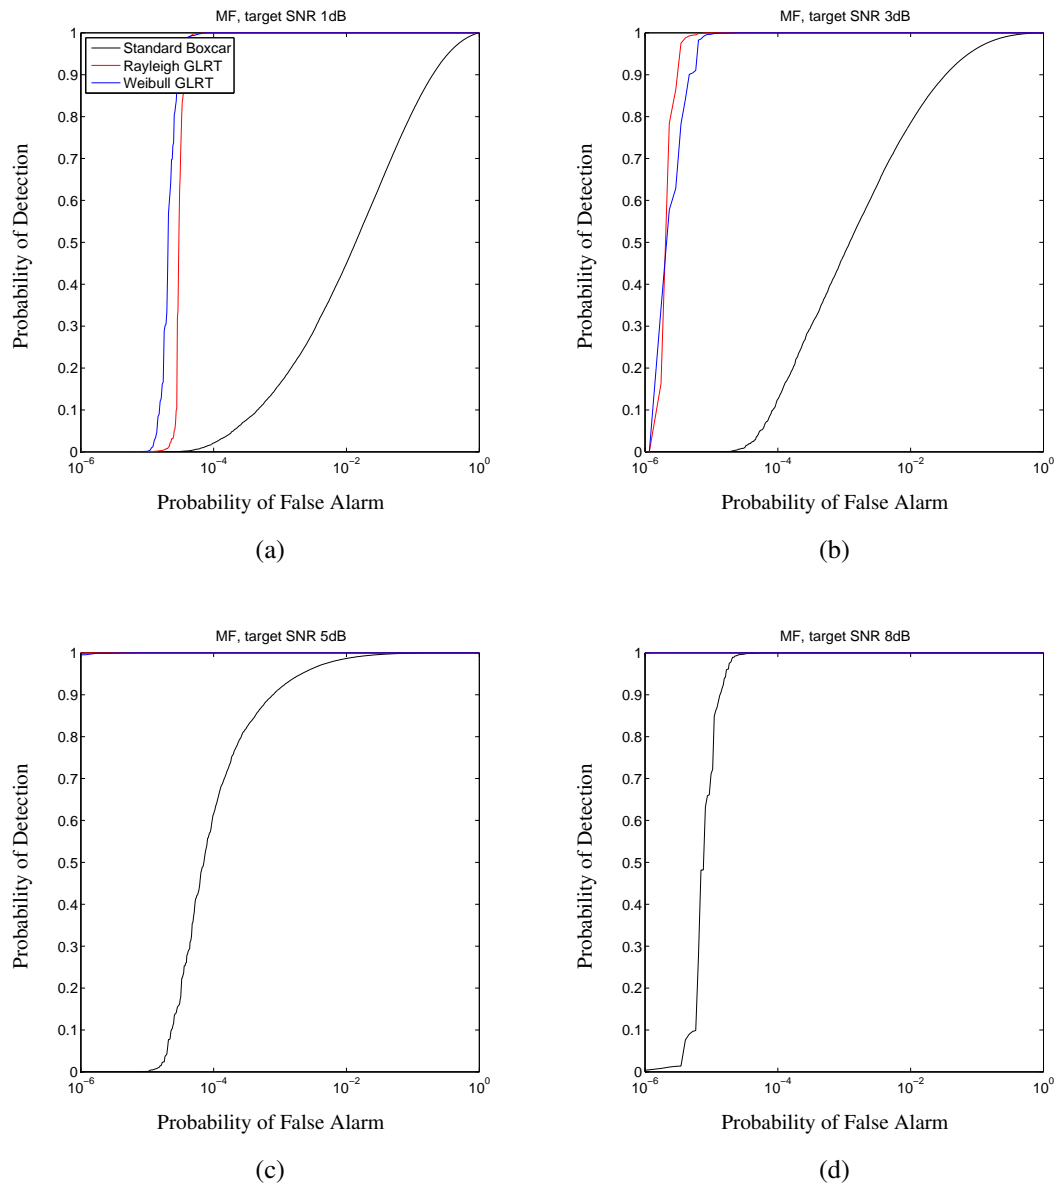


Figure B.8: ROC curves for various target strengths in the medium frequency band against Swerling type 0 simulated targets. Processed by the multi-sample GLRT detectors

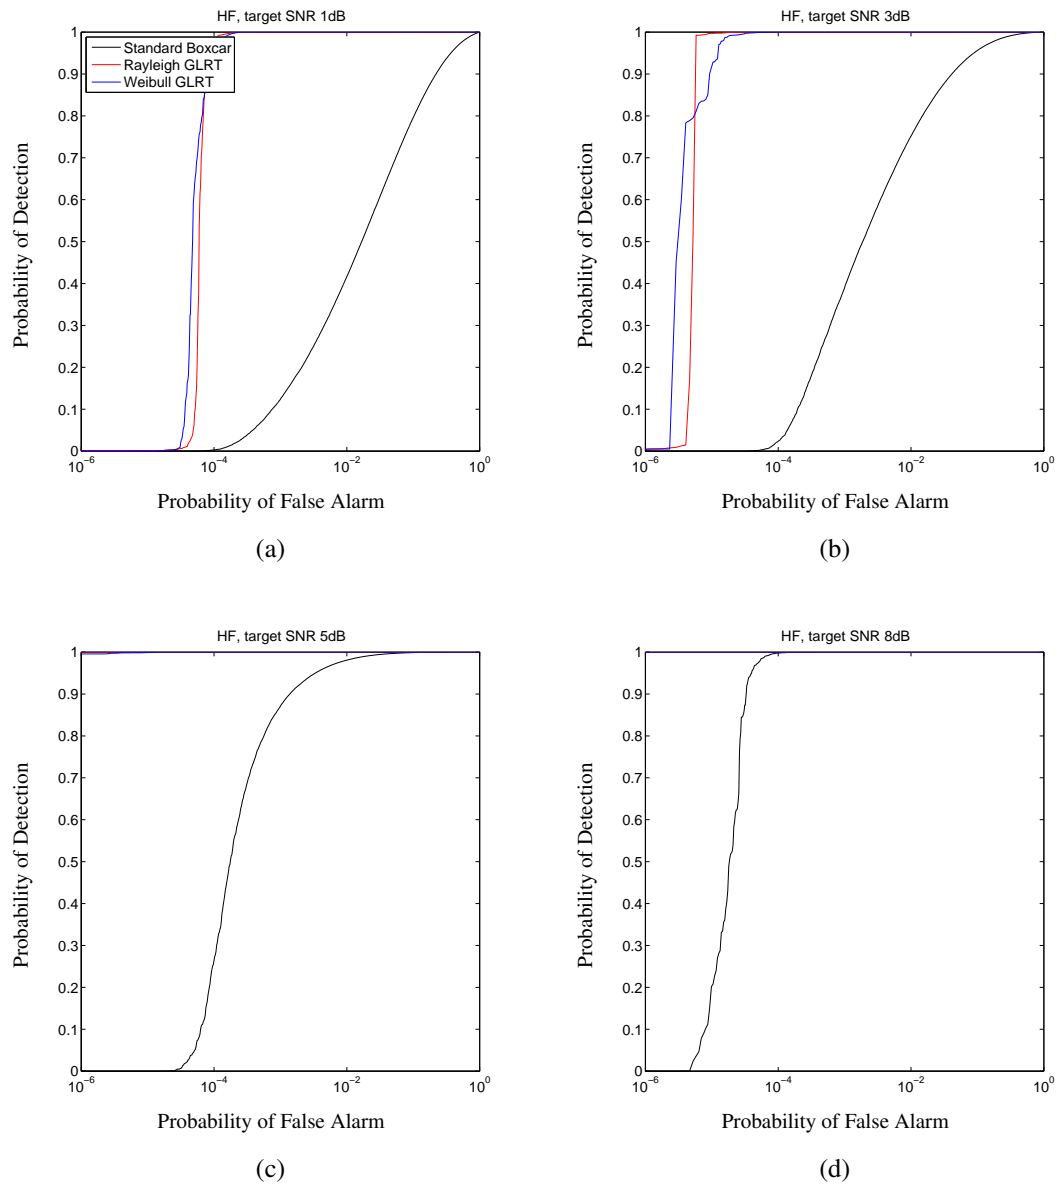


Figure B.9: ROC curves for various target strengths in the high frequency band against Swerling type 0 simulated targets. Processed by the multi-sample GLRT detectors

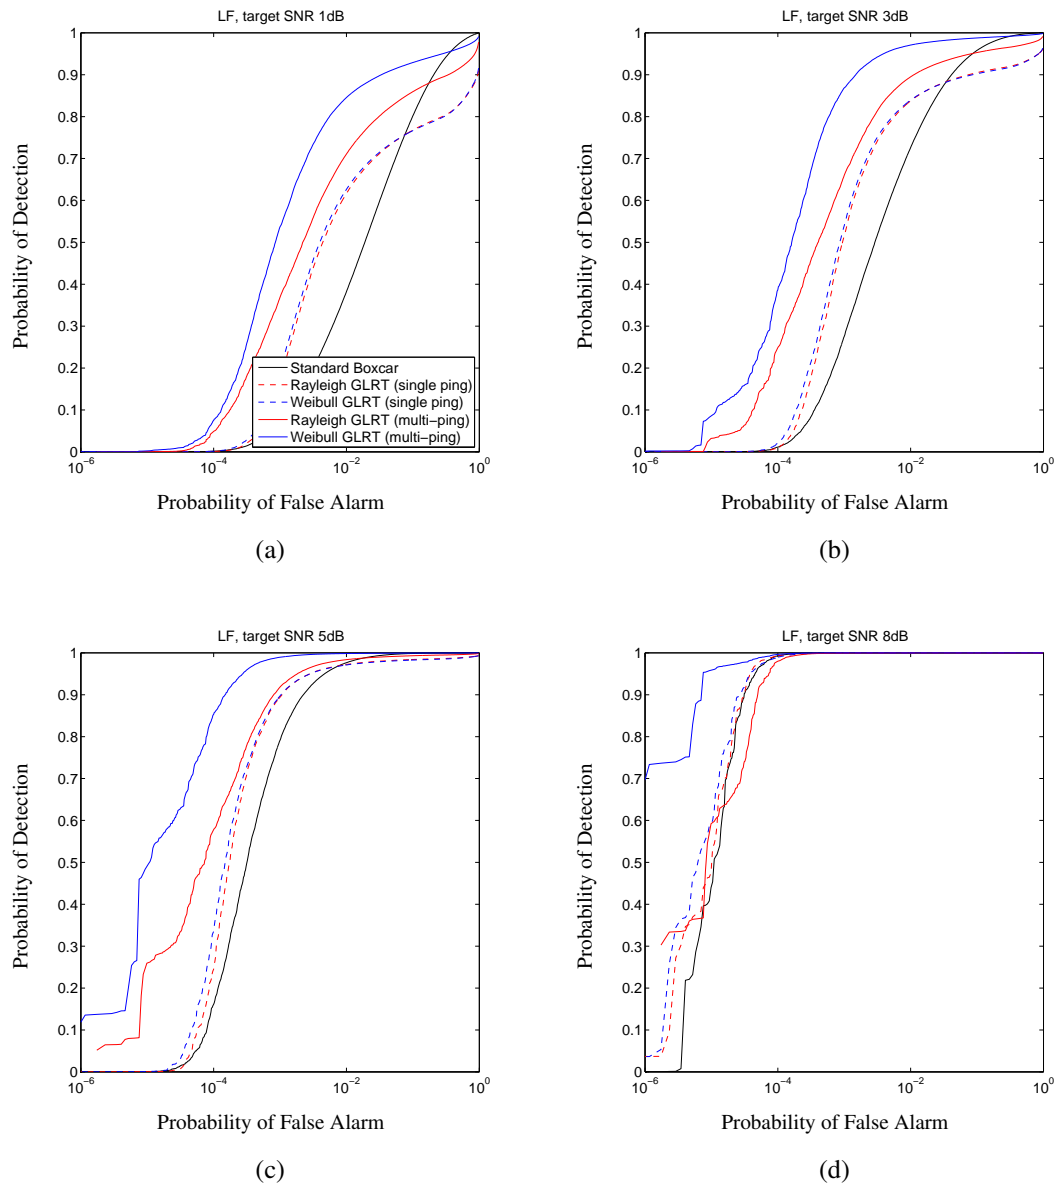


Figure B.10: ROC curves for various target strengths in the low frequency band against Swerling type 0 simulated targets. Processed by the multi-ping single-sample GLRT detectors

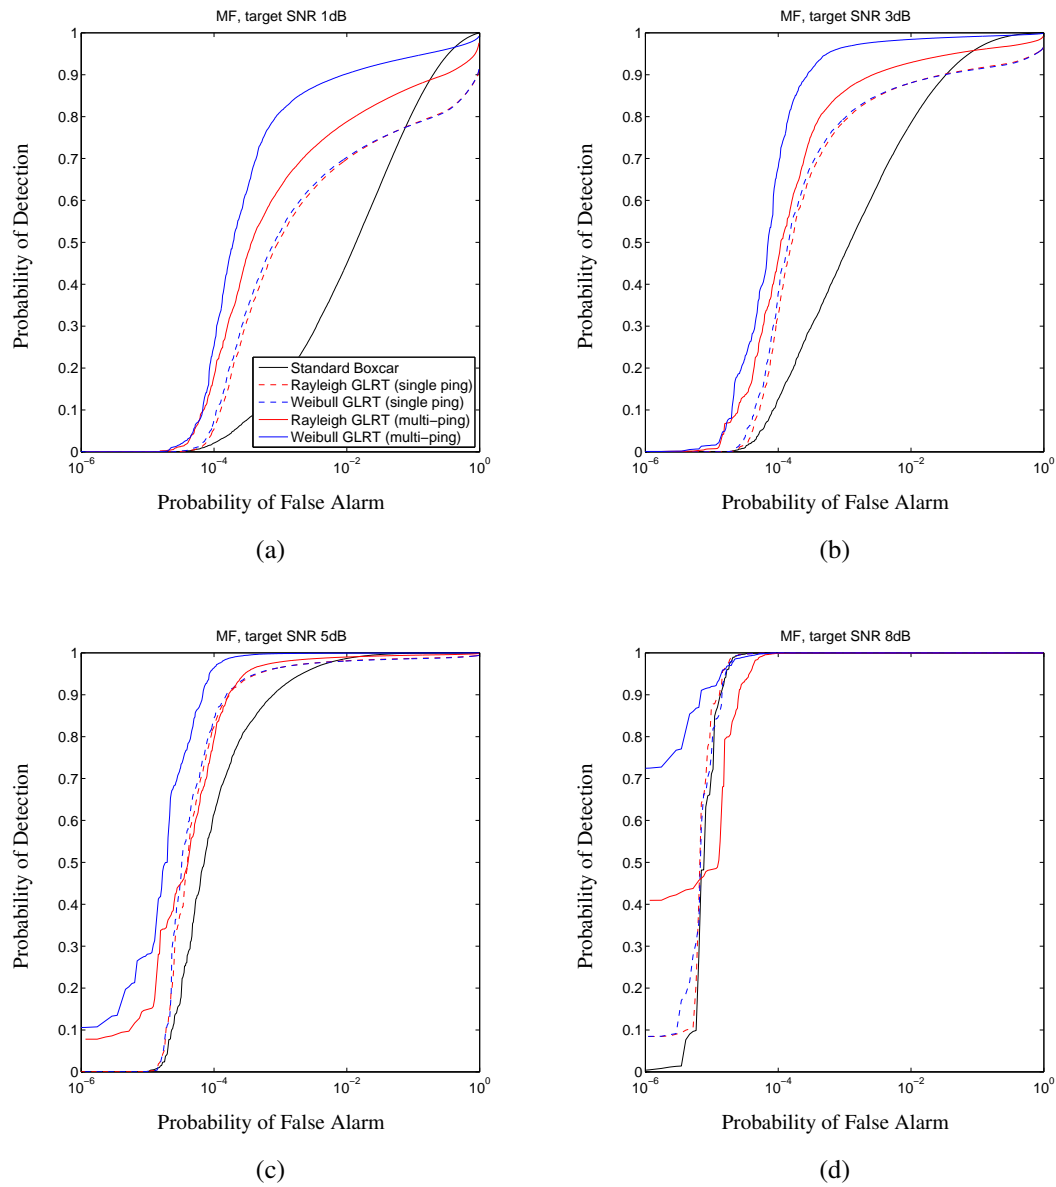


Figure B.11: ROC curves for various target strengths in the medium frequency band against Swerling type 0 simulated targets. Processed by the multi-ping single-sample GLRT detectors

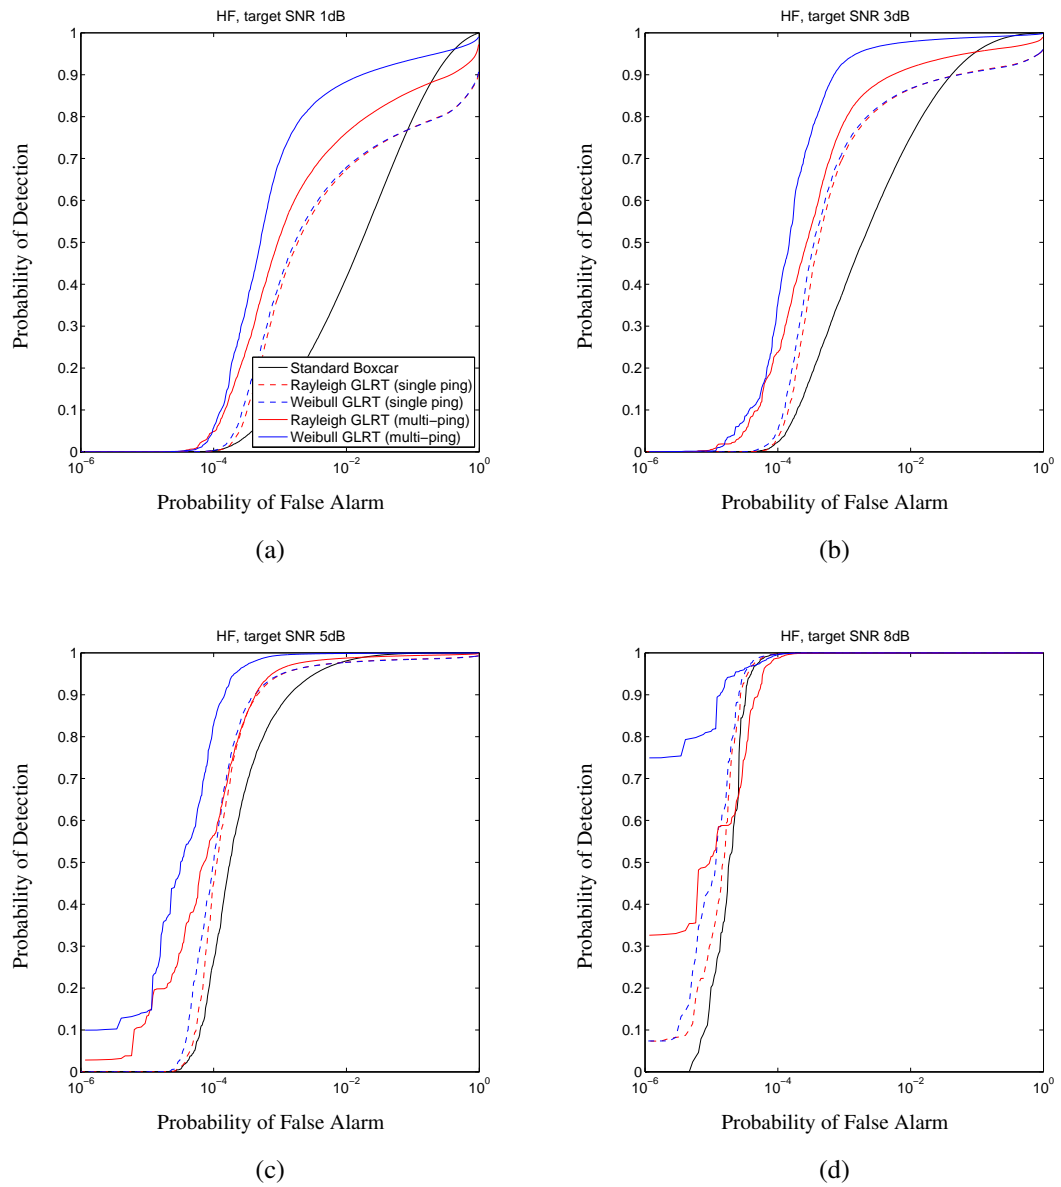


Figure B.12: ROC curves for various target strengths in the high frequency band against Swerling type 0 simulated targets. Processed by the multi-ping single-sample GLRT detectors

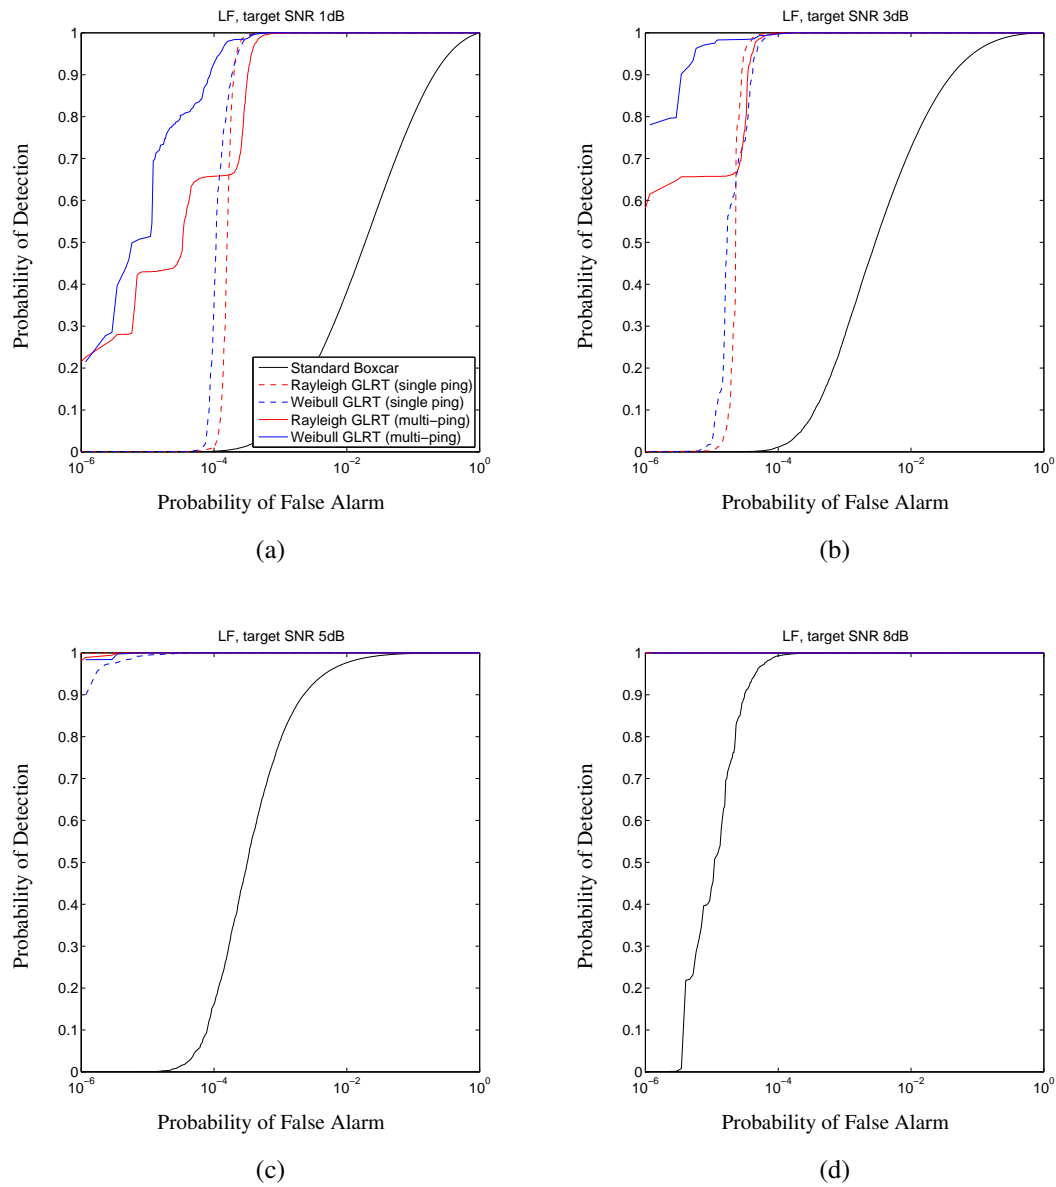


Figure B.13: ROC curves for various target strengths in the low frequency band against Swerling type 0 simulated targets. Processed by the multi-ping multi-sample GLRT detectors

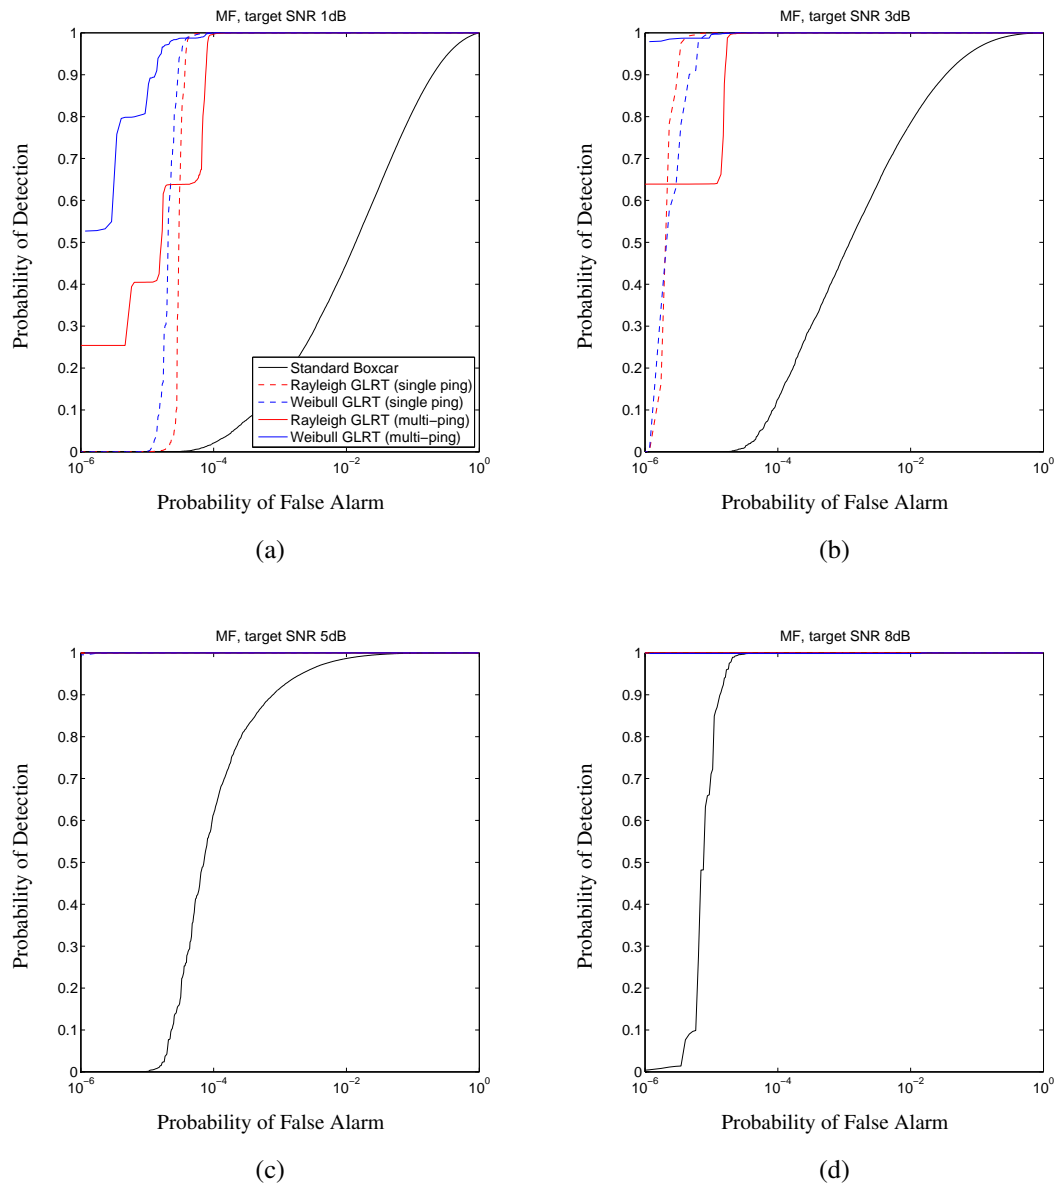


Figure B.14: ROC curves for various target strengths in the medium frequency band against Swerling type 0 simulated targets. Processed by the multi-ping multi-sample GLRT detectors

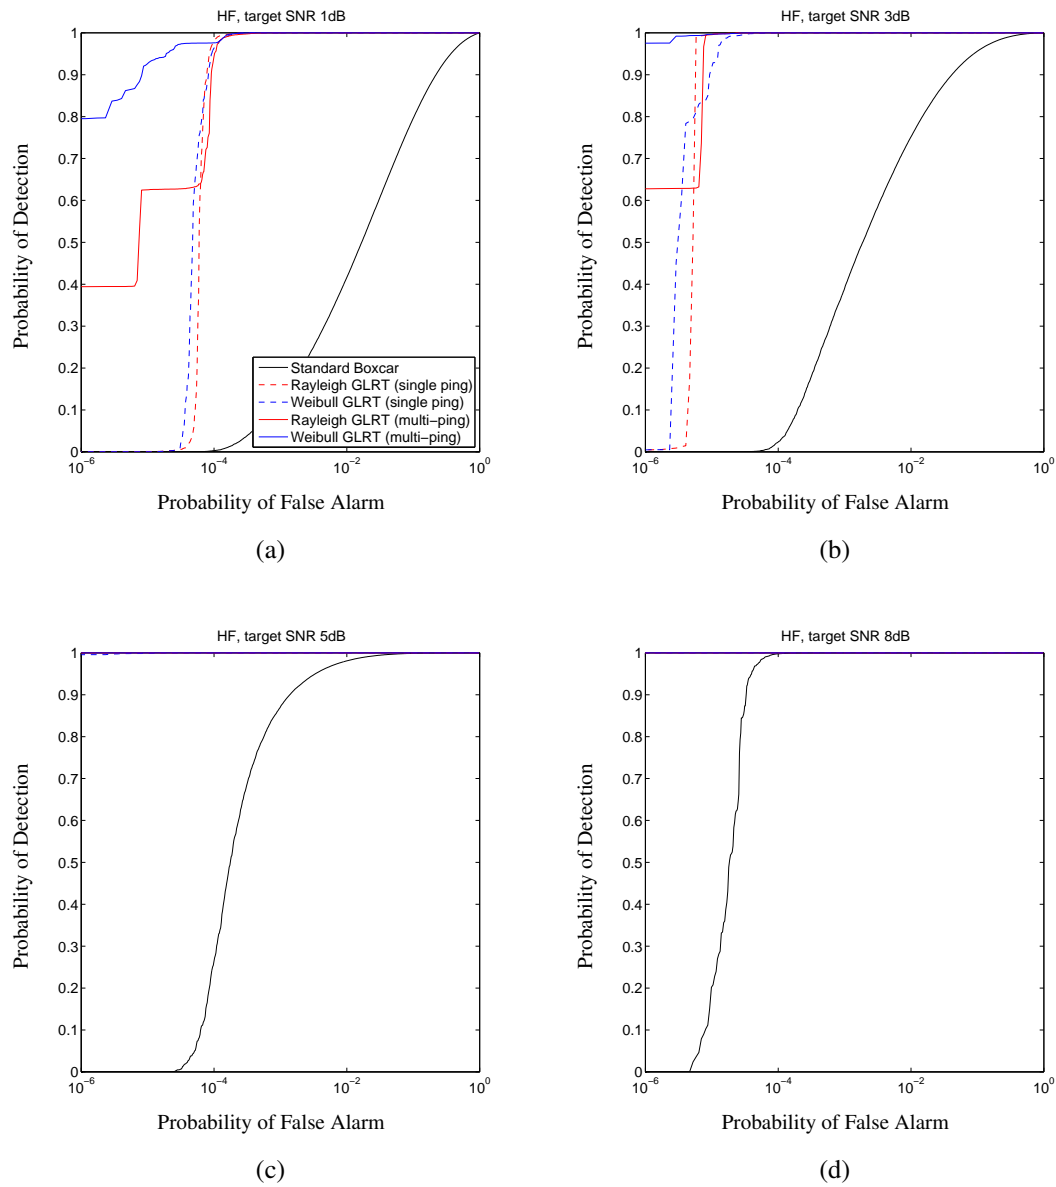


Figure B.15: ROC curves for various target strengths in the high frequency band against Swerling type 0 simulated targets. Processed by the multi-ping multi-sample GLRT detectors

Appendix C

Spatial Performance

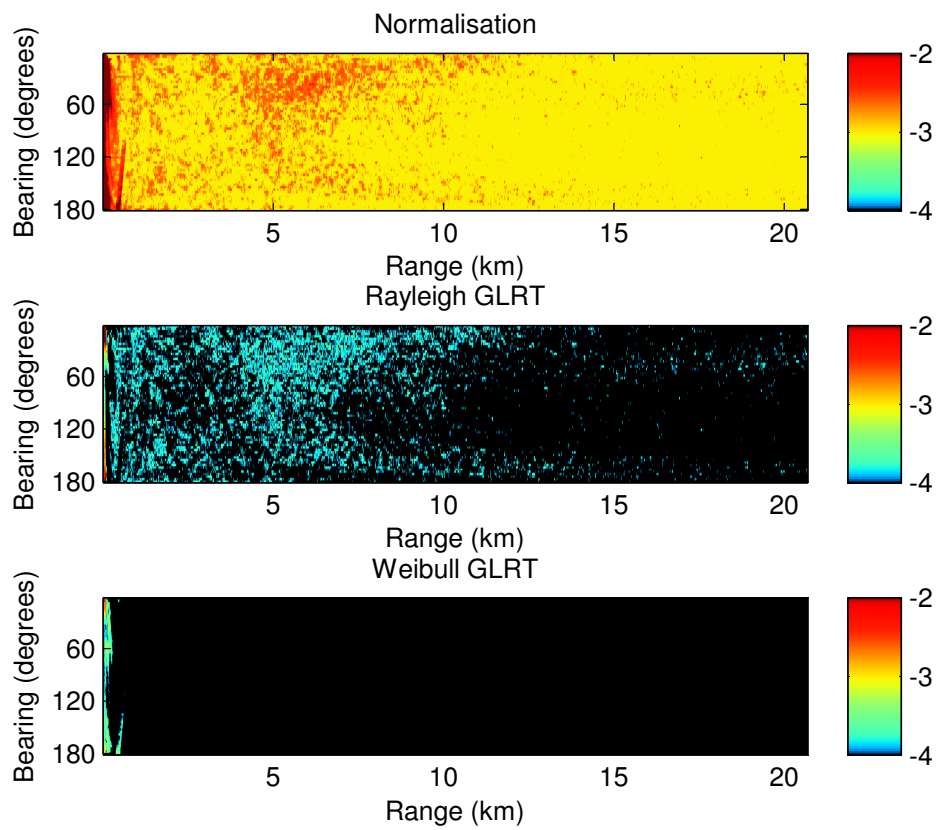


Figure C.1: Theoretical probability of false alarm (\log_{10}) for normalisation, Rayleigh GLRT, and Weibull GLRT detectors in the Low Frequency band, against a 3dB target at 50% probability of detection

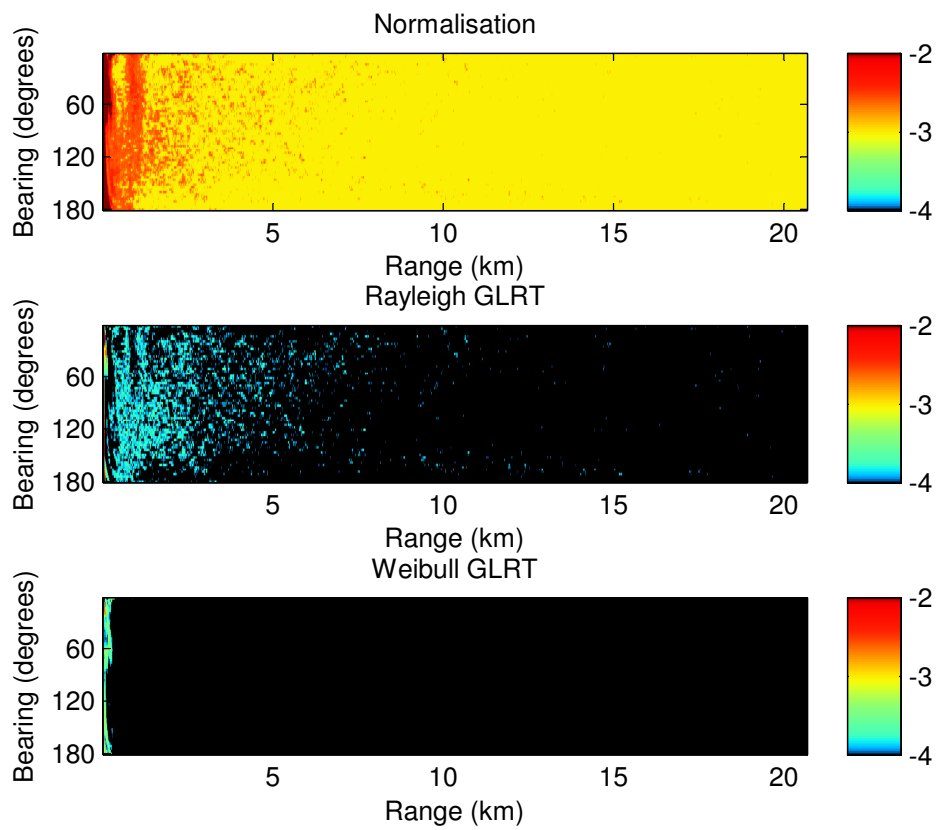


Figure C.2: Theoretical probability of false alarm (\log_{10}) for normalisation, Rayleigh GLRT, and Weibull GLRT detectors in the Medium Frequency band, against a 3dB target at 50% probability of detection

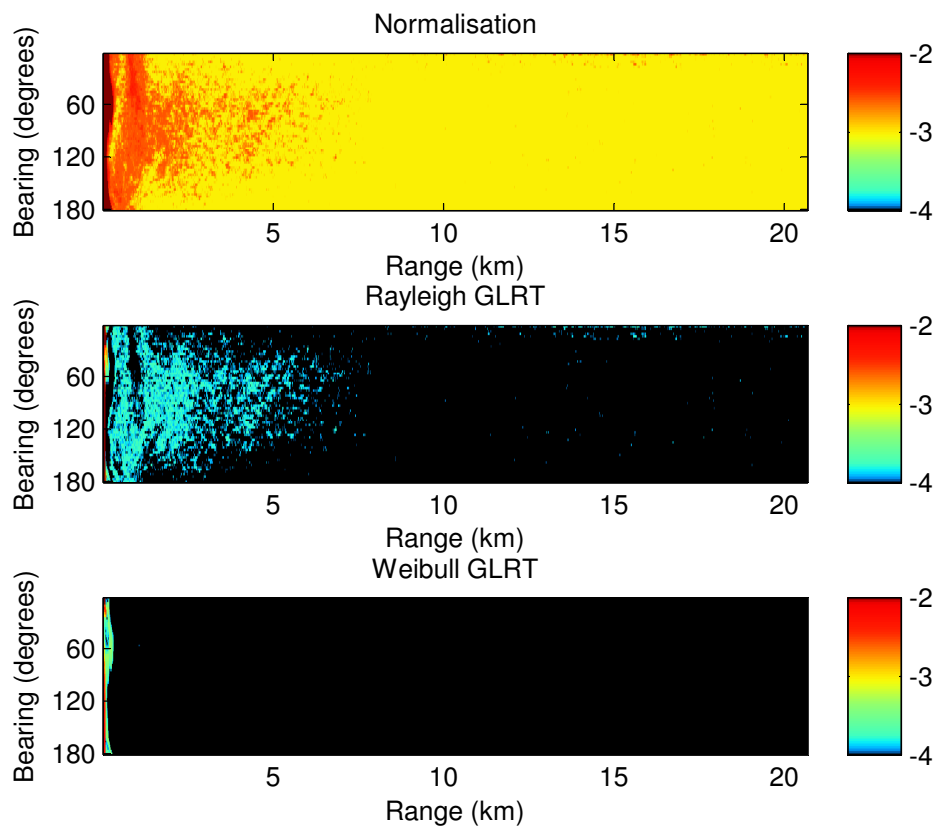


Figure C.3: Theoretical probability of false alarm (\log_{10}) for normalisation, Rayleigh GLRT, and Weibull GLRT detectors in the High Frequency band, against a 3dB target at 50% probability of detection

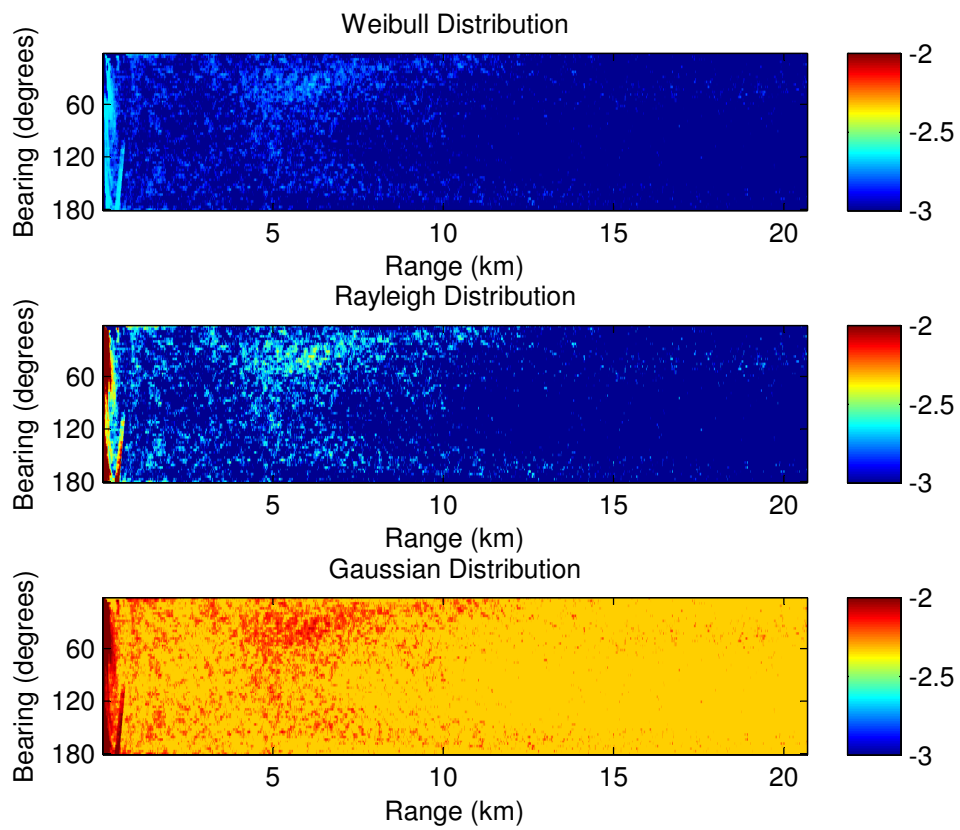


Figure C.4: Theoretical probability of false alarm (\log_{10}) for Weibull, Rayleigh, and Gaussian normalisation in the Low Frequency band, when selecting a probability of false alarm of 10^{-3} in K-distributed noise

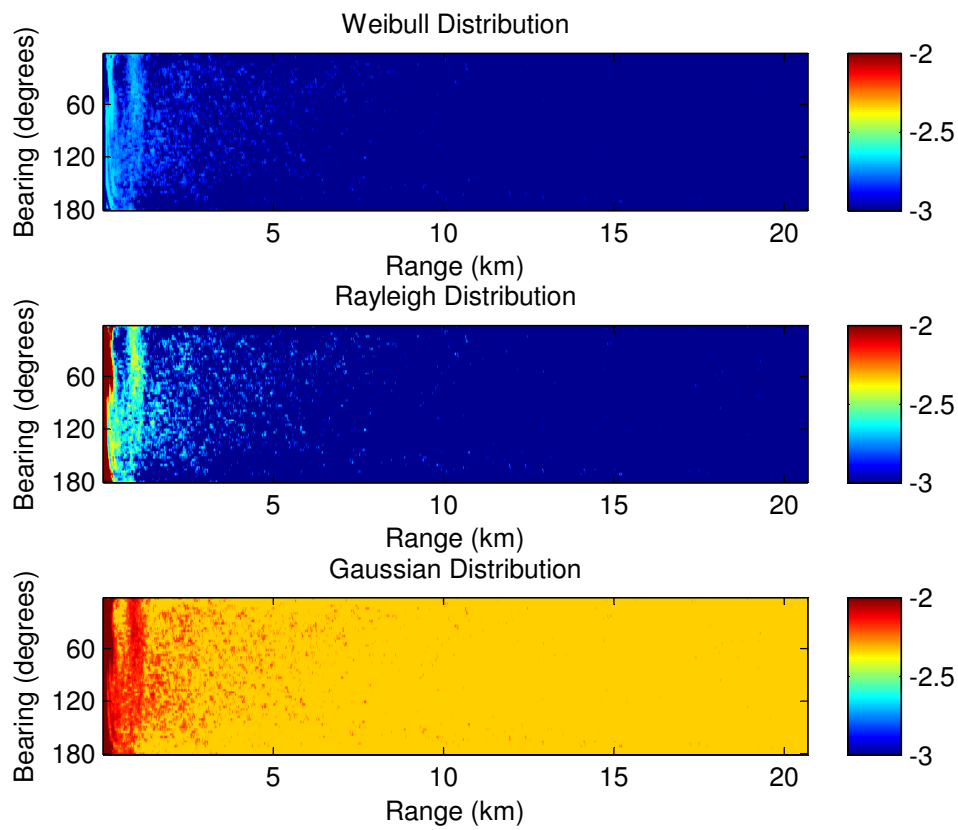


Figure C.5: Theoretical probability of false alarm (\log_{10}) for Weibull, Rayleigh, and Gaussian normalisation in the Medium Frequency band, when selecting a probability of false alarm of 10^{-3} in K-distributed noise

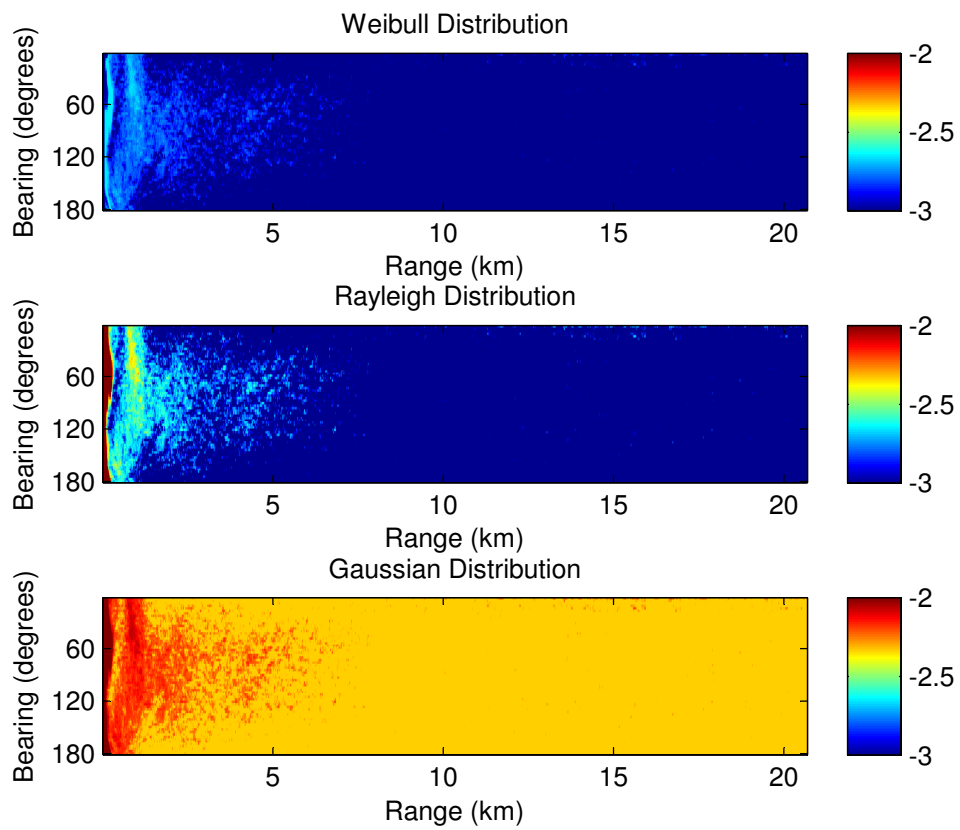


Figure C.6: Theoretical probability of false alarm (\log_{10}) for Weibull, Rayleigh, and Gaussian normalisation in the High Frequency band, when selecting a probability of false alarm of 10^{-3} in K-distributed noise




Figure C.7: Theoretical probability of false alarm (log10) averaged across all beams, for Weibull, Rayleigh, and Gaussian normalisation in the Low Frequency band, when selecting a probability of false alarm of 10^{-3} in K-distributed noise




Figure C.8: Theoretical probability of false alarm (log10) averaged across all beams, for Weibull, Rayleigh, and Gaussian normalisation in the Medium Frequency band, when selecting a probability of false alarm of 10^{-3} in K-distributed noise




Figure C.9: Theoretical probability of false alarm (log10) averaged across all beams, for Weibull, Rayleigh, and Gaussian normalisation in the High Frequency band, when selecting a probability of false alarm of 10^{-3} in K-distributed noise

References

- [1] I. S. 754-2008, *IEEE Standard for Floating-Point Arithmetic*. Institute of Electrical and Electronics Engineers, Inc., 2008, []1–58. DOI: 10.1109/IEEESTD.2008.4610935.
- [2] C. S. van Aartsen, “The response model - an operational model for towed array shape and attitude estimation from measurements by heading and depth sensors, validated on trial data,” *IEE Colloquium on Heading Sensors for Sonar and Marine Applications*, []5/1–5/10, Jan. 1994.
- [3] D. A. Abraham and A. P. Lyons, “Novel physical interpretations of K-distributed reverberation,” *IEEE Journal of Oceanic Engineering*, vol. 27, no. 4, []800–813, Oct. 2002.
- [4] —, “Reliable methods for estimating the K-distribution shape parameter,” *IEEE Journal of Oceanic Engineering*, vol. 35, no. 2, []288–302, Apr. 2010.
- [5] —, “Reverberation envelope statistics and their dependance on sonar bandwidth and scattering patch size,” *IEEE Journal of Oceanic Engineering*, vol. 29, no. 1, []126–137, Jan. 2004.
- [6] D. Abraham, “Broadband detection in K-distributed reverberation,” in *Sensor Array and Multichannel Signal Processing Workshop Proceedings, 2002*, Aug. 2002, []53–57. DOI: 10.1109/SAM.2002.1190998.
- [7] —, “Statistical normalization of non-rayleigh reverberation,” in *OCEANS '97. MTS/IEEE Conference Proceedings*, vol. 1, Oct. 1997, []500–505. DOI: 10.1109/OCEANS.1997.634415.

- [8] D. Abraham, J. Gelb, and A. Oldag, "Background and clutter mixture distributions for active sonar statistics," *Oceanic Engineering, IEEE Journal of*, vol. 36, no. 2, [231–247, Apr. 2011, ISSN: 0364-9059. DOI: 10.1109/JOE.2010.2102150.
- [9] D. Abraham and Prior, "Estimating independent-scatterer density from active sonar reverberation," in *8th European Conference on Underwater Acoustics*, Jun. 2006, [181–186.
- [10] M. Abramowitz and I. A. Stegun, *Handbook of Mathematical Functions: with Formulas, Graphs, and Mathematical Tables*, ninth Dover printing, tenth GPO printing, ser. Dover books on mathematics. New York: Dover, Jun. 1964, ISBN: 0486612724.
- [11] M. A. Ainslie and J. G. McColm, "A simplified formula for viscous and chemical absorption in sea water," *The Journal of the Acoustical Society of America*, vol. 103, no. 3, [1671–1672, Mar. 1998. DOI: 10.1121/1.421258.
- [12] V. M. Albers, *Underwater Acoustics Handbook*. Pennsylvania State University Press, 1960.
- [13] K. Alexander, "Environmental exemptions for the navys mid-frequency active sonar training program," U.S. Congressional Research Service, Tech. Rep. RL34403, Mar. 2008.
- [14] J. D. Altringham, T. McOwat, and L. Hammond, *Bats: biology and behaviour*. Oxford University Press, USA, 1998.
- [15] T. W. Anderson and D. A. Darling, "Asymptotic theory of certain "goodness of fit" criteria based on stochastic processes," *The Annals of Mathematical Statistics*, vol. 23, no. 2, [193–212, 1952.
- [16] A. Baldacci and G. Haralabus, "Signal processing for an active sonar system for advanced sensor technology applications and environmental adaptive schemes," *European Signal Processing Conference*, Sep. 2006.
- [17] R. Barakat, "Weak-scatterer generalization of the K-density function with application to laser scattering in atmospheric turbulence," *J. Opt. Soc. Am. A*, vol. 3, no. 4, [401–409, Apr. 1986.
- [18] R. Bareš, D. Evans, and S. Long, "Environmentally adaptive noise estimation for active sonar," *UDT Europe*, Jun. 2009.

- [19] —, “Noise estimation in long-range matched-filter envelope sonar data,” *Oceanic Engineering, IEEE Journal of*, vol. 35, no. 2, [230–235, Apr. 2010, ISSN: 0364-9059. DOI: 10.1109/JOE.2009.2036947.
- [20] R. Bareš, “On combining estimates of the K-distribution shape parameter,” To be submitted to *IEEE Signal Processing Letters*, 2012.
- [21] R. Benton, “Performance of an active sub-band processing scheme for clutter, reverberation and ambient noise limited environments,” *International Symposium on Underwater Reverberation and Clutter*, Sep. 2008.
- [22] D. Blacknell and R. J. Tough, “Parameter estimation for the K-distribution based on $[z \log(z)]$,” *IEE Proceedings, Radar, Sonar and Navigation*, vol. 148, no. 6, [309–312, Dec. 2001.
- [23] A. P. Blake, D. Blacknell, and C. J. Oliver, “High resolution SAR clutter textural analysis,” *Recent Developments in Radar and Sonar Imaging Systems: What Next?*, *IEE Colloquium on*, [10/1–10/9, Dec. 1995.
- [24] N. Bouhlef and S. Sevestre-Ghalila, “Nakagami markov random field as texture model for ultrasound RF envelope image,” *Computers in Biology and Medicine*, vol. 39, no. 6, [535–544, Jun. 2009.
- [25] G. E. P. Box and M. E. Muller, “A note on the generation of random normal deviates,” *The Annals of Mathematical Statistics*, vol. 29, no. 2, [610–611, 1958.
- [26] M. Brysbaert, “Algorithms for randomness in the behavioral sciences: a tutorial,” *Behavior Research Methods, Instruments, & Computers*, vol. 23, no. 1, [45–60, 1991.
- [27] E. H. Buck and K. Calvert, “Active military sonar and marine mammals: events and references,” U.S. Congressional Research Service, Tech. Rep. RL33133, Feb. 2007.
- [28] J. Cabrera and G. Watson, “Simulation methods for mean and median bias reduction in parametric estimation,” *Journal of Statistical Planning and Inference*, vol. 57, no. 1, [143–152, Jan. 1997.
- [29] G. Casella and R. L. Berger, “Statistical inference,” in, 2nd ed. Duxbury Press, Jun. 2001, [375.

- [30] M. J. Chantler and J. P. Stoner, "Automatic interpretation of sonar image sequences using temporal feature measures," *IEEE Journal of Oceanic Engineering*, vol. 22, no. 1, [47–56, Jan. 1997.
- [31] G. L. Chenadec, J. Boucher, and X. Lurton, "Angular dependence of K-distributed sonar data," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 45, no. 5, [1224–1235, May 2007.
- [32] A. Chernobai, S. Rachev, and F. Fabozzi, "Composite goodness-of-fit tests for left-truncated loss samples," University of California, Santa Barbara, Tech. Rep., Jun. 2005.
- [33] T. C. Choinski and C.-H. Lee, "Massively parallel systems design for real-time embedded applications," Naval Undersea Warfare Center, Newport, Tech. Rep., Feb. 1994.
- [34] V. Choulakian and M. A. Stephens, "Goodness-of-fit tests for the generalized pareto distribution," *Technometrics*, vol. 43, no. 4, [478–484, Nov. 2001.
- [35] P. C. Chu, M. Cornelius, and M. Wegstaff, "Effect of suspended sediment on acoustic detection using the navys CASS-GRAB model," *OCEANS, 2005. Proceedings of MTS/IEEE*, [1–7, Sep. 2005.
- [36] T. Clancy, *The Hunt for Red October*. Tom Clancy, 1984.
- [37] D. Clark, I. Ruiz, Y. Petillot, and J. Bell, "Particle PHD filter multiple target tracking in sonar image," *Aerospace and Electronic Systems, IEEE Transactions on*, vol. 43, no. 1, [409–416, Jan. 2007, ISSN: 0018-9251. DOI: 10.1109/TAES.2007.357143.
- [38] D. Clark, B.-N. Vo, and J. Bell, "GM-PHD filter multitarget tracking in sonar images," *Signal Processing, Sensor Fusion, and Target Recognition XV*, vol. 6235, no. 1, I. Kadar, Ed., 2006. DOI: 10.1117/12.663522.
- [39] P. Corona, G. Ferrara, and M. Migliaccio, "Generalized stochastic field model for reverberating chambers," *IEEE Transactions on Electromagnetic Compatibility*, vol. 46, no. 2, [655–660, Nov. 2004.
- [40] F. N. David and N. L. Johnson, "The probability integral transformation when parameters are estimated from the sample," *Biometrika*, vol. 35, no. 1-2, [182–190, May 1948.

- [41] L. Devroye, *Non-Uniform Random Variate Generation*. Springer-Verlag, 1986, [127–82].
- [42] C. Dias, A. Samaranayaka, and B. Manly, “On the use of correlated beta random variables with animal population modelling,” *Ecological Modelling*, vol. 215, no. 4, [293–300, 2008].
- [43] C. R. Dietrich and G. N. Newsam, “Fast and exact simulation of stationary Gaussian processes through circulant embedding of the covariance matrix,” *SIAM Journal of Scientific Computing*, vol. 18, no. 4, [1088–1107, Jul. 1997].
- [44] C. Dolph, “A current distribution for broadside arrays which optimizes the relationship between beam width and side-lobe level,” *Proceedings of the IRE*, vol. 34, no. 6, [335–348, Jun. 1946, ISSN: 0096-8390. DOI: 10.1109/JRPROC.1946.225956].
- [45] P. Fearnhead and Z. Liu, “Efficient online inference for multiple changepoint problems,” Sep. 2006, [5–8].
- [46] A. Frantzis, “Does acoustic testing strand whales?”, *Nature*, vol. 392, no. 6671, [29–29, Mar. 1998. DOI: 10.1038/32068].
- [47] C. A. Furuti, *Map projections: properties: preserving directions*, Retrieved 31 January 2011, Aug. 2010. [Online]. Available: <http://www.progonos.com/furuti/MapProj/Normal/CartProp/Rhumb/rhumb.html>.
- [48] A. Gambardella, F. Nunziata, A. Sorrentino, G. Ferrara, and M. Migliaccio, “Single-look SAR images and detection of sea dark areas,” *Proceedings of Envisat Symposium 2007*, Apr. 2007.
- [49] Y. S. Gill and R. Khanna, “Analysis of data transmission techniques in rayleigh, rician and nakagami fading model under various modulation schemes,” *International Journal of VLSI and Signal Processing Applications*, vol. 1, no. 2, [108–111, May 2011].
- [50] O. Gomes, C. Combes, and A. Dussauchoy, “Parameter estimation of the generalized gamma distribution,” *Mathematics and Computers in Simulation*, vol. 79, [955–963, Feb. 2008].
- [51] Google Inc., *Google earth (version 6.0.1.2032)*, [Software], Retrieved 16 January 2010, Dec. 2010. [Online]. Available: <http://earth.google.com>.

- [52] —, *Google maps*, Retrieved 16 January 2010. [Online]. Available: <http://maps.google.com>.
- [53] E. Gould, N. C. Negus, and A. Novick, “Evidence for echolocation in shrews,” *Journal of Experimental Zoology*, vol. 156, no. 1, [19–37, 1964, ISSN: 1097-010X. DOI: 10.1002/jez.1401560103.
- [54] J. Groen, S. P. Beerens, and Y. Doisy, “Shape and doppler correction for a towed sonar array,” *OCEANS '04. MTT/IEEE TECHNO-OCEAN '04*, vol. 2, [613–620, Nov. 2004.
- [55] X. Hao, C. J. Bruce, C. Pislaru, and J. F. Greenleaf, “Characterization of reperfused infarcted myocardium from high-frequency intracardiac ultrasound imaging using homodyned K distribution,” *Ultrasonics, Ferroelectrics and Frequency Control, IEEE Transactions on*, vol. 49, no. 11, [1530–1542, Nov. 2002.
- [56] —, “Identification of reperfused infarcted myocardium from highfrequency intracardiac ultrasound images using homodyned K distribution,” *Ultrasonics Symposium, 2001 IEEE*, vol. 2, [1189–1192, 2001.
- [57] C. L. S. Howeth, *History of Communications-Electronics in the United States Navy*. United States Government Printing Office, 1963.
- [58] D. P. Hruska, “Improved techniques for statistical analysis of the envelope of backscattered ultrasound using the homodyned K distribution,” Master’s thesis, University of Illinois at Urbana-Champaign, 2007.
- [59] W. Hu, “Estimation of K-distribution parameters using $[z^r \log(z)]$,” *Radar Conference, 2009 IET International*, [1–4, Apr. 2009.
- [60] D. Iskander and A. Zoubir, “Estimation of the parameters of the K-distribution using higher order and fractional moments [radar clutter],” *Aerospace and Electronic Systems, IEEE Transactions on*, vol. 35, no. 4, [1453–1457, Oct. 1999, ISSN: 0018-9251. DOI: 10.1109/7.805463.
- [61] E. Jakeman, “On the statistics of K-distributed noise,” *Journal of Physics A: Mathematical and General*, vol. 13, [31–48, Jan. 1980.
- [62] E. Jakeman and P. N. Pusey, “A model for non-Rayleigh sea echo,” *IEEE Transactions on Antennas and Propagation*, vol. 24, no. 6, [806–814, Nov. 1976.

- [63] E. Jakeman and R. J. A. Tough, "Generalized K distribution: a statistical model for weak scattering," *J. Opt. Soc. Am. A*, vol. 4, no. 9, [1764–1772, Sep. 1987.
- [64] Q. Jiang, S. Wang, D. Ziou, and A. E. Zaart, "Automatic detection for ship targets in radarsat sar images from coastal regions," in *Vision Interface '99, Trois-Rivières, Canada*, May 1999, [131–137.
- [65] G. Jones, "Echolocation," *Current Biology*, vol. 15, no. 13, [R484 –R488, 2005, ISSN: 0960-9822. DOI: 10.1016/j.cub.2005.06.051.
- [66] I. R. Joughin, D. B. Percival, and D. P. Winebrenner, "Maximum likelihood estimation of K distribution parameters for SAR data," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 31, no. 5, [989–999, Sep. 1993.
- [67] R. E. Kalman, "A new approach to linear filtering and prediction problems," *Transactions of the ASME—Journal of Basic Engineering*, vol. 82, no. Series D, [35–45, 1960.
- [68] S. M. Kay, *Fundamentals of Statistical Signal Processing*. Prentice Hall PTR, 1993, vol. 1: Estimation Theory.
- [69] —, *Fundamentals of Statistical Signal Processing*. Prentice Hall PTR, 1998, vol. 2: Detection Theory.
- [70] E. J. Kelly, "An adaptive detection algorithm," *Aerospace and Electronic Systems, IEEE Transactions on*, vol. AES-22, no. 2, [115–127, Mar. 1986.
- [71] K. Kim, N. Neretti, and N. Intrator, "Video enhancement for underwater exploration using forward looking sonar," *Advanced Concepts for Intelligent Vision Systems*, vol. 4179, [554–563, Oct. 2006.
- [72] A. J. Kinderman and J. F. Monahan, "Generating random variables from the ratio of two uniform variates," *Proc. 9th Symposium on the Interface of Computer Science and Statistics*, [197–200, 1976.
- [73] C. Kleiber and S. Kotz, *Statistical Size Distributions in Economics and Actuarial Sciences*. Wiley, 2003.
- [74] W. C. Knight, R. G. Pridham, and S. M. Kay, "Digital signal processing for sonar," *Proceedings of the IEEE*, vol. 69, no. 11, [1451–1506, Nov. 1981.
- [75] D. E. Knuth, *Art of Computer Programming, Volume 1: Fundamental Algorithms*. Addison-Wesley, 1969, ISBN: 0-201-03802-1.

- [76] F. Kopsaftopoulos and S. Fassois, "Vibration based health monitoring for a lightweight truss structure: experimental assessment of several statistical time series methods," *Mechanical Systems and Signal Processing*, vol. 24, no. 7, [1977–1997, 2010, Special Issue: ISMA 2010, ISSN: 0888-3270. DOI: 10.1016/j.ymssp.2010.05.013.
- [77] M. J. Krzych, "Case study: using field programmable gate arrays in a beowulf cluster," Naval Undersea Warfare Center, Newport, Tech. Rep., Feb. 2005.
- [78] H. Levine, L. Bildsten, M. Brenner, C. Callan, S. Flatté, J. Goodman, M. Gregg, J. Katz, W. Munk, and P. Weinberger, "Active sonar waveform," JASON, MITRE Corporation, Tech. Rep., Jun. 2004.
- [79] J. H. Lienhard and P. L. Meyer, "A physical basis for the generalized gamma distribution," *Quarterly of Applied Mathematics*, vol. XXV, no. 3, [330–334, Oct. 1967.
- [80] H. W. Lilliefors, "On the kolmogorov-smirnov test for normality with mean and variance unknown," *Journal of the American Statistical Association*, vol. 62, no. 318, [399–402, Jun. 1967.
- [81] P. Lombardo and C. Oliver, "Estimation of texture parameters in K-distributed clutter," *Radar, Sonar and Navigation, IEE Proceedings -*, vol. 141, no. 4, [196–204, Aug. 1994.
- [82] S. Lopez-Estrada, R. Cumplido-Parra, and C. Torres-Huitzil, "A hybrid approach for target detection using cfar algorithm and image processing," in *Computer Science, 2004. ENC 2004. Proceedings of the Fifth Mexican International Conference in*, Sep. 2004, [108–115. DOI: 10.1109/ENC.2004.1342595.
- [83] "Lucien ira blake dies," *The New York Times*, May 1916.
- [84] J. B. MacQueen, "Some methods for classification and analysis of multivariate observations," *Proceedings of 5th Berkeley Symposium on Mathematical Statistics and Probability*, vol. 1, [281–297, 1967.
- [85] R. Mahler, "Multitarget bayes filtering via first-order multitarget moments," *Aerospace and Electronic Systems, IEEE Transactions on*, vol. 39, no. 4, [1152–1178, Oct. 2003, ISSN: 0018-9251. DOI: 10.1109/TAES.2003.1261119.

- [86] A. Mao, C. G. A. Harrison, and T. H. Dixon, "Noise in gps coordinate time series," *J. Geophys. Res.*, vol. 104, no. B2, [2797–2816, 1999, ISSN: 0148-0227. DOI: 10.1029/1998JB900033.
- [87] J. I. Marcum, "A statistical theory of target detection by pulsed radar," RAND Corporation, Tech. Rep. RM-754-PR, Dec. 1947.
- [88] G. Marsaglia, "Random variables and computers," *Information theory, statistical decision functions, random processes: Transactions of the Third Prague Conference*, [499–512, Jun. 1962.
- [89] G. Marsaglia and W. W. Tsang, "A simple method for generating gamma variables," *ACM Transactions on Mathematical Software*, vol. 26, no. 3, [363–372, Sep. 2000.
- [90] G. Marsaglia, "Random number generators," *Journal of Modern Applied Statistical Methods*, vol. 2, no. 1, [2–13, May 2003.
- [91] G. Marsaglia and W. W. Tsang, "The ziggurat method for generating random variables," *Journal of Statistical Software*, vol. 5, no. 8, [1–7, Oct. 2000, ISSN: 1548-7660.
- [92] M. Martin-Fernandez, R. Cardenes, and C. Alberola-Lopez, "Parameter estimation of the homodyned K distribution based on signal to noise ratio," *IEEE Ultrasonics Symposium*, Oct. 2006.
- [93] F. J. Massey, "The kolmogorov-smirnov test for goodness of fit," *Journal of the American Statistical Association*, vol. 46, no. 253, [68–78, Mar. 1951.
- [94] MathWorks, *R2010b documentation: acceptance-rejection methods*, Retrieved 31 January 2011, 2010. [Online]. Available: <http://www.mathworks.com/help/toolbox/stats/br5k9hi-1.html#br5k9hi-4>.
- [95] —, *Solution ID: 1-OYFMA*, Retrieved 26 February 2010, 2009. [Online]. Available: <http://www.mathworks.com/support/solutions/en/data/1-OYFMA>.
- [96] M. Matsumoto and T. Nishimura, "Mersenne twister: a 623-dimensionally equidistributed uniform pseudo-random number generator," *ACM Trans. Model. Comput. Simul.*, vol. 8, no. 1, [3–30, 1998, ISSN: 1049-3301. DOI: 10.1145/272991.272995.

- [97] M. Migliaccio, G. Ferrara, A. Gambardella, and D. Aiello, "The k-generalized model to characterize marine features in SAR images: first results," *Proceedings of SEASAR 2006*, Jan. 2006.
- [98] National Research Council, *Low-Frequency Sound and Marine Mammals: Current Knowledge and Research Needs*. National Academy Press, 1994.
- [99] W. Ng, J. Li, S. Godsill, and J. Vermaak, "A review of recent results in multiple target tracking," in *Image and Signal Processing and Analysis, 2005. ISPA 2005. Proceedings of the 4th International Symposium on*, Sep. 2005, [40–45]. DOI: 10.1109/ISPA.2005.195381.
- [100] C. J. Oliver, "Optimum texture estimators for SAR clutter," *Journal of Physics D: Applied Physics*, vol. 26, no. 11, [1824–1835, Nov. 1993. DOI: 10.1088/0022-3727/26/11/002.
- [101] K. Panta, B.-N. Vo, S. Singh, and A. Doucet, "Probability hypothesis density filter versus multiple hypothesis tracking," *Signal Processing, Sensor Fusion, and Target Recognition XIII*, vol. 5429, no. 1, I. Kadar, Ed., [284–295, 2004. DOI: 10.1117/12.543357.
- [102] A. Pesavento, H. Ermert, E. B. Zeitvogel, and J. Grifka, "High resolution imaging of generalized K-distribution parameters using maximum likelihood estimation for ultrasonic diagnosis of muscle after back surgery," *Ultrasonics Symposium, 1998. Proceedings., 1998 IEEE*, vol. 2, [1353–1356, Oct. 1998.
- [103] Photographer: PH3 C. Yebba, *Virin: dn-sc-85-06854*, U.S. Department of Defense, Retrieved 07 December 2010, Aug. 1983. [Online]. Available: <http://www.defenceimagery.mil>.
- [104] G. Pollon, "Statistical parameters for scattering from randomly oriented arrays, cylinders, and plates," *Antennas and Propagation, IEEE Transactions on*, vol. 18, no. 1, [68–75, Jan. 1970.
- [105] M. K. Prior, "Estimation of K-distribution shape parameter from sonar data: sample size limitations," *IEEE Journal of Oceanic Engineering*, vol. 34, no. 1, [45–50, Jan. 2009.
- [106] M. K. Prior, "A scatterer map for the malta plateau," *IEEE Journal of Oceanic Engineering*, vol. 30, no. 4, [676–690, Oct. 2005.

- [107] R. Raghavan, "A method for estimating parameters of K-distributed clutter," *Aerospace and Electronic Systems, IEEE Transactions on*, vol. 27, no. 2, [238–246, Mar. 1991, ISSN: 0018-9251. DOI: 10.1109/7.78298.
- [108] B. I. Raju and M. A. Srinivasan, "Statistics of envelope of high-frequency ultrasonic backscatter from human skin in vivo," *Ultrasonics, Ferroelectrics and Frequency Control, IEEE Transactions on*, vol. 49, no. 7, [871–882, Jul. 2002.
- [109] N. J. Redding, "Estimating the parameters of the K distribution in the intensity domain," DSTO Electronics and Surveillance Research Laboratory, Australia, Tech. Rep., Jul. 1999.
- [110] D. Reid, "An algorithm for tracking multiple targets," *Automatic Control, IEEE Transactions on*, vol. 24, no. 6, [843–854, Dec. 1979, ISSN: 0018-9286. DOI: 10.1109/TAC.1979.1102177.
- [111] R. J. Renomeron, D. V. Rabinkin, J. C. French, and J. L. Flanagan, "Small-scale matched filter array processing for spatially selective sound capture," *J. Acoust. Soc. Am*, vol. 102, [3208–3217, 1997.
- [112] J. R. Reynolds, "Submarines, sonar, and the death of whales: enforcing the delicate balance of environmental compliance and national security in military training," *William & Mary Environmental Law and Policy Review*, vol. 32, no. 3, [759–802, 2008.
- [113] M. A. Richards, *Generating swerling random sequences*, Retrieved 30 October 2010, Jan. 2008. [Online]. Available: <http://users.ece.gatech.edu/mrichard/GeneratingSwerlingRandomSequences.pdf>.
- [114] —, *Relationship between the gamma, erlang, chi-square, and swerling 3/4 probability density functions*, Retrieved 30 October 2010, Aug. 2007. [Online]. Available: <http://users.ece.gatech.edu/mrichard/Gamma,Erlang,&Chi-Square.pdf>.
- [115] —, *Rice distributions for RCS*, Retrieved 30 October 2010, Aug. 2010. [Online]. Available: <http://users.ece.gatech.edu/mrichard/Ricepowerpdf.pdf>.

- [116] M. Sezgin and B. Sankur, "Survey over image thresholding techniques and quantitative performance evaluation," *Journal of Electronic Imaging*, vol. 13, no. 1, [146–165, Jan. 2004.
- [117] P. M. Shankar, "A general statistical model for ultrasonic backscattering from tissues," *Ultrasonics, Ferroelectrics and Frequency Control, IEEE Transactions on*, vol. 47, no. 3, [727–736, May 2000.
- [118] J. Sijbers, A. den Dekker, A. J. D. Dekker, P. Scheunders, and D. V. Dyck, "Maximum-likelihood estimation of rician distribution parameters," *IEEE Transactions on Medical Imaging*, vol. 17, [357–361, 1998.
- [119] C. D. Sinclair, B. D. Spurr, and M. I. Ahmad, "Modified anderson darling test," *Communications in Statistics - Theory and Methods*, vol. 19, no. 10, [3677–3686, Jan. 1990.
- [120] R. W. Sinnott, "Virtues of the haversine," *Sky and Telescope*, vol. 68, no. 2, [159, 1984.
- [121] M. Skolnik, *Introduction to Radar Systems*, 3rd ed. McGraw Hill, 2002.
- [122] J. P. Snyder, *Flattening the Earth: Two Thousand Years of Map Projections*. University of Chicago Press, 1993.
- [123] W. Song and L.-C. Hsiao, "Generation of autocorrelated random variables with a specified marginal distribution," *Simulation Conference Proceedings*, [374–377, Dec. 1993. DOI: 10.1109/WSC.1993.718074.
- [124] M. Sonka, V. Hlavac, and R. Boyle, *Image Processing, Analysis, and Machine Vision*, 2nd ed. Thomson-Engineering, 1998.
- [125] M. A. Stephens, "Asymptotic results for goodness-of-fit statistics with unknown parameters," *The Annals of Statistics*, vol. 4, no. 2, [357–369, Mar. 1976.
- [126] —, "EDF statistics for goodness of fit and some comparisons," *Journal of the American Statistical Association*, vol. 69, no. 347, [730–737, Sep. 1974.
- [127] —, "Goodness of fit for the extreme value distribution," *Biometrika*, vol. 64, no. 3, [583–588, Dec. 1977.
- [128] —, "Tests of fit for the logistic distribution based on the empirical distribution function," *Biometrika*, vol. 66, no. 3, [591–595, Dec. 1979.

- [129] R. D. Strum and D. E. Kirk, *First principles of discrete systems and digital signal processing*. Boston, MA, USA: Addison-Wesley Longman Publishing Co., Inc., 1988, [521, ISBN: 0-201-09518-1.
- [130] Y. Sun and M. Farooq, “Fuzzy modelling of noise for sonar sensors in a realistic sea environment,” *Signal Processing and Its Applications, 1996. ISSPA 96., Fourth International Symposium on*, vol. 2, [821–822, Aug. 1996.
- [131] P Swerling, “Probability of detection for fluctuating targets,” *IRE Transactions on Information Theory*, vol. 6, no. 2, [269–308, Apr. 1960.
- [132] G. Turin, “An introduction to matched filters,” *Information Theory, IRE Transactions on*, vol. 6, no. 3, [311–329, Jun. 1960, ISSN: 0096-1000. DOI: 10.1109/TIT.1960.1057571.
- [133] T. Vincenty, “Direct and inverse solutions of geodesics on the ellipsoid with application of nested equations,” *Survey Review*, vol. 23, no. 176, [88–93, Apr. 1975.
- [134] B.-N. Vo, S. Singh, and A. Doucet, “Sequential monte carlo methods for multitarget filtering with random finite sets,” *Aerospace and Electronic Systems, IEEE Transactions on*, vol. 41, no. 4, [1224–1245, Oct. 2005, ISSN: 0018-9251. DOI: 10.1109/TAES.2005.1561884.
- [135] B.-N. Vo and W.-K. Ma, “The gaussian mixture probability hypothesis density filter,” *Signal Processing, IEEE Transactions on*, vol. 54, no. 11, [4091–4104, Nov. 2006, ISSN: 1053-587X. DOI: 10.1109/TSP.2006.881190.
- [136] A. D. Waite, *Sonar for practising engineers*, 3rd ed. Wiley, 2001.
- [137] Y. Wang, P. Li, and J. Zhang, “Fast random bit generation in optical domain with ultrawide bandwidth chaotic laser,” *Photonics Technology Letters, IEEE*, vol. 22, no. 22, [1680–1682, Nov. 2010, ISSN: 1041-1135. DOI: 10.1109/LPT.2010.2080265.
- [138] K. D. Ward, R. J. Tough, and S. Watts, *Sea Clutter: Scattering, the K Distribution and Radar Performance*, ser. 20. IET Radar, Sonar and Navigation, 2006.
- [139] T. Wood, D. Allwright, P. Bond, S. Long, and I. Moroz, “A new method for processing passive sonar data,” in *Information Fusion (FUSION), 2010 13th Conference on*, Jul. 2010, [1–7.

- [140] T. Wood, "Tracking in dense clutter with the PHD filter," *Proceedings of the IMA Mathematics in Defence Conference, Farnborough*, 2009.
- [141] Y. Yang, C. Wan, L. Zhou, and C. S. Chia, "Detection of signals in reverberation via beamspace adaptive matched filter," *OCEANS '04. MTTT/IEEE TECHNO-OCEAN '04*, vol. 2, [603–607, Nov. 2004.