Explainable machine learning for wheat above-ground biomass estimation under spatial cross-validation: sensor contributions and exploratory harvest forecasting

(This paper is a non-peer reviewed preprint submitted to EarthArXiv)

Francisco Zambrano[a], Abel Herrera[b], Mauricio Molina-Rocco[c]

[a] *Facultad de Medicina Veterinaria y Agronomía, Universidad de Las Américas, Santiago, Chile,*

[b] *Hemera Centro de Observación de la Tierra, Facultad de Ciencias, Ingeniería y Tecnología; Universidad Mayor, Santiago, Chile,*

[c] *Laboratorio de Suelos y Sostenibilidad. Departamento de Acuicultura y Recursos Agroalimentarios, Universidad de Los Lagos, Osorno, 5311157, Chile,*

Abstract

Global food security faces increasing challenges under climate change, making accurate monitoring of wheat (*Triticum aestivum*) critical. This study presents an explainable machine learning (ML) framework to estimate in-season wheat above-ground biomass (AGB) and to explore harvest forecasting one to four months ahead, integrating in-situ weather and soil moisture data with Sentinel-1, Sentinel-2, and PlanetScope imagery across four Mediterranean wheat field-seasons (2020 to 2023) in central Chile. We evaluated 41 models, comprising 40 recipe-algorithm combinations spanning eight sensor-combination recipes and five ML algorithms (Random Forest [RF], XGBoost, GLMnet, bagMLP, and KNN), plus a stacked ensemble, and obtained a preliminary, exploratory assessment of their spatial transferability with a leave-one-site-out cross-validation (LOSO-CV) scheme across the four site-seasons, with model transparency assessed via the DALEX framework.

For in-season estimation (stage 1), the best LOSO-CV configuration (RF using only PlanetScope predictors, *rec7*) reached R^2=0.78 (RMSE=7.87 t/ha, MAE=5.90 t/ha). A sensor-ablation analysis showed that Sentinel-2 reduced RMSE by 23.4% relative to a weather-only baseline, Sentinel-1 added essentially no further skill (+0.5%), and PlanetScope contributed a further 5.7% reduction. DALEX analysis of the stacked ensemble identified a Sentinel-2 SWIR-related chlorophyll index ($S2_{SWIR12}$-MCARI), two PlanetScope band/index predictors, a cumulative Sentinel-2 red-edge index, and the Sentinel-1 VH/VV backscatter ratio as the leading, comparably important predictors, spanning optical, SAR, and accumulated-thermal-time variables. As a secondary, exploratory objective, the same LOSO-CV framework was applied to a self-consistency analysis of

harvest-AGB forecasting one to four months ahead, using the stage 1 model's own spatial predictions as the target in the absence of independent harvest-date ground truth; self-consistency (RF, all predictors) was modest and decayed sharply with lead time (R^2 from 0.46 at one month to 0.10 at four months; RMSE from 4.5 to 6.5 t/ha), underscoring the need for both independent harvest-date measurements and larger multi-site networks to properly evaluate multi-month forecasting. This multi-sensor, cloud-resilient framework provides a promising, explainable basis for in-season AGB monitoring in Mediterranean wheat systems, although the four-site-season LOSO-CV evaluation should be read as a preliminary sensitivity analysis of spatial transferability rather than definitive proof of generalizability (see Discussion).

Keywords: above-ground biomass, Sentinel-1, Sentinel-2, PlanetScope, soil moisture, machine learning, precision agriculture

1. Introduction

As the world population is projected to reach 9.8 billion by 2050 (UN, 2024), food security is expected to face significant challenges (Bahar et al., 2020), which are further intensified by climate change (Chen, 2025). Wheat (*Triticum aestivum*) is the most consumed crop worldwide, and its production is key for food security (FAO, 2025). Furthermore, droughts significantly impact agriculture, as evidenced by the decrease in crop yield (Mishra et al., 2015; Naumann et al., 2021; Santini et al., 2022). In fact, global wheat production could drop by 13% by mid-century under climate change (Pequeno et al., 2024). Quantifying above-ground biomass (AGB) is of utmost importance since it allows us to monitor and optimize agricultural processes (Dietz et al., 2021). Proxies of vegetation productivity derived from vegetation indices can help monitor and predict the impact of drought at a regional scale (Zambrano, 2021; Zambrano et al., 2016, 2018). However, on a farm level, we must quantify this impact using crop yield as a productivity metric, which can aid in determining the economic impact of seasonal variations in crop yield. Generally, the estimation and monitoring of agricultural production is made by using three methods: i) crop growth models (Lobell and Burke, 2010; Wang et al., 2023), ii) surveys (Gardner et al., 1980), and iii) in-situ measurements. The models help predict how crops will grow, yield, and produce under different environmental and management conditions, which is useful for forecasting food supply and understanding how climate change and policy decisions affect global agriculture, but the downside is that they require many different variables to obtain meaningful results. The objective of the surveys is to collect timely and reliable data from the agricultural

*Corresponding author

sector in order to carry out statistical studies of crops, but it provides summaries at administrative units (e.g., states, districts) (Gardner et al., 1980). Additionally, field measurements provide reliable data, known as the ground truth, that can be used for comparison with models and surveys; however, this process is costly and time-consuming.

The use of empirical models that are based solely on remote sensing sources has become a straightforward option for the estimation of crop yield, especially for cereals (Lischeid et al., 2022). There are multiple remote sensing methods to estimate AGB, among which the methods based on multispectral imaging (Marshall et al., 2022; Segarra et al., 2022), Synthetic Aperture Radar (SAR) (Chao et al., 2019; Hu et al., 2024), LiDAR (Light Detection and Ranging) (Bates et al., 2021), hyperspectral (Yue et al., 2021), or combination of them (David et al., 2022; Li et al., 2024; Wang et al., 2016), stand out, offering promising results. Estimation methods based on SAR data focus on the measurement of the structure of the crop, soil moisture, and their temporal variations. An advantage of this technique is that it can penetrate the cloud cover, allowing a constant measurement in the study area; its disadvantage is the angle of observation to the crop and the speckle noise inherent to the signal (Lopes et al., 1990; Maghsoudi et al., 2012; Moran et al., 1997). Estimation methods based on LiDAR and hyperspectral data focus on the measurement of the crop structure and phenological monitoring of the crop (Brovkina et al., 2017; Luo et al., 2017; Wang et al., 2016). Its main disadvantage is the high cost of implementation, which limits its proliferation. Another common approach is the use of spectral indices (e.g., NDVI, EVI, etc.) derived from optical sensors such as Sentinel-2 (David et al., 2022) for specific dates. Furthermore, the use of accumulated spectral indices (Peroni Venancio et al., 2020) over the growing season has shown improvements in estimating biomass/yield compared with using indices just for specific dates. The main advantage is its simplicity and strong correlation with the observed biomass. Currently, the research about estimating biomass is moving forward integrating Sentinel 1 & 2 (Uribeetxebarria et al., 2023), LiDAR, hyperspectral, and Landsat missions (He et al., 2018; Zhang et al., 2023), which aids in better monitoring of the crop's growth stages (Shen et al., 2023).

Effective machine learning (ML) methods, along with remote sensing data, are increasingly used to estimate the AGB in various cereals, including wheat. In Southern Brazil, (Atkinson Amorim et al., 2022) collected images from unmanned aerial vehicles (UAV) and used random forest, support vector regressions, and neural networks to estimate AGB in wheat. Their results reached a R^2 of 0.90 and a root mean square error (RMSE) of 0.83 t/ha. In another study, (Liu et al., 2024) used spectral data obtained from UAV, phenological information, and multiple machine learning algorithms to estimate AGB on wheat, having results of an R^2

ranging from 0.84 to 0.91 and an RMSE of 1.69 to 2.11 t/ha. Considering another approach, (Marshall et al., 2022) used hyperspectral and Sentinel-2 data and applied random forest and partial least squares to estimate wheat AGB on different stages of the crop. Their results achieve an R^2=0.81 and RMSE=3.7 t/ha with the random forest model. Further, for different types of grass cover, Sentinel-1 has been tested to estimate AGB with varied results (David et al., 2022; Li et al., 2024; Nuthammachot et al., 2022). The benefits of SAR Sentinel-1 are that it can gather dependable data even on cloudy days, which is not possible with optical sensors like Sentinel-2 and Landsat. Additionally, Sentinel-1's backscatter is closely linked to the dielectric constant, making it a proxy of soil moisture (Fan et al., 2025; Zhu et al., 2024), and, therefore, the amount of water crops can absorb from the soil, which affects their growth. Thus, it is worthwhile to explore its usefulness for AGB prediction.

Here, we aim to develop and evaluate an explainable ML framework for estimating wheat AGB during the growing season from multi-sensor remote sensing (Sentinel-1, Sentinel-2, and PlanetScope) and in-situ weather and soil moisture data, and to obtain a preliminary, exploratory assessment of its spatial transferability under a leave-one-site-out cross-validation (LOSO-CV) scheme across four Mediterranean wheat field-seasons. As a secondary, exploratory objective, we evaluate whether the same predictors and modeling framework can reproduce, from earlier-season covariates, the framework's own spatial AGB predictions at harvest one to four months in advance. We defined three specific objectives: i) derive predictors from weather, soil moisture, Sentinel-1, Sentinel-2, and PlanetScope data; ii) apply ML algorithms together with explainable-AI techniques to estimate the spatiotemporal variation of wheat AGB through the growing season and identify which sensors and variables drive model performance; and iii) explore the feasibility of forecasting AGB at harvest at four lead times (one to four months before harvest) using the same modeling and LOSO-CV validation framework.

2. Materials and Methods

2.1. Study Area

The study was conducted during the 2020 to 2023 growing seasons (June to January) across four wheat field sites (*Triticum aestivum*) located in central Chile: i) La Cancha in the Metropolitan region, ii) Hidango in the O'Higgins region (two different fields in two seasons), and iii) Villa Baviera in the Maule region. Sampling took place in delineated fields at each site (Figure 1). In Hidango, separate fields were established for the 2021-2022 and 2022-2023 seasons; in Villa Baviera and La Cancha, sampling occurred during the 2020-2021

and 2022-2023 seasons, respectively. To capture spatial variability and ensure representativeness, sampling points were randomly selected and evenly distributed: five in La Cancha, Villa Baviera, and the northern Hidango field (2021 to 2022 season); and six in the southern Hidango field (2022 to 2023 season) (Figure 1).

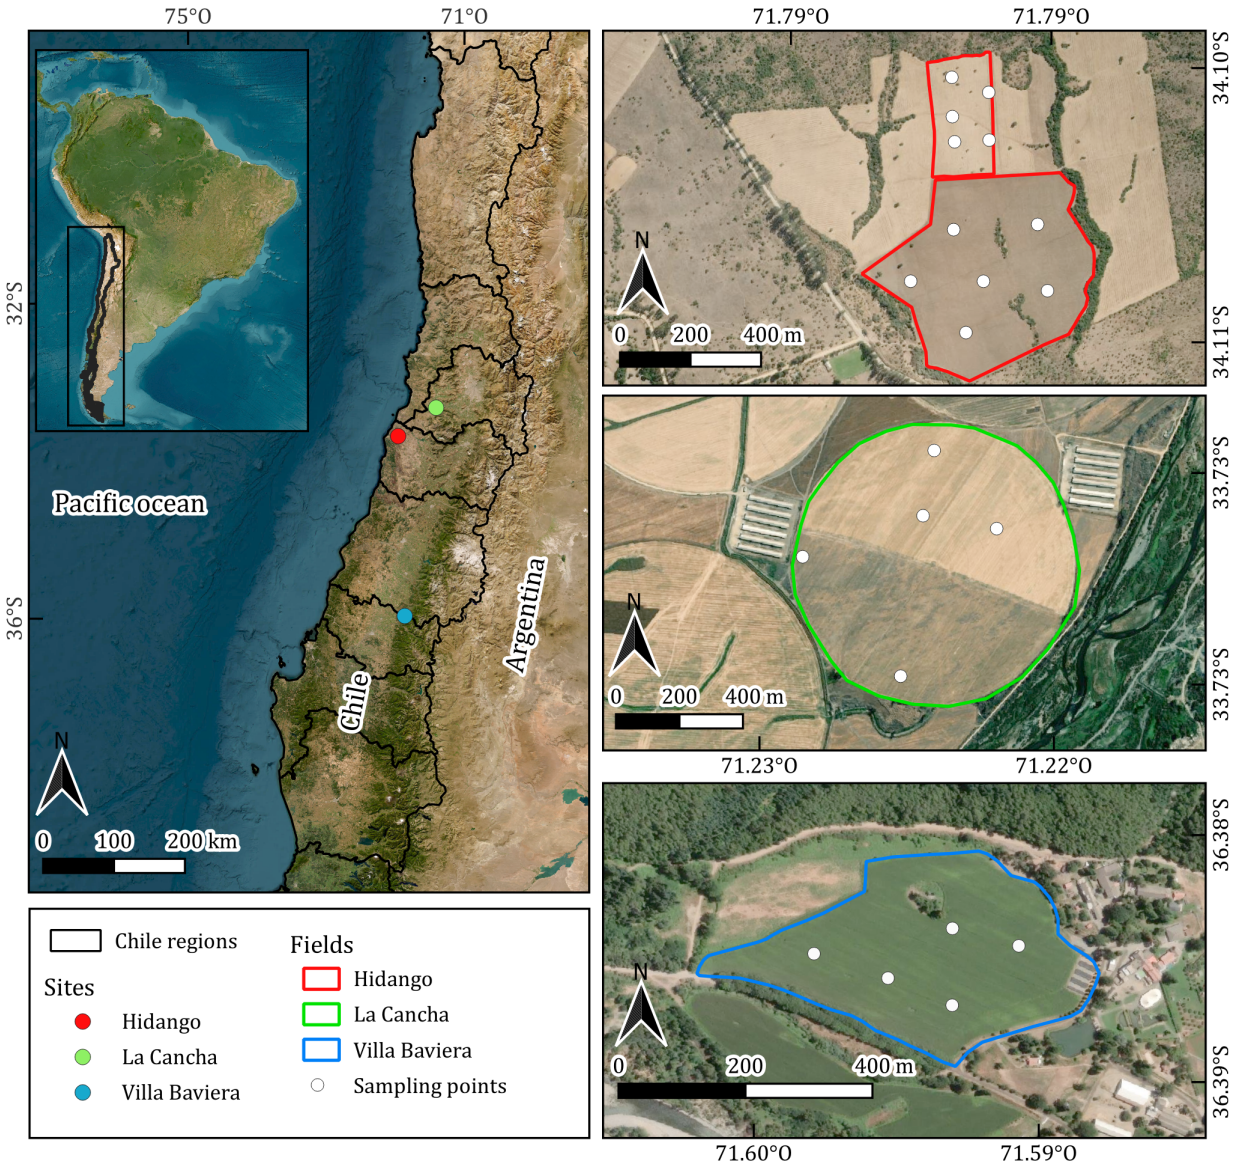

Figure 1: **Study Area.** The map on the left panel shows the location of the sites in Central Chile. The remaining panels illustrate the field boundaries and sampling points for each site. In Hidango, the northern field is from the 2021 to 2022 season, and the southern one from 2022 to 2023.

A variety of spring wheat was planted in Villa Baviera (Llareta-INIA; (Garrido et al., 2013)), while the other sites had winter wheat (Irafen-INIA and Pantera-INIA; (Del Pozo et al., 2023)). Spring wheat is typically sown in late winter or early spring and harvested within the same year, whereas winter wheat is sown in

autumn, requires vernalization, and is harvested the following summer (Liu and Chen, 2025; Zhao et al., 2022). According to the FAO–UNESCO Soil Map of the World (FAO, 2014); La Cancha is characterized by shallow, fine-textured Regosols on nearly flat terrain; Hidango presents fine-textured Luvisols on gently undulating landscapes; and Villa Baviera features deep, well-structured Nitisols on slightly sloped land. The climate in these regions is Mediterranean with winter rainfall (Csb) under the Köppen–Geiger classification (Beck et al., 2023); Hidango, however, exhibits an oceanic influence. Mean annual precipitation ranges from 700 to 900 mm, with average temperatures between 11 and 12 °C.

2.2. Data acquisition and processing

2.2.1. Field data

AGB, defined as the total dry mass of dead and living plant organs per unit ground area (Chave et al., 2014), was sampled at distinct wheat phenological stages across each season. Phenological stages were identified using the Zadoks scale, which describes 10 principal growth stages from sowing to ripening (Zadoks et al., 1974). At each sampling point (Figure 1), biomass from a 0.25 m^2 area was harvested, dried in a forced-air oven at 70°C for three days, weighed on a precision scale, and converted to t/ha. Sampling intervals were planned to capture variability across growth stages, ensuring representative data for each field and season. The temporal distribution of sampling dates across phenological stages is illustrated in Figure 2, with NDVI time series indicating canopy development over the season. A total of 153 AGB wheat samples were collected: 88 in Hidango (40 in 2021 to 2022 and 48 in 2022-2023), 35 in La Cancha, and 30 in Villa Baviera.

Additionally, soil moisture was measured using TEROS 12 sensors (METER Group, USA), which operate via a capacitance/frequency domain method to assess dielectric permittivity, achieving ±1% accuracy in typical mineral soils (Fragkos et al., 2024). The sensors were installed at a central point within each field's sampling area at depths of 15, 30, and 45 cm in La Cancha and Villa Baviera, and 15, 50, and 75 cm in Hidango in both seasons. Each sensor recorded data at 30-minute intervals, and daily values were computed by averaging the measurements from each day.

Soil water content (cm^3/cm^3) measured by the sensors was converted to water depth (mm) within the 0 to 55 cm using the specific depth of the profile at which the sensor was installed.

2.2.2. Weather Data

Precipitation and temperature data were obtained from automatic meteorological stations managed by the Red Agrometeorológica de Chile. This national network provides continuous, open-access weather data

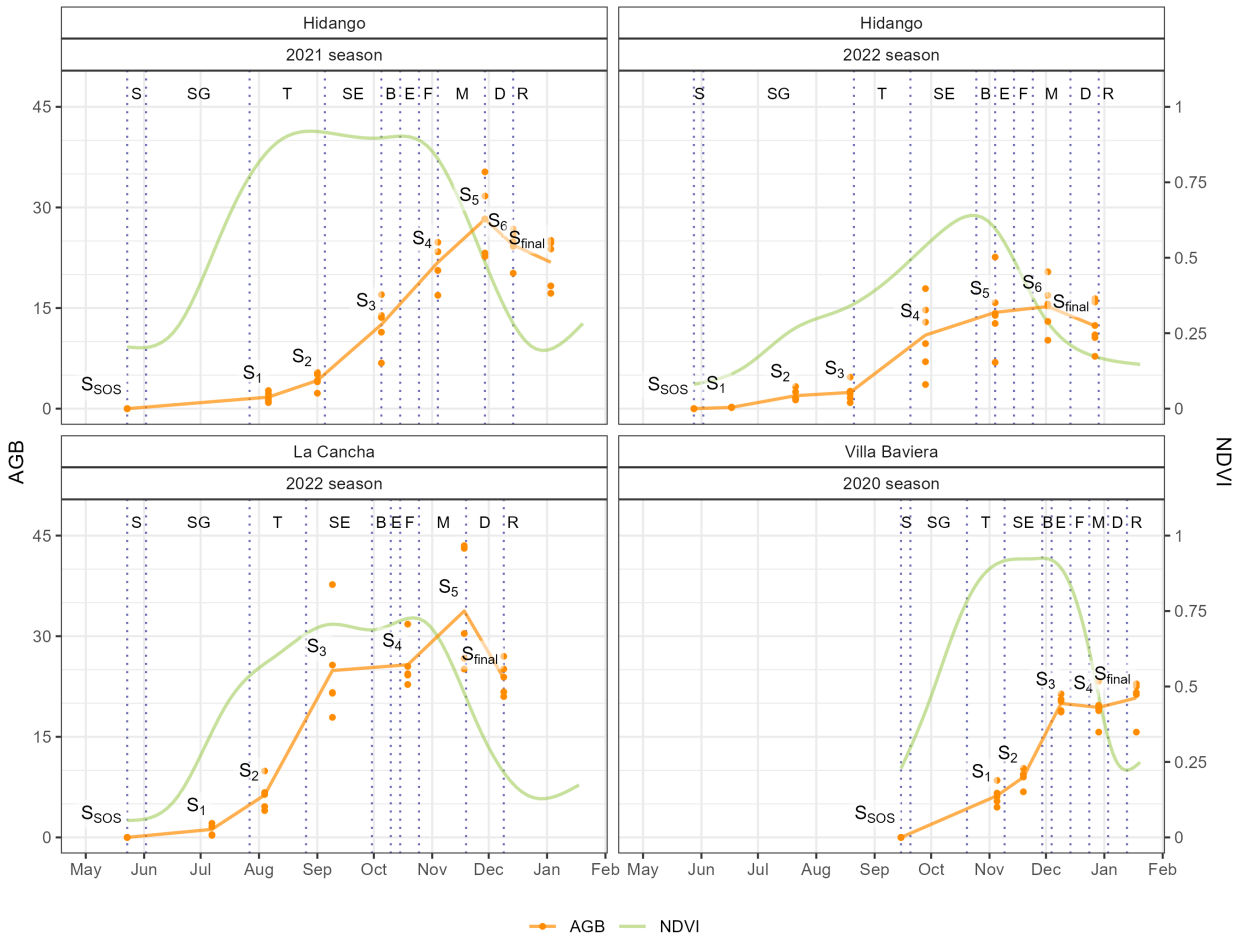

Figure 2: **Above-ground biomass (AGB) sampling dates across the seasons and fields.** Zadoks scale, with NDVI included to illustrate changes in canopy vigor throughout the seasons. S = Sowing, SG = Seeding Growth, T = Tillering, SE = Stem Elongation, B = Booting, E = Ear Emergence, F = Flowering, M = Milk Development, D = Dough Development, and R = Ripening. AGB sampling dates are labeled S_1, S_2, ..., S_n (subscripted), where n is the total number of samples collected during the season; these are distinct from the unsubscripted "S" (Sowing) Zadoks stage label above. S_{SOS} indicates the start of the season (SOS) when AGB = 0, and S_{final} represents AGB at harvest.

Table 1: Meteorological stations used for each field, including the station code and name, the operating institution (Inst.), its coordinates, and its distance from the corresponding field (Dist.).

Field	**Code**	**Name**	**Inst.**	**Lat (°S)**	**Lon (°W)**	**Dist. (km)**
Hidango	260	Hidango	INIA	34.10	71.80	1.03
La Cancha	50	Chocalan	FDF	33.73	71.21	1.16
Villa Baviera	103	Parral Norte	FDF	36.23	71.73	20.89

Inst.: institution operating the station; INIA: Instituto de Investigaciones Agropecuarias; FDF: Fundación para el Desarrollo Frutícola. Dist.: distance between the station and the corresponding field.

for agricultural use (http://agromet.inia.cl; http://agrometeorologia.cl), recording several meteorological variables at hourly intervals. Each station is equipped with a Texas Electronics TE525MM-L25 tipping-bucket rain gauge for precipitation and a Vaisala HMP60-L11 sensor for air temperature and relative humidity (Chacón C., 2019). For each study site, the nearest meteorological station was selected (Table 1), and hourly precipitation and air temperature data were downloaded starting from the sowing date.

Daily precipitation (mm) was cumulatively summed over the growing season to obtain daily accumulated precipitation (mm, PP). Daily minimum (Tmin) and maximum (Tmax) temperatures were used to compute Growing Degree Days (GDD):

$$GDD = \frac{(T_{\max} + T_{\min})}{2} - T_{base}$$

A base temperature of 2°C was applied for the early growth stages (sowing to tillering), while 6°C was used for reproductive stages (stem elongation to ripening) (Del Pozo et al., 1987). Negative GDD values were set to zero. Then, the accumulated GDD values were calculated by summing the daily GDD from the start of the season of each field (GDD).

2.2.3. Remote Sensing Data

Satellite time series imagery from three missions, Sentinel-1, Sentinel-2, and PlanetScope, were used to derive predictors for aboveground biomass (AGB) modeling. These datasets include radar and optical sensors, covering different spectral ranges (SAR, visible, NIR, SWIR) and having different spatial resolutions.

The Sentinel-1 (S1) mission, operated by the European Space Agency (ESA), consists of a constellation of C-band synthetic aperture radar (SAR) satellites that operate in several acquisition modes (Interferometric Wide [IW], Extra Wide [EW], and Stripmap [SM]), with spatial resolutions ranging from 5 to 40 m and a swath width up to 410 km. Data acquisition is independent of cloud cover or sunlight, enabling consistent

time-series generation under any weather condition. We used Level-1 Ground Range Detected (GRD) products from Sentinel-1A (S1A), the only satellite of the constellation operational during the study period, as Sentinel-1B stopped transmitting in 2022 and Sentinel-1C was not yet launched. GRD products are radiometrically calibrated, multi-looked, and terrain-corrected using a digital elevation model, then projected to the WGS84 datum (Table 2). These products are delivered in dual-polarization (VV, VH or HH, HV), depending on the acquisition mode and region, with a pixel spacing of 10 to 40 m (Torres et al., 2012).

Sentinel-2 (S2) is a dual-satellite optical mission carrying the MultiSpectral Instrument (MSI), which captures 13 spectral bands across the visible, near-infrared (NIR), and shortwave infrared (SWIR) regions. Spatial resolution varies by band (10, 20, or 60 m), and the system provides global coverage with a revisit time of ~5 days at the Equator. We used Level-2A surface reflectance products, which are atmospherically corrected (bottom-of-atmosphere, BOA) and orthorectified (Table 2). These products include cloud and shadow masks, aerosol optical thickness (AOT), and water vapor content layers, along with scene classification maps and pixel-level quality information (Drusch et al., 2012).

Table 2: Summary of satellite data used to derive predictors for AGB estimation. For Sentinel-2, paired wavelength values give the Sentinel-2A / Sentinel-2B central wavelengths, respectively; for Sentinel-1, the C-band wavelength is expressed in nm (5.546e7 nm ≈ 5.55 cm) for consistency with the optical sensors. Blank cells in the Revisit (d) column repeat the value given in the first row of each sensor block.

Dataset	**Collection**	**Bands**	**Wavelength (nm)**	**Res. (m)**	**Revisit (d)**
Sentinel-1	GRD	VV	5.546e7 (C-Band)	10	6
		VH	5.546e7 (C-Band)	10	
Sentinel-2	Level-2A	1 - Coastal	442.7 / 442.2	60	5
		2 - Blue	492.7 / 492.3	10	
		3 - Green	559.8 / 558.9	10	
		4 - Red	664.6 / 664.9	10	
		5 - RE1	704.1 / 703.8	20	
		6 - RE2	740.5 / 739.1	20	
		7 - RE3	782.8 / 779.7	20	
		8 - NIR	832.8 / 832.9	10	
		8A - RE	864.7 / 864.0	20	
		9 - Water vap.	945.1 / 943.2	60	
		10 - Cirrus	1373.5 / 1376.9	60	
		11 - SWIR1	1613.7 / 1610.4	20	
		12 - SWIR2	2202.4 / 2185.7	20	
PlanetScope	Level 3B SR	1 - C. Blue	441.5	3	1
		2 - Blue	490.0	3	
		3 - Green I	531.0	3	
		4 - Green	565.0	3	
		5 - Yellow	610.0	3	
		6 - Red	665.0	3	
		7 - Red Edge	705.0	3	
		8 - NIR	865.0	3	

PlanetScope (PS) is a commercial constellation of ~120 CubeSats operated by Planet Labs, which capture daily multispectral imagery at a nominal resolution of 3 m. The constellation includes different sensor types deployed over time, such as Dove-C (three- or four-band imagers with a split-frame NIR filter), Dove-R (four-band imagers with a butcher-block filter), and SuperDove (eight-band imagers including the visible, NIR, and Red Edge bands). We used Level-3B Surface Reflectance products, which are orthorectified and atmospherically corrected using Planet's in-house model based on global aerosol assumptions. Although not fully corrected for haze or thin cirrus clouds, the products are radiometrically calibrated, georeferenced to sub-pixel accuracy, and delivered with metadata and quality indicators (Roy et al., 2021). Due to the sensor transition that occurred during the study period, both Dove and SuperDove imagery were included. To ensure consistency across the dataset, only the four bands common to all sensor types (Blue, Green, Red, and NIR) were used.

S1 Ground Range Detected (GRD) imagery was accessed via Google Earth Engine; S2 Level-2A products were obtained from the Planetary Computer; and PS imagery was acquired through the Planet Explorer platform, excluding images with cloud cover over the fields. A total of 185 S1, 208 S2, and 180 PS images were used across the four fields in their respective seasons: Hidango 2021-2022 (S1: 52; S2: 60; PS: 47), Hidango 2022-2023 (S1: 56; S2: 56; PS: 40), La Cancha 2022-2023 (S1: 43; S2: 53; PS: 41), and Villa Baviera 2020-2021 (S1: 34; S2: 39; PS: 52).

2.2.4. Preprocessing of remote sensing data

For data preprocessing, Sentinel-1 backscatter values (VV and VH) were first converted from decibels to linear scale, and the VH/VV polarization ratio was then computed. No additional masking was applied to Sentinel-1, as its GRD products undergo standard preprocessing, including radiometric calibration, multi-looking, and orthorectification. For Sentinel-2 and PlanetScope, reflectance values were scaled by dividing by 10,000, and values outside the [0.01, 1] range were discarded. Low-quality Sentinel-2 pixels were identified using the Scene Classification Layer (SCL), excluding classes 0, 1, 2, 3, 8, 9, and 10. For PlanetScope, the Universal Data Mask (UDM) was used to retain only clear-sky pixels.

All variables were interpolated to create continuous daily time series, ensuring they aligned with AGB sampling dates and allowing for biomass estimation at consistent temporal offsets across different fields and seasons. For Sentinel-1, VV, VH, and VH/VV ratio were temporally extended by duplicating each acquisition across the days until the midpoint to the following observation. For Sentinel-2 and PlanetScope, spectral bands were interpolated using a Kalman smoothing approach (Moritz and Bartz-Beielstein, 2017), and

cumulative sums were computed from the start of each season. Daily values from all sources, and cumulative sums for Sentinel-2 and PlanetScope, were extracted at sampling locations and later used as input variables for AGB estimation. Extraction was performed at each sampling point's coordinates as the value of the raster cell containing the point (nearest-pixel extraction via the {terra} package, with no spatial averaging over a buffer); given the 0.25 m^2 plot size relative to the 3-60 m sensor resolutions, each sampling plot is fully contained within a single pixel for every sensor.

2.2.5. Selection and processing of VIs

A total of 53 vegetation indices (VIs) were selected to estimate AGB using the daily interpolated S2 and PS data (Table 3). Daily cumulative sums of these VIs were also computed from the start of each season to capture integrated vegetation dynamics. Furthermore, 36 of these 53 VIs originally derived from the B8 band were recalculated using the B8A band (available only for S2) to assess potential differences in vegetation response between these two near-infrared bands, resulting in a total of 89 VIs across both products. The indices were chosen based on their relevance for assessing vegetation status, biomass estimation, and yield prediction in wheat, forage, maize, and similar crops (Peng et al., 2023; Segarra et al., 2022; Uribeetxebarria et al., 2023; Wang et al., 2022; Wu et al., 2008; Zheng et al., 2017). Both daily values and cumulative VIs were used as predictor variables in the modeling framework.

Table 3 outlines the vegetation indices (VIs) chosen for the machine learning (ML) model after eliminating those that were highly correlated. To see all the other VIs that were used, please refer to the Supplementary Material Table S1.

2.3. Modeling with Machine Learning Algorithms

2.3.1. Overview of the modeling process

Figure 3 shows a general scheme for the modeling process. The dataset for modeling is composed of the in situ AGB data and the remote sensing predictors derived from S1, S2, and PlanetScope, as well as the weather variables. All this data is generated as a time series for each growing season in each site. Then a feature engineering process is applied to the data before running the models. Thereafter, the dataset is used in the first stage to run the ML models to estimate the spatio-temporal variation of AGB within the growing season. After achieving the initial results, we tested the ML prediction models of AGB at harvest by using covariates from four lead times: 1, 2, 3, and 4 months. The ML models for estimation and prediction pass through a hyperparameter optimization and an explainable ML process.

Table 3: Vegetation indices selected for AGB modeling after removing highly correlated variables (see Supplementary Material Table S1 for the full candidate set). The Dataset column indicates the source product(s) used to compute each index; an asterisk (*) marks indices for which the Sentinel-2 formula uses the B8A (red-edge) band in place of the standard NIR band B8 (Section "Selection and processing of VIs").

N°	Dataset	Name	Abbr.	Formula	Ref.
1	S2	OSAVI Variant	OSAVI_2	$\frac{1.16 \cdot (RE_3 - Red)}{(RE_3 + Red + 0.16)}$	Jin et al. (2020)
2	S2	-	MCARI/OSAVI_2	$MCARI/OSAVI_2$	Wu et al. (2008)
3	S2	SWIR 11 MCARI	MCARI_{SW11}	$((RE_1 - SW_1) - 0.2(RE_1 - G)) \cdot (RE_1/SW_1)$	Herrmann et al. (2011)
4	S2	SWIR 11 TCARI	TCARI_{SW11}	$3((RE_1 - SW_1) - 0.2(RE_1 - G) \cdot (RE_1/SW_1))$	Herrmann et al. (2011)
5	S2	SWIR 12 MCARI	MCARI_{SW12}	$((RE_1 - SW_2) - 0.2(RE_1 - G)) \cdot (RE_1/SW_2)$	Herrmann et al. (2011)
6	S2*	NDRE 3	NDRE3	$(NIR - RE_3)/(NIR + RE_3)$	Magney et al. (2017)
7	S2*	Water Index	WI	$NIR/Water\ vapour$	Jin et al. (2020)
8	S2-PS	MCARI	MCARI	$((RE_1 - R) - 0.2(RE_1 - G)) \cdot (RE_1/R)$	Daughtry (2000)
9	S2-PS	TCARI	TCARI	$3((RE_1 - R) - 0.2(RE_1 - G) \cdot (RE_1/R))$	Haboudane et al. (2002)
10	S2-PS	-	TCARI/OSAVI	$TCARI/OSAVI$	Haboudane et al. (2002)
11	S2-PS	-	MCARI/OSAVI	$MCARI/OSAVI$	Daughtry (2000)
12	S2*-PS	Simple Ratio	SR	NIR/R	Tucker (1979)
13	S2-PS	OSAVI	OSAVI	$\frac{1.16 \cdot (NIR - R)}{(NIR + R + 0.16)}$	Rondeaux et al. (1996)
14	S2-PS	CI-Red	CI_{red}	$(NIR/R) - 1$	Clevers and Gitelson (2013)
15	S2-PS	CVI	CVI	$(NIR/G) \cdot (R/G)$	Vincini et al. (2008)
16	S2-PS	EVI	EVI	$\frac{2.5(NIR - R)}{(NIR + 6R - 7.5B + 1)}$	Jiang et al. (2008)
17	S2-PS	NDVI	NDVI	$(NIR - R)/(NIR + R)$	Rouse et al. (1974)
18	S2-PS	GNDVI	GNDVI	$(NIR - G)/(NIR + G)$	Gitelson et al. (1996)
19	S2-PS	NDRE 1	NDRE1	$(NIR - RE_1)/(NIR + RE_1)$	Magney et al. (2017)
20	S2-PS	-	NDRE/NDVI	$NDRE_1/NDVI$	Raper and Varco (2015)
21	S2-PS	SIPI	SIPI	$(NIR - B)/(NIR - R)$	Peñuelas et al. (1994)

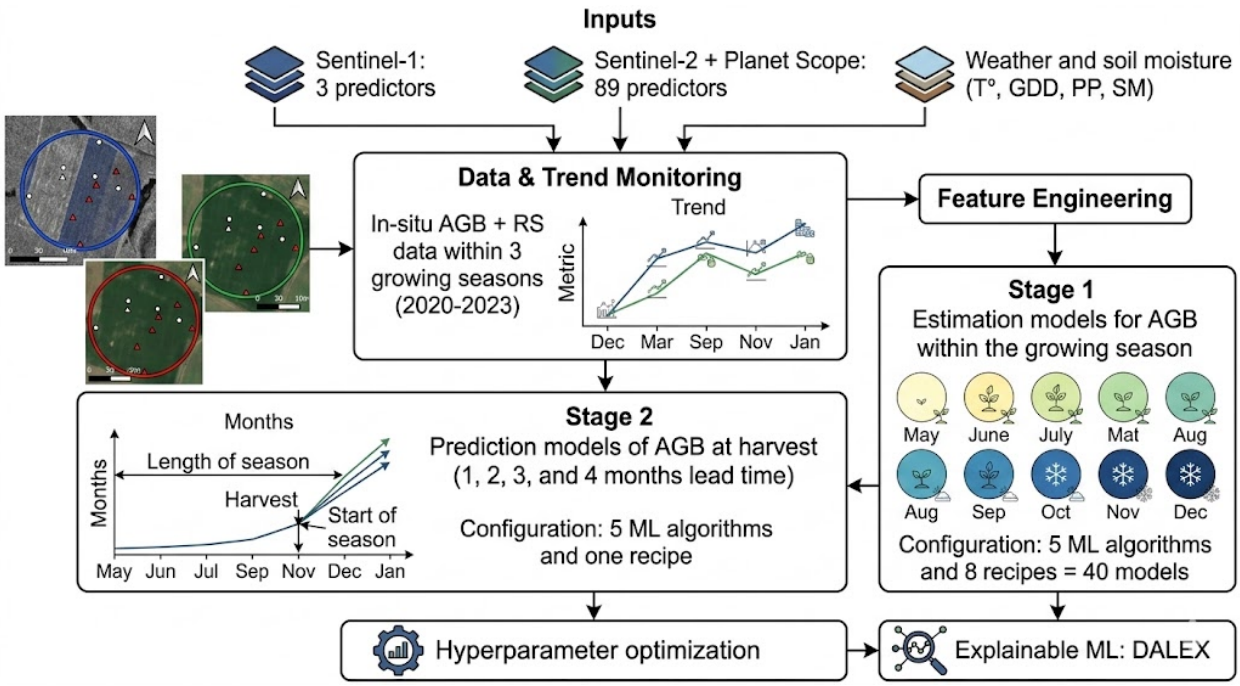

Figure 3: **Methodological scheme of the modeling process by machine learning of the AGB estimation within the growing season (stage 1) and the prediction at harvest at 1, 2, 3, and 4 month lead times (stage 2).** Input data (in-situ AGB and soil moisture, weather variables, and Sentinel-1, Sentinel-2, and PlanetScope predictors) are combined through a feature engineering step before being passed to the two modeling stages, both of which undergo hyperparameter optimization and explainable ML (DALEX) analysis.

We tested five machine learning (ML) algorithms to estimate spatio-temporal in-season and AGB (stage 1) and to predict AGB at harvest (stage 2): extreme gradient boosting (XGBoost; (Chen and Guestrin, 2016)), random forest (RF; (Ho, 1995)), regularized generalized linear regression (GLMnet; (Hastie et al., 2015)), single layer, feed-forward neural networks (bagMLP; (Breiman, 1996a)), and K-nearest neighbors (KNN; (Hechenbichler and Schliep, 2004)). We also tested an ensemble model which combines the predictions for the five models (Breiman, 1996b; Wolpert, 1992). These models were selected based on their robustness in regression tasks, capacity to handle high-dimensional data, interpretability, and for being state-of-the-art.

2.3.2. Defining dataset recipes for estimation and prediction models

The dataset used for the estimation modeling (stage 1) included covariates derived from soil moisture data, weather variables, remote sensing spectral information, and VIs. The response variable was AGB (t/ha), obtained from field measurements. We split the predictor variables into different subsets with the aim of evaluating independently how weather, Sentinel-1, Sentinel-2, and PlanetScope contribute to the estimation of AGB. We called each subset a recipe (rec). Thus, *rec1* used all the variables; *rec2* uses only Sentinel-1; *rec3* uses Sentinel-1 and weather; *rec4* uses only weather variables; *rec5* uses only Sentinel-2; *rec6* uses Sentinel-2 and weather; *rec7* uses only PlanetScope; and *rec8* uses PlanetScope and weather.

2.3.3. Feature engineering

Before running the ML modeling for estimation and prediction, we applied different preprocessing methods to the dataset. First, to fill missing data, we imputed it using k-nearest neighbor (Gower, 1971). All of the predictors were then normalized to have a standard deviation of one and a mean of zero. Also, we applied a filter to eliminate highly correlated variables (Pearson |r| 0.9), which removed one of the 43 candidate predictors ($S2_{SWIR11}$_MCARI) for *rec1* in the stage 1 (estimation) dataset, leaving 42 predictors. This correlation-based filter did not remove any variable from the stage 2 (prediction) dataset (Section 2.3.5), so all 43 candidate predictors were retained for *rec1* in stage 2. To document the residual multicollinearity in the full predictor set, we additionally computed variance inflation factors (VIF) on the 42 stage 1 predictors and iteratively removed variables with VIF $>$ 10, which reduced the set to 22 variables. This VIF audit was performed for transparency only and was not applied as a modeling step, since the regularized (GLMnet) and tree-based (RF, XGBoost) algorithms used here are tolerant of collinearity among spectral indices; the full audit is provided in the project repository. While collinearity among spectral indices does not bias point predictions from regularized or tree-based models, it does affect permutation-based variable importance (Section 2.3.8): permuting one of a set of mutually correlated predictors while holding the

others fixed generates combinations of values that are not jointly observed in the training data, which can distort the resulting importance and partial-dependence estimates. The importance results presented below (Section 3.1.2) should therefore be interpreted as indicating groups of mutually-informative, partly redundant predictors rather than as precisely localized rankings of individual variables; the VIF-reduced 22-variable set identified here provides a natural starting point for a future reanalysis using conditional permutation importance or VIF-reduced predictor sets (see Discussion, Study limitations).

2.3.4. Estimation models for AGB within the growing season

For the estimation model (stage 1), covariates from the same day as the AGB measured in situ were employed. Thus, we generate a modeling dataset by including the value of the covariates for each sample point. We used a combination of 40 models, corresponding to the eight recipes multiplied by the five ML algorithms. We then ran the models, tuned them through hyperparameter optimization (see section 2.3.6), and used them to estimate AGB spatially for each day and site during the growing season. This process yields a daily spatiotemporal dataset of AGB for each site.

2.3.5. Prediction models for AGB at harvest

From the estimation modeling (stage 1), we obtained the daily spatio-temporal AGB within the growing season. For prediction, we took the AGB of the harvest date per site as the response variable; this harvest-date AGB corresponds to the stage 1 model's spatial prediction (Section 2.3.4) for each site-season, rather than an additional in-situ destructive sample collected specifically at harvest. Consequently, the stage 2 models are trained and evaluated against a model-derived quantity rather than an independent ground-truth measurement: they test whether remote-sensing and weather covariates available one to four months before harvest are sufficient to reproduce the stage 1 model's own spatial output at the harvest date, not whether they can forecast actual, field-measured AGB at harvest. We therefore treat stage 2 throughout this study as an internal-consistency, proof-of-concept analysis of the modeling pipeline rather than as an independent validation of harvest-forecasting skill against ground truth; establishing the latter would require a dedicated destructive-sampling campaign at harvest, which was outside the scope of the present field campaigns (see Discussion, Study limitations). As covariates, we selected the remote sensing covariates (S1, S2, and PS variables in Table 3) and weather data, which lagged from one to four months before harvesting as lead time. To generate the dataset for prediction modeling, we took 100 random samples per site for the responses and covariates. These 100 samples per site-season capture within-field spatial variation in the modeled harvest-date AGB and the corresponding predictor surfaces around a single harvest date; because

the leave-one-site-out cross-validation scheme (Section 2.3.6) holds out all samples belonging to a given site-season as one block, this within-site resampling does not leak information across folds. However, because these 100 points are drawn from a smoothed, spatially-autocorrelated stage 1 prediction surface rather than from independent biological replicates, they are not 100 independent observations: the resulting dataset (400 points across the four site-seasons) carries substantial redundancy, so its nominal size should be regarded as an upper bound on the effective sample size, and the stage 2 performance metrics reported below (Section 3.2) may understate the uncertainty that would be obtained with independent harvest-date samples. In the modeling process, we used hyperparameter optimization to enhance the performance of our models and then applied them to predict aboveground biomass (AGB) at harvest for each lead time.

2.3.6. Hyperparameter optimization and evaluation

To tune each parameter per algorithm for the estimation and prediction modeling, we used hyperparameter optimization, defining ten random candidates per parameter combination (Table 4). For the estimation models (stage 1), performance was evaluated using a leave-one-site-out cross-validation (LOSO-CV) scheme, in which the data were partitioned into four folds, each corresponding to one of the four site-seasons (Hidango 2021 to 2022, Hidango 2022 to 2023, La Cancha 2022 to 2023, and Villa Baviera 2020 to 2021). For each recipe-algorithm combination, the ten hyperparameter candidates were fitted on the three remaining site-seasons and evaluated on the held-out site-season; the configuration with the highest mean R^2 across the four folds was selected, and its corresponding mean R^2, RMSE, and MAE across folds were retained as the final performance estimate. This LOSO scheme directly evaluates the spatial transferability of each model to a field/season not used for training, which is the most policy-relevant generalization scenario for operational AGB monitoring. We note that, because the held-out site-season is used both to select the best-performing hyperparameter configuration and to report that configuration's performance, this procedure does not provide a fully nested estimate of generalization error: the reported R^2, RMSE, and MAE for each recipe-algorithm combination may be optimistically biased relative to a strictly nested LOSO-CV, in which hyperparameter selection would be confined to the three training site-seasons only (e.g., via an inner cross-validation loop). We retain this non-nested scheme because, with only four site-seasons, an inner-loop split would leave too few independent units per fold; we therefore interpret the reported metrics primarily as a basis for relative comparison among recipes and algorithms, all of which are subject to the same selection procedure and thus the same potential bias, rather than as unbiased absolute estimates of out-of-sample performance. The complementary sensor-ablation analysis (Table 5), which uses fixed hyperparameters and is therefore not subject to this selection-based optimism, corroborates the recipe-level

ranking (Section 3.1.1), providing some reassurance that the qualitative conclusions are not solely an artifact of this leakage; we discuss this limitation further, including the need for nested cross-validation in future work, in the Discussion (Study limitations). For the prediction models (stage 2), the same LOSO-CV scheme was applied to the 100-random-samples-per-site-per-lead-time dataset, using recipe *rec1* (all predictors) combined with five algorithms (RF, XGBoost, GLMnet, bagMLP, and KNN). For each lead time and algorithm, the ten hyperparameter candidates were evaluated across the four LOSO folds and the configuration with the highest mean R^2 across folds was retained, with its mean R^2, RMSE, and MAE across folds reported as the final performance estimate; selected models were then refit on the complete dataset using the chosen hyperparameters. We used the mean absolute error (MAE), root mean square error (RMSE), and r-squared (R^2) metrics (see Supplementary Material Eqs. S1, S2 & S3).

2.3.7. Stacked ensemble model

In addition to the 40 individual recipe-algorithm combinations evaluated via LOSO-CV, we built a stacked ensemble model following (Breiman, 1996b) and (Wolpert, 1992). Using the {stacks} package (Couch and Kuhn, 2022), the tuned candidates from the stage 1 workflow set were supplied as ensemble members; a LASSO meta-learner (penalized linear regression) was fitted on their out-of-fold predictions to estimate blending weights, retaining only members with a non-zero coefficient. The selected members were then refit on the complete dataset (all four site-seasons, recipe *rec1*, all predictors) to produce the final ensemble model, which was used as the basis for the explainable ML analysis described below.

Table 4: Parameters per machine learning algorithm.

Parameter	**Description**	**Algorithm**
trees*	Number of trees in the ensemble (fixed at 1000).	XGBoost, RF
tree_depth	Maximum depth of the tree.	XGBoost, RF
min_n	Minimum data points in a node required to split.	XGBoost, RF
loss_reduction	Reduction in loss function required to split.	XGBoost
sample_size	Proportion of data exposed to the fitting routine.	XGBoost
learn_rate	Rate at which boosting adapts iteration-to-iteration.	XGBoost
mtry	Predictors randomly sampled at each split.	XGBoost, RF
penalty	Total amount of regularization.	GLMnet, bagMLP
mixture	Proportion of L1 regularization (lasso).	GLMnet
hidden units	Number of units in the hidden model.	bagMLP
epochs	Number of training iterations.	bagMLP
neighbors	Number of neighbors to consider (k).	KNN
weight_func	Kernel function type to weight distances.	KNN
dist_power	Parameter for Minkowski distance calculation.	KNN

RF: random forest; XGBoost: extreme gradient boosting; GLMnet: regularized generalized linear regression; bagMLP: single-layer, feed-forward neural networks; KNN: k-nearest neighbour.
*This parameter was not tuned.

2.3.8. Explainable Machine Learning with DALEX

To assess the contribution of each covariate to the model's performance, we performed a variable importance analysis. Further, to evaluate how the most important variables impact the model's predictions, we used partial dependence plots. For this, we used the DALEX (moDel Agnostic Language for Exploration and eXplanation). DALEX provides a model-agnostic framework for interpreting machine learning models by quantifying the impact of individual variables on model predictions. We computed permutation-based variable importance, which evaluates the increase in prediction error after permuting the values of a given variable while keeping all other variables fixed. This approach captures both main effects and interactions and is applicable across diverse model types. For estimation, we analyzed the ensemble model that used all the predictors (recipe *rec1*), and for prediction, we used the best model for each lead time. For each variable, we calculated the increase in root mean squared error (ΔRMSE) caused by the permutation. Higher ΔRMSE values indicate greater importance of the corresponding variable to model performance. The analysis was repeated over 50 random permutations to ensure stability of the importance rankings. The top predictors were identified by ranking variables based on the mean ΔRMSE. In addition, we calculated a "scaled dropout loss" metric to compare variable influence across lead times in the prediction models.

2.4. Software

We utilized the R programming language (R Core Team, 2025) for processing satellite data, as well as for data analysis, ML modeling, and visualization. To access satellite data, we employed the {earthdatalogin} package (Boettiger et al., 2023). For smoothing the time series of vegetation indices, we applied the {imputeTS} package (Moritz and Bartz-Beielstein, 2017). We used the {terra} package (Hijmans, 2023) and the {sf} package (Pebesma, 2018) for processing raster and vectorial data. For data science tasks and visualization, we utilized the {tidyverse} suite (Wickham et al., 2019). For machine learning modeling, we work with the {tidymodels} framework (Kuhn and Wickham, 2020). For the variable importance we used {DALEX} (Biecek, 2018). The {tmap} package (Tennekes, 2018) was employed for mapping purposes.

3. Results

3.1. Estimation model for AGB within the growing season

3.1.1. Selecting the best models

For AGB estimation, we tested 40 models generated by a combination of five machine learning algorithms and eight recipes, each evaluated using leave-one-site-out cross-validation (LOSO-CV) across the four site-

seasons. Across the 40 models, R^2 ranged from 0.44 to 0.78, MAE from 5.2 to 9.0 t/ha, and RMSE from 6.7 to 10.2 t/ha (Figure 4). RF and XGBoost were generally the best-performing algorithms, while GLMnet was the weakest overall. The recipe *rec7*, which uses only PlanetScope (PS) predictors, ranked first with RF (R^2=0.78±0.08, RMSE=7.87±2.08 t/ha, MAE=5.90±2.32 t/ha; mean±SD across the four LOSO folds) and second with XGBoost (R^2=0.77, RMSE=7.55 t/ha, MAE=5.75 t/ha); *rec8* (PS+weather) ranked third and fourth with XGBoost and RF, respectively. By contrast, *rec2* (S1 only) consistently ranked among the worst recipes regardless of algorithm (ranks 36-40 of 40, R^2=0.44-0.53), with its GLMnet configuration ranking last overall, indicating that Sentinel-1 backscatter alone carries little transferable information about AGB across sites and seasons. The recipe that incorporates all covariates (*rec1*) did not outperform the PlanetScope-only recipes; its best configuration (bagMLP) ranked sixth overall (R^2=0.75, RMSE=6.96 t/ha, MAE=5.34 t/ha). This pattern is corroborated by an independent sensor-ablation analysis (Table 5), in which adding Sentinel-1 to a Sentinel-2+weather model changed RMSE by only +0.5% (i.e., a marginal degradation), whereas adding PlanetScope to a Sentinel-1+Sentinel-2+weather model improved RMSE by 5.7%.

Figure 5 compares observed and out-of-fold predicted AGB for the top-ranked configuration (*rec7*_RF, PlanetScope only), its runner-up (*rec7*_XGBoost), and a representative all-predictors configuration (*rec1*_RF, rank 8 overall, R^2=0.74, RMSE=7.88 t/ha) for context. The PlanetScope-only models track the 1:1 line closely across the full biomass range, including the high-biomass observations (>25 t/ha) where the all-predictors model shows a tendency to underestimate.

As an independent check on the recipe ranking above, we ran a complementary sensor-ablation analysis with fixed-hyperparameter XGBoost models (fixed hyperparameters were used here, rather than the per-recipe tuned configurations underlying Figure 4, to isolate the marginal effect of each sensor group from hyperparameter-tuning variance) under the same LOSO-CV scheme, sequentially adding sensor groups to a weather-only baseline (Table 5). Adding Sentinel-2 to the weather-only model reduced RMSE by 23.4%, confirming the strong contribution of optical vegetation indices. Adding Sentinel-1 to the Sentinel-2+weather model produced a negligible change in RMSE (+0.5%, i.e., a slight degradation), while subsequently adding PlanetScope reduced RMSE by a further 5.7%. This independent analysis is consistent with the recipe ranking in Figure 4: optical sensors (Sentinel-2 and PlanetScope) drive most of the predictive performance, whereas Sentinel-1 SAR backscatter contributes little once optical and weather information are available. As Table 5 shows, the across-fold SD for each configuration (2.6-3.4 t/ha) is substantial relative to these mean differences, underscoring that with only four LOSO folds these ablation effects, particularly the near-zero

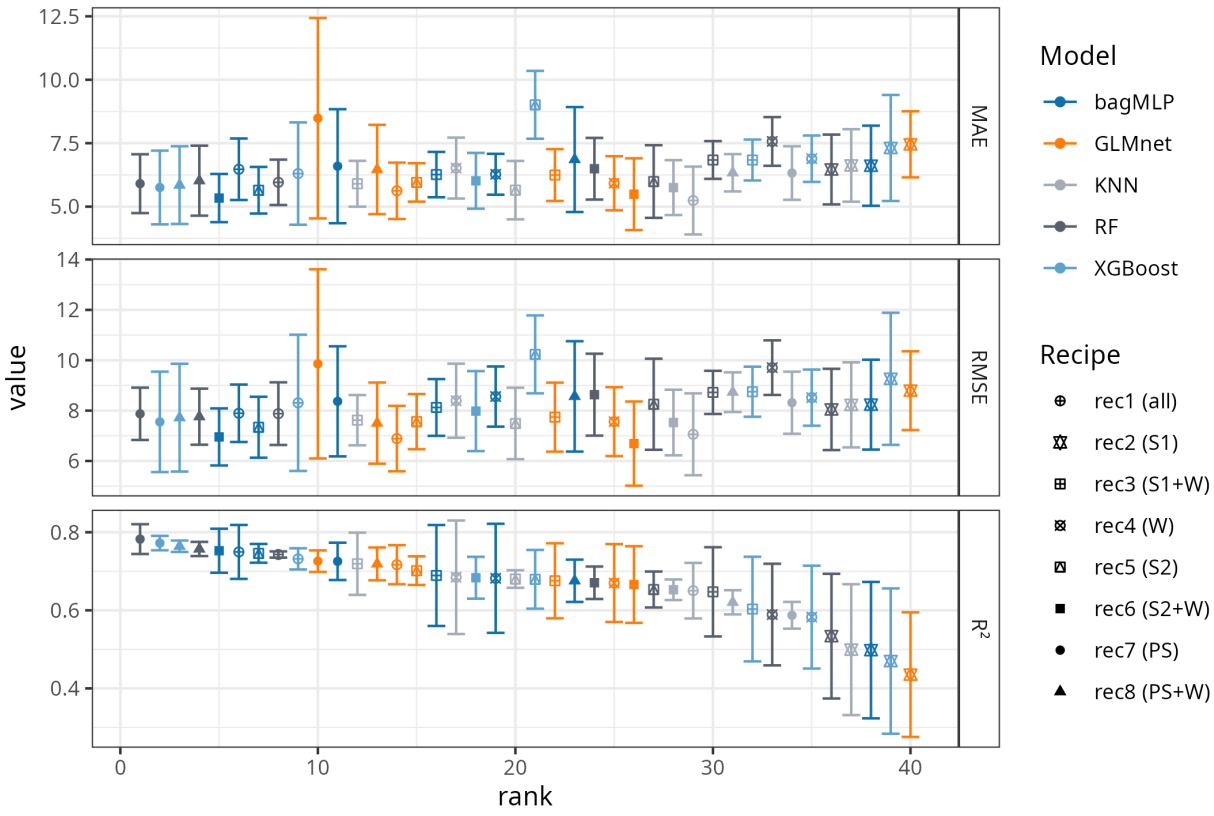

Figure 4: **Performance metrics (R^2, MAE, and RMSE) for the forty models (five ML algorithms and eight recipes), ranked by mean R^2 across the four LOSO-CV folds from best (rank 1) to worst (rank 40).** Points and error bars show the mean ±SD across the four folds for the best hyperparameter configuration of each recipe-algorithm combination. RF: random forest, KNN: k-nearest neighbor, XGBoost: extreme gradient boosting, GLMnet: generalized linear models with regularization, bagMLP: feed-forward neural networks. *rec1*: all predictors, *rec2*: Sentinel-1 (S1), *rec3*: S1+weather (W), *rec4*: weather only, *rec5*: Sentinel-2 (S2), *rec6*: S2+W, *rec7*: PlanetScope (PS), *rec8*: PS+W.

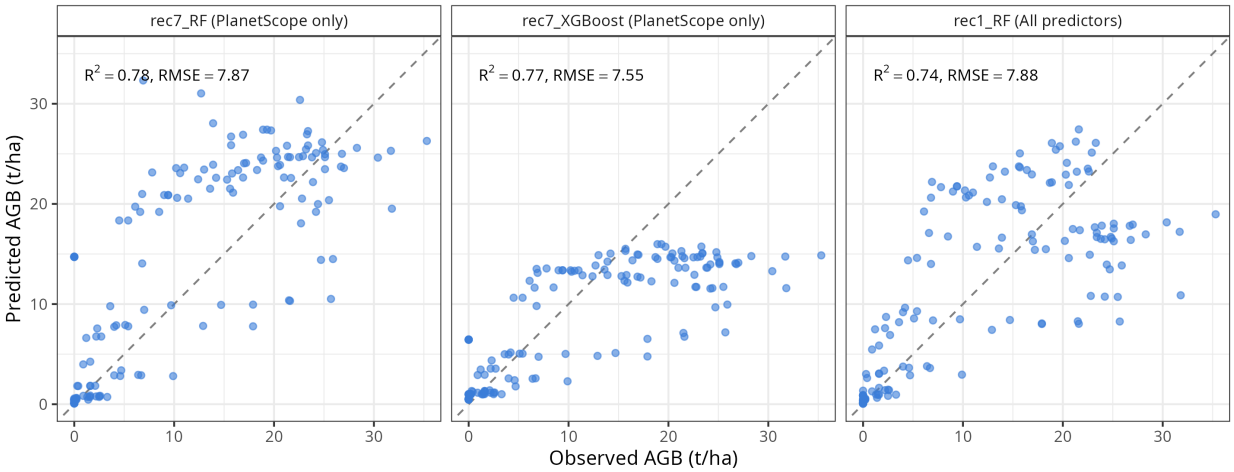

Figure 5: **Comparison of observed and out-of-fold predicted aboveground biomass (LOSO-CV) for the top-ranked models.** R^2 and RMSE correspond to the official ranking (mean across the four LOSO folds, best hyperparameter configuration).

Table 5: Sensor-ablation analysis. Each row reports the change in mean LOSO-CV RMSE (fixed-hyperparameter XGBoost) when a sensor group is added sequentially to the previous model. RMSE values are mean ± SD across the four LOSO folds.

Sensor added	**RMSE before (t/ha)**	**RMSE after (t/ha)**	**RMSE change (%)**
Sentinel-2 (vs. weather only)	9.96	7.63	−23.4
Sentinel-1 (vs. S2 + weather)	7.63	7.68	0.5
PlanetScope (vs. S1 + S2 + weather)	7.68	7.24	−5.7

Sentinel-1 contribution, should be interpreted as indicative rather than statistically definitive.

3.1.2. Explainable Machine Learning with DALEX

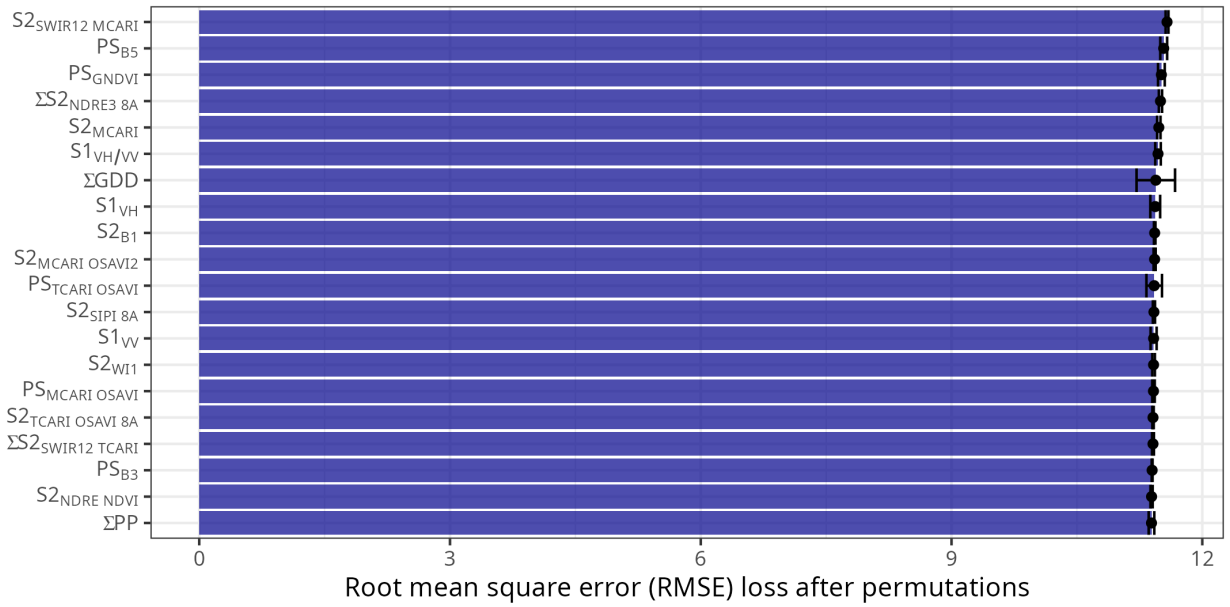

Figure 6: **Optical (Sentinel-2, PlanetScope), SAR (Sentinel-1), and accumulated thermal-time variables show comparable impact on the LOSO-tuned ensemble model performance.** The twenty most important variables, ranked by the increase in RMSE (dropout loss) after permutation in the stacked ensemble model (recipe *rec1*, all predictors). Error bars show ±1 standard deviation across 50 permutations. The symbol indicates cumulative indices computed from the start of the growing season.

For the LOSO-tuned ensemble (recipe *rec1*, all predictors), permutation-based variable importance shows a relatively flat ranking across the leading predictors (Figure 6): the Sentinel-2 SWIR-related chlorophyll index $S2_{SWIR12\_MCARI}$ ranks highest, closely followed by two PlanetScope bands/indices (PS_{B5}, PS_{GNDVI}), a cumulative Sentinel-2 red-edge index ($S2_{NDRE3\_8A}$), $S2_{MCARI}$, the Sentinel-1 backscatter ratio $S1_{VH/VV}$, GDD, $S1_{VH}$, $S2_{B1}$, and $S2_{MCARI/OSAVI2}$, all within a narrow range of dropout-loss values (11.4-11.6 RMSE units). Unlike the recipe-level ranking in Figure 4, where PlanetScope-based recipes clearly outperform the others, the ensemble-level importance indicates that optical (S2, PS), SAR (S1), and accumulated thermal-time predictors all contribute comparably once combined in a single model, consistent with the

residual collinearity documented by the VIF audit (Section 2.3.3).

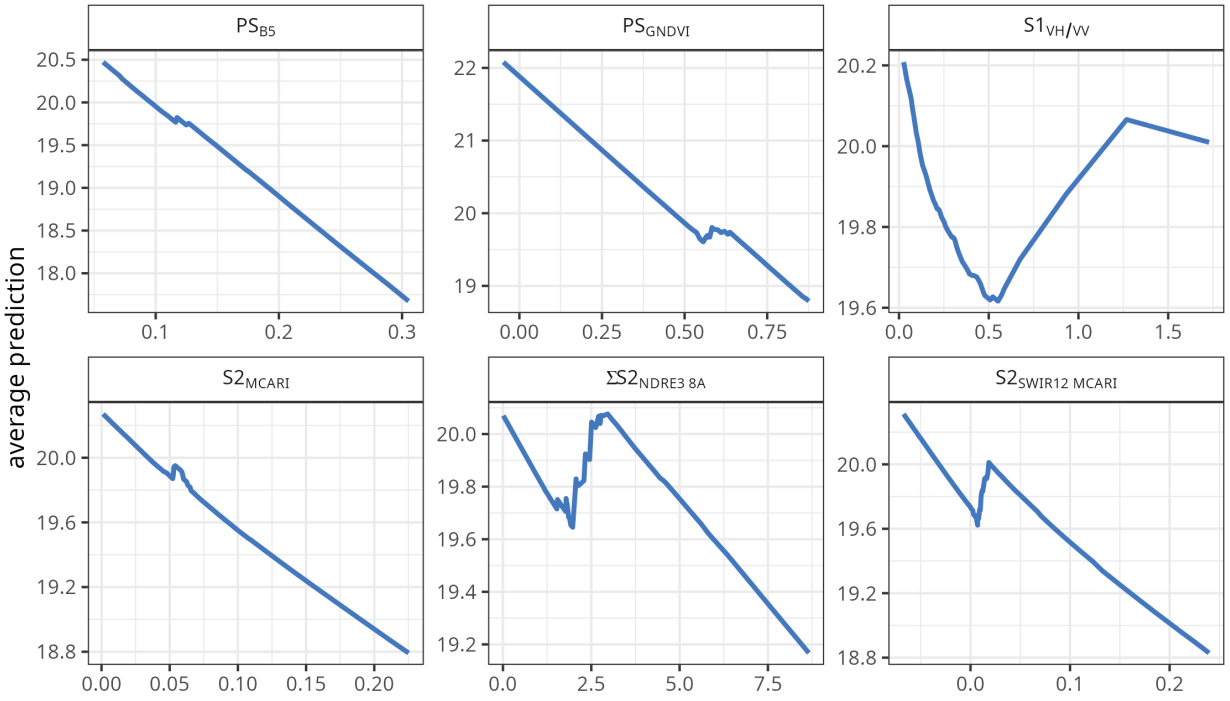

Figure 7: **Partial dependence plots for the six leading predictors of the LOSO-tuned ensemble model (recipe *rec1*, all predictors).** Each panel shows the average predicted AGB (t/ha, y-axis) as a function of the predictor value (x-axis), with all other predictors held at their observed distribution. The symbol indicates a cumulative index computed from the start of the growing season.

Partial dependence plots for the six leading predictors of the LOSO-tuned ensemble (Figure 7) reveal predominantly negative, near-monotonic relationships for the optical predictors: predicted AGB decreases by roughly 2 t/ha as PS_{B5} increases across its observed range, by a similar margin as PS_{GNDVI} increases from 0 to 0.8, and more gently for $S2_{MCARI}$ and $S2_{SWIR12\_MCARI}$. The cumulative red-edge index $S2_{NDRE3\_8A}$ shows a non-monotonic, oscillating profile with local maxima and minima spanning about 1 t/ha. The Sentinel-1 backscatter ratio $S1_{VH/VV}$ displays a shallow U-shape, with higher predicted AGB at both low (~0) and high (>1) ratio values and a minimum around 0.4 to 0.6. These patterns should be interpreted with caution: given the narrow, largely overlapping importance scores in Figure 6 and the residual collinearity among the 42 candidate predictors documented in Section 2.3.3, the marginal effect attributed to any single optical index likely reflects shared variance with co-varying spectral, SAR, and thermal-time predictors rather than an isolated physiological response. This collinearity-based caution applies with particular force to the Sentinel-1 VH/VV ratio: its comparatively high ensemble-level importance in Figure 6 reflects variance shared with co-varying optical and thermal-time predictors, and is therefore not in conflict with the sensor-ablation

result (Table 5), in which Sentinel-1 added essentially no marginal skill once those optical and weather predictors were already present.

3.1.3. Spatio-temporal variation of AGB within the growing season

Finally, we run the models to make a daily spatial estimation of biomass through the growing season at each site. The averaged daily AGB for each site is shown in Figure 8, along with the in-situ AGB measurements.

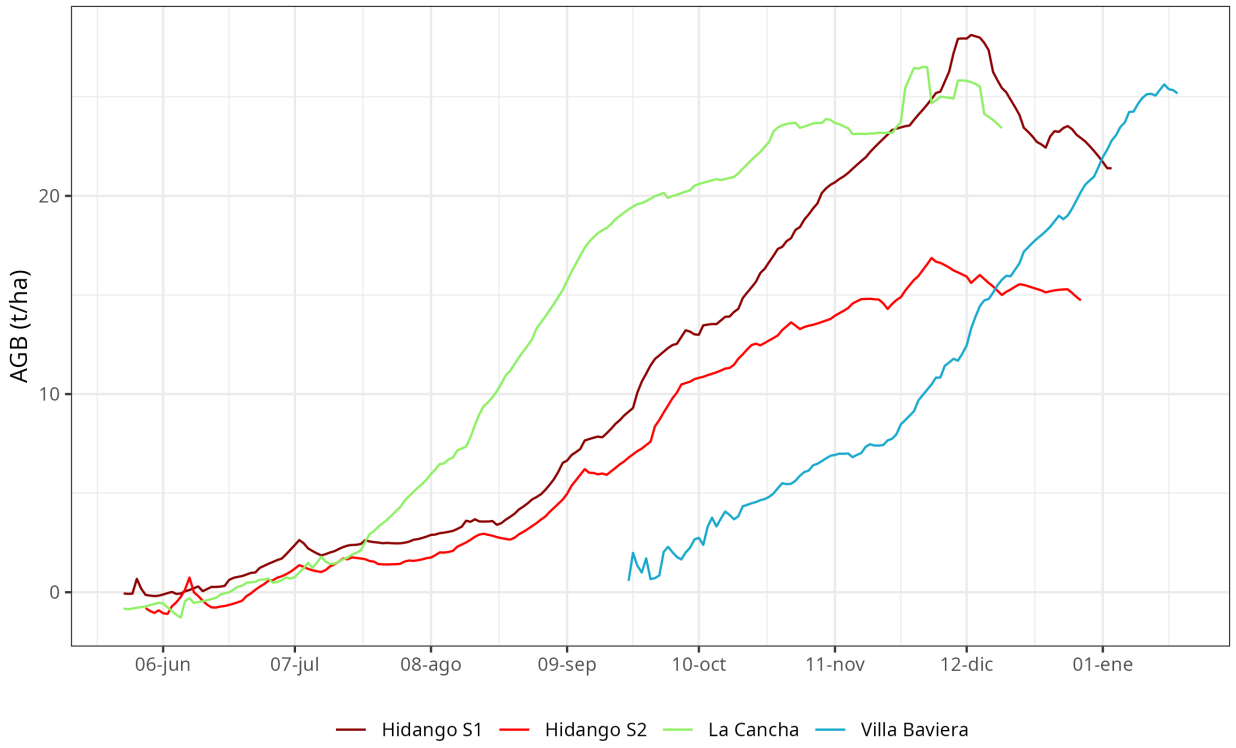

Figure 8: **Daily temporal variation of the averaged AGB per site (lines) and the averaged in situ measured AGB per site (points).** For a better comparison, the dates were converted to a common reference year. In the legend, "Hidango S1" and "Hidango S2" denote the two Hidango site-seasons (2021-2022 and 2022-2023, respectively); these season labels are unrelated to the Sentinel-1/Sentinel-2 sensor abbreviations used elsewhere in this manuscript.

The estimation adjusted well for Hidango in both seasons but underestimated in La Cancha from the middle of the season on. For Villa Baviera, the model estimates well for most of the season, but at the end shows an overestimation. In Figure 9, we show the monthly average of AGB per site. In Hidango season 2021-2022 and in La Cancha, the AGB goes from below 2 t/ha to up to 30 t/ha. In Hidango season 2022-2023, the AGB was lower, reaching 22.50 t/ha. In Villa Baviera, the wheat variety has a shorter season, from September to January; the AGB at harvest was up to 30 t/ha. The sites that showed the higher spatial variability were Hidango season 2022-2023 and La Cancha (Figure 8 and Figure 9).

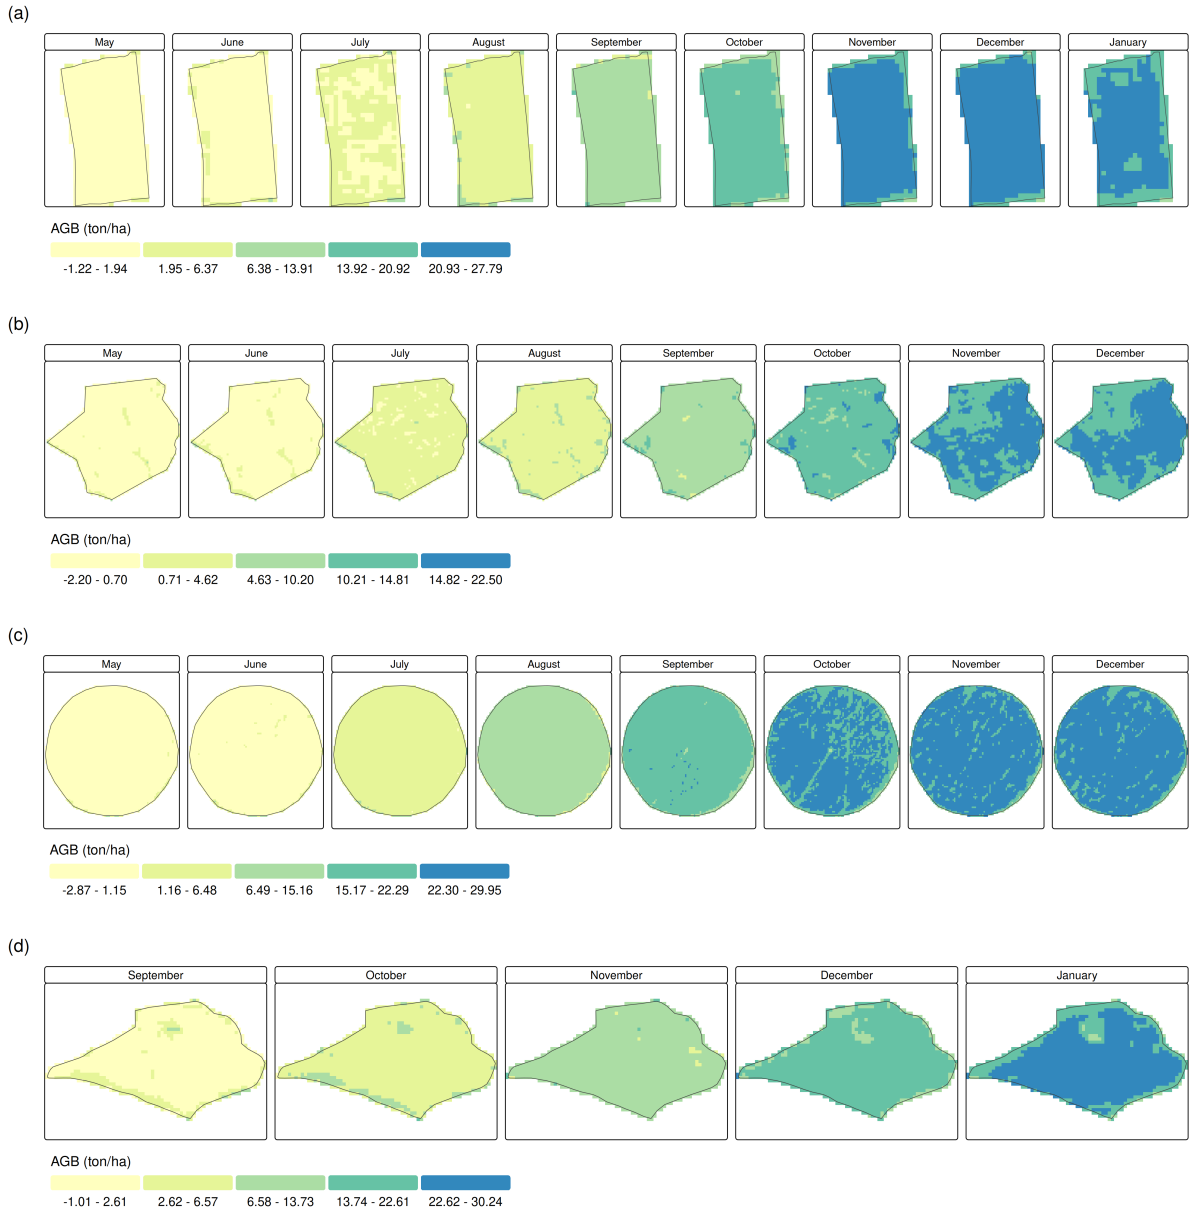

Figure 9: **Monthly averaged above-ground biomass (AGB) estimated with the stacked ensemble model (recipe *rec1*, all predictors) for the sites** (a) Hidango season 2021-2022, (b) Hidango season 2022-2023, (c) La Cancha season 2022-2023, and (d) Villa Baviera season 2020-2021.

3.2. Prediction models for AGB at harvest

As a secondary, exploratory objective, we used the same leave-one-site-out cross-validation (LOSO-CV) framework (recipe *rec1*, all predictors) to evaluate whether covariates available one to four months before harvest are sufficient to reproduce the stage 1 model's own spatial AGB prediction at the harvest date (Section 2.3.5), a self-consistency check on the modeling pipeline rather than validation against independent harvest-date measurements, comparing four algorithms (RF, XGBoost, KNN, and bagMLP); GLMnet was excluded from this comparison because its cross-fold RMSE was highly unstable (15.5 to 54.7 t/ha across leads). Random Forest (RF) was the most consistently competitive of the four across lead times and is used as the reference model below. For RF, R^2 decreased from 0.46 (RMSE=4.52 t/ha, MAE=4.21 t/ha) at a one-month lead to 0.25 (RMSE=5.07 t/ha, MAE=4.73 t/ha) at two months, 0.15 (RMSE=5.09 t/ha, MAE=4.73 t/ha) at three months, and 0.10 (RMSE=6.46 t/ha, MAE=6.13 t/ha) at four months (Figure 10). The other three models followed a broadly similar declining trend. These results indicate modest, rapidly decaying self-consistency between early-season covariates and the stage 1 model's later-season predictions under LOSO-CV, and should not be read as a direct estimate of skill at forecasting field-measured AGB at harvest.

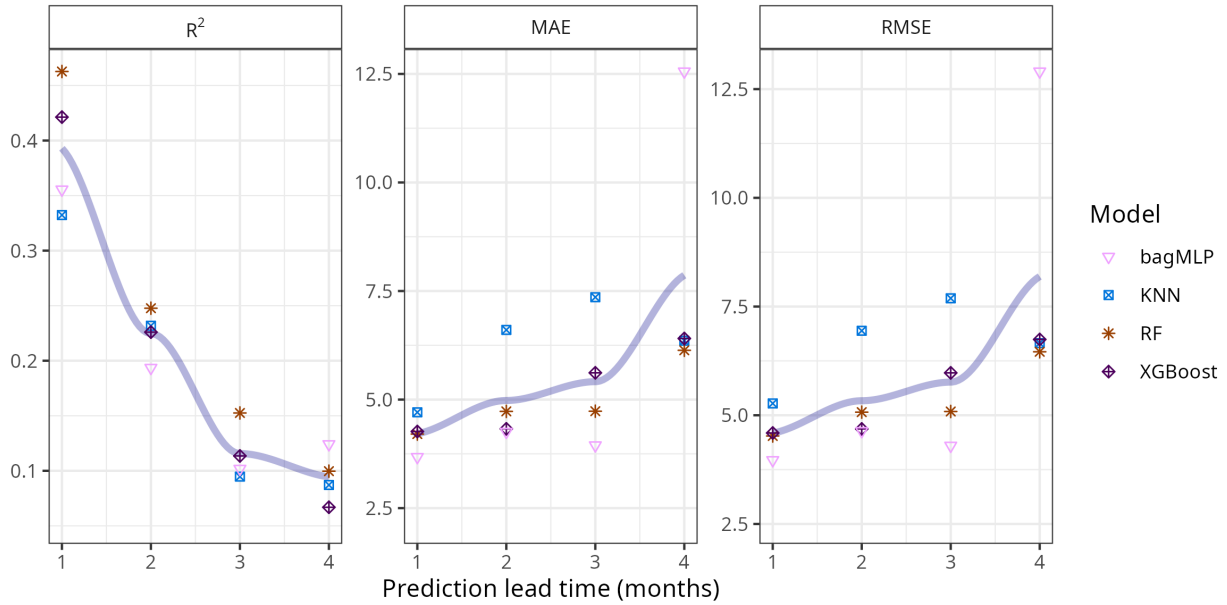

Figure 10: **LOSO-CV performance metrics for AGB-at-harvest forecasts one to four months ahead** (recipe *rec1*, all predictors; best hyperparameter configuration per fold, as in stage 1), for four algorithms: bagged multilayer perceptron (bagMLP), K-nearest neighbors (KNN), random forest (RF), and extreme gradient boosting (XGBoost). R^2 is dimensionless; MAE and RMSE are in t/ha. GLMnet was excluded due to highly unstable cross-fold RMSE (15.5 to 54.7 t/ha). The shaded line shows the overall trend (LOESS smooth) across the four models.

Table 6 lists the eight variables with the most consistent permutation-based importance across the four

Table 6: The eight most important variables (ranked by mean rank) for the RF prediction model (recipe rec1, all predictors) across the four LOSO-CV lead-time models. N: number of leads (of four) in which the variable was among the five most important. Mode rank: most frequent rank across the four leads. Mean rank: average rank across the four leads (lower values indicate more consistent importance). Scaled dropout loss: mean z-scored permutation-importance loss (RMSE increase after permutation) across the four leads.

Variable	**N**	**Mode rank**	**Mean rank**	**Scaled dropout loss**
SM	3	1	5.75	1.91
$\sum S2_{B1}$	2	2	8.25	1.23
$\sum S2_{SWIR12\_TCARI}$	1	5	9.75	−0.11
$\sum PS_{GRVI}$	3	1	11.25	3.24
PS_{B8}	0	8	11.75	−0.09
$S2_{NDRE\_NDVI}$	0	9	12.00	−0.12
PS_{B5}	2	2	12.75	0.56
PS_{TCARI_OSAVI}	3	5	13.00	0.45

lead-time RF models (recipe *rec1*). Soil moisture (SM) had the lowest (best) overall mean rank (5.75), reflecting its rank-1 importance at the four-month lead and rank-2/3 importance at the two- and three-month leads, despite not appearing among the top six predictors at the one-month lead. By contrast, the cumulative PlanetScope green-red vegetation index (PS_{GRVI}) showed a strongly lead-dependent pattern: it was by far the most important predictor at one-, two-, and three-month leads (permutation z-scores of 6.2, 4.1, and 4.0, respectively, several times larger than any other variable at those leads), but fell to near the bottom of the ranking (42nd of 43 variables) at the four-month lead, yielding a middling overall mean rank (11.25). The remaining variables in Table 6 ($S2_{B1}$, $S2_{SWIR12\_TCARI}$, PS_{B8}, $S2_{NDRE\_NDVI}$, PS_{B5}, PS_{TCARI_OSAVI}) each appeared among the top six predictors for at most one or two of the four leads. Three of these eight variables ($S2_{SWIR12\_TCARI}$, PS_{B8}, $S2_{NDRE\_NDVI}$) have negative scaled dropout-loss values, meaning that permuting them slightly reduced, rather than increased, prediction error on average, a pattern consistent with noise rather than genuine predictive signal, and a further indication of the instability discussed below. This pronounced rotation in variable importance with lead time, with no single predictor ranking consistently across all four horizons, further illustrates the instability of harvest forecasting under a four-site-season LOSO-CV scheme and reinforces the need for larger multi-site networks to support reliable multi-month forecasting (see Discussion).

4. Discussion

4.1. Models performance

For in-season AGB estimation, leave-one-site-out cross-validation (LOSO-CV) across four Mediterranean wheat field-seasons spanning three locations, two wheat habits (winter and spring), and distinct soil types offered a preliminary, exploratory assessment of spatial transferability: because each of the four LOSO

folds confounds geography with soil type and, for the Villa Baviera fold, wheat habit (see Study limitations below), these results should be read as an indicative sensitivity analysis rather than as definitive proof of transferability to new, unseen sites. Under this scheme, the best-performing configuration (RF using only PlanetScope predictors, recipe *rec7*) reached R^2=0.78 (RMSE=7.87 t/ha, MAE=5.90 t/ha); as discussed in Methods (Section 2.3.6), this and the other LOSO-CV performance estimates may carry some optimistic bias from the hyperparameter-selection procedure, so absolute values should be read with appropriate caution while the relative ranking among recipes and algorithms is more robust. This R^2 is broadly comparable to studies that used UAV imagery (Atkinson Amorim et al., 2022; Liu et al., 2024) or hyperspectral data (Marshall et al., 2022) (R^2=0.81-0.91), despite the more demanding cross-site validation applied here, since those studies typically relied on within-site or random data splits that do not test spatial transferability. The higher absolute RMSE and MAE relative to those single-site studies likely reflect the wider AGB range and greater inter-site/inter-season heterogeneity sampled here (AGB up to ~30 t/ha across winter and spring wheat).

Contrary to our initial expectation that the all-predictors recipe (*rec1*) or a Sentinel-1-plus-weather combination (*rec3*) would perform best, recipes built on PlanetScope alone (*rec7*) or with weather (*rec8*) consistently outranked *rec1*, while the Sentinel-1-only recipe (*rec2*) was consistently among the worst-performing (R^2=0.44-0.53). The complementary sensor-ablation analysis (Table 5) corroborates this pattern: adding Sentinel-2 to a weather-only baseline reduced RMSE by 23.4%, whereas adding Sentinel-1 to a Sentinel-2-plus-weather model left RMSE essentially unchanged (+0.5%), and subsequently adding PlanetScope produced a further 5.7% reduction. Under LOSO-CV, therefore, optical sensors (PlanetScope and Sentinel-2) carry most of the transferable AGB signal, whereas Sentinel-1 backscatter alone adds little once optical and weather data are available, a different picture from within-site evaluations, where SAR-derived indices are often reported among the strongest predictors (Wang et al., 2023).

At the ensemble level (recipe *rec1*, all predictors), permutation-based variable importance (Figure 6) shows a comparatively flat ranking: a Sentinel-2 SWIR-related chlorophyll index ($S2_{SWIR12\_MCARI}$) ranks highest, closely followed by two PlanetScope predictors (PS_{B5}, PS_{GNDVI}), a cumulative Sentinel-2 red-edge index ($S2_{NDRE3\_8A}$), $S2_{MCARI}$, the Sentinel-1 VH/VV ratio, GDD, $S1_{VH}$, $S2_{B1}$, and $S2_{MCARI\_OSAVI2}$, all within a narrow range of dropout-loss values. This indicates that, once combined in a single model, optical, SAR, and accumulated thermal-time predictors carry comparable, partly redundant information, consistent with the residual collinearity documented by the VIF audit (Section 2.3.3). The partial dependence profiles

for these leading predictors (Figure 7) are predominantly negative and, for the Sentinel-1 VH/VV ratio, non-monotonic (U-shaped); given this collinearity, we interpret these marginal effects with caution rather than as isolated physiological responses.

For the secondary, exploratory objective (stage 2), the same LOSO-CV scheme applied to the 1-4-month lead-time dataset (recipe *rec1*, RF as the most stable model across leads) yielded modest and rapidly decaying self-consistency between early-season covariates and the stage 1 model's harvest-date prediction: R^2 fell from 0.46 (RMSE=4.52 t/ha) at one month to 0.10 (RMSE=6.46 t/ha) at four months. Permutation importance for these RF models (Table 6) showed a pronounced rotation across lead times rather than a stable set of top predictors: the cumulative PlanetScope green-red index (PS_{GRVI}) dominated at one-, two-, and three-month leads but became unimportant at four months, while soil moisture (SM) became the leading predictor only at the four-month lead. This instability, together with the modest and rapidly decaying R^2, is consistent with the limited number of independent site-seasons (n=4) available to a leave-one-site-out scheme: each fold trains on only three site-seasons and is evaluated on a fourth, entirely unseen field, season, and (in most folds) wheat habit. Beyond this small-sample instability, a more fundamental limitation is that the stage 2 target is itself a stage 1 model output rather than an independent harvest-AGB ground truth (Section 2.3.5): even with a larger network of site-seasons, this design would remain a test of internal consistency between early- and late-season model behavior rather than of true forecasting skill against field-measured biomass. We therefore interpret these results as a proof of concept for the modeling pipeline rather than as evidence of operational forecasting skill, and as a call for both larger multi-site, multi-season networks and an independent harvest-date sampling campaign to properly evaluate multi-month AGB forecasting (see Study limitations).

4.2. Application for wheat yield management

The ability to estimate in-season AGB with policy-relevant accuracy (*rec7*_RF, R^2=0.78) can support irrigation, fertilization, and crop-insurance decisions, as well as yield monitoring by government agencies. However, our sensor-ablation results temper one common assumption in the literature: under LOSO-CV, Sentinel-1 backscatter did not improve on an optical-plus-weather baseline (+0.5% RMSE), so SAR should not be regarded as a source of additional accuracy when optical imagery is available. Its practical value is more likely as a substitute rather than a complement: maintaining partial model coverage during persistent cloud cover, when Sentinel-2 or PlanetScope observations are unavailable. Confirming this cloud-substitution role, and quantifying any accuracy trade-off relative to a fully-optical pipeline, would require evaluating

model performance under simulated optical data gaps. Such an analysis is beyond the scope of the present study but a natural direction for future work.

In contrast to UAV- or hyperspectral-based approaches, this framework offers lower equipment and operational costs. The best-performing configuration (*rec7*_RF) relies on PlanetScope imagery without requiring in-situ soil moisture, and a fully public alternative (Sentinel-2 plus weather, *rec6*) captures a substantial share of the achievable improvement (a 23.4% RMSE reduction over a weather-only baseline). This suggests two complementary deployment pathways: a fully public Sentinel-2-plus-weather pipeline that is, in principle, transferable to any region with public meteorological data, and a higher-accuracy PlanetScope-based pipeline that requires a commercial imagery subscription. In both cases, however, the framework's reliance on in-situ weather station data (discussed below) remains the main barrier to low-cost, large-scale deployment, and the LOSO-CV results indicate that performance at new, unmonitored sites would still need to be validated before operational use.

To place the stage 2 forecasting numbers in a decision-making context: at a one-month lead, an RMSE of 4.5 t/ha on AGB values reaching up to ~30 t/ha corresponds to a relative error of roughly 15%, which could support farm-level triage decisions, for example, flagging fields tracking substantially below their expected biomass trajectory for closer monitoring, supplemental irrigation, or adjusted fertilization timing, but would likely be too coarse for fine-grained input allocation (e.g., variable-rate fertilization) or for individual-field crop-insurance payouts, which typically require errors below ~10%. By four months, the near-zero R^2 (0.10) and RMSE of 6.5 t/ha indicate that field-level forecasts are not yet reliable enough to inform individual management decisions at this horizon; at that stage, the framework's value would be limited to regional-scale aggregation, where individual-field errors partly average out, supporting early-warning systems for below-average regional yields rather than field-specific actions.

4.3. Study limitations

The weather data used here come from stations within or near the study fields. Regional or fully scalable applications would need to evaluate performance using satellite-derived weather products (e.g., CHIRPS, ERA5). Soil moisture (SM) measurements, used as a predictor, require dedicated in-field sensors that many farmers and monitoring programs lack; while SM was not among the leading predictors for in-season estimation (stage 1), it emerged as the most important predictor for four-month-ahead forecasting (stage 2, Table 6), so its availability remains relevant for longer-horizon applications. Advances in satellite-based soil moisture estimation at the farm scale (Van Hateren et al., 2023; Zeyliger et al., 2022) could help address this

dependence and facilitate generalization.

Another limitation is that this study is based on three locations (four site-seasons) in central Chile, covering winter and spring wheat varieties. The leave-one-site-out cross-validation scheme used throughout this study is, by design, a conservative estimate of transferability to new sites, but with only four site-seasons each LOSO fold trains on just three of them, particularly limiting for the stage 2 forecasting models, where the rotation in variable importance across lead times (Table 6) likely reflects this small sample as much as genuine seasonal dynamics. Expanding the AGB dataset to more sites, seasons, and varieties is the most direct way to improve and properly evaluate generalization; however, in Chile, reliable and publicly available farm-scale wheat AGB data remain scarce.

A further methodological limitation concerns the hyperparameter-selection procedure described in Section 2.3.6: because the held-out LOSO fold is used both to select the best hyperparameter configuration and to report its performance, the reported R^2, RMSE, and MAE values, including the headline *rec7*_RF estimate (R^2=0.78), may be optimistically biased relative to a fully nested cross-validation, in which hyperparameter tuning would be confined to the training folds only. We were unable to implement a nested LOSO-CV here because, with only four site-seasons, an inner-loop split would leave too few independent units per fold to tune hyperparameters meaningfully. The fixed-hyperparameter sensor-ablation analysis (Table 5), which is not subject to this issue, corroborates the recipe-level ranking and the near-zero marginal contribution of Sentinel-1, suggesting that the qualitative, relative conclusions of this study are more robust than the absolute performance values. We recommend that studies with a larger number of site-seasons implement nested cross-validation, confining hyperparameter selection strictly to training data, to obtain unbiased generalization-error estimates.

A related limitation is that the experimental design confounds wheat habit with location: Villa Baviera is the only spring-wheat site-season, while the other three are winter wheat. As a result, the LOSO fold in which Villa Baviera is held out simultaneously tests transferability to a new location, a new soil type, and a different wheat habit, so it is not possible to fully disentangle whether model performance and the variable-importance patterns reported here reflect generalizable wheat phenology or site-specific (including habit-specific) effects. Given this extreme confounding in a four-fold design, we emphasize that the LOSO-CV results reported throughout this study, including the headline R^2=0.78 for the *rec7*_RF configuration, constitute a preliminary sensitivity analysis of spatial transferability rather than a definitive demonstration that the framework generalizes to new sites, soils, or wheat habits. Resolving this confound would require

additional spring-wheat site-seasons at locations other than Villa Baviera, which we identify as a priority for future data collection.

A further limitation specific to the secondary, exploratory forecasting objective (stage 2, Section 2.3.5) is that its response variable is the stage 1 model's own spatial prediction at the harvest date rather than an independent, field-measured AGB value. The stage 2 results (Section 3.2) therefore demonstrate the internal consistency of the modeling pipeline, i.e., the extent to which pre-harvest covariates can reproduce the stage 1 model's later-season output, but do not constitute evidence that the framework can forecast actual, field-measured biomass at harvest; a genuine test of forecasting skill would require a dedicated destructive-sampling campaign at the harvest date, independent of the stage 1 model. Relatedly, the 100 samples per site-season used for stage 2 modeling are drawn from a smoothed, spatially-autocorrelated stage 1 prediction surface rather than from independent biological replicates, so the nominal sample size (400 across four site-seasons) overstates the effective, independent information content of the stage 2 dataset and may understate the true uncertainty of the reported stage 2 metrics. We therefore present the stage 2 results strictly as an exploratory proof of concept, and recommend that future work pair this modeling framework with independent in-situ harvest-AGB measurements before drawing conclusions about real-world forecasting performance.

Other concerns include the speckle noise inherent to Sentinel-1 and its dependence on incidence angle, which may have contributed to its limited marginal contribution under LOSO-CV (Table 5) and could make consistent SAR-based performance across landscapes or farming systems challenging. Finally, the residual collinearity among cumulative vegetation indices, documented by the VIF audit (Section 2.3.3), affects not only the flat ensemble-level variable importance (Figure 6) but the reliability of permutation-based importance itself (Section 2.3.3): because permuting one of several correlated predictors generates unrealistic, out-of-distribution combinations, the rankings in Figure 6 and Table 6 should be read as identifying groups of mutually-informative predictors rather than precisely localized individual importances. Future work should re-run the explainability analysis on the VIF-reduced 22-variable set identified here, or apply conditional permutation importance methods designed for correlated predictors, before drawing fine-grained conclusions about individual variable contributions.

5. Conclusion

In this study, we evaluated 40 recipe-algorithm combinations, comprising eight sensor-combination recipes (weather, Sentinel-1, Sentinel-2, and PlanetScope, individually and combined) and five ML algorithms, plus a stacked ensemble, to estimate in-season wheat AGB (stage 1) under a preliminary leave-one-site-out spatial cross-validation across four Mediterranean field-seasons (see Discussion, Study limitations, for caveats regarding the confounding of geography with wheat habit and the hyperparameter-selection procedure). The best-performing model used only PlanetScope predictors with Random Forest (*rec7*_RF, R^2=0.78, RMSE=7.87 t/ha, MAE=5.90 t/ha); RF and XGBoost were generally the strongest algorithms, while the all-predictors recipe (*rec1*) and the Sentinel-1-only recipe (*rec2*) ranked lower. A sensor-ablation analysis confirmed that optical sensors (Sentinel-2, PlanetScope) drive most of the transferable AGB signal, while Sentinel-1 backscatter adds essentially no skill once optical and weather data are available. At the ensemble level, permutation-based variable importance showed a comparatively flat ranking among Sentinel-2 and PlanetScope spectral indices, the Sentinel-1 VH/VV ratio, and accumulated growing degree days, indicating substantial redundancy among the candidate predictors.

As a secondary, exploratory objective, we used the same LOSO-CV framework to test, as a self-consistency check, whether covariates available one to four months before harvest can reproduce the stage 1 model's own spatial AGB prediction at harvest (stage 2, recipe *rec1*, Random Forest as the most stable algorithm across leads). Self-consistency was modest and decayed sharply with lead time, from R^2=0.46 (RMSE=4.52 t/ha) at one month to R^2=0.10 (RMSE=6.46 t/ha) at four months, and the leading predictors rotated substantially across leads. A cumulative PlanetScope green-red index dominated at one to three months, while soil moisture became most important only at four months. This decay and instability reflect the limited number of independent site-seasons (n=4) available to a leave-one-site-out scheme and underscore that robust multi-month AGB forecasting will require substantially larger multi-site, multi-season networks. Because the stage 2 target is itself a stage 1 model output rather than an independent harvest-AGB measurement, these results demonstrate the internal consistency of the modeling pipeline rather than validated forecasting skill against ground truth; independent harvest-date sampling is needed to establish the latter.

Overall, these results show that an explainable, multi-sensor ML framework can estimate in-season wheat AGB with policy-relevant accuracy, with optical sensors (PlanetScope, Sentinel-2) as the primary drivers of skill and Sentinel-1 offering, at most, a complementary role under persistent cloud cover; however, given the four-site-season design, these LOSO-CV results should be regarded as a preliminary sensitivity analysis

rather than definitive proof of generalizability to new sites, soils, and wheat habits (see Discussion). Harvest forecasting at multi-month leads remains an open challenge that will benefit from larger and more diverse site networks. Future work should focus on (i) expanding the AGB dataset across more sites, seasons, and wheat varieties, including additional spring-wheat site-seasons, to enable robust multi-month forecasting and to disentangle geography from wheat habit in transferability assessments; (ii) testing satellite-derived weather and soil moisture products (e.g., ERA5, CHIRPS, SAR-based SM) to reduce dependence on in-situ stations; (iii) re-running the explainability analysis on the VIF-reduced predictor set or using conditional permutation importance to address residual collinearity; (iv) implementing nested cross-validation for hyperparameter selection as the number of site-seasons grows; and (v) collecting independent in-situ AGB measurements at harvest to validate the stage 2 forecasting framework against ground truth rather than against the stage 1 model's own predictions.

6. Acknowledgements

We thank Dr. Christian Alfaro from the Instituto de Investigaciones Agropecuarias (INIA) for his support and for providing access to the wheat fields. We also thank Professor Francisco Tapia from Universidad Mayor for his support and his knowledge about wheat development during this research. Finally, we thank Mr. Fabian Llanos, who coordinated the field trips that allowed us to collect the wheat biomass in-situ data.

7. Data and code availability

The data and code used for the analysis and modeling of above-ground biomass (AGB) in wheat fields is openly available at the following GitHub repository: https://github.com/FONDECYT-11190360/AGB_wheat_modelling. This repository contains data, R scripts and documentation for data preprocessing, model development, variable importance analysis, and visualization.

8. CRediT authorship contribution statement

Francisco Zambrano: Project administration, Funding acquisition, Formal analysis, Conceptualization, Methodology, Writing - Original Draft, Writing - Review & Editing, Visualization. **Abel Herrera**: Formal Analysis, Data curation, Visualization, Writing - Original Draft, Writing - Review & Editing. **Mauricio Molina-Rocco:** Conceptualization, Writing - Original Draft, Writing - Review & Editing.

9. Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

10. Funding Sources

Chile's National Research and Development Agency (ANID) funded this study through the grants Fondecyt de Iniciación N°11190360, FONDEF IdEA I+D ID21I10297, and the drought emergency FSEQ210022. We also like to thank the Instituto de Investigaciones Agropecuarias (INIA), the enterprises Ariztia, and Villa Baviera for giving access to the fields of wheat.

11. Declaration of generative AI and AI-assisted technologies in the writing process

During the preparation of this work, the authors used ChatGPT-3.5 and Microsoft Copilot to polish the writing of the manuscript. After using this service, the authors reviewed and edited the content as needed and took full responsibility for the publication.

12. Glossary

Table 7: Abbreviations and definitions

Abbreviation / Term	**Definition**
AGB	Above-Ground Biomass. The total mass of living plant material (excluding roots) in a given area, typically measured in tons per hectare (t/ha). It is the primary dependent variable being modeled and forecasted.
ML	Machine Learning. A subset of artificial intelligence that provides systems with the ability to automatically learn and improve from experience without being explicitly programmed.
SAR	Synthetic Aperture Radar. A type of radar used in satellite remote sensing (e.g., Sentinel-1) that transmits microwaves and records the backscattered signal, providing data related to surface roughness and moisture, independent of weather and light conditions.

Continued on next page

Table 7 – continued from previous page

Abbreviation / Term	Definition
S1	Sentinel-1. The European Space Agency (ESA) satellite mission providing C-band Synthetic Aperture Radar (SAR) data. Used for monitoring land and ocean surface dynamics.
S2	Sentinel-2. The European Space Agency (ESA) satellite mission providing high-resolution optical imagery (13 spectral bands). Used for land cover classification and vegetation monitoring.
PlanetScope	A constellation of Earth-observing satellites operated by Planet Labs. They provide daily, high-resolution optical imagery, complementing the temporal resolution of Sentinel missions.
In-Situ Data	Data collected directly at the field location, involving physical measurement. In this study, this primarily refers to collected weather variables and soil moisture measurements.
GDD	Growing Degree Days. A measure of heat accumulation used to predict the rate of plant development between growing stages (phenological stages). It is calculated based on daily maximum and minimum air temperatures.
SM	Soil Moisture. The water content of the soil, measured here with TEROS 12 capacitance sensors and used as both a predictor of crop development and an indicator of water stress.
RF	Random Forest. An ensemble learning method that builds multiple decision trees on bootstrapped samples and averages their predictions, widely used for its robustness and resistance to overfitting.
XGBoost	eXtreme Gradient Boosting. A highly optimized and efficient open-source implementation of the gradient boosting framework, popular for its speed and performance in structured data prediction.
GLMnet	Generalized Linear Model with regularization. A model that applies regularized regression (L1/L2 penalties) to linear models, often used for variable selection and preventing overfitting.

Continued on next page

Table 7 – continued from previous page

Abbreviation / Term	**Definition**
KNN	k-Nearest Neighbors. A non-parametric supervised learning method used for classification and regression. The output is based on the classification or mean value of its nearest neighbors.
bagMLP	Bagged Multilayer Perceptron. An ensemble learning technique where multiple Multilayer Perceptron (neural network) models are trained on different subsets of the data (bagging), and their predictions are averaged.
Cross-Validation	A statistical technique used to estimate how accurately a predictive model will perform in practice. It involves partitioning the data into subsets for training and testing.
LOSO-CV	Leave-One-Site-Out Cross-Validation. A spatial cross-validation scheme in which each fold corresponds to one entire site-season; models are trained on the remaining site-seasons and evaluated on the held-out one, directly testing spatial transferability.
VIF	Variance Inflation Factor. A diagnostic statistic that quantifies the degree of multicollinearity among predictor variables in a regression model; values above 10 are commonly considered indicative of problematic collinearity.
R^2	Coefficient of Determination. A statistical measure representing the proportion of the variance for a dependent variable that is explained by the independent variables in a regression model. A value closer to 1 indicates a better fit.
RMSE	Root Mean Square Error. A measure of the average magnitude of the errors. It tells you how concentrated the data is around the line of best fit. It has the same units as the dependent variable (t/ha).
MAE	Mean Absolute Error. The average of the absolute differences between predicted and observed values, expressed in the same units as the dependent variable (t/ha). Less sensitive to large outliers than RMSE.

Continued on next page

Table 7 – continued from previous page

Abbreviation / Term	**Definition**
DALEX	moDel Agnostic Language for Exploration and eXplanation. A specific framework or library used in this research to provide model-agnostic tools for explaining the behavior of complex ML models (contributing to XAI).

References

Atkinson Amorim, J., Schreiber, L., De Souza, M., Negreiros, M., Susin, A., Bredemeier, C., Trentin, C., Vian, A., De Oliveira Andrades-Filho, C., Doering, D., Parraga, A., 2022. Biomass estimation of spring wheat with machine learning methods using uav-based multispectral imaging. Int. J. Remote Sens. 43, 4758–4773. doi:10.1080/01431161.2022.2107882.

Bahar, N., Lo, M., Sanjaya, M., Van Vianen, J., Alexander, P., Ickowitz, A., Sunderland, T., 2020. Meeting the food security challenge for nine billion people in 2050: What impact on forests? Glob. Environ. Change 62, 102056. doi:10.1016/j.gloenvcha.2020.102056.

Bates, J., Montzka, C., Schmidt, M., Jonard, F., 2021. Estimating canopy density parameters time-series for winter wheat using uas mounted lidar. Remote Sens. 13, 710. doi:10.3390/rs13040710.

Beck, H., McVicar, T., Vergopolan, N., Berg, A., Lutsko, N., Dufour, A., Zeng, Z., Jiang, X., van Dijk, A., Miralles, D., 2023. High-resolution (1 km) köppen-geiger maps for 1901–2099 based on constrained cmip6 projections. Sci. Data 10. doi:10.1038/s41597-023-02549-6.

Biecek, P., 2018. Dalex: Explainers for complex predictive models in r. J. Mach. Learn. Res. 19, 1–5.

Boettiger, C., López, L., Panda, Y., Lind, B., 2023. earthdatalogin: Nasa "earthdata" login utilities.

Breiman, L., 1996a. Bagging predictors. Mach. Learn. 24, 123–140. doi:10.1007/bf00058655.

Breiman, L., 1996b. Stacked regressions. Mach. Learn. 24, 49–64.

Brovkina, O., Novotny, J., Cienciala, E., Zemek, F., Russ, R., 2017. Mapping forest aboveground biomass using airborne hyperspectral and lidar data in the mountainous conditions of central europe. Ecol. Eng. 100, 219–230. doi:10.1016/j.ecoleng.2016.12.004.

Chacón C., G., 2019. Descripción y usos de la Red de Agrometeorología INIA. Technical Report. Instituto de Investigaciones Agropecuarias. Santiago, Chile.

Chao, Z., Liu, N., Zhang, P., Ying, T., Song, K., 2019. Estimation methods developing with remote sensing information for energy crop biomass: A comparative review. Biomass Bioenergy 122, 414–425. doi:10.1016/j.biombioe.2019.02.002.

Chave, J., Réjou-Méchain, M., Búrquez, A., Chidumayo, E., Colgan, M., Delitti, W., Duque, A., Eid, T., Fearnside, P., Goodman, R., Henry, M., Martínez-Yrízar, A., Mugasha, W., Muller-Landau, H., Mencuccini,

M., Nelson, B., Ngomanda, A., Nogueira, E., Ortiz-Malavassi, E., Pélissier, R., Ploton, P., Ryan, C., Saldarriaga, J., Vieilledent, G., 2014. Improved allometric models to estimate the aboveground biomass of tropical trees. Glob. Change Biol. 20, 3177–3190. doi:10.1111/gcb.12629.

Chen, M., 2025. Food security or climate action. Nat. Clim. Change 15, 354–355. doi:10.1038/s41558-025-02297-y.

Chen, T., Guestrin, C., 2016. Xgboost: A scalable tree boosting system, in: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 785–794. doi:10.1145/2939672.2939785.

Clevers, J., Gitelson, A., 2013. Remote estimation of crop and grass chlorophyll and nitrogen content using red-edge bands on sentinel-2 and -3. Int. J. Appl. Earth Obs. Geoinformation 23, 344–351. doi:10.1016/j.jag.2012.10.008.

Couch, S., Kuhn, M., 2022. stacks: Tidy model stacking. R package version 1.0.1.

Daughtry, C., 2000. Estimating corn leaf chlorophyll concentration from leaf and canopy reflectance. Remote Sens. Environ. 74, 229–239. doi:10.1016/s0034-4257(00)00113-9.

David, R., Rosser, N., Donoghue, D., 2022. Improving above ground biomass estimates of southern africa dryland forests by combining sentinel-1 sar and sentinel-2 multispectral imagery. Remote Sens. Environ. 282, 113232. doi:10.1016/j.rse.2022.113232.

Del Pozo, A., García-Huidobro, J., Novoa, R., Villaseca, S., 1987. Relationship of base temperature to development of spring wheat. Exp. Agric. 23, 21–30. doi:10.1017/S0014479700001095.

Del Pozo, A., Méndez-Espinoza, A., Garriga, M., Estrada, F., Castillo, D., Matus, I., Lobos, G., 2023. Phenotypic variation in leaf photosynthetic traits, leaf area index, and carbon discrimination of field-grown wheat genotypes and their relationship with yield performance in mediterranean environments. Planta 258. doi:10.1007/s00425-023-04163-7.

Dietz, K., Zörb, C., Geilfus, C., 2021. Drought and crop yield. Plant Biol. 23, 881–893. doi:10.1111/plb.13304.

Drusch, M., Del Bello, U., Carlier, S., Colin, O., Fernandez, V., Gascon, F., Hoersch, B., Isola, C., Laberinti, P., Martimort, P., Meygret, A., Spoto, F., Sy, O., Marchese, F., Bargellini, P., 2012. Sentinel-2:

Esa's optical high-resolution mission for gmes operational services. Remote Sens. Environ. 120, 25–36. doi:10.1016/j.rse.2011.11.026.

Fan, D., Zhao, T., Jiang, X., García-García, A., Schmidt, T., Samaniego, L., Attinger, S., Wu, H., Jiang, Y., Shi, J., Fan, L., Tang, B.H., Wagner, W., Dorigo, W., Gruber, A., Mattia, F., Balenzano, A., Brocca, L., Jagdhuber, T., Wigneron, J.P., Montzka, C., Peng, J., 2025. A sentinel-1 sar-based global 1-km resolution soil moisture data product: Algorithm and preliminary assessment. Remote Sens. Environ. 318, 114579. doi:10.1016/j.rse.2024.114579.

FAO, 2014. World Reference Base for Soil Resources 2014: International Soil Classification System for Naming Soils and Creating Legends for Soil Maps. Technical Report. Food and Agriculture Organization of the United Nations. Rome.

FAO, 2025. Cereal supply and demand brief, july 2025 update.

Fragkos, A., Loukatos, D., Kargas, G., Arvanitis, K., 2024. Response of the teros 12 soil moisture sensor under different soils and variable electrical conductivity. Sensors 24, 2206. doi:10.3390/s24072206.

Gardner, B., Durost, D., Lin, W., Lu, Y.C., Nelson, G., Whittlesey, N., 1980. Measurement of US agricultural productivity: a review of current statistics and proposals for change. Technical Report.

Garrido, M., Román, L., Silva, P., Castellaro, G., García De Cortázar, V., Acevedo, E., 2013. Characterization of genetic coefficients of durum wheat (triticum turgidum l. ssp. durum) llareta-inia and corcolén-inia. Chil. J. Agric. Res. 73, 02–03. doi:10.4067/s0718-58392013000200002.

Gitelson, A., Kaufman, Y., Merzlyak, M., 1996. Use of a green channel in remote sensing of global vegetation from eos-modis. Remote Sens. Environ. 58, 289–298. doi:10.1016/S0034-4257(96)00072-7.

Gower, J., 1971. A general coefficient of similarity and some of its properties. Biometrics 27, 857. doi:10.2307/2528823.

Haboudane, D., Miller, J., Tremblay, N., Zarco-Tejada, P., Dextraze, L., 2002. Integrated narrow-band vegetation indices for prediction of crop chlorophyll content for application to precision agriculture. Remote Sens. Environ. 81, 416–426. doi:10.1016/s0034-4257(02)00018-4.

Hastie, T., Tibshirani, R., Wainwright, M., 2015. Statistical Learning with Sparsity: The Lasso and Generalizations. Chapman and Hall/CRC. doi:10.1201/b18401.

He, M., Kimball, J., Maneta, M., Maxwell, B., Moreno, A., Beguería, S., Wu, X., He, M., Kimball, J., Maneta, M., Maxwell, B., Moreno, A., Beguería, S., Wu, X., 2018. Regional crop gross primary productivity and yield estimation using fused landsat-modis data. Remote Sens. 10, 372. doi:10.3390/rs10030372.

Hechenbichler, K., Schliep, K., 2004. Weighted k-Nearest-Neighbor Techniques and Ordinal Classification. Technical Report. doi:10.5282/UBM/EPUB.1769.

Herrmann, I., Pimstein, A., Karnieli, A., Cohen, Y., Alchanatis, V., Bonfil, D., 2011. Lai assessment of wheat and potato crops by ven s and sentinel-2 bands. Remote Sens. Environ. 115, 2141–2151. doi:10.1016/j.rse.2011.04.018.

Hijmans, R., 2023. terra: Spatial data analysis.

Ho, T., 1995. Random decision forests, in: Proceedings of 3rd International Conference on Document Analysis and Recognition, pp. 278–282.

Hu, X., Li, L., Huang, J., Zeng, Y., Zhang, S., Su, Y., Hong, Y., Hong, Z., 2024. Radar vegetation indices for monitoring surface vegetation: Developments, challenges, and trends. Sci. Total Environ. 945, 173974. doi:10.1016/j.scitotenv.2024.173974.

Jiang, Z., Huete, A., Didan, K., Miura, T., 2008. Development of a two-band enhanced vegetation index without a blue band. Remote Sens. Environ. 112, 3833–3845. doi:10.1016/j.rse.2008.06.006.

Jin, X., Li, Z., Feng, H., Ren, Z., Li, S., 2020. Deep neural network algorithm for estimating maize biomass based on simulated sentinel 2a vegetation indices and leaf area index. Crop J. 8, 87–97. doi:10.1016/j.cj.2019.06.005.

Kuhn, M., Wickham, H., 2020. Tidymodels: a collection of packages for modeling and machine learning using tidyverse principles.

Li, H., Li, F., Xiao, J., Chen, J., Lin, K., Bao, G., Liu, A., Wei, G., 2024. A machine learning scheme for estimating fine-resolution grassland aboveground biomass over china with sentinel-1/2 satellite images. Remote Sens. Environ. 311, 114317. doi:10.1016/j.rse.2024.114317.

Lischeid, G., Webber, H., Sommer, M., Nendel, C., Ewert, F., 2022. Machine learning in crop yield modelling: A powerful tool, but no substitute for science. Agric. For. Meteorol. 312, 108698. doi:10.1016/j.agrformet.2021.108698.

Liu, L., Chen, X., 2025. Estimation of the impact of climate warming on spring wheat (triticum aestivum l.) phenology from observations and modelling in the arid region of northwest china. J. Agron. Crop Sci. 211. doi:10.1111/jac.70011.

Liu, T., Yang, T., Zhu, S., Mou, N., Zhang, W., Wu, W., Zhao, Y., Yao, Z., Sun, J., Chen, C., Sun, C., Zhang, Z., 2024. Estimation of wheat biomass based on phenological identification and spectral response. Comput. Electron. Agric. 222, 109076. doi:10.1016/j.compag.2024.109076.

Lobell, D., Burke, M., 2010. On the use of statistical models to predict crop yield responses to climate change. Agric. For. Meteorol. 150, 1443–1452. doi:10.1016/j.agrformet.2010.07.008.

Lopes, A., Touzi, R., Nezry, E., 1990. Adaptive speckle filters and scene heterogeneity. IEEE Trans. Geosci. Remote Sens. 28, 992–1000. doi:10.1109/36.62623.

Luo, S., Wang, C., Xi, X., Pan, F., Peng, D., Zou, J., Nie, S., Qin, H., 2017. Fusion of airborne lidar data and hyperspectral imagery for aboveground and belowground forest biomass estimation. Ecol. Indic. 73, 378–387. doi:10.1016/j.ecolind.2016.10.001.

Maghsoudi, Y., Collins, M., Leckie, D., 2012. Speckle reduction for the forest mapping analysis of multi-temporal radarsat-1 images. Int. J. Remote Sens. 33, 1349–1359. doi:10.1080/01431161.2011.568530.

Magney, T., Eitel, J., Vierling, L., 2017. Mapping wheat nitrogen uptake from rapideye vegetation indices. Precis. Agric. 18, 429–451. doi:10.1007/s11119-016-9463-8.

Marshall, M., Belgiu, M., Boschetti, M., Pepe, M., Stein, A., Nelson, A., 2022. Field-level crop yield estimation with prisma and sentinel-2. ISPRS J. Photogramm. Remote Sens. 187, 191–210. doi:10.1016/j.isprsjprs.2022.03.008.

Mishra, A., Ines, A., Das, N., Prakash Khedun, C., Singh, V., Sivakumar, B., Hansen, J., 2015. Anatomy of a local-scale drought: Application of assimilated remote sensing products, crop model, and statistical methods to an agricultural drought study. J. Hydrol. 526, 15–29. doi:10.1016/j.jhydrol.2014.10.038.

Moran, M., Inoue, Y., Barnes, E., 1997. Opportunities and limitations for image-based remote sensing in precision crop management. Remote Sens. Environ. 61, 319–346. doi:10.1016/S0034-4257(97)00045-X.

Moritz, S., Bartz-Beielstein, T., 2017. imputets: Time series missing value imputation in r. R J. 9, 207–218. doi:10.32614/RJ-2017-009.

Naumann, G., Cammalleri, C., Mentaschi, L., Feyen, L., 2021. Increased economic drought impacts in europe with anthropogenic warming. Nat. Clim. Change 11, 485–491. doi:10.1038/s41558-021-01044-3.

Nuthammachot, N., Askar, A., Stratoulias, D., Wicaksono, P., 2022. Combined use of sentinel-1 and sentinel-2 data for improving above-ground biomass estimation. Geocarto Int. 37, 366–376. doi:10.1080/10106049.2020.1726507.

Pebesma, E., 2018. Simple features for r: Standardized support for spatial vector data. R J. 10, 439–446. doi:10.32614/RJ-2018-009.

Peng, J., Zeiner, N., Parsons, D., Féret, J.B., Söderström, M., Morel, J., 2023. Forage biomass estimation using sentinel-2 imagery at high latitudes. Remote Sens. 15, 2350. doi:10.3390/rs15092350.

Pequeno, D., Ferreira, T., Fernandes, J., Singh, P., Pavan, W., Sonder, K., Robertson, R., Krupnik, T., Erenstein, O., Asseng, S., 2024. Production vulnerability to wheat blast disease under climate change. Nat. Clim. Change 14, 178–183. doi:10.1038/s41558-023-01902-2.

Peroni Venancio, L., Chartuni Mantovani, E., Do Amaral, C., Usher Neale, C., Zution Gonçalves, I., Filgueiras, R., Coelho Eugenio, F., 2020. Potential of using spectral vegetation indices for corn green biomass estimation based on their relationship with the photosynthetic vegetation sub-pixel fraction. Agric. Water Manag. 236, 106155. doi:10.1016/j.agwat.2020.106155.

Peñuelas, J., Gamon, J., Fredeen, A., Merino, J., Field, C., 1994. Reflectance indices associated with physiological changes in nitrogen- and water-limited sunflower leaves. Remote Sens. Environ. 48, 135–146. doi:10.1016/0034-4257(94)90136-8.

R Core Team, 2025. R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing. Vienna, Austria.

Raper, T., Varco, J., 2015. Canopy-scale wavelength and vegetative index sensitivities to cotton growth parameters and nitrogen status. Precis. Agric. 16, 62–76. doi:10.1007/s11119-014-9383-4.

Rondeaux, G., Steven, M., Baret, F., 1996. Optimization of soil-adjusted vegetation indices. Remote Sens. Environ. 55, 95–107. doi:10.1016/0034-4257(95)00186-7.

Rouse, J., Hass, R., Schell, J., Deering, D., Harlan, J., 1974. Monitoring the vernal advancement and retrogradation (green wave effect) of natural vegetation. Technical Report. Texas A&M University.

Roy, D., Huang, H., Houborg, R., Martins, V., 2021. A global analysis of the temporal availability of planetscope high spatial resolution multi-spectral imagery. Remote Sens. Environ. 264, 112586. doi:10.1016/j.rse.2021.112586.

Santini, M., Noce, S., Antonelli, M., Caporaso, L., 2022. Complex drought patterns robustly explain global yield loss for major crops. Sci. Rep. 12, 5792. doi:10.1038/s41598-022-09611-0.

Segarra, J., Araus, J., Kefauver, S., 2022. Farming and earth observation: Sentinel-2 data to estimate within-field wheat grain yield. Int. J. Appl. Earth Obs. Geoinformation 107, 102697. doi:10.1016/j.jag.2022.102697.

Shen, Y., Zhang, X., Yang, Z., Ye, Y., Wang, J., Gao, S., Liu, Y., Wang, W., Tran, K., Ju, J., 2023. Developing an operational algorithm for near-real-time monitoring of crop progress at field scales by fusing harmonized landsat and sentinel-2 time series with geostationary satellite observations. Remote Sens. Environ. 296, 113729. doi:10.1016/j.rse.2023.113729.

Tennekes, M., 2018. tmap: Thematic maps in r. J. Stat. Softw. 84, 1–39. doi:10.18637/jss.v084.i06.

Torres, R., Snoeij, P., Geudtner, D., Bibby, D., Davidson, M., Attema, E., Potin, P., Rommen, B., Floury, N., Brown, M., Traver, I., Deghaye, P., Duesmann, B., Rosich, B., Miranda, N., Bruno, C., L'Abbate, M., Croci, R., Pietropaolo, A., Huchler, M., Rostan, F., 2012. Gmes sentinel-1 mission. Remote Sens. Environ. 120, 9–24. doi:10.1016/j.rse.2011.05.028.

Tucker, C., 1979. Red and photographic infrared linear combinations for monitoring vegetation. Remote Sens. Environ. 8, 127–150. doi:10.1016/0034-4257(79)90013-0.

UN, 2024. World Population Prospects 2024: Summary of Results. Technical Report UN DESA/POP/2024/TR/NO. 9. United Nations. New York.

Uribeetxebarria, A., Castellón, A., Aizpurua, A., 2023. Optimizing wheat yield prediction integrating data from sentinel-1 and sentinel-2 with catboost algorithm. Remote Sens. 15, 1640. doi:10.3390/rs15061640.

Van Hateren, T., Chini, M., Matgen, P., Pulvirenti, L., Pierdicca, N., Teuling, A., 2023. On the potential of sentinel-1 for sub-field scale soil moisture monitoring. Int. J. Appl. Earth Obs. Geoinformation 120, 103342. doi:10.1016/j.jag.2023.103342.

Vincini, M., Frazzi, E., D'Alessio, P., 2008. A broad-band leaf chlorophyll vegetation index at the canopy scale. Precis. Agric. 9, 303–319. doi:10.1007/s11119-008-9075-z.

Wang, C., Nie, S., Xi, X., Luo, S., Sun, X., 2016. Estimating the biomass of maize with hyperspectral and lidar data. Remote Sens. 9, 11. doi:10.3390/rs9010011.

Wang, F., Yang, M., Ma, L., Zhang, T., Qin, W., Li, W., Zhang, Y., Sun, Z., Wang, Z., Li, F., Yu, K., 2022. Estimation of above-ground biomass of winter wheat based on consumer-grade multi-spectral uav. Remote Sens. 14, 1251. doi:10.3390/rs14051251.

Wang, H., Cheng, M., Liao, Z., Guo, J., Zhang, F., Fan, J., Feng, H., Yang, Q., Wu, L., Wang, X., 2023. Performance evaluation of aquacrop and dssat-substor-potato models in simulating potato growth, yield and water productivity under various drip fertigation regimes. Agric. Water Manag. 276, 108076. doi:10.1016/j.agwat.2022.108076.

Wickham, H., Averick, M., Bryan, J., Chang, W., McGowan, L., François, R., Grolemund, G., Hayes, A., Henry, L., Hester, J., Kuhn, M., Pedersen, T., Miller, E., Bache, S., Müller, K., Ooms, J., Robinson, D., Seidel, D., Spinu, V., Takahashi, K., Vaughan, D., Wilke, C., Woo, K., Yutani, H., 2019. Welcome to the tidyverse. J. Open Source Softw. 4, 1686. doi:10.21105/joss.01686.

Wolpert, D., 1992. Stacked generalization. Neural Netw. 5, 241–259. doi:10.1016/s0893-6080(05)80023-1.

Wu, C., Niu, Z., Tang, Q., Huang, W., 2008. Estimating chlorophyll content from hyperspectral vegetation indices: Modeling and validation. Agric. For. Meteorol. 148, 1230–1241. doi:10.1016/j.agrformet.2008.03.005.

Yue, J., Zhou, C., Guo, W., Feng, H., Xu, K., 2021. Estimation of winter-wheat above-ground biomass using the wavelet analysis of unmanned aerial vehicle-based digital images and hyperspectral crop canopy images. Int. J. Remote Sens. 42, 1602–1622. doi:10.1080/01431161.2020.1826057.

Zadoks, J., Chang, T., Konzak, C., 1974. A decimal code for the growth stages of cereals. Weed Res. 14, 415–421. doi:10.1111/j.1365-3180.1974.tb01084.x.

Zambrano, F., 2021. Four decades of satellite data for agricultural drought monitoring throughout the growing season in central chile, in: Drought. CRC Press.

Zambrano, F., Lillo-Saavedra, M., Verbist, K., Lagos, O., 2016. Sixteen years of agricultural drought assessment of the biobío region in chile using a 250 m resolution vegetation condition index (vci). Remote Sens. 8, 1–20. doi:10.3390/rs8060530.

912 Zambrano, F., Vrieling, A., Nelson, A., Meroni, M., Tadesse, T., 2018. Prediction of drought-induced
913 reduction of agricultural productivity in chile from modis, rainfall estimates, and climate oscillation indices.
914 Remote Sens. Environ. 219, 15–30. doi:10.1016/j.rse.2018.10.006.

915 Zeyliger, A., Muzalevskiy, K., Zinchenko, E., Ermolaeva, O., 2022. Field test of the surface soil moisture
916 mapping using sentinel-1 radar data. Sci. Total Environ. 807, 151121. doi:10.1016/j.scitotenv.2021.
917 151121.

918 Zhang, H., Zhang, Y., Liu, K., Lan, S., Gao, T., Li, M., 2023. Winter wheat yield prediction using integrated
919 landsat 8 and sentinel-2 vegetation index time-series data and machine learning algorithms. Comput.
920 Electron. Agric. 213, 108250. doi:10.1016/j.compag.2023.108250.

921 Zhao, Y., Wang, X., Guo, Y., Hou, X., Dong, L., 2022. Winter wheat phenology variation and its response
922 to climate change in shandong province, china. Remote Sens. 14, 4482. doi:10.3390/rs14184482.

923 Zheng, Y., Wu, B., Zhang, M., 2017. Estimating the above ground biomass of winter wheat using the
924 sentinel-2 data. Natl. Remote Sens. Bull. 21, 318–328. doi:10.11834/jrs.20176269.

925 Zhu, L., Dai, J., Liu, Y., Yuan, S., Qin, T., Walker, J., 2024. A cross-resolution transfer learning approach for
926 soil moisture retrieval from sentinel-1 using limited training samples. Remote Sens. Environ. 301, 113944.
927 doi:10.1016/j.rse.2023.113944.