Radar Drone Detection: Drone RCS Analysis

BSc Thesis

S. Abnaamar & E.I. Özdal

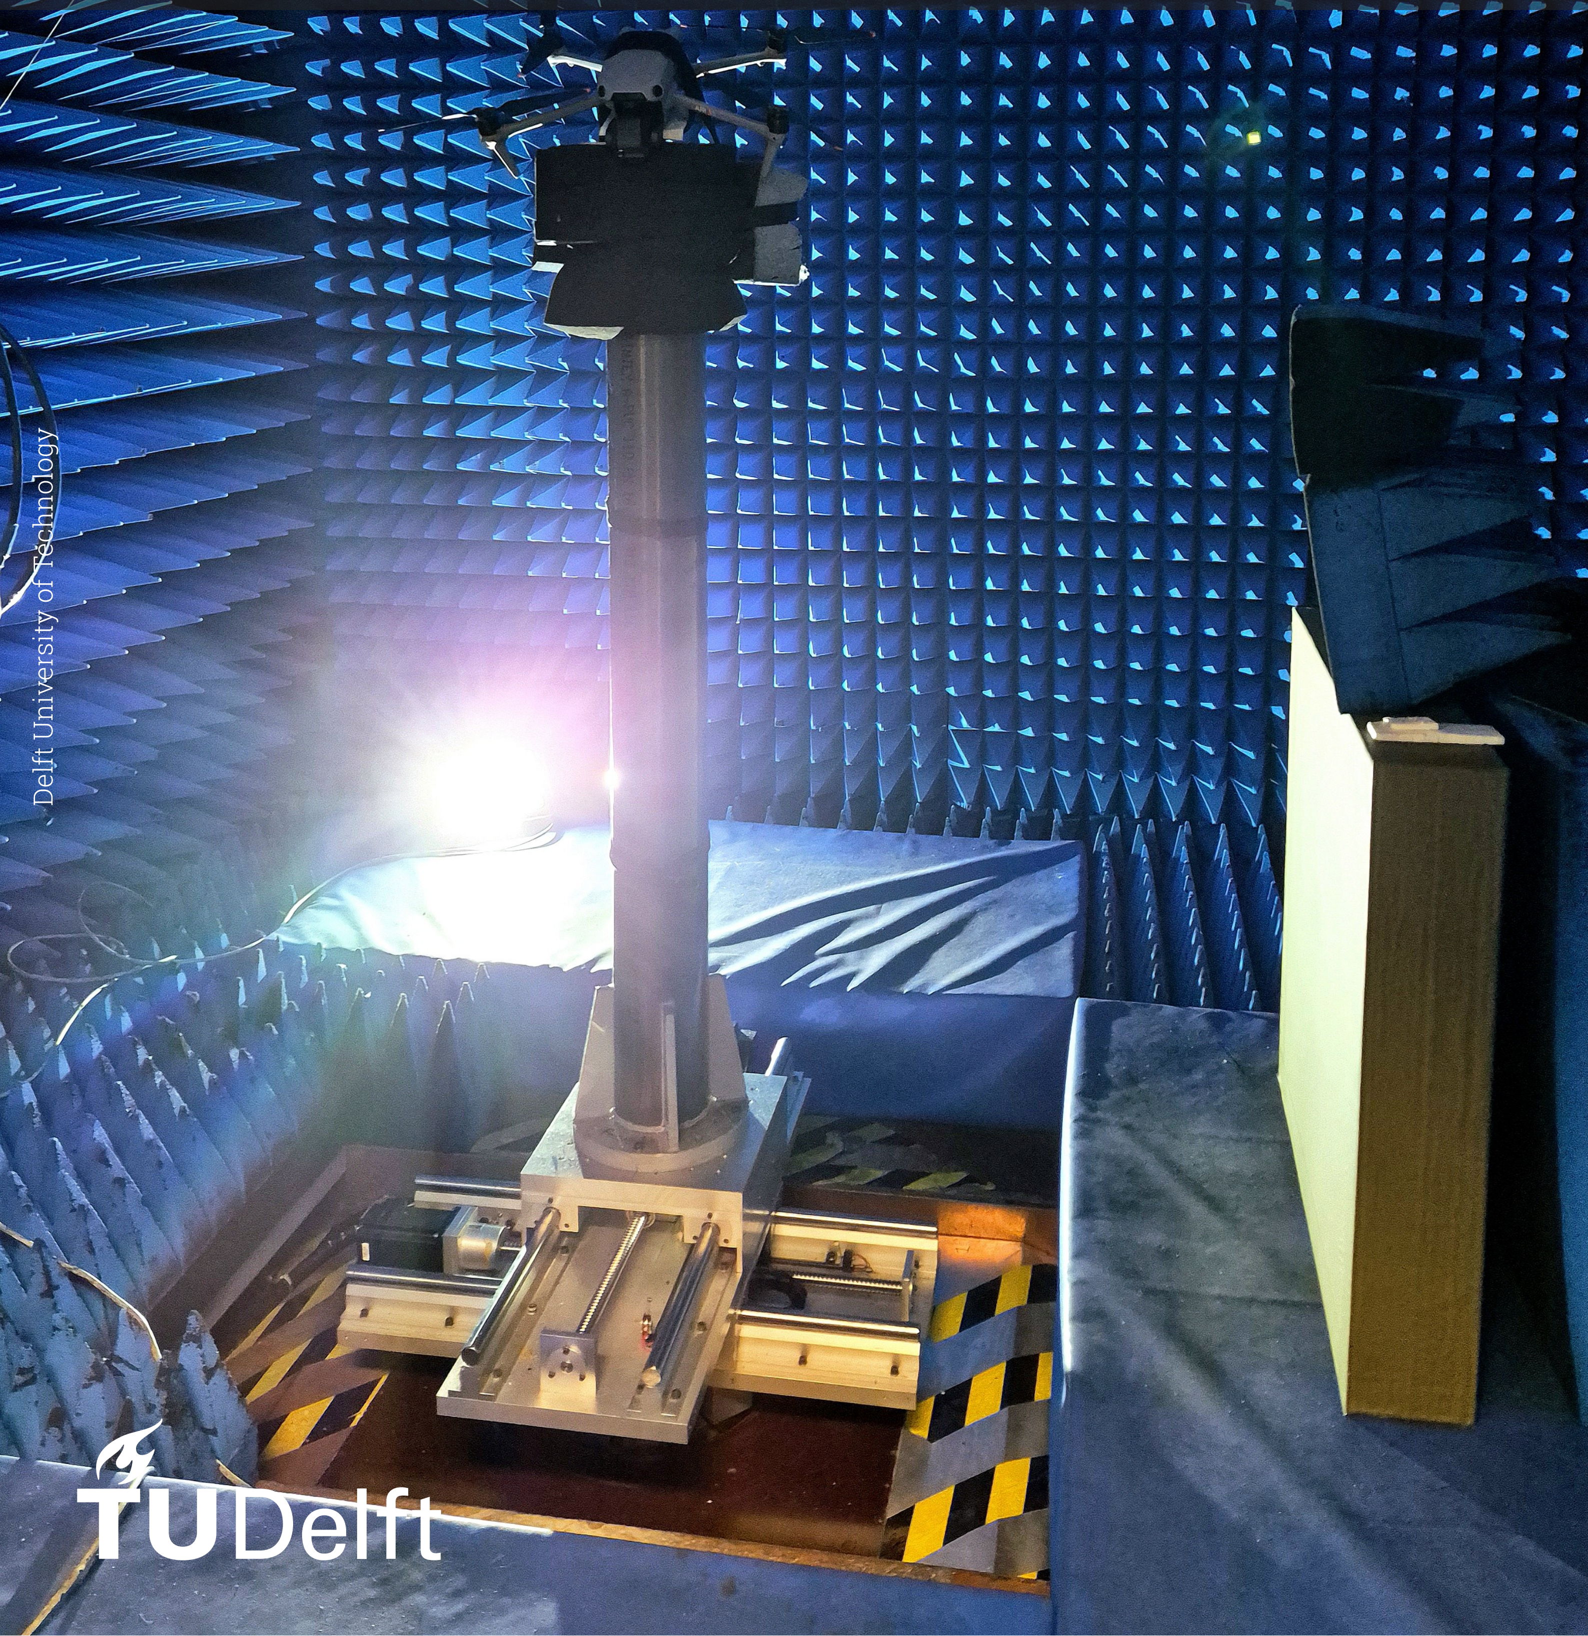

Delft University of Technology

Radar Drone Detection: Drone RCS Analysis

BSc Thesis

by

S. Abnaamar & E.I. Özdal

Supervisors: prof.dr. O. Yarovyi
dr. N.C. Kruse
Faculty: Faculty of Electrical Engineering, Mathematics and Computer Science, Delft

TU Delft

Abstract

The increasing accessibility of consumer drones raises safety and security risks, emphasising the necessity of reliable drone detection systems. This thesis characterises the radar cross-section of two commercial drones at 24 GHz frequency: the DJI Air 3S and DJI Neo, forming part of a broader drone detection and tracking system using frequency-modulated continuous wave radar. Measurements were performed in the Delft University Chamber for Antenna Tests anechoic chamber across a range of azimuth and elevation angles, using two different-bandwidth radar configurations. Here background subtraction and range-gating was utilised to process the results. The higher bandwidth dataset yielded an improved transmit-receive isolation and was prioritised for statistical modelling. Elevation angle configurations were pooled into a combined dataset, due to being considered a non-significant factor in forming the radar cross-section distribution. Five candidate models were compared, including the Swerling I and III target models, the gamma, Weibull, and log-normal distributions. Model selection was based on goodness-of-fit tests including Kolmogorov-Smirnov and Pearson chi-squared tests, along with the Akaike and Bayesian information criteria. Swerling III was deemed the optimal model for the DJI Air 3S, due to statistically performing the best and also being physically grounded; Swerling III captures a dominant scatterer with smaller adjacent scatterers, where the DJI Air 3S has a main body and smaller, protruding propeller arms. For the DJI Neo, the Weibull distribution was selected, while being statistically tied with Swerling III, as its shape parameter provides a measure for radar cross-section fluctuation behaviour, and the drone's compact geometry is insufficiently consistent with the physical implications of Swerling III.

Preface

This thesis reports on the development of a radar-based drone detection system commissioned by Dopplium. The project consists of three parallel subgroups: analysis of drone radar reflection characteristics, radar data processing, and drone target tracking. This work discusses the analysis and modelling of the radar cross-section (RCS) of low-RCS drone targets.

We are grateful to our direct project supervisor dr. N.C. Kruse for providing many sharp, targeted insights and repeatedly challenging our decision-making. We equally thank prof.dr. O. Yarovyi, dr. G. Joseph, and dr. F. Fioranelli for their tailored feedback at key stages of this work. We are also grateful to Dopplium and the Microwave Sensing, Signals and Systems (MS3) group for offering us the opportunity to undertake this project as part of their team.

S. Abnaamar & E.I. Özdal
Delft, July 2026

Contents

1 Introduction

The advancement of drone technology and a reduction in drone production costs have the potential to cause significant changes in modern society. Besides posing airspace disruption risks, drones can be used for illegal smuggling, espionage, targeted remote attack, and property damages [12]. Annually, hundreds of drone incidents cause great economic damages and societal turmoil throughout the globe [3]. Thus, an increase in drone accessibility and their technological advancements requires appropriate measures to mitigate safety hazards.

1.1. Project overview

To counteract hostile drone usages, one must first be able to detect when a drone is present in the region of interest and also be able to track it. Realising this is the main focus of our project; implementing a drone detection and tracking system through a 24 GHz Frequency-Modulated Continuous Wave (FMCW) radar. The project is split into three general parts; determining the radar cross-section (RCS) profile of the relevant drones; a signal processing part for range, velocity & angle; and a drone tracking aspect. These three project segments aim to collectively realise a drone detection system for practical use.

This thesis documents the first segment of this project, which is specifically concerned with finding the RCS profiles of the DJI Air 3S and DJI Neo drones and statistically modelling them. The radar cross-section of an object is a measure of the amount of power it scatters back to the radar for a given amount of transmitted power. It is a function of both the frequency of the wave that is transmitted, as well as the physical dimensions of the object.

For complex targets like drones and planes, the RCS can vary strongly when observing the object under different angles. This gives rise to RCS profiles. Presences of the multiple different scattering mechanisms can all constructively or destructively interfere with one another, where these RCS profiles often have a chaotic nature. In these profiles, even slight tilts can cause the object to change its RCS drastically. For this reason, RCS is modelled as a random variable.

1.2. State of the art

The literature states that the RCS of small drones has been found to generally be very low [9, 14, 30, 22]. This causes severe issues when it comes to detecting them and differentiating them from other moving targets. An accurate statistical model of the RCS of the drones chosen in this project segment will aid the signal processing subgroup in their analysis. This is because a more realistic profile and characterisation of the target allows them to appropriately calibrate their detection algorithms onto the physical model.

Furthermore, the maximal range at which an FMCW radar can unambiguously detect an object is inversely proportional to the bandwidth chosen for the frequency modulated chirps it sends out. However, The minimal distance between two targets that an FMCW radar needs to separate them is also inversely proportional to the bandwidth. In other words, increasing the bandwidth causes the maximum range to go down, while the resolution in target range gets better, and vice versa.
In reality, however, the maximum range the radar can detect is decided by the power the radar is able to receive back from the target, relating directly to its RCS. It is then impractical to sacrifice range resolution to get a higher unambiguous range, when a target at that range is unable to reflect enough power to even be visible by the radar. Knowing the RCS provides ground for choosing a bandwidth that will result in the best maximal range possible without sacrificing range resolution.

The RCS profiles of both rotary wing drones [10] and fixed-wing drones [5] have been studied at frequencies near 24 GHz. As explained in great detail by Ezuma et al. [10], an RCS measurement setup can be split into two types. The first being an outdoor measurement, which has the advantage of allowing both the transmitted and reflected waves to have entered the far field region by the time they reach the target and radar respectively. This is important as RCS is defined assuming the far-field condition holds, as this is standard for practical radar operation in the field.
The downside of such a setup is the difficulty of accounting for the presence of clutter. This clutter makes it difficult to ascertain how much the target itself contributed to the total power received by the radar. The second type, an indoor measurement, has the advantage of being able to control clutter and surface reflections, but has the disadvantage of often failing to be big enough to guarantee the far field condition holds for both transmitted and reflected waves. Nevertheless, the research on the measurement of RCS is predominantly focused on indoor setups [5, 10, 21, 22].

Methods for dealing with the near-field effects present in anechoic chamber measurements have been studied greatly. The methods range from constructing compact-range setups, where parabolic mirrors are used to focus the spherical wavefronts into plane waves [10], to correcting for the spherical wavefronts by means of back-projection [24]. Moreover, essentially all research done on drone RCS profiles in an anechoic chamber use a Vector Network Analyzer (VNA) instead of a radar to more precisely know the system losses and transmitted power.

Most studies have been conducted on the RCS profiles of drones at a carrier frequency of 24 GHz or higher, testing only the case of the drone facing the radar line of sight [10, 21, 22]. This means there are remarkably few studies on the effects of aspect angle on the RCS profile of a drone at high frequency.

Furthermore, the literature does not conform to consensus distributions for classifying the commercial drone RCS. Variations in choice include using Swerling models [15, 18], gamma or Weibull distributions [29, 32], log-normal distributions [9], the Rician distribution [33], among some others. Some even develop their own probability density function [11]. These models are generally assessed on performance through goodness-of-fit testing and statistical metric rankings. Goodness-of-fit tests often include the Cramér-von Mises distance [29, 30], Kolmogorov-Smirnov test [29, 30, 32], while the statistical metrics often include the Akaike Information Criteria [9, 10] and Bayesian Information Criteria [10].

1.3. Project scope & problem definition

This project aims to study the RCS profile of the DJI Air 3S and DJI Neo drones at a selection of different elevation angles, using a 24 GHz FMCW radar. This will include making measurements in an anechoic chamber. Due to the limited amount of time available for the project, the measurements will not include any measures for avoiding or reversing the near-field effects, and will instead focus on using calibration factors to compensate for the errors introduced by this more simplistic setup. Additionally, the measurements will be performed with the 24 GHz EV-Tinyrad24G radar, instead of a VNA.

Then, Swerling target models alongside three statistical models will be used to characterise the RCS distribution. The performance of these models will be evaluated using the Kolmogorov-Smirnov and chi-squared goodness-of-fit tests alongside Akaike & Bayesian Information Criteria as scoring metrics. The definitive model of the drone will be used to simulate the drone in a flight simulation to integrate into the radar processing subgroup. Additionally, a moving-propeller Doppler spread analysis will be performed for determining the maximal range for detection.

1.4. Structure of the thesis

To conclude the introduction, a general overview of the thesis will be provided.
Firstly, Chapter 2 presents the programme of requirements for the deliverables of this project. Then, Chapter 3 will discuss the design process used to develop this product, including the measurements setup, processing methodology, statistical models, and the maximum range evaluation method. Chapter 4 presents the results of this project, consisting of the RCS profiles of the drones under scrutiny and the performance of the chosen statistical models when applied to these profiles. Following this, a discussion of the validity of the results is presented Chapter 5, including comparisons with the aforementioned literature and an evaluation of the sources of error. Lastly, a conclusion of the thesis is provided in Chapter 6. Furthermore, two appendices discuss the ethical implications of this project (Appendix A) and the usage of generative AI throughout the project (Appendix B).

2
Programme of Requirements

This chapter provides an overview of the potential stakeholders for the product, alongside the project scoping and technical & functional requirements.

2.1. Product stakeholders

There are multiple categories into which you can compartmentalise the potential stakeholders with right to our project. These are listed below:

Table 2.1: Potential stakeholder analysis for the product (external stakeholders)

Stakeholder	Description
Governmental organisations	Drone detection technologies are of great interest to governmental organisations, especially military/defence departments. Drones are used increasingly more often in military/warfare contexts, which means this group of stakeholder would be interested in our project due to the relevance and importance of intercepting drone attacks/spying. Another application for the government would be, for instance, the smuggling of contraband into prisons.
Government officials or people at risk of assassination or spying	Government officials or people at risk of assassination or spying would be another potential stakeholder. Since this group is a highly valuable target, a direct drone attack would be plausible, which our technology aims to detect in order to prevent great harms. Since drones have the potency to kill these individuals, this stakeholder would highly value our project.
Drone users	In turn, another stakeholder would be such drone users. Whether it would be a hostile country, organisation, or individual, the potential of harm infliction boosts the importance of implementing drone detection measures.
Airport/airspace security	Furthermore, our project is relevant for airport/airspace security. Since drones can cause significant safety concerns or monetary losses for airspace infrastructure, this stakeholder can use our technology for implementing an additional surveillance layer for combating drone interferences.
Large-scale event organisers	Another stakeholder would be large-scale event organisers. A crowded gathering is vulnerable to a drone attack, which could cause great harms to the gathered masses, or contraband smuggling through drones. This stakeholder would, through our technology, add a layer of protection that mitigates aerial interference with the event.

Table 2.2: Potential stakeholder analysis for the product (institutional stakeholders)

Stakeholder	Description
Research or development organisations	Research or development organisations would also be greatly interested in our project. This is due to the incentive for the improvement, testing, and applications of drone detection technologies.
TU Delft, *Dopplium*, the company CEO & project supervisors, and the project students	Lastly, a stakeholder that is affected by this project includes the university-institute, involved company, company CEO, and project supervisors that facilitate this project. Along with the students working on this project, this group is directly involved with the project.

All in all, there are many parties and stakeholders that would be interested in the development of our drone detection technology. These have individually been listed and explored with right to their role in our project.

2.2. Project scoping

The project is based on detection and characterisation of small drones through radar-based detection, using Radar Cross Section (RCS) analysis and micro-Doppler signature modelling. Below is a table indicating the project scoping, describing the needs, objectives, and success measures for five stakeholders.

Table 2.3: Project scoping: stakeholders, needs, objectives, and success measures

Stakeholder	Needs	Objectives	Success Measure
Government / military	Reliable detection and surveillance	Minimise drone attacks and spying	Probability of detection, P_D, above a defined threshold
People at assassination or spying risk	Timely warnings	Reduce the time before drone threats are observed	Reduction in average detection time, t, for drone threats
Large-scale event organisers	Protection against drone interference and illicit smuggling	Maximise drone detection rate in crowded event spaces	Increase in average number of detections, N_{avg}, normalised to event size and duration
Airport / airspace security	Enhanced surveillance systems	Minimise drone interference in airspace	Reduction in drone-caused collisions
Research and development organisations	Realistic models for drone characteristics	Maximise understanding of drone RCS and micro-Doppler effects	Accuracy of statistical models measured through goodness-of-fit metrics and prediction error thresholds

2.3. List of functions

There are a number of functions our project product will handle. For the RCS analysis subgroup, functions 7 to 9 are relevant. All requirements have been listed below:

1. Detect small-sized drones
2. Track the path of these drones over time
3. Estimate the range (i.e., distance) the detected drone is located at
4. Estimate the velocity the drone is moving at
5. Estimate the angle the drone is located at
6. Provide a practical graphical user interface
7. Provide a dataset of drone radar cross-section measurements
8. Provide a realistic radar cross-section model of a particular drone
9. Provide electromagnetic simulations of the drone.

2.4. List of requirements

While the full system has an extensive list of additional requirements, this section will cover the requirements of the RCS analysis subgroup. These have been listed below.

Functional requirements:

1. The system shall use Frequency-Modulated Continuous Wave (FMCW) radar signals at a bandwidth of 24.000 to 24.250 GHz.
2. The system shall emit a maximum transmit power of 250mW.
3. The radar cross-section (RCS) measurement system shall be calibrated according to reference targets, including a dihedral, and two spheres of $r = 0.125\,\mathrm{m}$ and $r = 0.05\,\mathrm{m}$.
4. The radar cross-section (RCS) measurement system shall have an accuracy within $\pm 3\,\mathrm{dBsm}$ of the theoretical value for these reference targets.
5. The system shall store measurement recordings for at least seven days.
6. The system shall measure RCS values reaching as low as $-30\,\mathrm{dBsm}$.
7. The RCS fluctuations measured by the system shall be classified using Swerling models and two statistical distribution models.
8. The RCS model must include elevations of at least $0° - 30°$ in elevation with $5°$ resolution.
9. The RCS model must include ranges of $\pm 90°$ in azimuth with $3°$ or $5°$ resolution.

Reliability requirements:

10. The radar must be able to operate for half an hour continuously.
11. The RCS model must score below 1000 on the Akaike Information Criterion Test [9].
12. The radar shall be able to operate in light-interference weather conditions such as light rain (<2.5mm/h) and slight breezes (<11km/h).

Security/Ethics:

13. The system shall allow radar recordings to be deleted after analysis when they are no longer needed.
14. The system shall meet power transmission safety regulations of the Netherlands.
15. The system shall strictly operate at the given frequency band to prevent unintended interference with the environmental frequency-band infrastructure.

3 Materials & Methods

This chapter describes the measurements and methodology devised for determining the radar cross-section (RCS) profiles of the DJI Air 3S and DJI Neo. It analyses statistical characteristics, develops models and a simulation, and discusses implications of these results for the maximal range at which these drones would be able to be measured by the radar.

The RCS of an object can be inferred from the amount of power reflected back by the drone which is received by the radar. Using the radar equation, the RCS is given by:

$$\sigma = P_r \frac{(4\pi)^3 R^4}{P_t G_t G_r \lambda^2} \tag{3.1}$$

with P_r and P_t being the received and transmitted powers, R the distance between the drone and radar, λ the center wavelength that the radar operates at and G_r and G_t being the receiver and transmitter antenna gain.

The measurements that were conducted consisted of illuminating the drones with a 24 GHz radar and analysing the power it received back in the form of reflections. To minimise the power received by the radar originating from sources other than the drone (i.e. clutter, noise, and multipath reflections), all measurements were performed in the Delft University Chamber for Antenna Tests (DUCAT), which is the anechoic chamber of the Microwave Sensing and Systems group of the TU Delft. These measurements, taken at various azimuth and elevation angles for both drones, were then processed to extract the power that was received by the radar. The obtained results were calibrated using measurements taken of a collection of reference objects whose RCS could be analytically calculated.

After converting power to RCS using Equation 3.1, a statistical model was developed for the RCS of both drones for a number of distributions. A choice between these models was then made using various goodness-of-fit tests and likelihood scoring metrics. Lastly, these models are integrated into the radar processing subgroup, after which the maximum measurable range for the drones under scrutiny was analysed.

3.1. Overview of the setup and measurements

The measurements were done using the far-field configuration of the DUCAT, for which illustrations are provided in Appendix C. In this configuration, the DUCAT is set up with two pillars, separated by a distance of 3.7 meters. One of which having the capacity to be rotated remotely with the use of stepping motors. The object being studied was placed on this rotating pillar, while the radar was mounted on top of the other pillar fixed in angle. In addition, the drones were mounted on brackets that allowed them to be pitched forwards or backwards, allowing the study of the effect of different elevation angles on the RCS.

The radar used for these tests was the EV-Tinyrad24G radar [1]. Two sets of measurements were performed, labelled dataset 1 and 2, where different radar parameters are used in order to analyse their influence on the eventual RCS results. Table 3.1 summarises the parameters that were altered between the datasets.

Table 3.1: Radar parameters changed across the two measurement sets

	Dataset 1	**Dataset 2**
Bandwidth (MHz)	70	250
Chirp slope ($\mathrm{MHz}/\mu\mathrm{s}$)	0.345	1.23
Measurement time per drone orientation (s)	3	8

Tables 3.2 and 3.3 give an overview of all the measurements that were performed in the anechoic chamber for dataset 1 and 2 respectively.

Table 3.2: Measurements performed for dataset 1

Object	**Elevation**	**Azimuth**	**Azimuth step size**
No object	N/A	N/A	N/A
Sphere ($r = 5\,\mathrm{cm}$)	N/A	N/A	N/A
Sphere ($r = 12.5\,\mathrm{cm}$)	N/A	N/A	N/A
Dihedral reflector ($a = 0.142\,\mathrm{m}$)	$0°$	$0°$	N/A
DJI Air 3S	$0°$	$0° - 360°$	$2°$
	$5°$	$0° - 90°$	$4°$
	$10°$	$0° - 90°$	$4°$
	$15°$	$0° - 90°$	$4°$
	$20°$	$0° - 90°$	$4°$
	$25°$	$0° - 90°$	$4°$
	$30°$	$0° - 90°$	$4°$
	-$10°$	$0° - 90°$	$4°$
	-$20°$	$0° - 90°$	$4°$
	-$30°$	$0° - 90°$	$4°$
DJI Neo	$0°$	$0° - 90°$	$5°$
	-$30°$	$0° - 90°$	$5°$
	$30°$	$0° - 90°$	$5°$

Table 3.3: Measurements performed for dataset 2

Object	**Elevation**	**Azimuth**	**Azimuth step size**
No object	N/A	N/A	N/A
Sphere ($r = 5\,\mathrm{cm}$)	N/A	N/A	N/A
Sphere ($r = 12.5\,\mathrm{cm}$)	N/A	N/A	N/A
Dihedral reflector ($a = 0.142\,\mathrm{m}$)	$0°$	$0°$	N/A
DJI Air 3S	$0°$	$0° - 180°$	$3°$
	$5°$	$0° - 90°$	$5°$
	$10°$	$0° - 90°$	$5°$
	$15°$	$0° - 90°$	$5°$
	$20°$	$0° - 90°$	$5°$
	$23°$	$0° - 90°$	$5°$

As seen in the tables, both datasets contain one measurement without an object, effectively measuring the power coming from background scatterers, transmit-receive (Tx-Rx) leakage, interference, and noise. In addition, three reference targets were utilised. Namely, a metal sphere of radius $r = 5\,\mathrm{cm}$, a metal sphere of radius $r = 12.5\,\mathrm{cm}$, and a dihedral reflector with $a = 0.142\,\mathrm{m}$ & $b = 0.21\,\mathrm{m}$. These measurements were taken for both datasets, and the setup for the dihedral and a sphere are shown in Appendix C.2.

These reference measurements are necessary to determine loss factors introduced with this measurement setup, which are not included in the radar equation. The dihedral reflector is ideal for quantifying all the losses introduced by the radar electronics and any physical phenomena present in the setup. This is due to the dihedral reflector having a notably high RCS, meaning the power it reflects back is much higher than any background sources in the anechoic chambers.

However, precisely for this reason, it is ineffective for judging the effect of these background contributions on the total power received by the radar when a much weaker scatterer is illuminated. Table 3.4 shows the theoretical RCS of each of the three reference objects used, with the formulas for the RCS of a sphere and dihedral reflector in the optical region given by Equations 3.2 and 3.3. It is clear that the spheres have much lower RCS than the dihedral reflector, making them suitable for analysing these background contributions.

$$\sigma_{sphere} = \pi r^2 \tag{3.2}$$

$$\sigma_{dihedral} = \frac{8\pi a^2 b^2}{\lambda^2} \tag{3.3}$$

Table 3.4: Theoretical RCS values of the three reference objects at 24 GHz

Object	**RCS (**dBsm**)**
Sphere ($r = 5\,\mathrm{cm}$)	-21.1
Sphere ($r = 12.5\,\mathrm{cm}$)	-13.1
Dihedral reflector ($a = 0.142\,\mathrm{m}$, $b = 0.21\,\mathrm{m}$)	21.2

As for the drone measurements, a positive elevation was defined as the drone being pitched upwards. For every elevation angle that was tested, a sweep in azimuth was performed by rotating the supporting pillar at set increments. Here, an elevation of $0°$ represents the front of the drone pointing directly towards the radar, which would in reality correspond to a situation where the drone is diving towards the radar. For that reason, and due to the limited time available, more focus was put on positive elevation angles. The reasoning behind this was that these angles represented the radar looking at the bottom of the drone, which is the more common case for a radar situated on the ground.

Another consequence of time constraints was that a full $360°$ azimuthal sweep was only done for dataset 1 and at $0°$ elevation. For dataset 2, the fact that the DJI Air 3S is symmetric to the right and left of the radar's line of sight was used to justify only sweeping from $0°$ to $180°$. This symmetry argument was verified by taking test measurements at a small selection of negative azimuth angles. For the other elevation angles, in addition to the symmetry argument, the decision was made to only sweep from $0°$ to $90°$. The argument for this being that, given that the radar this project is based around is a short-range radar, it should be primarily used to detect drones which approach the radar, rather than being used for general surveillance. Hence, it was decided that the RCS values corresponding to the front of the drone were of greater interest than the back, and instead, the resolution in elevation was prioritised.

An important difference between the datasets is the absence of any measurements of the DJI Neo for dataset 2. Additionally, less elevation angles were tested for dataset 2, and the resolution in azimuth was decreased. The reason for this difference is that only half of the measurement time allocated to dataset 1 was available for dataset 2, considering the limited availability of the anechoic chamber.

For dataset 1, an additional measurements was performed with the drones' blades spinning. The drone was restrained to counteract it moving around in the chamber. This measurement will be used to gain insight into the effect of the rotating blades on the power distribution across Doppler velocity bins.

3.2. RCS determination methodology

Figure 3.1 shows a schematic representation of the hardware of the EV-Tinyrad24G radar, taken from the EV-Tinyrad24G User Guide [1]. The EV-Tinyrad24G has two transmit antennas, of which only one was used for this project. It also has four receive antennas forming four Rx channels. It transmits a frequency modulated chirp signal which reflects off of the target and is received by all four Rx channels. This received signal is then multiplied with the transmitted signal, which is also called mixing, and low-pass filtered to acquire a beat signal. This signal's frequency content is a one-to-one measure of the distance to the target. This is done by the ADF5904 receiver module, which passes the resulting beat signal to the 16-bit ADAR7251 Analog-to-Digital converter (ADC). The raw output of this ADC is the data which the user receives after a measurement is completed.

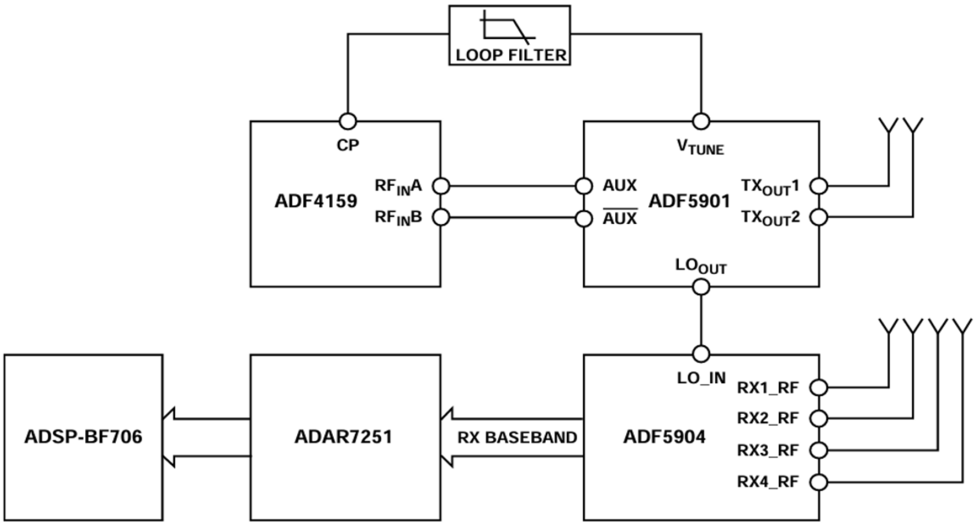

Figure 3.1: Block diagram showing the hardware of the EV-Tinyrad24G [1]

As explained previously, determining the RCS of an object mostly comes down to determining the power the radar received back from that object, given the amount of power the radar transmitted. Since the signal at the output of the ADC is the beat signal and not the reflected chirp received by the Rx antenna array, this conversion can be assigned a fixed power gain. This means that the power received by the antenna can be determined by calculating the power of the beat signal, and consequently reversing all gains and losses introduced by the mixing process and hardware components.

A single measurement consists of a large number of chirps. The signal used to determine received power was the average signal across all chirps, and channel 1 was used as the Rx channel. The reason for choosing a specific channel, in contrast with taking an average across all four channels, was that each channel seemed to have a different bias, and the averaged beat signal seemed to vary significantly between channels.

Simply treating the target as the source of all the power contained in the beat signal is insufficient, especially for the weaker scatterers. This is because this approach assumes that there are no other sources of power present in the measurements. Early on in the design process, this was discovered to not be the case. The FFT magnitude plots of the empty room measurement shown in Chapter 4 will indicate that there are, in fact, sources present in the measurements that are independent of the target. The most dominant of these sources being a signal at a range of about $1\ \mathrm{m}$, which was assumed to be the leakage between the transmit and receive antennas. This assumption will be explained in Chapter 4.

Due to the existence of background sources like the Tx-Rx leakage, these must first be subtracted from the beat signal before calculating its power. This is done in frequency domain, using the FFT magnitude of the measurement taken with no target. This approach is feasible due to the fact that these background sources are the same across all measurements. The subtraction of these sources is necessary because the power they contribute is comparable, or even greater than that of the weaker scatterers like the spheres and drones. This will be shown in the calibration results in Chapter 4. The choice to subtract the magnitudes of both FFTs was made due to it performing better during the calibration of the reference measurements than subtracting the FFTs normally.

An additional technique for minimising the effect of background sources is range-gating, which refers to only considering the power contained in a specific range bin or set of range bins. The utility of this technique becomes apparent when considering that the Tx-Rx leakage and the target occupy (in large part) different range bins. This makes it possible to avoid the main lobe of the leakage signal, while still including the ranges where the target's reflections are strongest.
The more range bins exist between the two signals, the better this technique works, as the main lobe of the target's reflection would overlap with a further, and therefore weaker, sidelobe of the leakage. Therefore, the contribution of the leakage signal to the power calculation will be much less. This is the reason a higher bandwidth was chosen for dataset 2, as this corresponds to a higher range resolution, meaning more range bins between the target and leakage. In other words, the main and sidelobes of both signals become narrower, causing their strength to decay faster as a function of range.

While this method does help minimise the influence that the leakage has over the eventual RCS results, it is not perfect. A large portion of power carried by the scattered waves originating from the target is contained in range bins outside of the bin the target resides in. Removing these other range bins cuts away a large portion of signal power, which has to be corrected for with a gain factor. This factor was chosen to be a static gain across all measurements.

Lastly, The mean of the beat signal was subtracted in time domain to reduce the leakage even further, as it is a relatively low frequency signal. The complete process for extracting the power received by the antenna from the raw Tinyrad data output is summarised as follows:

$$P_r = \frac{G_{cal}}{G_{system} N^2} \sum_{k=k_{min}}^{k_{max}} ||X[k]| - |S[k]||^2 \tag{3.4}$$

Where $X[k]$ and $S[k]$ are the N-point FFTs of the measured beat signal; $X[k]$ is averaged across all chirps and without DC offset, and $S[k]$ is the objectless beat signal. Moreover, k_{min} and k_{max} are the frequency bounds of the range bin of interest, G_{system} is the complete gain that the EV-Tinyrad24G hardware introduces, and G_{cal} is the calibration gain added to compensate for all losses. These losses include both losses from physical phenomena and the losses introduced by range-gating as explained previously.

An important loss associated to physical phenomena is the effect of the small size of the anechoic chamber. As explained in the introduction, this causes the radar to be in the near-field region of the waves scattering off of the target. This effect is seen as a source of error, and will be further discussed in Chapter 5. The calibration gain was tuned based off the reference measurements. The system gain can be acquired from the EV-Tinyrad24G User Guide and the Tinyrad calibration report [1, 13]. The total system gain amounts to $17\,\mathrm{dB}$.

The final step in extracting RCS results from the measurements is to convert this received power to RCS using Equation 3.1. For the objects whose RCS can be analytically determined, verifying the resulting RCS resumes itself to comparing the theoretical and measured values. The DJI Air 3S and DJI Neo, however, pose a complication; there exists no publicly available RCS profile for these drone models. To verify the sensibility of the drones' RCS results, they were compared to RCS profiles of comparable drone models that have already been studied at comparable frequencies.

Additionally, a sufficiently accurate electromagnetic model of the drone could be used to simulate its reflective properties at various angles, which serves as a more direct reference to compare with the RCS measurements. The complication with devising such a model is that common drone modelling techniques, specifically the thin-wire approximation, might not be valid for the high frequency electromagnetic radiation this project is concerned with. Due to the limited amount of time available for this project, it was decided that only the validity of the thin-wire approximation would be theoretically examined, leaving the modelling itself open for future work.

3.3. Statistical modelling of the drone RCS

After calibration, the measured drone RCS data can be used for constructing a distribution model required for the data processing subgroup to gauge the effectiveness of their detection algorithms. For this, two types of models are used: the radar-theory based Swerling models [36], and statistical distributions which are consistent with the literature on small drone RCS modelling. This model list selection is based on generally well-performing models commonly used in this field [5, 10, 18, 30, 32].

After a theoretical explanation of these models in Sections 3.3.1 & 3.3.2, the performance of each model has been evaluated using goodness-of-fit testing and statistical performance metrics in Section 3.3.3. Alternatives of these models and evaluation methods are provided in Section 3.3.4, before describing the integration of this RCS modelling into the radar processing subgroup's pipeline (Section 3.3.5).

3.3.1. Swerling target models

The Swerling target models I-IV are widely used in radar research, as they have a broad range of applications. More specifically, the selection of a Swerling model characterises both the scattering behaviour and temporal fluctuation rate of the target. Swerling I & II consider many equal RCS scatterers with low variance, while Swerling III & IV assume a single dominant scatterer with many smaller scatterers surrounding it. Moreover, the distinction between the odd- and even-numbered Swerling models lies with the fluctuation rate; the former in each pair signals a scan-to-scan variation in RCS, while the latter signals pulse-to-pulse variation. This describes whether each full measurement frame or each separate chirp respectively will count as an independent RCS data point. Below is a table providing an overview of the classification of the Swerling target models.

Table 3.5: Swerling target model classification, based on scattering behaviour and fluctuation rate

Model	Scatterer Type	Fluctuation Rate
Swerling I	Equal scatterers	Scan-to-scan
Swerling II	Equal scatterers	Pulse-to-pulse
Swerling III	Dominant scatterer	Scan-to-scan
Swerling IV	Dominant scatterer	Pulse-to-pulse

The distinction of fluctuation rate is based on whether the target's RCS changes considerably within a single pulse or merely between pulses, which ultimately affects the statistical variability of the RCS. This can be measured as a ratio between the chirp duration and the propeller blade rotation period $\tau_{\text{chirp}}/T_{\text{blade}}$, where $T_{\text{blade}} = (N_{\text{blades}} \cdot f_{\text{rotor}})^{-1}$. For the scan-to-scan models (Swerling I, III), the condition holds that $\tau_{\text{chirp}} \ll T_{\text{blade}}$. A chirp duration of $\tau_{\text{chirp}} = 224\,\mu\text{s}$ is used for the measurements, and the DJI Air 3S has $N_{\text{blades}} = 2$ blades per rotor. Since DJI does not publish rotor rotational speed, an approximation for f_{rotor} must be made based on typical multirotor drone propellers.

Similar drones typically operate at around $5\,\text{k} - 20\,\text{kRPM}$, depending on drone characteristics [26]. Therefore, the rotational speed is assumed to be $20\,\text{kRPM}$. This equates to $T_{blade} = (2 \cdot (20\text{k}/60))^{-1} = 1.5\,\text{ms}$, meaning $\tau_{\text{chirp}}/T_{\text{blade}} = 15\%$. Since Swerling II/IV assume pulse-to-pulse RCS fluctuation, the RCS must decorrelate on a timescale that is shorter than the chirp duration. In this case, the RCS is considered approximately constant over a single chirp, meaning Swerling II and IV are not considered as appropriate candidate models [36].

Instead, Swerling I and III are used as candidate models, both serving a distinct purpose; Swerling III models the DJI Air 3S target drone more appropriately, since the drone consists of a main body with smaller surrounding propellers (i.e., dominant scatterer surrounded with smaller, weaker scatterers). However, the measurement data might contain noise at a level such that the more complex Swerling III model might overfit the data. This is because the peak generated by the main drone body might not be sufficiently distinguishable with right to the noise level and surrounding scatterers. Hence, Swerling I was also included in the analysis.

3.3.2. Statistical distribution models

Moreover, three statistical distributions are utilised to model the drone RCS; the gamma, Weibull, and log-normal distributions. This selection is based on their prevalence in drone RCS modelling literature and their demonstrated performance [5, 30].

The gamma and Weibull distributions are closely related to each other; both belong to the generalised gamma family and reduce to exponential distributions under specific parameters [25]. The gamma distribution is described by three parameters; the curve shape k, its scale θ in the x-axis, and the location parameter loc, which in this case is set to zero due to the RCS being strictly non-negative. This RCS condition requires all used distributions to start at zero by physical constraint. The probability density function (PDF) of gamma is as follows:

$$f(x; k, \theta) = \frac{x^{k-1} e^{-x/\theta}}{\theta^k \cdot \Gamma(k)}, \quad x \geq 0 \tag{3.5}$$

where $\Gamma(k)$ is the gamma function with shape argument k, and x denotes the RCS random variable. Similarly, the Weibull distribution is specified by shape parameter β and scale parameter η, with its PDF being:

$$f(x; \beta, \eta) = \frac{\beta}{\eta} \left(\frac{x}{\eta}\right)^{\beta-1} e^{-(x/\eta)^\beta}, \quad x \geq 0 \tag{3.6}$$

The Weibull distribution is expected to perform particularly well due to its shape parameter β steering tail decay independently of the scale η. This flexible tail behaviour allows the distribution to capture different intensities of tail fluctuation behaviours [32].

To complete the standard modelling set used in the literature, the log-normal distribution is included. This distribution assumes that $\ln(x)$ is normally distributed, with its parameters being the log-mean μ and log-standard deviation σ. The log-normal PDF is:

$$f(x; \mu, \sigma) = \frac{1}{x \cdot \sigma\sqrt{2\pi}} \exp\left(-\frac{[\ln(x) - \mu]^2}{2\sigma^2}\right), \quad x \geq 0 \tag{3.7}$$

The reason for including this model is due to its ability to fit complex RCS targets with many contributing scatterers well. There is sufficient empirical evidence from other studies on RCS, where commercial drone measurements consistently have the log-normal as one of the best performing models [5, 9, 10]. This empirical observation stays consistent with highly varying scattering behaviour of different commercial drones, where the combined RCS fluctuations can give rise to approximately log-normally distributed RCS characteristics [10].

All three of these statistical distributions have $p = 2$ free parameters estimated from the measurement data via maximum likelihood estimation (MLE). On the other hand, the Swerling models contain only one hyperparameter m indicating the choice of Swerling model. The average RCS $\bar{\sigma}$ for the drone is treated as a fixed quantity, and is specified independent of the model fitting. Hence, $p = 1$ for Swerling. These parameter counts are accounted for in the model performance evaluations in Section 3.3.3.

3.3.3. Goodness-of-fit testing & likelihood scoring

To evaluate how well these candidate Swerling models and statistical distributions describe the RCS measurement data, two parallel assessment methods are used. Goodness-of-fit tests assess the statistical consistency of the measurement data with each theoretical model through hypothesis-testing. On the other hand, likelihood-based information criteria are applied to rank models on a common scale, accounting for model complexity to discourage overfitting. The goodness-of-fit tests include the Kolmogorov-Smirnov (KS) test and the Pearson chi-squared test, while the information criteria include the Akaike Information Criterion (AIC) and Bayesian Information Criterion (BIC).

Kolmogorov-Smirnov test

The KS test quantifies the maximum absolute deviation between the empirical cumulative distribution function (CDF) and the theoretical CDF. The reason for including the KS test is due to its ability to be used with smaller sample sizes, requiring no binning process, and its established use in drone RCS literature [30]. However, it assumes a continuous distribution under the null hypothesis, meaning the discrete and finite-resolution RCS measurements might introduce slight biases affecting the robustness of the test. For n observations, the KS test statistic D_n is defined as [23]:

$$D_n = \sup_x |F_n(x) - F(x)|, \tag{3.8}$$

where $F_n(x)$ symbolises the empirical CDF and $F(x)$ is the theoretical CDF. The null hypothesis H_0 states that the data follows the specified distribution $F(x)$. The evidence against H_0 can be assessed through either using a critical value determined by the significance level α, or by a p-value approach. For the former, H_0 would be rejected upon the condition $D_n > D_{\text{crit}}$, with D_{crit} commonly following a specified table [35]. The p-value, on the other hand, tracks how often a mismatch greater than size D_n can be observed under the specified distribution. These two statistics can be converted to one another.

However, since the model parameters are estimated from the same data used in this standard KS test, the fitted distribution is artificially closer to the data than an independently specified one. More specifically, the critical values become anti-conservative, causing H_0 to reject too rarely [19]. This leads to overly optimistic assessments of the model fit. Therefore, corrected p-values are acquired through the parametric bootstrap procedure described in Section 3.3.3.

Pearson Chi-Squared goodness-of-fit test

The second goodness-of-fit test used is the Pearson chi-squared test. The reason for implementing this test is due to how it can naturally be used on discretised RCS measurements, and it is particularly useful for detecting overall distribution shape differences. Additionally, it evaluates the same null hypothesis H_0 as the KS test. In this case, the data points are grouped into k bins, where the normalised squared deviations between observed and theoretical counts is summed:

$$\chi^2 = \sum_{i=1}^{k} \frac{(O_i - E_i)^2}{E_i}, \tag{3.9}$$

where χ^2 is the chi-squared statistic, O_i is the observed count in the ith bin, and E_i is the corresponding expected count under the theoretical model. This statistic asymptotically follows a chi-squared distribution with $(k-1-p)$ degrees of freedom. Here, p is the number of parameters that are estimated from the data [6]. As mentioned in Section 3.3.2, the Swerling target models have $p = 1$, while the Gamma, Weibull, and log-normal distributions have $p = 2$.

A disadvantage to using the chi-squared test is its binning effect; this test requires $n \geq 5$ samples per bin [6]. Since RCS distributions are often characterised with positively skewed models, such as Weibull, there might be sparse observations in the tails [32]. This risk is further increased considering the small sample sizes of the measurement dataset. As a result, the chi-squared bin condition might be violated in these extremities using linear binning. To account for this, a quantile-based method is applied for computing the chi-squared statistic χ^2, where each bin has an equal probability and expected counts across bins are approximately uniform [20].

Similar to the KS test, the estimation of parameters from the same data implies the standard critical values are anti-conservative. Hence why, the p-values are again corrected using the parametric bootstrapping procedure described in the following section.

Parametric bootstrap and Lilliefors correction

An issue with both goodness-of-fit tests is that the parameters are estimated from the same data distribution as the test evaluates. This means that the standard critical values no longer hold and p-values are overly optimistic. In the context of the KS test, this is called the Lilliefors correction [19]. A way to counteract this is through parametric bootstrapping.

This bootstrapping involves iteratively computing the observed test statistic D_n or χ^2 against the distribution fitted to the real data. In this case, $B = 999$ synthetic datasets with sample size n are drawn from the fitted distribution. Each of the synthetic datasets is then refitted with the new estimated parameters, recomputing the test statistics. Through this, a null distribution of the test statistic is built which remediates bias caused by estimating parameters. Finally, the corrected p-value is estimated by the proportion of the individual bootstrap statistics that are equal or greater than the observed statistic.

A choice of $B = 999$ replications is made, since this provides a p-value resolution of $1/999 \approx 0.001$. This is sufficiently precise to reliably read off at the significance level $\alpha = 0.05$ used in our application. A non-smooth number is deliberately used to prevent a p-value from hitting the exact rejection boundary, which is simply a measure to prevent ambiguity.

Akaike Information Criterion & Bayesian Information Criterion

While the goodness-of-fit tests measure the statistical consistency of the model with the measurement data, two statistical metrics are used for complementary rankings of models: the Akaike Information Criterion (AIC) and Bayesian Information Criterion (BIC). Both of these metrics function analogously, rewarding log-likelihood while discouraging overfitting with different penalty terms. For these two information criteria, a lower score indicates a better model.

The AIC is defined as:

$$\mathrm{AIC} = 2k - 2\ln(\hat{L}) \tag{3.10}$$

where k is the number of free parameters as mentioned in Section 3.3.2 and $\hat{L}$ is the maximised log-likelihood of the distribution. The term $2k$ introduces a penalty for model complexity; a fixed cost linear to the parameter count.
On the other hand, the BIC is defined as:

$$\mathrm{BIC} = k\ln(n) - 2\ln(\hat{L}) \tag{3.11}$$

where n is the number of samples. For this metric, the penalty term is $k\ln(n)$; the cost scales linearly with right to k, and logarithmically with right to sample size. Hence, the BIC is a more conservative metric than the AIC for larger datasets.

The AIC and BIC are included in the model analysis because the p-values provided by the goodness-of-fit tests are insufficient for comparing the models; they only indicate whether H_0 can be rejected for a given α. This in turn means that a higher p-value does not indicate a better model, but only states that there is insufficient evidence for rejecting it. Since the AIC and BIC output a score on a common scale which also penalises model complexity, the different candidate models can be fairly ranked based on performance [5].

RCS measurement data visualisation

For the purpose of visualising the RCS measurement data, histograms in Section 4.2.1 are implemented using the Rice rule for binning, where $k = \lceil 2n^{1/3} \rceil$, with n being the sample size [17]. This rule is applied to obtain a standard and reproducible bin count rather than tuning bin count arbitrarily by hand. Unlike the chi-squared test, visualisation has no constraint on the data point count per bin, so using linear binning is appropriate for visual representations of the data.

3.3.4. Alternative models & evaluation methods

While the selection of models and model evaluation methods cover a wide base, there were still many alternative options to analyse. This is because generally, the literature does not conform to a certain consensus distribution for classifying the drone RCS, and instead, varies between using Swerling models [15, 18], gamma distributions [29], log-normal distributions [9], or entirely different models. Given the time constraints of this project, a targeted choice has been made for concluding with the models described in Section 3.3.

Alternative models

If the project had spanned a longer time, three additional statistical models would have been looked into. Firstly, the Rician distribution describes a target with a single dominant scatterer alongside diffuse clutter returns. This has the same physical implication as Swerling III, yet is implemented through amplitude instead of power distribution, which would prove useful in Section 3.3.5 [33].
Secondly, the Nakagami-m model is a generalised model which approximates the Rician for $m > 1$, with m being the shape parameter that determines the severity of amplitude fluctuation. This provides the model with great flexibility for encompassing the scattering regimes [10].
Lastly, the chi-squared distribution functions quite similarly to the Swerling distributions. Compared to Swerling, the benefit with this model is its ability to have a non-integer degree of freedom, which allows it to capture intermediate scattering characteristics, at the expense of having an additional free parameter.

Alternative evaluation methods

As for the evaluation methods, three additional options were initially considered; the Anderson-Darling (AD) goodness-of-fit test, Cramér–von Mises (CVM) test, and Kullback-Leibler (KL) divergence.

The AD test is a weighted version of the KS test; it assigns greater weight to the distribution's tail deviations, meaning that it has an increased sensitivity to misfits in the tails. This would most likely prove useful for RCS detection performance [2].
Another test, the CvM test, integrates the squared deviation between the empirical CDF and theoretical CDF, which allows it to detect systematic misfits better [29]. The reason for not implementing this test, is due to its resemblance to the already-implemented KS and chi-squared tests; the KS test already describes the stand-alone deviation in CDFs, while chi-squared test indicates frequency discrepancies at bin-level. The CvM does not add entirely unique information about the distributions, thus it was not prioritised.
Lastly, the KL divergence is a measure for the information lost when the fitted distribution approximates the empirical distribution. Essentially, this complements the AIC and BIC, which already provide rankings of the models on a common scale. Hence, the KL divergence was deemed redundant enough to not implement alongside these metrics.

All in all, the chosen models provide sufficient variety for modelling the drone RCS for the purpose of this shorter project. Moreover, these models are tested through both hypothesis-testing as well as performance metric scoring.

3.3.5. RCS modelling for the radar processing subgroup

The definitive RCS model is provided to the radar processing subgroup in order to use the drone's true RCS values in a radar signal simulation. From this, the maximum detection range of the radar can be determined. The simulation includes range-Doppler and range-angle maps, which are 2D representations of range, radial velocity, and target angle. The intent is to model realistic behaviour of a 24 GHz FMCW radar with $N_{rx} = 4$ receive antennas, which are fed into the subgroup's processing pipeline, including Constant False-Alarm Rate (CFAR) and Direction-of-Arrival (DOA) estimations.

Signal model

Drone flights are simulated where the measured beat signal $y[n]$ is represented by the signal model:

$$y[n] = s_{\text{target}}[n] + s_{\text{clutter}}[n] + w[n] \tag{3.12}$$

Here, $s_{\text{target}}[n]$ is the signal formed by the simulated target using the obtained RCS model, $s_{\text{clutter}}[n]$ is the addition of static clutter in the simulation, and $w[n]$ is additive white Gaussian noise (AWGN). These individual components are elaborated upon in the following sections.

Target simulation

The target signal generation requires acquisition of the beat signal to obtain the target range $R(t)$. The beat signal for an FMCW radar is computed through

$$f_b = \frac{2S \cdot R(t)}{c} \tag{3.13}$$

where the chirp slope $S = 0.345\,\text{MHZ}/\mu\text{s}$, and c is the speed of light. The instantaneous range $R(t)$ is computed directly from the simulated trajectory with specified $x(t)$ and $y(t)$, such that $R(t) = \sqrt{x(t)^2 + y(t)^2}$. Moreover, the Doppler frequency f_d is required for synthesising the target in the range-Doppler domain. f_d is acquired through:

$$f_d = \frac{2v_r}{\lambda} \tag{3.14}$$

with $\lambda = c/f_c$, and carrier frequency $f_c = 24\,\text{GHz}$. The radial velocity v_r here is obtained through computing the range gradient and dividing it by the chirp duration as $\nabla R(t)/T_{\text{chirp}}$.

Finally, the received power at each chirp is determined by the radar range equation mentioned in Equation 3.1 through a calibration. Namely, a reference received power is defined for $R_{\text{ref}} = 1\,\text{m}$ & $\sigma_{\text{ref}} = 1\,\text{m}^2$. The received power is then scaled to any range and RCS value accordingly, as:

$$P_r(t) = P_{\text{cal}} \cdot \sigma(t) \cdot \left(\frac{R_{\text{ref}}}{R(t)}\right)^4 \tag{3.15}$$

The RCS $\sigma(t)$ is synthesised through obtained measurement results mentioned in Section 4.2.3. This Swerling III model can be synthesised through its equivalent Gamma distribution as $\sigma(t) \sim \Gamma\big(k = 2,\, \theta = \frac{\bar{\sigma}}{2}\big)$, where $\bar{\sigma} = 0.2257\,\text{m}^2$.

Flight trajectory simulation

Furthermore, flight trajectories were simulated in collaboration with the radar processing subgroup, tailoring to their requirements. This included two independent, simulated drone flights used to assess the radar's capability to detect multiple targets in a noisy, cluttered environment. For simplicity, the final obtained signal is approximated as an addition of the two drone flights: $s_{\text{target}} = s_1 + s_2$. This approximation is likely to decrease in performance or fail when drones' scattered signals interfere with each other through multi-pathing, or more generally because the simulation combines two independent single-target generations, instead of modelling the targets within a single environment. This environment would model simultaneous radar illuminations of both drones that occurs in reality, instead of separate, independent illuminations.

The main trajectory that was used describes the drones moving in range intervals and hovering stationary at set intervals. The first drone's flight consists of a pattern where the drone moves through range intervals spanning from $10 - 70\,\text{m}$, and hovers stationary every $10\,\text{m}$ interval. The transitions from moving to stationary are cosine-interpolated to have a more realistic, smoother transition. The second drone's flight has the same pattern with a lateral offset of $5\,\text{m}$.

Clutter model

To thoroughly test the limits of the radar system in realistic conditions, static clutter from fixed-position scatterers are included in the synthesis as a set of reflectors. These are placed in a matrix with known range R_c, angle θ_c, and RCS σ_c. The resulting clutter component for chirp index i, number of chirps n, and receiver channel m is computed as:

$$s_{\text{clutter}}[i,n,m] = A_c \cos(2\pi f_{b,c} \cdot t_n + \phi_{\text{drift}}(i) + \phi_0 + m\pi \sin\theta_c) \tag{3.16}$$

where A_c is acquired through taking $\sqrt{P_c}$ as $\sigma^2 \sim P$. Moreover, ϕ_0 is a uniform distribution initial phase, and $\phi_{\text{drift}}(i) = \sum_{k=0}^{i} \delta(k)$ with $\delta_k \sim \mathcal{N}(0, 0.01)$. This value is experimentally determined and used to model slight wind breezes that affect the received scatterings and Doppler velocities.

Noise model

Lastly, additive white Gaussian noise $w[n]$ is added to the simulation in order to represent thermal noise at the receiver. The power of this noise is defined relative to the SNR of a reference target fixed at range $R_{\text{ref,SNR}} = 30\,\text{m}$, which establishes the SNR as a parameter independent of the simulated trajectory. Thus, noise power σ_w^2 is formulated as:

$$\sigma_w^2 = \frac{P_{\text{cal}} \cdot \bar{\sigma} \cdot (R_{\text{ref}}/R_{\text{ref,SNR}})^4}{2 \cdot \text{SNR}_{\text{linear}}} \tag{3.17}$$

where $\bar{\sigma} = 0.2257\,\text{m}^2$ is the DJI Air 3S's mean RCS, and $\text{SNR}_{\text{linear}} = 10^{\text{SNR}_{\text{dB}}/10}$. The factor $1/2$ is required to account for the real-valued sampling used by the radar, as noise power is being split across positive and negative frequencies. More specifically, the noise variance per sample is σ_w^2 for real sampling, while for complex sampling this would be $\sigma_w^2/2$.

Final output and further pipeline in the radar processing subgroup

In summary, a signal model $y[n]$ is generated through synthesising a target beat signal $s_{\text{target}}[n]$ in an environment with static clutter $s_{\text{clutter}}[n]$ and additive Gaussian white noise $w[n]$.
This signal model is provided to the radar processing subgroup, which will analyse the results using CFAR algorithms, MVDR beamforming and Direction-of-Arrival angle estimation. They produce a range profile through applying range FFT, and then a Doppler FFT across chirps per block to form the range-Doppler maps.

3.4. Maximal detectable range

The maximal range for the RCS target detection is set by the minimum received power. This power received from the target with a certain RCS must still be distinguishable from the environment. In practice, a desired detection probability P_d and false alarm probability P_FA is chosen, which corresponds to a minimum SNR required.

Deciding on a desired detection probability is the task of the radar processing subgroup. For the sake of this analysis a standard value of 0.9 was chosen for P_d and 10^{-6} for P_FA. Chapter 3.3 of [27] shows that a fluctuating Swerling III target requires a minimum SNR of $17.2\,\text{dB}$ for these P_d and P_FA values. The minimum received power can then be calculated as follows:

$$P_{min} = SNR_{min} k T_0 B_n F \tag{3.18}$$

with k being Boltzmann's constant, T_0 being equal to $290\,\text{K}$, B_n the noise bandwidth contained in a single velocity bin, and F the noise figure of the receiver. Subsequently, the radar range equation was used to determine the maximal range at which the drones are still discernable from the background:

$$R = \left[K \frac{P_t G_t G_r \lambda^2 \sigma_{med}}{(4\pi)^3 P_{min}} \right]^{\frac{1}{4}} \tag{3.19}$$

With σ_{med} being the median RCS of the target.

The factor K was included to account for the effect of the spinning blades of these drones. The power received by the radar is in reality spread out over multiple velocity bins. Only the power contained in the central velocity bin is considered when determining whether a certain region in the range-Doppler maps counts as a detection. Therefore, a range-Doppler map was made of the measurement taken of the drone in the anechoic chamber with its blades spinning. The Doppler FFT was taken over the same number of chirps as the radar processing subgroup. The ratio between the power contained in the zero-velocity bin and the total power across all velocity bins was then used as a correction factor in the radar range equation.

4

Results

4.1. Measurement results

This section describes the measurement results obtained from dataset 1 & 2. First, the different reference and background measurements are analysed, where potential clutter and interference sources are identified. Then, drone configurations are analysed, and corresponding statistics are presented.

4.1.1. Background contributions

Figures 4.1a and 4.1b show the FFT magnitude of the beat signal, plotted against range, for the measurement where no target was present in the anechoic chamber. These results are important for analysing the validity of the drone RCS results. This is because any background source present in the measurements will interfere with the signal originating from the target. This will unavoidably generate an error in the calculated power received by the radar, even with subtraction of this background FFT.

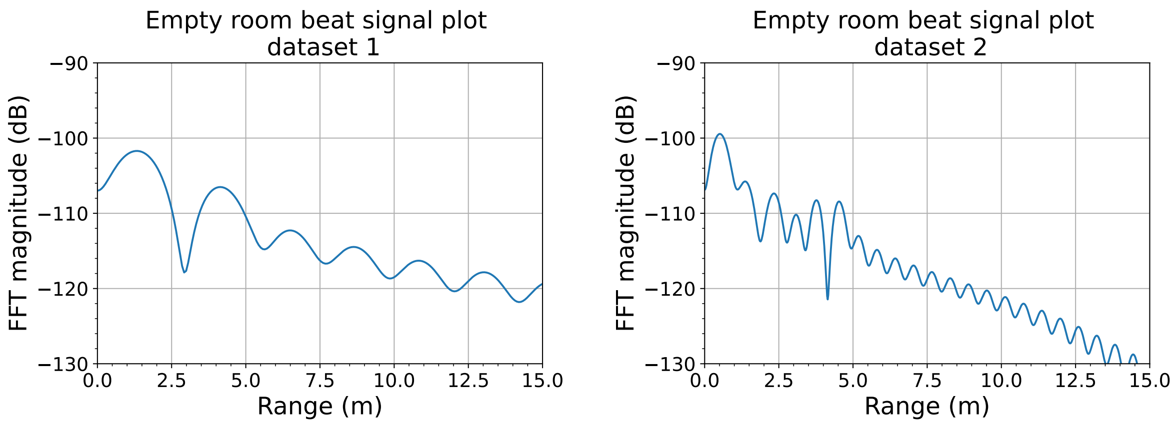

(a) Empty room measurement for dataset 1

(b) Empty room measurement for dataset 2

Figure 4.1: Empty room measurements performed for dataset 1 and 2, with an object or source located at approximately $1\,\mathrm{m}$

It can be seen that its shape is indicative of an object or source located at a distance of approximately $1\,\mathrm{m}$. This signature peak at low range appears in almost all measurements done in the anechoic chamber, and for both datasets. It also appears in the recordings done by the radar processing subgroup. Moreover, it always appears at the same position within each dataset, while shifting in position between the sets. This makes its position dependent on the radar parameters. The conclusion that this source represented the leakage between the transmit and receive antennas of the radar was based off of these observations.

Dataset 2 shows an additional source located near the rotating pillar. This signature appears in all measurements taken for dataset 2 and does not seem to move between measurements. The underlying reason for the presence of this additional scattering source remains unclear.

4.1.2. Calibration results

The beat signal plots for the reference measurements illustrate the utility of considering only a specific range bin of interest. They also show the difference between datasets 1 and 2. The beat signal plots for all three reference objects in both datasets are shown in Figures 4.2-4.4, where each plot is overlaid with the corresponding empty room measurement. The vertical dashed lines indicate the bounds of the range bin used for calculating the received power.

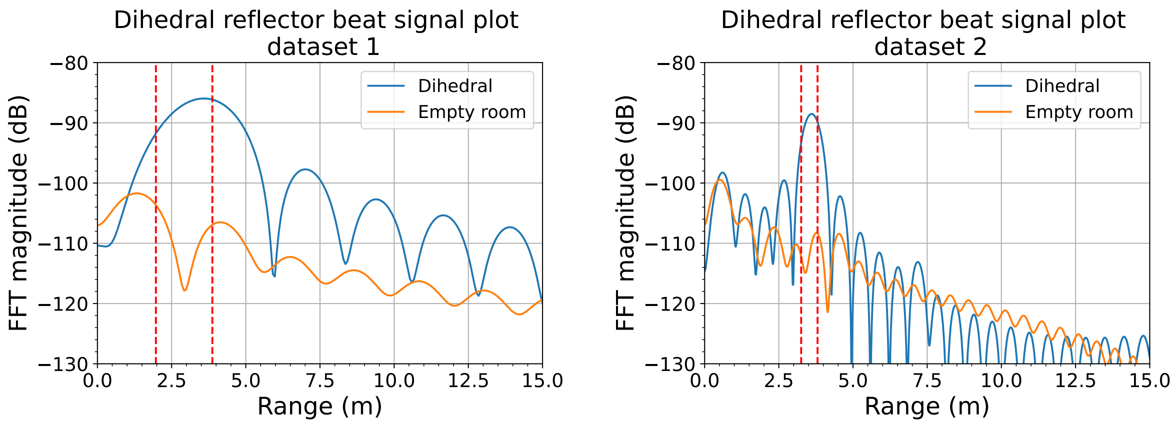

(a) Dihedral reflector measurement for dataset 1. The vertical dashed lines indicate the first range bin

(b) Dihedral reflector measurement for dataset 2. The vertical dashed lines now indicate the sixth range bin

Figure 4.2: Dihedral reflector measurements compared to the background measurement, for dataset 1 and 2, where the blue and orange lines represent the FFT magnitude of the dihedral and empty room measurement, respectively

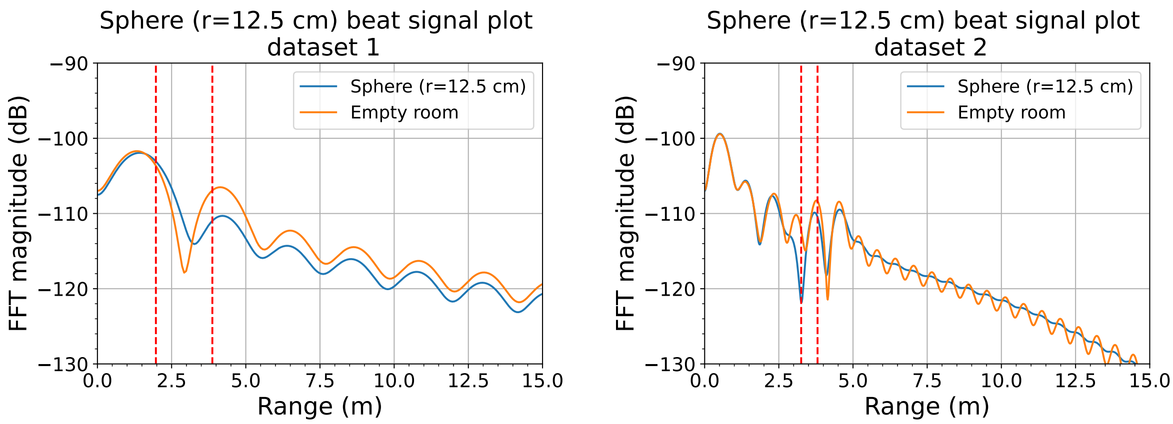

(a) Sphere ($r = 12.5\,\mathrm{cm}$) measurement for dataset 1

(b) Sphere ($r = 12.5\,\mathrm{cm}$) measurement for dataset 2

Figure 4.3: Measurements of the sphere of $r = 12.5\,\mathrm{cm}$ compared to the background measurement, for dataset 1 and 2, where the blue and orange lines represent the FFT magnitude of the sphere and empty room measurement, respectively

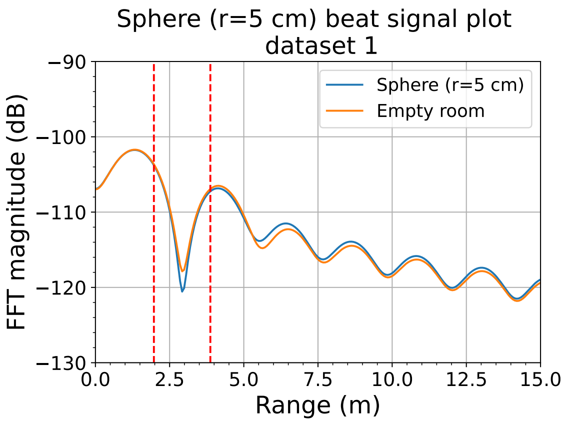

(a) Sphere ($r = 5\,\mathrm{cm}$) measurement for dataset 1

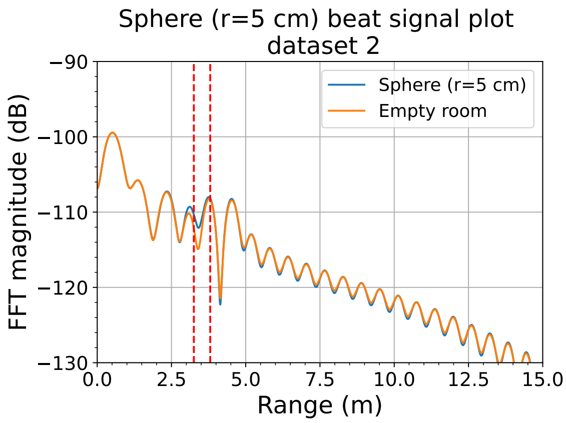

(b) Sphere ($r = 5\,\mathrm{cm}$) measurement for dataset 2

Figure 4.4: Measurements of the sphere of $r = 5\,\mathrm{cm}$ compared to the background measurement, for dataset 1 and 2, where the blue and orange lines represent the FFT magnitude of the sphere and empty room measurement, respectively

For dataset 1, the range bin of interest coincides partly with the main lobe and first sidelobe of the leakage signal. In this region, the scattered wave coming from an object such as the spheres is easily masked by the leakage as can be seen in Figure 4.3. Dataset 2 was collected using a higher bandwidth, and therefore a higher range resolution. This can be seen in the figures corresponding to dataset 2, as all main and sidelobes are much narrower. The signal originating from the target now overlaps with a much weaker leakage sidelobe.

Table 4.1 shows the comparison between the theoretical RCS values from Table 3.4 and the measured RCS values of the three reference objects for both dataset 1 and 2. The calibration gain factor was chosen to be $5.1\,\mathrm{dB}$ for dataset 1 and $7.6\,\mathrm{dB}$ for dataset 2. These gains were derived by aligning the measured RCS values as much as possible with the theoretical values, while prioritising the dihedral reflector and larger sphere.

For both datasets, two of the three reference objects were able to be calibrated within $1\,\mathrm{dB}$, with the third being off by 4 to $5\,\mathrm{dB}$. For dataset 2, the object not aligning with the others is the small sphere. This is to be expected due to its low RCS, making it harder to isolate from the background. On the contrary, for dataset 1, the larger sphere is the object that does not align, which is unexpected. This, together with the target being situated close to the main lobe of the leakage, is the reason the results from dataset 2 were judged to be more reliable than dataset 1 for the statistical analysis.

Table 4.1: Comparison of measured vs theoretical RCS values of reference objects

Reference object	Theoretical RCS (dBsm)	Measured RCS dataset 1 (dBsm)	Error dataset 1 (dBsm)	Measured RCS dataset 2 (dBsm)	Error dataset 2 (dBsm)
Dihedral reflector	21.2	-20.0	-1.2	20.5	-0.7
Sphere ($r = 12.5\,\mathrm{cm}$)	-13.1	-8.6	+4.5	-14.0	-0.9
Sphere ($r = 5\,\mathrm{cm}$)	-21.0	-21.0	0.0	-17.3	+3.7

4.1.3. DJI Air 3S RCS results

With the calibration factor determined for both datasets, the drone measurements were processed. The results for an elevation of $0°$ are shown in Figures 4.5, 4.6 and 4.7 in the form of a polar plot where an azimuth of $0°$ corresponds to the front of the drone.

Dataset 1

The first results shown here represent the RCS profile of the DJI Air 3S according to dataset 1.

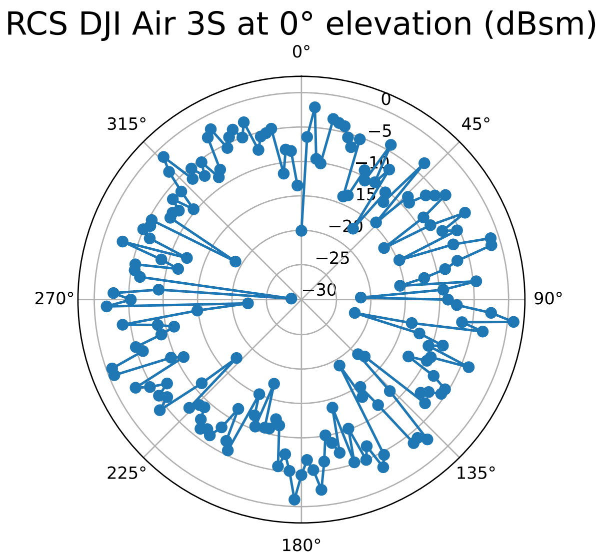

Figure 4.5: Polar radar cross-section plot for the DJI Air 3S at $0°$ elevation, according to dataset 1

The plot illustrates not only the strong dependence of drone RCS on azimuth angle, but also the highly stochastic nature of drone RCS. Slight variations in azimuth can cause fluctuations in the RCS on a scale of multiple order of magnitudes. For this reason, analysing the statistical characteristics of these profiles provides more information. This chapter shows only the plots for elevation angles of $0°$. The other plots can be viewed in Appendix D. The statistical characteristics of the RCS profiles of each elevation angle tested for dataset 1 is shown below in Table 4.2.

Table 4.2: Radar cross-section statistics per elevation angle (dataset 1)

Elevation (deg)	n	**Mean (dBsm)**	**Median (dBsm)**	**Std (m^2)**	**Min (dBsm)**	**Max (dBsm)**	**95th Percentile (dBsm)**
-30	24	-10.07	-10.69	8.350×10^{-2}	-19.87	-5.20	-5.79
-20	24	-8.41	-9.76	1.100×10^{-1}	-19.39	-4.30	-4.70
-10	24	-8.65	-8.72	7.350×10^{-2}	-21.59	-5.53	-5.62
0	180	-6.29	-7.26	1.962×10^{-1}	-28.61	0.89	-2.06
5	24	-9.07	-10.05	9.217×10^{-2}	-16.80	-4.27	-4.67
10	24	-7.62	-7.30	9.080×10^{-2}	-20.35	-4.84	-4.98
15	24	-8.66	-9.32	8.752×10^{-2}	-21.62	-5.10	-5.29
20	24	-8.79	-9.75	8.739×10^{-2}	-17.37	-4.46	-5.26
25	24	-7.69	-7.42	9.436×10^{-2}	-23.28	-3.77	-5.56
30	24	-8.73	-9.49	9.917×10^{-2}	-16.96	-4.64	-4.83

Elevation (deg)	Positive azimuth RCS (dBsm)	Negative azimuth RCS (dBsm)
±15	-1.11	-5.61
±30	-9.40	-6.33
±45	-5.74	-7.79
±60	-4.43	-3.06
±75	-4.28	-6.67
±90	-8.51	-3.84

Table 4.3: Comparison between selected positive and negative azimuth angle for purposes of verifying the symmetry of the RCS profile

Dataset 2

The next results are those of the DJI Air 3S according to dataset 2. These measurements went up only to an azimuth angle of $180°$, as an argument of symmetry was made. This argument can now be scrutinised.

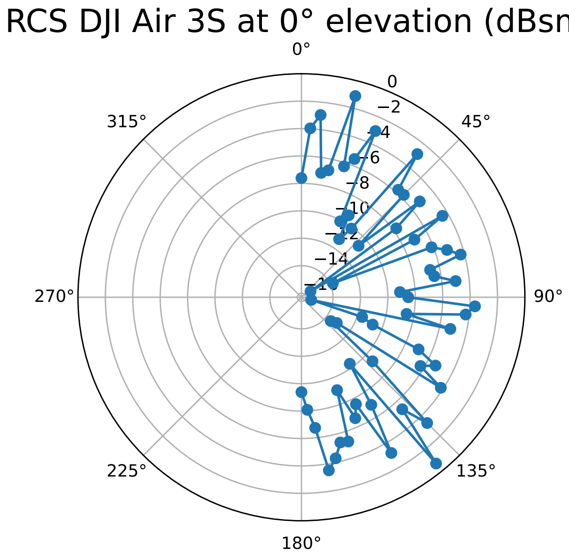

Figure 4.6: Polar radar cross-section plot for the DJI Air 3S at $0°$ elevation, according to dataset 2

Comparing the two half-spans in azimuth of $0° - 180°$ and $180° - 360°$, a deviation of on average $3\,\mathrm{dB}$ is observed between a measurement in one half compared to its mirrored counterpart, as shown in Table 4.3. Given the stochastic nature of the RCS measurements, and a highly-likely misalignment in the propellers' instantaneous blade configurations, the discrepancy cannot be attributed to physical asymmetry of the drone. Therefore, the drone's geometric symmetry along its longitudinal axis is assumed to hold.

The statistical characteristics of this RCS profile, shown in Table 4.4, can now be compared with those of dataset 1. While the results for each elevation are not exactly the same , they are never off by more than $3\,\mathrm{dB}$ when compared to their dataset 1 counterparts. An exception to this being the maximal en minimal RCS results, which are highly prone to outliers caused by the stochastic nature of the RCS.

Table 4.4: Radar cross-section statistics per elevation angle for the DJI Air 3S (dataset 2)

Elevation (deg)	n	Mean (dBsm)	Median (dBsm)	Std (m^2)	Min (dBsm)	Max (dBsm)	95th Percentile (dBsm)
0	61	-5.98	-6.65	1.669×10^{-1}	-15.56	-0.69	-2.93
5	19	-7.27	-7.70	1.252×10^{-1}	-14.77	-3.17	-3.72
10	19	-7.69	-9.13	1.470×10^{-1}	-26.22	-3.16	-3.56
15	19	-5.95	-7.04	1.602×10^{-1}	-12.48	-2.12	-2.49
20	19	-6.42	-6.87	1.584×10^{-1}	-14.99	-2.12	-2.23
23	19	-6.90	-6.99	1.130×10^{-1}	-14.66	-2.90	-4.02

4.1.4. DJI Neo results

Lastly, the results for DJI Neo are shown, for which only dataset 1 contained measurements. Additionally,these measurements only spanned from $0°$ to $90°$ in azimuth

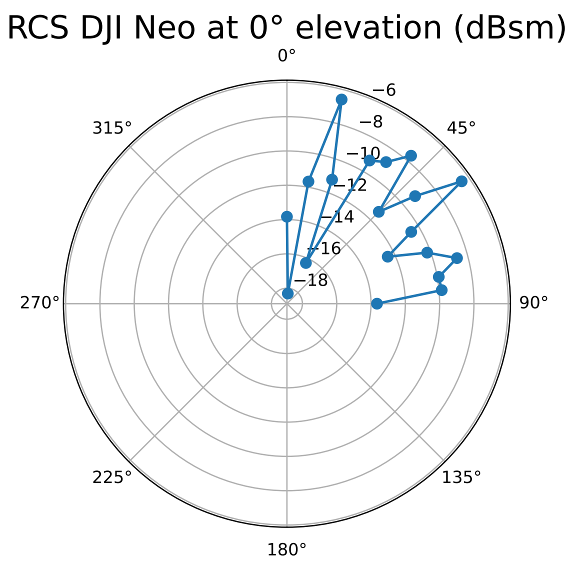

Figure 4.7: Polar radar cross-section plot for the DJI Neo at $0°$ elevation

The statistical characteristics of the RCS profile for the DJI Neo is shown below in Table 4.5. As expected, its RCS results are lower than that of the DJI Air 3S.

Table 4.5: Radar cross-section statistics per elevation angle for DJI Neo

Elevation (deg)	n	Mean (dBsm)	Median (dBsm)	Std (m^2)	Min (dBsm)	Max (dBsm)	95th Percentile (dBsm)
-30	19	-9.38	-10.10	6.534×10^{-2}	-16.53	-5.47	-6.46
0	19	-9.62	-10.20	6.645×10^{-2}	-18.74	-5.88	-5.90
30	19	-10.22	-10.50	5.727×10^{-2}	-21.20	-6.35	-7.35

4.2. DJI Air 3S modelling results

This section describes the modelling results obtained for the RCS profiles of the DJI Air 3S. More specifically, RCS distributions, statistics, and model fits are provided and analysed.

4.2.1. Dataset selection & elevation pooling

The second dataset described in Section 3.1 is used for modelling the DJI Air 3S. This is due to the improved Tx-Rx isolation compared to the first dataset, mentioned in Section 4.1.2. The trade-off with regards to this choice is the significantly lower sample size and no negative elevations being measured; dataset 1 has sample size $n = 387$ while dataset 2 has $n = 156$.
Inspecting the RCS distribution per measurement setup, the following plots are obtained:

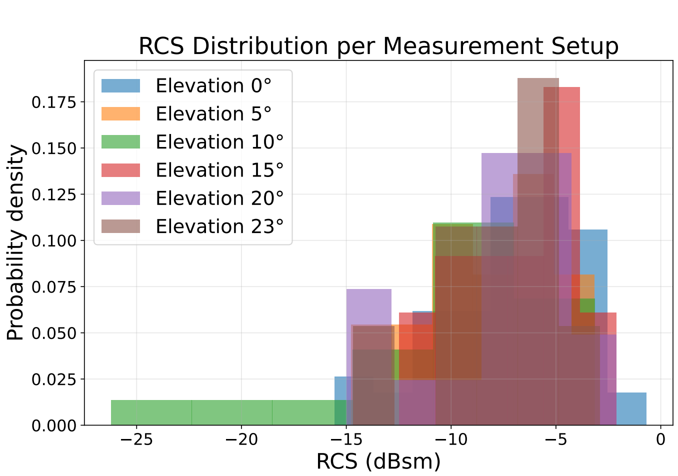

(a) Estimated probability density of the measured RCS values for each measurement setup, with histogram bins following Rice's rule

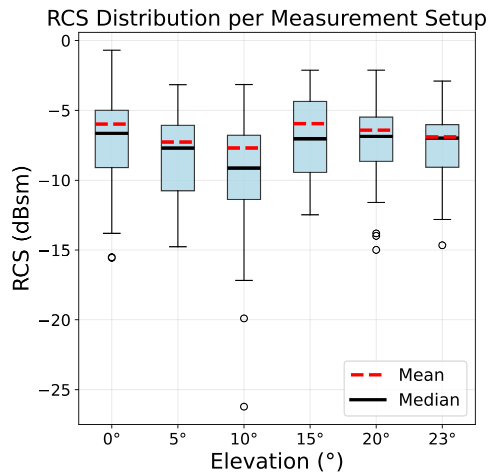

(b) Boxplot comparison of the RCS distributions across all measurement setups, with $\Delta\mu > 2\text{dB}$

Figure 4.8: Distribution of measured DJI Air 3S RCS values for the different measurement setups. The histogram and boxplots illustrate substantial overlap between elevations, with mean differences below 2dB, and noteworthy outliers at $10°$

Here, it can be observed that the spread of RCS values within an elevation is comparable to the spread between elevations. This is further confirmed in Table 4.4, where it can be seen that the mean varies from -5.95dBsm to -7.69dBsm, or less than 2dB, while the IQR spans $3 - 5\text{dB}$ across all elevations.

It is also observed that the $10°$ elevation data has significant outliers reaching -20 and -26.22dBsm. These two points are likely caused by drone configurations where cancellations take place due to near-field effects, resulting in these deeper fade outliers in RCS.

Regardless, the RCS behaviour remains consistent throughout the rest of the elevation configurations, with statistical tests not indicating a noteworthy deviation. Thus, elevation angle can be treated as a non-significant factor and pooling all measurements is justified. This justification can be explained due to the drone geometry changing insufficiently over $0° - 23°$ elevation to notice significant changes, unlike with the azimuth sweeps. Hence, the different configurations are pooled for modelling the RCS.

4.2.2. Candidate model analysis & performance

After pooling the different measurement configurations into one distribution, the candidate models explained in Section 3.3 can be analysed. This results in the following distribution:

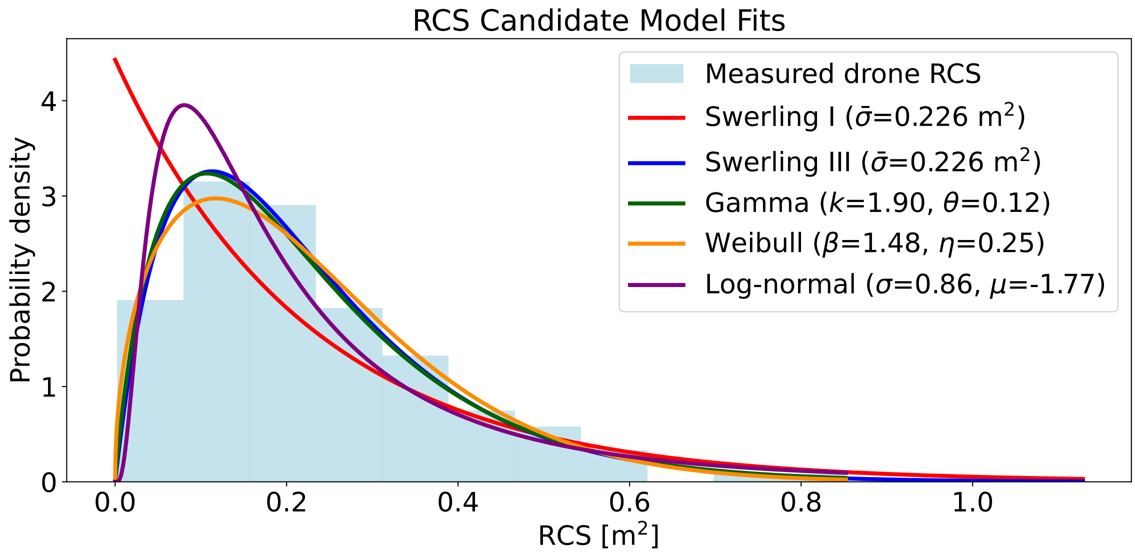

Figure 4.9: RCS measurement distribution with the five candidate models, where their respective parameters mentioned in Section 3.3 are specified. It seems that the Swerling III and gamma models fit the data the best, followed by Weibull, and that Swerling I and log-normal significantly deviate from the data distribution

Here, it can be observed that Swerling I, using the average RCS value of $\overline{\sigma} = 0.226\mathrm{m}^2$ misses the general shape of the data due to its simplicity, where too much probability mass is present near zero. The log-normal distribution with parameters $\sigma = 0.86$ and $\mu = -1.77$ captures the general shape of the data, yet overshoots its peak significantly. Moreover, Weibull ($\beta = 1.48, \eta = 0.25$) only slightly undershoots the peak. Lastly, the gamma distribution ($k = 1.90, \theta = 0.12$) and Swerling III model ($\overline{\sigma} = 0.226\mathrm{m}^2$) both seem to fit the model precisely, capturing its peak as well as the tail behaviour.

The performance of the candidate models are also measured using the goodness-of-fit testing and information criteria metrics described in Section 3.3.3. The results of these tests are as follows:

Table 4.6: Goodness-of-fit and information criteria results for RCS distribution models

Model	p_{KS}	p_{χ^2}	AIC	BIC
Swerling I	$0.0013^{\dagger}$	$0.0089^{\dagger}$	-150.36	-147.32
Swerling III	0.653	0.915	$-181.85^{\ddagger}$	$-178.80^{\ddagger}$
Gamma	0.160	0.940	-180.09	-173.99
Weibull	0.693	0.983	-181.55	-175.45
Log-normal	$< 0.001^{\dagger}$	0.227	-152.86	-146.76

† Reject H_0 at $\alpha = 0.05$. ‡ Lowest AIC or BIC.

Here, it can be observed that Swerling I is rejected by both the KS and chi-squared tests under $\alpha = 0.05$, and performs the worst on the AIC and BIC along with the log-normal. The Swerling I model's performance is expected to fail due to its oversimplification of the measurement data. The log-normal, however, gets overwhelmingly rejected by the KS test, while the chi-squared test fails to reject. Since the log-normal fit both overshoots the peak significantly, and has this peak also slightly shifted to the left, probability mass accumulated too significantly, skyrocketing the maximum CDF deviance mentioned in Equation 3.8. On the other hand, the chi-squared test that works through binning fails to observe this issue, and as a result fails to reject.

The gamma, Weibull, and Swerling III models unanimously fail to be rejected, so a choice will be made through the information criteria and specific model properties in the following section.

4.2.3. Final DJI Air 3S model choice

Even though the gamma distribution seemed to be a better fit in Figure 4.9, it scores lower than the Weibull distribution, and considering they serve a similar purpose, the gamma distribution will not be chosen. Furthermore, while the Swerling III scores slightly better than the Weibull distribution (-181.85 and -178.80 compared to -181.55 and -175.45), they are still very competitive. For this reason, the final choice between these two models is made based on model characteristics.

The Weibull distribution has an additional shape parameter that increases its flexibility for modelling RCS fluctuation behaviours, which is less desirable due to the low sample size of $n = 156$. Additionally, it fails to describe the target as well as the Swerling III model; Swerling III implicitly categorises the drone target as a single dominant scatterer with surrounding smaller scatterers, as mentioned in Section 3.3.1. Another advantage of Swerling III is its rigidity due to having one fewer free parameter, which reduces the risk of fitting onto noise artefacts. Thus, Swerling III is chosen as a definitive model for modelling the DJI Air 3S drone at a bandwidth of $B = 20.000 - 20.250\,\mathrm{GHz}$.

4.3. DJI Neo modelling results

In a similar manner as the DJI Air 3S, the modelling results for the DJI Neo are described in this section. The first dataset is used, since no recordings of the DJI Neo were made in the second dataset.

4.3.1. Configuration difference analysis

To analyse whether the pooling argument holds, the RCS distributions of the three elevation setups are compared in the following figure:

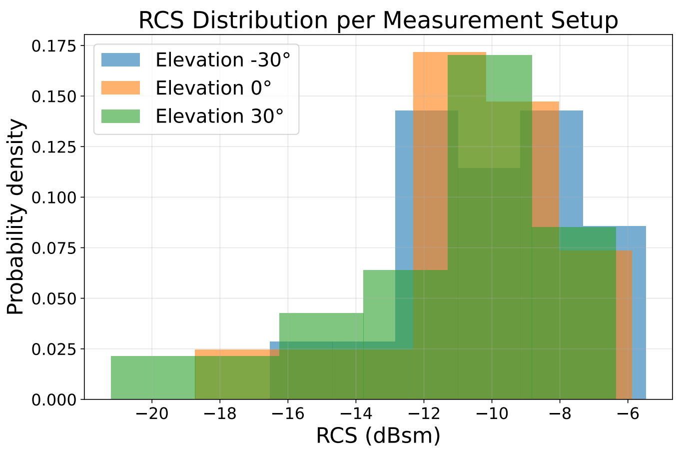

(a) Estimated probability density of the measured RCS values for each measurement setup, with histogram bins following Rice's rule

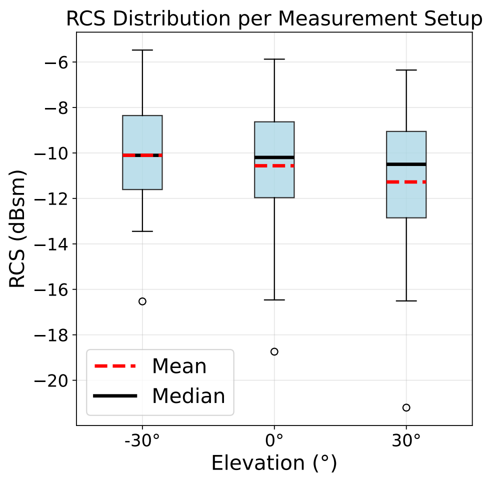

(b) Boxplot comparison of the RCS distributions across all measurement setups, with $\Delta\mu > 1\mathrm{dB}$

Figure 4.10: Distribution of measured DJI Neo RCS values for the different measurement setups. The histogram and boxplots illustrate substantial overlap between elevations, with mean differences below $2\mathrm{dB}$, and noteworthy outliers at $10°$

Again, the different configurations are consistent across the three elevation angles. As listed in Table 4.5, the mean between configurations varies from -9.38 to $-10.22\,\mathrm{dBsm}$, equating to a spread below $1\,\mathrm{dB}$. Variability over elevation is also consistently around $0.06\,\mathrm{m}^2$. Conclusively, the same pooling argument can be made for modelling the DJI Neo.

4.3.2. Candidate model analysis & performance

The five candidate models are fitted to the pooled distribution of the DJI Neo RCS measurements and shown in the following figure:

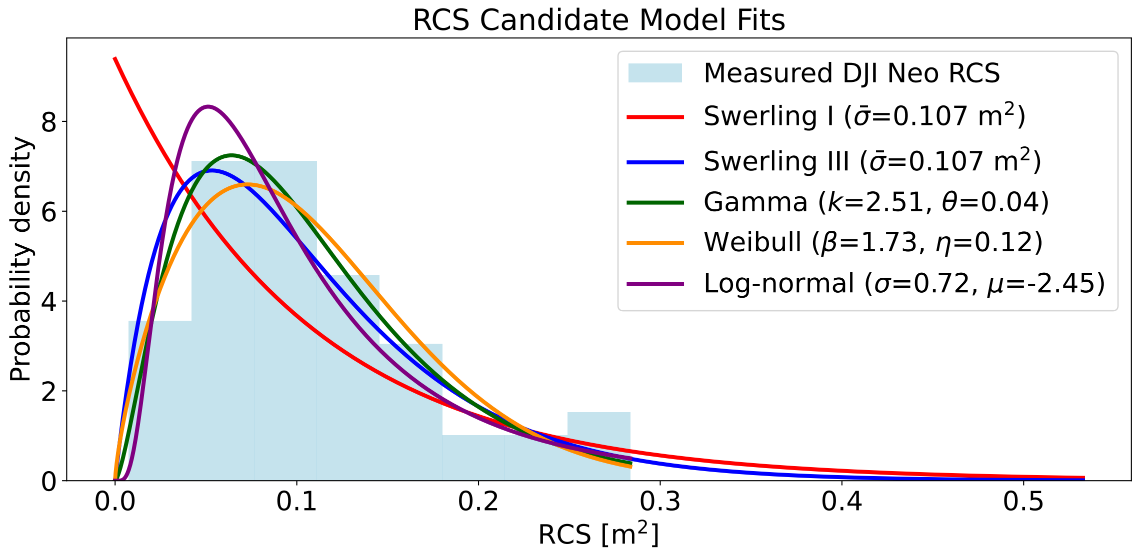

Figure 4.11: DJI Neo RCS measurement distribution with the five candidate models, where their respective parameters mentioned in Section 3.3 are specified. Similar to Figure 4.9, Swerling I and log-normal seem to miss the general outline of the distribution, while the other three distributions generally capture this shape. The limited sample size of $n = 57$ seems to introduce irregularities, such as a smeared peak and bump at the end of the right-side tail

The visual results of this figure are comparable with that of the DJI Air 3S. Swerling I fails to describe the general shape of the distribution, and log-normal overshoots the peak significantly while also shifting leftward. Swerling III, gamma, and Weibull, however, all capture the outline of the peak and tails pretty well, though the limited sample size of $n = 57$ limits the resolution of the distribution. More specifically, the histogram lacks resolution to evidently localise the modal peak, this peak seems to have been smeared out over two bins. There also seems to be an increased bump at the end of the right-end tail. These would have likely smoothened out by increasing resolution in azimuth or intermediate elevations, and by using the radar parameters of dataset 2.

The goodness-of-fit and information criteria results are shown in the following table:

Table 4.7: Goodness-of-fit and information criteria results for the DJI Neo RCS distribution models

Model	p_{KS}	p_{χ^2}	**AIC**	**BIC**
Swerling I	$0.004^{\dagger}$	$0.0046^{\dagger}$	-139.28	-137.24
Swerling III	0.519	0.651	-159.12	$-157.07^{\ddagger}$
Gamma	0.801	0.874	-158.67	-154.58
Weibull	0.834	0.787	$-159.31^{\ddagger}$	-155.22
Log-normal	0.166	0.667	-150.99	-146.90

† Reject H_0 at $\alpha = 0.05$. ‡ Lowest AIC or BIC.

Swerling I is rejected by both the KS and chi-squared test at $\alpha = 0.05$, and scores the worst on the two information criteria, verifying its inadequacy. On the other hand, log-normal is not rejected by both tests, while visually it distinctly overshoots. The test results can be attributed to the lower sample size not fully encapsulating the deviations this model has with regards to the RCS distribution. Evidently, its AIC and BIC are the second lowest of the candidates with a noticeable margin.

The remaining models, Swerling III, gamma, and Weibull, all pass the goodness-of-fit tests. These models have a low difference in AIC scores of below 0.7, and BIC scores differ below 2.5. There is also a conflict; Weibull performs best on the AIC and Swerling III ranks first in the BIC, due to having one less free parameter. This implies that these statistical comparison methods by themselves are not pronounced enough to fully motivate a model choice. Therefore, to also be in line with Section 4.2.3, the distinction will be made with regards to model characteristics.

4.3.3. Final DJI Neo model choice

Provided Swerling III, gamma, and Weibull are near-indistinguishable at this sample size using the statistics, the definitive model choice is made through physical motivation and model characteristics.

Swerling III implies the target is rendered as a dominant scatterer with smaller, adjacent ones. This physically described the DJI Air 3S well, given its main body and distinctly protruding propeller arms. However, the DJI Neo does not distinguishably have a main-scatterer body and adjacent-scatterer propeller geometry; it is much more compact and integrated into a meshed shape. Therefore, Swerling III lacks its physical grounding for the DJI Neo, as compared to the DJI Air 3S, and is excluded.

Both gamma and Weibull have two free parameters, and similar AIC and BIC scores. Weibull has a slightly lower score on the AIC (-159.31 vs. -158.67) as well as the BIC (-155.22 vs. -154.58), and an advantage with regards to parameter characteristics. Its shape parameter β provides great insights into the shape and tail behaviour of the RCS distribution, which means this parameter is a physically interpretable measure of fluctuation characteristics for comparison [32]. Thus, the Weibull distribution is selected as the RCS model for the DJI Neo at a bandwidth of $B = 20.000 - 20.070\,\mathrm{GHz}$.

4.3.4. RCS modelling for the radar processing subgroup results

The RCS model of the DJI Air 3S was successfully integrated into a radar signal simulation. Working in collaboration with the radar processing subgroup, targets were shown to be detectable for ranges that are consistent with expectations given the provided $\bar{\sigma} = 0.2257\,\mathrm{m}^2$. Here, a realistic SNR value of $-20\,\mathrm{dB}$ was used provided by the radar processing subgroup.

Figure 4.12 displays the received power across range for the two simulated drone flights. They follow the R^{-4} decay rate of the radar equation, and display a spread in power. This spread around the general trend-line indicates the RCS fluctuations from the fitted distribution $\Gamma(k = 2, \theta = \bar{\sigma}/2)$ mentioned in Section 3.3.5.

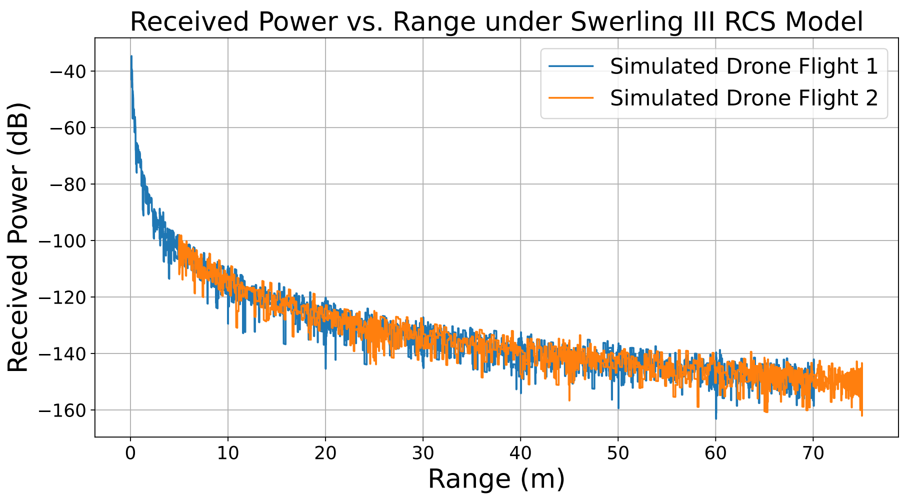

Figure 4.12: The received power of two simulated drone flights as a function of range using the fitted Swerling III RCS model mentioned in Section 3.3.5. The R^{-4} decay is clearly visible, implying a correct implementation of the radar equation. The spread across the general trend-line reflects the per-block RCS fluctuations from the fitted distribution

The RCS variance is shown to be non-negligible with a spread of approximately $10\,\mathrm{dB}$, which signifies the importance of modelling the target as a Swerling **III** distribution, compared to using a fixed σ-value. On the other hand, the radar processing subgroup identified the detection performance to be overly optimistic, which is likely caused by the primitive static clutter model consisting of a sparse set of fixed-position scatterers with slight phase drift. Nevertheless, the concept of radar signal simulation using the DJI Air 3S drone's RCS model has been demonstrated, with room for improving simulation aspects such as the clutter modelling.

4.4. Doppler spread and maximal range results

Finally, the effect of rotating propeller blades on the power contained within the zero-velocity bin was examined. Figure 4.13a shows the velocity plot for the measurement where the drone blades were set to spin while Figure 4.13b shows the velocity plot for a static drone without rotating blades. This second plot is made using the $0°$ elevation and azimuth measurement of the DJI Air 3S. Both plots represent the cross-section of their corresponding Range-Doppler spectra at a range of $3.7\,\mathrm{m}$.

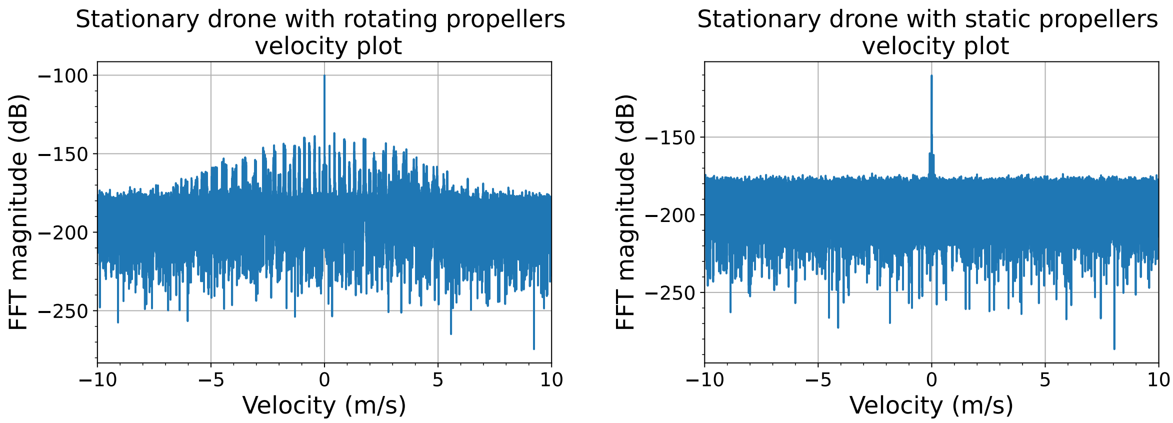

(a) FFT magnitude vs. velocity plot for stationary DJI Air 3S with rotating propellers

(b) FFT magnitude vs. velocity plot for stationary DJI Air 3S with static propellers

Figure 4.13: A comparison of FFT magnitude vs. velocity plots of rotating and static propellers

The power, originally stored in the zero-velocity bin, has now become spread out over multiple velocity bins. Applying the same power calculation as described in Section 3.2, namely subtracting the background measurement and range gating, the power held by each velocity bin was calculated and summed across all velocity bins to get the total received power. The ratio between this total power and the power contained only in the zero-velocity bin was then found to be 98.5%. Which indicates a small fraction of power is spread across velocity bins. With a median RCS of -$9\,\mathrm{dB}$, the resulting maximal range calculated was $162\,\mathrm{m}$.

5 Discussion

This chapter covers the discussion of the validity of the RCS results which were presented in the previous chapter and an evaluation of the methods used throughout this project.

5.1. Verification of RCS results

As mentioned in Chapter 3, verifying the acquired RCS profiles is non-trivial. However, it is possible to assess the plausibility and/or reliability of these results by comparing them with the RCS profiles of comparable drones and analysing all the sources of error present in the calculation. In absence of an electromagnetic model to aid this assessment, an analysis is given of the validity of the commonly used thin-wire approximation for a frequency of 24 GHz.

5.1.1. Drone RCS comparison

Table 5.1 shows the RCS properties of rotary blade drones that have been studied in relevant literature [31].

Table 5.1: RCS statistical characteristics of comparable drone models at 26 GHz [31]

Drone model	Mean (dBsm)	Std (dBsm)	Max (dBsm)
DJI Matrice M100	-10.5	6.8	15.8
Walkera Voyager 4	-12.6	6.1	8.1
DJI Mavic Pro	-16.8	6.1	2.3
DJI F450	-17.1	5.9	3.3
Parrot AR.Drone	-19.5	6.2	4.6

The DJI Matrice M100, Walkera Voyager 4, DJI Mavic Pro, and DJI F450 all resemble the DJI Air 3S in either size, material, appearance, or a combination. The Parrot AR.Drone has a high resemblance to the DJI Neo. While these results may seem to indicate the wide range of plausible RCS characteristics a drone can have, in reality it is a testiment to the dependence of RCS on target material. The DJI Matrice M100 and Walkera Voyager 4 being made of carbon-fiber, while the DJI Mavic Pro and DJI F450 are made of plastic. With the DJI Air 3S also being made of plastic, a lower RCS profile would be the expectation based off of other drone models. However, the results for the DJI Air 3S presented in the previous chapter tend to be on the higher end, with a notable mean RCS of around 5- $7\,\mathrm{dB}$ for both datasets at an elevation of $0°$. As mentioned in the introduction, there is little research done on other elevation angles at this frequency, making it infeasible to compare these results with those of other drones.

5.1.2. Sources of error

While comparing the results with other drone models can give an indication of their plausibility, the RCS of a three-dimensional body is still highly dependent on its shape. This makes a comparison with other drones, no matter how many characteristics they share in common with the DJI Air 3S or DJI Neo, an indication of inaccuracy in the measured RCS at best. For this reason, the best way to analyse the accuracy of the results presented in this thesis, is to critically evaluate the sources of error inherent to the devised measurement setup and methodology.

Measurement setup

The introduction briefly mentioned the fact that the standard way of conducting indoor RCS measurements, is using a VNA [10]. The benefits of using a VNA as transmitter and receiver are, for one, that a VNA has a much higher accuracy when it comes to knowing its own gains and losses. After all it is an instrumentation device. Secondly, it allows for a choice of which type of antenna to use to best suit the measurements being carried out, allowing for the transmitted waves to be more focussed, reducing the effects of multipath propagation and clutter.

Additionally, an FMCW radar necessarily contains a mixing step, where the signal received by the receive antenna is multiplied with the transmitted signal. This operation, even though it is modelled by the EV-Tinyrad24G user guide [1] as a constant gain, will result in a power gain that depends on the signal received by the antenna. Modelling it as a constant gain is therefore an approximation, which necessarily introduces an error that would not have existed if a VNA was used. Quantifying this error, however, is difficult, which is why it is merely mentioned as a source of uncertainty in the accuracy assessment of the measurements.

The reason for not utilising a VNA during our measurements is, like most decisions made regarding the measurement setup, the substantial lack of time that was available for preparing and conducting the measurements themselves. This is due to the anechoic chamber being reserved for a different project which ran parallel with this one.

The effect of near-field measurements

The Fraunhofer-criteria states that the condition which has to hold for a wave to be in the far-field region is that the distance from the source $z \gg D^2/\lambda$, where D is the largest dimension of the reflecting or transmitting object.

For the waves travelling from the radar to the target, D is the $48\,\mathrm{mm}$ length of the EV-Tinyrad24G's patch transmit antenna. With a wavelength of $12.5\,\mathrm{mm}$, the ratio D^2/λ equals $0.184\,\mathrm{m}$. This makes the distance of $3.7\,\mathrm{m}$ between the target and radar approximately 20x as large as this ratio. Thus, the waves propagating from the radar to the target can be considered plane waves by the time they reach the target. For the waves scattering off the target this is observed to not be the case. Taking the $325\,\mathrm{mm}$ width of the DJI Air 3S as D, the resulting ratio D^2/λ equals $8.45\,\mathrm{m}$, which is well beyond the length of the anechoic chamber.

This is quite different to the situation described in studies utilising compact-range measurement setups, such as [10]. A compact-range setup involves focusing the spherical wavefronts originating from the transmit antenna into plane waves. For the setup described in this project, however, these waves pass the Fraunhofer-criteria. Therefore, a compact range setup as described in [10] would not be of much use here.

Where the impact of the nature of these spherical reflected waves was felt most was in the usage of the different Rx channels of the radar. The spherical nature of the received waves causes the phase of all interfering scattering sources on the target and/or background to vary between Rx channels. This causes the beat signals corresponding to each Rx channel to vary in frequency content. Thereby achieving different performances when the methods devised in Chapter 3 are applied to them.

Compensating for these near-field effects is far beyond the scope of this project due to its complexity. In absence of such a compensation, it can be stated that due to the definition of RCS relying on the reception of plane waves, any RCS measurements performed without accounting for these near-field effect return merely approximations of the RCS. Nonetheless, the near-field effects were seen as a source of error and included in the calibration gain. For the dihedral reflector of dataset 1, taking into account the loss due to range-gating, there seemed to be no additional loss associated to any near-field effects. However, for dataset 2, an additional loss of approximately $3\,\mathrm{dB}$ had to be included in the calibration gain factor on top of the range-gating loss. It is possible this could have been an error introduced by the spherical nature of the received wave.

Range-gating
Another source of error is the assumption that the power lost by range-gating the signal can be reversed with a static gain for all signals. This is clearly incorrect, as different targets will have their power be distributed differently across range bins. This could be seen for all reference objects that were tested. Table 5.2 shows the calculated received power of the reference measurements with and without range-gating for dataset 1 to illustrate this. These calculations were done with background subtraction and without calibration and system gain factor.

Table 5.2: Received power calculated with and without range-gating (dataset 1)

Reference objects	**Received power without range-gating (dB)**	**Received power with range-gating (dB)**
Dihedral reflector	-53.2	-58.6
Sphere ($r = 12.5\,\mathrm{cm}$)	-78.2	-87.2
Sphere ($r = 5\,\mathrm{cm}$)	-80.9	-99.6

As is evident, the lost power varies greatly between the objects. The same analysis is not included for the reference measurements of dataset 2 for the sake of brevity, as they follow the same pattern. Having said this, the extent to which range-gating negatively impacts the reliability of RCS measurements on its own can be disputed. Range-gating or band-limiting appears to be a standard practice in the literature, appearing in [10].

Background sources
It is tempting to assume the presence of background sources can be attributed to a failure in properly setting up the measurement setup in the anechoic chamber. In reality, the presence of background sources like the Tx-Rx leakage is unavoidable when using a continuous wave radar. Subtracting background power is in fact a technique applied in the literature, as is the case in [10]. Here the background power is subtracted vectorially. For the measurements taken for this project, however, subtracting background power vectorially caused the weaker scatterers to gain a large amount of power. Therefore, the background was subtracted in magnitude. This way of subtracting background contributions is an approximation as it disregards the phase of both signals entirely. It remains unclear why doing a vector subtraction caused such a significant error in the results.

5.1.3. Validity of thin-wire approximation

While electromagnetic modelling of drone propellers in the UHF band has been studied extensively [4], the research on modelling for higher frequency bands is sparse. A common method used to model drone propellers in the UHF band is the use of the thin-wire approximation.

The thin-wire approximation consists of reducing the three-dimensional drone blades to a network of thin current-carrying wires. According to [7], there are four main assumptions central to the thin-wire approximation, quoted directly:

- Transverse currents can be neglected relative to axial currents on the wire.
- The circumferential variation in the axial current can be neglected.
- The current can be represented by a filament on the wire axis.
- The boundary condition on the electric field need be enforced in the axial direction only.

For these assumptions to hold, that the radius of the wires must be significantly smaller than both the wavelength transmitted by the radar and the length of the wire [7]. To map the drone blade to a thin wire, it can be approximated as an elliptical cylinder with a and b as semi-minor and semi-major axis. The mapping $r = (a + b)/2$, discussed in chapter 9 of [4], can then be used to find the radius of the corresponding thin wire. Taking a to be half the propellers' thickness gives a length of $1\,\mathrm{mm}$. Taking b to be half the propellers' chord length gives a length of $9\,\mathrm{mm}$. Mapping this to the thin wire results in a radius of $5\,\mathrm{mm}$. Given a wavelength of $12.5\,\mathrm{mm}$, the thin wire radius is two-fifths of the wavelength. With this the conclusion is drawn that the thin-wire approximation does not hold at a frequency of 24 GHz or higher.

5.2. Model choice interpretation

For the DJI Air 3S, the two best-performing models (Swerling III and Weibull) are almost indistinguishable with regards to statistics, as seen in the AIC gap of 0.3 and BIC span of 3.35 with $n = 156$. The argument for choosing Swerling III over Weibull was therefore made on grounds of the drone's geometry and model flexibility, rather than a decisive statistical margin.

The model choice for the DJI Neo uses an even smaller sample size of $n = 57$, which further abrades the ability of AIC and BIC to provide decisive, robust scores. This is reflected in the conflict of test results mentioned in Section 4.3.2. The deciding factor for this drone was reversed compared to the DJI Air 3S, where this time the choice favoured Weibull's flexibility over Swerling III's target scattering implications. While this may seem like an inconsistency in selection approaches, it reflects how the drone's specific geometry is considered, and that the model families have differences in applications highly relevant for consideration in the final choice.

Furthermore, the log-normal distribution is reported in literature to be a consistent high-performer for modelling commercial drone RCS [5, 10], yet it performs as one of the two worst models for both drones. While this could simply be attributed to different scattering behaviours caused by drone model geometries, it could also be a characteristic of the difference in frequency band between this work and the cited literature. The relative size of the wavelength compared to the drone might distort the peak shape of the plot such that it becomes less pronounced than in lower-frequency studies, being a potential explanation of the leftward shift and overshoot in both drone RCS distributions.

Overall, there are certain limitations that apply to the modelling results. The most apparent one is sample sizes, where performance tests do not show strong discriminations between different models that all visually fit the data well, especially for the DJI Neo. There is also a systematic uncertainty with regards to calibration. The calibration of the datasets still contains slight errors, even after prioritising two of three calibration targets and using the mentioned calibration techniques. This propagates into fitted parameters of the models as well as the average RCS. However, this error most likely affects all models approximately equally, implying AIC and BIC metrics are not distinctly sensitive to this error. Instead, it affects the parameter values more than the relative comparison of models.

Moreover, the pooling of elevation setups was only validated over a span of $0° - 23°$ for dataset 2, and for three discrete configurations of the DJI Neo. Different elevations, particularly speculated as you approach $\pm 45\,°$ and beyond, are likely to conflict with this argument, as the drone geometry starts to significantly alter. In such cases, elevation angle would have to be treated as an explicit factor in the model, instead of pooling configurations. Two approaches can be taken to capture effects of elevation changes; either a distribution can be fit per elevation for obtaining a list of model parameters, or the parameters can be modelled as a function of elevation angle.

5.3. Maximum range estimation

The fraction of the total power that was spread out across neighbouring velocity bins was calculated to be 1.5%. While being an extremely small fraction, this result is actually expected based on the fact that the propellers of the DJI Air 3S are made of a plastic composite. Plastic propellers reflect up to $20\,\mathrm{dB}$ less power than metal or carbon propellers [28]. This validates the comparatively low signal power present in the adjacent velocity bins, as these are reflections originating from the spinning blades of the drone.

The calculated maximum range can be compared to the empirical findings of the radar processing subgroup. By taking measurements of a drone flying away from them, they discovered the drone was no longer detectable past $50\,\mathrm{m}$. While it is true that this discrepancy is partially explained by the fact that they use a different desired probability of detection, it is also worth noting that with a more standard median drone RCS of -$20\,\mathrm{dB}$, the calculated maximum range was around $80\,\mathrm{m}$. This is another indication that the measured RCS results for the DJI Air 3S (and DJI Neo) are likely higher than they should realistically be.

6 Conclusions

The aim of this thesis was to characterise the radar cross-section of two consumer drones: the DJI Air 3S and DJI Neo. To achieve this, anechoic chamber measurements were gathered, where statistical models were fitted to the resulting distributions. Using a selected drone model, integration into a radar signal simulation for the radar processing subgroup was performed. Furthermore, an ethical assessment and analysis of the project and its implications is provided in Appendix A.

6.1. Anechoic chamber measurements

Two measurement sets were performed in the anechoic chamber. The first set configured a bandwidth $B = 70\,\mathrm{MHz}$, while the second set used $B = 250\,\mathrm{MHz}$ with an increase in individual measurement duration. The second dataset yielded a significantly improved transmit-receive (Tx-Rx) isolation. This signifies that the bandwidth is a critical factor in measurement reliability for the close-range anechoic setup used. Calibration using a dihedral reflector and two spheres achieved a deviation within $1\,\mathrm{dB}$ for two out of three objects in both datasets, with the third object deviating $4 - 5\,\mathrm{dB}$.

Both drones were measured across a range of elevations. For dataset 1 and 2, the DJI Air 3S was measured across $\pm 30°$ and $0° - 23°$ in elevation, and the DJI Neo was measured across $\pm 30°$ elevation in dataset 1. This was proven to be a non-significant factor in analysing the RCS distribution. The mean RCS values between elevations for a drone varied below $2\,\mathrm{dB}$, while within-elevation values spanned $3 - 5\,\mathrm{dB}$. This result implied pooling the different elevations into one combined RCS distribution is a reasonable operation, simplifying the analysis. Similarly, different azimuthal sweeps were performed, which characterise the distribution. Since the drone is geometrically symmetric, an argument was made to limit the sweep to one side of the drone. This was further limited to one quadrant to maximise the use of the limited available measurement time.

6.2. Modelling and model evaluation

Using the pooled RCS distributions, five candidate models were evaluated using the KS and chi-squared goodness-of-fit tests with parametric bootstrap correction, along with AIC and BIC metrics. Both Swerling I and log-normal were consistently rejected by the tests, or ranked lowest on the metrics for both drones. Swerling III was selected for the DJI Air 3S over the Weibull distribution due to a physical argument: the drone consists of a dominant body with protruding propeller arms, serving the Swerling III scattering model's nature of describing a dominant scatterer surrounded by adjacent, smaller scatterers.
The DJI Neo, however, does not have this physical grounding due to its compact, meshed geometry. For this drone, Weibull was selected, as its shape parameter β provides a physically meaningful and interpretable measure of the RCS fluctuation, especially regarding its tail behaviour. All in all, the statistical performance requirement of the AIC being below 1000 for the selected models set in Section 2.4 was satisfied with a clear margin.

6.3. Model integration into the radar processing pipeline

Then, the Swerling III model used for the DJI Air 3S was successfully integrated into the radar processing subgroup's simulation pipeline. The received power observed in the simulation followed the expected R^{-4} decay rate of the radar range equation. Moreover, the per-block RCS fluctuations resulted in a spread of approximately $10\,\mathrm{dB}$ around the general trend-line. This was a gained advantage of using the drone RCS model, compared to assuming a fixed RCS. The detection performance was overly optimistic, speculated to have been caused by the simplified static clutter model.

6.4. Programme of requirements reflection

A set of technical and functional requirements was compiled within the subgroup and across the entire system. Two out of three technical functions of the subgroup have been met. Datasets of both drones' radar cross-section measurements have been compiled, and an RCS model of both these drones have been formed. Sadly, there was insufficient time to build electromagnetic simulations for the drones.

The functional requirements have also only partially been met. The requirements regarding bandwidth operation region, power transmission & regulations, have all been met. The drones have also been modelled accordingly using the specified Swerling models, or statistical distribution models based on the state-of-the-art. The criterion of the AIC values being below 1000 was based on inspection of an RCS study [9]. For the selected models, the AIC values are significantly below this, only reaching -150, which suggests excessive conservativeness in model performance expectation and actual results.

However, some requirements have not been met. The reference target calibration had shown $4.5\,\mathrm{dBsm}$ deviance from the theoretical value, as opposed to the set $\pm 3\,\mathrm{dBsm}$ limit. This is determined to be leakage interference. Moreover, only some RCS models are within the targeted resolution, as seen in Section 3.1; in dataset 2 for the DJI Air 3S, the elevation only reaches $23°$ due to setup limitations, and the DJI Neo only has three separate elevation angles. However, due to the symmetry argument justifying the reduced sweep in azimuth, the azimuth resolution was deviated from deliberately.

6.5. Broader implications of this project & potential applications

The work this thesis delivers is publicly unavailable RCS profiles for two commercial drones: the DJI Air 3S and DJI Neo, specifically at a bandwidth of $24\,\mathrm{GHz}$, exceeding most comparable studies in the literature. The statistical models derived provide ready-made characterisation for integration into radar detection algorithms, based around short-range systems fixated on ground.

6.6. Lessons learned

The most challenging aspect of this project was facing a highly specialised field without prior background or experience in RCS analysis and modelling. A take-away from this is that there could be no consensus in the literature regarding a topic (i.e., drone RCS model selection), and that studies involve a wide range in frequency-band, drone types, and applied methodology. The exposure to a specialised field equipped us with knowledge and experience of how to find the silver-lining, among a vast field of literature.

6.7. Future work

Electromagnetic modelling of the drones was scoped, but not completed. This left a theoretical question of the thin-wire approximation's validity at $24\,\mathrm{GHz}$. Through completion of the EM-modelling, a direct physical simulation reference would be acquired for the measured RCS profiles. This enables assessment of near-field errors and calibration uncertainties, as well as micro-Doppler effects of spinning drone blades.
Similarly, many alternative models and evaluation methods were described in Section 3.3.4. A natural next step for the modelling analysis would be implementing the selection of these to further establish the choice of appropriate drone RCS model.

References

[1] Analog Devices. *EV-TINYRAD24G User Guide: Evaluating the TinyRad 24 GHz Demonstration Platform*. UG-1709. Rev. 0. [Online]. Available: https://www.analog.com/media/en/technical-documentation/user-guides/ev-tinyrad24g-ug-1709.pdf. Accessed: Jun. 12, 2026. Analog Devices. Mar. 2020.

[2] T. W. Anderson and D. A. Darling. “A Test of Goodness of Fit”. In: *Journal of the American Statistical Association* 49.268 (1954), pp. 765–769. DOI: 10.2307/2281537.

[3] J. J. Andersson and S. Simon. *Minding the drone gap: Drone warfare and the EU*. European Union Institute for Security Studies Brief 2024/17. Accessed: 2026-06-12. Oct. 2024. URL: https://www.iss.europa.eu/publications/briefs/minding-drone-gap-drone-warfare-and-eu.

[4] Constantine A. Balanis. *Antenna Theory: Analysis and Design*. John Wiley & Sons, 2005.

[5] S. Björklund and H. Hernnäs. “Statistical analysis of the radar cross section of two small fixed-wing drones using typical flights”. In: *IET Radar, Sonar & Navigation* 18.1 (2024), pp. 125–136. DOI: 10.1049/rsn2.12522.

[6] Sorana D. Bolboacă et al. “Pearson-Fisher Chi-Square Statistic Revisited”. In: *Information* 2.3 (2011), pp. 528–545. DOI: 10.3390/info2030528.

[7] G. J. Burke and A. J. Poggio. *Numerical Electromagnetics Code (NEC) - Method of Moments*. Tech. rep. Lawrence Livermore National Laboratory, 1981.

[8] CO2 Emissiefactoren Nederland. *Elektriciteit – stroom onbekend (gridmix)*. Accessed: 2026-06-07. 2026. URL: https://co2emissiefactoren.nl/factoren/2026/11/52/elektriciteit-stroom-onbekend-gridmix/.

[9] Martins Ezuma, Mark Funderburk, and Ismail Guvenc. “Compact-Range RCS Measurements and Modeling of Small Drones at 15 GHz and 25 GHz”. In: *arXiv preprint arXiv:1911.05926* (2019).

[10] Martins Ezuma et al. “Radar Cross Section Based Statistical Recognition of UAVs at Microwave Frequencies”. In: *arXiv preprint arXiv:2102.11954* (2021). DOI: 10.48550/arXiv.2102.11954.

[11] R. Guay, G. Drolet, and J. R. Bray. “Measurement and Modelling of the Dynamic Radar Cross-Section of an Unmanned Aerial Vehicle”. In: *IET Radar, Sonar & Navigation* 11.7 (2017), pp. 1155–1160. DOI: 10.1049/iet-rsn.2016.0520.

[12] L. Harriss and Z. Mir. *Misuse of civilian drones*. POSTnote 610, UK Parliament Parliamentary Office of Science and Technology. Jan. 2020. DOI: 10.58248/PN610. URL: https://post.parliament.uk/research-briefings/post-pn-0610/.

[13] Inras GmbH. *TinyRad-021200050271 Calibration Report*. Tech. rep. Calibration report for the 24-GHz TinyRad MIMO radar system. Linz, Austria: Inras GmbH, Aug. 2023.

[14] Ioannis K. Kapoulas et al. “Small Fixed-Wing UAV Radar Cross-Section Signature Investigation and Detection and Classification of Distance Estimation Using Realistic Parameters of a Commercial Anti-Drone System”. In: *Drones* 7.1 (2023), p. 39. DOI: 10.3390/drones7010039.

[15] Byung-Kwan Kim et al. “Drone Detection with Chirp-Pulse Radar Based on Target Fluctuation Models”. In: *ETRI Journal* 40.2 (2018). DOI: 10.4218/etrij.2017-0090.

[16] Ioan E. Lager et al. *Bachelor Graduation Project EE (EE3L1): BAP–Q4 Variant*. Provisional version, June 7, 2025, Academic Year 2025–2026. Delft University of Technology, Faculty of Electrical Engineering, Mathematics and Computer Science. 2025.

[17] David M. Lane. *Online Statistics Education: A Multimedia Course of Study*. Chapter 2: Graphing Distributions. Rice University. URL: http://onlinestatbook.com/ (visited on 06/07/2026).

[18] S. Liaquat, N. M. Mahyuddin, and I. H. Naqvi. “Drone Detection Using Swerling-I Model with L-Band/X-Band Radar in Free Space and Raining Scenario”. In: *Proceedings of the 12th International Conference on Robotics, Vision, Signal Processing and Power Applications*. Ed. by N. S. Ahmad, J. Mohamad-Saleh, and J. Teh. Vol. 1123. Lecture Notes in Electrical Engineering. Singapore: Springer, 2024. DOI: 10.1007/978-981-99-9005-4_53.

[19] Hubert W. Lilliefors. “On the Kolmogorov–Smirnov Test for Normality with Mean and Variance Unknown”. In: *Journal of the American Statistical Association* 62.318 (1967), pp. 399–402. DOI: 10.2307/2283970.

[20] Henry B. Mann and Abraham Wald. “On the Choice of the Number of Class Intervals in the Application of the Chi Square Test”. In: *The Annals of Mathematical Statistics* 13 (1942), pp. 306–317. DOI: 10.1214/aoms/1177731569.

[21] R. Nakamura and H. Hadama. “Characteristics of ultra-wideband radar echoes from a drone”. In: *IEICE Communications Express* 6.9 (2017), pp. 530–534. DOI: 10.1587/comex.2017XBL0079.

[22] Ryohei Nakamura, Hisaya Hadama, and Akihiro Kajiwara. “Ultra-wideband radar reflectivity of a drone in millimeter wave band”. In: *IEICE Communications Express* 7.9 (2018), pp. 341–346. DOI: 10.1587/comex.2018xbl0084.

[23] NIST/SEMATECH. *Kolmogorov-Smirnov Goodness-of-Fit Test*. 2012. URL: https://www.itl.nist.gov/div898/handbook/eda/section3/eda35g.htm (visited on 06/10/2026).

[24] J. W. Odendaal and J. Joubert. “Radar Cross Section Measurements Using Near-Field Radar Imaging”. In: *IEEE Transactions on Instrumentation and Measurement* 45.6 (1996), pp. 948–954. DOI: 10.1109/19.543991.

[25] Pedro Luiz Ramos. “Bayesian and classical inference for the generalized gamma distribution and related models”. Ph.D. dissertation. São Carlos, Brazil: University of São Paulo, 2018. DOI: 10.11606/T.104.2019.tde-03012019-143627.

[26] Retekdrones. *How Many Revolutions per Minute (RPM) a Drone Motor Needs*. 2025. URL: https://retekdrones.com/blog/how-many-revolutions-per-minute-rpm-a-drone-motor-needs/ (visited on 06/13/2026).

[27] Mark A. Richards, James A. Scheer, and William A. Holm, eds. *Principles of Modern Radar: Basic Principles*. SciTech Publishing, 2010. ISBN: 978-1-891121-52-4. DOI: 10.1049/SBRA021E.

[28] M. Ritchie et al. “Micro-drone RCS analysis”. In: *2015 IEEE Radar Conference* (2015). DOI: 10.1109/RadarConf.2015.7411926.

[29] Massimo Rosamilia et al. “Radar Detection Performance Prediction Using Measured UAVs RCS Data”. In: *IEEE Transactions on Aerospace and Electronic Systems* (2022), pp. 1–16. DOI: 10.1109/TAES.2022.3227224.

[30] Massimo Rosamilia et al. “RCS Measurements of UAVs and Their Statistical Analysis”. In: *2022 IEEE International Workshop on Metrology for AeroSpace (MetroAeroSpace)*. 2022, pp. 179–184. DOI: 10.1109/MetroAeroSpace54187.2022.9856394.

[31] V. Semkin et al. “Analyzing Radar Cross Section Signatures of Diverse Drone Models at mmWave Frequencies”. In: *IEEE Access* 8 (2020), pp. 48958–48969. DOI: 10.1109/ACCESS.2020.2979339.

[32] Weiqiang Shi, Xiao-Wei Shi, and Le Xu. “Radar Cross Section (RCS) Statistical Characterization Using Weibull Distribution”. In: *Microwave and Optical Technology Letters* 55.6 (June 2013), pp. 1355–1358. DOI: 10.1002/mop.27557.

[33] X. Song et al. “Dominant-plus-Rayleigh Models for RCS: Swerling III/IV versus Rician”. In: *IEEE Transactions on Aerospace and Electronic Systems* 49.3 (July 2013), pp. 2058–2064. DOI: 10.1109/TAES.2013.6558040.

[34] Statistics Netherlands (CBS). *Energy consumption of private dwellings; dwelling type and region*. Accessed: 2026-06-07. 2024. URL: https://www.cbs.nl/nl-nl/cijfers/detail/81528NED.

[35] M. A. Stephens. “EDF Statistics for Goodness of Fit and Some Comparisons”. In: *Journal of the American Statistical Association* 69.347 (1974), pp. 730–737. DOI: 10.2307/2286009.

[36] P. Swerling. *Probability of detection for complex targets*. Research Memorandum RM-1217. Santa Monica, CA, USA: RAND Corporation, 1954. URL: https://www.rand.org/content/dam/rand/pubs/research_memoranda/2006/RM1217.pdf.

[37] U.S. House of Representatives, Committee on Homeland Security. *Counter-UAS Authorities: Assessing the Government's Response to the Evolving Threat of Unmanned Aircraft Systems*. 119th Congress, Hearing, Figure 1: "Drone Threat Progression Abroad and in the Homeland". 2025. URL: https://www.congress.gov/119/chrg/CHRG-119hhrg61404/CHRG-119hhrg61404.pdf (visited on 06/07/2026).

[38] P. Wendt, A. Voltes-Dorta, and P. Suau-Sanchez. "Estimating the costs for the airport operator and airlines of a drone-related shutdown: An application to Frankfurt International Airport". In: *Journal of Transportation Security* 13.3–4 (2020), pp. 93–116. DOI: 10.1007/s12198-020-00212-4.

A

Ethical Implications of the Product and Engineering Process

The product developed during this project has significant potential to impact various fields through improving UAV detection technology. As drones become increasingly accessible and technologically potent, challenges arise with regards to safety and security. Hence, the development of countermeasures, starting with reliable drone detection, becomes increasingly important. [37]

This product impacts military & defense organisations the most, due to their emphasised safety management responsibilities. Proper drone detection allows for timely interception or tracking, which ultimately limits offensive drone users' capabilities.
This also raises concerns for recreational or institutional drone operators. Legitimate drones, such as delivery drones or police inspection drones, could be misidentified as threats and instigate interventions harmful for these users.

Likewise, airport operators and airspace security organisations also significantly benefit from refined drone detection systems. Drone activity greatly disrupts these organisations, where flights become heavily delayed, or airspaces close down temporarily. Such incidents can cause major economic losses; a drone disruption of Gatwick Airport lasting 33 hours caused around €55.8 million in losses. [38] Consequently, organisations can timely respond to drone incidents through a more robust detection system, increasing safety and operational continuity.

The development process of the RCS modelling essentially consisted of two aspects: conducting drone measurements in a controlled laboratory environment and performing computational analyses. Through their electrical energy consumption, the project accumulates a carbon footprint. In recent years, Dutch households consumed 2500 kWh on average according to Statistics Netherlands (CBS). [34] This equates to around 610 kg CO_2e using the Dutch grid electricity emission factor 0.244 kg CO_2e/kWh. [8]

Comparatively, twelve hours of measurements, and approximately forty hours of computation time accrued 7.0 kWh, or 1.71 kg CO_2e. Here, the power consumption was generously estimated as 250 W and 100 W respectively. This means our project left a carbon footprint equivalent to 0.3% of an Dutch household's electricity consumption, which suggests the environmental impact of the development process is considered insignificant.

Moreover, privacy and safety concerns commonly raised with such projects are considered negligible for our project. All measurements were conducted within the isolated anechoic chamber. This means no external RF signals could have been unwittingly collected. Since no biometric, personal, or other third-party information was collected, this aspect raised no concern.

While the product can be used for defensive or (airspace) security reasons, it could also be utilised maliciously. The information that the detection system outputs can reveal activity of legitimate UAV operations, e.g. of police forces, security, or inspection drones. In the hands of unauthorised individuals, this technological capability can be utilised to avoid law-enforcement activities, circumvent surveillance, or conduct other illicit practices. Through publishing the RCS signatures and classification procedure, we actively lower constraints for these adversaries to build their own systems, as documentation on the field also grants opportunities for developing countermeasures.

Despite these concerns, we found that the benefits of this technology outweighed its potential misuse. As such, our focus remained on technical performance. Retrospectively, there are two aspects that required attention: the implications of our detection flaggings, and data collection. Our drone detections do not explicitly analyse the behaviour of the target, meaning legitimate drones alike can very well be incorrectly flagged. For this reason, it would have been more ethically considerate to analyse this and pose potential solutions.

Similarly, there was no attempt to limit data collection to the targets of interest. This poses privacy concerns that require valid consideration. All in all, our future projects would benefit from reflections on ethical reflections for the design aspects, to further solidify justification of deploying the system.

If this concept demonstrator were to be developed into a real-world application, certain measures should be taken to ensure ethically responsible use. Due to the radio-frequency transmissions of the system, interference with existing radio-frequency systems must be minimised to prevent disrupting the nearby existing electromagnetic infrastructure. Therefore, installation and operation of this system must comply with local jurisdiction, and applicable (inter)national standards for EM-spectrum use. In the European system, this means complying with the EU U-Space framework (Commission Implementing Regulation 2021/664). Our radar operates in the open band around 24 GHz within a power limit that does not require a license.

Lastly, the system should exclusively be installed for well-justified safety and security purposes, such as the aforementioned airport surveillance or civil protection intents. Access to the detection data must also be restricted to authorised and trustworthy individuals. This is due to the importance of proper data interpretations, where an understanding of the system's behaviour is required to distinguish false positives from actionable flaggings, and keep decisions responsible. To ensure proper responsibility, there must be a clear, standardised procedure before action.

In summary, there are both beneficial and malicious ways our results or product could be utilised. Based on the considerations above, the potential societal benefits outweigh the possible malicious use cases, and our development process does not raise significant ethical concerns.

B

Usage of Generative AI

The growing use of Artificial Intelligence is a topic that requires appropriate handling in the context of academia. While it can serve as a useful tool to support one's work, it can also be (mis)used to produce content that is either incorrect or not genuinely understood by the user. This appendix aims to provide transparency on the usage of generative AI throughout the project.

B.1. Initial stages of the project & rule-setting

Initially, a convention was followed to restrict AI-usage within the project, as a natural follow-up of the Bachelor End Project introductory presentation of dr.ing. I.E. Lager as well as guidelines set in the BAP manual. [16] This convention was agreed upon to ensure full autonomy and independence over the end result. However, we quickly realised that the use of AI, if limited to tedious tasks and general support, was a tool that can greatly boost productivity. Thus, the personal ruleset regarding AI use was revised after the first week.

For the purpose of state-of-the-art review, AI was used to supplement our own research into the literature. After depleting our own ideas for search queries and reading through the relevant papers, AI was prompted to locate any other important sources in the literature that were overlooked. Through this process, many papers were discovered that deepened our understanding of the field and topics.

B.2. Development process of the project

AI also played a role in the development process. Specifically, a single array broadcasting issue in Python easily consumed two hours of time better spent on advancing the system. From there on out, AI was inquired to resolve syntax errors. Similarly, AI helped accommodate certain plotting preferences. For repetitive visualisation tasks, AI was asked to replicate styles of our other plots, add legend indices, or adhere to very specific requests, such as generating a function for mapping a recording file to a specific label:

```
def make_label(filepath):
    stem = Path(filepath).stem
    mapping = {
        "results_0elev0azi_-90deg_symmetry-argument": r"Elevation 0° LHS",
        "results_0elev0azi_180deg":  r"0°",
        "results_0elev5azi0_90deg":  r"5°",
        "results_10elev0azi_90deg":  r"10°",
        "results_15elev0azi_90deg":  r"15°",
        "results_20elev0azi_90deg":  r"20°",
        "results_23elev0azi_90deg":  r"23°",
    }
    return mapping.get(stem, stem)
```

In this `make_label()` function, for instance, the different measurement result files are mapped to labels indicating their elevation angle. This was, for example, used in Figure 4.8b.

Lastly, AI also helped in the early conceptual thinking processes. After our supervisor proposed directions for what to implement, we compared this to our own initial ideas. Here, we would use AI to list the differences between approaches, and these differences would form an idea on how to move forward in the analysis.

B.3. Final stages of the project

For the final stages of this project, specifically the writing phase, AI has mainly been used for three purposes: translating structured data texts into LaTex format & turning source metadata into ready-made BibTex citations, resolving specific grammatical uncertainties in writing, and a general cohesion check throughout the text.

After figuring out the proper structure for IEEE-based tables, AI was asked to translate table data into LaTex format. Similarly, turning papers into BibTex citation entries is a tedious task. So, the relevant metadata for each source was provided to AI to generate an IEEE citation in BibTex format, which was then verified manually to correct any hallucination errors.

Moreover, specific grammatical questions or odd sentence structures were submitted to AI for suggestions. AI was also used to detect minor errors in the text such as double spacings, incorrect punctuation, and duplicate mentions throughout the pages. Since this was difficult to reliably detect through manual review, AI helped smoothen this process.

B.4. Closing words regarding AI usage

Overall, we greatly value independent, self-built works. This can, however, cause clashes with AI usage with regards to authenticity and development. We notice, once a prompt has been submitted and the AI model returns a response, the chain of one's own creativity arriving at a solution is interrupted. Hence, we tried setting limits to AI use to certain cases, mainly being tedious tasks and enriching own thoughts and ideas.

C
Measurement setup in the anechoic chamber

This appendix aim to provide visuals with regards to the measurements performed in the DUCAT anechoic chamber. Figures of the room, reference objects, the radar, and drones have been provided.

C.1. Inside the room & general setup overview

Pictures of the room have been provided in Figures C.1C.1a and C.1C.1b. A zoom-in of the radar, as seen from the perspective of the target, is provided in Figure C.2.

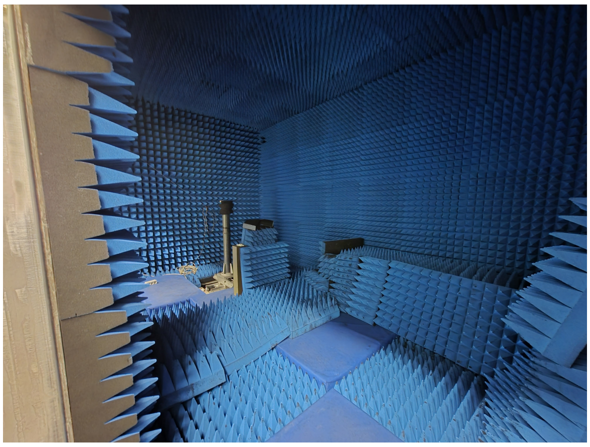

C.1a The radar-side view of the anechoic chamber setup, facing the target

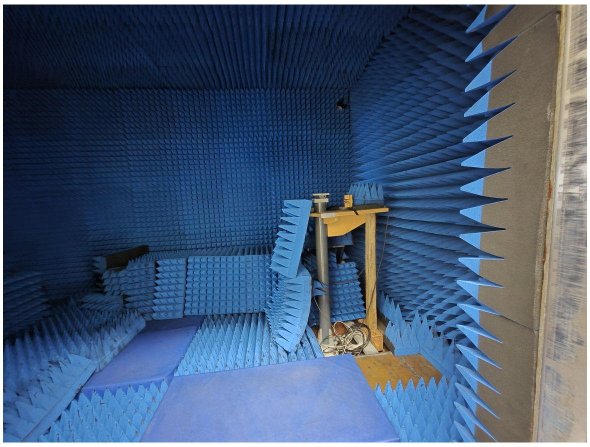

C.1b A view of the fixed pillar, on which the radar is mounted

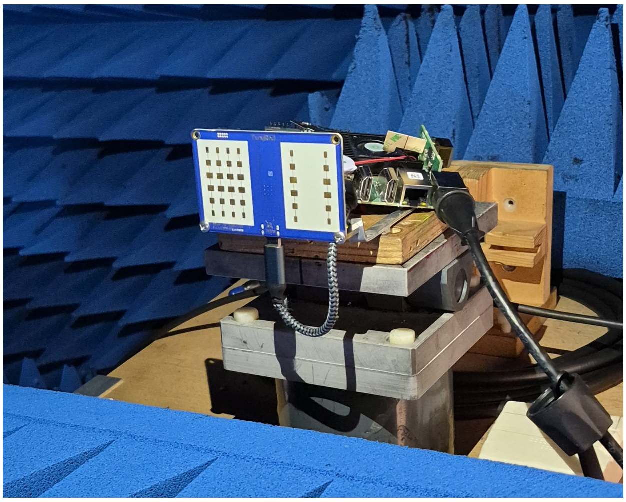

Figure C.2: A view of the radar; the shown side faces the target.

C.2. Reference targets

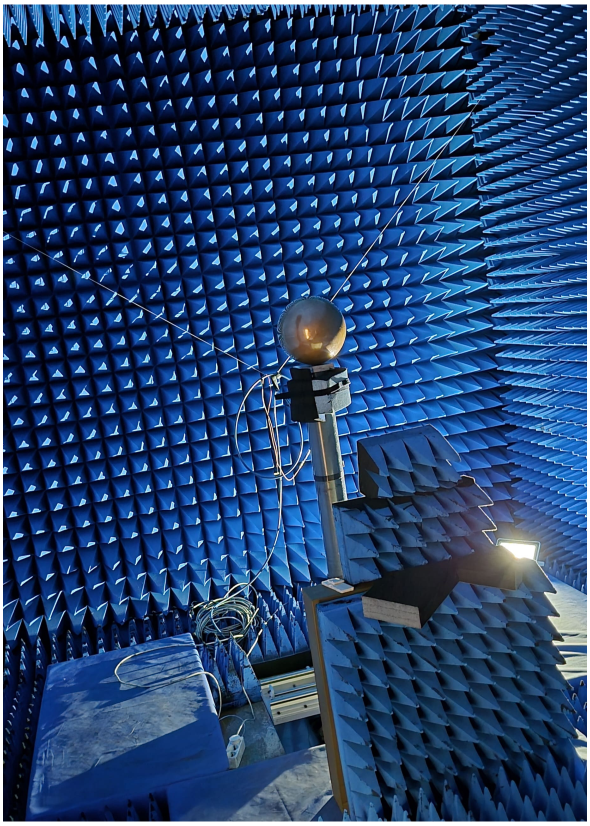

C.3a A view of one of the sphere measurements

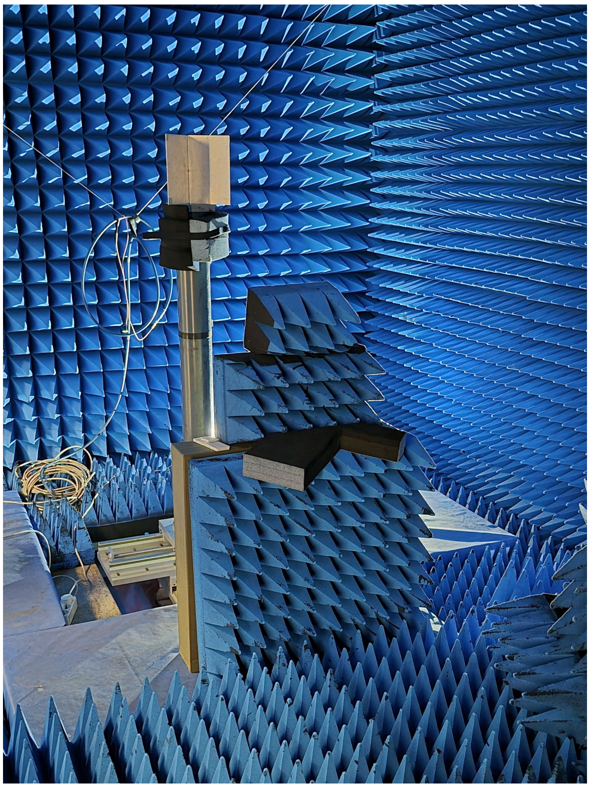

C.3b A view of the dihedral measurement

C.3. Drone measurements

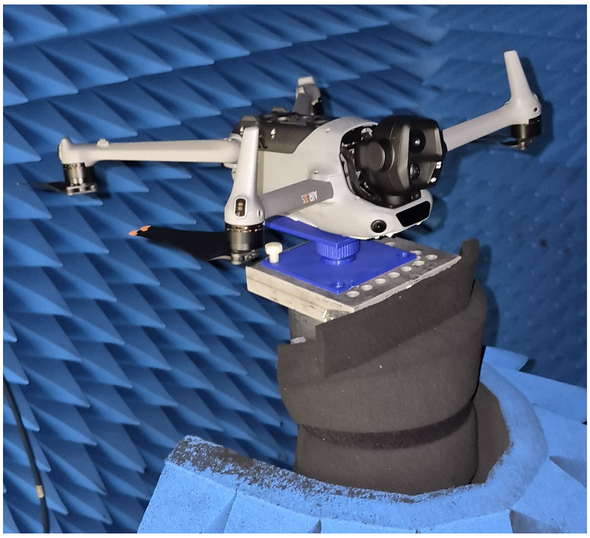

C.4a A view of the DJI Air 3S measurement

C.4b A view of the DJI Neo measurement

D

RCS plots for non-zero elevation angles

D.1. DJI Air 3S dataset 1

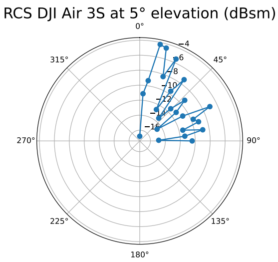

(a) RCS polar plot for the DJI Air 3S at 5° elevation according to dataset 1

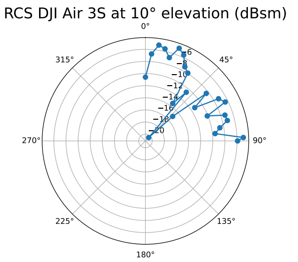

(b) RCS polar plot for the DJI Air 3S at 10° elevation according to dataset 1

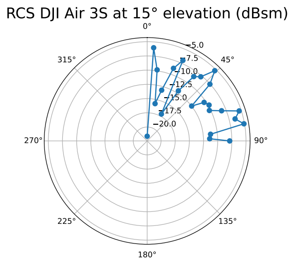

(c) RCS polar plot for the DJI Air 3S at 15° elevation according to dataset 1

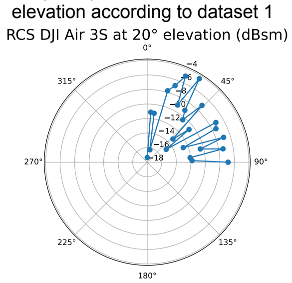

(d) RCS polar plot for the DJI Air 3S at 20° elevation according to dataset 1

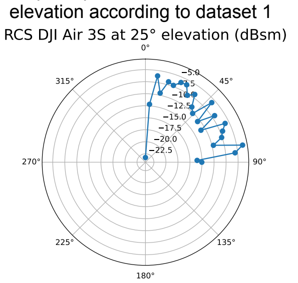

(e) RCS polar plot for the DJI Air 3S at 25° elevation according to dataset 1

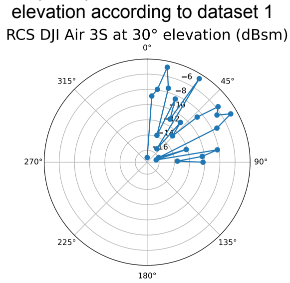

(f) RCS polar plot for the DJI Air 3S at 30° elevation according to dataset 1

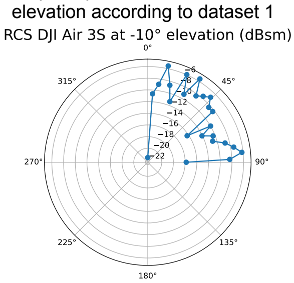

(g) RCS polar plot for the DJI Air 3S at -10° elevation according to dataset 1

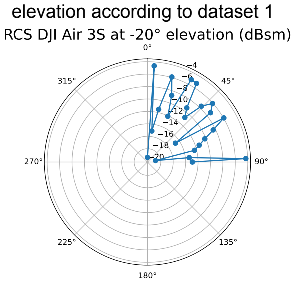

(h) RCS polar plot for the DJI Air 3S at -20° elevation according to dataset 1

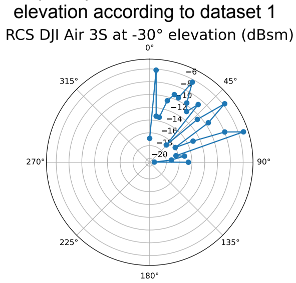

(i) RCS polar plot for the DJI Air 3S at -30° elevation according to dataset 1

D.2. DJI Air 3S dataset 2

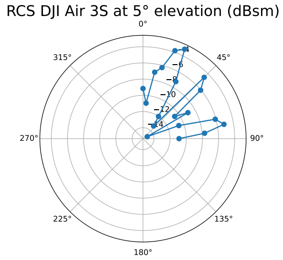

(a) RCS polar plot for the DJI Air 3S at $5°$ elevation according to dataset 2

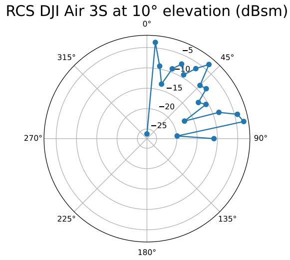

(b) RCS polar plot for the DJI Air 3S at $10°$ elevation according to dataset 2

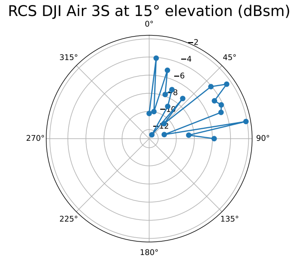

(c) RCS polar plot for the DJI Air 3S at $15°$ elevation according to dataset 2

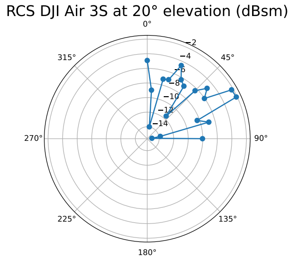

(d) RCS polar plot for the DJI Air 3S at $20°$ elevation according to dataset 2

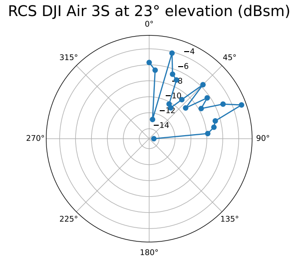

(e) RCS polar plot for the DJI Air 3S at $23°$ elevation according to dataset 2

D.3. DJI Air Neo

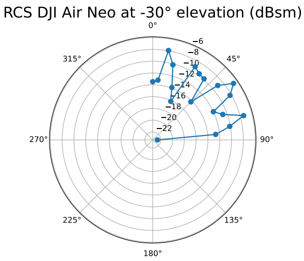

(a) RCS polar plot for the DJI Air Neo at $-30°$ elevation

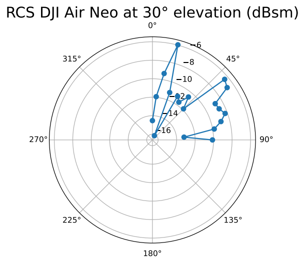

(b) RCS polar plot for the DJI Air Neo at $30°$ elevation