UNIVERSITY OF OKLAHOMA
GRADUATE COLLEGE

A PROPAGATION MODEL FOR MILLIMETER-WAVE RADAR PERFORMANCE ASSESSMENT IN AGRICULTURAL SCENARIOS

A THESIS

SUBMITTED TO THE GRADUATE FACULTY

in partial fulfillment of the requirements for the

Degree of

MASTER OF SCIENCE

By

ELISE A MCGOLDRICK

Norman, Oklahoma

2026

A PROPAGATION MODEL FOR MILLIMETER-WAVE RADAR PERFORMANCE ASSESSMENT IN AGRICULTURAL SCENARIOS

A THESIS APPROVED FOR THE
SCHOOL OF ELECTRICAL AND COMPUTER ENGINEERING

BY THE COMMITTEE CONSISTING OF

Dr. Nathan Goodman (Chair)

Dr. Justin Metcalf

Dr. Russell Kenney

Acknowledgments

I would like to begin by expressing my deepest gratitude to my mother, Kara Martinsen, and my father, Aaron McGoldrick. Their unwavering love, encouragement, and belief in me have been the foundation upon which this achievement was built. From the earliest days of my education through the completion of this thesis, they have offered patience, support, and countless sacrifices that words cannot fully repay. I am endlessly grateful for everything they have given me.

I would also like to thank the rest of my family for their constant encouragement and support throughout this journey. Their presence, both near and far, reminded me of what matters most and gave me the strength to persevere through the challenges of this work.

I owe a profound debt of gratitude to my advisor, Dr. Nathan Goodman, whose mentorship has shaped not only this thesis but my growth as a scholar. His guidance, patience, and generosity with his time over the years have been invaluable. He challenged me to think more critically, pushed me to strive for excellence, and offered steady encouragement even during the most difficult stages of this process. I am truly fortunate to have had the opportunity to learn from him.

Table of Contents

List of Tables

List of Figures

Abstract

Grain carts decouple a combine's harvesting rate from a truck's hauling schedule, but cart operators currently have only a rough, largely visual sense of how full the cart is at any given moment. A millimeter-wave (mmWave) radar mounted above the cart and scanning down across the bin as it fills could provide a continuously updated, through-dust estimate of pile height and shape, but the radar return in this environment is shaped not only by the grain pile itself but by the cart's canopy bars, the falling grain stream from the unloading auger, and airborne dust, all of which must be understood before a fill level estimate can be trusted. This thesis develops a physics-based propagation and signal simulation of a grain cart filling scene and validates the simulated radar returns against measured data from a real grain cart.

The simulation models an FMCW phased array radar illuminating a time-evolving scene composed of a Gaussian height field grain pile, three fixed cylindrical canopy bars, a falling grain chute, and a volumetric dust cloud, with line-of-sight self-shadowing of the pile computed at every time step. An antenna self-localization algorithm was developed that exploits the bars' known geometry as fixed reference reflectors, allowing the system to estimate its own antenna position from the radar returns themselves rather than assuming a perfectly known, static mount.

Comparing simulated and measured range and azimuth imagery shows that the model reproduces the dominant geometric behavior of the real scene. The canopy bar returns appear at matching locations in both datasets, and the pile return broadens and shifts as fill level increases in a manner consistent with the simulated line-of-sight obstruction model, which shows a two-phase pattern: visibility of the pile first rises as a small pile comes into clearer view of the antenna, then falls as a larger pile increasingly

shadows itself.

The comparison also identifies the model's largest gap: an elevated, structured noise floor in the measured data, attributable to unmodeled backscatter from the cart's metal floor and walls, is shown to be a more significant source of simulated-versus-measured mismatch than dust, the falling grain return, or the canopy bars, each of which is also shown to differ from the real scene in specific, explainable ways. The antenna self-localization algorithm achieves sub-meter three-dimensional position accuracy on average, with no evidence of cumulative drift, though accuracy is uneven across axes.

Together, these results demonstrate that a bar-anchored, physics-based propagation model can capture the key geometric behavior of a dynamic, cluttered agricultural radar scene, while also identifying adding cart floor and wall backscatter to the signal model as the clearest priority for closing the remaining gap between simulation and measured reality.

Chapter 1

Introduction and Background

1.1 Radar Agriculture Applications

Agriculture runs on timely information, from national and continental crop inventories down to individual field and plant conditions, to guide efficient, sustainable decisions. Optical and infrared sensing, the conventional tool for gathering that information, is routinely blocked by cloud cover, darkness, dust, foliage, and other obscurants, often at the exact moments the information matters most. It is this insensitivity to obscurants, more than any single frequency band or platform, that has made radar a recurring choice across nearly every spatial and temporal scale relevant to agriculture. On the continental-scale, satellite SAR crop inventories are updated annually, field-scale soil and moisture mapping are updated over a growing season, and plant or bin-scale volumetric quantity estimation are updated over minutes. The remainder of this section surveys representative work at the wide-area and small-area ends of this spectrum before Section 1.2 narrows to the specific plant/bin-scale sensing problem this thesis addresses.

A substantial body of work uses spaceborne SAR to classify crop type over large areas, motivated by the fact that optical sensors are blocked by cloud cover during exactly the growth stages that matter most for classification accuracy. Shang et al. [1] combined multi-year, multi-sensor, multi-site Landsat, RADARSAT-1, and ASAR data across five Canadian provinces and found that overall classification accuracies of at least 85% were achievable using multi-temporal satellite data, with the key operational finding being that adding two ASAR images in VV/VH polarization to a single optical image could substitute for the late-season optical acquisitions that are often unavailable due to cloud cover, allowing crop inventories to be produced earlier in the growing season. Their results also showed a synergistic effect from combining multiple radar frequencies (C-band ASAR and L-band ALOS PALSAR), since C-band primarily captures the top of the canopy while L-band penetrates deeper toward the stalk, together characterizing a larger volume of the crop canopy than either frequency alone [1].

Tan et al. [2] extended this line of work specifically to the early-season classification problem, using RADARSAT-2 polarimetric SAR and a particle-filter-based dynamic classification procedure driven by Growing Degree Day–normalized crop maturity rather than calendar date, to make the model portable across regions and seasons. Their method achieved over 90% classification accuracy for canola, corn, soybean, and wheat by the end of July, outperforming a multi-layer perceptron trained on the same polarimetric features, and demonstrated that canola in particular could be distinguished with very high accuracy well before the growing season concluded [2].

McNairn et al. [3] addressed a related but distinct problem: rather than identifying what crop is present, their work aimed to infer land management practices, specifically tillage and soil conservation practices, by combining optical crop-residue estimates with SAR-derived surface roughness. Using a multi-angular roughness index computed from RADARSAT-1 imagery at two incidence angles, they were able to separate smooth and rough soil surface classes with high correlation to ground-measured RMS height, while a Landsat-derived residue index (MSACRI) correctly classified all tested fields with greater than 60% residue cover. Their conclusion was that soil roughness alone is insufficient to distinguish conservation from conventional tillage without residue information and illustrates a recurring theme in this literature: single-sensor radar data is often necessary but not sufficient, and the strongest results come from fusing radar with a second, complementary data source.

Where the above studies concern what crop is present or how the soil beneath it has been managed, other radar applications target the physical condition of the standing crop itself. Han et al. [4] used dual-polarization Sentinel-1 imagery to monitor corn lodging following a wind and rain event, building a classification model from the ratio between plant height before and after lodging, derived from VH and VV+VH polarimetric indices. Their model achieved 97%, 100%, and 83% classification accuracy for mild, moderate, and severe lodging respectively. They also note that surface water and weed cover can be mixed into the backscatter signal in severely lodged, disturbed fields. This observation is a useful illustration of how ground clutter unrelated to the primary target of interest can degrade a radar-based agricultural model, a challenge structurally similar to the cart-floor and wall backscatter problem identified in this thesis.

Radar's cloud and light-independence also makes it particularly suited to time-critical monitoring during and immediately after weather disasters, when optical imagery is often unavailable for days. Boryan et al. [5] developed an operational Sentinel-1 flood monitoring procedure for the U.S. Department of Agriculture, applied to Hurricane Harvey in 2017, that combined a simple log-transform thresholding technique with the USDA's Cropland Data Layer to estimate inundated cropland and pasture acreage within 48 hours of the flood event. Validated against manually-derived ground reference data, the method achieved greater than 95% producer accuracy, demonstrating that even a comparatively simple SAR processing pipeline can deliver operationally useful results under time pressure when the alternative, optical imagery, is simply unavailable [5].

Moving from continental and field scales down toward the plant and bin scale that this thesis addresses, a smaller body of work trades coverage area for fine spatial and temporal resolution by mounting radar on or near a single field, vine, or bin rather than on a satellite platform. Closer in scale and platform to the system modeled in this thesis, Sanjiwani et al. [6] developed a mobile, ground-vehicle-mounted Stepped-Frequency Continuous Wave (SFCW) ground-penetrating radar system integrated with GPS for georeferenced soil water content mapping. Their system reconstructed the radar return via frequency-domain convolution and an IFFT, then estimated soil water content from the reflection coefficient using the Topp equation, achieving 96.5% accuracy against gravimetric ground truth on a single sample and a mean error of 3.54% across ten field locations [6]. This work is a useful point of comparison for the antenna self-localization and scene-scanning architecture in this thesis, since it demonstrates a similar pattern of pairing a scanning radar with a positioning system to build a spatially resolved map from a moving or repositioned platform, albeit at UHF/microwave frequencies suited to subsurface penetration rather than the mmWave, above-surface sensing problem addressed here.

The application closest in spirit to this thesis's own scanning FMCW architecture is the work of Henry et al. [7], who used a ground-based 24 GHz FMCW radar to estimate the quantity of grapes on grapevines for precision viticulture, in the presence of natural clutter from leaves, wood, and irrigation hoses. Their system scanned a 3D volume in front of each vine, converted the resulting echo distribution into voxels, and applied a contour-detection algorithm together with a novel "spread factor" parameter to classify which echoes were attributable to grapes rather than clutter, since water content in grapes produces a distinctively high dielectric permittivity and therefore strong reflectivity at microwave frequencies [7]. The resulting volume estimator was shown to scale linearly with measured grape mass across six grapevines of two varieties.

This work is directly analogous to the grain-cart problem addressed in this thesis in several respects: both use a scanning ground-based FMCW radar to estimate a volumetric quantity of interest in the presence of fixed and incidental clutter, and both rely on a return-magnitude-based classification step to separate the target of interest from structural and environmental clutter sources.

At the hardware and signal-processing level, Ahmad and Soraghan [8] investigated the application of MIMO radar concepts, borrowed from telecommunications, to agricultural SAR imaging. Their preliminary work focused on speckle noise reduction as a necessary pre-processing step, comparing Frost, Kuan, Lee, and SRAD filters on real SAR imagery of an agricultural region in Peninsular Malaysia and finding the Frost filter to give the best qualitative result, and proposed a bistatic MIMO configuration with multiple transmit antennas on a moving platform and multiple receive antennas on a stationary platform as a path toward higher-resolution agricultural radar imagery in future work [8]. While preliminary, this reflects a broader trend of adapting radar architectures originally developed for other domains — MIMO from telecommunications, phased-array beamforming from defense and automotive radar — into agricultural sensing problems, a trend this thesis's own generic phased-array FMCW model is also part of.

Outside of agriculture proper, mmWave FMCW radar has an established track record for range finding inside extremely dusty, enclosed industrial volumes, a sensing environment that shares much of its character with the enclosed, dusty grain cart addressed in this thesis. A widely cited mining sensor study by Brooker et al. [9] compared laser, acoustic, microwave, and mmWave radar candidates for measuring rock and backfill levels inside underground ore passes and stopes. Environments with dust levels exceeding 10 g/m^3, strong air currents, and target reflectivity variations of 30–40 dB showed that mmWave FMCW radar was the only technology able to combine a sufficiently narrow beam with adequate dust penetration at ranges beyond about 60 m. The resulting deployed systems achieved range accuracy within roughly 0.2 m at 117 m range in an ore pass, and produced full 3D point cloud reconstructions of fill surfaces in large underground stopes, including resolving differences in the angle of repose between different fill materials [9]. This mining sensor study is the closest documented analogue to the grain cart problem: both involve a fixed or near-fixed radar looking down into an enclosed, dusty bin at a granular pile whose surface shape and angle of repose are the primary quantities of interest, and the mining work's success with FMCW mmWave radar under comparable or worse dust conditions supports the feasibility of the architecture modeled in this thesis.

Taken together, this body of work shows that radar has been applied across nearly every spatial and temporal scale relevant to agriculture, from continental crop inventories updated annually, to field-scale soil and moisture mapping updated over a growing season, to plant-scale volumetric quantity estimation updated over minutes. The grain cart sensing problem addressed in this thesis sits at the fastest and smallest end of this spectrum, a fixed, minutes-long scan of a single enclosed bin, and is most directly comparable to the ground-based FMCW viticulture work of Henry et al. [7] and the mining bin-level sensing of Brooker et al. [9], rather than to the satellite SAR literature. Prior work on grain monitoring specifically has otherwise relied on combine-mounted yield sensors, load cells, or capacitive and ultrasonic level sensors rather than mmWave imaging radar, leaving the enclosed, dynamically filling grain cart, with its combination of structural bars, a moving fill stream, and a growing self-shadowing pile, without a published radar propagation model prior to this work. The recurring challenge across nearly all of the work surveyed above, separating a target return from structural, environmental, or platform-induced clutter, is the same challenge this thesis's LOS-gating, dust model, and clutter-anchored antenna localization are built to address, and it is the specific, minutes-long, bin-scale version of that problem that Section 1.2 now motivates in detail.

1.2 Motivation

Across the studies surveyed above, one operational advantage recurs regardless of platform, frequency, or target: radar can measure through cloud cover, darkness, dust, foliage, and other obscurants that routinely interrupt optical and infrared agricultural sensing. This is the same underlying advantage whether the obstruction is a monsoon cloud bank blocking a satellite pass over a flooded field [5], a canopy of leaves and irrigation hoses in front of a grapevine [7], or a wall of airborne dust inside an underground ore pass [9]. It is this insensitivity to obscurants, rather than any single frequency band or platform, that makes radar a recurring choice across such a wide range of agricultural and industrial sensing problems, and it is the same property that motivates its use in the grain cart scenario modeled in this thesis. A mmWave radar mounted above a filling cart can, in principle, see through the airborne grain dust and low light conditions that make optical or LIDAR-based fill sensing unreliable during active unloading.

During grain harvest, a combine unloads its hopper into a moving grain cart so that harvesting can continue without interruption. The cart driver then ferries the load to a waiting truck or wagon at the edge of the field. This three role job between the combine operator, the cart driver, and the truck driver is what allows modern harvest operations to run continuously rather than stopping every few minutes to unload. A grain cart can raise overall harvest efficiency substantially by keeping the combine picking instead of sitting at the field edge [10, 11], and carts in common use today range from roughly 650 to over 1,400 bushels in capacity and ride on flotation tires or tracks so they can be driven into soft or muddy ground that would normally bog down a truck [11]. A grain cart's job is to decouple the combine's harvesting rate from the truck's hauling schedule. The cart driver pulls alongside the moving combine, matches its speed, and receives the unloaded grain through an auger spout while the combine continues cutting [10], continually judging how much capacity remains, anticipating where the combine will need to unload next, and deciding when to break off and deliver the load to a truck, all while watching the field for ruts, wet spots, and other hazards [11]. Today this fill level judgment is almost entirely visual or experiential.

Despite this central role in the harvest workflow, the cart operator has only a rough sense of how full the cart is at any moment. Knowing how much grain is in the cart is typically a visual estimate through a cab window by the cart operator, supplemented at best by a simple capacitive or load based level sensor. Under filling the grain cart wastes trips to the truck, and over filling risks spillage or overloading the cart's structure. Commercial grain carts are sold with various scale and camera accessories, but a radar based, through dust fill sensor is not a mainstream feature of the product line [11]. A radar based fill level sensor mounted above the cart, scanning down across the bin as it loads, could give the operator a continuously updated picture of pile height and shape. However, the radar return in this environment is affected not just by the pile itself but by the cart's structural canopy bars, the falling grain stream from the unloading auger, and the suspended dust cloud. All of these scene characteristics must be understood and modeled before a fill level estimate can be trusted. The gap between current fill level estimation and implementing a mmWave radar is the practical motivation for the scenario modeled in this thesis.

1.3 Objective

This thesis addresses the need for a radar based cart fill level estimation by building a physics based propagation and signal simulation of the grain cart filling scene, then validating the simulated radar returns against real world data. This work has main three objectives. First, to construct a time evolving geometric model of a grain cart filling scenario, including the pile, canopy bars, and grain chute/auger stream. Secondly, to simulate the mmWave radar return from each scene component, accounting for LOS obstruction by the growing pile and attenuation from airborne dust. Lastly, to validate the simulated range profiles against measured radar imagery from a real grain cart, characterizing where the model succeeds and where it diverges.

Chapter 2 reviews relevant background on grain cart operation, mmWave radar fundamentals, and propagation modeling in cluttered and dusty environments. Chapter 2 also describes the overall system and simulation architecture, along with details about the grain pile and scene geometry model, then presents the radar signal model and processing chain. Chapter 3 presents simulation results and validation against measured data. Chapter 4 concludes with future work.

Having motivated the grain-cart sensing problem, this section steps back to survey how radar has been applied elsewhere in agriculture, situating this thesis's own scanning FMCW architecture within that broader body of work. Beyond mining and industrial bin-level sensing, radar has been applied across a wide range of agricultural monitoring problems, spanning crop identification, soil condition assessment, crop structural health, disaster response, and precision-agriculture quantity estimation. These applications differ substantially in platform (satellite, ground vehicle, or fixed ground-based scanner), frequency, and target, but together they establish radar, and particularly SAR, as an established tool in agricultural remote sensing, distinct from the enclosed-bin, mmWave, short-range sensing problem addressed in this thesis.

A substantial body of work uses spaceborne SAR to classify crop type over large areas, motivated by the fact that optical sensors are blocked by cloud cover during exactly the growth stages that matter most for classification accuracy. Shang et al. combined multi-year, multi-sensor, multi-site Landsat, RADARSAT-1, and ASAR data across five Canadian provinces and found that overall classification accuracies of at least 85% were achievable using multi-temporal satellite data, with the key operational finding being that adding two ASAR images in VV/VH polarization to a single optical image could substitute for the late-season optical acquisitions that are often unavailable due to cloud cover, allowing crop inventories to be produced earlier in the growing season [1]. Their results also showed a synergistic effect from combining multiple radar frequencies (C-band ASAR and L-band ALOS PALSAR), since C-band primarily captures the top of the canopy while L-band penetrates deeper toward the stalk, together characterizing a larger volume of the crop canopy than either frequency alone [1].

Tan et al. extended this line of work specifically to the early-season classification problem, using RADARSAT-2 polarimetric SAR and a particle-filter-based dynamic classification procedure driven by Growing Degree Day-normalized crop maturity rather than calendar date, to make the model portable across regions and seasons [2]. Their method achieved over 90% classification accuracy for canola, corn, soybean, and wheat by the end of July, outperforming a multi-layer perceptron trained on the same polarimetric features, and demonstrated that canola in particular could be distinguished with very high accuracy well before the growing season concluded [2].

McNairn et al. addressed a related but distinct problem: rather than identifying what crop is present, their work aimed to infer land management practices, specifically tillage and soil conservation practices, by combining optical crop-residue estimates with SAR-derived surface roughness. Using a multi-angular roughness index computed from

RADARSAT-1 imagery at two incidence angles, they were able to separate smooth and rough soil surface classes with high correlation to ground-measured RMS height, while a Landsat-derived residue index (MSACRI) correctly classified all tested fields with greater than 60% residue cover [3]. Their conclusion – that soil roughness alone is insufficient to distinguish conservation from conventional tillage without residue information – illustrates a recurring theme in this literature: single-sensor radar data is often necessary but not sufficient, and the strongest results come from fusing radar with a second, complementary data source.

Where the above studies concern what crop is present or how the soil beneath it has been managed, other radar applications target the physical condition of the standing crop itself. Han et al. used dual-polarization Sentinel-1 imagery to monitor corn lodging following a wind and rain event, building a classification model from the ratio between plant height before and after lodging, derived from VH and VV+VH polarimetric indices [4]. Their model achieved 97%, 100%, and 83% classification accuracy for mild, moderate, and severe lodging respectively, and their discussion of error sources – noting that surface water and weed cover can be mixed into the backscatter signal in severely lodged, disturbed fields – is a useful illustration of how ground clutter unrelated to the primary target of interest can degrade a radar-based agricultural model, a challenge structurally similar to the cart-floor and wall backscatter problem identified in this thesis (Section **??**).

Radar's cloud- and light-independence makes it particularly suited to time-critical monitoring during and immediately after weather disasters, when optical imagery is often unavailable for days. Boryan et al. developed an operational Sentinel-1 flood monitoring procedure for the U.S. Department of Agriculture, applied to Hurricane Harvey in 2017, that combined a simple log-transform thresholding technique with the USDA's Cropland Data Layer to estimate inundated cropland and pasture acreage within 48 hours of the flood event [5]. Validated against manually-derived ground reference data, the method achieved greater than 95% producer accuracy, demonstrating that even a comparatively simple SAR processing pipeline can deliver operationally useful results under time pressure when the alternative – optical imagery – is simply unavailable [5].

Closer in scale and platform to the system modeled in this thesis, Sanjiwani et al. developed a mobile, ground-vehicle-mounted Stepped-Frequency Continuous Wave (SFCW) ground-penetrating radar system integrated with GPS for georeferenced soil water content mapping [6]. Their system reconstructed the radar return via frequency-domain convolution and an IFFT, then estimated soil water content from the reflection coefficient using the Topp equation, achieving 96.5% accuracy against gravimetric ground truth on a single sample and a mean error of 3.54% across ten field locations [6]. This work is a useful point of comparison for the antenna self-localization and scene-scanning architecture in this thesis, since it demonstrates a similar pattern of pairing a scanning radar with a positioning system to build a spatially resolved map from a moving or repositioned platform, albeit at UHF/microwave frequencies suited to subsurface penetration rather than the mmWave, above-surface sensing problem addressed here.

The application closest in spirit to this thesis's own scanning FMCW architecture is the work of Henry et al., who used a ground-based 24 GHz FMCW radar to estimate the quantity of grapes on grapevines for precision viticulture, in the presence of natural clutter from leaves, wood, and irrigation hoses [7]. Their system scanned a 3D volume in front of each vine, converted the resulting echo distribution into voxels, and applied a contour-detection algorithm together with a novel "spread factor" parameter to classify which echoes were attributable to grapes rather than clutter, since water content in grapes produces a distinctively high dielectric permittivity and therefore strong reflectivity at microwave frequencies [7]. The resulting volume estimator was shown to scale linearly with measured grape mass across six grapevines of two varieties. This work is directly analogous to the grain-cart problem addressed in this thesis in several respects: both use a scanning ground-based FMCW radar to estimate a volumetric quantity of interest in the presence of fixed and incidental clutter, and both rely on a return-magnitude-based classification step (spread factor vs. this thesis's LOS-gated reflectivity and array gain, Section **??**) to separate the target of interest from structural and environmental clutter sources.

Finally, at the hardware and signal-processing level, Ahmad and Soraghan investigated the application of MIMO radar concepts, borrowed from telecommunications, to agricultural SAR imaging [8]. Their preliminary work focused on speckle noise reduction as a necessary pre-processing step, comparing Frost, Kuan, Lee, and SRAD filters on real SAR imagery of an agricultural region in Peninsular Malaysia and finding the Frost filter to give the best qualitative result, and proposed a bistatic MIMO configuration with multiple transmit antennas on a moving platform and multiple receive antennas on a stationary platform as a path toward higher-resolution agricultural radar imagery in future work [8]. While preliminary, this reflects a broader trend of adapting radar architectures originally developed for other domains – MIMO from telecommunications, phased-array beamforming from defense and automotive radar – into agricultural sensing problems, a trend this thesis's own generic phased-array FMCW model (Section 2.2) is also part of.

Taken together, this body of work shows that radar has been applied across nearly every spatial and temporal scale relevant to agriculture, from continental crop inventories updated annually, to field-scale soil and moisture mapping updated over a growing season, to plant-scale volumetric quantity estimation updated over minutes. The grain cart sensing problem addressed in this thesis sits at the fastest and smallest end of this spectrum – a fixed, minutes-long scan of a single enclosed bin – and is most directly comparable to the ground-based FMCW viticulture work of Henry et al. [7] and the mining bin-level sensing of Brooker et al. [9] discussed above, rather than to the satellite SAR literature, though the recurring challenge across nearly all of this work – separating a target return from structural, environmental, or platform-induced clutter – is the same challenge this thesis's LOS-gating, dust model, and clutter-anchored antenna localization are built to address.

Outside of agriculture, mmWave FMCW radar has an established track record for range finding inside extremely dusty, enclosed industrial volumes. A widely cited mining sensor study by Brooker et al. compared laser, acoustic, microwave, and mmWave radar candidates for measuring rock and backfill levels inside underground ore passes and stopes. Environments with dust levels exceeding 10 g/m^3, strong air currents, and target reflectivity variations of 30–40 dB showed that mmWave FMCW radar was the only technology able to combine a sufficiently narrow beam with adequate dust penetration at ranges beyond about 60 m [9]. The resulting deployed systems achieved range accuracy within roughly 0.2 m at 117 m range in an ore pass, and produced full 3D point cloud reconstructions of fill surfaces in large underground stopes, including resolving differences in the angle of repose between different fill materials [9].

The mining sensor study [9] is the closest documented analogue to the grain cart problem. Both involve a fixed or near-fixed radar looking down into an enclosed, dusty bin at a granular pile whose surface shape and angle of repose are the primary quantities of interest, and the mining work's success with FMCW mmWave radar under comparable or worse dust conditions supports the feasibility of the architecture modeled in this thesis. Prior work on grain monitoring more specifically has otherwise relied on combine-mounted yield sensors, load cells, or capacitive and ultrasonic level sensors rather than mmWave imaging radar, leaving the enclosed, dynamically filling grain cart with its combination of structural bars, a moving fill stream, and a growing self-shadowing pile without a published radar propagation model prior to this work.

Chapter 2

System Model and Simulation

2.1 Scenario Description

With the motivation and prior literature established, this section introduces the specific grain cart geometry and radar configuration that the rest of the chapter's model is built around. The simulated grain cart measures approximately 7.72 m long, 4.17 m wide, and 2.95 m tall, consistent with the larger end of commercially available dual-auger grain carts, which range up to roughly 1,400+ bushel capacity [11]. The radar antenna is mounted near the top of the cart, offset from center along the cart's length, and looks down and along the cart's interior.

The scenario unfolds over multiple discrete time steps, each representing a fixed interval of grain flow from the unloading auger spout, allowing the pile to grow progressively and the LOS obstruction pattern to evolve. This geometry directly mirrors the real sequence of events when a combine unloads its hopper into the cart while both vehicles continue moving across the field. An unload typically takes on the order of one to two minutes from a full hopper, during which the pile inside the cart visibly grows and shifts as the auger stream continues to land [10]. The simulation covers a short window early in this process, chosen to keep the simulation computationally tractable while still capturing the early, fast changing portion of pile growth, when the rate of change in pile height and LOS obstruction is largest.

Table 2.1 summarizes the core radar parameters used throughout the simulation. The simulated center frequency and bandwidth are achievable with existing single chip

Parameter	Value
Center Frequency (F_0)	50 GHz
Operating Wavelength (λ)	6 mm
Bandwidth (B)	4 GHz
Pulse Width (T_p)	5 ns
ADC Sample Rate (F_{adc})	5 GHz
Array Size	32 $\times$ 32 elements
Element Spacing	0.5λ
Element Gain	$\approx$ 3 dB
Transmit Power (P_t)	3 W

Table 2.1: Radar parameters used throughout the MATLAB simulation.

automotive and industrial mmWave transceivers operating in the 76–81 GHz band with up to 5 GHz of bandwidth, which strengthens the practical relevance of the simulation's parameter choices even though the exact center frequency differs.

2.2 Radar and Array Fundamentals

Before the individual scene components can be described, this section reviews the array and waveform theory, range resolution, beamforming, and the steering vector, that the rest of the model's radar return calculations depend on. The two parameters that define a millimeter-wave FMCW radar's performance are its center frequency and its operating bandwidth. For a linear frequency modulated (LFM) chirp waveform and matched filtering for range compression, range resolution is set by bandwidth as

$$\Delta R = \frac{c}{2B} \tag{2.1}$$

where ΔR is the range resolution, c is the speed of light, and B is the radar's operating bandwidth. Millimeter-wave automotive radars typically employ several GHz of bandwidth. For example, using 4 GHz as the operating bandwidth gives roughly 3.75 cm of range resolution in this system, fine enough to resolve pile growth on the order of a few centimeters per scan.

The simulation is intentionally generic rather than tied to specific commercial components. The simulation specifies the parameters that matter physically for the propagation and imaging problem at hand: center frequency, bandwidth, array configuration, and waveform type. By building the propagation and clutter analysis around a generic FMCW phased array architecture rather than a single vendor's hardware, the resulting conclusions about LOS shadowing, dust attenuation, and bar/chute clutter should generalize across the broad family of mmWave radars that share this same basic architecture, rather than being artifacts of one device's particular design. The generic radar modeled here forms a steerable beam electronically rather than by mechanically rotating an antenna, which is what allows it to scan the cart's interior in both elevation and azimuth without any moving parts. The basic idea is that a wave arriving at an N-element array from direction $\hat{u}$ reaches each element at a slightly different time, and therefore with a slightly different phase. If the array deliberately applies a compensating phase shift to each element before summing the signals, it can make the contributions from one particular direction add constructively while contributions from other directions tend to cancel. The compensating phase shift required at each element is captured by the steering vector,

$$\mathbf{v}(\hat{u}) = \left[e^{j2\pi F_0 \hat{u}\cdot\mathbf{x}_1/c},\ e^{j2\pi F_0 \hat{u}\cdot\mathbf{x}_2/c},\ \ldots,\ e^{j2\pi F_0 \hat{u}\cdot\mathbf{x}_n/c}\right] \tag{2.2}$$

where $\mathbf{x}_n$ is the position vector of the n-th element, F_0 is the center frequency, and $\hat{u}$ is a unit vector pointing toward the desired look direction. This same steering vector applies whether the array is transmitting a beam toward direction $\hat{u}$ or combining received signals from direction $\hat{u}$. Figure 2.1 shows the antenna's coordinate frame and the array geometry that this steering vector describes. Applying $\mathbf{v}(\hat{u})$ as a set of per-element weights and summing the result produces constructive addition corresponding to direction $\hat{u}$, while energy from other directions partially cancels. The narrower the resulting mainlobe, the better the array can distinguish between two closely spaced directions. Two features control how narrow that mainlobe is. The first feature is aperture size: a physically larger array (more elements, or wider spacing between them) produces a narrower beam for the same wavelength. The second feature is amplitude tapering. Applying a smoothly varying weight across the array trades a slightly wider mainlobe for substantially lower sidelobes, which matters in this scenario because a strong sidelobe response toward the metal canopy bars could otherwise be mistaken for a return from the pile itself, a risk made worse by the fact that many features in the scene reflect energy back toward the array, including the distributed, extended surface of the pile itself.

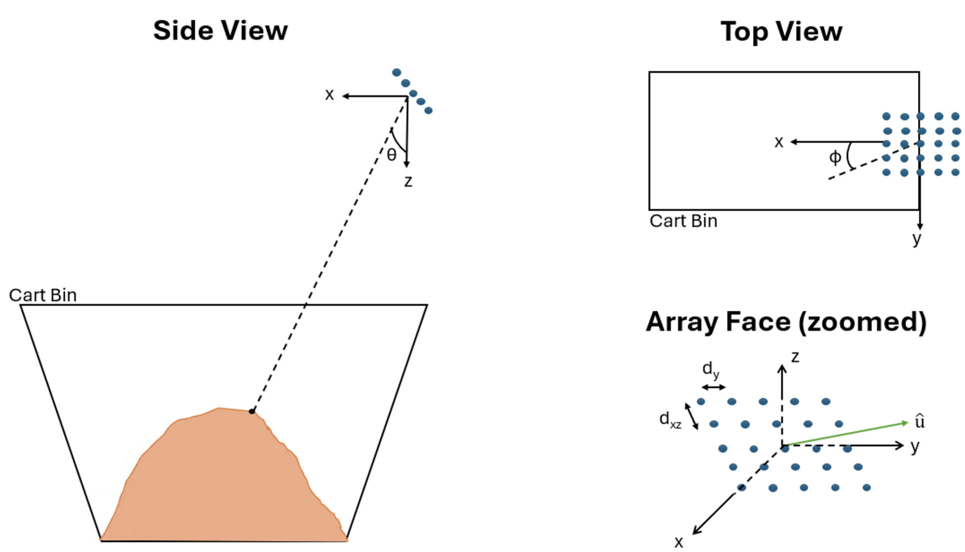

Figure 2.1: Antenna's coordinate frame and the array geometry.

In the grain cart simulation, this generic beamforming method is applied twice at every steering angle: once to compute how strongly the array's commanded beam illuminates each true scatterer location, and again using the known positions of the canopy bars to work backward and estimate where the antenna itself must be. Both uses rely on the same steering-vector relationship above, which is why the simulation can remain agnostic to which physical radar eventually implements it.

2.3 Scene Components

With the array theory in place, this section introduces the four physical scattering sources, the pile, canopy bars, grain spout, and dust cloud, that together make up the simulated scene. The four scattering sources populate the scene, each generated and triangulated separately before being merged into a single combined mesh for line-of-sight and visualization purposes.

The grain pile surface, the primary target of interest, is represented as a point cloud from a single Gaussian height field model and triangulated via 2D Delaunay triangulation to produce a continuous mesh surface usable for both rendering and LOS ray casting [12].

Three parallel cylindrical canopy bars span the cart's width at a fixed height of 2.94 m above the cart base. Each canopy bar is generated between explicit start and end points, sweeping a circular cross section of radius $R = 0.1$–0.15 m along its centerline. These bars serve multiple roles: they are structural clutter sources that must be distinguished from the pile return, but they can also serve as fixed reference targets for the antenna self-localization algorithm. Each bar's triangulated mesh is produced independently and then merged, with face index offsets applied so the three meshes don't share any vertex and face IDs.

The grain spout delivers new material from the combine and is modeled as a tilted stream originating near the cart's top and landing on the pile surface. Its triangulated geometry is produced using the current pile point cloud as an input, so the stream's landing point can be referenced relative to the pile's evolving footprint. This represents, in simplified geometric form, the unloading auger spout through which grain transfers from the combine hopper into the cart [11].

A volumetric dust cloud fills the upper cart volume above the pile, generated by the falling and settling grain. The dust cloud is represented not as a triangulated surface but as a 3D grid of voxels spanning the cart's interior up to the antenna's height.

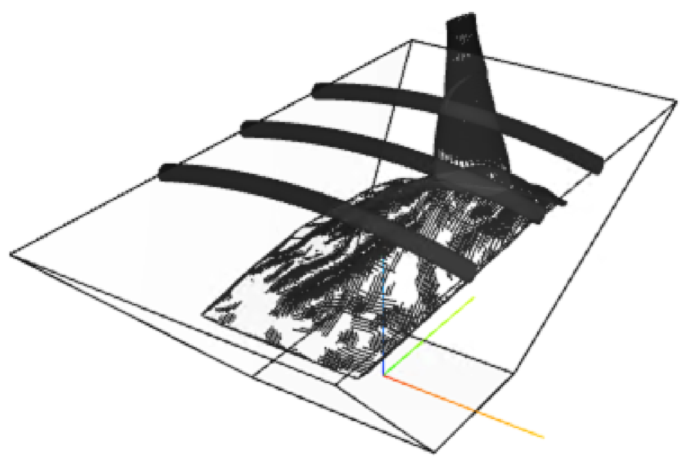

Figure 2.2: The complete scene with all four scattering sources.

2.4 Pile Growth Model

Having introduced the pile as one of the four scene components, it is import to show how its volume, footprint, and surface texture evolve at each simulated time step. Grain volume accumulates at a constant volumetric flow rate $V = 0.07\ \mathrm{m}^3/\mathrm{s}$ over each time step of duration Δt, consistent with the auger flow rates that empty a combine's hopper into the cart in roughly the one to two minutes typical of a real unload during active harvesting [10]. At time step t, the accumulated pile volume is

$$V_p^{(t)} = V_p^{(t-1)} + V\Delta t \tag{2.3}$$

The pile volume is converted to a footprint and height using a simplified conical growth model, solved iteratively from one time step to the next. Several of the numerical constants in this section, including the pile's shape factor scaling, were not derived analytically but were tuned empirically, adjusted iteratively until the simulated pile's growth and surface texture qualitatively resembled a real grain pile, since no ground-truth pile geometry dataset was available to fit them against directly. The pile's effective base area is computed from the current volume, using the previous time step's effective radius $r_p^{(t-1)}$ to account for the base spreading according to the angle of repose,

$$A_p^{(t)} = \frac{V_p^{(t)}}{h_0^{\text{init}} + \dfrac{V_p^{(t)}}{\pi \left(r_p^{(t-1)}\right)^2}}, \qquad r_p^{(t)} = \sqrt{\frac{A_p^{(t)}}{\pi}} \tag{2.4}$$

where h_0^{init} is the pile's seed height: a small, fixed starting height assigned before any grain volume has accumulated, which keeps (2.4) well defined at the very first time step rather than dividing by an initial radius of zero. A candidate central height, the height of the pile's peak at the current time step, is then computed from the updated volume and radius,

$$\tilde{h}^{(t)} = \frac{3V_p^{(t)}}{\pi \left(r_p^{(t)}\right)^2} \tag{2.5}$$

capped at a maximum of 2 m, consistent with the cart's finite physical height. The candidate central height is perturbed by a random factor to represent natural variation in real grain flow,

$$h_0^{(t)} = \min\Big(\tilde{h}^{(t)},\, h_{\max}\Big) \cdot (1 + 0.05\,\xi)\,, \qquad \xi \sim \mathcal{U}(-1, 1) \tag{2.6}$$

giving a roughly $\pm 5\%$ random variation on the central height at every time step. The angle of repose is set to 35°, a typical value for free-flowing granular agricultural material [13].

The pile's spatial spread is controlled by a shape factor that shrinks in magnitude as the pile matures. A stability factor,

$$s^{(t)} = \min\left(1,\ \frac{V_p^{(t)}}{0.1}\right) \tag{2.7}$$

saturates to 1 once roughly 0.1 m^3 of grain has accumulated, which in practice happens within the first few time steps. Equation (2.7) feeds into the shape factor,

$$\varphi^{(t)} = 1.6\left(1 + 0.1\left(1 - s^{(t)}\right)\eta\right), \qquad \eta \sim \mathcal{U}(0, 1) \tag{2.8}$$

so that the pile's spread carries more randomness while $s^{(t)}$ is still small and settles toward a fixed spread once $s^{(t)} \to 1$.

The pile surface is evaluated over a grid of $N_{cx} \times N_{cy}$ scattering cells spanning the cart's width and length, with grid coordinates

$$
\begin{aligned}
x_i &= -\frac{W_{\text{cart}}}{2} + (i-1)\frac{W_{\text{cart}}}{N_{cx}-1}, \quad i = 1, \ldots, N_{cx} \\
y_j &= -\frac{L_{\text{cart}}}{2} + (j-1)\frac{L_{\text{cart}}}{N_{cy}-1}, \quad j = 1, \ldots, N_{cy}
\end{aligned}
\tag{2.9}
$$

so that the grid is centered on the cart's footprint and spans its full width W_{cart} and length L_{cart}. For a grid point at position $\mathbf{x} = (x_i, y_j)$, the height field's peak amplitude and its added base offset are both driven by the growth model's output $h_0^{(t)}$ from (2.6),

$$
z_0(\mathbf{x}) = h_0^{(t)} \exp\left(-(\mathbf{x} - \boldsymbol{\mu})^T \left(\mathbf{P}^{(t)}\right)^{-1} (\mathbf{x} - \boldsymbol{\mu})\right) \tag{2.10}
$$

A random ripple term is added to break up the otherwise perfectly smooth Gaussian shape,

$$
z_1(\mathbf{x}) = z_0(\mathbf{x}) + \sqrt{\sigma_r \, z_0(\mathbf{x})}\, n(\mathbf{x}), \qquad n(\mathbf{x}) \sim \mathcal{N}(0,1) \tag{2.11}
$$

so that regions of the pile with greater local height receive proportionally larger fluctuations than the shallow edges. The rippled surface is then smoothed by convolving with a 2D Hanning window,

$$
z_2(\mathbf{x}) = (z_1 * h_{\text{smooth}})(\mathbf{x}), \qquad h_{\text{smooth}} = \frac{\mathbf{w}_H \mathbf{w}_H^T}{\sum \mathbf{w}_H \mathbf{w}_H^T} \tag{2.12}
$$

which suppresses the highest frequency component of the random ripple while leaving the broader Gaussian shape intact. Consequently, the simulated pile is not a perfectly smooth idealized cone but has a slightly textured surface representative of loose grain. A fixed base height offset equal to $h_0^{(t)}$ is then added to every point, giving the final surface $z(\mathbf{x}) = z_2(\mathbf{x}) + h_0^{(t)}$.

Because the cart's interior walls slope inward toward the floor, the pile surface must be clipped to the cart's actual physical envelope rather than extending through it. The cart's four sloped interior surfaces are represented as unit normal vectors,

$$\hat{\mathbf{n}}_w = \frac{(-\sin\theta_w,\ 0,\ \cos\theta_w)}{\sqrt{\sin^2\theta_w + \cos^2\theta_w}}, \qquad \hat{\mathbf{n}}_l = \frac{(0,\ -\sin\theta_l,\ \cos\theta_l)}{\sqrt{\sin^2\theta_l + \cos^2\theta_l}} \tag{2.13}$$

where $\theta_w = \tan^{-1}(H_{\text{cart}}/0.5W_{\text{cart}})$ and $\theta_l = \tan^{-1}(H_{\text{cart}}/0.5L_{\text{cart}})$. Since each component vector already satisfies $\sin^2\theta + \cos^2\theta = 1$, both normals are unit vectors by construction and the normalization has no numerical effect. A pile grid point is retained only if it lies on the interior side of all four sloped surfaces, checked via the sign of its dot product with each of the four normal vectors, and points that fail this test are discarded before the pile is passed on to the rest of the scene.

2.5 LOS Obstruction Analysis

An import part of the scenario is to describe how the model determines which parts of the grain surface is actually visible to the antenna at each time step. For each pile surface point, visibility from the antenna is computed by testing for line-of-sight between the antenna position and every point on the pile, using a receiver site placed at each candidate pile point [14]. This produces a binary visibility value at every pile point, where a value of 1 indicates an unobstructed line of sight from the antenna to that point, and a value of 0 indicates the point is obstructed. This visibility mask is then used to gate a randomly generated complex reflectivity,

$$\alpha_p = \sqrt{0.5}\,(n_1 + jn_2) \tag{2.14}$$

assigned to each pile point, producing the LOS-visible reflectivity α_p that is saved for use in signal generation. Pile points that are not visible from the antenna are simply excluded from contributing to the radar return, which is the model's representation of self-shadowing.

Figure 2.3 shows a representative example of the line-of-sight result for one antenna position and time step, with the antenna looking down on the pile. The red points mark pile locations obstructed from the antenna, whether by the canopy bars, the falling grain stream, the cart wall, or the pile's own geometry, and the green points mark locations with a clear line of sight. In this particular example, most of the back of the pile is obstructed by the falling grain from the auger spout, a portion of the front of the pile is obstructed by the cart's side wall, and a smaller number of obstructed points on the side of the pile arise from self-shadowing by the pile's own surface. Out of the 7318 points on the grain pile in this example, only 48.3% were visible to the antenna.

As the pile grows across multiple time steps, points near its base and on the far side from the antenna progressively lose visibility, since the rising pile surface increasingly blocks its own lower and far-side regions from the antenna's vantage point as the antenna's viewing angle grows shallower relative to the pile. This self-shadowing behavior is similar to the LOS modeling challenges documented in underground ore-pass and stope radar sensing, where growing or irregular fill surfaces and partial blockages were identified as a major source of degraded target discrimination that the sensor design had to account for [9].

2.6 Dust Cloud Model

Turning from the solid scene geometry to the airborne environment, this section reviews the mmWave dust attenuation literature that the voxel-based dust model, described in the following subsection, is built on. Various groups have produced work that has

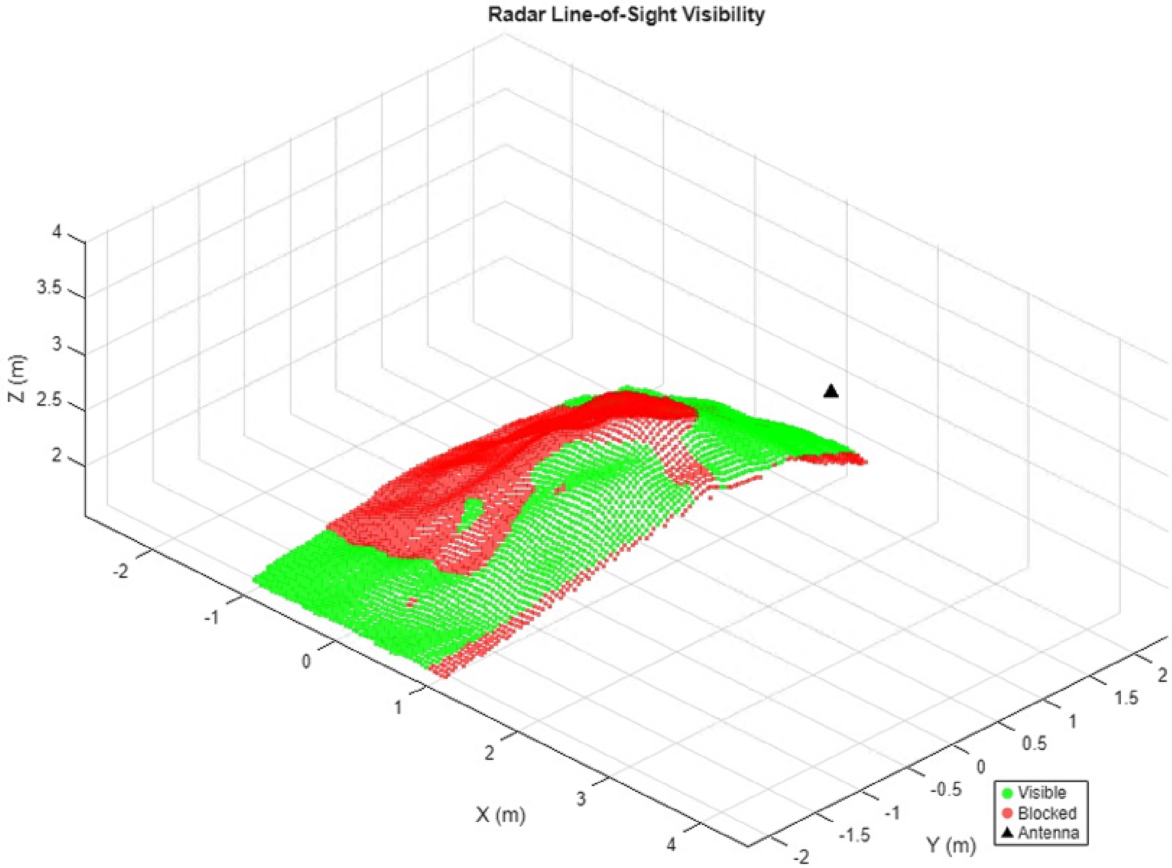

Figure 2.3: The line-of-sight results when the antenna is looking down on the pile. The red points indicated where the line-of-sight is obstructed and the green points shows where line-of-sight is visible to the antenna.

measured and modeled mmWave attenuation through suspended dust and sand, and several of the governing equations are directly transferable to the dust voxel model used in the scene simulation of this work. The starting point for nearly all of this literature is single-particle scattering theory. The attenuation coefficient of a population of particles is obtained by integrating the per-particle scattering cross section $\sigma_t(a)$ over the particle size distribution $N(a)$. Ahmed's foundational treatment expresses the specific attenuation for spherical particles as

$$\alpha_s = (1.143 \times 10^6)\,[N_d F G''] \int_0^\infty a^3 p(a)\,da \quad [\text{dB/km}] \tag{2.15}$$

where N_d is the particle number density (m^{-3}), F is the frequency in GHz, a is the particle radius, $p(a)$ is the normalized probability density function of particle size, and G'' is the imaginary part of the Mosotti–Lorentz–Lorenz factor $G = (\epsilon - 1)/(\epsilon + 2)$, with ϵ the complex relative permittivity of the dust material [15]. Because N_d is essentially impossible to measure directly in the field, Ahmed substitutes the optical-visibility relation

$$V_0 = \frac{15}{\alpha_0}\ \text{km}, \qquad \alpha_0 = 2\pi N_d \int_0^\infty a^2 p(a)\,da\ (4.343 \times 10^3)\ \text{dB/km} \tag{2.16}$$

so that visibility, an easily measured field quantity, stands in for the otherwise unmeasurable particle number density N_d. This lets attenuation be re-expressed entirely in terms of visibility V_0 and an effective radius,

$$a_e = \frac{\int a^3 p(a)\,da}{\int a^2 p(a)\,da} \tag{2.17}$$

Physically, a_e represents the single particle size that would produce the same attenuation as the full, real distribution of particle sizes if every particle were replaced by a sphere of that one radius. This yields the compact linear form $\alpha = \eta a_e$ (dB/km) and $\phi = \xi a_e$ (deg/km), where η and ξ depend only on frequency, permittivity, and visibility. The practical implication for the dust model used in this thesis is that it does not need to track an explicit particle count, as Ahmed's formulation does. The model can instead be driven by a visibility equivalent attenuation-per-voxel parameter, provided that parameter is benchmarked against a stated effective particle radius and visibility rather than chosen arbitrarily.

Elabdin et al. [16] extended this approach from the Rayleigh approximation to full Mie scattering, which removes the restriction that particle radius must be much smaller than a wavelength, a restriction that becomes questionable above roughly 37 GHz. Their attenuation expression for spherical particles, using Collin's Mie-based total cross-section efficiency factor, reduces after substituting the visibility relation to a frequency band specific formula,

$$A_d = \frac{a_e F}{V}\left[x_c + y_c a_e^2 F^2 + z_c a_e^3 F^3\right] \quad \text{dB/km} \tag{2.18}$$

where F is frequency in GHz, V is visibility in km, a_e is the equivalent particle radius in meters, and x_c, y_c, z_c are constants that depend only on the real and imaginary parts of the dust's complex permittivity. Elabdin et al. [16] record specific versions of this formula for the S, X, Ku, K, Ka, and W microwave bands and validate the Ku-band and Ka-band formulas against measured duststorm attenuation in Sudan and Saudi Arabia, finding values within roughly 0.01–0.04 dB/km of measured data across a range of visibilities from 0.6 to 5.6 km.

The follow up study by Elsheikh et al. [17] then tested this Mie-based model against several published particle size distributions, including exponential, normal, and three lognormal fits. The paper found that representing the real spread of particle sizes, rather than assuming a single equivalent radius, reduced the error between predicted and measured attenuation by roughly 10–50% depending on visibility, with the log-normal and discrete distribution fits each outperforming the single-equivalent-radius approximation, particularly as visibility deteriorated toward zero.

There are two complications from the broader literature that affect how the voxel attenuation parameter below is interpreted. First, the controlled dust tunnel measurements of Wikner [18] found two-way mmWave transmission loss across one meter of dense (≈3000 mg/m^3) dust to be only about 0.02 dB at 94 GHz and 0.1 dB at 217 GHz. Wikner used the Hashin–Shtrikman dielectric mixing bounds

$$\frac{\epsilon_{\min}}{\epsilon_{\max}} = \epsilon_e + \frac{1-f}{\dfrac{1}{\epsilon_i - \epsilon_e} + \dfrac{f}{3\epsilon_e}} \tag{2.19}$$

to bound the dust and air mixture permittivity before applying a Rayleigh attenuation formula structurally identical to Ahmed's. The result from (2.19) is one to two orders of magnitude lower than the values reported in [16, 17] for atmospheric duststorms at comparable visibility, a discrepancy attributable to the much shorter path length, different particle composition, and possibly the limits of instrumentation sensitivity in the tunnel test [18].

Second, the field study in [19] shows that charged sand and dust storms do more than attenuate amplitude. Using the extinction cross section for a single charged particle, where k_0 is wavenumber, ϵ_e is the complex permittivity of charged sand from the Maxwell Garnett mixing formula, σ is surface charge density, and θ_c is the charge distribution angle, they derive both an amplitude attenuation rate A and a phase-shift rate

$$\beta(t_m) = 2\pi k_0^2 \int_0^{\infty} \mathrm{Im}[S(t_m)]\, N(a)\, da \tag{2.20}$$

and show experimentally that the phase term causes significant pulse to pulse decoherence, reducing the adjacent pulse correlation coefficient and degrading coherent processing far more than the roughly 3.3 dB amplitude loss alone would suggest [19]. The dust model below does include a random phase term at each voxel, but this phase is drawn once and held fixed for the duration of a given time step, giving the dust cloud a realistic, speckle-like spatial texture rather than representing any change in phase between successive pulses. Kang et al.'s [19] decoherence mechanism is fundamentally different. It describes how the phase response of charged dust particles drifts from one pulse to the next within a single measurement, which degrades any processing that depends on comparing pulses coherently over time. Because the dust model below holds its phase fixed within a time step, it has no representation of this pulse-to-pulse phase drift, creating a resulting gap between simulated and real world detection performance under charged-dust conditions.

Theoretical models incorporating realistic particle size distributions predict attenuation that scales strongly with visibility and the assumed effective particle radius, ranging from roughly 0.01 to several dB/km depending on conditions and frequency band [15–17]. Recorded radar returns from dust tunnel and mining environments, by contrast, show attenuation closer to 0.01–0.5 dB under strong dust conditions [9, 18]. Grain dust inside a cart is denser and coarser than atmospheric duststorm aerosol, so the model's chosen 1 dB attenuation value, described next, accounts for worse possible dust conditions.

Dust in the simulated scene is modeled as a 3D grid of voxels filling the volume above the pile, spanning the cart's footprint in x (length) and y (width), and extending from the pile's current maximum height up to the antenna's height in z. The voxel grid spacing in each dimension is tied to the radar's own angular resolution rather than chosen arbitrarily. The spatial resolution in the x and y directions are each set to one-half of the array's elevation and azimuth beamwidth (2.30), so that each voxel is finer than the radar's own resolution cell. The dust grid is always resolved relative to what the radar beam itself could plausibly distinguish, rather than using a fixed physical voxel size unrelated to the sensor's own resolution.

Each voxel is assigned a random complex reflection coefficient,

$$\alpha_i = \alpha_v e^{j2\pi\psi_i} \tag{2.21}$$

where α_v is the assigned reflection coefficient magnitude for each voxel and $\psi_i \sim \mathcal{U}(0,1)$ is a random phase draw. The transmission coefficient shares this same phase term,

$$\gamma_i = \gamma_v e^{j2\pi\psi_i} \tag{2.22}$$

where the magnitude terms α_v and γ_v are derived from an assumed two-way attenuation of 1 dB, treated as a selectable model parameter, spread evenly across the number of voxels a ray would cross over the full dust cloud height. For each pile scatterer, the actual number of voxels along that scatterer's specific line-of-sight path determines its individual two-way attenuation due to dust, so that scatterers farther from the antenna, whose paths cross more dust voxels, are attenuated more than nearer ones, even though the assumed dust density per unit path length is uniform.

The dust cloud's own backscattered return is handled separately from this pile attenuation term. Because dust fills a volume rather than sitting at a single range, the dust voxels' contributions must account for self-attenuation along the path from the antenna to each voxel and back. Voxels are sorted by range from the antenna, and a cumulative one-way transmission factor is built via a cumulative product of the sorted γ_i values, so that a given voxel's effective return strength is reduced by the cumulative attenuation transmitted along its line-of-sight path, representing the physical fact that dust closer to the radar partially shadows dust farther away.

2.7 LFM Waveform Design

Having covered the scene geometry and dust attenuation, this section turns to the second key resolution parameter, the transmitted chirp waveform, and how it sets the radar's range resolution. The array's steered response and beamwidth set the radar's angular resolution. The second key resolution quantity is range resolution, set by the bandwidth B of the FMCW/LFM waveform, independent of the array geometry. The grain cart simulation's 4 GHz bandwidth, yielding $\Delta R \approx 3.75$ cm, per (2.1), is consistent with this same resolution-bandwidth relationship, simply scaled to a far shorter operating range where finer absolute resolution is both achievable and necessary to resolve centimeter-scale pile growth between time steps.

The transmitted waveform is a linear frequency modulated (LFM) chirp, chosen because pulse compression by means of a matched chirp gives fine range resolution without requiring the short, high-peak-power pulses that a simple unmodulated pulse of the same resolution would need. The chirp rate is set by the ratio of bandwidth to pulse width,

$$\gamma = \frac{B}{T_p} \tag{2.23}$$

Rather than transmitting a sharp-edged rectangular pulse, the waveform's envelope is built by convolving a rectangular pulse of duration T_p with a smoothing window $w(t)$ of duration $0.1\,T_p$,

$$a_p(t) = \mathrm{rect}\left(\frac{t}{T_p}\right) * w(t) \tag{2.24}$$

then normalizing to unit peak amplitude. This convolution smooths the pulse's rising and falling edges, which suppresses the range sidelobes that a hard-edged pulse would otherwise produce after matched filtering. With the pulse envelope $a_p(t)$ and chirp rate γ fixed, the complex baseband transmitted waveform is

$$x(t) = a_p(t)\, e^{-j2\pi(0.5B_{tot})t + j\pi\gamma t^2} \tag{2.25}$$

where B_{tot} is the total absolute pulse bandwidth achieved once the envelope smoothing has set the final pulse length, and the $-0.5B_{tot}$ term centers the chirp's instantaneous frequency sweep symmetrically around the carrier rather than starting it entirely on one side of baseband. The matched filter is simply the time-reversed complex conjugate of this waveform, which is the filter that maximizes signal-to-noise ratio (SNR) for a known waveform buried in white noise and is the standard tool used to compress the long chirp back down into a narrow range-resolved peak at the receiver [20].

In the discrete-time implementation, the continuous pulse duration T_p corresponds to a pulse length in ADC samples,

$$L_p = \mathrm{round}\left(\frac{T_p}{T_{adc}}\right) \tag{2.26}$$

Two further implementation details matter for correctly interpreting later range profiles. First, a minimum delay before the first valid ADC sample is enforced,

$$t_{min} = T_{adc}\,(1.05\, L_p) \tag{2.27}$$

to skip past the direct leakage of the transmitted pulse into the receiver before any genuine target echo could arrive. Second, thermal noise is added to every received signal as complex white Gaussian noise, with the noise power of a single discrete-time ADC sample set by the standard receiver noise expression,

$$P_n = k_B T F_{adc} F_n \tag{2.28}$$

where k_B is Boltzmann's constant, $T = 290$ K is the assumed system noise temperature, F_{adc} is the ADC sample rate, and $F_n = 10$ is the assumed receiver noise figure, a deliberately conservative value intended to represent a low-cost, worst-case receiver front end rather than an idealized low-noise design [21].

2.8 Received Signal Formation

With the waveform and array theory both defined, this section combines them into the full received-signal model, showing how each scene component's reflectivity, array gain, and propagation delay are assembled into a range profile at every steering angle. The array's steered response toward any given scatterer is given by

$$G(\hat{u}) = \left| \sum_n w_n \cdot v_n(\hat{u}_{steer}) \cdot v_n^*(\hat{u}_{scatter}) \right|^2 \tag{2.29}$$

where w_n is the real valued per-element amplitude taper weight. This is the discrete equivalent of the array factor used throughout phased array radar theory. Each scatterer's contribution to the received signal is weighted by the array's gain in the direction of the scatterer, given the array's current commanded steering direction. A scatterer sitting exactly along the steered look direction receives full array gain, while one lying off that direction is illuminated only by whatever gain remains in the sidelobe structure.

The element spacing is fixed at half a wavelength, $d = 0.5\lambda$, to avoid grating lobe ambiguity, so for the N_y-element linear dimension in elevation, the array's beamwidth follows the standard small-aperture approximation for the 3 dB beamwidth,

$$BW_{3dB} \approx \frac{\lambda}{N_y d} \tag{2.30}$$

This is the same beamwidth relation that underlies a design tradeoff documented broadly in the radar literature. Brooker et al. [9] offer one concrete example of this tradeoff in practice: for a fixed, practical-sized antenna aperture, microwave radar and millimeter-wave radar produce dramatically different beamwidths from the same physical aperture size, and it was this beamwidth versus aperture relationship that drove their selection of millimeter-wave frequencies for narrow-beam sensing inside an enclosed industrial bin. The same underlying engineering tradeoff motivates the 50 GHz operating frequency chosen for the grain cart simulation in this thesis. The dust voxel spacing is itself tied to this same beamwidth. The spatial resolution of the dust grid in x and y is set to one-half of the array's elevation and azimuth beamwidth, so that each voxel is finer than the radar's own resolution cell.

Each of the four scene components contributes a separate term to the total received signal before they are summed together, and each term follows the same general pattern. A per-scatterer complex reflectivity, weighted by the array's steered gain pattern toward that scatterer, per (2.29), modulated onto the chirp waveform at that scatterer's own round-trip delay. For the pile, the array's gain pattern toward every visible pile point is computed following the same construction, using $v_n(\hat{u}) = e^{j2\pi F_0 \hat{u}\cdot\mathbf{x}_n/c}$ as the array response vector built from each pile point's unit direction vector, with w_n the 2D Hanning amplitude taper. The same construction is repeated for the bars, auger chute, and dust from their own unit direction vectors. Each scatterer's two-way propagation delay is

$$\tau = \frac{2R}{c} \tag{2.31}$$

where R is the scatterer's range from the antenna position. The frequency-domain representation of the transmitted waveform, delayed by τ, is

$$W_{delay}(f) = X(f)\, e^{-j2\pi f(\tau - t_{min})} \tag{2.32}$$

where $X(f) = \mathcal{F}\{x(t)\}$ is the Fourier transform of the transmitted waveform from (2.25). The time-domain received contribution from the pile alone is the inverse Fourier transform of the product of the pile's per-scatterer reflectivity, its array gain, and this delayed waveform spectrum, summed over all visible pile points. The bar return, grain-spout return, and dust return are built identically, each using its own reflectivity, gain pattern, and range-dependent delay.

Isolating each component's individual contribution to the final range profile is not simply a consequence of summing the four returns together in the time domain. Summing them first would blend the four contributions together indistinguishably. Instead, each component's return is matched-filtered separately before being combined, which is what allows its individual range profile to be examined in isolation from the other three. This treatment also means the model treats the four scatterer types as statistically independent and non-interacting. Multiple scattering between the dust and the pile surface is not represented. This decomposition into independently attenuated component returns is similar to how recent radar performance studies in dusty and sandy environments separate amplitude attenuation, phase distortion, and decoherence effects rather than treating dust as a single scalar loss term [19].

For every combination of elevation (θ) steering angle and azimuth (ϕ) steering angle, the simulation recomputes the array's steering vector toward the commanded look direction,

$$\hat{u}_{steer} = [\sin(\phi),\, \cos(\phi)\cos(\theta),\, -\cos(\phi)\sin(\theta)] \tag{2.33}$$

and recalculates the four gain patterns toward every scatterer using this new steering vector, then forms the total received signal in the time domain. Matched filtering is performed by convolving the total received signal with the time reversed conjugate matched filter to obtain a clean range profile indexed by sample number. This converts directly to physical range using

$$R_a = \frac{c\,t_{tf}}{2} \tag{2.34}$$

The location of the strongest peak in this range profile at each steering angle is recorded as the range at which the dominant scatterer along that beam direction sits, and is converted into an estimated 3D point by projecting that range outward from the antenna along the current steering direction. Repeating this across the full elevation and azimuth grid at a given time step builds up a 3D point cloud estimate of whichever surface happens to dominate the return along each beam direction. Because the simulation always records only the single strongest peak per steering angle, a steering direction whose beam could contain both the pile and a nearby bar will report whichever of the two has the larger combined reflectivity and gain product.

2.9 Antenna Position Estimation

The final section of the system model turns from computing the radar return to the inverse problem. Using the bar and edge returns just described to estimate where the antenna itself actually is. Separating a target return from structural and environmental clutter is well studied in forward-looking and non-side-looking radar. The spatial steering vector fitting (SSVF) framework of Jiang, Liao, and Qu [22] offers a useful point of reference for how a phased array's known steering-vector geometry can be leveraged to separate a target return from structural clutter. The framework itself was developed for airborne ground-moving-target detection rather than for a static bin: the grain cart modeled in this thesis is a fixed, non-moving radar platform, where the antenna does not translate through space and clutter returns do not shift with platform motion. In their original formulation, the array elements are spaced along a single dimension, and the spatial-domain steering vector for the i-th range gate and k-th Doppler channel is a vector of phase progressions across the N array elements. The array modeled in this thesis, however, is a planar array with separate element spacings d_y along the array's y-axis and d_{xz} along its x and z-axis, so the single-dimension steering vector must be extended to account for both spacings independently. For element (n, m) at position

$$\mathbf{x}_{n,m} = n\, d_y\, \hat{y} + m\, d_{xz}\, \hat{x} \tag{2.35}$$

the steering response toward look direction $\hat{u}(\theta_k, \phi_i) = [\sin\phi_i,\ \cos\phi_i \cos\theta_k,\ -\cos\phi_i \sin\theta_k]$, using the same convention as (2.33), is

$$S_s(\theta_k, \phi_i)\big|_{n,m} = \exp\left[j\frac{2\pi}{\lambda}\left(n\, d_y \cos(\Psi_{i,k}) + m\, d_{xz} \sin(\phi_i)\right)\right] \tag{2.36}$$

where $\cos(\Psi_{i,k}) = \cos(\phi_i)\cos(\theta_k)$ is the direction-cosine term that scales the elevation-dimension spacing d_y, matching Jiang et al.'s original single-spacing formulation term-for-term, while the additional $\sin(\phi_i)$ term scales the azimuth-dimension spacing d_{xz} and only appears once the array is treated as genuinely two-dimensional [22]. This is structurally the same object as the steering vector $\mathbf{v}(\hat{u}) = e^{j2\pi F_0 \hat{u}\cdot\mathbf{x}/c}$ introduced in Section 2.2, with $\mathbf{x}$ generalized here to the two-dimensional element position $\mathbf{x}_{n,m}$. Setting $d_y = d_{xz} = d$, as is done throughout the rest of this simulation, collapses (2.36) back to the single-spacing case.

The relevance of Jiang et al.'s work to this thesis is conceptual rather than a direct numerical implementation. The radar is not airborne and does not face the Doppler-domain range dependence that motivates the rest of their framework. The radar does, however, face an analogous problem in the spatial domain. The bar, chute, and dust returns are clutter sources whose steering response must be characterized and separated from the pile return at every antenna position and time step. The use of the three known bar locations as calibration anchors, described below, plays a similar structural role to Jiang et al.'s use of known reference geometry – a fixed, well-characterized reference is used to correct for a geometry-dependent distortion in the quantity of actual interest [22].

Because the canopy bars sit at known, fixed locations for a given cart, their radar returns can be used to recover the antenna's actual position at each time step. The process begins by combining the three known bars and an arbitrary number of known cart top edges into one unified list of line segments, each defined by a start point, an end point, and a corresponding unit direction vector.

Working with full segments rather than fixed midpoints is the key to estimating the antenna's true position. A fixed center point is only informative about a bar's range and angle from one specific spot, whereas treating each bar and cart edge as an extended line lets the algorithm use the full spread of detections along that line's length. A single point gives only one range and angle pair to constrain the antenna's unknown 3D position, so estimating that position from single-point reflectors alone leaves the solution underconstrained unless enough independent points are combined. A line segment, by contrast, supplies a continuum of valid reference locations along its length, so a detection landing anywhere from one end of the bar to the other still contributes a geometrically consistent constraint, rather than being penalized for not matching a single fixed center. This is also why the projection in the residual calculation is clamped to the segment's true extent. Clamping converts the line into a finite, physically bounded constraint rather than an infinite line, so a detection that would otherwise project beyond the bar's actual end is still correctly penalized for the distance to that true endpoint.

Peak detection is performed by flattening the 3D range, elevation, and azimuth data cube for the current time step into a 2D range and angle bin array, sorting all angle bins by their maximum range profile value in descending order, and walking down that sorted list to accept the strongest peaks first, skipping any candidate peak that falls within a minimum separation of a peak already accepted. This suppression step prevents the same physical reflector from being counted multiple times across adjacent, highly correlated steering angles. This process continues until the target number of reference peaks has been found or the candidate list is exhausted, with a warning issued if fewer than the target number of distinct reference peaks could be isolated. If the algorithm consistently falls short of the target number of reference peaks, that indicates some bars or edges are not producing distinguishable returns at the tested SNR or steering resolution.

Each accepted peak's steering angle is converted into a unit look direction vector using the same steering-angle convention introduced in (2.33), and is initially assigned to whichever known segment minimizes a combined angular and range matching score against that segment's midpoint,

$$s = \varepsilon_a + 0.1\,\varepsilon_r \tag{2.37}$$

The 0.1 weighting on range error relative to angle error reflects that angular mismatch is treated as the more reliable indicator of which physical bar or edge a given peak belongs to. The angle error is calculated as

$$\varepsilon_a = \cos^{-1}(\hat{\mathbf{u}}_p \cdot \hat{\mathbf{u}}_e) \tag{2.38}$$

where $\hat{\mathbf{u}}_p$ is the peak's unit direction vector and $\hat{\mathbf{u}}_e$ is the unit vector from the currently assumed antenna position to the segment's midpoint. The range error is found by

$$\varepsilon_r = |r_p - \|\mathbf{r}_m - \mathbf{r}_A\|| \tag{2.39}$$

where r_p is the peak's estimated range, $\mathbf{r}_m = (x_m, y_m, z_m)$ is the line segment's midpoint, and $\mathbf{r}_A = (x_A, y_A, z_A)$ is the currently assumed antenna position.

With every peak assigned to a segment, the antenna position itself is recovered by nonlinear least squares, minimizing the total mismatch between each peak's predicted 3D location and its closest point on its assigned line segment [23]. A peak whose predicted location overshoots past the true end of a bar is penalized for the distance to that bar's actual endpoint, rather than being credited for lying on the bar's infinite extension.

The optimizer is initialized at the antenna's known ground truth position for a given time step, since this position is available directly from the simulation but would not be known in a real-world deployment. This choice validates that the residual formulation and segment-clamping logic converge correctly and remain stable at the true solution, rather than testing the algorithm's ability to recover antenna position from an uninformed or perturbed starting guess. Convergence behavior from more realistic, imperfect initial guesses is not evaluated in this thesis and is noted as a direction for future work in Chapter 4.

Conceptually, using a small number of fixed, known reflectors as calibration anchors, together with extended known line segments, mirrors how clutter spectrum compensation methods for non-side-looking radar use a known spatial steering vector reference to correct for range before further processing [22]. In both cases, a structure whose true geometry is known in advance is used to correct a quantity that cannot be measured directly.

Chapter 3

Results

3.1 Noise Floor and Clutter

This chapter compares the simulated range and azimuth imagery against measured data from the real grain cart, beginning with the overall background level against which every other scatterer in the scene must be distinguished before turning to the individual scattering sources and the antenna localization results in the sections that follow. The most immediately visible difference between the simulated and measured range and azimuth maps is the noise floor. The simulated maps show a nearly uniform, low level dark blue background outside of the discrete scatterer returns, whereas the measured maps show a textured, speckled background with substantially elevated clutter across the entire frame, including faint arc like banding that radiates outward from the antenna position. This is consistent with the underlying model assumptions. The simulation's thermal noise term is generated from a single, idealized receiver noise figure and ADC bandwidth, and the model has no explicit term for energy scattered or multipath reflected off the cart's metal floor and side walls.

In the physical setup, the steel cart walls and floor are large, continuous, electrically conductive surfaces directly in the radar's field of view. They produce multipath and diffuse scattering that contaminate every range bin, not just the bins occupied by the bars or pile. The simulation's combined mesh includes the cart shell geometrically for visualization and LOS purposes, but the radar signal model sums only four explicit scattering terms: the pile, bars, grain shoot, and dust. It does not include a cart floor or cart wall backscatter term. The elevated and structured noise floor in the measured data is therefore best explained as un-modeled ground and wall clutter superimposed on top of whatever genuine thermal noise is present, rather than as a discrepancy in the thermal noise parameters themselves. This is a primary, identified source of simulated versus measured divergence, and is a natural next addition to the signal model.

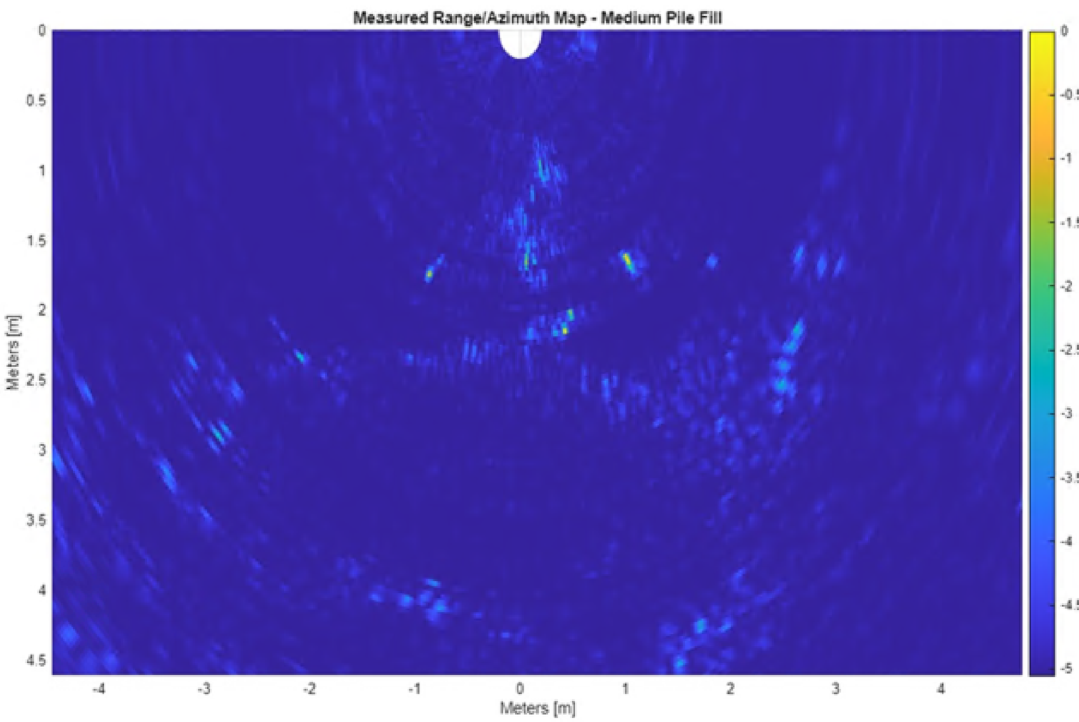

Figure 3.1: Measured range and azimuth map at medium pile fill. Measured data shows an elevated, speckled clutter floor attributable to ground and cart wall multipath, with the bar returns visible near the top of the frame as a banded arc of varying intensity.

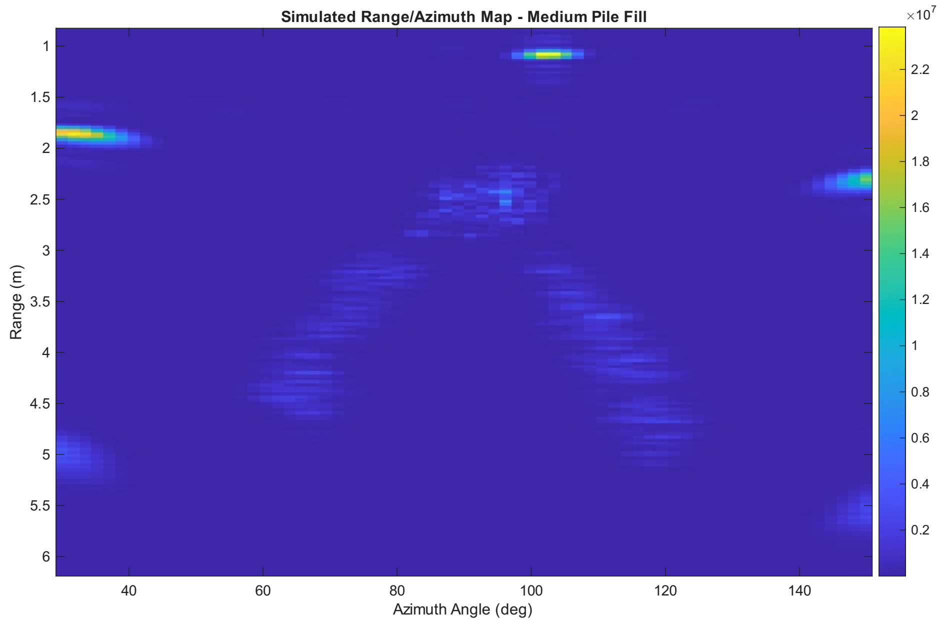

Figure 3.2: Simulated range and azimuth map at medium pile fill. Simulated data shows the same bar arc with uniform intensity along its length, a comparatively low and uniform noise floor, and the pile return concentrated in the central beam directions.

3.2 Falling Grain Stream Returns

With the background clutter difference established, this section turns to the first discrete scattering source in the scene, the falling grain stream from the unloading auger, comparing its appearance in the measured imagery against the current simulated result at medium pile fill. A second clear difference is the strength and extent of the falling grain return. In the measured plots, the falling stream produces a visually prominent cluster of moderate to strong returns concentrated just below the antenna's near range blind zone, spanning a noticeably broader vertical extent that is consistent with the radar illuminating the grain stream along its full visible length as it falls from the auger spout into the pile. In the simulated plots shown in Figure 3.2, the grain shoot return still appears as a comparatively faint, narrow feature rather than a broad curtain, visible near the top center of the map at close range and near broadside azimuth. At no time step does it approach the spatial extent or intensity seen in the real data.

This is explained directly by how the grain shoot is modeled. The simulation represents the falling stream as a single tilted cylinder of fixed radius with a fixed reflectivity vector, and at each steering angle the radar illuminates only whatever thin slice of that cylinder's surface falls within the current beam, rather than the full, continuously falling column of grain that a real radar return integrates over. Because the steering loop records only the single strongest range profile peak per angle, and because the chute's modeled radar cross section is small relative to the pile and bars, the chute return is easily "out competed" by neighboring scatterers in most beam positions, leaving only a thin trace where the beam happens to align almost exactly with the stream's narrow geometric extent.

The real falling grain stream, by comparison, is not a smooth cylinder but a turbulent, spreading curtain of grain with much greater angular and radial extent and a correspondingly larger effective radar cross section, which is why it appears as a broad, persistent feature in the measured imagery rather than a thin line. This difference is a known limitation of the current grain shoot model rather than a processing artifact, and is a natural candidate for refinement.

3.3 Support Bar Returns

This section examines the three canopy support bars, the strongest and most spatially consistent feature in both datasets. The three support bars produce the strongest,

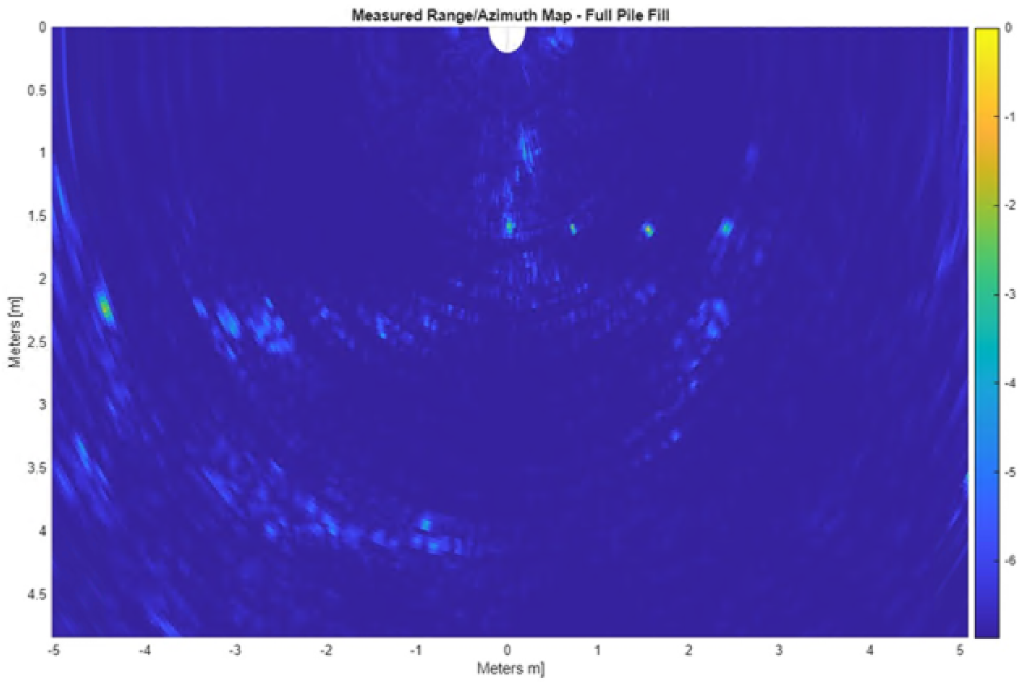

Figure 3.3: Measured range and azimuth map at full pile fill. The pile return broadens and extends to a wider range of azimuth angles, consistent with the pile's larger footprint.

most spatially consistent feature in both simulated and measured data: a high intensity, curved and banded return near the top of each map corresponding to the bars' fixed height and position across the cart. This is the one feature that validates well qualitatively. Both datasets place the dominant bar return in the same near range, broadside azimuth region, and both show the expected curved banding pattern as azimuth sweeps away from broadside.

The one notable difference is amplitude consistency along each bar's length. In the simulated maps the bar returns are essentially uniform in intensity across the full steering sweep, while in the measured maps the bar returns visibly brighten and dim at different azimuth angles, giving the bar arcs a "spotty" or beaded appearance rather

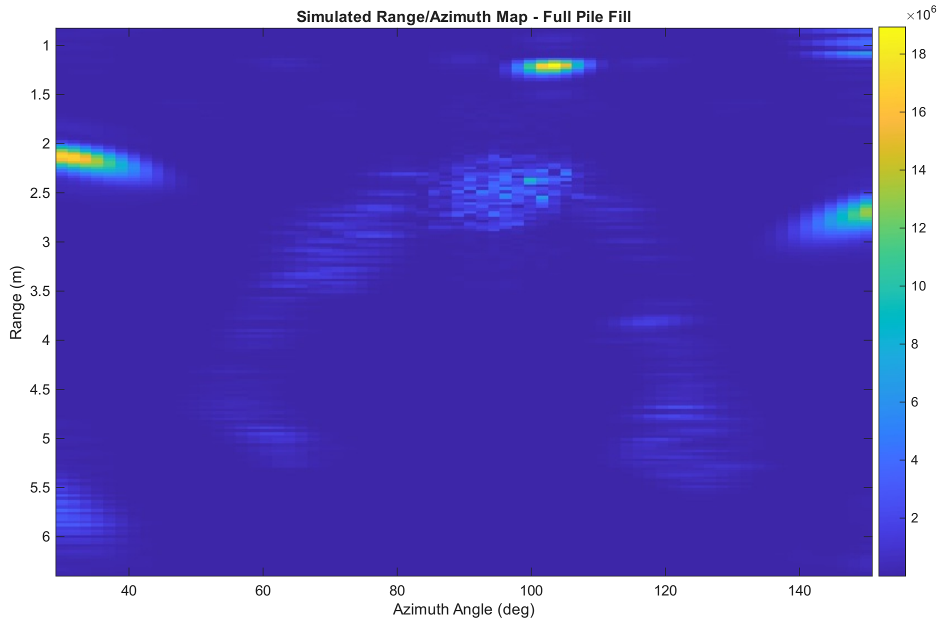

Figure 3.4: Simulated range and azimuth map at full pile fill. The simulated map's pile return takes on a wider arc shape at full fill rather than the tighter, more localized cluster seen at medium fill.

than a continuous bright band. This is consistent with the bar reflectivity model, which assigns every bar point the same fixed magnitude reflectivity, implicitly treating the bar as a perfectly uniform cylindrical reflector along its entire length. Real structural bars are not perfectly uniform, smooth cylinders. They have surface irregularities, welds, mounting brackets, paint and rust variation, and small deviations from a perfect circular cross section, all of which cause the true radar cross section to vary with aspect angle and position along the bar's length. The uniform RCS assumption in the current bar model is therefore the most likely explanation for why the simulated bar arcs look artificially smooth and constant in brightness compared to the more variable, weathered appearance of the bar returns in the measured data.

To examine how sensitive the bar return is to the assumed bar geometry, the bar radius was reduced in a follow-up simulation run, shown in Figure 3.5. Because each bar's cross section is swept along its centerline at radius R (Section 2.3), shrinking R directly shrinks the physical area, and therefore the number of scattering points, sampled along the bar's surface at each time step. With fewer points on a thinner cylinder contributing to the array's steered response at any given angle, the resulting bar return in Figure 3.5 is visibly smaller and less intense than the wider-bar case shown in Figures 3.2 and 3.4. The bar arc appears thinner and more fragmented, and the peak return level near broadside is reduced, since less reflecting surface falls within the beam at any one steering angle. This behavior is the expected consequence of the bar model's reflectivity being distributed per unit surface area rather than assigned as a single lumped return per bar, and it offers a useful secondary lever, alongside the non-uniform RCS treatment discussed above, for bringing the simulated bar arc's amplitude and apparent thickness closer to the more variable, weathered appearance seen in the measured data.

3.4 Pile Returns

Having covered the background clutter and the two structural clutter sources, this section turns to the grain pile itself, the primary target of interest, tracing how its simulated and measured returns change across low, medium, and full fill levels. The grain pile itself is the hardest feature to compare directly, because it is also the feature most affected by the noise floor and clutter differences discussed above. In the measured maps, the pile return is partially obscured by the elevated ground and wall clutter, making it difficult to cleanly separate pile from background by eye, whereas in the

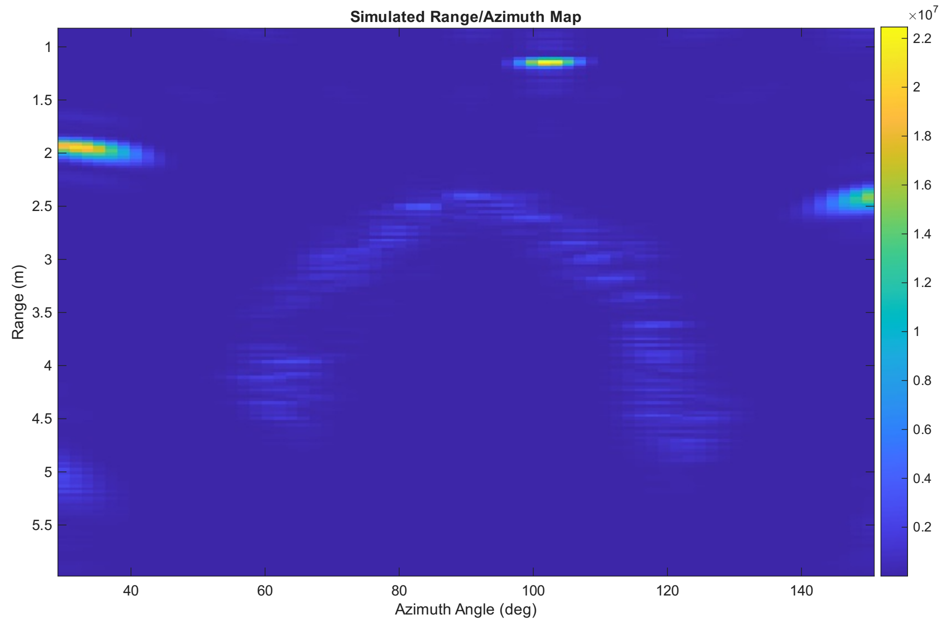

Figure 3.5: Simulated range and azimuth map with reduced canopy bar radius. With the bar radius made smaller, fewer scattering points are swept along each bar's surface, and the resulting bar-arc return near the top of the frame is correspondingly smaller and fainter than the wider-bar case shown in Figures 3.2 and 3.4.

simulated maps the pile appears as a comparatively well defined, higher contrast cluster of points against the simulation's much lower noise floor.

Despite this, a consistent qualitative trend holds across both datasets and across the low, medium, and full fill comparisons. As fill level increases, the region of strong returns broadens and shifts in range, consistent with the pile's growing footprint and height progressively occupying more of the radar's field of view and changing which surface points have line-of-sight to the antenna. The specific shape of the simulated pile return also changes noticeably between fill levels, appearing as a tighter, more localized cluster at low fill levels and broadening into a wider, arc like distribution at full fill levels, which is a direct consequence of the single Gaussian height field model. As the pile volume input grows, both the effective base radius and the LOS visible footprint expand, so the set of pile points actually contributing to the return spreads out accordingly. This pile shape progression is a useful qualitative validation point in its own right, and is shown explicitly across all three fill levels side by side rather than only at one representative time step.

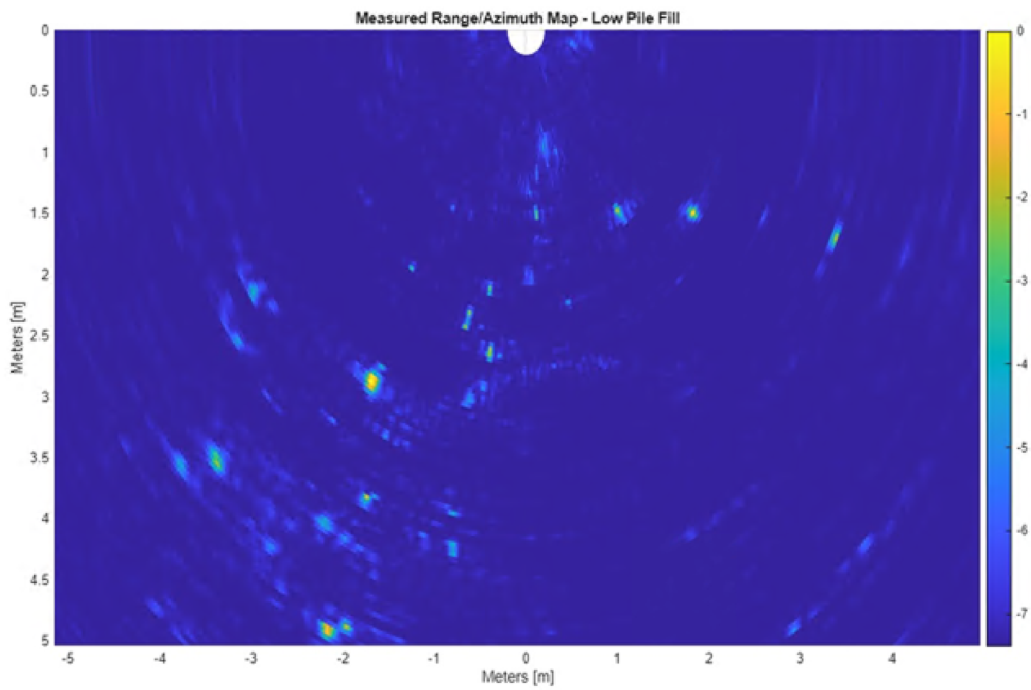

Figure 3.6: Measured range and azimuth map at low pile fill. The measured data shows a small diffuse, low-amplitude return in the same general region, partially obscured by clutter.

Because the pile return in Figures 3.2 and 3.4 is shown alongside the grain chute's narrow trace, isolating the pile's own shape from that of the shoot is easier when the shoot is removed from the scene entirely. Figure 3.8 shows the same medium-fill scene simulated without the grain chute cylinder present, which removes the chute's arc

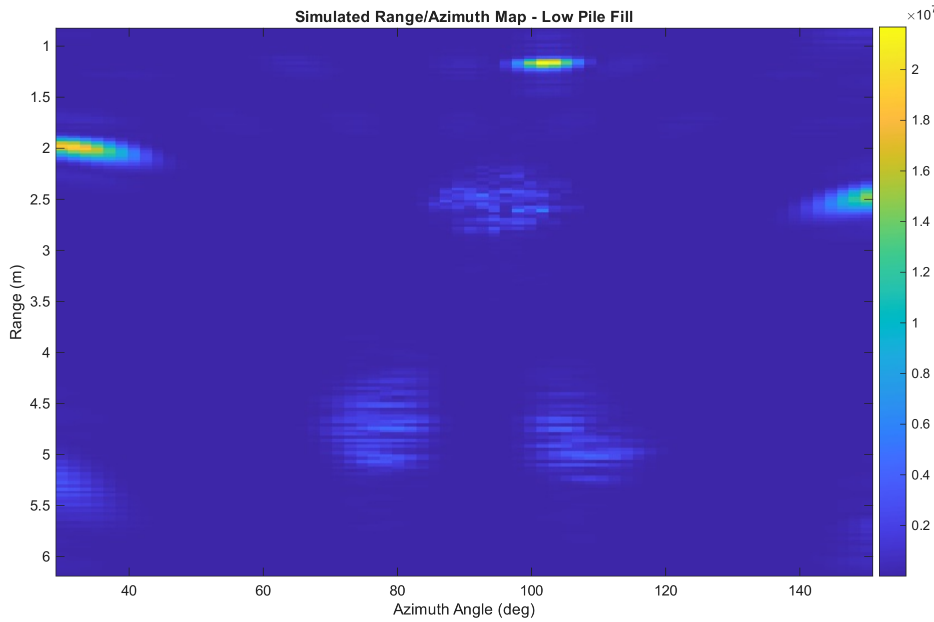

Figure 3.7: Simulated range and azimuth map at low pile fill. The simulated pile return appears as a small, tightly localized cluster relative to the medium and full-fill cases.

feature from the center of the map and leaves only the pile and bar returns. Viewed this way, the pile's arc-shaped broadening with fill level, described above, is easier to attribute unambiguously to the pile surface itself rather than to any interaction between the pile and shoot returns competing for the same steering angles, and this no-chute view is a useful companion figure to Figures 3.2, 3.4, 3.6, and 3.7 for that reason.

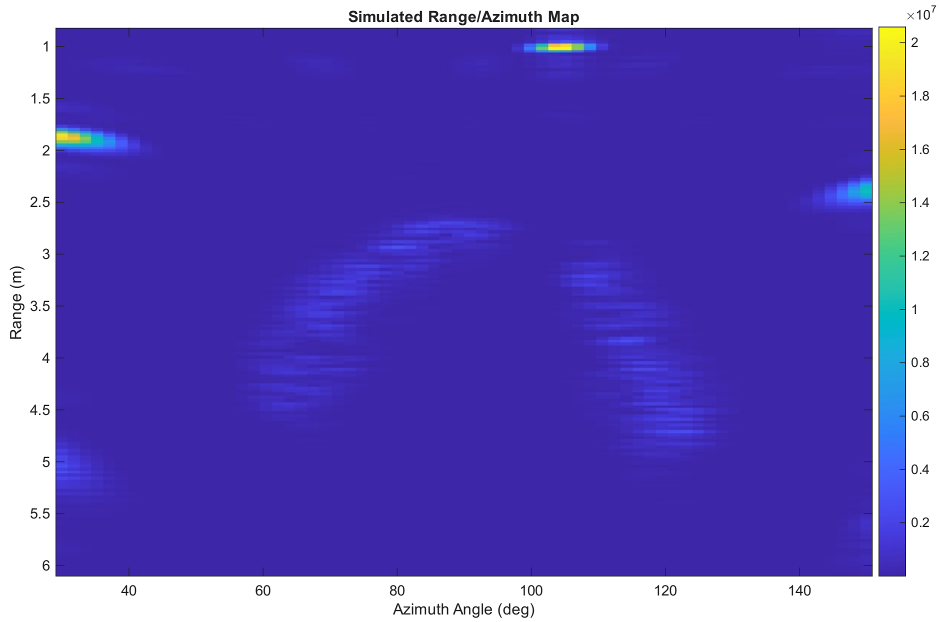

Figure 3.8: Simulated range and azimuth map with the grain chute cylinder removed from the scene. With the chute's trace absent, the pile return's central arc and the bar returns near the top of the frame are isolated from any chute related contribution, giving a cleaner view of the pile only return shape discussed above.

3.5 Line of Sight Obstruction

This section moves from the range-azimuth imagery to the underlying line-of-sight geometry that shapes the pile return discussed above, examining how self-shadowing evolves as the pile grows across the full run. The radar line-of-sight visibility plots, Figures 3.9-3.10, show the expected qualitative behavior. As fill level increases from medium to full, the blocked (red) region grows and the boundary between visible (green) and blocked terrain shifts further from the antenna, consistent with a taller, larger pile increasingly shadowing its own far side and base from the antenna's vantage point.

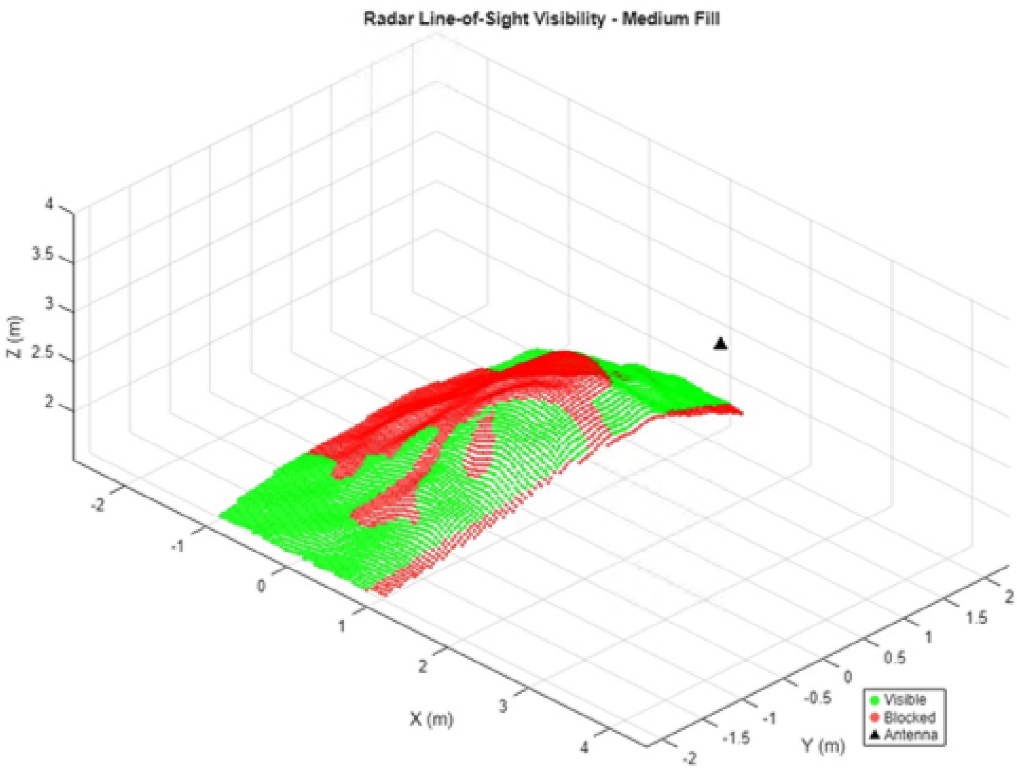

Figure 3.9: Simulated radar line-of-sight visibility at medium fill. Green points are visible from the antenna and red points are blocked by self shadowing as the pile obstructs its own lower and far-side surface. At medium fill the blocked region is fragmented into many small patches, consistent with the pile's surface ripple texture being large relative to the pile's overall size.

The X and Y axes in Figure 3.9 and 3.10 correspond directly to the cart's length and width dimensions defined, with Z representing height above the cart base. The antenna marker (black triangle) sits near the corner of the plotted volume corresponding to its actual mounted position, offset from center along the cart's length. The visibility computation records the total, visible, and blocked pile point counts at every time step. Figure 3.11 plots the resulting visible point fraction across the full 40 time step

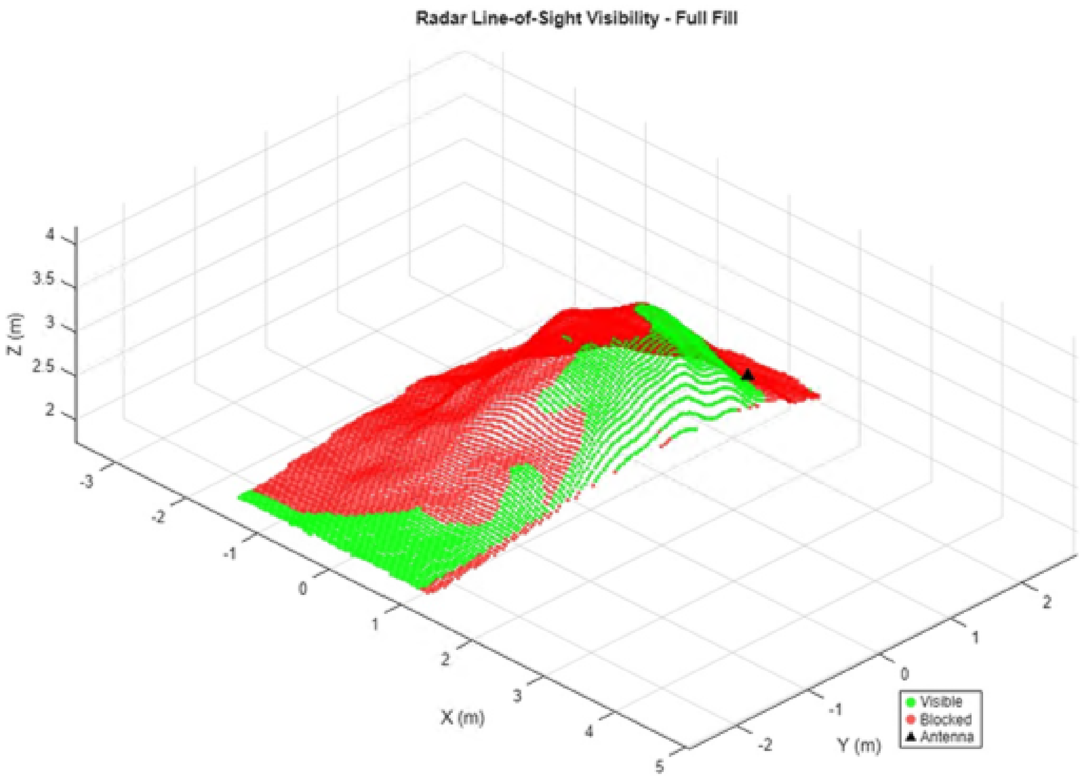

Figure 3.10: Simulated radar line-of-sight visibility at full fill. Green points are visible from the antenna and red points are blocked by self shadowing as the pile obstructs its own lower and far-side surface. At full fill the blocked region consolidates into one large, contiguous patch as a single dominant slope shadow comes to dominate over the fine ripple texture.

simulation run. The trend in Figure 3.11 is not the simple monotonic decline that Figure 3.9 and 3.10 alone might suggest. It has two distinct phases. In the first phase, time steps 1–28, visible fraction rises essentially monotonically from 0% to its run maximum of 56.6%, even as the pile is actively growing throughout this entire phase. This is consistent with the early pile being very small and low relative to the antenna's vantage point. At time step 1 the pile has only 681 candidate points and none are visible, most plausibly because the pile is still low enough that the antenna's line of sight to it grazes at a shallow, easily obstructed angle rather than because the pile is

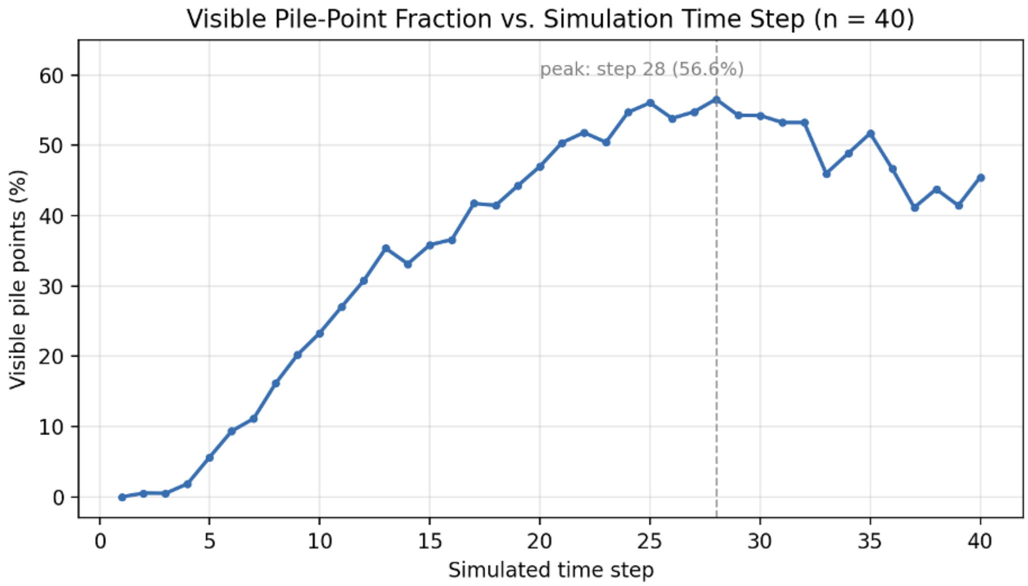

Figure 3.11: Visible pile point fraction across all 40 simulated time steps. The pile starts with 0% visibility at step 1 and rises steadily to a peak of 56.6% visible at step 28, before declining, with some local fluctuation. to 45.5% by the final step.

already self shadowing at this size. As the pile grows taller and broader through this phase, more of its surface comes into direct, unobstructed view of the antenna faster than the pile's own growing bulk can begin shadowing itself, so visible fraction climbs.

In the second phase, time steps 28–40, visible fraction declines with notable local fluctuation rather than a clean monotonic drop from the 56.6% peak down to a run minimum of 41.2% at step 37, before partially recovering to 45.5% by time step 40. This is the phase where the pile has grown large enough that self shadowing begins to dominate over the earlier visibility gain effect. Once the pile is tall and wide enough, its own near side surface increasingly blocks the antenna's view of its far side and base, exactly the mechanism illustrated for the medium and full fill cases in Figure 3.9 and 3.10. The fluctuation within this declining phase is consistent with the random ripple and per time step alteration terms in the pile surface model. The smoothing filter ripple factor and the randomness factor applied to pile height and shape at every time step locally reshapes the pile surface enough to expose or reshadow small regions from one time step to the next, superimposed on top of the larger underlying self shadowing trend rather than reversing it.

This two phase structure shows that visible fraction is not a simple substitute for "how much the pile has grown," but the result of two competing geometric effects. Increasing exposure to the antenna as the pile rises into clearer view, and increasing self obstruction as the pile becomes large enough to shadow itself. That trade off against each other over the course of the fill cycle, with the crossover occurring around step 28 in this run.

This growing blocked region is also the direct cause of the broadening pile arc shape observed in the simulated range and azimuth maps of Figures 3.2 and 3.4. Once a given pile region loses line-of-sight, its scatterers are excluded from the alpha reflectivity coefficient and no longer contribute to the radar return, regardless of how much grain is physically present there. In other words, the pile shape change documented is not purely a consequence of the pile growing larger in the Gaussian height field model. It is equally a consequence of which portion of that larger pile remains visible to the antenna at each fill level, making Figures 3.9 and 3.10 a necessary companion to, not just a parallel illustration alongside, Figures 3.1–3.8.

Figure 3.12 shows the simulated LOS obstructed view range and azimuth map for the same obstructed geometry discussed above. Two non-pile scatterers remain clearly identifiable as strong, well-localized returns: the support bars, appearing as bright spots near 30°, 100°, and 148° azimuth at roughly 1–2 m range, and the back edge of the cart, visible as fainter paired returns near 4.5–5 m range at the outer edges of the azimuth sweep.

Both non pile scatterers remain strong and well-localized in this view despite the loss of line-of-sight to the pile surface itself. Which is the expected result given the

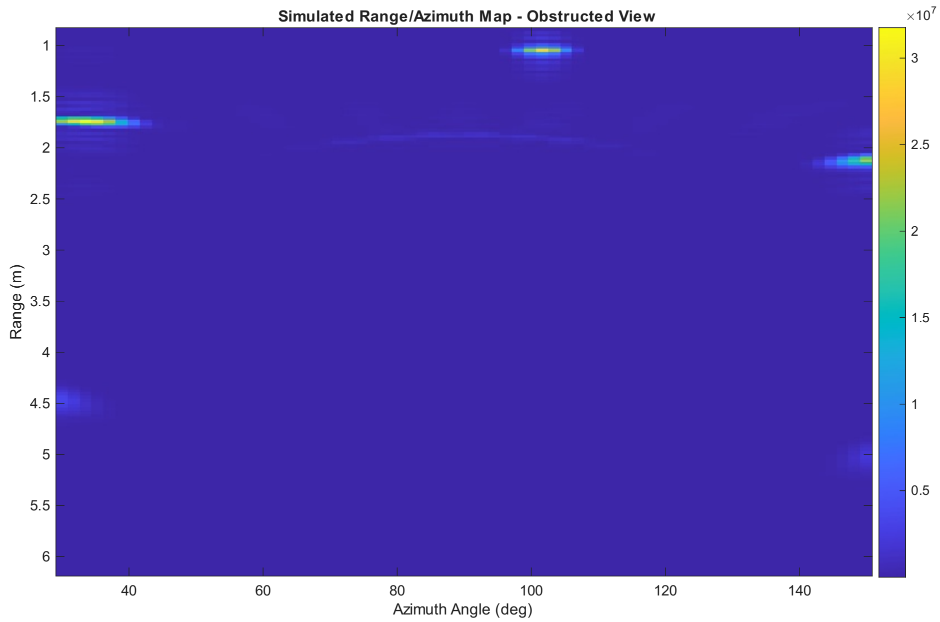

Figure 3.12: Simulated LOS obstructed view range and azimuth map. The bright spots near 30°, 100°, and 148° azimuth at roughly 1–2 m range correspond to the support bars. The fainter paired returns near 4.5–5 m range, at the outer edges of the azimuth sweep, correspond to the back edge of the cart.

scene geometry rather than a contradiction of the LOS model. The bars and the back edge of the cart both sit above and along paths that are largely independent of the pile's own self shadowing, so their visibility to the antenna is governed mainly by the cart geometry rather than by how much of the pile surface is currently blocked. This is a useful internal consistency check on the simulation itself. The model correctly continues to render the bar returns in the obstructed view region, confirming that the LOS gating is applied only to the pile's own reflectivity and not inadvertently suppressing the other three scene components.

3.6 Antenna Position Estimation

The final results section shifts from the scattering source comparisons above to the antenna self-localization algorithm introduced in Section 2.9, quantifying how accurately it recovers the antenna's true position from the bar and edge returns alone. Section 2.9 introduced the antenna self localization problem and the estimation function implemented to solve it, anchored on the known, fixed geometry of the support bars and the cart top edges as well. To evaluate how well this estimation approach recovers the antenna's true position, the actual and estimated antenna coordinates were logged across 40 simulated time steps, along with the corresponding per step error.

At every time step, the antenna's actual position is randomly moved from its nominal mounted location before the estimation algorithm is run against it. This random movement is a deliberate modeling choice, not simulation noise to be minimized away. A radar antenna mounted on a grain tractor auger is not rigidly fixed in the same way a lab antenna would be. The cart flexes, vibrates, and shifts slightly as it is towed across the field and as grain is unloaded into it, so the antenna's true position at any given moment will differ slightly from its nominal mounted location. Randomizing the actual antenna position independently at every time step is intended to mimic this real world non stationary, so that the estimation algorithm is tested against a moving target rather than against a single fixed, idealized position it could simply be tuned to.

Table 3.1 summarizes the per axis and overall 3D position error across all 40 time steps. Errors are reported as estimated minus actual, in meters. The Y-axis estimate is by far the most accurate and most consistent of the three. Its RMSE (0.118 m) and standard deviation (0.075 m) are both roughly a third of the corresponding X and Z axis values, indicating the algorithm recovers the antenna's Y coordinate both more precisely and more repeatably than its X or Z coordinate. The Y axis error is

Axis	Mean Error (m)	MAE (m)	RMSE (m)	Std. Dev. (m)	Max |Error| (m)
X	0.014	0.391	0.439	0.439	0.909
Y	-0.091	0.102	0.118	0.075	0.214
Z	-0.210	0.299	0.394	0.333	0.933
3D distance	-	0.573	0.602	0.184	0.989

Table 3.1: Antenna position estimation error statistics across 40 simulated time steps. MAE = mean absolute error RMSE = root-mean-square error Std. Dev. = Standard Deviation

also the most clearly biased of the three. Its mean error is –0.091 m with a small standard deviation, meaning the estimator is consistently low by roughly the same amount nearly every time step, rather than scattering randomly above and below the true value.

The X and Z axis estimates show substantially larger errors, RMSE of 0.439 m and 0.394 m respectively, and with a wider spread, standard deviations of 0.439 m and 0.333 m. With maximum absolute errors approaching 0.93 m on both axes. Unlike the Y axis, the mean X axis error is close to zero (0.014 m), indicating the X axis estimate is not systematically biased in either direction Its larger RMSE comes from high per step variance rather than a consistent offset. The Z axis error sits in between, with a moderate negative mean bias (–0.210 m) combined with substantial spread.

Figure 3.14 plots the full 3D Euclidean position error at each of the 40 time steps. The error fluctuates between roughly 0.22 m and 0.99 m with no clear upward or downward trend across the run, averaging 0.573 m (median 0.588 m, RMSE 0.602 m). The absence of a trend across time steps suggests the estimator's accuracy does not systematically degrade or improve as the simulation progresses, and that the per step error is dominated by the same few percentage points of geometry peak detection error at every time step rather than by any cumulative drift.

These results indicate that the bar and edge anchored estimation approach recovers the antenna's Y position reliably but is noticeably less precise along X and Z,

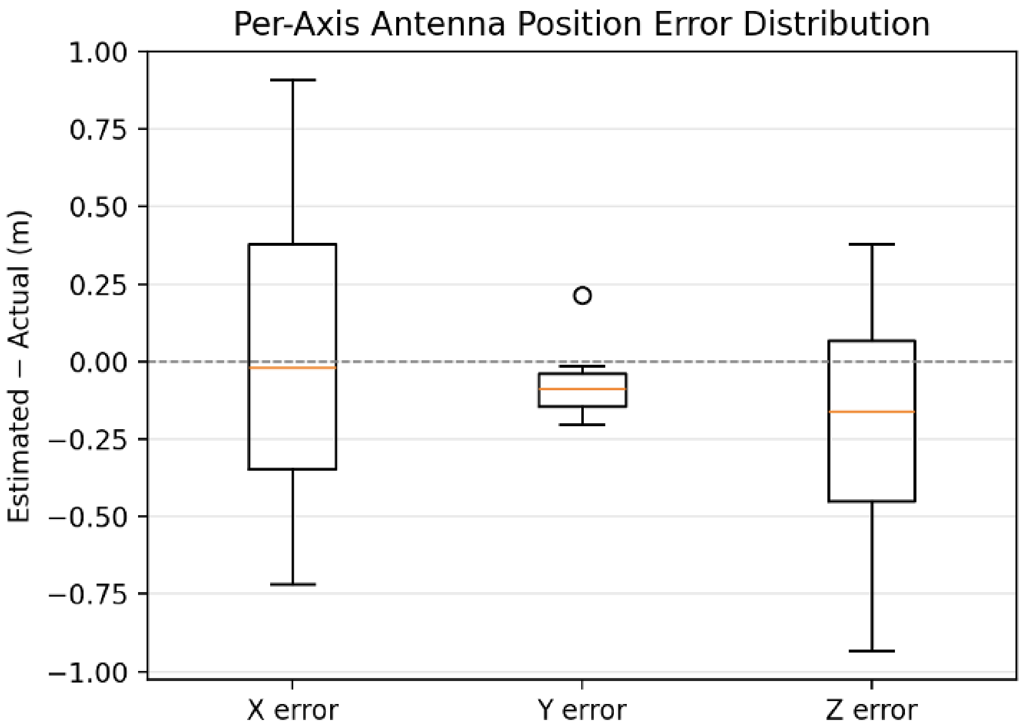

Figure 3.13: Distribution of per axis antenna position error across 40 time steps. The Y axis error is tightly clustered around a small negative bias. The X and Z axis errors show wider spread and larger maximum deviations.

with sub meter but not sub decimeter accuracy in the worst cases. Given that this estimated position feeds directly into the pile height and range profile calculations described, propagating this 3D error distribution forward into an estimated pile height error budget.

The results in Table 3.1 characterize the antenna self-localization algorithm's accuracy over 40 simulated time steps at a single, growing pile height. To assess whether this accuracy holds as the grain pile geometry changes, since a taller or shorter pile changes both the visibility of the bar and edge reference targets and the multipath/occlusion environment the radar operates in, the same estimation procedure was additionally run as a 100-trial Monte Carlo study at three representative pile heights. At each trial, the antenna's actual position was again randomly perturbed from its nominal mounted

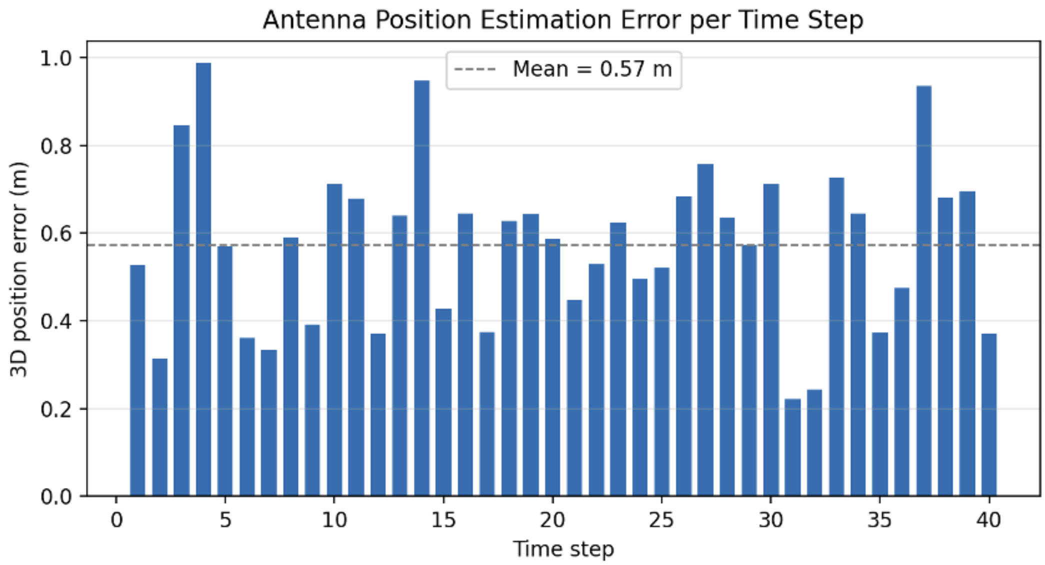

Figure 3.14: 3D antenna position estimation error at each of 40 simulated time steps, with the run mean (0.573 m) marked. Error fluctuates without a clear trend, consistent with a per step estimation error rather than a cumulative drift.

location, and the estimator was run against a freshly regenerated set of bar and edge returns for that trial's true position, rather than reusing a single fixed measurement. This tests the estimator against a new true position and a new radar realization at every trial, rather than testing only the optimizer's convergence from different starting guesses on identical data.

Table 3.2 summarizes the per-axis and 3D position error for the low pile height case across 100 trials, using the same procedure described above. At the low pile height,

Axis	Mean Error (m)	MAE (m)	RMSE (m)	Std. Dev. (m)	Max |Error| (m)
X	0.009	0.285	0.399	0.329	0.767
Y	0.040	0.092	0.173	0.023	0.150
Z	–0.118	0.120	0.256	0.290	0.864
3D distance	-	0.332	0.573	0.116	0.899

Table 3.2: Antenna position estimation error statistics for the low pile height, across 100 trials with a freshly regenerated true antenna position and radar realization at each trial. MAE = mean absolute error, RMSE = root-mean-square error, Std. Dev. = Standard Deviation.

the estimator is close to unbiased on the X axis (mean error 0.009 m) but shows a moderate negative bias on Z (–0.118 m) and a small positive bias on Y (0.040 m). The X and Z axes carry most of the error and the most spread, while the Y axis is by far the most accurate and most consistent of the three, with a MAE of only 0.092 m and the smallest maximum error of any axis (0.150 m). Overall 3D distance error averaged 0.332 m (RMSE 0.573 m), noticeably lower than either the medium or full pile height cases discussed below.

Table 3.3 summarizes the per-axis and 3D position error for the medium pile height case across 100 trials, using the same procedure described above. At the medium pile

Axis	Mean Error (m)	MAE (m)	RMSE (m)	Std. Dev. (m)	Max |Error| (m)
X	0.035	0.467	0.579	0.564	0.988
Y	–0.125	0.176	0.255	0.098	0.237
Z	0.243	0.217	0.378	0.421	0.928
3D distance	-	0.646	0.732	0.218	0.956

Table 3.3: Antenna position estimation error statistics for the medium pile height, across 100 trials with a freshly regenerated true antenna position and radar realization at each trial. MAE = mean absolute error, RMSE = root-mean-square error, Std. Dev. = Standard Deviation.

height, error grows on every axis relative to the low pile height case. The X axis shows both a larger MAE (0.467 m vs. 0.285 m) and substantially more spread, and now carries the largest maximum error of the three axes (0.988 m). The Z axis similarly grows in both mean bias (0.243 m vs. –0.118 m) and spread. The Y axis remains the most accurate and consistent axis, as it was at the low pile height, though its own error grows as well. Overall 3D distance error averaged 0.646 m (RMSE 0.732 m), roughly double the low pile height's mean 3D error.

Table 3.4 summarizes the per-axis and 3D position error for the full pile height case across 100 trials. At the full pile height, error increases sharply relative to both the low and medium cases on every axis. Overall 3D distance error averaged 1.532 m (RMSE

Axis	Mean Error (m)	MAE (m)	RMSE (m)	Std. Dev. (m)	Max |Error| (m)
X	1.021	1.021	1.024	0.078	1.442
Y	0.527	0.570	0.711	0.479	1.564
Z	0.909	0.909	0.910	0.024	0.997
3D distance	-	1.532	1.543	0.189	2.093

Table 3.4: Antenna position estimation error statistics for the full pile height, across 100 trials with a freshly regenerated true antenna position and radar realization at each trial. MAE = mean absolute error, RMSE = root-mean-square error, Std. Dev. = Standard Deviation.

1.543 m), roughly 2.4 times the medium pile height's mean 3D error and 4.6 times the low pile height's.

Error increases monotonically with pile height across all three cases, from a mean 3D error of 0.332 m at the low pile height, to 0.646 m at the medium pile height, to 1.532 m at the full pile height. A plausible explanation is that as the pile grows taller, it presents a larger scattering surface directly in or near the radar's line of sight to the bar and edge reference targets, adding more clutter and multipath reflection into each range profile. The resulting increase in noise around the true bar and edge returns would be expected to make the peak-detection less reliable, directly increasing the position error the optimizer converges to. This is consistent with the change in error character observed between the low and medium pile heights, where growth in error is accompanied by growth in standard deviation on the X and Z axes (0.329 m to 0.564 m, and 0.290 m to 0.421 m, respectively) rather than by a shift in mean error alone. The Y axis follows the same trend at a smaller scale, remaining the most accurate and consistent axis at both pile heights while still degrading somewhat as pile height increases.

3.7 Limitations

Having presented the comparison results, this final section of the chapter states the model's key simplifications explicitly, both to bound how the preceding results should be interpreted and to set up the prioritized future work roadmap in Chapter 4. Several simplifications in the current model are worth stating explicitly, both because they bound how the results and should be interpreted because each one is a concrete, addressable target for the future work.

The grain chute is modeled as a single rigid cylinder of fixed radius and fixed reflectivity, which produces a clean, symmetric ring shaped return in the simulated imagery rather than the broad, diffuse, turbulent curtain return seen in the measured data. The three support bars are likewise assigned a single uniform reflectivity along their entire length, which produces an artificially smooth, constant brightness bar arc in simulation, compared to the visibly weathered, spatially varying bar returns in the measured imagery. The cart floor and side walls are included in the combined triangulated mesh for visualization and line-of-sight purposes but contribute no backscatter term to the radar signal model. This is the single largest identified source of simulated versus measured mismatch and is therefore a component already present geometrically in the model but not yet connected to the signal chain.

The dust cloud is represented purely as a magnitude attenuation and reflection term with a flat, round number two-way attenuation rather than one derived from an estimated grain dust particle size distribution and the Mie-scattering literature reviewed. Separately, the LOS visibility calculation considers only self obstruction of the pile by its own surface. The pile point cloud is subsampled to a small number of target points per time step for computational tractability, which limits the spatial detail available to the range profile and antenna localization calculations relative to the full pile surface. The antenna is treated as effectively static within each time step, with only a single random position movement applied per step to emulate real world mounting vibration and platform motion. There is no continuous motion or vibration modeled within a single 0.5 s time step, only a discrete jump between steps.

Chapter 4

Conclusion

This thesis set out to make three contributions. A combined line-of-sight obstruction and volumetric dust attenuation propagation model for a dynamic agricultural radar scene. Three independent antenna self localization algorithms anchored on the cart's structural bars and a quantitative comparison between simulated and measured radar returns identifying the dominant sources of image degradation in this environment. All three were demonstrated, with the results in Chapter 3 also clarifying their practical scope and limits.

The combined LOS obstruction and dust model was shown to reproduce the expected qualitative self shadowing behavior as the pile grows with visible point fraction first rising as the small early pile comes into clearer view of the antenna, then falling as the larger, later pile increasingly shadows itself, across the full 40 time step run. The dust component of the model, while functioning as intended, could not be independently validated against the measured data in this thesis, because the measured imagery's un modeled ground and wall clutter is large enough to obscure any comparably sized dust signature. This is reported as a finding in its own right rather than treated as a successful validation.

An antenna self localization algorithms were implemented, with the line segment and nonlinear least squares functions evaluated quantitatively across 40 simulated time steps. It achieves sub meter 3D position accuracy on average (RMSE 0.602 m) with no evidence of drift over time, and notably better precision on the antenna's Y coordinate than on X or Z.

The quantitative simulated versus measured comparison identified un modeled cart floor and wall backscatter as the dominant source of image quality mismatch between the two datasets. Secondary mismatches were identified and explained for the grain shoot return, the support bar returns, and the dust return. Together, these findings give the propagation model in this thesis a specific, prioritized roadmap for improvement.

4.1 Future Work

The clearest near term extension is adding a cart floor and cart wall backscatter term to the radar signal model, using the combined triangulated mesh's existing cart shell geometry, which the comparison in Chapter 3 identifies as the single highest leverage addition for closing the simulated versus measured gap. Once this is in place, an empty cart calibration scan would let the floor and wall return be characterized and subtracted, which would in turn make it possible to finally isolate and validate the dust model's contribution against real data, something the current measured dataset's clutter floor does not permit.

A second priority is replacing the grain chute's rigid cylinder model with a distributed, time varying scatterer. Conceptually similar to the existing dust voxel model. To better capture the broad, turbulent return observed in the measured data. A similar refinement could replace the bars' uniform reflectivity assumption with a spatially varying one, informed by the measured bar return intensity pattern itself, to better match the real weathered appearance of the bar arcs.

Finally, on the dust side specifically, the attenuation assigned should be replaced with a value derived from an estimated grain dust particle size distribution and benchmarked against the Mie-scattering formulas reviewed, and the model's magnitude only treatment of dust should eventually be extended to include the phase distortion and decoherence mechanism documented by Kang et al., which the current voxel model has no analogue for and which may matter more than amplitude attenuation alone for coherent processing.

Reference List

[1] J. Shang, H. McNairn, C. Champagne, and X. Jiao, "Contribution of multi-frequency, multi-sensor, and multi-temporal radar data to operational annual crop mapping," in *Proceedings of the IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, Agriculture and Agri-Food Canada. Boston, MA, USA: IEEE, 2008, pp. III–378–III–381.

[2] W. Tan, A. Sinha, Y. Li, L. Ma, and J. Li, "Early-season crop classification with RADARSAT-2 polarimetric synthetic aperture radar imagery," in *Proceedings of the IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*. IEEE, 2020, pp. 4167–4170.

[3] H. McNairn, A. Smith, E. Huffman, I. Jarvis, A. Pacheco, and E. Gauthier, "A national inventory of land management practices: Estimating soil conservation practices using optical and radar imagery," in *Proceedings of the IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, Agriculture and Agri-Food Canada. IEEE, 2004, pp. 1629–1632.

[4] D. Han, H. Yang, G. Yang, and C. Qiu, "Monitoring model of corn lodging based on Sentinel-1 radar image," in *Proceedings of the IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*. Beijing, China: IEEE, 2017.

[5] C. G. Boryan, Z. Yang, A. Sandborn, P. Willis, and B. Haack, "Operational agricultural flood monitoring with Sentinel-1 synthetic aperture radar," in *Proceedings of the IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, USDA National Agricultural Statistics Service. IEEE, 2018, pp. 5831–5834.

[6] M. S. Sanjiwani, M. W. R. Dzulizar, A. A. Pramudita, H. H. Ryanu, V. Mappadang, and A. R. Dwinanda, "Ground vehicle SFCW radar for soil water content sensing and mapping," in *Proceedings of the 2025 International Conference on Converging Technology in Electrical and Information Engineering (ICCTEIE)*. Bandung, Indonesia: IEEE, 2025, pp. 141–146.

[7] D. Henry, H. Aubert, T. Véronèse, and É. Serrano, "Remote estimation of intra-parcel grape quantity from three-dimensional imagery technique using ground-based microwave FMCW radar," *IEEE Instrumentation & Measurement Magazine*, vol. 20, no. 3, pp. 20–24, Jun. 2017.

[8] N. S. Ahmad and J. J. Soraghan, "A preliminary study towards the implementation of MIMO radar system for agriculture applications," in *Proceedings of the 35th International Conference on Infrared, Millimeter, and Terahertz Waves*. Kangar, Malaysia: University of Strathclyde / Universiti Malaysia Perlis, 2010.

[9] G. Brooker, S. Scheding, R. Hennessy, and M. Bishop, "Millimetre wave radar sensors for mining applications," in *Radar 2003*. IEEE, 2003, pp. 212–218.

[10] Iowa Agriculture Literacy Foundation, "Grain cart: What is it and why do farmers use them?" https://iowaagliteracy.wordpress.com/2018/11/14/grain-cart-what-is-it-and-why-do-farmers-use-them/, 11 2018, accessed: [05/01/26].

[11] Demco Products, "How to run a grain cart," https://www.demco-products.com/blog/how-to-run-a-grain-cart, accessed: [05/01//26].

[12] MathWorks, "delaunaytriangulation," https://www.mathworks.com/help/matlab/ref/delaunaytriangulation.html, mATLAB Documentation. Accessed: [03/15/25].

[13] N. N. Mohsenin, *Physical Properties of Plant and Animal Materials*. New York, NY, USA: Gordon and Breach Science Publishers, 1970.

[14] MathWorks, "txsite.los," https://www.mathworks.com/help/antenna/ref/txsite.los.html, mATLAB Antenna Toolbox Documentation. Accessed: [03/15/25].

[15] A. S. Ahmed, "Role of particle-size distributions on millimetre-wave propagation in sand/duststorms," *IEE Proceedings H – Microwaves, Antennas and Propagation*, vol. 134, no. 1, pp. 55–59, 2 1987.

[16] Z. Elabdin, M. R. Islam, O. O. Khalifa, H. E. A. Raouf, and M. J. E. Salami, "Development of mathematical model for the prediction of microwave signal attenuation due to duststorm," in *Proceedings of the International Conference on Computer and Communication Engineering (ICCCE) 2008*. Kuala Lumpur, Malaysia: IEEE, May 2008, pp. 1156–1160.

[17] E. A. A. Elsheikh, M. R. Islam, A. H. M. Z. Alam, A. F. Ismail, K. Al-Khateeb, and Z. Elabdin, "The effect of particle size distributions on dust storm attenuation prediction for microwave propagation," in *Proceedings of the International Conference on Computer and Communication Engineering (ICCCE)*, Kuala Lumpur, Malaysia, 5 2010, pp. 1–5.

[18] D. Wikner, "Millimeter-wave propagation measurement through a dust tunnel," U.S. Army Research Laboratory, Adelphi, MD, USA, Tech. Rep. ARL-TR-4399, 3 2008.

[19] H. Kang, X. Zeng, H. Zhao, J. Li, Q. Guo, and M. Martorella, "Impact of sand and dust storms on radar performance: Attenuation and phase distortion analysis with field experiments," *IEEE Geoscience and Remote Sensing Letters*, vol. 23, p. 3501105, 2026.

[20] M. I. Skolnik, *Introduction to Radar Systems*, 3rd ed. New York: McGraw-Hill, 2001.

[21] M. A. Richards, *Fundamentals of Radar Signal Processing*, 2nd ed. New York: McGraw-Hill Education, 2014.

[22] H. Jiang, G. Liao, and Y. Qu, "Compensation of clutter spectrum for forward-looking radar based on the spatial steering vector fitting," in *Proceedings of 2009 IET International Radar Conference*. Xi'an, China: National Key Laboratory of Radar Signal Processing, Xidian University, 2009.

[23] MathWorks, "lsqnonlin," https://www.mathworks.com/help/optim/ug/lsqnonlin.html, mATLAB Optimization Toolbox Documentation. Accessed: [07/25/24].