UNIVERSITY OF OKLAHOMA

GRADUATE COLLEGE

SOFTWARE-DEFINED COHERENT-ON-RECEIVE PROCESSING FOR MAGNETRON-BASED COMMERCIAL MARINE RADAR

A THESIS

SUBMITTED TO THE GRADUATE FACULTY

in partial fulfillment of the requirements for the

Degree of

MASTER OF SCIENCE

By

Daniel Jensen
Norman, Oklahoma
2026

SOFTWARE-DEFINED COHERENT-ON-RECEIVE PROCESSING FOR MAGNETRON-BASED COMMERCIAL MARINE RADAR

A THESIS APPROVED FOR THE
SCHOOL OF ELECTRICAL AND COMPUTER ENGINEERING

BY THE COMMITTEE CONSISTING OF

Dr. Justin Metcalf, Chair

Dr. Jessica Ruyle

Dr. Nathan Goodman

Acknowledgments

First and foremost, I am incredibly thankful to my parents, Matthew and Carolina, whose unwavering encouragement and belief in me made this journey possible. Their example has been a constant source of motivation. I would like to extend my sincere appreciation to my committee members, Dr. Nathan Goodman and Dr. Jessica Ruyle, for their time, insight, and valuable feedback. Their expertise and thoughtful perspectives have significantly strengthened this work. I am especially grateful to my committee chair, advisor, and mentor, Dr. Justin Metcalf, for his continuous support and dedication throughout this process. His guidance has been instrumental not only in this research but also in my growth as an engineer and researcher. I would also like to thank my coworkers, Brian Stade and Ahmed Ahmed, for their help with all the experiments included in this work. Their efforts and camaraderie made this experience more rewarding. Finally, I would like to acknowledge everyone else who, in ways both big and small, contributed to my success throughout this journey. Your support has not gone unnoticed and is deeply appreciated.

This material is based upon work supported by the U.S. Geological Survey under Grant/Cooperative Agreement No. G23AC00569 and the Department of Energy Integrated Energy Systems Office through the Pacific Northwest National Laboratory under Standard Research Subcontract No. 814000. Portions of this work were included in [1], to be published in the 2026 IEEE Radar Conference.

Table of Contents

List of Tables

List of Figures

Abstract

Commercial navigational marine radar systems are attractive for maritime and oceanographic use due to their high power and low cost. However, their magnetron-based transmitters generate random phase variations, and their receive subsystems are susceptible to intermediate frequency drift. This thesis develops a software-based coherent-on-receive signal processing framework that enables coherent Doppler processing on unmodified magnetron-based radar platforms. To resolve the random pulse-to-pulse phase variations, the study proposes two distinct phase correction algorithms: a Gradient Ascent optimization and a computationally efficient Phase Difference Averaging technique. The proposed algorithms are evaluated in comparison with an existing technique rooted in matched filtering principles. Furthermore, the research mitigates non-stationary intermediate frequency by utilizing the proposed power spectral density and template matching estimation techniques. The efficacy of the framework is rigorously tested through experimental data collections in radar environments, including stationary ground targets, dynamic platform motion, and true sea clutter. Experimental results confirm that signal processing algorithms can extract coherent Doppler information on moving targets, including unmanned aerial vehicles, helicopters, thrown projectiles, and birds, even when obscured by dominant environmental clutter. Ultimately, this work develops and validates advanced adaptive signal processing techniques that overcome the inherent hardware limitations of affordable marine

radars, rendering them viable for advanced detection applications.

Chapter 1

Introduction

1.1 Overview

Radar is an electromagnetic remote sensing technology used to detect, track, and image targets or areas of interest within complex environments. Similar to other electromagnetic sensing modalities, such as infrared, LiDAR, and even the human visual system, radar operates by transmitting electromagnetic (EM) waves and analyzing the signals reflected from objects within a scene. In most radar systems, this energy is radiated into free space via a directive antenna, which also enables estimation of the spatial location of objects relative to the radar. By precisely measuring the propagation time of the transmitted signal, the system can determine the distance, or range, to a target. Advancements in radar technology have enabled precise control of signal frequency and phase, paving the way for coherent processing techniques that enable a system to estimate the velocity of objects in a phase-coherent system. Radar systems are capable of reliable operation in conditions that degrade other sensing modalities, including darkness, fog, rain, and snow. Marine radar systems, in particular, are typically shipborne and are used to detect and monitor objects such as vessels, landmasses, buoys, precipitation, and sea ice. These systems are essential for safe maritime navigation, especially in low-visibility conditions, and are regulated under international law.

In particular, commercial off-the-shelf (COTS) marine radar units are widely deployed in demanding operational environments due to their cost-effectiveness and high peak power output. Most COTS marine radar systems rely on a magnetron as their radio-frequency (RF) power source, which is a cheap, compact, and power-efficient device capable of generating radio frequency (RF) energy. However, unlike other high-power RF sources used in radar systems, the magnetron operates as a power oscillator rather than a power amplifier, which does not natively allow for amplitude or phase modulation of the transmitted signal [2]. Consequently, each transmitted pulse exhibits a random initial phase, rendering the system inherently non-coherent. This lack of coherence prevents the direct application of coherent integration techniques and limits the radar's ability to perform Doppler processing for velocity estimation [3]. Doppler processing and other advanced signal processing techniques increase the operational capability of a radar system dramatically, so enabling coherent capabilities in COTS systems is highly desirable.[1]

Early approaches to attaining coherence in a magnetron-based system focused on injecting a low-power coherent reference signal into the magnetron to control the output phase [4], [5]. These approaches achieved pulse-to-pulse coherence over bandwidths close to 200 MHz by injecting a low-noise oscillator signal strictly timed to coincide with magnetron startup. The resulting phase-locked magnetron emission exhibited phase stability sufficient for coherent pulse-Doppler processing while maintaining the magnetron's high output power [4], [5]. However, for modern off-the-shelf marine radars, the transmitter is often sealed, and the subsystems are proprietary, making such modifications impractical. In addition, the

[1]Disclaimer: portions of this section were included in and have been accepted for publication in [1]

injection-locking system requires a stable reference oscillator and RF components capable of handling the magnetron's high output power, thereby increasing the system's complexity and cost.

Coherent-on-receive (COR) processing can offer a more practical approach to achieving coherence in magnetron-based systems without the same expensive and invasive hardware modifications. In COR systems, the transmitter operates with random phase variations, which the signal processing, not the hardware, estimates and compensates for. One approach sampled the transmitted signal directly, enabling synchronous analog-to-digital converters to digitize the transmit signal and the received echoes for COR processing [6]. In [7], a high-power marine radar is modified with a tunable synthesizer and high- and low-gain channels to enable COR processing. In both examples, a digitized transmit pulse model is created and used to account for the magnetron-induced phase offset. Implementations based on this principle have demonstrated the capability for Doppler processing and moving-target detection by sampling and processing transmit and receive signals with appropriate phase corrections. While not explicit on the specific processing and hardware modifications, others have also used a modified X-band marine radar for coherent signal processing [8], [9], [10].

A key distinction of this work from existing studies is the complete absence of subsystem hardware modifications, which introduces previously unknown experimentally observed obstacles. Beyond the random initial phase of the magnetron, a critical secondary issue is the presence of intermediate-frequency (IF) instabilities within the receiver chain [11]. Although the precise source of this time-varying frequency drift is difficult to identify in proprietary systems, both literature and experimental observations suggest that it arises from instability in local oscillators [12]. Unlike the constant phase offset introduced by the magnetron, IF drift induces time-varying phase modulation within individual pulses, necessitating additional compensation strategies. Effective mitigation of both the random transmit phase and IF drift with adaptive signal processing techniques is essential for realizing a practical coherent-on-receive marine radar system.

This thesis develops and evaluates a signal processing framework to enable coherent-on-receive operation in unmodified COTS marine radar systems. The work includes experimental characterization of the radar platform and the development of techniques to address both magnetron-induced phase offsets and time-varying IF drift. The proposed methods are evaluated using multiple coherence metrics and validated through extensive experiments involving both stationary and moving targets in ground and maritime clutter environments. By demonstrating reliable coherence restoration and enabling meaningful Doppler measurements, this work establishes a practical pathway for extending the capabilities of COTS marine radar systems for advanced sensing applications.

1.2 Contributions

This thesis makes the following contributions:

- Development of two novel algorithms, Gradient Ascent and Phase Difference Averaging, for estimating and compensating the initial transmit phase in magnetron-based pulsed radar systems

- Introduction of two estimation techniques for compensating time-dependent intermediate-frequency instabilities in COTS marine radar systems without hardware modification

- Comprehensive quantitative evaluation of the proposed methods, alongside existing techniques, using novel coherence metrics across diverse radar environments

1.3 Thesis Outline

The remainder of this document is organized into the following chapters:

Chapter 2 reviews fundamental radar concepts, including hardware characteristics, signal models, range measurements, noise definitions, signal-to-noise interactions, and clutter. It also introduces the basics of Doppler processing and discusses the role of coherence in radar performance. The content in this chapter creates the foundation upon which the material in this thesis is built.

Chapter 3 describes the COTS radar system used in this study, including data acquisition and system interfacing. It outlines experimental procedures for characterizing system noise and time delay behaviors, and defines the creation of an in-situ transmit pulse model. Finally, emphasis is placed on the digital downconversion process and diagnosis of the time-varying IF drift.

Chapter 4 presents the development of coherency algorithms designed to compensate for constant phase errors and restore coherence. Techniques such as phase difference estimation, gradient-based optimization, and matched filtering are described, along with adaptive digital downconversion methods for mitigating IF drift. Performance metrics meant to encapsulate coherence of the system, including phase variance, angular variance, and Doppler characteristics, are also introduced.

Chapter 5 evaluates the proposed methods using experimental data from stationary targets and clutter environments. Assessment and optimization of algorithm effectiveness is made in situations with and without dominating ground clutter. Performance is analyzed with range-Doppler visualizations, coherence metrics, and comparative studies. Observed limitations, including sensitivity to

IF estimation uncertainty, are also discussed.

Chapter 6 extends the analysis to non-stationary clutter in both ground and maritime environments. Scenarios include varying target dynamics and sea clutter conditions.

Chapter 7 concludes the thesis by summarizing the key findings and contributions, reiterating the scientific contributions of the work. It also describes limitations of the current approach and provides recommendations for future work.

Chapter 2

Radar Principles

2.1 Radar Hardware

2.1.1 Radar Structure

Modern radar systems vary widely in their hardware implementations, but most share a common set of core components and subsystems designed to perform specific functions. This section provides a highly simplified description of key elements in a typical radar system, following the signal path from generation to digitization. This discussion is not exhaustive and omits many additional filtering, frequency conversion, and RF device stages that may be present in practical systems. This work focuses primarily on pulsed radar systems, in which the system transmits a pulse of RF energy over a defined duration, followed by a listening interval during which echoes are received before the next pulse is transmitted.

Transmit Subsystem

The transmit subsystem begins with a waveform generator, which produces a representation of the desired phase-modulated signal at low frequencies, typically near DC. This representation is commonly referred to as the baseband signal. Modern waveform generators exercise precise control over frequency and phase, enabling the synthesis of a wide array of arbitrary waveforms [13].

However, baseband signals are not suitable for direct transmission and must be mixed to the signal carrier frequency prior to transmission into the environment. This process is accomplished using an RF mixer and a stable local oscillator, which provides the reference signal for frequency conversion. The mixing process inherently produces undesirable image frequencies and intermodulation products, which must be removed using appropriate filtering. After application of an analog filter, the signal reaches a high power amplification stage to produce signals of the necessary strength. In systems with a magnetron power oscillator, the functions of waveform generation, frequency conversion, and amplification are effectively combined within a single device, making it an attractive option for cost-efficient applications.

After amplification, the signal is routed through a circulator, a non-linear device that passes energy between ports in a unidirectional manner (i.e., Port 1 to Port 2, Port 2 to Port 3, and Port 3 to Port 1), while maintaining high isolation against power flow in the reverse direction. This protects sensitive receiver components from high transmit power while allowing much weaker return echo signals to pass to the receiver. Finally, the signal is radiated by the antenna, which shapes the spatial distribution of the transmitted energy according to its radiation pattern. Antenna transmission is inherently non-ideal, leading to the formation of sidelobes that radiate energy outside the main beam.

Receive Subsystem

After transmission, the radiated energy propagates through the environment, interacting with objects and undergoing reflection, scattering, and attenuation, with some portion of energy returning to the radar system. The receive antenna concentrates EM energy density from free space and transfers it to the circulator, which routes the signal to the receive subsystem. The receive chain sometimes includes a limiter circuit to protect downstream components from high-power signals. The received signal is then amplified, typically by several orders of magnitude with a low-noise amplifier (LNA) designed to minimize the introduction of additional noise. The amplified signal is subsequently downconverted to an intermediate frequency with a stable local oscillator and filtered to suppress image frequencies. The final stage of the receive subsystem is the creation of in-phase and quadrature data from the IF signal as discussed in Section 2.1.2. This receive architecture is commonly referred to as a superheterodyne receiver [3].

2.1.2 Demodulation Principles

The motivation for downconversion and baseband processing is rooted in the Nyquist sampling theorem, which states that the sampling rate must be greater than or equal to twice that of the highest frequency present in the signal. High sampling rates can raise demands on power consumption, data throughput, and computational load, so received data is typically shifted to lower frequencies prior to digitization, allowing for significantly reduced sampling rates. Frequency translation is achieved through "mixing," a non-linear operation that in ideal contexts acts as multiplication by a sinusoidal signal. Using the trigonometric identity

$$\sin(\omega_{RF})\cos(\omega_{LO}) = \frac{\sin(\omega_{RF}+\omega_{LO})+\sin(\omega_{RF}-\omega_{LO})}{2}, \tag{2.1}$$

which contains both sum and difference frequency components. By filtering out the unwanted image frequency, in this case $\omega_{LO}+\omega_{RF}$, the signal can be shifted in frequency. These principles underlie upconversion to RF, downconversion to IF, and subsequent conversion to baseband. In practice, perfect mixing of two analog frequencies is unattainable, producing intermodulation products that corrupt the signal and must be mitigated through additional filtering.

Analog Demodulation

After downconversion to IF, the received signal is represented as a real-valued voltage waveform centered at the intermediate frequency. Analog demodulation generates data products by splitting the signal into two parallel channels of equal length and mixing them with sine and cosine reference signals, respectively. Each signal is low-pass filtered to remove the image frequencies, which produces two analog signals that are 90° out of phase, called the in-phase (I) and quadrature (Q) signals. Together they form the complex representation of the signal

$$x(t) = I(t) + jQ(t) = e^{j\theta(t)} \tag{2.2}$$

which is widely referred to as IQ data. IQ representation enables efficient signal processing by capturing both amplitude and phase information. Unlike real-valued signals, which must have symmetric spectral content, complex-valued signals can be used to represent an asymmetric frequency spectrum. In this way, the signal spectrum can be centered at DC, and the received information can be fully recreated with a sampling rate equal to the bandwidth, although traditionally the sampling rate is chosen by some margin above the minimum threshold. This approach significantly reduces power consumption, computational burden, and financial cost on the system and is a widely implemented practice.

Digital Demodulation

While analog demodulation is a critical element of many radar systems, it can also be responsible for many practical challenges that must be addressed. In an ideal context, the I and Q channels have identical amplitude and phase responses and are precisely 90° out of phase from one another. This is impossible in practice, leading to a disparity in amplitude and phase often called IQ imbalance [14]. An alternative approach is to sample the IF directly and perform the final demodulation in the digital domain. Although this requires a higher sampling rate (exceeding twice the IF plus half the signal bandwidth), it is advantageous in many cases because of the improved flexibility and removal of analog imperfections.

Digital downconversion is typically implemented in two ways. The first technique is a digital frequency shift with a complex exponential ($e^{-j2\pi f_{IF}t}$), followed by a lowpass filter. Because real-valued signals have symmetric spectra, the lowpass filter is required to remove the mirrored frequency components previously located around $-f_{IF}$. An alternative method employs the Digital Hilbert Filter [15] and was not used in this work. An optional final step in digital demodulation is decimation, which reduces the sampling rate after the signal has been shifted to baseband. Since the newly downconverted signal spectrum is now centered at DC, the sampling rate needs only to satisfy $f_s \geq \beta$, where β is the signal bandwidth, to fully encapsulate the recorded information. Decimation can significantly reduce data throughput and processing requirements, but was not used in this study.

2.2 Basic Radar Signal Processing

2.2.1 Transmit/Receive Signal Model

Transmit Signal Model

The operational modes and capabilities of a radar system are strongly influenced by the signal waveform. A single pulse of an arbitrary waveform modulated in amplitude and phase can be expressed as

$$\bar{x}(t) = a(t)cos[2\pi F_c t + \theta(t)] = a(t)Re\{e^{j(2\pi F_c t + \theta(t))}\} \tag{2.3}$$

where a(t) denotes the pulse amplitude envelope, F_c is the carrier frequency, and $\theta(t)$ represents the phase modulation applied to the transmitted waveform. In many systems, the transmitter operates in saturation, such that increases in input power do not translate to a proportional increase in output power. Consequently, it is common to have no amplitude modulation, and $a(t)$ is often approximated as a rectangular pulse with a duration of τ_{pw} [3]. For phase-modulated waveforms, $\theta(t)$ may vary rapidly with time, and is manipulated to improve or specialize in a wide array of parameters. Although there was no intentional phase modulation applied in this system, the use of a magnetron transmitter ensures that the initial phase of the transmitted signal is random. Accordingly, the phase modulation is modeled with $\theta(t) = \phi_m$ where ϕ_m is the phase offset of the mth pulse relative to an arbitrary reference.

Return Signal Model

The received signal corresponding to a single transmitted pulse can be modeled as a time-delayed and amplitude-scaled version of the transmitted signal:

$$\bar{y}(t) = a(t-\tau)e^{j2\pi F_c(t-\tau)+\theta(t-\tau)} + N(t) \tag{2.4}$$

where τ is the round-trip propagation delay associated with a target at range R_0, given by $\tau = 2R/c$. F_c is the carrier frequency, $a(t)$ reflects amplitude scaling governed by the radar range equation (Section 2.2.5), and $N(t)$ traditionally represents an additive noise term as described in Section 2.2.4. This model neglects target motion, which is addressed in Section 2.3.3. As in the transmit model, the phase modulation term is a constant written with $\theta(t) = \phi_m$.

2.2.2 Range Quantities

In pulsed radar systems, it is often impractical to simultaneously transmit and receive EM energy. Although a circulator provides high isolation between transmit and receive paths, it is an imperfect device that allows leakage of high-power transmit energy into the receiver. This leakage can dominate the receiver input during transmission, effectively preventing the detection of target echoes within a certain interval. The corresponding space traversed by reflected EM waves during this time is known as the blind range and is given by

$$R_{Blind} = \frac{\tau_{pw} c}{2}. \tag{2.5}$$

If a target is located within this range, it can easily be missed due to saturation of the receive subsystem by the high-power transmit leakage.

Another key parameter is the range resolution, which describes the ability to distinguish between objects near one another in range. For a simple, unmodulated pulse, the radar resolution is found with

$$\Delta R = \frac{\tau_{pw} c}{2}. \tag{2.6}$$

In a system utilizing pulse compression waveforms, this quantity is dependent on the bandwidth of the signal rather than the pulsewidth and can be found with

$$\Delta R \approx \frac{c}{2\beta} \tag{2.7}$$

where β is the bandwidth of the pulse. However, the hardware used in this study used only simple pulse waveforms, for which the range resolution is equal to the blind range.

A third system-specific parameter is the unambiguous range. Pulsed radar systems emit pulses at the pulse repetition frequency (PRF), with corresponding pulse repetition interval (PRI) $T_{PRI} = 1/\text{PRF}$. If a return echo arrives after a subsequent pulse has been transmitted, ambiguity arises in determining its origin, and the system falsely assumes the signal is an echo from the second radiated pulse. Thus, the maximum range for which echoes can be uniquely associated with a single pulse is then called the unambiguous range and written with

$$R_{ua} = \frac{cT_{PRI}}{2}. \tag{2.8}$$

There exist methods for mitigating this ambiguity, including staggered PRFs, but none were implemented with the hardware in this study. Fortunately, the hardware used in this study provides an unambiguous range of 54.3 km, which significantly exceeds the maximum detectable range of interest, so the possibility for range ambiguities was ignored.

2.2.3 Pulsed Radar Data Structure

Radar signal processing is often restricted to echoes within a predetermined set of ranges. The blind range, for example, is typically useless and might be omitted from the processed data. The measured data from the ranges of interest are compiled for pulses that will be processed together into a coherent processing interval (CPI). For the sequence of pulses within a CPI, received data is organized into discrete range bins, each corresponding to a specific distance interval.

The resulting data structure can then be represented as a two-dimensional matrix. One dimension, referred to as fast-time, corresponds to samples within a single pulse and maps directly to the individual range bins. The second dimension, known as slow-time, represents the data contained within an individual range bin between successive pulses in the CPI. In this work, the range bin size is defined as the round-trip interval traveled by a propagating signal in the time between two adjacent samples. It is determined by the sampling period $T_s = 1/F_s$ of the fast-time signal and written with

$$R_{\text{bin}} = \frac{cT_s}{2}. \tag{2.9}$$

The range bin spacing can be significantly finer than the system's range resolution, which aids in resolving targets that otherwise would be located at or near resolution boundaries.

2.2.4 Noise

Thermal Noise

In every radar system, the EM echoes from targets in the scene must be distinguishable from the noise of the system. Noise is defined as unwanted electromagnetic energy that interferes with a receiver's ability to detect a desired signal. The dominant noise source is typically thermal noise, generated by the constant random motion of electrons. These random motions produce power variations that are modeled by an additive zero-mean Gaussian random process [16]. The noise power spectral density at frequencies lower than 100 GHz with temperatures larger than 50 K can be approximated with $S_n(F) = kT_0$ W/Hz, where k is Boltzmann's constant (k $= 1.38 \times 10^{-23}$J/K), and T_0 is the temperature in Kelvin. When passing through an amplifier, both thermal noise and the signal are amplified, scaling the noise power by G_r, which represents the receive gain of the system. Each RF component contributes noise to the system, which is often lumped together into the noise figure term F_n. The total noise power is obtained by integrating the spectral density over the noise-equivalent bandwidth β_n, resulting in

$$N(t) = kT_0 G_r F_n \beta_n. \tag{2.10}$$

However, for the equipment in this study, the noise figure of the receive subsystem and its overall receiver gain were unknown, and so the true thermal noise calculation could not be performed.

Quantization Noise

When digitizing any signal, it is important to also consider the effects of quantization of a continuous-time process into discrete samples. In [17] the quantization error of a system can be approximated as white Gaussian noise if the quantizer does not overload, the probability distribution of the error is uniform over the range of quantization error, the error sequence is uncorrelated with the input sequence, the error sequence is a sample sequence of a stationary random process, and the random variables of the error process are uncorrelated. If these criteria are met, the white-noise model of quantization error has a power spectral density equal to its variance and can be written with

$$P_{qn} = \sigma_{qn}^2 = \frac{2^{-2(B-1)} X_{max}^2}{12} \qquad |\omega| \leq \pi \tag{2.11}$$

where B is the bit resolution of the ADC and X_{max} is the full-scale input range [17]. There are many examples of situations where these assumptions are valid, but experimental and theoretical studies concluded that the measured correlation between input and error decreases and the error samples become uncorrelated when the quantization step size is small relative to the input signal and the signal traverses multiple quantization bins between samples [18], [19], [20]. These conditions may be met initially by measured radar return signals, but as target range increases, the received signal amplitudes can diminish enough to change the statistical properties of the noise.

The statistical characterization of quantization noise becomes very difficult when those assumptions fail, with some proposed representation methods described in [21]. As a result, metrics and evaluations conducted in this study attempt to avoid target ranges and radar cross-section (RCS) values at which this quantization noise is not easily described. In most radar systems, it is customary to set the quantization noise significantly below the amplified thermal noise floor to reduce the effect on overall noise. However, for the equipment in this study, setting the quantization noise floor at such a level was not possible due to the required input signal power ranges and physical hardware limitations.

2.2.5 Signal-to-Noise Ratio

The relationship between radar system hardware characteristics and the maximum detectable range of a target is modeled by the radar range equation. For isotropic radiation, the power density at range R is $P_t/(4\pi R^2)$. If the antenna has a transmit gain of G, an effective area of A_e, and the radiated energy reflects off an object with a RCS of σ, a simple version of the radar-range equation can be modeled with

$$P_r = \frac{P_t G}{4\pi R^2} \cdot \frac{\sigma}{4\pi R^2} \cdot A_e = \frac{P_t G A_e \sigma}{(4\pi)^2 R^4} \tag{2.12}$$

where P_t and P_r are the peak transmit and receive powers. For monostatic systems, the effective aperture is $A_e = \frac{G_r \lambda^2}{4\pi}$ where λ is the wavelength of the transmitted signal and G_r is the receive gain of the antenna. The signal power loss in a system is denoted L_s, and the attenuation of the signal in the atmosphere is represented by L_a. A received radar echo must then be distinguished from the noise power by some algorithm, which is highly dependent on the signal-to-noise ratio (SNR). The SNR is a unitless measure of the received signal power, calculated with the radar range equation, divided by the noise power of the system, and can be written as

$$SNR = \frac{P_r}{P_n} = \frac{P_t G_t G_r \lambda^2 \sigma}{(4\pi)^3 R^4 L_s N(t)}. \tag{2.13}$$

Traditionally, $N(t)$ is written assuming that thermal noise forms the noise floor using Equation (2.10), but in the case of this system, the noise floor is set by the quantization error of the ADC. As mentioned earlier, quantization noise is difficult to characterize under the conditions of this work, so an estimation of the noise level was determined experimentally with the methods described in Section 3.1. Although no fixed SNR threshold guarantees detection, it strongly influences the performance of a detector in a radar system. The principles of detection are not examined in this work, but it is important to know that the SNR plays a large role in the relationship between the probability of detection and the probability of false alarm. An example of this dependence for a simple threshold detector in the presence of complex Gaussian noise is illustrated in Figure 2.1.

2.3 Radar Coherence

Coherence in a pulsed radar system refers to the ability to preserve phase relationships between successive transmitted pulses. When the system can maintain or accurately recover the initial phase of each transmitted waveform, it enables

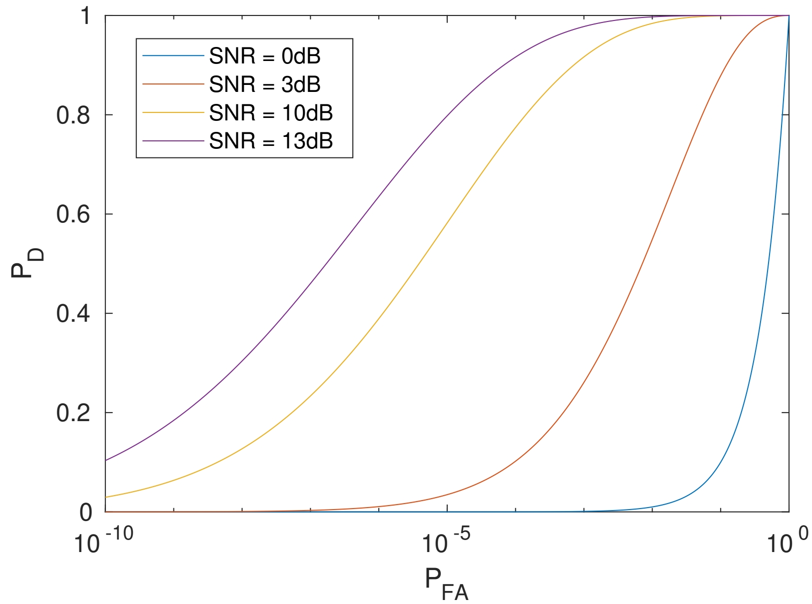

Figure 2.1: Example probability of detection vs. probability of false alarm for multiple SNR values

the extraction of additional information beyond amplitude-only measurements. As described in [22], coherent radar systems are generally classified into two categories.

Class I radar systems use a pulsed oscillator with a coherent reference oscillator to create a coherent pulse train. These systems do not ordinarily produce transmitted waveforms on transmission of the signal, but instead correct the variability in phase with a detector circuit using the coherent reference oscillator in the receive subsystem. In contrast, Class II systems utilize a stable master oscillator to generate pulses with a consistent initial phase across transmissions. This architecture inherently supports coherent processing and can be extended to systems that apply intentional phase modulation, such as pulse compression techniques.

The radar system considered in this study does not strictly conform to either class. It exhibits non-coherent pulse emissions and lacks a dedicated coherent reference oscillator. Instead, coherence is achieved through software-based compensation of phase instabilities applied directly to sampled radar echo data.

2.3.1 Noncoherent Processing

In a fully noncoherent radar system, signal processing capabilities are limited due to the absence of pulse-to-pulse phase information. In such systems, the phase information of IQ data is discarded, and detection and tracking rely solely on the magnitude of the measured data samples. SNR improvements can be achieved through noncoherent integration, where some function of the signal magnitudes (log-magnitude, magnitude-squared, or magnitude itself) is summed across a coherent processing interval. The output for noncoherent integration of the magnitude of the radar measurements, as described in [3], is given by

$$Y = \sum_{n=0}^{N-1} |Ae^{j\phi} + w[n]|^2 = NA^2 + \sum_{n=0}^{N-1} |w[n]|^2 + \sum_{n=0}^{N-1} 2Re\{Ae^{j\phi} \cdot w^*[n]\} \tag{2.14}$$

where A is the signal amplitude, N is the number of integrated pulses, $e^{j\phi}$ is the signal phase, and $w[n]$ is the additive noise. The presence of the cross-term involving both signal and noise complicates separation into distinct signal and noise components. Consequently, the non-coherent integration gain, G_{NC}, of a signal cannot be written simply, but is stated in [23] to be within the range of

$$G_{NC} = N^{\alpha} \qquad 0.5 < \alpha < 1. \tag{2.15}$$

Noncoherent integration is often the final stage of processing prior to detection in such systems and does not provide any information regarding target velocity.

2.3.2 Coherent Integration

The integration gain is substantially easier to quantify in a fully coherent radar system where the phase information is retained. The summation of N coherent complex voltage samples yields

$$Z = \sum_{n=0}^{N-1} \{Ae^{j\phi} + w[n]\} = NAe^{j\phi} + \sum_{n=0}^{N-1} w[n] \tag{2.16}$$

which is cleanly separable into signal and noise components [3]. Assuming the noise samples are independent, zero-mean, and have variance σ_n^2, the power of the summed noise terms is simply $N\sigma_n^2$. The integrated signal power is calculated with N^2A^2, formed from the summed complex voltage samples. The resulting SNR after coherent integration is

$$SNR_{\text{out}} = \frac{N^2A^2}{N\sigma_n^2} = N\frac{A^2}{\sigma_n^2} = N(SNR_{\text{in}}). \tag{2.17}$$

Thus, coherent integration provides an improvement in SNR proportional to the number of pulses integrated in the CPI.

This result, however, requires that corresponding samples from separate pulses have the same initial transmit phase of ϕ. If the initial transmit phase changes between pulses without compensatory measures, the integration gain suffers as $\sum_{n=0}^{N-1} Ae^{j\phi_m} \neq NAe^{j\phi}$. Additionally, the integration gain is only valid if the noise component $w[n]$ is fully independent of the signal $Ae^{j\phi}$, or the noise power calculation changes. The signal $Ae^{j\phi}$ is deterministic, so its integrated power level does not change with non-independent noise, but the integrated noise power is found by calculating the variance of Z, given by

$$\begin{aligned}\mathbb{E}[Z^2] - \mathbb{E}[Z]^2 &= \mathbb{E}\left[N^2A^2e^{j2\phi} + 2NAe^{j\phi}\sum_{n=0}^{N-1} w[n] + \left(\sum_{n=0}^{N-1} w[n]\right)^2\right] \\ &\quad - \mathbb{E}\left[NAe^{j\phi} + \sum_{n=0}^{N-1} w[n]\right]^2 \\ &= N^2A^2e^{j2\phi} + \mathbb{E}\left[2NAe^{j\phi}\sum_{n=0}^{N-1} w[n]\right] + N\sigma_n^2 - N^2A^2e^{j2\phi} \\ &= \mathbb{E}\left[2NAe^{j\phi}\sum_{n=0}^{N-1} w[n]\right] + N\sigma_n^2.\end{aligned} \tag{2.18}$$

This resulting noise power contains an additional term that depends on the signal-noise relationship and the length of the CPI, and whose expected value is likely nonzero. While thermal noise is independent of the signal, quantization noise violates this assumption when the signal amplitude approaches the quantization step size [17]. Consequently, objects in the radar environment that might emerge from beneath a thermal noise floor with coherent integration don't necessarily exhibit the same behavior when the noise floor is set by quantization error effects.

2.3.3 Doppler Processing

A leading advantage of a coherent radar system is the ability it creates to estimate target velocity through Doppler processing. This capability enables improved clutter suppression, enhanced moving target detection, micro-Doppler analysis, and improved target discrimination. Unlike noncoherent integration, which produces a single trace of summed measurements, coherent Doppler processing yields a range-Doppler map (RDM), representing target energy as a function of both range and velocity [24]. An example RDM is shown in Figure 2.2.

Doppler processing is based on the principles of traveling waves. When there is relative motion between a wave source and an observer, the spacing of wave

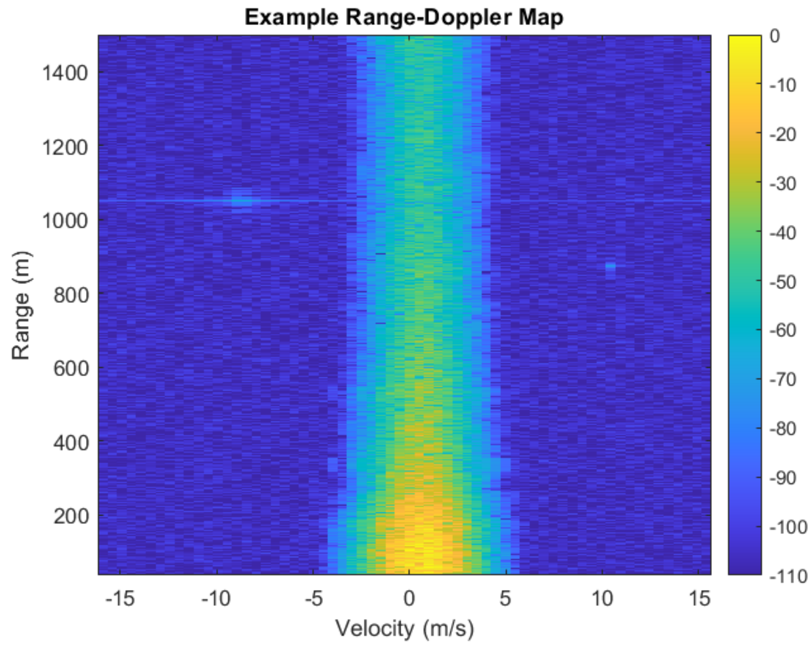

Figure 2.2: Example range-Doppler map with strong clutter and two moving targets

fronts changes proportionally to the radial relative velocity, thereby changing the observed frequency of the wave. That change is called the Doppler frequency shift, which can be calculated for non-relativistic objects detected by radar using

$$F_d = \frac{2v}{c} F_c = \frac{2v}{\lambda} \tag{2.19}$$

where λ is the wavelength and v is the radial velocity component between target and observer. Within the duration of a single pulse, the Doppler-induced phase shift exists in theory, but measurement of that shift is an impossible task in practice. Therefore, Doppler estimation relies on phase differences between successive pulses, which requires coherence over the CPI.

Phase History

The principles of pulse-Doppler processing can be illustrated using the signal model from Section 2.2.1. If an object is assumed to move with constant radial velocity over a CPI, its range can be written as a function of time with $R(t) = R_0 - vt$, where R_0 is the initial range and v is the radial velocity. The receive signal model from Equation (2.4), ignoring pulse modulation and noise terms, can be written with

$$y_m(t) = A\left(t - \frac{2R(t)}{c}\right) \exp\left[j\left(\frac{-4\pi}{\lambda_0}R(t) + 2\pi F_c t\right)\right]. \tag{2.20}$$

Proper downconversion removes the $j2\pi F_0 t$ term, and if the amplitude scaling is absorbed into a constant k, the signal as a function of m pulses radiated every T_{PRI} seconds becomes

$$y_m\left(t = mT_{PRI} + \frac{2R_0}{c}\right) = k\ \exp\left[-j\frac{4\pi}{\lambda_0}\left(R_0 - \frac{2vR_0}{c}\right)\right] \exp\left[j\frac{4\pi v}{\lambda_0} mT_{PRI}\right]. \tag{2.21}$$

Under the stop-and-hop assumption, it is assumed that the target does not move while the mth pulse is in the air. In this assumption $2vR_0 << c$, so the $\frac{2vR_0}{c}$ term approaches zero, and the round-trip phase term can be assumed not to change between pulses. Thus, the second exponential term is responsible for the phase change between pulses and can be written with $\frac{4\pi v}{\lambda_0}$ every T_{PRI} seconds. This is known as the phase history of a series of pulses and is equal to a $2\pi F_d$ radian shift between each pulse, forming the basis for velocity estimation.

Range-Doppler Representation

The information provided by the phase history of a coherent pulse train can then be used to formulate a representation of the targets in the scene using the discrete Fourier transform (DFT) applied along the slow-time axis. This operation converts the sequence of measured data into its constituent frequencies. In this way, the sequential phase shifts across slow-time for individual range bins are concentrated into frequency bins that represent the energy concentrated at those Doppler frequencies, forming the range-Doppler map.

Since the phase shifts of a moving target between adjacent pulses are equal to its Doppler shift in radians, the PRF of the system functions as the sampling rate of that Doppler shift. The Nyquist Theorem then states that the unambiguous measurable Doppler shift of a target must fall within

$$-\frac{F_{PRF}}{2} < F_d < \frac{F_{PRF}}{2}. \tag{2.22}$$

Objects moving at speeds exceeding the unambiguous Doppler shifts will then alias to some frequency within the allowable range, and their true velocity will be obscured.

The Doppler frequency resolution is determined by the properties of the DFT and is given by

$$\Delta F_d = \frac{F_{PRF}}{N} \tag{2.23}$$

where N is the number of pulses in a CPI. Thus, increasing N produces a finer velocity sampling grid, but does not alter the measurable velocity bounds. The final output of conventional range-Doppler processing is therefore a range-Doppler map characterized by these fixed bounds and the corresponding velocity resolution.

2.4 Clutter

Clutter consists of electromagnetic radar returns from physical objects in the environment, including terrain, vegetation, sea surfaces, and weather phenomena. These returns can obscure targets of interest and degrade detection performance. Unlike noise, clutter arises from real physical scatterers in the scene, rather than random or error-based quantities that corrupt a received signal. This distinction is important because in a coherent system, the clutter returns sum coherently in addition to those from the target of interest.

Depending on the application, clutter may be either a nuisance or a source of useful information. For example, a radar altimeter, SAR, and ground-mapping radar systems all make use of radar echoes from the land, while these echoes interfere with other radar applications. In this study, both ground and sea clutter are considered. As the name might suggest, ground clutter comprises the signal returns of objects in the scene located on the ground. When radiating from a stationary platform, ground clutter has no velocity relative to the radar system and can be mitigated using relatively simple filtering techniques. In contrast, sea clutter is inherently dynamic and exhibits complex statistical behavior, requiring more sophisticated modeling and suppression methods [25]. Both clutter types are analyzed in this work as potential signals of interest and as sources of interference, each presenting unique challenges for detection and processing.

Chapter 3

Radar System Characterization

3.1 Experiment Hardware

The physical radar system used for experimentation in this work was the Furuno FAR 1528, with the specifications listed in [26]. A key advantage of this work is the use of a commercially available hardware system. However, a direct consequence of this choice is that many of the hardware pieces within the RF receive chain were integrated into a proprietary printed circuit board (PCB), for which limited documentation and vendor support are available. As a result, the receive subsystem after the circulator is treated as a ”black box” system whose characteristics are assumed to be fixed but unknowable until the intermediate frequency stage of the superheterodyne architecture is reached.

At the IF output, an SHP-50+ lumped LC high-pass filter (41–800 MHz) is connected in series. The filtered analog signal is then digitized by an M4i-4481-x8 Spectrum Instruments analog-to-digital converter (ADC). The Furuno antenna unit is controlled remotely by a processing unit, which also provides a 2760 Hz trigger signal to the ADC for synchronization purposes. Digitized samples are transferred to a host computer via direct memory access (DMA) over a PCIe interface using the manufacturer-provided driver library. A block diagram of the system is shown in Figure 3.1. Two antennas were used for data collection. The

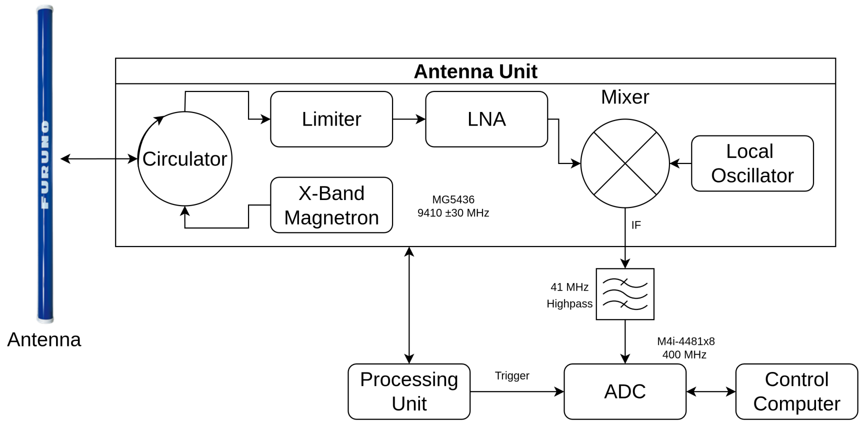

Figure 3.1: System block diagram,
reprinted from [1]

first was a Furuno XN20AF horizontally polarized slotted waveguide antenna, with a vertical beamwidth of 20°, a horizontal beamwidth of 1.23°, and a mainlobe gain of 30 dB. The second was a dish antenna provided by the United States Geological Survey (USGS), with approximately 1.35° beamwidth in both the vertical and horizontal planes. Detailed information regarding sidelobe structure and relative sidelobe levels was not available for either antenna.

The operational frequency of the Furuno FAR 1528 is specified at 9410 ± 30 MHz with a peak transmit power of 25 kW. The pulsewidth of the system is software-selectable and was fixed at 120 ns for all experimentation, corresponding to a range resolution of 17.99 m. The PRF of the system was 2760 Hz, yielding an unambiguous velocity of 21.98 m/s. Since all observed targets were assumed to fall within this velocity range, Doppler aliasing was ignored. Radar echo measurements were sampled by the M4i-4481x8 at 400 MS/s via the IF coaxial output, resulting in a range bin spacing of 0.3747 m.

3.2 M4i.4481-x8 Computer Interface

Following propagation through the receive subsystem, the analog radar return signal is digitized at the input of the M4i.4481-x8 ADC into a 14-bit signed integer representation of voltage. These samples are transferred to the control computer using DMA in accordance with the protocols and driver libraries described in [27]. Upon transfer to the RAM of the control computer, each digitized measurement is stored as a 16-bit integer in little-endian format, with appropriate handling of the two unused bits. The raw ADC codes are converted to double-precision floating-point voltage values using

$$V_{in} = ADC_{code} \times \frac{V_{\max}}{ADC_{\max}} \tag{3.1}$$

where ADC_{code} is the signed 14-bit integer value, $V_{\max}$ is the user-defined full-scale voltage, and $ADC_{\max} = 2^{14} - 1$. After conversion, the data is arranged in a two-dimensional matrix of fast-time (range) and slow-time (pulse index) raw samples as described in Section 2.2.3. All subsequent processing is performed using double-precision floating-point arithmetic.

3.3 Noise Floor Considerations

Characterizing the system noise floor presented a challenge in this experimental setup. In conventional radar systems, the noise floor is dominated by thermal noise, which can be modeled as additive white Gaussian noise (AWGN). Under these conditions, coherent integration yields an SNR improvement proportional to the number of pulses in the CPI, as described in Section 2.3.2. However, the system design did not permit the quantization noise floor to be set significantly below the thermal noise level. Consequently, the effective noise floor is dominated by quantization noise, which is not independent of the signal under small-signal amplitude conditions. Although Equation (2.11) predicts a quantization noise power of approximately -103 dB, its underlying assumptions are violated when signal amplitudes approach the quantization step size, rendering the model inaccurate. Thus, to estimate the effective noise floor, a region of data containing no

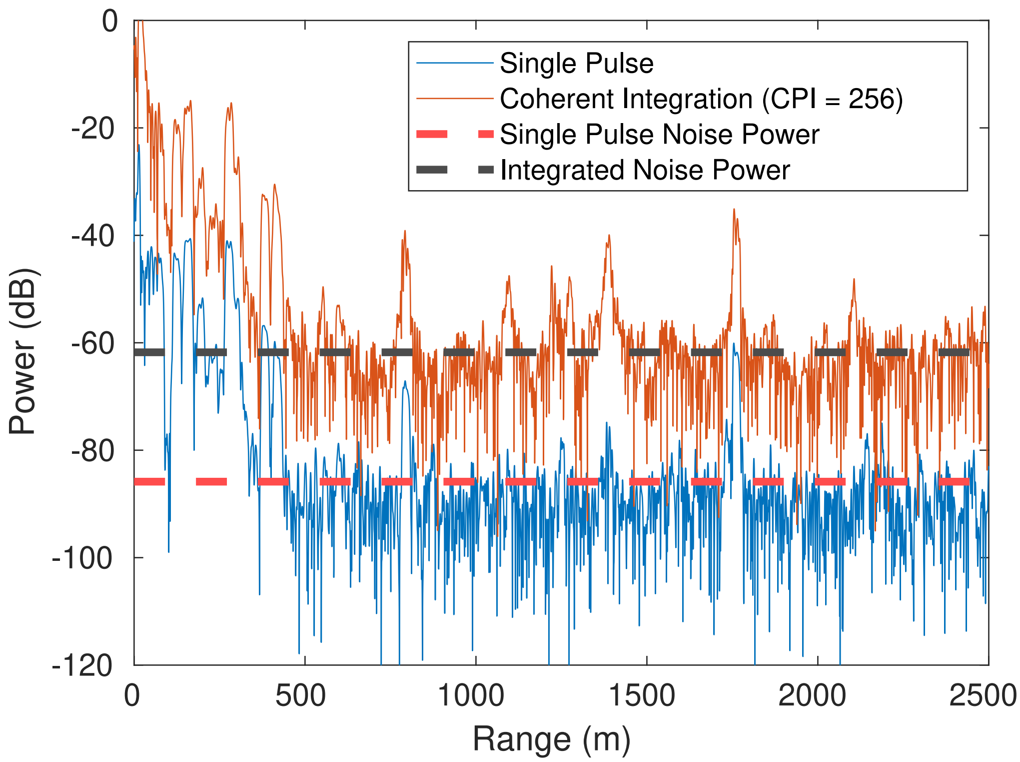

Figure 3.2: Power level comparison of single pulse vs. 256 coherently integrated pulses

discernible scatterers was analyzed, and the variance of the ensemble was found to be $\sigma_n^2 = -85.94$ dB. This value was observed to depend on the digital filter bandwidths, so measurements in this study were conducted using a 40 MHz bandpass filter and a low-pass filter with $f_c = 20$ MHz.

The measured power values of a single pulse, compared with those obtained after coherent integration over 256 pulses, are presented in Figure 3.2. The coherence of the pulses in this CPI was validated using the methods described in Section 4.3. As expected, strong target returns increased by approximately 24 dB, consistent with coherent integration of 256 pulses. However, the noise floor exhibited a similar increase, indicating that the noise components are also summing coherently. An estimated integrated noise trace, obtained by adding a theoretical coherent gain of $\log_{10}(256)$ to the single-pulse noise estimate (-85.94 dB), closely matches the measured results in regions devoid of strong scatterers. This behavior indicates that the noise floor of the system is the result of quantization error and scales similarly to coherent integration regardless of the proposed signal processing methods. As a consequence, weak target returns from small or distant objects do not emerge from the noise floor with coherent integration, representing a fundamental limitation of the system.

3.4 Transmit Pulse Modeling

Accurate compensation of magnetron-induced phase variations using the proposed methods requires a high-fidelity representation of the transmitted waveform. In an ideal context, the demodulated transmit signal follows the model in Equation (2.3), with a perfectly rectangular amplitude envelope and no phase modulation. In practice, real radar systems exhibit finite rise and fall times during which the system is not transmitting at peak power, referred to as radar pulse smearing [28]. As explained in Section 3.5.2, there can also exist frequency offsets introduced by imperfect downconversion techniques. These situationally dependent deviations from simulated signal representations prevent artificial creation of accurate waveform models, necessitating alternative transmit signal model development.

A more robust approach is to derive the transmit signal model directly from measured data. This approach was utilized by researchers in [29] who created a template of their transmitted waveform using an oscilloscope to measure the characteristics of their signal for use in downstream signal processing. In this study, the transmitted waveform is estimated on a pulse-by-pulse basis using digitized radar measurements. This approach improves robustness to fluctuating waveform characteristics that would otherwise degrade the accuracy of a fixed template.

Referring back to Section 2.1.1, when a radar transmits a high-power pulse, the circulator protects the sensitive receive components from extremely high-power voltages. Although the circulator provides isolation between transmit and receive paths, it is not ideal, typically offering 20-30 dB of isolation [30]. As a result, a portion of the transmitted signal leaks directly into the receive chain, which can be exploited to construct a high-fidelity model of the transmitted waveform for each pulse.

The pulsewidth of this radar system, 120 ns, and the sampling rate of 400 MS/s correspond to an indexed pulsewidth of approximately 48 samples. Initial attempts to isolate transmit power leakage estimated the location visually with non-coherently integrated radar data, plotted in Figure 3.3. While visual estimation was an accurate and adaptable technique, it was time-consuming and impractical to use in any operational system. Thus, it was desirable to create a programmatic methodology to determine the transmit leakage sample range. The baseband representation of a simple pulse waveform without phase-modulation is a constant rectangular pulse of comparable sample width. As a result, to enable automated detection, the leakage region is estimated via convolution of a rectangular window with the non-coherent sum of baseband signals. Mathematically, the calculation is given by

$$\begin{aligned} x[n] &= \sum_{i=1}^{N_p} |s_{bb,i}[n]| \\ n_{\text{Leakage}} &= \arg\max_{n \in x} (x[n] * \mathbf{1}_N) \end{aligned} \tag{3.2}$$

where $s_{bb,i}[n]$ denotes the baseband sample of the ith pulse, N_p is the number of pulses, and $\mathbf{1}_N$ is a vector of ones with length equal to the number of samples in a transmitted pulse. The index n_{Leakage} corresponds to the final sample of the leakage region, and the full set of leakage indices includes the preceding 47 samples. The resulting leakage location estimate is overlaid on the noncoherent sum in Figure 3.3, demonstrating accurate localization of the transmit leakage region. A marker representing when true radiation begins, based on the work in

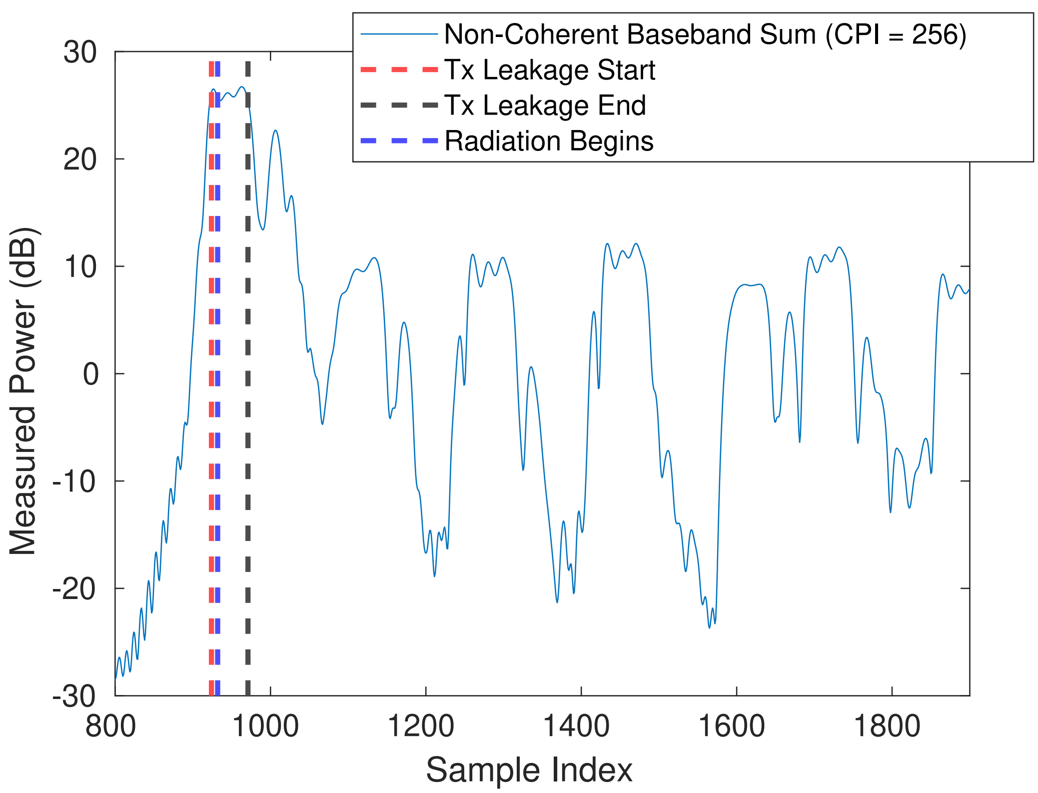

Figure 3.3: Transmit pulse model index extraction example

Section 3.6, is included in the figure, showing that the transmit pulse model could be contaminated by near-in clutter.

3.5 Digital Downconversion

3.5.1 Initial Filtering

When sampled by the ADC, the radar return signals remain centered at an intermediate frequency that should be brought to baseband prior to the application of coherence processing algorithms. In real-world scenarios, significant quantities of spectral content often exist outside the frequency band of interest, which can degrade or even saturate the important radar signals. For magnetron-based systems, common sources of such interference include power supply switching noise, harmonic distortion, intermodulation products, and internal reflections [31], [32]. Regardless of the origin of these undesired spectral components, appropriate filtering, either analog or digitally, is required. The IF of the Furuno FAR 1528 system is specified as 60 MHz with a signal bandwidth of 8.33 MHz, corresponding to a rectangular pulsewidth of 120 ns. These radar return echoes are observable in the measured spectrum, superimposed on a bed of corruptive spectral content, with an example of the unfiltered spectrum shown in Figure 3.4. It can be seen that the dominant contents of the frequency spectrum are concentrated at low frequencies, but that there also exist structured components representing the IF signals centered near 60 MHz. To suppress the strong low-frequency content, an analog high-pass filter with cutoff frequency $f_c = 41$ MHz was placed in series with the IF signal path, as described in Section 3.1. In addition, a 40 MHz digital bandpass filter centered at 60 MHz was applied to further isolate the signals of interest prior to downstream processing. Note that the bandpass filter significantly

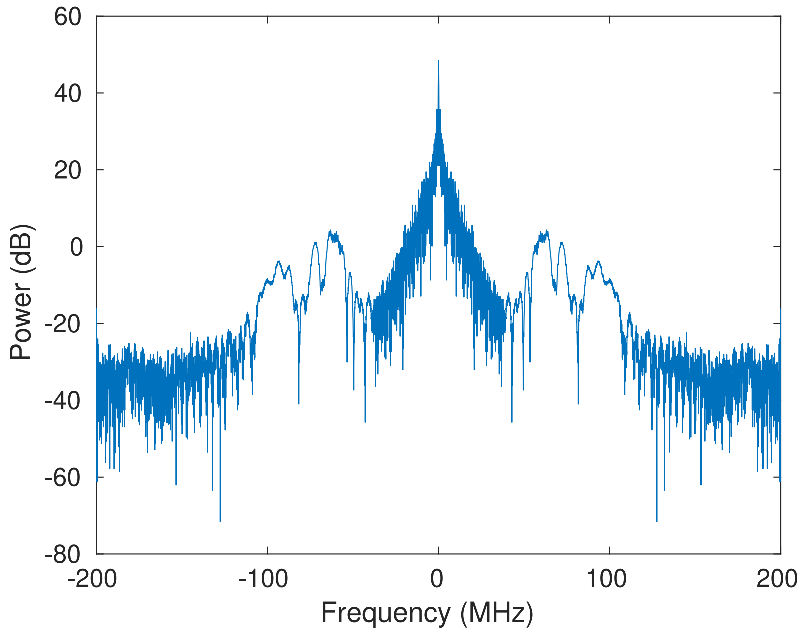

Figure 3.4: Unfiltered frequency spectrum

exceeds the signal bandwidth for reasons described in Section 3.5.2. The filtered frequency spectrum of a single pulse is shown in Figure 3.5.

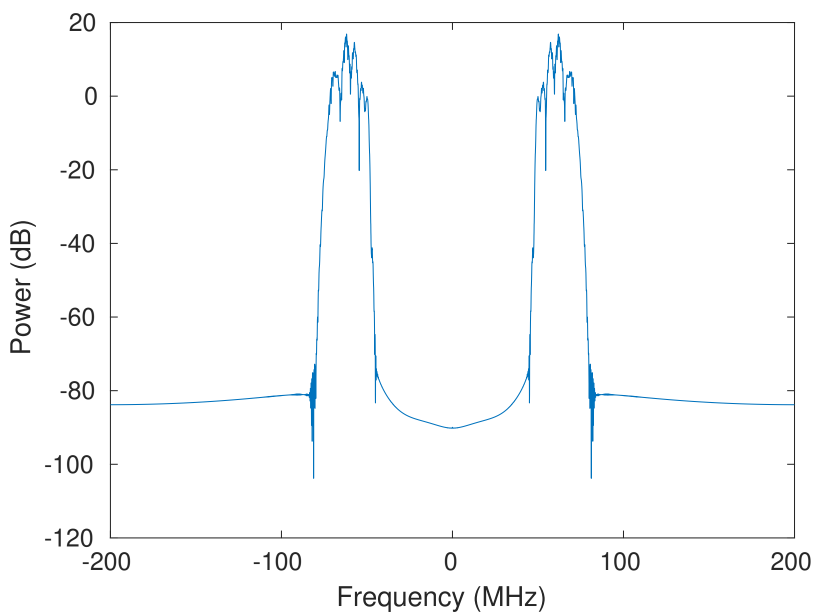

Figure 3.5: Filtered frequency spectrum

3.5.2 Pulse-Dependent Intermediate Frequency

As discussed in Section 2.1.2, digital downconversion can be achieved by multiplying the signal by a time-dependent complex exponential of the form $e^{-j2\pi f_{IF}t}$, followed by digital low-pass filtering. While this process is straightforward in conventional coherent radar systems, it was observed that the IF in this system is not constant and exhibits drift on the order of several megahertz over time. Figure 3.6 illustrates the fast-time spectral content of multiple pulses within a long CPI, clearly demonstrating this variation in IF. A key implication of this behavior is

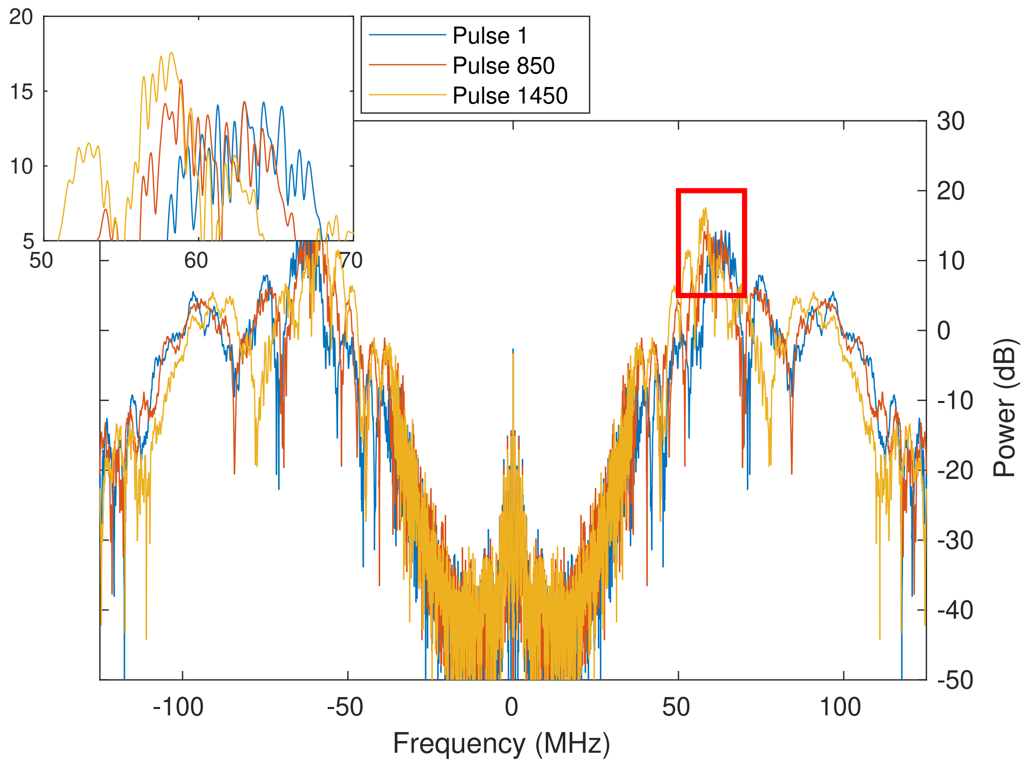

Figure 3.6: Fast-time frequency comparison,
reprinted from [1]

that simple downconversion using a fixed-frequency complex exponential, as described in Section 2.1.2, would fail to accurately convert the signal to baseband. When a measured return signal is improperly downconverted to baseband, the newly demodulated signal model for a signal point target can be expressed with

$$s_{bb}(t) = q(t - \tau)e^{j2\pi f_{err}t}e^{j\phi} \tag{3.3}$$

where f_{err} denotes the mismatch between the assumed and true IF. Assuming the IF remains constant over the duration of a single pulse, this new signal model introduces a linearly time-varying phase term. More importantly, this phase variation differs from pulse to pulse, degrading coherence across the CPI. Methods for estimating and compensating for this IF drift are presented in Section 4.1. Without correction, this loss of phase alignment reduces coherence and limits the effectiveness of the coherent processing techniques to create a pseudo-coherent radar system. An example of the baseband fast-time spectral content without IF drift correction is shown in Figure 3.7. It can be seen that there is a clear and repetitive structure in the drifting of the IF, a structure that remained largely consistent for all experiments conducted in this study.

IF Drift Source Identification

Determining the origin of the observed IF drift was critical for selecting an appropriate compensation strategy. If the drift originated in the transmit chain, successive pulses would propagate through the air at different carrier frequencies, invalidating conventional Doppler processing. While advanced techniques exist to address frequency-agile systems [33], they are computationally intensive and require accurate estimates of the frequency misalignments. Conversely, if the drift stems from components within the receive chain, the mitigation techniques outlined in Section 4.1 can compensate for the shifts, provided the IF estimates are accurate.

To identify the source of the drift, a WR-90 waveguide section with coupling

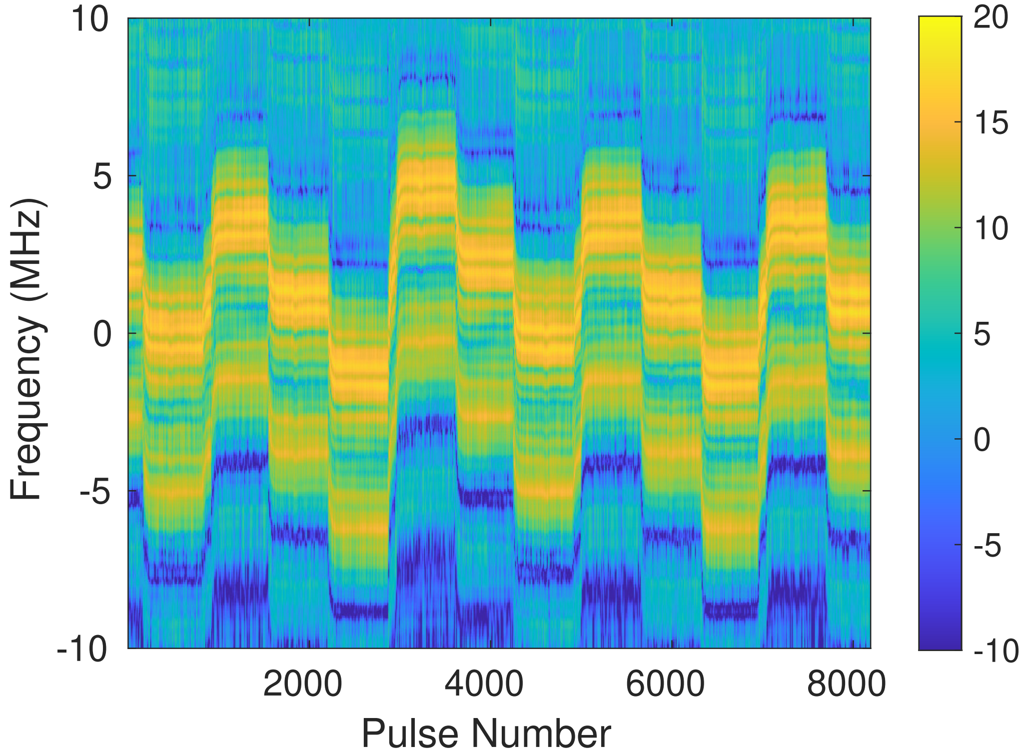

Figure 3.7: Fast-time frequency baseband content without IF drift compensation

into SMA was placed before the antenna, enabling direct measurement of the transmitted signal. An image of the waveguide coupler is shown in Figure 3.8. The high-power transmit waveform was attenuated to protect sensitive downstream RF components and mixed to a measurable IF frequency using a ZX05-14-S+ MiniCircuits mixer with a 9.355 GHz local oscillator generated by a Vaunix LMS-183 DX signal generator. The resulting IF signal was digitized using the same M4i.4481-x8 ADC employed for receiver measurements. This configuration allowed for direct observation of the transmit signal's spectral content prior to any receive subsystem distortions. The results, shown in Figures 3.9 and 3.10, indicate that the transmitted signal does not possess any of the frequency-drifting components present in the receive data over a long CPI. Therefore, it is clear that the IF variability originates within the receive subsystem, confirming that agile

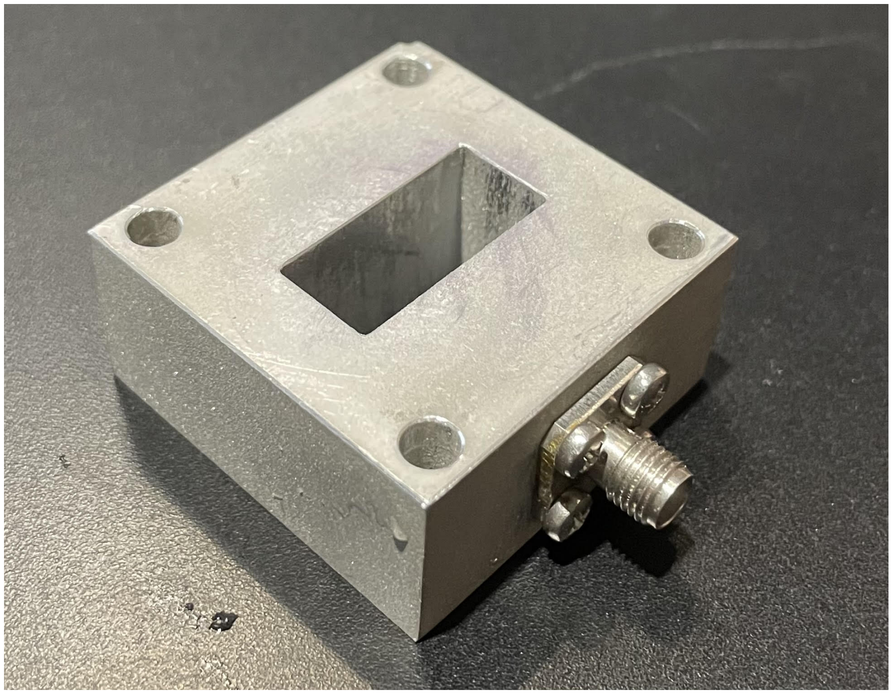

Figure 3.8: Waveguide coupler (designed by Brian Stade)

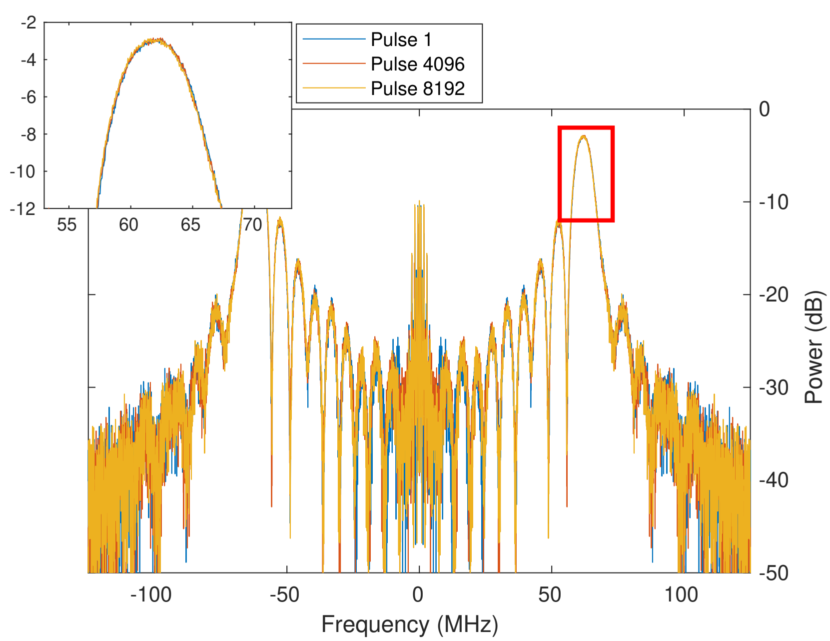

Figure 3.9: Transmit pulse fast-time frequency pulse comparison

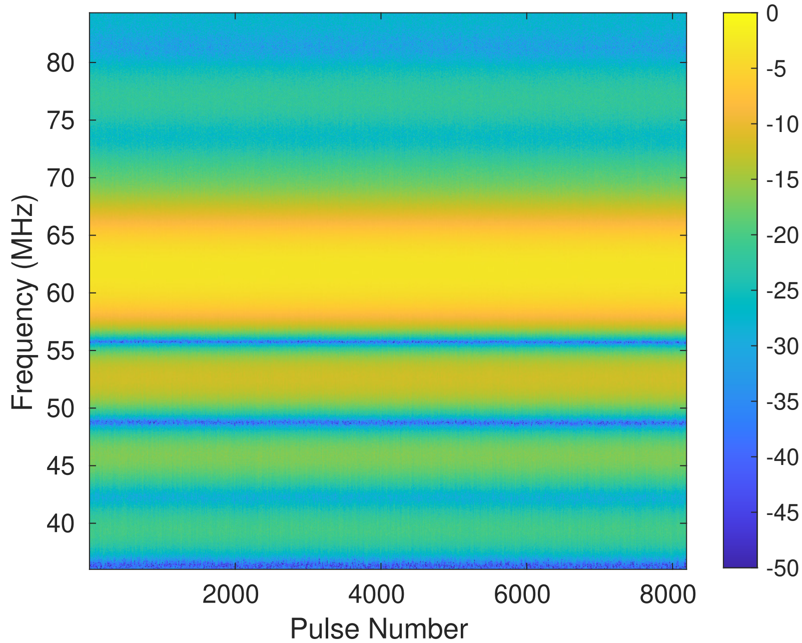

Figure 3.10: Transmit pulse fast-time frequency vs. pulse number

carrier frequency compensation techniques are unnecessary.

3.6 Trigger Timing Delay

Another fundamental question relating to the system characterization is whether the estimated transmit pulse model, obtained in Section 3.4, arises solely from transmit power leakage, or includes contributions from near-in clutter. Although the main beam of the antenna in these experiments and a deployed system would be directed toward open space, sidelobe radiation could possibly reflect from nearby clutter, contaminating the measured signal.

Furthermore, the system does not provide a direct measurement of the precise transmission time of each pulse. The ADC is triggered by a signal from the processing unit to ensure synchronization between pulses, but this trigger preceded radiation by an unknown, but consistent, delay. As a result, it was necessary to determine the absolute timing of transmitted pulses in relation to experimentally measured data. It should be mentioned that the identification of a transmit model, shown in Equation (3.2), operated independently of the delay, but that knowledge of the delay was essential for accurate range identification of targets.

To determine this delay, the antenna phase center was positioned at multiple known distances from a large stationary reflector. Measurements were conducted using a truck-mounted configuration, where the antenna was fixed to a steel structure securing the system to the vehicle. An image of the truck-mounted system with the attached USGS antenna is shown in Figure 3.11. In this way, the radar system was easily transported, allowing for simple data collection at several range bin locations for the same large scatterer. Care was taken to select distances larger than the blind range of the system and spanning several range resolution lengths. For each measurement location, the measured offset in range was calculated with

$$R_{\text{offset}} = R_{\text{measured}} - \frac{kc}{2F_s}, \tag{3.4}$$

Figure 3.11: Truck mounted setup

where k is the sample index corresponding to the peak response in coherently integrated matched filtered data. For IQ data yet to be matched filtered, the range offset becomes

$$R_{\text{unmatched}} = R_{\text{offset}} + \frac{(l_{tx} - 1)c}{2F_s} \tag{3.5}$$

where l_{tx} is the length, in samples, of the transmit pulse model.

Using a large semi-truck as a reflector and a Revasri Laser Rangefinder for ground-truth measurements, the results summarized in Table 3.1 were obtained. The values in the table were found using a 10 ft. coaxial cable connecting the processing unit and ADC, and 256 pre-trigger samples, both of which affect the

Table 3.1: Ground Truth vs. Measured Range/Sample Offset

Ground Truth (m)	**Radar-Measured**	**Range Offset**	**Sample Offset**
43.85	411.0904	-367.2424	980
61.35	426.8295	-365.4815	975
80.25	447.0655	-369.8655	979
97.75	465.8025	-368.0545	982
116.70	483.4153	-366.7173	979
134.25	501.0281	-366.7801	979
143.65	513.3946	-366.7466	979

trigger time delay. Under these conditions, the average range offset was found to be -366.83 m, corresponding to a delay of approximately 979 samples in matched-filtered data. For data without matched filtering, the range delay was found using Equation (3.5) to be -349.26 m, equal to a sample delay of 932 samples. This unmatched filtered delay equates to an experimental time delay of 2.33 μs from when the ADC begins collecting samples to when the pulse is actually radiated.

A comparison of the incoherent sums of 256 pulses from several distances included in Table 3.1 is presented in Figure 3.12. The corrected range markers were obtained by adding the experimentally determined range offset value to the uncorrected range axis. It is clear that they coincide with the beginning of the strong return pulses from the semi-truck, validating the estimated delay. A marker is also included in Figure 3.3, where it can be seen that the majority of measured data contained in a transmit pulse comes after radiation begins, presenting the possibility for corruption of the transmit pulse model by near-in clutter.

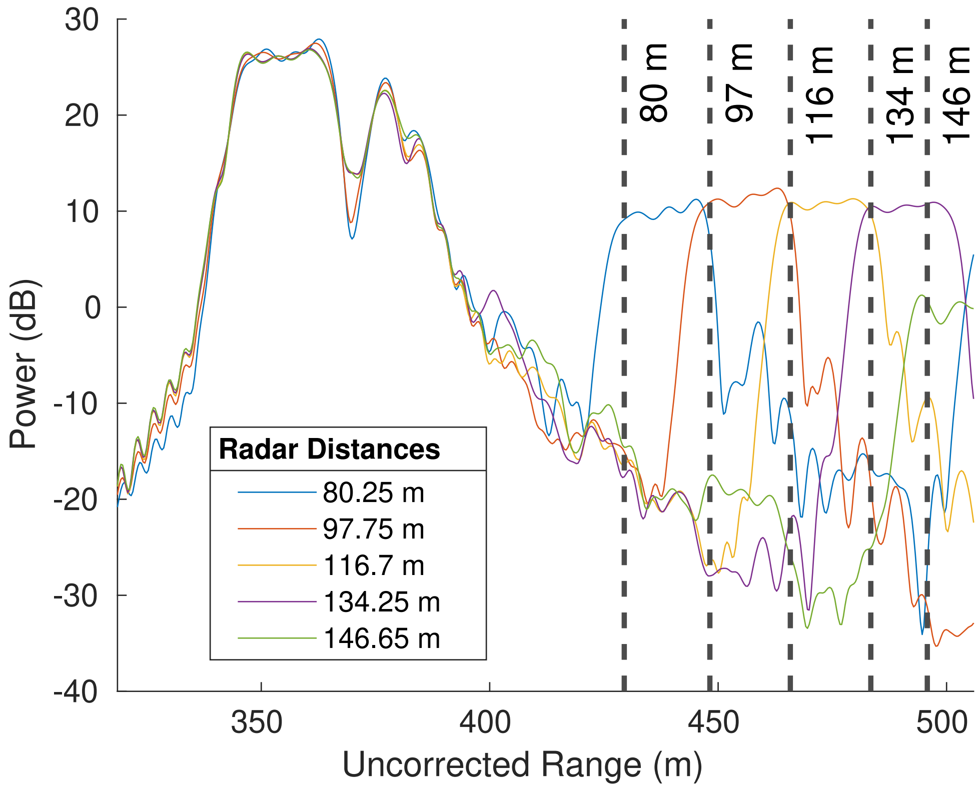

Figure 3.12: Non-coherent sum (CPI = 256) of large semi-truck at various ranges

Chapter 4

Signal Processing

4.1 Adaptive Digital Downconversion

The misalignment of fast-time spectral content between individual pulses presents a significant obstacle to coherence, particularly for targets at longer ranges. However, if an accurate estimate of the IF can be produced for each pulse within a CPI, a pulse-dependent frequency shift can be applied to remove time-dependent variability in the baseband signals. Regrettably, the frequency-domain characteristics of real-world radar echoes are highly scene-dependent and difficult to characterize statistically in a generalizable manner. To address this challenge, this work developed the following two pulse-by-pulse IF estimation techniques that operate without prior knowledge of the contents or properties of the radar scene.

4.1.1 Power Spectral Estimation

The first proposed method estimates the IF by examining the power spectrum of an individual signal. The method is also denoted in this work as the PSD IF estimation technique. In this approach, the IF is approximated as the frequency bin corresponding to the maximum spectral magnitude, and can be found with

$$f_{\mathrm{IF,est}} = \arg\max_{f} \left|\mathcal{F}\{s(t)\}\right|^2 \tag{4.1}$$

where $\mathcal{F}\{\}$ denotes the discrete Fourier transform operation and $s(t)$ represents the fast-time voltage samples of a single, initially filtered pulse. An example of the resulting IF estimates is shown in Figure 4.1, overlaid on the fast-time spectral content across a CPI of 8192 pulses.

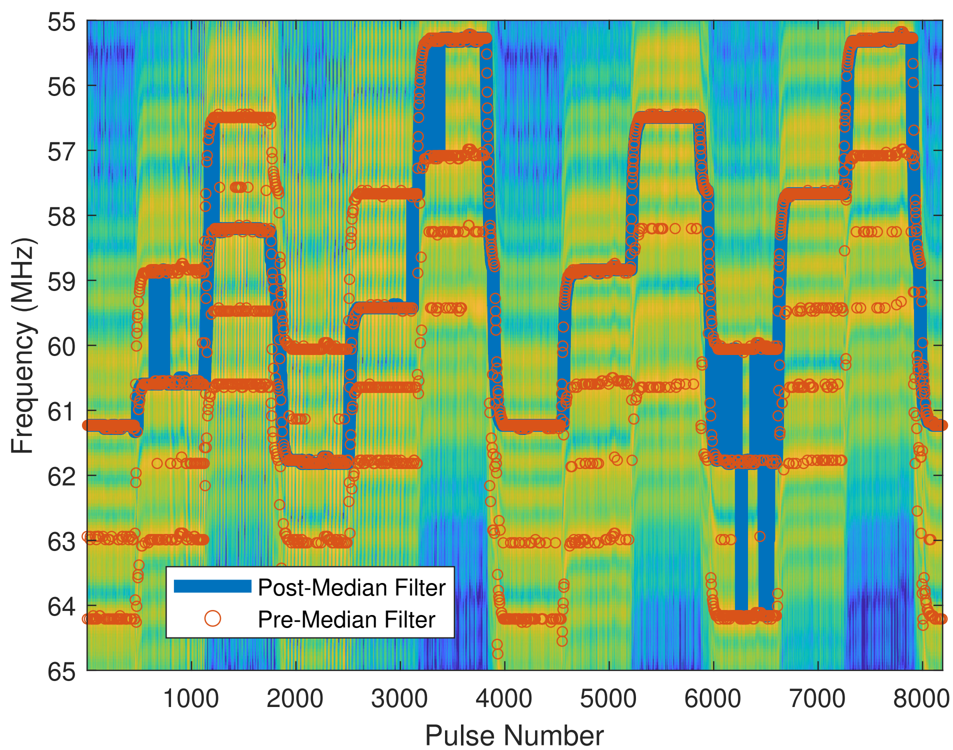

Figure 4.1: Power spectral estimates of drifting IF (pre and post median filter), reprinted from [1]

However, the in-band spectral distribution of real-world radar returns is sensitive to small deviations in the radar scene, and in some cases has been observed not to possess a centralized main lobe of spectral energy. When this energy is not concentrated in a dominant lobe, the algorithm performance can become unstable and subject to large discontinuities between adjacent frequency estimates. Since the underlying system and visible fast-time spectral distribution do not exhibit such abrupt IF jumps, these discontinuities can be attributed to estimation error. Such large estimation errors wreak havoc on the measured coherence of the system and must be mitigated.

A classic choice for reduction of outliers is the median filter, which can be applied to the IF estimates to produce a more robust and consistent set of frequency offsets, as shown in Figure 4.1. However, it can be seen that in the displayed dataset, the outliers are particularly common, resulting in final IF estimation values that maintain sharp discontinuities even after filtering. These residual artifacts are unacceptable for a coherent system, motivating the development of an alternative estimation method.

4.1.2 Template Matching Estimation

The second approach estimates pulse-to-pulse IF variation by quantifying relative frequency shifts rather than independently estimating the IF for each pulse. This method assumes that the fast-time spectral content of the radar scene remains approximately stationary over the CPI. Violations of this assumption, such as with an excessively long CPI or highly dynamic objects within the radar scene, can degrade IF estimation performance.

The algorithm proceeds as follows. First, the received signals are converted to a pseudo-baseband representation with non-adaptive digital downconversion and subsequent low-pass filtering. A reference template is then constructed from the conjugate, frequency-reversed, fast-time spectral content of an arbitrary pulse within the CPI. Next, the template is convolved with the spectral representation of each pulse, and the frequency bin corresponding to the maximum of the convolution is identified. This frequency bin provides an estimate of the relative IF shift between the selected template and the ith pulse. The absolute IF of the reference pulse is then estimated using the PSD method described in Equation (4.1). The final adaptive IF estimate for each pulse is then obtained by adding the estimated relative offset to the template pulse IF estimate. Formally, the method is expressed as

$$S_i[f] = \mathcal{F}\{s_i[t]\}, \tag{4.2}$$

$$f_{\text{IF,est}} = \arg\max_f (S_i[f] * S^*_{\text{temp}}[-f]) \quad + \quad \arg\max_f |S_{\text{temp}}[f]|^2 \tag{4.3}$$

where $S_{\text{temp}}[f]$ denotes the discrete spectral content of the reference pulse and $S_i[f]$ corresponds to the spectral content of the ith pulse. This template matching approach provides a more robust estimate of the IF compared to the PSD-based method. An example of its performance is shown in Figure 4.2, using the same dataset as in Figure 4.1. The results demonstrate a significant reduction in estimation variability and discontinuities. Although this method greatly improves estimation stability, it remains susceptible to occasional errors. As with the PSD-based approach, applying a moving median filter to the sequence of IF estimates enhances robustness and ensures smoother pulse-to-pulse frequency tracking.

4.2 Coherency Algorithms

The inherent incoherency of a radar system using a magnetron arises from a random and uncontrollable initial phase associated with each transmitted pulse. This randomness is incorporated into the transmit signal model with

$$\bar{x}(t) = A(t)e^{j(2\pi F_c t + \phi_m)} \tag{4.4}$$

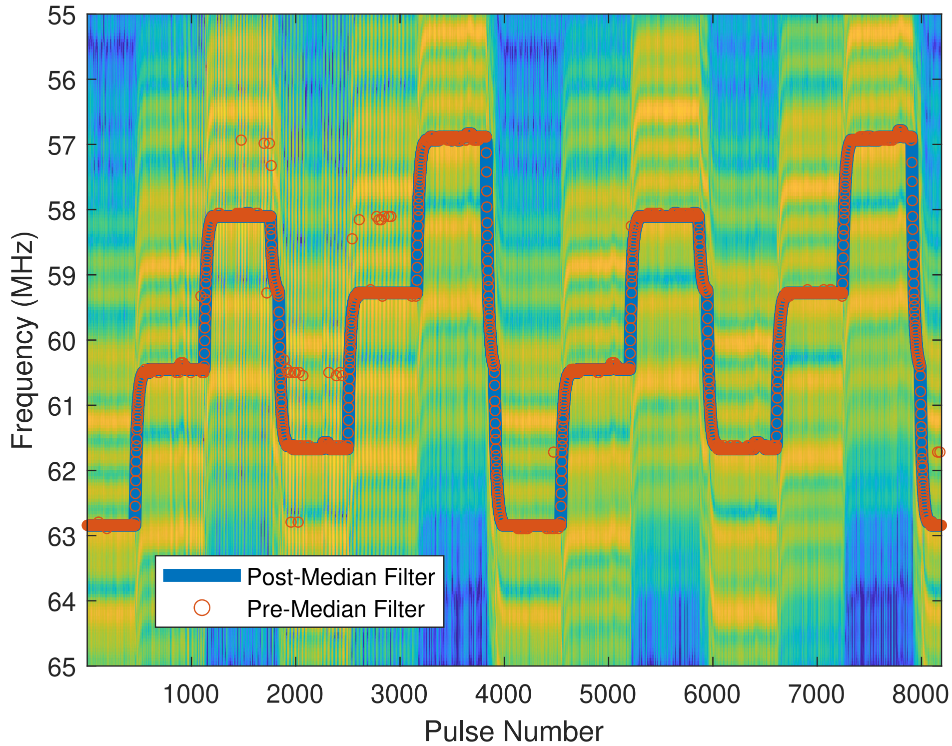

Figure 4.2: Template matching estimates of drifting IF (pre and post median filter),
reprinted from [1]

where ϕ_m is the initial phase for the mth pulse. If ϕ_m can be accurately estimated and compensated with signal processing, the radar system can effectively operate coherently despite the magnetron's intrinsic incoherency. Using a transmit pulse model derived via the methods described in Section 3.4, many techniques for phase correction are possible, but this work examines only three.

4.2.1 Gradient Ascent

Gradient ascent is a machine learning optimization technique used to maximize a multivariate function by iteratively updating parameters in the direction of steepest ascent. Analogous to fitting a model in linear regression, it utilizes the gradient to find the set of variables that maximizes a user-defined objective, often called a cost function. The gradient is written as a vector of partial derivatives with respect to each parameter included in the multivariate function and is equal in size to the set of variables. The learning rate of the algorithm, γ, is a typically small value that represents the step size taken in the direction of each partial derivative on every iteration of the algorithm. Proper selection of γ is critical, without definitive rule-sets for how it should be chosen. A large γ accelerates convergence but may prevent the algorithm from converging or lead to overshooting of the true maximum. Conversely, a small learning rate will be more likely to be caught in local maxima and will approach the steady-state more slowly, raising computational cost.

The algorithm update rule for the ith parameter is given by

$$\theta_{i,\text{new}} = \theta_{i,\text{old}} + \gamma \frac{\partial z}{\partial \theta_i} \tag{4.5}$$

where $\theta_{i,\text{new}}$ is the new estimate of the ith variable for an iteration, $\theta_{i,\text{old}}$ is the previous estimate of the ith variable ($\theta_{i,\text{new}}$ of the iteration before), and $\frac{\partial z}{\partial \theta_i}$ is the partial derivative of the function z with respect to the variable θ_i. This iterative process repeats for a set number of iterations or until a convergence threshold is met. The cost function can be chosen to maximize a wide variety of measurements and metrics, making the gradient ascent algorithm a very useful tool.

In a coherent radar system, the baseband phase of returns from transmit leakage and stationary clutter should remain constant across pulses. Accordingly, the cost function in this work is designed to maximize the coherent summation of M samples across N pulses using a constant phase correction per pulse, given by

$$\mathrm{z}(\bar{\theta}) = \sum_{m=1}^{M} \left| \sum_{n=1}^{N} s_{nm} e^{j\theta_n} \right| \tag{4.6}$$

where s_{nm} represents the mth sample from the nth pulse of the transmit model waveform, θ_n is the phase-corrective term for the nth pulse, and $\bar{\theta}$ is the vector of all phase corrections within the CPI. The gradient is then found by computing the partial derivative of each θ_n. This is derived from the maximized function in the following steps:

$$\frac{\partial z}{\partial \theta_i} = \frac{\partial}{\partial \theta_i} \sum_{m=1}^{M} \left| \sum_{n=1}^{N} s_{nm} e^{j\theta_n} \right| \tag{4.7}$$

which can be rewritten using the Euclidean Norm as

$$\frac{\partial z}{\partial \theta_i} = \sum_{m=1}^{M} \frac{\partial}{\partial \theta_i} \sqrt{\left(\sum_{n=1}^{N} s_{nm} e^{j\theta_n} \right) \left(\sum_{n=1}^{N} s_{nm} e^{j\theta_n} \right)^*}. \tag{4.8}$$

Applying the Chain rule, the partial derivative becomes

$$\frac{\partial z}{\partial \theta_i} = \sum_{m=1}^{M} \frac{\frac{\partial}{\partial \theta_i} \left(\sum_{n=1}^{N} s_{nm} e^{j\theta_n} \right) \left(\sum_{n=1}^{N} s_{nm} e^{j\theta_n} \right)^*}{2\sqrt{\left(\sum_{n=1}^{N} s_{nm} e^{j\theta_n} \right) \left(\sum_{n=1}^{N} s_{nm} e^{j\theta_n} \right)^*}}. \tag{4.9}$$

Since the partial derivative is only affected by terms containing θ_i, the numerator of Equation (4.9) can be simplified using the Product Rule to

$$-j \sum_{n=1}^{N} s_{nm} s_{im}^* e^{j(\theta_n - \theta_i)} + j \sum_{n=1}^{N} s_{nm}^* s_{im} e^{j(\theta_i - \theta_n)} \tag{4.10}$$

which is the sum of two conjugates and thus can be rewritten with

$$2 \ Re \left\{ j \sum_{n=1}^{N} s_{nm}^* s_{im} e^{j(\theta_i - \theta_n)} \right\}. \tag{4.11}$$

Putting everything together, the final partial derivative can be written with

$$\frac{\partial z}{\partial \theta_i} = \sum_{m=1}^{M} \frac{Re \left\{ j \sum_{n=1}^{N} s_{nm}^* s_{im} e^{j(\theta_i - \theta_n)} \right\}}{\left| \sum_{n=1}^{N} s_{nm} e^{j\theta_n} \right|}. \tag{4.12}$$

With this formulation, the gradient ascent algorithm can be implemented directly. In this work, it was decided to use $\gamma = 0.05$ and 100 iterations. The convergent behavior of the algorithm is illustrated in Figure 4.3, where the phase corrective terms and cost function are plotted as a function of iteration for a CPI of 1024 pulses. It can be seen that the cost function increases monotonically as the phase corrections migrate toward values that maximize coherent integration. As expected, the convergence becomes asymptotic because as variables near their steady-state values, their gradient decreases until there is a negligible effect on subsequent iterations. The resulting phase corrections approximate the conjugate of the magnetron-induced phase offsets and can be compensated by applying a pulse-dependent phase rotation to every sample in each respective pulse, resulting in a coherent system.

However, the gradient ascent algorithm has notable limitations. The algorithm cannot distinguish between local and global maxima, so if it begins near or moves initially towards a local maximum, it can become entrenched in that value when there exists a much higher maximum elsewhere in the vector space. It is also difficult to choose a learning rate prior to experimentation, and the optimal step size is not a constant value for every scenario. Perhaps the most important drawback of the gradient ascent approach for this work is the computational complexity of the algorithm. Even with modest iteration counts and large learning rates, the method took significantly longer than other coherence algorithms and was too computationally intensive for real-time implementation.

4.2.2 Phase Difference Averaging

The Phase Difference Averaging (PDA) algorithm is a computationally efficient alternative to the gradient ascent approach that leverages the same assumption of

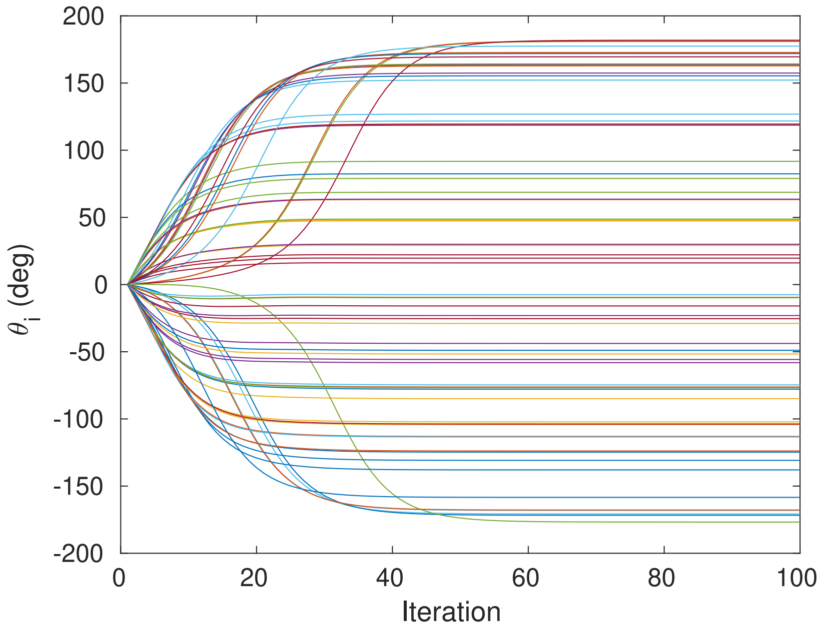

(a) Phase corrective terms

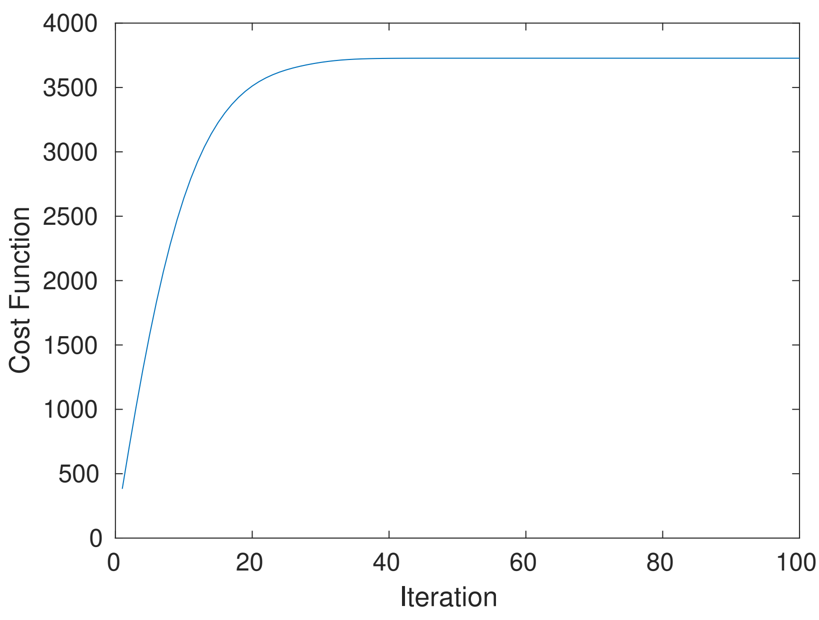

(b) Cost function $z(\bar{\theta})$

Figure 4.3: Gradient ascent convergence indicators as function of iteration

phase consistency across pulses for stationary returns and transmit power leakage. The PDA algorithm also estimates and then compensates for the magnetron initial phase offset, but does not require iterative optimization, making it well-suited for real-time or pseudo-real-time. In this work, pseudo-real-time will be defined as computationally swift enough to handle the throughput of measured data, regardless of the required latency.

The algorithm operates on the baseband representation of the transmit signal model, including considerations taken for non-stationary IF. For each pulse, the phase difference relative to a chosen reference pulse is computed for all samples. The name "phase difference averaging" is slightly misleading because, rather than calculating the mean, the final estimated phase offset is found with the median of the phase differences for each sample. Formally, the algorithm is written with

$$\tilde{\theta}_i = \text{med}\{\angle s_{i,m} - \angle s_{\text{ref},m}\} \quad m = 0, 1, ..., M - 1 \tag{4.13}$$

where $s_{i,m}$ is the transmit signal model of the ith pulse, $s_{\text{ref},m}$ is the signal model of the reference pulse, and M is the number of samples in the transmit pulse model. The resulting $\tilde{\theta}_i$ values provide per-pulse phase corrections, which are applied via multiplication by $e^{-j\tilde{\theta}_i}$.

This approach yields phase correction estimates that closely match those obtained via gradient ascent while significantly reducing computational burden. A comparison of the phase corrective values between the PDA and gradient ascent algorithm for 128 pulses is shown in Figure 4.4. It is shown that the two methods of estimating the phase corrective terms produce nearly identical results, despite the PDA algorithm executing significantly quicker. For example, with a CPI of 128 pulses, the PDA algorithm was approximately 121 times faster than the gradient ascent computation. In another test with a CPI of 1024 pulses, the PDA algorithm speed improvement increased to over 470 times faster. This substantial efficiency gain makes the PDA algorithm a practical solution for real-time imple-

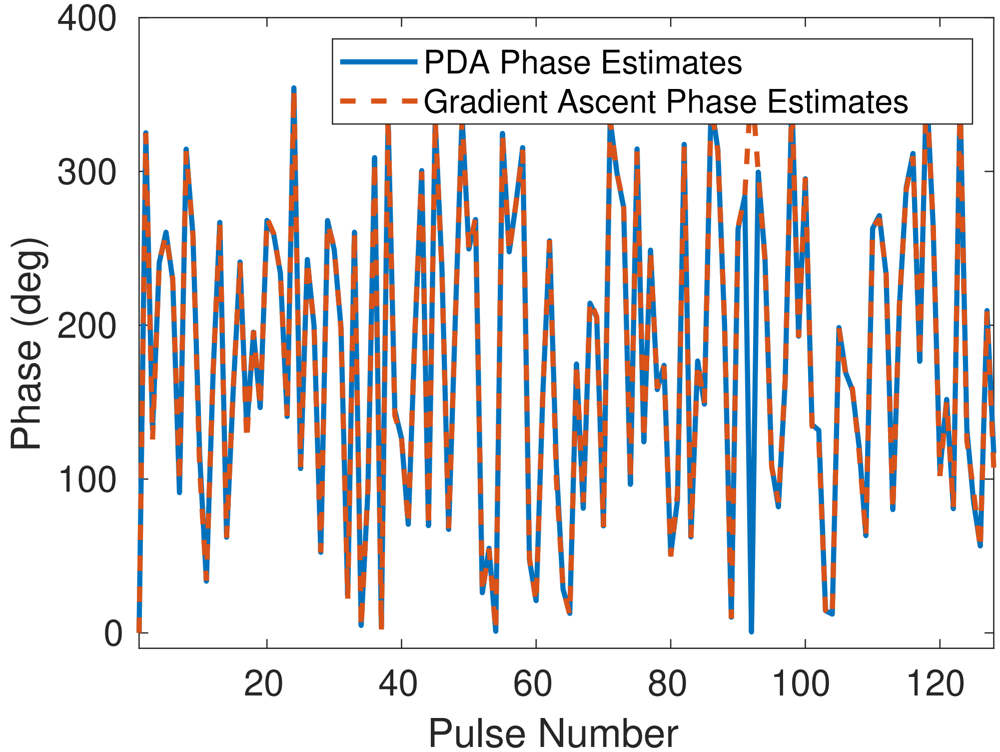

Figure 4.4: Comparison of phase corrective terms from PDA and gradient ascent algorithms

mentations.

4.2.3 Coherent-On-Receive Matched Filtering

The third approach is a coherent-on-receive (COR) matched filtering technique implementing the ideas proposed in [7]. Similar to the other methods, this technique requires accurate knowledge of a transmit pulse model, but unlike previous methods, it does not explicitly estimate the magnetron-induced phase offsets. Instead, it removes them implicitly in a computationally efficient manner using a convolution of the conjugated, time-reversed transmit signal model. The mathematical justification for the method is as follows. If the transmitted pulse model is written as

$$s(t) = q(t)e^{j2\pi F_c t}e^{j\phi_m} \tag{4.14}$$

then the return signal from the radar environment with delay τ can be written with

$$r(t) = q(t-\tau)e^{j2\pi F_c(t-\tau)}e^{j\phi_m} \tag{4.15}$$

where q(t) is the pulse amplitude envelope, F_c is the carrier frequency, and ϕ_m is the pulse-dependent initial phase offset. After demodulation, the baseband signals become

$$s_d(t) = q(t)e^{j\phi_m} \tag{4.16}$$

and

$$r_d(t) = q(t-\tau)e^{-j2\pi F_c\tau}e^{j\phi_m}. \tag{4.17}$$

If $s_d(t)$ is time-reversed and conjugated, it can be convolved with $r_d(t)$ to produce

$$r_d(t) * s_d^*(-t) = q(t-\tau)e^{-j2\pi F_c\tau}e^{j\phi_m} * q(-t)e^{-j\phi_m} \tag{4.18}$$

where the ϕ_m terms can be simplified to obtain the final result of

$$r_d(t) * s_d^*(-t) = e^{-j2\pi F_c\tau}q(t-\tau) * q(-t), \tag{4.19}$$

demonstrating that the phase term ϕ_m cancels completely, provided the transmit pulse model is accurate.

This method is computationally efficient and inherently robust, as each pulse is processed independently without a reference pulse. Each pulse is made coherent with information from its own adaptive transmit model, resulting in a radar system devoid of magnetron-induced phase offsets. Take, for example, a situation without proper downconversion and some remaining f_{err} on the baseband signal. In this scenario, using the same variables as Equations (4.14 - 4.19), the algorithm performance is written

$$s_d(t) = q(t)e^{j(2\pi f_{err}t+\phi_m)}, \quad r_d(t) = q(t-\tau)e^{-j2\pi F_c\tau}e^{j(2\pi f_{err}t+\phi_m)} \tag{4.20}$$

so

$$\begin{aligned} r_d(t) * s_d^*(-t) &= q(t-\tau)e^{-j2\pi F_c\tau}e^{j(2\pi f_{err}t+\phi_m)} * q(-t)e^{-j(-2\pi f_{err}t+\phi_m)} \\ &= e^{-j2\pi F_c\tau}\left(e^{j2\pi f_{err}t}q(t-\tau)\right) * \left(q(-t)e^{j2\pi f_{err}t}\right). \end{aligned} \tag{4.21}$$

This shows that the COR algorithm remains effective at handling the magnetron-induced ϕ_m term despite the presence of residual time-dependent frequency errors, provided those errors are present in both the transmit signal model and receive data. An additional benefit of COR matched filtering is the associated signal processing gain, whereby the energy of a received pulse is concentrated in range. While less pronounced than in pulse compression waveforms such as LFM or phase-coded signals, this gain still improves detectability in conventional simple pulse systems.

4.3 Coherence Metrics

There is a clear distinction between fully coherent and non-coherent range-Doppler maps, but assessing the degree to which an RDM is coherent becomes much more difficult to observe. Consequently, it is necessary to establish well-defined metrics to quantitatively and qualitatively evaluate coherence. In this study, coherence was assessed using a combination of numerical measures and observational criteria.

4.3.1 Phase Variance

The first metric is based on evaluating the phase relationships across pulses of selected range bins containing a stationary scatterer. Assuming the signal returns reflect exclusively from stationary objects, the phase of returns in a coherent system within a given range bin would remain constant across the CPI. The phase variance is therefore computed as

$$Var(\phi) = E\{(\phi - \bar{\mu}_m)^2\} \tag{4.22}$$

where ϕ is the phase term for the range bin of interest and $\bar{\mu}$ is the mean of the ensemble of ϕ. Prior work has used the standard deviation of phase for strong stationary scatterers to quantify the coherence of their pseudo-coherent system [7], providing a useful benchmark for comparison. This metric was widely used throughout this study.

Phase Variance Instabilities

Despite its utility, phase variance is an imperfect metric, prone to several inherent limitations. First, the assumption that a range bin contains only stationary scatterers will never be met in practice, especially at higher frequencies. In an attempt to mitigate this issue, experiments in this study considered physically large and high RCS targets (metal sheets, large buildings, radio towers, large jumbo-trons, etc) that might better resist movement. The wavelength of the carrier frequency in this particular system is only 0.0319 m, which means that even minuscule movements from the scatterers under consideration will result in measurable phase shifts in the received data. In these cases, the phase variance will be artificially inflated regardless of the coherence achieved by the proposed algorithms.

A second limitation arises most often in long CPI scenarios, where the non-stationary characteristics of the IF drift become a critical issue. The standard phase variance metric produces favorable results for highly coherent range bins, but the metric produces unrepresentative results as coherence degrades due to the large discontinuities caused by phase wrapping. This phase wrapping also affected highly coherent datasets, but could be avoided with strategic range bin choices, phase unwrapping, or a constant phase shift and modulus operation across pulses. However, in range bins that were not highly coherent, there was no simple solution to the phase wrapping issue. As an example, Figure 4.5 demonstrates the issues created by phase wrapping from a single range bin of a large scatterer 1760 m from the antenna, across a long CPI. It can be seen that when the phase shifts

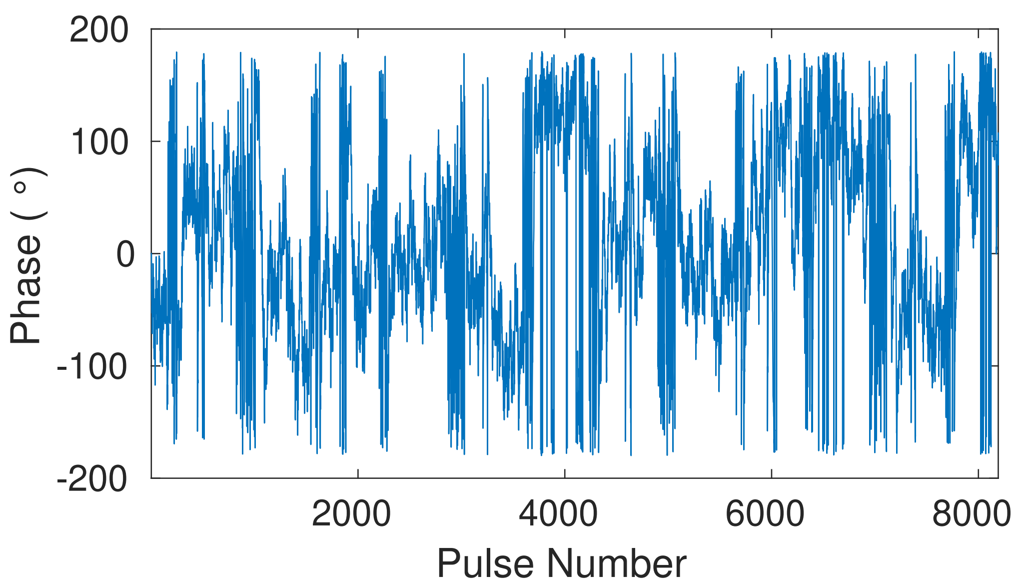

Figure 4.5: Long CPI phase vs. pulse number with wrapping

between ±180°, the traditional variance values give an inaccurate portrayal of the coherence of the system. Furthermore, small IF estimation errors introduce pulse-to-pulse phase inconsistencies that grow with range, creating range bins where the

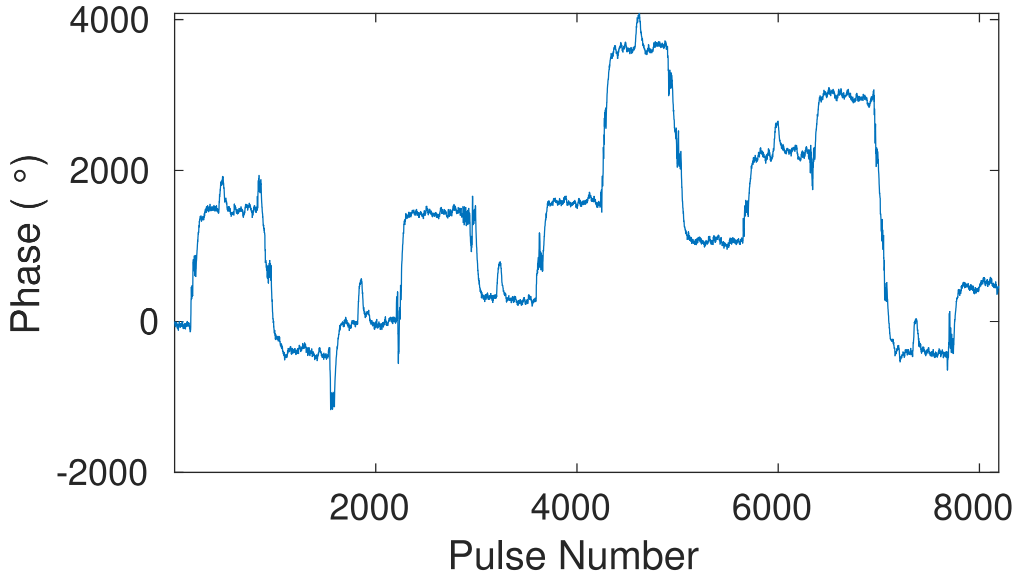

Figure 4.6: Long CPI phase vs. pulse number unwrapped

phase changes enough that previous unwrapping solutions are ineffective. As an example, the unwrapped phase of Figure 4.5 is shown in Figure 4.6. It is shown in the unwrapped version that there are periods of consistent phase stability interrupted by large shifts in phase estimates, and it was found that those periods coincide with the stable and transitory periods of the IF drift. Attempts were made to quantify IF estimation errors using gradient ascent and linear regression on long-CPI phase comparisons without success. In the absence of adequate compensatory measures, it was clear that the development of more robust coherence metrics was necessary. Nevertheless, phase variance remained a useful metric for many experimental conditions and was used extensively in testing and analysis.

4.3.2 Angular Variance

To address the limitations associated with wrapping in standard phase variance, a second metric based on directional statistics was introduced. Unlike linear statistics, directional statistics attempts to account for the periodic nature of angular data, which is a non-Euclidean space. The typical arithmetic mean $\frac{1}{n}\sum_{i=1}^{n}\theta_i$ does not always produce satisfactory results from a geometric standpoint. For example, angles of $[-179°, 179°]$ would arithmetically have a mean of 0° while a truer geometric mean would be 180°. To account for this, the circular mean, as defined by [34], is found with

$$C = \sum_{i=1}^{n} \cos(\theta_i), \quad S = \sum_{i=1}^{n} \sin(\theta_i), \quad R^2 = C^2 + S^2 \quad (R \geq 0) \tag{4.23}$$

and

$$\cos\bar{\theta} = C/R, \quad \sin\bar{\theta} = S/R \tag{4.24}$$

where θ_i is the ith angle measurement, n is the number of samples and $\bar{\theta}$ is the sample circular mean. Traditionally, the circular variance is then found with

$$V = 1 - R/n \tag{4.25}$$

where V is the sample circular variance. However, the conventional circular variance did not fully encapsulate the grouping of angular data in a way that was satisfactory for this project. For example, if a set contains the angular measurements shown in Figure 4.7, the circular mean would be zero and the circular variance would have the maximum value of 1, even though the data is not a maximally dispersed distribution [34]. There exist several alternative algorithms for calculating angular dispersion, specifically for deviations in wind direction [35], [36], that inspired the proposed metric for measuring coherency of a radar system, which is denoted as angular variance. In this work, the angular variance of

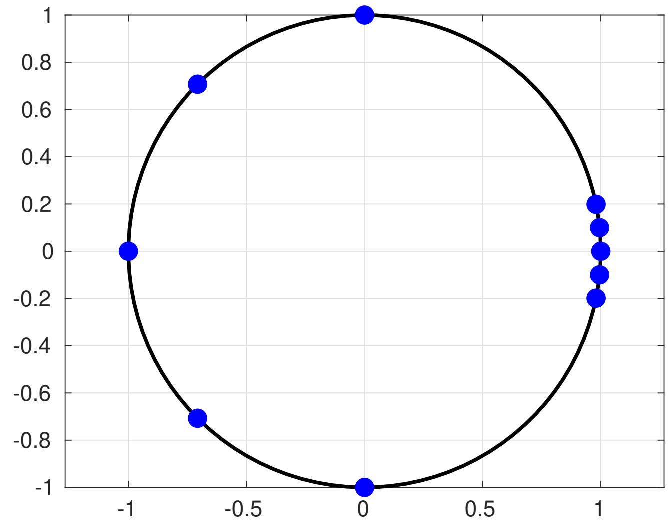

Figure 4.7: Angular phase variance example unit circle

a sample set blends linear variance with the circular mean and is written with

$$\sigma_\theta = \frac{1}{n-1}\sum_{i=1}^{n}\Delta_i^2 \tag{4.26}$$

where Δ_i is defined as the minimum of $|\theta_i - \bar{\theta}|$, $2\pi - |\theta_i - \bar{\theta}|$, and $|\theta_i - \bar{\theta}| - 2\pi$. In this way, the angular variance mitigates the discontinuities introduced by phase wrapping and provides a more stable measure of dispersion for angular data. As such, it was used extensively in this study, particularly in scenarios where traditional phase variance proved unreliable.

Angular Variance Disadvantages

The angular variance is not without its own limitations. As with traditional phase variance, the motion of scatterers relative to the antenna during a CPI increases angular dispersion regardless of the performance of the proposed algorithms. Additionally, the calculation of the circular mean may yield unintuitive results when phase wrapping is not present. For example, in a dataset containing angle measurements of $[0°, 0°, 90°]$, the circular mean would be 26.57° while the arithmetic mean would be 30°, which is a more representative central value. As a result, the standard phase variance was preferred in highly coherent datasets, while angular variance was employed when wrapping effects were significant. It should be noted that using a minimized angular difference term for traditional phase variance would assist with some of the phase wrapping issues, but was mostly unnecessary for the highly coherent range bin analysis conducted in this study.

4.3.3 Peak Doppler Sidelobe Levels

A third metric evaluates coherence based on phase variation in a single range bin indirectly through Doppler-domain characteristics. For a stationary scatterer, the upsampled Doppler response should approximate a sinc function consistent with the Fourier transform of a rectangular window. The peak sidelobe levels of a sinc function occur at -13.2 dB relative to the main lobe, irrespective of rectangular window length. Thus, by analyzing Doppler slices corresponding to selected range bins, the deviation of peak sidelobe levels from this theoretical value provides a quantitative assessment of coherence.

An example of a coherent, partially-coherent, and non-coherent sidelobe structure using the same range bin is shown in Figure 4.8. The examples shown used a CPI of 256 pulses with an upsampling factor of eight. Special care was taken to choose a point at which the IF was drifting heavily within the CPI, so the coherent trace uses the template matching IF estimation and the COR method, the partially coherent trace did not use adaptive digital downconversion but did use the COR method, and the noncoherent trace is formed without any coherency algorithms or adaptive digital downconversion. Note that the peak sidelobe levels

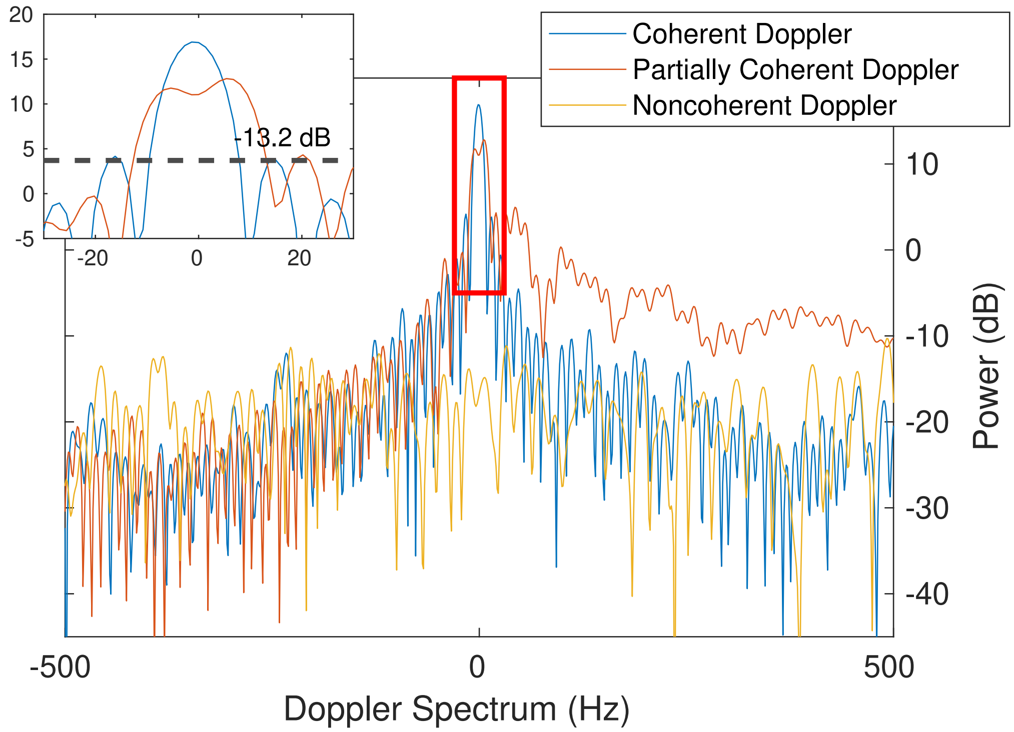

Figure 4.8: Comparison of Doppler sidelobes for coherent, partially coherent, and noncoherent results

of the coherent trace are very close to the -13.2 dB boundary, with a measured level of -12.7 dB, a small difference that can be attributed to slight misestimations of the transitory IF. The peak sidelobe levels of the partially coherent trace appear at similar power levels to that of the coherent trace, but the mainlobe peak is much lower, resulting in a peak sidelobe level of only -7.86 dB. There is no main lobe in the noncoherent data, and so the sidelobe level metric is not applicable. While this method is also sensitive to scatterer motion, it provides a complementary perspective on system performance.

4.3.4 Moving Target Energy Signature

A primary objective of this work was to enable reliable detection and characterization of moving targets within varied clutter environments for aeroecological studies. Coherent processing facilitates the separation of moving targets, such as birds, insects, or other small objects, from other powerful radar signal returns. A variety of moving targets were analyzed, including drones, projectiles, vehicles, and birds present in the scene. Many of these targets exhibited low RCS values that would likely be obscured by sea clutter if measured with non-coherent processing methods. However, coherent Doppler processing allowed these objects to be detected despite high-power contaminating clutter, proving the performance of the coherency algorithms.

Beyond simple detection of moving objects in the scene, the energy distribution of the targets was examined as they traversed the radar beam. The SNR of each target was particularly useful in the evaluation, as a highly coherent radar system would more tightly focus the energy of a moving object on the RDM. Additionally, certain targets exhibit micro-Doppler signatures arising from internal motion or rotation, providing further insight into target characteristics and reinforcing the benefits of coherent processing.

Chapter 5

Stationary Targets and Clutter

5.1 Coherency Algorithm Evaluation

After defining the coherency algorithms theoretically, their performance was evaluated using the metrics introduced in Section 4.3. To assess coherence, the radar system was deployed on the third floor of the Advanced Radar Research Center (ARRC) at the University of Oklahoma and equipped with the XN20AF fan-beam antenna. The antenna was manually directed toward several large, stationary scatterers with an unobstructed line of sight from the radar position. Data was collected for several seconds while the radar was pointed at each target. The recorded data was processed in MATLAB using the gradient ascent, PDA, and COR methods described in Section 4.2 for a 500-pulse CPI.

Pulses of interest were strategically selected from regions in the dataset in which the IF exhibited minimal drift, ensuring that the adaptive digital down-conversion would not skew results between the different targets. Precise range information for the selected scatterers was not available, primarily because the mainlobe beamwidth of the XN20AF antenna illuminated a large area (approximately 5.1 x 83.6 m^2 for only the nearest target), encompassing a substantial portion of each structure. As a result, the displayed ranges were determined using radar fast-time information rather than direct physical measurements. To validate these choices, a Google Earth measurement project was created to confirm that the selected radar fast-time ranges corresponded to the intended stationary scatterers, which is shown in Figure 5.1.

Figure 5.1: Large stationary target locations (Google Earth)

The phase standard deviation (square root of linear phase variance) across pulses is given in Table 5.1. It is evident from these results that the COR method consistently outperformed the other coherency methods, followed closely by the

Table 5.1: Phase Standard Deviation for Specified Target

Target	Range (m)	COR	PDA	Gradient Ascent
Nat. Weather Center	237	0.807°	1.421°	0.908°
Stephenson Research	365	1.075°	1.227°	1.176°
Stephenson Life Sciences	570	1.653°	2.642°	1.727°
Baseball Stadium	1396	3.965°	4.077°	4.084°
Sam Noble Museum	1760	4.737°	4.753°	4.758°

gradient ascent algorithm. This behavior is to be expected because, given a reasonable transmit pulse model, the COR method should remove the magnetron-induced phase offset without the risk of estimation error. The gradient ascent algorithm is expected to perform at least as well as, and potentially better than, the PDA method due to its iterative optimization of the phase corrective terms. This superior performance is reflected at the shorter range targets, although the PDA algorithm marginally outperforms gradient ascent in some of the longer range datasets. This outcome is encouraging because the PDA technique, which is the fastest and simplest algorithm, achieved phase variance values similar to those of the more computationally intensive methods.

The stationary targets were selected to be large buildings or structures whose likelihood of movement during the data collection was minimized. However, despite this precaution, it is clear that there exists non-zero phase variance for every target, and that the variance grows with increasing range. One explanation for this could be movement of the radar system itself. The FAR 1528 system was designed to rotate, so there was no locking mechanism to prevent azimuthal movements during the radiation process. Even a very small movement of the antenna during transmission would induce a Doppler shift on the transmitted signal that would affect pulse-to-pulse coherence. Another possibility comes from the area illuminated by the antenna's main beam, which was very large and grew with increasing range. If a scatterer in the range bin of interest exhibits any movement, such as trees swaying, people walking, or vehicles driving, the corresponding radar echoes will contain unequal phase shifts equivalent to the Doppler shift of those scatterers. A larger range has a larger illuminated area, which translates to an increased possibility that the measured data contains radar echoes from moving targets. The distance required for measured phase shifts between pulses is very small, 0.21 mm for even the highest measured standard deviation, so it is also possible that the targets themselves were shifting at those scales during the CPI, presenting a third potential culprit for the nonzero phase standard deviation. All possibilities would increase phase variance across pulses, regardless of the performance of the coherency algorithms. Nevertheless, the measured phase variance values remain extremely low, demonstrating effective performance of the coherence algorithms over a wide range of target distances.

Using the same experimental setup, short-CPI coherency was also examined by evaluating the peak Doppler sidelobe levels for the selected stationary scatterers, with the results summarized in Table 5.2. A perfectly coherent radar system

Table 5.2: Peak Doppler Sidelobe Levels for Specified Target (dB)

Target	**Range (m)**	**COR**	**PDA**	**Gradient Ascent**
Nat. Weather Center	237	-11.431	-11.434	-11.159
Stephenson Research	365	-12.311	-11.999	-12.016
Stephenson Life Sciences	570	-12.375	-12.346	-12.202
Baseball Stadium	1396	-11.079	-11.130	-11.049
Sam Noble Museum	1760	-9.782	-9.819	-9.757

observing a fully stationary scene is expected to produce peak sidelobe levels of approximately -13.2 dB. However, the same factors that affect phase and angular variance also degrade the reliability of this metric. With this in mind, it can be

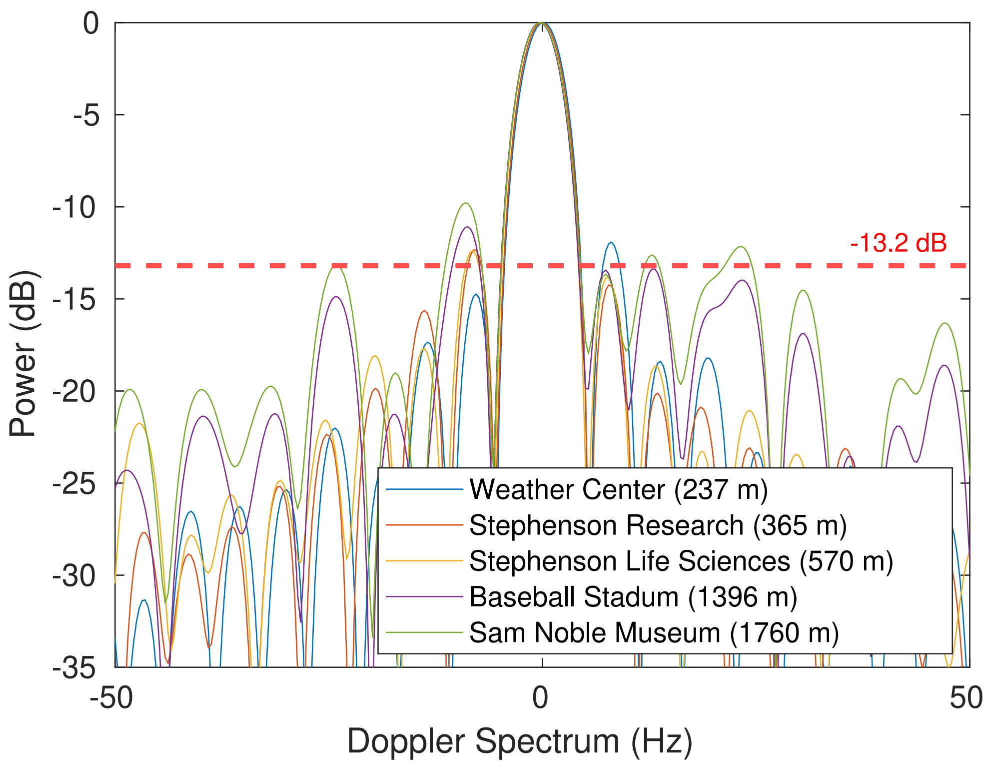

Figure 5.2: Normalized doppler sidelobes of COR method with different targets

seen that the measured peak Doppler sidelobe levels do not consistently reflect the trends observed in Table 5.1. In practice, the sidelobe metric did not reliably indicate which coherency technique was superior and was consistently observed to be a less robust measure of coherency. Furthermore, the largest observed difference in sidelobe levels between coherency algorithms, 0.275 dB, was extremely small, corresponding to only 6.5% difference in measured power. Consequently, this metric was not used in subsequent analyses, although it provided a coarse method of validating the coherence of the stationary targets. A normalized representation of the peak Doppler sidelobe levels, obtained using the COR method, is shown in Figure 5.2 for the scatterers included in Table 5.2. Note that the sidelobe levels increase with increasing range, following the pattern shown by the phase variance evaluation.

5.2 Adaptive Digital Downconversion Evaluation

5.2.1 Dominant Ground Returns

To evaluate the effectiveness of the IF estimation and compensation techniques, system coherence was assessed over a long CPI in which significant IF drift was present. This experiment utilized the same setup described in Section 5.1, using the Furuno XN20AF antenna, but extending the CPI length to 8192 pulses. In this way, the IF drift was comparable in length and structure to that shown in Figure 3.7. The system setup and experimental capabilities of this work did not make it possible to obtain ground truth about the proper values of the IF drift, complicating direct analysis of algorithm performance. As a result, algorithm performance was assessed indirectly through coherence metrics, and was thus subjected to the same limitations presented in Section 4.3.

A comparison of the phase standard deviation between non-adaptive digital downconversion, the PSD method (Section 4.1.1), and the template matching method (Section 4.1.2) is presented in Table 5.3. A median filter length of 51 pulses was applied in each case. The results demonstrate that the template matching

Table 5.3: Long CPI Standard Deviation of Phase for Specified Target

Target	Range (m)	No Adaptive	PSD	Template
Ground Plane	47	12.532°	2.552°	1.415°
Ground Plane	106	13.899°	6.065°	1.500°
Nat. Weather Center W.	237	10.931°	3.743°	1.223°
Stephenson Research S.	365	10.784°	6.129°	2.033°
Stephenson Research N.	429	12.944°	5.118°	2.341°

technique consistently yields the lowest phase variation across all targets, indicating superior phase stability and IF estimation accuracy. The PSD-based approach provides moderate improvement over non-adaptive downconversion but remains less precise than the template matching method. In contrast, the absence of adaptive digital downconversion results in substantial degradation of phase stability, often rendering the phase variance metric ineffective.

A zoomed-in RDM comparison between the non-adaptive digital downconversion and template matching method approaches is shown in Figure 5.3. The results clearly illustrate the significant impact of accurate IF estimation and compensation, with large spreading in Doppler of the stationary targets when this compensatory measure is neglected. However, while the phase variance evidence suggests that the non-adaptive case creates a fully incoherent result, Figure 5.3a suggests otherwise. There is still some semblance of structure in the RDM, whereas in a fully noncoherent system, there would be a constant spread of energy across all Doppler bins. For reference, an example of a fully non-coherent RDM is shown is Figure 5.4.

The influence of median filter length on IF estimation performance was also investigated. To arrive at an optimized median filter length, phase stability was evaluated across several representative range bins as a function of median filter length. Due to the large quantity of phase stability calculations and the manual tuning required for phase wrapping, angular variance was used in place of standard phase variance. The results are shown in Figure 5.5 for filter lengths ranging from 1 to 151. For clarity, the true variance curves for each target in the figure were shifted by its mean to emphasize the trends between phase stability and median filter length. The results indicate that as the median filter length increases, the phase variance tends to grow for the included targets. However, in many of the traces, very short filter lengths exhibit significantly elevated variance levels, suggesting reduced estimation accuracy. This behavior is expected, as in-

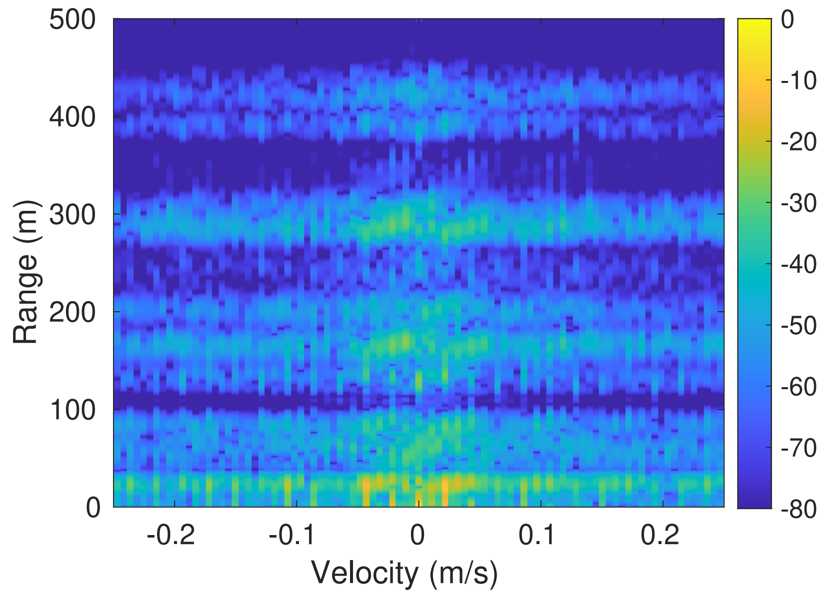

(a) No adaptive digital downconversion

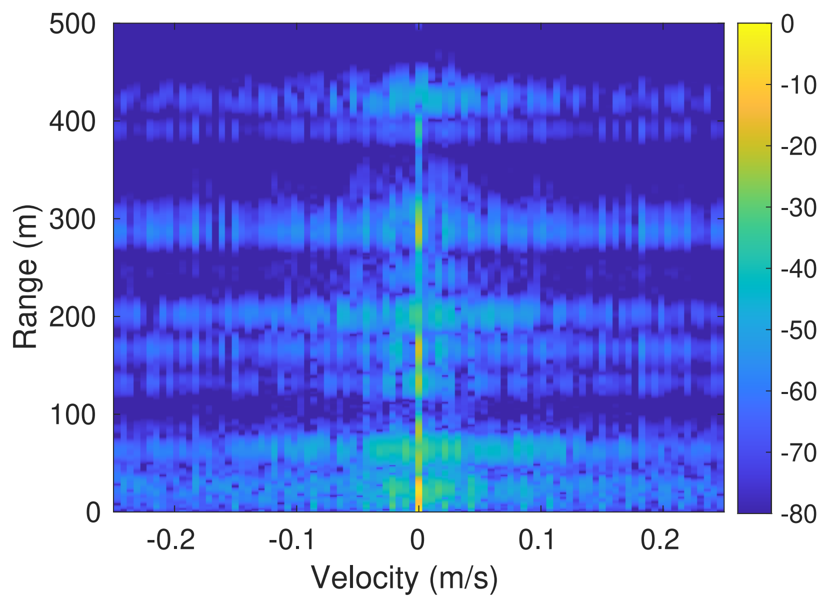

(b) Template matching algorithm

Figure 5.3: RDM comparison of adaptive and non-adaptive downconversion methods, reprinted from [1]

creasing the median filter length raises the likelihood of incorporating unrelated IF estimates from neighboring pulses, particularly for those occurring near tran-

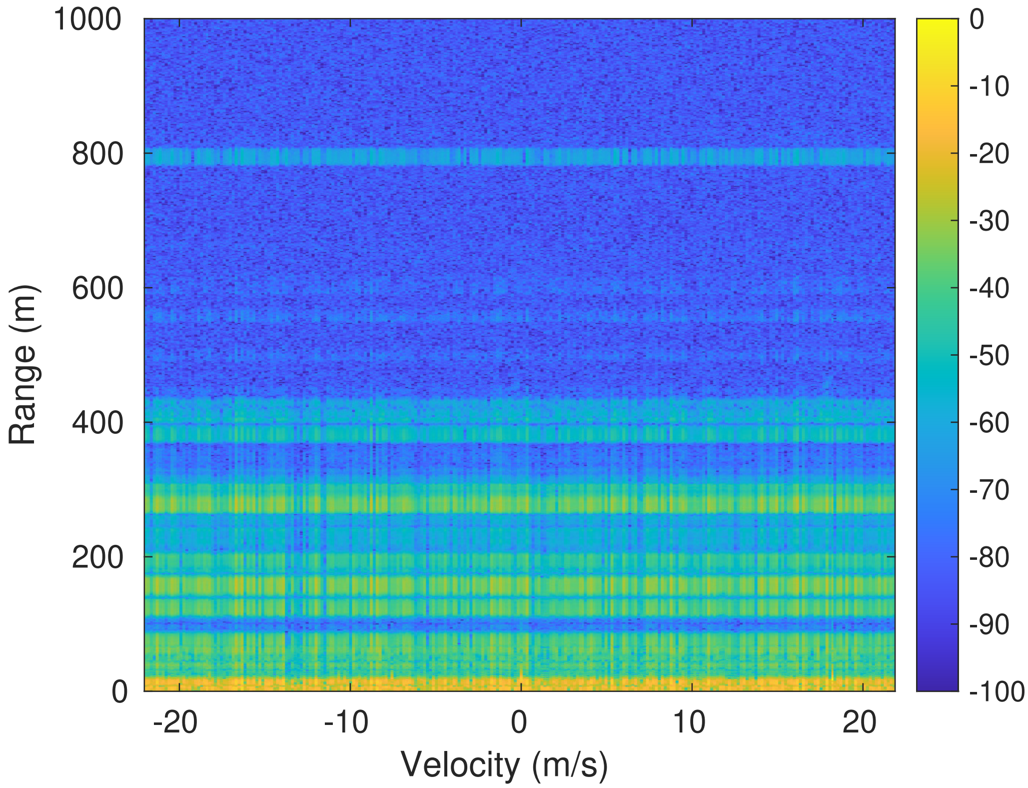

Figure 5.4: Example of fully non-coherent RDM with several strong scatterers

sitional IF drift periods. However, as discussed in Section 4.1, the IF estimation techniques are also prone to large outliers, which significantly degrade adaptive digital downconversion performance. These observations highlight a trade-off in filter design between suppressing outliers and preserving local accuracy. Based on this analysis, the median filter lengths used in this study were selected within the range of 11 to 51 pulses. It was also observed that even-length filters consistently degraded performance, likely due to asymmetric bias relative to the pulse index, and were therefore excluded.

Overall, the results strongly support the effectiveness of the proposed adaptive digital downconversion techniques. However, it is important to consider the parameters involved in creating these specific datasets. The Furuno XN20AF

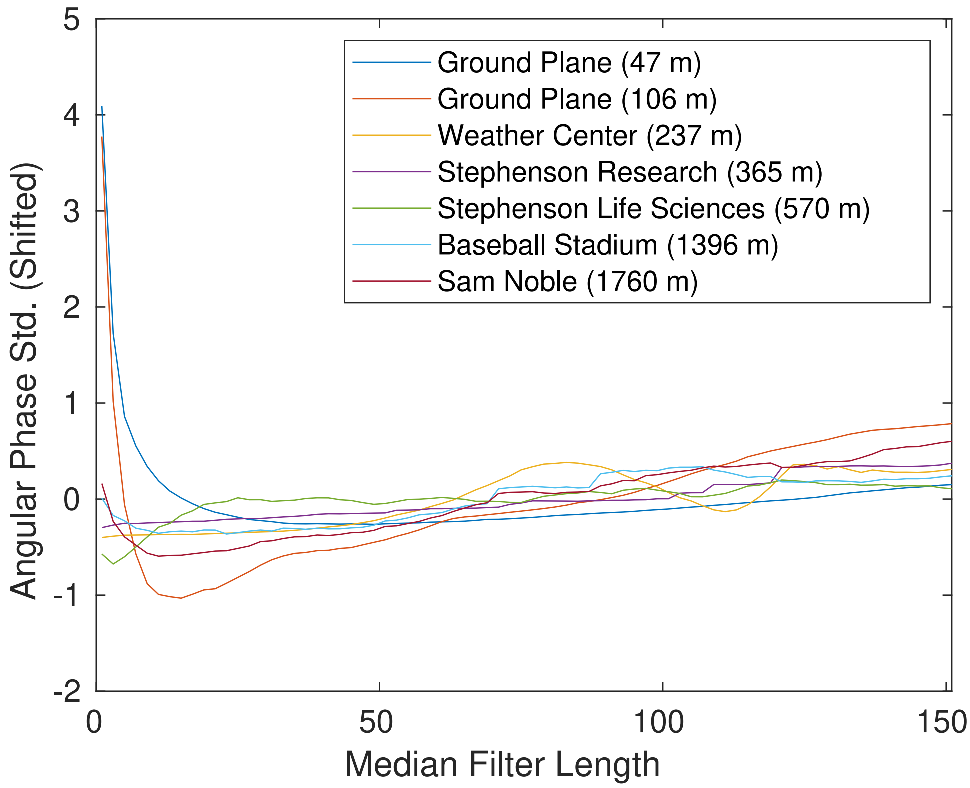

Figure 5.5: Angular phase variance of targets vs. median filter length

antenna radiates in a fan-beam shape, illuminating a large spatial region and resulting in a large quantity of energy reflected back at the radar. This can be seen in Figure 5.3b as the strong energy signature reflected by multiple ground-based targets located at 0 m/s. These dominant reflections supplement the IF estimation and improve the performance of the estimation. In scenarios lacking strong stationary scatterers or in environments with strong moving targets, the IF estimation performance experiences some degradation. As a result, further evaluation was conducted to assess algorithm performance in situations without dominant ground clutter returns.

Table 5.4: Long CPI Phase Standard Deviation (Angular) with Minimal Ground Clutter

Range (m)	PSD	PSD No Med.	Template	Template No Med.
91.1	12.502°	12.621°	2.276°	8.253°
109.8	13.007°	13.543°	5.221°	13.083°
141.3	11.315°	11.485°	2.980°	5.288°
181.8	13.371°	12.880°	3.486°	3.456°
199.4	12.146°	11.994°	4.847°	8.444°

5.2.2 Minimal Stationary Clutter

The performance of the proposed adaptive digital downconversion techniques was further evaluated in scenarios with minimal stationary clutter. For these experiments, the USGS dish antenna was directed at slightly elevated angles toward the experimental targets. Data collection utilized the same truck-mounted system described in Section 3.6, enabling flexible positioning of the radar across multiple range bins. The USGS antenna has a small vertical beamwidth, which, combined with the elevated angle, significantly minimized the clutter returns from the ground. This resulted in receiver data containing substantially less energy, making IF estimation more challenging. The system was positioned at several locations near the National Weather Center, and the angular phase standard deviation measurements, computed over an 8192 pulse CPI, are presented in Table 5.4.

A clear degradation in phase stability is observed for long-CPI experiments conducted in low-clutter environments, despite the targets being at relatively short ranges. In this dataset, the PSD method demonstrated little success, even when combined with median filtering, and failed to achieve meaningful coherence according to the metrics available. Examples of some actual estimates produced by the PSD and template matching methods with and without median filtering are presented in Figures 5.6 and 5.7. In these results, it can be seen that the

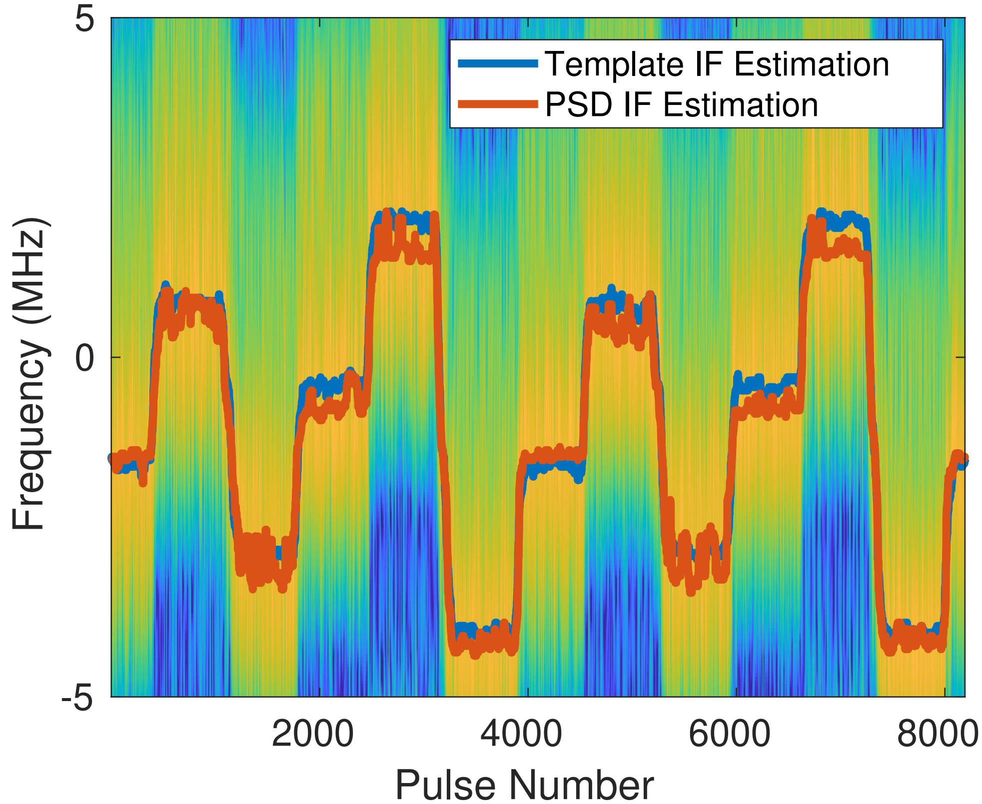

Figure 5.6: IF estimates with minimal ground clutter (with median filter)

PSD algorithm exhibits visibly more erratic behavior than the template matching approach, so for the remainder of this work, only the template matching method was considered.

The frequency response of the receiver subsystem is unaffected by the objects in the radar environment, so the true IF drift is expected to retain the same general structure as observed in Figures 4.1 and 4.2, characterized by periods of relative stability interspersed by brief transitional intervals. It is also clear that even with median filtering, there is a diminished estimation accuracy of the proposed methods, particularly within the stable portions of the IF drift pattern. This observation motivated the development of additional strategies to improve

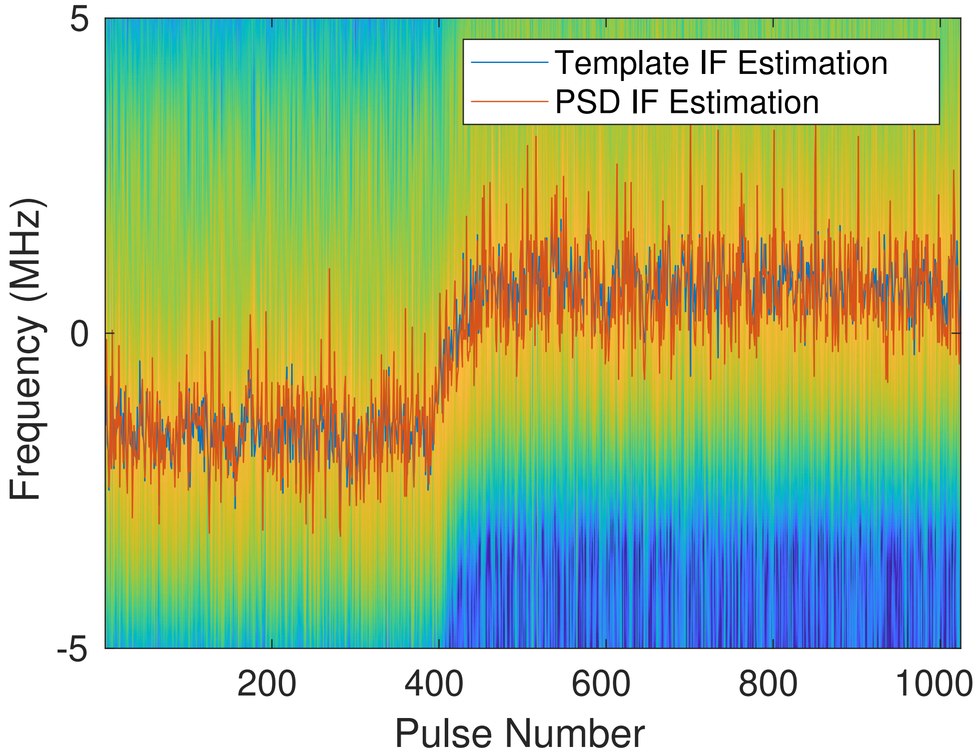

Figure 5.7: IF estimates with minimal ground clutter (no median filter)

IF estimation that were previously unnecessary.

5.2.3 IF Pattern Decision Framework

As shown in Table 5.4 and Figure 5.6, the proposed IF estimation techniques suffer in the absence of strong stationary clutter, resulting in reduced system coherence. In an attempt to address this limitation, additional estimation strategies were developed. Analysis of datasets containing dominant ground clutter revealed that stable regions of IF drift consistently contained identical frequency values, which produced the most reliable coherence metrics. In contrast, datasets with minimal clutter or strong moving objects demonstrated large variation, even in relatively stable periods of IF drift. Since the behavior of the receive chain is unaffected by the radar scene, the IF structure shown in stationary clutter-dominated datasets was assumed to better approximate the true IF drift. This observation motivated the development of a classification framework to distinguish between stable and transitional IF regions.

In the proposed framework, each pulse is classified as belonging to either a stable or transitory portion of the IF drift. Pulses identified as stable are grouped, and a constant IF estimate value of the mean of their respective estimates is assigned to the entire group. The IF estimation performance was observed to be a volatile value, with small changes in the radar environment and system operating parameters having substantial consequences. One such parameter was the median filter length of the original IF estimation, which for this investigation was chosen to be 21 pulses. Another is the window size used in each classification scheme, chosen for both methods to be 51 pulses.

Variance-Based IF Classification

The first classification scheme categorizes an IF estimate based on the variance of the surrounding IF estimates. The underlying assumption is that estimates from stable regions exhibit low variance, whereas those in transitional regions display higher variability. Thus, the variance-based method is simple, and an individual IF estimate is considered part of a stable section if

$$E\{(f_{est,w} - \bar{\mu}_f)^2\} \leq \sigma^2_{\text{threshold}} \tag{5.1}$$

where $f_{est,w}$ are the IF estimates within a window w and $\sigma^2_{\text{threshold}}$ is some threshold variance value above which it can be assumed that an IF estimate is not in a stable section. Consecutively stable IF estimates are grouped together, and their mean is chosen as the estimate for the entire ensemble. An example of the results

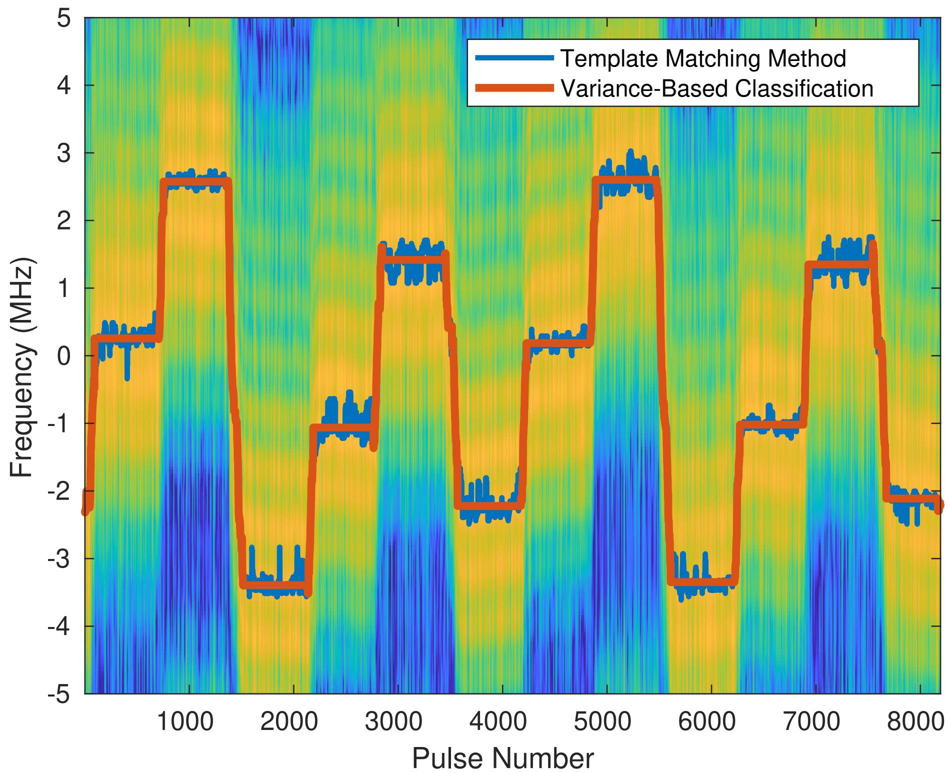

Figure 5.8: Variance-based IF classification comparison to template matching method

of the variance-based IF classification method can be seen in Figure 5.8 overlaid on the estimates of the normal template matching method. In general, the stable sections of the IF structure are successfully clustered to the same value and achieve an improved degree of coherence. However, the classification encounters some difficulty near the borders of sections in the drift structure, classifying some transitory estimates as part of the stable section and vice versa.

The information in Figure 5.8 was created using a $\sigma^2_{\text{threshold}}$ that was chosen specifically for this dataset, which required direct adjustment of the threshold. It was desirable to identify a threshold value that would yield consistent performance across multiple datasets without the need for manual tuning. To establish a baseline for optimizing $\sigma^2_{\mathrm{threshold}}$, the variance of stable IF regions was computed across a collection of datasets with minimal ground clutter. The resulting values were then averaged to obtain an empirically derived initial estimate of $\sigma^2_{\mathrm{exp}} = 5.61 \times 10^9$. The coherence of each range bin listed in Table 5.4 was then evaluated

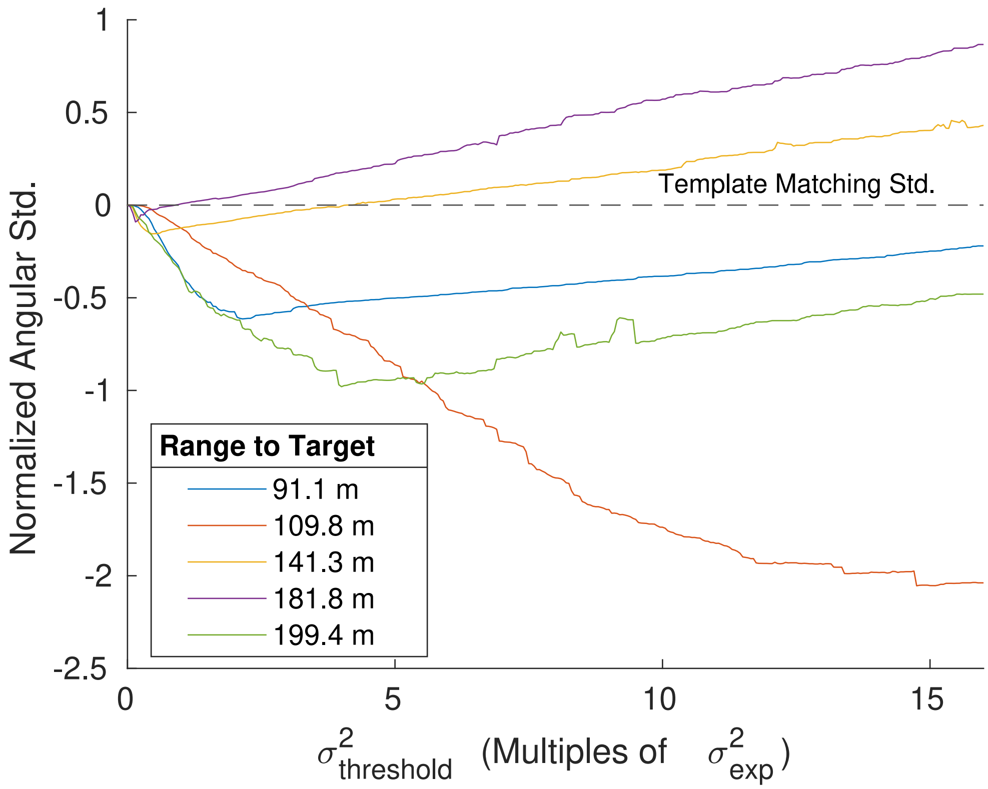

Figure 5.9: Variance-based IF classification coherency vs. threshold

using the variance-based IF classification as a function of $\sigma^2_{threshold}$, with the results shown in Figure 5.9. To produce an aligned figure, each trace was normalized by the corresponding measured angular standard deviation value included in Table 5.4. In this way, negative values indicate lower phase variance and an improved performance by the variance-based classification scheme relative to the template matching method.

Classification of stable IF regions does not universally improve performance, particularly when $\sigma^2_{threshold}$ is chosen poorly. Additionally, it was observed that the method enhances coherence within stable segments of the IF drift, but that it can degrade estimation accuracy during transitional intervals. This degradation is most pronounced in pulses near the boundaries of stable regions, which may be incorrectly grouped and assigned the value of the group instead of retaining their more accurate original estimate. Despite these limitations, the results in Figure 5.9 demonstrate that the method can produce a largely beneficial result. Choosing an optimal $\sigma^2_{threshold}$ was difficult, but analysis of the results in Figure 5.9 led to the choice of $\sigma^2_{threshold} = 1.2 * \sigma^2_{exp} = 6.7320 * 10^9$ for the remainder of this study.

Linear Regression IF Classification

The second classification approach distinguished sections of the IF structure by using the slope of a linear regression fit to local IF estimates. The premise is that IF estimates from transitional regions will exhibit larger slope magnitudes, while slopes from stable regions will be significantly smaller. Again, the framework is simple, classifying a pulse as stable if

$$\frac{d}{dn}\left(\mathcal{LR}\{f_{w,est}\}\right) \leq m_{\text{threshold}} \tag{5.2}$$

where n is the pulse number, w is the window of surrounding pulses, $\mathcal{LR}\{\}$ is the linear regression operator, and $m_{\text{threshold}}$ is the chosen slope threshold value. IF estimates classified as stable are again grouped, and their mean is chosen to be the estimate for each group. The coherence of each range bin listed in Table 5.4 was then evaluated using the linear-regression based IF classification as a function of $m_{\text{threshold}}$, with the results shown in Figure 5.10. Again, the traces were aligned by normalizing to zero with the measured angular variance of the template matching

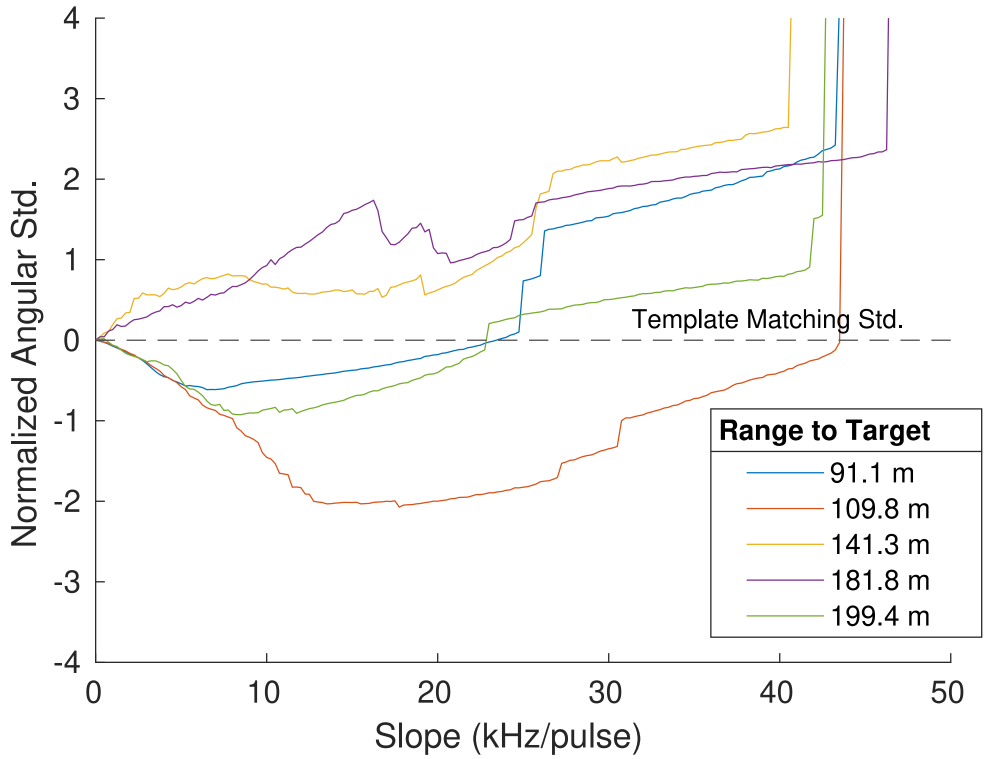

Figure 5.10: Linear regression based IF classification coherency vs. threshold

method found in Table 5.4. It is clear that the linear regression method does not produce consistent effects for each dataset and is seen to perform generally worse than the variance-based classification scheme. However, it was still important to select a generalizable threshold, which, based on the displayed results, was selected to be $m_{\text{threshold}} = 6.5$ kHz/pulse for the remainder of this study.

IF Classification Comparison

The performance of the classification methods was compared against the original template matching approach for both minimal and dominant clutter environments. The angular variance of both situations is presented in Table 5.5 for a CPI of 8192, and the same median filter and threshold values used in the previous sections. The results show that the IF classification techniques can improve

Table 5.5: IF Classification Phase Standard Deviation (Angular) with Minimal Ground Clutter

Range (m)	Template	Variance Based	Linear Regression Based
Minimal Ground Clutter (USGS Antenna)			
91.1	2.276°	1.831°	1.662°
109.8	5.221°	5.042°	4.382°
141.3	2.980°	2.866°	3.758°
181.8	3.486°	3.504°	4.018°
199.4	4.847°	4.374°	4.092°
Dominant Ground Clutter (Furuno Antenna)			
106	1.114°	1.275°	1.317°
237	1.058°	1.549°	1.663°
365	1.953°	3.013°	3.215°
429	2.238°	3.504°	3.743°
567	3.594°	4.897°	5.126°

coherence in low-clutter scenarios, but increase the angular variance in situations with strong stationary clutter environments. This outcome is expected, as the classification process is an imperfect estimator that can misidentify pulses. This misclassification of pulses can lead to incorrect grouping and assignment of IF values, reducing coherence. When the initial IF estimates are already accurate, as they are in ground clutter dominant scenes, the additional processing can be detrimental rather than beneficial.

Chapter 6

Non-Stationary Targets and Clutter

6.1 Measurements with Platform Motion

The final objective of this project is the deployment of the pseudo-coherent radar system on a large buoy in the Pacific Ocean. Depending on environmental conditions and operational modes, such a system is expected to encounter significant levels of non-stationary clutter, so it was essential to evaluate system performance under similar conditions. Prior to conducting expensive field work in true maritime environments, preliminary testing was performed locally to assess coherence in the presence of substantial non-stationary clutter. In particular, it was necessary to demonstrate that coherence could be achieved when moving objects dominate the received radar returns, as may be the case upon deployment on a buoy.

There is an absence of naturally occurring non-stationary clutter phenomena in the vicinity of the University of Oklahoma, so non-stationary conditions were emulated by using a moving platform. The truck-mounted system setup was critical in this data collection, as the radar system was fully contained and operable from within the vehicle. For this experiment, the USGS antenna was mounted facing directly behind the truck while the vehicle was driven forward. Under these conditions, nearly all observed radar returns exhibit Doppler shifts, with the exception of transmit power leakage and near-in clutter located on the truck itself. An RDM taken from the resulting dataset is shown in Figure 6.1 for a stable IF CPI of 256 pulses. During this data collection, a large building located behind

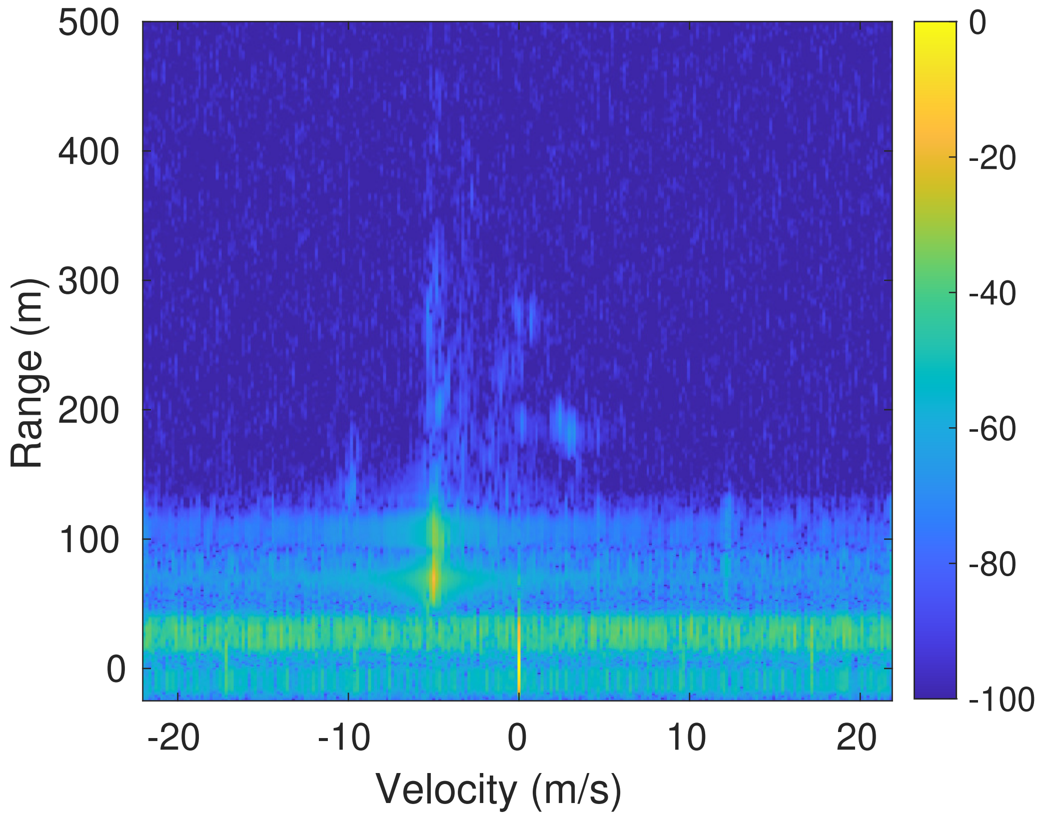

Figure 6.1: RDM of moving platform with stable IF (CPI = 256)

the radar system served as a dominant scatterer, appearing in Figure 6.1 as a strong, distributed target located at roughly 5 m/s. Additionally, several objects with positive Doppler velocities (i.e., moving towards the radar) are observed. Their polarity can be explained by the dish antenna backlobes and sidelobes that likely illuminated targets, such as trees and a chain link fence, outside the primary beam. This dataset provided an important initial validation that coherent processing is achievable in environments dominated by non-stationary clutter.

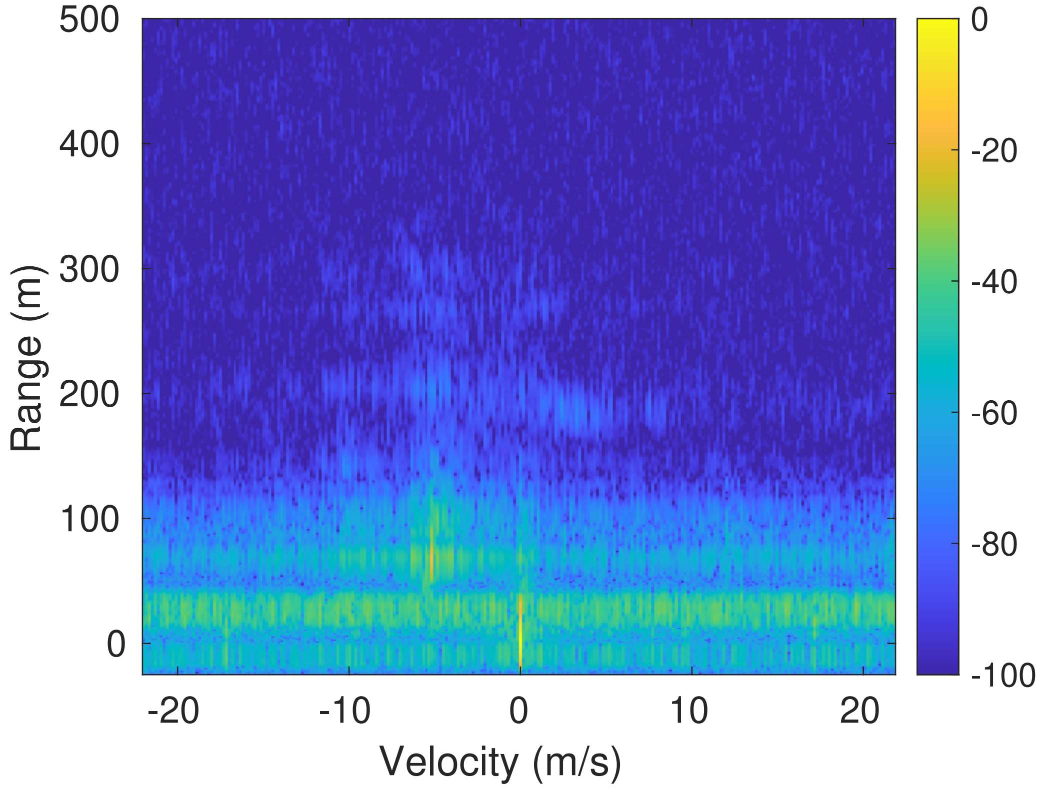

Figure 6.2: RDM of moving platform with transitory IF (CPI = 256)

However, similarly to the effects observed in minimal ground clutter scenarios, the IF estimation techniques suffered in the presence of non-stationary clutter, with an example RDM shown in Figure 6.2. This RDM was created by choosing a CPI that directly straddles two stable IF sections from the same dataset and general time period as that of Figure 6.1. It can be seen that the misestimations of IF result in Doppler smearing, reducing object SNR, and complicating what would have been detection and classification. However, despite this degradation, the data retains discernible moving target information that would be impossible in a fully non-coherent system.

6.2 Sea Clutter Performance

To further evaluate system performance, experiments were conducted in a true sea clutter environment. To address this need, the truck-mounted system was deployed in Galveston, Texas, where several data collection experiments were performed. Datasets were collected at Galveston Seawall at Diamond Beach and

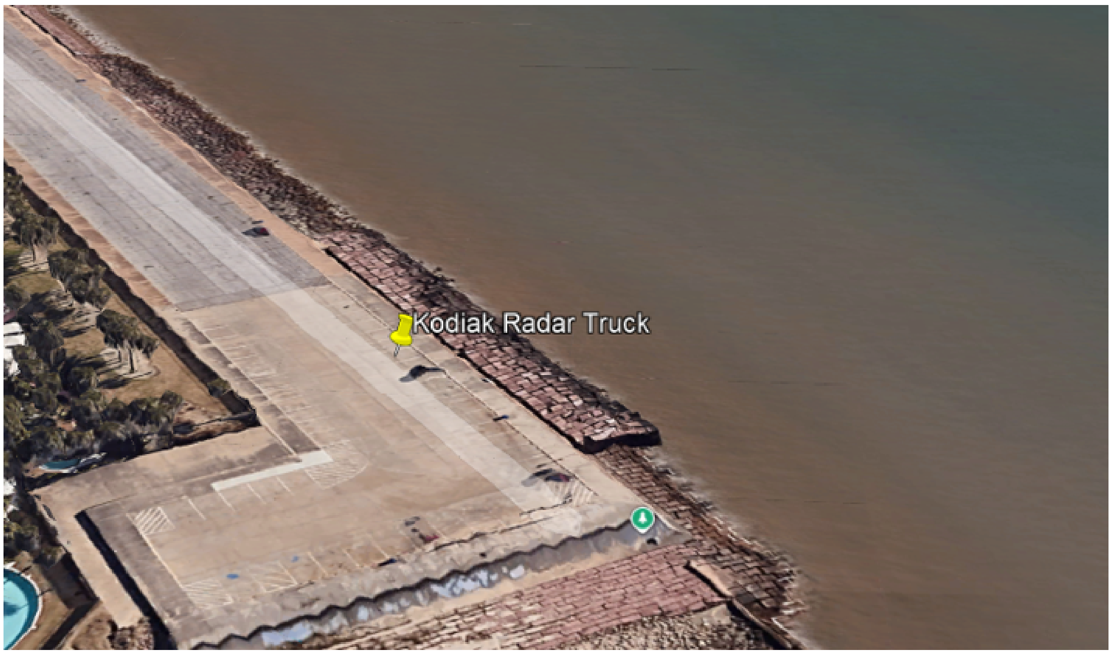

Figure 6.3: Galveston Seawall location

Galveston East Beach, pictured in Figures 6.3 and 6.4. In both cases, the radar system was positioned such that reflections from stationary objects between the radar and shoreline were within the blind range of the system. At the seawall location, ocean waves directly approached the radar with wind speeds of approximately 3-4 m/s. At the East Beach location, wave motion included a significant tangential component, resulting in Doppler shifts whose polarity depended on antenna pointing direction.

An important objective of this data collection was the characterization of sea clutter across multiple elevation angles to evaluate the degree to which it might

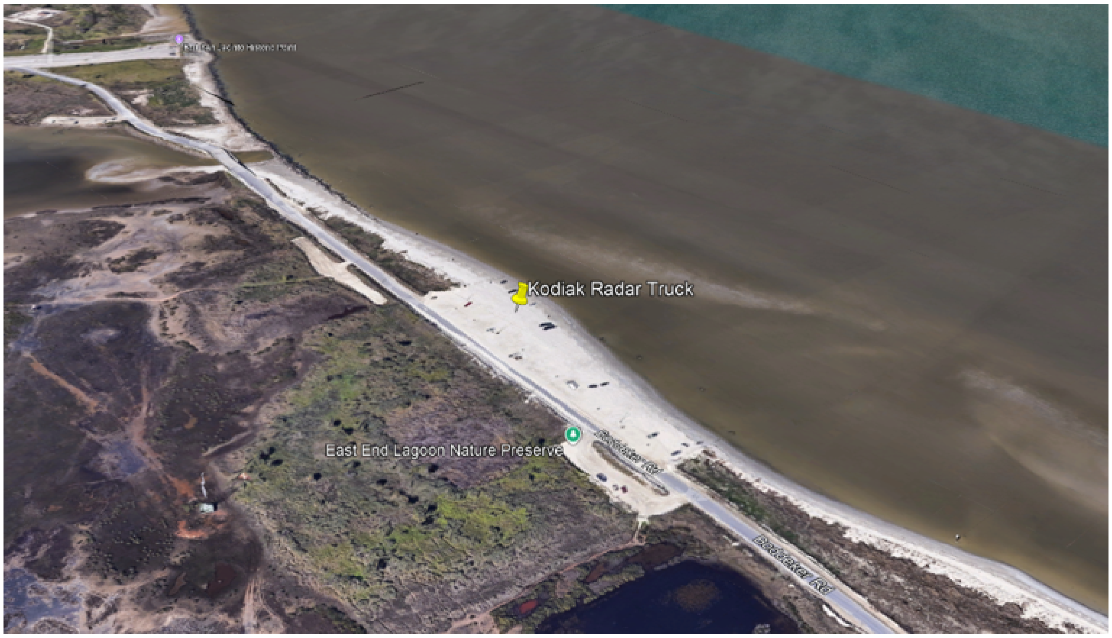

Figure 6.4: Galveston East Beach location

interfere with the operating modes of a deployed system. Radar data was acquired with the USGS dish antenna at elevation angles of 0°, 5°, 10°, 15°, 20°, 25°, 30°, 45°, 60°, 75°, and 90° for a period of several seconds. Representative RDMs for a stable IF CPI of 256 pulses at 0° and 5° elevation are shown in Figures 6.5 and 6.6, respectively. The results of Figure 6.6 demonstrate that the narrow vertical beamwidth of the USGS dish antenna effectively suppresses sea clutter, a phenomenon that was consistent for all datasets collected with elevation angles above 0°. However, at 0°, the sea clutter is sufficiently strong to obscure low-RCS targets unless they are well separated in range or Doppler.

Quantitative evaluation of coherence in the presence of real sea clutter was also performed by directing the antenna toward a stationary scatterer located beyond a region of breaking waves. Finding such a target proved challenging, and only a single suitable dataset was obtained. The selected scatterer was a permanent pier approximately 1.7 km from the radar, as shown in Figure 6.7. For a 500 pulse

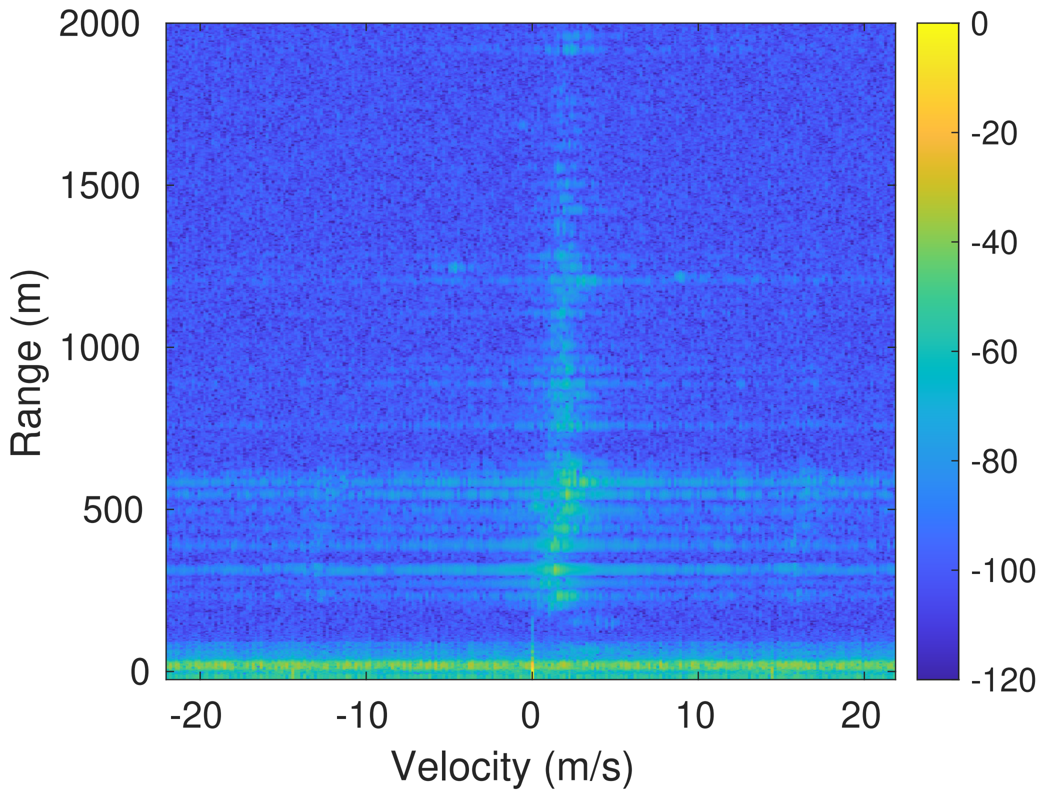

Figure 6.5: Galveston Seawall with 0° elevation angle (CPI = 256)

CPI selected from a stable IF region, the measured angular standard deviation was 6.891°. By comparison, a large building at a similar range in a ground clutter environment exhibited a standard deviation of 4.737° (Section 5.1). This increase is likely due to increased motion of the pier induced by wave action. It is a reasonable assumption that the large building referenced in Section 5.1 would be less likely to suffer the same caliber of movement as a wooden pier in the ocean. Importantly, a standard deviation of 6.891° corresponds to an approximate range displacement of only 0.30 mm, which is physically sensible in the given environment. Despite these limitations, the short CPI coherence indicates that the system maintains a high degree of coherence in sea clutter conditions.

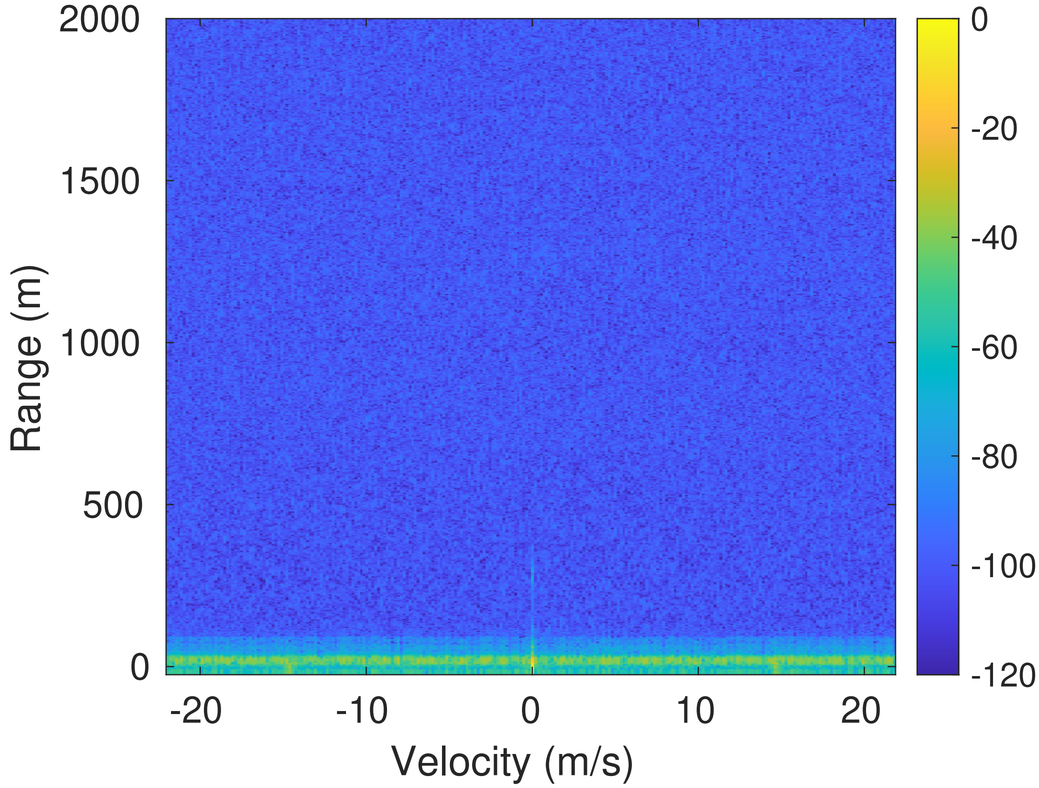

Figure 6.6: Galveston Seawall with 5° elevation angle (CPI = 256)

6.3 Moving Target Analysis

Another primary goal of this work was to develop and examine the system's ability to resolve targets in Doppler. Accordingly, experiments were conducted with a variety of moving targets under diverse environmental conditions. Ground truth velocity measurements were unavailable, so performance was assessed through analysis of target SNR and Doppler energy signature. In theory, improved coherence is expected to produce higher peak target power and reduced energy spreading in Doppler. However, it should be noted that many targets presented in this work were spinning or contained smaller moving parts that contributed their own micro-Doppler effects, which can be mistaken for Doppler spreading, regardless of coherence.

Figure 6.7: Galveston pier location relative to radar system

6.3.1 Targets in Ground Clutter

Experiments in ground clutter were conducted using the Furuno XN20AF antenna, which has a narrow horizontal beamwidth of 1.23°. It was initially difficult to catch the moving targets in the narrow beam, so objects were intentionally thrown across the azimuthal direction and were only fully illuminated for a brief period of time. As a result, a CPI of 512 pulses was selected to avoid range bin migration and ensure the target was illuminated throughout its duration. The resulting RDMs for a thrown baseball, frisbee, and football are shown in Figures 6.8–6.10. In each case, the target is clearly resolved with a concentrated main lobe, indicating successful coherent processing. Additionally, each object displays slightly altered behavior in the surrounding Doppler bins, which indicates successful measurement of the micro-Doppler signatures of each object. The normalized Doppler slices for these targets are presented in Figure 6.11, where the Doppler energy signature

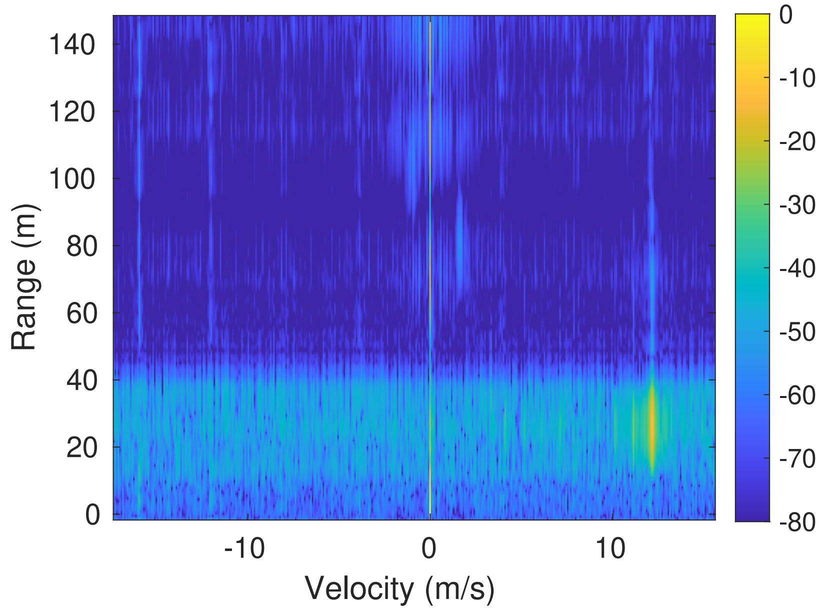

Figure 6.8: Thrown baseball (CPI = 512)

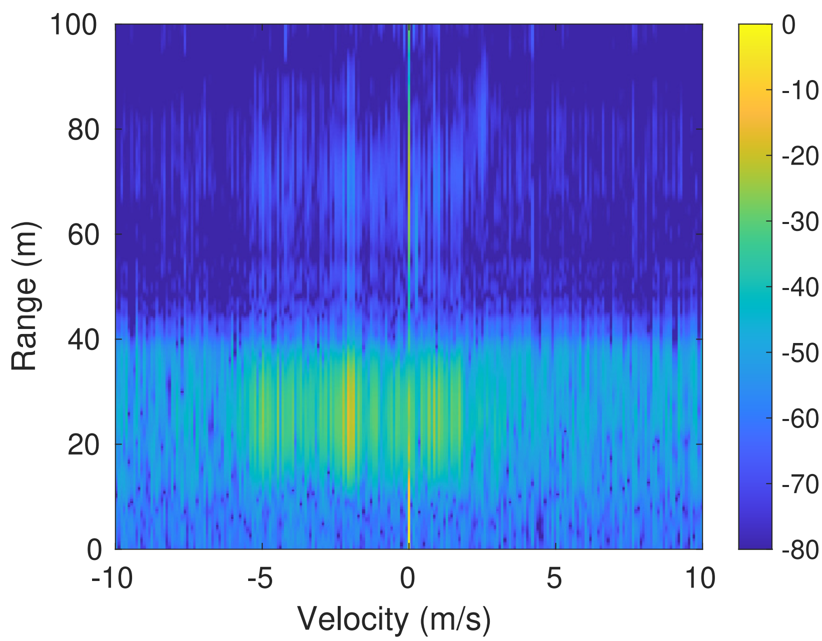

Figure 6.9: Thrown frisbee (CPI = 512)

can be more closely examined. The baseball, which most closely resembles a point scatterer, exhibits minimal Doppler structure in the surrounding bins. The

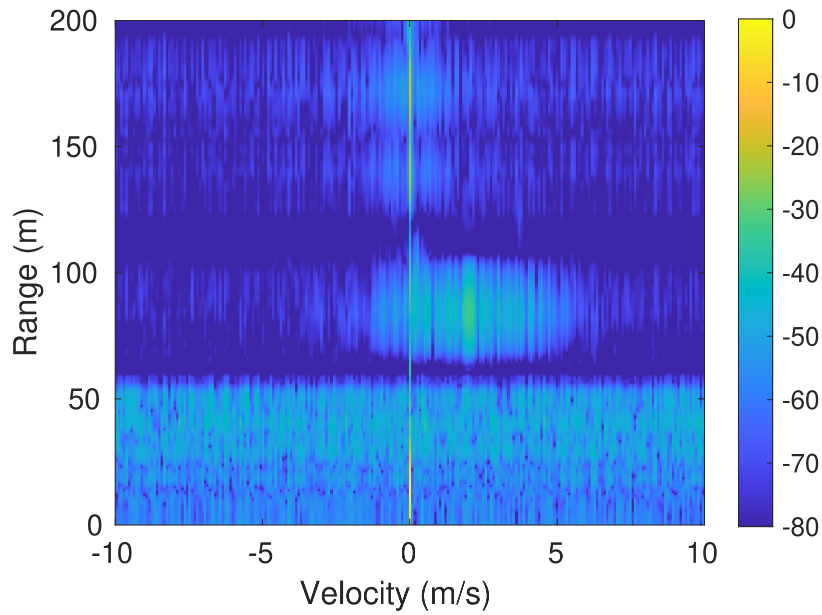

Figure 6.10: Thrown football (CPI = 512)

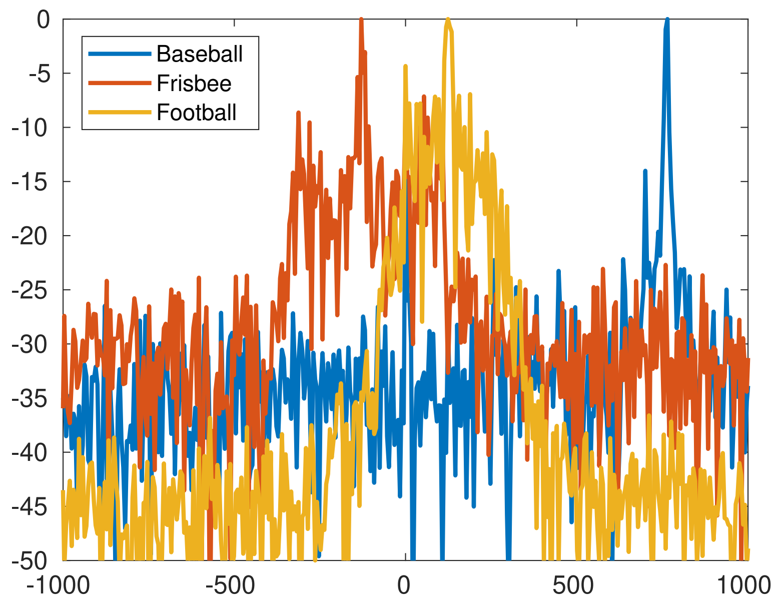

Figure 6.11: Normalized doppler energy signature of thrown objects

frisbee, which contains physical elements spinning with a wide radius, displays a clear Doppler spread related to the orientation and angular velocity of the object.

The football also exhibits a distinct Doppler signature that differs from both the baseball and frisbee, and cannot be attributed to smearing from IF misestimation. Collectively, these results confirm the effectiveness of the proposed methods and highlight the potential for future work on the use of pseudo-coherent radar systems in the measurement and classification of the micro-Doppler signatures for objects in motion.

6.3.2 Targets in Sea Clutter

For experiments involving moving targets in the presence of sea clutter, the USGS antenna was used. Each dataset consisted of several minutes of recorded radar returns, such that there was a more diverse array of moving target data. Due to the relatively narrow vertical and horizontal beamwidths of the antenna, targets with meaningful tangential velocity traversed the main beam in short periods of time. Consequently, the coherent processing interval (CPI) was again selected to be 512 pulses. The moving targets considered in this portion of the study included a small remote-controlled quadcopter, as well as birds and a search helicopter that presented themselves as targets of opportunity in the radar scene. The corresponding RDMs for each target are presented in Figures 6.12-6.14. The UAV and helicopter datasets were collected at the Galveston Seawall location, while the bird observations were made at the Galveston East Beach site.

In each case, the Doppler signatures of the targets are clearly distinguishable from the surrounding environment. The quadcopter is particularly prominent, likely due to its motion away from the radar system while the surface waves move in the opposite direction. The bird dataset contains multiple targets exhibiting both positive and negative velocities, reflecting a range of motion relative to the radar. The sea clutter is noticeably reduced in this dataset, on account of the more

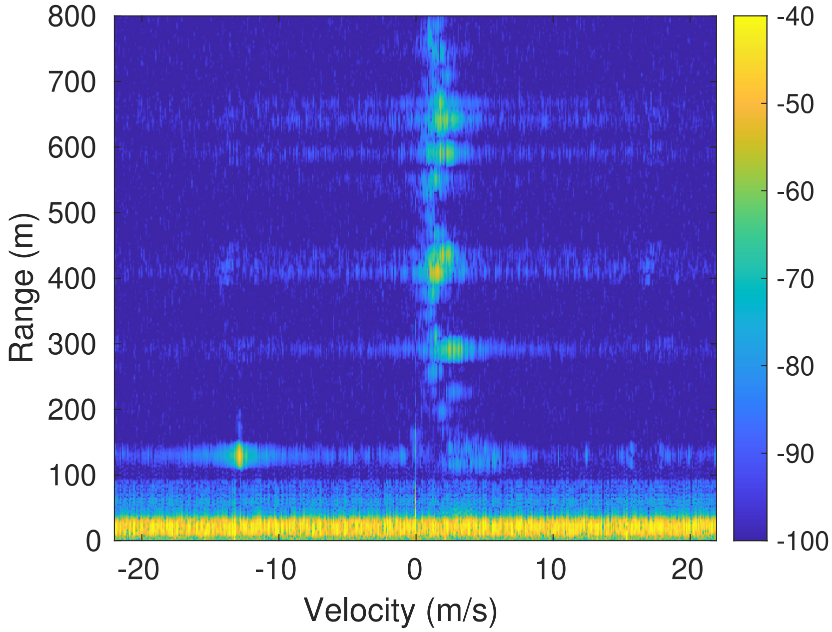

Figure 6.12: Quadcopter drone over the ocean (CPI = 512)

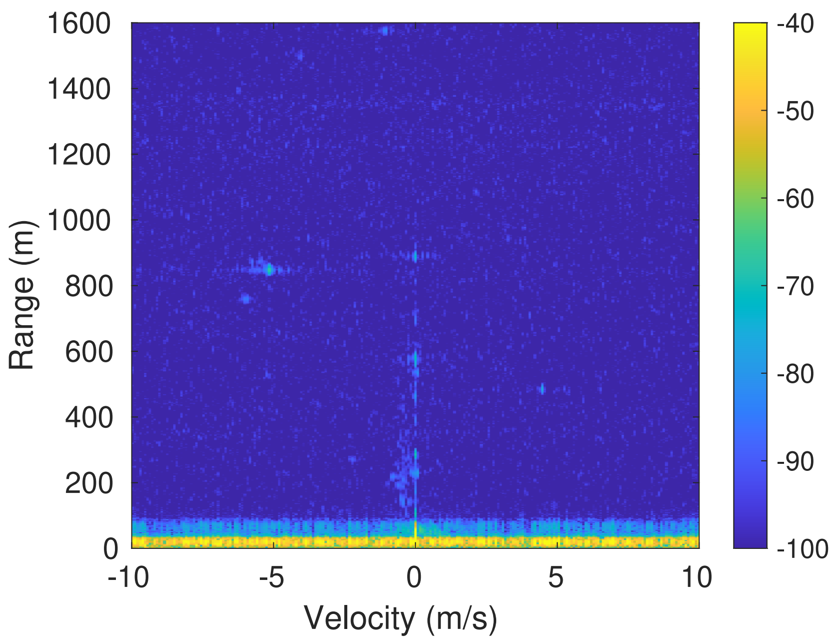

Figure 6.13: Gathering birds over the ocean (CPI = 512)

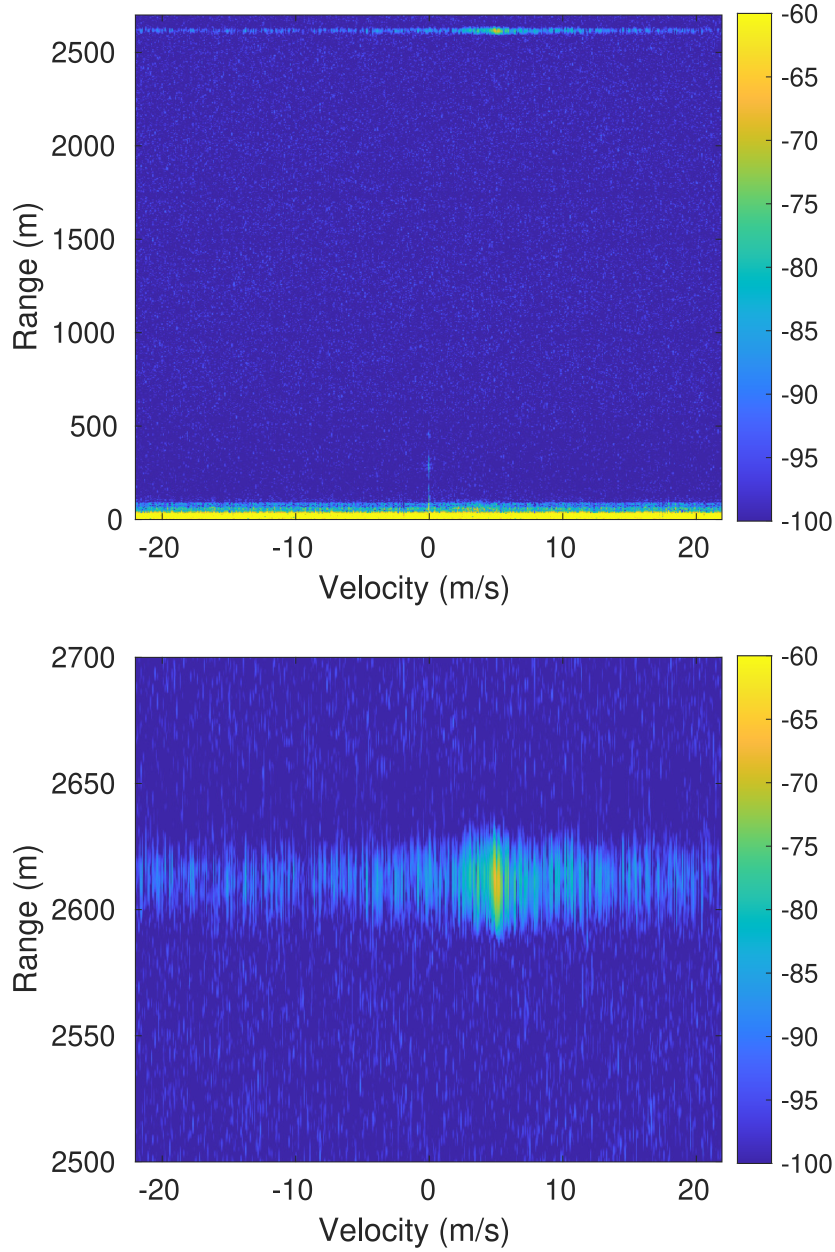

Figure 6.14: Rescue helicopter flying over the ocean (CPI = 512)

sheltered sea conditions at the East Beach location and the antenna orientation, which reduced the radial velocity component of the waves. Nevertheless, the radar system resolved many birds in the scene that are low in RCS and located at diverse ranges within the same CPI. The helicopter dataset differs slightly in that it was collected with the antenna elevated to 5°, spatially filtering nearly all sea clutter from the measured data. However, the elevated angle is an important operational mode for the system, where, despite being the most distant target considered in this study, the helicopter reflections were concentrated into a distinct peak. It is important to note that the moving targets observed in the Galveston sea clutter datasets do not directly overlap with the strongest sea clutter returns in range-Doppler space. Further investigation is therefore warranted to assess target detectability in scenarios where significant overlap occurs. Such analysis would likely require precise knowledge of both antenna beam geometry and target position.

6.3.3 SNR Evaluation and IF Misestimation Effects

The results presented in Section 6.3.2 demonstrate strong system performance under favorable conditions, but do not consider the impact of misestimation in the IF drift. Each example is selected from a stable portion of the IF drift, where the assumption of a constant IF is valid and estimation errors are minimized. This manual selection of CPI is infeasible on a deployed system, where such tuning of large quantities of data is impossible. As a result, it is necessary to evaluate system performance under general conditions, including periods of IF transition.

As discussed in Section 5.2.2, reduced stationary clutter adversely affects IF estimation accuracy, but the presence of strong moving targets introduces further challenges. It was consistently observed that strong moving scatterers altered the fast-time spectral content such that the main-lobe shifted from pulse to pulse. An example is illustrated in Figure 6.15 for the thrown baseball. This behavior violates a key assumption of the template matching approach, which relies on an essentially consistent fast-time frequency content over the CPI. Similarly, the

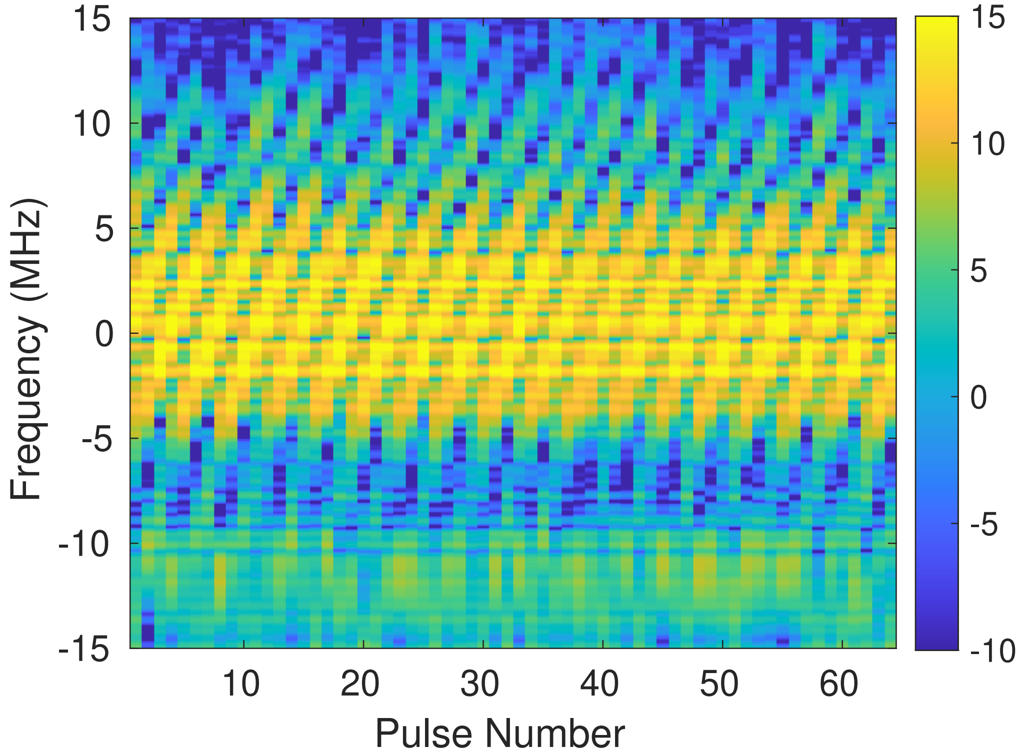

Figure 6.15: Fast-time frequency content with strong moving target

PSD estimation method correctly identifies the shifting spectral peak, resulting in inconsistent IF estimation. The degradation in IF estimation accuracy leads to Doppler spreading of target energy, similar to that in Figure 6.2, reducing the effective SNR.

To quantify this effect, the targets introduced earlier were evaluated in terms of SNR under different IF conditions and compensation strategies, using a CPI of 256 pulses. The system noise floor for this CPI length was experimentally determined to be -61.85 dB, and SNR values were found by subtracting that noise level from the peak target power. The results are summarized in Table 6.1, where each column corresponds to a different combination of IF stability and compensation strategy. They include a CPI selected from a stable IF section without adaptive digital downconversion, a stable IF section using the template matching method,

Table 6.1: SNR of Moving Targets vs. IF Estimation Technique and Drift Structure

Target	**No Adaptive, Stable**	**Template, Stable**	**No Adaptive, Transitory**	**Template, Transitory**
Baseball	91.54 dB	90.73 dB	88.93 dB	89.73 dB
Football	64.95 dB	65.28 dB	58.44 dB	61.19 dB
Frisbee	85.26 dB	83.61 dB	83.67 dB	83.92 dB
Quadcopter	57.20 dB	56.73 dB	54.84 dB	56.30 dB
Small Bird	33.78 dB	29.31 dB	24.93 dB	21.97 dB
Helicopter	43.15 dB	36.61 dB	45.98 dB	39.51 dB

a transitory IF section without adaptive digital downconversion, and a transitory IF section using the template matching method. To obtain the transitory CPI, the stable CPI was simply shifted by 256 to ensure an IF transition is included in the new interval. This selection process created some unavoidable inconsistencies where a target was more heavily in the beam for different IF drift sections, which could account for the transitory sections exhibiting higher SNR values than those of the stable portions. Regardless, it is evident that with dominant ground clutter, the template matching technique overcomes the issues presented by the shifting spectral content and outperforms non-adaptive digital downconversion in the transitory bands. It was found that non-adaptive digital downconversion is superior in the stable sections, where it can be assumed that the IF has no drift. This effect is more pronounced in sea clutter environments, where the absence of strong stationary scatterers further complicates IF estimation and reduces the effectiveness of adaptive techniques. Interestingly, in some instances, the lack of stationary clutter degrades the template matching performance enough that it is exceeded by the non-adaptive digital downconversion, even during a transitional period.

Several considerations should be noted when interpreting these results. The

SNR analysis is based on a limited dataset and may not fully capture system behavior across all operational scenarios. Furthermore, SNR is highly dependent on target orientation and trajectory relative to the antenna beam, sensitizing results to rapid positional changes. The reported SNR values are also referenced to the system noise floor and do not account for interference from clutter, which would reduce the effective signal-to-interference ratio that determines detection and tracking relationships. Despite these limitations, the results demonstrate that the proposed signal processing framework is capable of resolving moving targets across a range of conditions, including clutter environments, target dynamics, and IF stability. This performance supports the overall objective of enabling coherent processing and meaningful Doppler analysis using magnetron-based radar systems.

Chapter 7

Conclusions and Future Work

7.1 Conclusion

This thesis presented the design, implementation, and evaluation of a coherent-on-receive signal processing framework for magnetron-based marine radar systems. By addressing both random transmit phase variations and receiver-induced intermediate frequency drift, the proposed methods successfully enabled coherent processing using unmodified commercial hardware. The results demonstrate that software-based compensation techniques can effectively restore phase coherence across both short and long coherent processing intervals.

Through the development and comparison of multiple coherency algorithms, it was shown that reliable phase correction can be achieved under a range of operating conditions. Among the methods considered, the matched filtering (COR) approach consistently provided the highest degree of phase stability when an accurate transmit pulse model was available, while the gradient ascent and phase difference alignment methods offered competitive performance with varying implementation complexity. Adaptive digital downconversion was shown to be essential for maintaining coherence over long CPIs, with the template matching approach demonstrating superior robustness compared to power spectral estimation techniques. Although performance degradation was observed during periods of rapid IF variation and in environments dominated by non-stationary clutter, the overall framework improved coherence and enabled meaningful Doppler analysis.

Experimental validation confirmed that the proposed system is capable of resolving a wide range of targets, including low radar cross-section objects, in both ground-based and maritime environments. The ability to generate well-defined range-Doppler signatures for moving targets, including drones, helicopters, projectiles, and biological scatterers, demonstrates the effectiveness of the approach in achieving the primary objectives of this work.

Despite these successes, several limitations were identified. These include sensitivity to IF estimation errors and reduced performance in scenarios with strong non-stationary clutter or limited stationary reference returns. These observations highlight important areas for continued investigation, including the development of more robust IF estimation strategies. Nevertheless, this work demonstrates the feasibility of transforming low-cost, widely available marine radar systems into capable pseudo-coherent sensing platforms through advanced signal processing alone. This capability has significant implications for expanding the use of marine radar in scientific and operational applications, particularly in aeroecology, where financially accessible and high-performing sensing solutions are increasingly valuable.

7.2 Scientific Impact

This thesis addresses an underexplored problem space and introduces several contributions to the existing body of work on coherent processing with commercial marine radar systems. Prior efforts to achieve coherence with magnetron-based systems have largely relied on invasive hardware modifications to the subsystems within. In contrast, this work demonstrates that high levels of coherence can be achieved through software-based approaches.

Specifically, the proposed power spectral density and template matching IF estimation methods provide practical solutions to compensating receiver-induced frequency drift without requiring hardware changes. In addition, this work presents a detailed characterization of algorithm performances across a variety of radar environments, including scenarios with dominant stationary clutter, minimal clutter, and highly dynamic conditions. The comparison of established and newly proposed coherency algorithms further contributes to the understanding of phase correction techniques, including the identification of approaches that achieve near-equivalent performance with reduced computational complexity.

Another key contribution is the quantitative evaluation of coherence using multiple complementary metrics, addressing a gap in prior literature where qualitative assessments were more common. Collectively, these contributions culminate in the development of a low-cost pseudo-coherent radar system suitable for future scientific investigations, particularly in the study of aeroecological phenomena.

7.3 Future Work

Several avenues for future research emerge from the findings of this thesis. Although the proposed signal processing framework demonstrates strong performance, further improvements in IF estimation remain an important objective, particularly in environments characterized by minimal stationary clutter or highly non-stationary returns. Additional investigation into the effects of residual phase and frequency misalignment on downstream processing tasks, such as target detection and tracking, would further enhance system understanding. Moreover, the analysis and classification of micro-Doppler signatures of moving targets in pseudo-coherent radar data represents an unexplored space with promising potential for future research branching directly from the work done in this thesis.

References

[1] D. S. Jensen, A. Ahmed, and J. G. Metcalf, "End-to-end signal processing pipeline for coherent-on-receive marine radar," in *2026 IEEE Radar Conference (RadarConf'26)*, Phoenix, AZ, USA: IEEE, 2026, pp. 1–6.

[2] M. I. Skolnik, *Radar handbook*. New York, NY: McGraw-Hill, 1970. [Online]. Available: `https://cds.cern.ch/record/98948`.

[3] M. A. Richards, *Fundamentals of Radar Signal Processing*, 2nd. McGraw-Hill Education, 2014.

[4] R. E. Nyswander, *Method and apparatus for magnetron coherence*, Google Patents, U.S. Patent 6,914,949, 2005.

[5] P. A. Devito, "Some properties of an injection-locked pulsed magnetron," Defense Technical Information Center, Tech. Rep., 1973.

[6] S. Kristoffersen, K. V. Hoel, O. Thingsrud, and E. B. Kalveland, "Digital coherent processing to enhance moving targets detection in a navigation radar," en, in *2014 International Radar Conference*, Lille, France: IEEE, Oct. 2014, 1–6, ISBN: 978-1-4799-4195-7. DOI: `10.1109/RADAR.2014.7060295`. [Online]. Available: `http://ieeexplore.ieee.org/document/7060295/`.

[7] G. E. Smith et al., "High power coherent-on-receive radar for marine surveillance," in *2013 International Conference on Radar*, 2013, pp. 434–439. DOI: `10.1109/RADAR.2013.6652028`.

[8] K. Hessner and J. L. Hanson, "Extraction of coastal wavefield properties from x-band radar," in *2010 IEEE International Geoscience and Remote Sensing Symposium*, 2010, pp. 4326–4329. DOI: `10.1109/IGARSS.2010.5650134`.

[9] H. Chien, H.-Y. Cheng, and Y. Zhong, “Validating of coherent-on-receive marine radar in nearshore using drifter cluster,” in *OCEANS 2017 - Aberdeen*, 2017, pp. 1–4. DOI: 10.1109/OCEANSE.2017.8084883.

[10] M. Aljohani, A. Mrebit, L. Lo Monte, and M. C. Wicks, “Radar imaging using pseudo-coherent marine radar technology,” *IET Radar, Sonar & Navigation*, vol. 14, no. 6, pp. 905–916, DOI: https://doi.org/10.1049/iet-rsn.2019.0154.

[11] A. W. Doerry, “Radar receiver oscillator phase noise,” Sandia National Laboratories (SNL-NM), Albuquerque, NM (United States), Tech. Rep., Apr. 2018. DOI: 10.2172/1528837. [Online]. Available: https://www.osti.gov/biblio/1528837.

[12] T. Parker, “Characteristics and sources of phase noise in stable oscillators,” in *41st Annual Symposium on Frequency Control*, 1987, pp. 99–110. DOI: 10.1109/FREQ.1987.201009.

[13] H.-f. Zhang et al., “High-speed arbitrary waveform generator based on fpga,” in *2013 IEEE Nuclear Science Symposium and Medical Imaging Conference (2013 NSS/MIC)*, 2013, pp. 1–5. DOI: 10.1109/NSSMIC.2013.6829760.

[14] .

[15] M. Shao, P. Zhang, and Z. Zhang, “Realization of digital down conversion in pulse radar receiver,” in *2016 IEEE MTT-S International Microwave Workshop Series on Advanced Materials and Processes for RF and THz Applications (IMWS-AMP)*, 2016, pp. 1–4. DOI: 10.1109/IMWS-AMP.2016.7588435.

[16] J. C. Curlander and R. N. McDonough, “Synthetic aperture radar: Systems and signal processing,” 1991. [Online]. Available: https://api.semanticscholar.org/CorpusID:264163199.

[17] A. V. Oppenheim and R. W. Schafer, *Discrete-Time Signal Processing*, 3rd. USA: Prentice Hall Press, 2009, ISBN: 0131988425.

[18] W. R. Bennett, “Spectra of quantized signals,” *The Bell System Technical Journal*, vol. 27, no. 3, pp. 446–472, 1948.

[19] B. Widrow, "A study of rough amplitude quantization by means of nyquist sampling theory," *IRE Transactions on Circuit Theory*, vol. 3, no. 4, pp. 266–276, 1956. DOI: 10.1109/TCT.1956.1086334.

[20] B. Widrow and I. Kollár, *Quantization Noise: Roundoff Error in Digital Computation, Signal Processing, Control, and Communications*. Cambridge University Press, 2008.

[21] R. Gray, "Quantization noise spectra," *IEEE Transactions on Information Theory*, vol. 36, no. 6, pp. 1220–1244, 1990. DOI: 10.1109/18.59924.

[22] G. B. Jordan, "Comparison of two major classes of coherent pulsed radar systems," *IEEE Transactions on Aerospace and Electronic Systems*, vol. AES-11, no. 3, pp. 363–371, 1975. DOI: 10.1109/TAES.1975.308087.

[23] M. A. Richards, "Noncoherent integration gain, and its approximation," *Georgia Institute of Technology, Technical Memo*, 2010.

[24] D. Koks, "How to create and manipulate radar range–doppler plots," Cyber Electronic Warfare Division, Defense Science and Technology Organisation, Edinburgh, SA 5111, Australia, Tech. Rep. DSTO–TN–1386, 2014. [Online]. Available: https://apps.dtic.mil/sti/tr/pdf/ADA615308.pdf.

[25] B. Stade, "Sea clutter modeling for the evaluation of maritime radar detection capabilities," Master's thesis, University of Oklahoma, Norman, OK, 2025.

[26] *Specifications of Marine Radar FAR-1528*, Furuno, Camas, Washington, 2026. Accessed: Apr. 2, 2026. [Online]. Available: https://www.furuno.it/docs/SPECIFICATIONSSP_FAR1528_MARINE_RADAR.pdf.

[27] *M4i-4481x8 datasheet*, Spectrum Instruments, Grosshansdorf, Germany, 2026. Accessed: Apr. 2, 2026.

[28] M. Streßer et al., "On the interpretation of coherent marine radar backscatter from surf zone waves," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 60, pp. 1–14, 2022. DOI: 10.1109/TGRS.2021.3103417.

[29] J. Horstmann, J. Bödewadt, R. Carrasco, M. Cysewski, J. Seemann, and M. Streer, "A coherent on receive x-band marine radar for ocean observations," *Sensors*, vol. 21, no. 23, 2021, ISSN: 1424-8220. DOI: 10.3390/s21237828. [Online]. Available: https://www.mdpi.com/1424-8220/21/23/7828.

[30] M. N. Anwar Tarek, M. Roman, and E. A. Alwan, "Power efficient rf self-interference cancellation system for simultaneous transmit and receive," in *2021 IEEE International Symposium on Antennas and Propagation and USNC-URSI Radio Science Meeting (APS/URSI)*, 2021, pp. 113–114. DOI: 10.1109/APS/URSI47566.2021.9704722.

[31] S.-T. Han, D. Kim, J. Kim, and J.-R. Yang, "Noise suppression and precise phase control of a commercial s-band magnetron," *IEEE Access*, vol. 8, pp. 145 881–145 886, 2020. DOI: 10.1109/ACCESS.2020.3013651.

[32] I.-K. Baek et al., "Origin of sideband and spurious noises in microwave oven magnetron," *IEEE Transactions on Electron Devices*, vol. 64, no. 8, pp. 3413–3420, 2017. DOI: 10.1109/TED.2017.2714339.

[33] R. G. Mattingly, A. F. Martone, and J. G. Metcalf, "Techniques for mitigating the impact of intra-cpi waveform agility," *IEEE Transactions on Radar Systems*, vol. 2, pp. 24–40, 2024. DOI: 10.1109/TRS.2023.3343192.

[34] N. I. Fisher, "Descriptive methods," in *Statistical Analysis of Circular Data*. Cambridge University Press, 1993, 15–38.

[35] G. R. Ackermann, "Means and standard deviations of horizontal wind components," *Journal of Climate and Applied Meteorology*, vol. 22, no. 5, pp. 959–961, 1983.

[36] R. J. Yamartino, "A Comparison of Several 'Single-Pass' Estimators of the Standard Deviation of Wind Direction.," *Journal of Applied Meteorology*, vol. 23, no. 9, pp. 1362–1366, Sep. 1984.

Appendix A

List of Acronyms and Abbreviations

ADC	Analog-to-digital converter
$CFAR$	Continuous False Alarm Rate
COR	Coherent-On-Receive
$COTS$	Commercial Off the Shelf
CPI	Coherent Processing Interval
DFT	Discrete Fourier Transform
EM	Electromagnetic
$GMTI$	Ground-Moving Target Indicator
IF	Intermediate Frequency
IQ	In-Phase and Quadrature
LNA	Low Noise Amplifier
PCB	Printed Circuit Board
PDA	Phase Difference Averaging
PRF	Pulse Repetition Frequency
PRI	Pulse Repetition Interval
$RADAR$	Radio Detection and Ranging
RCS	Radar Cross Section
RDM	Range-Doppler Map
RF	Radio Frequency

SAR	Synthetic Aperture Radar
SNR	Signal-to-Noise Ratio
TWT	Traveling-Wave-Tube
$*$	Linear Convolution