A millimeter-wave radar system for plot-scale canopy phenotyping in breeding trials

by

Garrett Peterson

B.S., Kansas State University, 2015

M.S., Kansas State University, 2020

AN ABSTRACT OF A DISSERTATION

submitted in partial fulfillment of the
requirements for the degree

Doctor of Philosophy

Mike Wiegers Department of Electrical and Computer Engineering
Carl R. Ice College of Engineering

KANSAS STATE UNIVERSITY
Manhattan, Kansas

2026

Abstract

Plant breeding programs are increasingly limited not by the rate at which new genotypes can be generated but by the rate at which their physical traits can be evaluated in the field, a disparity commonly described as the phenotyping bottleneck. Structural traits such as canopy height and density are strongly tied to yield yet remain slow and labor-intensive to acquire, and the established sensing modalities used to automate their measurement have limitations. Optical imaging and LiDAR are sensitive to ambient conditions and interact primarily with the outermost canopy surface. Radar offers a complementary alternative by measuring range directly and independently of illumination. It can also partially penetrate a canopy. Compact, plot-scale, millimeter-wave radar instruments designed for phenotyping, together with a processing pipeline to turn their raw output into per-plot traits need to be developed.

This dissertation presents the design, implementation, and field evaluation of a compact 120 GHz stepped-frequency continuous-wave (SFCW) radar for high-throughput soybean phenotyping. The instrument developed is a low-power, portable sensor built around a highly integrated millimeter-wave transceiver and a custom embedded acquisition system, with a processing pipeline that converts its raw range-time acquisitions into per-plot canopy height and structural density measures. It employs a continuous coherent calibration technique that cancels the transmit-receive leakage that often limits compact continuous-wave radars, preserving the dynamic range needed to detect weak canopy and ground returns. The processing pipeline removes the system's characteristic artifacts and applies a particle filter with a backward smoother to track the ground through canopy shadowing, a region-based detector to extract the canopy relative to that ground, and a radar integrated energy measure as a proxy for canopy structural density, all within a calibrate-once human-in-the-loop workflow.

The system was deployed in a soybean breeding trial at Kansas State University across the 2025 growing season. Over five measurement sessions spanning emergence to maturity, the radar tracked canopy development and produced height estimates in agreement with manual ground truth measurements, including close agreement under a fully closed canopy. An experiment established that the height estimate is both repeatable and consistent with hand measurement, and the integrated-energy density measure agreed with an independent optical canopy-cover estimate at a Pearson correlation of $r = 0.837$. The semi-automated workflow transferred a single tuned parameter set across measurement batches within a growth stage and acquired continuous along-row height profiles at a sampling density that manual measurement cannot practically match, roughly eight times faster than hand measurement at equivalent density, while concentrating the demand on skilled labor into a single tuning step. The results demonstrate that proximal millimeter-wave radar is a viable and complementary instrument for high-throughput structural phenotyping.

A millimeter-wave radar system for plot-scale canopy phenotyping in breeding trials

by

Garrett Peterson

B.S., Kansas State University, 2015

M.S., Kansas State University, 2020

A DISSERTATION

submitted in partial fulfillment of the
requirements for the degree

DOCTOR OF PHILOSOPHY

Mike Wiegers Department of Electrical and Computer Engineering
Carl R. Ice College of Engineering

KANSAS STATE UNIVERSITY
Manhattan, Kansas

2026

Approved by:

Major Professor
Don Gruenbacher

Copyright

© Garrett Peterson 2026.

Abstract

Plant breeding programs are increasingly limited not by the rate at which new genotypes can be generated but by the rate at which their physical traits can be evaluated in the field, a disparity commonly described as the phenotyping bottleneck. Structural traits such as canopy height and density are strongly tied to yield yet remain slow and labor-intensive to acquire, and the established sensing modalities used to automate their measurement have limitations. Optical imaging and LiDAR are sensitive to ambient conditions and interact primarily with the outermost canopy surface. Radar offers a complementary alternative by measuring range directly and independently of illumination. It can also partially penetrate a canopy. Compact, plot-scale, millimeter-wave radar instruments designed for phenotyping, together with a processing pipeline to turn their raw output into per-plot traits need to be developed.

This dissertation presents the design, implementation, and field evaluation of a compact 120 GHz stepped-frequency continuous-wave (SFCW) radar for high-throughput soybean phenotyping. The instrument developed is a low-power, portable sensor built around a highly integrated millimeter-wave transceiver and a custom embedded acquisition system, with a processing pipeline that converts its raw range-time acquisitions into per-plot canopy height and structural density measures. It employs a continuous coherent calibration technique that cancels the transmit-receive leakage that often limits compact continuous-wave radars, preserving the dynamic range needed to detect weak canopy and ground returns. The processing pipeline removes the system's characteristic artifacts and applies a particle filter with a backward smoother to track the ground through canopy shadowing, a region-based detector to extract the canopy relative to that ground, and a radar integrated energy measure as a proxy for canopy structural density, all within a calibrate-once human-in-the-loop workflow.

The system was deployed in a soybean breeding trial at Kansas State University across the 2025 growing season. Over five measurement sessions spanning emergence to maturity, the radar tracked canopy development and produced height estimates in agreement with manual ground truth measurements, including close agreement under a fully closed canopy. An experiment established that the height estimate is both repeatable and consistent with hand measurement, and the integrated-energy density measure agreed with an independent optical canopy-cover estimate at a Pearson correlation of $r = 0.837$. The semi-automated workflow transferred a single tuned parameter set across measurement batches within a growth stage and acquired continuous along-row height profiles at a sampling density that manual measurement cannot practically match, roughly eight times faster than hand measurement at equivalent density, while concentrating the demand on skilled labor into a single tuning step. The results demonstrate that proximal millimeter-wave radar is a viable and complementary instrument for high-throughput structural phenotyping.

Table of Contents

List of Figures xi

List of Tables xiii

Acknowledgments xiv

Dedication xv

1 Introduction 1

1.1 The Phenotyping Bottleneck 1

1.2 Manual Phenotyping Measurements 3

1.3 Limitations of Existing Sensing Modalities 4

1.4 Radar as a Complementary Modality 5

1.5 Problem Statement 6

2 Literature Review 9

2.1 High-Throughput Phenotyping Platforms 9

2.2 Optical and Photogrammetric Approaches 10

2.3 LiDAR for Structural Phenotyping 11

2.4 Radar in Agriculture 12

2.5 Millimeter-Wave Sensing 12

2.6 Traditional Radar Systems 13

2.7 Traditional Radar Signal Processing 15

3 Radar Phenotyping System 20

3.1 Requirements . . . 21
3.2 System Architecture . . . 21
3.3 Continuous Coherent Calibration . . . 23
3.4 Surface Testing . . . 26
3.5 Breeding Trial Data Collection . . . 29
3.5.1 Early Growth - June 27^{th} . . . 30
3.5.2 Canopy Development - July 16^{th} . . . 33
3.5.3 Mid-Season Canopy - July 29^{th} . . . 37
3.5.4 Reproductive Growth - August 15^{th} . . . 40
3.5.5 Late Season - August 21^{st} . . . 40
3.6 Summary . . . 42

4 Phenotyping Radar Processing Pipeline . . . 44
4.1 Overview . . . 45
4.2 Range–Time Image Formation . . . 46
4.3 Calibration-Sweep Removal and Segment Structure . . . 46
4.4 Cleanup Chain . . . 47
4.4.1 Fixed-Clutter and Artifact Suppression . . . 47
4.4.2 Multilook Averaging . . . 49
4.4.3 Per-Sweep Normalization . . . 50
4.4.4 Cleanup Parameters . . . 51
4.5 Tracking and Detection . . . 51
4.5.1 Ground Tracking Particle Filter . . . 51
4.5.2 Canopy Region Detection . . . 54
4.6 Integrated Canopy Energy . . . 57
4.7 Evaluation . . . 57
4.7.1 Workflow Demonstration . . . 58
4.7.2 Integrated Energy and Optical Canopy Cover . . . 59

4.7.3 Measurement Cost and Information Density 60
4.7.4 Value of the Spatial Profile 62
4.7.5 Cross-Stage Limitations 63
4.8 Summary 63

5 Conclusions and Future Work 65
5.1 Summary of Contributions 65
5.2 System-Level Conclusions 67
5.3 Limitations 68
5.4 Future Work 69
5.4.1 Automated Parameter Selection 69
5.4.2 Platform Stabilization and Self-Referencing 69
5.4.3 Broader Validation 70
5.4.4 Multi-frequency and Polarimetric Extension 70
5.4.5 Integration into Operational Phenotyping 70

Bibliography 71

A Radar Schematic 76

List of Figures

1.1 Representative layout of a soybean breeding trial 2

1.2 Soybean anatomy through growth stages and canopy closure 3

2.1 Ideal radar block diagrams 14

2.2 CW distance radar operations 15

2.3 Range-time image formation example 17

2.4 Particle filter example operation 18

3.1 Phenotyping radar 20

3.2 120 GHz phenotyping research radar 21

3.3 Continuous coherent calibration flow diagram[34] 25

3.4 Radar multi surface test range-time image 27

3.5 Surface test sample video frames 28

3.6 Overhead soybean season growth 30

3.7 Early Growth - June 27^{th} - Measurement path route 31

3.8 Early Growth - June 27^{th} - First 2 row (1811.1-1911.1) panorama 31

3.9 Early Growth - June 27^{th} - Range-time images 32

3.10 Canopy Development - July 16^{th} - Range-time images 34

3.11 Canopy Development - July 16^{th} - Plot 1811 panoramic reconstructions . . . 34

3.12 Canopy Development - July 16^{th} - Plot 1919 range-time image 36

3.13 Canopy Development - July 16^{th} - Plot 1919 row 4 reconstructed panorama . 36

3.14 Mid-Season Canopy - July 29^{th} - Row 1917.2 reconstructed panoramic 37

3.15 Mid-Season Canopy - July 29^{th} - Row 1917.2 repeated pass range-time image 38

3.16 Mid-Season Canopy - July 29^{th} - Row 1917.2 height measurement vs. radar estimate comparison 39
3.17 Reproductive Growth - August 15^{th} - Row 1919.1 41

4.1 Processing pipeline flow diagram 45
4.2 Starting range-time image example 47
4.3 Slow-time filter example 48
4.4 Mirror subtract example 49
4.5 Range band nulling example 49
4.6 Multilook averaging example 50
4.7 Normalization example 50
4.8 Ground tracker example 54
4.9 Region detection example 56
4.10 Bad kernel size example 58
4.11 Tracking/Detection applied to large data sets 59
4.12 Radar integrated energy vs. video canopy estimate 61

List of Tables

3.1 Typical Operating Conditions 23

3.2 Multi surface test points of interest 26

3.3 Experiment summary table 29

3.4 Early Growth - June 27^{th} - Early plot heights 31

3.5 Canopy Development - July 16^{th} - Height results (plot averages) 35

3.6 Canopy Development - July 16^{th} - Measurement results (genotype averages) 35

3.7 Mid-Season Canopy - July 29^{th} - Row 1917.2 repeated pass statistics 38

4.1 Ground tracker example parameters 55

4.2 Region detection example parameters 57

4.3 July 16^{th} - Radar Energy, Video Canopy (Plot Averages) 60

4.4 July 16^{th} - Radar Energy, Video Canopy (Genotype Averages) 61

Acknowledgments

I would like to thank the members of my committee — Dr. Don Gruenbacher, Dr. William Kuhn, Dr. William Schapaugh, and Dr. Dan Flippo — for their support throughout the development and testing of this work. The many discussions I had with them shaped the direction of the project, and the physical resources and field access they provided made it possible to design, build, and evaluate the radar system presented here. I am grateful for their time, expertise, and generosity.

I also acknowledge the use of Claude (Anthropic) in the course of this work. It was used to assist with the development of the underlying data-processing software, providing drafting and language-editing assistance, and figures for portions of this document. All AI-assisted material was reviewed, tested, and edited by the author, who takes full responsibility for the content of this dissertation.

Dedication

To my beautiful and supportive family, who gave me inspiration and motivation whenever I needed it on this journey.

Chapter 1

Introduction

1.1 The Phenotyping Bottleneck

Modern plant breeding programs rely on the evaluation of phenotypic traits across large populations to identify and select elite varieties. The ability of breeding programs to produce genetic improvements is not only limited by the rate at which new genotypes can be generated, but also by the rate at which the physical traits can be evaluated in the field. Advances in genomics have driven the cost of generating genetic diversity steadily downward, so that breeding programs can now produce candidate genotypes faster than their resulting phenotypes can be characterized. The result is a structural imbalance with genotypes being cheaply and quickly produced but the phenotypes (traits under selection) slow and expensive to acquire.

This imbalance is widely described as the phenotyping bottleneck and has become a central constraint on the throughput of modern breeding programs. Typical breeding trials evaluate many genotypes side by side under common field conditions. Genotypes are grown in a small managed unit called a plot. As field conditions vary spatially from soil, drainage, and micro-climates across a field the genotypes are replicated strategically appearing in several plots at different positions to average and lessen the impact of environmental conditions. A plot is typically made up of several adjacent crop rows planted with fixed spacing and plots

are grouped into blocks separated by alleys that provide access for planting, management, and measurement. The number of plots in a trial is the number of genotypes multiplied by the number of replicates. Modern breeding programs may evaluate thousands or tens of thousands of plots in a single season. Figure 1.1 illustrates a representative breeding trial layout of five genotypes with four replications over 20 plots.

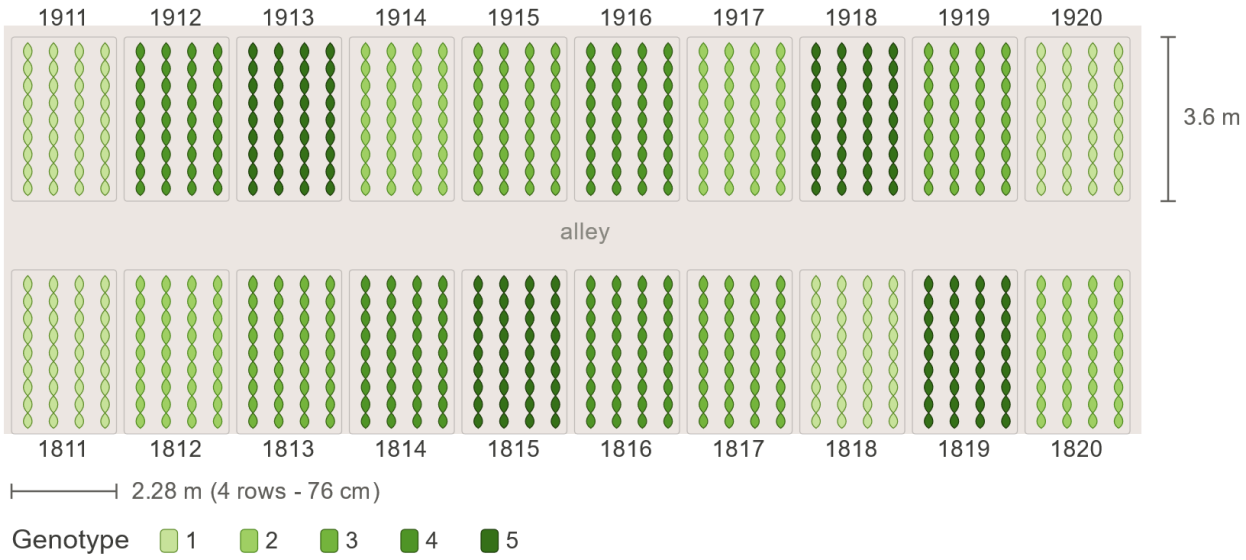

Figure 1.1: Representative layout of a soybean breeding trial

A common trait of interest is the canopy height and its development over the growing season which is informative on broader structural features of the crops. Canopy height can correlate with biomass accumulation, lodging risk, and yield potential. The canopy height trajectory over the growing period can reflect underlying differences in developmental timing among genotypes. Capturing the trajectory requires repeated measurements at successive growth stages, which multiplies the per-plot burden by the number of measurement dates. As modern breeding programs often contain thousands of plots, the measurement burden quickly compounds.

The crop itself can change substantially between growth stages. Soybean development for example proceeds through a sequence of vegetative and reproductive stages, conventionally denoted as V_x and R_x stages. In early vegetative stages the plant is short and sparse with limited leaf area so that both the underlying ground is easily observed from above between

adjacent rows. As the plant advances into the reproductive stages it grows taller and broader, often to the point when the canopy fills in and closes any gaps between adjacent rows. Figure 1.2 shows an illustration of soybeans through several growth stages. The closing of the upper canopy is what can make the ground increasingly difficult to observe as the season progresses.

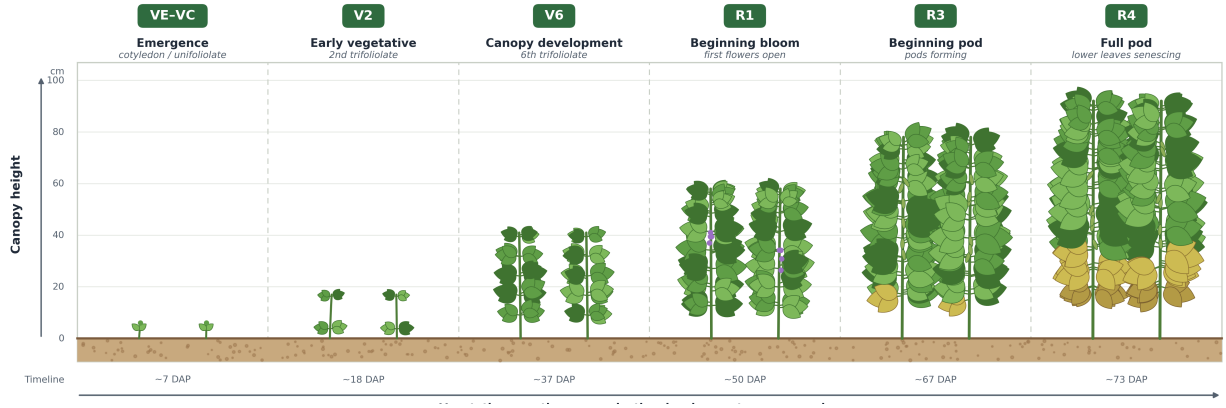

Figure 1.2: Soybean anatomy through growth stages and canopy closure

1.2 Manual Phenotyping Measurements

Despite the scale of modern breeding trials, canopy height is still commonly measured by hand. In the standard procedure a technician walks each plot and records height with a measuring stick, typically averaging several readings to obtain a representative plot value. The method is direct and requires no special instrumentation, but has several intrinsic costs that become limiting at scale.

First, the procedure is slow. A careful per-plot measurement can take several minutes, and this rate scales linearly with number of plots. A trial of several hundred plots measured at several growth stages represents a substantial labor and time commitment. Second, each measurement is labor-intensive and demands skilled personnel throughout as consistent judgment about where the canopy top lies and how to average across an uneven plot requires training. Finally, manual height measurements are subject to technician to technician variation with different technicians making different decisions about the representative points and canopy boundaries. This introduces variance in the evaluation that is unrelated to the

genotype.

Biomass measurements are regularly obtained in a destructive manner through cutting and taking wet and dried weights. This removes plants from the trial and can introduce more variation as the same plants cannot be evaluated continuously through the season.

These factors together mean that manual structural phenotyping does not scale well to plot counts and measurement intervals needed by modern breeding programs. The time required to measure every plot at every growth stage is a key resource that an automated or semi-automated alternative should aim to conserve.

1.3 Limitations of Existing Sensing Modalities

Due to the bottleneck and limitations previously described, a range of remote and proximal sensing modalities have been proposed to automate canopy characterization and relieve the manual measurement burden, many offer genuine advantages but also carry limitations that constrain their applicability. A more complete survey of these approaches is given in Chapter 2.

As elaborated, optical imaging approaches such as RGB photogrammetry are inexpensive and widely deployed on both ground-based and aerial platforms. They recover canopy geometry only indirectly and their performance depends strongly on favorable and consistent conditions. Variable environmental conditions such as lighting from cloud cover, dust/debris obstruction, and wind-induced canopy motion all degrade the reconstructed geometry. Additionally optical sensors interact primarily with the outermost canopy surface and cannot penetrate dense foliage to characterize internal canopy structure or sub-canopy features.

Light detection and ranging (LiDAR) can provide dense three-dimensional point clouds and has demonstrated accurate measurement of structural traits such as plant height and canopy volume. It is similarly sensitive to occlusion due to dense foliage and performance can be affected by ambient conditions. Like optical imaging it does not readily report on the vertical distribution of structures beneath the canopy surface.

The common theme of these established methods is a combination of environmental sen-

sitivity and inability to see below the canopy surface. As modern breeding programs require large scales and timed measurement intervals, environmental conditions can be variable and make it difficult to characterize canopies in breeding trials.

1.4 Radar as a Complementary Modality

Radar sensing offers a complementary set of properties to optical imaging and LiDAR modalities previously described. A radar measures range directly, rather than inferring geometry from images and its operation is independent of ambient illumination conditions. Radar measurements are relatively tolerant to environmental conditions such as dust and fog that degrade optical performance. As breeding programs operate on tight seasonal timelines, this insensitivity to minor environmental conditions is valuable.

Radar can also provide partial penetration in a vegetation canopy, returning energy from both the canopy and underlying ground with a single measurement. The return depends on the dielectric properties of the examined material, which are influenced by plant biomass and water content, providing a physical link to traits of interest that are difficult to assess optically. While an optical sensor reports primarily on the canopy surface, a radar return could encode information about the vertical distribution of scattering structures throughout the canopy.

These advantages are most fully realized at millimeter-wave frequencies. At 120 GHz the wavelength is approximately 2.5 mm, comparable in scale to leaf and stem structures of the crop. This relative scale provides more sensitivity to canopy architecture. At higher frequencies large absolute bandwidths are more easily attained which enables a finer range resolution to resolve smaller targets. A 5 GHz bandwidth provides a theoretical range resolution of 3 cm, which is sufficient to resolve structural features within a crop. High atmospheric attenuation at 120 GHz, which is often a drawback for long-range systems, is acceptable for the short standoff distances in proximal phenotyping, and can even be desirable to limit interference in settings with multiple concurrent measurement systems. The cost of operating at a higher frequency is reduced canopy penetration, a limitation

revisited in Chapters 3 and 5.

While radar does not replace optical imaging or LiDAR sensing it does complement by providing an alternative direct measurement, greater environmental condition tolerance, and a degree of canopy penetration that the established modalities lack. Realizing this potential for breeding-trial phenotyping however requires new instrumentation systems and a processing methodology to convert raw output into meaningful plot traits.

1.5 Problem Statement

Although the physical basis for radar interaction with vegetation is established, there is limited prior work on radar systems specifically designed and optimized to measure breeding-relevant phenotypic traits at the spatial and temporal resolution that modern breeding programs require. As elaborated in Chapter 2, satellite synthetic aperture radar operates at field or regional scales but lacks the resolution to resolve individual breeding plots. Proximal vegetative-sensing radar studies have generally been confined to lower frequencies and limited in scope. A compact, proximal, plot-scale, millimeter-wave instrument for canopy phenotyping remains unaddressed in prior work.

Even given a suitable radar instrument, raw data acquisitions produced are not directly interpretable as phenotypes. Often in radar systems fixed clutter, antenna-coupling artifacts, sweep-to-sweep drift, and multiplicative speckle will limit the performance. To obtain canopy measurements an accurate ground and canopy region estimate need to be extracted from a noisy two-dimensional image in which the canopy edge is diffuse and intermittent due to inter-row gaps and alleys common in breeding-trial layouts. As the crop under evaluation develops through the growing season, parameters of a processing methodology may need to adjust. Converting raw range-time images into per-plot canopy height and structural density in a scalable, repeatable fashion is therefore a non-trivial signal-processing problem.

This dissertation addresses these gaps. It presents the design, implementation, and field evaluation of a custom 120 GHz stepped-frequency continuous-wave (SFCW) radar for high-throughput soybean phenotyping. It also uses the data collected from the demonstration

to develop a processing pipeline that converts the raw radar range-time datasets into per-plot canopy height and an associated structural density measure within a calibrate-once, human-in-the-loop workflow. The specific contributions are as follows:

1. The design and implementation of a compact 120 GHz SFCW radar sensor for field-level phenotyping and custom embedded acquisition system in a compact, low-power form factor suitable for breeding-trial deployment

2. A field evaluation of the system in a soybean breeding trial at Kansas State University across the 2025 growing season, demonstrating tracking of canopy development and estimation of plant height in agreement with manual ground-truth measurements across multiple growth stages

3. A processing methodology that converts raw radar acquisitions into per-plot canopy height and structural density. The pipeline includes:

   (a) a system-aware cleanup chain that forms range-time images and suppresses calibration sweeps, fixed clutter, slow drift, speckle, and sweep-to-sweep variation

   (b) a segment-aware particle-filter ground tracker with backward smoother that follows the ground through canopy shadowing and reports a per-chirp uncertainty

   (c) a region-based canopy detector that extracts the canopy envelope relative to the tracked ground

   (d) an integrated-energy measure over the ground-to-canopy region as a proxy for structural density

4. A demonstration and evaluation of a calibrate-once-and-scale-to-field workflow, in which an expert tunes the detector parameters on a small number of representative plots and then applies the same configuration to additional plots without intervention, including an assessment of the workflow's measurement cost relative to manual methods at equivalent spatial sampling density, the value of the resulting sub-plot height profile, and the cross-stage re-tuning limitation.

The remainder of the dissertation is organized as follows. Chapter 2 surveys high-throughput phenotyping platforms and the optical, LiDAR, and radar sensing modalities relevant to canopy characterization, and reviews the radar and signal-processing background on which this work builds. Chapter 3 describes the 120 GHz SFCW radar system, including its requirements, architecture, continuous coherent calibration, and field data collection. Chapter 4 presents the processing pipeline that converts raw radar data into range-time images, per-plot canopy height and structural density, and evaluates it on field data. Chapter 5 summarizes the contributions, draws system-level conclusions, discusses limitations, and outlines directions for future work.

Chapter 2

Literature Review

This chapter reviews the body of work on which this dissertation builds. It begins with the high-throughput phenotyping platforms that motivate automated canopy sensing, then surveys the current dominant sensing modalities (LiDAR and optical imaging) and their limitations. It next turns to radar, its established agricultural uses, particular advantages of millimeter wave sensing, and an overview of relevant architectures. The chapter closes with radar signal-processing background by surveying the techniques reported in literature and developing a foundation for the processing pipeline presented in Chapter 4.

2.1 High-Throughput Phenotyping Platforms

The phenotyping bottleneck described in Chapter 1 as a result of disparity between the rate at which new genotypes can be generated and the rate at which their phenotypes can be measured has driven investment in high-throughput phenotyping (HTP) platforms[1;2]. The goal of any HTP platform is to acquire phenotypic measurements across plots in a rapid non-destructive manner with sufficient repeatability that the resulting traits can be used for selection in breeding trials. Literature distinguishes platforms primarily by how the sensor is carried and at what scale it can operate.

Ground based platforms mount sensors on carts, tractors, gantries, field robots, or by

hand and traverse plots at close range. They offer high spatial resolution due to the proximal range and stable sensor geometry at the cost of throughput limited by ground travel speed. Recent work has demonstrated autonomous field robots using optical sensors capable of high throughput via continuous operation in row crops of large trials[3;4].

Aerial platforms like unmanned aerial vehicles (UAVs) are able to traverse a field in a faster manner than most ground based platforms, but are generally able to carry fewer sensors due to payload limitations and can have a limited resolution as a result of the distance typically operated from the crops. The distance/resolution problem can be resolved some by operating closer to the crops but this can cause interactions directly with the crops that cause motion and can reduce the value of the data collected[5].

Reviews of the field consistently identify structural traits, particularly plant height, canopy cover, and biomass as high-value targets for automation via HTP as these traits are labor intensive to measure manually and strongly tied to yield and agronomic performance[1;6]. The work described in this dissertation is demonstrated in a ground-based proximal manner targeting plot-level structural traits with spatial resolution that satellite approaches cannot reach.

2.2 Optical and Photogrammetric Approaches

Optical sensing encompasses RGB cameras, multispectral and hyperspectral imagers, and thermal infrared sensors deployed on both ground and aerial platforms. RGB cameras are relatively inexpensive and lightweight, while multispectral/hyperspectral imagers and thermal infrared sensors can be both costly and physically larger systems. The RGB cameras can provide canopy-cover estimation through greenness-based segmentation[7] and structure from motion photogrammetry that allows three-dimensional canopy geometry to be reconstructed from overlapping images to provide height and volume estimates[8]. The multispectral/hyperspectral imagers and thermal infrared sensors allow for richer information in the form of vegetation indices derived from reflectance related to chlorophyll content[9] or more direct evaporation measurements via leaf temperature[10].

The primary weakness of optical approaches is their dependence on imaging conditions. Photogrammetric height reconstructions have been shown to be sensitive to canopy motion[5] and can depend on stable illumination[11] which can be difficult to maintain in real field environments with cloud cover, shadowing from the platform, and dust or fog. Additionally readings can be sensitive to day-to-day solar radiance levels which can lead to measurement repeatability problems across sessions. Optical sensors also interact primarily with the canopy surface and cannot characterize internal structure. Greenness-based canopy cover segmentation will be used as a reference measurement in this dissertation to compare against radar derived structural density in Chapter 4.

2.3 LiDAR for Structural Phenotyping

Light detection and ranging (LiDAR) measures distance directly by using the time of flight of emitted laser pulses or via phase comparison. They produce point measurements that can build point clouds through motion or scanning. As it measures geometry rather than inferring it, LiDAR has become one of the most accurate proximal modalities for structural phenotyping, with demonstrated success in estimation of plant height, canopy volume, and above-ground biomass across a range of crops[12;13]. It can be deployed on both ground and aerial systems.

Comparative studies have evaluated LiDAR against UAV photogrammetry for plant-height estimation and generally find LiDAR more robust as a result of reconstruction errors that affect image-based methods[14]. The main limitation of LiDAR for phenotyping is its lack of canopy penetration. As the laser interacts predominantly with the outer most foliage and measurements largely only describe the upper surface, carrying little information about sub-canopy structure or the underlying ground. Performance can additionally be impacted by ambient environmental conditions that affect laser light such as dust and fog. Additionally these systems can be costly and heavy for high-grade versions which may limit the scope of deployment.

2.4 Radar in Agriculture

Radar has a long history of vegetation monitoring through remote sensing. Most of this operates at scales and geometries far removed from plot-level phenotyping though. The largest body of agricultural radar literature concerns satellite based synthetic aperture radar (SAR). C-band Sentinel-1 and multi-polarization systems such as RADARSAT-2 have been used extensively for crop classification, large scale biomass estimation, and soil moisture retrieval. They operate by exploiting the dependence of back-scatter on aggregate canopy architecture, surface roughness, and dielectric constant[15;16]. The long time temporal evaluation of back-scatter through a growing season can track crop growth stage and biomass accumulation[17]. However, these systems typically operate at a field or regional scale and their coarse spatial resolution does not allow for evaluation of individual breeding plots which makes them unsuitable for the per-plot measurements that breeding programs require.

At the field level, ground-penetrating radar (GPR) has been applied to non-destructive root phenotyping and soil-moisture characterization[18;19]. GPR directs energy downward into the ground and characterizes the sub-surface. It can evaluate root traits and soil state rather than the above ground canopy structure that is the subject of this work.

Closest to the present application are proximal continuous-wave radar systems aimed directly at vegetation. Frequency modulated continuous-wave (FMCW) radar has been explored for vegetation sensing in a limited number of studies including ultra-wideband work on wheat canopy height that reported encouraging first results[20;21]. This line of work establishes the physical viability of proximal radar canopy sensing but has remained at lower microwave frequencies, leaving millimeter-wave plot-scale regime targeted here largely unexplored.

2.5 Millimeter-Wave Sensing

Millimeter-wave frequencies confer several advantages for close-range agricultural measurement. The first is favorable interaction between wavelength and plant structure. V-Band

(e.g 40-75 GHz) yields wavelengths as short as 4 mm, while 120 GHz is approximately 2.5 mm, comparable to dimensions of leaves and stems. This enhances sensitivity to canopy architecture relative to longer-wavelength systems that interact with the canopy as a bulk medium. The second is resolution which is inversely proportional to swept bandwidth. The ease of obtainable wide bandwidths at higher frequencies due to being smaller fractionally directly enable fine range resolution. A 5 GHz swept bandwidth yields a theoretical range resolution of 3 cm. This resolution is sufficient to resolve macro-structural canopy features. At higher frequencies atmospheric absorption would also increase so that low power units can share spectrum and make licensing easier.

The growth of automotive radar has made highly integrated transceivers at millimeter wave frequencies widely available. This has allowed the investigation of their alternative use in short-range sensing[22]. The normal trade-off with millimeter-wave radars is the high atmospheric attenuation and reduced penetration ability. For proximal phenotyping the higher attenuation is largely benign as standoff distances are short, and can even be advantageous as it can suppress interference between concurrent systems. The reduced penetration is a genuine limitation that constrains ground tracking under closed canopies as discussed in Chapters 3 and 5. The 120 GHz band thus represents a trade of penetration depth for an increase in structural sensitivity.

2.6 Traditional Radar Systems

This section reviews radar architectures relevant to compact proximal sensing and develops the operating principles on which the system described in Chapter 3 are based. Ranging radars fall into two categories: pulsed and continuous-wave families. Pulsed systems transmit a pulse and time the returning echo. To achieve a fine range resolution very short pulses are required to realize a wide instantaneous bandwidth. All radar systems have their range resolution determined by Equation 2.1[23].

$$\Delta r = \frac{c}{2 \cdot BW} \tag{2.1}$$

Continuous-wave (CW) architectures transmit continuously and recover range from the frequency/phase relationship difference between the transmitted and received signals. As the transmitted waveform no longer relies on short pulses the power can be distributed over the sweep duration allowing for lower peak power systems. A device known as a mixer is used which multiplies the waveforms from two inputs and through multiplication in the time domain, the mixer output resolves as the difference (and summation, although this is commonly filtered away) in frequency and phase. In a real world environment with many returns from multiple targets at different distances the result will contain many frequencies of different amplitudes superimposed on each other. Utilizing a Fourier transform on the result provides the frequency components of the signal that correspond to distance. An example of ideal CW radar system diagrams can be seen in Figure 2.1.

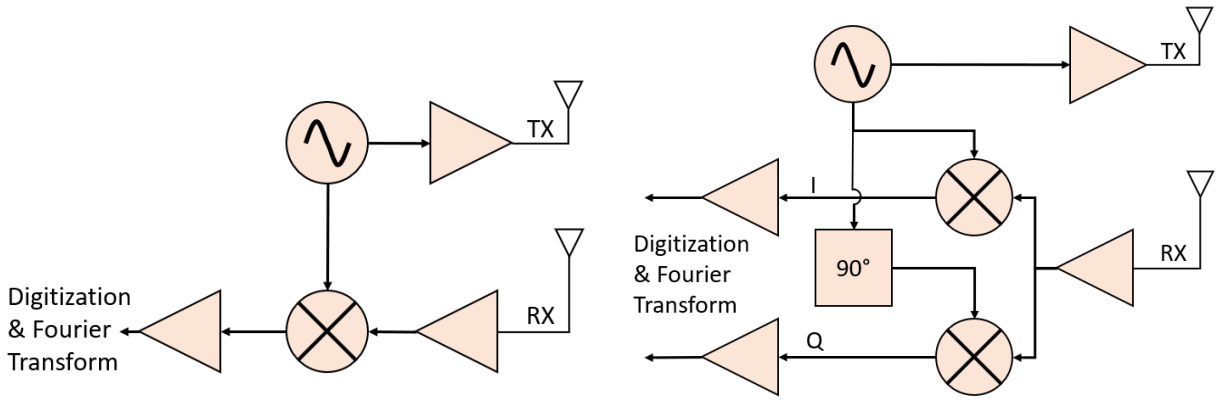

(a) Ideal Radar - Only Real

(b) Ideal Radar - Quadrature

Figure 2.1: Ideal radar block diagrams

In frequency-modulated continuous-wave (FMCW) radars the transmit frequency is swept with slope k Hz per second and travels out to reflect off a target and return some time (τ) seconds later. The return frequency can be compared with a mixer to the current transmit frequency ($f_{tx}+\tau k$) to obtain a difference in frequency (Δf). A plot showing the transmitted and received frequency for a given time delay can be seen in Figure 2.2a with annotation. The resulting difference in frequency is proportional to target distance/range R shown in Equation 2.2.

$$R = \frac{c\tau}{2} = \frac{c\Delta f}{2k} \implies \Delta f = \frac{2Rk}{c} \tag{2.2}$$

Stepped-frequency continuous-wave (SFCW) radars rely on a measured phase change at discrete frequency steps to resolve a distance. They operate by taking evenly spaced frequency steps and looking at the phase of the received reflection at each frequency point. By stepping through frequencies the phase changes due to inverse relationship between frequency and wavelength. Figure 2.2b shows an example of different stepped frequencies over a fixed amount of time/distance and the Figure 2.2c shows the resulting phases plotted for each frequency linearly spaced. The new purple waveform plotted on top of the phase returns shows the resulting signal generated by the mixer whose frequency is related to the time/distance.

Both these modes of operation result in a frequency output that is proportional to the distance to the reflecting target.

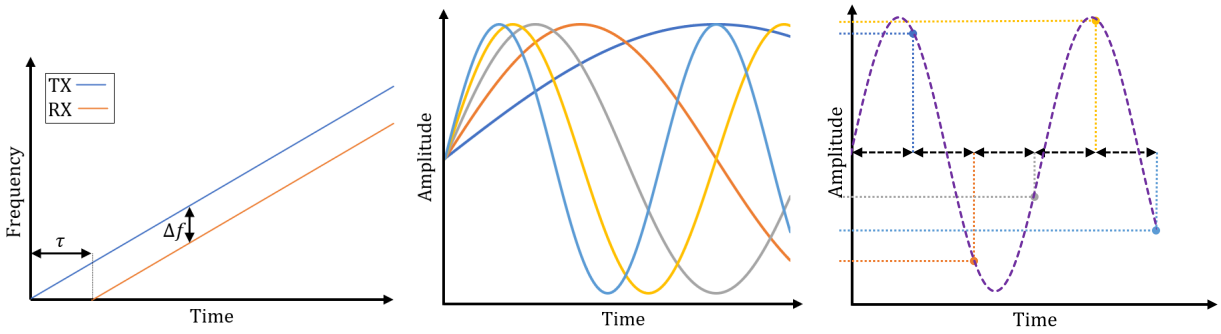

(a) FMCW frequency plot (b) SFCW amplitude plot (c) SFCW Equi-distant Phases

Figure 2.2: CW distance radar operations

2.7 Traditional Radar Signal Processing

Converting raw radar data acquisitions into interpretable measurements requires a sequence of established signal processing operations. This section surveys those operations as reported in radar and tracking literature and builds a conceptual toolkit to be used in later chapters.

For FMCW and SFCW radar systems, ranging information is encoded in frequency content of the mixer output and the discrete Fourier transform (DFT) or Fast Fourier Transform (FFT) is key to resolving distance. To provide the best results, a system like that shown in Figure 2.1b should be used to provide a complex input to the DFT. By providing complex inputs to the transform, the resulting output does not produce the conjugate-symmetric (two-sided) spectrum that a real-only input would. Equation 2.3 shows the result of an Euler expansion of a real only signal (output from Figure 2.1a) which provides two terms that resolve through the Fourier transform into two discrete tones (f_0 and $-f_0$) while Equation 2.4 (output from Figure 2.1b) results in just one term and corresponding tone (f_0).

$$S_{Baseband} = \cos(2\pi f_0 t) = \frac{1}{2}e^{j2\pi f_0 t} + \frac{1}{2}e^{-j2\pi f_0 t} \tag{2.3}$$

$$\begin{aligned} S_I + jS_Q &= \cos(2\pi f_0 t) + j\sin(2\pi f_0 t) \\ &= \left(\frac{1}{2}e^{j2\pi f_0 t} + \frac{1}{2}e^{-j2\pi f_0 t}\right) + j\left(\frac{1}{2j}e^{j2\pi f_0 t} - \frac{1}{2j}e^{-j2\pi f_0 t}\right) \\ &= e^{j2\pi f_0 t} \end{aligned} \tag{2.4}$$

Additionally it is common for FMCW and SFCW systems to utilize both an increasing and decreasing frequency sweep/"chirp". The results are similar for both but result in a "negative" frequency in the decreasing ramp due to the negative slope term. In an IQ system this is realized at the DFT output.

Often a windowing function can be applied to the sampled data to soften rectangular time domain windows that produce side-lobes in the transformed result. The trade-offs among window choices are cataloged by Harris[24]. By stacking successive range profiles as the sensor advances in time (and often physically in motion) a two dimensional range-time image can be constructed. Figure 2.3 shows a range-time image formation from raw I/Q samples through Fourier transform to the final image.

Clutter and artifact suppression in range-time images are key to deal with issues present

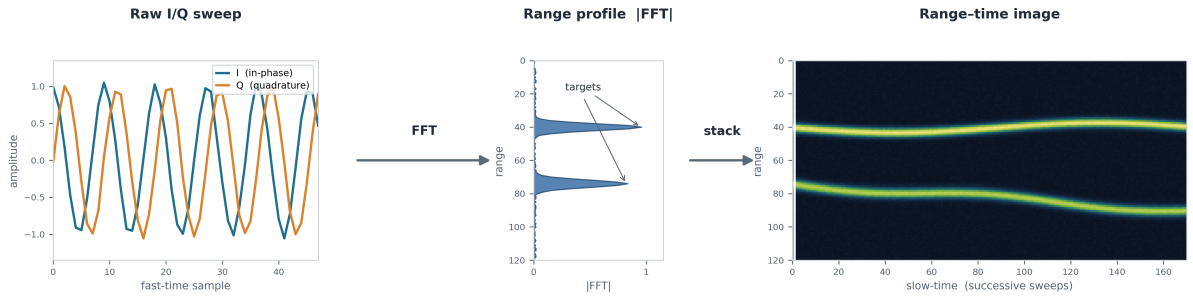

Figure 2.3: Range-time image formation example

in many radar systems. These can result from stationary reflections and slow varying drift in components from thermal changes. Standard approaches to stationary clutter removal estimate a slowly varying background and subtract it. This is effectively a high-pass filtering along the slow-time axis or by mean/median subtraction within affected range bands[23]. Speckle is caused by multiplicative interference from rough surfaces. This is conventionally reduced through multilook averaging which involves incoherently averaging adjacent chirps. This reduces temporal resolution but resolves targets better by reducing variance[25]. These techniques will be used as a foundation and adapted to the segment structure imposed by our systems periodic calibration process in the cleanup chain of Chapter 4.

Extracting continuous surfaces from a noisy range-time image is a tracking problem. The simplest approach looks for the largest peaks if the surface of interest produces a strong continuous reflection, but fails when the return is intermittently weakened or obscured. A Kalman filter is optimal for linear-Gaussian problems, but represents the state as a single Gaussian random variable and so cannot capture the multimodal evidence that arises when two strong targets are present. Recursive Bayesian state estimation provides a principled alternative by propagating a probability distribution over target position from one measurement to the next. Particle filters address the regime well by representing the state distribution with a population of weighted samples. A particle filter accommodates non-linear dynamics and multimodal observation likelihoods and exposes an uncertainty estimate at

each point as a natural by-product of the weighted population[26;27]. Figure 2.4 shows an example of a particle filter traversing an ambiguous measurement period. However, forward filtering alone is causal and cannot use later evidence to resolve prior ambiguity. A backward smoothing pass re-weights and re-samples the stored forward trajectories using subsequent measurements to produce a non-causal, lower-variance estimate[28]. The ground tracker described in Chapter 4 is built upon these ideas.

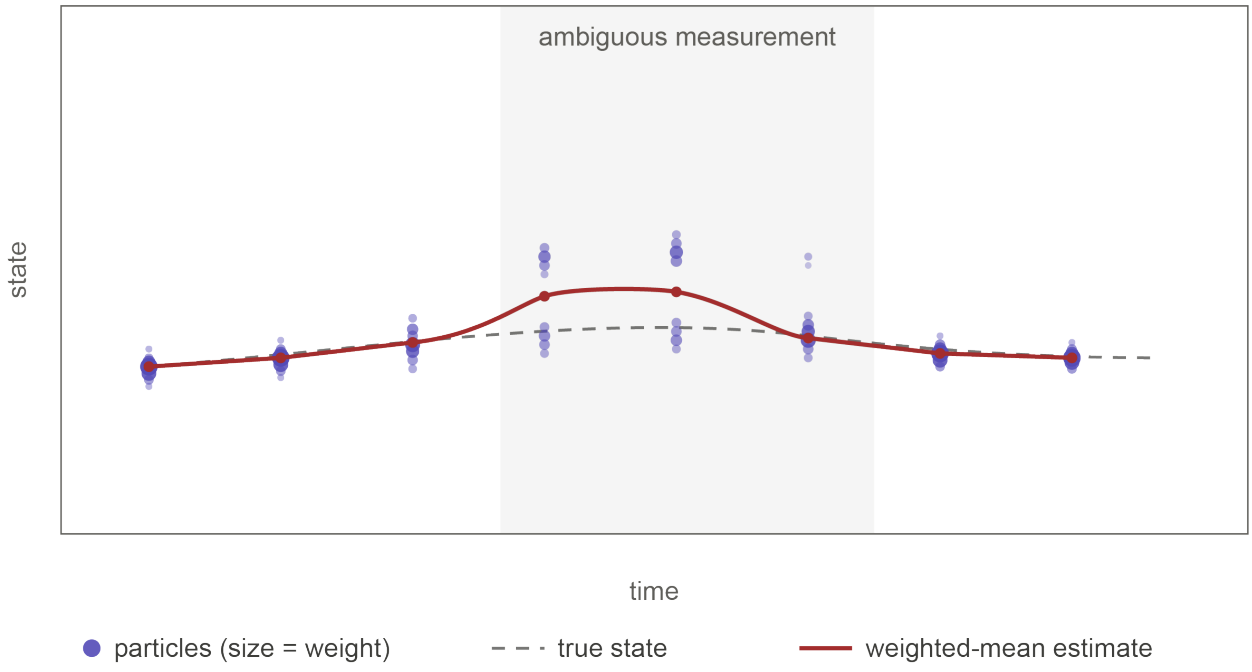

Figure 2.4: Particle filter example operation

Where a distributed radar return must be identified instead of a single value, region-based image processing is required. Foundational image analysis tools such as amplitude thresholding to form binary masks, morphological closing to bridge small gaps in connected regions, and connected-component analysis to isolate and filter candidate regions are required[29]. Chapter 4 will combine and apply these primitives together with the ground tracker also developed so that canopy regions can be identified in the range-time image.

Together these four families of processing: range-time image formation, background/speckle suppression, Bayesian surface tracking, and region based detections constitute the methodological foundation of this processing pipeline. Chapter 4 develops each into a specific segment aware algorithm required by the developed 120 GHz SFCW system and evaluates

them on field data.

Chapter 3

Radar Phenotyping System

In order to develop a radar system capable of measuring crop development with sufficient resolution to resolve phenotypic traits, a system with a wide bandwidth was needed. A high frequency millimeter wave was desired to maximize the interaction of the radar illumination with the crop. Based on the commercial available components, 120 GHz was selected as the frequency of operation. The bare electronics and packaged radar can be found in Figure 3.1.

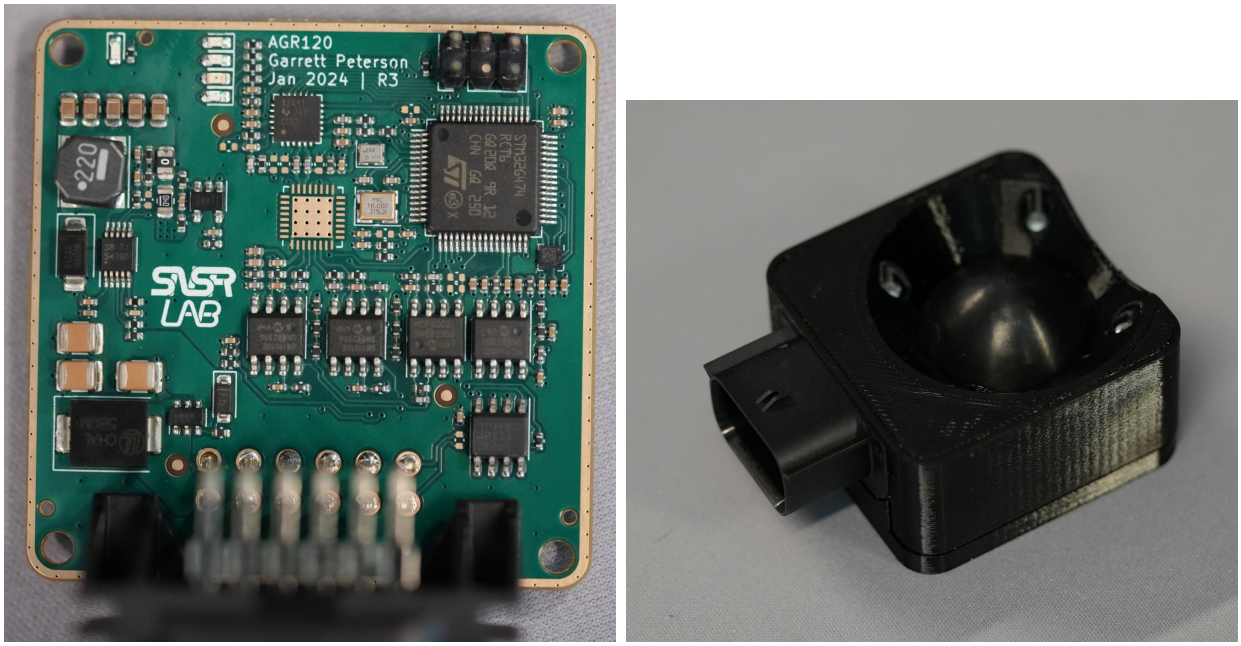

(a) Radar Electronics (b) Packaged Radar

Figure 3.1: Phenotyping radar

3.1 Requirements

An effective phenotyping radar sensor needs to meet several requirements. First it needs to be compact to allow for easy deployment on several platforms. It should be small enough to be hand-carried or deployed on a drone. Second it should have a low power consumption to again allow for flexible deployment where power may be limited. Third it should allow for selective illumination of crops to allow for characterization of only an individual row or singular plot in some cases, which implies a required narrow beamwidth. Fourth it should provide a robust communication method that is easy to integrate with external systems but also provide a sufficient rate of transfer to allow for data to be collected that will enable high-throughput.

3.2 System Architecture

Based on the requirements described in the previous section, a 120 GHz SFCW radar system was developed. The design combines a commercial millimeter-wave transceiver, a general-purpose microcontroller, and a compact analog signal chain on a single board, packaged for field deployment. A system diagram is shown in Figure 3.2.

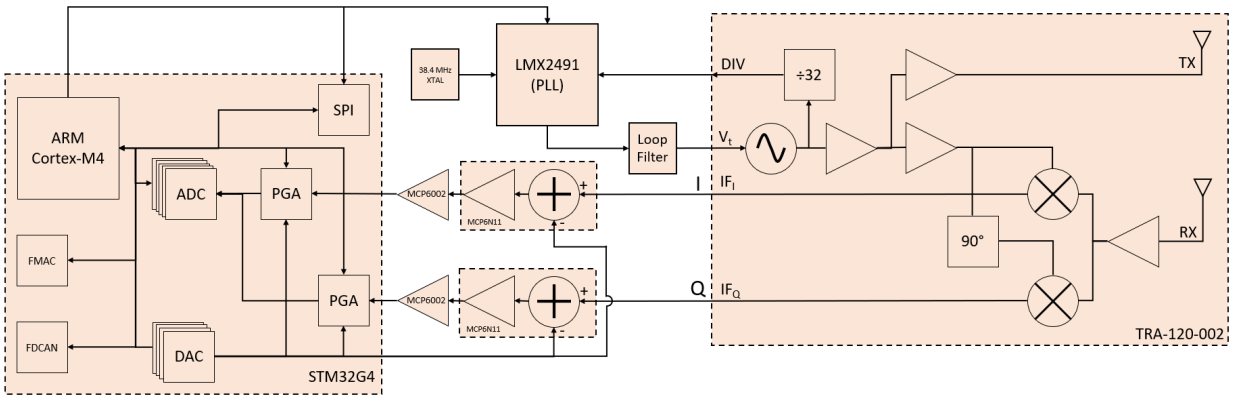

Figure 3.2: 120 GHz phenotyping research radar

At the core of the RF front end is the indie Semiconductor TRA-120-002 transceiver[30], which integrates most of the high-frequency components, including the voltage-controlled

oscillator (VCO), mixers, RF amplifiers, frequency dividers, and antennas, into a single package. It was selected both because it is one of the only commercially available options in the 120 GHz band and because of this high level of integration. Pairing it with a Texas Instruments LMX2491 phase-locked loop (PLL) enables both FMCW and SFCW operation. The SFCW mode was ultimately chosen to overcome the short-range limitations discussed in the following section. Once configured, the LMX2491 steps through frequencies with general-purpose input/output (GPIO) pin toggling, and operating at 120 GHz exploits the small millimeter wavelengths to achieve the strong canopy interaction described in Chapter 2.

To keep the design compact, these core integrated circuits (ICs) are paired with an STM32G474 microcontroller. The STM32G474 provides a rich set of analog peripherals such as an integrated temperature-compensated voltage reference, several high-speed 12-bit analog-to-digital converters (ADCs), a pair of high-speed 12-bit digital-to-analog converters (DACs), and programmable-gain amplifiers (PGAs). Together these give a large, controllable dynamic range. On the receive side, a Microchip MCP6N11 instrumentation amplifier converts the differential output of the TRA-120-002 to a single-ended signal, which a Microchip MCP6002 operational amplifier amplifies further before it is fed into the microcontroller's integrated ADCs.

Power is supplied by an integrated DC-DC converter built around the TI TPS54260, accepting a 6–48 VDC input. To reduce cost and size, all components are carried on a single printed circuit board (PCB) measuring 50 × 50 mm. Communication uses a CAN-FD (Controller Area Network with Flexible Data-rate) bus, chosen for its error-correcting signaling and noise immunity and for its easy integration into existing agricultural implement systems; the nominal bus speed is 500 kBaud. A complete schematic is provided in Appendix A.

To narrow the antenna beam, a dielectric lens focuses it to approximately 20° in the E-plane and 10° in the H-plane[31]. While this makes the sensor more sensitive to angular alignment, it improves the signal-to-noise ratio for targets within the beam and reduces clutter from off-axis scatterers such as adjacent crop rows and surrounding soil, so that only

the crop of interest is examined. Packaged in an additively manufactured housing with the focusing lens, the final system measures 58 × 58 × 30 mm and weighs only 60 grams.

The system connects through a single Molex MX120G connector that carries power, transmits and receives data, and provides an addressing scheme so that multiple units can share one communication bus. Three address pins, shorted to adjacent grounds in different combinations, yield eight unique addresses. Nominal power consumption is 1.5 W across the input voltage range.

The data presented in this dissertation were collected with a fixed set of radar parameters in the SFCW mode described above. Because the system is intended for proximal, close-range sensing, a 3.84 m unambiguous range was used, corresponding to a 128-step frequency ramp. Each measurement cycle consists of three stages: an increasing frequency ramp, a short constant-frequency pause, and a decreasing frequency ramp. The dwell time at each step is 10 µs, allowing for signal transit time, PLL/DAC settling, and ADC sampling. The typical operating parameters are summarized in Table 3.1.

Table 3.1: Typical Operating Conditions

Parameter	Typical Value	Unit
Bandwidth	5	GHz
Range Resolution	3	cm
Unambiguous Range	3.84	m
Frequency Steps	128	
Frequency Step Dwell	10	uS
Antenna Beam-width 3dB (E-Plane)	20	degrees
Antenna Beam-width 3dB (H-Plane)	10	degrees

3.3 Continuous Coherent Calibration

One key difference from the ideal radar system diagrams in Figure 2.1 is the addition of a feedback path in the I/Q outputs from the TRA-120-002. This feedback allows for correction signals to be fed into the chain to reduce the impact of large system leakages that are present in most FMCW and SFCW systems. These leakages result in large part due to mixer isolation

between the RF and LO ports but also due to coupling between the transmit/receive antennas and on the circuit board traces. It is difficult to design large isolation between the antennas and circuit board traces. The result of this is false low-frequency (close-range) targets. This is problematic due to the proximal sensing desired as crops to be evaluated are also in a relatively close range bin and much weaker in signal strength. One hardware solution to this issue is adding a DC blocking capacitor in the signal chain post conversion which creates an analog high pass filter. This blocks any DC components and very low frequency returns but causes a frequency response that can also filter out close range targets limiting the systems effective blind range which is undesirable in a close-range sensor.

Another solution is to create the necessary feedback signal to cancel undesired components, the raw radar data is low pass filtered to select low frequency components that need to be removed. This filter is implemented digitally inside the microcontroller as a Finite Impulse Response (FIR) filter. This filter allows for greater control and dynamic optimization in comparison to the analog DC blocking capacitor. Once the undesired signals have been isolated, an inverse copy appropriately scaled for amplifier chain gains is created through the integrated DACs inside the microcontroller and fed back into the chain at the instrumentation amplifiers as a reference. By feeding back this signal in the chain ahead of much of the gain in the system, large signals that would saturate the system if high gains were implemented can be effectively canceled out. This allows an iterative approach of removing these large signals followed by increasing gain through the integrated PGAs. By achieving larger gains without saturating the ADCs small signals that would normally reside below the quantization noise floor can be amplified to inside the ADCs dynamic range. As the undesired signal can drift in time the correction signals must be periodically re-calculated. This work was originally developed by the author[32;33] in a less robust manner and further refined[34]. Figure 3.3 shows a flow diagram of the process.

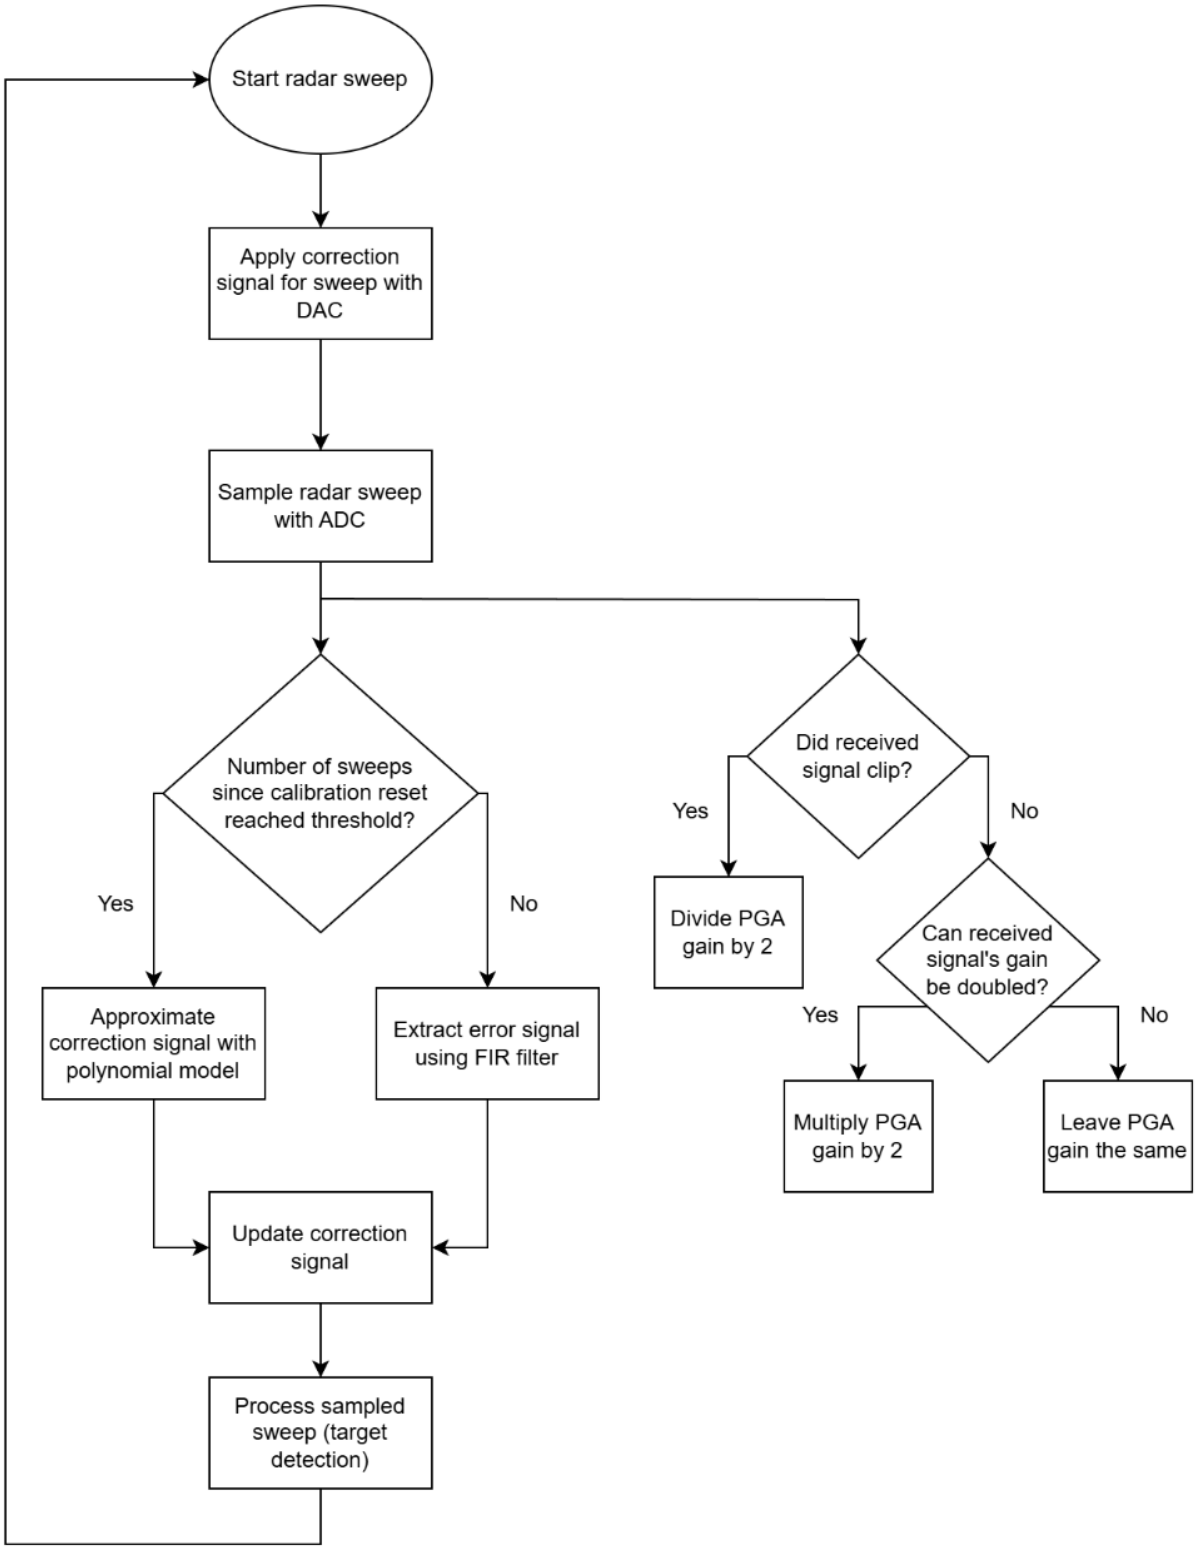

Figure 3.3: Continuous coherent calibration flow diagram[34]

3.4 Surface Testing

To evaluate the radar performance on different targets, data was collected directing the radar at a variety of surfaces. Figure 3.4 shows processed data as a range-time image with different targets. This was performed by hand-carrying the radar while synchronous co-located video was taken. Table 3.2 describes points of interest in the range-time image with timestamps and references to Figure 3.5 which shows representative sample frames from the video for different surfaces/targets. This test demonstrates the radar's capability to distinguish different targets.

Table 3.2: Multi surface test points of interest

Annotation	Start (s)	Stop (s)	Description	Reference
A	0	9	Mostly barren soil with a few small (3-4cm) crops	Fig 3.5a
B	9	17	Grass alley 10cm tall	Fig 3.5b
C	17	27	Crop Row traversing perpendicular crop at 20	Fig 3.5c
D	27	41	Grass alley	Fig 3.5b
E	41	58	Gravel surface	Fig 3.5d
F	58	66	Short grass (5cm) transition to field with residue	Fig 3.5e
G	66	77	Gravel Surface with height change at 74	Fig 3.5d
H	77	88	Tall Grass 40cm tall (with seeding)	Fig 3.5f
I	88	100	Medium Grass 20cm tall	Fig 3.5b
J	100	136	Bare soil with corn stalk residue	Fig 3.5g
K	136	142	Short grass with gravel trailing gravel transition	Fig 3.5e
L	142	155	Side view tall grass with close targets	Fig 3.5h
M	155	170	Gravel with trailing height change	Fig 3.5i

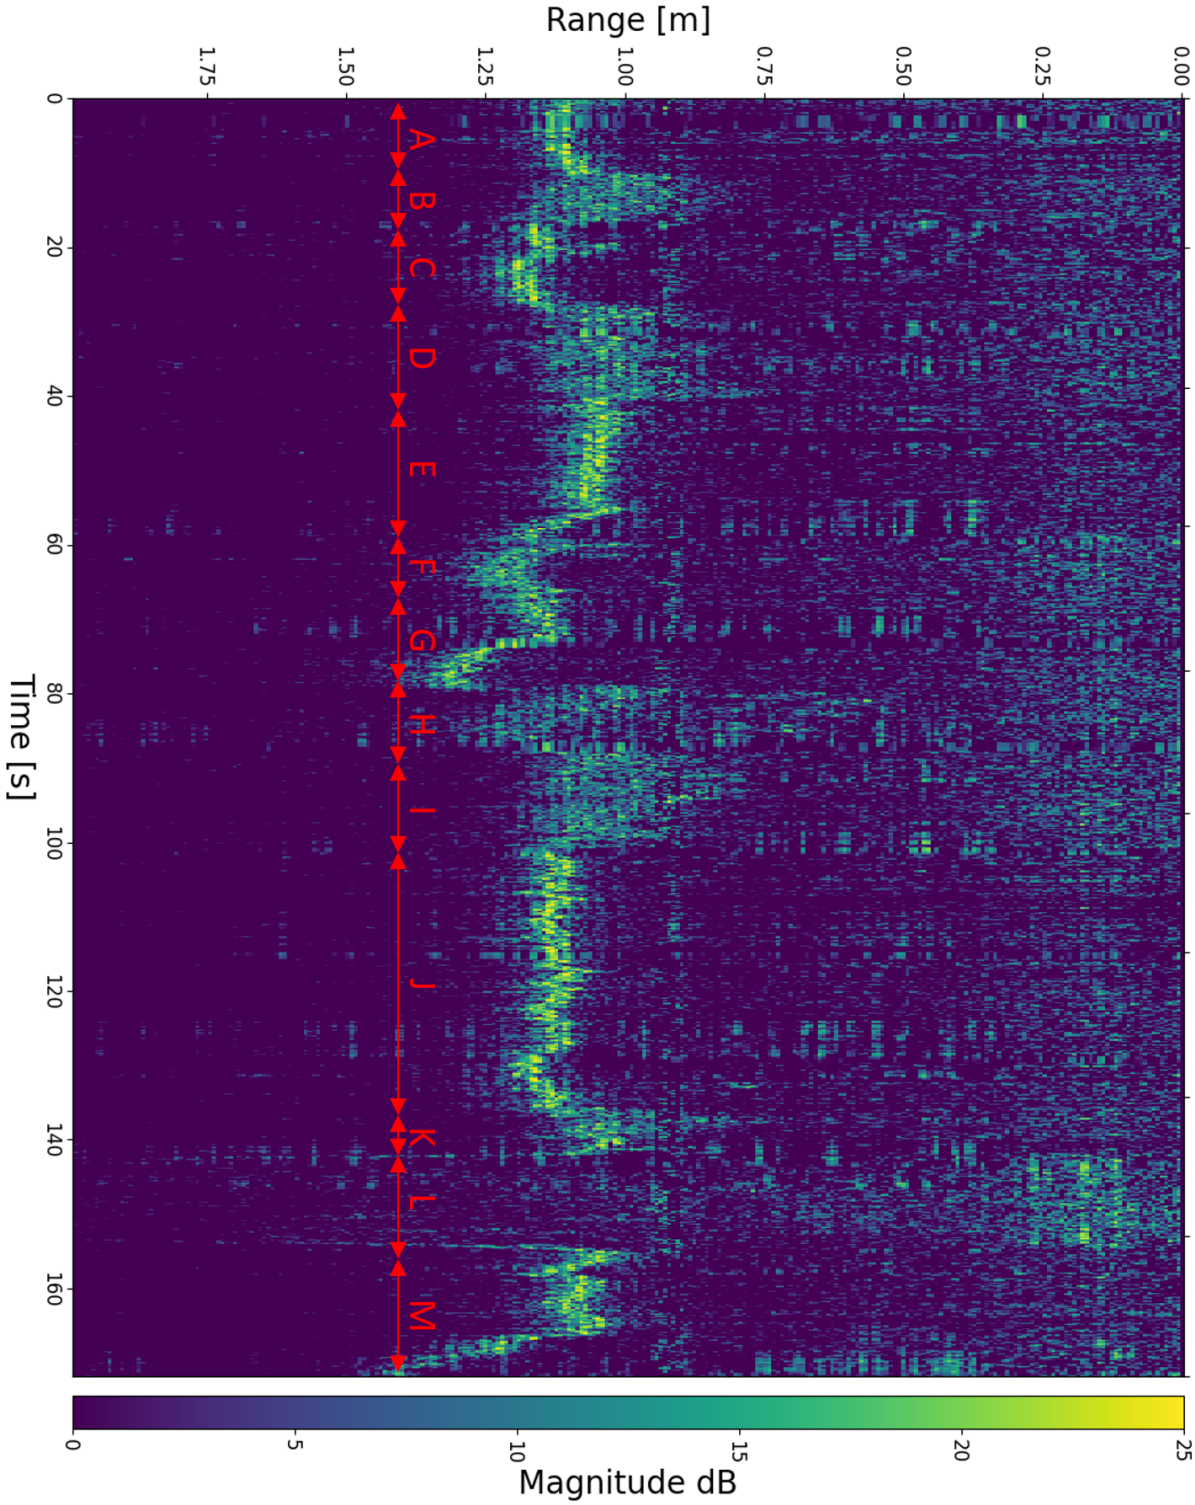

Figure 3.4: Radar multi surface test range-time image

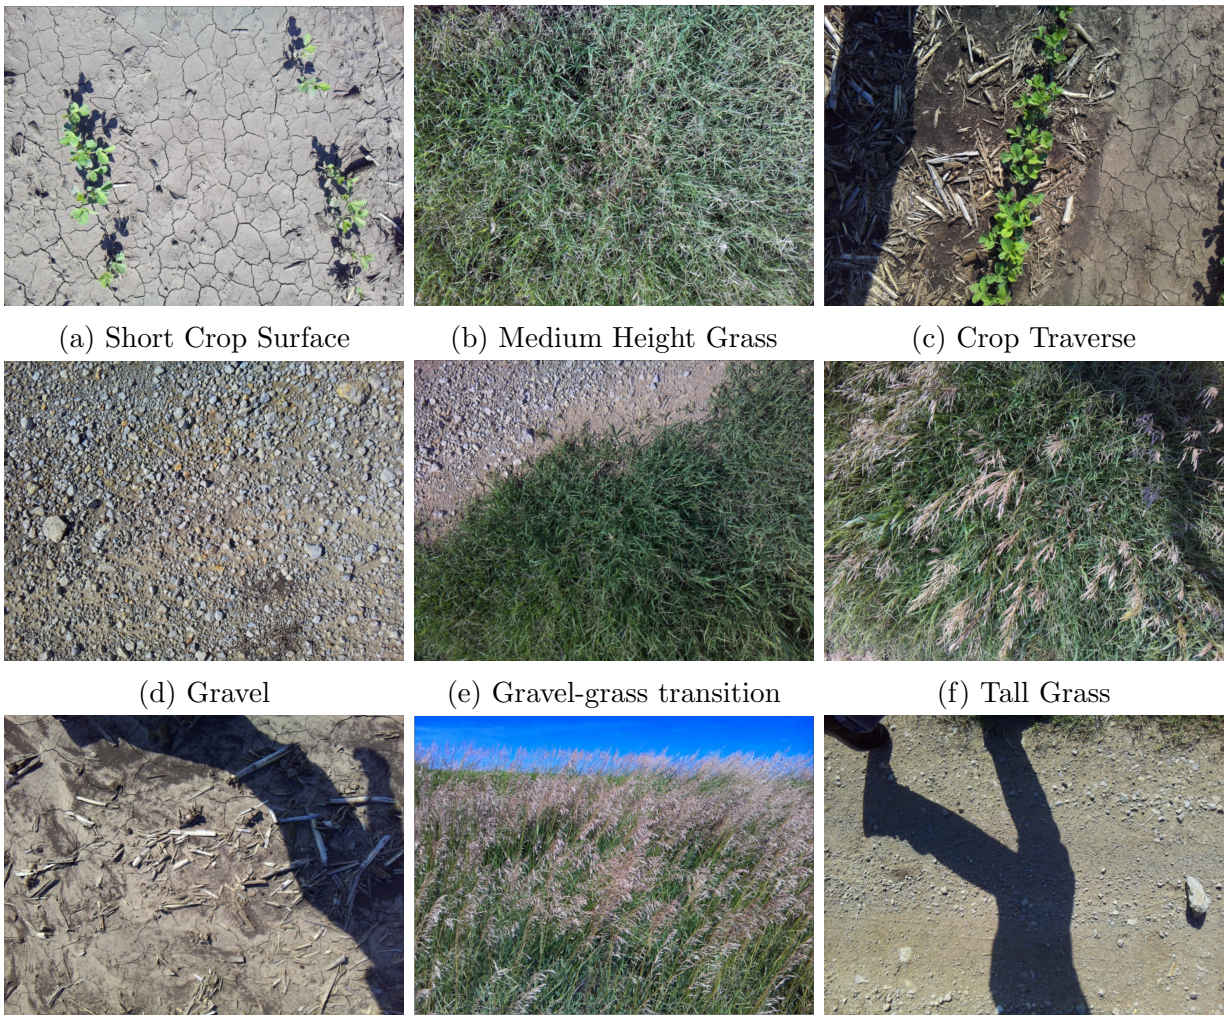

(a) Short Crop Surface (b) Medium Height Grass (c) Crop Traverse

(d) Gravel (e) Gravel-grass transition (f) Tall Grass

(g) Soil with corn residue (h) Side view grass (i) Dirt road (with shadow)

Figure 3.5: Surface test sample video frames

3.5 Breeding Trial Data Collection

The system was deployed in a breeding trial experiment conducted at Kansas State University's Agronomy North Farm in Manhattan, Kansas during the 2025 growing season. The experiment, planted on 6/9/2025 consisted of 20 soybean plots representing five genotypes arranged in four replications. Each plot consisted of four rows, 3.6m in length, with 76 cm row spacing. The trial included soybean genotypes varying in maturity, plant height, canopy architecture, and other morphological traits relevant to the breeding program. The map of the field can be found in Figure 1.1. Measurements were collected at multiple growth stages throughout the growing season to capture crop development dynamics.

The radar sensor was used in a hand carried downward-facing configuration directly above the crops at a heights of one to two meters from the ground. The sensor had a narrow beam illuminating a ground area of approximately 0.3 m^2 at a 1.75 meter distance and a 0.1 m^2 area at 1.0 meter. The experimental scanning routine moved parallel following the row from above.

Manual ground truth measurements were collected concurrently with radar acquisitions to provide reference data for evaluating the radar's phenotyping capability. Plant height was measured by hand at various points to obtain the canopy average. Synchronous co-located video footage was also captured with the radar data to identify anomalies.

The experiment datasets comprise five radar measurement sessions collected over two months spanning 6/27/25-8/21/25. A summary of the datasets are provided in Table 3.3.

Table 3.3: Experiment summary table

Date	Growth Stage	Description
6/27/2025	First and Second Trifoliolate (V1—V2)	Early growth stage
7/16/2025	5^{th} Trifoliolate (V6)	Canopy Development, heights
7/29/2025	Beginning Bloom (R1)	Starting flowering, heights
8/15/2025	Beginning Pod (R3)	Starting seed development, heights
8/21/2025	Full Pod (R4)	Mature growth, heights

The radar measurements collected across the five sessions reveal a progression of canopy development throughout the growing season. Each range-time plot represents a transects

along rows, with time and equivalently, distance on the horizontal axis, range from the radar on the vertical axis, and color intensity representing the magnitude of the reflected signal. The following subsections describe the radar data at each collection date and relate them to concurrent physical height measurements. A representative image of each growth stage and radar view can be found in Figure 3.6

Figure 3.6: Overhead soybean season growth

3.5.1 Early Growth - June 27^{th}

At this earliest measurement date, the soybean canopy was in its initial vegetative stages with limited leaf area and short plant height. As the canopy is small and sparse at this stage, both the ground surface and canopy fall within the main beam and are clearly visible continuously in the range-time image. Figure 3.7 shows a path traversed row by row starting with the first row of plot 1811 and continuing until finishing in the last row of plot 1820. Figure 3.8 shows a reconstructed panorama of the first two rows traversed.

The complete range-time image in Figure 3.9a shows strong continuous returns corresponding to the soil surface starting at 1.1 m. The soil surface has been annotated on this plot using the tracker derived in Chapter 4. The ground distance is shown to have variation in distance over time, this is due to the hand carried nature of this experiment. The canopy size at this growth stage has a fairly weak return but this can be seen by time windowing and looking closer at just a few rows. Figure 3.9b shows an annotated range-time image of the first four rows, traversing the first row of plot 1811(1811.1), through the first row of 1911(1911.1), then the second row of 1911(1911.2), and finally to the second row of 1811(1811.2). The average canopy height estimates of these first four plots shown can be

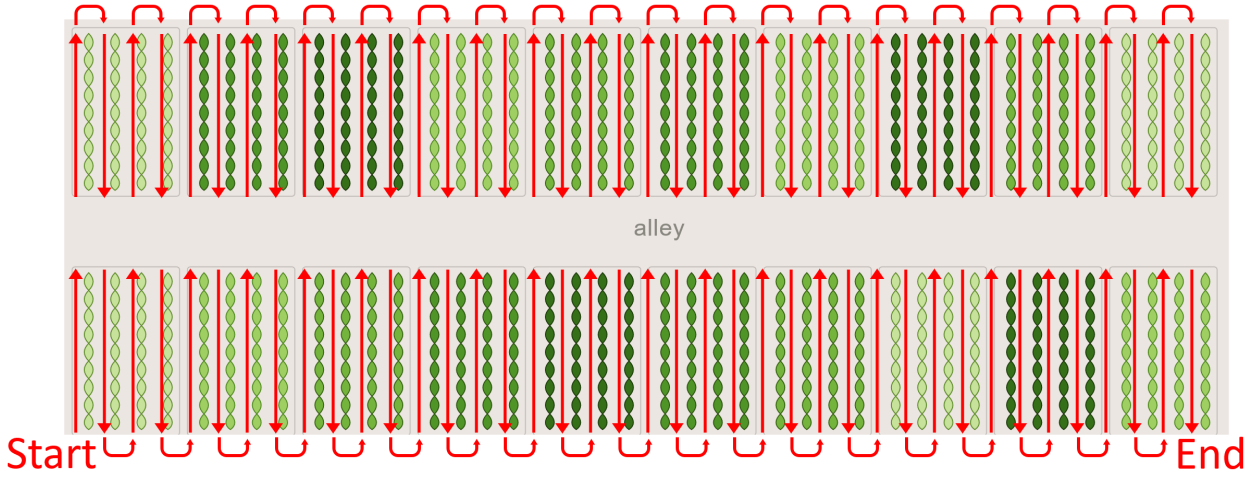

Figure 3.7: Early Growth - June 27^{th} - Measurement path route

found in table 3.4.

Figure 3.8: Early Growth - June 27^{th} - First 2 row (1811.1-1911.1) panorama

Table 3.4: Early Growth - June 27^{th} - Early plot heights

Plot	Row	Height Estimate(cm)
1811	1	12
1911	1	5
1911	2	4
1811	2	7

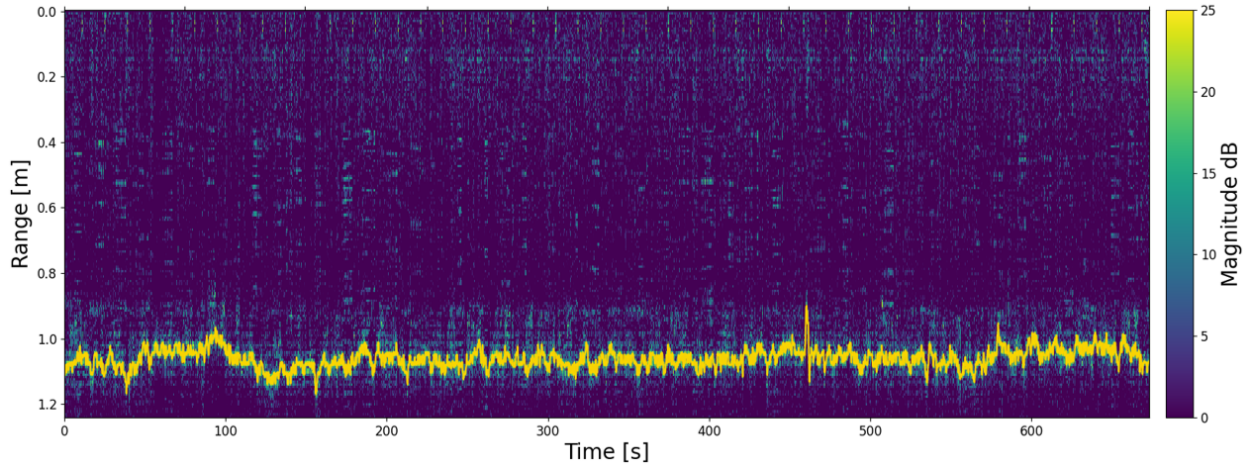

(a) Full range-time plots

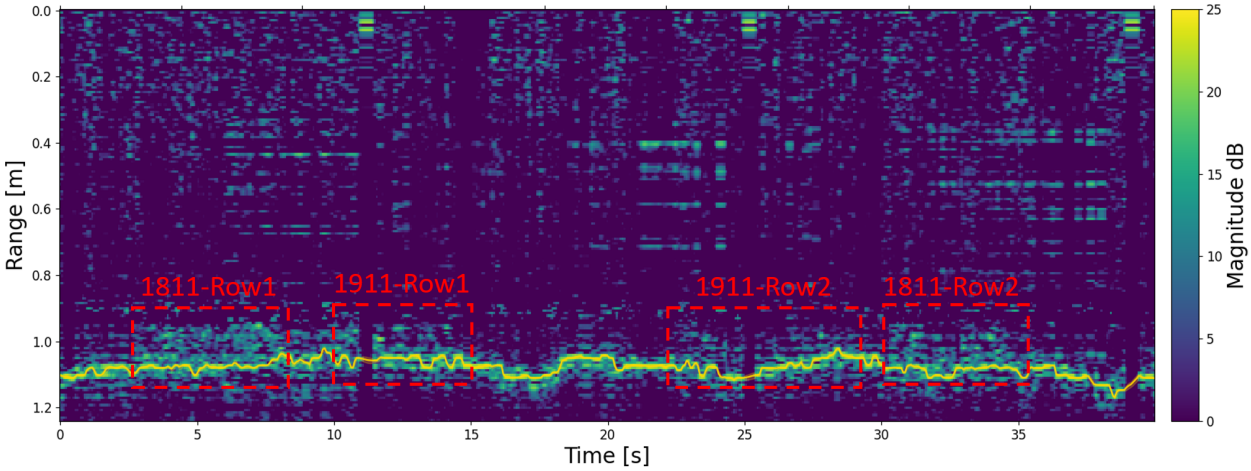

(b) First 4 rows (1811.1-1911.1-1911.2-1811.2)

Figure 3.9: Early Growth - June 27^{th} - Range-time images

3.5.2 Canopy Development - July 16^{th}

By mid-July the canopy had developed substantially since the first measurement, with increased leaf area beginning to limit ground visibility within the row. Figure 3.10a shows a data set that traversed the lower sections (18xx plots) 40 rows. The rows now exhibit a much more visible repeated canopy structure separated by turns in the empty alleys. Figure 3.10b zooms in to show the first eight rows annotated and composed of plots 1811 and 1812. A notable increase in the intensity and vertical extent of the canopy return compared to the June 27 data can be seen. Ground returns within the plot rows are often attenuated heavily as the developing foliage absorbs and scatters most of the incident signal before it reaches the soil surface. However, because this is a breeding trial with bare alleys, the ground surface remains visible in between the rows of the range-time image. At this stage, the ground range estimate needs to start being derived from these inter-row gaps rather than purely from penetration through the canopy.

In the zoomed in range-time image a few other noticeable height changes can also be seen, a dip in the measured canopy height in the middle of row 1811.2, and an overall lower height for 1811.4. Figure 3.11 shows reconstructed panoramic images of plot 1811. In Figure 3.11b you can see a decrease in the canopy size around the middle of the row. Figure 3.11d shows a smaller and less dense canopy which can also be seen with a stronger ground return in the corresponding range-time image. Table 3.5 shows radar height estimate results for the individual 20 plots. Table 3.6 shows the results averaged across genotypes.

While evaluating the row data some showed higher standard deviations in height estimates. Upon further investigation the rows identified showed significant growth inhibition. Plot 1919 row 4 (1919.4) shows heavy inhibition in its middle due to residue from prior cropping, preventing emergence. Figure 3.12 shows a zoomed in range-time image of plot 1919 in which the damaged row 4 (1919.4) can be seen. Figure 3.13 shows the reconstructed panorama of this row.

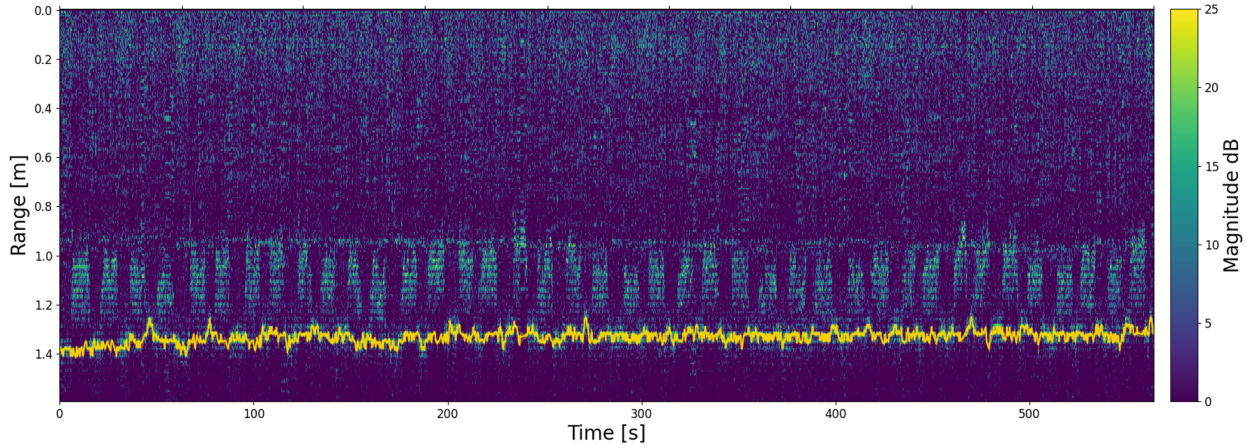

(a) Full range-time plots

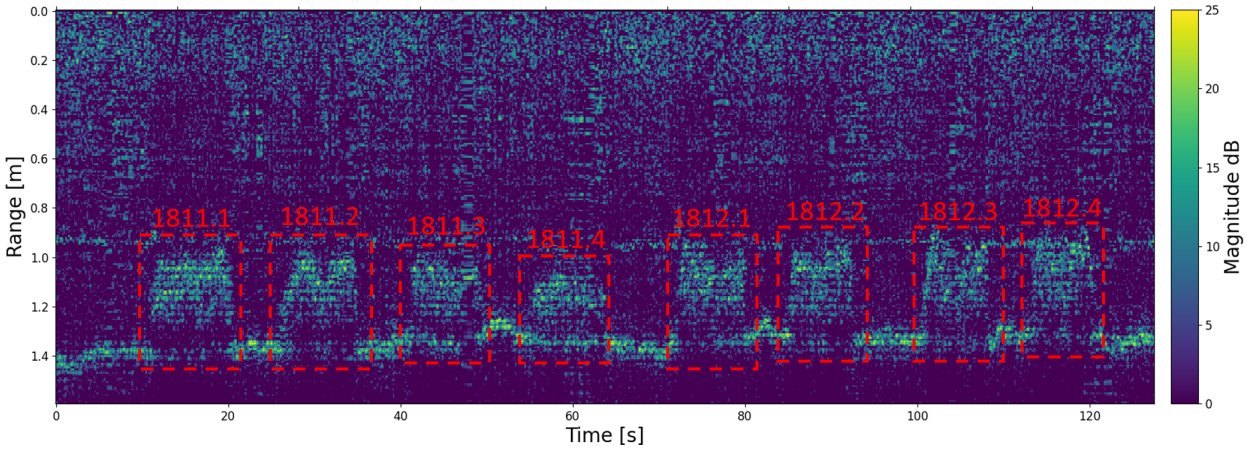

(b) First 2 plots (1811-1812)

Figure 3.10: Canopy Development - July 16^{th} - Range-time images

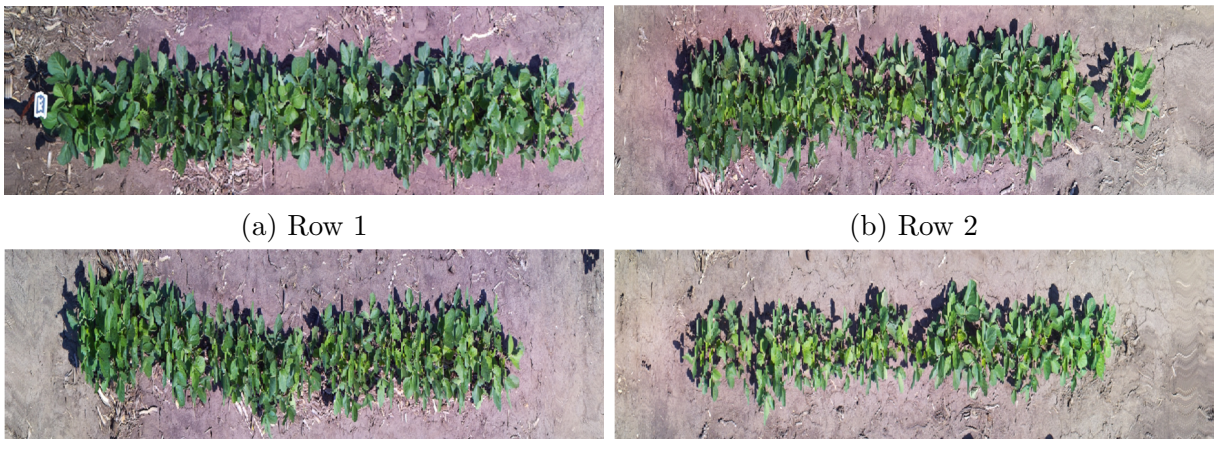

(a) Row 1

(b) Row 2

(c) Row 3

(d) Row 4

Figure 3.11: Canopy Development - July 16^{th} - Plot 1811 panoramic reconstructions

Table 3.5: Canopy Development - July 16^{th} - Height results (plot averages)

Genotype	Plot Number	Average Height (cm)	Maximum Height (cm)	Minimum Height (cm)	Height Standard Deviation (cm)
KS4525NS	1811	30.9	35.5	24.6	4.1
KS4525NS	1818	30.5	31.9	27.6	1.7
KS4525NS	1911	29.8	32.6	27.0	2.5
KS4525NS	1920	28.5	31.1	26.3	1.7
RM3.2	1812	34.4	36.1	33.5	1.0
RM3.2	1820	32.0	36.3	27.7	3.3
RM3.2	1914	31.4	34.6	27.6	2.5
RM3.2	1917	32.4	35.1	29.6	2.1
RM4.2	1813	30.2	32.2	29.2	1.2
RM4.2	1817	27.4	29.3	24.2	1.9
RM4.2	1915	28.6	30.8	26.4	1.9
RM4.2	1919	29.4	34.6	22.0	5.1
K20-1313	1814	32.7	35.3	29.3	2.2
K20-1313	1816	30.2	32.8	26.2	2.5
K20-1313	1912	32.2	36.8	30.0	2.7
K20-1313	1916	33.4	35.2	30.9	1.6
K19-2402	1815	32.6	34.2	29.6	1.8
K19-2402	1819	32.9	37.0	28.7	3.2
K19-2402	1913	34.5	38.1	30.4	2.7
K19-2402	1918	37.1	40.2	30.7	3.8

Table 3.6: Canopy Development - July 16^{th} - Measurement results (genotype averages)

Genotype	Average Height (cm)	Maximum Height (cm)	Minimum Height (cm)	Height Standard Deviation (cm)
KS4525NS	30.0	35.5	24.6	2.5
RM3.2	32.5	36.3	27.6	2.2
RM4.2	28.9	34.6	22.0	2.5
K20-1313	32.1	36.8	26.2	2.2
K19-2402	34.3	40.2	28.7	2.9

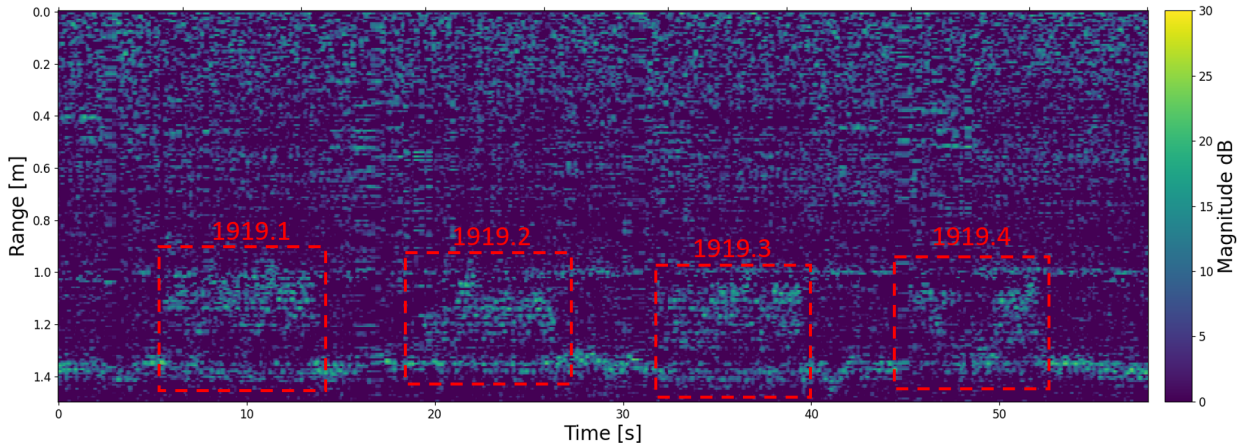

Figure 3.12: Canopy Development - July 16^{th} - Plot 1919 range-time image

Figure 3.13: Canopy Development - July 16^{th} - Plot 1919 row 4 reconstructed panorama

3.5.3 Mid-Season Canopy - July 29^{th}

By late July the range-time images (Figure 3.15 show a single dominant scattering layer corresponding to the upper canopy surface. The canopy reflection is broader in range compared to earlier dates, indicating a more developed vertical canopy structure with scattering contributions from multiple leaf layers. The range-time image exhibits a relatively uniform canopy return with a radar-estimated average height of 53 cm across the plots. The increased canopy density is also reflected in a higher total integrated back-scattered power across the canopy range bins, consistent with greater biomass and leaf area intercepting the radar beam.

To evaluate the repeatability of height measurements and their accuracy compared to manual measurements, a representative row (1917.2) shown in Figure 3.14 was repeatedly measured. The range-time image in Figure 3.15 shows the annotated results of four repeated pass over the row using the ground tracker (yellow) and canopy detection (red) developed in Chapter 4.

Figure 3.14: Mid-Season Canopy - July 29^{th} - Row 1917.2 reconstructed panoramic

Table 3.7 shows a summary of statistics for the four passes compared to hand measurements. The hand measurements were taken every 30 cm along the row and Figure 3.16 shows a comparison of the first pass height estimated derived from the ground tracker and canopy detector compared with the hand measurements. The radar estimate is the solid line and the manual measurement in circular points with a smoothed dashed-line interpolation. The distance along the x axis for the radar measurements was scaled to the time window, this assumed a constant rate of travel and accurate detection window. This comparison plot shows

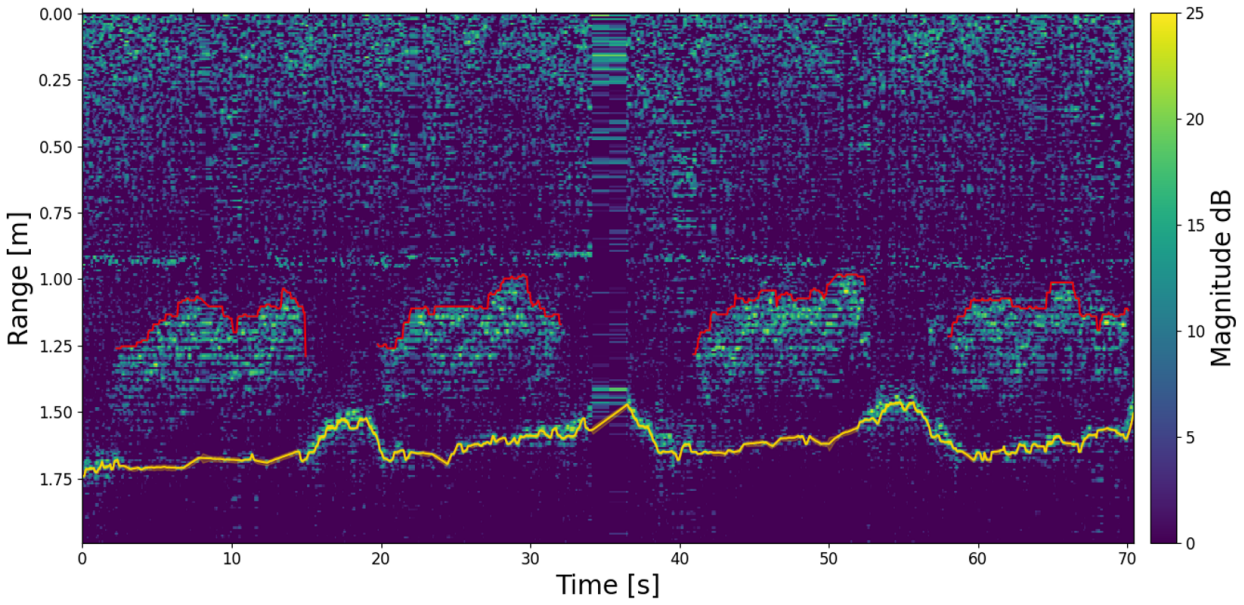

Figure 3.15: Mid-Season Canopy - July 29^{th} - Row 1917.2 repeated pass range-time image

a real advantage of the radar based height measurements in that it samples the height much more often. The comparison plot and statistics show a high degree of agreement between the manual and radar measurements.

Table 3.7: Mid-Season Canopy - July 29^{th} - Row 1917.2 repeated pass statistics

Measurement	Average Height (cm)	Maximum Height (cm)	Minimum Height (cm)	Within-Row Std. Deviation (cm)
Manual	56.3	63.5	43.2	6.2
First Pass	52.6	63.3	32.7	5.9
Second Pass	55.1	63.0	44.4	4.4
Third Pass	53.6	65.0	38.6	6.3
Fourth Pass	53.5	64.9	35.8	6.3
Mean of Radar Passes	53.7	64.1	37.9	5.7

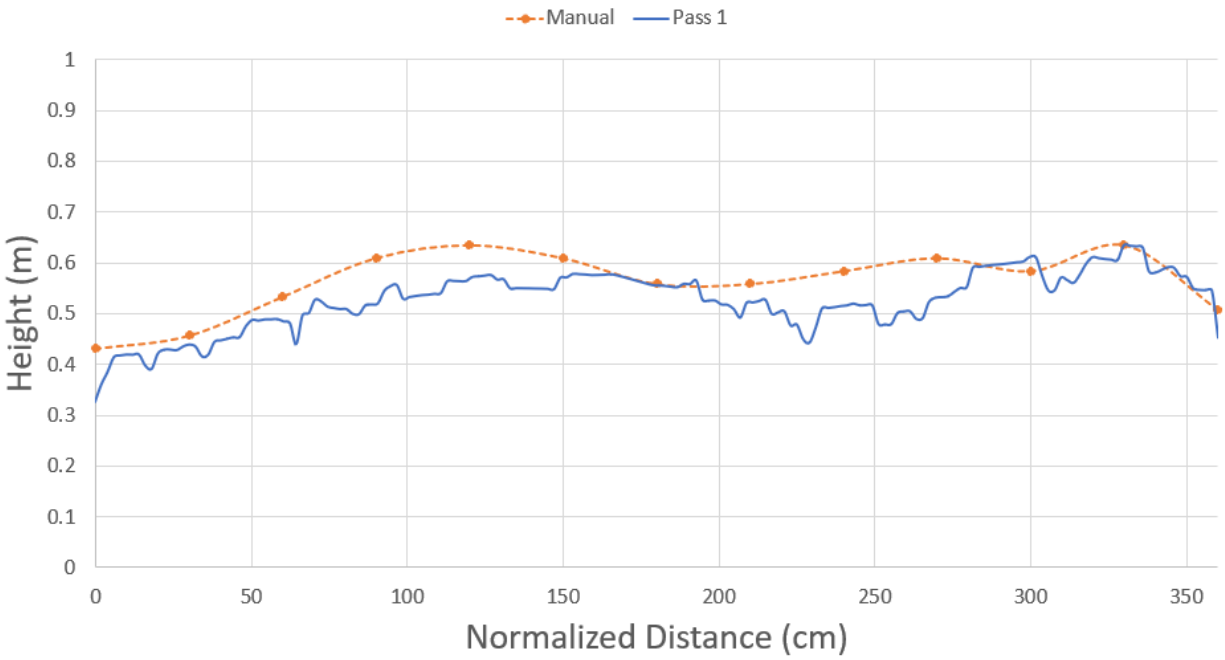

Figure 3.16: Mid-Season Canopy - July 29^{th} - Row 1917.2 height measurement vs. radar estimate comparison

3.5.4 Reproductive Growth - August 15^{th}

At this collection date the soybean plots were in their reproductive stages and beginning to form pods, with the canopy at or near its maximum height and density. Ground returns are effectively absent across the entire row transect, as the fully closed canopy prevents enough signal from reaching the soil surface to produce reflections in the range-time image noise floor. The radar-estimated average canopy height is 82 cm against a manual measurement of 80 cm.

Figure 3.17 revisits the damaged row 1919.4 at this later growth stage. The reconstructed panoramic in Figure 3.17a shows the still apparent damaged section near the middle of the row and also the canopy closure as the adjacent rows completely obscure the ground between at this point. Figure 3.17b shows the range-time image for two repeated passes over the damaged row with the damaged sections annotated with red boxes.

3.5.5 Late Season - August 21^{st}

The final session captured at full pod development with full canopy maturity. The range-time images at this date are similar to the August 15 dataset, with strong continuous canopy returns but no discernible ground reflection within the row. The radar-estimated average canopy height is 100 cm. At this growth stage, some plots exhibited early signs of senescence, which may be reflected in subtle changes to the back-scatter intensity as leaf water content begins to decline.

(a) Reconstructed panorama

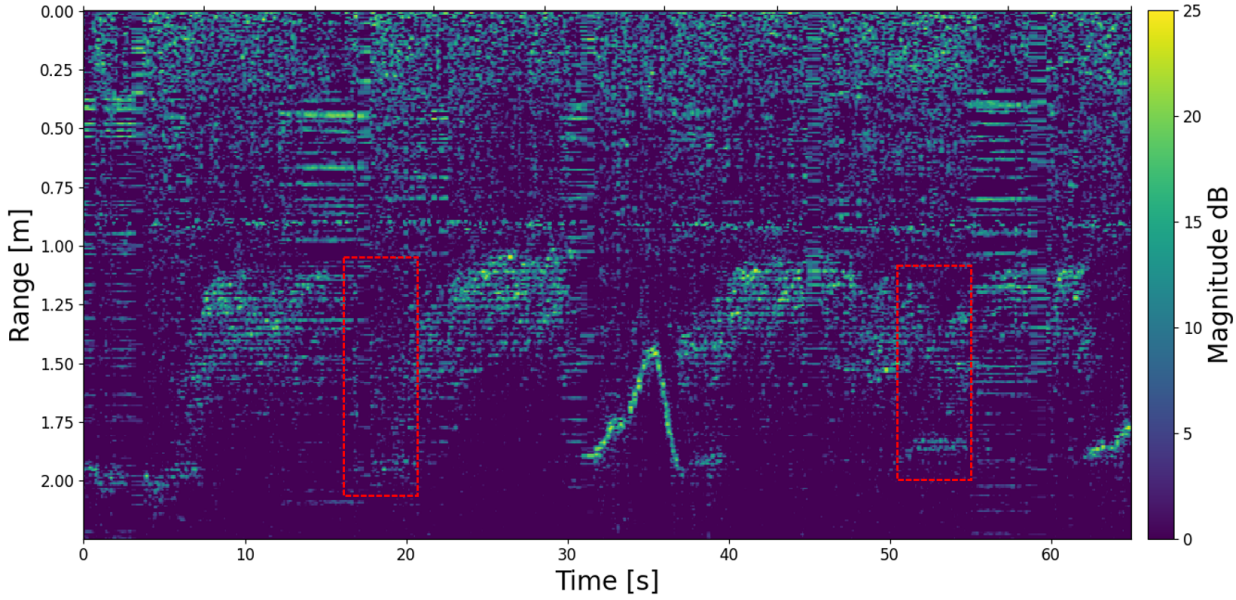

(b) Range-time image

Figure 3.17: Reproductive Growth - August 15^{th} - Row 1919.1

3.6 Summary

This chapter presented the design of a compact 120 GHz SFCW phenotyping radar and its deployment across the 2025 growing season in a soybean breeding trial. Five measurement sessions, spanning 6/27 to 8/21, captured the crop from the early vegetative stages through full pod development, and together they demonstrate both the capability of the system and the way its resulting images change as the canopy develops.

Two distinct measurement regimes emerged over the season. In the early stages (June 27 and July 16), the canopy was sparse enough that the radar resolved both the ground surface and the canopy directly, and the ground return could be tracked continuously beneath the crop. As the canopy closed the gaps between adjacent rows in the reproductive stages (August 15 and 21), the ground return within the row was lost, and the ground reference had to be recovered from the bare inter-row alley gaps characteristic of a breeding-trial layout rather than from penetration through the canopy. The radar-estimated average canopy height tracked this development monotonically across the season, from a few centimeters at emergence to roughly 30 cm in mid-July, 53 cm in late July, 82 cm in mid-August, and 100 cm at full maturity. Where direct manual comparison was available in the closed-canopy regime, the radar estimate agreed closely with hand measurement (82 cm against 80 cm on August 15).

Beyond tracking seasonal development, the July 29 repeated-pass study on row 1917.2 quantified the measurement's repeatability and accuracy. Four passes over the same row produced average heights within a few centimeters of one another and of the manual measurement reference, and the dense radar sampling along the row reproduced the spatial height profile measured by hand while sampling it far more frequently and without the operator-to-operator judgment inherent to the manual method. This combination of repeatable measurements across passes and consistency with ground truths is the central evidence that the system measures a real, stable phenotypic trait rather than an artifact of any single transect.

The season's data also showed the radar's sensitivity to localized agronomic features.

The growth inhibition in row 1919.4, identified initially by elevated height variance and confirmed in the reconstructed panoramas, was tracked through subsequent measurement dates, demonstrating that the system resolves not only plot and genotype level differences but also within-row anomalies relevant to a breeding program.

The primary limitation of the current deployment is that, once the canopy closes and the ground return is lost, the absolute height estimate depends on the stability of the sensor height above the ground. For the hand-carried configuration used here this was sufficient, but it motivates a more stable or self-referencing platform for fully closed-canopy operation, a direction revisited in Chapter 5. The range time images in this chapter and the ground and canopy annotations shown throughout are produced by the processing pipeline developed next in Chapter 4, which formalizes the image creation, tracking, and detection methods used here to produce per-plot canopy height and structural density.

Chapter 4

Phenotyping Radar Processing Pipeline

The radar system from Chapter 3 produces raw stepped-frequency acquisitions, and Section 3.5 showed that with correction these acquisitions reveal the gross progression of the canopy across the season. Processing a per-plot canopy height and a structural density measure automatically requires a more sophisticated processing chain. This chapter develops that processing chain, which converts a raw acquisition into a cleaned image, tracks the ground through partial obscurement, detects the canopy relative to that ground, and integrates the canopy return into a density measure. The chain is wrapped in a calibrate-once, human-in-the-loop workflow in which an experienced technician tunes a small set of physically interpretable parameters on a few representative plots, after which the same configuration is applied without further intervention to the remaining plots acquired under comparable conditions. The expert judgment that would otherwise be spent measuring every plot is instead spent once and amortized across the field. The process augments manual phenotyping and multiplies the throughput of a single expert.

4.1 Overview

The pipeline is organized in four stages. The first is the range-time formation. The second is a system-aware process that corrects the non-idealities present in the system described in Chapter 3. It removes the periodic calibration sweeps inserted by the continuous coherent calibration process, reduces speckle through multilook averaging, suppresses fixed clutter and slow-drift artifacts, and normalizes sweep-to-sweep amplitude variation. The third and fourth stages apply two automated detectors to the image. A particle filter with backward smoother that tracks the ground and reports a per-chirp uncertainty, and a region-based detector that extracts the canopy envelope with corresponding top surface relative to the tracked ground and yields an explicit canopy presence/absence decision. Finally a structural density measure is developed as the integrated canopy energy, which is computed from the region between the tracked ground and the detected canopy top.

A design principle followed throughout is segment awareness. Removing the calibration sweeps leaves the slow-time axis discontinuous, partitioning each acquisition into contiguous segments separated by calibration gaps. Most operations work strictly within these segments and don't combine samples across a boundary. The ground tracker carries continuity across the gaps by initializing each new segment from the smoothed final estimate of the segment that preceded it. Figure 4.1 shows the processing flow. In the following sections (4.2, 4.3, 4.4) a sample dataset will be shown walking through the range-time image formation and cleanup stages.

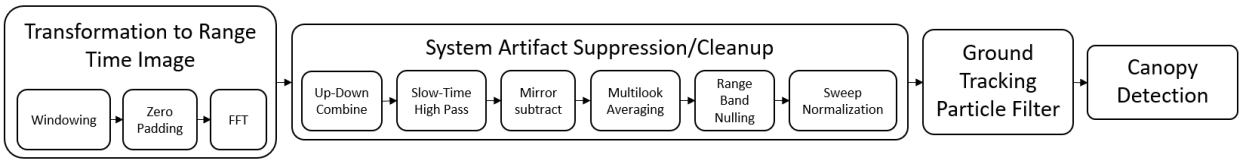

Figure 4.1: Processing pipeline flow diagram

4.2 Range–Time Image Formation

Each measurement consists of an up-down frequency ramp acquisition, producing two complex baseband sequences after down-conversion. An optional windowing is applied to each sequence to reduce the spectral leakage that the rectangular sweep envelope can produce[24], followed by zero-padding with an integer factor. A factor of four will be used in the example images in this chapter. The zero-padding interpolates the range axis onto a finer grid for the downstream detectors without altering the underlying range resolution, which remains fixed by the swept bandwidth from Chapters 2 and 3. A discrete Fourier transform then yields a complex range profile for each ramp.

Because the transform output is complex, targets from the up-ramp appear at positive frequencies while targets from the down-ramp appear at negative frequencies. The down-ramp profile is mirrored to align with the up-ramp, and the two are combined coherently; this improves the signal-to-noise ratio and suppresses ramp-asymmetric artifacts. The combined profile is converted to decibels, and successive profiles are stacked along the slow-time axis to form the range-time image $S[x, y]$, where y indexes range bin and x indexes sweep. The raw and interpolated range-time images are shown in Figures 4.2a and 4.2b.

4.3 Calibration-Sweep Removal and Segment Structure

As described in Chapter 3, the continuous coherent calibration procedure periodically inserts calibration sweeps so that the correction waveform can be re-estimated as the undesired leakage drifts in time. These calibration sweeps do not carry valid target data and are detected and excised. Their removal leaves the slow-time axis discontinuous with the original measurement partitioned into contiguous segments separated by calibration gaps. Subsequent processing follows the segment-aware approach introduced in Section 4.1, operating within individual segments and never combining samples across a boundary. The same principle is carried into the ground tracker of Section 4.5.1, which initializes each new segment from the

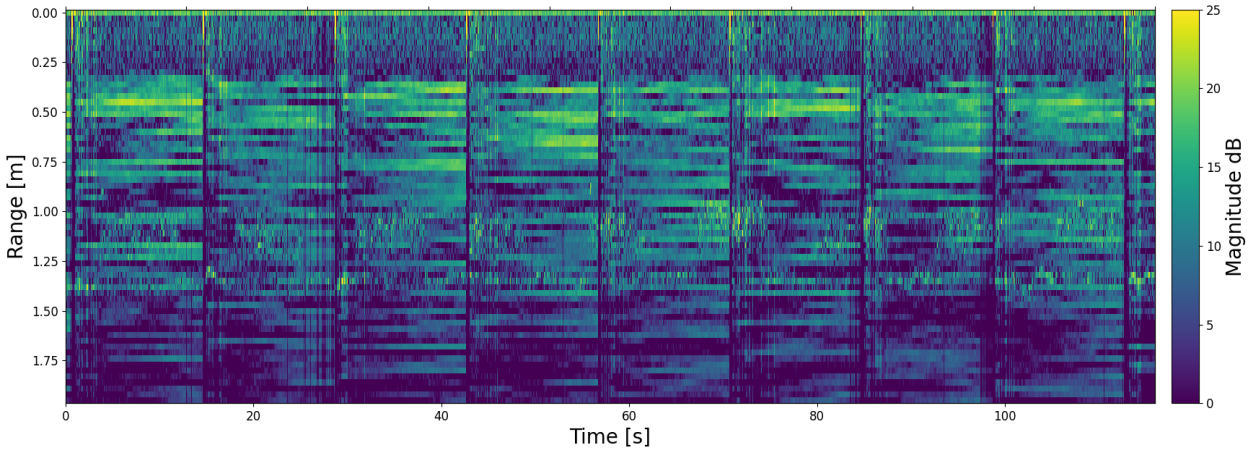

(a) Raw FFT

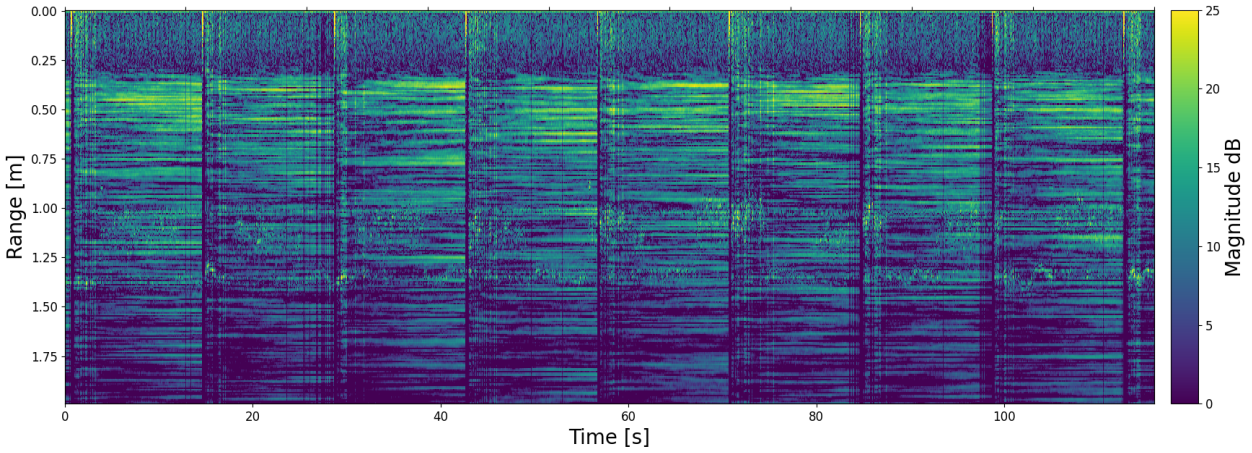

(b) Interpolated FFT

Figure 4.2: Starting range-time image example

smoothed final estimate of the segment that preceded it.

4.4 Cleanup Chain

4.4.1 Fixed-Clutter and Artifact Suppression

Three suppression stages address artifacts with distinct signatures, applied in sequence. The first removes a broadband slow drift that affects many range bins. A Gaussian blur of the image along slow time is computed and subtracted from the original. This in effect is a slow-time high pass filter. The slow-time standard deviation is chosen large compared to the row-feature spacing, so that the subtracted background captures system drift without

removing real along-row variation. Figure 4.3 shows the result.

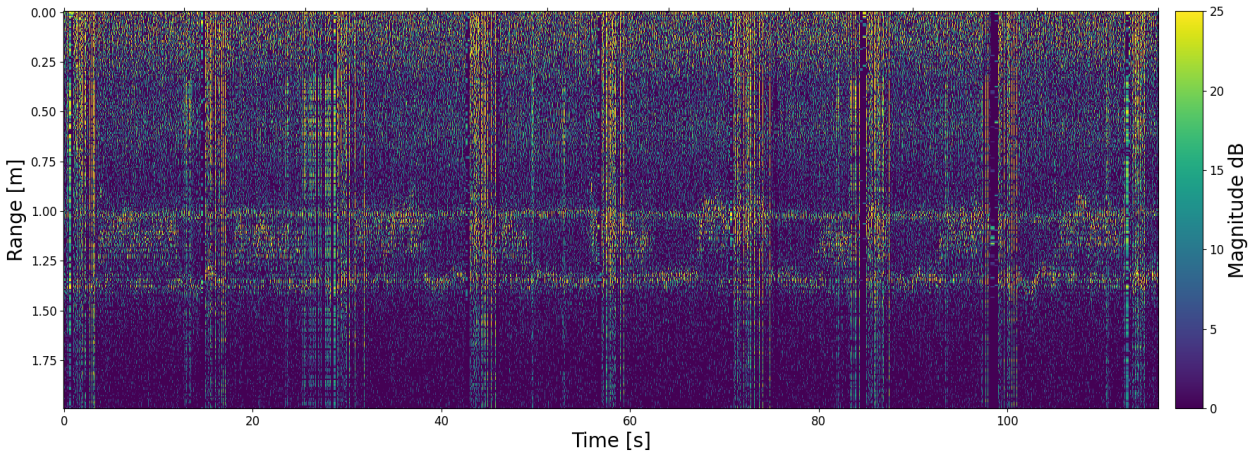

Figure 4.3: Slow-time filter example

The second stage exploits the fact that a purely real-valued baseband signal has a hermitian-symmetric spectrum. This means that at every bin in the positive frequency there is a magnitude-matched twin in the negative spectrum. While the up/down ramp coherent combination in the formation puts genuine targets exclusively in the positive-frequency range, this means that anything that shows up with equal amplitude in both halves is not a target but system clutter. This is a combination of DC offset from imperfect ADC centering, near-DC 1/f noise and IF feedthrough, even-order mixer distortion, and any front-end leakage or transmit-receive coupling whose spectrum is symmetric about zero frequency. To remove these, the linear amplitude of the negative frequency spectrum is mirrored and subtracted from the positive range. Figure 4.4 shows the impact of this

The third stage targets a range band where the known false targets reside, removing them as a horizontal band. A background is estimated within the affected band and subtracted, computed independently for each range bin inside the band. It is calculated as mean in slow time which removes a purely stationary band. Real targets vary on a faster time scale than the background and survive the subtraction, while the stationary artifact does not. A cosine taper at the band edges suppresses ringing from the abrupt range mask. Figure 4.5 shows the result of nulling an artificial target band at the 1.025m distance.

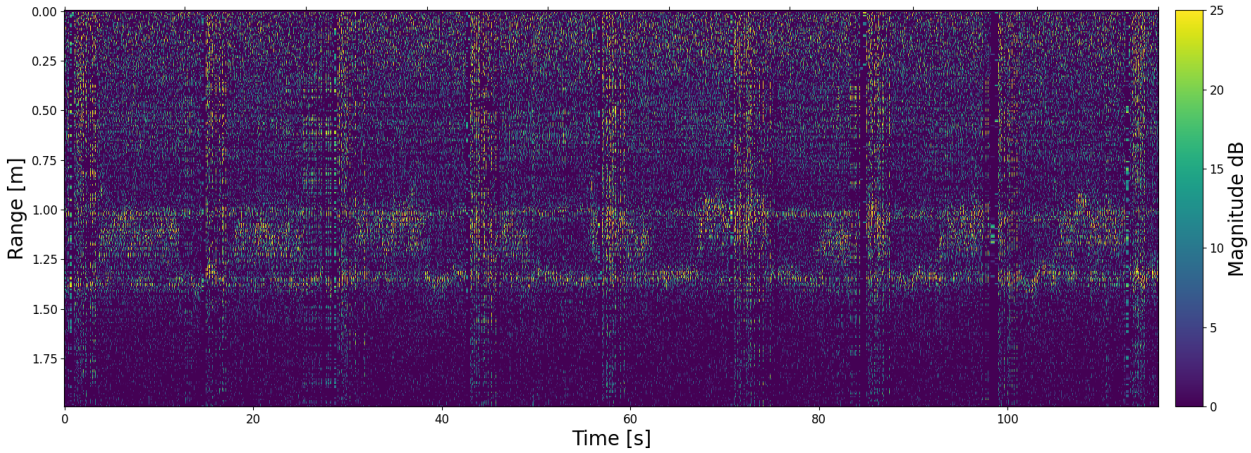

Figure 4.4: Mirror subtract example

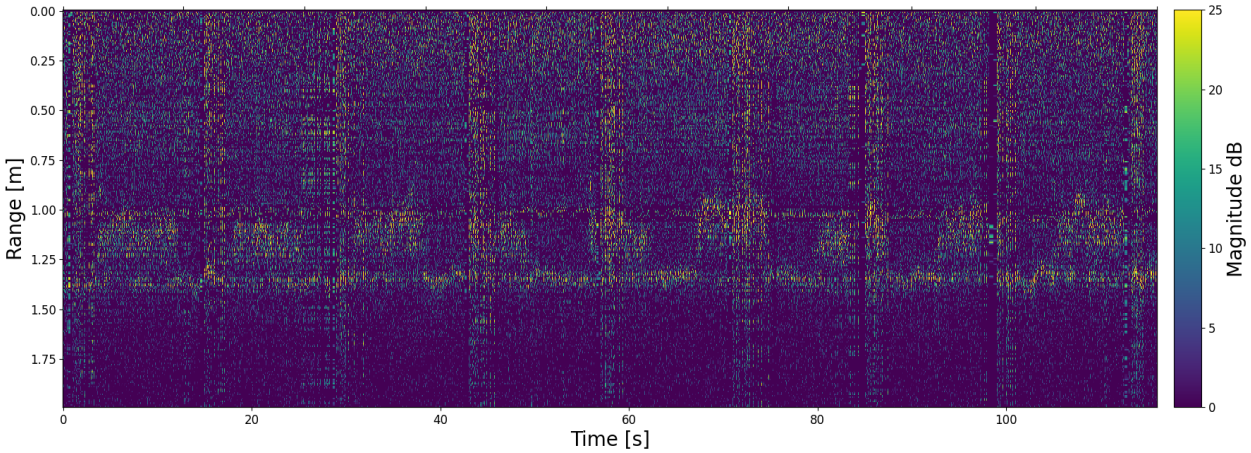

Figure 4.5: Range band nulling example

4.4.2 Multilook Averaging

Speckle from the rough scattering surfaces of soil and vegetation appears as small gaps in what should be continuous returns. It is reduced by multilook averaging of adjacent sweeps, performed in linear power. The averaging is implemented as a slow-time average, whose size is chosen so that leakage across segment boundaries is negligible. The trade-off is a reduction in speckle against effective slow-time resolution. The averaging width is selected so that the post-multilook sampling still resolves the spatial features of interest along the row. Figure 4.6 shows the result of an adjacent sample averaging to increase in signal-to-noise ratio at the cost of temporal resolution.

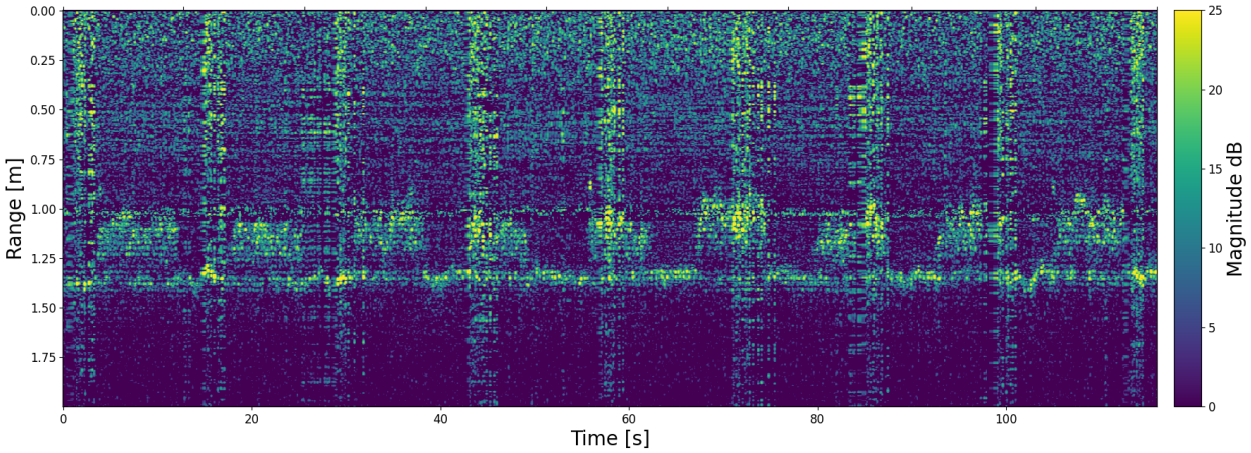

Figure 4.6: Multilook averaging example

4.4.3 Per-Sweep Normalization

Sweep-to-sweep amplitude drift, introduced by temperature and analog gain changes associated with the calibration chain, appears as vertical banding in the range-time image. For each sweep the mean noise power can be determined by evaluating a range where no targets will be present. A good choice of window is for ranges beyond the ground distance. Once the mean noise floor is determined, the sweep is normalized by this value. Figure 4.7 shows the impact of normalization. Since this will scale all the amplitudes the displayed absolute magnitudes are adjusted in this figure but the range displayed remained constant.

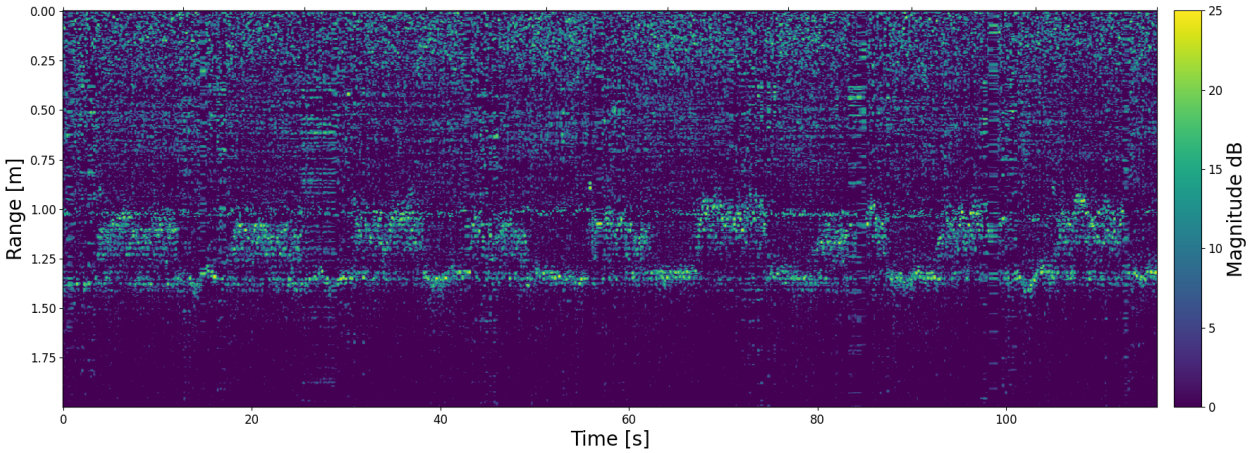

Figure 4.7: Normalization example

4.4.4 Cleanup Parameters

Unlike the detector parameters in the following sections, the cleanup parameters like the slow-time Gaussian widths, the affected range band for nulling, and the multilook averaging count are not dependent on the crop or its growth stage. They are determined primarily by the travel speed of the sensor along the row. As long as the measurement approach is consistent they can be held fixed across acquisitions.

4.5 Tracking and Detection

The cleaned range-time images contain two structures of interest. The ground return should be defined as singular value at each point in slow time whose position varies smoothly across the acquisition, while the canopy is a diffuse region above the ground. Because the two have fundamentally different forms, they are handled by different algorithms. A particle filter with a backward smoother processes the ground, and a region-based detector was used for the canopy.

4.5.1 Ground Tracking Particle Filter

The ground return is generally a strong reflection in unobstructed chirps and traces a smooth horizon across slow time, but it is weakened and can be completely obscured when canopy coverage is dense. A generic peak-picking algorithm would frequently mistake the stronger canopy return for the ground in these regions. A particle filter is a well-suited alternative. It represents the ground position at each chirp by a population of weighted samples rather than a single estimate, which lets it carry several competing hypotheses through an ambiguous period instead of committing to the brightest bin, and it exposes a per-chirp uncertainty as a natural by-product of the weighted spread[26;27]. The general behavior of such a filter through an ambiguous measurement period is illustrated in Chapter 2 (Figure 2.4).

Each particle represents a candidate range-bin position for the ground at chirp x, and the estimate is confined throughout to a search band (r_{low}, r_{high}) that brackets the physically

plausible ground depth for the platform. Every particle is clipped back into this band after it is moved, so the band acts as a hard limit on where the ground can be. The way the population is initialized depends on whether a prior estimate is available. For the first segment of a data set, with no previous information, the particles are drawn uniformly across the search band. For every later segment the population is instead seeded as a tight Gaussian centered on the smoothed final position of the preceding segment, which propagates continuity across the calibration gaps even though no observations bridge them.

For each subsequent chirp the filter performs three operations: resampling, propagation, and observation reweighting. Resampling is done by drawing a new particle set from the previous chirp's weighted population by systematic resampling. A single uniform random offset is drawn and combined with a set of evenly spaced strata, and the corresponding samples are selected from the cumulative weight distribution. This concentrates the population on the high-weight samples from the previous step and discards low-weight ones, while the even stratification keeps the resampling variance low relative to drawing each sample independently. Propagation is accomplished by perturbing each resampled particle by an independent Gaussian step in range, encoding that the ground varies smoothly between adjacent chirps. With y_t the ground estimate at chirp t, y_{t-1} the resampled position, and $\mathcal{N}(0, \sigma^2_{\text{mot}})$ the random-walk step shown in Equation 4.1.

$$y_t = y_{t-1} + \mathcal{N}\left(0, \sigma^2_{\text{mot}}\right). \tag{4.1}$$

Finally an observation reweighting assigns each propagated particle a weight from the image amplitude S at its range bin. Using S_{floor} for a lower clamp on the amplitude and α for a likelihood-shaping scale.

The clamp S_{floor} ensures that particles sitting in sub-noise regions all receive effectively equal weight rather than being separated by noise fluctuations, and α sets how sharply the weight responds to amplitude. A small α makes the filter commit aggressively to the brightest bin, while a large α spreads weight more evenly and keeps competing hypotheses alive. In practice the weights are computed in the log domain and normalized by subtracting the per-

chirp maximum before exponentiation, so that the exponent in Equation 4.2 is a numerically stable soft maximum over the scaled amplitudes. The proportionality in Equation 4.2 is resolved by normalizing the weights to sum to one.

$$w \propto \exp\left(\frac{\max(S, S_{\text{floor}})}{\alpha}\right), \tag{4.2}$$

S_{floor} acts as a lower clamp on the amplitude, so that particles in regions below the noise floor receive equal weight rather than being discriminated by noise, and α is a likelihood-shaping parameter.

An additional depth bias addresses a physical assumption for a downward-looking radar, that the ground is the farthest viable return. When canopy and ground returns are comparable in strength the farther of the two should be preferred. A linear ramp $\beta \cdot r_{frac}$ is added to the amplitude before exponentiation, where $r_{frac} \in [0, 1]$ is the normalized position of the particle within the search band (0 at r_{low}, 1 at r_{high}) and β is a user-set gain. The ramp is gated on particles already above the noise floor through the indicator $\mathbf{1}[S > S_{floor}]$ which is 1 when the argument is true and zero otherwise. This biases the choice among genuine returns without lifting sub-noise samples into the high-weight tail. This yields the depth-biased weighting shown in Equation 4.3.

$$w \;\propto\; \exp\left(\frac{\max(S,\, S_{\text{floor}}) \;+\; \beta\, r_{frac} \mathbf{1}[S > S_{\text{floor}}]}{\alpha}\right), \tag{4.3}$$

The three operations constitute a causal forward pass. The resulting weighted population at each chirp reflects only the evidence up to that chirp. Since the data set is processed post-acquisition, a backward pass then refines the forward estimates using future evidence. Working from the final chirp backward, the stored forward particles at chirp t are re-weighted by their transition probability to the already-smoothed mean at chirp $t+1$. Using the same Gaussian motion model used in propagation evaluated between each forward particle and its next smoothed position. The smoothed estimate at t is taken as the resulting weighted mean. This is a mean-targeted simplification of the canonical marginal particle smoother[28] rather than summing transitions over the full forward–backward particle pair distribution,

it targets only the smoothed mean of the following chirp, which is inexpensive and sufficient for a single smooth ground estimate.

The smoother returns a smoothed mean $y_g(x)$, taken as the final ground estimate, and a standard deviation $\sigma_g(x)$ computed as the spread of the smoothed population. The deviation is reported as a per-chirp uncertainty band. This uncertainty widens in regions where the return is ambiguous, primarily where canopy shadowing weakens the ground amplitude and the weighted population spreads across more range bins.

The tracker is run independently within each segment, with the smoothed final position of one segment handed forward to initialize the next as described above. The parameters σ_{mot}, S_{floor}, α, and β, together with the prior range estimate, constitute the tracker parameters. Of these, σ_{mot} is fixed primarily by the measurement implementation. As it reflects the height stability of the platform during a pass, while S_{floor}, α, and β depend on canopy intensity and are growth-stage dependent. These are set as part of the per-stage calibration. Figure 4.8 shows the smoothed track overlaid on the cleaned range-time image and Table 4.1 lists the parameters used.

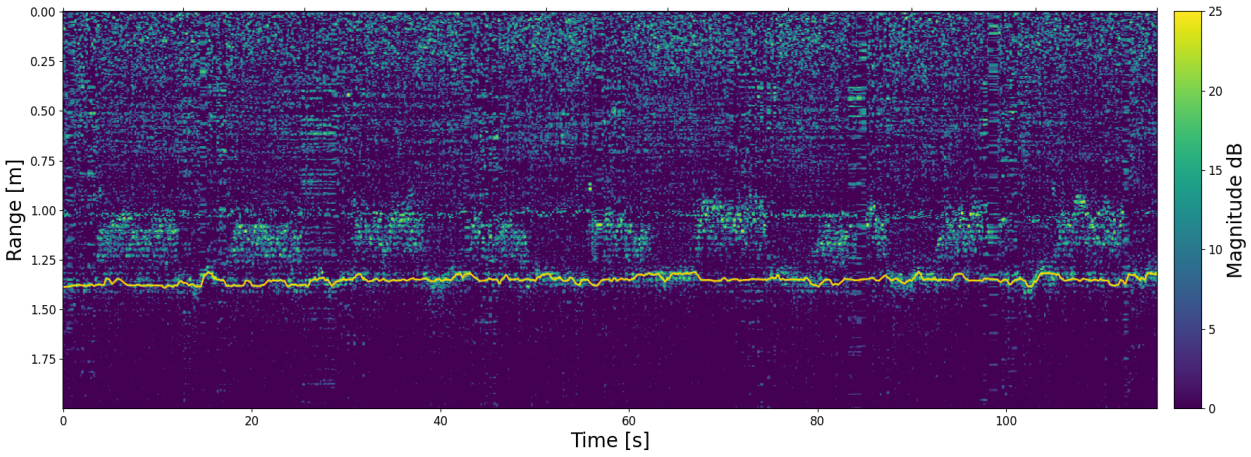

Figure 4.8: Ground tracker example

4.5.2 Canopy Region Detection

Canopy returns appear as a region of higher-amplitude returns above but near the ground, with intensity and extent that vary across the plot. The canopy is not a continuous horizon

Table 4.1: Ground tracker example parameters

Parameter	Value
σ_{mot}	0.1 m
S_{floor}	3 dB
α	2
β	100

at a single range, so a per-chirp canopy surface must be defined relative to the canopy region as a whole rather than picked out as a single peak. The detector operates in three steps.

First, a search band is defined for each chirp x relative to the tracked ground. It is bounded on the far side (nearest the ground) by a small standoff h_{min} that excludes the ground return and its immediate skirt, and on the near side (nearest the radar) by the maximum expected canopy height h_{max} for the growth stage. Only bins between these two limits are eligible, and chirps for which the ground estimate is undefined are not searched. Defining the band relative to the tracked ground rather than to a fixed range makes the detector insensitive to the overall platform height and to slow vertical drift of the ground across a plot.

Second, within the search band, a binary mask of canopy candidates is formed by thresholding the image at τ dB. The mask is then morphologically closed to bridge small dropouts inside an otherwise continuous canopy for gaps where a few interior bins momentarily fall below the threshold. The closing uses a structuring element of independent slow-time and range dimensions, k_t and k_r, whose shape is selectable as rectangular and elliptical. The rectangular element closes uniformly in both directions while the elliptical element bridge gaps more selectively along the two axes. After closing, the search-band window is re-applied so that the closing operation cannot grow candidates outside the band. In the example shown a rectangular kernel was used.

Finally, connected components are extracted from the closed mask and components whose slow-time extent falls below a minimum width w_{min} are discarded. This removes brief bright flecks that are not sustained canopy. For each chirp that retains at least one surviving component, the canopy top $y_c(x)$ is taken as the shallowest range bin (nearest the radar)

among the surviving components in that chirp. Chirps with no surviving component are flagged canopy-absent and excluded from downstream height computation.

Two additional refinements can extend this baseline detector. The first is a per-chirp thinness gate to remove detections that fall below a thickness threshold, which helps to remove detections of narrow false target bands. The second is a two-pass hysteresis mode for noisier images. This uses a stricter seed threshold τ_{seed} to first mark high-confidence canopy elements. The seed components are cleaned in a similar manner as the previously discussed masks. Then the more permissive τ mask is then labeled into connected components, and only those components that contain at least one seed bin are retained. The effect is that a bright, confident core anchors each region while its true extent is recovered at the lower threshold, so that an isolated patch of intermediate-amplitude return is admitted only when it is connected to a confident seed.

The parameters τ, w_{min}, h_{max}, h_{min}, k_t, and k_r govern the baseline detector, with τ_{seed} and the seed-cleanup widths added when using the two-pass mode. As with the ground-tracker parameters of the previous section these depend on the crop and growth stage and are set during the per-stage calibration. Figure 4.9 shows a closed detected canopy region with the resulting canopy-top trace using the parameters in Table 4.2.

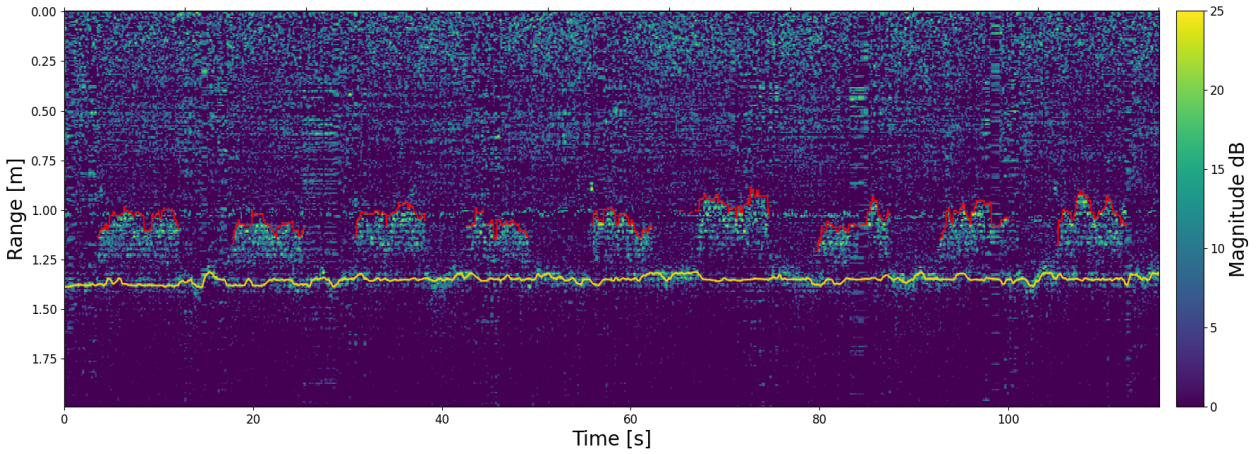

Figure 4.9: Region detection example

Table 4.2: Region detection example parameters

Parameter	Value
k_t	4
k_r	3
τ	3 dB
$w_{\min}$	20 bins
$h_{\max}$	0.5 m

4.6 Integrated Canopy Energy

A new metric the radar "Integrated Canopy Energy" is proposed as a metric related to the structural density of the canopy. It is evaluated over the region between the tracked ground and the detected canopy top for each chirp x and is defined as

$$E(x) = \sum_{\substack{y \\ y_c(x) \le y \le y_g(x) - \delta}} 10^{S[x,y]/10}\, \Delta r, \tag{4.4}$$

where $y_c(x)$ is the canopy top from Section 4.5.2, $y_g(x)$ is the smoothed ground estimate from Section 4.5.1, δ is a small ground skirt that excludes ground-bleed contributions, Δr is the range bin width (calculated as the radar range resolution and scaled by interpolation), and $S[x, y]$ is the cleaned magnitude in dB. The integration is performed in linear power so that bright bins contribute in proportion to their physical scattering cross-section. Chirps for which the canopy region is absent (y_c undefined) or collapsed ($y_g - \delta < y_c$) are flagged invalid and excluded from any downstream per-plot averaging. The per-plot integrated energy is the mean of $E(x)$ across the valid chirps inside each plot segment.

4.7 Evaluation

This section evaluates the pipeline for three areas. First, the transfer of a single tuned parameter set across measurement batches acquired under comparable conditions. Second, the agreement of the integrated-energy density measure with an independent optical canopy-cover estimate. Third, the processing throughput of the semi-automated workflow relative

to manual measurement. Cross-stage limitations are then discussed.

4.7.1 Workflow Demonstration

Because the processing pipeline was developed after data collection, the cleanup and detection parameters were determined by tuning on the first few rows of a large dataset to act as representative plots then applied to the remainder of that set and an additional set taken on the same day without further intervention. Figure 4.10 shows a poorly performing parameter set in which an over-sized morphological kernel ($k_r = 5$), combined with calibration artifacts not fully removed in cleanup, produced a poor canopy estimate that over estimated the canopy envelope and began to track false targets. It is included to show the failure mode that the tuning step is intended to eliminate.

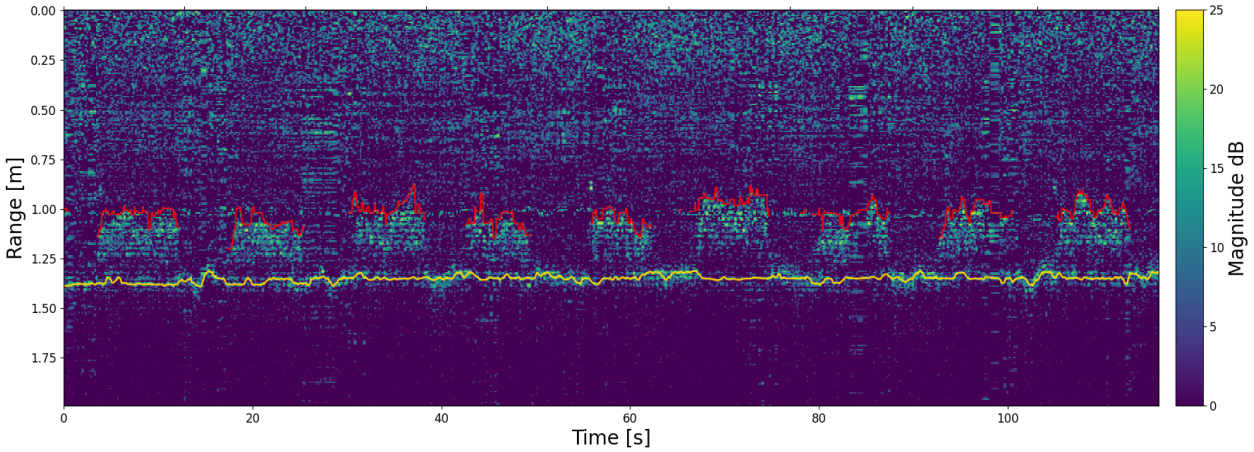

Figure 4.10: Bad kernel size example

With the parameters from the prior Table 4.2, the tracker and canopy detector were applied to all rows in the dataset with no parameter adjustment and to a second dataset containing the remaining rows. Figure 4.11 shows the ground and canopy estimates overlaid on the cleaned range–time images for both batches.

The successful application of a single tuned parameter set to both batches, with no observed need for re-tuning, supports the calibrate-once-scale-to-the-field claim of Section 4.1 within a single growth stage.

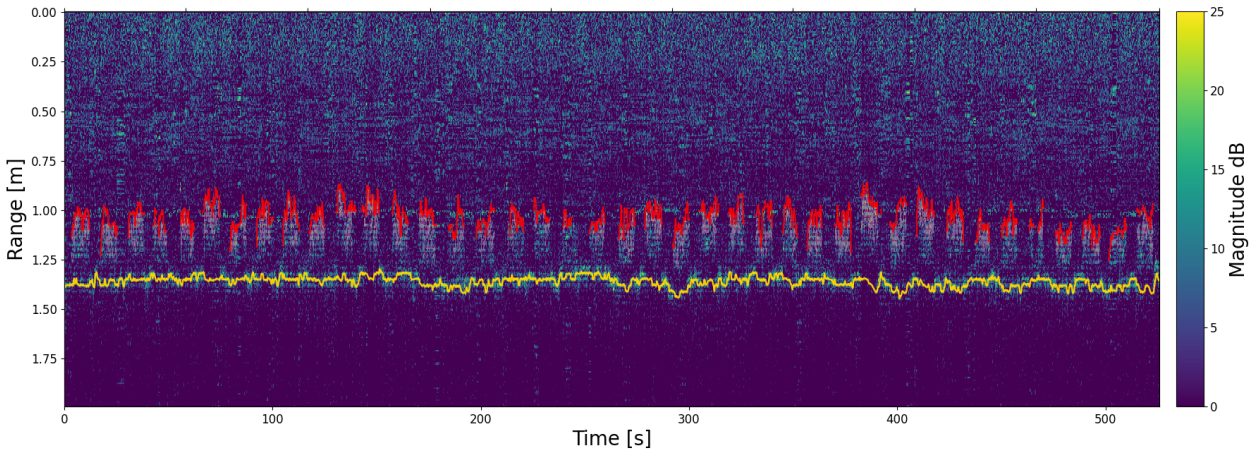

(a) First 40 rows, Plots 18xx

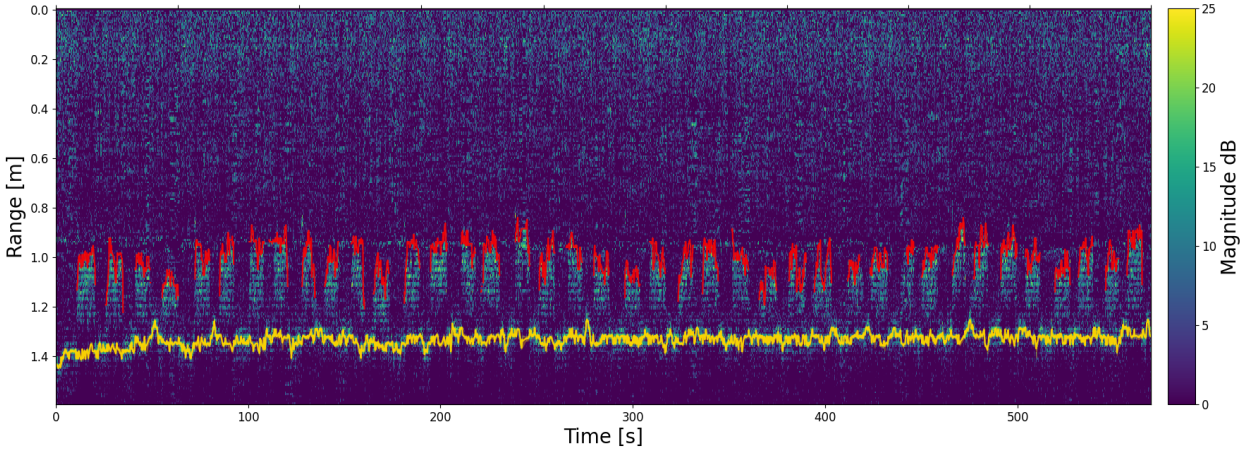

(b) Second 40 rows, Plots 19xx

Figure 4.11: Tracking/Detection applied to large data sets

4.7.2 Integrated Energy and Optical Canopy Cover

To assess whether the integrated-energy phenotype $E(x)$ of Section 4.6 captures structural canopy density in a way that agrees with an independently derived measurement, it is compared against a canopy-coverage estimate extracted from synchronous, co-located video. Canopy cover is extracted from each video frame by greenness-based image segmentation over a defined area centered on the crop row, producing a continuous estimate of canopy-coverage percentage that is related to canopy density. This is the same optical reference introduced in Section 3.5.

Comparing the per-plot average (across rows) integrated energy with the per-plot average video canopy coverage in a middle growth stage from the July 16^{th} data set yields a Pearson

correlation of $r = 0.837$ when assessed over the entire breeding trial. Table 4.3 shows the averaged data per-plot, Table 4.4 shows them for genotypes, and Figure 4.12 shows this relationship across the plots.

Table 4.3: July 16^{th} - Radar Energy, Video Canopy (Plot Averages)

Plot Number	Genotype	Average Radar Integrated Energy	Average Video Canopy Estimate (%)
1811	KS4525NS	0.822	41.5
1812	RM3.2	0.887	45.6
1813	RM4.2	0.752	36.5
1814	K20-1313	0.870	46.4
1815	K19-2402	0.793	45.4
1816	K20-1313	0.734	39.7
1817	RM4.2	0.589	31.4
1818	KS4525NS	0.782	40.0
1819	K19-2402	0.831	46.5
1820	RM3.2	0.815	42.1
1911	KS4525NS	0.672	38.4
1912	K20-1313	0.722	43.6
1913	K19-2402	0.767	44.6
1914	RM3.2	0.646	37.0
1915	RM4.2	0.516	34.2
1916	K20-1313	0.704	43.5
1917	RM3.2	0.728	40.0
1918	K19-2402	0.827	50.6
1919	RM4.2	0.590	33.6
1920	KS4525NS	0.557	35.4

4.7.3 Measurement Cost and Information Density

In conventional breeding trials the measuring protocol for plot height involves taking a single representative height measurement with a height stick. This is fast, and the workflow presented here does not beat it on elapsed time, but for the twenty-plot demonstration trial the two are comparable. They do not however, produce the same measurement. The manual protocol returns one value per plot, whereas the radar returns a continuous height profile along the row at an effective along-row sample spacing of approximately 1.5 cm, or roughly 240 samples per 3.6 m row.

Table 4.4: July 16^{th} - Radar Energy, Video Canopy (Genotype Averages)

Genotype	Average Height (cm)	Radar Integrated Canopy Energy	Video Canopy Estimate (%)
KS4525NS	30.0	0.708	38.8
RM3.2	32.5	0.769	41.2
RM4.2	28.9	0.612	33.9
K20-1313	32.1	0.757	43.3
K19-2402	34.3	0.804	46.8

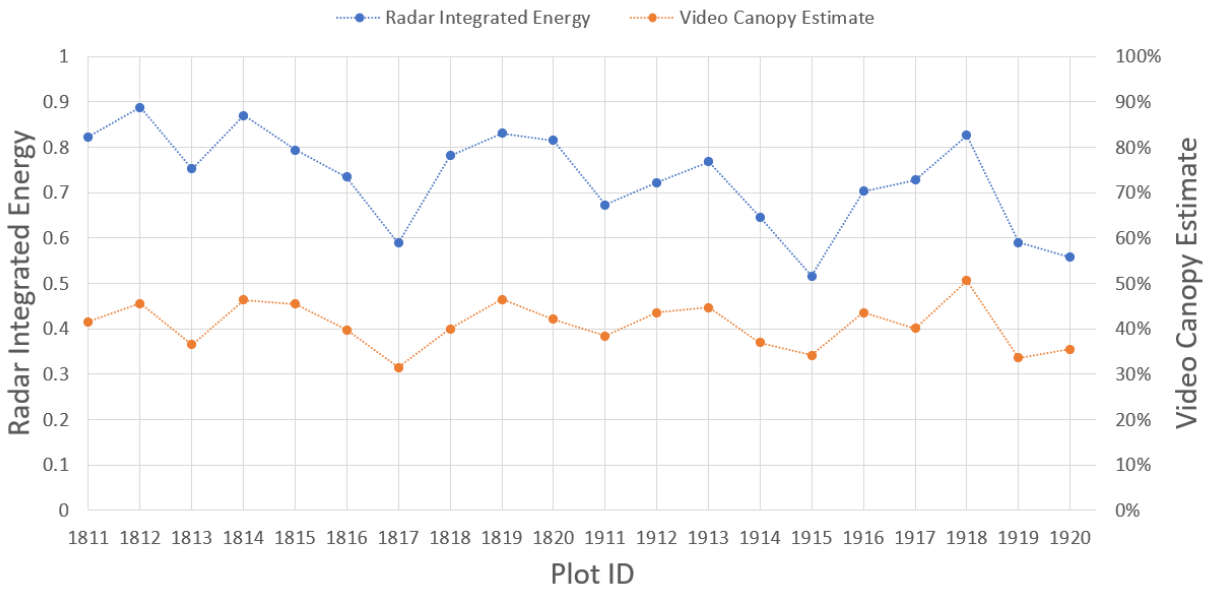

Figure 4.12: Radar integrated energy vs. video canopy estimate

Acquiring a comparable profile by hand across a whole trial is infeasible. The ground-truth transect of Section 3.5.3 was collected at 30 cm intervals, thirteen readings per row, at a working rate of approximately two minutes per row. Extended to the eighty rows of the demonstration trial it implies 160 minutes of continuous skilled labor that scales linearly with trial size. The radar acquisition covered forty rows in nine minutes, an average of 14 seconds per row including travel to the adjacent row, giving roughly 19 minutes for eighty rows. Adding the single tuning session on a small number of representative rows yields approximately 20 minutes end to end. At matched spatial sampling density the workflow is therefore about eight times faster, and it delivers a sample spacing an order of magnitude finer than the 30 cm reference it is being compared against.

The second benefit is the precision of the plot-level height itself. The within-row height variation recorded in Table 3.7 has a standard deviation of 6.2 cm, with values ranging from 43.2 to 63.5 cm along a single 3.6 m row. A single height-stick reading is one sample from that distribution, so the conventional protocol carries a sampling uncertainty on the plot mean of the same order, roughly 6 cm. The four repeated radar passes over the same row returned mean heights of 52.6, 55.1, 53.6, and 53.5 cm, a spread of 1.04 cm, and this includes detection and traversing variation rather than sampling variation alone. The dense profile thus estimates the plot mean more precisely than a single manual reading, which matters directly for genotype evaluation.

Finally, the labor the two protocols consume is of different kinds. Manual scoring requires a trained scorer present for every plot, within the narrow window in which the trial must be scored, and introduces inter-rater variation across a season. The workflow presented here requires an untrained operator to traverse the rows and concentrates expert judgment into one tuning session that is then amortized across the field.

4.7.4 Value of the Spatial Profile

The dense height profiles support analyses that a single per-plot sample cannot. Three are relevant to a breeding program.

First, growth development. The profile is registered along the row, so repeated sessions across the season produce a within-plot growth trajectory rather than a sequence of averages. The rate of growth and its uniformity can be resolved at sub-plot scale.

Second, plot health and outlier detection. Stand establishment failures, gaps, lodging, and edge effects appear as localized departures in the height profile. Figure 3.16 shows within-row structure of this kind, with height varying around 20 cm along a single row. A single reading either misses such a feature or lands on it and biases the plot value with no indication that it has done so. The profile makes the feature visible and allows the affected segment to be flagged or excluded before the plot average is computed.

Third, within-plot variability as a phenotype. The along-row standard deviation and the

shape of the height distribution are measures of plot uniformity that a single per-plot value cannot provide. Because the profile is spatially registered, sub-plot data can also be carried into the spatial models used to correct for field heterogeneity in trial analysis.

4.7.5 Cross-Stage Limitations

The demonstration above applied a single calibration within a single growth stage. The workflow is not expected to generalize across stages without re-tuning, because several detector parameters are tied to physical canopy properties that change through the season. The maximum expected canopy height $h_{\max}$ that defines the search band of Section 4.5.2 must move upward as the canopy grows. A band tuned for an early stage will exclude the canopy of a later stage entirely. The detection threshold τ depends on the strength of the canopy return relative to background, which varies with leaf area, water content, and incidence geometry. The minimum component width $w_{\min}$ depends on stand density and on the row-spacing intermittency typical of the stage. The depth-bias strength β of the tracker depends on how aggressively the canopy obscures the ground return, which changes substantially between sparse early canopies and dense late ones. A single fixed parameter set is therefore not expected to perform well across stages, and the workflow assumes per-stage re-tuning. Removing this dependence by automating the parameter-selection step is a direction of future work and is discussed further in Chapter 5.

4.8 Summary

This chapter presented a processing pipeline that converts raw 120 GHz SFCW range-time datasets into per-plot canopy height and an associated structural density measure. The pipeline is organized around a calibration-aware cleanup chain that includes calibration-sweep detection and excision with segment-aware downstream processing, multilook averaging, fixed-clutter and slow-drift suppression, and per-sweep normalization. The cleanup chain feeds two automated detectors; a particle filter with a backward smoother that tracks

the ground and reports a per-chirp uncertainty; and a region-based detector that extracts the canopy top relative to the tracked ground. With the detected ground and canopy an integrated-energy measure was derived to provide a proxy for structural density. Wrapped in a calibrate-once, human-in-the-loop workflow, the pipeline was shown to transfer a single tuned parameter set across measurement batches within a growth stage, to produce an integrated-energy measure that agrees with an optical canopy-cover reference, and to deliver a sub-plot height profile at a spatial sampling density that manual measurement cannot practically match, at roughly an eighth of the elapsed time required to collect a comparable profile by hand. The principal limitation of this approach is the need to re-tune the detector parameters across growth stages which motivates the automated parameter-selection work discussed in Chapter 5.

Chapter 5

Conclusions and Future Work

This dissertation set out to address two gaps identified in Chapter 1. First the absence of a compact, plot-scale, millimeter-wave radar instrument designed for canopy phenotyping. Second the absence of a processing methodology to convert such an instrument's raw acquisitions into per-plot phenotypic traits in a scalable and repeatable way. The preceding chapters presented a 120 GHz stepped-frequency continuous-wave radar built for plot-scale phenotyping, its field deployment across a soybean growing season, and a processing pipeline. The pipeline turns raw data into range-time images, per-plot canopy height, and provides a structural density measure. This chapter summarizes the contributions, draws conclusions at the system level, states the limitations, and outlines directions for future work.

5.1 Summary of Contributions

The work in this dissertation produced the following contributions, which together correspond to the goals set out in Section 1.5.

A compact 120 GHz SFCW radar sensor with embedded acquisition system was designed and implemented in a low-power, portable form factor suitable for breeding-trial deployment, as described in Chapter 3. The instrument pairs a highly integrated millimeter-wave transceiver with a microcontroller and analog components to form an acquisition chain, and

communicates data over a robust CAN-FD interface intended for integration with agricultural implements or platforms.

A continuous coherent calibration technique was applied to mitigate the transmit-receive leakage that often limits compact continuous-wave radars. By estimating the undesired low-frequency components and feeding back a scaled inverse correction ahead of the receiver gain stages, the system cancels the large coupling that would otherwise saturate the receiver, preserving the dynamic range needed to detect the weak canopy and ground returns characteristic of proximal phenotyping. An initial surface discrimination test demonstrated the instrument's ability to distinguish smooth surfaces from plant biomass across a range of targets.

The system was evaluated in a soybean breeding trial at Kansas State University over the 2025 growing season. Five measurement sessions spanned most of the crops development. The radar tracked canopy development and produced crop height estimates in agreement with manual ground truth measurements across multiple growth stages. This included close agreement in the closed-canopy regime (82 cm radar against 80 cm manual on August 15^{th}). An experiment further established that the height estimate is repeatable across passes and consistent with the hand-measured spatial profile along the row.

The processing methodology of Chapter 4 begins with a cleanup chain that forms range-time images and removes the system's non-idealities such as periodic calibration sweeps, fixed clutter and antenna-coupling artifacts, slow drift, speckle, and sweep-to-sweep amplitude variation. This yields a cleaned image suitable for automated trait extraction. The chain is segment-aware, respecting the discontinuities that the calibration process imposes on the slow-time axis.

Operating on that cleaned image, a particle filter with a backward smoother follows the ground through regions where canopy shadowing weakens or obscures it, reporting a per-chirp uncertainty that widens precisely where the return is ambiguous. A region-based canopy detector operating relative to the tracked ground extracts the canopy top and provides an explicit canopy presence decision, and a radar integrated energy measure derived from the region between ground and canopy top serves as a proxy for canopy structural density.

Finally, a calibrate-once, human-in-the-loop workflow was demonstrated and evaluated, in which an expert tunes parameters on a small number of representative rows and the same configuration is applied without further intervention to subsequent measurements. The radar integrated energy measure agreed with an independent optical canopy cover estimate at a Pearson correlation of $r = 0.837$ across the breeding trial, the single tuned parameter set transferred across measurement batches within a growth stage without re-tuning, and the workflow acquired along-row height profiles roughly eight times faster than hand measurement at equivalent spatial sampling density, yielding a plot-mean height estimate several times more precise than a single height-stick reading and concentrating the demand on skilled labor into a single tuning step.

5.2 System-Level Conclusions

Taken together, the results support the central premise of this work: that a proximal millimeter-wave radar is a viable and complementary instrument for high-throughput structural phenotyping. Three conclusions can be drawn at the system level.

First, the radar measures real, stable phenotypic traits rather than artifacts of a particular transect. The combination of repeatability and agreement with manual ground truth, observed in the July 29^{th} repeated-pass study, is the strongest single piece of evidence for this. This study showed the radar crop height measurement is both reproducible and was corroborated by manual measurements. The monotonic tracking of canopy height across the season and the agreement of the radar integrated energy measure with an optical canopy cover reference reinforce the same conclusion for the density measure.

Second, the radar's behavior is governed by two distinct measurement regimes divided by canopy closure. While the canopy is open, the radar can resolve both the ground and the canopy directly and the ground can be tracked continuously beneath the crop. However, once the canopy closes the within-row ground return is lost and the ground reference must be recovered from the bare inter-row alleys that the breeding-trial layout provides. The processing pipeline, utilizing the segment-aware ground tracker with its uncertainty estimate

allows a single instrument to operate usefully across both regimes. The recognition that breeding-trial inter-row alley gaps furnish a ground reference even under a closed canopy is what makes late-season operation possible.

Third, the value of the system is realized through augmentation rather than replacement. The workflow does not remove the expert; it concentrates the expert's judgment into a one-time tuning step and amortizes it across the field, multiplying the throughput of a single expert. The gain is not in elapsed time against conventional single-point scoring, which is already fast, but in the spatial resolution of the resulting measurement and in the class of labor the measurement requires. This frames both the source of the demonstrated throughput advantage and the origin of the principal limitation discussed next.

5.3 Limitations

The conclusions presented above are subject to several limitations, the most consequential of which was encountered directly in the course of the work.

The dominant limitation is the per growth stage re-tuning requirement. A single tuned parameter set transfers across measurements within a growth stage but it does not generalize across stages. Several detector parameters are tied to physical canopy properties that change through the season; the search-band height bound must rise as the canopy grows, the detection threshold depends on canopy return strength, the minimum component width depends on stand density and gap alley size, and the tracker's depth bias depends on how aggressively the canopy obscures the ground. The "calibrate-once" claim therefore holds best within a stage and a measurement campaign, but not across the season.

A second limitation concerns absolute height accuracy once the canopy has closed. With the within-row ground return lost, the absolute height estimate depends on the stability of the sensor height above the ground. For the hand-carried demonstration this was provided only by the operator. The agreement with manual measurement in the closed-canopy regime was sufficient for this deployment, but absolute accuracy in that regime is tied to platform stability rather than to a continuously observed ground.

Third, the reduced penetration inherent to operating at 120 GHz is a constraint. The short millimeter wavelength that gives the system its sensitivity to canopy architecture is the same property that prevents the beam from penetrating a dense canopy to the soil. This forces the reliance on inter-row gaps for the ground reference under closed canopies. This is a direct trade of penetration depth for structural sensitivity, and it bounds the conditions under which a direct sub-canopy ground return is available.

Finally, the field evaluation was conducted on a single crop (soybean) at a single site over a single season, with the radar integrated-energy density measure validated against an optical proxy rather than against a destructive biomass ground truth.

5.4 Future Work

The limitations above map directly onto the most promising directions for future work.

5.4.1 Automated Parameter Selection

Removing per-growth-stage re-tuning requirement would be a significant new contribution. As the detector parameters are physically interpretable and tied to canopy properties that themselves evolve predictably through the season, the tuning step is a strong candidate for automation. A parameter-selection method that adapts the search band, detection threshold, minimum component width, and depth bias from the data, either through a calibrated schedule keyed to growth stage, an optimization against the available ground-reference points, or a learned model would convert the present calibrate-once-per-stage workflow into a calibrate-once-per-season option. Establishing classical signal-processing baselines of the kind developed in Chapter 4 could provide the reference and training data for a learned model.

5.4.2 Platform Stabilization and Self-Referencing

The closed-canopy absolute-height limitation motivates a more stable or self-referencing platform. Mounting the sensor on a ground vehicle, gantry, or stabilized boom would fix the

sensor height compared to a hand-carried deployment left to the operator. An auxiliary range reference from a second downward sensor over a known target, or an inertial/height-tracking aid could supply a continuous ground reference independent of canopy penetration. Both options would address the regime in which the current system is limited.

5.4.3 Broader Validation

Extending the evaluation across additional crops, sites, growing seasons, and validating the radar integrated energy measure against destructive biomass and leaf-area measurements rather than an optical proxy would establish the generality of the density measure and the robustness of the height estimate beyond the single trial reported here. Repeated deployment across growing seasons would also test the stability of any automated parameter selection scheme.

5.4.4 Multi-frequency and Polarimetric Extension

The penetration/sensitivity trade suggests exploring complementary bands. Pairing the millimeter-wave sensor with a lower frequency channel that penetrates the canopy more readily could recover a sub-canopy ground return directly under closed canopies but would need to maintain a large absolute bandwidth (resolution). Polarimetric measurements could also add sensitivity to canopy orientation and structure beyond the scalar density measure used here. Both would extend the measurement capabilities of a proximal radar phenotyping system.

5.4.5 Integration into Operational Phenotyping

Finally, the CAN-FD interface and compact form factor were chosen with implement integration in mind. Deploying the sensor as one node among several on an autonomous or operator-driven field platform, and integrating the processing pipeline into a near-real-time workflow, would move the system from a demonstrated capability toward a routine tool in a breeding program's high-throughput phenotyping pipeline.

Bibliography

[1] José Luis Araus and Jill E. Cairns. Field high-throughput phenotyping: the new crop breeding frontier. *Trends in Plant Science*, 19(1):52–61, 2014. ISSN 1360-1385. doi: https://doi.org/10.1016/j.tplants.2013.09.008.

[2] Robert T. Furbank and Mark Tester. Phenomics – technologies to relieve the phenotyping bottleneck. *Trends in Plant Science*, 16(12):635–644, 2011. ISSN 1360-1385. doi: https://doi.org/10.1016/j.tplants.2011.09.005.

[3] Jason DeBruin, Thomas Aref, Sara Tirado Tolosa, Rebecca Hensley, Haley Underwood, Michael McGuire, Chinmay Soman, Grace Nystrom, Emaa Parkinson, Catherine Li, Stephen Patrick Moose, and Girish Chowdhary. Breaking the field phenotyping bottleneck in maize with autonomous robots. *Communications Biology*, Volume 8 - 2025, 2025. doi: 10.1038/s42003-025-07890-7.

[4] David Deery, Jose Jimenez-Berni, Hamlyn Jones, Xavier Sirault, and Robert Furbank. Proximal remote sensing buggies and potential applications for field-based phenotyping. *Agronomy*, 4(3):349–379, 2014. ISSN 2073-4395. doi: 10.3390/agronomy4030349. URL https://www.mdpi.com/2073-4395/4/3/349.

[5] Glenn Slade, Karen Anderson, Hugh A. Graham, and Andrew M. Cunliffe. Repeated drone photogrammetry surveys demonstrate that reconstructed canopy heights are sensitive to wind speed but relatively insensitive to illumination conditions. *International Journal of Remote Sensing*, 46(1):24–41, 2025. doi: 10.1080/01431161.2024.2377832. URL https://doi.org/10.1080/01431161.2024.2377832.

[6] Asheesh K. Singh, Arti Singh, Soumik Sarkar, Baskar Ganapathysubramanian, William Schapaugh, Fernando E. Miguez, Clayton N. Carley, Matthew E. Carroll, Mariana V.

Chiozza, Kevin O. Chiteri, Kevin G. Falk, Sarah E. Jones, Talukder Z. Jubery, Seyed V. Mirnezami, Koushik Nagasubramanian, Kyle A. Parmley, Ashlyn M. Rairdin, Johnathon M. Shook, Liza Van der Laan, Therin J. Young, and Jiaoping Zhang. *High-Throughput Phenotyping in Soybean*, pages 129–163. Springer International Publishing, Cham, 2021. ISBN 978-3-030-73734-4. doi: 10.1007/978-3-030-73734-4_7. URL https://doi.org/10.1007/978-3-030-73734-4_7.

[7] Lei Li, Qin Zhang, and Danfeng Huang. A review of imaging techniques for plant phenotyping. *Sensors*, 14(11):20078–20111, 2014. ISSN 1424-8220.

[8] Leonardo Volpato, Francisco Pinto, Lorena González-Pérez, Iyotirindranath Gilberto Thompson, Aluízio Borém, Matthew Reynolds, Bruno Gérard, Gemma Molero, and Francelino Augusto Rodrigues. High throughput field phenotyping for plant height using uav-based rgb imagery in wheat breeding lines: Feasibility and validation. *Frontiers in Plant Science*, Volume 12 - 2021, 2021. ISSN 1664-462X. doi: 10.3389/fpls.2021.591587.

[9] Jinru Xue and Baofeng Su. Significant remote sensing vegetation indices: A review of developments and applications. *Journal of Sensors*, 2017:1353691, 2017. doi: 10.1155/2017/1353691.

[10] Hamlyn G. Jones, Rachid Serraj, Brian R. Loveys, Lizhong Xiong, Ashley Wheaton, and Adam H. Price. Thermal infrared imaging of crop canopies for the remote diagnosis and quantification of plant responses to water stress in the field. *Functional Plant Biology*, 36(10-11):978–989, 2009. doi: 10.1071/FP09123.

[11] Yuxiang Wang, Zengling Yang, Gert Kootstra, and Haris Ahmad Khan. The impact of variable illumination on vegetation indices and evaluation of illumination correction methods on chlorophyll content estimation using UAV imagery. *Plant Methods*, 19(1):51, 2023. doi: 10.1186/s13007-023-01028-8.

[12] James D. C. Walter, James Edwards, Glenn McDonald, and Haydn Kuchel. Estimating

biomass and canopy height with lidar for field crop breeding. *Frontiers in Plant Science*, Volume 10 - 2019, 2019. ISSN 1664-462X. doi: 10.3389/fpls.2019.01145.

[13] Shangpeng Sun, Changying Li, Andrew H. Paterson, Yu Jiang, Rui Xu, Jon S. Robertson, John L. Snider, and Peng W. Chee. In-field high throughput phenotyping and cotton plant growth analysis using lidar. *Frontiers in Plant Science*, Volume 9 - 2018, 2018. ISSN 1664-462X. doi: 10.3389/fpls.2018.00016.

[14] Simon Madec, Fred Baret, Benoît de Solan, Samuel Thomas, Dan Dutartre, Stéphane Jezequel, Matthieu Hemmerlé, Gallian Colombeau, and Alexis Comar. High-throughput phenotyping of plant height: Comparing unmanned aerial vehicles and ground lidar estimates. *Frontiers in Plant Science*, Volume 8 - 2017, 2017. ISSN 1664-462X. doi: 10.3389/fpls.2017.02002. URL https://www.frontiersin.org/journals/plant-science/articles/10.3389/fpls.2017.02002.

[15] Amanda Veloso, Stéphane Mermoz, Alexandre Bouvet, Thuy Le Toan, Milena Planells, Jean-François Dejoux, and Eric Ceschia. Understanding the temporal behavior of crops using sentinel-1 and sentinel-2-like data for agricultural applications. *Remote Sensing of Environment*, 199:415–426, 2017. ISSN 0034-4257. doi: https://doi.org/10.1016/j.rse.2017.07.015.

[16] Heather McNairn, Jiali Shang, Xianfeng Jiao, and Catherine Champagne. The contribution of alos palsar multipolarization and polarimetric data to crop classification. *IEEE Transactions on Geoscience and Remote Sensing*, 47(12):3981–3992, 2009. doi: 10.1109/TGRS.2009.2026052.

[17] Dipankar Mandal, Vineet Kumar, Debanshu Ratha, Subhadip Dey, Avik Bhattacharya, Juan M. Lopez-Sanchez, Heather McNairn, and Yalamanchili S. Rao. Dual polarimetric radar vegetation index for crop growth monitoring using sentinel-1 sar data. *Remote Sensing of Environment*, 247:111954, 2020. ISSN 0034-4257. doi: https://doi.org/10.1016/j.rse.2020.111954.

[18] Xiuwei Liu, Xuejun Dong, and Daniel I. Leskovar. Ground penetrating radar for underground sensing in agriculture: a review. *International Agrophysics*, 30(4):533–543, 2016. ISSN 0236-8722. doi: 10.1515/intag-2016-0010. URL https://doi.org/10.1515/intag-2016-0010.

[19] Xiuwei Liu, Xuejun Dong, Qingwu Xue, Daniel I. Leskovar, John Jifon, John R. Butnor, and Thomas Marek. Ground penetrating radar (gpr) detects fine roots of agricultural crops in the field. *Plant and Soil*, 423(1-2):517–531, December 2017. ISSN 1573-5036. doi: 10.1007/s11104-017-3531-3.

[20] Daniel Gomez-Garcia, Fernando Rodriguez-Morales, Stephen Welch, and Carl Leuschen. High-throughput phenotyping of wheat canopy height using ultrawideband radar: First results. *IEEE Geoscience and Remote Sensing Letters*, 19:1–5, 2022. doi: 10.1109/LGRS.2020.3039179.

[21] Garrett Peterson. Miniature k-band radar for agricultural remote sensing. In *2021 IEEE 21st Annual Wireless and Microwave Technology Conference (WAMICON)*, pages 1–4, 2021. doi: 10.1109/WAMICON47156.2021.9443626.

[22] Ali Gharamohammadi, Anand G. Dabak, Zigang Yang, Amir Khajepour, and George Shaker. Volume-based occupancy detection for in-cabin applications by millimeter wave radar. *Remote Sensing*, 16(16), 2024. ISSN 2072-4292. doi: 10.3390/rs16163068. URL https://www.mdpi.com/2072-4292/16/16/3068.

[23] Mark A. Richards. *Fundamentals of Radar Signal Processing*. McGraw-Hill, New York, NY, 2nd edition, 2014.

[24] Fredric J. Harris. On the use of windows for harmonic analysis with the discrete fourier transform. *Proceedings of the IEEE*, 66(1):51–83, 1978. doi: 10.1109/PROC.1978.10837.

[25] John C. Curlander and Robert N. McDonough. *Synthetic Aperture Radar: Systems and Signal Processing*. John Wiley & Sons, New York, NY, 1991.

[26] Christophe Andrieu and Arnaud Doucet. Particle filtering for partially observed gaussian state-space models. *Journal of the Royal Statistical Society: Series B(Statistical Methodology)*, 64(4):827–836, 2002. doi: 10.1111/1467-9868.00363.

[27] M. Sanjeev Arulampalam, Simon Maskell, Neil Gordon, and Tim Clapp. A tutorial on particle filters for online nonlinear/non-gaussian bayesian tracking. *IEEE Transactions on Signal Processing*, 50(2):174–188, 2002. doi: 10.1109/78.978374.

[28] Mark Briers, Arnaud Doucet, and Simon Maskell. Smoothing algorithms for state-space models. *Annals of the Institute of Statistical Mathematics*, 62(1):61–89, 2010. doi: 10.1007/s10463-009-0236-2.

[29] Rafael C. Gonzalez and Richard E. Woods. *Digital Image Processing*. Pearson, New York, NY, 4th edition, 2018.

[30] Indie Semiconductor/Silicon Radar. Tra-120-002 datasheet, . URL https://downloads.ffo.indiesemi.com/datasheets/Datasheet_TRA_120_002_V0.81.pdf.

[31] Indie Semiconductor/Silicon Radar. Tra-120-002 lens information, . URL https://siliconradar.com/datasheets/Datasheet_Collimator_Lens_V1.1.pdf.

[32] Garrett Peterson and William Kuhn. Coherent calibration cancellation for stepped frequency radar systems. In *2018 IEEE Radar Conference (RadarConf18)*, pages 1594–1597, 2018. doi: 10.1109/RADAR.2018.8378805.

[33] Garrett Peterson. A stepped frequency continuous wave short range radar for agriculture measurements with coherent cancellation calibration. Master's thesis, Kansas State University, Manhattan, KS, 2020.

[34] Dane Thompson. Design of stepped frequency continuous wave radars using continuous coherent calibration. Master's thesis, Kansas State University, Manhattan, KS, 2024.

Appendix A

Radar Schematic

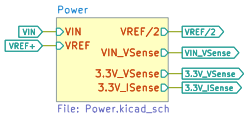

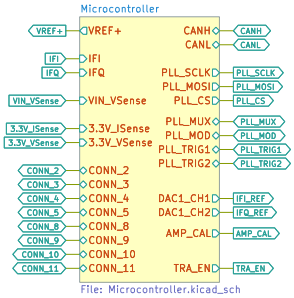

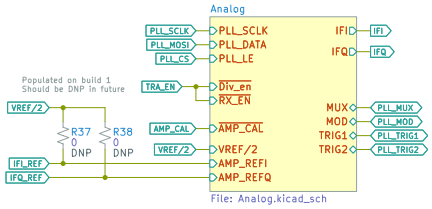

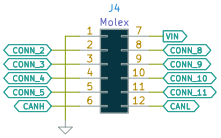

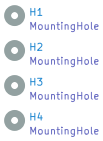

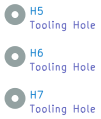

Garrett Peterson
Kansas State University

Sheet: /
File: GPM_Radar.kicad_sch

Title: AGR120 Top

Size: A4	Date: 2023-06-08	**Rev: 2**
KiCad E.D.A. 10.0.1		Id: 1/4

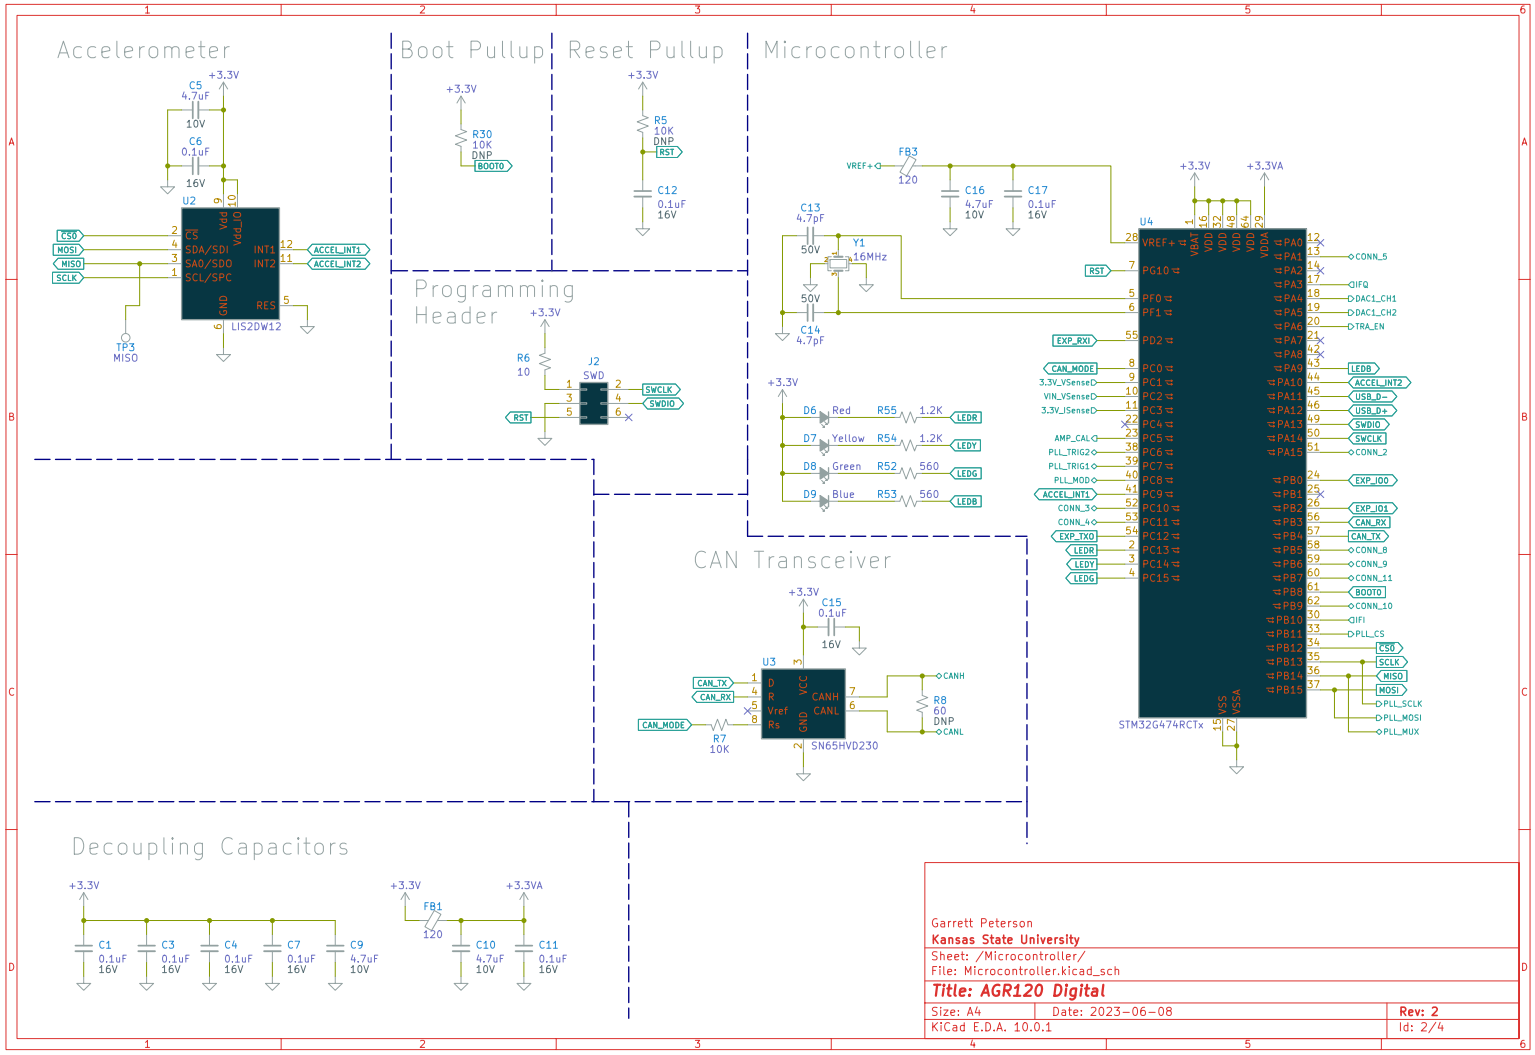
Accelerometer
Boot Pullup
Reset Pullup
Microcontroller
Programming Header
CAN Transceiver
Decoupling Capacitors
LIS2DW12
STM32G474RCTx
SN65HVD230
SWD
Garrett Peterson
Kansas State University
Sheet: /Microcontroller/
File: Microcontroller.kicad_sch
Title: AGR120 Digital
Size: A4
Date: 2023-06-08
Rev: 2
KiCad E.D.A. 10.0.1
Id: 2/4

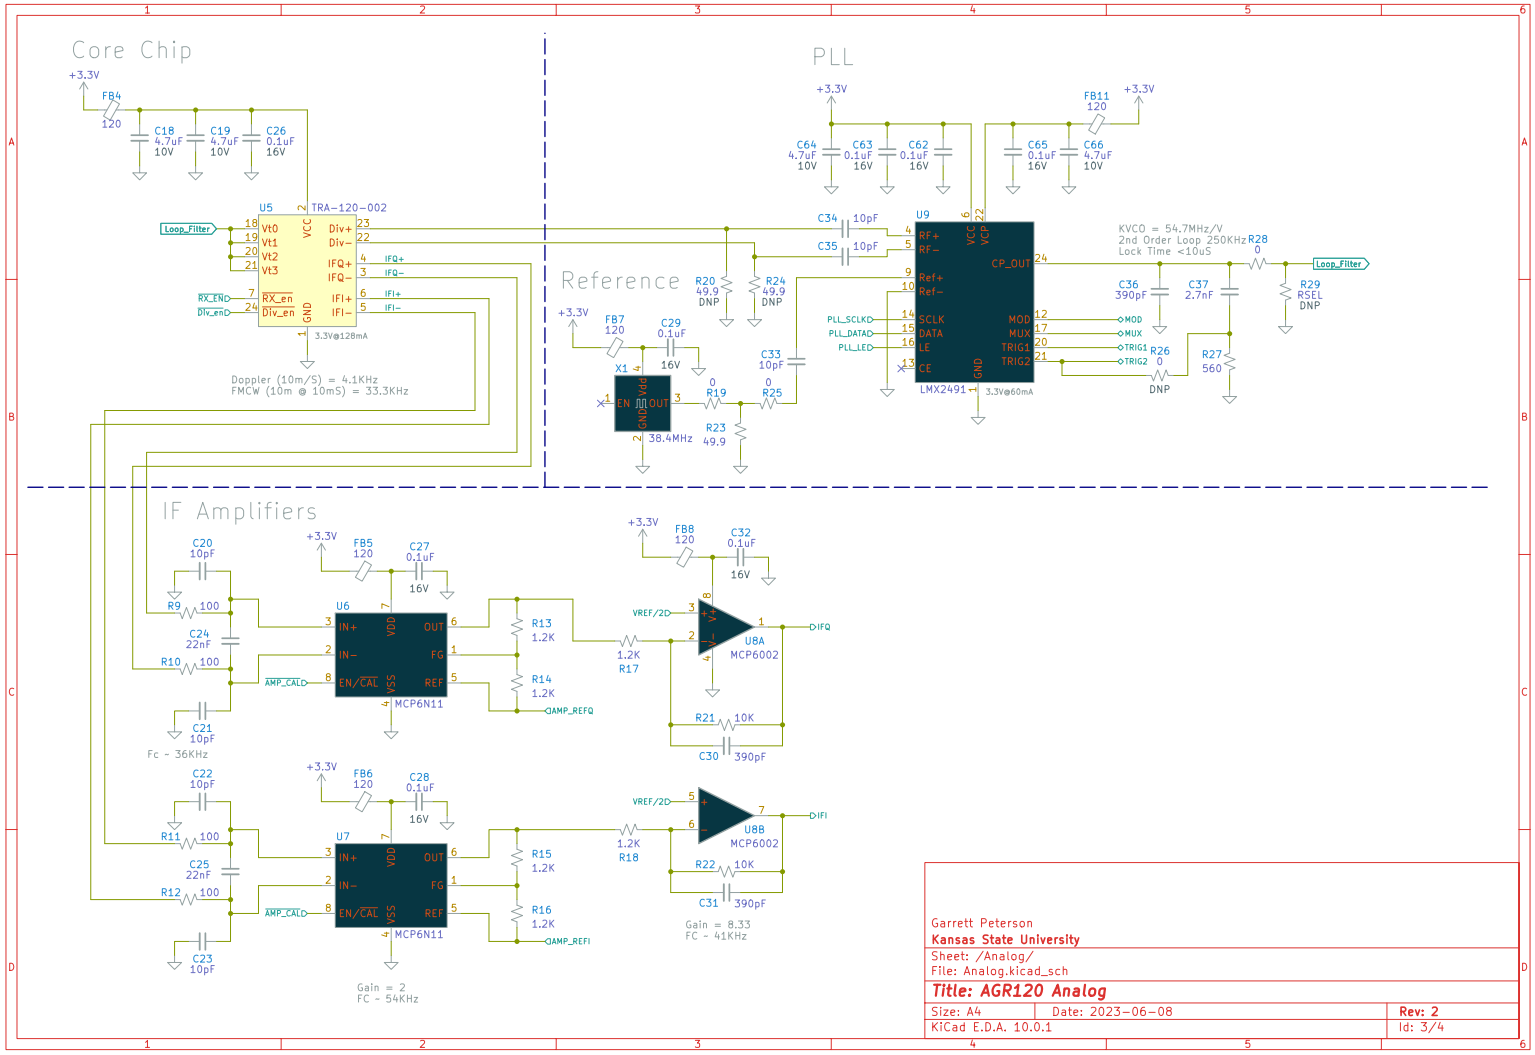

Garrett Peterson		
Kansas State University		
Sheet: /Analog/		
File: Analog.kicad_sch		
Title: AGR120 Analog		
Size: A4	Date: 2023-06-08	**Rev: 2**
KiCad E.D.A. 10.0.1		Id: 3/4

3.3V Switching Supply

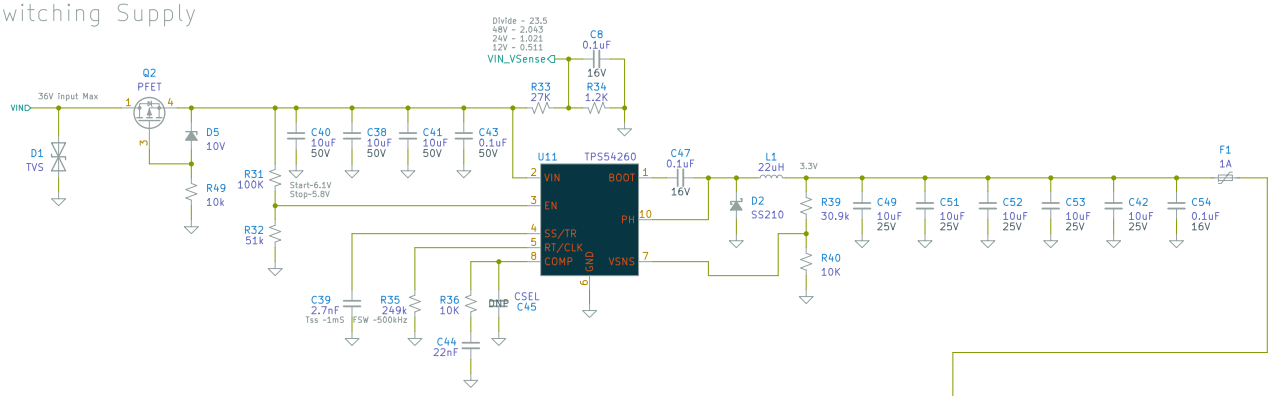

Reference Voltages

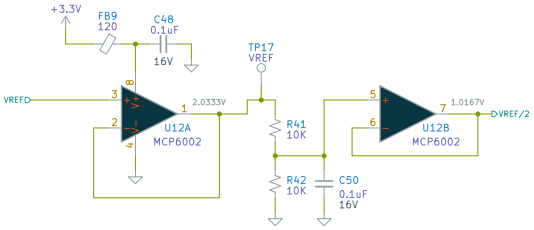

3.3V Current Sense

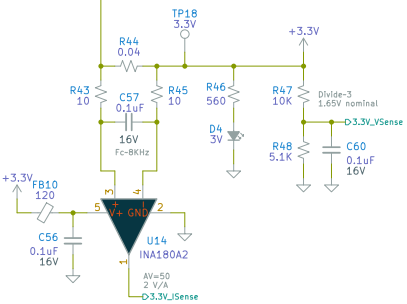

Garrett Peterson
Kansas State University

Sheet: /Power/
File: Power.kicad_sch

Title: AGR120 Power

Size: A4	Date: 2023-06-08	**Rev: 2**
KiCad E.D.A. 10.0.1		Id: 4/4