

Evaluating Digital Agriculture Recommendations with Causal Inference

Ilias Tsoumas^{*1,2}, Georgios Giannarakis^{*1}, Vasileios Sitokonstantinou¹, Alkiviadis Koukos¹,
Dimitra Loka³, Nikolaos S. Bartsotas¹, Charalampos Kontoes¹, Ioannis Athanasiadis²

¹BEYOND Centre, IAASARS, National Observatory of Athens

²Wageningen University & Research

³Hellenic Agricultural Organization ELGO DIMITRA

{i.tsoumas, giannarakis, vsito, akoukos, nbartsotas, kontoes}@noa.gr,

{ilias.tsoumas, ioannis.athanasiadis}@wur.nl,

{dimitra.loka}@elgo.gr

Abstract

In contrast to the rapid digitalization of several industries, agriculture suffers from low adoption of smart farming tools. Even though recent advancements in AI-driven digital agriculture can offer high-performing predictive functionalities, they lack tangible quantitative evidence on their benefits to the farmers. Field experiments can derive such evidence, but are often costly, time consuming and hence limited in scope and scale of application. To this end, we propose an observational causal inference framework for the empirical evaluation of the impact of digital tools on target farm performance indicators (e.g., yield in this case). This way, we can increase farmers' trust via enhancing the transparency of the digital agriculture market, and in turn accelerate the adoption of technologies that aim to secure farmer income resilience and global agricultural sustainability against a changing climate. As a case study, we designed and implemented a recommendation system for the optimal sowing time of cotton based on numerical weather predictions, which was used by a farmers' cooperative during the growing season of 2021. We then leverage agricultural knowledge, collected yield data, and environmental information to develop a causal graph of the farm system. Using the back-door criterion, we identify the impact of sowing recommendations on the yield and subsequently estimate it using linear regression, matching, inverse propensity score weighting and meta-learners. The results revealed that a field sown according to our recommendations exhibited a statistically significant yield increase that ranged from 12% to 17%, depending on the method. The effect estimates were robust, as indicated by the agreement among the estimation methods and four successful refutation tests. We argue that this approach can be implemented for decision support systems of other fields, extending their evaluation beyond a performance assessment of internal functionalities.

Introduction

The increasing global population and the changing climate are putting pressure on the agricultural sector, demanding the sustainable production of adequate quantities of nutritious food, feed and fiber. Nowadays, many industries enjoy automated and effective decision-making via harnessing the

data that digitalization generates. However, the agricultural sector experiences limited adoption of precision agriculture and smart farming technologies (Gabriel and Gandorfer 2022). This might seem odd at first sight, given the surge of sophisticated digital tools that utilize Artificial Intelligence (AI) techniques and combine remote sensing data with data from Internet of Things (IoT) sensors to offer agricultural information of great detail (Sharma et al. 2020; Nanushi et al. 2022; Choumos et al. 2022). Yet farmers are skeptical about the effectiveness and actual contribution of these tools to their revenues and daily work (Lowenberg-DeBoer and Erickson 2019; Lioutas, Charatsari, and De Rosa 2021).

Traditionally, quantifying the impact of a service would require the design and execution of a randomized experiment (Boruch 1997). Nevertheless, field experiments for the evaluation of digital agriculture tools are seldom done since they are costly and time-consuming, requiring specialized designs and follow-up experiments for any changes in the product (Vaessen 2010; Diggle and Chetwynd 2011). In addition, field experiments can jeopardize the crops and hence the farmers' livelihood; and if potential damages are not covered, no prudent farmer would want to participate. As a result, the providers of digital agriculture tools often resort to unproven promises that unavoidably create customer mistrust. An observational causal inference framework (Pearl 2009) can fill this gap by emulating the experiment we would have liked to run (Hernán and Robins 2016).

Related Work

Causal inference with observational data has been the subject of recent work across diverse disciplines, including ecology (Arif and MacNeil 2022), public policy (Fougère and Jacquemet 2019), and Earth sciences (Massmann, Gentine, and Runge 2021; Runge et al. 2019). In agriculture, it has been used to identify and estimate the effect of agricultural practices on various agro-environmental metrics (Qian and Harmel 2016; Deines, Wang, and Lobell 2019; Giannarakis et al. 2022a,b). Using causal inference to test digital agriculture can provide reliable insights of superior socioeconomic impact than of those inferred by naive descriptive studies, including transparent benefits for the farmers, increased reliability, and honest pricing.

^{*}These authors contributed equally.

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

According to Adelman (1992), the comprehensive evaluation of decision support systems has three facets: i) the subjective evaluation that assesses the system from the perspective of the end-user, ii) the technical evaluation that assesses the performance of the system’s internal functionalities and iii) the empirical evaluation that experimentally assesses the impact of the system (Adelman 1992). Subjective evaluation has been widely practiced for decision support (Zhai et al. 2020) and recommender systems (Pu, Chen, and Hu 2012), retrieving user feedback (e.g., on usability, accessibility etc.) through surveys and questionnaires. From the perspective of the AI-based expert systems, the technical evaluation is based on predictive, classification and ranking accuracy metrics (Schröder, Thiele, and Lehner 2011). Technical evaluation metrics extending beyond accuracy are used within the context of recommender systems (e.g., coverage, serendipity) (Ge, Delgado-Battenfeld, and Jannach 2010). These metrics capture the recommendation quality as perceived by the user, connecting the technical and subjective evaluation concepts.

Interestingly, empirical evaluation methods, and in particular with regards to the impact assessment of digital agriculture tools, have been seldom employed. From the perspective of agricultural economics, tangential questions have been studied (Müller 1974; Roberts et al. 2009; Schimmelpfennig and Ebel 2016; McFadden, Rosburg, and Njuki 2022), but without approaching the question from a farm system standpoint, hence not leveraging available structural knowledge and reaping its benefits (Cinelli, Forney, and Pearl 2020). Thus, we propose a framework for the empirical evaluation of digital agriculture recommendations with causal inference. In this context, we designed and implemented a recommendation system for the optimal sowing of cotton. The system was tested in a real-world case study by providing it to a local agricultural cooperative and monitoring the results.

For several arable crops (e.g., cotton, maize, chickpea) sowing time is of great importance. Mistimed sowing can lead to suboptimal plant emergence and adversely affect the crop yield (Huang 2016; Nielsen et al. 2002; Richards, Maphosa, and Preston 2022). The agro-climatic conditions for optimal sowing have been extensively studied (Freeland Jr et al. 2006; Boman and Lemon 2005). Soil temperature, soil moisture and ambient temperature, during the first days after sowing, play a crucial role in proper germination and emergence and ultimately determine the yield and its quality (Bradow and Bauer 2010; Bauer, May, and Camberato 1998). The University of California (UC IPM 2004) and the Texas Tech University (Barbato, Seshadri, and Mulligan 2011) have developed digital tools that provide daily recommendations for optimal cotton sowing, using air temperature forecasts from the National Oceanic and Atmospheric Administration (NOAA). The accuracy of the University of California’s cotton planting tool was studied by comparing the temperature forecasts to ground truth measurements from two weather stations (Munier et al. 2004). The tool forecasted the correct planting conditions (Kerby, Keeley, and Johnson 1989) 75% of the time, showing that weather forecast models can provide solid sowing recommendations.

To the best of our knowledge, there are no works that evaluate the effectiveness of any type of decision support or recommendation system in the agricultural sector through causal reasoning and beyond their predictive accuracy (Luma-Osmani et al. 2020; Pasquel et al. 2022). The contributions of this work are summarized as follows: i) the design and implementation of the first empirical evaluation framework for digital agriculture based on causal inference; ii) the development of a high-resolution, knowledge-based recommendation system for the optimal sowing of cotton using weather predictions, which was operationally used in a real-world case study; iii) the identification of the causal effect of sowing recommendations on yield, its subsequent estimation with different methods (linear regression, matching, inverse propensity score weighting and meta-learners), and the use of refutation tests to evaluate the robustness of estimates. Due to its accurate weather forecasts and customized rules, we find that the system offered effective sowing recommendations. It increased the cotton yield of the farmers that followed the recommendations by a factor that ranged from 12% to 17%, depending on the estimation method used.

Agricultural Recommendation System

In this work, we design, implement and evaluate a knowledge based recommendation system (Aggarwal 2016) for optimal cotton sowing. The recommendations are based on satisfying specific environmental conditions, as retrieved from the related literature, which would ensure successful cotton planting. The system is operationally deployed using high resolution weather forecasts. Sec. 1 of the Appendix contains an algorithmic presentation of the system.

According to literature, the minimum daily-mean soil temperature for cotton germination is 16°C (Bradow and Bauer 2010). Soil or ambient temperatures lower than 10°C result in less vigorous and malformed seedlings (Boman and Lemon 2005). As a general rule for cotton, agronomists recommend daily-mean soil temperatures higher than 18°C for at least 10 days after sowing and daily-maximum ambient temperatures higher than 26°C for at least 5 days after sowing. We summarize the conditions for optimal cotton sowing in Table 1 (Freeland Jr et al. 2006; Boman and Lemon 2005). Using these conditions and Numerical Weather Predictions (NWP) we implement a recommendation system that advises on whether any given day is a good day to sow or not.

Type of Temperature	Statistic	Condition	Condition Priority
soil (0-10 cm)	mean	>18°C	optimum
ambient (2 m)	max	>26°C	optimum
soil (0-10 cm)	mean	>15.56°C	mandatory
soil (0-10 cm)	min	>10°C	mandatory
ambient (2 m)	min	>10°C	mandatory

Table 1: Optimal conditions for sowing cotton. All conditions refer to the period from sowing day to 5 days after, except the first soil condition that refers to 10 days after.

Open-access high-resolution NWP forecasts are rarely available. For this reason, we implement the WRF-ARW model (Skamarock et al. 2019) with a grid resolution of 2 km. This enables us to reach a high spatio-temporal resolution for parameters that are crucial during the cotton seeding period; soil and ambient temperature are retrieved at an hourly rate for the forthcoming 2.5 days. Ideally, 10-day predictions at a 2 km spatial resolution should be available every morning, as it is required by the conditions in Table 1. However, this would demand an enormous amount of computational power. To simulate the desired data, we combine the 2.5-day high resolution forecasts with the GFS (NCEP and USDOC 2015) 15-day forecasts that are given on a 0.25 degrees (roughly 25 km) spatial resolution.

$$a_i = \frac{GFS_{day=i}}{GFS_{day=1}}, i \in \{3, \dots, 10\} \quad (1)$$

$$ART_j = \begin{cases} WRF_{day=j} & , j \in \{1, 2\} \\ WRF_{day=1} \cdot a_j & , j \in \{3, \dots, 10\} \end{cases} \quad (2)$$

Eq. (1) shows how we extract the 10-day weather trend factor using GFS forecasts. We calculate the percentage change between each forecast (for $day = 3$ to $day = 10$) and the corresponding next day ($day = 1$) forecast. Eq. (2) shows how we produce the artificial (ART) 10-day forecasts at 2 km spatial resolution. We keep the original WRF forecasts for the next two days and for the rest we apply the respective 10-day trend factor to the next day WRF forecast.



Unfavorable \dashv \dashv Favorable

Figure 1: Optimal sowing map for a given day. The black circle at the center depicts the GFS grid point that represents the entire black-lined box. The white circles depict the 144 ART grid points for the same area.

We generate ART forecasts in order to provide recommendations that can vary up to the field-level, which would have been impossible with GFS forecasts alone. This is depicted in Figure 1. In order to evaluate the quality of our ART forecasts, we compared them with measurements from the nearest operational weather station in the area of interest for the critical sowing period, from 15/4/2021 to 15/5/2021. We have limited our comparison to the maximum and minimum ambient temperatures, as there were no soil tempera-

ture measurements available. It is worth noting that the nearest grid point of GFS to the station is only 0.87 km away, however the maximum distance can be up to 12 km away. On the other hand, the equivalent grid point of ART is 1.41 km away, which incidentally is the maximum possible distance between any location and the nearest ART point.

Initially, we compared the next day forecasts of GFS against their ART (or WRF) equivalent. The comparison analysis revealed a Mean Absolute Error (MAE), between the two forecasts and the station for maximum ambient temperature, equal to 2.39°C (GFS) versus 1.48°C (ART), and for minimum ambient temperature 1.52°C (GFS) versus 1.74°C (ART). Overall, WRF appears to behave well and slightly better than GFS. This difference is expected to be greater for other locations in the grid, as for this particular case the station happened to be very close to the GFS grid point. Furthermore, we calculated the MAE and Root Mean Squared Error (RMSE) of all daily 5-day forecasts of ART against the ground station for a period of interest (for graphical comparisons, see Appendix Figure 1). For the maximum temperature we found $MAE = 2.41$, $RMSE = 3.11$, whereas for the minimum temperature we found $MAE = 2.75$, $RMSE = 3.70$.

Real-World Case Study. We combined the ART weather forecasts and the conditions listed in Table 1 to produce a recommendation system in the form of daily maps over the fields of the farmers of the cooperative (Figure 1). The sowing recommendation maps were served through the website of the cooperative that farmers visited on a daily basis. The cooperative collected and provided for each field: its geo-referenced boundaries, the sowing date, the seed variety, the harvest date, the precise final yield, and for a subset of the fields the yield of the previous year. We then combined this data with publicly available observations from heterogeneous sources, such as satellites (Sentinel-2), weather stations and GIS maps, to engineer an observational dataset that enables a causal analysis for studying the impact of the recommendation system on the yield.

Causal Evaluation Framework

Notation and Terminology. We encode the farm system in the form of a Directed Acyclic Graph (DAG) $G \equiv (V, E)$ where V is a set of vertices consisting of all relevant variables, and E is a set of directed edges connecting them (Pearl 2009). The directed edge $A \rightarrow B$ indicates causation from A to B , in the sense that changing the value of A and holding everything else constant will change the value of B . We are using Pearl's *do*-operator to describe interventions, with $\mathbb{P}(Y = y | do(T = t))$ denoting the probability that $Y = y$ given that we intervene on the system by setting the value of T to t . Following popular terminology, we name the variable T , of which we aim to estimate the effect, as *treatment* and the variable Y , which we want to quantify the impact of T on, as *outcome*. The parents of a node are its *direct causes*, while a parent of both the treatment and outcome is referred to as a *common cause* or *confounder*. Our end goal is to account for exactly the variables $Z \subseteq V$ that will allow us to estimate the Average Treatment Effect

(ATE) of the treatment on outcome, as shown in Eq. (3).

$$\text{ATE} = \mathbb{E}[Y|do(T = 1)] - \mathbb{E}[Y|do(T = 0)] \quad (3)$$

Problem Formulation. We thus aim to develop a causal graph G whose vertices V capture the relevant actors of the system we study, and edges E indicate their relationships. The system recommendations should be part of the graph, along with cotton yield and the agro-environmental conditions that interfere in this physical process.

Because the end goal is the evaluation of the recommendation system and its actual impact on yield, we designate as *treated* the fields that farmers sowed on a day that was seen as favorable by the system, i.e., the corresponding 10-day ART forecasts satisfied the appropriate conditions, and as *control* the fields that were sown on a non-favorable day. Because the system outputs 4 levels of recommendation ranging from 0 (bad) to 3 (good), we define a day as favorable when all conditions are satisfied, i.e., when the system outputs the highest recommendation value that is 3. A day is then defined as non-favorable when the system outputs any other recommendation value. Binarizing the treatment in that way allows for greater flexibility in estimator selection and easier interpretation. Formally,

$$T = \begin{cases} 1 & , \text{ if recommendation of sowing date} \in \{3\} \\ 0 & , \text{ if recommendation of sowing date} \in \{0,1,2\} \end{cases}$$

Beyond the recommendation system, multiple factors influence the decision to sow or not. This is precisely the challenge we aim to address by employing a graphical analysis and explicitly modeling the farm system structure. The ATE we aim to estimate captures the difference between what the average yield would have been if we intervened and forced farmers to follow the recommendation by sowing on a favorable day, and the average yield if we forced them to defy the recommendation by sowing on an unfavorable day. Such an estimand is of primary significance for the farmers, but also for proving the reliability and therefore accelerating the adoption of smart farming tools. Given that confounding factors are controlled for, we henceforth refer to the ATE as the *(average) causal effect of following the recommendation* in the sense described above.

Cotton Domain Knowledge and Graph Building. Cotton yield and quality are ultimately determined by the interaction between the genotype, environmental conditions and management practices throughout the growing season. Nevertheless, the first pivotal steps for a profitable yield are a successful seed germination and emergence which are greatly dependent on timely sowing (Bauer, May, and Camberato 1998; Bradow and Bauer 2010).

Emergence and germination mediate the effect of T on Y ; however, Crop Growth (CG) was not observed. We thus turned to the popular Normalized Difference Vegetation Index (NDVI) in order to obtain a reliable proxy of CG , and specifically used the trapezoidal rule across NDVI values from sowing to harvest (Eklundh and Jönsson 2015). Even though in the case of cotton, trapezoidal NDVI is not linearly correlated with yield (Dalezios et al. 2001; Zhao et al.

2007), it is correlated with early season Leaf Area Index (LAI) (Zhao et al. 2007), which in turn is a good indicator of early season crop growth rate (Virk, Snider, and Pilon 2019). Furthermore, seed germination and seedling emergence are greatly dependent on soil moisture. Hence, soil moisture SM is a confounder for the relation $T \rightarrow Y$ that we study. As a SM proxy, we used the Normalized Difference Water Index (NDWI) at sowing day which is highly correlated with soil moisture in bare soil (Casamitjana et al. 2020).

Agricultural management practices before sowing (AbS) comprise tilling operations for preparing a good seedbed. Practices during sowing (AdS) include a sowing depth of 4 – 5 cm and an average distance of 0.91 m between rows and 7.62 cm between seeds. After sowing practices (AaS) comprise basic fertilization, irrigation and pest management. It is reasonable to think that all aforementioned practices are a result of a common cause that we can define as Agricultural Knowledge (AK), capturing the skills and experience of a farmer. We possess no quantitative information on the agricultural knowledge or the practices followed by each farmer. However, the farmer’s cooperative is not large, and aims for consistent, high-quality produce. As a result, they have developed highly consolidated routines for interacting with their crops: this includes common practices, homogeneous fertilizer application, and jointly owned machinery. We thus note that even if we do not have numerical data on AbS , AdS , AaS , the cooperative directors do not observe significant differences across fields and for the purposes of our study these variables are considered to be constant.

At the same time, it is rational to assume that the agricultural knowledge (AK) of any farmer interacts with crops exclusively through management practices. Because of the aforementioned condition, the influence of AK on the system is nullified and we hence omit it from the graph. While we note that the above limit the external validity of our results (Calder, Phillips, and Tybout 1982), by assuming that agricultural practices are constant for all farmers and that AK only interacts with the system through them, we implicitly control for all of them (Huntington-Klein 2021).

Apart from soil moisture, soil and ambient temperatures at the time of sowing and for 5-10 days after, affect seed germination, seedling development and final yield (Virk, Snider, and Pilon 2019; Boman and Lemon 2005; Varco 2020). Low temperatures result in reduced germination, slow growth and less vigorous seedlings that are more prone to diseases and sensitive to weed competition (Bradow and Bauer 2010). This knowledge is incorporated in the sowing recommendations, in the form of numerical rules, and consequently in the treatment T . We thus added in the graph the weather forecast WF (variables listed in Table 1) as a parent node of T . We also had access to the weather on the day of sowing WS (min & max ambient temperature in °C) from a nearby weather station, influencing WF , T , and CG .

Topsoil (0-20 cm) properties SP (% content of clay, silt and sand) and organic carbon content SoC (g C kg⁻¹) also affect cotton seed germination and seedling emergence due to differences in water holding capacity and consequently in soil temperature and aeration, drainage and seed-to-soil

contact (Varco 2020). Data on *SP* and *SoC* were retrieved from the European Soil Data Centre (ESDAC) (Ballabio, Panagos, and Monatanarella 2016; de Brogniez et al. 2015). Both variables were included in the graph as confounders of *T* and *CG*. Seed variety also determines seed germination, emergence and final yield (Snider, Pilon, and Virk 2020). Seed mass and vigor (Liu et al. 2015; Snider, Pilon, and Virk 2020) are related to the seed variety (*SV*); we hence added the latter as a confounder for *T* and *Y*. In this case, we had 13 different cotton SVs.

Id	Variable Description	Source
T	Treatment	Farmers' Cooperative, RS
WF	Weather forecast	GFS, WRF
WS	Weather on sowing day	Nearest weather station
WaS	Weather after sowing	Nearest weather station
CG	Crop Growth	NDVI via Sentinel-2
SM	Soil Moisture on sowing	NDWI via Sentinel-2
SP	Topsoil properties	Map by ESDAC
SoC	Topsoil organic carbon	Map by ESDAC
SV	Seed Variety	Farmers' Cooperative
G	Geometry of field	Farmers' Cooperative
AdS	Practices during sowing	Farmers' Cooperative
AbS	Practices before sowing	Farmers' Cooperative
AaS	Practices after sowing	Farmers' Cooperative
HD	Harvest Date	Farmers' Cooperative
Y	Outcome (Yield)	Farmers' Cooperative

Table 2: Farm system variable identifier, description and source. RS = Recommendation System.

The geometrical properties of the field (perimeter to area ratio, *G*) were also considered, as border effects can play a minor role on crop growth, confounding the effect of *T* on *Y* (Green 1956). Since temperature is the primary environmental factor controlling plant growth (Bange and Milroy 2004; Hatfield and Prueger 2015), temperature fluctuations were observed throughout the growing season from the nearest weather station, constituting a parent variable *WaS* (min & max ambient temperature in °C) of crop growth *CG*. Lastly, the Harvest Date (*HD*) mediates the effect of *CG* on *Y*, influencing both yield potential and quality (Dong et al. 2006; Bange, Caton, and Milroy 2008). Table 2 summarizes the variables' description, abbreviation and source.

Causal Graph. Figure 2 displays the final causal graph *G*. We note that, in reality, it is impossible to account for all factors interacting in the system in order to claim that the estimated effect will not contain any bias. However, because the selection of variables is deeply rooted on well-understood agro-environmental interactions, bias is expected to be minimized, in the sense that no important interactions are left unaccounted for. Furthermore, we extensively test the reliability of effect estimates through multiple refutation checks.

Effect Identification and Estimation. Because the calculation of causal effects requires access to counterfactual values that are by definition not observed (Holland 1986), observational methods rely on identification techniques and assumptions that aim at reducing causal estimands such as

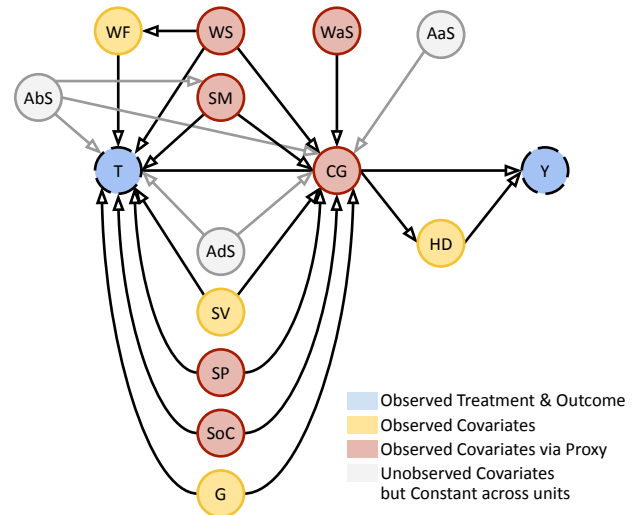


Figure 2: Graph of the farm system, encoding the causal relations between the relevant agro-environmental actors.

$\mathbb{P}(Y = y|do(T = t))$ to statistical ones, such as $\mathbb{P}(Y = y|T = t)$. The back-door criterion is a popular identification method that solely relies on a graphical test to infer whether adjusting for a set of graph nodes $Z \subseteq V$ is sufficient for identifying $\mathbb{P}(Y = y|do(T = t))$ from observational data. Formally, a set of variables Z satisfies the back-door criterion relative to an ordered pair of variables (T, Y) in a DAG G , if no node in Z is a descendant of T and Z blocks every path between T and Y that contains an arrow into T . After (if) we have obtained a back-door adjustment set to condition on, we can proceed with estimating the ATE of interest. The back-door criterion already provides a formula for the interventional distribution. Given a set of variables Z satisfying the back-door criterion we can identify the causal effect of T on Y as $\mathbb{P}(y|do(t)) = \sum_z \mathbb{P}(y|t, z)\mathbb{P}(z)$.

In our study, ATE estimation is done with several methods of varying complexity. To check covariate balance and as a method prerequisite, we model the propensity scores $\mathbb{P}(T = 1|Z = z)$, i.e., the probability of receiving treatment given features (Rosenbaum and Rubin 1983). Linear regression and distance matching are selected as baseline estimation methods. The popular Inverse Propensity Score (IPS) weighting is also used (Stuart 2010). We finally apply modern machine learning methods, i.e., the baseline T-learner and the state-of-the-art X-learner (Künzel et al. 2019).

Refutation Methods. One of the biggest challenges in causal inference pertains to model evaluation. Given the fact that ground truth estimates are not observed, we resort to performing robustness checks and sensitivity analyses of estimates, in line with recent research (Sharma and Kiciman 2020; Cinelli and Hazlett 2020). We perform the following tests: i) Placebo treatment, where the treatment is randomly permuted and the estimated effect is expected to drop to 0; ii) Random Common Cause (RCC), where a random confounder is added to the dataset and the estimate is expected to remain unchanged; iii) Random Subset Removal (RSR),

where a subset of data is randomly selected and removed and the effect is expected to remain the same; iv) Unobserved Common Cause (UCC), where an unobserved confounder acts on the treatment and outcome without being added to the dataset, and the estimates should remain relatively stable. The Placebo, RCC and RSR tests are bootstrapped to generate confidence intervals and p-values (DiCiccio and Efron 1996). The UCC returns a heatmap of new ATE estimates depending on the strength of unobserved confounding.

Experiments, Results and Discussion

The sowing period lasted from early April to early May 2021, the harvest took place from mid to late September, and the yield per hectare ranged from 1,250 to 6,960 kg (Figure 2 of Appendix contains relevant histograms). The dataset consists of 171 fields (51 treated and 120 control). Variables that registered intra-field values (NDVI, NDWI) were averaged at the field-level. For the experiments, we are using the popular doWhy (Sharma and Kiciman 2020) and Causal ML (Chen et al. 2020) Python libraries. Our implementation and dataset are publicly available (Tsoumas 2023).

Applying the back-door criterion on graph G (Figure 2), the following adjustment set of nodes Z was found to be sufficient for identifying the ATE:

$$Z = \{WS_{MIN, MAX}, SOC, SM, G, SP_{SILT, CLAY, SAND}, ABS, ADS, SV_{1-13}\} \quad (4)$$

All variables in Z are numerical, including the one-hot encoded vectors of the categorical SV_{1-13} variable of seed variety. AbS and AdS are constant, and thus excluded for estimation purposes. We scale the data by subtracting the mean from each variable and dividing by its standard deviation. The treatment T is binarized, with 1 indicating that a farmer sowed on a favorable day, and 0 indicating the opposite.

Propensity modeling is a prerequisite of IPS weighting. We thus begin by discussing the propensity model that is fit. Given the relatively small dataset size, logistic regression is used on the scaled back-door adjustment set Z for classifying each field into the treatment/control group. We subsequently trim the dataset by removing all rows with extreme propensity scores (< 0.2 or > 0.8) to aid the overlap assumption (Imbens and Rubin 2015). The resulting distribution of propensity scores can be seen at Figure 3. The model scores 0.81 in accuracy, 0.64 in F1-score, and 0.88 in ROC-AUC. After trimming extreme propensity scores, a subset of 48 treated and 37 control units remains. There is decent overlap between the propensity score distributions of the treatment and control group, indicating that they are comparable and enabling reliable propensity-based ATE estimation.

Table 3 and Figure 4 show the results of the ATE estimation per method, alongside the corresponding 95% confidence intervals and p-values. Besides Linear Regression, other methods do not provide confidence intervals by default. For matching, IPS, and meta-learners confidence intervals and the resulting p-values are hence bootstrapped. Both the T-learner and X-learner use a Random Forest for modeling the outcome Y .

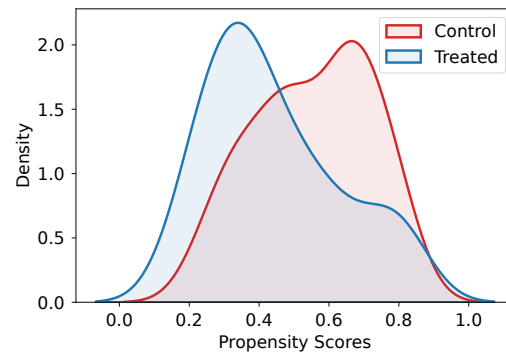


Figure 3: Distribution of propensity scores for the control and treatment group after trimming extreme scores.




Figure 4: ATE point estimates and 95% confidence intervals for all estimation methods.

All methods detect a significant ATE at 95% confidence level, with point estimates ranging from 372 to 546 kilograms of cotton per hectare. For context, the average observed yield is 3,145 kg/ha. We thus infer that the causal effect of following the sowing recommendation on yield is significantly positive, driving a yield increase ranging from 12% to 17%, depending on the estimation method used.

Of central importance are the refutation tests we run after having estimated the recommendation impact. Table 3 features analytic results for all method / refutation test combinations. All estimation methods are robust against performing the following data manipulations and re-estimating the ATE: randomly permuting the treatment (Placebo test), adding a confounder (RCC test), sampling a subset of data (RSR test) and creating unobserved confounding (UCC test). Specifically, Placebo ATE estimates do not differ significantly from 0, while RCC and RSR estimates do not differ significantly from the already obtained ATE. For the UCC test, the mean ATE estimates are reduced yet remain positive, despite unobserved confounding of significant magnitude. Confidence intervals and p-values are bootstrapped (1000 iterations).

The results indicate that the recommendation system’s advice drove a net increase in yield that was deemed both significant and robust from a statistical perspective. By utilizing

Causal Effect Estimation				Refutations						
				Placebo		RCC		UCC	RSR	
Method	ATE	CI	p-value	Effect*	p-value	Effect*	p-value	Effect*	Effect*	p-value
Linear Regression	546	(211, 880)	0.0015	-25.74	0.39	546	0.49	85	543	0.45
Matching	448	(186, 760)	0.0060	50.82	0.39	432	0.40	116	438	0.48
IPS weighting	471	(138, 816)	0.0010	38.82	0.40	470	0.40	113	462	0.45
T-Learner (RF)	372	(215, 528)	0.0240	9.26	0.49	373	0.46	-	353	0.42
X-Learner (RF)	437	(300, 574)	0.0050	5.10	0.50	430	0.37	-	409	0.36

Table 3: Results of Average Treatment Effect estimation. Includes point estimates, 95% confidence intervals, and four refutation tests. For the Placebo, RCC and RSR refutations, the new ATE estimate is reported (denoted as Effect*), alongside the respective p-value (< 0.05 indicates a failed test). The UCC column reports the mean ATE estimate of the corresponding heatmap (for full heatmaps and details see Sec. 2 of Appendix). Numbers are in cotton kg/ha, rounded to the nearest integer.

the theory of graphical causal models, the analysis transparently puts forward its assumptions and explicitly incorporates domain knowledge in it. Combined with accurate and performant systems, such analyses can benefit the reliability and adoption of digital agriculture as well as farmers’ trust. The provision of information on the actual impact expected from a recommendation system may also enable a cost-benefit analysis on behalf of the farmer, by simply comparing the digital tool cost to the expected yield gain.

Even though the analysis is transparent, it is as good as the causal assumptions it makes and the DAG it develops. Our graph is consistent with agri-environmental knowledge on cotton, however there is always a possibility that bias exists, either due to a missed confounder or due to a missed interaction between observed variables. The robustness checks we performed were all successful; noting that when we add strong unobserved confounding, the UCC test estimates become volatile - an expected behavior to a certain degree.

Given the homogeneous management practices among farmers in our data, we remark that external validity of estimates is low, as results cannot be expected to generalize to other farms that might follow different routines. Nevertheless, it is not uncommon for farmers to follow similar practices in other regions or even entire countries. While the transfer of effect estimates warrants caution, the same does not hold for the proposed framework itself. Given relevant data and knowledge, a graph-based empirical evaluation of an agricultural recommendation system can normally proceed. If consensus among estimation methods in terms of ATE significance is reached, the tool is deemed beneficial; otherwise, more work is required. All in all, this system equips farmers with a provably valuable tool based on cotton knowledge and weather forecasts. It contributes to a successful growing season and lowers the likelihood of farmers resorting to expensive actions, e.g., replanting a field.

For the growing season of 2022, the recommendation system was deployed at national scale and extended to two other crops (maize, sunflower). These new pilot applications will allow us to practically test the external validity of our results across different seasons, crops and locations. Moreover, given the developed causal graph G , the crop growth (CG) variable that is sufficiently captured through its NDVI proxy, mediates the effect of T on Y . The front-door crite-

riion (Pearl 2009) might thus provide an alternative identification method for the ATE, and we plan on exploring it in collaboration with domain experts. Finally, the more growing seasons the recommendation system has seen, the more data are obtained. Going beyond ATE estimation by learning Conditional ATEs and using causal machine learning methods for providing personalized effect estimates is another next step. Due to the rich and well-established domain knowledge, we finally believe that the potential of causal reasoning in agriculture extends far beyond effect identification. Fitting Structural Causal Models and performing counterfactual inference can enable a greater understanding of the farm system and supercharge decision support tools.

Most generally, the essential condition that allowed us to utilize causal inference for empirically evaluating an agricultural recommendation system is the ample, long-established domain knowledge that exists. Decision support systems are being used on multiple fields (Marakas 2003) such as medical decision making (Sutton et al. 2020) or forest and fire management (Segura, Ray, and Maroto 2014; Martell 2015). The aforementioned fields possess accumulated domain knowledge on the interactions a good system exploits; the same way we possess information on environmental conditions related to cotton planting. We thus expect graphical approaches to be valuable for the empirical evaluation of decision support systems of diverse domains.

Conclusion

In this study, we design, implement, and test a digital agriculture recommendation system for the optimal sowing of cotton. Using the collected data and leveraging domain knowledge, we evaluate the impact of system recommendations on yield. To do so, we utilize and propose causal inference as an ideal tool for empirically evaluating decision support systems. This idea can be upscaled to other digital agriculture tools as well as to different fields with well-established domain knowledge. This paradigm is in principle different to decision support systems that frequently use black-box algorithms to predict variables of interest, but are oblivious to the evaluation of their own impact. In that sense, this work comes to the defence of the farmer, by introducing an AI framework for elaborating on the assumptions, reliability, and impact of a system before discussing service fees.

Acknowledgements

We thank the Orchomenos Agriculture Cooperative for collaboration and data provision, and Corteva Hellas for support. This research was funded by the EU H2020 program through the “EuroGEO Showcases: Applications Powered by Europe” (e-shape) project (Grant agreement ID: 820852). It was also funded by the General Secretariat for Research and Innovation (GSRI, Greece) under the Action ERANETS 2021A [Call ID: 037KE - A/A MIS 4888] (Project Number: T12ERA5-00075) and through the 2019-2020 BiodivERsA joint call for research proposals [BiodivClim ERA-Net CO-FUND programme, MICROSERVICES project].

References

- Adelman, L. 1992. *Evaluating decision support and expert systems*. Wiley-Interscience.
- Aggarwal, C. C. 2016. Knowledge-based recommender systems. In *Recommender systems*, 167–197. Springer.
- Arif, S.; and MacNeil, M. A. 2022. Utilizing causal diagrams across quasi-experimental approaches. *Ecosphere*, 13(4): e4009.
- Ballabio, C.; Panagos, P.; and Monatanarella, L. 2016. Mapping topsoil physical properties at European scale using the LUCAS database. *Geoderma*, 261: 110–123.
- Bange, M.; and Milroy, S. 2004. Impact of short-term exposure to cold night temperatures on early development of cotton (*Gossypium hirsutum* L.). *Australian journal of agricultural research*, 55(6): 655–664.
- Bange, M. P.; Caton, S. J.; and Milroy, S. P. 2008. Managing yields of high fruit retention in transgenic cotton (*Gossypium hirsutum* L.) using sowing date. *Australian Journal of Agricultural Research*, 59(8): 733–741.
- Barbato, L.; Seshadri, S.; and Mulligan, K. 2011. Temperature Forecast and Cotton Seed Planting Interactive Map. <https://agriculture.basf.us/crop-protection/products/seeds/cotton-planting-forecasting.html>. Accessed: 2021-03-15.
- Bauer, P. J.; May, O. L.; and Camberato, J. J. 1998. Planting date and potassium fertility effects on cotton yield and fiber properties. *Journal of production agriculture*, 11(4): 415–420.
- Boman, R.; and Lemon, R. 2005. Soil temperatures for cotton planting. *Texas Coop. Ext. Bull. SCS-2005-17*.
- Boruch, R. F. 1997. *Randomized experiments for planning and evaluation: A practical guide*, volume 44. Sage.
- Bradow, J. M.; and Bauer, P. J. 2010. Germination and seedling development. In *Physiology of cotton*, 48–56. Springer.
- Calder, B. J.; Phillips, L. W.; and Tybout, A. M. 1982. The concept of external validity. *Journal of consumer research*, 9(3): 240–244.
- Casamitjana, M.; Torres-Madroño, M. C.; Bernal-Riobo, J.; and Varga, D. 2020. Soil moisture analysis by means of multispectral images according to land use and spatial resolution on andosols in the Colombian Andes. *Applied Sciences*, 10(16): 5540.
- Chen, H.; Harinen, T.; Lee, J.-Y.; Yung, M.; and Zhao, Z. 2020. Causalml: Python package for causal machine learning. *arXiv preprint arXiv:2002.11631*.
- Choumos, G.; Koukos, A.; Sitokoustantinou, V.; and Kontoes, C. 2022. Towards Space-to-Ground Data Availability for Agriculture Monitoring. *arXiv preprint arXiv:2205.07721*.
- Cinelli, C.; Forney, A.; and Pearl, J. 2020. A Crash Course in Good and Bad Controls. *SSRN Electronic Journal*.
- Cinelli, C.; and Hazlett, C. 2020. Making sense of sensitivity: Extending omitted variable bias. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 82(1): 39–67.
- Dalezios, N.; Domenikiotis, C.; Loukas, A.; Tzortzios, S.; and Kalaitzidis, C. 2001. Cotton yield estimation based on NOAA AVHRR produced NDVI. *Physics and Chemistry of the Earth, Part B: Hydrology, Oceans and Atmosphere*.
- de Brogniez, D.; Ballabio, C.; Stevens, A.; Jones, R.; Montanarella, L.; and van Wesemael, B. 2015. A map of the topsoil organic carbon content of Europe generated by a generalized additive model. *European Journal of Soil Science*, 66(1): 121–134.
- Deines, J. M.; Wang, S.; and Lobell, D. B. 2019. Satellites reveal a small positive yield effect from conservation tillage across the US Corn Belt. *Environmental Research Letters*, 14(12): 124038.
- DiCiccio, T. J.; and Efron, B. 1996. Bootstrap confidence intervals. *Statistical science*, 11(3): 189–228.
- Diggle, P. J.; and Chetwynd, A. G. 2011. Experimental design: agricultural field experiments and clinical trials. In *Statistics and Scientific Method*. Oxford University Press.
- Dong, H.; Li, W.; Tang, W.; Li, Z.; Zhang, D.; and Niu, Y. 2006. Yield, quality and leaf senescence of cotton grown at varying planting dates and plant densities in the Yellow River Valley of China. *Field Crops Research*, 98(2-3): 106–115.
- Eklundh, L.; and Jönsson, P. 2015. TIMESAT: A software package for time-series processing and assessment of vegetation dynamics. In *Remote sensing time series*, 141–158. Springer.
- Fougère, D.; and Jacquemet, N. 2019. Causal inference and impact evaluation. *Economie et Statistique/Economics and Statistics*.
- Freeland Jr, T. B.; Pettigrew, B.; Thaxton, P.; and Andrews, G. L. 2006. Agrometeorology and cotton production. *World Meteorological Organization*.
- Gabriel, A.; and Gandorfer, M. 2022. Adoption of digital technologies in agriculture—an inventory in a European small-scale farming region. *Precision Agriculture*, 1–24.
- Ge, M.; Delgado-Battenfeld, C.; and Jannach, D. 2010. Beyond accuracy: evaluating recommender systems by coverage and serendipity. In *Proceedings of the fourth ACM conference on Recommender systems*, 257–260.
- Giannarakis, G.; Sitokoustantinou, V.; Lorilla, R. S.; and Kontoes, C. 2022a. Personalizing Sustainable Agriculture with Causal Machine Learning. *arXiv preprint arXiv:2211.03179*.
- Giannarakis, G.; Sitokoustantinou, V.; Lorilla, R. S.; and Kontoes, C. 2022b. Towards assessing agricultural land suitability with causal machine learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1442–1452.
- Green, J. M. 1956. Border Effects in Cotton Variety Tests 1. *Agronomy Journal*, 48(3): 116–118.
- Hatfield, J. L.; and Prueger, J. H. 2015. Temperature extremes: Effect on plant growth and development. *Weather and climate extremes*, 10: 4–10.
- Hernán, M. A.; and Robins, J. M. 2016. Using big data to emulate a target trial when a randomized trial is not available. *American journal of epidemiology*, 183(8): 758–764.
- Holland, P. W. 1986. Statistics and causal inference. *Journal of the American statistical Association*, 81(396): 945–960.
- Huang, J. 2016. Different sowing dates affected cotton yield and yield components. *International Journal of Plant Production*, 10(1): 63–83.
- Huntington-Klein, N. 2021. *The effect: An introduction to research design and causality*. Chapman and Hall/CRC.

- Imbens, G. W.; and Rubin, D. B. 2015. *Causal inference in statistics, social, and biomedical sciences*. Cambridge University Press.
- Kerby, T.; Keeley, M.; and Johnson, S. 1989. Weather and seed quality variables to predict cotton seedling emergence. *Agronomy Journal*, 81(3): 415–419.
- Künzel, S. R.; Sekhon, J. S.; Bickel, P. J.; and Yu, B. 2019. Metalearners for estimating heterogeneous treatment effects using machine learning. *Proceedings of the National Academy of Sciences*, 116(10): 4156–4165.
- Lioutas, E. D.; Charatsari, C.; and De Rosa, M. 2021. Digitalization of agriculture: a way to solve the food problem or a trolley dilemma? *Technology in Society*, 67: 101744.
- Liu, S.; Remley, M.; Bourland, F.; Nichols, R.; Stevens, W.; Jones, A. P.; and Fritschi, F. 2015. Early vigor of advanced breeding lines and modern cotton cultivars. *Crop Science*, 55(4): 1729–1740.
- Lowenberg-DeBoer, J.; and Erickson, B. 2019. Setting the record straight on precision agriculture adoption. *Agronomy Journal*, 111(4): 1552–1569.
- Luma-Osmani, S.; Ismaili, F.; Raufi, B.; and Zenuni, X. 2020. Causal Reasoning Application in Smart Farming and Ethics: A Systematic Review. *Annals of Emerging Technologies in Computing (AETiC)*, Print ISSN, 2516–0281.
- Marakas, G. M. 2003. *Decision support systems in the 21st century*, volume 134. Prentice Hall Upper Saddle River.
- Martell, D. L. 2015. A review of recent forest and wildland fire management decision support systems research. *Current Forestry Reports*, 1(2): 128–137.
- Massmann, A.; Gentine, P.; and Runge, J. 2021. Causal inference for process understanding in Earth sciences. *arXiv preprint arXiv:2105.00912*.
- McFadden, J. R.; Rosburg, A.; and Njuki, E. 2022. Information inputs and technical efficiency in midwest corn production: evidence from farmers' use of yield and soil maps. *American Journal of Agricultural Economics*, 104(2): 589–612.
- Müller, J. 1974. On sources of measured technical efficiency: the impact of information. *American Journal of Agricultural Economics*, 56(4): 730–738.
- Munier, D.; Goodell, P.; Strand, J.; et al. 2004. Accuracy of cotton-planting forecasts assessed in the San Joaquin Valley. *California agriculture*, 58(3): 164–168.
- Nanushi, O.; Sitokonstantinou, V.; Tsoumas, I.; and Kontoes, C. 2022. Pest presence prediction using interpretable machine learning. *arXiv preprint arXiv:2205.07723*.
- NCEP, N., NWS; and USDOC. 2015. NCEP GFS 0.25 Degree Global Forecast Grids Historical Archive. <https://doi.org/10.5065/D65D8PWK>. Accessed: 2022-05-20.
- Nielsen, R. L.; Thomison, P. R.; Brown, G. A.; Halter, A. L.; Wells, J.; and Wuethrich, K. L. 2002. Delayed planting effects on flowering and grain maturation of dent corn. *Agronomy Journal*.
- Pasquel, D.; Roux, S.; Richetti, J.; Cammarano, D.; Tisseyre, B.; and Taylor, J. A. 2022. A review of methods to evaluate crop model performance at multiple and changing spatial scales. *Precision Agriculture*, 1–25.
- Pearl, J. 2009. *Causality*. Cambridge university press.
- Pu, P.; Chen, L.; and Hu, R. 2012. Evaluating recommender systems from the user's perspective: survey of the state of the art. *User Modeling and User-Adapted Interaction*, 22(4): 317–355.
- Qian, S. S.; and Harmel, R. D. 2016. Applying statistical causal analyses to agricultural conservation: A case study examining P loss impacts. *JAWRA Journal of the American Water Resources Association*, 52(1): 198–208.
- Richards, M. F.; Maphosa, L.; and Preston, A. L. 2022. Impact of Sowing Time on Chickpea (*Cicer arietinum* L.) Biomass Accumulation and Yield. *Agronomy*, 12(1): 160.
- Roberts, M. J.; Schimmelpfennig, D.; Livingston, M. J.; and Ashley, E. 2009. Estimating the value of an early-warning system. *Applied Economic Perspectives and Policy*, 31(2): 303–329.
- Rosenbaum, P. R.; and Rubin, D. B. 1983. The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70(1): 41–55.
- Runge, J.; Bathiany, S.; Bollt, E.; Camps-Valls, G.; Coumou, D.; Deyle, E.; Glymour, C.; Kretschmer, M.; Mahecha, M. D.; Muñoz-Marí, J.; et al. 2019. Inferring causation from time series in Earth system sciences. *Nature communications*, 10(1): 1–13.
- Schimmelpfennig, D.; and Ebel, R. 2016. Sequential adoption and cost savings from precision agriculture. *Journal of Agricultural and Resource Economics*, 97–115.
- Schröder, G.; Thiele, M.; and Lehner, W. 2011. Setting goals and choosing metrics for recommender system evaluations. In *UCER-STI2 workshop at the 5th ACM conference on recommender systems, Chicago, USA*, volume 23, 53.
- Segura, M.; Ray, D.; and Maroto, C. 2014. Decision support systems for forest management: A comparative analysis and assessment. *Computers and Electronics in Agriculture*, 101: 55–67.
- Sharma, A.; Jain, A.; Gupta, P.; and Chowdary, V. 2020. Machine learning applications for precision agriculture: A comprehensive review. *IEEE Access*, 9: 4843–4873.
- Sharma, A.; and Kiciman, E. 2020. DoWhy: An end-to-end library for causal inference. *arXiv preprint arXiv:2011.04216*.
- Skamarock, W. C.; Klemp, J. B.; Dudhia, J.; Gill, D. O.; Liu, Z.; Berner, J.; Wang, W.; Powers, J. G.; Duda, M. G.; Barker, D. M.; et al. 2019. A description of the advanced research WRF model version 4. *NCAR: Boulder, CO, USA*, 145: 145.
- Snider, J. L.; Pilon, C.; and Virk, G. 2020. Seed Characteristics and Seedling Vigor. In *Cotton Seed and Seedlings*, 9–23. The Cotton Foundation.
- Stuart, E. A. 2010. Matching methods for causal inference: A review and a look forward. *Statistical science: a review journal of the Institute of Mathematical Statistics*, 25(1): 1.
- Sutton, R. T.; Pincock, D.; Baumgart, D. C.; Sadowski, D. C.; Fedorak, R. N.; and Kroeker, K. I. 2020. An overview of clinical decision support systems: benefits, risks, and strategies for success. *NPJ digital medicine*, 3(1): 1–10.
- Tsoumas, I. 2023. Code&Data. <https://github.com/Agri-Hub/AAAI23-Eval-AgriRecommendations>. Accessed: 2023-02-10.
- UC IPM. 2004. Weather Products: Cotton Planting Forecast. ipm.ucanr.edu/WEATHER/cottonforecast.html. Accessed: 2022-05-10.
- Vaessen, J. 2010. Challenges in impact evaluation of development interventions: opportunities and limitations for randomized experiments. IOB Discussion Papers 2010.01, Universiteit Antwerpen, Institute of Development Policy (IOB).
- Varco, J. J. 2020. Soil Climatic and Edaphic Effects on Cotton Germination and the Final Stand. In *Cotton Seed and Seedlings*, 43–53. The Cotton Foundation.
- Virk, G.; Snider, J. L.; and Pilon, C. 2019. Physiological contributors to early season whole-crop vigor in cotton. *Crop Science*, 59(6): 2774–2783.
- Zhai, Z.; Martínez, J. F.; Beltran, V.; and Martínez, N. L. 2020. Decision support systems for agriculture 4.0: Survey and challenges. *Computers and Electronics in Agriculture*, 170: 105256.
- Zhao, D.; Reddy, K. R.; Kakani, V. G.; Read, J. J.; and Koti, S. 2007. Canopy reflectance in cotton for growth assessment and lint yield prediction. *European Journal of Agronomy*, 26(3): 335–344.