CHALMERS
UNIVERSITY OF TECHNOLOGY

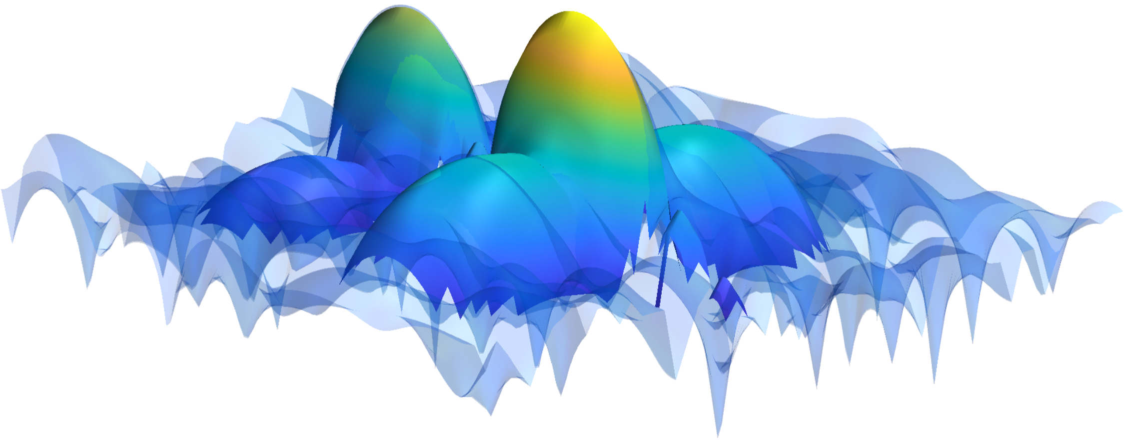

OMP-Based Target Detection for Pulse-Doppler Radar

Improved Detection of Closely Spaced Targets and Targets in Presence of Strong Clutter

Master's thesis in Electrical Engineering

TESS RIMARCUK, ADAM UDÉN

DEPARTMENT OF ELECTRICAL ENGINEERING
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2026
www.chalmers.se

Master's thesis 2026

OMP-Based Target Detection for Pulse-Doppler Radar

Improved Detection of Closely Spaced Targets and Targets in Presence of Strong Clutter

TESS RIMARCUK, ADAM UDÉN

Department of Electrical Engineering
Division of Signal Processing and Biomedical Engineering
Chalmers University of Technology
Gothenburg, Sweden 2026

OMP-Based Target Detection for Pulse-Doppler Radar
Improved Detection of Closely Spaced Targets and Targets in Presence of Strong Clutter
TESS RIMARCUK, ADAM UDÉN

Supervisor: Björn Hallberg, Saab AB
Examiner: Tomas McKelvey, Department of Electrical Engineering

Master's Thesis 2026
Department of Electrical Engineering
Division of Signal Processing and Biomedical Engineering
Chalmers University of Technology
SE-412 96 Gothenburg
Telephone +46 31 772 1000

Cover: Sparse reconstruction of noisy pulse-Doppler radar data

Typeset in LaTeX
Printed by Chalmers Reproservice
Gothenburg, Sweden 2026

OMP-Based Target Detection for Pulse-Doppler Radar
Improved Detection of Closely Spaced Targets and Targets in Presence of Strong Clutter
TESS RIMARCUK, ADAM UDÉN
Department of Electrical Engineering
Chalmers University of Technology

Abstract

Radar systems transmit radio-frequency electromagnetic waves and receive the reflected signals from objects within their field of view. A surveillance radar aims to detect and estimate the state of targets in the surveillance volume. Traditional detection algorithms struggle in environments containing closely spaced objects or heavy clutter, resulting in false alarms or missed detections. Greedy compressed sensing algorithms, such as Orthogonal Matching Pursuit (OMP), address this through an iterative process that identifies and subtracts the strongest component in the residual signal, enabling detection of weaker targets masked by stronger ones. In this Master's thesis, OMP was implemented and evaluated on both real and simulated pulse-Doppler radar data, focusing on scenarios involving closely spaced targets and clutter-rich environments. The performance of the OMP detector was compared to a conventional Cell-Averaging Constant False Alarm Rate (CA-CFAR) detector. A proposed hybrid of the two detectors, denoted CA-OMP, was also evaluated. The results indicate that the OMP detector can resolve closely spaced targets and targets in the presence of strong clutter more effectively than CA-CFAR, particularly in simulated data. However, on real data it is unable to fully suppress detected signal components, resulting in false alarms and degraded detection performance across the surveillance space. The CA-OMP detector exhibited improved detection performance on real data compared to the CA-CFAR detector, while mitigating most of the incorrect detections generated by the OMP algorithm.

Keywords: Pulsed- Doppler radar, Target detection, Clutter handling, Orthogonal Matching Pursuit, Greedy algorithms

Acknowledgements

We would like to express our sincere gratitude to our supervisor, Björn Hallberg, for their guidance, valuable feedback, and continuous support throughout this thesis work. We would also like to thank our examiner and co-supervisor Tomas McKelvey for insightful discussions and assistance during the project.

Furthermore, we are grateful to Saab Surveillance for providing the opportunity and resources necessary to carry out this research.

Tess Rimarcuk and Adam Udén, Gothenburg, May 2026

List of Acronyms

Below is the list of acronyms that have been used throughout this thesis listed in alphabetical order:

AWGN	Addative White Gaussian Noise
CA-CFAR	Cell Averaging - Constant False Alarm Rate
CFAR	Constant False Alarm Rate
CPI	Coherent Processing Interval
CRLB	Cramér-Rao Lower Bound
CUT	Cell Under Test
CW	Continuous Wave
DFT	Discrete Fourier Transform
FoM	Figure of Merit
OMP	Orthogonal Matching Pursuit
PRF	Pulse Repetition Frequency
PRI	Pulse Repetition Interval
RCS	Radar Cross Section
RMSD	Root Mean Square Deviation
SIR	Signal to Interference Ratio
SNR	Signal to Noise Ratio

Contents

List of Figures

List of Tables

1

Introduction

Radar, originally an acronym for RAdio Detection And Ranging, is a common technology used in applications as varied as traffic speed measurements, weather monitoring and prediction, and autonomous driving systems [1, Ch. 1.1]. Radar technology also sees widespread use in military technology, where it is utilized for surveillance, navigation, and weapon guidance, among others. A key part of radar technology is the signal processing. The primary uses of a radar signal processor are detection, tracking, and imaging [1, Ch. 1.2]. Detection and tracking are the operations most commonly associated with radar and generate target indications on a radar screen that can be followed. Imaging is a somewhat different operation used for generating images of an area. This paper will focus on the detection part of the signal processor.

When the radar is used for detection and tracking, detection is the initial step in which the algorithm determines whether the received data contains signals possibly corresponding to targets. In the following tracking stage, the motion of the detected targets is estimated over time, providing target trajectories and filtering out irrelevant detections. Finally, the output of the tracking step is presented to the radar operator. As detection is the first stage in the chain, it provides the foundation for the tracking stage, making its performance critical to the effectiveness of the overall system.

Detecting isolated targets in low-noise conditions is a relatively straightforward task for most radar detection algorithms. However, such conditions rarely reflect real-world scenarios. In real radar data challenges arise due to the presence of multiple reflective objects within the field of view of the radar. Target responses from closely spaced objects might overlap, making them appear as a single target instead of two. Similarly, if the data contains targets with varying signal strengths, weaker targets are likely to be masked by a stronger target's main lobe or sidelobes, making reliable detection difficult. Furthermore, reflections from the ground and surrounding terrain, known as clutter, introduce further complexity. These reflections often look like many closely spaced target responses with potentially high amplitude. The sidelobes from such clutter signals can obstruct and hide the responses from targets of interest, thereby reducing detection performance. A robust radar detector must be able to distinguish between closely spaced targets, detect targets of widely varying signal strengths, and targets in the presence of strong clutter.

1.1 Purpose

This thesis aims to implement and evaluate an Orthogonal Matching Pursuit (OMP) algorithm for radar target detection in pulse-Doppler radar systems. The implementation will be evaluated using both MATLAB-simulated radar data and real-world data. Based on the results, the suitability of the detector for practical applications will be discussed, and its performance will be compared with a standard Cell Averaging-Constant False Alarm Rate (CA-CFAR) detector.

1.2 Limitations

To limit the scope of the project and to reduce the computational burden, the algorithm will be developed for use on two dimensional data containing only range and Doppler information. The direction-of-arrival dimension can be considered a possible extension for future research.

Furthermore, both simulated and real-world radar data will be obtained from a stationary ground-based radar with slowly moving targets such as birds and cars.

1.3 Related Research

Sparse, greedy reconstruction algorithms such as OMP are extensively studied in both signal and image processing, with numerous extensions and variants proposed [2], [3], [4], [5], [6], [7].

However, to the authors' knowledge, existing studies either rely completely on simulated radar data or use experimental data obtained in highly controlled experimental settings. As a result, the detection performance of the OMP algorithm on real radar data obtained from a complex environment is not yet well established.

In addition to target detection, OMP-like methods have been used for clutter suppression in passive and noise radar applications [8], [9].

2

Theory

A radar is a device that emits radio-frequency electromagnetic waves and receives the returns, also called echoes, from targets [1, Ch. 1.1]. A target in this setting is defined as any object of interest to the radar operator. The returning echoes are attenuated, time delayed, and phase-shifted versions of the output signal. By analyzing the modifications of the returning signal, the radar system is able to determine the distance between the radar and the target, referred to as the range, and in most cases also the direction and velocity of the reflecting object. The range is calculated by the time it takes for the signal to travel to the target and back. The direction, typically described by azimuth and elevation angles, is estimated by antenna characteristics or array processing techniques. Lastly, velocity is determined from the frequency shift of the signal caused by the Doppler effect. This thesis considers only two dimensions, range and velocity, and consequently estimation of direction will not be considered.

Radar systems are developed with specific performance requirements in mind [1, Ch. 1.3]. For instance they can be designed to achieve a specific range resolution, detection sensitivity, or velocity accuracy. These parameters are strongly affected by the type of signal modulation that is applied to the transmitted signal, called the radar waveform. There are two general classes of waveforms: the continuous wave (CW) and the pulsed wave. CW radars transmit and receive simultaneously, whereas pulsed radars transmit short pulses and only receive between pulse transmissions. In addition to selecting between a CW or a pulsed waveform, it is common to apply phase or frequency modulation to the signal to further increase performance and enable velocity estimation. This thesis will utilize a pulse-Doppler type radar.

2.1 Pulse-Doppler Radar

In a pulse-Doppler radar, the received signal from multiple pulses is processed to form a two dimensional map called the range-Doppler map. One of the dimensions corresponds to Doppler frequency while the other corresponds to range. In this map, reflections of objects will appear as peaks over the noise floor, making it a convenient representation for detection. Consequently, before detection one of the steps is to convert the data into the range-Doppler domain.

2.1.1 Fast-time and Slow-time

A pulse-Doppler radar operates by transmitting a fixed number of pulses, referred to as a pulse train, at the Pulse Repetition Frequency (PRF). When not transmitting, the antenna receives. The data received from one pulse train is called a Coherent Processing Interval (CPI) and is processed by reshaping the time series over the entire CPI. The time series is cut into sections corresponding to the Pulse Repetition Interval (PRI), which is the inverse of the PRF. Each such section is then stacked on top of one another, creating a complex-valued data matrix with the dimensions fast-time and slow-time, as shown in Figure 2.1. The fast-time dimension corresponds to each cut-out section and is sampled at the system's sampling rate. The slow-time dimension can also be understood as the pulse index in the pulse train and is sampled at the PRF.

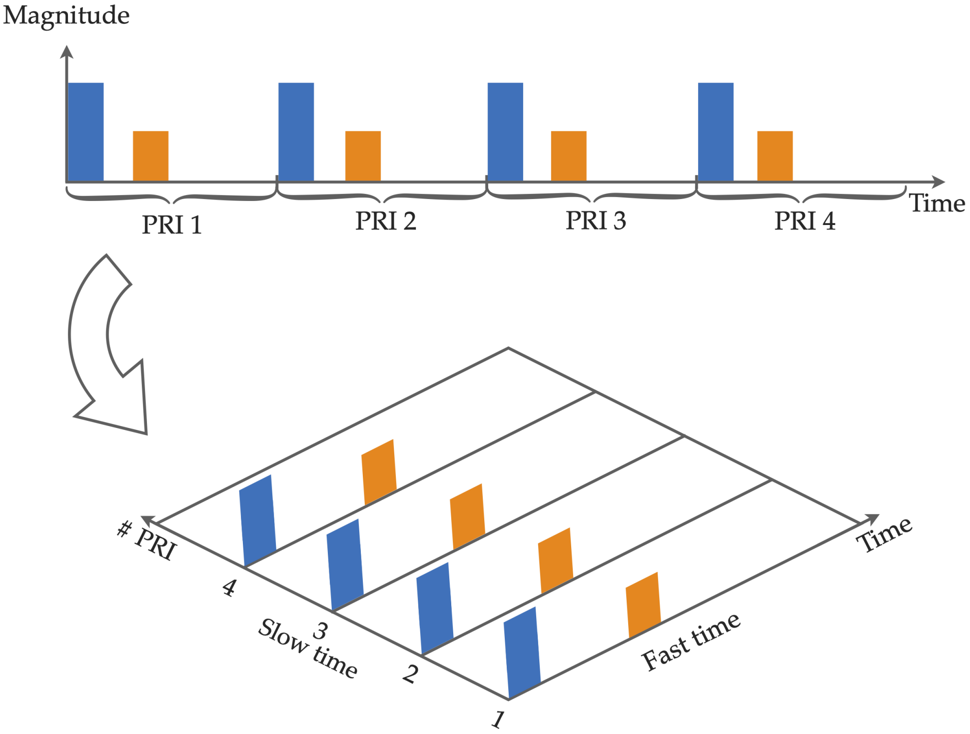

Figure 2.1: The received signal from two targets are shown in the upper figure. The response from the blue and orange target are repeated every PRI. The depiction below it shows how the received signal is cut up into slow- and fast-time domains.

2.1.2 Range and Doppler Equations

In the fast-time dimension, the distance between the receiver and the target can be determined. This is achieved by measuring the time τ between the transmitted pulse and the received echo. This time corresponds to the pulse traveling twice the distance between the receiver and the target R at the speed of light, c, resulting in

$$R = \frac{c\tau}{2}. \tag{2.1}$$

In the slow-time dimension, the radial velocity of the target can be determined. Assuming that the target moves slowly enough such that its echo remains at the same delay position in the fast-time dimension, the samples in this dimension can be understood as sampling one specific point on the echo waveform. If the target is stationary, this will result in a constant amplitude and phase, but if the target is moving the phase will increase or decrease linearly in time due to the Doppler shift. The Doppler frequency shift f_D can thereby be determined by Fourier transforming in the slow-time dimension. The Doppler frequency can be translated to the target's radial velocity v_D using

$$v_D = \frac{f_D\lambda}{2} \tag{2.2}$$

where λ is the wavelength of the transmitted wave.

2.1.3 Pulse Compression and Fourier Transformation

The pulses used in a pulse-Doppler system are in general sinusoidal signals oscillating at a carrier frequency in the radio frequency range with some frequency modulation. A longer pulse duration means the pulse contains more energy as well as longer integration time for the receiver. The more energy a pulse contains, the further it is able to travel and still be detected, resulting in a longer detection range. However, since a pulse of duration T traveling at the speed of light will be Tc long, the echoes from targets separated by a distance smaller than the range resolution $\Delta R = Tc/2$ will interfere with each other, making detection more difficult [1, Ch. 2.1]. The range resolution can be improved by utilizing a shorter pulse.

To maximize both range resolution and detection range, an approach called pulse compression is used. Pulse compression is usually performed using a matched filter and allows for a long pulse with high energy to be transmitted while at the same time improving the range resolution. This is achieved through the transmission of a modulated pulse. A common choice of modulation is the linear frequency modulation, or chirp, where the phase is modulated according to

$$f(t) = e^{j2\pi(-\frac{Bt}{2}+\frac{Bt^2}{2T})} \tag{2.3}$$

where B is the bandwidth, t is the time, and T is the pulse duration. The received signal is then correlated with a reference of the transmitted pulse. The more the received signal matches the reference in phase, the higher the amplitude of the resulting output will be [10, Ch. 18]. Figure 2.2 shows the received pulses of two targets before and after matched filtering. In Figure 2.2a the pulses are completely separated, while the pulses interfere in Figure 2.2b. After pulse compression, it is possible to distinguish the pulses from each other in both cases. When using pulse compression the range resolution is instead determined by the bandwidth through $\Delta R = c/2B$ [1, Ch. 4].

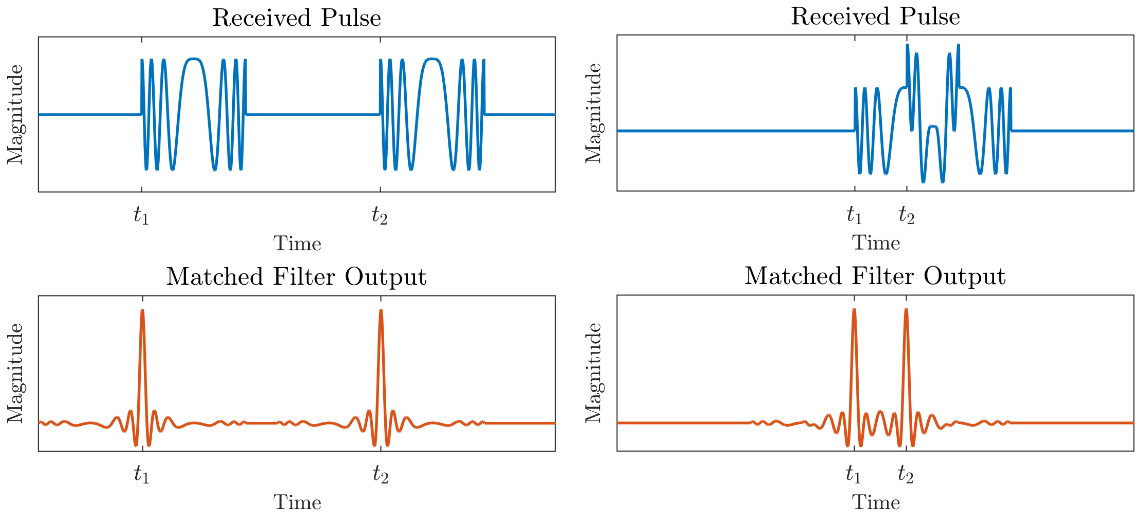

(a) The pulse returns from two well separated targets.

(b) The pulse returns from two targets that are interfering.

Figure 2.2: The received pulse before and after applying a matched filter. Further, t_1 and t_2 are the times the pulses return.

The output after matched filtering will not be an ideal impulse. Instead, it will consist of a main lobe flanked by several smaller peaks called sidelobes. For a linear frequency modulated chirp that follows Equation 2.3, the output after pulse compression will resemble Figure 2.3a. The main lobe width determines the range resolution, and the sidelobes will extend to neighboring fast-time samples. In many cases sidelobes affect the detection of targets located nearby in range. Sidelobes from strong targets may have larger amplitudes than main lobes from weaker targets, resulting in the weaker targets disappearing in the sidelobe structure. Furthermore, the sidelobes may also be stronger than the detection threshold, leading to false detections.

In the slow-time dimension, a discrete Fourier transform (DFT) is applied to extract the velocity information. The Fourier transform of an infinite duration complex exponential is the Dirac delta function. In practice, the signal will not be infinite but observed over a finite interval. Mathematically, this corresponds to multiplying the exponential by a rectangular window, which consists of ones within the finite interval and zeros outside. This results in a sinc-shaped output after Fourier transforming, seen in Figure 2.3b. Hence, sidelobes can impact the detection performance negatively as well.

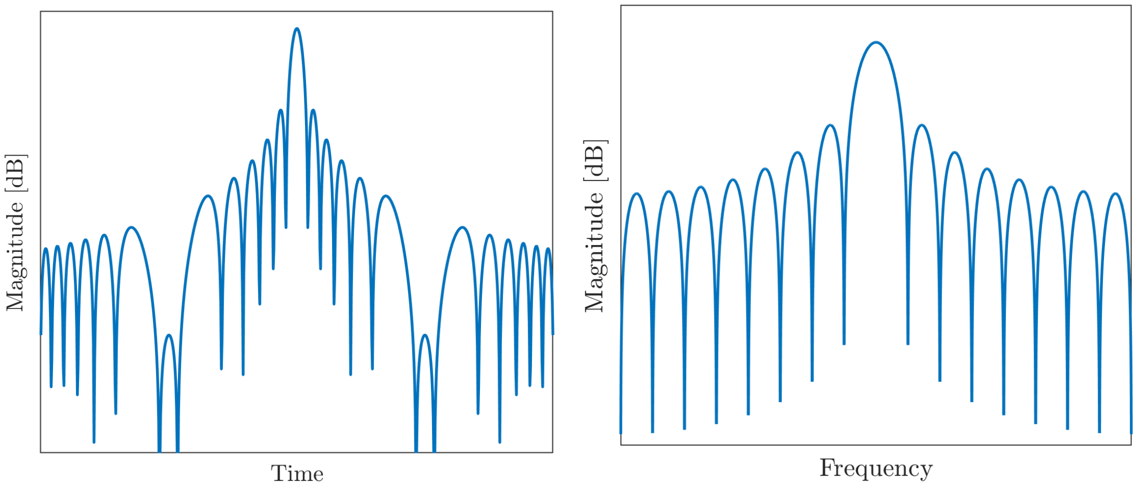

(a) Output after pulse compression of the chirp.

(b) Output after Fourier transforming a complex exponential.

Figure 2.3: There are visible sidelobes in both the range and Doppler directions after pulse compression and Fourier transforming.

A common solution to suppress sidelobes is to replace the rectangular window with a tapered window before computing the DFT. The tapered window will gradually reduce the signal towards the edges instead of abruptly truncating it. The signal will still be zero outside the interval. A tapered window can also be applied to the matched filter before pulse compression to suppress the sidelobe levels in the range dimension. However, applying a tapered window will widen and reduce the amplitude of the main lobe, reducing the resolution and signal-to-noise ratio (SNR). The consequences of applying the commonly used Hann window in the range and Doppler- dimensions are displayed in Figures 2.4.

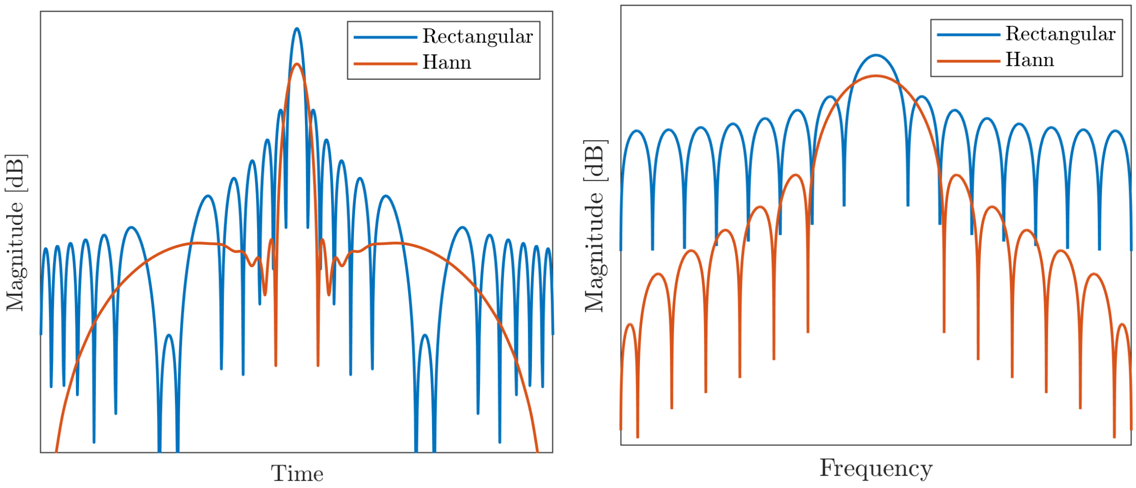

(a) Output after pulse compression with a rectangular window (blue) and Hann window (orange).

(b) Output after Fourier transforming a sinusoid with a rectangular window (blue) and Hann window (orange).

Figure 2.4: Comparison of the signals using a rectangular and Hann window.

The result after applying pulse compression in the fast-time dimension and the discrete Fourier transform in the slow-time dimension is the range Doppler map.

2.2 Detection Theory for Radars

As previously mentioned, the first step in the radar signal processing chain is detection. This step consists of accurately determining whether there are targets of interest present in the returning signal. The echoes sensed by the radar receiver will not only contain target returns, but also thermal receiver noise, returns from objects not of interest, and other disturbing signals. The task of the signal processor is therefore to decide between two hypotheses: the null hypothesis and the non-null hypothesis. The null hypothesis is that the received signal contains only noise. The non-null, or alternative, hypothesis is that at least one target is present in addition to the noise. If the processor determines that a target is present, further processing must determine its range and radial velocity [1, Ch. 6.2].

2.2.1 Noise and Clutter

The returning signal always contains some form of noise. Noise is generated both by external sources and by the receiver itself. The externally generated noise is primarily caused by solar and cosmic noise, and the internally generated noise is predominantly thermal noise from electronic components [1, Ch. 2.4]. In order for a target signal to be detectable, the amplitude of its peak must be larger than the peaks of the noise floor after processing. The strength of a target signal depends on several factors such as the radar cross-section (RCS) and the range to the scattering object. The RCS characterizes the reflective properties of the target and is governed by material composition and geometry of the object.

In addition to the inherent scattering strength of the object, the signal strength is inversely proportional to the range of the target following the expression $1/R^4$, where R is the distance to the target [11, Ch. 1]. The noise floor has the same average power across all ranges. Consequently, there will be a distance at which the magnitude of the target signal is not large enough to be distinguished from the noise. This phenomenon is depicted in Figure 2.5.

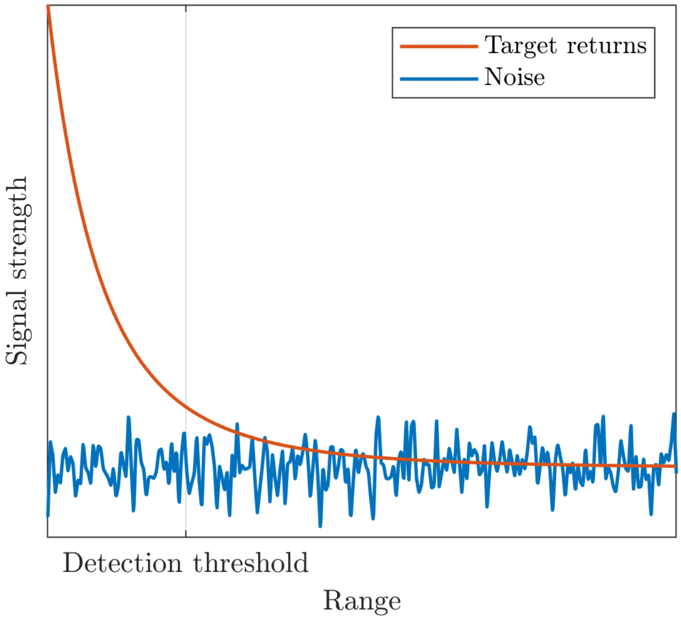

Figure 2.5: The signal strength from a target return (orange) decreases rapidly with range and will after some threshold be lost in the noise (blue).

Furthermore, all echoes sensed by the radar are not scattered off of targets of interest. The radar pulse scatters off of all objects it encounters such as terrain, buildings, birds, etc. Depending on the application, these returns can be both wanted or unwanted. All echoes from objects not of interest are called clutter. Unlike noise, the signal from clutter is not random. Instead, clutter looks like a response from a real target. This can make it difficult to determine whether a signal originates from clutter or a target, and the clutter signal may also interfere and mask echoes from interesting targets.

2.2.2 SNR and SIR

In order to evaluate system performance, ratios that compare the level of a desired signal to the level of undesired signals are used. Two such ratios are the Signal-to-Noise Ratio (SNR) and Signal-to-Interference Ratio (SIR).

The SNR is defined as the signal power, P_s, divided by the noise power P_n

$$\mathrm{SNR} = \frac{P_\mathrm{s}}{P_\mathrm{n}}. \tag{2.4}$$

For deterministic cases, the signal power can be defined as the squared amplitude of the signal peak for a specific time or the average signal power over time.

Similarly, the SIR is calculated as the signal power divided by the interference power P_i [1, Ch. 1.4.1]

$$\mathrm{SIR} = \frac{P_\mathrm{s}}{P_\mathrm{i}}. \tag{2.5}$$

The interference power can originate from the sidelobes of nearby targets or clutter, among other sources.

2.2.3 Probability of Detection and False Alarm

Two performance metrics used to evaluate the radar detection performance are the probability of detection, P_D, and the probability of false alarm, P_{FA}. The first represents the probability that the radar correctly identifies the presence of a target, while the other is the probability that it detects a target when there is none there. There is a fundamental trade-off between these two metrics. Increasing the sensitivity of the detector results in a higher P_D and more targets detected, but also increases the P_{FA} and the amount of false detections. In the same manner, decreasing the sensitivity of the detector results in a decrease of both the P_{FA} and the P_D. The appropriate sensitivity is a design choice that depends on the operator's requirements and the environment in which the radar is operated. To address this trade-off, there are adaptive detectors designed to keep a constant false alarm rate, called CFAR detectors.

2.2.4 Cell Averaging-Constant False Alarm Rate- CA-CFAR

Constant False Alarm Rate, often called CFAR, is a detector property commonly desired in a radar detector. Its primary purpose is to maintain a constant P_{FA} in the presence of noise and clutter with unknown and varying strength [1, Ch. 6.5]. The signal strength threshold for detection is not kept constant but is instead varied to maintain the CFAR property, regardless of the local noise and clutter strength.

One common version of a CFAR detector is called the Cell Averaging CFAR detector (CA-CFAR). The CA-CFAR detector adapts the detection threshold based on a local estimate of the mean noise and interference power. The data sample being evaluated is called the cell-under-test, or CUT. To determine whether the CUT contains a target the algorithm estimates the local noise level using neighboring cells, referred to as training cells. This noise level is computed by averaging the signal power in the training cells and represents an estimate of the background noise or clutter level in the vicinity of the CUT. To prevent main lobe or sidelobe interference from a potential target in the CUT the cells closest to it are excluded from the noise averaging. These cells are called the guard cells. There are several ways to choose the training and guard cells in 2-dimensional CA-CFAR. One common choice is the cross depicted in Figure 2.6.

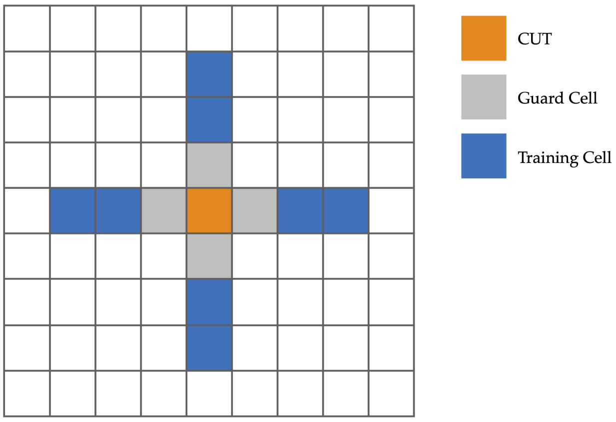

Figure 2.6: The choice of window model for the CA-CFAR. The orange cell is the CUT, the gray cells are the guard cells and the blue cells the training cells.

In CA-CFAR the threshold is determined by multiplying the estimate of the noise level, σ^2, with a constant, α, as seen in equation 2.6,

$$T = \alpha \cdot \sigma^2. \tag{2.6}$$

For a desired false alarm probability of P_{FA}, α is calculated according to

$$\alpha = N((P_{FA})^{-1/N} - 1) \tag{2.7}$$

where N is the total number of training cells used [1, Ch. 6.5.3].

This type of detector operates on each cell independently and has some known limitations [1, Ch. 6.5.4]. The first limitation comes from the fact that the CA-CFAR concept assumes isolated targets. When two targets are spaced such that they are present in each other's training cells it is possible that the estimated background becomes so large that one or both targets are missed. The second limitation emanates from the CA-CFAR assumption that the interference in the training cells and CUT is homogeneous. Interference from clutter can in general be very heterogeneous, which can lead to both false alarms and missed detections in regions where the clutter interference changes rapidly.

2.3 Greedy Algorithms and the Sparse Signal Model for Radar

In many radar applications it is assumed that the data is sparse and can be described by Equation 2.8

$$\boldsymbol{x} = \boldsymbol{S\theta} + \boldsymbol{n}. \tag{2.8}$$

Here, $\boldsymbol{x} \in \mathbb{C}^{M\times 1}$ is the measurement vector, $\boldsymbol{S} \in \mathbb{C}^{M\times N}$ is a dictionary scaled by the amplitudes $\boldsymbol{\theta} \in \mathbb{C}^{N\times 1}$ with $P \ll M$ non-zero elements, and $\boldsymbol{n} \in \mathbb{C}^{M\times 1}$ is additive noise [12]. The dictionary $\boldsymbol{S}$ has steering vectors as columns, each of which describe the echo of a point scatterer with some position and radial velocity. As a consequence of the sparse model, several algorithms have been proposed with the purpose of solving the problem

$$\hat{\boldsymbol{\theta}} = \arg\min_{\boldsymbol{\theta}} ||\boldsymbol{\theta}||_0, \text{ s.t. } ||\boldsymbol{x} - \boldsymbol{S}\boldsymbol{\theta}||_2^2 \leq \eta, \tag{2.9}$$

i.e. the sparsest solution for a given error tolerance η. Algorithms attempting to solve this problem are called compressive sensing algorithms and include regularization algorithms, iterative thresholding algorithms, Bayesian algorithms and greedy algorithms [12]. The greedy algorithms are often quite computationally efficient compared to the other approaches, making them popular in real-time processing systems. They work by choosing the best option at each intermediate step without reconsidering previous steps. One algorithm in the greedy family is Orthogonal Matching Pursuit, abbreviated OMP [13].

2.3.1 Orthogonal Matching Pursuit

Greedy algorithms for compressed sensing have been studied extensively. An early development was the CLEAN algorithm, initially developed in radio astronomy to more accurately reconstruct images of the sky obtained with radio interferometers [14]. However, the technique was later adopted for radar applications [15] with various developments [16], [17] to account for extended (not point-like) targets and pulse compressed data. These developments closely resemble the matching pursuit algorithm [18] for sparse signal recovery. These algorithms work by iteratively finding the strongest steering vector in the dictionary and removing its contribution to the signal. The convergence of the Matching Pursuit algorithm was quickly improved by an orthogonalization procedure known as Orthogonal Matching Pursuit [13]. In the OMP algorithm the residual signal after each iteration is made orthogonal to all previously found steering vectors. OMP is theoretically able to reliably recover the signal with an error that is only a small factor worse than the minimal attainable error with the same number of terms [19], [20]. OMP and OMP-like methods have also been successfully applied to detection in simulated radar data [2], [3]. And several generalizations and developments have been proposed such as selecting several steering vectors per iteration [4], employing look ahead strategies for steering vector selection [5], improved runtime by employing randomized steering vector selection [6] or refining the steering vector selection through Newton refinements [7].

The OMP algorithm in its most basic form is shown in Algorithm 1. First, the residual is initialized as the data $\tilde{\boldsymbol{x}} = \boldsymbol{x}$, the matrix of found steering vectors $\boldsymbol{F}$ as empty and the model order P as zero. Then the dictionary $\boldsymbol{S}$ is correlated with the data producing a correlation vector $\boldsymbol{E} = \boldsymbol{S}^\dagger \tilde{\boldsymbol{x}}$, where $(\cdot)^\dagger$ denotes the conjugate transpose, and the steering vector with the largest correlation $\boldsymbol{s}'$ is chosen. If the correlation strength of this steering vector exceeds a predetermined threshold T, the steering vector is appended to the matrix of found steering vectors and the residual is updated by projecting the space spanned by $\boldsymbol{F}$ out of the data, otherwise the algorithm is terminated. The matrix that projects the data out of the space spanned by $\boldsymbol{F}$ can formally be written as $\boldsymbol{P}_{\boldsymbol{F}}^{\perp} = \mathbf{1} - \boldsymbol{F}(\boldsymbol{F}^{\dagger}\boldsymbol{F})^{-1}\boldsymbol{F}^{\dagger}$, where $\mathbf{1}$ is the identity matrix.

Algorithm 1 Basic OMP Algorithm

```
Require: data $\boldsymbol{x}$, dictionary $\boldsymbol{S} = [\boldsymbol{s}_1, \dots, \boldsymbol{s}_K]$ and threshold $T$
Initialize $\tilde{\boldsymbol{x}} = \boldsymbol{x}$, $\boldsymbol{F} = [\cdot]$, $P = 0$
while running do
    $\boldsymbol{E} = \boldsymbol{S}^{\dagger}\tilde{\boldsymbol{x}}$
    $l = \underset{i=1,\dots,K}{\arg\max}|\boldsymbol{E}_i|$
    if $|\boldsymbol{E}_l| > T$ then
        $\boldsymbol{s}' \leftarrow \boldsymbol{s}_l$
        $\boldsymbol{F} \leftarrow [\boldsymbol{F}, \boldsymbol{s}']$
        $P \leftarrow P + 1$
        $\tilde{\boldsymbol{x}} \leftarrow \boldsymbol{P}_{\boldsymbol{F}}^{\perp}\boldsymbol{x}$
    else
        running ← false
    end if
end while
```

2.4 Evaluation Methods and Performance Metrics

The following section describes the theory behind some performance metrics commonly used for radar target detectors.

2.4.1 Root Mean Square Deviation

The Root Mean Square Deviation (RMSD), also know as Root Mean Square Error, is defined according to Equation 2.10. It measures the difference between the estimated values, e, and their corresponding true values, t. Each deviation is squared to avoid positive and negative values canceling each other out before summation.

$$\text{RMSD} = \sqrt{\frac{1}{n}\sum_{i=1}^{n}|e_i - t_i|^2} \tag{2.10}$$

where n is the number of observations.

2.4.2 Cramér-Rao Lower Bound

The Cramér-Rao Lower Bound, often abbreviated CRLB, determines the lowest theoretical variance for an unbiased estimator [1, Ch. 7.1]. An estimator is considered unbiased if the expected value of the estimated parameters equals the true parameters. Let $p_{\boldsymbol{x}}(\boldsymbol{x}|\theta)$ denote the joint probability density function of the data $\boldsymbol{x}$ given the parameter θ. The CRLB can then be expressed as

$$\sigma_{\hat{\theta}}^2 \geq \frac{1}{\mathbb{E}\left[\left(\frac{\partial}{\partial\theta}\ln p_{\boldsymbol{x}}(\boldsymbol{x}|\theta)\right)^2\right]}.$$

For a complex-valued, time delayed signal in Additive White Gaussian Noise (AWGN) after matched filtering the CRLB for the time delay parameter τ evaluates to

$$\sigma_{\hat{\tau}}^2 = \frac{1}{8\pi^2\chi B_{\text{rms}}^2}\ \text{s}^2$$

where χ is the SNR and B_{rms}^2 is the root-mean-square bandwidth of the signal [1, Ch. 7.2]. For a linear frequency modulated signal $B_{\text{rms}}^2 \approx B^2/12$ where B is the bandwidth in equation 2.3. Converting time delay to range R, see equation 2.1, the CRLB in range for a linear frequency modulated signal after matched filtering becomes

$$\sigma_{\hat{R}}^2 = \frac{3c^2}{8\pi^2\chi B^2}.$$

For the sinusoidal signal in AWGN describing the Doppler shift f_D, the frequency CRLB after a discrete Fourier transformation is described by

$$\sigma_{\hat{f_D}}^2 = \frac{6}{(2\pi T_s)^2 N^3 \chi}$$

where N is the number of pulses, T_s is the PRI and χ is the SNR [1, Ch. 7.2]. The CRLB for Doppler frequency can be converted to the Doppler velocity through equation 2.2.

2.4.3 Figure of Merit

Estimations from a radar detector are not always correct. Some detections are false alarms where the detector has decided a target is present although only noise or other interfering signals are. Another error occurs when the detector fails to identify targets in the signal. The Figure of Merit (FoM) can be used to evaluate the performance of the detector and catch the influence of these errors. If N_d is the number of correct detections, N_{fa} the number of false alarms and N_t the number of existing targets the equation 2.11 describes the FoM [21].

$$\text{FoM} = \frac{N_d}{N_{fa} + N_t} \tag{2.11}$$

3

Methods

This chapter details the methods and procedures used to develop, adapt, and evaluate the OMP algorithm. The work consisted of three main parts. Firstly, alterations beyond the basic implementation of the OMP detector were introduced to improve parameter estimation accuracy, clutter handling, and computational efficiency. Secondly, the tunable parameters of both the CA-CFAR and OMP detector were investigated in order to achieve good detection performance and comparable operating conditions. A proposed hybrid algorithm combining the two detectors was also presented. Finally, the test cases aming to evaluate the OMP detector are described along with a motivation for each test.

3.1 OMP Implementation

The implementation of the OMP algorithm closely resembled the pseudocode of the basic OMP algorithm in Algorithm 1, but with some modifications. The complete implementation is described in detail in the following sections.

3.1.1 Choice of Dictionary and Steering Vectors

A fundamental part of the OMP algorithm is the choice of dictionary and steering vectors. As described in Section 2.1, the return from a target, $s(t)$, was modeled as a time-delayed and Doppler shifted copy of the reference pulse,

$$s(t; \tau, f_D) = e^{j2\pi f_D t} f(t - \tau) \tag{3.1}$$

where τ corresponded to the range of the target according to Equation 2.1, f_D to the radial velocity of the target according to Equation 2.2, and $f(t)$ denoted the pulse train reference. Since the steering vectors constituting the dictionary should describe the expected signal from a target, the steering vectors were chosen as time-discretized versions of Equation 3.1 for different τ and f_D. Time was discretized according to $t = k\Delta t_s + l\Delta t_f$ for $k = 0, \ldots, N_s - 1$ and $l = 0, \ldots, N_f - 1$. Here, Δt_s was the sampling interval in slow-time (the PRI), N_s was the number of pulses per CPI, Δt_f was the sampling interval in fast-time, and N_f the number of fast-time samples. Thus, element (k, l) of the steering vector took the form

$$s_{k,l}(\tau, f_D) = e^{j2\pi f_D(k\Delta t_s + l\Delta t_f)} f\Big((k\Delta t_s + l\Delta t_f) - \tau\Big).$$

It was assumed that the Doppler-induced phase variation during a pulse was negligible, i.e. $2\pi f_D l \Delta t_f \ll 1$ such that $e^{j2\pi f_D l \Delta t_f} \approx 1$. Furthermore, range-bin migration was assumed negligible such that $f(t)$ was periodic in Δt_s. This made the steering vectors separable in the fast-time and slow-time dimensions,

$$s_{k,l}(\tau, f_D) \approx e^{j2\pi f_D k \Delta t_s} f\Big(l\Delta t_f - \tau\Big) \tag{3.2}$$

which was rewritten using the Kronecker product. The Kronecker product [22, Eq. (2)] of $\boldsymbol{A} \in \mathbb{C}^{m\times n}$ and $\boldsymbol{B} \in \mathbb{C}^{p\times q}$ is a $pm \times qn$ matrix defined as

$$\boldsymbol{A} \otimes \boldsymbol{B} = \begin{bmatrix} a_{11}\boldsymbol{B} & \dots & a_{1n}\boldsymbol{B} \\ \vdots & \ddots & \vdots \\ a_{m1}\boldsymbol{B} & \dots & a_{mn}\boldsymbol{B} \end{bmatrix}.$$

Using the definition above, the steering vectors could be written as

$$\boldsymbol{s}(\tau, f_D) = \boldsymbol{r}(\tau) \otimes \boldsymbol{d}(f_D) \tag{3.3}$$

for the range steering vector $\boldsymbol{r}(\tau) \in \mathbb{C}^{N_f \times 1}$ with elements $r_l(\tau) = f(l\Delta t_f - \tau)$ and the Doppler steering vector $\boldsymbol{d}(f_D) \in \mathbb{C}^{N_s \times 1}$ with elements $d_k(f_D) = e^{j2\pi f_D k \Delta t_s}$.

The dictionary $\boldsymbol{S}$, consisting of steering vectors described by Equation 3.3 for the time delays $\tau = i\Delta t_f,\ i = 0, \dots, N_f - 1$ and Doppler frequencies $f_D = m/(N_s \Delta t_s),\ m = 0, \dots, N_s - 1$, enabled the correlation step with data $\boldsymbol{x}$, $\boldsymbol{S}^\dagger \boldsymbol{x}$, of the OMP algorithm to be implemented efficiently using fast Fourier transforms, see Appendix A.1 for details. The resulting calculations corresponded exactly to conventional signal processing of pulse-Doppler radar, i.e. pulse compression along the fast-time dimension and discrete Fourier transforming along the slow-time dimension. Furthermore, the separable nature of the steering vectors in Equation 3.3 enabled the projection operation of the OMP algorithm to be performed more efficiently, see Appendix A.2 for details.

3.1.2 Sub-bin Parameter Estimation

With the dictionary described above, the parameters τ and f_D of a steering vector could only be determined on the range-Doppler grid described by $\tau = i\Delta t_f,\ i = 0, \dots, N_f - 1$ and $f_D = m/(N_s \Delta t_s),\ m = 0, \dots, N_s - 1$. However, targets were not guaranteed to be located exactly on the grid points. The true target position could deviate by up to half a bin in each direction from the grid point. In order to achieve sub-bin resolution, the range-Doppler grid was more densely sampled at points with high correlation strength.

Finer sampling was accomplished by generating finely spaced steering vectors and calculating the correlation strength explicitly using the inner product. For each candidate coarse grid point, two independent one-dimensional sweeps were performed in each dimension. In each sweep, N evenly spaced steering vectors within $\pm\Delta t_f/2$ in the time-delay dimension and $\pm(2N_s\Delta t_s)^{-1}$ in the Doppler dimension of the coarse grid estimate were generated. For each steering vector $\boldsymbol{s}$, the correlation $\boldsymbol{s}^\dagger\boldsymbol{x}$ was evaluated directly. Figure 3.1 shows both the finely sampled grid (blue) and the coarse grid (orange) along with the improved estimates.

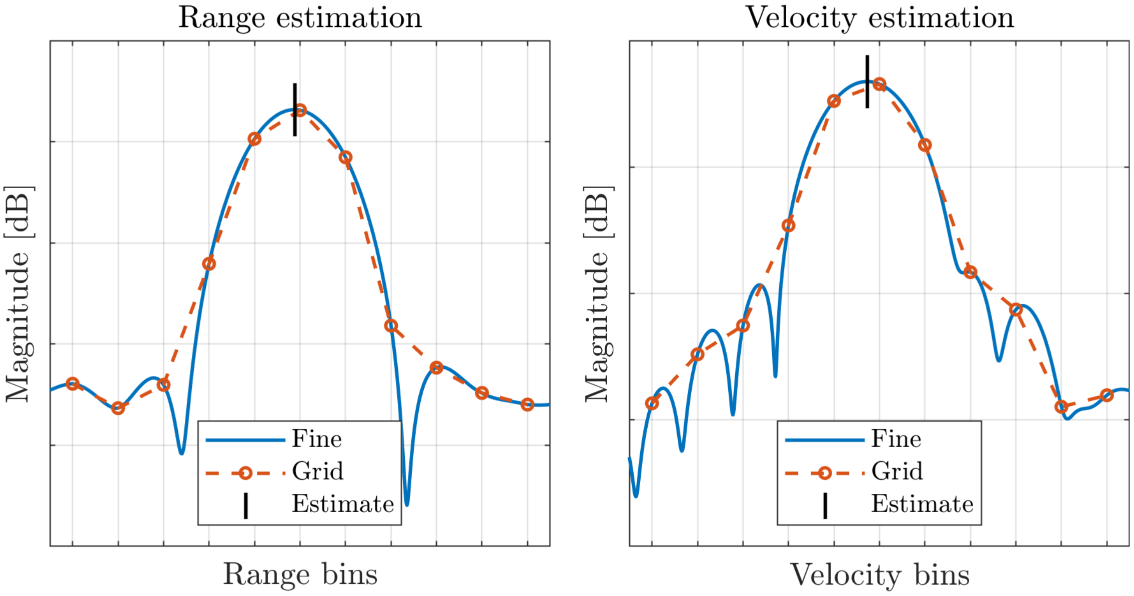

Figure 3.1: The finely sampled grid (blue) and the coarse grid (orange) along with the improved estimates.

In order to determine which points to resolve more finely, the $n_{test} \geq 1$ strongest points in the coarse range-Doppler map were chosen. However, since several of the largest values in the range-Doppler map corresponded to the main lobe or sidelobes of a single target, a clustering approach was employed. Instead of finding the n_{test} largest values in the range-Doppler map, the k largest values for $k \geq n_{test}$ were found. These values were then clustered such that only the strongest point of each contiguous cluster was kept. The coarse grid was resolved more finely around the n_{test} strongest of the clustered points, see Figure 3.2. The left subfigure shows the result if only the $n_{test} = 5$ largest values are chosen to sample more finely, the middle subfigure shows the $k = 100$ strongest points, while the right subfigure shows the $n_{test} = 5$ retained values after clustering.

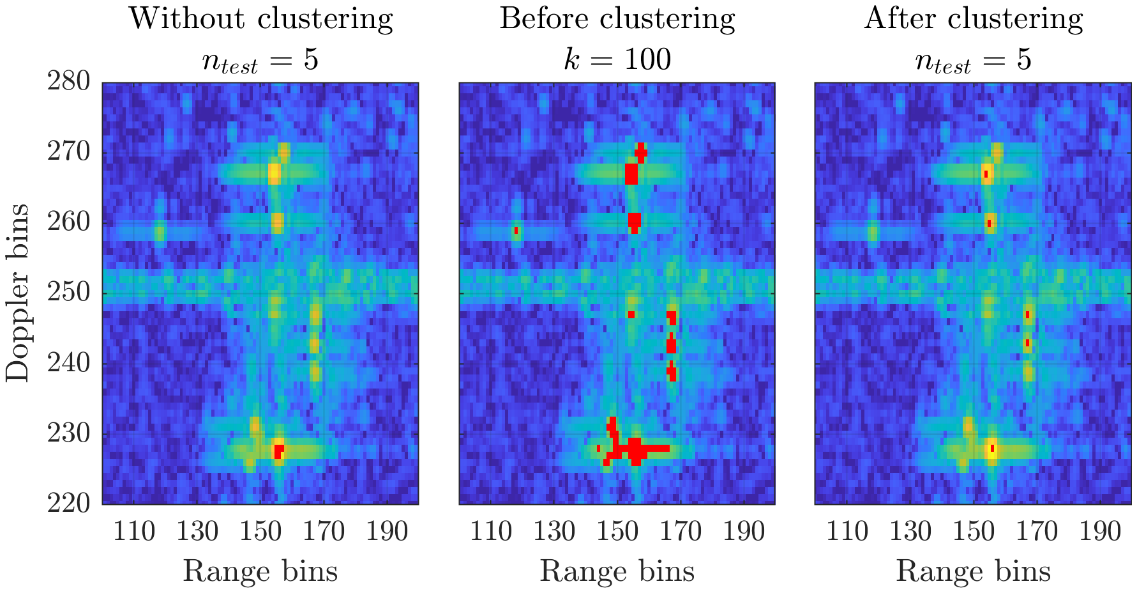

Figure 3.2: The left subfigure shows the result if only the $n_{test} = 5$ largest values are chosen to sample more finely, the middle subfigure shows the $k = 100$ strongest points, while the right subfigure shows the $n_{test} = 5$ retained values after clustering. The corresponding selected points in each case are highlighted in red.

3.1.3 Clutter Handling

The OMP algorithm was further modified to improve clutter handling. Clutter affected both the estimation of the steering vector parameters in the correlation step and the amplitude estimation in the projection step.

Regarding the parameter estimation in the correlation step, the possibility of windowing was incorporated. As explained in Section 3.1.1 and Appendix A.1, with the previously described dictionary, the correlation step corresponded to standard range-Doppler processing, which enabled windowing techniques to be used in both the range and Doppler dimensions. This reduced the interference from clutter and nearby targets when estimating the target position parameters.

Furthermore, two operational modes, *clutter removal* and *detection*, were added to the algorithm. In the *clutter removal* mode, the OMP algorithm was constrained to find targets in the zero-Doppler and its directly adjacent bins. In the *detection* mode, the algorithm was instead restricted from detecting targets in the zero-Doppler bin and its three nearest neighbors on each side.

3.1.4 Final Pseudo-Code

The complete implementation of the adapted OMP algorithm is seen as pseudocode in Algorithm 2.

Algorithm 2 Adapted OMP Algorithm

```
Require: Data $\boldsymbol{x}$, dictionary $\boldsymbol{S} = [\boldsymbol{s}_1, \ldots, \boldsymbol{s}_K]$, parameters: $T$, $k$, $n_{test}$, $N$
Initialize $\tilde{\boldsymbol{x}} = \boldsymbol{x}$, $\boldsymbol{F} = [\ ]$, $P = 0$, running = true
while running do
    $\boldsymbol{E} \leftarrow$ range_doppler_process($\tilde{\boldsymbol{x}}$)
    if mode == $\textit{clutter removal}$ then
        $\boldsymbol{E} \leftarrow$ keep_only_clutter_bins($\boldsymbol{E}$)
    else
        $\boldsymbol{E} \leftarrow$ mask_clutter_bins($\boldsymbol{E}$)
    end if
    $\mathcal{C} \leftarrow$ find_largest_points($\boldsymbol{E}, k$)
    $\mathcal{C} \leftarrow$ cluster($\mathcal{C}, n_{test}$)
    $E_{\max} \leftarrow 0$, $\hat{\boldsymbol{s}} \leftarrow$ null
    for $c \in \mathcal{C}$ do
        $(\tau, f_D) \leftarrow$ fine_sample_estimate$\Big(\boldsymbol{E}(c), N\Big)$
        $\boldsymbol{s} \leftarrow$ get_steering_vector($\tau, f_D$)
        if $|\boldsymbol{s}^\dagger \tilde{\boldsymbol{x}}| > E_{\max}$ then
            $E_{\max} \leftarrow |\boldsymbol{s}^\dagger \tilde{\boldsymbol{x}}|$
            $\hat{\boldsymbol{s}} \leftarrow \boldsymbol{s}$
        end if
    end for
    if $E_{\max} > T$ then
        $\boldsymbol{F} \leftarrow [\boldsymbol{F}, \hat{\boldsymbol{s}}]$
        $P \leftarrow P + 1$
        $\tilde{\boldsymbol{x}} \leftarrow$ project_out($\boldsymbol{F}, \boldsymbol{x}$)
    else
        running $\leftarrow$ false
    end if
end while
```

3.2 Tuning of the Detectors

The performance for both the OMP and CA-CFAR detectors was dependent on the selection of several tunable parameters. These parameters affected important characteristics of the detector, such as sensitivity to noise and clutter, detection accuracy, and the ability to detect weak targets.

In general, the operating conditions and radar system configuration determine the appropriate parameter values. System-related factors, such as bandwidth, pulse duration, sampling frequency and pulse waveform, influence how a detector responds to target echoes and interference. Environmental factors like clutter density, target distribution, and noise level also affect suitable parameter settings.

In practice, parameter tuning usually involves trade-offs and balancing competing aims. For example, parameter settings that improve detection sensitivity typically raise the probability of false alarm, while parameters that provide robustness in clutter-rich environments tend to reduce resolution. In addition, in computationally demanding scenarios, performance is sometimes sacrificed for computational efficiency.

3.2.1 The OMP Detector

For the OMP detector, an important tunable parameter was the number of fine-grid points in the sub-bin parameter estimation, described in Section 3.1.2. In order to determine a suitable number, a single target in AWGN was simulated, and tests with a varying number of fine-grid points were conducted. The number of fine-grid points was set to $N = 200$ for all cases. Unless otherwise noted, the number of largest points found in the coarse grid before clustering, k, was set to 100. The number of points sampled more finely was set to $n_{test} = 10$ in all test cases.

To suppress false alarms caused by sidelobes or residual energy after signal removal, nearby detections could be grouped together. The grouping was conducted by finding all detections separated by at most a fixed range and velocity difference, and retaining only the detection with the strongest signal strength in each such connected cluster. The range and velocity threshold for the grouping were set to range and velocity resolution. Larger thresholds would remove detections further away from the true target, but increase the risk of removing a correct detection when targets are closely spaced.

Furthermore, the effect of varying the number of clutters removed in the *clutter removal* mode was investigated, and whether this depended on the choice of windowing used. After investigating this the choices that yielded the best FoM were used in the final evaluation tests.

3.2.2 The CA-CFAR Detector

Two of the key parameters to calibrate for the CA-CFAR detector were the number of training and guard cells. More training cells lead to more stable noise estimates with reduced variance, which generally improves detection performance in homogeneous environments with sparsely distributed targets. However, in non-homogeneous environments, a large number of training cells increases the risk that the estimate becomes biased due to other signal responses, potentially causing missed detections. On the other hand, more guard cells reduce target energy leakage contaminating the noise estimate, but it also reduces the available training cells and therefore increases the variance.

In the CA-CFAR detector used, the training and guard cells were arranged in accordance with the configuration shown in Figure 2.6. The number of training and guard cells in both range and velocity dimensions were tuned to maximize FoM on simulated multi-target scenarios. In range, 1 guard cell and 10 training cells were used, and in velocity, 5 guard cells and 10 training cells were used.

In the vicinity of an isolated target, the CA-CFAR detector, like the OMP detector, often identified a cluster of neighboring cells instead of a single detection. To address this, the algorithm can be configured to keep only the cell with the highest signal strength within each group of connected detections and remove the rest. Detections were considered connected if they were in neighboring bins.

To enable a fair comparison between the CA-CFAR detector and the OMP detector, both algorithms were calibrated to achieve the same probability of false alarm. The constant α in formula 2.6 controlled the probability of false alarm for CA-CFAR separately for the range and velocity dimensions. Since the threshold in both dimensions had to be met for a detection, the combined P_{FA} was slightly lower than the P_{FA} in each dimension. The total probability of false alarm, P_{FA}^{tot}, was estimated by generating data containing only noise and determining the probability that the detector falsely identified a target over many iterations. Using the probability of false alarm obtained from the CA-CFAR detector, the detection threshold for the OMP detector was selected to achieve the same. This was done by evaluating $\boldsymbol{s}^{\dagger}\boldsymbol{x}$ on noise-only data several times and setting the threshold to the $1 - P_{FA}^{tot}$ percentile of the gathered values.

3.3 The proposed CA-OMP Detector

The OMP algorithm implementation, see Algorithm 2, used a constant energy threshold T as the detection threshold. The appropriate value of T was selected for a specific probability of false alarm given either a known noise distribution or sufficient noise-only training data. In general, however, the noise background is not homogeneous, see Section 3.2.2, motivating the use of an adaptive threshold. Furthermore, by modifying the greedy procedure and selecting several steering vectors per iteration the computational efficiency can be improved [4]. Based on these two observations, a modified detection algorithm, denoted CA-OMP, was proposed. The method is a hybrid between the CA-CFAR detector and the OMP algorithm, see Algorithm 3.

In each iteration of the CA-OMP algorithm, a CA-CFAR detector was applied on the residual data $\tilde{\boldsymbol{x}}$, which resulted in several detections on the coarse grid. Next, the parameter estimates for all detections were improved with the finer sampling described in Section 3.1.2. Steering vectors corresponding to the refined parameter estimates were then created and appended to the matrix of previously found steering vectors. Finally, all previously found steering vectors were projected out of the data, and the procedure was repeated until the CA-CFAR detector generated no detections.

Algorithm 3 Proposed CA-OMP Algorithm

```
Require: data $\boldsymbol{x}$, dictionary $\boldsymbol{S} = [\boldsymbol{s}_1, \ldots, \boldsymbol{s}_K]$, Parameters: $N$
Initialize $\tilde{\boldsymbol{x}} = \boldsymbol{x}$, $\boldsymbol{F} = [\cdot]$, $P = 0$, running = true
while running do
    $\boldsymbol{E} \leftarrow$ range_doppler_process($\tilde{\boldsymbol{x}}$)
    $\mathcal{D} \leftarrow$ run_CA-CFAR($\boldsymbol{E}$)
    if length($\mathcal{D}$) $> 0$ then
        for $d \in \mathcal{D}$ do
            $(\tau, f_D) \leftarrow$ fine_sample_estimate$\Big(\boldsymbol{E}(d), N\Big)$
            $\boldsymbol{s} \leftarrow$ get_steering_vector($\tau, f_D$)
            $\boldsymbol{F} \leftarrow [\boldsymbol{F}, \boldsymbol{s}]$
            $P \leftarrow P + 1$
        end for
        $\tilde{\boldsymbol{x}} \leftarrow$ project_out($\boldsymbol{F}, \boldsymbol{x}$)
    else
        running $\leftarrow$ false
    end if
end while
```

3.4 Simulations of Radar Data

Assessing the detectors' performance using real-world data comes with difficulties due to the absence of a ground truth that distinguishes targets, clutter and noise. Additionally, conducting controlled experiments, such as scenarios with only one or a few signal responses in noise or targets positioned at specific distances from each other, is both time-consuming and difficult to implement in practice. As a result, simulated radar data is a good addition to complement real measurements and properly evaluate the detectors.

Radar data were simulated as signal responses following formula 3.2 embedded in AWGN. For simulations that aimed to mimic the real data scenario, the signal strength was randomized but decreasing in range, R, according to the $1/R^4$ -relationship described in Section 2.2.1. In these simulations, the system parameters were set to match those of real data. Furthermore, the post-processing SNR of targets was limited to the range 16 - 80 dB, chosen to resemble values observed in real-world data. Clutter responses were modeled the same way as target responses, but clutter was restricted to zero velocity and with a power restricted between 10 -100 dB compared to the noise power. Clutter signals were placed densely at random ranges. Generally there were approximately as many clutter signals as there were range bins. Moreover, the clutter strength decreased with range according to the same relationship as the targets. In contrast to clutter, target velocities were randomized according to a standard normal distribution, making low velocities more probable than high. However, targets were not allowed to have velocities very close to zero, since the *detection*-mode described in Section 3.1.3 was unable to detect targets in that region. The phases of the target and clutter signals were drawn from a uniform distribution between 0 and 2π.

3.5 Acquisition of Real Radar Data

Saab Surveillance provided pulse-Doppler radar data for this thesis. The dataset consisted of a single measurement collected prior to the start of the project. Table 3.1 contains the relevant system parameters of the data. Note that the ratio between bandwidth and sampling frequency was 1.25.

Table 3.1: System parameters for the real data.

Sampling frequency	10 MHz
PRF	5 kHz
Pulse length	2 µs
Pulse bandwidth	8 MHz
Number of pulses per CPI	500
Length of measurement	100 CPIs with 1 s periodicity

Figure 3.3 shows one CPI from the collected dataset, together with commentary indicating the most likely source of the observed signals. A road existed within the field of view of the radar at an approximate linear range of 2.7 km, meaning strong detections at this range are probably cars. In addition, visual observation confirmed that birds, mostly swifts, were very common at most ranges.

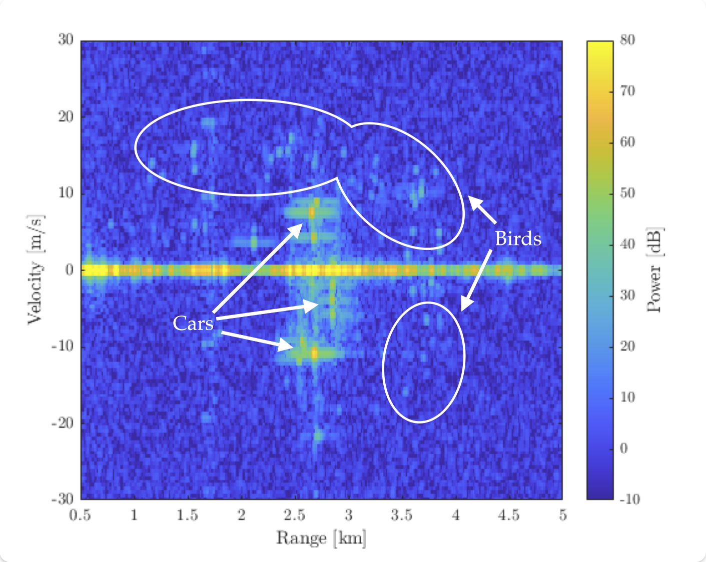

Figure 3.3: One CPI of the collected data. The strong peaks between 2.5 and 3 km are most likely cars. The weaker signals in the entire map are most probably birds. The strong band at 0 m/s is stationary clutter.

3.6 Test Cases

To evaluate the performance of the OMP algorithm, a number of test cases were created and analyzed. The tests ranged from simple scenarios involving one or a few targets in the presence of AWGN, to more complex and realistic scenarios containing multiple targets and clutter. Some test cases were applicable only for simulated data, while others were applicable to real data as well. In several of the test cases, the performance of the OMP algorithm was compared to that of the CA-CFAR detector and in some tests the CA-OMP detector too.

Each test was repeated a large number of times and the reported results correspond to the mean over all Monte Carlo iterations.

Some tests counted the number of correct, incorrect and missed detections. A detection was considered incorrect if the detected position was further away from the true position than the range or velocity resolution in the respective dimensions. A target was deemed missed if there was no detection within the aforementioned tolerance. All targets that were not assigned missed were considered correct.

Both the OMP and the CA-CFAR detector were calibrated for a false alarm probability of $\sim 1.55 \cdot 10^{-5}$ for all tests.

3.6.1 Test Case 1: Signal Removal at an Offset Position

In this test case, data was simulated containing one target with randomized range, velocity and phase. During signal removal, the steering vector was intentionally chosen with an offset in either range or velocity relative to the true target parameters. The removed energy was calculated. The purpose of this test was to evaluate the OMP algorithm's sensitivity to errors in the range and velocity estimation.

3.6.2 Test Case 2: Detection of One Target in Varying SNRs

For a set of SNRs, data containing a single target in AWGN was generated. In each iteration, the target's range, velocity and phase were randomized. The OMP algorithm was constrained to perform a single detection. The purpose of this test was to determine the fundamental detection ability of the OMP algorithm and its dependence on SNR.

3.6.2.1 Test Case 2a: Probability of Detection

In the first variant of Test Case 2, the probability of detection was calculated as a function of SNR. This test was carried out for the CA-CFAR detector as well as the OMP detector.

3.6.2.2 Test Case 2b: RMSD

In the second version of Test Case 2, the RMSD for the estimated target parameters was calculated and compared to the CRLB for different SNRs. Since the CA-CFAR detector operated on a coarser grid than the OMP detector a direct comparison of RMSD was not meaningful. Therefore, this test was only conducted for the OMP detector.

3.6.3 Test Case 3: Detection of Two Targets in Close Proximity

In this test case, data containing two targets was simulated. The separation between the targets was varied in either range or velocity while the other parameter was the same for the two targets. In each iteration, the position in the dimension that was not investigated and the phases of the targets were randomized. The number of detections was limited to two.

For each target separation the RMSD and energy removed was calculated. Additionally, the RMSD was compared to the CRLB of one target. When the two targets stop interfering with each other the RMSD should in theory be able to reach this CRLB.

The aim of this test was to investigate how the proximity of two targets, both in range and velocity, affected the OMP algorithm's possibility to correctly identify and remove the target signals.

The test was conducted under the following circumstances:

- Targets with equal strengths,
- Targets with different strengths,
- Rectangular and Hann windowed data.

3.6.4 Test Case 4: OMP vs. CA-CFAR Detector for Two Targets in Close Proximity

This test aimed to compare the detection performance of the OMP and CA-CFAR algorithms in scenarios involving two closely spaced targets. The targets were separated in either range or velocity and the separation was varied throughout the test. The phases of the targets were randomized for every iteration.

For every separation the algorithms' ability to detect both targets after clustering was assessed by counting the number of detections closer to the true position than the range and velocity resolution. Detections further away from the true targets were not counted.

3.6.5 Test Case 5: Multiple Targets in Presence of Clutter

In this scenario, data containing multiple targets and clutter were used. The OMP algorithm was executed iteratively and terminated when the residual energy fell below a threshold corresponding to a specified probability of false alarm.

The resulting detections were compared to those obtained using the CA-CFAR detector described in Section 3.2.2.

3.6.5.1 Test Case 5a: Simulated Data

For simulated data the true locations of the targets were known. After tuning the detectors the following performance metrics were determined and compared for the OMP, CA-CFAR and CA-OMP detectors:

- Number of correctly detected targets,
- Number of incorrect targets/false alarms,
- Number of missed targets,
- FoM.

3.6.5.2 Test Case 5b: Real Data

For real data, the results obtained using different detectors were analyzed and discussed. Since the true positions and number of targets were unknown, only qualitative comparisons could be made.

4

Results and Interpretations

Presented below are the results following the method and discussion in Sections 3.2 and 3.6. Following after each result is a short interpretation of it. The results are divided into two sections. Firstly, Section 4.1 displays the results from tests done to optimize the performance of the detection algorithms. Secondly, Section 4.1 contains the results of all test cases aiming to evaluate the OMP algorithm.

4.1 Tuning of the Detectors

The result of varying the number of grid points, N, in the fine parameter estimation, see Section 3.1.2, is shown below. Figure 4.1 presents the RMSD of the parameter estimates using the coarse grid (orange) and the fine grid (blue), along with the corresponding CRLB (black). In range the RMSD of the fine grid plateaued near the CRLB at around $N = 500$ points, and in velocity after about $N = 800$ points. Figure 4.2 depicts the removed energy using the parameter estimates of the coarse grid (orange) and the fine-grid (blue) for different N. For $N \gtrsim 200$ the energy removal was close to the original signal SNR of 50 dB. Since the average energy removal on the coarse grid was only ∼5 dB the estimation on the finer grid was needed. Due to the rapid increase in run-time with more fine grid points and the fact that the removed energy stabilized after $N = 200$, the number of fine-grid points in all following tests was set to $N = 200$.

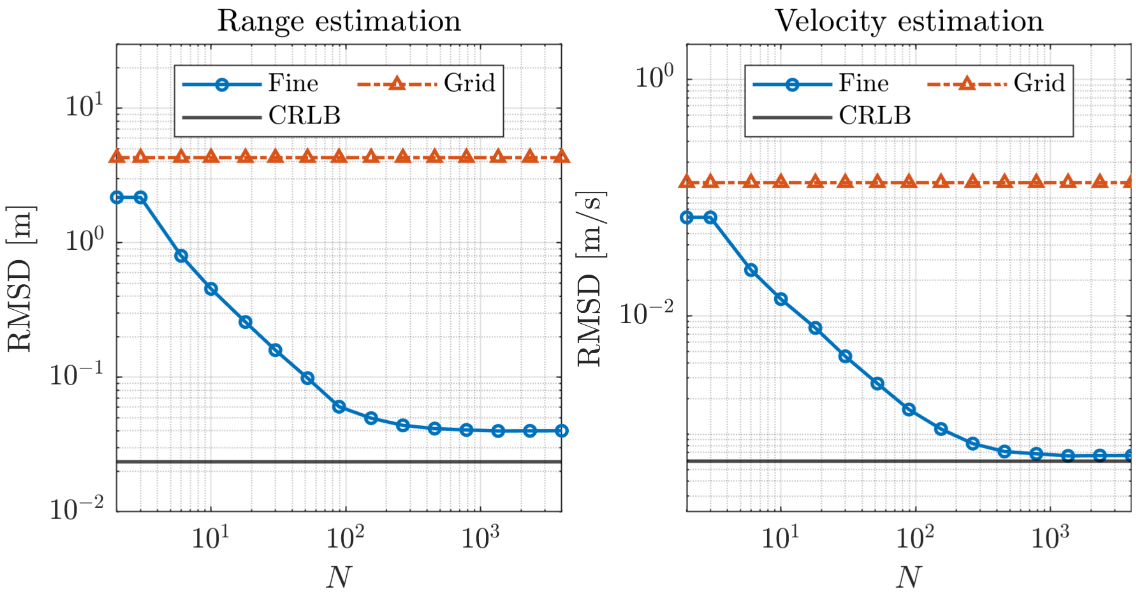

Figure 4.1: RMSD of the parameter estimates using the coarse grid (orange) and the fine grid (blue), along with the corresponding CRLB (black). The RMSD plateaued near the CRLB for $N \sim 500$ in range and $N \sim 800$ in velocity.

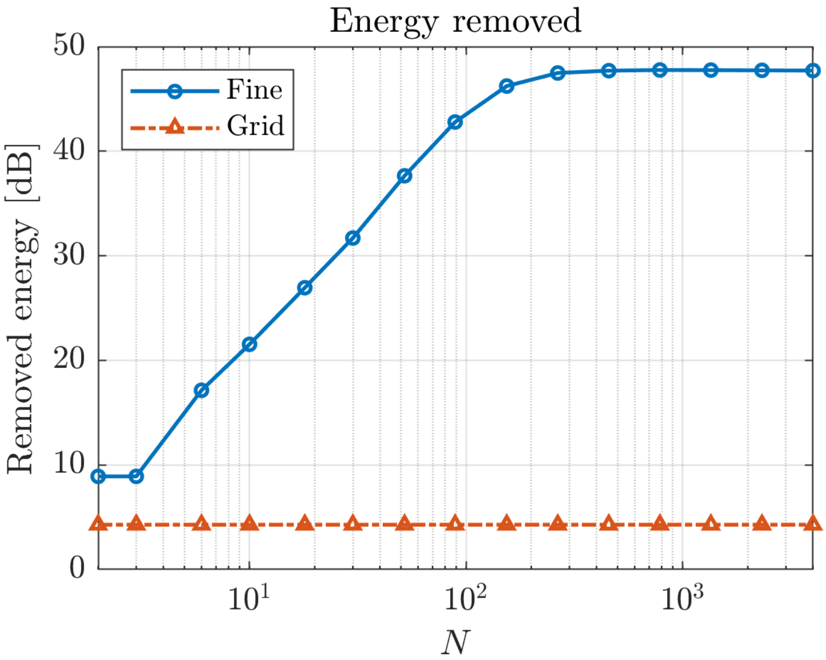

Figure 4.2: Removed energy using the parameter estimates of the coarse grid (orange) and the fine-grid (blue) for different N. The signal had an SNR of 50 dB. The removal approaches the maximum 50 dB after $N \sim 200$.

To determine a sufficient number of clutter signals to be removed by the OMP algorithm when using the *clutter removal mode* a test was conducted. In this test the algorithm removed between 0 and 400 clutter signals, in increments of 50, before searching for targets in the data. Following each run, the number of correct, incorrect, and missed targets were noted and the FoM calculated. The generated data contained 500 clutter signals and 40 target signals. All reported results were averaged over 100 Monte Carlo iterations.

Initially, the test was conducted with Hann windowing during the range-Doppler processing. Table 4.2 presents the detection results from clutter-free maps. The results after clutter removal, shown in Table 4.1, indicate that increasing the number of removed clutter scatterers prior to detection does not significantly improve the number of correct and missed detections, but decreases the number of false alarms. Since the CA-CFAR detector was calibrated to maximize the FoM, the number of removed clutter signals by the OMP algorithm was selected using the same criterion. Based on these results, the number was set to 250 in following tests.

Table 4.1: Performance metrics of the OMP detector for different number of clutters removed prior to detection. The range-Doppler processing used Hann windowing.

Clutters Removed	FoM	Correct	Incorrect	Missed
0	0.5888	38.76	29.01	1.24
50	0.8587	38.83	6.24	1.17
100	0.8934	38.87	4.30	1.13
150	0.9123	38.89	3.32	1.11
200	0.9133	38.91	3.34	1.09
250	0.9193	38.89	2.94	1.11
300	0.9186	38.90	2.96	1.10
350	0.9076	38.89	3.60	1.11
400	0.9089	38.90	3.39	1.10

Table 4.2: Performance metrics of OMP for a clutter-less map using Hann windowing in the range-Doppler processing.

FoM	Correct	Incorrect	Missed
0.9143	38.90	3.29	1.10

To evaluate if windowing with a tapered window was necessary to achieve better performance, the same experiment was repeated using a rectangular window in the range-Doppler processing instead. The results for different numbers of removed clutter signals are presented in Table 4.3, while Table 4.4 shows the corresponding results in clutter-less maps. The obtained FoM was lower for the rectangular windowing compared to the Hann windowing in maps both with and without clutter. This reduction is due to the larger number of incorrect detections, since the numbers of correct and incorrect detections are largely independent of the windowing function. This behavior is expected, as the OMP detector is more susceptible to detecting target sidelobes in rectangular windowed data than in Hann windowed.

A significant improvement in FoM can be observed when increasing the number of removed clutter signals from 0 to 50 for both windowing techniques. In contrast to the Hann windowed version, where the improvement is due to a reduction in incorrect detections, the increase is mainly caused by more correct detections for the rectangular windowed version. For Hann windowed processing, the larger number of incorrect detections when no clutter is removed is caused by detections of clutter sidelobes. For rectangular windowed processing the reason for fewer found targets is probably due to clutter sidelobes interfering with the target signals.

For both windowing functions, the number of incorrect detections begins to increase again after a certain amount of clutter signals have been removed, resulting in a corresponding decrease in FoM. The effect is more pronounced for rectangular windowed data. A likely explanation is that the algorithm eventually starts to identify residual components of previously removed clutter signals. Consequently, the clutter suppression becomes less accurate and introduces additional artifacts into the processed data that might leak out into the target detection region.

It should also be noted that clutter removal before target detection, for Hann windowed processing, can achieve performance comparable to that obtained from detection in a clutter free map. This is not the case for rectangular windowed processing. Since the clutter signals are densely spaced, they will interfere which hinders accurate parameter estimation. As a result, steering vectors corresponding to clutter signals are often generated with slight offset from the clutter signals' true position, leading to non-ideal energy removal. Because all clutter energy cannot be completely eliminated, it will continue to leak into the detection region, where the OMP detector might detect clutter sidelobes as false targets. This explains the higher number of incorrect targets compared to the clutter free case. For Hann windowed data the clutter sidelobes will be sufficiently suppressed to not leak into the detection area.

As the results obtained using Hann windowed data were significantly better, the Hann window was selected as the standard for future tests.

Table 4.3: Performance metrics of OMP as a function of the number of clutters removed for rectangular windowed data.

Clutters Removed	FoM	Correct	Incorrect	Missed
0	0.4473	30.88	29.60	9.12
50	0.5950	38.38	24.95	1.62
100	0.5890	38.46	26.23	1.54
150	0.6414	38.50	22.53	1.50
200	0.6630	38.50	21.38	1.50
250	0.6463	38.51	23.41	1.49
300	0.6180	38.53	26.11	1.47
350	0.5828	38.51	29.70	1.49
400	0.5673	38.50	31.46	1.50

Table 4.4: Performance metrics of OMP for a clutter-less map with rectangular windowed data.

FoM	Correct	Incorrect	Missed
0.7808	38.57	11.68	1.43

4.2 Test Cases

An overview of the different test cases can be seen in Table 4.5.

Table 4.5: Overview of the conducted tests and evaluation methods.

Test	Real or simulated	Varied parameter	Investigated	OMP,CA-CFAR or CA-OMP
Test 1	Simulated	Location of steering vector	Removed energy	OMP
Test 2a	Simulated	SNR	Probability of detection	OMP and CA-CFAR
Test 2b	Simulated	SNR	RMSD, CRLB	OMP
Test 3	Simulated	Separation between two targets	RMSD, CRLB, removed energy	OMP
Test 4	Simulated	Separation between two targets	Correctly detected targets	OMP and CA-CFAR
Test 5a	Simulated	Randomized target and clutter location	FoM, correct, incorrect and missed targets	All
Test 5b	Real	Randomized target and clutter location	CA-CFAR and OMP differences	All

4.2.1 Test Case 1

Figures 4.3 and 4.4 show the removed energy as a function of the range or velocity offset for varying bandwidth-sampling frequency ratios, denoted as oversampling rate. The target had an SNR of 50 dB, and the results were averaged over 250 Monte Carlo iterations.

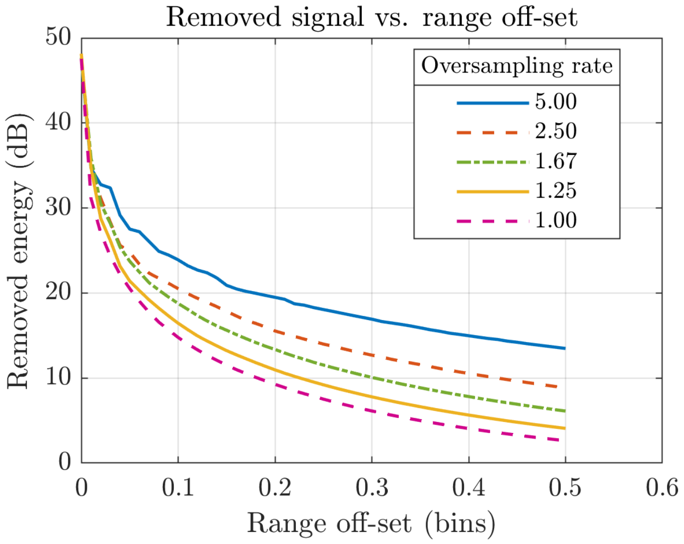

Figure 4.3: The removed energy using a steering vector offset in range from the true position. The simulated target had an SNR of 50 dB.

In Figure 4.3 the removed energy decreases rapidly with offsets in range. At an offset of 0.5 bins, the curves, listed in descending order, are at 13.5 dB, 8.9 dB, 6.1 dB, 4.1 dB, and 2.6 dB. Thus, the energy removal decreases rapidly for all oversampling rates, although the decrease is slower for larger oversampling rates. For very small offsets, the energy removal is close to the target SNR of 50 dB, which corresponds to the maximum removable energy.

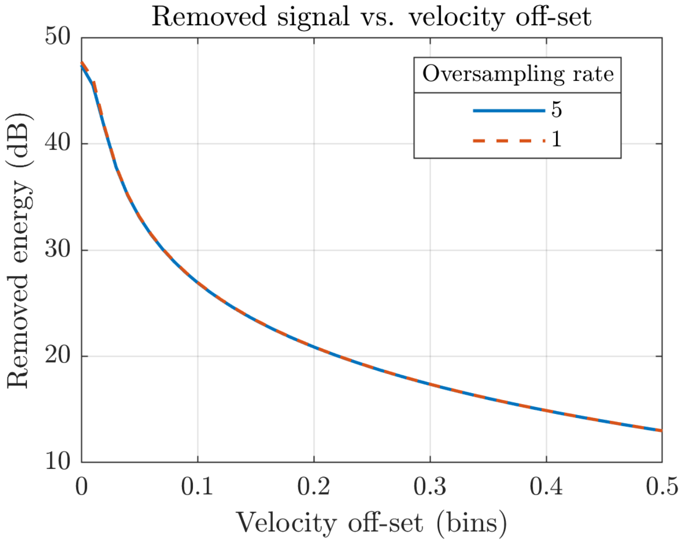

Figure 4.4: The removed energy using a steering vector offset in velocity from the true position. The simulated target had an SNR of 50 dB.

In contrast, in the velocity dimension the oversampling ratio has very little effect on the removed energy, as shown in Figure 4.4. This indicates that oversampling improves robustness to range mismatch but not to velocity mismatch.

4.2.2 Test Case 2a

Figure 4.5 shows the probability of detection for different SNRs between 12.6 dB and 20.0 dB for the OMP and the CA-CFAR detectors. The detection probability was calculated as the average over 500 Monte Carlo trials of a single target in noise with a random position and phase. The curves follow each other closely, and both have 0 % detection probability at around 12.6 dB and 100 % at approximately 18.5 dB, even though the OMP detector has a slightly higher detection probability in the transition region.

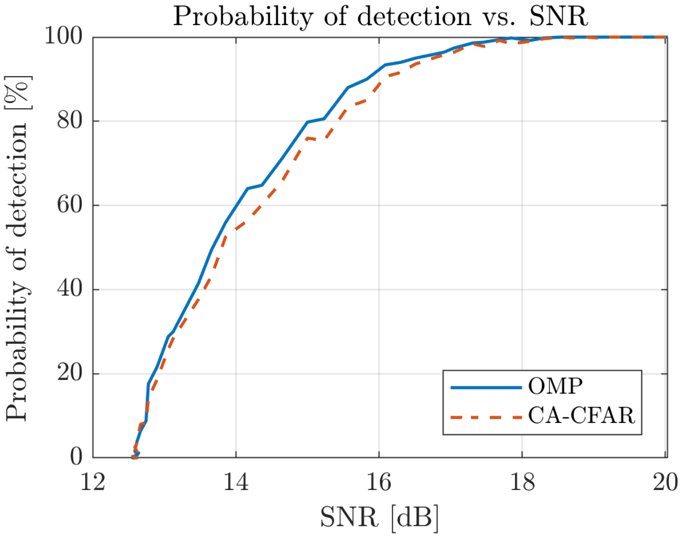

Figure 4.5: The OMP and CA-CFAR detectors' probability of detection for a single target as a function of SNR. The curves follow each other closely, but with a slightly higher detection probability for the OMP detector in the transition region.

4.2.3 Test Case 2b

The SNR span in Test 2b stretched from 22.5 dB to 62.4 dB. Since the SNRs are above the threshold for 100 % detection probability the RMSD can be obtained for each detection in both the range and velocity directions, see Figure 4.6. The RMSDs were averaged over 500 iterations, and compared to the CRLB. For large SNRs, the RMSD approached the CRLB in both range and velocity.

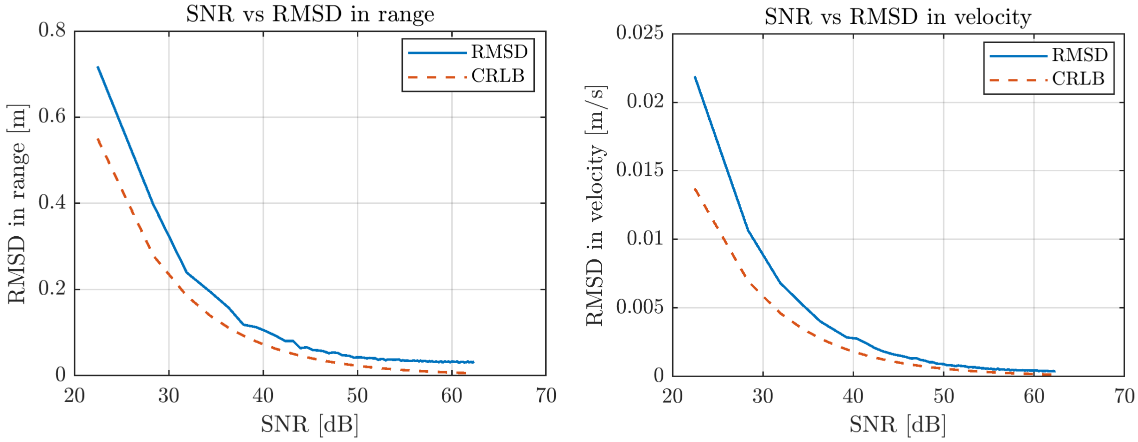

(a) RMSD in the range dimension and CRLB for different SNRs.

(b) RMSD in the velocity dimension and CRLB for different SNRs.

Figure 4.6: RMSD and CRLB for different SNRs. $P_{FA} \approx 1.55 \cdot 10^{-5}$

4.2.4 Test Case 3

In this section, the average RMSD and removed energy over 250 Monte Carlo simulations for two separated targets are shown in three different cases.

4.2.4.1 Equally Strong Targets with Rectangular Window

Figures 4.7 and 4.8 were generated when both targets had an SNR of 50 dB and a rectangular window was used in both the fast-time pulse compression and the slow-time FFT. The CRLB shown was calculated for a single target with an SNR of 50 dB. When the targets were separated in range, see Figure 4.7, the RMSD was near the CRLB and the removed energy was close to the maximum 53 dB for separations larger than the pulse length of 300 m. For separations smaller than this, the RMSD varied rapidly around 0.2 m and the removed energy around 25 dB until the targets were indistinguishable. When the targets were separated in velocity, see Figure 4.8, some sidelobe interference was visible for all separations in both the RMSD and removed energy.

It is clear that interference between targets greatly decreases the effectiveness of the OMP algorithm; even for separations up to 300 m the removed energy is diminished.

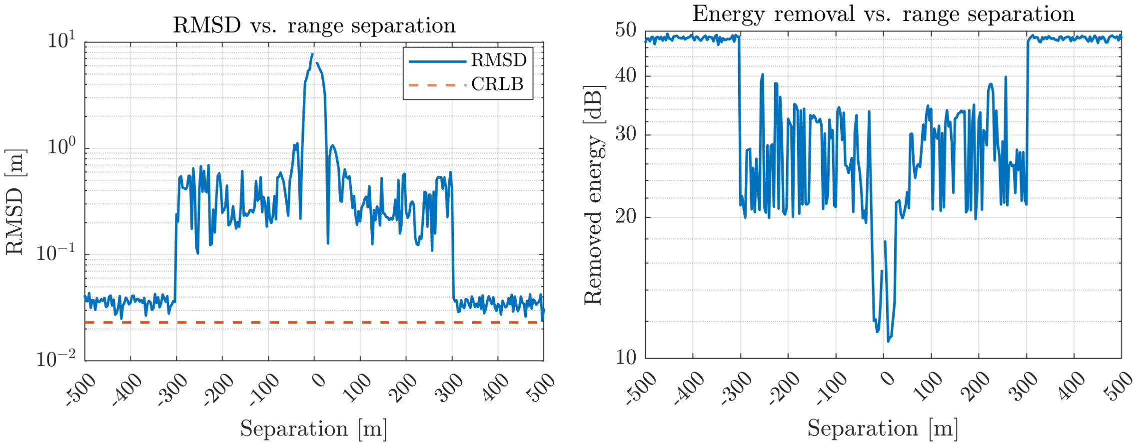

(a) RMSD in the range dimension as a function of range separations.

(b) Removed energy as a function of range separations.

Figure 4.7: RMSD and removed energy for two equally strong targets of 50 dB separated in range using rectangular windowed pulse compression. The RMSD was near the CRLB for separations larger than the pulse length of approximately 300 m, which also resulted in a removed energy close to the maximum 53 dB. Both RMSD and removed energy varied rapidly in the interference region between the two targets.

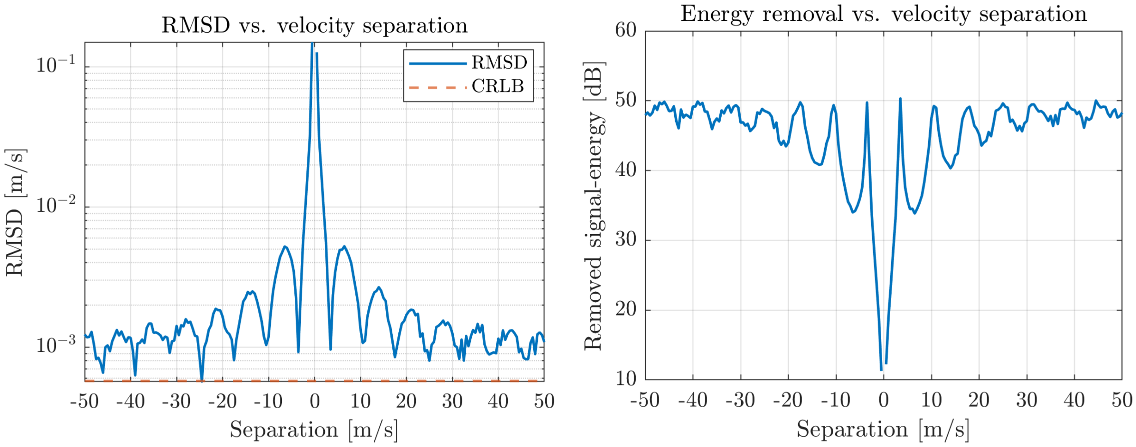

(a) RMSD in the velocity dimension as a function of velocity separations.

(b) Removed energy as a function of velocity separations.

Figure 4.8: RMSD and removed energy for two equally strong targets of 50 dB separated in velocity using rectangular windowed FFT. The RMSD was near the CRLB for large separations resulting in a removed energy close to the maximum 53 dB, however some interference between the targets was visible for all separations.

4.2.4.2 Equally Strong Targets with Hann Window

Figures 4.9 and 4.10 were also generated with both targets having an SNR of 50 dB, but a Hann window was applied in the processing of both the fast-time and the slow-time dimensions. The CRLB shown was calculated for a single target without windowing applied to facilitate comparison between the processing methods. In the range dimension, see Figure 4.9, both RMSD and removed energy exhibited a sidelobe structure similar to that after Hann windowed matched filtering, see Figure 2.4. When the separation was larger than 300 m, such that the targets were completely separated, the RMSD was slightly above the unwindowed CRLB. This was also seen in the removed energy, where the removed energy was approximately 47 dB instead of the optimal 53 dB. In the velocity dimension, see Figure 4.10, the interference between the targets was very small for separations larger than 3 m/s, resulting in an RMSD slightly above the CRLB and a removed energy of around 49 dB.

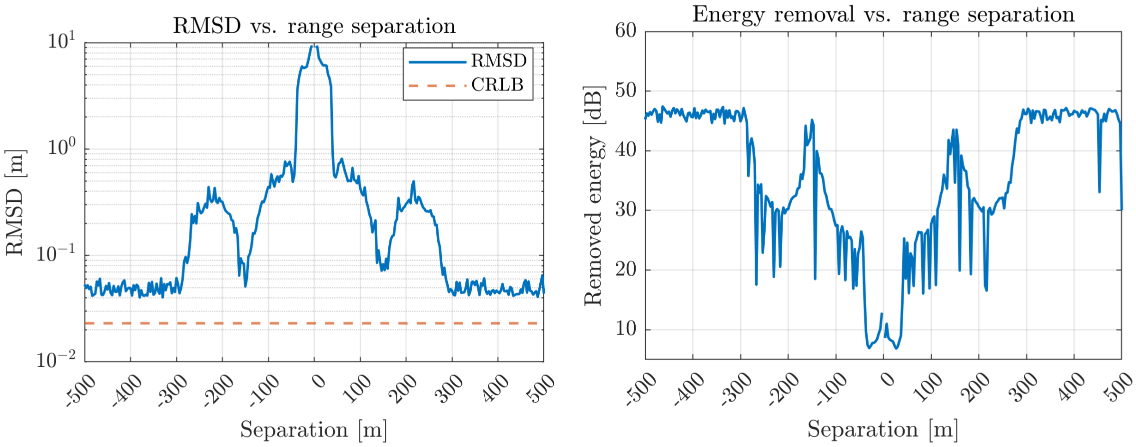

(a) RMSD in the range dimension as a function of range separation for Hann windowed pulse compression.

(b) Removed energy as a function of range separation for Hann windowed pulse compression.

Figure 4.9: RMSD and removed energy for two equally strong targets of 50 dB separated in range using Hann windowed pulse compression. For separations larger than the pulse length of 300 m the RMSD was slightly above the unwindowed CRLB, resulting in a removed energy of slightly less than the maximum 53 dB.

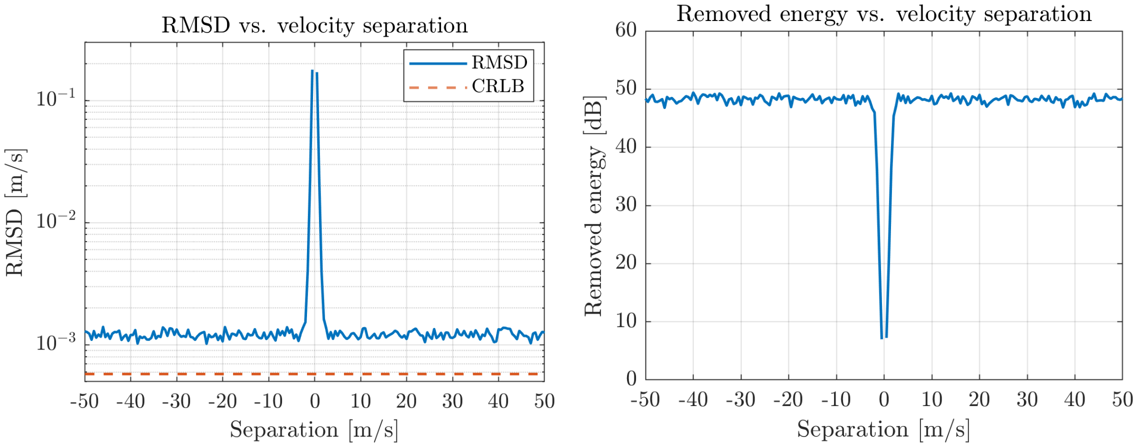

(a) RMSD in the velocity dimension as a function of velocity separation for Hann windowed data.

(b) Removed energy as a function of velocity separation for Hann windowed data.

Figure 4.10: RMSD and removed energy for two equally strong targets of 50 dB separated in velocity, using a Hann windowed FFT. The RMSD was slightly above the unwindowed CRLB with negligible interference for separations larger than approximately 3 m/s. The removed energy was stable around 49 dB for separations larger than 3 m/s.

When comparing the performance with or without Hann windowing, it is clear that applying a Hann window before pulse compression or Fourier transforming suppresses the interference between the closely spaced targets compared to the use of a rectangular window, generally resulting in more removed energy in the interference region. This effect was especially pronounced in the velocity dimension, where target interference became insignificant after approximately 3 m/s. However, the increased performance in the interference region also resulted in a slightly worse performance when the targets were completely separated.

4.2.4.3 Strong Target and Weak Target with Rectangular Window

In Figures 4.11 and 4.12 the RMSD and removed energy are shown for the case when one target had an SNR of 50 dB and the other 30 dB. The CRLB shown was the combined CRLB for both targets when completely separated. In the range dimension, see Figure 4.11, the RMSD approached the CRLB after approximately 60 m. The removed energy was close to 45 dB for several separations, which can be compared to the maximum 50 dB. However, some interference was detectable, especially slightly below 300 m. In the velocity dimension, see Figure 4.12, the interference is small for separations larger than 3 m/s with an RMSD close to the CRLB and a removed energy of 46 dB.

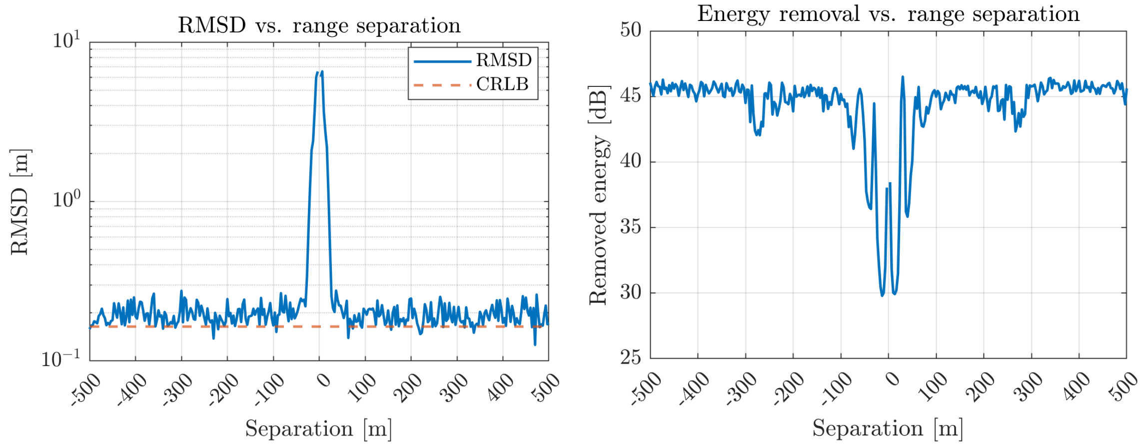

(a) RMSD in the range dimension as a function of range separation when one target is stronger than the other.

(b) Removed energy as a function of range separation when one target is stronger than the other.

Figure 4.11: RMSD and removed energy when one target had an SNR of 50 dB and the other of 30 dB separated in range using rectangular windowed pulse compression. The RMSD was close to the CRLB for separations larger than approximately 60 m. The removed energy was a few decibels less than the maximum 50 dB.

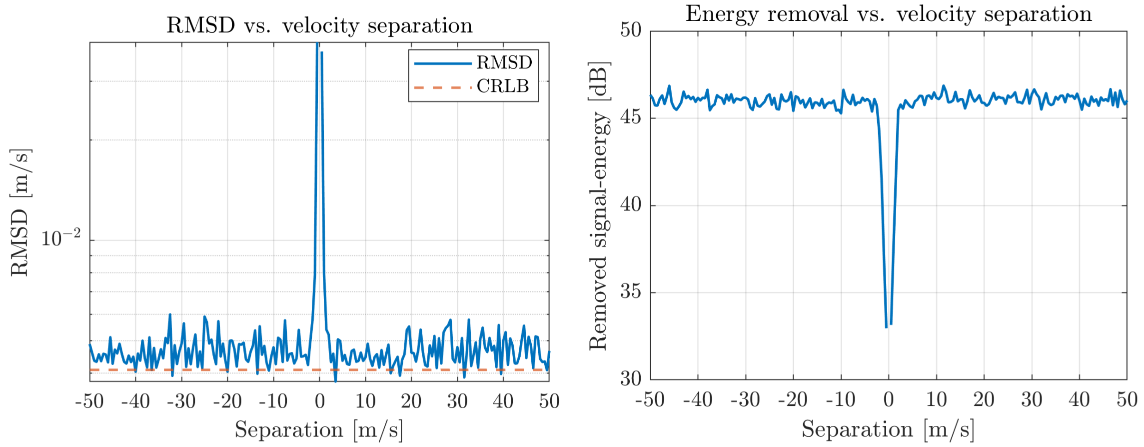

(a) RMSD in the velocity dimension as a function of velocity separation when one target is stronger than the other.

(b) Removed energy as a function of velocity separation when one target is stronger than the other.

Figure 4.12: RMSD and removed energy when one target had an SNR of 50 dB and the other of 30 dB separated in velocity using a rectangular windowed FFT. The RMSD was close to the CRLB for separations larger than approximately 3 m/s and the removed energy was stable at around 46 dB.

4.2.5 Test Case 4

The number of targets detected for two targets with different separations in either range or velocity, averaged over 300 Monte Carlo iterations, is seen in Figure 4.13 and 4.14. Results for both the OMP detector and the CA-CFAR detector were computed. In addition, the number of missed targets were noted for every separation. This only showed that for every instance where two targets or more were detected no targets where missed. This was therefore not included in the plots.

In Figure 4.13, both targets had an SNR of 50 dB. For the CA-CFAR detector, the targets masked each other when in each other's training region, resulting in almost zero detections for separations between ∼25 m to ∼170 m. For reference, the training region was 165 m. For separations between ∼25 m and ∼280 m, the OMP detector detected both targets, but false identification of residual peaks led to an average number of detections slightly above two. In the velocity direction, the same masking behavior was seen even though the CA-CFAR detector on average detected at least one of the targets between separations of 0 m/s and 6.9 m/s. The training region in the velocity dimension was ∼6.75 m/s. Furthermore, the OMP detector detected exactly two targets for separations larger than ∼1.1 m/s.

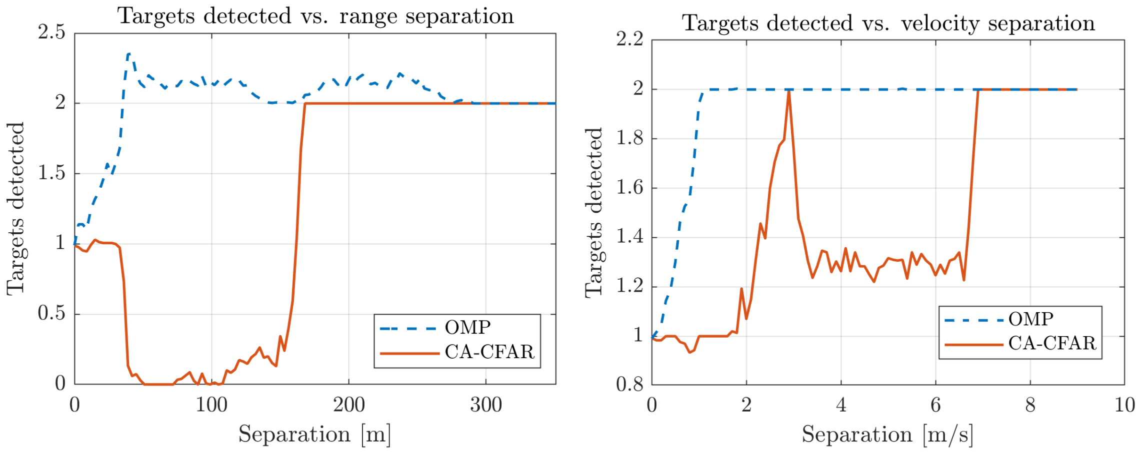

(a) Mean number of detections for two targets separated in range.

(b) Mean number of detections for two targets separated in velocity.

Figure 4.13: Mean number of detections for two targets using the OMP detector (blue) and the CA-CFAR detector (orange) as a function of the separation in range and velocity. Both targets had an SNR of 50 dB. For the CA-CFAR detector, the targets masked each other in range when in each other's training regions resulting in almost zero detections for separations between ∼25 m to 170 m. For the same separations, the OMP detector detected both targets, but left some residuals resulting in slightly more detections than two. In the velocity direction, the same general behavior was seen even though the CA-CFAR detector on average detected at least one of the targets.

In Figure 4.14, one target had an SNR of 50 dB and the other of 30 dB. In range, the CA-CFAR detector missed one target for separations up to ∼200 m and detected both targets reliably after 330 m. The OMP detector was able to detect both targets after ~25 m. In velocity, the CA-CFAR detector detected both targets reliably for separations larger than 8 m/s and missed one target for separations smaller than this, except for some unreliable detections of two targets around 2 m/s. The OMP detector detected both targets reliably when separations were larger than 1 m/s.

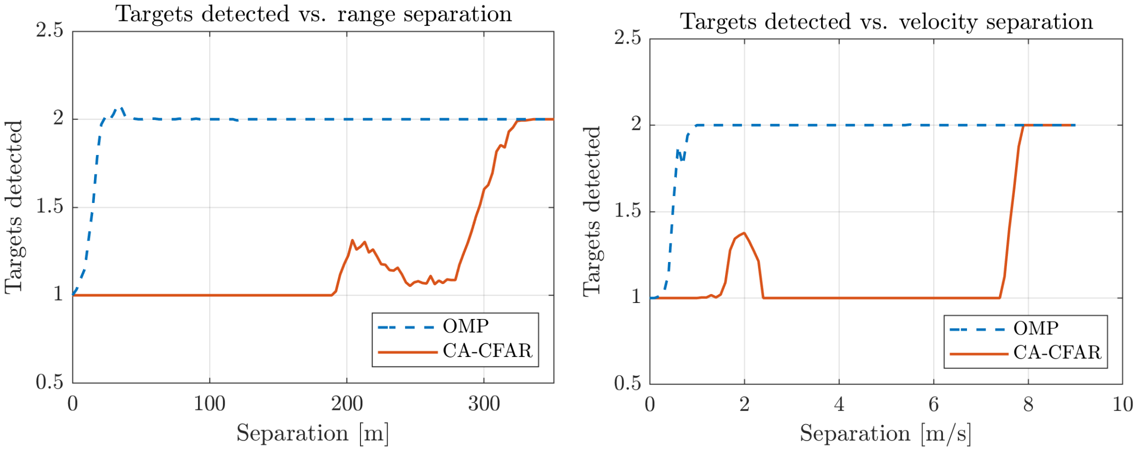

(a) Mean number of detections for two targets separated in range.

(b) Mean number of detections for two targets separated in velocity.

Figure 4.14: Mean number of detections for two targets using the OMP detector (blue) and the CA-CFAR detector (orange) as a function of separation in range and velocity. One target had an SNR of 50 dB and the other of 30 dB.

4.2.6 Test Case 5a

The performance metrics were shown in Table 4.6 were averaged over 300 Monte Carlo iterations. Each iteration contained 40 target signals. Based on the tuning results of Section 4.1, the data was Hann windowed and the algorithm allowed to find 250 clutters before detecting targets for both the OMP and CA-OMP detectors.

Table 4.6: Performance metrics of the OMP, CA-CFAR and CA-OMP detectors for Hann windowed, clutter suppressed data.

Detector	FoM	Correct	Incorrect	Missed
OMP	0.8929	38.90	4.56	1.10
CA-CFAR	0.6696	30.65	5.92	9.35
CA-OMP	0.8044	37.63	6.97	2.37

4.2.7 Test Case 5b

In the following section, the detection performance on the real radar data of the CA-CFAR detector, the OMP detector for two different choices of the parameter k, and the adapted CA-OMP detector is compared. Since no ground truth is available for the real radar data, the classification of correct and false detections is based on a qualitative assessment of the range-Doppler maps.

In Figure 4.15 the removal of a signal from a presumed car is shown. The left column shows the cross-section of the range-Doppler transformed data (solid blue) and the range-Doppler transformed fitted steering vector (dashed orange) in the Doppler direction (top) and range dimension (bottom). Visually, the steering vector fitted the data well, even though some mismatch was noticeable. In the right column, the range-Doppler transformed data before subtracting the steering vector (solid blue), and after subtracting the steering vector (dashed orange) is shown. The mainlobe was reduced by 18.3 dB to the level of the range side lobes, which were largely unaffected in strength.

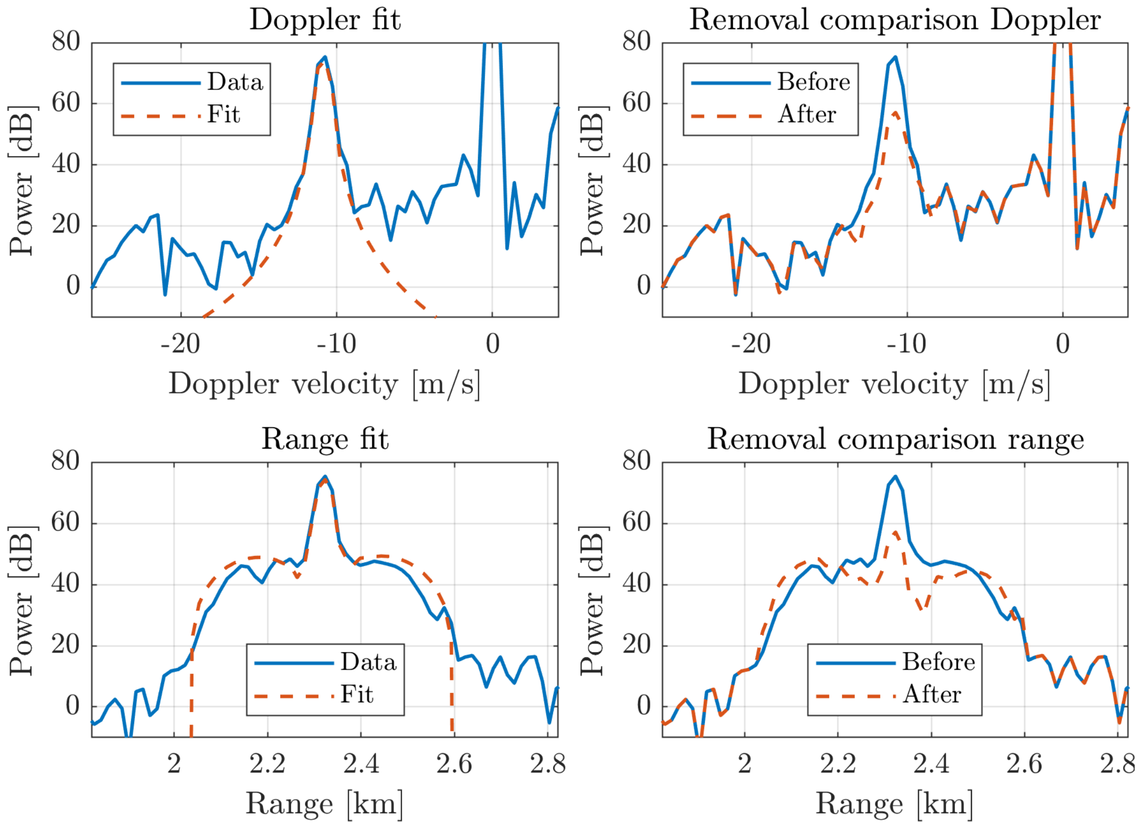

Figure 4.15: Signal subtraction of signal originating from a presumed car. Cross-section of the range-Doppler transformed data (solid blue) and the range-Doppler transformed fitted steering vector (dashed orange) in the Doppler direction (top) and range dimension (bottom) shown in the left column. In the right column, the range-Doppler transformed data before subtracting the steering vector (solid blue), and after subtracting the steering vector (dashed orange) is shown.

This can be compared to the removal of a signal originating from a bird, see Figure 4.16. In this case, the side lobes were at the level of the noise floor, and the signal was almost completely removed.

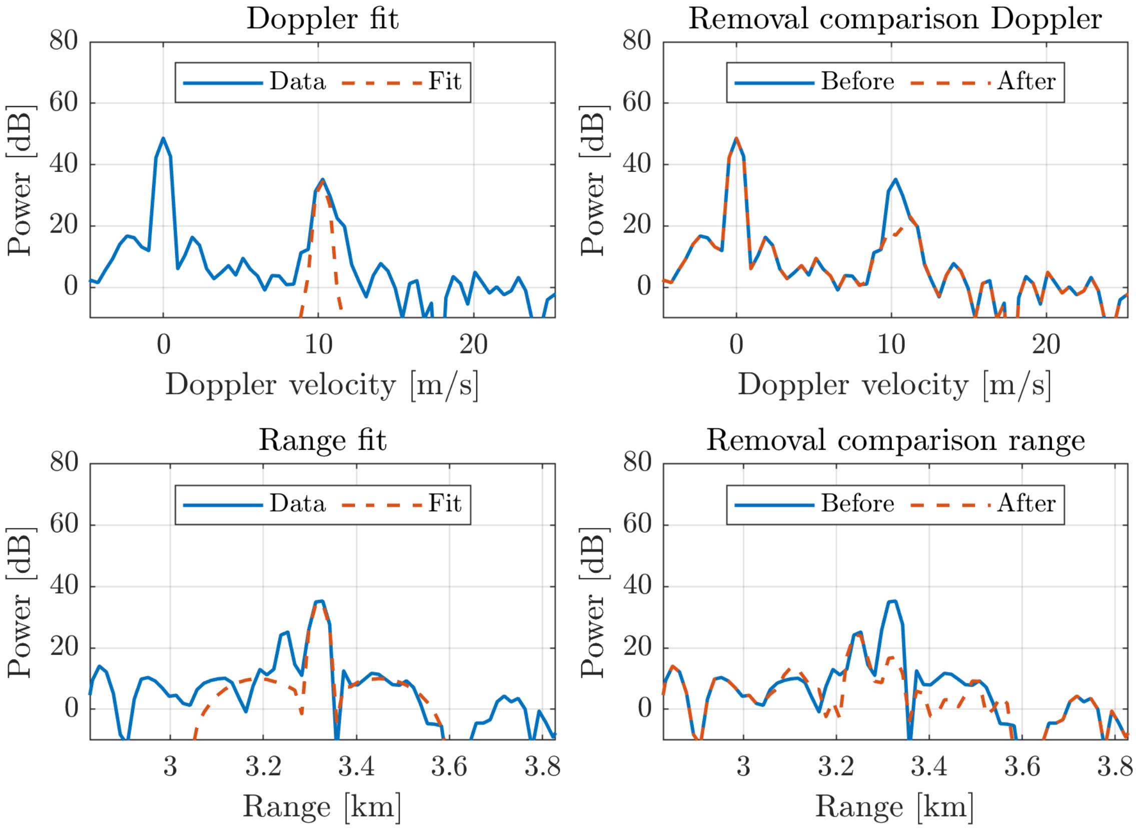

Figure 4.16: Signal subtraction of signal originating from a presumed bird. Cross-section of the range-Doppler transformed data (solid blue) and the range-Doppler transformed fitted steering vector (dashed orange) in the Doppler direction (top) and range dimension (bottom) shown in the left column. In the right column, the range-Doppler transformed data before subtracting the steering vector (solid blue), and after subtracting the steering vector (dashed orange) is shown.

Figure 4.17 shows the detections registered by the four different detectors in the 60th CPI zoomed in on the part of the data where vehicle detections are expected.

For the CA-CFAR detector, all detections are believed to correspond to true targets. However, the strong target approximately 2.7 km away traveling with a speed of approximately 4 m/s was missed by the detector. This miss was likely caused by masking from the stronger targets at around the same range with slightly higher velocities. The same is true for the missed target approximately 2.85 km away with a speed of -6 m/s.

The CA-OMP detector detected all targets identified by the CA-CFAR detector, plus a few more. The CA-OMP detector managed to detect both clear misses made by the CA-CFAR detector. In addition, the CA-OMP detector also detected a target around 2.6 km away with a speed of -10 m/s. It is unclear whether this corresponds to a true target masked by the other nearby targets, or the result of a detection on an artifact left by the projection procedure.

The OMP detector with $k = 100$ produced significantly more detections than the previous two methods. Most of these detections are believed to be the result of imperfect signal removal, effectively resulting in false alarms. This is especially clear for the almost straight lines of detections in the range dimension. At the same time, it can be noted that the "background" in the relevant piece of data is significantly higher than the noise floor, around 30 dB compared to 0 dB. While this could be the result of side lobes from the strong targets alone, it is also possible that there are other weak targets, such as birds, present in the region that ideally also should be detected.

Finally, the OMP detector with $k = 1500$ generated fewer detections than the OMP detector with $k = 100$. The detections were also more spread out, resulting in detections of slightly weaker targets, such as the target located in the upper right region missed by the $k = 100$ counterpart. However, with $k = 1500$ some strong targets were missed, such as the target 2.85 km away with a speed of -6 m/s. In contrast to the $k = 100$ detector, the $k = 1500$ detector produced few detections in the region with an elevated background level.

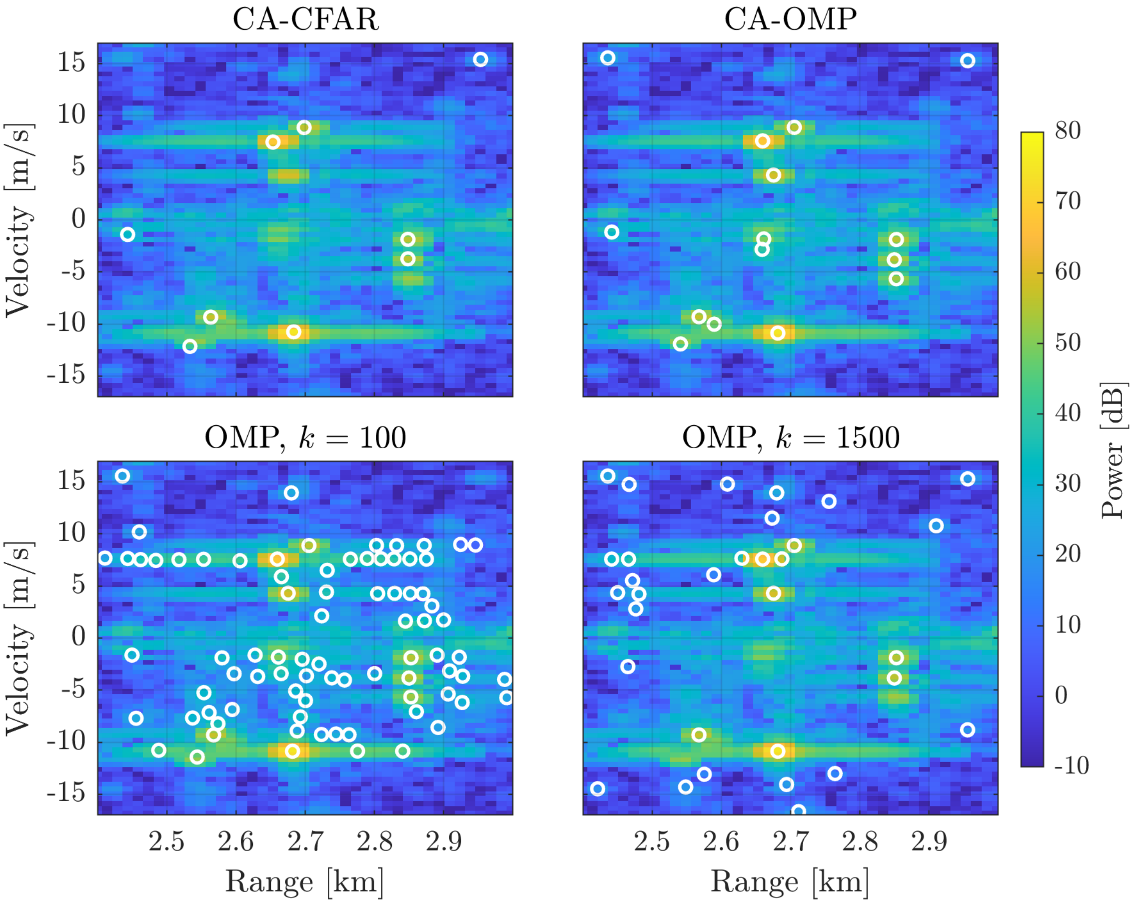

Figure 4.17: Detections by the CA-CFAR detector, the CA-OMP detector and the OMP detector with $k = 100$ and $k = 1500$ on the part of the data mainly containing cars.

Figure 4.18 shows the detections registered by the four different detectors in the

60th CPI zoomed in on the part of the data where detection of birds are considered probable.

The detections produced by the CA-CFAR detector and the CA-OMP detector resemble the detections on the part of the data containing vehicles in the sense that they identified the same targets, but CA-OMP had a few extra detections. All detections are believed to be correct, but both algorithms likely missed a few targets, probably due to masking by the ground clutter.

The OMP detector with $k = 100$ generated detections for some strong targets missed by the CA-CFAR and CA-OMP detectors, but simultaneously missed some weaker but more isolated targets detected by the previous two detectors. In contrast, its $k = 1500$ counterpart detected both the strong and the weak targets successfully. However, it is unclear whether the very closely spaced detections generated by the $k = 1500$ detector correspond to very close actual targets, or are a consequence of poor signal removal.

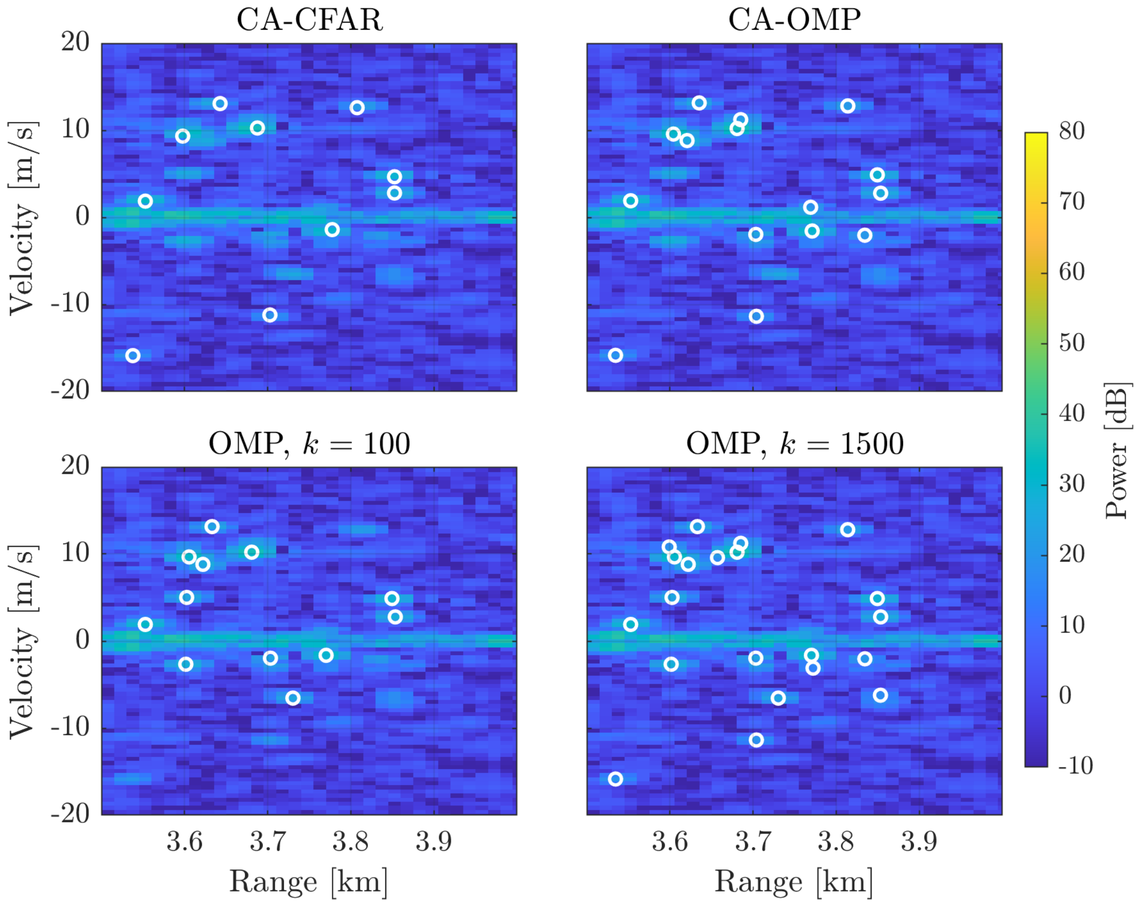

Figure 4.18: Detections by the CA-CFAR detector, the CA-OMP detector, and the OMP detector with $k = 100$ and $k = 1500$ on the part of the data containing echoes from birds.

5

Discussion

The objective of this thesis was to investigate the potential of using OMP for target detection in pulse-Doppler radar applications. This detector's performance and parameter estimation accuracy was compared to those of a more conventional CA-CFAR detector. Particular emphasis was placed on scenarios involving closely spaced targets and interference from clutter.

5.1 Test Case Performance for Simulated Data

The results from the test cases created to evaluate the performance are presented in Section 4.2. How well the OMP detector functioned under varying factors, such as SNR and separation between targets, could only be directly assessed for simulated data. The following sections interpret the current detection performance, both in terms of parameter estimation as well as energy removal, for the OMP detector.

5.1.1 Isolated Targets

Test Case 1 showed the relationship between parameter estimation error and removed energy. Assuming a target with an SNR of 50 dB, and that the ratio between bandwidth and sampling frequency was 1.25, an offset less than ~ 0.007 bins in range and ~ 0.029 bins in velocity is needed for the residuals to fall below the OMP detection threshold observed in Test Case 2a of 12 dB. With the radar parameters used, one range bin is ~14.99 m and one velocity bin is ~0.488 m/s, meaning that the required offsets correspond to ~10 cm and ~0.014 m/s. These values are small compared to the usual radar length scales, and can be compared to the CRLBs which provide the theoretically lowest possible standard deviations, in this case 2.3 cm and 5.8×10^{-4} m/s. If the amount of removed energy is less than what is required for residuals below the detection threshold, the algorithm might encounter problems with detection of the residuals instead of new targets.

Notably, as seen in Test Case 2b, the RMSD for a single target with 50 dB SNR is ~5 cm in range and ~0.0009 m/s in velocity. These offsets are less than the offsets required for sub-detection threshold residuals. However, the RMSD plateaued at around 50 dB without reaching the CRLB in both range and velocity. This indicates that the primary source of errors in the parameter estimation transitioned from noise-related uncertainty to an algorithmic limitation. This limitation is likely directly restricted to the number of fine samples evaluated in the sub-bin parameter estimation step N. It also indicates that removal of stronger signals, which require even more suppression for sub-detection threshold residuals, may become difficult as approximately the same RMSD can be expected for 80 dB SNR as for 50 dB.

Furthermore, in Test Case 2a, a comparison against a CA-CFAR detector was made. The results from this showed that an OMP and a CA-CFAR detector achieved very similar probabilities of detection for an isolated target in AWGN when both detectors were calibrated for the same probability of false alarm. The OMP detector, however, potentially displayed a slightly higher detection probability in the transition region between failed and 100 % successful detection. This behavior can likely be explained in cases when the target response spreads across several neighboring bins due to side lobe leakage or off-grid positioning. The OMP algorithm essentially interpolates the signal energy in the sub-bin parameter estimation, whereas the CA-CFAR detector operates on each bin individually.

The results from these simplified, single-target simulations indicate that the residual energy remaining after projection is unlikely to be detectable and pose a problem for isolated targets in AWGN. However, it is not these simple scenarios that cause detection difficulties for conventional detectors, but the cases when there are multiple targets in close proximity to each other.

5.1.2 Multiple Targets

Beyond the simulations involving a single target, additional experiments were conducted using two targets in order to investigate how interference between targets affects the detection performance. The results from Test Case 3 and 4 demonstrate that the effectiveness of detection depends on the spacing between the targets, the applied windowing function and the relative targets' strengths.

5.1.2.1 Target Interference with a Rectangular Window

In the case where a rectangular window was applied and the targets were of equal strength, strong sidelobes in both range and velocity caused substantial interference between targets. This interference caused errors in the parameter estimation of the first detected target, which consequently degraded the estimation and removal of the second target. This resulted in substantial signal energy left after subtraction. Figure 4.7b from Test Case 3 showed that the removed energy varied between approximately 40 dB and 11 dB when the targets were separated less than ~300 m. This is equivalent to between ~13 dB and ~42 dB of the original 53 dB signal energy remained after both detections. Notably, 300 m distance corresponds approximately to the pulse length in range. As discussed in Section 5.1.1, residual signal energy levels below ~12 dB are required to prevent residual detections entirely. Thus interfering signals pose a high risk of false alarms when two targets are closely spaced and a rectangular window is applied.

The RMSD mirrored the trend observed for the energy removal. In range, the RMSD remained elevated for separations below ∼300 m, varying between ∼1 m and ∼0.1 m. Once separations exceeded this region, the RMSD approached the value observed for the single target case. For target separations below ∼40 m, the RMSD increased to ∼8 m. Although these range errors are smaller than the range resolution, the surplus of false detections near the target crowds the detection map, which makes it unclear which detection is correct. Furthermore, numerical instabilities may arise in the projection step if the steering vectors become close to parallel. A similar effect was observed in the velocity dimension. Here, the residual energy gradually decreased with increasing separations, with sidelobe-like oscillations. When the velocity separation exceeded ∼10 m/s, the remaining signal energy fell below ∼12 dB.

When the data was rectangular windowed and the targets had unequal strength, namely one target had an SNR of 30 dB and the other of 50 dB, the total energy was ∼50 dB. In this case, achieving residual levels below 12 dB required only ∼60 m separation in range and ∼2 m/s separation in velocity. This is substantially less than in the equal strength case, but expected. Since the weaker target cannot perturb the strong target significantly, the first parameter estimation will be more accurate. Therefore the energy subtraction step will be more effective, allowing more precise detection of the second target too.

5.1.2.2 Target Interference with a Hann Window

The application of a Hann window significantly reduced interference effects overall, particularly in the velocity dimension. By suppressing sidelobes, disturbance from neighboring targets was reduced, improving the separability of closely spaced signals. However, for residual levels below ∼12 dB the target spacing in range still needed to exceed ∼300 m, whereas only ∼2 m/s separation was needed in velocity. Even though the spacing required for residuals below the detection threshold was the same regardless of windowing function, the average removed energy was slightly larger for the Hann window than for the rectangular window. The improved performance in the interference region came at the cost of reduced peak power and degraded resolution, as reflected by the RMSD deviating further from the CRLB, in Figures 4.9a and 4.10a, and somewhat less removed energy outside of the interference region, in Figures 4.9b and 4.10b. These observations illustrate the fundamental trade-off from radar signal processing that is associated with tapered windows.

5.1.2.3 Detector Comparison

In cases where residual energy is falsely identified as targets, clustering becomes beneficial. Instead of keeping many closely spaced detections, the clustering will only retain the detection with the strongest signal strength. Since the RMSD remains below the range and velocity resolution, as long as both targets are not in the same resolution cell, the strongest detection is generally sufficient for accurate localization.

Even though clustering can help in removing unwanted detections, problems with interference still remain. Certain target separations cause the targets to interfere in a way that prevents successful target detection of either one or both targets. Test Case 4 compares the performance of the OMP and CA-CFAR algorithms by evaluating how many, of two, targets were correctly identified as a function of range and velocity separation. Clustering was allowed for both algorithms. Unlike Test Case 3, where the OMP algorithm was forced to make two detections, OMP only identified a target if the correlation energy exceeded the threshold set to achieve a certain false alarm rate in Test Case 4. Consequently, the OMP detector cannot consistently detect both targets for very close separations, and especially separations less than the range and velocity resolutions since they are clustered into one.

For equal-strength targets separated in range, the OMP detector usually identified more than two targets for separations between $\sim$40 m and $\sim$300 m. This behavior was expected, as this region exhibited elevated RMSDs in figure 4.7a and suboptimal removed energy in figure 4.7b. In velocity, OMP started to detect both targets from $\sim$1.1 m/s. In contrast, the CA-CFAR detector only began to consistently detect both targets at separations larger than $\sim$160 m and $\sim$6.8 m/s. Since the combined length of training cells and guard cells amounts to $\sim$165 m and $\sim$6.84 m/s in each direction, the reduced detection performance is likely caused by the targets occupying each other's training cells. As previously discussed, this interferes with the noise estimate and hinders detection capability. Decreasing the number of guard and training cells decreases the length of the problematic regions but may decrease overall detection performance.

In the case that the targets had different signal strengths, the OMP detector's tendency to identify more than two targets is less pronounced compared to the equal strength scenario. This indicates that the stronger target's parameters could be estimated more precisely and the signal energy removed with higher accuracy, revealing the weaker target and improving detection performance. For the CA-CFAR detector, however, the masking of the weak target caused by the strong target significantly reduced its ability to resolve two distinct targets compared to the equal strength case. Consistent two-target detection was only achieved at separations of $\sim$320 m and $\sim$7.9 m/s. This indicates that it is not primarily that the signals occupy each other's training cells that cause missed detections, but rather that the weaker target became obscured by the sidelobes of the stronger target.

5.2 Performance on Complex Environments

The following section discusses the performance of the investigated detectors in complex environments and analyzes the factors that limit their detection capability. Furthermore, how the different design choices influence detection in both simulated and real data is considered and the potential for improvement is discussed.

5.2.1 Data From Simulated Environments

As seen in Table 4.6, the OMP detector had the best detection performance overall, evaluated in FoM, with the highest number of correct detections and the least number of incorrect and missed targets. On the other hand, the CA-CFAR detector had the lowest FoM, primarily due to the most missed detections. These results are in line with the expected results since the OMP detector achieves very high suppression of isolated targets in simulated data. Consequently, the OMP detector was also anticipated to successfully resolve close targets with varying strength, resulting in a low number of missed detections. It is believed that the incorrect detections are primarily the result of detection of residuals from very closely spaced targets. For the CA-CFAR detector instead, the number of missed detections is believed to be the result of the targets masking each other and the incorrect detections are the result of false alarms on strong noise spikes.

Unsurprisingly, the CA-OMP detector had a performance between the two other detectors, with a significantly lower number of missed detections than the CA-CFAR detector due to the iterative removal of signals. However, CA-OMP also had slightly more incorrect detections than CA-CFAR since the CA-OMP detector risks detecting on residual signals. It is also expected that the number of correct detections by the CA-OMP detector is slightly less than the OMP detector since the CA-OMP detector detects on the coarse grid where leakage of energy can decrease the probability of detection slightly, as seen in Figure 4.5.

5.2.2 Data From Real Environments

A major weakness of the OMP algorithm is its sensitivity to parameter estimation errors. Accurate estimates of both the true range and velocity are essential not only for the effectiveness of the energy removal, as seen in Figures 4.3 and 4.4, but also for accurate matching between steering vectors and data. In practice, strong responses in the coarse grid do not necessarily correspond to the strongest peaks in a more finely sampled grid since off-grid targets distribute their energy across multiple coarse-grid bins through spectral leakage. In theory, an interpolated grid can be created through zero-padding the Fourier transforms used in the range-Doppler processing. Unfortunately, this becomes computationally infeasible for the accuracy required for adequate energy removal. As a consequence, a method where only a few points were sampled more finely was implemented, see Section 3.1.2. The parameter affecting which points were resolved more finely, k, led to different behaviors of the OMP detector.

For large k, the size of the clusters before discarding all but the strongest point in each cluster becomes large. This risks causing signal energy from different targets, sidelobes, and residuals to become connected within a single cluster. If the retained point for the whole cluster has a relatively low energy in the fine grid, large regions of data can become unavailable for sampling more finely. Conversely, for smaller values of k, the clusters become smaller, enabling the detector to finely sample points that are more closely spaced.

If the signal removal was perfect, the two parameter choices are expected to perform similarly since high values in the coarse grid would correspond with a target and not a residual with high power. This is supported by the fact that the choice of k does not affect the detector performance on simulated data. In practice, however, imperfect signal cancellation often causes residuals with high correlation strengths in the rough grid. As a result, when k is small the algorithm tends to become trapped in residuals and is unable to explore points that are weak in the coarse grid but relatively strong in the fine grid. In contrast, when k is large the algorithm risks missing targets that are connected to residuals in the coarse grid, but becomes more exploratory of points with weak strength in the coarse grid.

The above behavior can also be seen when evaluating the OMP algorithm for $k = 100$ and $k = 1500$, see Figures 4.17 and 4.18. For the smaller k the algorithm generates several detections in regions around the strong signals from vehicles, most of which are assumed to originate from detection of residuals, but fails to detect some of the birds with a weaker signal strength in the coarse grid. When k is large, OMP misses some of the strong targets in the vehicle region but is able to detect most of the birds, even though some of the birds have weak signal strengths in the coarse grid.

In comparison to the OMP algorithm, the CA-CFAR detector provides more predictable detection behavior. It detects both strong and weak targets when sufficiently separated, but suffers from masking causing missed detections when the targets are densely spaced, as evident in Figures 4.17 and 4.18. It can be noted that it may be possible to improve the CA-CFAR detector's performance on the specific piece of data by altering the number of guard or training cells. The parameters used for the CA-CFAR detector were tuned to maximize the FoM in simulated data replicating the real radar data. Even though the performance may be improved specifically for the scenario at hand, the same general conclusions would still hold.

The main issue for the OMP algorithm is the incomplete removal of strong signals in real data, see Figure 4.15. The residuals left after removing a strong target tended to be on the level of the target's sidelobes at $\sim$45 dB, which is still strong compared to other relevant signals, such as the echoes from birds at $\sim$30 dB. Furthermore, the residuals are no longer similar to a return from a point target and several iterations are needed to decrease the signal enough such that other targets are stronger. This often results in the detector getting "trapped" in the residuals left by a strong target and may lead to several false alarms if the detections are spaced further away than the clustering threshold.

However, it can be observed that the residuals left are approximately flat, see Figure 4.15, and resemble an elevated background noise level. Based on this observation, a hybrid of the CA-CFAR and the OMP algorithm, the CA-OMP algorithm was proposed. Since the CA-CFAR detector adapts to local background levels, the predominantly flat residuals would be less likely to trigger multiple detections. This can also be seen in Figure 4.17 as a smaller number of false detections for the CA-OMP detector compared to the OMP detector with $k = 100$. Meanwhile, the decreased strength in very strong peaks can still enable weaker targets that would otherwise be masked to be detected, which was also shown in Figure 4.17.

5.3 Energy Removal and Steering Vectors

As previously noted, one of the most important aspects of the OMP procedure is the effective cancellation of targets. However, the effectiveness of the removal is dependent on three different mechanisms: correct parameter estimation, correct amplitude estimation, and appropriate steering vectors.

5.3.1 Parameter estimation

In simulated data, where generated signals and steering vectors are created similarly, the primary cause of poor energy removal is attributed to incorrect range and velocity estimates. This is supported by Figures 4.3 and 4.4, where even small deviations lead to significantly reduced energy removal. As seen in Figure 4.1, the sub-bin parameter estimation implemented approaches the CRLB when the target is isolated, which also results in almost complete energy removal, see Figure 4.2. However, when targets are closely spaced, interference decreases the accuracy of the estimates, see Chapter 4.2.4. This is a direct consequence of the greedy procedure where only one signal at a time is found and the parameter estimates are kept for all following iterations. In order to estimate the location of interfering signals more accurately, the combined response must be modeled as a combination of all involved signals. Since the number of signals is not known a priori, this is not possible directly. However, more complex methods that relax the greedy procedure by reevaluating previously found estimates after each new signal found have been proposed, such as NOMP [7].

5.3.2 Amplitude estimation

It is not only the parameter estimation step that is affected by interference when other signals are present, but also the amplitude estimation of the found steering vectors in the projection step. In theory, the continuous re-estimation of amplitudes in the projection step enables the algorithm to account for interfering amplitudes. However, this is no longer the case when some parts of the data, e.g. the clutter, is unavailable for the detector to resolve. Ideally, the entire received signal, including targets and clutter, would be sufficiently described by the sparse signal model 2.8. This would enable the detector to account for interference from both clutter and targets, while potential clutter-related detections could be filtered out after completion of the OMP algorithm.

In practice, however, it was not possible to sufficiently resolve the clutter, indicating that these returns are badly described as single-point scatterers. Despite this limitation, the small difference in FoM between the Hann-windowed case without clutterers and the Hann-windowed case with more than 150 clutterers removed, see Tables 4.1 and 4.2, indicates that the clutter removal mode still provided adequate suppression of the interfering clutter amplitudes.

5.3.3 Steering Vector Suitability

In the real radar data, target removal was much less efficient than in simulated data. This may partially be explained by the effects discussed above since the targets are densely spaced in the real scene, see Figure 4.17. However, the significantly decreased energy removal indicates that there are additional causes besides target density. As indicated in Figure 4.15, this is believed to be due to mismatch between steering vectors and received echoes.

The mismatch may be the result of several factors. One cause may be a discrepancy between the modeled returns and the true scattering behavior. The returns modeled according to Equation 3.3 are based on several assumptions such as:

- Point-like targets without internal structure,
- Constant radar cross-section during each CPI,
- Constant radial velocity during each CPI,
- Negligible multi-path scattering,
- Negligible hardware-induced distortion.

In practice, some or all of the above assumptions may be violated, resulting in slight mismatches which degrade target signal removal.

Furthermore, the steering vectors were further simplified by assuming negligible intra-pulse Doppler-induced phase variation and range-bin migration. Without these assumptions, the efficient implementations of both the correlation and projection steps cannot be applied directly, increasing the computational burden significantly.

During the project, some attempts to improve the match between steering vector and data were taken. One approach that was tested was by using the steering vector model in Equation 3.3 but with the bandwidth, pulse length and a variable tapering optimized to match the return from one specific signal in the data as closely as possible. The parameters were fitted to two different returns available in the data. The first return was from a radar reflector placed on a mast approximately 9 km away. The second return was from a presumed car that was relatively isolated from other targets, even though the return from a completely isolated car was not found in the data. In both cases, the removed energy was slightly increased for the target the parameters were adapted to, but the overall removal for other targets was decreased. This approach was primarily intended as an attempt to counteract the effects of potential hardware-induced distortions by effectively modeling the true target returns as perturbations of the reference pulse.

The direction believed most straightforward, and potentially most promising, to improve the match between steering vectors and data is to model each found target as the sum of several nearby point scatterers. The exact implementation and effectiveness of this procedure remains to be studied, however.

5.4 Outlook

During the project, several areas for research and development of the algorithm was identified. The most promising direction is believed to be the development of steering vectors that better model the recorded data. As mentioned in Section 5.3.3, several simplifications and approximations were used to arrive at the form of the steering vectors used in this thesis. A careful analysis of which of these assumptions have a significant impact on the signal model, and how an improved signal model can be incorporated in the algorithm has the potential to refine the signal removal step, which would enhance the overall detector performance.

Furthermore, it may be fruitful to further investigate the steering vector selection and stopping criterion of the OMP algorithm. In the CA-OMP detector, both of these mechanisms were replaced with a traditional CA-CFAR detector applied iteratively to the residual. However, other approaches are possible and may improve the performance. One interesting approach would be to use a generalized likelihood ratio test on partially overlapping partitions of the data.

Additionally, this thesis did not include a rigorous analysis of the computational complexity of the algorithms. However, it is clear that the OMP procedure requires larger computational resources than the traditional CA-CFAR algorithm. Both runtime and memory footprint may, in practice, be limiting factors of real-time deployment of the algorithms.

Finally, it may be interesting to evaluate the performance of the OMP algorithm on another dataset. The dataset used in this thesis was a pulse-Doppler radar with the parameters given in Table 3.1, adapted for detection of targets close to the radar. In particular, the pulse length, 2.0 µs, was quite short and the bandwidth, 8.0 MHz, quite large compared to the sampling frequency, 10 MS/s. As seen in Figure 4.3, a larger oversampling rate renders the energy removal less sensitive to the parameter estimation. A longer pulse length may also decrease the hardware induced distortions of the signal giving the electrical components more time to accurately construct and modulate the signal. In general, the radar parameters can affect both the hardware-induced distortions and how appropriate some of the assumptions made are regarding the propagation and reflection of the radar signal.

6

Conclusion

The aim of this Master's thesis was to evaluate the OMP algorithm as a detector for pulse-Doppler radar. The algorithm was also compared to a traditional CA-CFAR detector and a hybrid of the OMP and CA-CFAR detector was proposed. The main hypothesis was that the iterative signal estimation and removal in OMP could improve the resolution of interfering signals, leading to improved detector performance in clutter-heavy environments and scenarios with closely spaced targets.

In scenarios with an isolated target in AWGN, OMP and CA-CFAR achieved similar probabilities of detection across different SNRs. However, in cases with multiple interfering targets, the OMP detector showed some advantages. Compared to the CA-CFAR detector, OMP consistently resolved closely spaced targets at smaller separations than CA-CFAR, although imperfect signal removal caused more false alarms in high-interference regions. Overall, the OMP algorithm achieved the highest FoM for multi-target simulated data, followed by CA-OMP and finally CA-CFAR.

For real data, the energy removal step was significantly less effective. This resulted in several false alarms when the detector identified residuals as targets, and hindered the detection of weak targets. This behavior could partially be controlled through the parameter k, where large k generated more detections of weak targets at the expense of occasionally missing strong targets, and small k improved the detection of closely spaced targets at the cost of more false alarms and missed detections of weak targets. In contrast, the CA-OMP detector provided good detection coverage over the whole range-Doppler map and was able to resolve closely spaced targets better than the standard CA-CFAR detector.

The main challenge in applying the OMP detector on real data is believed to be the mismatch between the steering vectors and the true target returns. This mismatch likely stems from assumptions and approximations made when constructing the steering vectors. Future work could therefore focus on improving the signal model by investigating which assumptions meaningfully affect the steering vector fit to actual data. Another area that can be explored is alternative detection and stopping criteria, for example via a localized generalized likelihood ratio test. Lastly, the computational complexity of the OMP algorithm could be examined to determine the feasibility of real-time implementation.

Bibliography

[1] M. A. Richards, *Fundamentals of Radar Signal Processing (3rd Edition)*. McGraw Hill LLC, 2022, ISBN: 9781260468717.

[2] M. B. Guldogan and O. Arikan, "Detection of sparse targets with structurally perturbed echo dictionaries," *Digital Signal Processing*, vol. 23, no. 5, pp. 1630–1644, 2013, ISSN: 1051-2004. DOI: `https://doi.org/10.1016/j.dsp.2013.03.003`. [Online]. Available: `https://www.sciencedirect.com/science/article/pii/S1051200413000699`.

[3] J. Zhu, T. Zhao, T. Huang, and D. Zhang, "Analysis of random pulse repetition interval radar," in *2016 IEEE Radar Conference (RadarConf)*, 2016, pp. 1–5. DOI: `10.1109/RADAR.2016.7485193`.

[4] J. Wang, S. Kwon, and B. Shim, "Generalized orthogonal matching pursuit," *IEEE Transactions on Signal Processing*, vol. 60, no. 12, pp. 6202–6216, 2012. DOI: `10.1109/TSP.2012.2218810`.

[5] S. Chatterjee, D. Sundman, and M. Skoglund, "Look ahead orthogonal matching pursuit," in *2011 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2011, pp. 4024–4027. DOI: `10.1109/ICASSP.2011.5947235`.

[6] J. Wen, C. Li, Q. Shu, and Z. Zhou, "Randomized orthogonal matching pursuit algorithm with adaptive partial selection for sparse signal recovery," *SIAM Journal on Imaging Sciences*, vol. 18, no. 2, pp. 1028–1057, 2025. DOI: `10.1137/24M1648624`. eprint: `https://doi.org/10.1137/24M1648624`. [Online]. Available: `https://doi.org/10.1137/24M1648624`.

[7] S. N. H. Shah, S. Semper, A. U. Khan, C. Schneider, and J. Robert, "Newtonized orthogonal matching pursuit for high-resolution target detection in sparse ofdm isac systems," *IEEE Transactions on Vehicular Technology*, vol. 74, no. 10, pp. 16 137–16 151, 2025. DOI: `10.1109/TVT.2025.3573165`.

[8] X. Bai, J. Han, J. Zhao, Y. Feng, and R. Tao, "Clutter cancellation in passive radar using batch-based clean technique," *EURASIP Journal on Advances in Signal Processing*, vol. 2021, Aug. 2021. DOI: `10.1186/s13634-021-00769-9`.

[9] M. Ankel, R. Jonsson, T. Bryllert, L. Ulander, and P. Delsing, "Bistatic noise radar: Demonstration of correlation noise suppression," *IET Radar, Sonar Navigation*, vol. 17, Nov. 2022. DOI: `10.1049/rsn2.12345`.

[10] M. Parker, *Digital Signal Processing 101 (Second Edition)*. Newnes, 2017, ISBN: 978-0-12-811453-7. DOI: `https://doi.org/10.1016/B978-0-12-811453-7.00018-4`. [Online]. Available: `https://www.sciencedirect.com/science/article/pii/B9780128114537000184`.

[11] G. W. Stimson, H. D. Griffiths, C. J. Baker, and D. Adamy, *Stimson's Introduction to Airborne Radar (3rd Edition)*. Institution of Engineering and Technology (The IET), 2014, ISBN: 978-1-61353-022-1. [Online]. Available: `https://app.knovel.com/hotlink/toc/id:kpSIARE004/stimson-s-introduction/stimson-s-introduction`.

[12] G. Li, *Advanced Sparsity-Driven Models and Methods for Radar Applications*. Institution of Engineering and Technology, 2021, ISBN: 978-1-83953-075-3.

[13] Y. Pati, R. Rezaiifar, and P. Krishnaprasad, "Orthogonal matching pursuit: Recursive function approximation with applications to wavelet decomposition," in *Proceedings of 27th Asilomar Conference on Signals, Systems and Computers*, 1993, 40–44 vol.1. DOI: `10.1109/ACSSC.1993.342465`.

[14] J. Högbom, "Aperture Synthesis with a Non-Regular Distribution of Interferometer Baselines," *Astron. Astrophys. Suppl.*, vol. 15, pp. 417–426, Jun. 1974.

[15] J. Tsao and B. Steinberg, "Reduction of sidelobe and speckle artifacts in microwave imaging: The clean technique," *IEEE Transactions on Antennas and Propagation*, vol. 36, no. 4, pp. 543–556, 1988. DOI: `10.1109/8.1144`.

[16] R. Bose, A. Freedman, and B. Steinberg, "Sequence clean: A modified deconvolution technique for microwave images of contiguous targets," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 38, no. 1, pp. 89–97, 2002. DOI: `10.1109/7.993231`.

[17] H. Deng, "Effective clean algorithms for performance-enhanced detection of binary coding radar signals," *IEEE Transactions on Signal Processing*, vol. 52, no. 1, pp. 72–78, 2004. DOI: `10.1109/TSP.2003.820075`.

[18] S. Mallat and Z. Zhang, "Matching pursuits with time-frequency dictionaries," *IEEE Transactions on Signal Processing*, vol. 41, no. 12, pp. 3397–3415, 1993. DOI: `10.1109/78.258082`.

[19] J. Tropp, "Greed is good: Algorithmic results for sparse approximation," *IEEE Transactions on Information Theory*, vol. 50, no. 10, pp. 2231–2242, 2004. DOI: `10.1109/TIT.2004.834793`.

[20] J. A. Tropp and A. C. Gilbert, "Signal recovery from random measurements via orthogonal matching pursuit," *IEEE Transactions on Information Theory*, vol. 53, no. 12, pp. 4655–4666, 2007. DOI: `10.1109/TIT.2007.909108`.

[21] W. Wang, R. Wang, R. Jiang, H. Yang, and X. Wang, "Modified reference window for two-dimensional cfar in radar target detection," *The Journal of Engineering*, vol. 2019, no. 21, pp. 7924–7927, 2019. DOI: `https://doi.org/10.1049/joe.2019.0687`. eprint: `https://ietresearch.onlinelibrary.wiley.com/doi/pdf/10.1049/joe.2019.0687`. [Online]. Available: `https://ietresearch.onlinelibrary.wiley.com/doi/abs/10.1049/joe.2019.0687`.

[22] J. Brewer, “Kronecker products and matrix calculus in system theory,” *IEEE Transactions on Circuits and Systems*, vol. 25, no. 9, pp. 772–781, 1978. DOI: `10.1109/TCS.1978.1084534`.

[23] H. Lev-Ari, “Efficient solution of linear matrix equations with applications to multistatic antenna array processing,” *Communications in Information and Systems*, vol. 5, no. 1, pp. 123–130, 2005, ISSN: 1526-7555. DOI: `10.4310/CIS.2005.v5.n1.a5`.

A

Appendix 1

A.1 Correlation step

One of the most time consuming parts of the OMP-algorithm execution is the correlation step, i.e. the calculation

$$\boldsymbol{E} = \boldsymbol{S}^{\dagger}\tilde{\boldsymbol{x}} = \begin{bmatrix} \boldsymbol{s}_1^{\dagger} \\ \boldsymbol{s}_2^{\dagger} \\ \vdots \\ \boldsymbol{s}_K^{\dagger} \end{bmatrix} \tilde{\boldsymbol{x}} = \begin{bmatrix} \boldsymbol{s}_1^{\dagger}\tilde{\boldsymbol{x}} \\ \boldsymbol{s}_2^{\dagger}\tilde{\boldsymbol{x}} \\ \vdots \\ \boldsymbol{s}_K^{\dagger}\tilde{\boldsymbol{x}} \end{bmatrix}.$$

However, the calculation can be sped up through the choice of steering vectors. The steering vectors are the Kronecker product [22, Eq. 2] of a vector describing the response in the fast time (range) dimension, $\boldsymbol{r}_i$, and a vector describing the response in slow time (doppler), $\boldsymbol{d}_i$, such that

$$\boldsymbol{s}_i = \boldsymbol{r}_i \otimes \boldsymbol{d}_i.$$

The range response is modeled as a time-delayed version of the pulse reference $c(t)$ sampled with the fast time spacing Δt_f for some time-delay τ_i,

$$\boldsymbol{r}_i = \Big[c\Big(0\Delta t_f - \tau_i\Big),\ c\Big(1\Delta t_f - \tau_i\Big),\ \ldots,\ c\Big((N_f - 1)\Delta t_f - \tau_i\Big)\Big]^T,$$

and the Doppler response is modeled as a complex sinusoid at some Doppler-frequency f_i sampled with the slow time spacing Δt_s,

$$\boldsymbol{d}_i = \Big[e^{j2\pi f_i 0\Delta t_s},\ e^{j2\pi f_i 1\Delta t_s},\ \ldots,\ e^{j2\pi f_i (N_s-1)\Delta t_s}\Big]^T.$$

The full steer vector evaluates to

$$\boldsymbol{s}_i = \begin{bmatrix} e^{j2\pi f_i 0\Delta t_s} c\Big(0\Delta t_f - \tau_i\Big) \\ e^{j2\pi f_i 1\Delta t_s} c\Big(0\Delta t_f - \tau_i\Big) \\ \vdots \\ e^{j2\pi f_i (N_s-1)\Delta t_s} c\Big(0\Delta t_f - \tau_i\Big) \\ e^{j2\pi f_i 0\Delta t_s} c\Big(1\Delta t_f - \tau_i\Big) \\ \vdots \\ e^{j2\pi f_i (N_s-1)\Delta t_s} c\Big((N_f - 1)\Delta t_f - \tau_i\Big) \end{bmatrix}.$$

The vectorized data, formed by stacking the the columns of the original data on top of each other, can similarly be written

$$\tilde{\boldsymbol{x}} = \begin{bmatrix} \tilde{x}\Big(0\Delta t_s + 0\Delta t_f\Big) \\ \tilde{x}\Big(1\Delta t_s + 0\Delta t_f\Big) \\ \vdots \\ \tilde{x}\Big((N_s-1)\Delta t_s + 0\Delta t_f\Big) \\ \tilde{x}\Big(0\Delta t_s + 1\Delta t_f\Big) \\ \vdots \\ \tilde{x}\Big((N_s-1)\Delta t_s + (N_f-1)\Delta t_f\Big) \end{bmatrix}.$$

Thus

$$\begin{aligned} \boldsymbol{s}_i^\dagger \tilde{\boldsymbol{x}} &= \begin{bmatrix} e^{j2\pi f_i 0\Delta t_s} c\Big(0\Delta t_f - \tau_i\Big) \\ e^{j2\pi f_i 1\Delta t_s} c\Big(0\Delta t_f - \tau_i\Big) \\ \vdots \\ e^{j2\pi f_i (N_s-1)\Delta t_s} c\Big(0\Delta t_f - \tau_i\Big) \\ e^{j2\pi f_i 0\Delta t_s} c\Big(1\Delta t_f - \tau_i\Big) \\ \vdots \\ e^{j2\pi f_i (N_s-1)\Delta t_s} c\Big((N_f-1)\Delta t_f - \tau_i\Big) \end{bmatrix}^\dagger \begin{bmatrix} \tilde{x}\Big(0\Delta t_s + 0\Delta t_f\Big) \\ \tilde{x}\Big(1\Delta t_s + 0\Delta t_f\Big) \\ \vdots \\ \tilde{x}\Big((N_s-1)\Delta t_s + 0\Delta t_f\Big) \\ \tilde{x}\Big(0\Delta t_s + 1\Delta t_f\Big) \\ \vdots \\ \tilde{x}\Big((N_s-1)\Delta t_s + (N_f-1)\Delta t_f\Big) \end{bmatrix} \\ &= e^{-j2\pi f_i 0\Delta t_s} c^\star\Big(0\Delta t_f - \tau_i\Big)\tilde{x}\Big(0\Delta t_s + 0\Delta t_f\Big) + \\ &+ e^{-j2\pi f_i 1\Delta t_s} c^\star\Big(0\Delta t_f - \tau_i\Big)\tilde{x}\Big(1\Delta t_s + 0\Delta t_f\Big) + \dots \\ &+ e^{-j2\pi f_i (N_s-1)\Delta t_s} c^\star\Big(0\Delta t_f - \tau_i\Big)\tilde{x}\Big((N_s-1)\Delta t_s + 0\Delta t_f\Big) + \dots \\ &+ e^{-j2\pi f_i 0\Delta t_s} c^\star\Big(1\Delta t_f - \tau_i\Big)\tilde{x}\Big(0\Delta t_s + 1\Delta t_f\Big) + \dots \\ &+ e^{-j2\pi f_i (N_s-1)\Delta t_s} c^\star\Big((N_f-1)\Delta t_f - \tau_i\Big)\tilde{x}\Big((N_s-1)\Delta t_s + (N_f-1)\Delta t_f\Big) \\ &= \sum_{k=0}^{N_f-1} c^\star\Big(k\Delta t_f - \tau_i\Big)\Bigg(\sum_{l=0}^{N_s-1} e^{-j2\pi f_i l\Delta t_s}\tilde{x}\Big(l\Delta t_s + k\Delta t_f\Big)\Bigg). \end{aligned}$$

If we write $f_i = \xi_i/(N_s\Delta t_s)$ and $\tilde{x}\Big(l\Delta t_s + k\Delta t_f\Big) = \tilde{X}\Big(l,k\Big)$ we get

$$\sum_{k=0}^{N_f-1} c^\star\Big(k\Delta t_f - \tau_i\Big)\Bigg(\sum_{l=0}^{N_s-1} e^{-j2\pi\frac{\xi_i}{N_s}l}\tilde{X}\Big(l,k\Big)\Bigg).$$

and identify the second sum as element ξ_i of the discrete Fourier transform of $\tilde{X}(l,k)$ along the first dimension. Let

$$\sum_{l=0}^{N_s-1} e^{-j2\pi\frac{\xi_i}{N_s}l}\tilde{X}\Big(l,k\Big) = \mathcal{F}_1\Big(\tilde{X}\Big)(\xi_i,k)$$

and write

$$\sum_{k=0}^{N_f-1} c^\star\Big(k\Delta t_f - \tau_i\Big)\mathcal{F}_1\Big(\tilde{X}\Big)(\xi_i,k).$$

This can be identified as the cross-correlation between c and $\mathcal{F}_1(\tilde{X})$ along the second dimension for the time delay τ_i. In total we see that the correlation step $\boldsymbol{S}^\dagger\tilde{\boldsymbol{x}}$ can be calculated by taking the Fourier transform along the first dimension of the raw slow-time-fast-time data $\tilde{\boldsymbol{X}}$ and cross-correlating the data with the reference pulse. If we let $\tau_i = \zeta_i \Delta t_f$ both of these operations can be implemented efficiently with the Fast Fourier Transform, and corresponds exactly to the standard processing of pulse-Doppler radar data

$$\boldsymbol{S}^\dagger\tilde{\boldsymbol{x}} = \mathcal{F}_2^{-1}\Big\{\mathcal{F}_2^*(c)\cdot\mathcal{F}_2\Big(\mathcal{F}_1(\tilde{\boldsymbol{X}})\Big)\Big\}$$

A.2 Projection Operation

In general, the projection matrix

$$\boldsymbol{P_F} = \boldsymbol{F}\Big(\boldsymbol{F}^\dagger\boldsymbol{F}\Big)^{-1}\boldsymbol{F}^\dagger$$

is unfeasible to calculate explicitly, as $\boldsymbol{F}$ contains steer vectors as columns, each of size $M = (\#\text{fast-time samples})\times(\#\text{slow-time samples})$, resulting in $\boldsymbol{P_F}$ being an $M\times M$ matrix which both requires a lot of memory and computational power to invert. However, the projection $\boldsymbol{P_F s}$ can also be interpreted as the vector spanned by $\boldsymbol{F}$ that is closest to $\boldsymbol{s}$ in the Euclidean sense,

$$\boldsymbol{P_F s} = \underset{\boldsymbol{y}\in\text{span}(\boldsymbol{F})}{\arg\min}||\boldsymbol{s}-\boldsymbol{y}||_2^2.$$

If we parametrize $\boldsymbol{y} = \boldsymbol{F\theta}$ we receive the least squares problem

$$\hat{\boldsymbol{\theta}} = \arg\min_{\boldsymbol{\theta}}||\boldsymbol{s}-\boldsymbol{F\theta}||_2^2. \tag{A.1}$$

The structure of the matrix $\boldsymbol{F}$ can be utilized to more efficiently solve the problem. The matrix $\boldsymbol{F}$ is a concatenation of all previously found steer vectors,

$$\boldsymbol{F} = [\boldsymbol{s}_1|\boldsymbol{s}_2|\dots|\boldsymbol{s}_P]$$

where each column vector is the Kronecker product ($\otimes$) of a vector describing range, $\boldsymbol{r}$ and a vector describing Doppler $\boldsymbol{d}$, $\boldsymbol{s} = \boldsymbol{r}\otimes\boldsymbol{d}$. The Kronecker product [22, Eq. (2)] of $\boldsymbol{A}$ $m\times n$ and $\boldsymbol{B}$ $p\times q$ is an $pm\times qn$ matrix defined as

$$\boldsymbol{A}\otimes\boldsymbol{B} = \begin{bmatrix} a_{11}\boldsymbol{B} & \dots & a_{1n}\boldsymbol{B} \\ \vdots & \ddots & \vdots \\ a_{m1}\boldsymbol{B} & \dots & a_{mn}\boldsymbol{B}\end{bmatrix}.$$

Using this, $\boldsymbol{F}$ can be rewritten as

$$\boldsymbol{F} = \Big[\boldsymbol{r}_1\otimes\boldsymbol{d}_1|\boldsymbol{r_2}\otimes\boldsymbol{d_2}|\dots|\boldsymbol{r}_P\otimes\boldsymbol{d}_P\Big] = \boldsymbol{R}*\boldsymbol{D}$$

for $\boldsymbol{R}\equiv[\boldsymbol{r}_1|\boldsymbol{r}_2|\dots|\boldsymbol{r}_P]\in\mathbb{C}^{M_R\times P}$ and $\boldsymbol{D}\equiv[\boldsymbol{d}_1|\boldsymbol{d}_2|\dots|\boldsymbol{d}_P]\in\mathbb{C}^{M_D\times P}$ and where $*$ denotes the Khatri-Rao product [22, Eq. (14)]

$$\boldsymbol{A}*\boldsymbol{B}\equiv(\boldsymbol{A}_{ij}\otimes\boldsymbol{B}_{ij})_{ij}.$$

Furthermore, the Kronecker product of two vectors is related to the outer product as $\boldsymbol{v}_1 \otimes \boldsymbol{v}_2 = \text{vec}\ (\boldsymbol{v}_2 \boldsymbol{v}_1^T)$ where vec denotes the vectorization of a matrix by stacking the columns on top of each other [22, T3.4]. Thus $\boldsymbol{s} = \boldsymbol{r} \otimes \boldsymbol{d} = \text{vec}\ (\boldsymbol{d}\boldsymbol{r}^T)$. The least squares problem A.1 can now be rewritten as

$$\hat{\boldsymbol{\theta}} = \arg \min_{\boldsymbol{\theta}} ||\text{vec}\ (\boldsymbol{d}\boldsymbol{r}^T) - (\boldsymbol{R} * \boldsymbol{D})\boldsymbol{\theta}||_2^2. \tag{A.2}$$

If we assume the vector $\boldsymbol{\theta}$ can be written as the diagonal of a square diagonal matrix $\boldsymbol{\Theta}$ such that $\boldsymbol{\theta} = \text{vecd}\ \boldsymbol{\Theta}$. The problem A.2 then has the solution [23, Eq. (9)]

$$\text{vecd}\ \hat{\boldsymbol{\Theta}} = \left[(\boldsymbol{R}^\dagger \boldsymbol{R}) \odot (\boldsymbol{D}^\dagger \boldsymbol{D})\right]^{-1} \text{vecd}\ \left(\boldsymbol{D}^\dagger \boldsymbol{d}\boldsymbol{r}^T (\boldsymbol{R}^\dagger)^T\right)$$

where $\odot$ refers to element-wise multiplication.

The above solution only requires manipulation of matrices of sizes $M_R \times P$ and $M_D \times P$ which are much smaller than manipulation of the $M \times P$, where $M = M_R M_D$ matrices used in the straight-forward approach using Matlabs mldivide operator on Equation A.1.

DEPARTMENT OF ELECTRICAL ENGINEERING
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden
www.chalmers.se

CHALMERS
UNIVERSITY OF TECHNOLOGY