Data Selection in Binary Hypothesis Testing

by

Charles K. Sestok IV

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

December 2003

Author .....
Department of Electrical Engineering and Computer Science
December 9, 2003

Certified by .....
Alan V. Oppenheim
Ford Professor of Electrical Engineering
Thesis Supervisor

Accepted by .....
Arthur C. Smith
Chairman, Department Committee on Graduate Students

Data Selection in Binary Hypothesis Testing

by

Charles K. Sestok IV

Submitted to the Department of Electrical Engineering and Computer Science
on December 9, 2003, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

Traditionally, statistical signal processing algorithms are developed from probabilistic models for data. The design of the algorithms and their ultimate perfomance depend upon these assumed models.

In certain situations, collecting or processing all available measurements may be inefficient or prohibitively costly. A potential technique to cope with such situations is data selection, where a subset of the measurements that can be collected and processed in a cost-effective manner is used as input to the signal processing algorithm. Careful evaluation of the selection procedure is important, since the probabilistic description of distinct data subsets can vary significantly. An algorithm designed for the probabilistic description of a poorly chosen data subset can lose much of the potential performance available to a well-chosen subset.

This thesis considers algorithms for data selection combined with binary hypothesis testing. We develop models for data selection in several cases, considering both random and deterministic approaches. Our considerations are divided into two classes depending upon the amount of information available about the competing hypotheses. In the first class, the target signal is precisely known, and data selection is done deterministically. In the second class, the target signal belongs to a large class of random signals, selection is performed randomly, and semi-parametric detectors are developed.

Thesis Supervisor: Alan V. Oppenheim
Title: Ford Professor of Electrical Engineering

Acknowledgments

A great deal of work went into this thesis, and I was fortunate enough to have many people helping me along the way. Foremost, my work would not have taken the course it did without my thesis advisor Alan Oppenheim. Working with Al has been a tremendous opportunity and privilege. His guidance and friendship has been incredibly valuable. I can't fully express in words how much I value Al's advice and constant encouragement to think creatively.

I benefitted siginficantly from the active involvement of my thesis readers - David Karger and Charles Rohrs. Charlie and David's suggestions were invariably thought-provoking, and often lead to fruitful approaches to the problems I had posed. Their advice and effort significantly improved the quality of my research.

Throughout my time at MIT, I benefitted from the interest and active collaboration of several of our research group's sponsors, visitors, and faculty. I couldn't envision a better research sponsor than Stephen Blatt. Participation in the SensIT program with him was a terrific experience for me. Stephen gave a very useful application perspective that helped me to frame my thesis research. Ram Zamir, during his sabbatical visit to MIT, gave me stimulating suggestions for the thesis. I also benefitted from my experience as a TA of 6.432 with Greg Wornell. Thinking through the fundamentals of detection and estimation (and making exam problems) lead to some interesting lines of research.

I'm grateful for the financial support of a National Science Foundation Graduate Research Fellowship, the Texas Instruments Leadership University Consortium, the DARPA SensIT program, and the Army Research Lab Advanced Sensors Collaborative Technology Alliance.

My time at MIT was enriched by participation in the Digital Signal Processing Group. As a graduate student, the support of fellow students is important, and DSPG had a large and energetic group of graduate students to learn, think, and have fun with. My colleagues - Richard Barron, Maya Said, Wade Torres, Chris Hadjicostis, Stark Draper, Nick Laneman, Everest Huang, Emin Martinain, Andrew Russell, Huan

Yao, Albert Chan, Yonina Eldar, Brian Chen, Mike Lopez, Petros Boufounos, Ashish Khisti, Sourav Dey, Uri Erez, Vijay Divi - were always there to talk through a research problem, give advice, or blow off steam when necessary. DSPG's excellent staff - Darla Chupp, Alecia Batson, Janice Zaganjori, Dianne Wheeler, and Tricia Mulcahy - made sure that everything in the group ran smoothly, and we could focus our efforts on research. Last but not least, Vanni Aliberti provided sage advice and kept the critical computer system running smoothly.

Finally, my family's support and encouragement has been boundless during my time at MIT. I wouldn't have accomplished all that I have without the love of my parents Charles and Nancy, grandparents George and Shirley Dixon, and Edith Sestok, and my brother Evan.

Contents

1 Introduction **13**

1.1 Binary Hypothesis Testing for Signal Detection 14

1.2 Data Selection Problem Formulation 16

1.3 Applications and Connections to Previous Research 18

1.4 Outline 20

2 Data Selection in Deterministic Signal Detection **23**

2.1 Traditional Matched Filtering 23

2.2 Restricted Matched Filter 26

2.3 Example 28

2.4 Greedy Algorithms for the RMF Search 33

2.4.1 Expressions for Incremental SNR Changes During Greedy RMF Searches 34

2.4.2 Conditions for An Exact Solution with Maximum Energy Selection 36

2.4.3 Worst-Case Performance of the Greedy Algorithm 40

2.4.4 Nearly Tight Example of the Forward Greedy Algorithm . . . 42

2.4.5 Nearly Tight Example of the Backward Greedy Algorithm . . 45

2.4.6 Performance Comparison for the Greedy Algorithm 46

2.5 Dynamic Programming Solutions for Banded Matrices 50

2.5.1 Fragment Decomposition of Selected Subsets 51

2.5.2 Example of Dynamic Programming for Two-Measurement Subsets 55

2.5.3 Fragment Notation for the Dynamic Programming Algorithm 56
2.5.4 Optimization Procedure for the Dynamic Programming Algorithm . . . 58
2.5.5 Complexity of Dynamic Programming Algorithm . . . 62
2.5.6 Extension to General Banded Matrices . . . 63
2.6 Summary . . . 68

3 Data Selection in Random Signal Detection 69
3.1 Random Signal Models . . . 70
3.1.1 General Signal Model . . . 70
3.1.2 Notation . . . 71
3.2 Likelihood Ratio Test . . . 71
3.2.1 Example - Detecting a Sinusoidal Signal . . . 73
3.2.2 General Properties of the Likelihood Ratio for Even Signals . 78
3.3 Semi-parametric Detector . . . 82
3.3.1 Properties of the Likelihood Ratio for r . . . 83
3.3.2 Low-Complexity Algorithm to Calculate the Detector Threshold 85
3.3.3 Robustness Properties of the Threshold Detector . . . 92
3.4 Summary . . . 97

4 Conclusions 99
4.1 The Restricted Matched Filter . . . 99
4.2 Randomized Selection and Robust Detection . . . 100

A Proof of Theorem 7 103

B Proof of Theorem 8 105

C Proof of Lemma 5 109
C.1 Proof for Smooth Decision Regions . . . 109
C.2 Alternative Proof for Gaussian Noise Densities . . . 111

List of Figures

2-1 Whitened Matched Filter Detector based on $\mathbf{x}$. 24

2-2 Receiver operating characteristics for best and worst detectors. The best SNR is 1.384, and the worst is 1.179. 31

2-3 Receiver operating characteristics for best and worst detectors generated for an RMF instance with $N = 16$, $K = 8$, a constant signal, and a randomly generated covariance matrix. The best SNR is 1.8085, and the worst is 0.1174. 32

2-4 Target signal for greedy search examples. 47

2-5 SNR versus K for Forward Greedy selection algorithm, maximum signal energy selection, and exhaustive search selection. 48

2-6 $K = 5$ ROC for Forward Greedy selection algorithm, maximum signal energy selection, and exhaustive search selection. 49

2-7 The top decision tree shows the entire set of feasible node choices for an RMF instance with $N = 4$, $K = 2$, and $b = 1$. The bottom decision tree shows the simplification achieved by applying the dynamic programming relations in Theorems (3) and (4). 61

2-8 The top decision tree shows the entire set of feasible node choices for an RMF instance with $N = 5$, $K = 3$, and $b = 1$. The bottom decision tree shows the simplification achieved by applying the dynamic programming relations in Theorems (3) and (4). 62

3-1 One dimensional conditional densities for the signal model defined in equation (3.11). The top plot shows the density for H_0 with a solid line. The densities for H_1 with $A = 1,\ 2,\ 4,$ and 6 become progressively wider. The second plot shows the corresponding log-likelihood ratios, $L(x|K = 1)$. In all plots, $\sigma^2 = 1$. 76

3-2 The top plot shows the two dimensional likelihood ratio when $A = 1$ and $\sigma^2 = 1$. The bottom plot shows the decision region boundary curves for $P_F = 10^{-4}$. In both cases, the noise variance is $\sigma^2 = 1$. The solid curve shows the boundary when $A = 1$ and the dashed curve shows the boundary when $A = 6$. 79

3-3 Normalized radial densities for $K = 10,\ 50,$ and 101. Each radial density arises from a Gaussian density for $\mathbf{x}_g$ with covariance matrix $\mathbf{I}$. The plotted curves are $\frac{P_{r|K,H}(r|K,H_0)}{\max(p_{r|K,H}(r|K,H_0))}$. Each appears similar to a Gaussian with standard deviation near $\frac{1}{2}$. 86

3-4 Approximation of $\log a_K(r)$ by a parabola. The plot displays a case when $K = 5$. The parabolic approximation is designed to fit well enough to determine thresholds necessary to achieve $P_F = 10^{-4}$. . . . 87

3-5 Approximation accuracy. The ratio of the actual false alarm rate, P_F, to the desired value $\tilde{P}_F$ is shown for $K = 5, 10, 20, 50, 100, 150, 200$. The desired false alarm rate varies from $\tilde{P}_F = 10^{-2}$ to $\tilde{P}_F = 10^{-12}$. . . 90

3-6 The receiver operating characteristic of the example problem is shown for $K = 5,\ 10,$ and 25. 91

C-1 Illustration of $A(\mathbf{x})$. In Case 1, the transformation produces $A(\mathbf{x})$ by translation parallel to the x_1 axis. In situations where such a translation does not intersect $J_-(c)$, $A(\mathbf{x})$ is symmetric with $\mathbf{x}$ about the x_2 value in the midpoint of the region common to $J_-(c)$ and $J_+(c)$. . . . 112

C-2 Situation if $\ell_+ > \ell_-$. The pont $\mathbf{x}'$ belongs to the parallelogram formed by $J_-(c)$ but does not belong to the parallelogram formed by $J_+(c)$. . 115

C-3 Case 1 of Lemma 2 proof. In this situation $A(\mathbf{x})$ can increase the x_2 component of the point. As a consequence, the point $\mathbf{x}'$ on the parallelogram formed by $J_+(c)$ is not in the parallelogram defined by $J_-(c)$. This contradicts the convexity of R. 116

C-4 Case 2 of the proof of Lemma 2. The point $\mathbf{x}'$ contradicts the convexity of R. 117

Chapter 1

Introduction

Data selection algorithms identify a subset of data for subsequent signal processing. A variety of hardware architectures amenable to data selection have been proposed. Specific examples include multiple-antenna wireless communication systems [24], and wireless sensor networks [19]. In these systems, collecting data can provide a large portion of the operating cost, so collecting only a subset of available measurements can yield significant resource savings. Additionally, if the data has sufficient redundancy, a small subset can yield performance close to that possible with the full data set.

In practice, selection algorithms should be tuned to a particular underlying signal processing task. In this chapter, we provide the necessary background for analyzing data selection algorithms in binary hypothesis testing, a mathematical formulation of signal detection. First, we summarize binary hypothesis testing, and develop notation to describe data selection. Second, we discuss potential applications for data selection, and research in a variety of statistical signal processing problems that is related to data selection.

1.1 Binary Hypothesis Testing for Signal Detection

Binary hypothesis testing is a mathematical formulation useful for developing decision-making algorithms. Binary hypothesis tests are procedures to determine whether a measurement is consistent with one of two states of the world, called hypotheses. They are denoted by H_0 and H_1. Conventionally, H_0 is called the null hypothesis, and H_1 is called the target hypothesis. The data that serves as the basis for the decision is collected in an N-dimensional random vector $\mathbf{x}$, and is described by an *a priori* conditional density under each hypothesis: $p_{\mathbf{x}|H_0}(\mathbf{x}|H_0)$ and $p_{\mathbf{x}|H_1}(\mathbf{x}|H_1)$. The conditional notation will be used whether the hypothesis is regarded as a parameter or a random variable.

Based upon the data $\mathbf{x}$ and the conditional probability densities, binary hypothesis testing consists of criteria for deciding whether the state of the world is H_0 or H_1. A decision rule is a mapping from the sample space for $\mathbf{x}$ to one of the two hypotheses. It is denoted by $\hat{H}(\mathbf{x})$, and it takes a value in the set $\{H_0, H_1\}$. Although $\hat{H}(\mathbf{x})$ can be a random variable in the most general cases, the discussion in this thesis focuses on situations where it is a deterministic mapping.

The decision rule $\hat{H}(x)$ cannot be evaluated without an optimality criterion. A variety of criteria have been advanced as a basis for judging the performance of decision rules. In this thesis, we focus on the well-known Neyman-Pearson criteria. This formulation of the decision problem regards the state of the world as an unknown, deterministic quantity. It attempts to strike a balance between errors regardless of which state actually holds. If H_0 is true, the case $\hat{H}(\mathbf{x}) = H_1$ is called a type I error. In the terminology established by a RADAR analogy, this error is also called a false alarm. When H_1 holds, the error $\hat{H}(x) = H_0$ is referred to as a type II error or a miss.

A more careful definition of the performance criteria requires detailed notation. Assume that the decision rule $\hat{H}(\mathbf{x})$ is a deterministic mapping from elements in the sample space for $\mathbf{x}$ to the two hypotheses. Then, the sample space can be divided

into two sets, given by

$$\begin{aligned} \hat{H}_0 &= \{\mathbf{x}|\hat{H}(\mathbf{x}) = H_0\} \\ \hat{H}_1 &= \{\mathbf{x}|\hat{H}(\mathbf{x}) = H_1\}. \end{aligned} \tag{1.1}$$

The performance criteria are established in terms of the probabilites for these two events. Specifically, the false alarm and miss probabilities are

$$\begin{aligned} P_F &= \Pr[\hat{H}_1|H_0] \\ P_M &= \Pr[\hat{H}_0|H_1]. \end{aligned} \tag{1.2}$$

The detection probability is $P_D = 1 - P_M$. By an examination of extreme cases, we see that there is an inherent tradeoff between P_M and P_F. If $\hat{H}_0$ is the empty set, $P_D = 1$ because the decision rule always says that a target is present when H_1 holds. Likewise, $P_F = 1$ for the same reason under H_0. Conversely, if $\hat{H}_1$ is the empty set, then both P_F and P_D are zero. A decision rule that randomly chooses between H_0 and H_1 without any use of $\mathbf{x}$ can get any values of (P_F, P_D) on the line $P_F = P_D$. Certainly, a decision procedure using $\mathbf{x}$ can achieve a more attractive tradeoff between miss and false alarm probabilities.

The Neyman-Pearson formulation of the decision problem attempts to get the best P_D at a fixed maximum P_F. The optimization problem is to find a decision rule such that P_D is maximized subject to the constraint $P_F \leq P'$. Such decision rules are useful in situations where there is a maximum tolerable false alarm rate. The formulation of the detector depends on the parameter P'. A curve expressing the relationship between P_F and P_D is called the receiver operating characteristic (ROC).

The solution to the basic Neyman-Pearson formulation of the binary hypothesis testing problem is expressed in the likelihood ratio test (LRT). Since P_F and P_D depend on the event $\hat{H}(\mathbf{x}) = H_1$ under different hypotheses, they can be expressed as integrals over the same set in sample space, ignoring the potential for randomness in the decision rule. A Lagrange multiplier argument leads to the LRT. The test is

denoted

$$L(\mathbf{x}) = \frac{p_{\mathbf{x}|H}(\mathbf{x}|H_1)}{p_{\mathbf{x}|H}(\mathbf{x}|H_0)} \underset{\hat{H}(\mathbf{y})=H_0}{\overset{\hat{H}(\mathbf{y})=H_1}{\gtrless}} \eta. \tag{1.3}$$

The test compares the likelihood ratio $L(\mathbf{x})$, the ratio of the *a priori* conditional densities to a threshold, declaring $\hat{H}(\mathbf{x}) = H_1$ for all $\mathbf{x}$ where the ratio exceeds η. In certain situations, the test is defined in terms of a threshold test on the log-likelihood $\ell(\mathbf{x}) = \log(L(\mathbf{x}))$. In either situation, a detector designer sets the threshold in order that the test have a specific false alarm probability, P_F.

The implementation of the LRT depends upon the specific densities $p_{\mathbf{x}|H}(\mathbf{x}|H_0)$ and $p_{\mathbf{x}|H}(\mathbf{x}|H_1)$. In many cases, the signal detection problem is translated into a hypothesis testing problem by assuming plausible models for these densities. In a signal detection model, we are are searching for a target signal in a set of noisy measurements. In this case, under the null hypothesis H_0, $\mathbf{x} = \mathbf{n}$, where $\mathbf{n}$ is a random vector representing noise. Frequently, we shall assume that $\mathbf{n}$ is a zero-mean Gaussian random vector with a known covariance matrix, written as $\mathbf{\Lambda}$. Under the signal hypothesis H_1, the data is modeled by $\mathbf{x} = \mathbf{s} + \mathbf{n}$, where $\mathbf{s}$ is a representation of the signal. Much of the variation in detection based upon binary hypothesis testing comes from various models for $\mathbf{s}$: whether it is a deterministically known vector, a random vector, or a random vector with an element of uncertainty in its probability density, either parameterized or unparameterized.

1.2 Data Selection Problem Formulation

The research presented in this thesis examines binary hypothesis testing in the context of an additional constraint on the decision algorithm, intended to account for resource limits in the hardware implementation of the test. Assume that the details of the model for $\mathbf{s}$ are fixed, either as a deterministic or random vector. The data selection constraint changes the detection problem to one where only a subset of the available data can be examined in order to make a decision. The detector for a fixed subset choice is an LRT; the critical issue involves choice of the best such subset. This

section develops a notation to formalize the problem.

The entire set of data that can be read or collected given unlimited hardware resources is given by the random vector $\mathbf{x}$. The selection model permits transmission of a subset of the data, along with indices identifying the measurements. We develop notations that describe the selection, with the index of the selected samples within $\mathbf{x}$ either preserved or suppressed. As a source of the identity-preserving notation, consider the set of diagonal matrices with boolean entries, such that the weight of the main diagonal is constrained to be K. An example may be

$$\mathbf{G} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix}. \tag{1.4}$$

The notation for a selected data set wtih the identity of the selected measurements preserved is

$$\mathbf{x}_g = \mathbf{G}\mathbf{x}. \tag{1.5}$$

Using the matrix displayed in equation (1.4), the selected data is $\mathbf{x}_g = [x_1\ 0\ x_2\ x_3\ 0]^T$. Observation of this vector indicates the values of the selected measurements, as well as their index in the total data set $\mathbf{x}$.

A compact notation for selection can be denoted by matrix multiplication between $\mathbf{x}$ and a non-square boolean matrix. Let $\tilde{\mathbf{G}}$ be a $K \times N$ matrix restricted such that each row has a single non-zero entry, and each column contains at most one non-zero entry. An example corresponding to the same selection measurements shown in equation (1.4) is

$$\tilde{\mathbf{G}} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \end{bmatrix}. \tag{1.6}$$

Likewise, the selected data vector is denoted by

$$\tilde{\mathbf{x}}_g = \tilde{\mathbf{G}}\mathbf{x}. \tag{1.7}$$

This alternative notation is not unique; the selected subset is unchanged by any permutation of the rows of $\tilde{\mathbf{G}}$. Given only $\tilde{\mathbf{x}}_g$, the identity of the data in the original vector $\mathbf{x}$ cannot be determined.

In this notational framework, we can precisely define a detection with data selection problem. The subset size constraint restricts the number of non-zero entries in the selection matrix $\mathbf{G}$ or $\tilde{\mathbf{G}}$ to be K. Subject to this constraint, we must choose a way to get $\mathbf{G}$ that produces the best P_D operating point for a fixed maximum value of P_F. In subsequent chapters, we consider deterministic and random techinques for choosing $\mathbf{G}$. They will also consider a variety of models for $\mathbf{s}$. Varying assumptions about the signal model and selection procedure lead to several distinct problems springing from the subset selection constraint in detection.

1.3 Applications and Connections to Previous Research

The data selection problem formulated in this chapter modifies the traditional formulation of binary hypothesis testing with an additional constraint. The utility of data selection constraint depends upon the hardware characteristics of the system implementing the hypothesis test. Specifically, data selection can be applied to distributed signal processing systems, where signal measurement and processing are separated by a communication network. In many such systems, the communication network may form the bottleneck determining the overall performance [5]. This occurs if the majority of the energy consumption or time delay in the signal processing algorithm arises from congestion it injects in the network. Data selection reduces this congestion, thereby reducing the implementation cost of the algorithm. Additionally, in some situations, many of the measurements are redundant. Discarding them may not

seriously impair the performance of the statisitical signal processing algorithm. Even though the diminishing returns to selecting ever larger data subsets does not always hold, selection remains useful whenever its efficiency benefit outweighs the cost of lost information in the discarded data.

A variety of hardware architectures amenable to data selection have been proposed. Specific examples of hardware systems for which data selection has been suggested include multiple-antenna wireless communication systems [15, 24], and wireless sensor networks [19, 32] for applications such as target detection [7], tracking [3, 45], and classification [20]. In these systems, communicating data can provide a large portion of the operating cost, so collecting a subset of available measurements can yield significant resource savings. Additionally, if the sensors or antenna elements are sufficiently abundant, a small subset of the data can yield performance close to that possible with all the data.

Data selection is potential technique for improving the efficiency of a distributed signal processing system. It can complement several related lines of research in the design of distributed signal processing hardware and algorithms [34].

Careful design of the hardware topology of the distributed signal processing system is sometimes feasible. In situations where fine-grained control of the hardware and software is available, careful joint design of the hardware communication infrastructure and the assignment of algorithms to the distributed processing elements can benefit the system performance [12, 13].

Information theoretic analysis of combined signal quantization and processing provides an alternative approach to understanding the impact of commuication constraints on distributed signal processing. In fact, data selection can be viewed as a form of a quantization for spatially-distributed data. There are several strands of research considering the joint compression and processing of information. One such technique formulates the problem as a source coding problem using the tools of information theory [9]. If the ultimate signal processing application is estimation, the class of joint compression and processing problems are referred to as the CEO problem. The dependence of the CEO problem on network topology is considered by [11].

Additionally, there has been research into information theoretic performance bounds on joint compression and classification [1].

A second approach to the problem of quantization in detection has been termed distributed detection [38, 29]. This formulation of the problem considers the joint design of local quantizers at a set of spatially separated sensors and a global decision rule based on the quantized data, frequently taking the data rate available to each sensor as fixed. Viewed from this perspective, data selection is a compression rule where unequal data rates are allocated to each sensor. In cases where data from separate sensors is correlated, it has been found that unequal allocation of data rate between sensors can perform better than equal allocations [4]. Additionally, some research has examined data-dependent selection in the distributed detection model [35].

Distributed signal processing systems are not the only potential motivation for data selection algorithms. In a variety of applications, selection can improve the efficiency of an underlying signal processing algorithm. As an example, consider a signal approximation problem in source coding. With a redundant signal set, it may be of interest to find an accurate representation of a signal in which a certain fraction of the expansion coefficients are zero. In this situation, the reduced number of coefficients limits the cost of storing or communicating the approximate signal. There are a wide variety of algorithms to accomplish this task, such as matching pursuit [21]. Data selection procedures have been applied to many other signal processing tasks, including filter design [22, 36, 42], statistical regression [23], estimation [18], and feature selection for data mining [6, 26, 27].

1.4 Outline

The remainder of the thesis applies the data selection constraint to two useful binary hypothesis testing models. Chapter 2 discusses detection of a deterministic target signal in correlated Gaussian noise. The data selection constraint alters the familar matched filter solution to the problem, instead producing an algorithm we call the

restricted matched filter. Chapter 3 considers detection of a stochastic signal. The chapter focuses on randomized selection rules and considers the robustness of the resulting detector to uncertainties in the probabilistic model for the target signal.

Chapter 2

Data Selection in Deterministic Signal Detection

A classic problem in detection theory involves detecting a deterministic signal in Gaussian noise. The well-known solution to this problem, the matched filter, has applications such as receivers for communications and RADAR systems. This chapter considers the matched filter algorithm when combined with a data selection constraint. It specifies the signal modeling assumptions that lead to the matched filter detector, describes the data selection constraint, and considers solutions to the restricted problem.

2.1 Traditional Matched Filtering

As an example of the application of binary hypothesis testing, consider a familiar model for the detection of a target. Under hypothesis H_1, the target produces a known signal $\mathbf{s}$, representing the measurements of the energy it radiates. For example, it may represent the samples of a return from a target to a RADAR receiver or an antenna array. The measurements in $\mathbf{s}$ are corrupted by colored noise $\mathbf{n}$, representing the affect of background interference and measurement noise. For simplicity, $\mathbf{n}$ is modeled as a zero-mean, Gaussian random vector with covariance matrix $\mathbf{\Lambda}$. The density of a Gaussian random vector $\mathbf{n}$ with mean $\mathbf{m}$ and covariance matrix $\mathbf{\Lambda}$ is

denoted as

$$p_{\mathbf{n}}(\mathbf{n}) = \mathcal{N}(\mathbf{n}, \mathbf{m}, \mathbf{\Lambda}). \tag{2.1}$$

The noise $\mathbf{n}$ is independent of the hypothesis H. The measurement model for each hypothesis is

$$\mathbf{x} = \mathbf{n} \quad H = H_0 \tag{2.2}$$

$$\mathbf{x} = \mathbf{s} + \mathbf{n} \quad H = H_1. \tag{2.3}$$

With this model, the conditional probability densities of $\mathbf{x}$ under both hypotheses are

$$p_{\mathbf{x}|H}(\mathbf{x}|H_0) = \mathcal{N}(\mathbf{x}, 0, \mathbf{\Lambda}) \tag{2.4}$$

$$p_{\mathbf{x}|H}(\mathbf{x}|H_1) = \mathcal{N}(\mathbf{x}, \mathbf{s}, \mathbf{\Lambda}). \tag{2.5}$$

In this binary hypothesis test model, the vector $\mathbf{s}$ and covariance matrix $\mathbf{\Lambda}$ are known to the receiver. Additionally, we require that the covariance matrix $\mathbf{\Lambda}$ is positive-definite. Though we motivate this model using RADAR systems, it applies to other situations such as detecting information-bearing signals in a pulse-amplitude modulation communication system.

The likelihood ratio test for the signal model described by equations (2.2) and (2.3) is typically referred to as the whitened matched filter [31]. The detector structure can be decomposed into three elements, a whitening transformation, a projection, and a threshold operation.

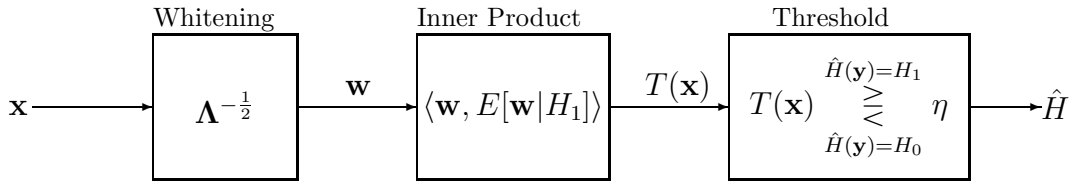

Figure 2-1: Whitened Matched Filter Detector based on $\mathbf{x}$.

The first block of the detector whitens the observations. The output is

$$\mathbf{w} = \mathbf{\Lambda}^{-\frac{1}{2}}\mathbf{x}. \tag{2.6}$$

Here, the matrix $\mathbf{\Lambda}^{-\frac{1}{2}}$ is the unique, postive-definite square-root matrix of $\mathbf{\Lambda}^{-1}$. One can verify that the covariance of $\mathbf{w}$ is the identity matrix. In principle, other linear transformations also produce white output; this transformation is shown for concreteness.

The second block in Figure 2-1 is an inner product with the conditional mean of the whitened observations $E[\mathbf{w}|H_1]$. This produces a sufficient statistic for the detection, the linear function $T(\mathbf{x})$. The final block compares $T(\mathbf{x})$ to a threshold to determine if the detector's decision is $\hat{H} = H_1$ or $\hat{H} = H_0$. The threshold η determines the (P_F, P_D) operating point on the ROC.

An alternative view of the whitening transformation and the inner product provides additional insight into the computation of $T(\mathbf{x})$. Since the whitening transformation is applied to the random vector $\mathbf{x}$, it is also applied to the mean, yielding $E[\mathbf{w}|H_1] = \mathbf{\Lambda}^{-\frac{1}{2}}\mathbf{s}$. Using this relation, the sufficient statistic for detection becomes

$$T(\mathbf{x}) = \mathbf{s}^T\mathbf{\Lambda}^{-1}\mathbf{x}. \tag{2.7}$$

This form of the inner product produces a simple analysis of the general hypothesis testing problem in terms of a scalar, Gaussian hypothesis testing problem.

The probability density for $T(\mathbf{x})$, conditioned on either hypothesis, is a scalar Gaussian density since the statistic is a linear combination of jointly Gaussian random variables. The conditional variance is the same for both hypotheses. The conditional densities, however, are distinguishable by their means. Under $H = H_0$, $T(\mathbf{x})$ is zero-mean. Under $H = H_1$, the mean is

$$E[T(\mathbf{x})|H_1] = \mathbf{s}^T\mathbf{\Lambda}^{-1}\mathbf{s}. \tag{2.8}$$

The parameter in equation (2.8) generalizes the notion of signal-to-noise ratio devel-

oped for detecting known signals in white Gaussian noise. In fact, when $\mathbf{\Lambda} = \sigma^2\mathbf{I}$, the conditional expectation $E[T(\mathbf{x})|H_1]$ reduces to the traditional value for signal-to-noise ratio. Throughout the remainder of this thesis, we refer to this conditional expectation as SNR. As in the white noise case, this parameter fully determines the ROC. Thus, for a fixed test threshold in η, the detector performance is

$$\begin{aligned} P_F &= Q(\eta) \\ P_D &= Q(\eta - \text{SNR}), \end{aligned} \tag{2.9}$$

where the Q function is given by

$$Q(\alpha) = \frac{1}{\sqrt{2\pi}} \int_{\alpha}^{\infty} e^{-\frac{u^2}{2}} du. \tag{2.10}$$

Two basic properties of the Q function are invertibility and monotonicity. The function is invertible over the interval $0 \leq \alpha \leq 1$. Additionally, it is monotonically decreasing, so $Q(\alpha_1) > Q(\alpha_2)$ for all argument pairs satisfying $\alpha_1 < \alpha_2$. These properties are important for analyzing the subset selection problem for this detector.

2.2 Restricted Matched Filter

The general data selection problem in section 1.2 generalizes the Neyman-Pearson criteria. In addition to a constraint on the maximum false alarm rate, it fixes the size of the data subset $\tilde{\mathbf{x}}_g$ available to the detector. Fortunately, for the hypothesis test considered in this chapter, the subset size constraint does not affect the form of the detector once the selected subset is fixed. We denoted the selected subset through the $\mathbf{G}$ notation developed in section 1.2. The conditional densities under H_0 and H_1 for the selected measurements in $\tilde{\mathbf{x}}_g$ are still Gaussian, though the application of $\mathbf{G}$ alters the mean and covariance. We refer to a whitened matched filter designed according to the statistics of a subset of the data as a restricted matched filter (RMF).

For a fixed subset indicated by $\mathbf{G}$, the RMF is designed according to equation (2.7) using the covariance and the conditional mean of the random vector $\tilde{\mathbf{x}}_g$. The

new mean and covariance are

$$E[\tilde{\mathbf{x}}_g|H_1] = \tilde{\mathbf{s}}_g \tag{2.11}$$

$$\mathbf{\Lambda}_{\tilde{\mathbf{x}}_g} = \tilde{\mathbf{G}}\mathbf{\Lambda}\tilde{\mathbf{G}}^T. \tag{2.12}$$

The conditional densities under H_0 and H_1 are

$$p_{\tilde{\mathbf{x}}_g|H}(\tilde{\mathbf{x}}_g|H_0) = \mathcal{N}(\tilde{\mathbf{x}}_g, 0, \mathbf{\Lambda}_{\tilde{\mathbf{x}}_g}) \tag{2.13}$$

$$p_{\tilde{\mathbf{x}}_g|H}(\tilde{\mathbf{x}}_g|H_1) = \mathcal{N}(\tilde{\mathbf{x}}_g, \tilde{\mathbf{s}}_g, \mathbf{\Lambda}_{\tilde{\mathbf{x}}_g}). \tag{2.14}$$

Since the probabilistic description of the hypothesis test after data selection has the same structure as the original test, the optimal detector is the whitened matched filter determined by $\tilde{\mathbf{s}}_g$ and $\mathbf{\Lambda}_{\tilde{\mathbf{x}}_g}$. Thus, for a fixed subset of measurements, represented by a particular instance of $\mathbf{G}$, the performance of the associated whitened matched filter is given by the quadratic form

$$\text{SNR}(\mathbf{G}) = \tilde{\mathbf{s}}_g^T \mathbf{\Lambda}_{\tilde{\mathbf{x}}_g}^{-1} \tilde{\mathbf{s}}_g. \tag{2.15}$$

Throughout the remainder of the thesis, the notation in equation (2.15) represents the SNR quadratic form for a subset of sensors indicated either symbolically as $\mathbf{G}$ or explicitly, such as $\{x_1, x_3, x_8\}$.

The equations (2.9) and (2.15) can evaluate the performance of the detector for any choice of the selection matrix $\mathbf{G}$ or $\tilde{\mathbf{G}}$. The following theorem provides a criteria for the best subset of measurements given constraints on the false alarm probability and the selected subset size.

Theorem 1 (Restricted Matched Filter) *Consider the binary hypothesis testing model in equations (2.4 - 2.5). With a fixed false alarm rate of the detector, P_F, and a fixed number of non-zero entries in $\mathbf{G}$, K, the selection matrix that maximizes P_D under these constraints maximizes* $\text{SNR}(\mathbf{G})$ *given in equation (2.15).*

Proof. Since the Q-function is invertible, the constraint on the false alarm rate fixes

the detector threshold at $\eta(P_F) = Q^{-1}(P_F)$. The detection probability depends on $\eta(P_F)$ and SNR($\mathbf{G}$). Increasing SNR($\mathbf{G}$) increases P_D since the derivative $Q'(\alpha) < 0$ for all $\alpha \in \Re$. The covariance matrix $\mathbf{\Lambda}$ is positive definite, so SNR($\mathbf{G}$) is finite for any selection matrix with K non-zero entries. There are a finite number of feasible selection matrices, so SNR($\mathbf{G}$) has an absolute maximum under the constraints, and the matrix that maximizes SNR also maximizes P_D. The optimal $\mathbf{G}$, however, is not necessarily unique. $\square$.

Theorem 1 shows that selecting the data subset yielding the best RMF is a combinatorial optimization problem. There are $\begin{pmatrix} N \\ K \end{pmatrix}$ selection matrices that satisfy the subset size constraint. Optimization of SNR over this finite set produces the best detector. Similar optimization problems have been considered by researchers in other fields, and are frequently called combinatorial feature selection problems [6]. A variety of heuristic solutions to related optimization problems have been suggested [23, 27]. Common approaches involve tests for local optima, branch and bound search [27], and incremental searches that add measurements to maximize the change in SNR, referred to as greedy algorithms.

2.3 Example

Before analyzing specific optimization algorithms, it is important to establish the impact of subset selection on the detector performance. In some situations, the gap between the best and worst values of SNR($\mathbf{G}$) can be large, indicating the need for careful selection of the data subset processed by the RMF.

In the case where the $\mathbf{n}$ is white, the optimal subset selection can be made by inspection. If the covariance matrix satisfies $\mathbf{\Lambda} = a\mathbf{I}$, the signal-to-noise ratio for any subset is

$$\text{SNR}(\mathbf{G}) = \frac{1}{a}\|\tilde{\mathbf{s}}_g\|^2 = \frac{1}{a}\sum_i G_{ii}s_i^2. \tag{2.16}$$

The maximum energy selection rule, setting $G_{ii} = 1$ for the K largest values of

s_i^2, maximizes SNR($\mathbf{G}$), as indicated by Welborn [43]. This result generalizes to any diagonal covariance matrix by substituting the SNR metrics of the individual measurements for s_i^2.

For covariance matrices that are not diagonal, the maximum energy selection rule no longer guarantees an optimum, and the covariance matrix can have a large effect on the distribution of SNR($\mathbf{G}$) over the possible values of $\mathbf{G}$ with K non-zero entries. There is an interesting trade-off between the energy of the selected signal $\tilde{\mathbf{s}}_g$ and the effect of the covariance $\mathbf{\Lambda}_{\tilde{\mathbf{x}}_g}$. This section illustrates the trade-off with examples of the RMF highlighting the impact of $\mathbf{\Lambda}_{\tilde{\mathbf{x}}_g}$ on the detector performance. In order to specify the detection problem precisely, we must define the mean vector $\mathbf{s}$ under $H = H_1$ and the noise covariance $\mathbf{\Lambda}$. Assume that $N = 16$ and the mean vector is

$$\mathbf{s}^T = \begin{bmatrix} 1 & \dots & 1 \end{bmatrix}. \tag{2.17}$$

Since the target signal is constant, the signal energy is $\|\tilde{\mathbf{s}}_g\|^2 = K$ for any selected subset. The variation in SNR($\mathbf{G}$) comes from the covariance alone.

In the first example, we choose the covariance matrix as a Toeplitz matrix defined by the first-order covariance sequence

$$K_{nn}[\alpha] = c^{-|\alpha|}. \tag{2.18}$$

The parameter c^{-1} determines the decay rate of the covariance function. In the example simulations, we've taken $c = 0.95$. The covariance matrix for is defined by the relation

$$[\mathbf{\Lambda}]_{ij} = K_{nn}[|i - j|]. \tag{2.19}$$

Comparison of the SNR quadratic forms for two specific subsets reveals the impact of subset selection on the covariance matrix $\mathbf{\Lambda}_{\tilde{\mathbf{x}}_g}$. The first subset selects the beginning eight measurements in $\mathbf{x}$, as defined by our arbitrary indexing. The measurements are

$$\tilde{\mathbf{x}}_{g,1}^T = \begin{bmatrix} x_1 & x_2 & \dots & x_8 \end{bmatrix}. \tag{2.20}$$

This set produces a performance metric SNR = 1.242.

The second subset selects measurements from the odd-indexed sensors. The second measurement vector is

$$\tilde{\mathbf{x}}_{g,2}^T = \begin{bmatrix} x_1 & x_3 & \dots & x_{15} \end{bmatrix}. \tag{2.21}$$

This sensor selection essentially changes the decay rate of the covariance function to c^{-2}. The second set of measurements decorrelate much faster than those in $\mathbf{G}_1$. The performance metric for the second set of measurements is $\text{SNR}(\mathbf{G}_2) = 1.430$. The metric in the second case has been increased by about 15 percent by a different choice of data.

This result can be explained in terms of the eigenvalues and eigenvectors of the covariance matrices $\mathbf{\Lambda}_{\tilde{\mathbf{x}}_g}$. The covariance matrices can be diagonalized

$$\mathbf{\Lambda}_{\tilde{\mathbf{x}}_g} = \mathbf{U}\mathbf{D}\mathbf{U}^T. \tag{2.22}$$

The matrix $\mathbf{U}$ has normalized eigenvectors of $\mathbf{\Lambda}_{\tilde{\mathbf{x}}_g}$ as columns and $\mathbf{D}$ is a diagonal matrix of the corresponding eigenvalues. These matrices are dependent on the sensor measurements selected by $\mathbf{G}$ that produce $\tilde{\mathbf{x}}_{g,1}$ and $\tilde{\mathbf{x}}_{g,2}$.

Substituting the result of equation (2.22) into the expression for the performance metric SNR in equation (2.15), produces an alternative equation for the performance metric

$$\text{SNR}(\mathbf{G}) = \sum_{i=1}^{K} \frac{\langle \mathbf{u}_{\mathbf{G},i}, \tilde{\mathbf{s}}_g \rangle^2}{d_i(\mathbf{G})}. \tag{2.23}$$

Here, $\mathbf{u}_{\mathbf{G},i}$ is an eigenvector of $\mathbf{\Lambda}_{\tilde{\mathbf{x}}_g}$ and $d_i(\mathbf{G})$ is the corresponding eigenvalue. The result indicates that there is a tradeoff between the projection of the conditional mean vector $\tilde{\mathbf{s}}_g$ onto the eigenvectors and the corresponding eigenvalues. In order to maximize SNR, the subset selection should produce a mean that has its largest projection along eigenvectors with minimal eigenvalues. The subset selection impacts both the mean vector and the eigenvalues and eigenvectors of the selected data, so general selection rules that focus only on signal energy or noise correlation can perform

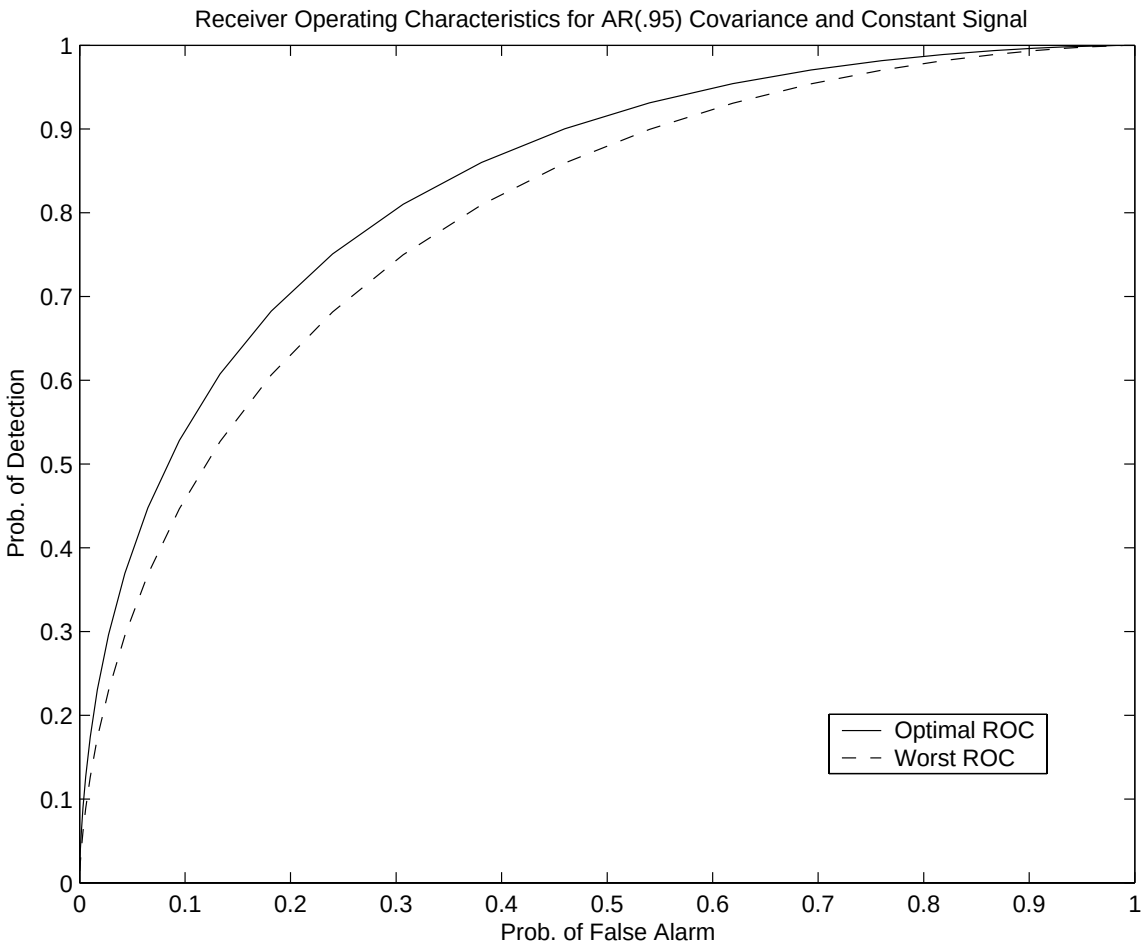

Figure 2-2: Receiver operating characteristics for best and worst detectors. The best SNR is 1.384, and the worst is 1.179.

suboptimally.

The example illustrates how the noise covariance influences the overall performance of the detector, but does not exhibit the potential magnitude of the effect. The best and worst ROCs for this example, displayed in Figure 2-2, show a small worst-case gap in P_D. Instances of the RMF with a different covariance matrix can produce a large gap between the best and worst ROCs. The ROCs in Figure 2-3 result when the covariance matrix is replaced with a randomly generated covariance matrix. In the instance tested, the maximum and minimum SNR are 1.8085 and 0.1174, respectively. There is a dramatic gap in P_D for many values of P_F. Although performance gaps this large do not occur in all cases, the example shows that selec-

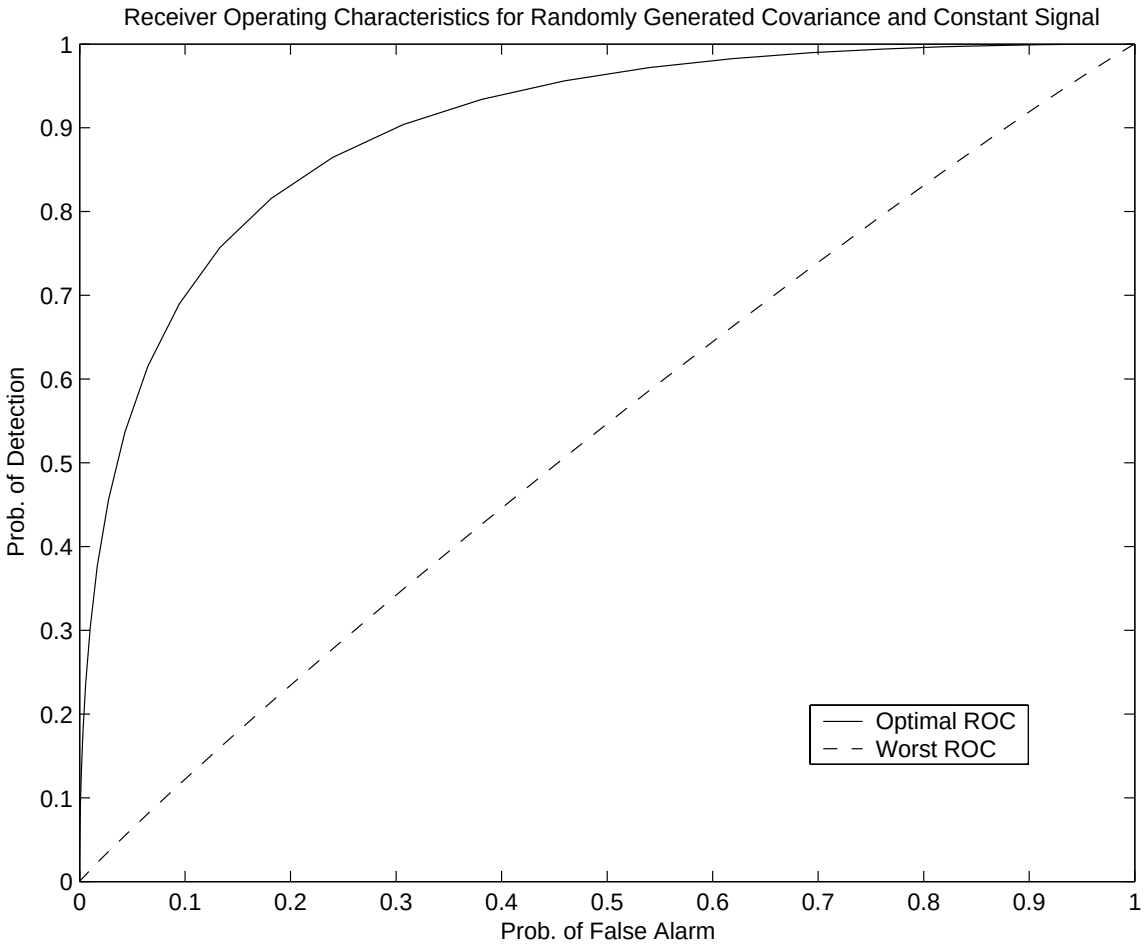

Figure 2-3: Receiver operating characteristics for best and worst detectors generated for an RMF instance with $N = 16$, $K = 8$, a constant signal, and a randomly generated covariance matrix. The best SNR is 1.8085, and the worst is 0.1174.

tion algorithms can make a significant impact on RMF performance under the right circumstances.

2.4 Greedy Algorithms for the RMF Search

As indicated in the prior example, a straightforward selection rule such as signal energy maximization may not always identify the optimal subset of measurements for the RMF. In situations where $\mathbf{\Lambda}$ is not diagonal, more complicated heuristic rules for data selection can improve the performance of the selected subset. A variety of heuristic approaches have been proposed for data selection for filter design [22], regression [23], and feature selection [27]. The work in these areas has considered several heuristic rules for maximizing the quadratic form in equation (2.15), providing algorithms that also address the search for a good RMF. For example, Miller [23] suggests a local search algorithm. The algorithm exchanges a selected datum for one not selected, terminating when none of the potential exchanges with the currently selected subset increase the SNR. Narendra and Fukunaga [27] discuss a branch-and-bound heuristic. They propose discarding data from the original set of N measurements, at each stage dropping the measurement that least reduces SNR. They suggest an exhaustive search of the possible K measurement subsets, and rely on the fact that many subsets may not be examined in detail because the search can be terminated early in cases when dropping a measurement reduces SNR below the best K measurement subset previously examined.

The greedy selection rule forms a critical component of these heuristic search algorithms. In each stage of the algorithms, measurements are added to optimize the incremental change in SNR given the current state of the subset. The overall performance of the algorithms depends upon the greedy selection rule as well as the termination criteria. An analysis of the greedy selection procedure, however, can illuminate certain aspects of the performance of these algorithms, since it produces the initial subsets examined by the more detailed heuristic search algorithms.

This section focuses on forward and backward greedy selection algorithms. The

forward greedy algorithm resembles the local search optimization. It builds a selected subset by adding measurements that maximize the incremental increase in SNR, terminating after selecting a K element subset. Frequently, this algorithm serves as the initialization for the local search procedure. The backward greedy algorithm discards measurements, minimizing incremental loss in SNR. It terminates after discarding $N-K$ measurements. This procedure selects the first subset inspected by the branch and bound algorithm.

2.4.1 Expressions for Incremental SNR Changes During Greedy RMF Searches

In order to use a greedy approach to the RMF optimization problem, we require an incremental expression for the SNR, highlighting the effect of adding or subtracting a measurement from the selected subset. With such an expression, the most beneficial incremental change to a specific selected subset can be identified.

An incremental expression for the quadratic form can be derived using a block decomposition of the signal vector $\mathbf{s}$ and the covariance matrix $\mathbf{\Lambda}$. Consider a block decomposition [14] of the covariance matrix given by

$$\mathbf{\Lambda} = \begin{bmatrix} \mathbf{\Lambda}_K & \mathbf{\Lambda}_B \\ \mathbf{\Lambda}_B^T & \mathbf{\Lambda}_R \end{bmatrix} \tag{2.24}$$

where $\mathbf{\Lambda}_K$ is the upper-left $K \times K$ block of the matrix, $\mathbf{\Lambda}_R$ is the lower-right $(N-K) \times (N-K)$ block, and $\mathbf{\Lambda}_B$ is the $K \times (N-K)$ off-diagonal block. When $\mathbf{\Lambda}_K$ is invertible, the inverse of the matrix can be expressed as

$$\mathbf{\Lambda}^{-1} = \begin{bmatrix} \mathbf{\Lambda}_K^{-1} + \mathbf{\Lambda}_K^{-1}\mathbf{\Lambda}_B(\mathbf{\Lambda}_R - \mathbf{\Lambda}_B^T\mathbf{\Lambda}_K^{-1}\mathbf{\Lambda}_B)^{-1}\mathbf{\Lambda}_B^T\mathbf{\Lambda}_K^{-1} & -\mathbf{\Lambda}_K^{-1}\mathbf{\Lambda}_B(\mathbf{\Lambda}_R - \mathbf{\Lambda}_B^T\mathbf{\Lambda}_K^{-1}\mathbf{\Lambda}_B)^{-1} \\ -(\mathbf{\Lambda}_R - \mathbf{\Lambda}_B^T\mathbf{\Lambda}_K^{-1}\mathbf{\Lambda}_B)^{-1}\mathbf{\Lambda}_B^T\mathbf{\Lambda}_K^{-1} & (\mathbf{\Lambda}_R - \mathbf{\Lambda}_B^T\mathbf{\Lambda}_K^{-1}\mathbf{\Lambda}_B)^{-1} \end{bmatrix}. \tag{2.25}$$

If we apply a block decomposition a vector $\mathbf{s}^T = [\mathbf{s}_K^T \mathbf{s}_R^T]$, the quadratic form between

$\mathbf{s}$ and $\mathbf{\Lambda}$ can be written as

$$\begin{aligned}\mathbf{s}^T\mathbf{\Lambda}^{-1}\mathbf{s} &= \mathbf{s}_K^T\mathbf{\Lambda}_K^{-1}\mathbf{s}_K + \mathbf{s}_K^T\mathbf{\Lambda}_K^{-1}\mathbf{\Lambda}_B(\mathbf{\Lambda}_R - \mathbf{\Lambda}_B^T\mathbf{\Lambda}_K^{-1}\mathbf{\Lambda}_B)^{-1}\mathbf{\Lambda}_B^T\mathbf{\Lambda}_K^{-1}\mathbf{s}_K \\ &\quad -2\mathbf{s}_K^T\mathbf{\Lambda}_K^{-1}\mathbf{\Lambda}_B(\mathbf{\Lambda}_R - \mathbf{\Lambda}_B^T\mathbf{\Lambda}_K^{-1}\mathbf{\Lambda}_B)^{-1}\mathbf{s}_R + \mathbf{s}_R^T(\mathbf{\Lambda}_R - \mathbf{\Lambda}_B^T\mathbf{\Lambda}_K^{-1}\mathbf{\Lambda}_B)^{-1}\mathbf{s}_R.\end{aligned} \tag{2.26}$$

This equation relates the two quadratic forms of different dimensions, $\mathbf{s}^T\mathbf{\Lambda}^{-1}\mathbf{s}$ and $\mathbf{s}_K^T\mathbf{\Lambda}_K^{-1}\mathbf{s}_K$, and can be used to evaluate the impact of incremental changes in the selected data.

The expression (2.26) can be used to determine the change in SNR from adding or deleting a measurement from the selected subset. Let $\mathbf{G}$ denote a subset, and consider the change in SNR from adding measurement i, such that $\mathbf{G}_{ii} = 0$. Let $\mathbf{z}_i = \tilde{\mathbf{G}}\mathbf{\Lambda}\mathbf{e}_i$, where the only non-zero entry of the vector vector $\mathbf{e}$ occupies the ith row, i.e. the entries of the vector are Kroenecker deltas $\mathbf{e}_{ij} = \delta_{ij}$. Thus, $\mathbf{z}_i$ contains the elements of the ith column of $\mathbf{\Lambda}$ corresponding to the measurements selected by $\mathbf{G}$. This vector corresponds to $\mathbf{\Lambda}_B$ in equation (2.26). The covariance matrix for the selected subset $\mathbf{\Lambda}_{\tilde{\mathbf{x}}_g}$ corresponds to $\mathbf{\Lambda}_K$, and so forth. Making the proper substitutions into the equation for the quadratic form, the SNR is

$$\mathrm{SNR}(\{\mathbf{G} \cup i\}) = \tilde{\mathbf{s}}_g^T\mathbf{\Lambda}_{\tilde{\mathbf{x}}_g}^{-1}\tilde{\mathbf{s}}_g + \frac{\left(\mathbf{s}_i - \mathbf{z}_i^T\mathbf{\Lambda}_{\tilde{\mathbf{x}}_g}^{-1}\tilde{\mathbf{s}}_g\right)^2}{\mathbf{\Lambda}_{ii} - \mathbf{z}_i^T\mathbf{\Lambda}_{\tilde{\mathbf{x}}_g}^{-1}\mathbf{z}_i}. \tag{2.27}$$

The forward SNR increment, the change in SNR due to adding measurement i to the subset identified in $\mathbf{G}$, is given by

$$\Delta^+(\mathbf{G}, i) = \frac{\left(\mathbf{s}_i - \mathbf{z}_i^T\mathbf{\Lambda}_{\tilde{\mathbf{x}}_g}^{-1}\tilde{\mathbf{s}}_g\right)^2}{\mathbf{\Lambda}_{ii} - \mathbf{z}_i^T\mathbf{\Lambda}_{\tilde{\mathbf{x}}_g}^{-1}\mathbf{z}_i}. \tag{2.28}$$

The SNR increment when a measurement is discarded from a selected subset has a similar form. For a subset where $\mathbf{G}_{jj} = 1$, let $\mathbf{G}'(j)$ indicate the selection matrix formed when j is discarded. The reduction in SNR is $\Delta^+(\mathbf{G}'(j), j)$ using the form of equation (2.28). For convenience, the SNR reduction is also written as $\Delta^-(\mathbf{G}, j)$. An

equivalent form for this reduction,

$$\Delta^-(\mathbf{G}, j) = \frac{\left(\mathbf{e}_j^T \mathbf{\Lambda}_{\tilde{\mathbf{x}}_g}^{-1} \tilde{\mathbf{s}}_g\right)^2}{\mathbf{e}_j^T \mathbf{\Lambda}_{\tilde{\mathbf{x}}_g}^{-1} \mathbf{e}_j}, \tag{2.29}$$

is derived in [27]. Note that in this equation, $\tilde{\mathbf{x}}_g$ and $\mathbf{\Lambda}_{\tilde{\mathbf{x}}_g}$ are determined by the $\mathbf{G}$ at the beginning of the stage, including the index for the measurement eventually discarded.

The greedy algorithms determine the measurement to select or discard by computing the SNR increments for each possible measurement. The forward greedy algorithm begins with no selected measurements and adds the unselected measurement that maximizes $\Delta^+(\mathbf{G}, i)$; the backward greedy algorithm begins with $\mathbf{G} = \mathbf{I}$, and removes the measurement that minimizes $\Delta^-(\mathbf{G}, i)$. These stages are repeated until a subset of K measurements remains.

The worst-case computational complexity for a stage of the forward and backward greedy algorithms are bounded by the same quantity. Equation (2.29) for the SNR reduction shows that the computation of $\Delta^-(\mathbf{G}, j)$ for all indices remaining in the subset $\mathbf{G}$ is dominated by the matrix inversion of $\mathbf{\Lambda}_{\tilde{\mathbf{x}}_g}$. The computational complexity of this operation for a subset of size K is $O(K^3)$ for a general positive-definite covariance matrix. For the forward greedy algorithm, the complexity for calculating $\Delta^+(\mathbf{G}, i)$ for each i not included in the subset $\mathbf{G}$ is dominated by the computation of K quadratic forms given by $\mathbf{z}_i^T \mathbf{\Lambda}_{\tilde{\mathbf{x}}_g}^{-1} \mathbf{z}_i$. The computational complexity of this operation is also $O(K^3)$. The total computational complexity of these algorithms is the sum of the complexities for several stages. For both algorithms, the upper bound on the total computational complexity to select a subset of size K is $O(K^4)$.

2.4.2 Conditions for An Exact Solution with Maximum Energy Selection

The computational complexity of the greedy selection algorithms exceeds the complexity of the maximum energy selection rule by a significant margin. In cases where

the maximum energy selection rule performs well, the added complexity of the greedy selection algorithm is unnecessary. In one instance, where $\mathbf{\Lambda} = \sigma^2\mathbf{I}$, the maximum energy rule selects the optimal subset. Quantifying the robustness of this result identifies situations where the greedy selection algorithms are unnecessary.

In this section, we determine a condition on $\mathbf{s}$ and $\mathbf{\Lambda}$ that guarantees that the maximum energy subset optimizes $\mathrm{SNR}(\mathbf{G})$ even though $\mathbf{\Lambda}$ is not a diagonal matrix. Essentially, the result shows that whenever the covariance matrix is similar enough to $\mathbf{I}$, the maximum energy subset optimizes the RMF. In the cases identified in this section, use of the detailed greedy selection heuristic is unnecessary.

There are a number of ways to quantify the similarity between two matrices. We use the condition number as a way to measure similarity between $\mathbf{\Lambda}$ and the identity matrix [17]. The condition, for positive definite matrices, is defined as the ratio

$$\kappa(\mathbf{\Lambda}) = \frac{d_{\max}(\mathbf{\Lambda})}{d_{\min}(\mathbf{\Lambda})}, \tag{2.30}$$

where $d_{\max}(\mathbf{\Lambda})$ and $d_{\min}(\mathbf{\Lambda})$ are the maximum and minimum eigenvalues of the covariance matrix. The condition number obeys the inequality

$$\kappa(\mathbf{\Lambda}) \geq \kappa(\mathbf{I}) = 1. \tag{2.31}$$

In a global sense, if $\kappa(\mathbf{\Lambda}) \approx 1$, the covariance matrix behaves similarly to a scaled version of the identity matrix. In such situations, it is reasonable that the maximum signal energy subset, the optimal choice for the RMF for white noise, remains optimal. A criteria indicating when this subset is the optimal choice for low-condition covariances can be developed.

As a first step, we consider the relationship between the condition number of the overall covariance matrix $\mathbf{\Lambda}$, and the covariance matrix for any set of selected data, $\mathbf{\Lambda}_{\tilde{\mathbf{x}}_g}$. A variational view of eigenvalues is useful for bounding the condition of the selected covariance matrix. One way to describe the maximal and minimal

eigenvalues of a matrix, $d_{\max}$ and $d_{\min}$, is by reference to the gain R, defined by

$$R = \frac{\mathbf{v}^T \mathbf{\Lambda} \mathbf{v}}{\mathbf{v}^T \mathbf{v}}. \tag{2.32}$$

Sometimes this quantity is called the Rayleigh-Ritz ratio [17]. Since $\mathbf{\Lambda}$ is a positive definite matrix, the gain satisfies the inequalities

$$d_{\min}(\mathbf{\Lambda}) \leq R \leq d_{\max}(\mathbf{\Lambda}). \tag{2.33}$$

Thus, for covariance matrices, the maximum and minimum eigenvalues give the largest and smallest values for the gain.

The bounds on gain in equation (2.33) also apply to the eigenvalues of $\mathbf{\Lambda}_{\tilde{\mathbf{x}}_g}$. For a selected covariance matrix, we can generate all possible gains by matrix multiplication between $\mathbf{\Lambda}$, and a vector $\mathbf{v}$ with entries where $\mathbf{G}_{ii} = 0$. This vector is unrestricted in the dimensions where $\mathbf{G}_{ii} = 1$. For example, if $N = 3$, and the RMF is formed by selecting the first two pieces of data, $\mathbf{\Lambda}_{\tilde{\mathbf{x}}_g}$ is the upper left 2×2 block of $\mathbf{\Lambda}$. The restriction $v_3 = 0$ produces a situation where

$$\frac{\tilde{\mathbf{v}}_g^T \mathbf{\Lambda}_{\tilde{\mathbf{x}}_g} \tilde{\mathbf{v}}_g}{\tilde{\mathbf{v}}_g^T \tilde{\mathbf{v}}_g} = \frac{\mathbf{v}^T \mathbf{\Lambda} \mathbf{v}}{\mathbf{v}^T \mathbf{v}}. \tag{2.34}$$

Thus, the gain for $\mathbf{\Lambda}_{\tilde{\mathbf{x}}_g}$ obeys the bounds

$$d_{\min}(\mathbf{\Lambda}) \leq d_{\min}(\mathbf{\Lambda}_{\tilde{\mathbf{x}}_g}) \leq R \leq d_{\max}(\mathbf{\Lambda}_{\tilde{\mathbf{x}}_g}) \leq d_{\max}(\mathbf{\Lambda}). \tag{2.35}$$

An additional set of useful inequalities relates the value of a quadratic form to global properties of the signal vector and matrix involved. Since the quadratic form can be represented as

$$\mathbf{s}^T \mathbf{\Lambda}^{-1} \mathbf{s} = \sum_{i=1}^{N} \frac{(\mathbf{u}_i^T \mathbf{s})^2}{d_i(\mathbf{\Lambda})}, \tag{2.36}$$

the maximum and minimum possible values for the quadratic form for fixed $d_{\max}$,

$d_{\min}$, and $\|\mathbf{s}\|$ are

$$\begin{aligned} \mathbf{s}^T\mathbf{\Lambda}^{-1}\mathbf{s} &\geq \tfrac{\|\mathbf{s}\|^2}{d_{\max}(\mathbf{\Lambda})} \\ \mathbf{s}^T\mathbf{\Lambda}^{-1}\mathbf{s} &\leq \tfrac{\|\mathbf{s}\|^2}{d_{\min}(\mathbf{\Lambda})}. \end{aligned} \tag{2.37}$$

Similar bounds apply to the selected signal and covariance matrix $\tilde{\mathbf{s}}_g$ and $\mathbf{\Lambda}_{\tilde{\mathbf{x}}_g}$. From the inequalities in equation (2.33), we can further see that the selected quadratic form obeys the inequalities

$$\begin{aligned} \tilde{\mathbf{s}}_g^T\mathbf{\Lambda}_{\tilde{\mathbf{x}}_g}^{-1}\tilde{\mathbf{s}}_g &\geq \tfrac{\|\tilde{\mathbf{s}}_g\|^2}{d_{\max}(\mathbf{\Lambda}_{\tilde{\mathbf{x}}_g})} \geq \tfrac{\|\tilde{\mathbf{s}}_g\|^2}{d_{\max}(\mathbf{\Lambda})} \\ \tilde{\mathbf{s}}_g^T\mathbf{\Lambda}_{\tilde{\mathbf{x}}_g}^{-1}\tilde{\mathbf{s}}_g &\leq \tfrac{\|\tilde{\mathbf{s}}_g\|^2}{d_{\min}(\mathbf{\Lambda}_{\tilde{\mathbf{x}}_g})} \leq \tfrac{\|\tilde{\mathbf{s}}_g\|^2}{d_{\min}(\mathbf{\Lambda})}. \end{aligned} \tag{2.38}$$

The maximum energy subset, denoted by $\mathbf{G}^*$, consists of the data with the K largest magnitude signal measurements. The corresponding selected signal vector is denoted $\tilde{\mathbf{s}}_{g^*}$ and the corresponding selected covariance matrix is $\mathbf{\Lambda}_{\tilde{\mathbf{x}}_{g^*}}$. When $\mathbf{\Lambda} = \mathbf{I}$, the optimal RMF always corresponds to the maximum energy subset. This follows because the SNR reduces to $\|\tilde{\mathbf{s}}_g\|^2$ and the maximum energy subset by definition satisfies the inequality $\|\tilde{\mathbf{s}}_{g^*}\|^2 \geq \|\tilde{\mathbf{s}}_g\|^2$ for $\tilde{\mathbf{s}}_g$ formed by selecting any K signal samples.

The maximum energy subset remains the source for the optimal RMF if the inequality

$$\tilde{\mathbf{s}}_{g^*}^T\mathbf{\Lambda}_{\tilde{\mathbf{x}}_{g^*}}^{-1}\tilde{\mathbf{s}}_{g^*} \geq \tilde{\mathbf{s}}_g^T\mathbf{\Lambda}_{\tilde{\mathbf{x}}_g}^{-1}\tilde{\mathbf{s}}_g \tag{2.39}$$

holds for all K sample subsets. The inequalities in equation (2.38) provide a useful bounds on the quadratic forms. Since $\tilde{\mathbf{s}}_{g^*}^T\mathbf{\Lambda}_{\tilde{\mathbf{x}}_{g^*}}^{-1}\tilde{\mathbf{s}}_{g^*} \geq \frac{\|\tilde{\mathbf{s}}_{g^*}\|^2}{d_{\max}(\mathbf{\Lambda})}$ and $\frac{\|\tilde{\mathbf{s}}_g\|^2}{d_{\min}(\mathbf{\Lambda})} \geq \tilde{\mathbf{s}}_g^T\mathbf{\Lambda}_{\tilde{\mathbf{x}}_g}^{-1}\tilde{\mathbf{s}}_g$, the inequality

$$\frac{\|\tilde{\mathbf{s}}_{g^*}\|^2}{d_{\max}(\mathbf{\Lambda})} \geq \frac{\|\tilde{\mathbf{s}}_g\|^2}{d_{\min}(\mathbf{\Lambda})} \tag{2.40}$$

provides a condition for the SNR of the maximum energy subset to exceed the SNR of another subset. If the inequality holds for every arrangement of $\mathbf{G}$ with K elements, the inequality gives a sufficient condition for the maximum energy subset to yield the best RMF. These arguments prove the following theorem.

Theorem 2 *Let* $\mathbf{s}$ *be an* N*-dimensional vector and* $\mathbf{\Lambda}$ *be a positive definite matrix with condition number* $\kappa(\mathbf{\Lambda})$*. The maximum energy subset given by* $\mathbf{G}^*$ *is the* K

measurement subset such that $\|\tilde{\mathbf{s}}_{g^*}\|^2 \geq \|\tilde{\mathbf{s}}_g\|^2$ *for any other* K *measurement set. If the condition*

$$\kappa(\mathbf{\Lambda})\frac{\|\tilde{\mathbf{s}}_g\|^2}{\|\tilde{\mathbf{s}}_{g^*}\|^2} \leq 1 \tag{2.41}$$

holds for all selection matrices $\mathbf{G} \neq \mathbf{G}^*$*, then* $\mathbf{G}^*$ *maximizes* $\mathrm{SNR}(\mathbf{G})$ *over all* K *element subsets.*

This theorem relates global properties of the signal and covariance matrix to the RMF solution. Specifically, if the signal energy in the maximum energy subset exceeds the signal energy in any other K-sample subset, the RMF remains the maximum energy subset, as in the case of white noise. The exact solution to the selection problem does not change significantly if $\mathbf{\Lambda}$ resembles the identity matrix sufficiently closely.

Interestingly, the bound does not make any assumptions on the structure of the covariance matrix. A drawback, however, occurs when the signal has energy spread approximately evenly in its samples. In this case, the gap in signal energy between $\tilde{\mathbf{s}}_{g^*}$ and other subsets can be very small for large N. In the pathological case where the signal is constant, the bound in Theorem 2 is no guide at all.

Finally, note that the theorem analyzes only the case of a maximum energy subset selection rule. It does not account for the exact selection rules in the heuristic greedy algorithms. An analysis of their performance, however, reveals a similar form for their worst-case approximation ratios.

2.4.3 Worst-Case Performance of the Greedy Algorithm

The greedy algorithm relies upon a heuristic rule for choosing a subset of data. As such, it is not guaranteed to produce the optimal answer. This section characterizes the worst-case performance of the greedy algorithm, relating it to the condition number of the covariance matrix $\mathbf{\Lambda}$ and the properties of $\mathbf{s}$.

The bounds on the Rayleigh-Ritz ratio given in equations (2.33) and (2.38) can be used to determine lower bounds on the SNR achieved by the greedy selection rule. Let the SNR values for the subset selected by the greedy algorithm and the

optimal subset be denoted by $\text{SNR}_{\text{greedy}}$ and SNR_{opt}. The signal vectors for these subsets are $\mathbf{s}_{\text{greedy}}$ and $\mathbf{s}_{\text{opt}}$. The performance bound can be expressed in terms of an approximation ratio α. The relation

$$\text{SNR}_{\text{opt}} = \alpha \text{SNR}_{\text{greedy}} \tag{2.42}$$

defines α. If the greedy algorithm finds the optimal solution the ratio is $\alpha = 1$.

The approximation ratio can be bounded using the linear algebra techniques from the previous section. A large value of α, however, does not always lead to a significant loss in detection probability. The relationship between $\text{SNR}(\mathbf{G})$ and the RMF detector performance depends upon the actual value of SNR_{opt}. As indicated in equation (2.9), the detection probability is $P_D = Q(t - \text{SNR}(\mathbf{G}))$ for a detector with decision threshold t. This function is non-linear; it changes rapidly from 0.97 to 0.02 over the interval $[-2, 2]$. If SNR_{opt} is large, large values of α may not seriously impact detector performance. However, if α is such that $\text{SNR}_{\text{greedy}}$ falls in the interval $[-2, 2]$, the loss in performance can be significant.

The bounds on α are derived with the SNR bounds in equation (2.38). From these inequalities, we can conclude that SNR_{opt} satisfies

$$\text{SNR}_{\text{opt}} \leq \frac{\|\mathbf{s}_{\text{opt}}\|^2}{d_{\min}(\mathbf{\Lambda})}. \tag{2.43}$$

In this expression, $\mathbf{s}_{\text{opt}}$ serves as a placeholder for the signal vector corresponding to the optimal RMF subset. Additionally, the performance of the greedy algorithm can be bounded from below by

$$\text{SNR}_{\text{greedy}} \geq \frac{\|\mathbf{s}_{\text{greedy}}\|^2}{d_{\max}(\mathbf{\Lambda})}. \tag{2.44}$$

The notation $\mathbf{s}_{\text{greedy}}$ is a placeholder for the signal vector entries selected by the greedy algorithm. Combining these two inequalities, we find the following bound on

the approximation ratio

$$\alpha \leq \frac{\|\mathbf{s}_{\text{opt}}\|^2}{d_{\min}(\mathbf{\Lambda})\text{SNR}_{\text{greedy}}} \leq \frac{\|\mathbf{s}_{\text{opt}}\|^2}{\|\mathbf{s}_{\text{greedy}}\|^2}\frac{d_{\max}(\mathbf{\Lambda})}{d_{\min}(\mathbf{\Lambda})}. \tag{2.45}$$

In terms of the condition number, the approximation ratio is bounded by

$$\alpha \leq \frac{\|\mathbf{s}_{\text{opt}}\|^2}{\|\mathbf{s}_{\text{greedy}}\|^2}\kappa(\mathbf{\Lambda}). \tag{2.46}$$

The overall bounds on the approximation performance of the greedy algorithm are

$$1 \leq \alpha \leq \kappa(\mathbf{\Lambda})\frac{\|\mathbf{s}_{\text{opt}}\|^2}{\|\mathbf{s}_{\text{greedy}}\|^2}. \tag{2.47}$$

In this bound, the quantity $\frac{\|\mathbf{s}_{\text{opt}}\|^2}{\|\mathbf{s}_{\text{greedy}}\|^2}$ cannot be determined exactly without knowing the optimal subset selection. The ratio between maximum energy and minimum energy for K element subsets of $\mathbf{s}$ can be used as a bound on this term. In some cases, this term can be large. In other cases, the condition for $\mathbf{\Lambda}$ dominates the bound. Unfortunately, there are instances where the greedy algorithm selects a subset where α is close to the upper bound in equation.

2.4.4 Nearly Tight Example of the Forward Greedy Algorithm

In this subsection, we construct an example where the forward greedy selection algorithm nearly meets the worst-case performance bound in equation (2.47). The strategy in this section is to fix $\mathbf{\Lambda}_{\text{greedy}}$ and $\mathbf{s}_{\text{greedy}}$, and construct an overall signal $\mathbf{s}$ and covariance $\mathbf{\Lambda}$ such that forward greedy algorithm selects the subset indentified by $\mathbf{s}_{\text{g}}$, and the resulting α is proportional to $\kappa(\mathbf{\Lambda})$, which can be chosen arbitrarily large.

Consider a covariance matrix that is block diagonal. It has the form

$$\mathbf{\Lambda} = \begin{bmatrix} \mathbf{\Lambda}_{\text{greedy}} & \mathbf{0} \\ \mathbf{0} & \mathbf{\Lambda}_{\text{o}} \end{bmatrix}. \tag{2.48}$$

This example is constructed specifically so that the forward greedy algorithm to choose the $K = N/2$ measurements corresponding to $\mathbf{\Lambda}_{\text{greedy}}$. The covariance matrix $\mathbf{\Lambda}_{\text{o}}$ will be constructed so that the approximation ratio α is proportional to $\kappa(\mathbf{\Lambda})$.

In order to construct this example, we require conditions that guarantee that the greedy algorithm will select measurements associated with $\mathbf{\Lambda}_{\text{greedy}}$ rather than $\mathbf{\Lambda}_{\text{o}}$. Recall that the greedy algorithm selects a measurement that maximizes the quantity given in equation (2.28) using the covariance matrix for the measurements it selected previously. Consider applying the greedy algorithm to the K measurements that we intend for it to select. The minimum value of Δ_i^+ when the greedy algorithm acts on the restricted set is denoted by $\Delta_{\min}$. The condition

$$\frac{s_j^2}{[\mathbf{\Lambda}_{\text{o}}]_{jj}} < \Delta_{\min} \text{ for all } j \in \mathbf{G}_{\text{o}}. \tag{2.49}$$

guarantees that the forward greedy algorithm selects no measurements in $\mathbf{G}_{\text{o}}$. At the first stage of the algorithm the increment for any measurement in $\mathbf{G}_{\text{o}}$ is $\Delta = \frac{s_j^2}{[\mathbf{\Lambda}_{\text{o}}]_{jj}}$. The condition in equation (2.49) ensures that the first measurement selected is the first selected when the greedy algorithm runs only on the measurements associated in $\mathbf{G}_{\text{greedy}}$. Since the measurements in the two subsets are uncorrelated, the increments for any of the measurements in $\mathbf{G}_{\text{o}}$ do not change as measurements from $\mathbf{G}_{\text{greedy}}$ are selected. None of the measurements in the optimal subset are selected at any stage of the forward greedy algorithm because they are less than $\Delta_{\min}$.

Thus, to give worst-case performance, our construction requires that the condition (2.49) be satisfied, and that the resulting values of $\text{SNR}(\mathbf{G}_{\text{greedy}})$ and $\text{SNR}(\mathbf{G}_{\text{o}})$ yield the approximation ratio $\alpha = \kappa(\mathbf{\Lambda}) \frac{\|\mathbf{s}_{\text{o}}\|^2}{\|\mathbf{s}_{\text{greedy}}\|^2}$.

Based upon intuition from the bounds, we require that the selected signal vector $\mathbf{s}_{\text{greedy}}$ falls in the subspace spanned by eigenvectors associated with the maximum

eigenvalue $d_{\max}(\mathbf{\Lambda})$, and that $\mathbf{s}_{\mathrm{o}}$ falls in the subspace spanned by eigenvectors with eigenvalue $d_{\min}(\mathbf{\Lambda})$. It is possible to construct such an instance where these conditions are established. Consider the covariance matrix

$$\mathbf{\Lambda}_{wc} = \begin{bmatrix} \mathbf{I}_K & \mathbf{0} \\ \mathbf{0} & \mathbf{\Lambda}_{\mathrm{o}} \end{bmatrix}. \tag{2.50}$$

With properly chosen signal $\mathbf{s}^T = [\mathbf{s}_{\mathrm{greedy}}^T \; \mathbf{s}_o^T]$, the greedy algorithm will select a subset that will meet the worst-case bound on the approximation ratio with equality. For this example, let $\mathbf{s}_{\mathrm{o}} = \frac{1}{\sqrt{N/2}}[1 \; 1 \; \ldots \; 1]^T$ and $\mathbf{s}_{\mathrm{greedy}} = \mu \mathbf{s}_{\mathrm{o}}$, with $\mu > 1$. With the proper choice of $\mathbf{\Lambda}_{\mathrm{o}}$, this example meets the worst case approximation ratio for the greedy algorithm.

The condition, expressed in equation (2.49), for the greedy algorithm to select measurements in $\mathbf{s}_{\mathrm{greedy}}$ rather than in $\mathbf{s}_o$ reduces to

$$\frac{\mathbf{s}_{\mathrm{o},i}^2}{\mathbf{\Lambda}_{\mathrm{o},ii}} < \min_j \mathbf{s}_{\mathrm{greedy},j}^2. \tag{2.51}$$

Since $\mu^2 \mathbf{s}_{\mathrm{o},i}^2 = \mathbf{s}_{\mathrm{greedy}_i}^2$, the condition is satisfied if the diagonal entries of $\mathbf{\Lambda}_{\mathrm{o}}$ satisfy $\mathbf{\Lambda}_{\mathrm{o,ii}} > \frac{1}{\mu^2}$. The eigen-decompostion of this matrix is $\mathbf{\Lambda}_{\mathrm{o}} = \mathbf{Q}\mathbf{D}\mathbf{Q}^T$, so the diagonal entries have the form $\mathbf{\Lambda}_{\mathrm{o},ii} = \sum_{j=1}^{K} (\mathbf{Q}_{ij})^2 \mathbf{D}_{jj}$. If we choose $\mathbf{Q}_{i1} = \mathbf{s}_{\mathrm{o},i} = \frac{1}{\sqrt{K}}$, a lower bound on the diagonal entry is

$$\mathbf{\Lambda}_{\mathrm{o},ii} \geq \min_{j \neq 1} \mathbf{D}_{jj} \left(\frac{K-1}{K} \right) \tag{2.52}$$

because $\sum_j \mathbf{Q}_{ij}^2 = 1$, and $\mathbf{Q}_{i1}$ is fixed. Thus, if the lower bound on the diagonal entries exceeds $\frac{1}{\mu^2}$, then so does each entry. The condition on the entries of $\mathbf{D}$ becomes $\min_{j \neq 1} \mathbf{D}_{jj} > \frac{K}{(K-1)\mu^2}$. Thus, we can pick μ such that this condition is satisfied simultaneously with the upper bound on the eigenvalues.

Additionally, since (2.52) does not restrict $\mathbf{D}_{11}$ for $\mathbf{\Lambda}_{\mathrm{o}}$, we can choose this eigenvalue as small as we wish. Since $\mathbf{s}_{\mathrm{o}}$ is the eigenvector for $\mathbf{D}_{11}$, the gap between the performance of the greedy algorithm and the optimum subset can be arbitrarily large.

Thus, for a choice of $\mathbf{D}_{11}$ small enough, $\alpha = \frac{\|\mathbf{s}_o\|^2}{\|\mathbf{s}_{\text{greedy}}\|^2}\kappa(\mathbf{\Lambda}_{wc})$. The ratio of the two signal vector magnitudes is $\frac{1}{\mu^2}$, so the approximation ratio is directly proportional to the condition of the covariance matrix, which can be made arbitrarily large by choosing small values for $\mathbf{D}_{11}$.

2.4.5 Nearly Tight Example of the Backward Greedy Algorithm

An instance of the RMF problem where the backward greedy algorithm achieves performance proportional to $\kappa(\mathbf{\Lambda})$ can be constructed. The strategy for constructing the $\mathbf{\Lambda}$ and $\mathbf{s}$ instances is similar to the strategy for the forward greedy worst case. A small submatrix that is ill conditioned is inserted into $\mathbf{\Lambda}$. The vector $\mathbf{s}$ and the remaining entries of $\mathbf{\Lambda}$ are chosen to insure that the backwards greedy algorithm will discard enough measurements so that the remaining covariance matrix $\mathbf{\Lambda}_K$ is no longer ill-conditioned.

Consider the following instance of the RMF problem. The covariance matrix is

$$\mathbf{\Lambda} = \begin{bmatrix} 1 & 0 & \frac{\rho+\tau}{2} & \frac{\rho-\tau}{2} \\ 0 & 1 & \frac{\rho-\tau}{2} & \frac{\rho+\tau}{2} \\ \frac{\rho+\tau}{2} & \frac{\rho-\tau}{2} & \frac{\xi+1}{2\xi} & \frac{\xi-1}{2\xi} \\ \frac{\rho+\tau}{2} & \frac{\rho-\tau}{2} & \frac{\xi-1}{2\xi} & \frac{\xi+1}{2\xi} \end{bmatrix}. \tag{2.53}$$

The variable ξ is a free parameter that will control $\kappa(\mathbf{\Lambda})$ and the optimal SNR. This variable can take on any value satisfying $\xi > 1$.The parameters ρ and τ are defined as

$$\rho = \sqrt{\frac{\xi-1}{\xi}} \quad \tau = \frac{1}{\sqrt{2\xi}}. \tag{2.54}$$

The condition of $\mathbf{\Lambda}$ is ξ, and the matrix is positive definite for any value of $\xi > 1$. The signal vector in this instance is

$$\mathbf{s}^T = \frac{1}{\sqrt{2}}\left[\frac{1}{\rho}\ \frac{1}{\rho}\ \frac{1}{2}\ -\frac{1}{2}\right]. \tag{2.55}$$

The covariance matrix and signal in equations (2.53-2.55) are chosen precisely so that the first stage of the greedy algorithm will discard the third measurement, and the remaining submatrix is well-conditioned regardless of ξ. Recall that the SNR increments for discarding a measurement can be written as $\Delta^-(\mathbf{G}, i) = \frac{(\mathbf{z}_i^T \mathbf{\Lambda}_{\tilde{\mathbf{x}}_g}^{-1} \tilde{\mathbf{s}}_g)^2}{\mathbf{z}_i^T \mathbf{\Lambda}_{\tilde{\mathbf{x}}_g}^{-1} \mathbf{z}_i}$. For the instance described, $\mathbf{G} = \mathbf{I}$, and the vector $\mathbf{\Lambda}^{-1}\mathbf{s}$ is

$$\mathbf{\Lambda}^{-1}\mathbf{s} = \begin{bmatrix} \frac{1}{\rho} + \frac{1}{1-\rho^2} + \frac{\tau_o\sqrt{\xi}}{2(1-\tau_o^2)} \\ \frac{1}{\rho} + \frac{1}{1-\rho^2} - \frac{\tau_o\sqrt{\xi}}{2(1-\tau_o^2)} \\ -\frac{1}{(1-\rho^2)} + \frac{\xi}{2(1-\tau_o^2)} \\ -\frac{1}{(1-\rho^2)} - \frac{\xi}{2(1-\tau_o^2)} \end{bmatrix}, \tag{2.56}$$

where $\tau_o = \frac{\tau}{\sqrt{\xi}}$. Using the values for ρ and τ from equation (2.54), the third entry of $\mathbf{\Lambda}^{-1}\mathbf{s} = 0$ for any value of $\xi > 1$. The backward greedy algorithm always drops this measurement first, independent of ξ. The remaining entries of the covariance matrix are well-conditioned for large values of ξ. This example produces an approximation ratio for the backwards greedy selection algorithm that is proportional to the condition number of $\mathbf{\Lambda}$.

2.4.6 Performance Comparison for the Greedy Algorithm

The performance bounds on the greedy algorithm show that it can perform poorly when the condition of $\mathbf{\Lambda}$ is relatively large. In practice, the algorithms often perform well even when $\kappa(\mathbf{\Lambda})$ is large. As an example, consider the signal shown in Figure 2-4. The signal is a damped sinusoid, and the covariance is a symmetric Toeplitz matrix generated by the sequence $12 * (.995)^i$, for $i = 0, \ldots, 14$. The condition for this covariance matrix is $\kappa(\mathbf{\Lambda}) = 5774.8$.

Despite the large condition number in this instance, the forward greedy algorithm performs well. For K ranging over the entire range of subset sizes, the forward greedy algorithm selects a subset with SNR close to the optimal selected by exhaustive search. Figure 2-5 shows the performance of the forward greedy algorithm and the maximum signal energy selection algorithm. The forward greedy algorithm does almost as well

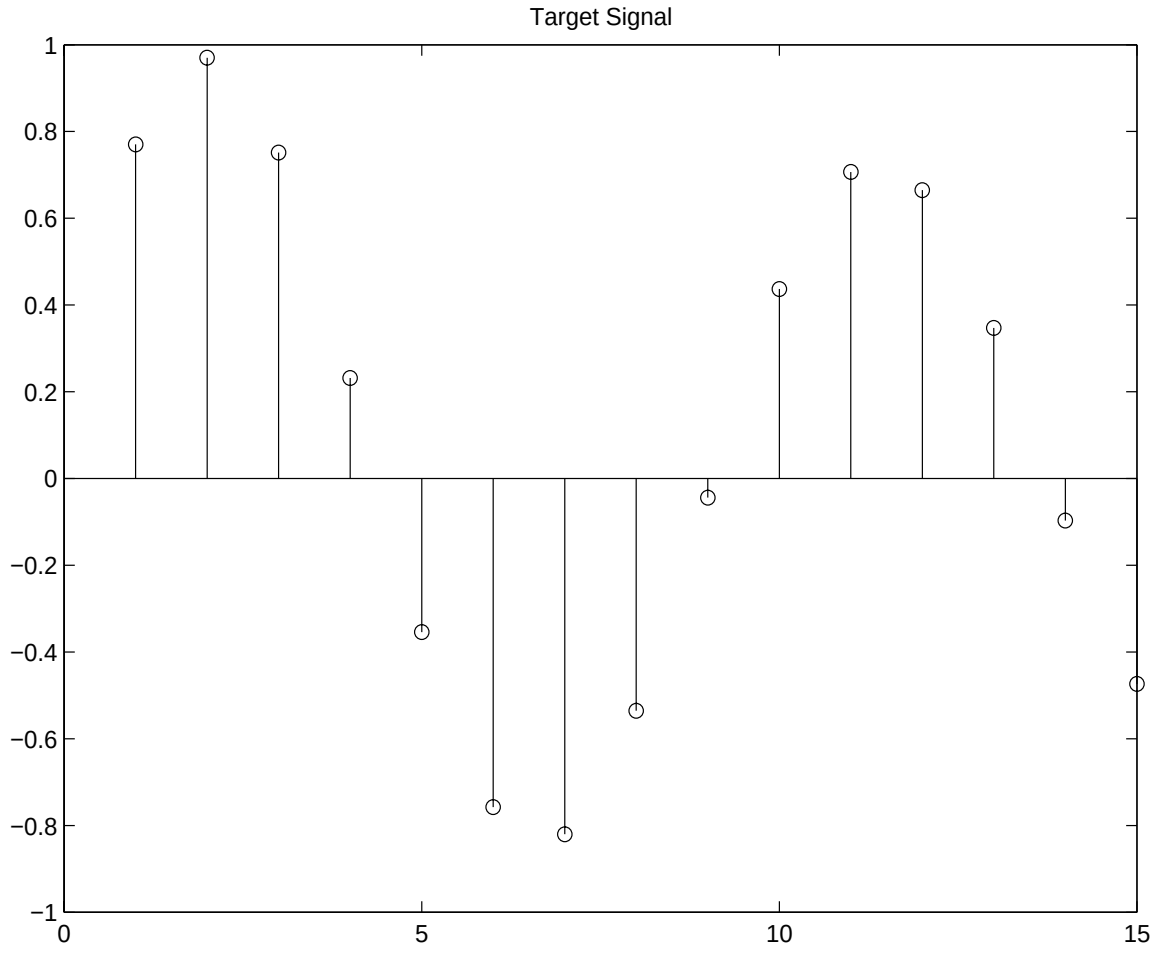

Figure 2-4: Target signal for greedy search examples.

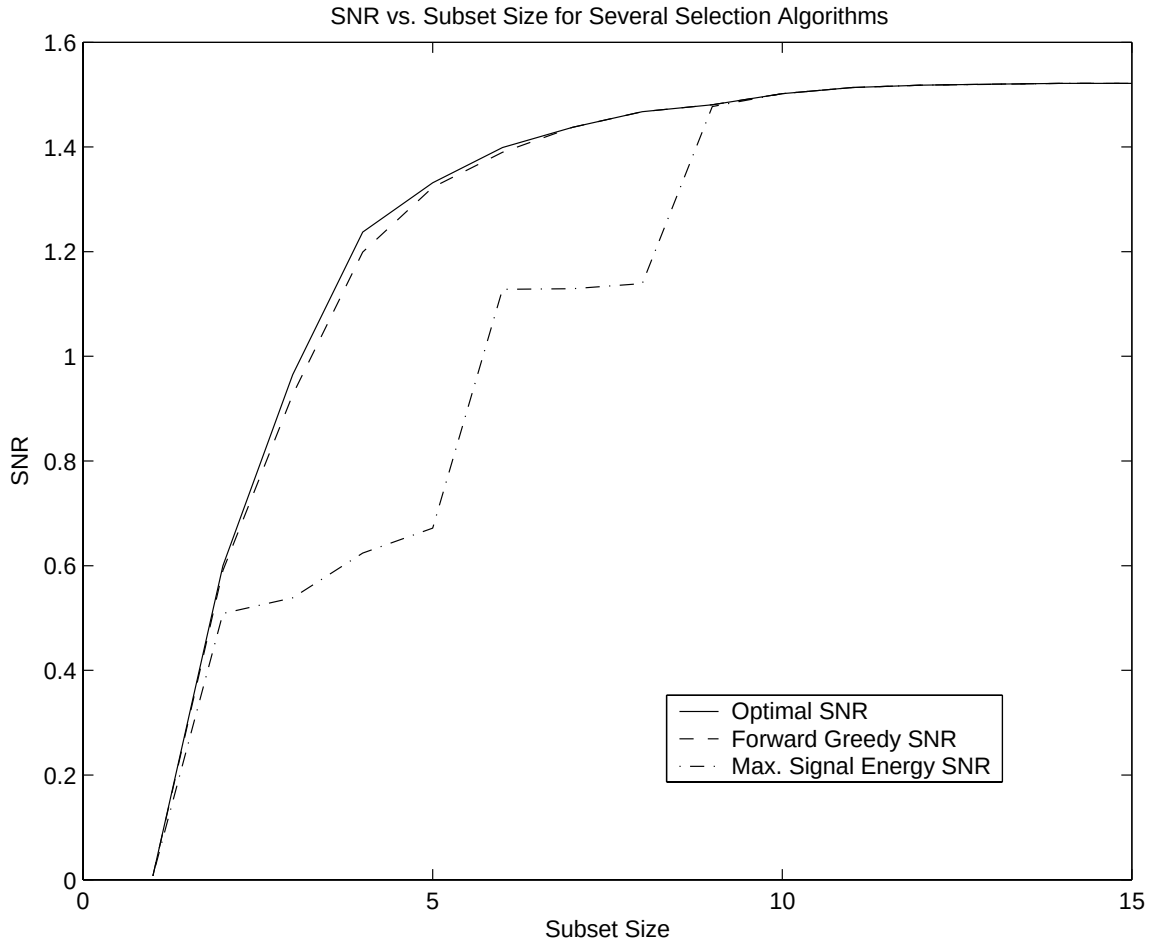

Figure 2-5: SNR versus K for Forward Greedy selection algorithm, maximum signal energy selection, and exhaustive search selection.

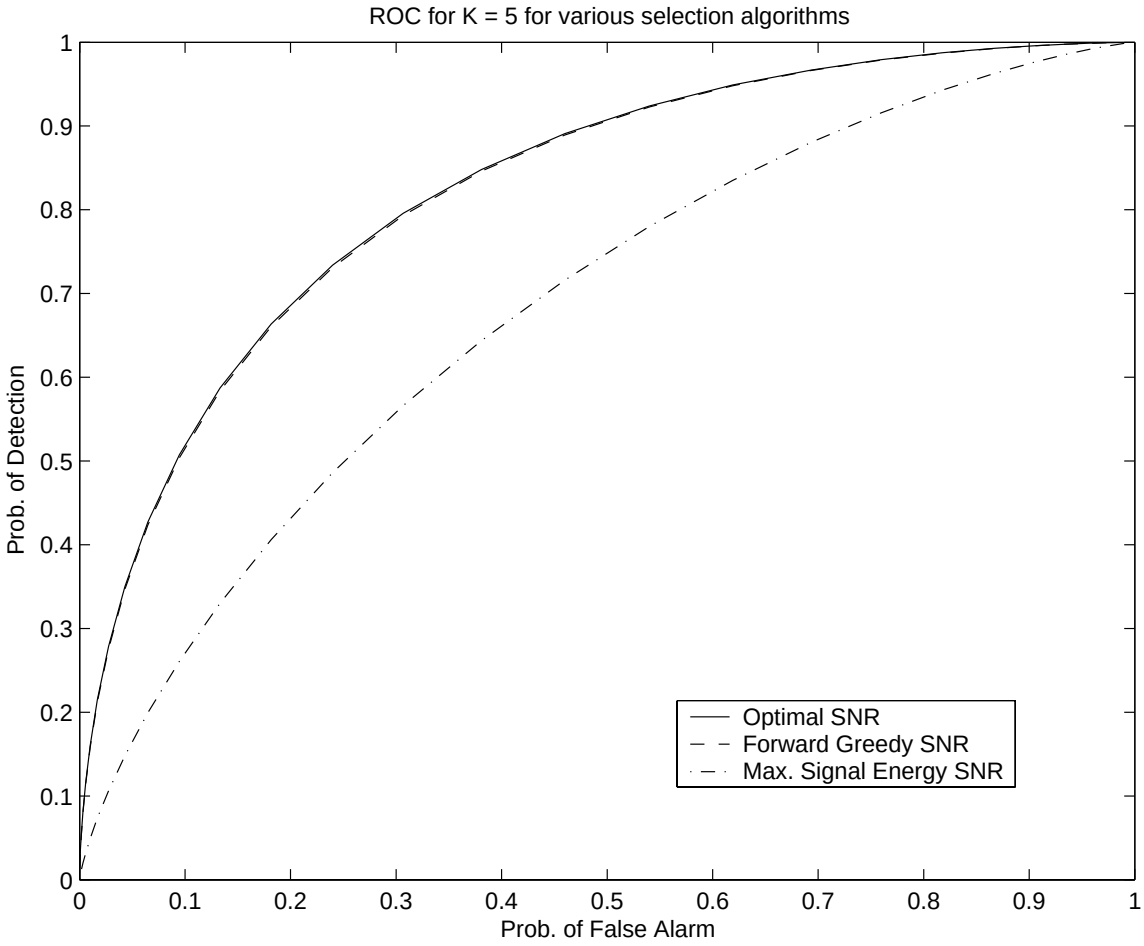

Figure 2-6: $K = 5$ ROC for Forward Greedy selection algorithm, maximum signal energy selection, and exhaustive search selection.

as the exhaustive search, and significantly better than the maximum signal energy selection rule. The resulting ROCs for the three selection algorithm for subsets of $K = 5$ are shown in Figure 2-6.

2.5 Dynamic Programming Solutions for Banded Matrices

As indicated in the previous section, selection problems involving noise covariance matrices possessing condition numbers similar to $\kappa(\mathbf{I})$ yield simplified solutions to the search for an optimal RMF data set. This principle applies to properties of a matrix beyond the condition number. Another way in which the matrix can resemble an identity matrix arises from some structural properties of the matrix. In this subsection, we focus on the class of banded matrices. The non-zero entries of these matrices cluster near the main diagonal. Since most entries far from the main diagonal are zero, the structural constraints give the matrix a resemblance to the identity matrix. Although the locations of non-zero entries resemble $\mathbf{I}$, the structural constraint does not restrict $\kappa(\mathbf{\Lambda})$.

Banded covariance matrices arise when measurements are correlated with a number of neighbors. These matrices can model a significant set of noise processes. For example, the background interference experienced by a set of sensors arranged in a uniform linear array may be plausibly described by a spatial random process with a banded covariance matrix. It is possible to assume that the background noise measurements at neighboring sensors are correlated, while the noise measurements at two distantly separated sensors are independent.

The optimization algorithms developed in this section depend upon the size of the neighborhood of correlated measurements. Assume that a particular indexing of the data vector $\mathbf{x}$ confines all non-zero entries of the correlation matrix to a band $2b + 1$ entries wide, centered on the main diagonal. The correlation distance of the matrix bounds the distance between two correlated measurements. Its precise defintion is

Definition 1 *For a correlation matrix* $\mathbf{\Lambda}$*, the correlation distance is*

$$b = \max |i - j| \text{ where } [\mathbf{\Lambda}]_{ij} \neq 0. \tag{2.57}$$

Note that this condition does not require that all measurements with indices separated by less than b be correlated.

The banded structure of $\mathbf{\Lambda}$ enables us to decompose SNR($\mathbf{G}$) into the sum of a contributions from measurement subset fragments separated by more than b inidices. The best subset can be determined by a dynamic programming algorithm. The following sections develop this algorithm for the case $b = 1$, and then extend it to larger correlation distances.

2.5.1 Fragment Decomposition of Selected Subsets

When $b = 1$, the measurement subset in $\mathbf{x}_g$ can be built from several groups of consecutive measurements, which we shall refer to as fragments. For example, consider the selected subset

$$\mathbf{x}_g^T = \begin{bmatrix} x_1 & 0 & x_3 & x_4 \end{bmatrix}. \tag{2.58}$$

In this example, the selected subset consists of $\{x_1, x_3, x_4\}$. For this subset, the two fragments are $\{x_1\}$ and $\{x_3, x_4\}$. The selection matrix $\mathbf{G}$ can be written as the sum of the two selection matrices for the fragments, i.e.

$$\mathbf{G} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} + \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}. \tag{2.59}$$

Finally, the covariance matrix for the selected data is

$$\mathbf{\Lambda}_{\tilde{\mathbf{x}}_g} = \begin{bmatrix} [\mathbf{\Lambda}]_{11} & 0 & 0 \\ 0 & [\mathbf{\Lambda}]_{33} & [\mathbf{\Lambda}]_{34} \\ 0 & [\mathbf{\Lambda}]_{43} & [\mathbf{\Lambda}]_{44} \end{bmatrix}. \tag{2.60}$$

Note that measurements from separate fragments of consecutive measurements are uncorrelated.

The decomposition can be extended to an arbitrary number of fragments. For an arbitrary selection matrix $\mathbf{G}$, there will be f fragments. The fragments, expressed as N-vectors are denoted by $\mathbf{x}_{\{g,i\}}$ for $i = 1, \ldots, f$, and selection matrices identifying the fragments are expressed as $\mathbf{G}_1, \mathbf{G}_2, \ldots, \mathbf{G}_f$. For example, in equation (2.58), the first fragment is

$$\mathbf{x}_{\{g,1\}}^T = \begin{bmatrix} x_1 & 0 & 0 & 0 \end{bmatrix}. \tag{2.61}$$

The corresponding selection matrix is

$$\mathbf{G}_1 = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}. \tag{2.62}$$

In this notation, any selected vector can be expressed via

$$\mathbf{x}_g = \sum_{i=1}^{f} \mathbf{x}_{\{g,i\}}, \tag{2.63}$$

as the sum of all of its fragment vectors. Likewise, the selection matrix for the entire subset is

$$\mathbf{G} = \sum_{i=1}^{f} \mathbf{G}_i. \tag{2.64}$$

In these examples, the fragment order is assigned by sorting the fragments according to their first measurement index. Thus, the fragment containing the element assigned

the smallest index is $\mathbf{G}_1$ and so forth.

In compact form, each fragment can be expressed as a w-dimensional vector, where w is the number of consecutive measurements in the fragment. These fragments are denoted by $\tilde{\mathbf{x}}_{\{g,i\}}$. In this notation, the first fragment in equation (2.58) is

$$\tilde{\mathbf{x}}_{\{g,1\}} = \left[\begin{array}{c} x_1 \end{array} \right]. \tag{2.65}$$

The corresponding selection matrices are denoted by $\tilde{\mathbf{G}}_i$ for $i = 1, \ldots, f$. A formal definition of subset fragments is

Definition 2 *Consider an instance of the RMF where $\mathbf{\Lambda}$ has correlation distance b. A correlated subset fragment, denoted by $\mathbf{x}_{\{g,i\}}$, is a set of measurements $\{x_{j_1}, x_{j_2}, \ldots, x_{j_a}\}$ such that the indices satisfy $j_1 < j_2 < \ldots < j_a$ and $|j_{i+1} - j_i| \leq b$ for every pair of consecutive indices.*

The fragment notation provides a useful tool for expressing the SNR of any subset $\mathbf{G}$ in a convenient form. Two lemmas justify this expression. First, any subset $\mathbf{G}$ can be written as a union of a number of fragments. Second, once the decomposition into fragments is determined, SNR($\mathbf{G}$) of the entire subset can be written as the sum of fragment SNRs. The first lemma is:

Lemma 1 *Consider any instance of the RMF where the covariance matrix $\mathbf{\Lambda}$ has correlation distance $b < N$. Any K element subset $\mathbf{G}$ possesses a unique decomposition $\mathbf{G} = \sum_{i=1}^{f} \mathbf{G}_i$ where each subset $\mathbf{G}_i$ represents a correlated fragment. The indices of $\mathbf{G}$ are ordered such that for any pair of measurements $x_l \in \mathbf{G}_i$, $x_m \in \mathbf{G}_{i+1}$, the measurement indices satisfy $l + b < m$.*

Proof. The subset decomposition can be constructed from a sorted list of the measurement indices denoted by $\{j_1, j_2, \ldots, j_K\}$. Expressed in ascending order for any subset $\mathbf{G}$, this list is unique. From the index list, the subset fragments can be constructed from the vector of index spacings $\mathbf{f} = [(j_2 - j_1), (j_3 - j_2), \ldots, (j_K - j_{K-1})]^T$. The elements of the vector are positive integers. Using $\mathbf{f}$, the desired subset decomposition can be determined. Initialize the subsets by assigning j_1 to the first subset

fragment $\mathbf{G}_1$. To determine the assignment of the selected measurement $(i+1)$, compare $\mathbf{f}_i$ to b. If $\mathbf{f}_i \leq b$, the measurement belongs to the same subset fragment as x_{j_i}. If $\mathbf{f}_i > b$, the measurement belongs to the next subset fragment.

The correctness of this algorithm can be proven by induction. It produces the correct decomposition for$K = 1$ or $K = 2$ elements. To carry out the induction proof, it is sufficient to consider adding the elements of the subset in ascending numerical order. Any such subset can be built in that order. If an element with index exceeding any in $\mathbf{G}$ is added, it is included in the most recently created fragment if it belongs there. Otherwise, it is included as a newly created fragment. In both cases, the new fragment decompostion is correct. □

A correlation distance $b = 1$ implies that the x_i measurements are uncorrelated unless they have consecutive indices. This property simplifies the calculation of the SNR for any selected subset of measurements.

Lemma 2 *For an RMF instance such that* $\mathbf{\Lambda}$ *has a correlation distance* $b < N$*, consider a subset* $\mathbf{G} = \sum_{i=1}^{f} \mathbf{G}_i$*, where* $\mathbf{G}_i$ *are correlated fragments ordered as in Lemma 1. The* SNR *for this subset is*

$$\mathrm{SNR}(\mathbf{G}) = \sum_{i=1}^{f} \mathrm{SNR}(\mathbf{G}_i) = \sum_{i=1}^{f} \tilde{\mathbf{s}}_{\{g,i\}}^{T} \mathbf{\Lambda}_{\tilde{\mathbf{x}}_{\{g,i\}}}^{-1} \tilde{\mathbf{s}}_{\{g,i\}}. \tag{2.66}$$

Proof. As a result of Lemma 1, measurements in distinct subset fragments are separated by more than b indices. Thus, the entries of $\mathbf{\Lambda}_{\tilde{x}_g}$ corresponding to measurement pairs in distinct fragments must be 0. If the entries of the selected covariance matrix correspond to $[\tilde{\mathbf{x}}_g]_i$ in sorted order, the resulting matrix is block-diagonal. The result in equation 2.66 follows from the block-diagonal structure of $\mathbf{\Lambda}_{\tilde{x}_g}$. □

Lemmas 1 and 2 provide a technique for expressing SNR in a simplified form. For an arbitrarily chosen covariance matrix, one that is not banded, the SNR contribution of an individual measurement depends on the entire subset. This can be seen by

examining by $\Delta^+(\mathbf{G}, i)$ or $\Delta^-(\mathbf{G}, i)$, defined in equations (2.28) and (2.29). The SNR cannot be expressed as a linear combination of a metric depending on a single entries of $\mathbf{G}$. In the case where the covariance matrix is banded, however, the fragment decomposition shows that the SNR expression can be simplified. The SNR becomes the sum of fragment SNRs, each one unaffected by the identity of the other fragments included in $\mathbf{G}$. The fragment decomposition thus retains some of the convenience of the SNR expression for diagonal matrices.

2.5.2 Example of Dynamic Programming for Two-Measurement Subsets

The fragment notation provides a tool to express the SNR for RMF instances with banded covariance matrices in a convenient form. A dynamic programming algorithm based upon this description can determine the optimal choice of measurements. This subsection gives an example for the case of $K = 2$ that illustrates the general ideas in the optimal algorithm developed in the following sections.

Consider choosing the optimal two-element subset for an RMF instance. In this case, a brute-force search for the optimal subset requires computation of the SNR for $O(N^2)$ pairs of measurements.

In situations where the correlation distance satisfies $b = 1$, the structure of the covariance matrix can be exploited to improve the complexity of finding the optimal RMF subset. In such situations, a set of measurements with consecutive indices forms a correlated fragment. Otherwise, the subset is composed of two distinct fragments. This observation, and the expression for SNR in equation 2.66 combine to simplify the search.

Consider a two-element subset with the first element fixed at a particular index i. If the subset elements are not consecutive, the subset SNR can be written

$$\text{SNR} = \text{SNR}(\{i\}) + \text{SNR}(\{i + j\}), \tag{2.67}$$

where $j > 1$. Since the first element of the subset is fixed at i, optimizing over all of

the subsets composed of two fragments is reduced to choosing the index of the second element according to

$$a = \arg \max_{j \in [2, N-i]} \text{SNR}(\{i + j\}). \tag{2.68}$$

For $K = 2$, the overall best subset constrained to have first index i can be determined by comparing $\text{SNR}(\{i, i+1\})$ and $\text{SNR}(\{i, a\})$. Finally, the best RMF subset of two measurements can be determined by calculating the best subsets with fixed first element for $i = 1$ to $i = N - 1$ and searching the list for the best subset.

Equation 2.68 indicates that the proposed optimization procedure for $K = 2$ utilizes results for single element subsets. Specifically, a list of the maximum SNR element with index greater than i, for $i = 1, \ldots, N-1$ simplifies the optimization in the equation from a search through a list to a look-up. The computation necessary to constructed this table is dominated by an $O(N \log N)$ sort of the single element SNRs. In fact, this sort dominates the entire computation of the best two element subset since computation of all the single element SNRs and optimization of fixed two element subsets outlined above are $O(N)$ computations.

The re-use of intermediate subsets in this search reduces the computation from $O(N^2)$ for a brute force search to $O(N \log N)$ once the structure of the covariance matrix is used. The bootstrap procedure in this example, computing intermediate solutions using prior results is characteristic of dynamic programming. It can be expanded to computing optimal subsets for any value of K. In order to determine the best RMF subset for $K = 3$, for example, intermediate results from $K = 1$ and $K = 2$ are necessary.

2.5.3 Fragment Notation for the Dynamic Programming Algorithm

The example for two element subsets can be generalized to subsets of any size. An alternative to the **G** notation for subsets that highlights the importance of correlated fragments aids the algorithm's description. In the example, the optimization algorithm identified the fragments of a subset by the index of the first element and the

size of the subset. For $b = 1$, these numbers completely identify a subset fragment. As an alternative to the **G** notation, a subset fragment can be identified as $v_{l,c}$. The fragment so denoted has l consecutive elements, beginning with element c. An example is $v_{3,5} = \{5, 6, 7\}$. This notation for a subset fragment can identify the initial subsets used in the first stage of the optimization algorithm. For an RMF instance of size N, values of l and c must satisfy $l + c - 1 \leq N$ to guarantee that all indices in $v_{l,c}$ correspond to real measurements.

The fragment notation can represent any measurement subset via the standard notation of set theory. Generally, a subset V can be written as $V = \cup_{i=1}^{f} v_{l_i,c_i}$. Note that some combinations of $\{v_{l_i,c_i}\}$ do not produce a valid subset. They may contain multiple fragments that contain adjacent or duplicate measurements. This defect is avoided if, for every pair of fragments v_{l_1,c_2} and v_{l_2,c_2} such that $c_1 < c_2$ the indices satisfy $l_1 + c_1 < c_2$. If this condition is satisfied, Lemma 2 may be restated:

Lemma 3 *Consider an instance of the RMF with correlation distance $b = 1$. For a set $V = \cup_{i=1}^{f} v_{l_i,c_i}$ such that for every pair of fragments v_{l_1,c_1}, v_{l_2,c_2} satisfies $l_1 + c_1 < l_2$,*

$$\mathrm{SNR}(V) = \sum_{i=1}^{f} \mathrm{SNR}(v_{l_i,c_i}). \tag{2.69}$$

Furthermore, any subset of measurements can be written in this form for an appropriate choice of l_i, c_i and f.

Proof. The condition $l_1 + c_1 < c_2$ guarantees that elements in distinct subsets are separated by at least one element that is not selected. Thus, the fragments v_{l_i,c_i} are all independent. The remainder of the proof is identical to the original formulation of this lemma. □

This notation establishes the first part of the optimization algorithm. A specific fragment is easily identified by its length and the index of its first measurement.

2.5.4 Optimization Procedure for the Dynamic Programming Algorithm

Using the new notation, the search for the optimal RMF subset can be reformulated in a convenient form. Given an instance of the RMF with correlation distance $b = 1$, and the associated measure $SNR(v_{l,c})$, for every correlated fragment, we seek a subset of elements V^* such that

$$\begin{aligned} &1.\ l_i + c_i < l_j\ \forall\ (v_{l_i,c_i}, v_{l_j,c_j})\ \in V^* \text{ where } l_i < l_j \\ &2.\ \textstyle\sum_{v_{l_i,c_i} \in V^*} l_i = K \end{aligned} \tag{2.70}$$

where V^* maximizes

$$\sum_{v_{l_i,c_i} \in V_o} \text{SNR}(v_{l_i,c_i}) \tag{2.71}$$

over all V_o satisfying the constraints.

The remaining stages of the optimization require two sets of state information. The first state consists of the best subset with the size and first index fixed. The ith entry of the second state table is the best subset with the first index no less than i. These definitions are sufficient for useful for defining the actual optimization procedure used.

Dynamic programming algorithms operate by computing optimal solutions to small subproblems and combining them. Subsets of measurements with a lower bound on their minimum element form a fruitful set of subproblems for the optimal RMF search. Notation for these subsets is established in the following defintion.

Definition 3 *Let $V^{K,c}$ be a subset of K measurements satisfying the following two constraints. First, each measurement $x_i \in V^{K,c}$ satisfies*

$$i \geq c. \tag{2.72}$$

Second, the metric $\text{SNR}(V^{K,c})$ *satisfies*

$$\text{SNR}(V^{K,c}) \geq \text{SNR}(V_o) \tag{2.73}$$

for every K element subset V_o that satisfies the constraint in equation (2.72).

Note that the optimal RMF V^* for a K element subset is $V^{K,1}$.

A second useful set of subproblems are measurement subsets with the first fragment fixed.

Definition 4 *Let $U^{K,l,c}$ be a subset of K measurements satisfying the following three constraints. First, the measurements represented by the fragment $v_{l,c}$ are included in $U^{K,l,c}$. Second, for any measurement $x_i \in U^{K,l,c}$ and $x_i \notin v_{l,c}$, the index satisfies $i \geq c + l + 1$. Third, the metric $\mathrm{SNR}(U^{K,l,c})$ satisfies*

$$\mathrm{SNR}(U^{K,l,c}) \geq \mathrm{SNR}(V_o) \tag{2.74}$$

for any K element subsets satisfying the first two constraints.

The subset $U^{K,l,c}$ is determined by optimization over all K measurement subsets with a common first fragment.

The subset $U^{K,l,c}$ satisfies more restrictive conditions than $V^{K,c}$. It is possible, however, to relate the two groups of subsets. The following Theorem relates $U^{K,l,c}$ and $V^{K,c}$.

Theorem 3 *For any instance of the RMF problem with $b = 1$, consider the subset $U^{K,l,c}$ where the indices satisfy $K \leq N$, $l < K$, and $c \leq N - K$. This subset satisfies*

$$U^{K,l,c} = v_{l,c} \cup V^{K-l,c+l+1}. \tag{2.75}$$

If the index $l = K$, then the subset is $U^{K,K,c} = v_{K,c}$ and

$$c \leq N - K + 1. \tag{2.76}$$

Proof. In the case $l = K$, the subset $U^{K,l,c}$ consists of a single fragment. The final index of the fragment is $c + K - 1$. Since it can be no larger than N, the first index is restricted to $c \leq N - K + 1$.

When $l < K$, the subset $U^{K,l,c}$ consists of at least two fragments separated by at least one unselected measurement. The maximum value of c occurs when $U^{K,l,c}$ is composed of two fragments separated by only one unselected measurement. The total length of this arrangement of measurements is $K+1$, so the first index is restricted to $c \leq N-K$.

The general form for any subset with the first fragment fixed is $Y = v_{l,c} \cup R$, where the remaining elements in the subset are represented by R. This is a $K-l$ measurement subset that with its first index satisfying $i \geq c+l+1$. The SNR metric for the subset is

$$\mathrm{SNR}(Y) = \mathrm{SNR}(v_{l,c}) + \mathrm{SNR}(R) \tag{2.77}$$

since $v_{l,c}$ is independent from all fragments composing R. The upper bound on its SNR is

$$\mathrm{SNR}(R) \leq \mathrm{SNR}(V^{K-l,c+l+1}).$$

The maximum SNR is achieved when $R = V^{K-l,c+l+1}$. $\square$

This theorem simplifies calculation of $U^{K,l,c}$ as long as the values of $V^{K-i,c+j}$ are available for values of i and j greater than one. The dynamic programming recursion is completed by the second key Theorem, relating $V^{K,c}$ to $U^{K,l,c}$.

Theorem 4 *For any instance of the RMF problem where $b = 1$,the subset $V^{K,c}$ satisfies*

$$V^{K,c} = \arg\max_{j \geq c} \max_{l=1,\ldots,K} \mathrm{SNR}(U^{K,l,j}) \tag{2.78}$$

where c is restricted to the interval $1 \leq c \leq N-K+1$.

Proof. Denote the first fragment of the subset $V^{K,c}$ by $v_{l,j}$. The indices for this fragment satisfy $j \geq c$ and $1 \leq l \leq K$ as a consequence of the definition of $V^{K,c}$. Since $V^{K,c}$ maximizes SNR for all K element subsets with first measurement bounded by c, it also maximizes SNR for any subset with initial fragment $v_{l,j}$ as long as $j \geq c$. Thus $V^{K,c} = U^{K,l,j}$. $\square$

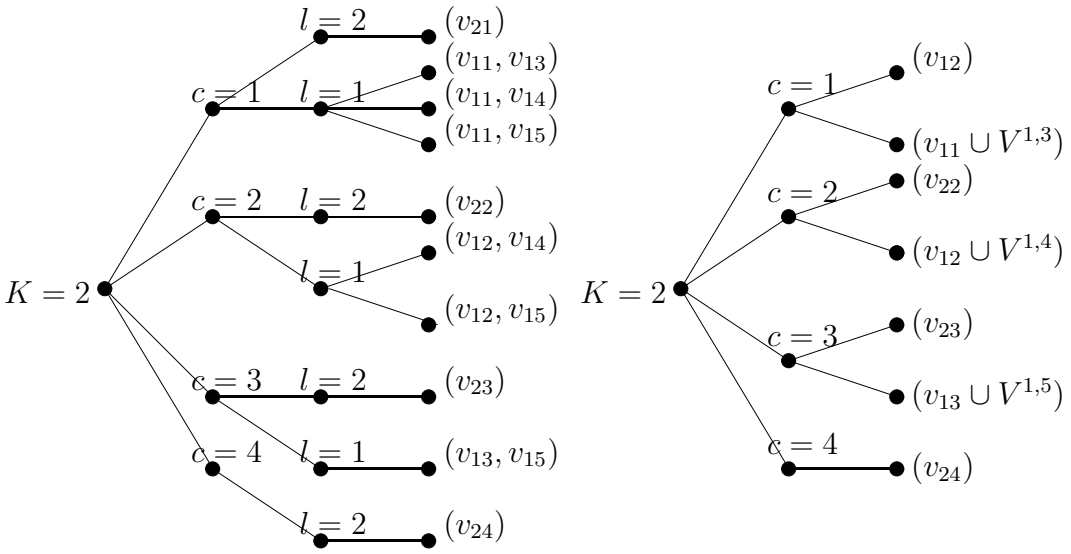

Figure 2-7: The top decision tree shows the entire set of feasible node choices for an RMF instance with $N = 4$, $K = 2$, and $b = 1$. The bottom decision tree shows the simplification achieved by applying the dynamic programming relations in Theorems (3) and (4).

The combination of Theorems 3 and 4 suggests a technique for calculating $V^{K,c}$ in the bootstrapped fashion characteristic of dynamic programming. Given $V^{i,c}$ for $1 \leq i \leq K$, $U^{K+1,l,j}$ can be calculated using Theorem 3. Maximization of the SNR metric over these subsets for $j \geq c$ yields $V^{K+1,c}$. Figures 2-7 and 2-8 display the decision trees for $N = 4$, $K = 2$ and $N = 5$, $K = 3$. The large trees show all possible fragment decompositions, sorted by c and l of each fragment. The small trees show the simplification in the search for the best subset due to the application of dynamic programming.

As is usual for dynamic programming, the algorithm determines the best subset for each value of K and the associated SNR. The final output of the algorithm is the tradeoff between subset size and SNR, and can be used to balance the costs of collecting and processing the data with the accuracy of the detector decision.

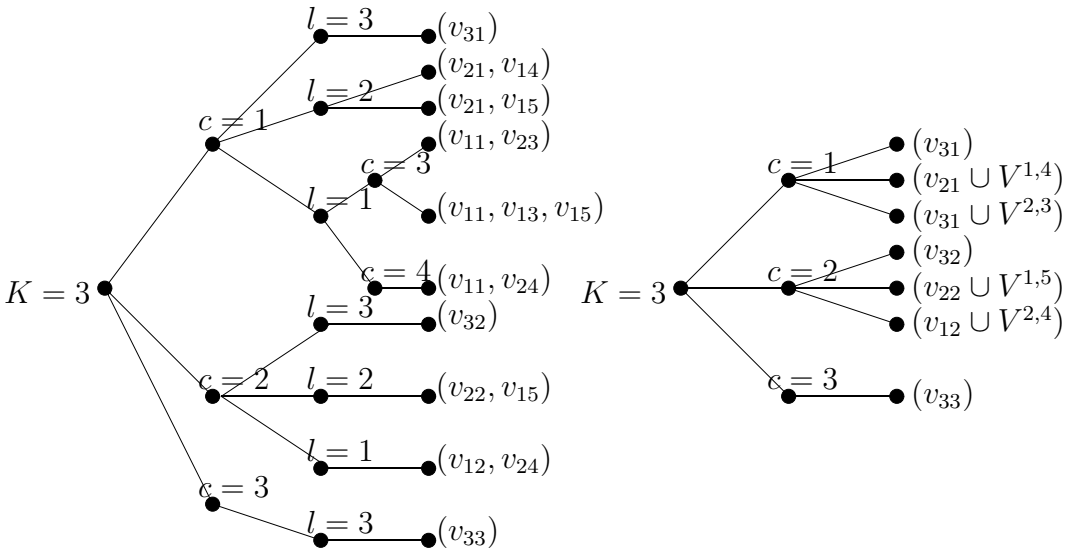

Figure 2-8: The top decision tree shows the entire set of feasible node choices for an RMF instance with $N = 5$, $K = 3$, and $b = 1$. The bottom decision tree shows the simplification achieved by applying the dynamic programming relations in Theorems (3) and (4).

2.5.5 Complexity of Dynamic Programming Algorithm

In order to evaluate the compuational and storage complexity of the dynamic programming algorithm, it suffices to determine the computation and storage required to complete a particular stage of the algorithm, and then sum up as K ranges from 1 to N.

For fixed K and c, the computation arises from maximizing $U^{K,l,c}$ over the possible values of l. These operations correspond to preserving the best leaf nodes from the second level groupings of the bottom decision trees in Figures (2.5.4) and (2.5.4). There are $N - K + 1$ possible starting points for this fragment, and the possible lengths of these fragments can range from 1 to K. Thus, the computation at stage K is bounded by (N - K + 1)(K). Summing from $K = 1$ to N gives an upper bound on the computation involved in the dynamic programming optimization:

$$C_{\text{opt}}(N) \leq \frac{1}{6} N(N+1)(N+2). \tag{2.79}$$

For each stage of the dynamic programming algorithm, a sort is required to determine $V^{K,c}$. For each stage, this contribues $O(N \log N)$ to the complexity. The dynamic programming computation dominates the computation from N sorts.

In addition to the optimization and sorting, the algorithm requires computation to calculate each of the SNR values for every correlated fragment. There are $N - K + 1$ fragments of K elements when $b = 1$. For each fragment, the complexity of calculating SNR is dominated by the $K \times K$ matrix inversion, and is bounded by $O(K^3)$. The total initialization complexity is $C_{\text{init}}(N) = O(N^5)$. The initialization costs exceed the actual cost of performing the optimization procedure.

The dynamic programming approach produces a polynomial time and space complexity search. This significant simplification over the $\begin{pmatrix} N \\ K \end{pmatrix}$ possibilites in the brute-force approach indicates the usefulness of the banded covariance matrix constraint.

2.5.6 Extension to General Banded Matrices

This section describes the dynamic programming algorithm applied to RMF instances with correlation distance $b > 1$. In such situations, the dynamic programming algorithm presented previously can be generalized to find the optimal subset for the RMF, at the cost of a more complicated indexing scheme for the correlated fragments and a more expensive initialization procedure. This section describes the changes in the notation and the algorithm required for the generalization, and compares the complexity of the results with those for the case when $b = 1$.

The first modification to the algorithm is a more detailed indexing scheme for subset fragments. For the purposes of the $b = 1$ algorithm, subset fragments are identified by two indices, the smallest element index in the fragment, denoted by c, and the total length of the fragment, denoted by l. If $b > 1$, however, two fragments composed of different numbers of measurements can have the same l and c indices. For example, if $b = 2$, the fragments $\{x_1, x_3\}$ and $\{x_1, x_2, x_3\}$ have indices $l = 3$ and $c = 1$. In order to identify the fragments for the purposes of the dynamic programming optimization algorithm, a third index, w, indicating the number of measurements in

the fragment, is required. Thus, we can identify by vertices $v_{w,l,c}$ the best fragment of w measurements, with minimum measurement index c, covering length l. In the new index scheme, $\{x_1, x_2, x_3\}$ has $w = 3$ and $\{x_1, x_3\}$ has $w = 2$. Note that the mapping between the fragments and the three indices is not invertible. There may be many distinct fragments that can be mapped to each (w, l, c) index set. During initialization, however, only the best fragment for each valid arrangement of (w, l, c) needs to be stored.

The restrictions between fragment indices necessary for them to satisfy the linear SNR decomposition must also be modified to account for general values of b. A pair of fragments $\mathbf{x}_{\{g,i\}}$ and $\mathbf{x}_{\{g,j\}}$ are not independent if any measurement in fragment i is correlated with an element in fragment j. Assume that fragment i has indices (w_i, l_i, c_i) and fragment j has indices (w_j, l_j, c_j). Without loss of generality, we can assume that $c_i < c_j$, so that fragment i starts first. Under these conditions, the two fragments are independent if and only if

$$c_i + l_i + b < c_j. \tag{2.80}$$

Initializing the dynamic programming algorithm becomes more complex as the correlation distance increases. The SNR values for each independent fragment must be calculated although only a fraction of them are used in the subsequent stages of the algorithm. Of the multiple fragments that are described by the same (w, l, c) triple, the one that has the maximum SNR is assigned to $v_{w,l,c}$ and used in the remainder of the algorithm.

After the initialization period, the dynamic programming algorithm can proceed. As in the case when $b = 1$, the algorithm generates the best subsets of fixed size $w = K$ satisfying $c \leq i$ for $i = 1, \ldots, N - w + 1$. The algorithm leverages past solutions to compute the best arrangements for the current stage. For a general $b > 1$, the state information for the algorithm is not changed, but the add-compare-select operation must be modified.

The state information needed for the algorithm $V^{w,c}$ is a table of the best w

measurement subsets with minimum element index $i \geq c$. Its definition is unchanged from the $b = 1$ case previously described. The definition of the intermediate state $U^{K,l,c}$, however, must be modified to account for the modifications in the indexing scheme for fragments.

The state $U^{K,w,l,c}$ is the optimal K element subset with the first fragment constrained to be $v_{w,l,c}$. The conditions in the definition for $U^{K,w,l,c}$ change to reflect the requirement that $v_{w,l,c}$ is the first fragment of the subset. Additionally, there are several constraints on the values of (w, l, c) for this subset. These bounds are summarized in the following Lemma.

Lemma 4 *Consider an RMF optimization problem with a banded covariance matrix possessing a correlation distance $b > 1$. For a subset $U^{K,w,l,c}$ satisfying definition 3, the relationship between the indices (w, l, c) take on two cases.*

In the first case, $U^{K,w,l,c} = v_{w,l,c}$, the indices satisfy

$$w = K \tag{2.81}$$

$$c \leq c_{\max} = N - K + 1 \tag{2.82}$$

$$l \leq l_{\max} = \min\{b(K-1) + 1, N - c + 1\}. \tag{2.83}$$

If $U^{K,w,l,c} = v_{w,l,c} \cup R$, where R is a non-empty subset of data, the indices satisfy

$$w < K \tag{2.84}$$

$$c \leq c'_{\max} = N - K + b + 1 \tag{2.85}$$

$$l \leq l'_{\max} = \min\{b(K-1) + 1, N - (c + b + K) + (w + 1)\}. \tag{2.86}$$

Proof. In both cases, the requirements on c and l guarantee that the largest measurement index in $U^{K,w,l,c}$ satisfies $i_{\max} \leq N$. This is a necessary condition since N is the largest index value.

In the first case, the subset U is composed of a single fragment. The minimum length of the fragment is $l \geq K$, achieved when all selected measurements are adjacent.

Additionally, for fixed c and l, the index of the final measurement is $i_{\max} = c + l + 1$. The upper bound on $i_{\max}$ yields the necessary condition

$$c + l - 1 \leq N. \tag{2.87}$$

The maximum value of c corresponds to the minimum value of l. Rearranging this equation and substituting $l = K$ gives the condition (2.82). Additionally, solving for l gives the bound $l \leq N - c + 1$. This bound is not always tight, especially if N is large and c is small. The bound on maximum correlation distance limits the difference between consecutive elements to b indices. There are $K - 1$ such pairs, so the bound given by the correlation distance is $l \leq b(K-1) + 1$. Combining the two upper bounds on l yields the condition (2.83).

In the second case, the subset U is composed of multiple fragments. Thus, there is a gap of at least b unselected measurements between the final element of $v_{w,l,c}$ and the first element of the remaining subset elements in R. In this situation, the maximum element is $i_{\max} = c + l + b + \text{length}(R) - 1$. Following the techinque from the first case, isolating c or l on the left hand side of the inequality equivalent to $i_{\max} \leq N$ and minimizing the right hand side produces these upper bounds on both indices

$$\begin{aligned} c &\leq N - (K + b) + 1 \\ l &\leq N - (K + c + b) + (w + 1). \end{aligned} \tag{2.88}$$

The first inequality follows since $l \geq w$ and $R \geq K - w$. Once again, combining the bound on l due to correlation distance with the second inequality in (2.88) produces the necessary condition (2.86). □

Theorems 3 and 4 require changes that reflect the indexing scheme for fragments with $b > 1$. The update of Theorem 3 relies on the conditions for a valid decomposition of $U^{K,w,l,c}$ from Lemma 4. The generalization of the theorem uses similar proof to Theorem 3

Theorem 5 *For any instance of the RMF problem, consider the subset $U^{K,w,l,c}$ with*

indices satisfying the conditions (2.84)-(2.86). This subset satisfies

$$U^{K,w,l,c} = v_{w,l,c} \cup V^{K-l,c+l+b}. \tag{2.89}$$

If the indexes satisfy (2.81)-(2.83), the subset is $U^{K,w,l,c}$.

Theorem 4 can be changed to account for the three indices that identify subset fragments. The updated theorem is:

Theorem 6 *For any instance of the RMF problem, the subset* $V^{K,c}$ *satisfies*

$$V^{K,c} = \arg \max_{j \geq c} \max_{1 \leq w \leq K} \max_{l \geq w} U^{K,w,l,j} \tag{2.90}$$

where the indices in the maximizations satisfy the constraints in Lemma 4.

The complexity of the dynamic programming optimization increases when $b > 1$. In the add-compare-select operations specified in equation (2.90), an upper bound on the number of terms the maximization examines is $(N - K + 1)K^2$ since there are $N - K + 1$ values of j to search over, and there are at most K^2 (w, l) index pairs for each value of j. Summation for $K = 1, \ldots, N$ shows that the computational complexity of the optimization algorithm satisfies

$$C_{\text{opt}} = O(N^4). \tag{2.91}$$

This exceeds the complexity for the case when $b = 1$ by a factor of N.

In order to perform the dynamic programming algorithm, an initialization procedure is necessary. As in the case when $b = 1$, the initialization computes SNR values for each independent subset. Unfortunately, the initialization procedure suffers from the curse of dimensionality that is familiar to dynamic programming in situations with many states. Since all elements no more than b indices apart are correlated, the initialization algorithm must examine all potential subsets of an interval of b measurements. Thus, a lower bound on the initialization complexity is $C_{\text{init}} \geq 2^b$. If b is large, the initialization procedure is no longer efficient.

2.6 Summary

This chapter examines data selection for detecting known signals in colored Gaussian noise. The traditional detector developed from the likelihood ratio test, the whitened matched filter, now depends on the subset selected. The performance of a restricted matched filter, a detector designed using a specific subset of measurements, depends upon a quadratic form that resembles signal-to-noise ratio. The search for the best RMF requires optimization of the SNR measurement, a problem considered in feature selection, data mining, and Artificial Intelligence research.

Two exact algorithms to find the optimal RMF are considered. In situations where the covariance matrix has a low condition number, the maximum energy subset leads to the optimal RMF. Additionally, for banded covariance matrices, dynamic programming produces the optimal subset selection.

Finally, heurisitc algorithms for optimization of SNR are evaluated. The forward and backward greedy selection algorithms, important for initialization of the local search and branch-and-bound optimization algorithms, are discussed. The condition number of the covariance matrix controls the worst-case performance of these algorithms. Worst-case instances of the optimization problem are shown for both greedy algorithms.

Chapter 3

Data Selection in Random Signal Detection

In detector design, the resulting algorithms depend strongly upon the *a priori* signal models used to describe data collected under the null and target hypotheses. In the case Chapter 2 discusses, the assumption of a deterministic target signal in the presence of additive Gaussian noise leads to the linear matched filter as a result of the likelihood ratio test. If other assumptions about the target signal are made, the likelihood ratio test generally produces non-linear mappings prior to the decision threshold. Furthermore, if the target signal cannot be described exactly in a concise model, the likelihood ratio test no longer is the optimal detector. Detectors for cases where the target signal is described by membership in a set of probability distributions with one or more unknown parameters are called non-parametric or semi-parametric detectors. When there is a lack of structure in the target signal model, and the uncertainty in the signal cannot be captured by a collection of unknown parameters, robust detection theories are often applied. A robust detector maximizes the worst-case performance of the detector over the uncertainty class of potential target signals.

This chapter discusses the interaction of data selection and detection when the target signals are specified by probability densities and less-structured uncertainty classes. We focus on the technique of randomized data selection and its use in non-parametric and robust detection scenarios. In this situation, we find that the robust

detectors are generated by the familiar square-law detector for sinusoids with unknown phase.

3.1 Random Signal Models

3.1.1 General Signal Model

In order to derive useful properties of the likelihood ratio test in the presence of random selection, we impose restrictions on the statistical model for the target signature. To balance the generality of the signal model with its special statistical structure, we assume that the probability density of the target signal is symmetric about the origin of the sample space. We shall refer to random vectors that satisfy this condition as even random vectors or even-symmetric signals. The precise definition of an even random vector is given in Definition 5.

Definition 5 *An N-dimensional random vector $\mathbf{s}$ is referred to as even if, for every $\mathbf{s}_o \in \Re^N$, its probability density function satisfies $p_{\mathbf{s}}(\mathbf{s}_o) = p_{\mathbf{s}}(-\mathbf{s}_o)$.*

This signal model establishes a useful structure on the probability density of the signal, enabling us to determine key properties of the likelihood ratio test. Additionally, the signal model is broad enough to model many interesting target signatures. For example, a sinusoid with an unknown, uniformly distributed phase satisfies the condition in Definition (5), as does a zero-mean, Gaussian random vector with a known covariance matrix.

The general binary hypothesis test for signals in additive Gaussian noise obeys the following statistical model:

$$\begin{aligned} H_0: \quad \mathbf{x} &= \mathbf{n} \\ H_1: \quad \mathbf{x} &= \mathbf{s} + \mathbf{n}. \end{aligned} \tag{3.1}$$

Here, we assume that $\mathbf{n}$, is an N-dimensional, zero-mean, white Gaussian random vector with covariance $\mathbf{\Lambda} = \sigma^2\mathbf{I}$. The signal vector $\mathbf{s}$ has an even-symmetric probabil-

ity density. Finally, we assume that $\mathbf{s}$ and $\mathbf{n}$ are independent random vectors. This model describes the statistics of the data without randomized selection.

While data selection algorithms accounting for many aspects of the network's state can be useful in practice, we choose a generic approach requiring limited *a priori* information and communication overhead. Specifically, we consider a randomized data selection strategy. This approach leads to useful algorithms in distinct fields such as estimation, hardware failure modeling, low power design [36], and theoretical computer science [25].

3.1.2 Notation

In our randomized selection rule, the decision to select measurement x_i depends on the outcome of an indicator random variable denoted g_i. The random variable is independent of all other indicator random variables and from other physically measurable quantities available to the detector. In our model, each measurement in the current time slot is selected with probability γ_g, i.e. g_i has the probability mass function

$$p_{g_i}(g) = \begin{cases} \gamma_g, & g = 1 \\ (1-\gamma_g), & g = 0. \end{cases} \tag{3.2}$$

This selection rule reduces the expected complexity of the detector implementation by a factor of γ_g since expected subset size is $E[K] = N\gamma_g$. Prior to discussing the specific detection problems, we examine the signal statistics for $\mathbf{x}_g$. For notational convenience, we will denote the conditional density of $\mathbf{x}_g$ given $\mathbf{G}$ by the expression

$$p_{\mathbf{x}_g|\mathbf{G}}(\mathbf{x}|\mathbf{G}). \tag{3.3}$$

3.2 Likelihood Ratio Test

In the presence of randomized data selection, the detector has access to the indicator random variables in $\mathbf{G}$ and processes the subset of the available data contained in $\mathbf{x}_g$.

The likelihood ratio for detectors with randomized selection can be expressed as

$$\begin{aligned} L(\mathbf{x}_g, \mathbf{G}) &= \frac{p_{\mathbf{x}_g,\mathbf{G}|H}(\mathbf{x}_g,\mathbf{G}|H_1)}{p_{\mathbf{x}_g,\mathbf{G}|H}(\mathbf{x}_g,\mathbf{G}|H_0)} \\ &= \frac{p_{\mathbf{x}_g|\mathbf{G},H}(\mathbf{x}_g|\mathbf{G},H_1)}{p_{\mathbf{x}_g|\mathbf{G},H}(\mathbf{x}_g|\mathbf{G},H_0)} \frac{p_{\mathbf{G}|H}(\mathbf{G}|H_1)}{p_{\mathbf{G}|H}(\mathbf{G}|H_0)} \\ &= L(\mathbf{x}_g|\mathbf{G}). \end{aligned} \tag{3.4}$$

The simplification in the likelihood occurs since the indicator random variables are independent of the hypotheses H_i.

Since conditioning upon $\mathbf{G}$ does not affect the selected data in $\mathbf{x}_g$, the detection problem based upon $\mathbf{x}_g$ and $\mathbf{G}$ reduces to an unconditional detection problem for the data associated with the non-zero indicator random variables. For example, if three pieces of data are available, there are eight possible arrangements of the indicator random variables. If measurements 1 and 2 are selected in time slot m, the detector must make a decision $\hat{H}$ based upon the joint densities

$$\begin{aligned} p_{\mathbf{x}_g|\mathbf{G},H}(\mathbf{x}|1,2,H_0) &= p_{x_1,x_2|H}(x_1,x_2|H_0) \\ p_{\mathbf{x}_g|\mathbf{G},H}(\mathbf{x}|1,2,H_1) &= p_{x_1,x_2|H}(x_1,x_2|H_1). \end{aligned} \tag{3.5}$$

Likewise, if measurements 2 and 3 are selected, the decision $\hat{H}$ is determined from $p_{x_2,x_3|H}(x_2,x_3|H_0)$ and $p_{x_2,x_3|H}(x_2,x_3|H_1)$.

Based upon (3.4), the likelihood ratio test for $\mathbf{x}_g$ and $\mathbf{G}$ reduces to the comparison of $L(\mathbf{x}_g|\mathbf{G})$ to a fixed threshold. While the test is optimal under the Neyman-Pearson detection criteria, it poses some practical problems. First, determining the threshold can become computationally complex when there is a large amount of data available for selection. The threshold that achieves a desired false alarm rate $\tilde{P}_F$ is determined by inverting the equation

$$\tilde{P}_F(\eta) = \sum_{\mathbf{G}} p_{\mathbf{G}|H}(\mathbf{G}|H_0)\Pr(L(\mathbf{x}_g) > \eta|\mathbf{G}, H_0). \tag{3.6}$$

If N samples of data are available, there are 2^N terms in the summation. Although it may be possible to approximate this function well by discarding terms with low

$p_{\mathbf{G}|H}(\mathbf{G}|H_0)$, determining the functional form of such an approximation may be troublesome. The functions of the threshold η given by $\Pr(L(\mathbf{x}_g) > \eta|\mathbf{G}, H_0)$ may not be easily parameterized. Second, since the threshold η is constant while $\mathbf{G}$ fluctuates, the conditional false alarm rate $P_F(\mathbf{G}, \eta) = \Pr(L(\mathbf{x}_g) > \eta|\mathbf{G}, H_0)$ fluctuates as well. In a situation where actions taken following a false alarm are costly, however, this fluctuation may not be desirable since it is induced by the random data selection rather than an information-bearing signal.

Faced with the practical difficulties of solving equation (3.6) for η, a suboptimal yet tractable alternative seems desirable. A reasonable approach fixes the conditional false alarm rate

$$P_F(\mathbf{G}, \eta) = \tilde{P}_F \tag{3.7}$$

for each realization of $\mathbf{G}$. Similar procedures have been discussed in [40] in a two-sensor situation. This constraint eliminates the fluctuations in the conditional false alarm rate and may simplify the implementation of the resulting detector since the constraint has a constant rather than exponential number of terms. It does, however, require the detector to adapt the test threshold to the arrangement of $\mathbf{G}$. In the remainder of the paper, we will focus on detectors designed with randomized data selection and the constraint imposed by equation (3.7).

In the remainder of this section, we discuss detector adaptation from several perspectives. First, we analyze the example problem of detecting a sinusoidal signal, a familiar target signature that satisfies the condition in Definition (5). Second, we demonstrate some properties of the likelihood ratio test for the general detection problem. Finally, we discuss semi-parametric detection and robust for any signal with an even-symmetric density.

3.2.1 Example - Detecting a Sinusoidal Signal

We consider detection of a sinusoidal signal in the presence of randomized data selection. Our analysis illustrates the difficulties associated with detection in the presence of uncertainty in the target signal and the challenges imposed by the fixed false alarm

requirement from equation (3.7).

Consider a set of data generated by sampling a signal at several locations, denoted by $v_i,\ i = 1, \ldots, N$. We shall assume that these locations can be modeled by a set of independent, identically distributed uniform random variables over an interval significantly larger than the sinusoid's wavelength.

Let H_0 denote the state in which the sinusoid is absent, and H_1 denote the state when it is present. The i-th measurement under each hypothesis is given by

$$\begin{aligned} H_0: \quad x_i &= n_i \\ H_1: \quad x_i &= A\cos(2\pi\tfrac{v_i}{\lambda} + \phi) + n_i. \end{aligned} \tag{3.8}$$

The random variable n_i is a zero-mean Gaussian random variable with variance σ^2. The probability density for $\mathbf{x}_g$, conditioned upon $\mathbf{G}$ and H_0 is white Gaussian. In order to determine the likelihood ratio and the resulting receiver operating characteristic, we also need the probability density for $\mathbf{x}_g$ conditioned upon $\mathbf{G}$ and H_1. This conditional density depends, in turn, on the joint density of

$$w_i = 2\pi\frac{v_i}{\lambda} + \phi \tag{3.9}$$

for the selected data in $\mathcal{S}$. The probability density for the signal is a function of the joint density of the phase random variables. Since the signal and noise are independent under H_1, the overall conditional density for $\mathbf{x}_g$ is the convolution of the signal density and the noise density. The determination of the joint density for the phase random variables is a key step in this calculation.

Since $\{v_i\}$ are independent and uniform over a large interval, we can approximate $\{w_i\}$ as independent, identically distributed uniform random variables over the region $[-\pi, \pi)$. Using this model, we can analyze the form of the likelihood ratio test for the model suggested in equation (3.8). Here, we assume that the base-station knows A exactly. The signal is $\mathbf{c}$, where $\mathbf{c}$ is a K-dimensional random vector. Each entry takes the form $c_i = A\cos(w_i)$. Based upon our approximation, the probability density for

$\mathbf{c}$ is

$$p_{\mathbf{c}|K}(\mathbf{c}|K) = \prod_{i=1}^{K} \frac{u(A - |c_i|)}{\pi\sqrt{A^2 - c_i^2}}, \tag{3.10}$$

where $u(\cdot)$ denotes the unit step function. This density is non-zero over the K-dimensional hypercube of side A. For fixed K, we denote[1] the randomly selected data by $\mathbf{x}_K$. This random vector lists the selected data contiguously, rather than with zeros as in $\mathbf{x}_g$. For notational convenience, we assume that measurements 1 to K are selected[2], so $\mathbf{x}_K = [x_1 \; x_2 \; \ldots \; x_K]^T$. The resulting signal model is

$$\begin{aligned} H_0: \quad \mathbf{x}_K &= \mathbf{n} \\ H_1: \quad \mathbf{x}_K &= \mathbf{c} + \mathbf{n}. \end{aligned} \tag{3.11}$$

Based on these probability density functions, we can construct the likelihood ratio test for fixed values of K and A. The conditional density under H_0 is Gaussian with zero mean. Under H_1, the conditional density is the convolution of the Gaussian with the density for $\mathbf{c}$ given in (3.10). The conditional density for $\mathbf{x}_K$ under H_1 can be written in terms of a one-dimensional convolution since both conditional densities are separable. The conditional density is

$$\begin{aligned} &p_{\mathbf{x}_K|K,H_1}(\mathbf{x}|K, H_1) = p_{\mathbf{c}|K}(\mathbf{x}_K|K) * p_{\mathbf{n}|K}(\mathbf{x}_K|K) \\ &= \textstyle\prod_{i=1}^{K} \int_{-\infty}^{\infty} \frac{u(A-|a_i|)}{\pi\sqrt{A^2-a_i^2}} \frac{1}{\sqrt{2\pi}\sigma} e^{\left(-\frac{(x_i-a_i)^2}{2\sigma^2}\right)} da_i \\ &= \textstyle\prod_{i=1}^{K} p_{x|H}(x_i|H_1). \end{aligned} \tag{3.12}$$

Since $\mathbf{x}_K$ under H_0 is a white, Gaussian random vector with variance σ^2, the likelihood ratio is

$$\begin{aligned} L(\mathbf{x}_K|K) &= \textstyle\prod_{i=1}^{K} \int_{-\infty}^{\infty} \frac{u(A-|a_i|)}{\pi\sqrt{A^2-a_i^2}} e^{\left(-\frac{(x_i-a_i)^2}{\sigma^2}\right)} e^{\left(-\frac{a_i^2}{2\sigma^2}\right)} da_i \\ &= \textstyle\prod_{i=1}^{K} L(x_i). \end{aligned} \tag{3.13}$$

[1]When applied to a vector, the subscript K indicates its dimension. This does not contradict our earlier notation, where the subscript of a scalar random variable indicated the identity of the measurement. The dimension subscript is always attached to a vector, not a scalar.

[2]This notation does not reduce the applicability of the analysis since our modeling assumptions make the measurements statistically indistinguishable. Their joint statistics depend only on K and not on the measurement identifiers.

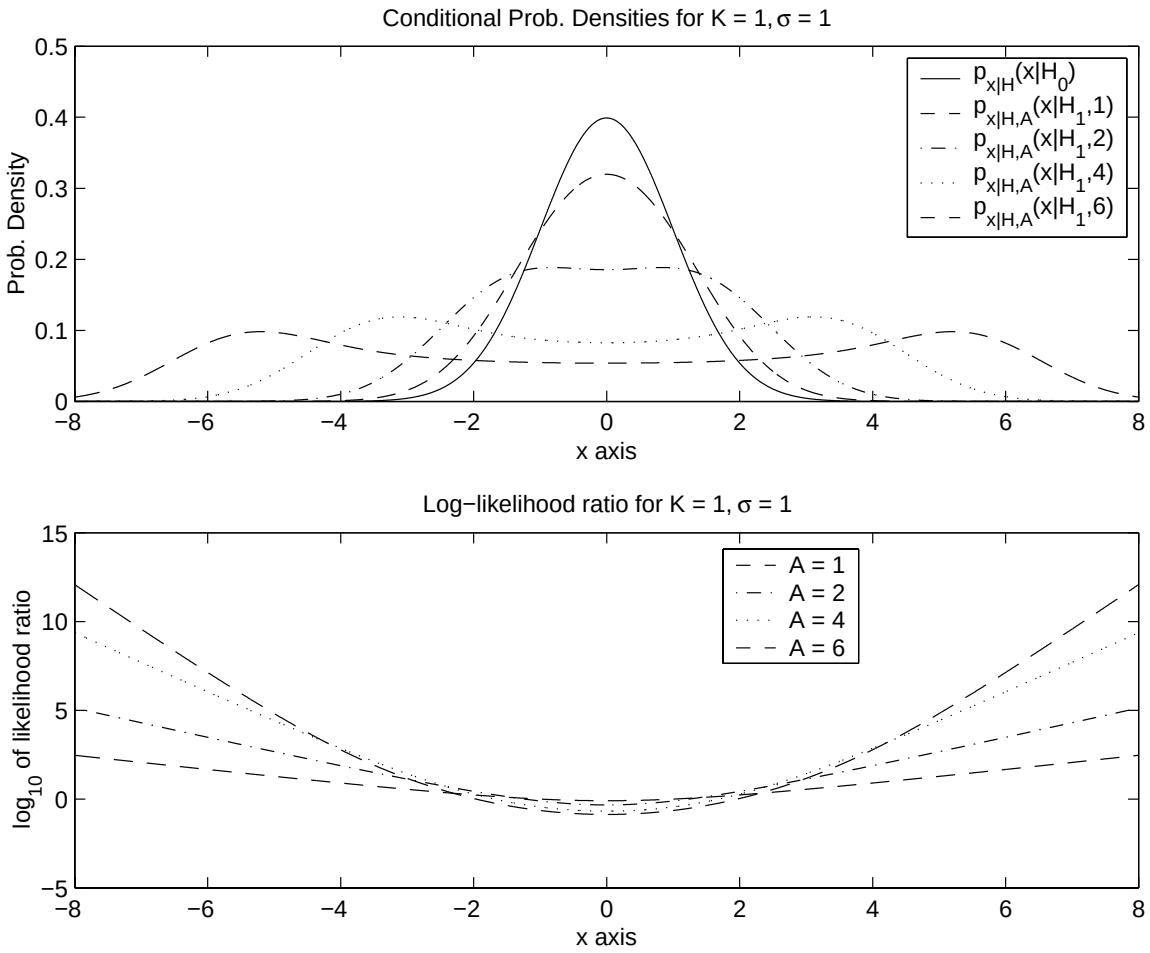

Figure 3-1: One dimensional conditional densities for the signal model defined in equation (3.11). The top plot shows the density for H_0 with a solid line. The densities for H_1 with $A = 1,\ 2,\ 4,$ and 6 become progressively wider. The second plot shows the corresponding log-likelihood ratios, $L(x|K = 1)$. In all plots, $\sigma^2 = 1$.

Analysis of the likelihood ratio test derived from equation (3.13) for $K = 1$ and $K = 2$ provides useful intuition about the general properties of the detector. The one-dimensional conditional densities, $p_{x|H}(x|H_1)$ and $p_{x|H}(x|H_0)$, and the associated log-likelihood ratios are shown in Figure 3-1. When $K = 1$, $L(x)$ is symmetric and increasing, so the likelihood ratio test from equation (1.3) simplifies to a threshold test of the form

$$|x| \underset{\hat{H}=H_0}{\overset{\hat{H}=H_1}{\gtrless}} t. \tag{3.14}$$

Since the detector compares the magnitude of the received data with a threshold, the implementation is simple.

Typically, the performance of a detector is shown by an operating characteristic, which plots the detection probability P_D as a function of the false alarm probability P_F. Both P_D and P_F can be calculated by integrating, respectively, the conditional densities $p_{\mathbf{x}_K|K,H}(\mathbf{x}_K|K, H_1)$ and $p_{\mathbf{x}_K|K,H}(\mathbf{x}_K|K, H_0)$ over the $\hat{H} = H_1$ decision region. Thus, the operating characteristic is generated as the threshold in equation (3.14) ranges over $0 \leq t < \infty$. It can be shown that the operating characteristic calculated from the likelihood ratio test gives the maximum achievable P_D for each false alarm rate $0 \leq P_F \leq 1$.

For $K = 1$, the detector described by equation (3.14) has an important universality property over the set of binary hypothesis tests for $A > 0$. The threshold that achieves a certain P_F can be determined in terms of the Q-function [44] and σ^2. Since the threshold can be determined without knowledge of the wave amplitude A, the test in equation (3.14) is a uniformly most powerful (UMP) test [31]. For such a test, the decision regions that maximize P_D subject to a constraint on P_F are invariant to the actual value of the parameter A. The actual value of P_D, however, does depend on A.

The contrast between the likelihood ratio test for $K = 1$ and $K = 2$ indicates some implementation challenges in the presence of uncertain signal models and random data selection. When $K > 1$, the likelihood ratio test for $\mathbf{x}_K$ is not a function of the received data magnitude $||\mathbf{x}_K||$, as shown in Figure 3-2. Since the likelihood

ratio is increasing in all directions, the likelihood ratio test will declare $\hat{H} = H_0$ in a simply connected region containing the origin. Outside this region, it will declare $\hat{H} = H_1$. Thus, the two-dimensional test determines a closed curve, expressed in polar coordinates as $r(\theta)$, that gives the boundary between the decision regions for $\hat{H} = H_0$ and $\hat{H} = H_1$. Since, $r(\theta)$ is not constant, the implementation of the likelihood ratio test is more complicated in two dimensions than in one.

For situations where $K > 1$, there is not a UMP detector. In order to determine $r(\theta)$ properly, we require $p_{\mathbf{x}_K|H}(\mathbf{x}|H_0)$, $p_{\mathbf{x}_K|H}(\mathbf{x}|H_1)$, and the desired value of P_F. As shown in Figure 3-2, the detector requires A to determine the decision regions in the likelihood ratio test.

Finally, the likelihood ratio test's decision regions depend on the value of K. The shape of the decision regions varies as K changes, as they did when K increased from 1 to 2. Evidently, larger values of K lead to more complicated decision regions. For example, the decision regions for $K = 2$ can be complicated sets in the (x_1, x_2) plane.

The difficulty in determining the decision regions under uncertainty in A and K makes the exact likelihood ratio test on $\mathbf{x}_K$ challenging to implement. First, the fluctuation in K means that the detector must be able to quickly adapt the decision regions for each time slot. Second, potential uncertainties in the target signal density prevent the detector from determining the exact likelihood ratio test. These challenges in the example detection problem persist for the general even signal model.

3.2.2 General Properties of the Likelihood Ratio for Even Signals

The sinusoid detection example illustrates several properties of the likelihood ratio in white, Gaussian noise. This section generalizes these properties to signals with even-symmetric probability densities. The qualitative behavior of the resulting decision regions is illustrated, and the prospects for practical implementation are discussed.

Following the derivation of equations (3.12) and (3.13), we can calculate the conditional likelihood ratio for an arbitrary signal with an even-symmetric probability

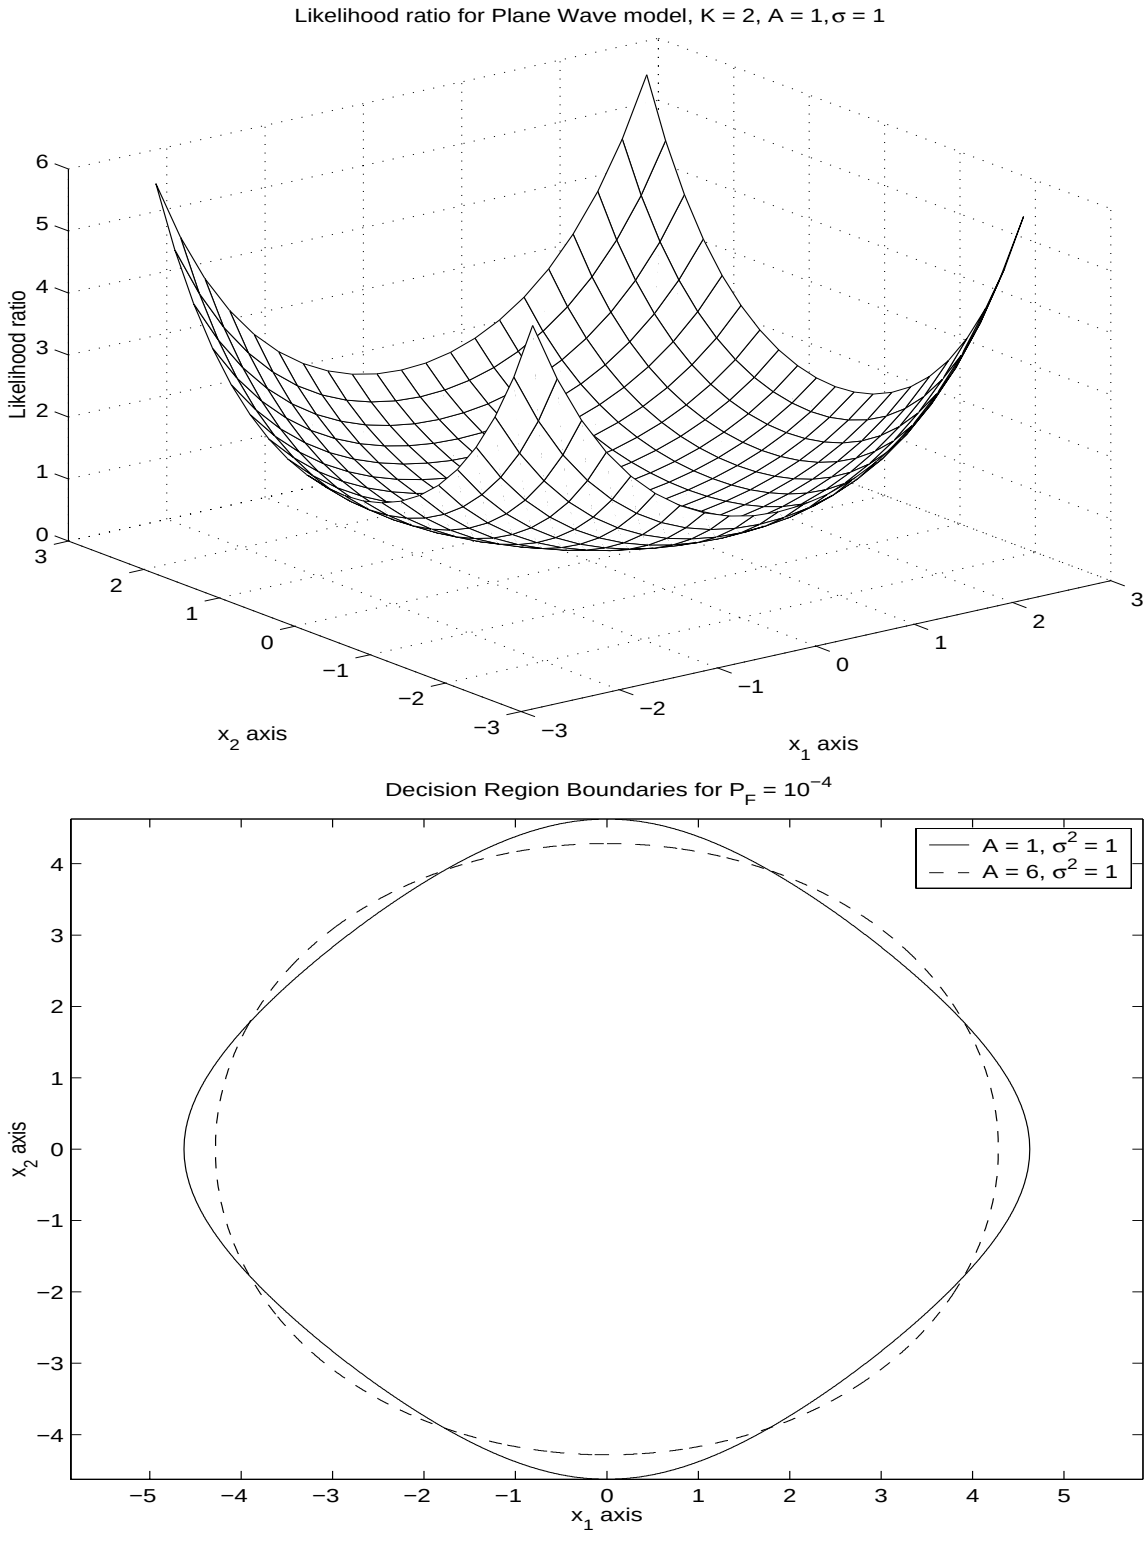

Figure 3-2: The top plot shows the two dimensional likelihood ratio when $A = 1$ and $\sigma^2 = 1$. The bottom plot shows the decision region boundary curves for $P_F = 10^{-4}$. In both cases, the noise variance is $\sigma^2 = 1$. The solid curve shows the boundary when $A = 1$ and the dashed curve shows the boundary when $A = 6$.

density. The binary hypothesis test's signal model is given by equation (3.1). The resulting expression for the conditional likelihood ratio is

$$L(\mathbf{x}_g|\mathbf{G}) = \int_{\mathbf{a}_K} p_{\mathbf{s}_g|\mathbf{G}}(\mathbf{a}_K|\mathbf{G}) e^{\left(-\frac{1}{2\sigma^2}\mathbf{a}_K^T\mathbf{a}_K\right)} e^{\left(\frac{1}{\sigma^2}\mathbf{x}_g^T\mathbf{a}_K\right)} d\mathbf{a}_K. \tag{3.15}$$

In the integral, the variable $\mathbf{a}_K$ is a K-dimensional vector, and the density $p_{\mathbf{s}_g|\mathbf{G}}(\mathbf{a}_K|\mathbf{G})$ denotes the joint density for the selected signal measurements.

The likelihood ratio is easily expressed in Cartesian coordinates. Its qualitative description, however, is easiest in generalized, K-dimensional spherical coordinates. When $K > 3$, the spherical coordinates can be determined via induction. In general, the transformation between spherical and Cartesian coordinates is expressed as

$$\begin{aligned}
x_1 &= r\sin(\theta)\textstyle\prod_{j=3}^{K}\sin(\phi_j) \\
x_2 &= r\cos(\theta)\textstyle\prod_{j=3}^{K}\sin(\phi_j) \\
x_3 &= r\cos(\phi_3)\textstyle\prod_{j=4}^{K}\sin(\phi_j) \\
&\vdots \quad\quad \vdots \\
x_{K-1} &= r\cos(\phi_{K-1})\sin(\phi_K) \\
x_K &= r\cos(\phi_K).
\end{aligned} \tag{3.16}$$

The domain of the radius is $r \geq 0$, and the domain of the angular variables is $\theta \in [0, 2\pi)$ and $\phi_i \in [0, \pi)$ for $i = 3, 4, \ldots, K$.

Using spherical coordinates, the boundary between the decision regions of the likelihood ratio test can be described. In the two-dimensional example, the curve dividing the decision regions is denoted by $r(\theta)$. In higher dimensions, we indicate the boundary surface by $r(\boldsymbol{\Theta})$, where the argument $\boldsymbol{\Theta} = [\theta\ \phi_3\ \ldots\ \phi_K]^T$ is a vector containing all the angular variables.

Without specific knowledge of the signal probability density $p_{\mathbf{s}_g|\mathbf{G}}(\mathbf{a}|\mathbf{G})$, the integral in equation (3.15) cannot be evaluated. The properties of even-symmetric signals, however, enable us to discover qualitative properties of the likelihood ratio. In spherical coordinates, we denote the likelihood ratio, conditioned upon $\mathbf{G} = \mathbf{G}$, by $L(r, \boldsymbol{\Theta}|\mathbf{G})$. In Appendix I, we prove the following theorem, showing that the likeli-

hood ratio test for an arbitrary even signal produces decision regions similar to those for the example.

Theorem 7 *Consider a detection problem of the class defined in equation (3.1). Let* $r = ||\mathbf{x}_g||$ *and* $\boldsymbol{\Theta} = [\theta \ \phi_3 \ \ldots \ \phi_K]^T$. *Then, the likelihood ratio* $L(r, \boldsymbol{\Theta}|\mathbf{G})$ *given by (3.15) increases monotonically without bound for any fixed* $\boldsymbol{\Theta}$.

As a consequence of this theorem, we can describe the likelihood ratio test in terms of $r(\boldsymbol{\Theta})$, a closed surface containing the origin. The interior of the surface is the decision region $\hat{H} = H_0$, and the remainder of the sample space is the decision region $\hat{H} = H_1$. In order to determine $r(\boldsymbol{\Theta})$, consider a fixed threshold η for the likelihood ratio test. Assuming that the probability densities contain no point masses, the subset of sample space satisfying the condition $L(r, \boldsymbol{\Theta}|\mathbf{G}) \leq \eta$ composes the decision region for H_0. Likewise, the condition $L(r, \boldsymbol{\Theta}|\mathbf{G}) > \eta$ determines the decision region for H_1. For a fixed vector $\boldsymbol{\Theta}_o$, Theorem 7 implies that the function of r given by $L(r, \boldsymbol{\Theta}_o|\mathbf{G})$ is strictly monotonically increasing. Thus, there is a unique solution to the equation $L(r, \boldsymbol{\Theta}_o|\mathbf{G}) = \eta$ if $\eta > L(0, \boldsymbol{\Theta}_o|\mathbf{G})$. The set of solutions generated as $\boldsymbol{\Theta}_o$ varies defines $r(\boldsymbol{\Theta})$, the boundary between the decision regions. Since the absolute minimum of the likelihood ratio occurs at $\mathbf{x}_g = \mathbf{0}$, the origin of the sample space is always included in the decision region for H_0, if it is non-empty. Hence, the qualitative description of the likelihood ratio test for the sinusoidal signal generalizes to any signal with an even-symmetric probability density.

Unfortunately, the implementation difficulties of the likelihood ratio test for the sinusoidal signal generalize as well. The detector must cope with fluctuations in the size of the selected subset and uncertainties in the *a priori* signal model. To implement the likelihood ratio test in all cases using randomized data selection, the detector should be able to determine the decision regions for each realization of $\mathbf{G}$. Additionally, for a fixed realization of the indicator random variables, determining these decision regions depends on the exact density $p_{\mathbf{x}_g|\mathbf{G},H}(\mathbf{x}_g|\mathbf{G}, H_1)$ and the desired false alarm probability $\tilde{P}_F$. If, as in the sinusoid detection model, there are unknown parameters in the signal model, or the signal is difficult to model *a priori*, the decision

regions that achieve $\tilde{P}_F$ and maximize P_D are difficult to determine.

3.3 Semi-parametric Detector

Techniques from the theory robust statistics have been applied to detector design in situations without precise *a priori* models [31, 30, 39]. Application of two such techniques can combat the challenges identified in the previous section. One technique introduced to cope with uncertainty in a signal model is invariance [37]. If a signal belongs to a class that is closed under some transformation, it is useful to design the detector so that its performance is also invariant to the transformation. The second challenge for the detector is the fluctuation in $\mathbf{G}$ and the associated task of rapidly updating the decision regions to satisfy the constraint in equation (3.7). This constraint enforces a constant false alarm rate (CFAR) condition on the detector: random fluctuations in the size of the selected data subset do not cause changes in the false alarm rate.

In this section we propose a low-complexity, semi-parametric detector addressing the implementation challenges of the exact likelihood ratio test. The detector is invariant to rotation of the target signal probability density, and maintains the CFAR property. It does use the noise variance as a known parameter, however.

The intuition behind our semi-parametric detector arises from the sinusoidal signal example when $K = 1$. In this case, the one-dimensional sample space simplifies the decision region and leads to a UMP detector. Even though we cannot find a UMP detector for $K > 1$, we can determine a test that has a weaker universality property over the class of even-symmetric random vectors. In this case, it is possible to determine a scalar-valued function of $\mathbf{x}_K$ so that there is a UMP test for the resulting random variable.

In this section, we use a detector based upon the scalar test statistic $r = ||\mathbf{x}_g||$. Note that this statistic is rotationally invariant. We analyze the properties of the likelihood ratio for r, and show that this statistic leads to a semi-parametric detector resembling the detector derived for $K = 1$ earlier. Based upon knowledge of the

Gaussian conditional density $p_{r|\mathbf{G},H}(r|\mathbf{G}, H_0)$, we design a threshold test

$$r \underset{\hat{H}=H_0}{\overset{\hat{H}=H_1}{\gtrless}} t, \tag{3.17}$$

where t is chosen to achieve a desired $\tilde{P}_F$. We show that this test has the maximum possible P_D for any set of decision regions based on r that achieve the false alarm rate $\tilde{P}_F$. Additionally, we show a method to calculate t that accounts for fluctuations in the selected subset and maintains the CFAR property.

3.3.1 Properties of the Likelihood Ratio for r

In order for the intuition based on the scalar detector to work, two conditions must be satisfied. First, the likelihood ratio test for r given $\mathbf{G}$ must reduce to a threshold test described in equation (3.17). This condition guarantees that the operating point of the test maximizes P_D over all decision regions that achieve the desired $\tilde{P}_F$, regardless of the target signal's density. Second, there ought to be an easy way to determine the proper threshold based upon a desired false alarm probability $\tilde{P}_F$. This allows the detector to be implemented with the required false alarm rate in real time. This subsection generates exact results for both these conditions, and the following subsection describes an approximate technique to calculate the test threshold.

In order to justify the simple threshold detector structure, we first verify that the likelihood ratio test simplifies to the form shown in equation (3.17). The likelihood ratio for r given $\mathbf{G}$ is

$$L(r|\mathbf{G}) = \frac{p_{r|\mathbf{G},H}(r|\mathbf{G}, H_1)}{p_{r|\mathbf{G},H}(r|\mathbf{G}, H_0)}. \tag{3.18}$$

Rather than calculate the ratio by direct computation of the densities $p_{r|\mathbf{G},H}(r|\mathbf{G}, H_0)$ and $p_{r|\mathbf{G},H}(r|\mathbf{G}, H_1)$, we can show the necessary result using the properties of $L(\mathbf{x}_g|\mathbf{G})$ established in the proof of Theorem 7. The following theorem, proven in Appendix II, establishes the validity of the threshold test on r.

Theorem 8 *For the statistical model established in equation (3.8), and a fixed value of K, let $r = ||\mathbf{x}_g||$. If the conditional densities $p_{\mathbf{x}_g|\mathbf{G},H}(\mathbf{x}_g|\mathbf{G}, H_i)$ for $i = 0, 1$ are continuous and positive, the likelihood ratio for $r = ||\mathbf{x}_g||$ increases monotonically without bound.*

In order to determine an appropriate CFAR detector, we require a rule for selecting a threshold t as a function of $\mathbf{G}$ and $\tilde{P}_F$. For a fixed t, the false alarm probability is

$$\begin{aligned} P_F &= Pr\{r > t|\mathbf{G}, H_0\} \\ &= \int_{r>t} p_{r|H,\mathbf{G}}(r|\mathbf{G}, H_0)dr \\ &= \int_{||\mathbf{x}_g||>t} p_{\mathbf{x}_g|\mathbf{G},H}(\mathbf{x}_g|\mathbf{G}, H_0)d\mathbf{x}_g. \end{aligned} \tag{3.19}$$

The density $p_{\mathbf{x}_g|\mathbf{G},H}(\mathbf{x}_g|\mathbf{G}, H_0)$ is a multi-variate Gaussian, so the integral (3.19) can be reduced the complementary distribution function (CDF) of a central χ^2 random variable of degree K [33]. In the case where $\sigma^2 = 1$, P_F is

$$P_F = \frac{1}{2^{K/2-1}\Gamma(K/2)} \int_t^\infty r^{K-1} e^{-r^2/2} dr. \tag{3.20}$$

This integral can be expressed using the incomplete gamma function denoted by $\Gamma(t^2/2, K/2)$. The false alarm rate in this situation is

$$P_F = 1 - \frac{\Gamma(t^2/2, K/2)}{\Gamma(0, K/2)}. \tag{3.21}$$

An exact determination of the test threshold that achieves a desired false alarm rate $\tilde{P}_F$ requires inversion of equation (3.21). In situations where the false alarm rate is always fixed to a single value, a lookup table for t versus K may be appropriate. If the application scenario requires that $\tilde{P}_f$ vary over time, the detector must be able to calculate the threshold numerically. Techniques to perform this calculation are described in [10].

3.3.2 Low-Complexity Algorithm to Calculate the Detector Threshold

In this subsection, we propose a threshold calculation based on an approximation of the conditional density for r under H_0 as a Gaussian with its parameters chosen as functions of K and $\tilde{P}_F$. The algorithm then uses the Gaussian approximation to calculate the threshold that would achieve $\tilde{P}_F$. This section argues for the plausibility of this approximation procedure, and evaluates its accuracy in determining t.

The algorithms for inverting equation (3.21) are iterative in nature. They may be computationally expensive to execute frequently. This section proposes an approximation appropriate for low false alarm rates that is based on the inverse of the Q-function. This can be accomplished without iterative algorithms using a rational approximation [8], and is less expensive to compute.

As shown in equation (3.20), the conditional density for r under H_0 is proportional to the term $q_K(r) = r^{K-1}e^{-r^2/2}$. For large values of r, the exponential decay dominates the behavior of this function. For values of r near zero, however, the r^{K-1} term dominates. The overall behavior is that of a sharply peaked function cresting at $r_{\max} = \sqrt{K-1}$. Figure 3-3 shows the normalized density for several values of K. Each peak resembles a Gaussian. Examining the logarithm of the function yields further insight on the resemblance. Taking the logarithm of $q_K(r)$ separates the terms in the function as

$$\log q_K(r) = (K-1)\log(r) - \frac{1}{2}r^2. \tag{3.22}$$

The function $\log q_K(r)$ is shown in Figure 3-4. Since the $\log(r)$ term grows slowly for large r, the apparent drop of the function near $r_{\max}$ resembles the quadratic $\frac{1}{2}r^2$. Specifically, the second derivative $q_K''(r_{\max}) = 2$ for all values of K. Consequently, the Gaussian with $\sigma = \frac{1}{2}$ approximates the behavior of $p_{r|H}(r|H_0)$ well near its peak at $\sqrt{K-1}$.

In order to estimate the density farther from the peak of $q_K(r)$, we extend the Gaussian approximation. Since the $\log(r)$ term in equation (3.22) varies slowly, it is possible to approximate $\log q_K(r)$ over several standard deviations by a quadratic

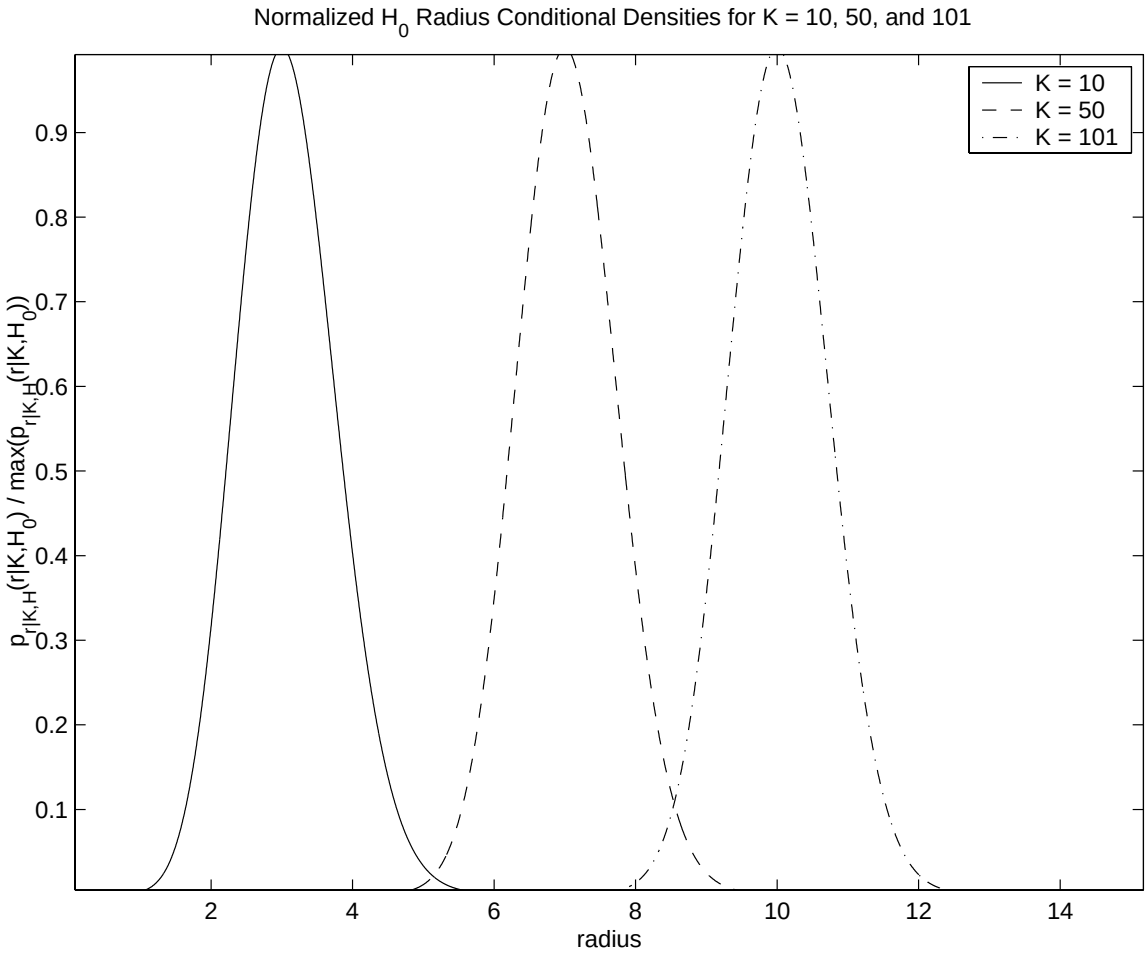

Figure 3-3: Normalized radial densities for $K = 10,\ 50,$ and 101. Each radial density arises from a Gaussian density for $\mathbf{x}_g$ with covariance matrix $\mathbf{I}$. The plotted curves are $\frac{P_{r|K,H}(r|K,H_0)}{\max(p_{r|K,H}(r|K,H_0))}$. Each appears similar to a Gaussian with standard deviation near $\frac{1}{2}$.

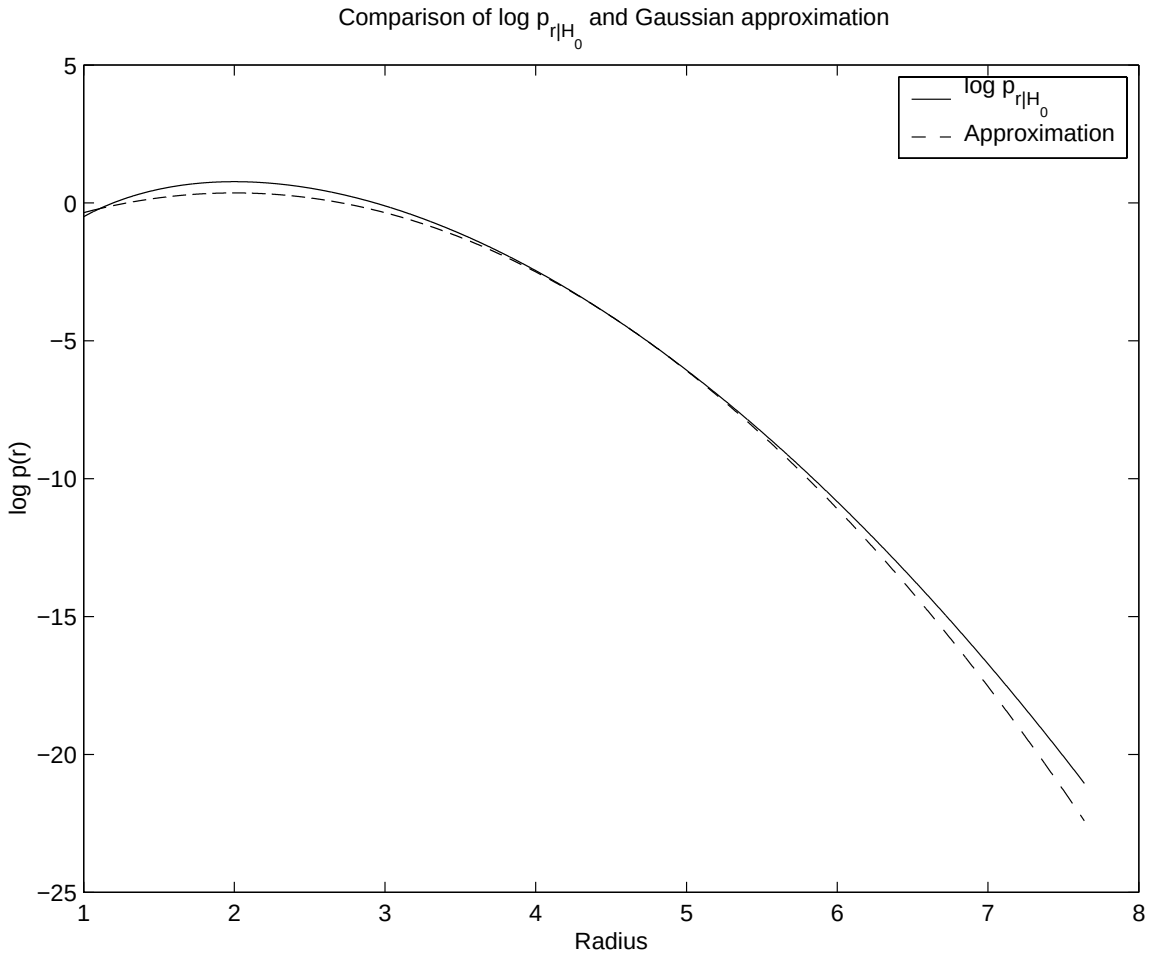

Figure 3-4: Approximation of $\log a_K(r)$ by a parabola. The plot displays a case when $K = 5$. The parabolic approximation is designed to fit well enough to determine thresholds necessary to achieve $P_F = 10^{-4}$.

with properly chosen peak and curvature. Essentially, the approximation of $\log q_K(r)$ resembles a Gaussian with its mean and variance adjusted to account for K and $\tilde{P}_F$. Figure 3-4 shows an example of approximating $\log a_K(r)$ with a parabola given by $\log g_K(r) = D - \frac{(r-\mu)^2}{2\sigma^2_{\text{approx}}}$. For properly chosen values of m, D, and σ^2_{approx}, the approximation can fit $\log a_K(r)$ closely over a particular interval. Since the conditional density decays quickly for large values of r, the interval where the approximation is accurate only needs to be a few standard deviations wide.

There are many possible ways to approximate $q_K(r)$ in order to determine an appropriate detector threshold. Our Gaussian approximation is designed to be accurate for small values of $\tilde{P}_F$, less than 10^{-1}. We suspect that this is a reasonable range

of operation for many detectors, since false alarms will initiate subsequent processing, expending power and communication resources. Detectors designed to operate efficiently would typically avoid a high false alarm rate.

In order to have a low false alarm rate, the detector threshold should at least be greater than $\sqrt{K-1}$, the peak location of $\log q_k(r)$. Our approximation is based upon determining a parabola that accurately fits $\log q_K(r)$ over an interval of r sufficiently large to suggest that the exact P_F is near $\tilde{P}_F$.

The Gaussian approximation has three free parameters, the mean m, variance σ^2, and the amplitude D. In our approximate fitting procedure, we assign $\mu = r_{\max}$, which guarantees that the peak of the approximate density coincides with the peak of $p_{r|\mathbf{G},H}(r|\mathbf{G}, H_0)$. Additionally, we will choose σ^2_{approx} and D so that the approximation intersects the true conditional density in two locations, denoted r_a and r_b. These fit points are chosen so that the interval $[r_a, r_b]$ is likely to contain the value of t that produces $\tilde{P}_F$.

The values r_a and r_b are determined by making an educated guess of the true threshold, and centering the fit points around it. The curvature of $\log q_K(r)$ at its peak is the same as a Gaussian with standard deviation $\sigma_{\text{init}} = 1/2$. Thus, we generate an initial guess for the threshold using a crude Gaussian approximation. The initial guess is

$$t_{\text{init}} = \mu + \sigma_{\text{init}} Q^{-1}(\tilde{P}_F). \tag{3.23}$$

Given this guess, the fit points are chosen with

$$\begin{aligned} r_a &= t_{\text{init}} - 0.01 \\ r_b &= t_{\text{init}} + 0.01. \end{aligned} \tag{3.24}$$

The choice of r_a and r_b is *ad hoc*. In general, however, it provides a close fit between $\log q_K(r)$ and the approximation on an interval extending several standard deviations past the true threshold. Since the true density is dominated by the $\exp(-r^2/2)$ term past its peak, the false alarm probability (3.20) is concentrated in the first several standard deviations past the threshold. The approximation does not fit accurately

far from t_{init}, however, this will not have a significant impact on the approximation accuracy for thresholds near r_a or r_b.

The parameters of $\log g_k(r)$ depend on the fit points and K. The fitting error between $\log q_K(r)$ and $\log g_K(r)$ is

$$E(r) = D - \frac{(r-\mu)^2}{2\sigma^2_{\text{approx}}} - (K-1)\log(r) + \frac{r^2}{2}. \tag{3.25}$$

The parameters are determined by requiring $m = \sqrt{K-1}$, $E(r_a) = 0$, and $E(r_b) = 0$. Solving for A and σ^2 yields

$$\begin{aligned} \sigma^2_{\text{approx}} &= \frac{r_b + r_a - 2\mu}{r_b + r_a - 2\frac{\mu^2}{r_b - r_a}\log\left(\frac{r_b}{r_a}\right)} \\ D &= \mu^2 \log(r_a) - \frac{r_a^2}{2} + \frac{(r_a-\mu)^2}{2\sigma^2_{\text{approx}}}. \end{aligned} \tag{3.26}$$

The coefficients produce the approximation $\log g_K(r) = D - \frac{(r-\mu)^2}{2\sigma^2_{\text{approx}}}$. In practice, this function fits $\log a_K(r)$ well over a range of r near t_{init}. Figure 3-4 shows an example of the curves for $K = 5$ and $P_F = 10^{-4}$.

The approximation gives a technique for determining the threshold to achieve a small $\tilde{P}_F$ for a given K. Adjusting for the constant factors in the integrals (3.20), the threshold is

$$t = \mu + \sigma_{\text{approx}} Q^{-1}\left(\tilde{P}_F \frac{2^{K/2-1}\Gamma(K/2)}{e^D}\right). \tag{3.27}$$

In order to verify the approximation accuracy, we compare the actual P_F versus $\tilde{P}_F$ over a range of exponentially spaced values from 10^{-2} to 10^{-12}. Each $\tilde{P}_F$ is one-tenth the previous value. Figure 3-5 shows the behavior of the approximation for several values of K. In general, the approximation appears to produce P_F values slightly smaller than $\tilde{P}_F$. The error increases with K. In general, the approximation is accurate to within a factor of 1.5 for desired false alarm rates between 0.01 and 10^{-12} and K between 5 and 200. A Matlab implementation of the approximate threshold calculation is nearly a factor of 13 faster than the exact threshold calculation based on inversion of the incomplete gamma function.

The actual performance of the threshold detector for the sinusoid detection prob-

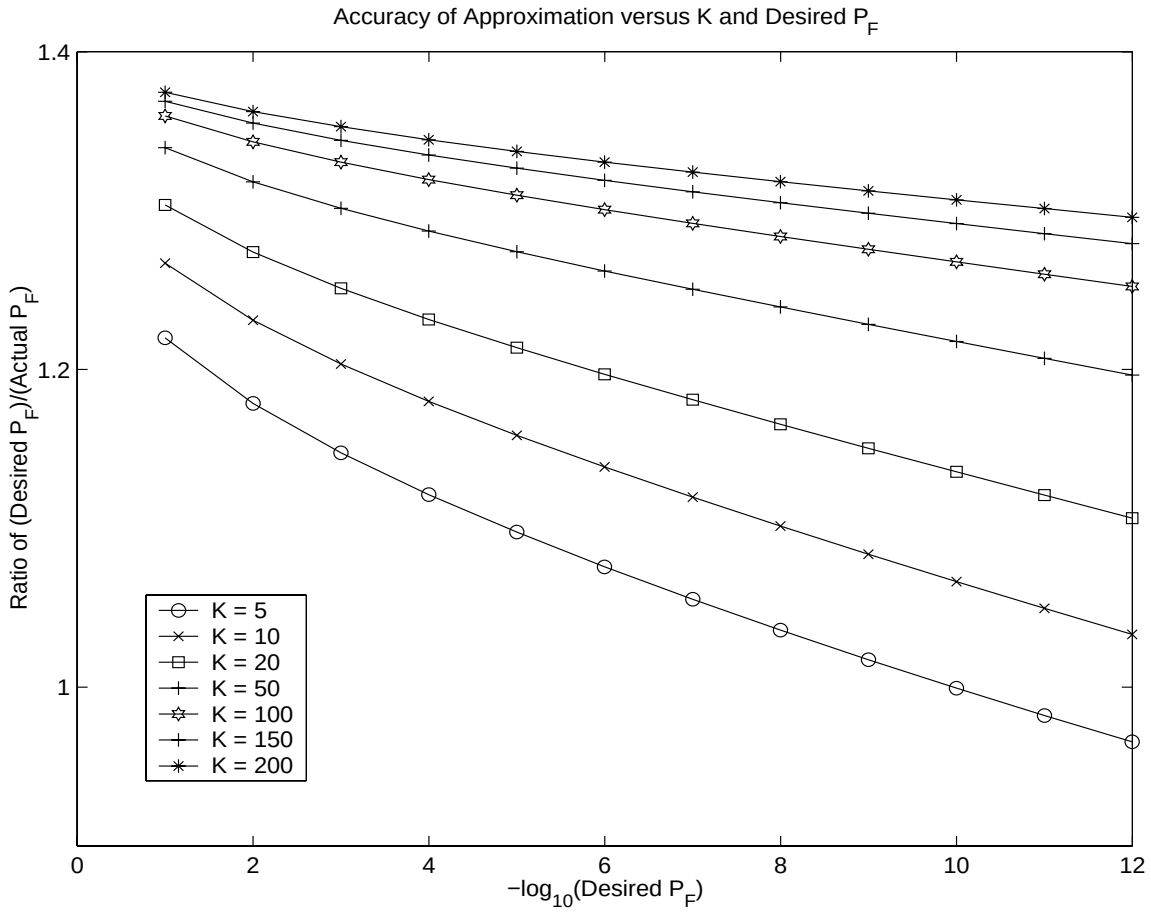

Figure 3-5: Approximation accuracy. The ratio of the actual false alarm rate, P_F, to the desired value $\tilde{P}_F$ is shown for $K = 5, 10, 20, 50, 100, 150, 200$. The desired false alarm rate varies from $\tilde{P}_F = 10^{-2}$ to $\tilde{P}_F = 10^{-12}$.

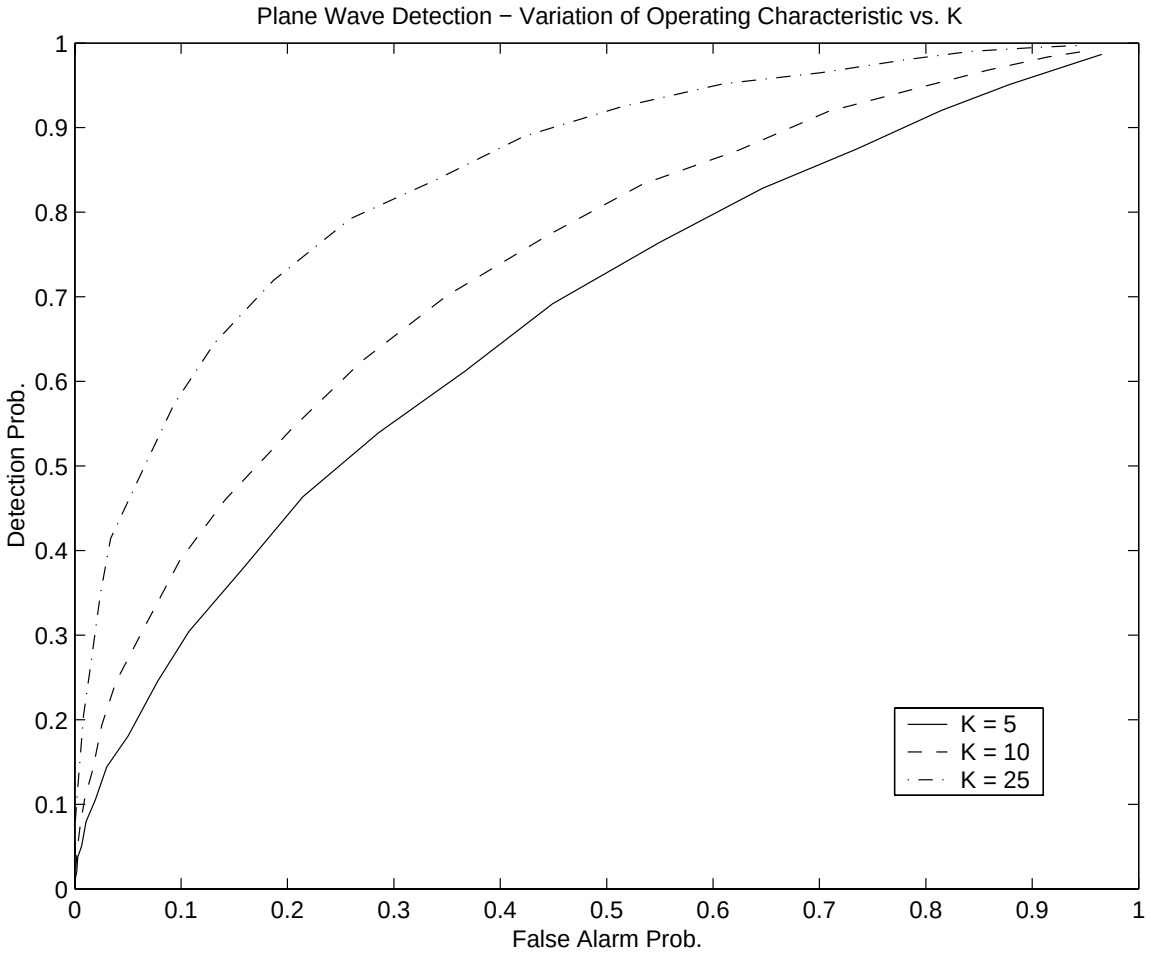

Figure 3-6: The receiver operating characteristic of the example problem is shown for $K = 5,\ 10,$ and 25.

lem described in part B is shown in Figure 3-6. The figure shows the performance of the detector for several values of K. In all cases the noise variance is $\sigma^2 = 1$, and the wave's amplitude is $A = 1$. The simulations shown in Figure 3-6 indicate that the performance of the detection algorithm improves as K or the ratio A/σ increase. These results are not particularly surprising, since increasing K gives the detector more raw data, and increasing A/σ improves the signal-to-noise ratio of each individual measurement. The results are useful primarily as a low-complexity baseline for the performance of detection schemes based upon more restrictive statistical models or data selection algorithms based on the details of such models.

3.3.3 Robustness Properties of the Threshold Detector

For a many sets of potential target signals, especially if they impose little structure on the signal density, it is unlikely that a single detector will maximize P_D at some fixed value of P_F for every potential signal density. Detectors that have this property are called uniformly most powerful (UMP) [31]. In the absence of a UMP detector, robustness criteria are used to design detectors that must operate with out precise *a priori* knowledge of the conditional probability densities for $\mathbf{x}$. A detector is robust if, for some set of potential detection algorithms, it maximizes the worst-case P_D at fixed P_F for any of the potential signal densities. This section determines a robust detector for signals with symmetric probability densities in white Gaussian noise with known variance.

We focus our attention on the even-symmetric signal densities $p_{\mathbf{s}}(\mathbf{s})$ that satisfy the following condition:

Definition 6 *A set of probability densities is distinguishable from zero if, for every density in the set,* $\Pr[||\mathbf{s}|| \geq \epsilon] = 1$ *for some* $\epsilon > 0$*. Here,* ϵ *is constant over the entire signal class, and does not depend on any individual density.*

The set of signals satisfying these two constraints is $\mathcal{S}_N$. The symmetry constraint gives the signal density useful structure, and is satisfied for many interesting target signals such as sinusoids with random phase or zero-mean Gaussians. The distinguishablilty constraint ensures that $\mathbf{s} + \mathbf{n} \neq \mathbf{n}$ with probability 1. Subject to these constraints, we allow both discrete and continuous densities for the target signal.

For the detection algorithm under consideration, we constrain the region $R = \{\mathbf{x} \in \Re^N | \hat{H}(\mathbf{x}) = H_0\}$ to a specific subset of all potential regions in $\Re^N$. A decision region is valid if it satisfies the following requirements.

Definition 7 *A decision region* R *is valid if*

1. $R = \rho(R_o)$ *where* ρ *is any rotation in* $\Re^N$.

2. $R_o = R_j \otimes \Re^{N-j}$.

3. R_j is an origin-symmetric, bounded, convex region in $\Re^j$.

The set of detectors satisfying these constraints is $\mathcal{R}_N$. Examples of such a region when $N = 2$ are the interior of an ellipse or the region $|x_1| < 2$. Without loss of generality, we will assume that R_j constrains $x_1, \ldots, x_j$ and the remaining variables in $\mathbf{x}$ are not constrained by the decision region.

For a specific signal density $p_{\mathbf{s}} \in \mathcal{S}_N$ and a specific decision region $R \in \mathcal{R}_{\mathcal{N}}$, the detection and false alarm probabilities are denoted by $P_D(R, p_{\mathbf{s}})$ and $P_F(R)$. They can be represented by the integrals

$$\begin{aligned} P_D(R, p_{\mathbf{s}}) &= 1 - \int_R p_{\mathbf{x}|H}(\mathbf{u}|H_1)d\mathbf{u} \\ &= 1 - \int_R (p_{\mathbf{s}} * p_{\mathbf{n}})(\mathbf{u})d\mathbf{u} \\ P_F(R) &= 1 - \int_R p_{\mathbf{x}|H}(\mathbf{u}|H_0)d\mathbf{u} \\ &= 1 - \int_R p_{\mathbf{n}}(\mathbf{u})d\mathbf{u}. \end{aligned} \tag{3.28}$$

For our purposes, the noise density is $p_{\mathbf{n}}(\mathbf{u}) = \mathcal{N}(\mathbf{u}; \mathbf{0}; \sigma^2\mathbf{I})$, i.e. white Gaussian noise.

The following theorem indicates that the threshold detector derived in the previous section is a robust detector over the class of signal densities and decision regions defined above.

Theorem 9 *Let $R' \in \Re^N$ be the H_0 decision region given by $R' = \{\mathbf{x} \in \Re^N | \mathbf{x}^T\mathbf{x} \leq \eta\}$. For any other decision region $R \in \mathcal{R}_N$ that satisfies $P_F(R) = P_F(R')$, the following condition holds:*

$$\min_{p_{\mathbf{s}} \in \mathcal{S}_N} P_D(R', p_{\mathbf{s}}) \geq \min_{p_{\mathbf{s}} \in \mathcal{S}_N} P_D(R, p_{\mathbf{s}}). \tag{3.29}$$

Thus, the spherical decision region is the maxmin robust decision region over $\mathcal{S}_N$ and $\mathcal{R}_N$ for any $0 < P_F < 1$.

In order to prove the theorem, we rely on a generalization of Lemma IV.2 from the paper by Gay *et. al.* [28]. They show that for a spherically symmetric, decreasing noise probability density function, the integral $\int_{\|\mathbf{u}\|<\eta} p_{\mathbf{n}}(\mathbf{u} - \mu\mathbf{a})d\mathbf{u}$ is a monotonically decreasing function of $|\mu|$ for any $\mathbf{a}$. Generalization of this result to any region in $\mathcal{R}_N$ is a key step towards proving Theorem 9.

In order to make the argument precise, we define the following notation. For any region $R \in \mathcal{R}_N$, let the function $f_R(\mathbf{a})$be

$$f_R(\mathbf{a}) = \int_R p_{\mathbf{n}}(\mathbf{u} - \mathbf{a}) d\mathbf{u}. \tag{3.30}$$

Since R is a convex, simply connected region, it is measureable, and the function $f_R(\mathbf{a})$ is well defined for all $\mathbf{a} \in \Re^N$.

An important step in the proof of Theorem 9 involves showing that $f_R(\mathbf{a}) \leq f_R(\mathbf{0})$ for any $\mathbf{a}$. The following lemma, proved in Appendix C, formalizes the result.

Lemma 5 *Consider the function $f_R(\mathbf{a})$ defined in equation (3.30). If the domain of integration in the expression satisfies $R \in \mathcal{R}_N$ and the probability density satisfies*

$$p_{\mathbf{n}}(\mathbf{u}_1) \geq p_{\mathbf{n}}(\mathbf{u}_2) \tag{3.31}$$

whenever $\|\mathbf{u}_1\| \leq \|\mathbf{u}_2\|$, $\mathbf{u}_1 = \mu \mathbf{u}_2$ for some constant $\mu \geq 0$, then

$$f_R(\mathbf{a}_1) \geq f_R(\mathbf{a}_2) \tag{3.32}$$

$\|\mathbf{a}_1\| \leq \|\mathbf{a}_2\|$, $\mathbf{a}_1 = \nu \mathbf{a}_2$ for some constant $\nu \geq 0$.

Armed with this result, we can determine the worst-case signal density for a specific decision region R. The detection probability can be minimized by placing all of the probability mass for the target signal at locations $\mathbf{x}_{\max} = \arg\max_{\|\mathbf{a}\|=\epsilon} f_R(\mathbf{a})$. The weaker the signal magnitude, the more difficult it is to detect. Because $f_R(\mathbf{a})$ is a decreasing function of $\|\mathbf{a}\|$, a maximum will occur on the constraint sphere $\|\mathbf{s}\| = \epsilon$, although it may not be unique. Since R is symmetric, the locations that satisfy this condition are also symmetric. A signal density satisfying our symmetry constraint is

$$p_{\mathbf{s}}(\mathbf{s}) = \frac{1}{2}(\delta(\mathbf{s} - \mathbf{x}_{\max}) + \delta(\mathbf{s} + \mathbf{x}_{\max})) = p_B(\mathbf{s}). \tag{3.33}$$

Note that this density takes the familiar form of a binary signal constellation. For any valid decision region R, it is possible to construct a binary signal constellation that

gives the worst case performance in terms of $P_D(R, p_{\mathbf{s}})$. The constellation depends upon the particular decision region.

For the spherical decision region $\Phi_t = \{\mathbf{x} | \|\mathbf{x}\| \leq t\}$, we can derive a useful invariance property for $P_D(\Phi_t, p_{\mathbf{s}})$. Due to the rotational symmetry of the decision region, the detection probability does not depend upon $p_{\mathbf{s}}$ if $\|\mathbf{s}\| = \epsilon$ with probability one.

Lemma 6 *Consider the decision region* $\Phi_t = \{\mathbf{x} | \|\mathbf{x}\| \leq t\}$. *For any signal density* $p_{\mathbf{s}} \in \mathcal{S}_N$ *satisfying the constraint*

$$\Pr(\|\mathbf{s}\| = \epsilon) = 1, \tag{3.34}$$

the detection probability satisfies

$$P_D(\Phi_t, p_{\mathbf{s}}) = P_D(\Phi_t, p_{U_\epsilon}) \tag{3.35}$$

where p_{U_ϵ} *is uniformly distributed on the sphere* $\|\mathbf{s}\| = \epsilon$.

Proof. For the decision region Φ_t, the function $f_{\Phi_t}(\mathbf{a})$ depends only on $\|\mathbf{a}\|$. If $\|\mathbf{a}\| = a$, we denote the function by the scalar argument $f_{\Phi_t}(a)$.

Using the expression in equation (3.28), the detection probability for a signal density $p_{\mathbf{s}}$ satisfying the constraint in equation (3.34) is given by

$$\begin{aligned}
P_D(\Phi_t, p_{\mathbf{s}}) &= 1 - \int_{\mathbf{u}} \int_{\mathbf{v} \in \Phi_t} p_{\mathbf{s}}(\mathbf{u}) p_{\mathbf{n}}(\mathbf{v} - \mathbf{u}) d\mathbf{u} d\mathbf{v} && (3.36)\\
&= 1 - \int_{\mathbf{u}} p_{\mathbf{s}}(\mathbf{u}) \left(\int_{\Phi_t} p_{\mathbf{n}}(\mathbf{v} - \mathbf{u}) d\mathbf{v} \right) d\mathbf{u} && (3.37)\\
&= 1 - \int_{\mathbf{u}} p_{\mathbf{s}}(\mathbf{u}) f_{\Phi_t}(\epsilon) d\mathbf{u} && (3.38)\\
&= 1 - f_{\Phi_t}(\epsilon) = P_D(\Phi_t, p_{U_\epsilon}). && (3.39)
\end{aligned}$$

The expression in equation (3.38) follows because the bracketed integral in the previous line depends upon $\mathbf{u}$ only through its magnitude. For our signal density, this magnitude is constant with probability one. □

As a result of this lemma, we can say several interesting things about the properties of the decision region Φ_t. First, the performance of Φ_t for any binary signal constellation (3.33) depends only on $\|\mathbf{x}_{\max}\|$. Also, since the spherical decision region is symmetric and convex, any such signal density is a worst-case density. Additionally, if $p_{\mathbf{s}}(\mathbf{s}) = p_{U_\epsilon}$ is uniformly distributed on the surface of a sphere of radius r, $P_D(\Phi, p_{U_\epsilon}) = P_D(\Phi, p_B)$ for any binary constellation at radius r. Finally, in Gaussian noise, the spherical decision region is the Neyman-Pearson decision region for the signal density p_{U_ϵ}.

We are finally ready to prove the Theorem 9. The proof of the theorem combines the observations about the spherical decision regions, uniform signal densities on the surface of a sphere, and the worst-case binary densities for arbitrary decision regions.

Proof. First, fix a decision region R that yields a false alarm rate $P_F(R)$. The worst-case binary signal density for this decision region is denoted $p_{(R,\min)}$. Consider a spherical decision region achieving the same false alarm rate. The following chain of inequalities result:

$$\begin{aligned} P_D(\Phi, p_{(R,\min)}) &= P_D(\Phi, p_{U_\epsilon}) \\ &\geq P_D(R, p_{U_\epsilon}) \\ &\geq P_D(R, p_{(R,\min)}). \end{aligned} \tag{3.40}$$

In equation (3.40), the equality follows from the symmetry properties of the spherical decision region. The first inequality follows because the spherical decision region is the Neyman-Pearson decision region for the signal density p_{U_ϵ}, and the second inequality follows because $p_{(R,\min)}$ is the worst-case signal density for the decision region R. The end result is that the worst-case detection probability for a spherical decision region exceeds the worst-case detection probability for any symmetric, convex decision region with an equivalent false alarm rate.

Thus far, the analysis has concentrated on the case when R is a bounded set that constrains $x_1, \ldots, x_N$. The definition of valid decision regions, however, allows for a situation where some dimensions are not constrained by R. Such regions have infinite volume. It is possible to apply the results to infinite, valid regions. In this case, the

values of the unconstrained variables $x_{j+1}, \ldots, x_N$ can be fixed arbitrarily. Viewed as a function of $x_1, \ldots, x_j$, the density $p_{\mathbf{n}}$ is a spherically symmetric function and retains the smoothness properites necessary to define the gradient in this subset of $\Re^N$. Thus, for $\mathbf{a} \in \Re^N$, components of the vector along $x_{j+1}, \ldots, x_N$ do not contribute to the derivative $\frac{df_R}{d\mathbf{a}}$, and the key results leading to the theorem still hold. In the case of an infinite valid region, note that it is possible to have signal densities p'_B that with $\|\mathbf{x}_{\max}\| > \epsilon$ but still satisfy $P_D(R, p_{(R,\min)}) = P_D(R, p'_B)$ by translating the points of the binary constellation along a direction in the unconstrained coordinates while preserving the symmetry of the density. In this situation, however, $P_D(\Phi, p'_B) > P_D(\Phi, p_{(R,\min)})$ since this perturbation will decrease f_{Φ_t}. Thus, for the target signal density p'_B, the inequality (3.40) still holds. $\square$

3.4 Summary

This chapter considered semi-parametric and robust detection with random data selection. A threshold test on the selected data vector magnitude maximizes the worst-case detector performance over the set of detectors with symmetric, convex decision regions. These properties generalize the minmax optimality of the square-law detector in Bayesian detection [28].

The chapter also presents an approximate algorithm for adapting the detector threshold to fluctuations in the selected subset size, an important challenge in implememtation of randomized selection algorithms. The approximation, which produces P_F values within a factor of 1.5 of the desired $\tilde{P}_F$ for a wide range of K and $\tilde{P}_F$ parameters, reduces the complexity of threshold calculation by a factor of nearly 13.

Chapter 4

Conclusions

This thesis has examined data selection algorithms for binary hypothesis testing. The data selection problem was examined at the two extremes of *a priori* knowledge about the target signal. For a fully known target signal, selecting the optimal data subset for the RMF reduces to a combinatorial optimization problem. At the other extreme, when the target signal resides in the broad class of stochastic signals with even symmetric densities, randomized sampling and robust detectors were considered. This chapter summarizes the contributions in both areas of research, and suggests further work.

4.1 The Restricted Matched Filter

The results in Chapter 2 define the RMF problem, and analyze a variety of solutions. In situations where the covariance matrix structure is constrained, efficient solutions of the problem exist. The low-condition and banded covariance matrices yield polynomial time solutions to the optimal RMF search. In the general problem, other approaches are necessary.

For unstructured covariance matrices, heuristic optimization rules incorporating greedy selection can be used to select a subset of measurements for the RMF. We showed that the worst-case behavior of these approaches depends on the condition number of the noise covariance matrix. The likelihood of encountering these patho-

logical cases in practice is not clear.

The combinatorial optimization problem formulated in Chapter 2 leads to a number of unresolved questions. First, the worst-case complexity of the problem is not known. It resembles a number of well-known NP-complete problems, such as the Densest Subgraph problem. A reduction between the RMF search and a well-known combinatorial optimization problem seems possible but has not yet been found. Second, a variety of heuristic search algorithms more complicated than the greedy search approach could yield better approximation performance for SNR($\mathbf{G}$). Greedy search heuristics that change the selected subset size by more than one measurement per stage, searches for local maxima in SNR($\mathbf{G}$), and simulated annealing are natural approaches to the problem. Third, there may be a set of reasonable constraints on the covariance matrix, in addition to banded structure, that guarantee efficient solutions to the optimization problem. Furthermore, approximation algorithms possessing good worst-case performance may be found [16, 41]. Such algorithms could be used confidently in with unstructured covariance matrices. Finally, it may be fruitful to consider the approximation behavior of these search algorithms using the subset size necessary to exceed a particular value of SNR($\mathbf{G}$) as the objective function.

4.2 Randomized Selection and Robust Detection

Chapter 3 considers detection of a stochastic signal with an uncertain probability density. In this extreme case, little is known about the exact structure of the density, so randomized data selection is used to control the expected subset size. Robust detectors for the signal model suggested are determined, and a threshold calculation approximation that copes with fluctuating subset size is derived.

The results of this chapter may be extendable to cases where the signal probability density has less uncertainty. In the case where the probabilty density is known exactly, a non-random sensor censoring approach has been suggested [35]. This algorithm uses data-dependent selection. In situations with some uncertainty in the signal density, selection algorithms between deterministic data-dependent and randomized selection

may work well.

The data selection procedure can be extended in a number of potentially interesting directions. First, there are many detection problems beyond binary hypothesis testing where data selection may provide implementation advantages. For example, we have not considered multiple hypothesis testing or many other significant detection problems.

Additionally, the selection algorithms discussed in this thesis can be combined with rate-distortion approaches to controlling communications in distributed signal processing systems. There are preliminary results indicating that selection can be combined with quantization in detection problems with correlated noise measurements [4].

Appendix A

Proof of Theorem 7

Theorem 7 *Consider a detection problem of the class defined in equation (3.1). Let* $r = ||\mathbf{x}||$ *and* $\mathbf{\Theta} = [\theta \ \phi_3 \ \ldots \ \phi_K]^T$. *Then, the likelihood ratio* $L(r, \mathbf{\Theta}|\mathbf{G})$ *given by (3.15) increases monotonically without bound.*

Proof. Using the spherical coordinate system defined in (3.16), a vector $\mathbf{x} \in \Re^N$ can be written

$$\mathbf{x} = r\mathbf{v}_{\mathbf{\Theta}}, \tag{A.1}$$

where $\mathbf{v}_{\mathbf{\Theta}}$ is a unit vector defined by the angular variables. In terms of these coordinates, the conditional likelihood ratio is

$$L(r, \mathbf{\Theta}|\mathbf{G}) = \int_{\mathbf{a}_K} p_{\mathbf{s}_g|\mathbf{G}}(\mathbf{a}_K|\mathbf{G}) e^{\left(-\frac{1}{2\sigma^2}\mathbf{a}_K^T\mathbf{a}_K\right)} e^{\left(r\frac{\mathbf{v}_{\mathbf{\Theta}}^T\mathbf{a}_K}{\sigma^2}\right)} d\mathbf{a}_K. \tag{A.2}$$

Expressed in spherical coordinates, the derivative is

$$\frac{dL}{dr} = \int_{\mathbf{v}_{\mathbf{\Theta}}^T\mathbf{a}_K \geq 0} p_{\mathbf{s}_g|\mathbf{G}}(\mathbf{a}_K|\mathbf{G}) e^{\left(-\frac{1}{2\sigma^2}\mathbf{a}_K^T\mathbf{a}_K\right)} \left(\frac{\mathbf{v}_{\mathbf{\Theta}}^T\mathbf{a}_K}{\sigma^2}\right) \left(2\sinh\left(r\frac{\mathbf{v}_{\mathbf{\Theta}}^T\mathbf{a}_K}{\sigma^2}\right)\right) d\mathbf{a}_K. \tag{A.3}$$

Each term in the integrand is positive over the region where $\mathbf{v}_{\mathbf{\Theta}}^T\mathbf{a}_K \geq 0$, thus the integrand is positive for every value of $r \geq 0$, and the likelihood ratio is always increasing. Similar analysis shows that the second derivative of the likelihood ratio is always positive.

Since $\frac{dL}{dr} > 0$ and $\frac{d^2L}{dr^2} > 0$ for all positive values of r, straightforward calculus

shows that $L(r, \mathbf{v}_{\boldsymbol{\Theta}}|\mathbf{G})$ satisfies the inequality

$$L(r, \boldsymbol{\Theta}|\mathbf{G}) > L(r_o, \boldsymbol{\Theta}|\mathbf{G}) + L'(r_o, \boldsymbol{\Theta}|\mathbf{G})(r - r_o) \tag{A.4}$$

for every $\boldsymbol{\Theta}$ and every pair $r \geq r_o \geq 0$. Thus, there is no upper bound on $L(r, \boldsymbol{\Theta}|\mathbf{G})$; the likelihood ratio diverges along the direction identified by the unit vector $\mathbf{v}_{\boldsymbol{\Theta}}$. □

Appendix B

Proof of Theorem 8

Theorem 8 *For the statistical model established in equation (3.8), and a fixed value of K, let $r = ||\mathbf{x}_g||$. If the conditional densities $p_{\mathbf{x}_g|\mathbf{G},H}(\mathbf{x}_g|\mathbf{G}, H_i)$ for $i = 0, 1$ are continuous and positive, the likelihood ratio for $r = ||\mathbf{x}_g||$ increases monotonically without bound.*

Proof.

For the random variable r, the likelihood ratio is given by

$$L(r|\mathbf{G}) = \frac{p_{r|\mathbf{G},H}(r|\mathbf{G}, H_1)}{p_{r|\mathbf{G},H}(r|\mathbf{G}, H_0)}. \tag{B.1}$$

We shall derive a monotonically increasing lower bound for $p_{r|\mathbf{G},H}(r|\mathbf{G}, H_1)$ in terms of the likelihood ratio for $\mathbf{x}_g$.

In order to make the argument precise, let

$$f(r) = \min_{||\mathbf{x}_g||=r} L(\mathbf{x}_g|\mathbf{G}). \tag{B.2}$$

Under the assumption that this is a well-defined function, the conditional densities for $\mathbf{x}$ satisfy

$$p_{\mathbf{x}_g|\mathbf{G},H}(\mathbf{x}_g|\mathbf{G}, H_1) \geq p_{\mathbf{x}_g|\mathbf{G},H}(\mathbf{x}_g|\mathbf{G}, H_0) f(r) \tag{B.3}$$

whenever $||\mathbf{x}_g|| = r$. Thus, the density $p_{r|\mathbf{G},H}(r|\mathbf{G}, H_1)$ satisfies

$$p_{r|\mathbf{G},H}(r|\mathbf{G}, H_1) \geq B_K r^{K-1} \exp(-r^2/2\sigma^2) f(r). \tag{B.4}$$

This inequality shows that the likelihood ratio for the random variable r is bounded below by $f(r)$, i.e. $L(r|\mathbf{G}) \geq B_K f(r)$. In the remainder of this section, we will prove the existence and continuity of $f(r)$ and demonstrate that $f(r)$ increases monotonically without bound.

Consider the minimization from equation (B.2) in spherical coordinates. The domain for the angular variables is a compact set; each ϕ_i lies in the interval $0 \leq \phi_i < \pi$, and θ lines in the interval $0 \leq \theta < 2\pi$. With fixed r, $L(r, \mathbf{\Theta}|\mathbf{G})$ is a positive, bounded function over a compact set, so the minimum exists [2].

The results of Theorem 7 show that $f(r)$ is an increasing function. For any $r \geq 0$ and $\epsilon \geq 0$, consider the value of

$$f(r+\epsilon) = \min_{||\mathbf{x}_g||=r+\epsilon} L(\mathbf{x}_g|\mathbf{G}). \tag{B.5}$$

Theorem 7 guarantees that $L((r+\epsilon), \mathbf{\Theta}) > L(r, \mathbf{\Theta})$ for any fixed $\mathbf{\Theta}$. Additionally, the likelihood ratio satisfies $L(r, \mathbf{\Theta}) \geq f(r)$ as a consequence of the definition of f. Combining these inequalities yields

$$L\left((r+\epsilon), \mathbf{\Theta}|\, \mathbf{G}\right) > f(r) \text{ for all } \mathbf{\Theta}, \text{ if } \epsilon > 0. \tag{B.6}$$

Since there is a strict inequality between $L((r+\epsilon), \mathbf{\Theta})$ and $f(r)$, $f(r+\epsilon)$ must satisfy

$$f(r+\epsilon) > f(r) \text{ for all } \epsilon > 0. \tag{B.7}$$

Thus, as a consequence of theorem 7, the bound function $f(r)$ is strictly monotonically increasing.

To prove that $f(r)$ is continuous, we must verify that for every $\epsilon > 0$, we can find

$\delta > 0$ such that

$$|f(r+\delta) - f(r)| < \epsilon. \tag{B.8}$$

The likelihood bound satisfies $f(r) = L(r, \boldsymbol{\Theta}|\mathbf{G})$ for some unit vector $\boldsymbol{\Theta}$ due to its definition as a minimization. Thus, using the fact that $L(r, \boldsymbol{\Theta}|\mathbf{G})$ is an increasing function of r, we can determine the inequality

$$f(r+\delta) = \min_{||\mathbf{x}_g||=r+\delta} L(\mathbf{x}_g|\mathbf{G}) < L((r+\delta), \boldsymbol{\Theta}_o|\mathbf{G}) \tag{B.9}$$

for any fixed $r > 0$ and $\delta > 0$. Using this inequality produces the following property of $f(r)$

$$|f(r+\delta) - f(r)| < |L((r+\delta), \boldsymbol{\Theta}_o|\mathbf{G}) - L(r, \boldsymbol{\Theta}_o|\mathbf{G})|. \tag{B.10}$$

Since $L(r, \boldsymbol{\Theta}_o|\mathbf{G})$ is the ratio of two positive, continuous functions of r, it is also a continuous function. Thus, for any value of $\epsilon > 0$, we can find a value of δ such that $|L((r+\delta), \boldsymbol{\Theta}_o|\mathbf{G}) - L(r, \boldsymbol{\Theta}_o|\mathbf{G})| < \epsilon$. This property, along with the inequality (B.10) establish that $f(r)$ is a continuous function of r.

Finally, we can establish that $f(r)$ has a lower bound similar to equation (A.4). In the proof of theorem 7, we showed that, with fixed $\boldsymbol{\Theta}$, $\frac{dL}{dr} > 0$ for any $r > 0$ and any $\boldsymbol{\Theta}$. Thus, we can guarantee that

$$g(r) = \min_{||x_g||=r} \frac{dL}{dr} \tag{B.11}$$

is a well-defined function because we are minimizing a function bounded from below over a compact set. Thus, for every $\boldsymbol{\Theta}$ and every pair $r \geq r_o > 0$, the inequality

$$L(r, \boldsymbol{\Theta}|\mathbf{G}) \geq L(r_o, \boldsymbol{\Theta}|\mathbf{G}) + g(r_o)(r - r_o) \tag{B.12}$$

holds. Minimizing both sides of the expression guarantees that $f(r)$ satisfies

$$f(r) \geq f(r_o) + g(r_o)(r - r_o). \tag{B.13}$$

The function $f(r)$ increases at least linearly, so it is not bounded. Since $L(r|\mathbf{G}) \geq B_K f(r)$, the theorem is proven. $\square$

Appendix C

Proof of Lemma 5

Lemma 5 *Consider the function $f_R(\mathbf{a})$ defined in equation (3.30). If the domain of integration in the expression satisfies $R \in \mathcal{R}_N$ and the probability density satisfies*

$$p_{\mathbf{n}}(\mathbf{u}_1) \geq p_{\mathbf{n}}(\mathbf{u}_2)$$

whenever $\|\mathbf{u}_1\| \leq \|\mathbf{u}_2\|$, $\mathbf{u}_1 = \mu\mathbf{u}_2$ for some constant $\mu \geq 0$, then

$$f_R(\mathbf{a}_1) \geq f_R(\mathbf{a}_2)$$

$\|\mathbf{a}_1\| \leq \|\mathbf{a}_2\|$, $\mathbf{a}_1 = \nu\mathbf{a}_2$ for some constant $\nu \geq 0$.

C.1 Proof for Smooth Decision Regions

In order to prove the lemma for regions of arbitrary shape, we examine the directional derivative of the function f_R. Given a unit vector $\mathbf{v}$ indicating the direction of displacement, the directional derivative of f_R along $\mathbf{v}$ is

$$\frac{df_R}{d\mathbf{v}} = \langle \nabla f_R, \mathbf{v} \rangle, \tag{C.1}$$

where ∇ denotes the gradient operator and $\langle , \rangle$ indicates an inner product of two vectors. In terms of the integral expression for $f_R(\mathbf{a})$, the derivative is

$$\frac{df_R}{d\mathbf{v}} = \int_R \langle \nabla p_{\mathbf{n}}(\mathbf{u}-\mathbf{a}), \mathbf{v} \rangle d\mathbf{u}. \tag{C.2}$$

Using the divergence theorem from vector calculus, the volume integral in equation (C.2) can be expressed as an equivalent surface integral. The resulting expression is

$$\frac{df_R}{d\mathbf{v}} = -\int_{S(R)} p_{\mathbf{n}}(\mathbf{u}-\mathbf{a}) \langle \mathbf{h}(\mathbf{u}), \mathbf{v} \rangle dS. \tag{C.3}$$

In this equation, the notation $S(R)$ denotes the surface of region R, $\mathbf{h}(\mathbf{u})$ is the outward-facing normal vector for the surface at the point $\mathbf{u}$, and dS is the surface area increment. The use of the divergence theorem requires auxiliary smoothness conditions on $S(R)$ and $p_{\mathbf{n}}$ to guarantee that the gradient vector and the surface normal vectors are well defined.

The properties we have assumed for R and $p_{\mathbf{n}}(\mathbf{u})$ enable us to show that the directional derivative of $f_R(\mathbf{a})$ is negative or zero for the direction specified by $\mathbf{a}$. In other words, along any straight line emanating from the origin of sample space, $f_R(\mu\mathbf{a})$ is a non-increasing function of μ over the domain $0 \leq \mu < \infty$.

To make the argument precise, we define the following notation. Let $S^+_{\mathbf{a}}(R)$ be the subset of $S(R)$ satisfying $\langle h(\mathbf{x}), \mathbf{a} \rangle > 0$ for all $\mathbf{x} \in S^+_{\mathbf{a}}(R)$. Similarly, we can define $S^-_{\mathbf{a}}(R)$ as the subset of $S(R)$ satisfying $\langle h(\mathbf{x}), \mathbf{a} \rangle < 0$. Due to the convexity of R, any point $\mathbf{x} \in S^+_{\mathbf{a}}(R)$ has $\langle \mathbf{h}(\mathbf{x}), \mathbf{a} \rangle > 0$. For points in $S^-_{\mathbf{a}}(R)$, the sign on the inner product is reversed.

Using this notation, we can write equation (C.3) as

$$\frac{df_R}{d\mathbf{a}} = -\int_{S^+_{\mathbf{a}}(R)} p_{\mathbf{n}}(\mathbf{u}-\mathbf{a}) \langle \mathbf{h}(\mathbf{u}), \mathbf{a} \rangle d\mathbf{u} - \int_{S^-_{\mathbf{a}}(R)} p_{\mathbf{n}}(\mathbf{u}-\mathbf{a}) \langle \mathbf{h}(\mathbf{u}), \mathbf{a} \rangle d\mathbf{u}. \tag{C.4}$$

For every point $\mathbf{u} \in S^+_{\mathbf{a}}(R)$, $-\mathbf{u}$ belongs to $S^-_{\mathbf{a}}(R)$ and satisfies $\langle \mathbf{h}(\mathbf{u}), \mathbf{a} \rangle = -\langle \mathbf{h}(-\mathbf{u}), \mathbf{a} \rangle$. Additionally, for $\mathbf{u} \in S^+_{\mathbf{a}}(R)$, $\|\mathbf{u}-\mathbf{a}\| < \|-\mathbf{u}-\mathbf{a}\|$. This implies that $p_{\mathbf{n}}(\mathbf{u}-\mathbf{a}) > p_{\mathbf{n}}(-\mathbf{u}-\mathbf{a})$ since $p_{\mathbf{n}}(\mathbf{x})$ is a decreasing function of $\|\mathbf{x}\|$. Thus, the directional derivative

can be written

$$\frac{df_R}{d\mathbf{a}} = -\int_{S_{\mathbf{a}}^+(R)} \left(p_{\mathbf{n}}(\mathbf{u}-\mathbf{a}) - p_{\mathbf{n}}(-\mathbf{u}-\mathbf{a})\right) \langle \mathbf{h}(\mathbf{u}), \mathbf{a}\rangle d\mathbf{u}. \tag{C.5}$$

The term $(p_{\mathbf{n}}(\mathbf{u}-\mathbf{a}) - p_{\mathbf{n}}(-\mathbf{u}-\mathbf{a}))$ is positive for any $\mathbf{a}$, so the directional derivative $\frac{df_R(\mathbf{a})}{d\mathbf{a}}$ is either negative or zero. This result applies for arbitrarily shaped origin-symmetric, convex regions and for spherically-symmetric, unimodal noise probability densities.

C.2 Alternative Proof for Gaussian Noise Densities

The use of the divergence theorem in the previous section's proof of the theorem imposes smoothness constraints on the decision region R. Its surface must be sufficiently smooth so that it posses a well-defined normal vector. In this section, we prove that $f_R(\alpha_1\mathbf{u}) > f_R(\alpha_2\mathbf{u})$ whenever $\alpha_2 > \alpha_1 > 0$ without making smoothness assumptions about $S(R)$. The proof, however, relies on properties of white Gaussian probability densities.

Rather than analyze the derivative directly, as in our previous proof, we will examine the function

$$\Delta(\mathbf{x}, \alpha_1, \alpha_2) = p_{\mathbf{n}}(\mathbf{x}-\alpha_1\mathbf{u}) - p_{\mathbf{n}}(\mathbf{x}-\alpha_2\mathbf{u}). \tag{C.6}$$

The definition allows us to express changes in the function $f_R(\alpha\mathbf{u})$ as

$$f_R(\alpha_1\mathbf{u}) - f_R(\alpha_2\mathbf{u}) = \int_R \Delta(\mathbf{x}, \alpha_1, \alpha_2) d\mathbf{x}. \tag{C.7}$$

The results in this section show that for every point $\mathbf{x}_1 \in R$ such that $\Delta(\mathbf{x}_1) < 0$, we can define a 1-1 function $A(\mathbf{x}_1)$ with the properties that $\Delta(A(\mathbf{x}_1), \alpha_1, \alpha_2) + \Delta(\mathbf{x}_1, \alpha_1, \alpha_2) \geq 0$ and $A(\mathbf{x}_1) \in R$. These properties of $\Delta(\mathbf{x}, \alpha_1, \alpha_2)$ guarantee that

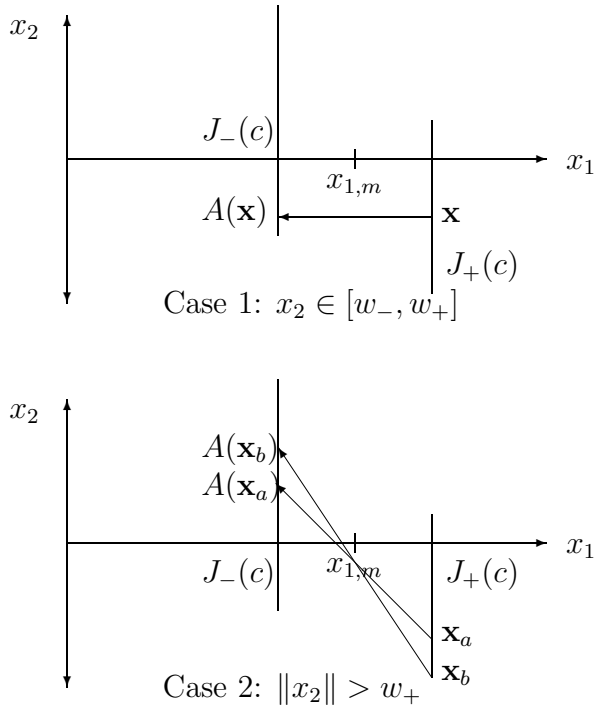

Figure C-1: Illustration of $A(\mathbf{x})$. In Case 1, the transformation produces $A(\mathbf{x})$ by translation parallel to the x_1 axis. In situations where such a translation does not intersect $J_-(c)$, $A(\mathbf{x})$ is symmetric with $\mathbf{x}$ about the x_2 value in the midpoint of the region common to $J_-(c)$ and $J_+(c)$.

$f_R(\alpha\mathbf{u})$ is a decreasing function of α. We shall first prove the result concerning Δ in $\Re^2$ and then generalize to $\Re^N$.

For the case of $N = 2$, we can set the coordinate system so that $\mathbf{u}$ is a unit vector pointing in the x_1 direction. In this case, because we have assumed that the probability density is $p(\mathbf{x}) = \mathcal{N}(\mathbf{x}; 0; \sigma^2\mathbf{I})$, the function $\Delta(\mathbf{x}, \alpha_1, \alpha_2)$ is zero the line $x_{1,m} = (\alpha_1+\alpha_2)/2$, negative to its right, positive to its left. We will define our mapping $A(\mathbf{x})$ by creating a correspondence between points on the line $x_1 = x_{1,m} - c$ and the line $x_1 = x_{1,m} + c$. The figure shows examples of the appropriate transformation.

An algebraic expression of the transformation requires a notation to distinguish the two cases. The transformation maps points from one slice of R to another by translation parallel to the x_1 axis if possible. If not, the points are reflected about a specific point between the slices. In order to describe $A()$, we define notion of slice,

reflection, and translation. By slices, we mean the sets

$$\begin{aligned} J_-(c) &= \{\mathbf{x}|\mathbf{x} \in R, x_1 = x_{1,m} - c\} \\ J_+(c) &= \{\mathbf{x}|\mathbf{x} \in R, x_1 = x_{1,m} + c\}. \end{aligned} \tag{C.8}$$

Note that due to the convexity and boundedness of R, the slices will be finite-length line segments parallel to the x_2 coordinate axis.

Whether any points are reflected, as in case 2 from Figure C.2, depends on the relative arrangements of the endpoints of these line segments. Let $x_{2,+}^{\max}$ and $x_{2,+}^{\min}$ be the maximum and minimum x_2 coordinate values for points in $J_+(c)$. Likewise, let $x_{2,-}^{\max}$ and $x_{2,-}^{\min}$ be the corresponding values for $J_-(c)$. Case 2 of $A()$ occurs when the interval $[x_{2,+}^{\min}, x_{2,+}^{\max}]$ is not contained in the interval $[x_{2,-}^{\min}, x_{2,-}^{\max}]$.

In order to define the case 2 in the figure properly, we need the x_2 coordinate of the reflection center. Let $\beta < \gamma < \phi < \chi$ be a sorted list of $x_{2,+}^{\max}$, $x_{2,+}^{\min}$ $x_{2,-}^{\max}$, and $x_{2,-}^{\min}$. Then, let $x_{2,m} = \frac{1}{2}(\gamma + \phi)$. This is the x_2 coordinate of the reflection center. In vector notation, the reflection center is

$$\mathbf{x}_m = \begin{bmatrix} x_{1,m} \\ \\ x_{2,m} \end{bmatrix}. \tag{C.9}$$

Finally, a number of later results refer to the boundary values of the intersection between $[x_{2,+}^{\min}, x_{2,+}^{\max}]$ and $[x_{2,-}^{\min}, x_{2,-}^{\max}]$ when case 2 occurs. We will identify the values γ and ϕ with either $J_+(c)$ or $J_-(c)$ as follows:

$$w_+ = \begin{cases} \gamma & x_{2,+}^{\min} = \beta \\ \phi & x_{2,-}^{\min} = \beta \end{cases} \quad w_- = \begin{cases} \phi & x_{2,+}^{\min} = \beta \\ \gamma & x_{2,-}^{\min} = \beta. \end{cases} \tag{C.10}$$

Using the definitions in (C.8) and (C.10), we can specify the mapping $A(\mathbf{x})$ in $\Re^2$.

The mapping is

$$A(\mathbf{x}) = \begin{cases} \begin{bmatrix} x_1 - 2c \\ x_2 \end{bmatrix} & x_2 \in [\gamma, \phi] \\ \\ \mathbf{x} + 2(\mathbf{x}_m - \mathbf{x}) = 2\mathbf{x}_m - \mathbf{x} & x_2 \notin [\gamma, \phi]. \end{cases} \tag{C.11}$$

The mapping translates parallel to the x_1 axis to get a direct map when possible. Otherwise, it matches the points in $J_+(c)$ with $|x_2|$ values that exceed $|w_+|$ with points in $J_-(c)$ with $|x_2|$ that exceeds w_-.

In order to prove the overall result of this note, we need to show that the mapping is well-defined, and that $\mathcal{N}(A(\mathbf{x}); \alpha_1\mathbf{u}; \sigma^2\mathbf{I}) > \mathcal{N}(\mathbf{x}; \alpha_2\mathbf{u}; \sigma^2\mathbf{I})$. We show that $A(\mathbf{x})$ is well defined before proving the inequality.

In order for the mapping defined in equation (C.11) to be well-defined, the length of the line segment $J_-(c)$ must exceed the length of $J_+(c)$. If not, a 1-1 mapping between the two sets is impossible. The following lemma guarantees this result.

Lemma 7 *Let ℓ_+ be the length of the line segment $J_+(c)$ and ℓ_- be the length of the line segment $J_-(c)$. If the set $R \in \Re^2$ is symmetric and convex, then $\ell_- \geq \ell_+$.*

Proof. As a consequence symmetry and convexity, R must contain the parallelograms formed by both $J_+(c)$, $J_-(c)$, and their reflections about the origin. If $\ell_+ > \ell_-$, then there must be a a point on the line $x_1 = x_{1,m} - c$ that is in the parallelogram formed by $J_+(c)$ and its reflection, but is not in the parallelogram formed by $J_-(c)$ and its reflection. Figure 2 shows this situation. The parallelograms are always nested as shown since $|x_{1,m} + c| > |x_{1,m} - c|$. This contradicts the convexity of R. □

In order to show prove that

$$p_{\mathbf{n}}(A(\mathbf{x}) - \alpha_1\mathbf{u}) \geq p_{\mathbf{n}}(\mathbf{x} - \alpha_2\mathbf{u}) \tag{C.12}$$

for all $\mathbf{x} \in J_+(c)$, we need to analyze the x_2 coordinates of the points $\mathbf{x}$ and $A(\mathbf{x})$. Their relative displacement from $\alpha_1\mathbf{u}$ and $\alpha_2\mathbf{u}$ along the x_1 axis is identical due to

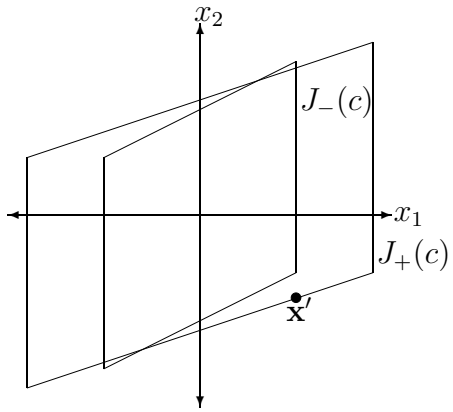

Figure C-2: Situation if $\ell_+ > \ell_-$. The pont $\mathbf{x}'$ belongs to the parallelogram formed by $J_-(c)$ but does not belong to the parallelogram formed by $J_+(c)$.

the definition of the sets $J_+(c)$ and $J_-(c)$. Since the probability density is white and Gaussian, it depends only on the norm of the argument. Thus, in a situation where the x_1 cooridnate is fixed, the density is greatest when $\|x_2\|$ is smallest. As a consequence of the symmetry and convexity of R, we can prove the following lemma, which is sufficient to guarantee the inequality in equation (C.12) is true.

Lemma 8 *For the mapping defined in equation (C.11), the x_2 coordinate value satisfies $|x_2| \geq |(A(\mathbf{x}))_2|$.*

Proof. We proceed with a proof by contradiction.

Without loss of generality, we can assume that $w_- > w_+$. This condition can be guaranteed by a proper choice of coordinate system. Every point $\mathbf{x} \in J_-(c)$ that is reflected by $A()$ satisfies $x_2 < w_+$. Additionally, every $\mathbf{y} \in J_-(c)$ satisfies $y_2 > w_+$, regardless of the behavior of $A()$.

Assume that $|x_2| < |A(\mathbf{x})_2|$ for some $\mathbf{x} \in J_+(c)$. This assumption implies that

$$|w_+| < |w_-| \tag{C.13}$$

since $A()$ maps points satisfying $x_2 < w_+$ to points satisfying $y_2 > w_-$.

Let $\xi = \arg\min_{J_+(c)} |x_2|$, $\psi = \arg\max_{J_+(c)} |x_2|$, and $\psi' = -\mathbf{I}\psi$. We will show that

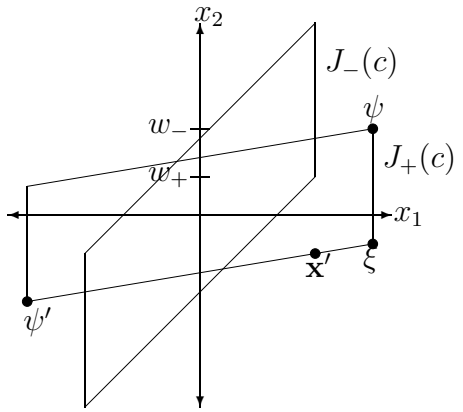

Figure C-3: Case 1 of Lemma 2 proof. In this situation $A(\mathbf{x})$ can increase the x_2 component of the point. As a consequence, the point $\mathbf{x}'$ on the parallelogram formed by $J_+(c)$ is not in the parallelogram defined by $J_-(c)$. This contradicts the convexity of R.

the line between ξ and ψ' does not intersect $J_-(c)$, contradicting the convexity of R.

Two cases will be considered separately. If there are points $\mathbf{x} \in J_+(c)$ that satisfy $x_2 \in [w_+, w_-]$, then $\xi_2 = \min_{J_-(c)} |x_2| < w_+$ and $\max_{J_+(c)} |x_2| = w_-$. In this case, $\psi_2' = -w_- < w_+$ as a consequence of equation (C.13), which followed from our initial assumption that $|x_2| < |A(\mathbf{x})_2|$. Thus, the $x_2 < w_+$ for any point on a line segment joining ξ and ψ'. This implies that the line segment does not intersect $J_-(c)$ since $x_2 > w_+$ for all $\mathbf{x} \in J_-(c)$. Thus, we arrive at a contradiction of the convexity of R.

In the case that there are no points $\mathbf{x} \in J_+(c)$ that satisfy $x_2 \in [w_+, w_-]$, we have $\max_{J_+(c)} |x_2| = w_+$ and $\min_{J_-(c)} |x_2| = w_-$. Thus, $\psi_2' = -w_+$. From the definition of our coordinate system, $w_- > w_+$. Only if $w_- > 0$ can this inequality hold as well as (C.13). If this inequality, a consequence of the assumption $|x_2| < |A(\mathbf{x})_2|$, is true, then we have $w_- > |w_+|$, implying that the line segment joining ξ and ψ' does not intersect $J_-(c)$. Once again, this contradicts the convexity of R.

Thus, if the assumption $|x_2| < |A(\mathbf{x})_2|$ holds, R cannot be a symmetric, convex set. $\square$

Using the inequality (C.12), we can show that $\Delta(A(\mathbf{x}), \alpha_1, \alpha_2) + \Delta(\mathbf{x}, \alpha_1, \alpha_2) \geq 0$,

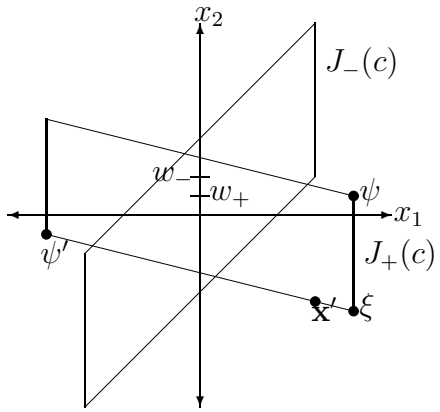

Figure C-4: Case 2 of the proof of Lemma 2. The point $\mathbf{x}'$ contradicts the convexity of R.

which proves that $f_R(\alpha\mathbf{u})$ is a non-increasing function of α. Since we have assumed that the noise density is white and Gaussian, the ratio

$$L(\mathbf{x}) = \frac{p_{\mathbf{n}}(\mathbf{x} - \alpha_1\mathbf{u})}{p_{\mathbf{n}}(\mathbf{x} - \alpha_2\mathbf{u})} \tag{C.14}$$

depends only on the x_1 coordinate of $\mathbf{x}$. Consequently, for a fixed value of c, $L(\mathbf{x}) = 1/L(\mathbf{y})$ for $\mathbf{x} \in J_+(c)$ and $\mathbf{y} \in J_-(c)$. The mapping from $J_+(c)$ to $J_-(c)$ exchanges the relative distances from $\alpha_1\mathbf{u}$ and $\alpha_2\mathbf{u}$ along the x_1 axis. This result allows us to re-write Δ as

$$\Delta(\mathbf{x}, \alpha_1, \alpha_2) = \begin{cases} (1 - 1/L(\mathbf{x}))p_{\mathbf{n}}(\mathbf{x} - \alpha_1\mathbf{u}) & \mathbf{x} \in J_-(c) \\ (L(\mathbf{x}) - 1)p_{\mathbf{n}}(\mathbf{x} - \alpha_2\mathbf{u}) & \mathbf{x} \in J_+(c). \end{cases} \tag{C.15}$$

Based on the equality (C.15), we can write

$$\Delta(A(\mathbf{x}), \alpha_1, \alpha_2) - \Delta(\mathbf{x}, \alpha_1, \alpha_2) = (1 - L(\mathbf{x}))(p_{\mathbf{n}}(A(\mathbf{x}) - \alpha_1\mathbf{u}) - p_{\mathbf{n}}(\mathbf{x} - \alpha_2\mathbf{u})). \tag{C.16}$$

For any $\mathbf{x} \in J_+(c)$, the likelihood ratio satisfies $L(\mathbf{x}) < 1$. Thus, combined with the inequality (C.12) with equation (C.16), we get the result $\Delta(A(\mathbf{x}), \alpha_1, \alpha_2) - \Delta(\mathbf{x}, \alpha_1, \alpha_2) > 0$. This proves that $f_R(\alpha\mathbf{u})$ is a non-increasing function of α for

the case $R \in \Re^2$.

Since every point in the the region R belongs to either $J_+(c)$ or $J_-(c)$ for only one value of c, the lemmas show that there is a the main theorem holds in 2-D. We can extend to $\Re^N$ by noting that any point $\mathbf{x} \in \Re^N$ can be represented uniquely as a sum of vectors parallel to $\mathbf{u}$ and perpindicular to $\mathbf{u}$. Formally, the vector space $\Re^N = \text{span}\{\mathbf{u}\} \oplus \text{span}\{\mathbf{u}\}^\perp$ is a direct sum of the span of $\mathbf{u}$ and its orthogonal complement. For a particular vector in $\text{span}\{\mathbf{u}\}^\perp$, we can reduce to a two-dimensional problem and apply the appropriate $A(\mathbf{x})$ to all points that apply. Since the direct sum decompositon is unique, the transformation is one-one over R in $\Re^N$, and the inequality

$$p_{\mathbf{n}}(A(\mathbf{x}) - \alpha_1\mathbf{u}) > p_{\mathbf{n}}(\mathbf{x} - \alpha_2\mathbf{u}) \tag{C.17}$$

for every $\mathbf{x} \in R$ such that $\langle \mathbf{x}, \mathbf{u} \rangle > \frac{1}{2}(\alpha_1 + \alpha_2)$. Note that although the points along the $\mathbf{u}$ axis are included in the two dimensional problem for every vector in $\text{span}\{\mathbf{u}\}^\perp$, for these points, the mapping $A()$ is invariant to the choice of the vector in the orthogonal complement.

Bibliography

[1] R. Ahlswede and I. Csiszar. Hypothesis testing with communication constraints. *IEEE Transactions on Information Theory*, 32(4):533–542, July 1986.

[2] Dimitri Bertsekas. *Nonlinear Programming*. Athena Scientific, Belmont, MA, 1999.

[3] Stephen Blatt. Distributed sensor fusion for sensor networks. In *Proc. 4th Annual Conf. on Information Fusion*, pages TuC3–TuC8, 2001.

[4] J-F. Chamberland and V. V. Veeravalli. Decentralized detection in sensor networks. *IEEE Transactions on Signal Processing*, 51(2):407–416, Feb. 2003.

[5] A. P. Chandrakasan, R. Amirtharajah, S. Cho, J. Goodman, G. Konduri, J. Kulik, W. Rabiner, and A. Wang. Design considerations for distributed microsensor networks. In *Proc. CICC*, pages 279–286, 1999.

[6] M. Charikar, V. Guruswami, R. Kumar, S. Rajagopalan, and A. Sahai. Combinatorial feature selection algorithms. In *Proc. IEEE Symposium on Foundations of Computer Science*, pages 631–640, Nov. 2002.

[7] J. Chen, K. Yao, and R. E. Hudson. Source localization and beamforming. *IEEE Signal Processing Magazine*, 19(3):30–39, March 2002.

[8] W. J. Cody. Rational chebyshev approximations for the error function. *Mathematical Computation*, 23:631–637, 1969.

[9] Neuhoff D. L and D. Marco. Distributed encoding of sensor data. In *Proc. IEEE Information Theory Workshop*, pages 108–110, Oct. 2002.

[10] A. R. DiDonato and A. H. Morris Jr. Computation of the incomplete gamma function ratios and their inverse. *ACM Transactions on Mathematical Software*, 12(4):377–393, Dec. 1986.

[11] S. C. Draper. *Successive Structuring for Souce Coding Algorithms for Data Fusion, Storage, and Buffering in Networks*. Ph.D. dissertation, Mass. Inst. of Technology, Department of Electrical Engineering and Computer Science, June 2002.

[12] D. Estrin and R. Govindan. Next century challenges: Scalable coordination in sensor networks. In *Proc. IEEE Mobicom*, pages 263–270, 1999.

[13] J. M. Glass and P. Kotiveeriah. The allocation problem in distributed signal processing. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 29(4):817–829, April 1981.

[14] Gene Golub and Charles Van Loan. *Matrix Computations*. Johns Hopkins University Press, Baltimore, MD, third edition, 1996.

[15] D. Gore, R. Nabar, and A. Paulraj. Selecting an optimal set of transmit antennas for a low rank matrix channel. In *Proc. ICASSP*, pages 2785–2788, 2001.

[16] Dorit S. Hochbaum, editor. *Approximation Algorithms for NP-Hard Problems*. PWS Publishing, Boston, MA, 1997.

[17] Roger A. Horn and Charles R. Johnson. *Matrix Analysis*. Cambridge University Press, Cambridge, UK, 1985.

[18] V. Krishnamurthy. Algorithms for optimal scheduling and management of hiden markov model sensors. *IEEE Transactions on Signal Processing*, 50(6):1382–1397, 2002.

[19] S. Kuman, F. Zhao, and D. Shepherd. Collaborative signal and information processing in microsensor networks. *IEEE Signal Processing Magazine*, 19(3):13–14, March 2002.

[20] D. Li, K. Wong, Y. Hu, and A. Sayeed. Detection, classification, and tracking of targets. *IEEE Signal Processing Magazine*, 19(3):17–29, March 2002.

[21] S. G. Mallat and Z. Zhang. Matching pursuits with time-frequency dictionaries. *IEEE Transactions on Signal Processing*, 41(12):3397–3415, Dec. 1993.

[22] D. Mattera, F. Palmieri, and S. Haykin. Efficient sparse fir filter design. In *Proc. IEEE Conf. Acoustics, Speech, and Signal Processing*, May 2002.

[23] A. J. Miller. *Subset Selection in Regression*. Chapman and Hall, New York, NY, 1990.

[24] A. Molisch, M. Win, and J. Winters. Capacity of MIMO systems with antenna selection. In *Proc. IEEE International Conf. on Communication*, pages 570–574, 2001.

[25] Rajeev Motwani and Prabhakar Raghavan. *Randomized Algorithms*. Cambridge Univ. Press, Cambridge, U. K., 1995.

[26] A. N. Mucciardi and E. E. Gose. A comparison of seven techniques for choosing subsets of pattern recogintion properties. *IEEE Transactions on Computers*, 20(9):1023–1031, Sept. 1971.

[27] P. M. Narendra and K. Fukunaga. A branch and bound algorithm for feature subset selection. *IEEE Transactions on Computers*, 26(9):917–922, Sept. 1977.

[28] D. Pastor, R. Gay, and A. Groenenboom. A sharp upper bound for the probability of error of the likelihood ratio test for detecting signals in white gaussian noise. *IEEE Transactions on Information Theory*, 48(1):228–238, Jan. 2002.

[29] G. Polychronopolous and J. N. Tsitsiklis. Explicit solutions for some simple decentralized detection problems. *IEEE Transactions on Aerospace and Electronic Systems*, 26(2):282–292, March 1990.

[30] H. V. Poor and Saleem A. Kassam. Robust techniques for signal processing: A survey. *Proceedings of the IEEE*, 73:431–481, March 1985.

[31] H. Vincent Poor. *An Introduction to Signal Detection and Estimation*. Springer-Verlag, New York, NY, second edition, 1994.

[32] G. J. Pottie and W. J. Kaiser. Wireless integrated sensor networks. *Communications of the ACM*, 43(5):51–58, May 2000.

[33] John G. Proakis. *Digital Communications*. McGraw-Hill, Boston, MA, 1995.

[34] V. Raghunathan, C Schrugers, S. Park, and M. Srivastava. Energy-aware wireless microsensor networks. *IEEE Signal Processing Magazine*, 19(3):40–50, March 2002.

[35] C. Rago, Peter Willett, and Y. Bar-Shalom. Censoring sensors: A low-communication-scheme for distributed detection. *IEEE Transactions on Aerospace and Electronic Systems*, 32(2):554–568, April 1996.

[36] M. R. Said and A. V. Oppenheim. Discrete-time randomized sampling. In *Proc. 8th Annual IEEE Conf. on Electronics, Circuits and Systems*, pages 1407–1411, 2001.

[37] Louis L. Scharf. *Statistical Signal Processing: Detection, Estimation, and Time Series Analysis*. Addison Wesley, Reading, MA, 1991.

[38] Robert R. Tenney and Nils R. Sandell. Detection with distributed sensors. *IEEE Transactions on Aerospace and Electronic Systems*, 17(4):501–510, July 1981.

[39] J. B. Thomas. Nonparametric detection. *Proceedings of the IEEE*, 58:623–631, May 1970.

[40] S. C. A Thomopolous and N. A. Okello. Distributed detection with consulting sensors and communication cost. *IEEE Transactions on Automatic Control*, 37:1398–1405, Sept. 1992.

[41] Vijay Vazirani. *Approximation Algorithms*. Springer-Verlag, New York, NY, 2001.

[42] Jennifer L. Webb and D. C. Munson Jr. Chebyshev optimization of sparse FIR filters using linear programming with an application to beamforming. *TranSP*, 44(8):1912–1922, Aug. 1996.

[43] M. Welborn. *Flexible Signal Processing Algorithms for Wireless Communications*. Ph.D. dissertation, Mass. Inst. of Technology, Department of Electrical Engineering and Computer Science, June 2000.

[44] John Wozencraft and Irwin Jacobs. *Principles of Communication Engineering*. Wiley, New York, NY, 1965.

[45] F. Zhao, J. Shin, and J. Reich. Information-driven dynamic sensor collaboration. *IEEE Signal Processing Magazine*, 19(3):61–72, March 2002.