remote sensing

Article

Genotype-Aware Prediction of Soybean Seed Composition from Multimodal UAV Imagery of the Standing Crop

Vasit Sagan [1,2,3,*], Kristen Rhodes [1], Sourav Bhadra [1], Haireti Alifu [1], Aviskar Giri [1,2], Ashutosh Pawar [1,2], Bishal Roy [1,2], Supria Sarkar [1] and Felix Fritschi [4]

[1] Department of Earth, Environmental and Geospatial Science, Saint Louis University, St. Louis, MO 63108, USA
[2] Geospatial Center, Saint Louis University, St. Louis, MO 63108, USA
[3] Department of Computer Science, Saint Louis University, St. Louis, MO 63108, USA
[4] Division of Plant Science and Technology, University of Missouri, Columbia, MO 65211, USA
* Correspondence: vasit.sagan@slu.edu

Highlights

What are the main findings?

- An end-to-end deep learning model estimates eight soybean seed composition traits from multimodal UAV imagery across reproductive growth stages, with sucrose, oil, and simple carbohydrates predicted most accurately, while protein and fiber were least responsive to imagery alone.
- Conditioning the model on genotype and growth stage substantially improves protein and oil prediction, and multispectral imagery alone carries most of the predictive power, with thermal and LiDAR adding little for most traits.

What are the implications of the main findings?

- Seed quality can be assessed non-destructively in the standing crop before harvest, giving breeding programs and growers trait-level information early enough to guide variety selection and in-season management.
- Because multispectral imagery provides most of the accuracy, scalable field-level monitoring is achievable with simpler, lower-cost UAV sensor payloads rather than full multimodal systems.

Abstract

Geospatial artificial intelligence (GeoAI) integrates multimodal remote sensing with deep learning to model complex agricultural systems at scale. Within this framework, accurate and non-destructive prediction of seed composition from in-season standing crops is essential for breeding and precision agriculture. This study developed an end-to-end convolutional neural network (CNN) framework to estimate eight seed traits (protein, oil, sucrose, fiber, starch, ash, complex and simple carbohydrates) from UAV-based multisensor imagery and associated genotype and phenological metadata. A total of 372 soybean samples were collected over two growing seasons (2020–2021) from two fields in Missouri, with UAV flights capturing multispectral (MSI), thermal (THR), and LiDAR (LDR) data at four time points spanning vegetative to reproductive growth stages. CNN models were trained in single- and multi-date configurations, incorporating genotype (GEN) and days after sowing (DAS) as additional features. The highest accuracy was achieved for sucrose (R^2 = 0.80), followed by simple carbohydrate (R^2 = 0.70) and starch (R^2 = 0.55), with notable gains from GEN and DAS. Multi-date models incorporating earlier acquisitions often matched or outperformed later or all-date combinations. Among modalities, MSI provided

Check for updates

Academic Editor: Jochem Verrelst

Received: 15 June 2026
Revised: 28 August 2026
Accepted: 31 August 2026
Published: 11 September 2026

the most robust estimates, with limited added value from LDR or THR. Unlike feature-based pipelines prone to multicollinearity, this image-to-trait approach enables automated, scalable prediction of soybean seed composition for in-season, field-level assessment.

Keywords: remote sensing; protein; oil; voxelization for canopy architecture; computer vision

1. Introduction

Seed composition refers to the nutritional value of seeds, determined by the chemical makeup of seed kernels as described by the proportional concentrations of protein, oil, carbohydrates, isoflavones, and minerals [1]. Non-destructive measurement of seed composition from standing crops helps to optimize crop management and breeding through data-driven approaches that provide insights into crop growth and productivity, guiding proactive farm management. For soybean (Glycine max), seed composition traits are key determinants of soybean meal quality [2,3]. The key components that make up soybean seed composition are protein, oil, fatty acids, carbohydrates, and other minor constituents like isoflavones and minerals [4,5]. These nutritional components, especially the relative proportions of protein and oil, define the quality, end-use value, and marketability of the soybean seed. The nutritional composition of soybeans, particularly the high protein content, has contributed to the increasing demand for soybeans as a global commodity crop [4]. Soybean protein provides a high-quality plant-based protein source as an alternative to animal protein [6]. With demand for energy soaring globally, new bioengineering methods are rapidly being developed to design soybeans with higher oil content to gain competitive advantage over other vegetable oil crops for fueling and supplementing future industrial activity [7,8].

For decades, soybean breeding has heavily focused on increasing yield. Indeed, yields since 1925 have increased by approximately a third of a bushel per acre per year, and in the past 30 years by about half a bushel per acre per year [9]. However, seed protein concentration has steadily declined [10]. The negative correlation between yield and protein concentration directly threatens the global competitiveness of U.S. soybean meal. Investing in technology is critical to ensure continued production of high-quality U.S. dehulled soybean meal and oil, and to not only maintain but also enhance the competitive edge of U.S. soybean in the international market [11]. These opposing trends underscore the need for innovative, non-destructive approaches to monitor and optimize seed composition in real time.

The agriculture and food industries have adopted various non-destructive and cost-effective technologies as alternatives to conventional wet chemistry methods for analyzing seed composition. These new instruments utilize different spectral signals or properties retrieved from the seed itself to rapidly predict and quantify chemical components [12,13]. Specifically, near-infrared spectroscopy (NIRS) is a powerful tool for assessing both major seed constituents like protein, oil, and fatty acids as well as micronutrients like minerals in seeds. Laboratory NIRS is a rapid and widely used method for high-throughput seed-composition prediction, although its accuracy depends on the calibration model, instrument, and target sample population [14]. Moreover, because NIRS is performed on harvested seed samples, it does not provide in-season information on the canopy conditions associated with composition variation across genotypes and environments [15]. The limitation of these techniques presents a major hindrance for the U.S. soybean value chain, particularly for breeding programs that have to process a very large number of samples, as well as for soy processors. Novel, non-destructive methods and technologies, including approaches that allow prediction of seed composition in the standing crop, are needed to accelerate variety development and to assist seed processors in ensuring delivery of high-quality soybean products [16]. It is crucial to assess soybean seed traits without relying solely on traditional lab-based procedures, as they are time-intensive and costly, making large-scale analysis difficult. It is also important for farmers to have in-season predictions of seed composition for better agronomic decisions and to optimize soybean quality.

Recent advances in unmanned aerial systems (UAS) and multisensory remote sensing, including multispectral imagery (MSI), thermal imagery (THR), and LiDAR data (LDR), provide new avenues for rapid, non-destructive assessment of crop condition and pre-harvest prediction of seed composition across agricultural fields [17]. These modalities characterize complementary canopy properties that may covary with subsequent seed composition. MSI captures spectral variation associated with canopy pigments, vigor, and senescence; THR characterizes canopy-temperature variation associated with plant water and heat responses; and LDR represents three-dimensional canopy structure and development. However, many of the current remote sensing approaches rely on manually constructed indices or textural metrics, which can be time-consuming to engineer and prone to multicollinearity [18], prompting a shift toward image-based deep learning that automatically learns discriminative features from raw imagery [19].

Non-destructive prediction of seed composition using remote sensing data employs two main approaches: spectral measurement of seeds in the laboratory and imaging of the crops in the field with remote sensing technologies over the course of the growing season. The latter is of particular interest because crop management practices (planting date, irrigation, fertilizer applications, crop rotation), abiotic conditions (drought, heat, and soil conditions) and biotic factors such as genotype and diseases play a vital role in seed development and composition [20]. Therefore, the ability to predict seed composition from in-season crops is expected to provide opportunities to improve yield and meal quality through proactive farm management. Within the last several years, non-destructive imaging techniques, mainly spectroscopy and hyperspectral imaging, have gained popularity among researchers as robust analytical techniques for the evaluation of the protein makeup and mineral nutrient uptake of soybean [21]. Deep learning, especially convolutional neural networks (CNNs) that utilize imagery as the direct input for machine learning-based modeling, offers strong techniques to extract predictive information from remotely sensed images [22]. These CNN-based methods align with the emerging concept of geospatial artificial intelligence (GeoAI), wherein advanced machine learning integrates with geospatial data to enable scalable field-level predictions [23]. Pre-harvest determination of seed composition allows for selection of soybean varieties with desirable nutritional (meal quality) profiles, as well as indicating the quality and potential applications for the soybeans post-harvest [24]. Recent studies demonstrated that spectral indices and features derived from high-resolution hyperspectral imagery and LiDAR data collected with UAVs were effective in predicting soybean seed protein and oil content prior to harvest [25]. Additionally, with its high spatial resolution and daily temporal resolution, PlanetScope satellite data proved effective for feature based prediction of seed composition at plot-scale [16]. Successes of these studies are attributed to the fact that remote sensing technology can capture the key phenological traits that determine seed composition. Taking advantage of spectral, spatial, and temporal features that are responsive to plant phenology and seed composition, UAV and PlanetScope data showed encouraging results for seed composition modeling with machine learning. Yet such feature-based pipelines often require considerable manual effort and can introduce multicollinearity, emphasizing the value of end-to-end CNNs that learn relevant features directly from imagery [19]. However, no studies to date have investigated end-to-end CNN approaches for seed composition prediction, which can be designed to automatically extract features without manual selection, advancing the scalability of remote sensing technologies.

Further research is needed to optimize remote sensing-based approaches for soybean seed composition modeling. Key knowledge gaps remain regarding how to best leverage different remote sensing modalities across critical crop growth stages with fully automated CNN approaches, how to characterize canopy architecture with computer vision in a meaningful way that quantifies plant light-use efficiency, and how fully automated CNN approaches can streamline seed composition modeling compared to conventional manual feature extraction. In addition, genotypes, planting dates, water status, and soil fertility may provide useful supplementary information, which requires systematic investigation.

In this study, UAV-based prediction refers to supervised, model-based inference of post-harvest seed-composition reference values from pre-harvest canopy observations and associated genotype and phenological metadata. The UAV sensors do not directly measure seed biochemistry. Instead, they observe spectral, thermal, and structural canopy characteristics that may covary with later seed composition through genotype, development, and environmental response. We therefore describe the task as image-to-trait prediction or pre-harvest prediction.

The overarching goal of this study was to evaluate how multisensor UAV data, growth stages (represented by days after sowing), and soybean genetic information can be integrated into an end-to-end CNN framework to predict seed composition. The specific research questions addressed were:

(1) How do different remote sensing modalities (i.e., multispectral, thermal, LiDAR) contribute to predicting soybean seed composition at different growth stages?
(2) What is the optimal combination and timing of remote sensing data collection for accurately predicting soybean seed composition prior to harvest?
(3) How can remote sensing-based seed composition models be improved with additional field data like genotype and planting date?
(4) How can LiDAR-derived point clouds be used to represent canopy architecture that reflects the formation of seed composition traits?

2. Study Area and Data

2.1. Study Area and Experimental Design

This study was conducted during the 2020 and 2021 growing seasons at two agricultural fields located at the Bradford Research Center (BRC) near Columbia, Missouri, USA (Figure 1). Missouri is characterized by a humid continental climate with distinct seasonal variations. A weather station at BRC recorded daily meteorological conditions throughout both seasons. In 2020, the highest monthly precipitation occurred in June (156.5 mm), and the lowest in November (8.4 mm). Average monthly temperatures peaked at 24.8 °C in July and dropped to 8.4 °C in November. In 2021, June recorded the highest monthly precipitation (296.6 mm), with the lowest again in November (6.4 mm); average monthly temperatures ranged from 24.0 °C in August to 6.4 °C in November.

The experimental design aimed at assessing the influence of genetic and environmental variability on soybean seed composition, which is critical for improving the generalizability of prediction models. Both experimental fields were rainfed and managed using standard herbicide applications during the pre- and post-emergence phases for weed control. Apart from the designated nitrogen-treatment plots in L2, no fertilizer was applied.

The H1G field (60 m × 83 m) was planted with soybeans on 6 June 2020 and 13 May 2021, using 15- and 30-inch row spacings at a seeding rate of 140,000 plants per acre. Two soybean genotypes were planted in each season, assigned to four replications, resulting in 90 viable plots in 2020 and 91 in 2021 (Table 1). Each plot measured approximately

20 ft × 20 ft, separated by 5-foot alleys to minimize edge effects (Figure 1). Notably, this field incorporated a rooting depth restriction experiment to simulate water-limited conditions. Six-meter-wide zones were excavated across the width of the field to depths of 0.30, 0.45, 0.60, 0.75, and 0.90 m. These zones were lined with impermeable plastic to restrict root penetration and subsequently backfilled with Mexico silt loam to match the original soil profile. Control zones with unrestricted rooting depth were maintained at the north and south ends of the field, as well as between treatment bands. During both years, two soybean cultivars and two row spacings were used in a fully replicated design, with four replications per treatment combination.

Figure 1. Overview of the study area and experimental setup. (**a**) The Bradford Research Center (BRC), University of Missouri, showing the two experimental fields: H1G and L2. (**b**) Location of Missouri within the United States, with red star indicating the location of the BRC. (**c**) Zoomed-in layout of the H1G field showing rooting depth treatment zones, where yellow lines mark different constrained soil depths. (**d**) Representative UAV image of the L2 soybean field, where yellow and red boxes indicate plots with 15″ and 30″ row spacings, respectively.

The L2 field (61 m × 220 m) was planted in four staggered rounds from 7 May to 16 June 2021, also in Mexico silt loam soils. The plot layout matched that of the H1G field, with dimensions of approximately 20 ft × 20 ft, 15- and 30-inch row spacings, and 5-foot alleys (Figure 2). The field included seven soybean cultivars, randomly distributed across 191 sample plots, with twelve replications per cultivar (Table 1). Additionally, 24 plots were assigned to a nitrogen treatment experiment to examine the impact of nutrient enrichment on seed composition. These plots used 30-inch spacing and included the cultivars 'P39A82S', P31A29L', and 'P39A45X'. Unlike H1G, the L2 field did not include rooting depth treatments. Weed management was performed using an early-season glufosinate application, supplemented by manual removal as needed.

Table 1. Seed sample distribution with respect to study sites and seed varieties.

BRC Experimental Field	Number of Seed Samples	Seed Varieties
H1G 2020	90	P29A85L, P44A37L
H1G 2021	91	P28A97L, P44A37L
L2 2021	191	P28A97L, P31A29L, P36A43L, P39A45X, P39A58X, P39A82S, P44A37L

Figure 2. Soybean plots from the H1G field: Soybean plants captured on 21 July 2021, during the V4–V5 vegetative stage, showing both 15-inch and 30-inch row spacings.

The rooting-depth restrictions in H1G and the nitrogen treatments in the 24-plot subset of L2 were components of the field experiments and were not used as input features. These treatments added environmental and phenotypic heterogeneity to the pooled dataset.

2.2. Data Collection

2.2.1. Seed Data Collection

At physiological maturity, depending on the planting date, soybean plots from both experimental fields were harvested using a small-plot research combine to collect seed samples and estimate grain yield. The harvested area per plot measured approximately 15 feet in length and 5 feet in width. Based on the row spacing (30″ or 15″), either the center two or center four rows within this area were harvested. Despite some sample loss due to equipment malfunction and deer damage, a total of 372 plots were successfully harvested over the two growing seasons (Table 1). From each harvested plot, a 60 mL seed subsample was retained to noninvasively assess seed composition traits, including protein, oil, carbohydrates, amino acids, and minerals.

2.2.2. NIR-Based Seed Composition Prediction

Seed composition traits were determined using a laboratory-based near-infrared (NIR) spectrometer (DA 7250 NIR Analyzer, PerkinElmer 201, Waltham, MA, USA) on 60 mL samples from each harvested plot. NIR spectroscopy enables rapid, non-destructive, and simultaneous prediction of multiple biochemical traits from intact seeds by detecting molecular vibrations in the 780–2500 nm range [14,26]. To quantify individual components, we used calibration curves previously developed at the University of Missouri, based on chemometric analysis of reference samples [27,28]. A Partial Least Squares (PLS) regression model was applied to relate the spectral transmittance and reflectance to known concentrations of seed constituents [27]. To reduce sample variability effects, the analyzer rotated each sample tray under a consistent light source and computed the spectral average per sample. Larger seed components (e.g., protein, oil) are generally estimated more accurately than compounds present at lower concentrations (<10% dry weight) [13].

From the forty-four compounds assessed by the model, we selected a subset including protein, oil, carbohydrates, sucrose (energy), and ash (minerals) for further analysis. To improve interpretability, carbohydrates were grouped into complex (e.g., fiber, starch) and simple (e.g., glucose, fructose) carbohydrates, based on physicochemical distinctions relevant to NIR modeling [29,30]. Specifically, simple carbohydrates consist of soluble sugars, while complex carbohydrates include structural polysaccharides. All values are expressed on a dry-weight basis. The DA 7250 PLS calibrations have shown strong cross-validation accuracy for these traits [27,28].

2.2.3. UAV Data Collection

Multiple field campaigns were conducted during the 2020 and 2021 growing seasons at the BRC to acquire multispectral, RGB, thermal infrared, and LiDAR datasets at key growth stages of soybean development (Table 2). All sensors were flown using a fleet of DJI Matrice 600 Pro (M600 Pro) hexacopters (DJI Technology Co. Ltd., Shenzhen, China), each equipped with a DJI 3A Pro Flight Controller, an inertial measurement unit (IMU), and real-time kinematic (RTK) GNSS receivers, enabling centimeter-level positioning accuracy (2–3 cm).

Table 2. UAV flight campaigns at the BRC over the growing seasons of 2020 and 2021.

Flight Date	Data Fusion	Field Site	Days After Sowing (DAS)	Growth Stage
24-Jul-20	MSI, Thermal, LiDAR	H1G	47	R5
19-Aug-20	MSI, Thermal, LiDAR	H1G	73	R6
4-Sep-20	MSI, Thermal, LiDAR	H1G	89	R7
21-Jul-21	MSI, Thermal, LiDAR	H1G & L2	34–74	V2–V5
10-Aug-21	MSI, Thermal, LiDAR	H1G & L2	56–96	R2–R5
24-Aug-21	MSI, Thermal, LiDAR	H1G & L2	70–110	R4–R6
9-Sep-21	MSI, Thermal, LiDAR	H1G & L2	84–124	R6–R7

Campaigns were scheduled to sample the temporal progression of canopy conditions from vegetative or early reproductive development through seed filling and senescence. Because sowing dates and crop development differed among fields and years, F1–F4 denote relative acquisition order rather than identical phenological stages across all field-year combinations; Table 2 reports the exact dates, days after sowing, and observed growth-stage ranges.

All flight plans were pre-programmed and automated, maintaining a consistent altitude of 50 m, flight speed of 3–5 m/s, and 40–70% forward and side overlap between adjacent flight lines to support orthomosaicking and 3D surface modeling. Prior to the first campaign, high-contrast ground control points (GCPs) were permanently installed across both the H1G and L2 fields to ensure accurate georeferencing and geometric correction. In addition, black and white reflectance reference panels were distributed along plot alleys before each flight to support radiometric calibration of the multispectral imagery and image registration; these panels were not used as thermal calibration targets.

Sensor payload configurations varied depending on the sensing modality. For multispectral, thermal, and RGB imagery, a MicaSense Altum-PT (MicaSense, Inc., Seattle, WA, USA) multispectral sensor, an ICI 8640 P thermal camera (Infrared Cameras Inc., Beaumont, TX, USA), and a Sony RX10 RGB camera (Sony Corporation, Tokyo, Japan) were co-mounted on a single M600 Pro platform, with each sensor corresponding to a specific imaging modality. The Altum-PT sensor captures data in five spectral bands (blue, green, red, red-edge, and near-infrared (NIR)) tailored for optimally capturing the spectral properties of vegetation. The ICI 8640 P sensor provides thermal infrared data with an accuracy of ±1 °C and a thermal sensitivity (NETD) of 0.02 °C. These sensors were mounted using a Gremsy gimbal (Gremsy, Ho Chi Minh City, Vietnam) and a custom 3D-printed ABS plastic tray for stabilization and modular payload integration.

Additionally, LiDAR data were collected using the Phoenix Scout-32 UAV LiDAR system (Phoenix LiDAR Systems, Los Angeles, CA, USA), which was mounted on a dedicated M600 Pro platform. This system integrates a Velodyne HDL-32 dual-return (Velodyne LiDAR, Inc., San Jose, CA, USA) LiDAR sensor capable of a 700,000 points-per-second scan rate and an accuracy of ±0.02 m, alongside a Sony A7R II RGB camera (Sony Corporation, Tokyo, Japan) used to generate colorized point clouds. All flight missions were conducted under consistent illumination conditions, and sensor deployment followed standardized configurations to ensure temporal consistency and cross-seasonal comparability.

3. Methodology

Figure 3 presents a graphical representation of the overall workflow employed in this study.

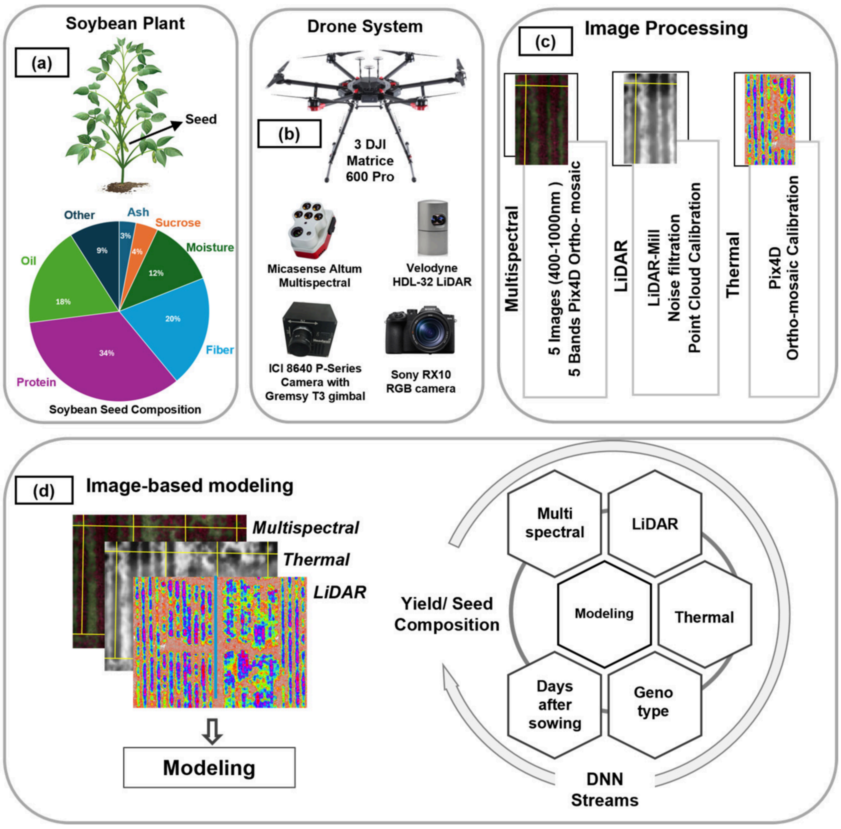

Figure 3. Workflow of the study outlining the data collection, preprocessing, and modeling stages. (**a**) Illustration of seed composition traits targeted for prediction. (**b**) UAV platform (DJI M600 Pro) and onboard sensors including multispectral, thermal, LiDAR, and RGB cameras. (**c**) Image preprocessing steps applied to multispectral, thermal, and LiDAR datasets. (**d**) Image-based modeling framework integrating multi-modal data and the overall modeling pipeline design.

3.1. UAV Data Processing

This section describes the preprocessing workflows for multispectral imagery, LiDAR point clouds, and thermal imagery.

3.1.1. Multispectral Imagery

Multispectral imagery was collected using a UAV-mounted MicaSense Altum-PT multispectral camera. Radiometric calibration was performed using Pix4Dmapper v4.7.5 (Pix4D SA, Lausanne, Switzerland), with calibration panel images captured in the field under flight-matched conditions. Additionally, the Pix4Dmapper workflow generated radiometrically corrected and mosaicked images for each of the five spectral bands. These individual bands were mosaicked into composite image files for each flight using ArcGIS Pro 3.0 (Esri, Redlands, CA, USA). Finally, orthorectification was performed using survey-grade, high-accuracy GCPs distributed across the test sites at BRC.

3.1.2. LiDAR Point Clouds

LiDAR point clouds were processed using the cloud-based LiDARMill v2 platform (Phoenix LiDAR Systems, Los Angeles, CA, USA). This platform performs multiple processing steps, including flight trajectory correction, noise filtering, point classification, and integration with a GNSS base station to produce point clouds in LAS format.

The workflow begins with trajectory smoothing, where onboard IMU data were fused with GNSS ground reference data to generate a smoothed trajectory file (SBET). The software further automatically detected and omitted turns and calibration maneuvers to reduce processing time and limit data to flight lines. Additionally, geometric observations across overlapping flight lines were used in the LiDAR Snap step to enhance alignment parameters and minimize offsets. Once point locations were refined, the software classified returns as either ground or non-ground points. Finally, a statistical outlier removal (SOR) algorithm was applied to eliminate outliers from the dataset, ensuring clean point clouds for further analysis.

3.1.3. Thermal Imagery

Thermal imagery was collected using an ICI 8640 P camera mounted on the UAV. The at-sensor radiometric temperature data were converted into surface temperature values (°C) using proprietary IR Flash software (ICI—Infrared Cameras Inc., Beaumont, TX, USA; https://buy.infraredcameras.com/products/ir-flash-pro-annual-subscription; accessed on 3 September 2026). This conversion step accounted for surface emissivity and atmospheric attenuation effects on the thermal imagery. Finally, orthorectification and mosaicking of the 32-bit TIFF-format thermal images were conducted using Pix4Dmapper software.

3.2. Modeling

3.2.1. Imagery Data Post-Processing

The data post-processing steps were consistent across imagery types (i.e., MSI and THR) and are summarized below.

A pipeline developed in Python v3.9.7 was used to handle mosaicked images for each field, in conjunction with corresponding plot boundary shapefiles created using ArcGIS Pro 3.0 (Esri, Redlands, CA, USA). These boundaries covered only the central harvested rows, avoiding the alleys to minimize edge effects and mixed pixels. The pipeline clipped each image based on the plot boundary shapes and automatically saved the extracted plot-level images. Since the ground sampling distance (GSD) varied for different image types, the resulting clipped images had varying pixel dimensions in width and height. To standardize the dataset, all images were uniformly resampled to 64×64 pixels. The nearest neighbor method was used for resampling multispectral imagery to preserve discrete spectral values. In contrast, bilinear interpolation was applied to thermal images, as it is well-suited for resampling continuous data. Various pixel sizes were tested during resampling, and 64×64 pixels was found to offer a suitable trade-off between preserving plot-level detail and minimizing processing time.

3.2.2. Voxel-Based Approach to Quantify Canopy Architecture

The LiDAR point cloud is usually represented by points consisting of 3-dimensional coordinates along with associated attribute information (e.g., RGB values, intensity) [31]. However, deep learning models require input data in structured n-dimensional matrix formats. Therefore, we converted the point cloud of a given plot into a voxel-based representation. In this representation, the LiDAR point cloud volume is divided into a collection of 3D regular cubes (voxels), each assigned an attribute value (e.g., RGB or intensity) [32]. Voxelization enables the creation of a 2D matrix along the z-axis (height dimension) and the generation of a structured 3D voxel grid [33].

The schematic process of voxelization is represented in Figure 4. First, the plot-level point cloud was clipped using the corresponding plot boundary shapefile via the Whitebox Geospatial Analysis Python API (whitebox v2.0.3). The laspy package (v2.0.3) was used to read the clipped point cloud and convert it into a 3D array. A statistical outlier removal (SOR) algorithm was applied using Open3D v0.13.0 to eliminate residual noise and outliers. Voxelization was also performed using Open3D, with LiDAR intensity values stored as the attribute of each voxel. The unit size of each voxel was set to 0.01 m^3, which we found to be a good balance between capturing the structural variability and maintaining manageable data volume. Finally, the voxels were converted into a 4D array where the first 3 dimensions represented x, y, and z, and the fourth dimension stored the LiDAR intensity values. The resulting arrays were resampled to 64 by 64 units in the horizontal plane, ensuring consistent dimensions across all samples. We did not change the limit of the z-axis as we wanted to capture the height variation in all the samples. However, the first and last horizontal layers were omitted as the first layer might contain point cloud from emergent weeds, while the last layer might still capture residual outliers despite the earlier filtering [34].

3.2.3. CNN for Multimodal Data Fusion

CNNs have been widely recognized for their ability to learn complex patterns from images, making them highly effective for both classification and regression tasks. Numerous studies have demonstrated the superiority of CNNs in computer vision applications, particularly in remote sensing and feature extraction tasks [35,36]. CNNs consist of stacked convolutional and pooling layers that automatically learn hierarchical spatial, spectral, and textural features at multiple scales. By leveraging these hierarchical feature representations, CNNs eliminate the need for hand-crafted feature engineering, enhancing model generalization and reproducibility across diverse datasets.

Let $x^{(i)}$ and $y^{(j)}$ denote the input and output feature maps of ith input and jth output in a convolution layer. The convolution operation is defined as:

$$y^{(i)} = f\left(b^{(j)} + \sum_i k^{(i)(j)} * x^{(i)}\right) \quad (1)$$

where $k^{(i)(j)}$ represents the convolutional kernel, $b^{(j)}$ is the bias, and $*$ denotes the convolutional operator. The activation function f, introduces non-linearity into the network, enhancing its ability to capture complex patterns in the data. Among available activation functions (e.g., Sigmoid, Tanh), we utilized ReLU due to its proven effectiveness in remote sensing applications [37]. ReLU applies a linear transformation for all positive values while setting negative values to zero. Therefore, Equation (1) can be expressed as:

$$y^{(i)} = max\left(0, b^{(j)} + \sum_i k^{(i)(j)} * x^{(i)}\right) \quad (2)$$

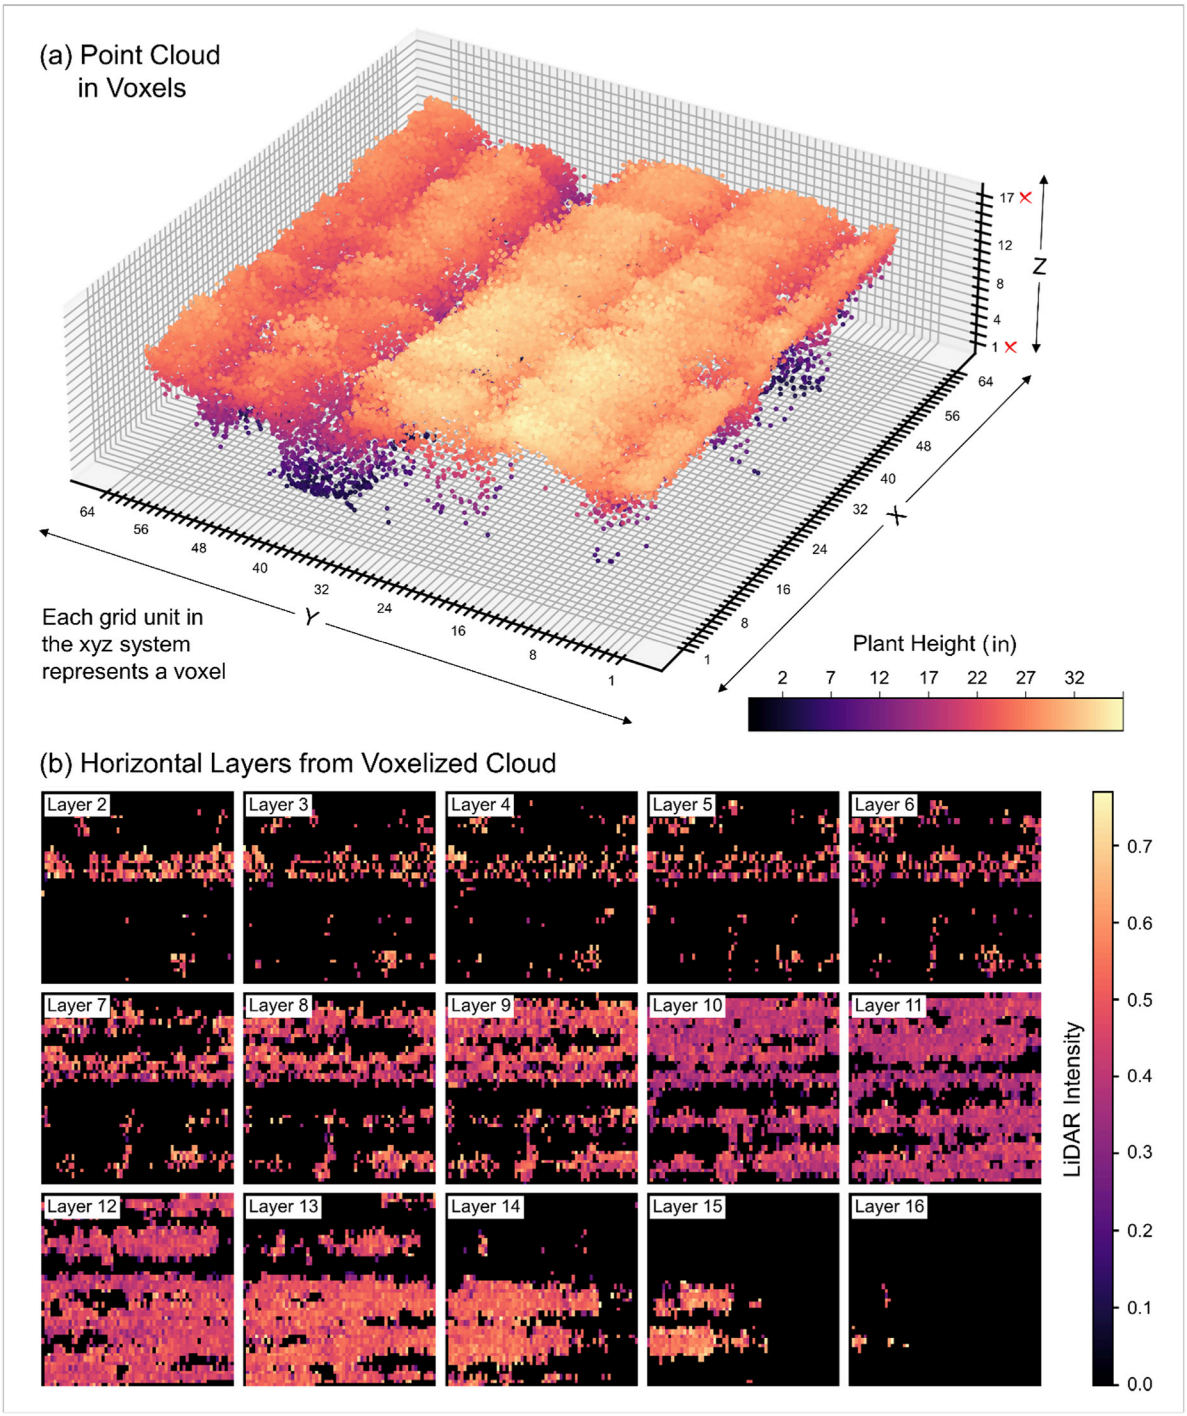

Figure 4. Schematic of the voxelization approach used in this study. (**a**) The 3D space is divided into cubic volume units, each assigned the average LiDAR intensity from corresponding points. (**b**) The voxels along the height (z-axis) are structured into 17 layers for CNN input, with the first and last layers removed to exclude soil-bed weeds and upper canopy noise.

Our dataset included three sensor modalities: spectral (MSI), structural (LDR), and thermal (THR), which were fused at the feature level in the CNN architecture (Figure 5). Feature-level fusion combines extracted features into a single vector for prediction [38], allowing the model to learn from each modality without post-processing. However, since input shapes varied by modality, we applied global average pooling (GAP) to standardize feature map sizes before concatenation. This operation averaged each feature map, making it independent of input dimensions. A final fully connected layer was used to predict seed composition traits.

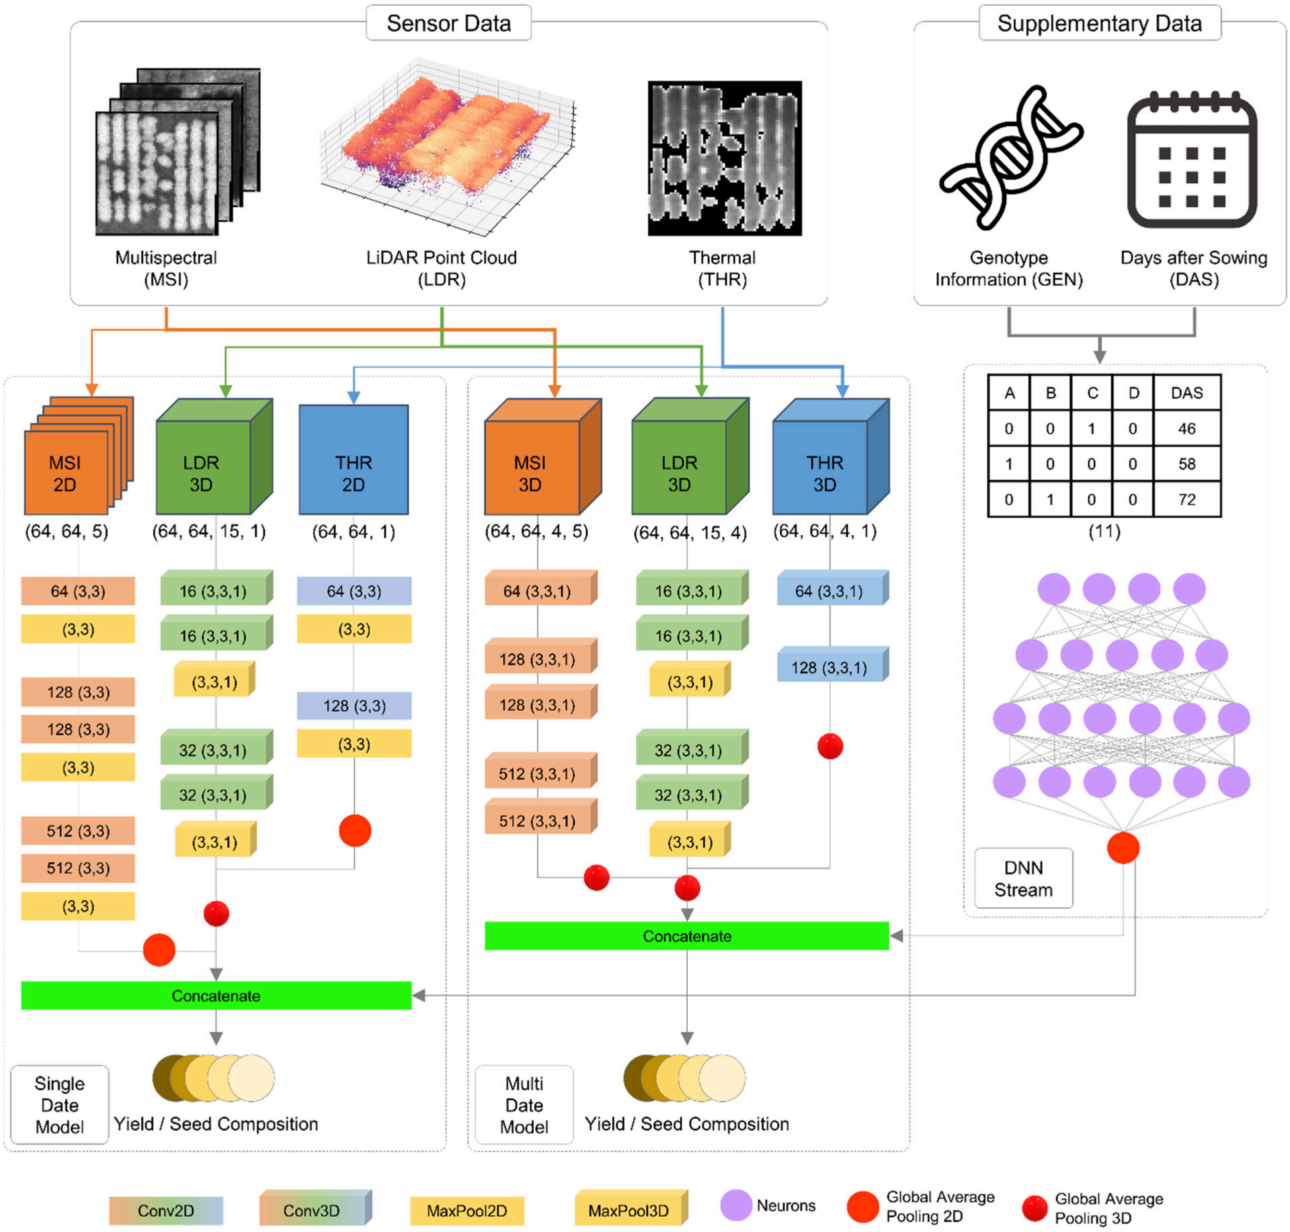

Figure 5. Overview of the image-to-trait modeling framework. Two model configurations were implemented: single-date models, which used UAV imagery from a single acquisition date, and multi-date models, which incorporated temporally stacked imagery across multiple dates. Multispectral, thermal, and LiDAR data were fused at the feature level, and both genotype (GEN) and days after sowing (DAS) were integrated into the model using a multi-layer perceptron. The convolutional blocks indicate the number of filters and kernel sizes in parentheses.

We implemented two types of fusion networks: single-date models and multi-date models. In the single-date configuration, input images had dimensions $W \times H \times C$, where W and H dictate image width and height, respectively, and C represents the number of spectral channels or bands. These models used 2D convolution and pooling layers to extract spatial and spectral features, enabling the models to learn from multiple sensor modalities while maintaining consistent data representation.

In contrast, multi-date models incorporated images from multiple dates, expanding the input to a four-dimensional tensor of size $W \times H \times T \times C$, where T represents the number of image acquisition dates. This configuration employed 3D convolution and pooling layers to simultaneously learn spatial ($W \times H$), spectral (C), and temporal (T) relationships within the data. This approach was applied to both the MSI and THR branches, where the temporal dimension captured multi-date imagery collected across different growth stages.

The LiDAR branch followed a slightly different design. For single-date models, voxelized LiDAR data were processed as a 4D array of shape $W \times H \times D \times C$, where D is the height from the ground and C is the number of channels (which is 1 and contains the LiDAR intensity values). A 3D convolutional network was used to extract structural features directly from the voxel grid. For the multi-date configuration, temporal stacking was introduced by converting the channel dimension C into a temporal dimension T, resulting in a 4D input of shape $W \times H \times D \times T$. Ideally, the input shape would be $W \times H \times D \times T \times C$, which can only be processed by a 4D convolution. Theoretically, a 4D convolution is possible [38], but we considered the overwhelming model complexity and chose not to use a 4D convolution.

The mathematical formulation of 2D convolution can be further extended from Equation (1) by considering p and q as the indices of the image matrix, and a and b as the corresponding indices of kernel k:

$$y[p,q] = \sum_{a=-\infty}^{\infty} \sum_{b=-\infty}^{\infty} k[a,b] * x[p-a, q-b] \tag{3}$$

The 3D convolution is similar to the 2D convolution, but instead of having a 2-dimensional kernel, there is a 3-dimensional kernel involved in the operation. Therefore, Equation (3) can be rewritten as:

$$y[p,q,r] = \sum_{a=-\infty}^{\infty} \sum_{b=-\infty}^{\infty} \sum_{c=-\infty}^{\infty} k[a,b,c] * x[p-a, q-b, r-c] \tag{4}$$

where the image matrix and kernel both have 3 indices, i.e., p, q, r, and a, b, c, respectively.

Figure 5 summarizes the modality-specific network branches. For the single-date models, MSI and THR inputs were processed using 2D convolutional layers, whereas voxelized LDR inputs were processed using 3D convolutional layers. For the multi-date models, temporally stacked MSI, THR, and LDR inputs were processed using 3D convolutional layers. Each branch used ReLU activations and pooling operations, followed by global average pooling to generate a fixed-length feature vector for each sensor modality. The filter counts and kernel dimensions used in each branch are provided in Figure 5.

Lastly, we integrated two ancillary non-image inputs: genotype (GEN) and days after sowing (DAS). GEN was one-hot encoded for eight soybean genotypes, producing a binary vector representation for each genotype class [39]. Because the encoded values were binary (0 or 1), no scaling was applied. DAS was represented as a single numeric feature indicating the number of days between sowing and image acquisition. The metadata features were processed through a four-layer multilayer perceptron (MLP), concatenated with the pooled sensor feature vectors, and passed to a fully connected regression layer to produce the trait estimate (Figure 5). GEN was included in both the single- and multi-date models, whereas DAS was included only in the multi-date models. GEN and DAS were the only ancillary non-image inputs; rooting-depth-restriction and nitrogen-treatment labels were not supplied to the network.

The evaluated input configurations provide incremental comparisons relative to an MSI baseline. Additional sensor and metadata inputs were progressively incorporated through the MSI, MSI + THR, MSI + LDR + THR, and metadata-enhanced configurations. This design allowed the complementary contribution of THR, LDR, GEN, and DAS to be assessed within the multimodal fusion framework.

3.2.4. Model Training

Deep neural networks rely on an optimizer to learn the mapping between extracted feature maps and the target variable. The principal idea behind this operation is gradient descent, which is a way to minimize an objective function $J(\theta)$ parameterized by a model's parameters $\theta \in \mathbb{R}^d$ and by updating the parameters in the opposite direction of the gradient $\nabla_\theta J(\theta)$ with respect to the parameters. However, numerous advancements have emerged beyond standard gradient descent, and we selected the Adam optimizer due to its widespread success and empirical performance superiority [40]. Adam requires the calculation of m_t and v_t, where m_t is the exponentially decaying average of past gradients, v_t is the exponentially decaying average of past squared gradients, and the gradient at any timestamp t can be denoted as g_t.

$$m_t = \beta_1 m_{t-1} + (1 - \beta_1) g_t \tag{5}$$

$$v_t = \beta_2 v_{t-1} + (1 - \beta_2) g_t^2 \tag{6}$$

where β_1 and β_2 are very small decay rates, recommended as 0.9 and 0.999, respectively, by the authors of Adam [41]. The m_t and v_t were found to be biased towards zero, due to the small decay rates, which are counteracted by computing the bias-corrected first and second moment estimates:

$$\hat{m}_t = \frac{m_t}{1 - \beta_1^t} \tag{7}$$

$$\hat{v}_t = \frac{v_t}{1 - \beta_2^t} \tag{8}$$

The $\hat{m}_t$ and $\hat{v}_t$ are then used to update the model parameters by the Adam update rule:

$$\theta_{t+1} = \theta_t - \frac{\eta}{\sqrt{\hat{v}_t} + \epsilon} \hat{m}_t \tag{9}$$

where η is the learning rate and ϵ is a smoothing term that avoids division by zero.

The networks were trained as regression models using the Adam optimizer and mean squared error (MSE) loss. The initial learning rate was 0.001, and the batch size was 32. Training was allowed to continue for a maximum of 500 epochs. Early stopping was applied independently to each trait and input configuration when the validation loss failed to improve for 20 consecutive epochs. A randomized validation subset was created from the training data, while the test set was used for final model evaluation. All models were trained on a workstation equipped with an Intel Xeon Platinum 8168 processor (2.7 GHz, 24 cores), 512 GB of RAM, and two NVIDIA Quadro P8000 GPUs.

3.2.5. Model Evaluation

The performance of the models was evaluated using two standard metrics: the coefficient of determination (R^2) and the root mean squared error (RMSE). These metrics quantify the accuracy of the model's predictions relative to actual observations. The equations are as follows:

$$R^2 = 1 - \frac{\sum_{i=1}^{n}(y_i - \hat{y}_i)^2}{\sum_{i=1}^{n}(y_i - \overline{y})^2} \tag{10}$$

$$RMSE = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2} \tag{11}$$

where n is the number of test samples, y_i represents the measured seed composition, $\hat{y}_i$ is the predicted value, and $\overline{y}$ is the mean of the observed values.

4. Results

4.1. Descriptive Statistics

A total of 372 soybean samples, collected over two growing seasons, were analyzed using the DA 7250 NIR Analyzer (PerkinElmer 201, Waltham, MA, USA) to determine seed composition (Table 3). Due to the varying concentration ranges among seed components, we used the coefficient of variation (CV), a unitless measure representing the standard deviation relative to the mean, to assess variability across traits. Sucrose (CV: 15.8%) and starch (CV: 12.93%) exhibited the greatest variability, whereas fiber (CV: 3.81%) and ash (CV: 2.75%) showed the least variability.

Table 3. Descriptive statistics of soybean seed-composition measurements on a dry-weight basis (%).

Component	Sample Size	Maximum	Minimum	Mean	SD	CV (%)
Protein	372	43.37	36.1	39.21	1.59	4.06
Oil	372	25.39	19.02	22.63	1.37	6.07
Comp. Carbs	372	11.97	8.46	10.28	0.64	6.24
Fiber	372	7.28	5.87	6.52	0.25	3.81
Sucrose	372	8.15	4.07	5.74	0.91	15.8
Ash	372	5.87	5.02	5.34	0.15	2.75
Starch	372	4.89	2.21	3.76	0.49	12.93
Simp. Carbs	372	1.66	0.9	1.09	0.1	8.89

Figure 6 presents both the actual and normalized distributions for all eight composition components. Consistent with typical soybean seed composition [42], protein and oil are the dominant constituents, followed by complex carbohydrates (fiber and starch), while ash and simple sugars are present in lower amounts. These results indicate that our NIR-derived estimates closely align with the expected quantitative trait distributions in soybean seeds. Additionally, Figure 7 compares seed composition across the eight soybean genotypes, highlighting genotype-specific differences in average concentrations.

4.2. Performance of Image-Based Models

Imagery-based models were developed using either single observations from UAV data (single date) or multi-date observations. Single-date models employed lighter architectures using 2D CNNs, whereas multi-date models used deeper 3D CNN architectures to effectively capture spatiotemporal information. Additionally, due to varying flight dates across experimental years and fields, we defined four distinct flight dates for modeling purposes (Table 4).

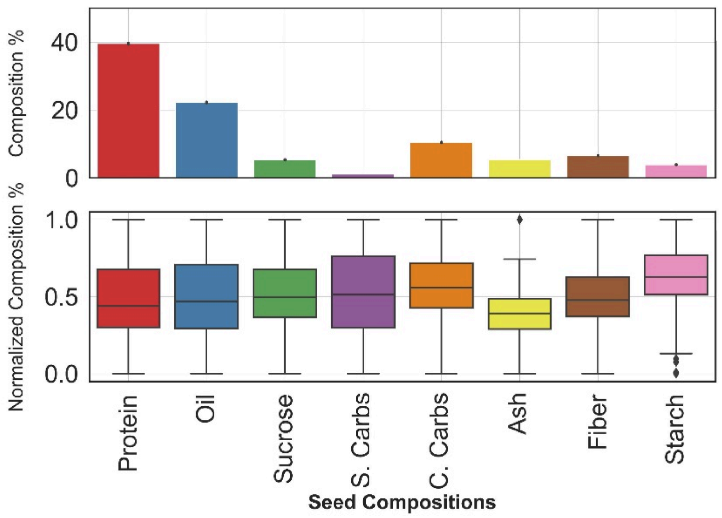

Figure 6. (**a**) Bar chart showing the distribution of soybean seed compositions as measured by the DA 7250 NIR Analyzer, indicating that protein and oil are the dominant components, followed by complex carbohydrates (fiber and starch). (**b**) Box plot showing the normalized distribution for the eight components, where each box represents the middle 50% of data points and the whiskers extend to the remaining data, excluding outliers.

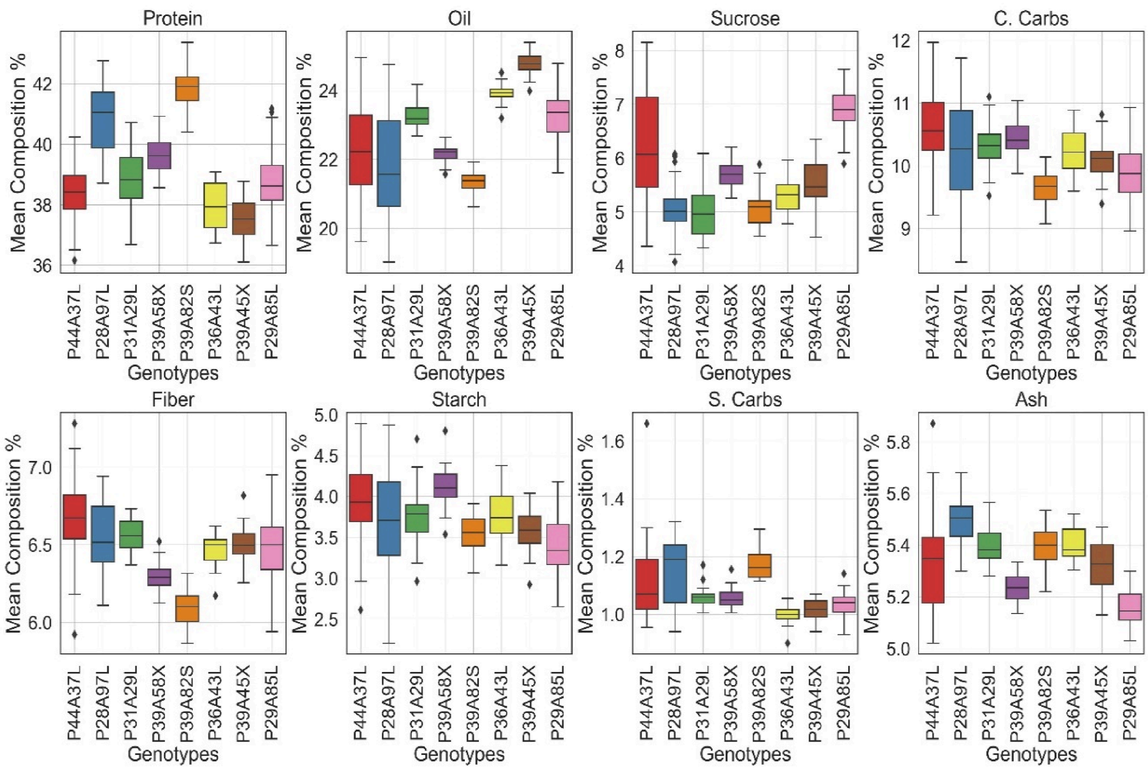

Figure 7. Average seed concentration per genotype for all samples and compositions, illustrating the variability across the eight soybean varieties.

Figures 8–10 present these input configurations as incremental, configuration-level comparisons. MSI served as the imagery baseline, allowing the effects of adding THR, LDR, GEN, and DAS to be examined within the evaluated model sequence. These comparisons should not be interpreted as complete attribution of predictive power to individual inputs.

Table 4. Flight dates considered in this study for single- and multi-date model configurations.

	2020	2021	
Flight Date	**H1G**	**H1G**	**L2**
F1	24 July	21 July	21 July
F2	24 July	10 August	10 August
F3	19 August	24 August	24 August
F4	4 September	9 September	9 September

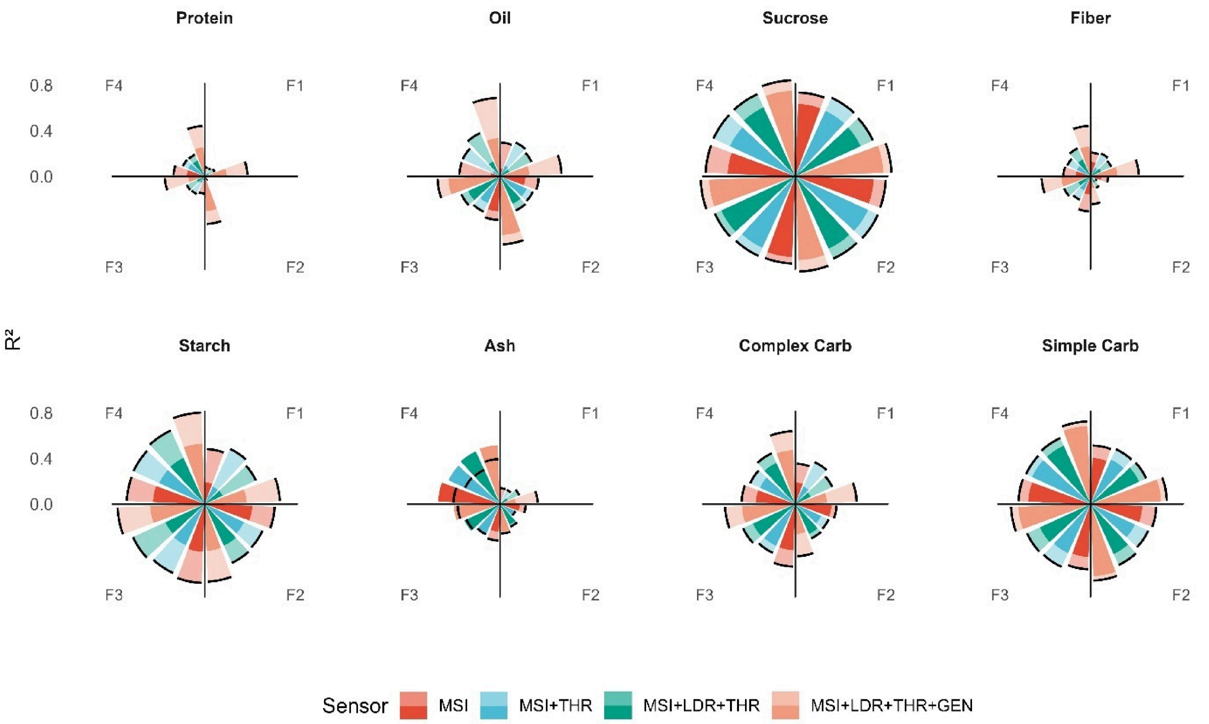

Figure 8. Modeling performance (R^2) of single-date image-based models for eight soybean seed composition traits across four UAV flight dates (F1–F4) and different sensor combinations. Each radial plot represents model performance (R^2) for a specific trait, evaluated using data from multispectral (MSI), thermal (THR), and LiDAR (LDR) imagery, with and without genotype (GEN) information. Bars extend radially from the center for each flight date, with sensor combinations color-coded as follows: MSI only (red), MSI + THR (blue), MSI + LDR + THR (green), and MSI + LDR + THR + GEN (orange). Faded (transparent) bars represent training performance, while darker bars represent test performance.

As the F1 stage represents a non-critical early growth phase with smaller canopy development, and due to limited data availability for the 2020 H1G field, the 24 July 2020 flight was used to represent both F1 and F2 stages to ensure consistency across model configurations and to maintain uniform temporal input dimensions for the 3D CNN architectures.

4.2.1. Single Date Models

Our results revealed distinct performance levels across traits. Sucrose exhibited the highest predictive ability (R^2: 0.59 to 0.77), followed by simple carbohydrates (0.40 to 0.68), ash (0.01 to 0.54), starch (0.17 to 0.53), oil (0.07 to 0.51), and complex carbohydrates (0.15 to 0.47). In contrast, predictions for protein (0.01 to 0.30) and fiber (0.03 to 0.26) were consistently lower. Notably, incorporating GEN information consistently enhanced model performance.

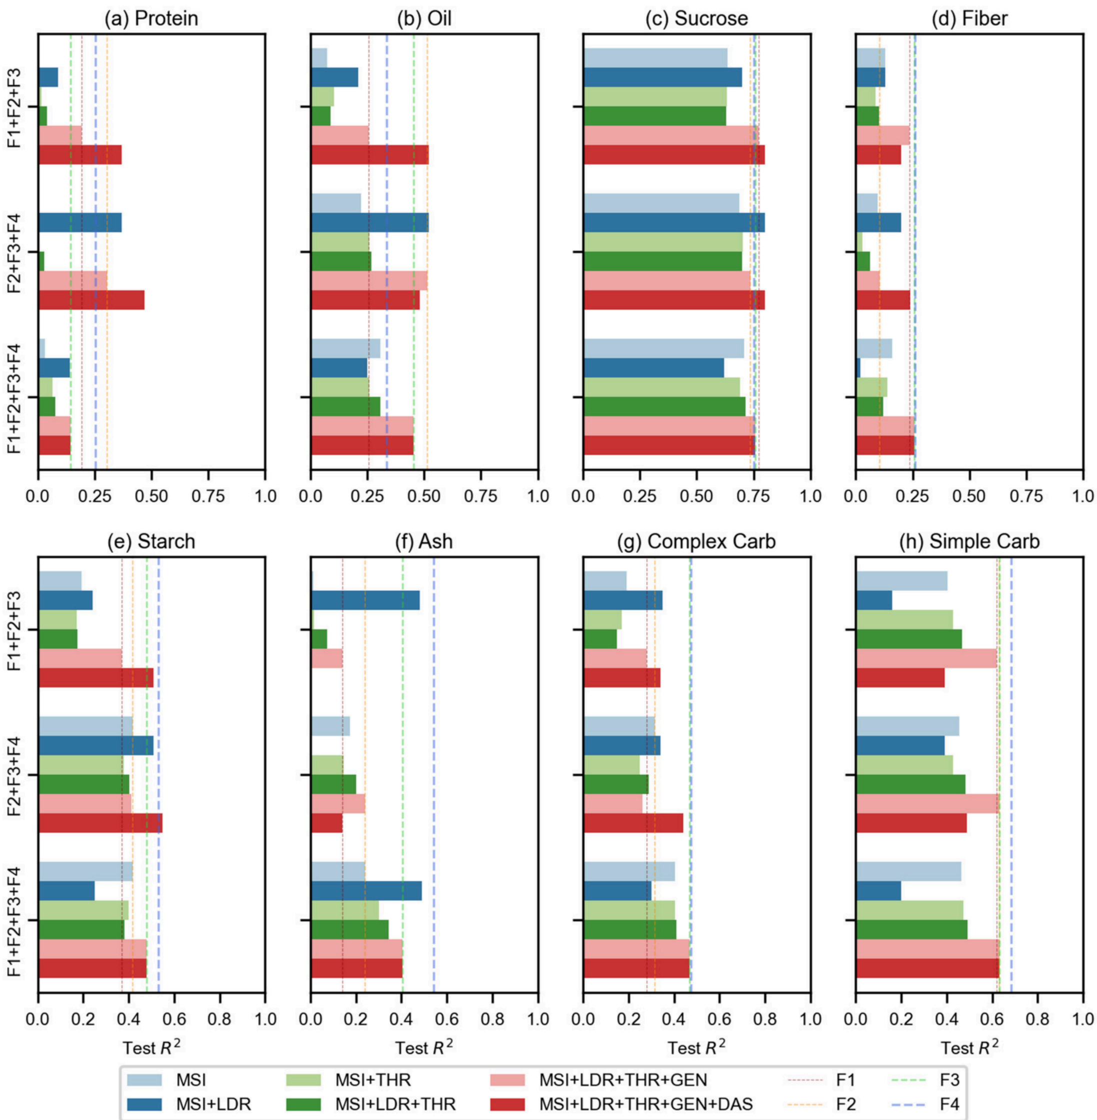

Figure 9. Multi-date model performance using multimodal UAV imagery to estimate soybean seed composition traits across four UAV flight dates. For each soybean seed composition trait, three multi-date configurations were evaluated by combining different subsets of the four temporal data points (F1, F2, F3, and F4). In each subplot, the bars represent results from various sensor combinations (MSI, THR, LDR, GEN), while lines indicate the best single-date model performance for the corresponding date.

We also examined the impact of UAV flight dates on model performance. Each single-date model was trained using data from one of four flight dates (F1–F4). Our findings indicate that some traits are best predicted at earlier growth stages, while others benefit from later-season imagery. Specifically, protein (R^2 = 0.30) and oil (R^2 = 0.51) performed best with F2 data, while sucrose peaked as early as F1 (R^2 = 0.77). In the absence of GEN information, oil (R^2 = 0.31) and sucrose (R^2 = 0.72) reached their highest accuracies at F3, yet most other traits performed better with the latest F4 data. Overall, incorporating GEN information improved model performance across all flight dates, underscoring the importance of both temporal selection (F1–F4) and GEN information in improving predictions.

In terms of sensor modalities, our results demonstrate that including GEN information significantly enhances model performance. With the sole exception of the ash model at F4, where MSI data alone was sufficient, every configuration showed improved accuracy when GEN information was incorporated into the fusion process. As shown in Figure 8, performance trends remained consistent over time, regardless of the sensor used. For example, protein predictions steadily improved with later-stage UAV data, while oil models increased in accuracy until F3 before declining at F4. Notably, in the absence of GEN information, the differences among sensor fusion configurations were minimal for a given date and trait, suggesting that MSI alone can provide robust input for the neural network, often yielding greater information gain than more complex inputs like LDR or THR data.

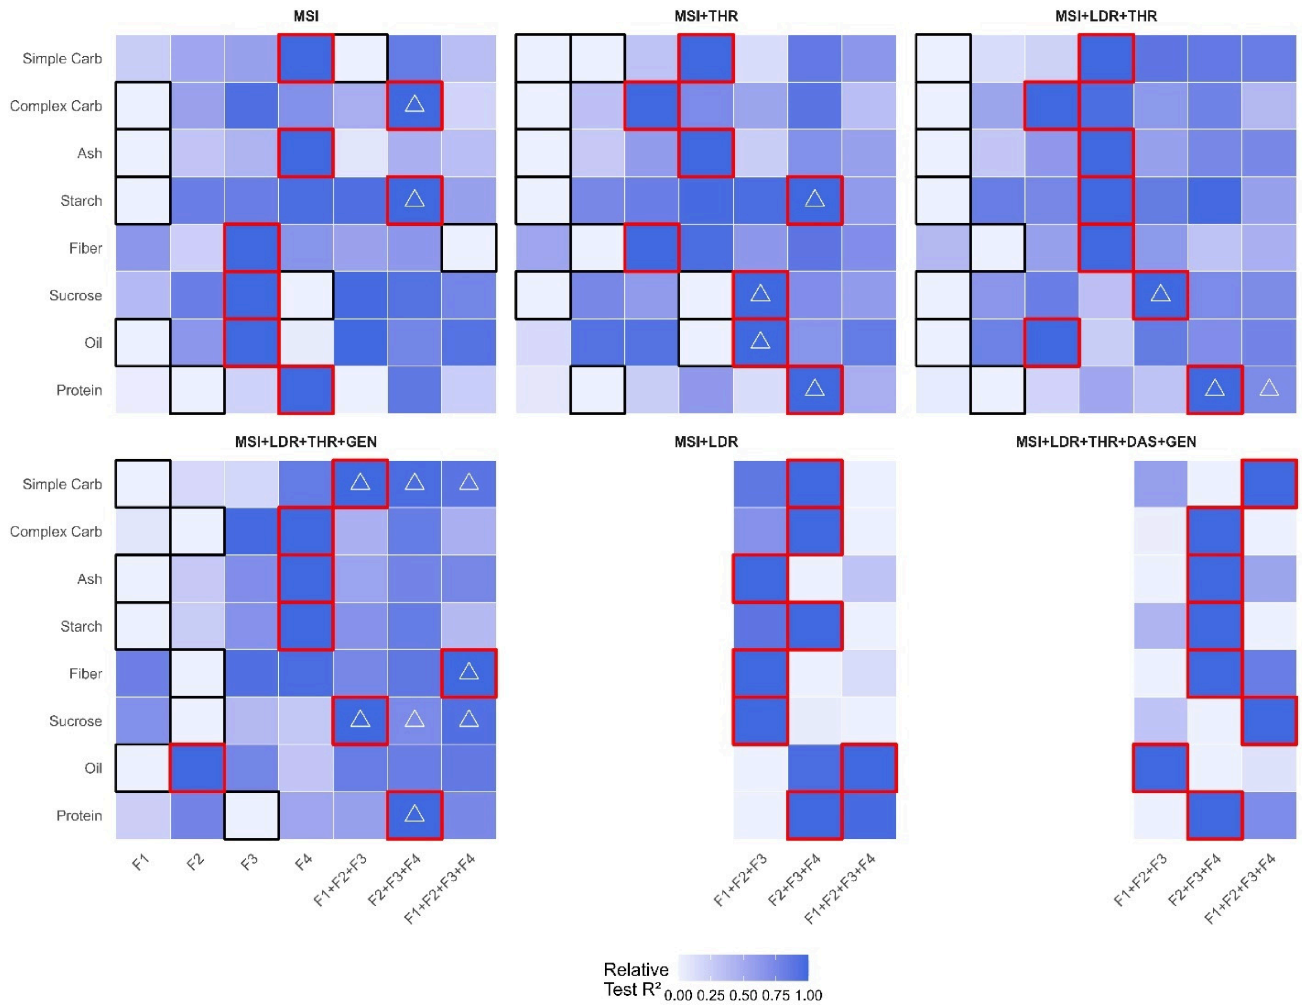

Figure 10. Relative test R^2 performance of multi-date models for eight soybean seed composition traits across sensor and temporal configurations. Heatmaps show model accuracy (R^2) by trait (rows) and flight configuration (columns) for all sensor combinations of MSI, THR, LDR, GEN, and DAS. Red borders highlight the best-performing configuration per trait, while triangles indicate multi-date models that outperformed all single-date models.

4.2.2. Multi-Date Models

We extended our analysis to multi-date image-based models that integrate UAV imagery from multiple flight dates to capture spatiotemporal dynamics. Three configurations were evaluated: (1) F1 + F2 + F3, which integrates data from the first three flight dates; (2) F2 + F3 + F4, which excludes the earliest flight and includes the final date; and (3) F1 + F2 + F3 + F4, which incorporates imagery from all four flight dates (Table 4). These configurations, along with varying sensor inputs and seed composition traits, provided a comprehensive evaluation of multi-date model performance.

Overall, the multi-date models exhibited performance patterns largely consistent with the single-date approach. Sucrose remained the highest predictive trait, achieving R^2 values ranging from 0.61 to 0.80, a modest improvement over the single-date performance (R^2: 0.59 to 0.77). The overall order for the remaining traits was similar, with simple carbohydrates (R^2: 0.35 to 0.70), ash (R^2: 0.06 to 0.44), starch (R^2: 0.16 to 0.55), oil (R^2: 0.20 to 0.52), and complex carbohydrates (R^2: 0.09 to 0.49). Notably, the prediction accuracy for protein improved considerably in the multi-date framework, with R^2 values ranging from 0.01 to 0.47 compared to 0.01 to 0.30 observed in single-date models.

Evidently, a detailed comparison of temporal configurations revealed that, while combining data from all four flight dates (F1 + F2 + F3 + F4) provides improved spatiotemporal information to the 3D models, this configuration did not necessarily yield superior performance compared to the three-date models (F1 + F2 + F3 or F2 + F3 + F4). In many sensor fusion categories, the four-date models yielded only slight, and sometimes negligible, improvements. For example, with the MSI + LDR + THR + GEN configuration, the fiber prediction achieved an R^2 of 0.27 using the 4-date combination, compared to R^2 values of 0.23 and 0.25 with other temporal configurations. Moreover, for several traits, the three-date models even outperformed the 4-date configuration. Notably, the F1 + F2 + F3 setup frequently outperformed the F2 + F3 + F4 combination, emphasizing the effectiveness of including earlier season imagery in the prediction of seed composition. For instance, oil prediction models built from F1 + F2 + F3 data achieved an R^2 of 0.52 with an RMSE of 4.31%, outperforming other temporal combinations, though for most traits, performance differences across configurations remained marginal (Figure 9).

Furthermore, the inclusion of GEN information consistently enhanced model performance across all temporal configurations. In addition, incorporating DAS information as a separate input stream further improved predictions for several traits. For example, protein prediction benefited notably from the simultaneous addition of GEN and DAS data across all multi-date configurations (Figure 9a), and oil and fiber models also showed noticeable improvements with the inclusion of DAS (Figure 9b,d). However, for traits like sucrose, starch, and complex carbohydrates, our models constructed solely on MSI performed comparably to those using additional sensor fusion inputs. These findings emphasize that while temporal depth from multi-date imagery is crucial, strategic sensor fusion and the integration of ancillary inputs such as GEN and DAS are key to optimizing predictive performance.

4.2.3. Cross-Configuration Comparison

Figure 10 summarizes relative test R^2 across the evaluated temporal and input configurations for the eight seed-composition traits. Multi-date configurations matched or exceeded the best single-date configurations for several traits, although the most informative temporal combination varied among traits. The F1 + F2 + F3 combination frequently produced the best or near-best performance, whereas including all four dates did not consistently improve accuracy. MSI-based configurations remained competitive across traits, while adding THR and LDR did not consistently improve performance. The inclusion of GEN and DAS improved several configurations, with the clearest gains observed for protein and oil.

5. Discussion

5.1. Contribution of UAV-Based Image Modalities for Seed Composition

Accurate prediction of soybean seed composition traits during crop growth is essential for enabling timely and precise decision-making in breeding and agricultural production systems. UAV-based remote sensing offers a powerful and scalable solution, uniquely capturing subtle structural, textural, and spectral variations in the canopy throughout critical soybean growth stages [43].

Compared with traditional feature-based approaches, imagery-based modeling enables rapid, repeatable, and non-destructive data collection at high spatial resolutions while allowing end-to-end models to learn representations without manually engineered features [44]. These characteristics may improve scalability, although model robustness and generalization still depend on the diversity and representativeness of the training and validation data.

Table 5 summarizes the potential canopy and metadata signals that may underlie the observed canopy-to-seed associations, together with the most informative acquisition timing and best reported test performance.

Table 5. Potential canopy and metadata signals, informative acquisition timings, and best reported test performance for each seed-composition trait.

Trait	Potential Canopy or Metadata Signals	Most Informative Timing Observed	Best Reported Test (R^2)
Protein	MSI spectral response potentially associated with canopy vigor and senescence; GEN and DAS provide genetic and developmental context	F2 among single-date models; multi-date integration improved peak performance	Single-date: 0.30; Multi-date: 0.47
Oil	MSI spectral response and reproductive phenology; GEN provides additional genetic context	F2 among single-date models; F1 + F2 + F3 produced the best reported multi-date performance	Single-date: 0.51; Multi-date: 0.52
Sucrose	MSI spectral response and early reproductive phenology	F1 among single-date models; multi-date integration provided a modest improvement	Single-date: 0.77; Multi-date: 0.80
Fiber	Potential canopy structural and later-season phenological information; DAS provides developmental context	Later reproductive acquisitions; F1 + F2 + F3 + F4 produced the highest reported multi-date performance	Single-date: 0.26; Multi-date: 0.27
Starch	MSI spectral response and later-season reproductive phenology	Later reproductive acquisitions	Single-date: 0.53; Multi-date: 0.55
Ash	MSI spectral response and potential canopy structural information from LDR	F4 in the single-date comparison	Single-date: 0.54; Multi-date: 0.44
Complex carbohydrates	MSI spectral response and reproductive phenology	Later reproductive acquisitions	Single-date: 0.47; Multi-date: 0.49
Simple carbohydrates	MSI spectral response and temporal canopy development	Later reproductive and multi-date acquisitions	Single-date: 0.68; Multi-date: 0.70

The trait-specific results indicate that the prediction pathway was not uniform across seed components. Sucrose and simple carbohydrates showed the strongest empirical associations with the evaluated inputs, whereas protein and fiber were estimated less consistently. Improvements following the inclusion of GEN and DAS also indicate that some of the predictive information originated from genotype and phenological metadata rather than from UAV imagery alone.

Our UAV data acquisition strategy was specifically designed to encompass critical development stages of soybean from R4 through R7. Prior research has clearly indicated the R5-R7 stages as essential phases for soybean seed development, marked by biosynthesis and accumulation of major seed components such as protein, oil, and carbohydrates [45]. During these stages, UAV-based remote sensing can effectively capture canopy structural dynamics, including leaf area expansion during R5 and R6, and leaf senescence occurring in late R7 through variations in spectral, textural, and structural canopy characteristics [46]. This temporal design allowed us to evaluate whether integrating canopy observations across multiple reproductive stages improved prediction relative to individual flight dates.

Among single-date models, sucrose was the most accurately predicted trait (R^2: 0.59–0.77), with peak performance achieved using early-season (F1) imagery. Protein and oil predictions were optimized with mid-season (F2) data (protein: R^2 = 0.30; oil: R^2 = 0.51). Including GEN information notably enhanced predictive accuracy across all single-date models, highlighting its role in capturing genotype-driven variations in seed composition. Later-season imagery (F3, F4) yielded improved accuracy for other traits (e.g., carbohydrates, starch, ash), reflecting pronounced canopy structural and physiological changes during crop senescence. These trait-specific findings are consistent with prior studies highlighting that seed composition components are differentially influenced by physiological processes across developmental stages. Specifically, oil and protein accumulation are most active during the R4-R5 reproductive phases, when biosynthesis peaks [43]. For other traits, later-date imagery consistently improved model performance, though the underlying cause remains unclear, as optimal growth stages may vary with genotype, environment, and management practices [47].

Additionally, the multi-date imagery models integrated UAV imagery from multiple flight dates to better capture temporal variations in soybean canopy characteristics. Three temporal combinations—early-to-mid-season (F1 + F2 + F3), mid-to-late-season (F2 + F3 + F4), and the comprehensive all-date combination (F1 + F2 + F3 + F4)—were evaluated. Overall, multi-date models consistently outperformed or performed comparably to single-date models. For instance, sucrose predictions improved modestly (R^2 increased from 0.77 to 0.80), while protein predictions improved substantially (R^2 increased from 0.30 to 0.47). Notably, the early-to-mid-season combination (F1 + F2 + F3) often yielded the best or near-best accuracy among the multi-date configurations. This pattern is consistent with the value of integrating canopy information across key reproductive stages [48]. Adding GEN and DAS generally improved model performance, indicating that these variables provided complementary genetic and phenological context beyond the UAV imagery alone.

Building on these findings, Figure 10 illustrates a perspective on how temporal and sensor-based variation influences predictive performance across eight soybean seed composition traits. Each heatmap illustrates the relative test R^2 across temporal configurations and sensor combinations, with red-bordered cells indicating the best performance per trait and triangles highlighting multi-date setups that outperformed all single-date models. Multi-date configurations generally matched or exceeded the best single-date configurations for several traits, although the most informative temporal combination varied among traits.

Predictions for protein, oil, and sucrose notably improved when multi-date imagery was augmented with GEN and DAS information, emphasizing the added value of integrating phenological and genetic context. Interestingly, while MSI and MSI + THR configurations performed well, adding LDR to the MSI + THR configuration did not consistently improve performance. This may indicate that LDR provided limited additional predictive information for these traits or introduced redundancy with information already represented by the other sensors. These results suggest that multi-date models may improve trait prediction by integrating canopy information observed across multiple developmental stages, particularly during the R5-R7 phases. Notably, the F1 + F2 + F3 configuration often outperformed the full four-date combination, potentially due to information saturation or noise introduced by later imagery. Overall, Figure 10 underscores the importance of strategically selecting sensor modalities and temporal configurations, especially when guided by biological timing, to optimize the performance of UAV-based seed composition models.

In comparison, single-date models offer practical simplicity and cost-efficiency, making them ideal for large-scale, routine applications, particularly in resource-limited settings. In contrast, multi-date models, though more computationally intensive, benefit in scenarios where precise trait prediction is crucial. The insights gained from integrating multimodal imagery, genotype, and temporal data may improve the precision of soybean seed composition prediction and field-level decision-making strategies.

5.2. Contribution of UAV-Based Sensor Fusion for Predicting Seed Composition

In the scope of GeoAI, multimodal remote sensing fusion offers a promising approach for understanding variations in crop health and yield [24]. Multimodal image fusion involves integrating data from multiple sensor types, including spectral, structural, and thermal modalities [49]. Our models were developed to predict soybean seed composition traits based on these UAV-based image modalities. Seed composition traits are closely linked to crop health, vigor, and environmental conditions [50]; therefore, we hypothesized that integrating multiple sensor types could enhance model performance. Specifically, we considered MSI to capture spectral responses at key wavelengths, LDR for detailed structural information, and THR to assess canopy temperature variation, a known indicator of physiological stress [51].

Our findings showed that in many cases, MSI alone outperformed other sensor combinations, especially when GEN information was included. This superior performance was likely because MSI data contained sufficient spatial and spectral variation per plot, which was effectively leveraged by the 2D CNN architecture. Such CNN advantages in managing spatial and spectral complexity are well-documented in remote sensing and computer vision literature [52]. Although LDR provided valuable canopy structural information, the inclusion of complex 3D voxel data increased model complexity without necessarily enhancing overall predictive accuracy. However, LDR notably improved ash content prediction, suggesting that morphological traits may be particularly relevant for certain seed components. Similarly, THR data provided physiological insights linked to plant stress and metabolic activity [53], though its individual contribution did not surpass MSI in predictive performance.

The limited and inconsistent gains from adding THR and LDR may reflect overlapping canopy information across modalities, sensor-specific noise, or increased model complexity relative to the available sample size. Because the current feature-level fusion approach concatenates the learned representations from all sensor branches, it does not explicitly reduce the influence of a modality that provides limited additional information for a particular trait, sample, or developmental stage. Future studies could evaluate attention or gating mechanisms that adaptively weight sensor contributions, as well as decision-level fusion that combines predictions from separately trained modality-specific models. These approaches may reduce the propagation of redundant or noisy information, although their benefits would need to be demonstrated through direct comparison and grouped validation.

The incorporation of GEN information emerged as a major strength of our imagery models. GEN information notably improved predictions, particularly for protein and oil content. This improvement reflected the experimental design, which included eight distinct soybean genotypes that introduced substantial variation in protein and oil content [54].

By encoding genotypes as one-hot vectors, the models effectively captured genetic contributions to trait variability, highlighting the potential benefits of combining UAV-based imagery with GEN information in seed composition modeling.

Importantly, the observed improvements following the addition of GEN and DAS do not establish how accurately these metadata variables would predict seed composition independently of UAV imagery. Similarly, the contributions of THR and LDR are conditional on their inclusion in MSI-based configurations. The present results therefore support configuration-level attribution rather than standalone estimates of each input's predictive capacity.

These findings underscore the importance of selecting appropriate sensor modalities and integrating supplementary management information. Future research across additional sites, diverse genotypes, and varied environments is needed to further elucidate the interactions among image modalities and genotype data, solidifying their practical application in digital agriculture.

The field treatments were not balanced across fields, years, and genotypes: rooting-depth restriction was confined to H1G, whereas nitrogen treatment was applied to a 24-plot subset of L2. Because these treatment labels were not included as model inputs and model performance was not evaluated separately by treatment, the present results do not establish prediction robustness under specific water- or nitrogen-related stresses. Future studies should use balanced treatment designs and treatment-stratified evaluation to quantify stress-specific prediction error.

5.3. Limitation and Future Work

The experimental scope should be considered when interpreting the findings. The dataset comprised 372 plots from two fields at one research center over the 2020 and 2021 growing seasons, and the genotypes were not uniformly represented across all field-year combinations. In addition, the F1–F4 design represented relative flight positions rather than identical phenological stages across fields and years. The present design supports comparison of the evaluated configurations within this heterogeneous field experiment. Future multi-environment studies that replicate genotypes across locations and years and align acquisitions by developmental stage would provide a stronger basis for separating genetic, environmental, and phenological contributions.

The rooting-depth restrictions and nitrogen treatments added useful environmental and phenotypic heterogeneity but were applied within specific field subsets. Because the current analysis pooled these experimental conditions, balanced treatment designs and treatment-stratified evaluation would help characterize prediction behavior under specific water- and nitrogen-related conditions. As the models are trained on values derived from NIR instead of wet chemistry data, the error from the NIR calibration acts as a 'noise floor.' This baseline uncertainty ultimately limits the maximum accuracy the UAV models can achieve.

Finally, the evaluated configuration sequence provides evidence about the incremental contribution of inputs within the multimodal framework but does not fully explain which bands, image regions, dates, or LiDAR structures drove individual predictions. Future work could combine repeated grouped validation with systematic input ablations and band-, spatial-, temporal-, and LiDAR-layer attribution. Together with attention-based or decision-level fusion, these analyses could support more interpretable and efficient trait-specific models. Furthermore, as this study focused on single model runs and evaluated performance primarily via R^2, future work should incorporate uncertainty estimates.

6. Conclusions

In this study, we evaluated whether multimodal UAV canopy observations, integrated with genotype and phenological metadata, could support pre-harvest prediction of soybean seed composition using an end-to-end convolutional neural network (CNN) architecture. Our results demonstrate that multispectral (MSI), thermal (THR), and LiDAR (LDR) canopy observations can be associated with post-harvest seed-composition reference values, although the sensors do not measure seed biochemical properties directly. Sucrose and oil emerged as the most accurately predicted traits, while fiber and ash remained less responsive to image-based modeling. The multi-date models generally outperformed single-date models, particularly when incorporating earlier season imagery (F1 + F2 + F3), suggesting the importance of phenological transitions during the R5-R7 stages. Adding GEN and DAS generally improved performance within the evaluated configurations, with the clearest gains observed in protein and oil.

- Sucrose exhibited the highest prediction accuracy (R^2 up to 0.80), with peak performance achieved using early-season imagery and minimal gain from complex sensor fusion.
- Oil and protein predictions improved notably with GEN inclusion, reaching R^2 values of 0.52 and 0.47 respectively in multi-date models.
- Adding GEN and DAS generally improved performance within the evaluated MSI-based configurations, particularly for protein and oil; however, metadata-only performance was not evaluated.
- The F1 + F2 + F3 temporal configuration frequently outperformed the full F1–F4 set, indicating potential performance degradation due to later imagery.
- MSI consistently contributed the most predictive power across traits, while LDR improved ash prediction but was redundant for others.

Collectively, these findings demonstrate the potential of temporally informed CNN models that integrate multimodal UAV observations with agronomic metadata for pre-harvest prediction of soybean seed composition. The framework provides a foundation for broader evaluation across environments, growing seasons, and soybean germplasm, as well as for its continued development toward breeding and precision-agriculture applications.

Author Contributions: Conceptualization, V.S., K.R. and F.F.; Methodology, V.S., K.R. and H.A.; Software, V.S.; Validation, V.S., K.R., S.B. and B.R.; Formal analysis, K.R., S.B., A.G., A.P., B.R. and S.S.; Investigation, K.R. and S.B.; Resources, V.S. and B.R.; Data curation, H.A.; Writing—original draft, V.S., K.R., S.B., Haireti Alifu, A.G., Supria Sarkar and F.F.; Writing—review & editing, V.S., A.P. and B.R.; Visualization, K.R., S.B., A.P., B.R. and S.S.; Supervision, V.S. and F.F.; Project administration, V.S. and F.F.; Funding acquisition, V.S. and F.F. All authors have read and agreed to the published version of the manuscript.

Funding: This work was supported by the United Soybean Board and the Foundation for Food & Agricultural Research (project numbers: 2120-152-0201 and 2331-201-0103), and partially by the U.S. Geological Survey under Grant/Cooperative Agreement No. G23AP00683 (GY23-GY27).

Data Availability Statement: The data presented in this study are available on request from the corresponding author.

Acknowledgments: The authors express their gratitude to the Remote Sensing Lab team at Saint Louis University and the ground data collection team at the University of Missouri in Columbia for their contributions.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Bellaloui, N.; Reddy, K.N.; Bruns, H.A.; Gillen, A.; Mengistu, A.; Zobiole, L.; Fisher, D.K.; Abbas, H.K.; Zablotowicz, R.M.; Kremer, R. Soybean seed composition and quality: Interactions of environment, genotype, and management practices. In *Soybeans: Cultivation, Uses and Nutrition*; Nova Science Publishers, Inc.: New York, NY, USA, 2011; pp. 1–42.
2. Hudson, K. Soybean Protein and Oil Variants Identified through a Forward Genetic Screen for Seed Composition. *Plants* **2022**, *11*, 2966. [CrossRef] [PubMed]
3. LeDell, E.; Poirier, S. H2o automl: Scalable automatic machine learning. In Proceedings of the AutoML Workshop at ICML, Online, 17–18 July 2020; Volume 2020.
4. Siamabele, B. The significance of soybean production in the face of changing climates in Africa. *Cogent Food Agric.* **2021**, *7*, 1933745. [CrossRef]
5. Medic, J.; Atkinson, C.; Hurburgh, C.R. Current Knowledge in Soybean Composition. *J. Am. Oil Chem. Soc.* **2014**, *91*, 363–384. [CrossRef]
6. He, J.; Evans, N.M.; Liu, H.Z.; Shao, S.Q. A review of research on plant-based meat alternatives: Driving forces, history, manufacturing, and consumer attitudes. *Compr. Rev. Food Sci. Food Saf.* **2020**, *19*, 2639–2656. [CrossRef] [PubMed]
7. Song, H.; Taylor, D.C.; Zhang, M. Bioengineering of Soybean Oil and Its Impact on Agronomic Traits. *Int. J. Mol. Sci.* **2023**, *24*, 2256. [CrossRef] [PubMed]
8. Declerck, F.; Hikouatcha, P.; Tchoffo, G.; Tédongap, R. Biofuel policies and their ripple effects: An analysis of vegetable oil price dynamics and global consumer responses. *Energy Econ.* **2023**, *128*, 107127. [CrossRef]
9. Aoyagi, W.S.A. *History of Soybean Yields and Yield Research: Extensively Annotated Bibliography and Sourcebook*; Soyinfo Center: Lafayette, CA, USA, 2021.
10. Assefa, Y.; Bajjalieh, N.; Archontoulis, S.; Casteel, S.; Davidson, D.; Kovács, P.; Naeve, S.; Ciampitti, I.A. Spatial Characterization of Soybean Yield and Quality (Amino Acids, Oil, and Protein) for United States. *Sci. Rep.* **2018**, *8*, 14653. [CrossRef] [PubMed]
11. Assefa, Y.; Purcell, L.C.; Salmeron, M.; Naeve, S.; Casteel, S.N.; Kovács, P.; Archontoulis, S.; Licht, M.; Below, F.; Kandel, H.; et al. Assessing Variation in US Soybean Seed Composition (Protein and Oil). *Front. Plant Sci.* **2019**, *10*, 298. [CrossRef] [PubMed]
12. Aulia, R.; Kim, Y.; Zuhrotul Amanah, H.; Muhammad Akbar Andi, A.; Kim, H.; Kim, H.; Lee, W.-H.; Kim, K.-H.; Baek, J.-H.; Cho, B.-K. Non-destructive prediction of protein contents of soybean seeds using near-infrared hyperspectral imaging. *Infrared Phys. Technol.* **2022**, *127*, 104365. [CrossRef]
13. Agelet, L.E.; Hurburgh, C.R. Limitations and current applications of Near Infrared Spectroscopy for single seed analysis. *Talanta* **2014**, *121*, 288–299. [CrossRef] [PubMed]
14. Venkatesan, S.; Masilamani, P.; Janaki, P.; Eevera, T.; Sundareswaran, S.; Rajkumar, P. Role of Near -Infrared Spectroscopy in Seed Quality Evaluation: A Review. *Agric. Rev.* **2020**, *41*, 106–115. [CrossRef]
15. Al-Amery, M.; Armstrong, P.R.; Geneve, R.L.; Hildebrand, D.F.; Lee, C.; Maghirang, E.B.; Sanches, M.F.; Vieira, R.D. Near-infrared spectroscopy used to predict soybean seed germination and vigour. *Seed Sci. Res.* **2018**, *28*, 245–252. [CrossRef]
16. Sarkar, S.; Sagan, V.; Bhadra, S.; Rhodes, K.; Pokharel, M.; Fritschi, F.B. Soybean seed composition prediction from standing crops using PlanetScope satellite imagery and machine learning. *ISPRS J. Photogramm. Remote Sens.* **2023**, *204*, 257–274. [CrossRef]
17. Zhang, C.H.; Kovacs, J.M. The application of small unmanned aerial systems for precision agriculture: A review. *Precis. Agric.* **2012**, *13*, 693–712. [CrossRef]
18. Dormann, C.F.; Elith, J.; Bacher, S.; Buchmann, C.; Carl, G.; Carré, G.; Marquéz, J.R.G.; Gruber, B.; Lafourcade, B.; Leitão, P.J.; et al. Collinearity: A review of methods to deal with it and a simulation study evaluating their performance. *Ecography* **2013**, *36*, 27–46. [CrossRef]
19. Zhu, X.X.; Tuia, D.; Mou, L.; Xia, G.S.; Zhang, L.; Xu, F.; Fraundorfer, F. Deep Learning in Remote Sensing: A Comprehensive Review and List of Resources. *IEEE Geosci. Remote Sens. Mag.* **2017**, *5*, 8–36. [CrossRef]
20. Turner, R.E.; Ebelhar, M.W.; Wilkerson, T.; Bellaloui, N.; Golden, B.R.; Irby, J.T.; Martin, S. Effects of Purple Seed Stain on Seed Quality and Composition in Soybean. *Plants* **2020**, *9*, 993. [CrossRef] [PubMed]
21. Sagan, V.; Maimaitiyiming, M.; Fishman, J. Effects of Ambient Ozone on Soybean Biophysical Variables and Mineral Nutrient Accumulation. *Remote Sens.* **2018**, *10*, 562. [CrossRef]
22. Aires, F.; Boucher, E.; Pellet, V. Convolutional neural networks for satellite remote sensing at coarse resolution. Application for the SST retrieval using IASI. *Remote Sens. Environ.* **2021**, *263*, 112553. [CrossRef]
23. Song, Y.; Kalacska, M.; Gašparović, M.; Yao, J.; Najibi, N. Advances in geocomputation and geospatial artificial intelligence (GeoAI) for mapping. *Int. J. Appl. Earth Obs. Geoinf.* **2023**, *120*, 103300. [CrossRef]
24. Maimaitijiang, M.; Sagan, V.; Sidike, P.; Hartling, S.; Esposito, F.; Fritschi, F.B. Soybean yield prediction from UAV using multimodal data fusion and deep learning. *Remote Sens. Environ.* **2020**, *237*, 111599. [CrossRef]

25. Dilmurat, K.; Sagan, V.; Maimaitijiang, M.; Moose, S.; Fritschi, F.B. Estimating Crop Seed Composition Using Machine Learning from Multisensory UAV Data. *Remote Sens.* **2022**, *14*, 4786. [CrossRef]
26. Roy, B.; Sagan, V.; Alifu, H.; Saxton, J.; Ghoreishi, D.; Shakoor, N. Soil Carbon Estimation From Hyperspectral Imagery With Wavelet Decomposition And Frame Theory. *IEEE Trans. Geosci. Remote Sens.* **2024**, *62*, 4511312. [CrossRef]
27. Xu, R.; Hu, W.; Zhou, Y.; Zhang, X.; Xu, S.; Guo, Q.; Qi, P.; Chen, L.; Yang, X.; Zhang, F.; et al. Use of near-infrared spectroscopy for the rapid evaluation of soybean [*Glycine* max (L.) Merri.] water soluble protein content. *Spectrochim. Acta Part A Mol. Biomol. Spectrosc.* **2020**, *224*, 117400. [CrossRef] [PubMed]
28. Li, X.; Zhang, L.; Zhang, Y.; Wang, D.; Wang, X.; Yu, L.; Zhang, W.; Li, P. Review of NIR spectroscopy methods for nondestructive quality analysis of oilseeds and edible oils. *Trends Food Sci. Technol.* **2020**, *101*, 172–181. [CrossRef]
29. Middelbos, I.S.; Fahey, G.C. 9—Soybean Carbohydrates. In *Soybeans*; Johnson, L.A., White, P.J., Galloway, R., Eds.; AOCS Press: Champaign, IL, USA, 2008; pp. 269–296.
30. Karr-Lilienthal, L.K.; Kadzere, C.T.; Grieshop, C.M.; Fahey, G.C. Chemical and nutritional properties of soybean carbohydrates as related to nonruminants: A review. *Livest. Prod. Sci.* **2005**, *97*, 1–12. [CrossRef]
31. Jiang, S.; Jiang, W.; Wang, L. Unmanned Aerial Vehicle-Based Photogrammetric 3D Mapping: A Survey of Techniques, Applications, and Challenges. *IEEE Geosci. Remote Sens. Mag.* **2021**, *10*, 135–171. [CrossRef]
32. Wang, L.; Xu, Y.; Li, Y.; Zhao, Y. Voxel segmentation-based 3D building detection algorithm for airborne LIDAR data. *PLoS ONE* **2018**, *13*, e0208996. [CrossRef] [PubMed]
33. Han, X.-F.; Jin, J.S.; Wang, M.-J.; Jiang, W.; Gao, L.; Xiao, L. A review of algorithms for filtering the 3D point cloud. *Signal Process. Image Commun.* **2017**, *57*, 103–112. [CrossRef]
34. Zeybek, M.; Şanlıoğlu, İ. Point cloud filtering on UAV based point cloud. *Measurement* **2019**, *133*, 99–111. [CrossRef]
35. Chai, J.; Zeng, H.; Li, A.; Ngai, E.W.T. Deep learning in computer vision: A critical review of emerging techniques and application scenarios. *Mach. Learn. Appl.* **2021**, *6*, 100134. [CrossRef]
36. Voulodimos, A.; Doulamis, N.; Doulamis, A.; Protopapadakis, E. Deep Learning for Computer Vision: A Brief Review. *Comput. Intell. Neurosci.* **2018**, *2018*, 7068349. [CrossRef] [PubMed]
37. Zeiler, M.D.; Ranzato, M.; Monga, R.; Mao, M.; Yang, K.; Le, Q.V.; Nguyen, P.; Senior, A.; Vanhoucke, V.; Dean, J.; et al. On rectified linear units for speech processing. In Proceedings of the 2013 IEEE International Conference on Acoustics, Speech and Signal Processing, Vancouver, BC, Canada, 26–31 May 2013; pp. 3517–3521. [CrossRef]
38. Choy, C.; Gwak, J.; Savarese, S. 4D Spatio-Temporal ConvNets: Minkowski Convolutional Neural Networks. In Proceedings of the 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Long Beach, CA, USA, 15–20 June 2019; pp. 3070–3079. [CrossRef]
39. Hancock, J.T.; Khoshgoftaar, T.M. Survey on categorical data for neural networks. *J. Big Data* **2020**, *7*, 28. [CrossRef]
40. Heusel, M.; Ramsauer, H.; Unterthiner, T.; Nessler, B.; Hochreiter, S. GANs Trained by a Two Time-Scale Update Rule Converge to a Local Nash Equilibrium. *arXiv* **2018**, arXiv:1706.08500.
41. Kingma, D.P.; Ba, J. Adam: A Method for Stochastic Optimization. *arXiv* **2017**, arXiv:1412.6980.
42. Basuchaudhuri, P. *Physiology of Soybean Plant*, 1st ed.; CRC Press: Boca Raton, FL, USA, 2020.
43. Wijewardana, C.; Reddy, K.R.; Bellaloui, N. Soybean seed physiology, quality, and chemical composition under soil moisture stress. *Food Chem.* **2019**, *278*, 92–100. [CrossRef] [PubMed]
44. Ma, L.; Liu, Y.; Zhang, X.; Ye, Y.; Yin, G.; Johnson, B.A. Deep learning in remote sensing applications: A meta-analysis and review. *ISPRS J. Photogramm. Remote Sens.* **2019**, *152*, 166–177. [CrossRef]
45. Dornbos, D.L., Jr.; McDonald, M.B., Jr. Mass and Composition of Developing Soybean Seeds at Five Reproductive Growth Stages. *Crop Sci.* **1986**, *26*, 624–630. [CrossRef]
46. Jiang, H.; Egli, D.B. Soybean Seed Number and Crop Growth Rate during Flowering. *Agron. J.* **1995**, *87*, 264–267. [CrossRef]
47. Du, Y.; Zhao, Q.; Chen, L.; Yao, X.; Zhang, H.; Wu, J.; Xie, F. Effect of Drought Stress during Soybean R2–R6 Growth Stages on Sucrose Metabolism in Leaf and Seed. *Int. J. Mol. Sci.* **2020**, *21*, 618. [CrossRef] [PubMed]
48. Gibson, L.R.; Mullen, R.E. Influence of Day and Night Temperature on Soybean Seed Yield. *Crop Sci.* **1996**, *36*, 98–104. [CrossRef]
49. Su, J.; Zhu, X.; Li, S.; Chen, W.-H. AI meets UAVs: A survey on AI empowered UAV perception systems for precision agriculture. *Neurocomputing* **2023**, *518*, 242–270. [CrossRef]
50. Triboi, E.; Triboi-Blondel, A.-M. Productivity and grain or seed composition: A new approach to an old problem—Invited paper. *Eur. J. Agron.* **2002**, *16*, 163–186. [CrossRef]
51. Elarab, M.; Ticlavilca, A.M.; Torres-Rua, A.F.; Maslova, I.; McKee, M. Estimating chlorophyll with thermal and broadband multispectral high resolution imagery from an unmanned aerial system using relevance vector machines for precision agriculture. *Int. J. Appl. Earth Obs. Geoinf.* **2015**, *43*, 32–42. [CrossRef]
52. Wang, Y.-H.; Su, W.-H. Convolutional Neural Networks in Computer Vision for Grain Crop Phenotyping: A Review. *Agronomy* **2022**, *12*, 2659. [CrossRef]

https://doi.org/10.3390/rs18183121

53. Sharkey, T.D. Effects of moderate heat stress on photosynthesis: Importance of thylakoid reactions, rubisco deactivation, reactive oxygen species, and thermotolerance provided by isoprene. *Plant Cell Environ.* **2005**, *28*, 269–277. [CrossRef]
54. Clemente, T.E.; Cahoon, E.B. Soybean Oil: Genetic Approaches for Modification of Functionality and Total Content. *Plant Physiol.* **2009**, *151*, 1030–1040. [CrossRef] [PubMed]