



NTNU – Trondheim
Norwegian University of
Science and Technology

Detection and Localization of High-Frequency Interception Signals using Compressive Sensing

Torjus Breisjøberg

Master of Science in Physics and Mathematics

Submission date: Januar 2013

Supervisor: Alex Hansen, IFY

Co-supervisor: Fredrik Hekland, KDS

Norwegian University of Science and Technology
Department of Physics

Abstract

We investigate detection and localization of high-frequency interception signals underwater using a fairly new framework known as Compressive Sensing. Within this framework we can sample at sub-Nyquist rate and still be able to reconstruct the signal. Compressive methods are proposed as an alternative to traditional methods to lever the computational burden by reducing the data rate by a significant amount. We use existing theory and develop methods to apply Compressive Sensing to the case of detection and localization of high-frequency interception signals. Provided simulations on synthetic and real data verifies the existence of this use.

Sammendrag

Vi undersøker deteksjon og lokalisering av høy-frekvente interception signaler under vann ved å bruke en relative ny metode kjent som Compressive Sensing. Compressive Sensing tillater oss å rekonstruere et signal selv om det er sterkt undersamplet. Reduksjonen i data raten gjør at Compressive Sensing kan være et reelt alternativ til vanlige metoder med tanke på prosesserings kraft. Vi bruker ekisterende teori og utvikler metoder for å lokalisere og bestemme peileretning på interception signaler ved bruk av Compressive Sensing. Simuleringer på syntetiske og ekte data blir brukt for å bekrefte bruksområdet.

Preface

During my stay as an employee at Kongsberg Defence and Aerospace, subdivision Naval Systems and Surveillance, I was given the opportunity to finalize my education in physics and mathematics at NTNU by completing a Master's Thesis in cooperation with KDA. With strong interest in signal processing and sonar systems, it was only natural to pursue these topics.

Compressive Sensing is a relatively new field and unknown for many. I have been given the opportunity and pleasure of investigating this exciting branch of signal processing over the course of twenty weeks. The concept of applying Compressive Sensing to detect and localize high-frequency interception signals for sonar applications is an intriguing idea and could not be left untouched. Since its proposal by my supervisor Fredrik Hekland, extensive work have been deployed to justify its existence with the general aim to be comparable to or supersede existing methods. A major challenge with real-time processing is the demand for efficient methods and algorithms. Effort has been directed into applying an efficient solution for detection and localization of high-frequency interception signals.

As of today, Compressive Sensing is still an active field of study, and there is an ever-growing significant amount of literature. The peculiar features of Compressive Sensing have been thoroughly examined in this thesis, with the objective gradually shifting over to an actual application of the theory. Attempting to get an overview and still restricting myself to focus on the important features have been a personal challenge; I love exploration and have learned far more than what is being presented in this humble paper. Of the challenges I have encountered on my journey, most of them been particularly insightful and enjoyable, some have been stressful and put me in despair and yet a few which left me wondering.

I would like to thank my supervisors Fredrik Hekland at KDS and Alex Hansen at NTNU for their support during this journey. I also wish to thank my superior Eileen Frydenberg for encouraging me to write this thesis. Finally, Kristin Omre at NTNU has earned my gratitude for her willingness and support with the administrative perspective of the thesis.

Contents

1	Introduction	1
1.1	Introduction	1
1.2	Organization of the Thesis	3
1.3	Notation	3
2	Compressive Sensing and Applications	4
2.1	Introduction to Compressive Sensing	4
2.2	Sparsity Revisited	7
2.3	Performance Guarantees	9
2.4	Restricted Isometry Property	12
2.5	Donoho-Tanner Phase Transition	15
2.6	Robust Reconstruction	17
2.7	Detection Theory	20
2.7.1	Classical Detection Theory	21
2.7.2	Compressive Detection Theory	23
2.7.3	Compressive Estimation	25
2.8	Direction of Arrival Estimation	25
2.8.1	The Narrowband Model	25
2.8.2	The Gürbüz-McClellan-Cevher Model	27
3	Detection and Localization of Interception Signals	29
3.1	Constructing the Basis	31
3.2	Random Vector Correlation	36

3.3	Handling Detections	37
3.4	Estimating Direction-of-Arrival	41
3.5	Scenario Challenges	41
3.5.1	Scenario 1: Passive Sonar	42
3.5.2	Scenario 2: Active Sonar	42
4	Simulations	43
4.1	Phase Transition Simulations	45
4.2	Random Vector Correlation Values	48
4.3	Basis Coverage	49
4.3.1	Coverage in (f, k) -space	49
4.3.2	Coverage in (f, f_1) -space	53
4.3.3	HFM and LFM correlation	56
4.4	Detector Simulations	59
4.4.1	Probability of Detection	59
4.4.2	Classification Performance	61
4.5	Direction-of-Arrival Simulations	65
4.5.1	Stretching the Design Frequency	65
4.5.2	Passive DOA Estimation Accuracy	70
4.5.3	Active DOA Estimation Accuracy	71
4.6	Passive Sonar Simulations	74
4.7	Active Sonar Simulations	78
4.8	Active and Passive Performance Using Real Data	82

5	Future work	87
6	Concluding remarks	87
A	Waveforms	89
B	Parametrization of the Phase Transition	90
C	Computing RMS Values	90
D	Metrics in Bearing Space	92

List of Figures

1.1	Anti-submarine warfare with active sonar	1
1.2	Common waveforms for anti-submarine warfare	2
2.1	Visualization of the ℓ_p -norm ball	8
2.2	Differences in solution by ℓ_1 and ℓ_2 minimization	10
2.3	Mapping from \mathbb{R}^N to \mathbb{R}^M	14
2.4	The strong and weak phase transitions	16
2.5	The RIP compared with DT phase transition	17
2.6	Phase transition for AMP based algorithms	18
2.7	Noise sensitivity phase transition	21
2.8	Compressive detector performance	24
2.9	Compressive beamforming	28
3.1	Phase matching and correlation	35
3.2	Random vector correlation probability distribution function	38
3.3	Expected random vector correlation magnitude	39
4.1	Reconstruction probability for LASSO	46
4.2	Reconstruction probability for OMP	46
4.3	Reconstruction probability for EMBP	47
4.4	LFM basis coverage in (f, k) -space	51
4.5	HFM basis coverage in (f, k) -space	52
4.6	LFM basis coverage in (f, f_1) -space	54
4.7	HFM basis coverage in (f, f_1) -space	55
4.8	LFM basis correlated with HFM correlation	57

4.9	HFM basis correlated with LFM correlation	58
4.10	Probability of detection	60
4.11	LFM detector classification error	63
4.12	HFM detector classification error	64
4.13	Exceeding the design frequency	67
4.14	Energy dissipation	68
4.15	Conventional beamforming compared with GMC	69
4.16	Passive direction-of-arrival estimation accuracy	72
4.17	Active direction-of-arrival estimation accuracy	73
4.18	Passive sonar example 1	76
4.19	Passive sonar example 2	77
4.20	Active sonar example 1	79
4.21	Active sonar example 2	81
4.22	Spectrogram of LFM pulse	82
4.23	Beamformer output of real data	83
4.24	Detection and DOA estimates from active sonar	85
4.25	Detection and DOA estimates from passive sonar	86
D.1	Metrics in bearing space	92

List of Tables

1	Overview of simulations	44
2	Size of detection basis	48
3	Random vector correlation	48

1 Introduction

1.1 Introduction

Underwater warfare heavily relies on the use of sonar¹ to detect, track and engage hostile targets. A sonar can operate in two basic modes: passive and active. Passive sonars are utilized on surface vessels and with more success on submarines with the purpose to detect other vessels based on their noise emissions. Unlike the completely silent operation of passive sonars, active sonars will actively ping the ocean in an attempt to detect targets in the returning echoes. Due to their noisy operation, active sonars are primarily found on frigates and other vessels intended for anti-submarine warfare, see figure 1.1 for an illustrative image.

An active sonar can in principle transmit any desired waveform, however, in practice only a small class of waveforms are used due to their superior theoretical capabilities in various situations. The three most commonly encountered waveforms are linear frequency modulation (LFM), hyperbolic frequency modulation (HFM) and continuous wave (CW). See figure 1.2 for plots of these waveforms and the corresponding one-sided Fourier spectrum. Without going into too many details regarding these waveforms, it is sufficient to mention that LFM and HFM exhibits high pulse compression² capabilities and can resolve closely spaced targets while CW offers much better capabilities for detecting Doppler shift induced by target movement. See appendix A for additional details.

¹Originally an acronym for SOund NAVigation and Ranging.

²Performing a matched filter with a waveform that permits pulse compression will result in a tall and narrow peak. Two closely spaced echoes can be resolved and the signal to noise ratio is increased.




Figure 1.1: A frigate sends out a ping (black wavefronts) which is reflected from the submarine (red wavefronts) and finally detected by the towed array.




Figure 1.2: The three most commonly encountered waveforms (a)-(c) and the corresponding frequency spectrum (d)-(e).

There are several distinct aspects with regard to this game of anti-submarine warfare. Firstly, it is advantageous for active sonar to emit the ping as loud as possible to enhance the returned echoes, often limited only by cavitation or the surroundings³. Secondly, the emitted waveform may reach a potentially high frequency as is the case for torpedoes which can operate in tens of kHz due to their limited size. Thirdly, knowing the direction of the pinger is vital for the survival of the submarine.

These considerations led most modern submarines to be equipped with a high-frequency interception array specialized for detecting such waveforms. The signal processing can be heavy; extremely high sampling rates combined with conventional beamforming to estimate the direction of arrival claims its toll on the available hardware.

We will attempt to address these issues by presenting a relatively new method known as Compressive Sensing where the number of samples can be dramatically decreased while offering a substantially better estimation of the direction-

³Active sonars can be notoriously difficult to operate in shallow waters due to reverberation. Increasing the transmission power may not even work under these conditions.

of-arrival (DOA). In particular we will consider a method where it is sufficient with only one sensor to be sampled at the Nyquist rate while all others can be substantially downsampled. The validity and efficacy of the method will be evaluated and simulation results presented. Two scenarios will be investigated using this method: passive sonar detecting a ping from a hostile pinger or an active sonar detecting return echoes. In both scenarios we will assume the signal-to-noise ratio to be relatively high.

1.2 Organization of the Thesis

In chapter 2 we will describe the framework of Compressive Sensing, including behavior under noise and applications for estimating direction-of-arrival. The following chapter 3 will depict a method to solve the aforementioned issues using the framework of Compressive Sensing, with the simulation results presented in chapter 4. Finally we do a small discussion of future work in chapter 5, only to be quickly followed up by concluding remarks in chapter 6.

1.3 Notation

The notation should be fairly straightforward, but to avoid any ambiguities we will detail the symbols used and their significance. Unless specified otherwise, scalars will be denoted in lower case, e.g. x , vectors \mathbf{f} in bold-face and matrices A in capitals. We use \mathbb{R} to represent the real numbers, \mathbb{C} the complex numbers and \mathbb{N} the natural numbers. The notation $A \in \mathbb{R}^{M \times N}$ represents a matrix A with M rows and N columns whose entries are real numbers. The transpose and conjugate transpose (Hermitian conjugate) is denoted as \cdot^T and \cdot^H respectively. The identity matrix is simply taken as \mathbf{I} and its size is determined from the context.

The length of a vector $\mathbf{x} = \{x_1, x_2, \dots, x_n\} \in \mathbb{R}^N$ in space ℓ_p is given by the associated norm $\|\mathbf{x}\|_p$ where $p \geq 1$. In most instances $p = 2$ and the norm is the usual Euclidean distance. In the context of Euclidean space, we let $\langle \cdot, \cdot \rangle$ denote the inner product. This thesis will also make heavily use of $p = 1$, or the Manhattan distance. We will also encounter the $p = 0$ “norm” which simply counts the number of non-zero entries in \mathbf{x} . Take care not to confuse the notation for the absolute value $|\cdot|$ with the norm.

As for probability distributions, we will let $\mathcal{N}(\mu, \sigma^2)$ denote the normal distri-

bution with mean μ and variance σ^2 , and $\mathcal{U}(a, b)$ is the uniform distribution with constant, non-zero value on the interval $[a, b]$ only. Random variables are written in upper case and the realization of a random variable is written in lower case. $\mathbb{E}[\cdot]$ represents the expectation value of a random variable.

2 Compressive Sensing and Applications

2.1 Introduction to Compressive Sensing

This short summary outlines the basics of Compressive Sensing. Subsequent sections will in detail investigate various aspects which is only briefly covered here.

For decades the Nyquist-Shannon sampling theorem [11, 12] stood as firm as a mountain in the field of signal processing. The theorem gave relief and support in stating what minimum rate a signal should be sampled for accurate reconstruction. However, it was also an impenetrable barrier for those who wished to do with fewer samples and still be able to reconstruct the signal accurately; the theorem did not allow any such extravagance.

The advent of Compressive Sensing shook the conception of the sampling theorem. Based on the groundbreaking work of Candès, Romberg, Tao [2, 3, 4, 5] and not to mention Donoho [6, 7, 8, 9, 10] a robust framework for sampling a broad class of signals with far fewer samples proved its existence. By no means did it invalidate the sampling theorem, but merely paved a path for signals whose underlying representation is known to be sparse.

Let $\{\boldsymbol{\psi}_n\}_{n=1}^N$ be a set of N orthonormal basis vectors for the space \mathbb{R}^N . Given an orthonormal representation basis $\Psi = \{\boldsymbol{\psi}_1, \dots, \boldsymbol{\psi}_N\}$, $\Psi \in \mathbb{R}^{N \times N}$, any signal $\mathbf{f} \in \mathbb{R}^N$ can be expressed by⁴

$$\mathbf{f} = \sum_{i=1}^N \boldsymbol{\psi}_i x_i = \Psi \mathbf{x}, \quad \mathbf{x} \in \mathbb{R}^N \quad (2.1)$$

The vector \mathbf{x} , often referred to as the coefficient vector, is said to be K -sparse when at most K entries of \mathbf{x} are non-zero. If all but the K largest entries of \mathbf{x}

⁴For easier theoretical treatment we will mostly consider real valued signals and orthonormal bases, although these restrictions can be greatly relaxed to include complex and over-complete bases.

are comparably small, \mathbf{x} is said to be approximately K -sparse. By keeping only the K largest entries and setting the others to zero we denote the approximate vector for \mathbf{x}_K .

In the framework of Compressive Sensing, sensing is accomplished by linear measurements of \mathbf{f} with respect to an orthonormal measurement basis $\Phi \in \mathbb{R}^{M \times N}$.

$$\mathbf{y} = \Phi \mathbf{f}, \quad \mathbf{y} \in \mathbb{R}^M \quad (2.2)$$

From these $M < N$ measurements Compressive Sensing guarantees exact recovery of x with high probability provided M is large enough and that Φ and Ψ are sufficiently incoherent [1].

The coherency between two sets of basis vectors is defined to be

$$\mu(\Phi, \Psi) = \sqrt{N} \max_{1 \leq m \leq M, 1 \leq n \leq N} |\langle \phi_m, \psi_n \rangle| \quad (2.3)$$

This is merely the largest correlation between any two elements of Φ and Ψ . As for how large or how small, linear algebra gives the answer $\mu \in [1, \sqrt{N}]$.

Sparsity and incoherence stand as the pillars of Compressive Sensing; both determine the number of samples needed for faithful reconstruction. Contemplate on their significance and one might see the workings behind Compressive Sensing. Sparsity indicates the actual information level is much lower than first anticipated, sparking a hope that one should only need about $\mathcal{O}(K)$ samples. This, however, requires all coefficients to be represented in the few samples taken; exactly the job of incoherence. Incoherence makes the measurement matrix dense in the representation basis, ensuring that the few samples taken of \mathbf{x} in the “dense domain” cover most of the sparse domain, thereby capturing the coefficients of \mathbf{x} in the sparse domain with high probability.

As a prime example consider the extreme case with maximum coherency between the measurement and representation bases. For simplicity choose $\Phi = \Psi = \mathbf{I}$ and thus $\mu = \sqrt{N}$. Each measurement yields $y_m = \phi_m \Psi \mathbf{x} = x_m$ and we see that exactly N measurements are necessary; nothing is gained as each measurement only conveys information about a single coefficient. Choose instead a measurement basis incoherent to the representation basis; e.g Φ consisting of orthonormalized vectors sampled independently and uniformly on the unit sphere. With high probability $\mu = \sqrt{2 \log N}$ [1]. Each measurement is guaranteed to contain information about almost all the coefficients and fewer measurements are needed to uniquely determine the solution. As for how many measurements, Candès and Romberg provides the follow bound [2]

$$M \geq C \cdot \mu^2 \cdot K \log N \quad (2.4)$$

where C is a (small) constant. More accurate and improved bounds are presented in section 2.4 and section 2.5.

Wind back a moment and embrace the role randomness played in the previous example. A random measurement basis is universal in the sense it will with high probability possess low coherence with any representation basis, thus *sensing with randomness ought to do the trick* and is shown to be near-optimal [3, 6]. This does not exclude deterministic designs however. In particular, it is shown that a Fourier basis and Dirac spikes are maximally incoherent with respect to each other [1], and in section 2.5 we briefly mention a particular combination of bases which exhibits extreme performance.

All that is left is to reconstruct the data with the following convex optimization program [4, 10] with $A \equiv \Phi\Psi \in \mathbb{R}^{M \times N}$

$$\min \|\tilde{\mathbf{x}}\|_1, \quad \text{subject to } \mathbf{y} = A\tilde{\mathbf{x}} \quad (\text{P1})$$

Of all possible solutions $\tilde{\mathbf{f}} = \Psi\tilde{\mathbf{x}}$ that are consistent with the data we choose the one whose coefficients has the smallest ℓ_1 norm. This program can be recast as a linear program making available a host of ever more efficient solution algorithms [13, 14]. This can be done by defining \mathbf{x}^+ and \mathbf{x}^- such that

$$\begin{aligned} x_i^+ &= \max\{x_i, 0\} \\ x_i^- &= \max\{-x_i, 0\} \end{aligned}$$

The problem now becomes

$$\min \mathbf{1}^\top \begin{pmatrix} \mathbf{x}^+ \\ \mathbf{x}^- \end{pmatrix}, \quad \text{subject to } (A \ -A) \begin{pmatrix} \mathbf{x}^+ \\ \mathbf{x}^- \end{pmatrix} = \mathbf{y}, \begin{pmatrix} \mathbf{x}^+ \\ \mathbf{x}^- \end{pmatrix} \geq 0$$

and can be solved efficiently.

Under the influence of noise in the sampling process, the following mixed program is proposed for reconstruction

$$\min \|\tilde{\mathbf{x}}\|_1, \quad \text{subject to } \|\mathbf{y} - A\tilde{\mathbf{x}}\|_2 \leq \epsilon \quad (\text{P2})$$

where ϵ bounds the amount of noise in the data. This is again a convex problem, referred to as *LASSO* (see [10]) and can be solved efficiently [14, 34].

There is much to gain by comparing with lossy compression. A signal is sampled at its full rate and *all* coefficients are computed to determine the K largest ones. These coefficients are then transmitted and the rest of data discarded;

a computational intensive method, yet common in image and audio processing. Compressive Sensing challenges this by measuring just what is needed and shifting the computational burden onto the receiver; only a few samples are measured and transmitted, and the receiver has the job with reconstruction. This does not replace traditional compression techniques, but paves way for applications where it can be expensive to perform high-rate sampling or transmitting large amounts of data, such as sonobuoys, distant sensor nodes and MRI.

2.2 Sparsity Revisited

This section attempts to give an intuitive explanation of why sparsity is vital for reconstruction and how ℓ_1 enters the picture as a sparsity-promoting function.

The following system of linear equations

$$\mathbf{y} = A\mathbf{x} \tag{2.5}$$

is in general ill-posed since $A \in \mathbb{R}^{M \times N}$ and $M < N$. To circumvent this issue another restriction on (2.5) is needed to obtain a unique solution. Following the principle of Occam's razor [15] to choose the solution that *describes the data best in the simplest possible way*, we enforce a sparsity constraint on (2.5) to obtain this unique solution among all possible solutions.

The function counting the number of non-zero components of a vector is the Hamming distance, often denoted for ℓ_0 "norm" (we abuse terminology and discard quotation marks for better readability, see [16] for why this is not a proper norm). The ℓ_p -norm is defined by [16]

$$\|\mathbf{x}\|_p \equiv \left(\sum_{i \in \mathbb{N}} |x_i|^p \right)^{\frac{1}{p}}$$

from which we obtain the ℓ_0 norm by taking the resulting limit as $p \rightarrow 0$

$$\|\mathbf{x}\|_0 = \lim_{p \rightarrow 0} \left(\sum_{i \in \mathbb{N}} |x_i|^p \right)^{\frac{1}{p}}$$

Unfortunately the following program

$$\min \|\tilde{\mathbf{x}}\|_0, \quad \text{subject to } \mathbf{y} = A\tilde{\mathbf{x}} \tag{P0}$$




Figure 2.1: The ℓ_p -norm ball (red) scaled to intersect with a linear constraint (blue). Sparse solutions are achieved in (a)-(c), while (d) and (e) produce non-sparse solutions.

requires a combinatorial search and is proven to be NP hard. Another sparsity-promoting function is desired.

For many decades the ℓ_1 norm has been used as a sparsity-promoting function with great success and important theoretical results supports this use [8]. This can be easily understood from figure 2.1 which shows the smallest ℓ_p -norm ball for a linear equality constraint. Intersection at an axis forces all other components to zero and sparsity is obtained. As indicated, $0 \leq p \leq 1$ produces sparse solutions while $p > 1$ does in general not. Despite ℓ_1 is bordering as a sparsity-promoting function, substitution of ℓ_0 with ℓ_1 is near-optimal for recovering sparse solutions after a result of Donoho [6]. Choosing ℓ_1 ensures that the resulting minimization program is convex and that we are guaranteed to arrive at the global optimum. Better reconstruction quality can be obtained by using IRLS⁵ with $p < 1$, however the program is no longer convex and the solver may only reach a local optimum.

A final example of the significant difference in solution between minimization of

⁵IRLS = Iterative Reweighted Least Squares

ℓ_1 and ℓ_2 is shown in figure 2.2. Two completely different solutions are obtained and it is clear that ℓ_1 enhances sparsity while ℓ_2 does not.

2.3 Performance Guarantees

The universality of random sensing and general performance of ℓ_1 minimization are important topics covered by this section.

The work of Donoho [6] and Candès and Tao [3] give important answers to these topics. Consider the coefficients of a *compressible signal*⁶ sorted in decreasing order of magnitude $|x_n| > |x_{n+1}|$. If

$$x_n \leq R \cdot n^{-1/p} \quad (2.6)$$

for some $R > 0$ and $0 < p < \infty$ then the following bound on reconstruction error by keeping only the K largest coefficients of \mathbf{x} holds

$$\|\mathbf{x} - \mathbf{x}_K\|_2 \leq C \cdot R \cdot K^{-r}, \quad r = \frac{1}{p} - \frac{1}{2} \quad (2.7)$$

for some constant C .

Following Donoho [6] let F be the class of signals \mathbf{f} for which the sparse representation obeys (2.6) for a given R and p . Let $I_M : F \mapsto \mathbb{R}^M$ be an information operator that samples M pieces of information about $\mathbf{x} \in \mathbb{R}^N$ in the following fashion

$$I_M(\mathbf{x}) = (\langle \boldsymbol{\xi}_1, \mathbf{x} \rangle, \dots, \langle \boldsymbol{\xi}_M, \mathbf{x} \rangle)$$

where the sampling kernels $\boldsymbol{\xi}_m$ are nonadaptive, i.e. independent of \mathbf{x} .


An approximate reconstruction is offered by an unspecified, possibly non-linear algorithm $P : \mathbb{R}^M \mapsto \mathbb{R}^N$, and the reconstruction performance is measured by the minimax ℓ_2 error

$$E_M(F) = \inf_{P_M, I_M} \sup_{\mathbf{x} \in F} \|\mathbf{x} - P_M(I_M(\mathbf{x}))\|_2$$


So far all possible methods of nonadaptive sampling are allowed and all possible methods of reconstruction are allowed. At this stage Donoho provides the following result (modification of Theorem 1 in [6])

$$E_M(F) \leq C \cdot R \cdot (M/\log N)^{-r} \quad (2.8)$$

⁶These are signals for which ordinary compression techniques works well due to the rapid decay of coefficients.



(a) ℓ_1 solution



(b) ℓ_2 solution

Figure 2.2: Differences in solution by ℓ_1 (a) and ℓ_2 (b) minimization. Exact solution is indicated by red crosses while blue circles shows the recovered solution. The recovered solution in (a) is exact.

for some constant C and sufficiently large N . Direct comparison with (2.7) seems to indicate that having $M \sim K \log N$ pieces of nonadaptive information provided by I_M is equivalent with direct knowledge of the K largest coefficients. A similar result was also discovered independently by Candès and Tao [3], which also confirmed the optimality of the information operator by arguing that one cannot, in general, design a set of K measurements that would allow significantly better reconstruction error by any method, no matter how intractable.

Earlier literature on sparse reconstruction suggested the required number of samples to be $M \sim K^2$ for accurate reconstruction, which is, ultimately, disappointingly small. The bound by Donoho, Candès and Tao greatly improved upon this for large systems.

In section 2.4 we will see that at least $M = 2K$ samples are required to avoid two vectors \mathbf{x}_1 and \mathbf{x}_2 , $\mathbf{x}_1 \neq \mathbf{x}_2$, to be mapped to the same output values, $A\mathbf{x}_1 \neq A\mathbf{x}_2$. This requirement can be elevated to $M = K + 1$ and *most* K -sparse vectors can still be distinguished, something that will become evident in the discussion on phase transitions in section 2.5.

A more surprising implication of (2.8) is that we do not know in advance which coefficients are likely to be the important ones, yet the *optimal* information operator is nonadaptive and depends on most of the class F and not on the specific signal. Indeed, if the information operator was changed to be adaptive in the sense that each sampling kernel ξ_m could depend on the previous samples, the following result (Theorem 2 in [6]) shows that this is of little help

$$E_M(F) \leq 2 \cdot E_M^{\text{Adaptive}}(F)$$

A similar result discouraging adaptive sampling schemes was discovered by Arias-Castro, Candès and Davenport [17].

A natural question arises on how to construct this optimal information operator. Equation (2.4) provides a clue by requiring low-coherence pairs, e.g. Fourier basis and spikes constitutes a low-coherence pair, wavelets and noiselets another. As already implied in section 2.1 random sensing provides a general answer. The following result by Donoho shows that sensing with random matrices sampled independently and uniformly on the unit sphere yields near-optimal information operators with overwhelming high probability (Theorem 5 in [6])

$$E_M(I_M^{\text{random}}, F) \leq C \cdot E_M(F)$$

where $C > 0$ is some constant and $E_M(I_M^{\text{random}}, F)$ is

$$E_M(I_M^{\text{random}}, F) \equiv \inf_{P_M} \sup_{\mathbf{x} \in F} \left\| \mathbf{x} - P_M \left(I_M^{\text{random}}(\mathbf{x}) \right) \right\|_2$$

This result is also supported by the work of Candes and Tao [3]. In addition Theorem 10 in [6] shows that nearly all random matrices satisfies (2.3).

Turn the attention onto the reconstruction algorithm P . A comforting result by Donoho is that ℓ_1 minimization (P1) is also near-optimal for reconstruction (Theorem 5 in [6])

$$E_M(P_M^{\ell_1}, I_M^{\text{random}}, F) \leq C \cdot E_M(F)$$

where $C > 0$ is some constant and $E_M(P_M^{\ell_1})$ is

$$E_M(P_M^{\ell_1}, I_M^{\text{random}}, F) \equiv \sup_{\mathbf{x} \in F} \left\| \mathbf{x} - P_M^{\ell_1} \left(I_M^{\text{random}}(\mathbf{x}) \right) \right\|_2$$

Finally, if \mathbf{x} has at most $C \cdot M / \log N$ non-zeroes for some constant C , then reconstruction by ℓ_0 and ℓ_1 minimization has the same unique solution (Theorem 8 in [6]). Compare this with (2.4).

The emphasis on random measurement matrices and recovery by ℓ_1 minimization so far is due to the two being near-optimal, universal and tractable. An exact quantification of 'near-optimal' is shown in section 2.5. Implementation wise there can be a penalty in both storage and computational performance by choosing random matrices over deterministic designs which can be computed in real time and efficient algorithms exists.

2.4 Restricted Isometry Property

This section outlines a popular tool that is related to incoherence and has its purposes for theoretical results on the robustness and stability of encoding/decoding pairs.

Not every combination of measurement and representation matrices can provide good compressive capabilities with an ℓ_1 solver. The process of acquiring compressive samples must also ensure that all relevant information is not only conserved (see the example in section 2.1), but also be well distanced in measurement space. As a result of this Candès and Tao proposed a key notion that is useful for studying the robustness of Compressive Sensing, dubbed *Restricted Isometry Property* (RIP)[1, 18].

Definition 2.1. For each integer $K = 1, 2, \dots$, define the restricted isometry constant $0 \leq \delta_K \leq 1$ of a matrix $A \in \mathbb{R}^{M \times N}$ as the smallest number such that

$$(1 - \delta_K) \|\mathbf{x}\|_2 \leq \|A\mathbf{x}\|_2 \leq (1 + \delta_K) \|\mathbf{x}\|_2 \quad (2.9)$$

holds for all K -sparse vectors \mathbf{x} .

An equivalent statement is to require all sub matrices $A_K^T A_K$ derived from A by taking K columns to be positive definite with eigenvalues in $[(1 - \delta_K)^2, (1 + \delta_K)^2]$ [20, 21]. A matrix is said to obey the RIP of order K if δ_K is not too close to one [1]. When it holds, the Euclidean length of any K -sparse vector is approximately conserved. This in turn implies no K -sparse vector can be located in the null space of A . If it did, no reconstruction algorithm, no matter how intractable, would be able to recover these vectors. In particular, Gaussian and other well behaved random matrices are shown to satisfy the RIP of order K given [1, 35]

$$M \geq C \cdot K \log \frac{M}{K} \quad (2.10)$$

where C is some constant.

To avoid the potential disastrous case of degeneracy, that is, two sparse vectors being mapped to the same values, δ_{2K} must be sufficiently smaller than one. In other wording, any pairwise distance between two K -sparse vectors must also be well preserved (see figure 2.3)

$$(1 - \delta_{2K}) \|\mathbf{x}_2 - \mathbf{x}_1\|_2 \leq \|A\mathbf{x}_2 - A\mathbf{x}_1\|_2 \leq (1 + \delta_{2K}) \|\mathbf{x}_2 - \mathbf{x}_1\|_2$$

The implications of RIP are delightful. Assume $\delta_{2K} < \sqrt{2} - 1$ and we are guaranteed tractable, robust and stable recovery (see e.g. [19, 20] and the references therein). Tractable recovery is provided in the sense that all K -sparse vectors are perfectly recovered by ℓ_1 minimization.

Let $\tilde{\mathbf{x}}$ be the solution from (P1), and \mathbf{x}_K be the vector \mathbf{x} with all but the largest K components set to zero. Robust recovery then follows from the next two bounds

$$\|\mathbf{x} - \tilde{\mathbf{x}}\|_1 \leq C \|\mathbf{x} - \mathbf{x}_K\|_1 \quad (2.11)$$

$$\|\mathbf{x} - \tilde{\mathbf{x}}\|_2 \leq C \frac{\|\mathbf{x} - \mathbf{x}_K\|_1}{\sqrt{K}} \quad (2.12)$$

for some constant C . The implications are astounding; \mathbf{x} is not necessarily a sparse signal, but (2.11) and (2.12) asserts that the quality of reconstruction




Figure 2.3: Mapping from \mathbb{R}^N to \mathbb{R}^M by a matrix that satisfies the RIP property.

is just as good as knowing the K largest components of \mathbf{x} in advance (in the event \mathbf{x} happens to be K -sparse, the recovered solution is of course exact).

Assume we were to measure the signal in addition to some noise $\|\mathbf{e}\|_2 \leq \epsilon$

$$\mathbf{y} = A\mathbf{x} + \mathbf{e} \quad (2.13)$$

and $\tilde{\mathbf{x}}$ is the solution by mixed norm minimization (P2), then the following bound holds for all $\mathbf{x} \in \mathbb{R}^N$

$$\|\mathbf{x} - \tilde{\mathbf{x}}\|_2 \leq C \frac{\|\mathbf{x} - \mathbf{x}_K\|_1}{\sqrt{K}} + C_1 \epsilon \quad (2.14)$$

where C and C_1 are some constants. The quality of reconstruction is bounded by two terms; the first term corresponds to near sparse data and the second term is just proportional to the noise level. This result assures Compressive Sensing degrades gracefully in the event of noise and provides stable reconstruction.

The RIP is a popular tool used to establish theoretical performance guarantees for a broad range of encoding/decoding pairs under influence of noise. However, while the implications of RIP are excellent, it suffers from two major drawbacks. To start with, it is difficult (NP hard) to tell whether a matrix obeys the RIP or not as it requires checking the eigenvalues of each sub matrix [19, 21] or a combinatorial search over all possible K -sparse vectors. A workaround is provided by using random matrices which almost always provides good restricted isometry constants, thus the emphasis on randomness. Secondly, and more importantly as shown in subsequent chapters, the RIP is *too* strict to be of practical use [21, 22, 24].

2.5 Donoho-Tanner Phase Transition

Several reconstruction algorithms, including ℓ_1 minimization, features a phase transition where the probability of recovery drops sharply from near unity to near zero. This section investigates this exotic and surprising behavior.

Donoho and Tanner spent their time counting faces of polytopes in high-dimensional geometry and that with great success. The surprising connection between random matrices, probability and high-dimensional geometry resulted in what is known as the *Donoho-Tanner phase transition* [22, 24]. As mentioned previously, replacing ℓ_0 with ℓ_1 is near-optimal; Donoho-Tanner phase transition shows for which matrices (P1) solves (P0). What is more surprising is that Donoho and Tanner investigated the phase transition for multiple probability distributions and discovered that the underlying distribution is irrelevant as long as it is well behaved; the phase transition is in this sense universal (see also [23]).

To continue and be able to interpret this result we must define the under- and over-sampling ratios respectively. The undersampling ratio is defined as

$$\delta = \frac{M}{N}, \quad \delta \in [0, 1] \quad (2.15)$$

while the oversampling ratio takes the form

$$\rho = \frac{K}{M}, \quad \rho \in [0, 1] \quad (2.16)$$

Although both are assumed in the limit $N \rightarrow \infty$ we will still use these for finite (and often very small) values of N .

There are two phase transitions of importance in the limit $N \rightarrow \infty$: the strong and the weak, both are depicted in figure 2.4. The strong phase transition is the boundary where exact recovery is guaranteed for all random matrices (no probability of failure), while the weak phase transition is where the probability for recovery is $\frac{1}{2}$. The rapid decay of the transition region (it decays with $N^{-\frac{1}{2}}$ as the problem size increases [24]) indicates that the weak phase transition boundary is the most interesting feature when designing Compressive Sensing applications. By accepting a negligible probability of failure, this phase transition boundary sets the limitation for reconstruction by (P1) and so replaces “the rule of thumb” (2.4) for a more accurate measure on the number of samples needed for reconstruction in the noise-less case. Approximate formulas are [25, 31]

$$M \geq 2K \log \frac{N}{M}, \quad M \ll N \quad (S1)$$




Figure 2.4: The strong phase transition (green) and the weak phase transition (black). S1 (blue) and S2 (red) are approximate formulas presented in the paper. Data with courtesy of J. Tanner [26].

for small values of δ and with a slight difference

$$M \geq 2K \log \frac{N}{K} \quad (\text{S2})$$

which holds more generally over a broad range of δ . See figure 2.4.

We take a step back to compare the RIP with the strong Donoho-Tanner phase transition as shown in figure 2.5 (note the change of vertical axis units). Despite the obvious pessimistic nature, the RIP still stands as a viable (and sometimes the only) tool for analysis of stability bounds under the influence of noise.

The phase transition does not only depend on the encoding matrix, but also on the recovery algorithm used. Over time multiple recovery algorithms have been proposed that can potentially outperform the Donoho-Tanner Phase Transition by exploiting different heuristics than (P1). We will briefly consider the *Approximate Message Passing* (AMP) algorithm originally proposed by Donoho, Maleki and Montanari [32] which yields a computationally attractive solution to (P2) in the presence of noise while being robust to the signal distribution. It is proven that in the large systems limit, $N \rightarrow \infty$, AMP does coincide with LASSO and thus possesses the same phase transition as described above [32, 33]. Better performance can be obtained by exploiting the underlying sig-




Figure 2.5: Comparing the RIP with the *strong* phase transition shows it is too pessimistic for practical purposes. See [27].

nal distribution (e.g. Bernoulli-Gaussian) through a Bayesian approach known as belief propagation (BP), see figure 2.6 [34].

Krzakala et al. recently showed that by using spatially-coupled matrices, the upper theoretical limit on exact reconstruction, $M = K + 1$ or $\rho = 1$, could be achieved in the noise-less case [28, 29]. Donoho et al. further improved on this work and also proved its robustness to noise [31]. However, enforcing A to be band diagonal implies no universal Φ exists and this approach has its uses for specialized designs only.

2.6 Robust Reconstruction

Rarely a signal goes undistorted by noise. This section will summarize and point out important aspects of noise. A major conclusion is the inefficiency of Compressed Sensing under the influence of noise.

Noise can be introduced into the process in two ways: the original signal is contaminated with noise before sampling or the measurements are contaminated by noise during transmission

$$\mathbf{y} = A(\mathbf{x} + \mathbf{z}) + \mathbf{w} \quad (2.17)$$




Figure 2.6: Phase transition for AMP based algorithms which exploits the underlying signal distribution compared with ℓ_1 minimization (dashed). Note the upper theoretical limit for exact reconstruction, $\rho = 1$, can be achieved in combination with a special class of spatially-coupled matrices. Data with courtesy of F. Krzakala [30].

Both forms of noise are shown to be essentially equivalent, with the exception the former can potentially have a much larger impact than the latter. Assume for a moment that $\mathbf{z} \sim \mathcal{N}(0, \sigma_0^2)$ and $\mathbf{w} \sim \mathcal{N}(0, \sigma^2)$ then (2.17) is equivalent to (2.13) with $\mathbf{e} \sim \mathcal{N}(0, \sigma^2 + \sigma_0^2/\delta)$. This is known as noise folding and the result is a significant lower signal-to-noise ratio (SNR) [37, 38]. As it is often desired to keep δ rather small, this can potentially degrade the reconstruction completely. It appears the only solution is to increase the number of measurements, which incidentally conflicts with the idea of Compressive Sensing.

Furthermore, how to handle sparse reconstruction under the influence of noise? Sparsity is crucial for reconstruction and algorithmic stability is required to handle noise. As elaborated by Xu et al. [36], these two desired properties for robust Compressive Sensing are fundamentally at odds with each other: sparse algorithms cannot be stable and vice versa. This implies ℓ_1 minimization cannot be stable, while ℓ_2 minimization, which is known to have strong stability properties, cannot produce sparse solutions (see section 2.2 and in particular figure 2.2). A trade off between sparsity and stability has to be made. Proposed is reconstruction by (P2) and under the assumption the RIP is satisfied, the stability bound (2.14) follows. Another, much more interesting bound also exists for reconstruction by (P2). Let $\mathbf{e} \sim \mathcal{N}(0, \sigma^2 \mathbf{I})$, and $M \geq C \cdot K \log \frac{N}{K}$ for some constant C , then with high probability the solution $\tilde{\mathbf{x}}$ has the following bound (see [38] and the references within)

$$\frac{1}{N} \|\tilde{\mathbf{x}} - \mathbf{x}\|_2^2 \leq C \cdot \frac{K\sigma^2}{M} \log N \quad (2.18)$$

where C is some constant. Note how the error scales with $\delta = M/N$. This effect is sometimes cited as a drawback of the compressive sensing framework and can be understood intuitively. If each measurement has a constant SNR, then taking more measurements should reduce the estimation error.

Candés and Davenport raise an important concern regarding (2.18) [38]. A randomly generated matrix A will be almost orthogonal to the signal \mathbf{x} , leading to a tremendous SNR loss. Compare (2.18) with an oracle which can design A with vectors localized to the support of \mathbf{x} . The oracle would be able to generate an estimate obeying

$$\mathbb{E} \left[\frac{1}{N} \|\tilde{\mathbf{x}} - \mathbf{x}\|_2^2 \right] \leq \left(\frac{K\sigma^2}{M} \right) \left(\frac{K}{N} \right)$$

Comparing these two shows that the performance of (2.18) is worse than what would be possible with an oracle by a factor of $(N/K) \log N$. For small values of K , this factor can be quite large. It is natural to ask whether (2.18) can

be improved upon. Specifically, is there a more intelligent choice of A and an accompanying efficient recovery algorithm that can give a fundamentally better bound? The answer is negative [38].

These results combined indicates Compressive Sensing performs subpar under noise.

Exact quantification of LASSO under the influence of noise has been studied by both Donoho et al and Wainwright. Wainwright considered signed support recovery and provides an exact inequality on the number of samples required for reconstruction, assuming one solves an equivalent variant of (P2) with a specified regularization parameter [39]. Of much more interest is the result by Donoho et al [40]. Let $\mathbf{y} = A\mathbf{x} + \mathbf{e}$ with $\mathbf{e} \sim \mathcal{N}(0, \sigma^2\mathbf{I})$. Consider the mean-squared error of reconstruction by (P2)

$$\text{MSE} = \frac{1}{N} \|\tilde{\mathbf{x}} - \mathbf{x}\|_2^2$$

From this we define the *noise sensitivity*

$$\text{NS}(\rho, \delta) = \inf_{\alpha} \sup_{\sigma > 0} \sup_G \frac{\text{MSE}(\rho, \delta, \alpha, \sigma, G)}{\sigma^2}$$

where α is the tuning parameter of the recovery algorithm, σ^2 is the noise variance and G denotes the marginal distribution of \mathbf{x} . This can be restated as choosing the least favorable distribution and noise level while tuning the algorithm optimally considering these choices; this is doing the best out of the absolutely worst. By choosing this error measure we ensure that any other choice of underlying probability distribution or noise variance will only lead to a more favorable result. See figure 2.7 for contour plots of NS. If the noise sensitivity is large, then the measurement noise may severely degrade the reconstruction.

2.7 Detection Theory

Here we will consider classical detection theory and compare with compressive detection theory.




Figure 2.7: Contour lines of the Noise Sensitivity in the (δ, ρ) plane. The black line is the Donoho-Tanner phase transition and corresponds with a noise sensitivity of infinity.

2.7.1 Classical Detection Theory

We wish to differentiate between the two hypotheses

$$\begin{aligned}\mathcal{H}_0 &: \mathbf{y} = \mathbf{n} \\ \mathcal{H}_1 &: \mathbf{y} = \mathbf{s} + \mathbf{n}\end{aligned}$$

where $\mathbf{s} \in \mathbb{R}^N$ is a known signal and $\mathbf{n} \sim \mathcal{N}(0, \sigma^2 \mathbf{I})$ is the noise. We are mainly interested in the *false alarm rate* and the *detection rate* probability measures

$$\begin{aligned}P_F &= \Pr(\mathcal{H}_1 \text{ given } \mathcal{H}_0 \text{ is true}) \\ P_D &= \Pr(\mathcal{H}_1 \text{ given } \mathcal{H}_1 \text{ is true})\end{aligned}$$

It is often desired to control the probability of false alarm rate, $P_F = \alpha$, and maximize the probability of detection, P_D , given this constraint. This is known as the Neyman-Pearson (NP) detector and the optimal rule can be shown to be the likelihood ratio test

$$\Lambda(\mathbf{y}) = \frac{f_1(\mathbf{y})}{f_0(\mathbf{y})} \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{>}} \eta \quad (2.19)$$

where η is chosen such that

$$P_F = \int_{\Lambda(\mathbf{y}) > \eta} f_0(\mathbf{y}) \, d\mathbf{y} = \alpha$$

For our hypotheses the likelihood functions are

$$f_0(\mathbf{y}) = (2\pi\sigma^2)^{-\frac{N}{2}} e^{-\frac{\|\mathbf{y}\|_2^2}{2\sigma^2}}$$

$$f_1(\mathbf{y}) = (2\pi\sigma^2)^{-\frac{N}{2}} e^{-\frac{\|\mathbf{y}-\mathbf{s}\|_2^2}{2\sigma^2}}$$

As it is easier to work with the logarithm of these functions we take the logarithm of (2.19) to obtain an equivalent test that simplifies to

$$t = \langle \mathbf{s}, \mathbf{y} \rangle \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{>}} \sigma^2 \log \eta + \frac{1}{2} \|\mathbf{s}\|_2^2 = \gamma$$

It can be shown that t is a sufficient statistic for our detection problem and contains enough information to differentiate between \mathcal{H}_0 and \mathcal{H}_1 . In the signal processing literature, t is often referred to as the matched filter. It is easy to see that t behaves like

$$t \sim \mathcal{N}(0, \sigma^2 \|\mathbf{s}\|_2^2) \quad \text{under } \mathcal{H}_0$$

$$t \sim \mathcal{N}(\|\mathbf{s}\|_2^2, \sigma^2 \|\mathbf{s}\|_2^2) \quad \text{under } \mathcal{H}_1$$

and thus

$$P_F = \Pr(t > \gamma | \mathcal{H}_0) = Q\left(\frac{\gamma}{\sigma \|\mathbf{s}\|_2}\right)$$

$$P_D = \Pr(t > \gamma | \mathcal{H}_1) = Q\left(\frac{\gamma - \|\mathbf{s}\|_2^2}{\sigma \|\mathbf{s}\|_2}\right)$$

where $1-Q(z)$ is the cumulative probability distribution of the standard normal distribution

$$Q(z) = (2\pi)^{-\frac{1}{2}} \int_z^\infty e^{-\frac{u^2}{2}} \, du$$

Returning to the constraint $P_F = \alpha$ and solving for γ yields

$$\gamma = \sigma \|\mathbf{s}\|_2 Q^{-1}(\alpha)$$

and the probability of detection is readily found to equal

$$P_D(\alpha) = Q\left(Q^{-1}(\alpha) - \sqrt{SNR}\right)$$

where we have defined $SNR = \frac{\|\mathbf{s}\|_2^2}{\sigma^2}$.

2.7.2 Compressive Detection Theory

The classical detection scheme can be extended to handle compressive measurements [41, 42, 43, 44, 45]. In this case the hypotheses are

$$\begin{aligned}\mathcal{H}_0 &: \mathbf{y} = \Phi \mathbf{n} \\ \mathcal{H}_1 &: \mathbf{y} = \Phi(\mathbf{s} + \mathbf{n})\end{aligned}$$

where $\Phi \in \mathbb{R}^{M \times N}$ is a known measurement matrix. The likelihood functions take upon the form

$$\begin{aligned}f_0(\mathbf{y}) &= \frac{e^{-\frac{1}{2}\mathbf{y}^\top(\sigma^2\Phi\Phi^\top)^{-1}\mathbf{y}}}{|\sigma^2\Phi\Phi^\top|^{\frac{1}{2}} \cdot (2\pi)^{\frac{M}{2}}} \\ f_1(\mathbf{y}) &= \frac{e^{-\frac{1}{2}(\mathbf{y}-\Phi\mathbf{s})^\top(\sigma^2\Phi\Phi^\top)^{-1}(\mathbf{y}-\Phi\mathbf{s})}}{|\sigma^2\Phi\Phi^\top|^{\frac{1}{2}} \cdot (2\pi)^{\frac{M}{2}}}\end{aligned}$$

and taking the logarithm of the likelihood ratio tests then gives

$$t = \mathbf{y}^\top(\Phi\Phi^\top)^{-1}\Phi\mathbf{s} \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \sigma^2 \log \eta + \frac{1}{2}\mathbf{s}^\top\Phi^\top(\Phi\Phi^\top)^{-1}\Phi\mathbf{s} = \gamma$$

From this we easily deduce the behavior of t

$$\begin{aligned}t &\sim \mathcal{N}\left(0, \sigma^2\mathbf{s}^\top\Phi^\top(\Phi\Phi^\top)^{-1}\Phi\mathbf{s}\right) \quad \text{under } \mathcal{H}_0 \\ t &\sim \mathcal{N}\left(\mathbf{s}^\top\Phi^\top(\Phi\Phi^\top)^{-1}\Phi\mathbf{s}, \sigma^2\mathbf{s}^\top\Phi^\top(\Phi\Phi^\top)^{-1}\Phi\mathbf{s}\right) \quad \text{under } \mathcal{H}_1\end{aligned}$$

and thus

$$\begin{aligned}P_F &= \Pr(t > \gamma | \mathcal{H}_0) = Q\left(\frac{\gamma}{\sigma\sqrt{\mathbf{s}^\top\Phi^\top(\Phi\Phi^\top)^{-1}\Phi\mathbf{s}}}\right) \\ P_D &= \Pr(t > \gamma | \mathcal{H}_1) = Q\left(\frac{\gamma - \mathbf{s}^\top\Phi^\top(\Phi\Phi^\top)^{-1}\Phi\mathbf{s}}{\sigma\sqrt{\mathbf{s}^\top\Phi^\top(\Phi\Phi^\top)^{-1}\Phi\mathbf{s}}}\right)\end{aligned}$$

Controlling the probability of failure by setting $P_F = \alpha$ allows us to solve for γ

$$\gamma = \sigma\sqrt{\mathbf{s}^\top\Phi^\top(\Phi\Phi^\top)^{-1}\Phi\mathbf{s}}Q^{-1}(\alpha)$$

and the probability of detection is readily found to equal

$$P_D(\alpha) = Q\left(Q^{-1}(\alpha) - \frac{\sqrt{\mathbf{s}^\top\Phi^\top(\Phi\Phi^\top)^{-1}\Phi\mathbf{s}}}{\sigma}\right)$$




Figure 2.8: Probability of detection P_D as a function of downsampling ratio δ . The false alarm probability has been set to $\alpha = 0.001$. Classical detection theory occurs at $\delta = 1$.

These expressions can be reduced in the case of Φ being an orthonormal measurement basis. Inserting $\Phi\Phi^T = \mathbf{I}_M$ allows us to write

$$t = \langle \Phi \mathbf{s}, \mathbf{y} \rangle$$

for the sufficient statistic and

$$P_D(\alpha) = Q\left(Q^{-1}(\alpha) - \frac{\|\Phi \mathbf{s}\|_2}{\sigma}\right)$$

for the detection probability. Assuming Φ is drawn on random, then with high probability

$$P_D(\alpha) = Q\left(Q^{-1}(\alpha) - \sqrt{M/N}\sqrt{SNR}\right)$$

as compared with the classical detector. The reduction in detection performance scales directly with the number of samples taken as noted in section 2.6. See figure 2.8 for comparison between compressive detector and classical detector. For sufficiently high signal-to-noise ratio the compressive detector is seen to be essentially equal to the classical detector in performance even for a significant downsampling ratio.

2.7.3 Compressive Estimation

While the theory on detection is all sound, the test $t \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{>}} \gamma$ requires knowledge of the signal level in advance and must be estimated from the data itself.

Specifically, suppose we observe the data $\mathbf{y} = \Phi \mathbf{s}$ and we wish to simultaneously estimate a set \mathcal{S} of linear functions $\langle \mathbf{s}_i, \mathbf{s} \rangle$ from the measurements \mathbf{y} . Assuming Φ is orthonormal, a natural estimator from the preceding section would be

$$\frac{N}{M} \langle \mathbf{y}, \Phi \mathbf{s}_i \rangle, \quad i = 1, 2, \dots, |\mathcal{S}|$$

Provided Φ is random and sufficiently well behaved, then Davenport et al [41] provides the following bound on estimation error with probability at least $1 - r$

$$\left| \frac{N}{M} \langle \Phi \mathbf{s}, \Phi \mathbf{s}_i \rangle - \langle \mathbf{s}, \mathbf{s}_i \rangle \right| \leq \kappa_r \frac{\|\mathbf{s}\|_2 \|\mathbf{s}_i\|_2}{\sqrt{M}}$$

where $\kappa_r = 2\sqrt{12 \log \left(\frac{4|\mathcal{S}|+2}{r} \right)}$. The bound is seen to grow only sub-linearly with the set size $|\mathcal{S}|$, however it decays rather slowly as a function of M .

2.8 Direction of Arrival Estimation

Direction of arrival (DOA) estimation is an important aspect in sonar processing. The number of sources is ultimately sparse in the bearing space and the framework of Compressive Sensing can be applied. The literature describes two models for estimating direction of arrival: the widely common narrowband model [46, 47, 48, 49] and a more general model by Gürbüz, McClellan and Cevher [50]. The subsequent chapters will detail both models.

2.8.1 The Narrowband Model

Assume a narrowband source emits a signal of the following form

$$x(t) = v(t)e^{i\omega t}$$

namely a plane wave with angular frequency ω modulated by a signal $v(t)$ whose bandwidth is small compared to ω such that the narrowband approximation

holds. Allowing the source to be located sufficiently distanced away from the observer ensures the validity of the plane wave model. The observer consists of L sensors in an array capable of measuring the attenuated and delayed signal. For simplicity we restrict ourselves to two dimensions such that the angle $\theta \in [-\pi, \pi)$ is sufficient to denote the arrival of the plane wave. The geometry of the array is of little importance; common models are the uniform linear array, the cylindrical array and random positioning. To obtain meaningful results we must assume the design frequency of the array equals or exceeds the plane wave frequency ω or grating lobes can be induced and give rise to incorrect estimates.

Neglect the attenuation factor⁷ and consider only the relative time-differences between sensors as the signal arrives onto the array. The l -th sensor measures the signal

$$\begin{aligned} y_l(t) &= x(t - \tau_l) \\ &= v(t - \tau_l)e^{i\omega(t-\tau_l)} \\ &\approx v(t)e^{i\omega t} \cdot e^{-i\omega\tau_l} \\ &= x(t)e^{-i\omega\tau_l} \end{aligned}$$

and τ_l is the time-difference for the signal to reach sensor l compared with a reference point in space and depends on the incident angle of arrival. A time-delay is seen to correspond to a phase shift. Collecting these phase delays gives rise to the array response vector⁸

$$\mathbf{a}(\theta) = [e^{-i\omega\tau_1}, \dots, e^{-i\omega\tau_L}]^\top$$

Adding noise $\mathbf{e}(t)$ to the equation the system model takes the following form for a single plane wave incident onto the array

$$\mathbf{y}(t) = \mathbf{a}(\theta)x(t) + \mathbf{e}(t)$$

This model can easily be extended to multiple sources impinging from different directions. Assume the number of sources to equal K and collect the array response vectors into the following fashion

$$A = [\mathbf{a}_1^\top, \dots, \mathbf{a}_K^\top]^\top$$

and the source signals into the following vector

$$\mathbf{x}(t) = [x_1(t), \dots, x_K(t)]^\top$$

⁷The attenuation factor is taken to be equal for all sensors in the array. This approximation is valid when the distance to the source is much greater than the spatial extent of the array.

⁸Another popular name is the steering vector as it “steers” the array to look in the direction of θ .

then the model changes accordingly as we incorporate multiple sources

$$\mathbf{y}(t) = A\mathbf{x}(t) + \mathbf{e}(t) \quad (2.20)$$

By no coincidence is the notation quite similar to the one used previously in this thesis. To employ the Compressive Sensing framework we need to know all the entries of A and the vector \mathbf{x} must be sparse. This is not the case as the source directions are unknown. To overcome these issues the bearing space is discretized into $N \gg K$ entries and the vector \mathbf{x} corresponds to these discrete bearings. Most of the entries in \mathbf{x} will be zero as the number of directions where a source exists are few. By discretizing the bearing space we are now able to populate the matrix A accordingly with one column vector for each discretized bearing in \mathbf{x} . As a result, the problem is now in a familiar form and can be solved by (P2) to determine the non-zero bearing locations.

A note or two is worth mentioning about the method just described. Unlike the emphasis on random vectors throughout this paper, this method uses a pair of Fourier⁹ and spike bases, which exhibits maximum incoherence as briefly mentioned earlier [1]. As (2.20) stands, only a single (complex) sample from each sensor is needed to obtain the DOA estimate. Extension to broadband sources can be easily made by solving for each frequency independently and then combining the results. This approach can be used in compressive beamforming [51], see figure 2.9.


2.8.2 The Gürbüz-McClellan-Cevher Model

This model works with any signal, not necessarily restricted to narrowband signals. For the method to work a reference signal must be known in advance. This will be the case for active sonars where the emitted ping will serve as the reference signal. For passive applications one sensor can be used as a reference sensor to obtain the reference signal. With a single source impinging onto the sensors the reference signal will be readily available, but in the event of multiple sources Gürbüz, McClellan and Cevher showed that this situation could be modeled as noise on the reference signal from the other sources.


Again the bearing space is discretized into N directions and the entries of \mathbf{x} will correspond with the discrete bearings. Consider a collection of T time samples from the l -th sensor

$$\mathbf{s}_l = [s_l(t_0), \dots, s_l(t_{T-1})]^T$$

⁹Take care to note that the representation basis will in general be overcomplete, complex and consist of the array response vectors.



(a) Conventional Beamformer



(b) Compressive Beamformer

Figure 2.9: Conventional beamforming compared with compressive beamforming in a linear, unscaled plot. An uniform linear array of 40 elements was used. Three broadband targets of opportunity at 60° , 70° and 108° . In addition a fourth, previously unknown target actively pings with HFM followed by CW at 150° . The compressive beamformer clearly produces narrower tracks, with the HFM slightly spread out.

The vector \mathbf{x} can then be related to the signal \mathbf{s}_l from the l -th sensor as follows

$$\mathbf{s}_l = \Psi_l \mathbf{x}$$

where the n -th column of the representation basis for the l -th sensor consists of time shifts of the reference signal $s(t)$ corresponding with the discretized bearings

$$[\Psi_l]_n = [s(t_0 + \tau_{l,n}), \dots, s(t_{T-1} + \tau_{l,n})]^\top$$

From each sensor M_l noisy compressive samples are taken with a random measurement basis Φ_l as previously described

$$\mathbf{y}_l = \Phi_l \Psi_l \mathbf{x} + \mathbf{e}_l$$

where \mathbf{e}_l denotes the noise in the compressive measurements. Finally data from each sensor is stacked

$$\begin{aligned} \mathbf{y} &= [\mathbf{y}_1^\top, \dots, \mathbf{y}_L^\top]^\top, & \mathbf{y} &\in \mathbb{R}^{LM_l \times 1} \\ \mathbf{e} &= [\mathbf{e}_1^\top, \dots, \mathbf{e}_L^\top]^\top, & \mathbf{e} &\in \mathbb{R}^{LM_l \times 1} \\ \Phi &= \begin{pmatrix} \Phi_1 & & 0 \\ & \ddots & \\ 0 & & \Phi_L \end{pmatrix}, & \Phi &\in \mathbb{R}^{LM_l \times LT} \\ \Psi &= [\Psi_1^\top, \dots, \Psi_L^\top]^\top, & \Psi &\in \mathbb{R}^{LT \times N} \end{aligned}$$

and setting $A = \Phi\Psi$ we arrive at the familiar form $\mathbf{y} = A\mathbf{x} + \mathbf{e}$.

The advantage of this method is the clear reduction in data rate. In passive applications one sensor is still required to be sampled at the Nyquist frequency to obtain the reference signal while all the other sensors can be sub-sampled. Compare this with the narrowband model where all sensors must be sampled at the Nyquist frequency for a faithful recovery of broadband signals. The intrinsic disadvantages of the Gürbüz-McClellan-Cevher method is the modeling of other sources as noise and the rather large system which must be solved. The inability of obtaining beam data directly may be another obstacle for certain applications.

3 Detection and Localization of Interception Signals

In this chapter we will outline the general strategy to detect and localize high-frequency interception signals using the theory from the preceding sections.

Two scenarios will be considered as briefly noted in the introduction section:

- Passive sonar detecting a ping from a hostile pinger.
- Active sonar detecting return echoes.

These two scenarios are very similar with two subtle differences. In the first scenario, we do not know the waveform in advance and we are mainly interested in direct path of propagation; all other echoes are of no interest to us (however multiple echoes can enable estimation of range to the pinger or be used for bi-static detection). The second scenario has the advantage of knowing the transmitted waveform, but the number of returned echoes can be significant, although still assumed to be sparse in time and bearing. In general the array configurations and properties (such as sampling frequency) would be different for each scenario, but for practical reasons and comparability we use the same aperture for both scenarios.

The actual interception array will in both scenarios be modeled as circular in the plane of interest and with a reference sensor in the middle. To avoid ambiguities in direction-of-arrival estimation at least three sensors in addition to the reference sensor is needed. We fix the number of sensors to equal 8 in the simulations, and will investigate the number of compressive samples required. We might as well utilize the concept of Compressive Sensing completely and sub-sample the reference sensor as well. Empirical testing indicates compressive beamforming achieves a better resolution while being robust when the design frequency is exceeded, encouraging to distance the sensors further apart than conventional use would dictate. Examples of this claim will be presented in the simulations section.

The general strategy for both scenarios is as follows:

- A basis of interception signals must be constructed.
- A detection must be made.
- An estimate of the signal origin must be found.

The general procedure will be to apply compressive detection combined with estimation and then use Gürbüz-McClellan-Cevher (GMC) method to estimate the direction of arrival. For the GMC method to work, the waveform parameters must be known. Subsequent sections will discuss the required steps in detail.

3.1 Constructing the Basis

The challenge with passive sonar is foremost to detect and estimate the waveform. This also applies to a certain degree for active sonar, albeit we possess clues about the waveform and the required basis size can be made significantly smaller. The reference sensor is proposed to be sub-sampled as this will reduce the computational burden if the set of interception signals is large.

Constructing a basis of the interception signals is a daunting task as the parameter space is multidimensional and infinite. The parameters that need to be estimated are

- Amplitude
- Phase factor
- Start frequency
- Band width
- Pulse duration
- Pulse type

The following strategies are proposed to reduce the parameter set as much as possible. To start with, a valid assumption will be to assume the pulse duration greatly exceeds whatever processing block length is in use. This ensure we are not troubled with the additional issue of estimating the onset of the pulse and we can expect better correlation between the measured waveform and our basis of interception signals.

Furthermore, employing a basis of complex signatures will capture both quadrature components of the waveform. Following the section on compressive estimation this will provide an estimate on both the amplitude and phase factor. This estimate will then be used to aid the detection procedure by providing both an estimate of the signal strength and noise level (which we consider to be everything not explained by the estimated waveform) so that the test $t \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{>}} \gamma$ can be done, see section 3.3 for detailed description.

Unfortunately, the other parameters cannot be estimated in a similar fashion. A multitude of strategies exists however, only limited by imagination. These range

from brute force approaches using a sufficiently dense basis of waveforms at various parameters to more elaborate divide-and-conquer approaches involving estimating one parameter and then use a second basis centered around this parameter. No matter which strategy we choose to pursue, the basis must be sufficiently dense to capture all possible waveforms, or the set of the most likely ones. This raises an important question: how dense should the basis be sampled? There is a clear trade-off between using longer block lengths for enhanced detection versus short blocks where the uncertainty principle ensures few basis vectors are needed. To answer this question we will first investigate the instantaneous frequency for HFM and its approximation LFM, described by two parameters f and k .

$$f(t) = \frac{f}{1 - kt} \approx f + fkt, \quad 0 \leq t \leq T$$

Depending on the value of k this either increases, stays constant or decreases in frequency. The initial frequency is limited by $0 < f < \frac{f_s}{2}$ where f_s is the sampling frequency. From this we can work out the bounds on k . This results in the following bounds

$$-\infty < k < \frac{f_s - 2f}{Tf_s} \quad (\text{HFM})$$

$$-\frac{1}{T} < k < \frac{f_s - 2f}{2fT} \quad (\text{LFM})$$

Note that both HFM and LFM are unbounded. HFM is unbounded from below when the final frequency $f_1 \rightarrow 0$ and LFM is unbounded from above when the initial frequency $f \rightarrow 0$. The divergence in k is only superficial; it is the choice of initial and final frequencies that has an impact as the phase is continuous and bounded despite the divergences in k . Yet we cannot ignore the divergence as it affects the sampling density. Restricting the ratio $\frac{f}{f_1}$ will provide finite bounds. A particular choice is $\frac{1}{2} < \frac{f}{f_1} < 2$, which we will use throughout the thesis. This seemingly restrictive choice is justified by the assumption of pulse length greatly exceeding the block length. The consequence of this assumption is exponential coverage as the pulse length increases, ensuring almost every possible pulse will in turn be detected. The bounds then take upon the following forms

$$-\frac{1}{T} < k < \min \left\{ \frac{1}{2T}, \frac{f_s - 2f}{Tf_s} \right\} \quad (\text{HFM})$$

$$-\frac{1}{2T} < k < \frac{1}{T} \quad (\text{LFM})$$

This should cover most waveforms encountered. We then return to the question about sampling the (f, k) space. The density hinges critically on the block

length used and we are interested in this relation. As a first approximation we treat the two parameters f and k independently. We can then obtain an estimate of the required density in f -space by setting $k = 0$ and correlating two LFM waveforms spaced Δf apart in frequency. This gives rise to the well-known sinc function.

$$\alpha = \frac{1}{N} \left| \sum_{n=0}^{N-1} e^{2\pi i \frac{\Delta f}{f_s} n} \right| = \left| \frac{\sin \frac{\pi N \Delta f}{2 f_s}}{N \sin \frac{\pi \Delta f}{2 f_s}} \right|$$

where we have used that the assumption that the pulse duration exceeds the block length, allowing us to insert $T = \frac{N}{f_s}$ where N is the block length. Using the approximation $\sin x \approx x - \frac{1}{3}x^3$ then yields for values of α close to unity

$$\frac{\Delta f}{f_s} \approx \frac{2\sqrt{3(1-\alpha)}}{\pi N}$$

The parameter α controls the minimum value of correlation between two waveforms and thus the loss in SNR due to misalignment in frequency between the two waveforms. Allowing a loss of maximum 3 dB is a common choice (analysis of how well a random vector correlates with the basis will be presented later) and corresponds to $\alpha = \frac{1}{\sqrt{2}}$. For this particular choice of α we obtain the approximation

$$\frac{\Delta f}{f_s} \approx \frac{1}{2N}$$

Thus, we must sample at twice this spacing, resulting in a spacing of $\Delta f \approx \frac{f_s}{N}$ in f -space, a well-known result from Fourier analysis. The corresponding analysis for k -space using LFM is slightly more complex.

$$\alpha = \frac{1}{N} \left| \sum_{n=0}^{N-1} e^{2\pi i \frac{f}{2} \Delta k \frac{n^2}{f_s^2}} \right|$$

There is no known closed form of this summation. By controlling the phase such that it does not exceed the corresponding linear phase in the previous case, we can assume the sums to be approximately equal

$$\frac{f}{2} \Delta k \frac{N^2}{f_s^2} \approx \frac{\Delta f}{f_s} N$$

Inserting $\Delta f = \frac{f_s}{2N}$ then results in

$$\Delta k^{LFM} \approx \frac{f_s^2}{f N^2}$$

Thus, we must sample at twice this spacing. By noting that k is bound by a limit that depends on N^{-1} , this means on the order of N grid points in k -space is sufficient to capture all LFM waveforms. A similar analysis can be done for the HFM. Again we employ the assumption of matching phases

$$-f \log \left[\frac{\left(1 - \left(k - \frac{\Delta k}{2}\right) \frac{N}{f_s}\right)^{\frac{1}{k - \Delta k/2}}}{\left(1 - \left(k + \frac{\Delta k}{2}\right) \frac{N}{f_s}\right)^{\frac{1}{k + \Delta k/2}}} \right] \approx \pm \frac{\Delta f}{f_s} N \quad (3.1)$$

For $k \rightarrow 0$ we expect the result to coincide with LFM. For $k \neq 0$ we apply a linear approximation

$$\Delta k^{HFM} = \Delta k^{LFM} + \left. \frac{\partial \Delta k^{HFM}}{\partial k} \right|_{k=0} k$$

Differentiating (3.1) with respect to k and setting $k = 0$ results in the exact expression

$$\left. \frac{\partial \Delta k^{HFM}}{\partial k} \right|_{k=0} = \frac{2(1 - \eta^2) \log\left(\frac{1+\eta}{1-\eta}\right) - 4\eta}{(1 - \eta^2) \log(1 - \eta^2) + 2\eta^2}$$

where we have set $\eta = \frac{\Delta k^{LFM} N}{2f_s}$. This expression can be greatly simplified by using the approximations $\log(1 + x) \approx x$ and $\log\left(\frac{1+x}{1-x}\right) \approx 2x + \frac{2}{3}x^3$. Keeping the terms to leading order in η then results in the final expression


$$\Delta k^{HFM} \approx \Delta k^{LFM} - \frac{8}{3} \frac{\Delta k^{LFM} N}{2f_s} k$$

We note the value of k for which $\Delta k^{HFM}(k) = 0$ is $k = \frac{3}{4T} > \frac{1}{2T}$, thus the step size will always be positive and non-zero on the defined interval. The approximations used above perform well as shown in figure 3.1. The correlation is seen to be generally slightly less than $\frac{1}{\sqrt{2}}$, however not by a large amount. Thus, about N grid points in k -space is sufficient to capture all HFM waveforms, albeit non-uniform sampling should be utilized. Denote k_0 for the starting point, then to cover the k -space neatly the next sample point $k_1 > k_0$ is related to k_0 by $k_1 = k_0 + \Delta k^{HFM}(k_0) + \Delta k^{HFM}(k_1)$. Solving for k_1 results in


$$k_1 = \frac{k_0 + \Delta k^{HFM}(k_0) + \Delta k^{LFM}}{1 - \Delta k^{HFM}(1) + \Delta k^{LFM}}$$

The opposite rule $k_1 < k_0$ is easily found to equal

$$k_1 = \frac{k_0 - \Delta k^{HFM}(k_0) - \Delta k^{LFM}}{1 + \Delta k^{HFM}(1) - \Delta k^{LFM}}$$



(a) Phase matching and resulting Δk



(b) Correlation using Δk

Figure 3.1: Plot (a) shows the resulting Δk from phase matching, while (b) shows the amount of correlation using the specified Δk . The term linear refers to the linear approximation. The parameters used to create these plots are $N = 1024$, $f = f_s/4$ and $f_s = 200$ kHz.

These formulas are prone to numerical instabilities however and care must be taken when dealing with small frequencies. To avoid this we suggest to start sampling at the largest possible value of k and work backwards. This will just slightly oversample the k -space, but not by a large amount.

To summarize; capturing most waveforms requires on the order of N points in f -space combined with N points in k -space, for a total of N^2 points. As the basis size is seen to grow quite rapidly, we might be forced to work with short block lengths. One issue that would be of interest is how well LFM and HFM correlates. This has direct impact on the estimation of pulse type and may reduce the basis size if these two waveforms correlate well. A proposed preliminary method is to remove highly correlated entries from the precomputed basis. Some simulations will be presented, but we leave the theory open for future research.

3.2 Random Vector Correlation

In section 3.1 we accepted a loss of maximum 3 dB for a waveform correlated with a basis of the same type. This loss should be compared with how well a random vector correlates with the basis, and in particular, ensure any correlation with a waveform exceeds random correlations by a significant amount. It should be obvious that increasing the block length should reduce the correlation with random vectors. On the other hand, increasing the number of basis vectors will tend to increase the maximum correlation factor as there are more possibilities. Realizing that this is a complex issue, we take on a probabilistic view and investigate this behavior.

To ease the analysis we assume the basis to be real and contain B vectors. Let the random vector be normal distributed $\mathcal{N}(0, \frac{1}{N})$ where N is the processing block length. By construction, the random vector will on average have unit norm and can be compared with results from section 3.1 and simulations. The correlation coefficient between this random vector and the b -th basis vector will also be normal distributed $C_b \sim \mathcal{N}(0, \frac{1}{N}), b \in \{1, \dots, B\}$ since the basis is normalized. The detection procedure chooses the largest correlation value in magnitude, Y , and we are interested in finding the distribution of this value from the independent correlation factors

$$Y = \max\{|C_1|, |C_2|, \dots, |C_B|\}$$

It is trivial to see that the random variable $|C_b|$ is distributed with a one-sided

normal distribution and has the probability distribution function

$$f_{|C|}(c) = \begin{cases} \sqrt{\frac{2N}{\pi}} e^{-c^2 \frac{N}{2}} & c \geq 0 \\ 0 & c < 0 \end{cases}$$

The cumulative probability distribution follows immediately by integration and is seen to be related to the error function

$$F_{|C|}(c) = \begin{cases} \operatorname{erf}\left(\sqrt{\frac{N}{2}}c\right) & c \geq 0 \\ 0 & c < 0 \end{cases}$$

The cumulative probability distribution of the maximum of correlation values is trivial to obtain. In fact,

$$F_Y(y; B) = P((C_1 \leq y) \cap (C_2 \leq y) \cap \dots \cap (C_B \leq y)) = F_{|C|}(y)^B$$

The probability distribution of Y is thus seen to be well concentrated and the variance is correspondingly small, see figure 3.2. In fact, as $N \rightarrow \infty$ the limiting probability distribution is a Dirac delta function, and we can use the expectation value as a measure for comparison. Since the random variable Y has a non-negative support we can determine the expectation value from


$$\mathbb{E}\{Y\} = \int_0^\infty 1 - F_Y(y; B) dy$$

The expected correlation magnitude is seen to depend on the block length and basis size as elaborated earlier, see figure 3.3. Values for the expected maximum correlation magnitude will be tabulated in the simulations section together with block lengths and the corresponding basis sizes.


3.3 Handling Detections

Here we combine the theory of compressive detection and estimation to handle detections of interception signals. Estimates of the test statistic and γ must be found to decide whether the interception signal is detected or not. We will assume a properly normalized complex basis of interception signatures is used for the treatment to follow. Note that the norm of the real and imaginary parts of the basis vectors separately are in general not equal, although we will assume this is the case.

Let $\{\mathbf{s}_i\}, \mathbf{s}_i \in \mathbb{C}^N$ denote the basis of interception signals whose vectors are constructed from the complex exponential. Assume the signal received on the



(a) $N=128$



(b) $N=256$

Figure 3.2: The probability distribution function (pdf) of the maximum correlation magnitude as a function of block length N and basis size B . In (a) $N = 128$ and in (b) $N = 256$.




Figure 3.3: Contour plot showing the expected random vector correlation magnitude. Note that both axes are logarithmic.

reference sensor consists of the i -th basis vector with a complex amplitude β_i hidden in Gaussian noise.

$$\mathbf{y} = \Phi (\Re\{\beta_i \mathbf{s}_i\} + \mathbf{n}), \quad \mathbf{n} \sim \mathcal{N}(0, \sigma^2 \mathbf{I})$$

where \Re denotes the real part and $\Phi \in \mathcal{R}^{M \times N}$ is an orthonormal measurement basis whose entries are normal distributed $\phi_{ij} \sim \mathcal{N}(0, \frac{1}{N})$. We compute the inner product between the compressive measurements and the basis vector

$$\tilde{t} = \langle \mathbf{y}, \Phi \mathbf{s}_i \rangle$$

This may resemble the test statistic, but this is not the case as \mathbf{s}_i is complex and the amplitude is unknown. Our goal is to provide estimates of the noise variance, $\hat{\sigma}^2$, and the complex amplitude, $\hat{\beta}_i$, from \tilde{t} , and then ultimately find estimates of the test statistic \hat{t} and level $\hat{\gamma}$ so that the test $\hat{t} \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\geq}} \hat{\gamma}$ can be performed.

The estimate of β_i follows directly from the chapter on compressive estimation

$$\hat{\beta}_i = 2 \frac{N}{M} \langle \mathbf{y}, \Phi \mathbf{s}_i \rangle = 2 \frac{N}{M} \tilde{t}$$

where the factor of two is entirely due to only the real part being represented in \mathbf{y} . Let \mathbf{y}^s denote \mathbf{y} in the noise-less case. An estimate of \mathbf{y}^s then follows

readily as

$$\hat{\mathbf{y}}^s = \Re\{\hat{\beta}_i \Phi \mathbf{s}_i\}$$

and after some manipulations we obtain an estimator for the signal norm

$$\|\hat{\mathbf{y}}^s\|_2^2 = \frac{1}{2} \Re \left\{ (\hat{\beta}_i^H)^2 \sum_m (\Phi \mathbf{s}_i)_m^2 + |\hat{\beta}_i|^2 \sum_m |\Phi \mathbf{s}_i|_m^2 \right\}$$

Although seemingly complex, this form has the advantage of not requiring any sums to be computed during runtime as both sums can be precomputed and stored in advance. The estimate of the test statistic t follows once $\hat{\mathbf{y}}^s$ is known

$$\begin{aligned} \hat{t} &= \langle \mathbf{y}, \hat{\mathbf{y}}^s \rangle \\ &= \hat{\beta}_i^H \tilde{t} \end{aligned}$$

where we used the definition of $\hat{\mathbf{y}}^s$ and $\hat{\beta}_i$ to arrive at the final form.

Assuming the residual $\mathbf{y} - \hat{\mathbf{y}}^s$ then represents the noise, a biased estimate of the noise variance is found by

$$\hat{\sigma}^2 = \frac{1}{M} \sum_m (\mathbf{y}_m - \hat{\mathbf{y}}_m^s)^2$$

as we can safely assume \mathbf{y} and \mathbf{y}^s both have a mean of zero. If the set of basis vectors is large, this can become a rather computational intensive task. Expanding the square allows us to write after some manipulations

$$\hat{\sigma}^2 = \frac{1}{M} \sum_m \mathbf{y}_m^2 + \frac{1}{M} \|\hat{\mathbf{y}}^s\|_2^2 - \frac{2}{M} \hat{t}$$

and the number of operations needed is dramatically reduced. Knowing $\hat{\sigma}$ and $\|\hat{\mathbf{y}}^s\|_2$ readily gives us an estimate for γ

$$\hat{\gamma} = \hat{\sigma} \|\hat{\mathbf{y}}^s\|_2 Q^{-1}(\alpha)$$

and finally the test $\hat{t} \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\geq}} \hat{\gamma}$ can be performed.

As a final note, only the compressive measurements of the basis is required to be stored, providing a small bonus to memory consumption.

3.4 Estimating Direction-of-Arrival

Once we have obtained a detection and the corresponding waveform we think resembles the actual data well enough, the next step is to estimate the direction of arrival by using the GMC method. The system matrix in GMC method can be poorly conditioned. To improve general reconstruction performance we propose to factorize the system matrix A using the singular value decomposition in the following way

$$A = U \cdot S \cdot V^T$$

By keeping the r largest eigenvalues we do not only improve the condition number, but also the reconstruction time as the system matrix is smaller. The new system then takes the form

$$\begin{aligned}\tilde{A} &= U_r^T U_r S_r V^T \\ \tilde{\mathbf{y}} &= U_r^T \mathbf{y}\end{aligned}$$

and we proceed with GMC as usual.

The choice of reconstruction algorithm should reflect the desired combination of computational effort, the number of compressive samples and sensors required and the number of source directions in case of echoes. Greedy algorithms excels in the case of passive sonar as we are then mainly interested in the direct path of propagation, which also carries most energy and will be picked by the first step of the algorithm. Greedy algorithms pursues another heuristics than ℓ_1 minimization and the phase transition will be correspondingly different. They are also simpler in nature and can be terminated rather quickly. The general procedure of a greedy algorithm is to minimize the residual by selecting vectors that contributes the most, from which they derive their name. In the case of active sonar, we propose an AMP-based based algorithm, although greedy algorithms can still be used.

3.5 Scenario Challenges

Each scenario has its challenges. Here we list the challenges and the proposed solutions.

3.5.1 Scenario 1: Passive Sonar

As noted earlier, passive sonar is mainly faced with the challenge of detecting and estimating the waveform. Having no clues on the waveform (except it can be CW, LFM or HFM) provides no reduction in the basis of interception signals. Requiring the basis to be sufficiently dense indicates of the order of N^2 vectors from sampling the (f, k) -space. As the processing block length, N increases, memory consumption and computational effort grows as N^3 in the detection step. Already at $N = 1024$ will most computers have exhausted all their memory. The computational limit is much more severe however. The number of blocks per second is $\frac{f_s}{N}$, and assuming N^3 operations per block this results in $f_s N^2$ operations per second. Simple mathematics reveals $N \sim 128$ marks the limit for a single core when we assume a sampling frequency of 200 kHz. Due to these considerations we proposed to subsample the reference sensor to reduce both memory usage and computational effort. Using a greedy algorithm instead of a fully fledged ℓ_1 minimization algorithm to obtain a DOA estimate should also improve the situation as the algorithm can be terminated almost immediately.

3.5.2 Scenario 2: Active Sonar

Even though the waveform is known, the assumption that the pulse greatly exceeds the block length used indicates a basis constructed from subdividing the waveform is still required. Only a single entry should be sufficient in the case of CW pulse, but to detect anything useful like a Doppler shift, a multitude of CW waveforms closely spaced around the transmitted frequency is needed. The basis for active sonar can and will be significantly smaller than what was the case for passive sonar.

We then perform compressed detection and all detections will be pursued since we are interested in multiple echoes. The detection step essentially works to filter out the possible basis vectors as the next step can be computational intensive. Estimating the phase of the detected basis vectors would certainly be advantageous, but in this case we cannot due to the possibility of multiple echoes of the same basis vector arriving simultaneously with different phases. To work around this issue we propose to use the complex basis vector directly in the GMC method. This should allow us to detect the same waveform, but with different phase and amplitude from multiple directions simultaneously. Output of the beamformer can be used to estimate the signal strength if needed. As

elaborated in section 2.6 on the topic of noise folding, adding noise directly into the system matrix resulted in an amplification of the noise inversely proportional to the downsampling factor. This could have severe effect on the direction-of-arrival estimation performance.

As a signal could imping on the array from multiple locations simultaneously, the reconstruction algorithm is required to fully complete, which can require a significant computational effort.

4 Simulations

An overview over simulations is presented in table 1. Unless specified otherwise, the default set of parameters for all simulations are

- Speed of sound is set to $c = 1500$ m/s.
- Sampling frequency is $f_s = 200$ kHz.
- Cylindrical array with 8 sensors and a reference sensor in the middle.
- Array spacing will be 8 times the conventional spacing.
- Processing block length is 256 samples.
- Reference sensor is downsampled by a factor of 4.
- Ten compressive samples per block length from every other sensor.
- Measurement bases are drawn on random to ensure good RIP constants.
- Probability of false alarm is set to $\alpha = 0.001$.
- GMC is solved using a greedy algorithm known as Orthogonal Matching Pursuit (OMP).
- In SVD precondition step, all eigenvalues smaller than a hundredth of the largest eigenvalue are discarded.
- Noise will not be added.
- Any noise added is normal distributed and uncorrelated.
- Bearing grid is uniform with a resolution of 1° .

Simulation description	Section
Reconstruction probability for ℓ_1 , OMP and AMP based algorithms.	section 4.1
Counting the number of basis vectors of a dense basis.	section 4.2
Correlating random vector with basis.	section 4.2
Basis coverage in (f, k) -space for LFM and HFM.	section 4.3.1
Basis coverage in (f, f_1) -space for LFM and HFM.	section 4.3.2
Basis correlation between LFM and HFM.	section 4.3.3
Probability of detection using detecting procedure.	section 4.4.1
Classification performance using detector procedure.	section 4.4.2
DOA accuracy improvement and robustness when exceeding the design frequency.	section 4.5.1
Accuracy of passive DOA estimation under noise and downsampling.	section 4.5.2
Accuracy of complex DOA estimation under noise and downsampling.	section 4.5.3
Examples of passive sonar usage.	section 4.6
Examples of active sonar usage.	section 4.7
Active and passive sonar usage with real data.	section 4.8

Table 1: Table of simulations done.

- Only one target will be assumed.

4.1 Phase Transition Simulations

There has been a lot of talk of reconstruction algorithms and the corresponding phase transitions. We will obtain reconstruction probability for three algorithms using Monte Carlo simulations and compare these with theoretical phase transitions. These algorithms are LASSO (ℓ_1 -minimization), Orthogonal Matching Pursuit (OMP, greedy algorithm) and EMBP (an AMP based solver [30]). The following steps were taken

1. Choose δ and ρ on random.
2. Construct orthonormal system matrix with random entries.
3. Reconstruction using the three algorithms.
4. Define success as relative reconstruction error less than 0.001.
5. Go to step 1 and repeat until sufficiently many trials are concluded.

On average 650 trials were obtained per data point. The theoretical phase transition for reconstruction by LASSO (P1) is seen to coincide well with simulated data as shown in figure 4.1. As the theoretical phase transition was derived assuming infinitely large system, this indicates $N = 256$ is already considered to be large enough in such sense. Comparing reconstruction probability for OMP in figure 4.2 with LASSO indicates OMP performs better for low δ 's. The theoretical phase transition for EMBP is seen to be slightly higher than what simulated data indicates in figure 4.3. One possible reason for this is the termination of the solver after maximum 1000 steps. EMBP is nonetheless seen to outperform both L1 and OMP.




Figure 4.1: Reconstruction probability for reconstruction by ℓ_1 minimization compared with theoretical phase transition.




Figure 4.2: Reconstruction probability for reconstruction by OMP. Theoretical phase transition is unavailable.




Figure 4.3: Reconstruction probability for reconstruction by EMBP compared with theoretical phase transition.

Block length	LFM	HFM
32	71	86
64	269	308
128	1052	1160
256	4149	4496
512	16485	17694
1024	65729	70185
2048	262534	279566
4096	1049362	1115908

Table 2: Number of basis vectors as function of block length and waveform.

Block length	LFM	HFM
32	0.465	0.477
64	0.382	0.387
128	0.305	0.307
256	0.238	0.239
512	0.182	0.183
1024	0.139	0.139
2048	0.105	0.105
4096	0.079	0.079

Table 3: Expected maximum correlation magnitude between random vector and the basis as a function of block length and waveform.

4.2 Random Vector Correlation Values

The number of basis vectors as function of block length and waveform for the passive case is tabulated in table 2. As seen, the number of basis vectors grows rapidly with increasing block length.

The tabulated values of the expected maximum correlation magnitude between a random vector and a given basis are given in table 3. The correlation value is seen to decrease with increasing block length, but since the basis size also grows with increasing block length, the decay is not as rapid as one could expect. Improved noise rejection capability follows with decreasing correlation value.

4.3 Basis Coverage

Passive detection of unknown interception signals require the basis to be sufficiently dense. The following three simulations ensures this is the case.

4.3.1 Coverage in (f, k) -space

We start by investigating the density in (f, k) -space. The steps taken are listed below as we run Monte Carlo simulations to verify the coverage. Both HFM and LFM will be considered and the resulting density plotted. Specifically, the following will be done for both HFM and LFM independently


1. A basis will be created using estimates from section 3.1.
2. Waveform with random parameters in (f, k) -space will be correlated with basis vectors.
3. The largest correlation coefficient will be recorded.
4. Go to step 2 and repeat until (f, k) -space is sufficiently covered.

If successful, this simulation aims to verify that treating f and k separately and using phase-matching are valid approximations. An important underlying assumption is that the pulse duration equals the block length used. Note that any combination of (f, k) values resulting in a final frequency outside $[0, \frac{f_s}{2}]$ will be removed from the construction of the basis.


Simulations results of coverage in (f, k) -space for LFM and HFM waveforms are shown in figure 4.4 and figure 4.5 respectively. Blue areas indicates combinations of (f, k) pairs that would result in a final frequency exceeding half the Nyquist frequency and thus is excluded. As indicated from the figures, the bases are shown to be sufficiently dense in the sense that any waveform in (f, k) -space is well correlated (typically less than 3 dB loss) with a basis vector. The high correlation factors indicate treating f and k separately and using phase-matching were reasonable approximations.

There is a peculiar break of symmetry for HFM at $f = 0.25f_s$. Despite thorough investigation, the origin behind this feature is still unknown. It has been hypothesized to be caused by aliasing, but no proof of this has been concluded. This feature will also be presented in subsequent plots, but only when HFM

is used as a basis, suggesting perhaps an error in the code. Double and triple checking the code and verifying the basis vectors with FFTs gave no conclusive answers. In addition, in figure 4.9b this feature take upon a slightly curved appearance.




(a) $N=128$




(b) $N=256$

Figure 4.4: Coverage of LFM basis in (f, k) -space. The blocklengths used are $N = 128$ and $N = 256$ respectively. Horizontal axis is in units of the sampling frequency.



(a) $N=128$



(b) $N=256$


Figure 4.5: Coverage of HFM basis in (f, k) -space. The blocklengths used are $N = 128$ and $N = 256$ respectively. Horizontal axis is in units of the sampling frequency.

4.3.2 Coverage in (f, f_1) -space


We then investigate the density in (f, f_1) -space. Recall the restriction $\frac{1}{2} < \frac{f}{f_1} < 2$. A question then naturally arises: how well does the basis correlate with waveforms which disobeys this restriction? We attempt to answer this question with the following procedure

1. A basis will be created using estimates from section 3.1.
2. Waveform with defined start and stop frequencies will be correlated with basis vectors.
3. The largest correlation coefficient will be recorded.
4. Go to step 2 and repeat until all start and stop frequencies are iterated over.

Simulation results of coverage in (f, f_1) -space for LFM and HFM waveforms are shown in figure 4.6 and figure 4.7 respectively. As indicated from the figures, the bases are shown to be sufficiently dense in the sense that any waveform in (f, f_1) -space is well correlated (typically less than 3 dB loss) with a basis vector when $\frac{1}{2} < \frac{f}{f_1} < 2$ is satisfied. The high correlation factors indicates treating f and k separately and using phase-matching were reasonable approximations. Outside the range $\frac{1}{2} < \frac{f}{f_1} < 2$ correlation is seen to be greatly reduced as anticipated. LFM is seen to correlate about as good as random noise (see table 3) outside this range, while HFM achieves a relatively high correlation and outperforms random noise. Implications are we might be able to detect HFM waveforms outside the defined range, but LFM waveforms will not.




(a) $N=128$




(b) $N=256$

Figure 4.6: Coverage of LFM basis in (f, f_1) -space. The blocklengths used are $N = 128$ and $N = 256$ respectively. Axes are in units of the sampling frequency.



(a) $N=128$



(b) $N=256$

Figure 4.7: Coverage of HFM basis in (f, f_1) -space. The blocklengths used are $N = 128$ and $N = 256$ respectively. Axes are in units of the sampling frequency.

4.3.3 HFM and LFM correlation


Although not covered by any previous theory, we will also investigate how well HFM and LFM correlates with each other.

1. An LFM basis will be created using estimates from section 3.1.
2. HFM waveform with defined start and stop frequencies will be correlated with basis vectors.
3. The largest correlation coefficient will be recorded.
4. Go to step 2 and repeat until all start and stop frequencies are iterated over.
5. Repeat this procedure with an HFM basis correlated with LFM waveforms.


We should expect high correlation between LFM and HFM waveforms. In particular, these waveforms coincide when $k = 0$ as they both reduce to CW. Deviations should thus appear as the initial and final frequency differ by a significant amount. As the bases are different, we should expect some variations depending on the basis used.

Simulation results of the LFM basis correlated with HFM waveforms is shown in figure 4.8. The corresponding plot for an HFM basis correlated with LFM waveforms is shown in figure 4.9. We expected LFM and HFM to correlate well when $k = 0$ and this is seen to be indeed the case as the largest deviations appear at the edges of the cone. Reduction in coverage is seen to increase as the block length increases. This is expected as longer correlation lengths will more efficiently filter out non-matching vectors. The effect of reduction near the edges is equivalent with tightening the bound on $\frac{f}{f_1}$. For sufficiently small block lengths we see only one basis is sufficient to capture both waveforms. Classification of pulse type can then be done by temporal means, e.g. tracking detections in a time-frequency plot.

Even outside the defined frequency range, HFM waveforms are seen to correlate notably better with an LFM basis than what random noise does (see table 3). The same cannot be stated for LFM waveforms correlated with an HFM basis where the gain is only marginally better.




(a) $N=128$




(b) $N=256$

Figure 4.8: Correlation between an LFM basis and HFM waveforms in (f, f_1) space. The blocklengths used are $N = 128$ and $N = 256$ respectively. Axes are in units of the sampling frequency.



(a) $N=128$



(b) $N=256$

Figure 4.9: Correlation between an HFM basis and LFM waveforms in (f, f_1) space. The blocklengths used are $N = 128$ and $N = 256$ respectively. Axes are in units of the sampling frequency.

4.4 Detector Simulations

The theory on compressive detection assumes the signal- and noise levels are known in advance. Here we investigate the detector performance by using estimates from section 3.3. The system performance hinges on the very important downsampling factor, yet not a word of it has been mentioned so far. This is about to change as we measure detector performance while tuning this parameter.

4.4.1 Probability of Detection

We begin by investigating the behaviour of probability of detection while the downsampling factor and noise levels are varied. The basis will consist of a single predefined waveform, and the signal is a combination of this waveform with additive Gaussian noise. Specifically, the following steps will take place

1. Construct the basis from one CW waveform.
2. Construct signal from the basis vector and additive Gaussian noise.
3. Create a random orthonormal measurement basis.
4. Obtain compressive measurements of basis and noisy signal.
5. Detection procedure as described in section 3.3.
6. Record if detection is successful or not.
7. Go to step 2 and repeat until sufficiently many trials are concluded.

Simulation results of the probability of detection as a function of SNR and downsampling ratio is shown in figure 4.10. Simulated results are compared with theoretical values from section 2.7.2. There is a relatively good comparison between simulated values and theoretical values, suggesting that the detection procedure performs well. However, some comments are in order.

There seem to be a slight tendency of over-detecting a signal when noise is present. This can be understood from tracing the estimation procedure. The noise will inherently force the estimation of signal level upwards since we attempt to fit with the best possible complex factor (over-fitting). As a result




Figure 4.10: Probability of detection as a function of signal-to-noise and downsampling ratio. A total of 10000 trials were performed for each data point. Theoretical values are shown for comparison using thin line width.

the noise level will be under-estimated and subsequently the probability of detection increases. This effect becomes more pronounced for low downsampling ratios as the noise will be less restrictive when fitting (a kind of uncertainty principle). This can be seen as the slight upshot to the left in the plot when noise is present. A direct consequence of the over-detecting is the increased probability of false alarms. When specifying the system this parameter should be adjusted accordingly.

The rapid decay in detection probability as function of downsampling ratio indicates the reference sensor should not be severely downsampled for practical applications. Some leeway are allowed however as we assume a high signal-to-noise ratio for incoming interception signals and that we are given several opportunities to detected the signal as it greatly exceeds the processing block length. Under these considerations a downsampling ratio around 0.2 – 0.3 may be acceptable. The performance implications follow immediately: memory consumption and computational effort reduced by the same factor for the detection step.

4.4.2 Classification Performance

The detector is also a classifier, attempting to differentiate between all the possible basis vectors. We then investigate the accuracy of classification while the downsampling factor and noise levels are varied. The basis will be created from section 3.1, and the signal consists of a random waveform (not necessarily in the basis, but of the same type) in additive Gaussian noise. The accuracy will be measured as root-mean-squared (see appendix C) distances in both f and k space between the strongest detection and actual waveform. The steps are as follows for HFM and LFM independently

1. Construct a basis following the rules in section 3.1.
2. Create a random waveform with additive uncorrelated noise.
3. Create a random orthonormal measurement basis.
4. Obtain compressive measurements of basis and noisy signal.
5. Detection procedure as described in section 3.3.
6. Choose the strongest detection.
7. Compute distances in (f, k) -space and record these values.
8. Go to step 2 and repeat until sufficiently many trials are concluded.

Take care to note that we set the detector performance parameter $\alpha = 1$ to ensure we always have a detection.


The obtained values of root-mean-squared classification error in (f, k) -space for LFM and HFM are shown in figure 4.11 and figure 4.12 respectively. There are three main contributions to the non-zero classification error at infinite SNR. Independent of the downsampling factor is the finite basis resolution. If this was the single source of error (and the only parameter) we would expect a frequency error around $\frac{f_s/2N}{\sqrt{3}} \approx 225$ Hz. The second effect involves the assumption of real and imaginary parts of complex basis vector contributing equally to the norm. Ensuring longer correlation lengths (that is, larger downsampling factor) should help reduce this error and this is indeed seen to be the case. Finally, using random measurement matrices may introduce large errors as elaborated in section 4.5.1.

At significant levels of noise, the simulated values for f error is found to border the theoretical RMS value, which lies just around $\frac{f_s/2}{\sqrt{6}} \approx 40$ kHz . Thus, little classification capabilities can be provided in a noisy environment. This is the manifestation that classification and parameter estimation is a difference and harder problem than detection.


Again we note the significant deterioration of detector performance for low downsampling factors. As such, the reference sensor should not be severely downsampled as mentioned in the previous section.

This is a two-parameter problem with high correlation between f and k and care must be taken when considering one parameter at time. To ease our understanding we assume for the moment we can treat them separately. From comparing the error in f between LFM and HFM we see that HFM performs slightly better than LFM. Since this is a two-parameter problem we should expect such outcome as LFM is more linearized than HFM and could possibly align better with multiple combinations of f and k than HFM.

Unlike f which is treated equally for both HFM and LFM, comparing the k plots with each other is difficult as they span different ranges. As such we comment individual features only. There is a curious behavior for high SNRs in k -space for LFM. The error is actually seen to increase for very low levels of noise as the number of compressive measurements taken increases. There is currently no good explanation for this behavior.




(a) f -space




(b) k -space

Figure 4.11: RMS classification error in (f, k) -space for LFM as a function of signal-to-noise and downsampling ratio. A total of 10000 trials were performed for each data point. Frequency axes are in units of the sampling frequency.



(a) f -space



(b) k -space

Figure 4.12: RMS classification error in (f, k) -space for HFM as a function of signal-to-noise and downsampling ratio. A total of 10000 trials were performed for each data point. Frequency axes are in units of the sampling frequency.

4.5 Direction-of-Arrival Simulations

There are several effects we wish to investigate with respect to bearing estimation. These include exceeding the design frequency and variance in bearing estimates.

4.5.1 Stretching the Design Frequency

We briefly noted that compressive sensing is apparently robust with regard to exceeding the design frequency and we wish to investigate this effect [51]. Knowledge from conventional beamforming states that the beam width decreases as the aperture size increases and thus the resolution increases. Exceeding the design frequency of the aperture, however, will introduce grating lobes and give rise to false targets. We compare this with compressive beamforming using the GMC method while the number of sensors is fixed and the aperture size increases. In detail, the following steps will be taken

1. Choose cylindrical sensor configuration.
2. Set sensor spacing.
3. Add narrowband target with random frequency and bearing.
4. Add noise on each sensor.
5. Obtain compressive measurements of pulse and noisy signal.
6. Obtain estimates using GMC method.
7. Compute performance indicators.
8. Go to step 2 and repeat until sufficiently many trials are concluded.

We note that it is physically impossible to avoid grating lobes and in particular uniform linear arrays are the most vulnerable. Any gain must thus appear from using a different sensor configuration. Conventional beamforming done on a cylindrical arrangement is also prone to grating lobes. As not all sensor will necessarily be in sync from the incoming wave, we expect the effect to be reduced (this we will denote as semi-grating lobes). We can expect compressive beamformer to perform even better by rejecting these semi-grating lobes in its quest for obtaining the sparsest solution of them all. Performance indicators for this simulation are

- RMS difference between strongest DOA estimate and actual bearing to verify the accuracy gain by increased aperture size.
- Measuring the average energy deposited in strongest DOA estimate versus other estimates to verify neglectable grating lobe levels.

The GMC method will be run until 90% of the data has been accounted for. By this procedure we intend to verify that compressive beamforming is robust and perform better when the design frequency is exceeded. See appendix D for details on how to add and subtract bearings in a consistent manner.

Simulation results on the root-mean-squared bearing error in degrees from exceeding the design frequency spacing is shown in figure 4.13. Two observations immediately comes to mind. First is the apparent minimum in bearing error at a sensor spacing of around 8 times the conventional spacing. This indicates we can only gain this much in estimation accuracy by increasing the aperture before the bearing estimate start to drift significantly, probably due to introduction of grating lobes or system instabilities. The gain in estimation accuracy from exceeding design frequency however is significant (about $\sqrt{8} \approx 3$).

The second observation is the level of estimation error. Under perfect conditions the estimation error should be bounded by grid resolution, in this case expected to equal $\frac{0.5^\circ}{\sqrt{3}} \approx 0.29^\circ$ on the average in absence of noise. After some investigation the main contribution to this discrepancy is found to stem from using random measurement matrices; most matrices perform well, but a few do not and significantly affect the final average estimation error. As the following analysis shows, the number of failing matrices can be truly neglectable, yet have a large impact: assume the root-mean-squared bearing estimation error for a failing matrix is $\frac{180^\circ}{\sqrt{3}}$ and for one that succeeds $\frac{0.5^\circ}{\sqrt{3}}$ (both under the assumption of uniform probability distribution, see appendix C). Then on the order of 1 out of 1000 matrices must fail to satisfy the requirements of compressive sensing in order to produce the results in figure 4.13. We note that increasing the number of compressive samples or choosing a matrix with better RIP constants is expected to improve this situation.

There is an apparent paradox with the minimum at a sensor spacing of 8 times the conventional spacing being achieved under noisy conditions. The exact mechanism behind this effect is currently under investigation.

In figure 4.14 we attempt to show the relation between the strongest bearing estimate and the additional (incorrect) estimates. There are several interesting features that requires a comment or two. First thing we note is that intro-




Figure 4.13: Average root-mean-squared error in degrees as the sensor spacing increases in units of design frequency spacing. SNR denotes the signal-to-noise ratio on each individual sensor. A total of 10000 trials were performed for each data point. Note the vertical axis is logarithmic.

duction of noise will increase the strength of false estimates. This is hardly a surprise however. Second thing to note is the clear tendency in increased strength of false estimates as the sensor spacing increases. This effect is the manifestation of grating lobes being introduced, not in the usual sense however. As noted earlier, these semi-grating lobes are likely to be rejected by compressive methods and should thus be considered more as increased noise in the system, which effect is to produce additional false estimates. As the graphs show, compressive beamforming is not immune to grating lobes, however the performance degrades gracefully and indicates robustness when exceeding the design frequency, especially for good signal-to-noise environments.

We then compare GMC with conventional beamforming under equal conditions in figure 4.15. The introduction of gating lobes for conventional beamformer is evident from the figure. As the frequency increases (the main lobe beam width becomes thinner) spurious peaks suddenly emerge out of nowhere, a certain sign of a grating lobe. Compressive beamforming is seen to behave quite differently. As the frequency increases, there is a tendency of the number of additional estimates and their strengths to increase. This is the quantitative effect behind figure 4.14. Note that a basis of LFM waveforms was used to obtain these results.






Figure 4.14: The average fraction of energy in additional estimates compared with the strongest estimate as the sensor spacing increases in units of design frequency spacing. SNR denotes the signal-to-noise ratio on each individual sensor. A total of 10000 trials were performed for each data point.

To summarize; exceeding the design frequency to increase the aperture size is seen to be advantageous in some situations, but in general care should be taken. In particular, a passive sonar is seen to benefit greatly from doing this, while an active sonar may struggle due to the sheer number of additional false estimates. This is not an issue, however, as one could simply utilize more sensors with fewer compressive measurements per sensor to increase the aperture size and still meet the requirements of Compressive Sensing.



(a) Conventional Beamforming



(b) GMC

Figure 4.15: Waterfall plot with a target located near 30° that pings with HFM which increases in frequency. Conventional beamformer in (a) shows the introduction of grating lobes as the frequency increases. GMC in (b) shows no sign of such behavior. The SNR is 10 dB on sensor level, sensor spacing is 8 times the conventional sensor spacing and GMC has been run until 90% of the energy is accounted for. Both plots are logarithmic.

4.5.2 Passive DOA Estimation Accuracy

Controlling the error in the bearing estimate is important. We investigate the relation between bearing estimation error and downsampling ratio under the influence of uncorrelated noise. This will form the basis for choosing an acceptable downsampling factor as more samples will generally improve the result. Direction-of-arrival estimates will be obtained with a known waveform to keep the detector performance out of the equation. Specifically, we do the following steps

1. Choose cylindrical sensor configuration.
2. Add narrowband target with random frequency and bearing.
3. Add noise on each sensor.
4. Create a random orthonormal measurement basis.
5. Obtain compressive measurements of pulse and noisy signal.
6. Obtain DOA estimate using GMC method.
7. Record squared difference between DOA estimate and true bearing.
8. Go to step 2 and repeat until sufficiently many trials are concluded.

We terminate the GMC method after one iteration, keeping only the strongest DOA estimate.

Simulation results are shown in figure 4.16. There is seen to be a general tendency for improved accuracy in the estimator as the number of compressive samples increases. This tendency is particularly evident when significant amount of noise is present. From (2.18) under the discussion of noise we see this behavior should indeed be the case. When less noise is present however, this tendency seems to be somewhat less pronounced, a phenomenon described in the following paragraph.

Assuming maximally sparsity in bearing space, the equation (S2) states a total of about 16 samples should be required in the noise-less case, which equates to 2 compressive samples per sensors in our case. This corresponds to the low end with $\delta \approx 0.04$ and $\rho \approx 0.06$. Comparing with figure 2.5 shows we are within the strong-phase transition region and should expect exact recovery. Increasing the

number of samples should thus provide no significant gain. Comparing with simulation results in figure 4.16 shows a significant improvement in estimation accuracy when going from 1 to 2 compressive samples per sensor as in accordance with theory. Going beyond this only results in small improvements in estimation accuracy.

By no means can we expect to get away with using only 16 compressive measurements in total. First and foremost, more samples would be required to handle noise. Second, the finite (and rather small) dimension of our system will necessarily result in a wide phase transition boundary (see e.g. [24]) and we can expect some matrices to fail to comply with the requirements of Compressive Sensing. This is hypothesized to be the main contribution to the rather large estimation error. Knowledge from the previous section showed that only a few matrices must fail in order to produce a large estimation errors.

We expect the total number of compressive measurements to be the decisive factor in satisfying the criteria of Compressive Sensing. Thus the number of sensors and measurements per sensor can be traded at will as long as the total number of compressive measurements is kept constant. The advantage of increasing the number of sensors is mirrored in improved DOA estimates that follows from a larger aperture size. The penalty is paid with a slight increase in additional computation being required when constructing the system matrix for the GMC method. The advantage greatly outweighs the penalty, however considerations from a hardware point of view, and perhaps economical as well, can make this decision more complex.

4.5.3 Active DOA Estimation Accuracy

We proposed using complex waveforms in GMC for active sonar as the phase is unknown. The following procedure will help with determine the efficiency of this solution.

1. Choose cylindrical sensor configuration.
2. Add narrowband target with random frequency and bearing.
3. Add noise on each sensor.
4. Create a random orthonormal measurement basis.
5. Obtain compressive measurements of pulse and noisy signal.




Figure 4.16: The root-mean-squared bearing error in estimation of DOA in the passive scenario as a function of number of compressive measurements per sensor. SNR denotes the signal-to-noise ratio on each individual sensor. A total of 10000 trials were performed for each data point. Note the vertical axis is logarithmic.




Figure 4.17: The root-mean-squared bearing error in estimation of DOA in the active scenario as a function of number of compressive measurements per sensor. SNR denotes the signal-to-noise ratio on each individual sensor. A total of 10000 trials were performed for each data point. Note the vertical axis is logarithmic.

6. Obtain DOA estimate using GMC method with complex waveform.
7. Record squared difference between DOA estimate and true bearing.
8. Go to step 2 and repeat until sufficiently many trials are concluded.

Again we terminate the GMC method after one iteration, keeping only the strongest DOA estimate.

Additional noise is introduced into the system matrix as we use a complex waveform when estimating direction-of-arrival. Initial understanding suggests this will be reflected in the accuracy of the estimator. Simulation results when using a complex waveform is shown in figure 4.17. This figure should be compared with figure 4.16 where the exact, and real, waveform is used for bearing estimation. The results are dire news indeed; even a high signal-to-noise ratio will require a substantial number of compressive measurements per sensor to obtain comparable accuracy in the estimator. Even more measurements would be required to handle multiple simultaneous incoming echoes. As a consequence an active aperture cannot be downsampled as drastically as in the passive case.

4.6 Passive Sonar Simulations

We construct a series of scenarios and display the performance of passive detection and localization in each scenario. As shown by earlier simulations it is sufficient to use a basis consisting of LFM waveforms only. The background noise will be normal distributed and set to 0 dB in all scenarios, broadband signatures will also be normal distributed and set to the appropriate level. Although unrealistic background noise and broadband noise is used, they have the advantage of being easy to work with. Several simplifications will be done, e.g. echoes and reverberation are not taken into account.

1. scenario

- Broadband target at -50° with strength 5 dB.
- Broadband target at 20° with strength 0 dB.
- Hidden pinger at 30° with strength 20 dB ($SNR \approx 10$ dB on sensor level). Emitted pulse is HFM starting at 2 kHz and ending at 4 kHz, lasting 5 seconds.

2. scenario

- Hidden pinger at 30° with strength 20 dB ($SNR = 20$ dB on sensor level). Emitted pulse is HFM starting at 25 kHz and ending at 30 kHz, lasting 300 ms.

The first scenario is intended to resemble a distant frigate pinging a very deep ocean. Simulation results are shown in figure 4.18. Despite the presence of two additional targets, these do not show up in the detection plot nor the DOA plot. From $T = 0$ s to $T = 5$ s only minor and random detections and DOA estimates with no structure are obtained. From the onset of the pulse at $T = 5$ s to its end at $T = 10$ s, a clear trace in the time-frequency plot can be seen. A corresponding trace in the DOA plot is also observed.


The DOA estimates are seen to be roughly centered around 30° as expected, however the variance in estimates is considerable. This stems from the relatively low pulse frequency combined with noise. Indeed, as the frequency of the pulse increases with time, the DOA estimates are seen to converge somewhat. Using the DOA estimates we are able to obtain a mean of 30.07° and a standard deviation of 6.0° . Compare this with figure 4.16.

The pulse type can barely be determined from the time-frequency plot: slight curvature indicates HFM. From the time-frequency plot, we are able to deduce the pulse start- and stop frequencies. These are found to be $f = 1.5$ kHz and $f_1 = 3.9$ kHz respectively. Despite the coarse resolution in frequency ($\Delta f = \frac{f_s}{N} \approx 800$ Hz), we are able to obtain reasonable values.


The second scenario is intended to resemble a torpedo homing in on its target. Simulation results are shown in figure 4.19. The frequency is into tens of kHz, and the pulse duration is relatively short. As in the first scenario, only minor and random detection and DOA estimates can be observed before the onset of the pulse. The pulse itself is clearly visible in the time-frequency detection plot and a clear trace is spotted in the corresponding DOA plot. As the frequency is higher and the noise lower than compared with the first scenario, the DOA estimates will also be correspondingly better localized around the mean.

The DOA estimates are seen to be roughly centered around 30° as expected, and the variance is small. As the pulse length is significant shorter and the increase in frequency is not as prominent as in the previous scenario (where frequency was doubled and the pulse length 50 times longer), it can be harder to observe the reduced variance with time. Using the DOA estimates we are able to obtain a mean of 30.40° and a standard deviation of 1.7° .

Determination of pulse type from the time-frequency plot is nearly impossible. However a slight curvature suggests it might be HFM. From the time-frequency plot, we are able to read of the pulse start- and stop frequencies. These are found to be $f = 25.0$ kHz and $f_1 = 29.7$ kHz respectively. As the frequency is relatively high, the coarse resolution in time-frequency plot does not affect these estimates considerably.




(a) Detections




(b) Direction-of-arrival estimates

Figure 4.18: Detection data in a time-frequency plot is shown in (a) and the DOA estimates in (b). Both plots are logarithmic, and the frequency axis in (a) is in units of the sampling frequency.



(a) Detections



(b) Direction-of-arrival estimates

Figure 4.19: Detection data in a time-frequency plot is shown in (a) and the DOA estimates in (b). Both plots are logarithmic, and the frequency axis in (a) is in units of the sampling frequency.

4.7 Active Sonar Simulations

We construct a series of scenarios and display the performance of active detection and localization in each scenario. As noted earlier a basis constructed from the transmitted waveform is sufficient. The background noise will be normal distributed and set to 0 dB in all scenarios. Although unrealistic background noise is used, it has the advantage of being easy to work with. Several simplifications will be done, e.g. complex issues such as reverberation and multipath echoes are nicely ignored or taken to be included into the background noise. From previous simulations we know the number of compressive measurements per sensor per block length must be increased, and we set this number to 30 for these simulation. We run GMC until 90% of the energy is counted for.


1. scenario

- Emitted pulse is HFM starting at 2 kHz and ending at 4 kHz, lasting 5 seconds.
- Target at -50° with strength 20 dB on sensor level. Delay = 0 s.
- Target at 100° with strength 20 dB on sensor level. Delay = 0 s.
- Target at 30° with strength 20 dB on sensor level. Delay = 1.7 s.


2. scenario

- Emitted pulse is HFM starting at 25 kHz and ending at 30 kHz, lasting 330 ms.
- Target at 30° with strength 20 dB on sensor level. Delay = 0.11 s.

The first scenario is intended to resemble a ping from a frigate, albeit the receiving sensor aperture is unlikely to be used in this case. Two targets are very close and equally distanced from the pinger, while the third target is slightly further away. Simulation results are shown in figure 4.20. The basis is seen to consist of very few vectors, in fact only three. Despite very strong returns (20 dB on sensor level from each target), acquiring reasonable DOA estimates is near impossible. The result of poor estimation accuracy originates from many sources: low ping frequency range, significant noise level per target, complex waveform in GMC method and still too few compressive measurements. The two first echoes which occurs simultaneously cannot be differentiated among the DOA estimates. The third echo, which has a delayed onset, only becomes visible when the other two echoes have faded.



(a) Detections




(b) Direction-of-arrival estimates


Figure 4.20: Detection data in a time-frequency plot is shown in (a) and the DOA estimates in (b). Both plots are logarithmic.

The second scenario is intended to resemble a torpedo homing in on its target and is more realistic than the previous scenario when considering the receiving aperture. Simulation results are shown in figure 4.21. The frequency is into tens of kHz, and the pulse duration is relatively short. With only one target and high frequency range we are able to detect and localize the target presence based on return echo. Using the DOA estimates we are able to obtain a mean of 29.74° and a standard deviation of 2.4° . This is twice the standard deviation in the corresponding passive scenario, yet the number of compressive measurements is three times larger.

There is an very interesting feature at -150° , which has been amplified by the logarithmic nature of the plot. This feature was also present in figure 4.20, albeit not that strongly. This 'shadow target' originates from a combination of the chosen aperture and waveform and always appears 180° shifted compared with the real target as the array configuration has some trouble with discriminating the 'left-right' direction of an incoming plane wave. In practice this effect is minimal; using more sensors, applying directivity on the sensors (which is always the case in the real world), and using sensors within $\sim \pm 60^\circ$ when estimating DOA or beamforming will all subdue this problem. We note that using more compressive measurements also reduces the presence of this 'shadow target'.



(a) Detections



(b) Direction-of-arrival estimates

Figure 4.21: Detection data in a time-frequency plot is shown in (a) and the DOA estimates in (b). Both plots are logarithmic.





Figure 4.22: Spectrogram of a real LFM pulse showing the linear increase in frequency from 3 kHz up to 8 kHz.

4.8 Active and Passive Performance Using Real Data


We then we apply real sonar data to verify the performance in practice. The receiving array is a linear uniform array consisting of 16 elements spaced 6.5 cm apart. Each element is sampled at 16 kHz. To enable the framework of Compressive Sensing, each sensor is downsampled by using a random orthonormal basis to produce 10 compressive measurements per block length (approximately 96% data reduction). One of the sensors is used as reference sensor and downsampled by a factor of 4.

The sonar data was recorded at Langstein in Stjørdal in 2010. The pinger is unfortunately positioned immediately in front of the array and we will not be able to discern its DOA. However, we should be able to pinpoint some of the possible echoes. The LFM ping sweeps from 3 kHz to 8 kHz during an interval of 0.3 s as seen in figure 4.22.

Output from conventional beamformer and compressive beamformer using the narrowband model are shown for comparison in figure 4.23. There seem to be a trace at 90° , most likely the pinger. Several prominent echoes can be spotted immediately after the ping and up to four seconds after the ping (corresponds to a range of approximately 3 km).



(a) Conventional beamformer




(b) Compressive beamformer

Figure 4.23: Output from conventional beamformer and compressive beamformer using real data. Data has been scaled and plotted in a logarithmic manner (dB) to empathize the dynamics.


For the active sonar usage, a basis was constructed using the LFM pulse, resulting in 80 basis vectors. The results from the detection procedure with the following DOA estimation using complex waveform are shown in figure 4.24. The dots to the left and right in the detection plot indicates possible echoes, discovered by observing detections with the same slope in time as the pulse itself has (red dots). Many more echoes like this are likely to exist; we only pointed out the most prominent ones. Comparing the positions of the dots with features in figure 4.23 indicates the detection performance is satisfactory. The corresponding DOA estimates are seen to somewhat resemble features figure 4.23, however not to a great satisfactory extent.

Despite this case is more suitable for active sonar, we attempt to compare with passive sonar procedure. A basis consisting of LFM pulses was created, resulting in 4152 basis vectors (see table 2). The results from the detection procedure with the following DOA estimation are shown in figure 4.25. The pulse itself is directly observable in the detection plot. From this we are able to deduce the pulse shape to be LFM, starting at around 3 kHz and ending around 8 kHz during roughly 0.3 seconds. After this we are unable to detect any visible echoes. However, there are some patches of high-frequency detections immediately after the pulse. These are also found in the detection plot for active sonar. As time progresses beyond 2 s, detections of low frequency components dominate the picture. Curiously enough, there seem to be traces of weak CW lines at approximately 0.5 kHz and 1 kHz of unknown origin.

Looking at the DOA estimate plot we are able to correlate several features with respect to the detection plot and figure 4.23. As the pinger is in near-field of the receiving aperture, we are unable to disclose its direction of arrival. This is seen as incoherent pattern during the pulse duration. A strong patch of high-frequency detections after 1 s is seen to correlate well with a target normal to the array at a range of approximately 1 km. This feature is also noticeable in the other plots. Subsequent DOA estimates are seen to be mainly centered around 90° , assumed to be due to the pinger vessel, or near the left side of the array (around 10°) which could be due to wharf activity or nearby road traffic.




(a) Detections




(b) Direction-of-arrival estimates

Figure 4.24: Detections using an active basis is shown in (a) and the corresponding DOA estimates in (b). Data has been scaled and plotted in logarithmic manner (dB) to empathize the dynamics.



(a) Detections



(b) Direction-of-arrival estimates

Figure 4.25: Detections using a passive basis is shown in (a) and the corresponding DOA estimates in (b). Data has been scaled and plotted in logarithmic manner (dB) to empathize the dynamics.

5 Future work

Several ideas came to mind, yet have so far been unchallenged. These involve investigating the following

- Correlation between LFM and HFM bases with aim to reduce basis size.
- Improving the detection procedure to produce better estimates.
- Model-based Compressive Sensing for enhanced detection and localization.
- Using small blocks for detections and combine to a longer block for DOA estimation.

6 Concluding remarks

Simulation results on basis coverage indicates the bases are sufficiently dense using the proposed quantification. For small block lengths, HFM and LFM bases are seen to correlate well and thus only one basis would be needed to capture both waveforms for passive applications. Active transmissions can utilize a significantly smaller basis size based on the transmitted ping.

The detection step can be computational intensive with large basis size. In essence it equates to the first step in a greedy recovery algorithm. To improve the situation we proposed to subsample the reference sensor. Simulations showed limited freedom in choosing the downsampling ratio as the detector performance quickly deteriorates. Estimation of signal strength and noise variance is seen to be slightly biased with the proposed method, leading to a slight over-detection which might need to be compensated for in real applications.

Passive direction-of-arrival estimation accuracy was verified to improve and be robust when sensor spacing is increased somewhat beyond the conventional spacing. Simulations showed what is conventionally perceived as grating lobes could be viewed as an increased level of noise in compressive systems. Passive applications allowed substantial downsampling of sensors; typically a factor of tenth to a hundredth as compared with conventional sampling schemes, depending on the robustness to noise.

To handle the event of multiple simultaneously incoming echoes we proposed to use the complex waveform in GMC method. This is particularly required for active direction-of-arrival estimation. Using complex waveform is seen and understood to correspond with additional noise, and consequently more compressive measurements are required as opposed to the passive case.

Compressive Sensing is an active field of research and a lot of effort is put in to improve the reconstruction efficiency. In the future we should expect better algorithms in leveraging the phase-transition and improving the reconstruction time. Greedy algorithms are particularly simple and efficient and little improvement can be expected for this class of algorithms. As such the considerations presented in this paper should be valid for the foreseen future.

Considering these aspects, Compressive Sensing stands as a strong alternative to traditional methods for detecting and localization of high-frequency interception signals, in particular for passive applications.

A Waveforms

Here we list the the instantaneous frequency evolution $f(t)$ and the resulting time-domain function for the phase $\phi(t)$ for LFM, HFM and CW respectively. Additional notes about performance is given.

The simplest waveform of them all is the Continuous Wave (CW) waveform.

$$\begin{aligned} f(t) &= f_0, \quad 0 \leq t \leq T \\ \phi(t) &= 2\pi f_0 t, \quad 0 \leq t \leq T \end{aligned}$$

The CW pulse is optimal in the sense of detecting Doppler shifts of radially moving targets, but possesses no pulse compression capabilities and poor range resolution. A waveform better suited for detection is the Linear Frequency Modulation (LFM). As the name implies, the LFM has an instantaneous frequency that increase linearly with time

$$\begin{aligned} f(t) &= f_0 + \frac{f_1 - f_0}{T}t, \quad 0 \leq t \leq T \\ \phi(t) &= 2\pi \left(f_0 t + \frac{1}{2} \frac{f_1 - f_0}{T} t^2 \right), \quad 0 \leq t \leq T \end{aligned}$$

where f_0 is the start frequency and f_1 the end frequency during the pulse transmission time T . The nature of LFM makes it excellent for detecting targets, however the echo from radially moving targets will not line up correctly with the matched form due to the Doppler shift. This is overcome by using the HFM waveform which is insensitive to Doppler shifts. HFM can be defined by

$$\begin{aligned} f(t) &= \frac{f_0}{1 - \frac{f_0}{f_1} \frac{f_1 - f_0}{T} t}, \quad 0 \leq t \leq T \\ \phi(t) &= -2\pi \frac{f_0 f_1 T}{f_1 - f_0} \log \left(1 - \frac{f_0}{f_1} \frac{f_1 - f_0}{T} t \right), \quad 0 \leq t \leq T \end{aligned}$$

While the HFM excels in detecting targets, it fails to provide any information about radial target movement.

Finally, these three waveforms are all related by observing the series expansions of a hyperbola

$$\frac{f_0}{1 - kt} = f_0 + f_0 kt + \dots$$

B Parametrization of the Phase Transition

The weak phase boundary curve admits the parametric expression [40]

$$\delta = \frac{2\phi(\tau)}{\tau + 2\phi(\tau) - 2\tau\Phi(-\tau)}$$

$$\rho = 1 - \frac{\tau\Phi(-\tau)}{\phi(\tau)}$$

for $\tau \geq 0$ and ϕ and Φ are the Gaussian and its integral

$$\phi(\tau) = \frac{1}{\sqrt{2\pi}} e^{-\tau^2/2}$$

$$\Phi(\tau) = \int_{-\infty}^{\tau} \phi(z) dz$$

C Computing RMS Values

The root-mean-squared is, as the name implies, the square root of the mean of the squares. In mathematical terms this is equivalent to

$$RMS = \sqrt{\frac{1}{N} \sum_{n=1}^N x_n^2}$$

where x_n is the n -th realization of the random variable X . In this form it can also be perceived as a biased estimator to the standard deviation of a zero-mean random variable. As such RMS is often frequently encountered, particularly in relation with simulations and measurements.

Knowing the underlying probability distribution $f_X(x)$ of X enables direction computation of the RMS once the distribution of $Y = X^2$ has been worked out.

$$RMS = \sqrt{\int y \cdot f_Y(y) dy} \tag{C.1}$$

From probability theory we have the following relation between the cumulative probability distributions $F_X(x)$ and $F_Y(y)$.

$$\begin{aligned} F_Y(y) &= P(X^2 \leq y) \\ &= P(|X| \leq \sqrt{y}) \\ &= P(-\sqrt{y} \leq X \leq \sqrt{y}) \\ &= F_X(\sqrt{y}) - F_X(-\sqrt{y}) \end{aligned}$$

Differentiating with respect to y results in the probability distribution

$$f_Y(y) = f_X(\sqrt{y})\frac{1}{2\sqrt{y}} + f_X(-\sqrt{y})\frac{1}{2\sqrt{y}}, \quad y > 0 \quad (\text{C.2})$$

A couple of worked examples then follows. Let $X \sim \mathcal{U}(0, a)$, i.e. drawn from a uniform probability distribution on $[0, a]$, and $Y = X^2$. By definition

$$f_X(x) = \begin{cases} \frac{1}{a} & 0 \leq x \leq a \\ 0 & \text{otherwise} \end{cases}$$

The relation (C.2) can be used immediately to result in

$$f_Y(y) = \begin{cases} \frac{1}{2a\sqrt{y}} & 0 \leq y \leq a^2 \\ 0 & \text{otherwise} \end{cases}$$

and the RMS follows trivially from (C.1)

$$RMS = \frac{a}{\sqrt{3}}$$

The next example considers the difference between two random variables drawn from the same distribution, i.e. $X = X_1 - X_2$ and $Y = (X_1 - X_2)^2$ with $X_1, X_2 \sim \mathcal{U}(0, a)$. Simple considerations shows X has a triangular distribution

$$f_X(x) = \begin{cases} \frac{a+x}{a^2} & -a \leq x \leq 0 \\ \frac{a-x}{a^2} & 0 \leq x \leq a \\ 0 & \text{otherwise} \end{cases}$$

Using (C.2) immediately results in

$$f_Y(y) = \begin{cases} \frac{1}{a\sqrt{y}} - \frac{1}{a^2} & 0 \leq y \leq a^2 \\ 0 & \text{otherwise} \end{cases}$$

and the RMS value is easily obtained from (C.1)

$$RMS = \frac{a}{\sqrt{6}}$$

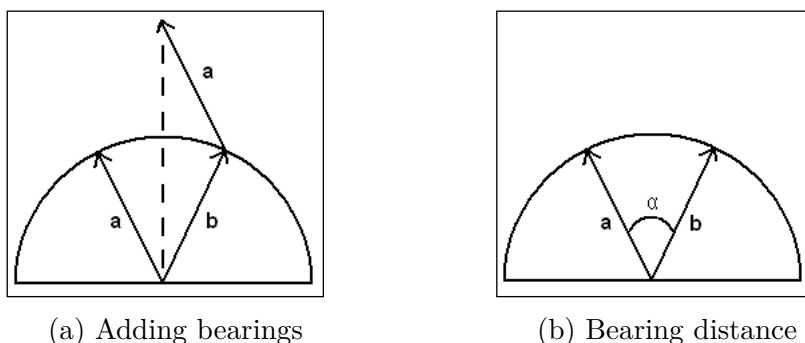


Figure D.1: Averaging bearings can be viewed as averaging of vectors (a), and distances in bearing can be understood as finding the angle between two vectors (b).

D Metrics in Bearing Space

Measuring distances and averaging in bearing space can be prone to errors if not done in a consistent manner. The proposed method involves replacing the bearing with a unit vector pointing in the correct direction. Averaging is done by averaging such vectors and then convert back to bearing space to obtain an average bearing. Measuring distances is also trivial and now corresponds with computing the inner product and then finding the angle, see figure D.1.

References

- [1] E. J. Candès and M.B. Walkin, “An Introduction to Compressive Sensing,” *IEEE Signal Processing Magazine*, **21**, March 2008.
- [2] E. J. Candès and J. Romberg, “Sparsity and Incoherence in Compressive Sampling,” *Inverse Problems*, **23** 969-985, 2006.
- [3] E. J. Candès and T. Tao, “Near-Optimal Signal Recovery From Random Projections: Universal Encoding Strategies?,” *IEEE Transactions on Information Theory*, **52**(12), December 2006.
- [4] E. J. Candès, J. Romberg and T.Tao, “Robust Uncertainty Principles: Exact Signal Reconstruction From Highly Incomplete Frequency Information,” *IEEE Transactions on Information Theory*, **52**(2), February 2006.

- [5] E. J. Candès, J. Romberg and T. Tao, “Stable Signal Recovery from Incomplete and Inaccurate Measurements,” *Communications on Pure and Applied Mathematics*, **LIX** 1207-1223, 2006.
- [6] D. L. Donoho, “Compressed Sensing,” *IEEE Signal Processing Magazine*, **52**(4) 1289-1306, 2006.
- [7] D. L. Donoho and M. Elad, “Optimally sparse representation in general (nonorthogonal) dictionaries via L_1 minimization,” *Proceedings of the National Academy of Sciences*, **100**(5) 2197-2202, 4th of March, 2003.
- [8] D. L. Donoho, “For Most Large Underdetermined Systems of Linear Equations the Minimal L_1 -norm Solution is also the Sparsest Solution,” *Communications on Pure and Applied Mathematics*, **59**(6) 797-829, 2006.
- [9] D. L. Donoho, “For Most Large Underdetermined Systems of Equations the Minimal L_1 -norm Near-Solution Approximates the Sparsest Near-Solution,” *Communications on Pure and Applied Mathematics*, **59**(7) 907-934, July 2006.
- [10] S. S. Chen, D. L. Donoho and M. A. Saunders, “Atomic Decomposition by Basis Pursuit,” *Society for Industrial and Applied Mathematics*, **43**(1) 129-159, 2001.
- [11] H. Nyquist, “Certain topics in telegraph transmission theory,” *Transactions of the A.I.E.E.*, 617-644, February 1928.
- [12] C. E. Shannon, “Communication in the presence of noise,” *Proc. Institute of Radio Engineers*, **37** 10-21, January 1949.
- [13] G. B. Dantzig, “Maximization of a linear function of variables subject to linear inequalities,” 1947. Published in T. C. Koopmans “Analysis of Production and Allocation,” *Wiley & Chapman-Hall*, 339-347, 1951.
- [14] S. Boyd and L. Vandenberghe, “Convex Optimization,” *Cambridge University Press*, 512-621, 2004.
- [15] “Occam’s razor,” *Wikipedia the free encyclopedia*, (cited 18-09-2012). Available online from http://en.wikipedia.org/wiki/Occam's_razor
- [16] “Norm (mathematics),” *Wikipedia the free encyclopedia*, (cited 18-09-2012). Available online from [http://en.wikipedia.org/wiki/Norm_\(mathematics\)](http://en.wikipedia.org/wiki/Norm_(mathematics))

- [17] E. Arias-Castro, E.J. Candès and M. A. Davenport, “On the Fundamental Limits of Adaptive Sensing,” *arXiv:1111.4646v4* [**math.ST**], 2012.
- [18] E. J. Candès and T. Tao, “Decoding by Linear Programming,” *IEEE Transactions on Information Theory*, **51**(12), December 2005.
- [19] J. Romberg and M. B. Wakin, “Compressed Sensing: A Tutorial,” *IEEE Statistical Signal Processing Workshop*, August 2007.
Available online from
<http://users.ece.gatech.edu/~justin/ssp2007>
- [20] A. Cohen, W. Dahmen, R. DeVore, “Compressed sensing and best k -term approximation,” *Journal of the American Mathematical Society*, **22** 211-231, 2009.
- [21] J. D. Blanchard, C. Cartis, J. Tanner, “Compressed Sensing: How Sharp is the Restricted Isometry Property,” *arXiv:1004.5026* [**cs.IT**], April 2010.
- [22] D. L. Donoho, J. Tanner, “Observed Universality of Phase Transitions in High-dimensional Geometry, with Implications for Modern Data Analysis and Signal Processing,” *arXiv:0906.2530* [**math.ST**], June 2009.
- [23] M. Bayati, M. Lelarge and A. Montanari, “Universality in Polytope Phase Transitions and Message Passing Algorithms,” *arXiv:1207.7321* [**math.PR**], July 2012.
- [24] D. L. Donoho, J. Tanner, “Precise Undersampling Theorems,” *Proceedings of the IEEE*, **98**(6), June 2010.
- [25] D. Donoho and J. Tanner, “Counting Faces of Randomly-Projected Polytopes when the Projection Radically Lowers Dimension,” *arXiv:math/0607364v2* [**math.MG**], September 2006.
- [26] J. Tanner, “Regular Polytopes and Cone,” (cited 27-09-2012). Available online from
rcos.maths.ed.ac.uk/polytopes.shtml
- [27] J. D. Blanchard, C. Cartis, J. Tanner and A. Thompson, “Phase Transitions for Greedy Sparse Approximation Algorithms,” *arXiv:1004.1821v1* [**cs.IT**], April 2010.

- [28] F. Krzakala, M. Mézard, F. Sausset, Y. Sun and L. Zdeborová, “Probabilistic Reconstruction in Compressed Sensing: Algorithms, Phase Diagrams, and Threshold Achieving Matrices,” *arXiv:1206.3953v1* [**cond-mat.stat-mech**], June 2012.
- [29] J. Barbier, F. Krzakala, M. Mézard and L. Zdeborová, “Compressed Sensing of Approximately-Sparse Signals: Phase Transitions and Optimal Reconstruction,” *arXiv:1207.2079v1* [**cs.IT**], July 2012.
- [30] F. Krzakala, “ASPICS: Applying Statistical Physics to Inference in Compressed Sensing,” (cited 30-10-2012). Available online from www.pct.espci.fr/~florent/ASPICS/ASPICS.html
- [31] D. Donoho, A. Javanmard and A. Montanari, “Information-Theoretically Optimal Compressed Sensing via Spatial Coupling and Approximate Message Passing,” *arXiv:1112.0708* [**cs.IT**], December 2011.
- [32] D. Donoho, A. Maleki and A. Montanari, “Message Passing Algorithms for Compressed Sensing,” *Proceedings of National Academy of Sciences*, **106** 18914-18919, 2009.
- [33] M. Bayati and A. Montanari, “The LASSO risk for gaussian matrices,” *arXiv:1008.2581v2* [**math.ST**], August 2012.
- [34] J. Vila and P. Schniter, “Expectation-Maximization Bernoulli-Gaussian Approximate Message Passing,” *Proceedings Asilomar Conference on Signals, Systems, and Computers*, November 2011.
- [35] R. Baraniuk, M. Davenport, R. DeVore and M. Wakin, “A Simple Proof of the Restricted Isometry Property for Random Matrices,” *Constructive Approximation*, February 2007.
- [36] H. Xu, C. Caramanis and S. Mannor, “Sparse Algorithms Are Not Stable: A No-Free-Lunch Theorem,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **34**(1) 187-193, January 2012.
- [37] E. Arias-Castro and Y. C. Eldar, “Noise Folding in Compressed Sensing,” *IEEE Signal Processing Letters*, **18**(8) 478-481, August 2011.
- [38] E. J. Candés and M. A. Davenport, “How well can we estimate a sparse vector?,” *arXiv:1104.5246v1* [**cs.IT**], April 2011.
- [39] M. J. Wainwright, “Sharp Thresholds for High-Dimensional and Noisy Sparsity Recovery Using ℓ_1 -Constrained Quadratic Programming (Lasso),” *IEEE Transactions on Information Theory*, **55**(5), May 2009.

- [40] D. L. Donoho, A. Maleki and A. Montanari, “The Noise-Sensitivity Phase Transition in Compressed Sensing,” *arXiv:1004.1218v1 [math.ST]*, April 2010.
- [41] M. A. Davenport, P. T. Boufounos, M. B. Wakin and R. G. Baraniuk, “Signal Processing With Compressive Measurements,” *IEEE Journal of Selected Topics in Signal Processing*, **4**(2) 445-460, April 2010.
- [42] A. Eftekhari, J. Romberg and M. B. Wakin, “Matched Filtering from Limited Frequency Samples,” *arXiv:1101.2713v2 [cs.IT]*, July 2012.
- [43] A. Eftekhari, J. Romberg and M. B. Wakin, “A Probabilistic Analysis of the Compressive Matched Filter,” *9th International Conference on Sampling Theory and Applications (SampTA 2011)*, Singapore, May 2011.
- [44] J. Haupt et al, “Compressive Sampling for Signal Classification,” *40th Asilomar Conference on Signals, Systems and Computers*, 1430-1434, November 2006.
- [45] J. Haupt and R. Nowak, “Compressive Sampling for Signal Detection,” *IEEE International Conference on Acoustics, Speech and Signal Processing*, III-1509 - III-1512, April 2007.
- [46] I. Bilik, “Spatial Compressive Sensing Approach For Field Directionality Estimation,” *IEEE Radar Conference*, 4-8th of May, 2009.
- [47] A. Gretsistas and M. D. Plumbley, “A Multichannel Spatial Compressed Sensing Approach for Direction of Arrival Estimation,” *Proceedings of the 9th international conference on Latent variable analysis and signal separation, Springer*, 458-465, 2010.
- [48] X. Ren, Q. Wan and H. Chen, “A Fast Method for DOA Estimation of Coherent Signal from Sparse Signal Reconstruction Perspective,” *7th International Conference on Wireless Communications, Networking and Mobile computing, IEEE*, 23-25th September 2011.
- [49] D. Malioutov, M. Çetin and A. S. Willsky, “A Sparse Signal Reconstruction Perspective for Source Localization With Sensor Arrays,” *IEEE Transactions on Signal Processing*, **53**(8), August 2005.
- [50] A. C. Gürbüz, J. H. McClellan and V. Cehver, “A Compressive Beamforming Method,” *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2008.

- [51] G. F. Edelmann and C. F. Gaumont, “Beamforming using compressive sensing,” *Journal of the Acoustic Society of America*, **130**(4), October 2011.