

# Towards assessing agricultural land suitability with causal machine learning

Georgios Giannarakis<sup>1</sup>

Vasileios Sitokonstantinou<sup>1,2</sup>

Roxanne Suzette Lorilla<sup>1</sup>

Charalampos Kontoes<sup>1</sup>

<sup>1</sup>BEYOND Center, IAASARS, National Observatory of Athens, Greece

<sup>2</sup>Remote Sensing Laboratory, National Technical University of Athens, Greece

{giannarakis, vsito, rslorilla, kontoes}@noa.gr

## Abstract

*Understanding the suitability of agricultural land for applying specific management practices is of great importance for sustainable and resilient agriculture against climate change. Recent developments in the field of causal machine learning enable the estimation of intervention impacts on an outcome of interest, for samples described by a set of observed characteristics. We introduce an extensible data-driven framework that leverages earth observations and frames agricultural land suitability as a geospatial impact assessment problem, where the estimated effects of agricultural practices on agroecosystems serve as a land suitability score and guide decision making. We formulate this as a causal machine learning task and discuss how this approach can be used for agricultural planning in a changing climate. Specifically, we extract the agricultural management practices of “crop rotation” and “landscape crop diversity” from crop type maps, account for climate and land use data, and use double machine learning to estimate their heterogeneous effect on Net Primary Productivity (NPP), within the Flanders region of Belgium from 2010 to 2020. We find that the effect of crop rotation was insignificant, while landscape crop diversity had a small negative effect on NPP. Finally, we observe considerable effect heterogeneity in space for both practices and analyze it.*

## 1. Introduction

One of the greatest challenges faced by humankind is producing and supplying food to a rapidly growing population in a climate changing world [47]. Given the increasing demand for natural resources, the expansion of croplands exerts substantial pressure on natural ecosystems. Ecosystem prosperity can be compromised by complex human-nature dynamics, risking human well-being itself [24].

The inclusion of complex causal relationships in environmental decision and policy making processes is key for policy implementation under sustainable management

regimes [32, 71]. Towards this direction, recent research utilizes environmental causal analyses to evaluate potential causes driving an observed or hypothesized change in specific target metrics [37]. In this paper, we discuss causal inference and machine learning in agricultural policy making that aims at increasing productivity to meet the global food requirements, while ensuring ecosystem resilience.

Predictive Machine Learning (ML) models can be inefficient for agricultural policy making since they are based on correlations and are not inherently capable of addressing causal questions that a policy maker is after [31]. Even when causal analyses are employed, they are usually based on localized experiments on few samples and cannot account for the spatial variability caused by changes in climatic conditions, management practices and other factors [18]. Policies are often horizontally implemented, leading to the lack of spatial targeting among areas with different ecological characteristics [16]. To mitigate these issues, the Common Agricultural Policy (CAP) of the European Union (EU) proposes the introduction of agri-environmental measures and eco-schemes that are tailored and optimized to the specificities of different regions [63].

Every policy measure aims at implementing an agricultural management practice to achieve a certain goal. In order to assess the efficiency of policy measures for agroecological resilience, policy makers need to know whether particular management practices can achieve a given set of objectives [41]. Specifically, it is required to know if and how practices are effective at influencing key target metrics, such as the supply of ecosystem services or climate change mitigation and adaptation [54]. There are several studies that approach policy evaluation in experimental settings [14]. Nevertheless, large-scale experiments on agricultural policies can be challenging and, thus, the only way to assess their impact is through observational data [19].

Land and crop information derived from satellite images, and environmental data derived from numerical simulations can both capture large areas with high frequency and at high spatial resolutions. Machine learning based causal infer-

ence on observed and simulated data is a promising complementary approach to field experiments, as it enables the detailed assessment of the heterogeneous effects of agricultural practices over millions of fields. The exploitation of this kind of data ensures high spatial representation, accounting for regional differences in climatic, soil and crop conditions. Nevertheless, it should be noted that causal inference on observational data is subject to biases from measurement, selection and confounding [49,52].

The evaluation of the agricultural practices of Crop Rotation - CR (growing different crops across a sequence of growing seasons) and Landscape Crop Diversity - LCD (increasing the number and evenness of crops grown in a landscape) can be used to allocate them in space. Evaluating such practices by their impact on agroecosystem productivity enables implementing resilient farming systems in the light of climate change and high consumer demand [21,23].

**Contributions.** We demonstrate the applicability of causal machine learning, specifically Double Machine Learning (DML), to estimate the impact that the LCD and CR agricultural practices had on Net Primary Productivity (NPP), which we use as an ecosystem service proxy of climate regulation [57]. By deriving treatment effects and analyzing their heterogeneity, we infer a data-driven and context-aware agricultural land suitability score for both practices. We thus propose a flexible observation-based agricultural land suitability framework that can be extended to include any management practice and agroecosystem metric that practitioners might have data on. Our approach, illustrated in Figure 1, was tested in Flanders for the period 2010-2020, where we studied the practices' effect heterogeneity and discussed results in the context of local characteristics.

## 2. Related Work

**Agricultural land suitability.** Understanding agricultural land suitability is critical for policy applications aiming to foster sustainable agricultural practices and ensure food security. According to Akpoti *et al.* [3], the analysis of land suitability depends on many factors, including the purpose of the assessment and data availability, and therefore there is no one-size-fits-all approach. Studied techniques are commonly categorized into traditional and modern approaches. Traditional approaches include quantitative and qualitative methods accounting for biophysical characteristics of land [17,22]. On the other hand, modern methods exploit a plethora of variables and, according to Mugiyo *et al.* [46], are categorized into i) computer-assisted overlay mapping, ii) geo-computation or machine learning and iii) Multi-Criteria Decision Making (MCDM). Jayasinghe *et al.* [42] used the Analytical Hierarchy Process (AHP), which is the most common MCDM method in literature, and the Decision-Making Trail and Evaluation Laboratory (DEMATEL) model to assess the suitability of land for tea

crops. Using DEMATEL they were able to visualize inter-relationships between factors in a causal diagram. Here, we explicitly frame agricultural land suitability as a causal machine learning task, attempting to spatially optimize a metric of interest via intervening on the land use; and to the best of our knowledge this is the first work that does so.

**Heterogeneous Treatment Effects.** The Average Treatment Effect (ATE) has traditionally been the main concern of causal inference. However, it does not necessarily convey all information needed, particularly in cases where treatment effects vary systematically with sample characteristics. This is frequently the case in applications, and treatment effect heterogeneity has been studied in diverse fields such as medicine, social sciences, and economics [25,29,36,38]. Such methods have also been researched in the context of earth and environmental sciences, where heterogeneity of effects is present. Effect magnitude is determined by various factors, including geography, climate and land cover. For instance, Serra-Burriel *et al.* [56] estimated the heterogeneous effects of wildfires and Deines *et al.* [18] found significant variability in the effects of conservation tillage on yield. Both studies used earth observations to uncover the heterogeneity of treatment effects.

**Causal Machine Learning.** While statistical methods can be used for the estimation of heterogeneous treatment effects [15,43,66], there has been a recent surge of machine learning methods tailored for the same task [11,20,34,39,64]. Such methods are suitable for high-dimensional data and are able to learn general functional forms, thus taking advantage of the ever-increasing volume of data. Causal machine learning was recently used in environmental studies utilizing both satellite observations and environmental data. The interaction of aerosols and clouds was studied in Jesson *et al.* [33] with causal neural networks [34]. In the context of agriculture, the investigation of the effects of conservation tillage on yield in Deines *et al.* [18] was carried out with causal forests [5,65].

By employing remote sensing data and numerical predictions, one can achieve large spatio-temporal coverage that is much needed for effect heterogeneity to manifest. We also took advantage of the scale of earth observations and climate data and combined them with data that reflect historical agricultural use. We were then able to derive the LCD and CR practices, and estimate their heterogeneous impact on the climate regulation for each land sample.

## 3. Data & Methods

We first provide a brief overview of the theory of Conditional Average Treatment Effects in Sec. 3.1. We then provide details on the data that were used in our approach, as illustrated in Figure 1. In Sec. 3.2, we describe the data derived from crop type maps, including i) the LCD and CR agricultural practices that are used as treatments in our anal-

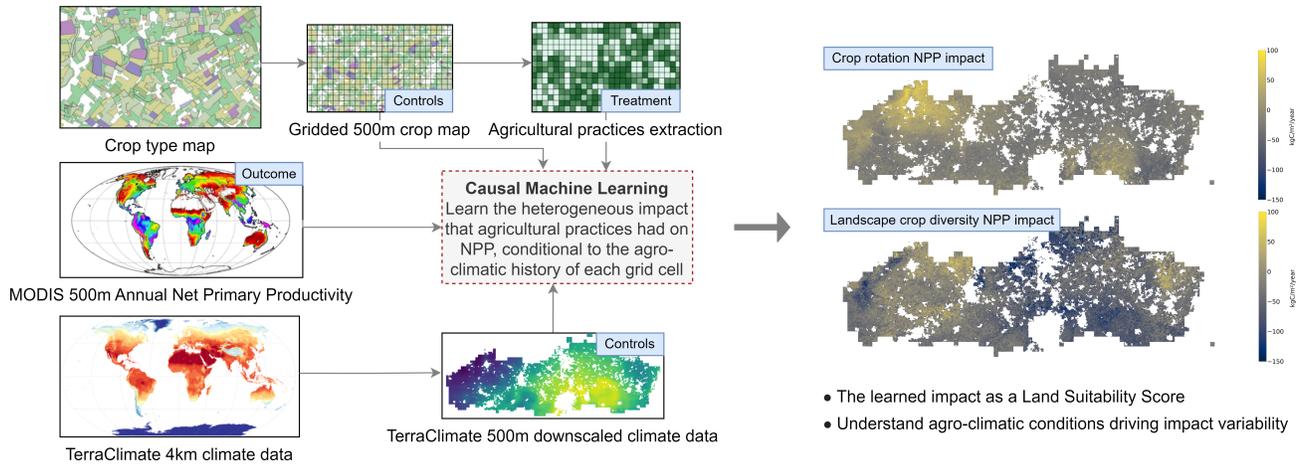


Figure 1. An overview of our process for learning the heterogeneous impact of Crop Rotation (CR) and Landscape Crop Diversity (LCD) on Net Primary Productivity (NPP). The study period extends from 2010 to 2020. First, the parcel-level crop type maps are transformed to gridded maps of 500m cells to match the pixel size of MODIS. For any given year and grid cell, by considering the percentage area of the cell that a crop type occupies, a crop abundance value for all major crop types is generated. Based on this, for each grid cell we compute i) the LCD practice using the Shannon diversity index and ii) the CR practice via quantifying the change between consecutive years in crop abundances within each cell, and summing for the whole study period. Our input data also include environmental variables from the TerraClimate dataset and MODIS NPP values for each cell. We run two separate analyses, one for each practice. Considering the value of each practice as the treatment, NPP as the outcome, and controlling for relevant crop abundances and environmental factors, the dataset is modeled with double machine learning. The heterogeneous impact of a practice per grid cell is studied and used as a land suitability score.

ysis and ii) the crop type abundance features that are used as controls. In Sec. 3.3, we provide information on the NPP, which is the outcome variable in our heterogeneous treatment effects analysis, and in Sec. 3.4 we describe the environmental factors that are also used as controls. The section concludes with Sec. 3.5 and the methodological setup.

### 3.1. Conditional Average Treatment Effects

**Terminology and Assumptions.** Using the potential outcomes framework [52], let  $Y(T)$  denote the value (potential outcome) of a random variable  $Y$  if we were to treat a unit with a binary treatment  $T \in \{0, 1\}$ . Given a vector of features  $X$  describing the units, we want to estimate:

$$\theta(x) = \mathbb{E}[Y(1) - Y(0)|X = x] \quad (1)$$

This heterogeneous treatment effect is referred to as the Conditional Average Treatment Effect (CATE) [2]. There are three important assumptions:

Overlap	$0 < \mathbb{P}(T = 1 X = x) < 1 \quad \forall x$
Unconfoundedness	$(Y(1), Y(0)) \perp\!\!\!\perp T X$
Consistency	$T = t \implies Y = Y(t)$

Overlap states that for all feature vectors in the population of interest, receiving treatment is possible but not certain. Unconfoundedness assumes that, conditional on  $X$ , potential outcomes are independent of the treatment. Consistency

postulates that, if the treatment is  $t$ , the observed outcome  $Y$  is actually the potential outcome under  $t$ . Provided these assumptions hold, CATE is identifiable from observational data and equals the following statistical estimand [33]:

$$\theta(x) = \mathbb{E}[Y|T = 1, X = x] - \mathbb{E}[Y|T = 0, X = x] \quad (2)$$

CATE can be similarly defined for continuous treatments, and by averaging over  $X$  it reduces to the standard ATE  $\mathbb{E}[Y(1) - Y(0)]$ . Variables included in the feature vector  $X$  are known as controls. Depending on the causal structure of the phenomenon under study, controlling for a variable might reduce or increase bias of effect estimates [13].

**Double Machine Learning.** Assuming unconfoundedness, Double Machine Learning [11] formulates the data generating process in terms of the Partially Linear Model [50]:

$$Y = \theta(X) \cdot T + g(X) + \varepsilon \quad (3)$$

$$T = f(X) + \eta \quad (4)$$

where  $\theta(X)$  is the CATE, and  $g, f$  are arbitrary functions (nuisance parameters). Notably, (4) keeps track of confounding as features  $X$  drive both the treatment  $T$  and outcome  $Y$ . CATE  $\theta(X)$  is then estimated using a two-stage estimation procedure. During the first stage, both the outcome  $Y$  and treatment  $T$  are separately predicted from features  $X$  using arbitrary ML models. Then, in the final stage, the CATE  $\theta(X)$  is estimated by solving (for the linear case)

$$\hat{\theta} = \arg \min_{\theta \in \Theta} \mathbb{E}[(\tilde{Y} - \theta(X) \cdot \tilde{T})^2] \quad (5)$$

over a model class  $\Theta$ , where  $\tilde{Y}$  are the residuals of the  $Y \sim X$  regression, and  $\tilde{T}$  are the residuals of the  $T \sim X$  regression. In (3) and (5) the linearity assumption can be dropped to allow for fully non-parametric CATE estimation.

### 3.2. Data derived from crop type maps

Co-designing practical and effective solutions to major challenges, such as climate change, requires understanding the pattern of interventions carried out by agents of change in agricultural ecosystems, i.e farmers [8]. LCD and CR are considered to be important management practices since they support synergistic improvements in crop yield, environmental health, and ecological sustainability [9, 21, 61]. Hence, such practices should be considered in strategies that promote sustainable agricultural production.

We exploited a series of yearly Land Parcel Identification System (LPIS) data to produce variables representing agricultural management practices [44]. LPIS is a geo-spatial database that contains the geometries of the parcels and the declared crop type by the farmers, as part of their application for CAP subsidies [51]. For each year, we extracted the proportion of the total area occupied by each crop type (crop abundance) per grid cell of the 500m MODIS square grid. We then calculated the Shannon diversity index ( $H'$ ) [45, 58] as a metric of LCD per grid cell. We finally took the mean of all years to end up with an average LCD value per cell for 2010-2020. For CR, we summed the absolute difference per crop type abundance per grid cell for two consecutive years; this procedure was repeated for all pairs of adjacent years with which we calculated the total (sum) rotations for the studied period. Both LCD and CR were used as treatments to test their effect on NPP, while crop type abundances for the major crop types in Flanders (grassland, maize, potato, wheat) were used as controls.

### 3.3. Net Primary Productivity

Net primary productivity (NPP) is the uptake flux, in which carbon from the atmosphere is sequestered by plants through the balance between photosynthesis and plant respiration. It is a fundamental ecological variable in biosphere functioning, the quantification of which is needed for assessing the carbon balance at regional and global scales [68]. As NPP has been widely used for measuring vegetation dynamics, this index is highly suitable for capturing environmental changes due to natural and anthropogenic factors [26, 73]. The MODIS NPP is produced by the US National Aeronautics and Space Administration (NASA) Earth Observing System. It is based on an energy budget approach, utilizing earth observations on the fraction of photosynthetically active solar radiation absorbed by the vege-

tation surface [48, 53]. The results have been validated as being able to capture spatio-temporal patterns across various biomes and climate regimes [70]. We used annual time series of MODIS NPP (MOD17A3HGF v006) gridded at 500m for 2010-2020 as the target metric (outcome) representing the ecosystem function of climate regulation over Flanders (North Belgium). The MOD17A3 annual product is derived by summing all 8-day net photosynthesis products (MOD17A2H), partly capturing seasonal variability.

### 3.4. Environmental factors

Variability in agricultural NPP arises due to environmental effects, with specific conditions favoring specific crop types. To uncover the diverse environmental conditions that support or inhibit climate regulation we used a series of parameters provided by Terraclimate [1]. The gridded meteorological data were produced through a climatically aided spatiotemporal interpolation of the WorldClim datasets to estimate monthly time series. Environmental factors used in this study included maximum and minimum temperature, actual evapotranspiration, climate water deficit, precipitation, soil moisture, downward surface shortwave radiation and vapor pressure. The idea behind the selection of the aforementioned variables relies on the wide recognition that temperature and precipitation directly affect NPP [12, 27, 68], the influence of water availability on soil productivity and vegetation growth [28], and the processes themselves (e.g. evapotranspiration) that relate to the ecosystem function of climate regulation [60, 67].

### 3.5. Methodological setup

We have thus created a population of  $N$  units indexed by  $i$ , where  $Y_i(T_i)$  is the observed outcome of unit  $i$  treated with  $T_i$  and  $X_i$  are the controls and features that generate systematic variation on the treatment effect. Specifically, the population consists of  $N$  grid cells exhaustively covering agricultural land within Flanders,  $Y_i$  is the observed NPP value of each cell,  $T_i$  is (the value of) the agricultural practice that was applied to the cell, and  $X_i$  refers to important characteristics of the specific grid cell, comprising environmental data and agricultural use (i.e., crop types cultivated within the grid cell). The CATE (1) is then the average impact that practice  $T_i$  had on NPP  $Y_i$  for the  $i$ -th grid cell, conditional to its characteristics  $X_i$ . Obtaining such local insights on where and why the impact magnitude of agricultural practices on important agro-ecosystem metrics differed is of primary interest to policy makers as it allows for targeted agricultural policy making.

All data mentioned in the previous subsections are retrieved as time series, with different spatio-temporal resolutions. Data were temporally aggregated over the period of study (2010-2020), since treating every cell-year combination as different units would potentially introduce interfer-

ence effects [62], where treatments (agricultural practices) during any year might also influence the outcome (climate regulation) of the next years. For the case of CR, temporal aggregation over the 2010-2020 period happens by summing all the rotations that happened in grid cell  $i$  from 2010 to 2020, while for LCD we consider the average Shannon diversity index over all years. This is the treatment  $T$  we are using for the analyses.

By adding the crop abundances to the feature vector  $X$ , we control for the four most dominant crop types mentioned in Sec. 3.2; the median abundance value of all other crop types is less than 2%. We also control for important environmental variables, listed in Sec. 3.4. Crop types are confounders for the causal relation under study and thus good controls as they drive the magnitude of both the agricultural practices and the NPP. As such, they should be included in the model. On the other hand, environmental data are not driving any practice, but they are driving the NPP values, thus making them strong candidates for good controls [13]. We finally binarize the treatment by letting the median CR value and median LCD value to be thresholds, over which we designate cells as treated units and the rest as control.

The assumption of unconfoundedness can't be tested, and within any observational study it is likely to be invalid. Even if the variables we controlled for will help with bias reduction, some bias from unobserved confounders might still be present. Nevertheless, we note that by being cautious with the selection of controls we try to avoid selection bias, and reported results provide a large-scale complementary approach to localized field experiments.

## 4. Experimental Results & Discussion

**Filtering.** For the experiments we used the EconML Python package implementing the Double Machine Learning method [6]. We derived CATE estimates for both LCD and CR at the 500m native NPP resolution. We ensured that fitting takes place over actual croplands by restricting the dataset to grid cells where the sum of crop abundances, given by the crop type maps, exceeded a threshold of 80%. To aid the overlap assumption, propensity scores were first estimated using a Gradient Boosting Propensity Model from the CausalML Python package [10], and units with extreme propensity scores ( $\leq 0.2$  or  $\geq 0.8$ ) were filtered out. Within DML, it is crucial to avoid overfitting first stage models; otherwise a portion of the outcome and treatment variability explained will be due to factors other than the controls. During the first stage of DML only, we therefore split the dataset to train and test (80-20) to evaluate the predictive performance and assess overfit.

**Double Machine Learning.** In first stage, for both the CR and LCD analyses, random forest regressors were selected to predict NPP from controls, outperforming Lasso and gradient boosting regression. To predict the binary

treatments themselves, logistic regression was used, outperforming random forest and gradient boosting classifiers.

All model selection procedures happened by performing 3-fold cross validation and a grid search for hyperparameter optimization. During the first stage only, the maximum absolute scaler was applied, as it retains the sparsity structure that is prevalent in crop abundance data. For the final stage regression, for both analyses, a Causal Forest [65] with 1000 trees was used with heterogeneity score as the splitting criterion. The final stage causal forest was fine-tuned based on the out of sample performance. Minimization of (5) happened with unscaled features to maintain interpretation in the original units of measurement. Fitting results for CR and LCD can be found in Table 1. First stage models captured a significant part of the variability of both the outcome and treatment variables in both cases. The difference between the training and test performance was marginal, indicating that models avoided overfitting.

Crop Rotation	Train	Test
$Y \sim X$ (Outcome Modeling)	0.75	0.74
$T \sim X$ (Treatment Modeling)	0.63	0.62
Landscape crop diversity	Train	Test
$Y \sim X$ (Outcome Modeling)	0.73	0.71
$T \sim X$ (Treatment Modeling)	0.59	0.62

Table 1. Performance ( $R^2$  for outcome modeling, F-1 score for treatment modeling) of ML models internally used by DML. Outcome  $Y$  is Net Primary Productivity, treatment  $T$  is crop rotation / landscape crop diversity,  $X$  is a vector of features.

### 4.1. CATE Results

Figure 2 illustrates the DML CATE estimates over the region of Flanders for LCD and CR. Intervening by applying a practice in a cell with a CATE estimate of e.g. 20, we would expect an increase of 20 kilograms of carbon for each squared meter of land within a year. From such results, a spatially explicit policy for sustainable agriculture can be extracted, by prioritizing for each cell the application of practices whose effect estimate (land suitability score) is high. Such analysis can be repeated for any agricultural practice and outcome metric of interest, while appropriate controls can be added using expert knowledge.

Our results indicated that in north-west areas that are characterised by high abundance of grasslands, CR increased NPP by approximately 100 kgC/m<sup>2</sup>/year (Figure 2a; square 1). This suggests that the contribution of grasslands in crop rotations is important in positively driving climate regulation [69]. Positive effects were also found on south-east areas where the dominating crop type is winter wheat followed by maize and grasslands (Figure 2a; square 2).

For LCD, negatively affected areas west (Figure 2b; square 3) and south-east (Figure 2b; square 5) that are dominated by grasslands and winter wheat revealed a decrease in NPP of over 100 kgC/m<sup>2</sup>/year. Due to the importance of the aforementioned crop types to the carbon cycle, one could expect positive LCD impacts on local climate regulation [7, 55]. However, as these crops appear clustered in space, diversifying such cropland regions led to a negatively affected NPP. By contrast, where multiple crop types exist, maize showed a significant role in carbon uptake as its existence in a diversified agricultural landscape indicated an NPP increase of approximately 50 kgC/m<sup>2</sup>/year (Figure 2b; square 4 and 6). In fact, maize has a high capacity in capturing large amounts of carbon from the atmosphere [40].

The ATE estimates for CR were 1.08 (95% CI [-20.35, 22.51]) and for LCD were -35.73 (95% CI [-58.73, -12.73]). The distribution of CATE estimates for both practices over the Flanders croplands can be seen in Figure 3. Besides the more granular CATE insights we are reporting, we also note that a LCD practice implemented over the entire Flanders cropland would have a weak yet significant negative impact on NPP. On the other hand, a CR practice implemented across the entire terrain would be able to bring multiple agro-environmental benefits without impacting the average NPP [30, 35].

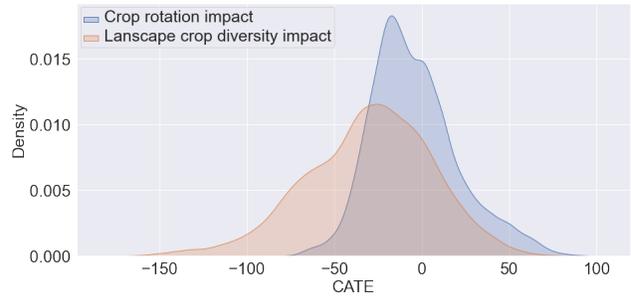
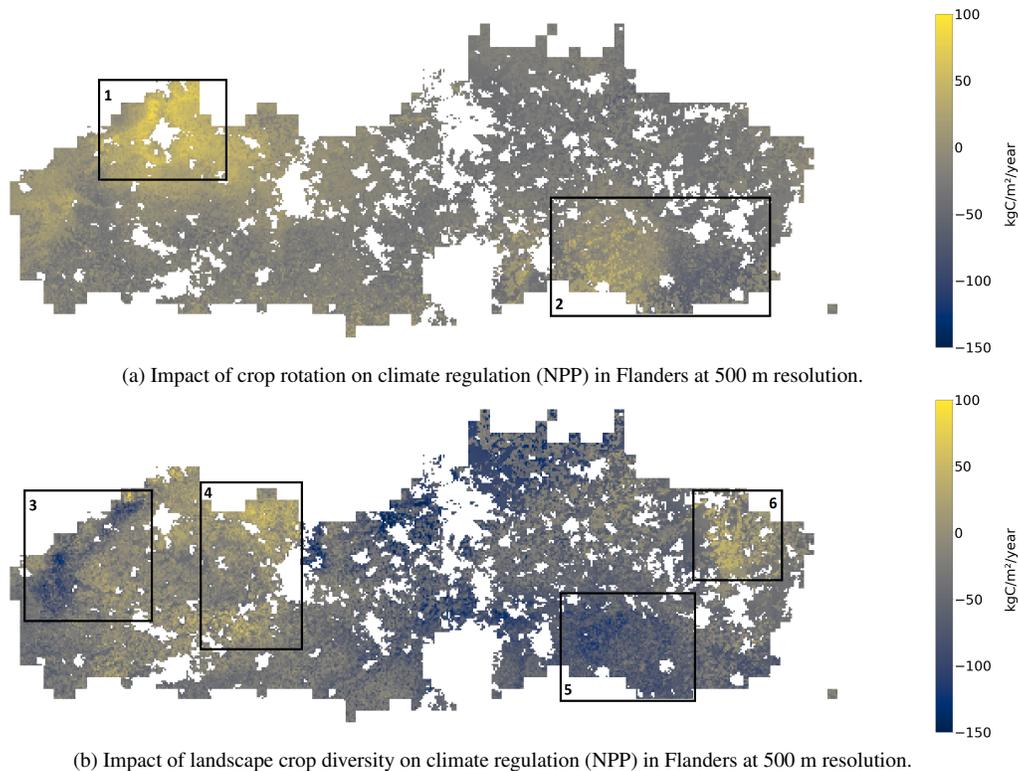


Figure 3. Distribution of CATE estimates for crop rotation and landscape crop diversity.

## 4.2. Heterogeneity Analysis

We used causal trees to analyze the heterogeneity of the model, accounting for all features in  $X$  [4, 6]. The tree successively splits on feature values that maximize the treatment effect difference across leaves. We constrained the tree depth to 2 to retain explainability.

In Figure 4a we see that minimum temperature was found to be the most important driver of effect heterogeneity for CR. Particularly, high minimum temperatures combined with low climate water deficit seem to benefit CR perfor-



(a) Impact of crop rotation on climate regulation (NPP) in Flanders at 500 m resolution.

(b) Impact of landscape crop diversity on climate regulation (NPP) in Flanders at 500 m resolution.

Figure 2. Estimated impact of crop rotation (a) and landscape crop diversity (b) on climate regulation (NPP) based on the heterogeneous treatment effect analysis for the Flanders region from 2010 to 2020. The annual NPP outcome variable is measured in kgC/m<sup>2</sup>/year.

mance in Flanders. Wheat abundance was also detected as a major contributor to CR heterogeneity, spatially coinciding with high CATE estimates as reported in Sec. 4.1.

LCD impact heterogeneity (Figure 4b) was found to be affected by maximum temperature, solar radiation and grassland abundance. LCD negatively affected NPP the least when maximum temperature was lower than 15.1°C, and radiation was lower than 119.9 W/m<sup>2</sup>. Additionally, in high maximum temperatures, diversifying crops in grassland-abundant cells appears more preferable than diversifying cells lacking grasslands.

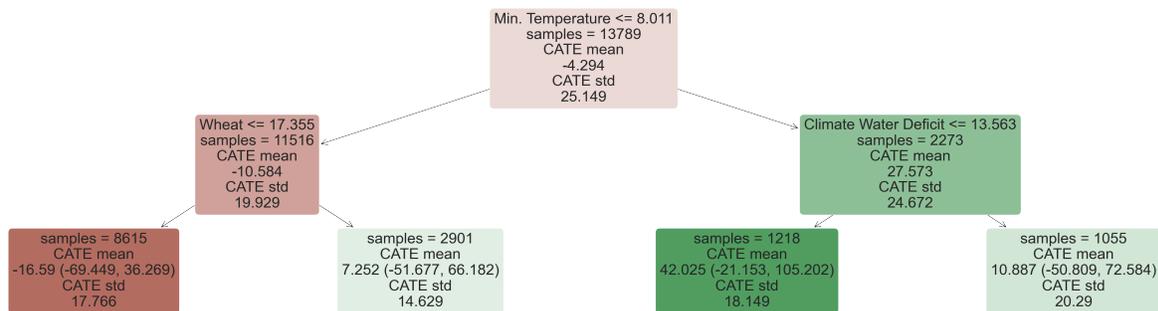
While trees are useful to detect the most important features in the sense described above, a complementary understanding of effect heterogeneity can be obtained by plotting CATE estimates as a function of selected features. In Figures 5 and 6, we see clear patterns that show maximum temperature driving effect heterogeneity for both practices. For the rest of the features, we report Spearman correlation coefficients between them and the estimated CATE in Figure 7. There are moderate correlations for almost all factors, indicating that the effect under study is indeed heterogeneous. The correlation values for environmental features are larger but comparable to the ones of crop features, highlighting that both control categories contribute to heterogeneity.

### 4.3. Agricultural land suitability & climate change

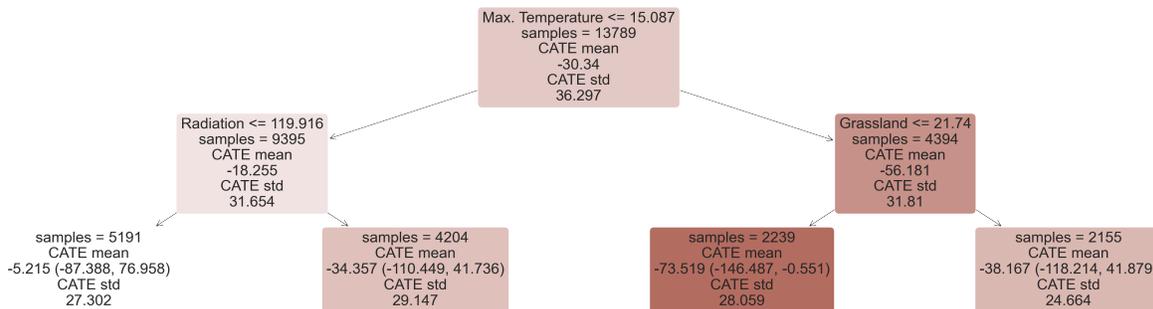
So far, we estimated the impact of agricultural practices on climate regulation from historical data (2010-2020). Nevertheless, important climatic variables are expected to change over the next decades as a result of climate change. The estimation of their trajectory has been the subject of numerous scientific studies [59]. In order to devise impactful agricultural policies, we need to evaluate the performance of practices in future climatic conditions.

In the context of causal machine learning, we learned the function  $\theta(x) = \mathbb{E}[Y(1) - Y(0)|X = x]$  where feature vector  $X$  contains multiple climate variables. In theory, we are able to derive the impact of a practice on an outcome metric under future climatic conditions by changing the relevant variables of feature vector  $X$  and re-calculating. However,  $\theta(x)$  was learned from historical data, and by definition future climatic conditions were not observed. By changing climate variables to match future projections, we are essentially extrapolating the function to points outside the observed feature space.

In Figure 5 for example, a quadratic trend between the CR CATE and maximum temperature is seen, allowing us to hypothesize that in slightly warmer conditions the benefit of CR would increase. Thus, our results provide preliminary



(a) Heterogeneity tree for crop rotation.



(b) Heterogeneity tree for landscape crop diversity.

Figure 4. Trees interpreting effect heterogeneity for both practices studied. To be read from top to bottom, going left if the Boolean condition at the top of each box is true, and right if it is false. Sample size for each leaf is reported, alongside the CATE mean, standard deviation and confidence intervals. Wheat and grassland abundance are in %, min. and max. temperature are in °C, climate water deficit is in mm and radiation is in W/m<sup>2</sup>.

insights on an ex-ante impact assessment for the agricultural practices studied in the context of climate change, enabling the formulation of scientific hypotheses for further study.

Future work includes incorporating climatic projections in the pipeline and comparing the impact learned from them to the impact learned from historical data. Further work also comprises the study of more agricultural practices, e.g., intercropping and minimum soil cover, but also the study of more target outcomes, such as erosion prevention and pollination potential. This way, we can move towards a synthesis of ecosystem service trade-offs and the formulation of a comprehensive suite of policies in response to climate change. Finally, the correlations between agricultural practices and NPP, and consequently the treatment impact, depend on the spatial scale of analysis [72]. Therefore, agricultural practices should be evaluated at different scales to gain insights towards optimizing agriculture land suitability from farm to landscape level and achieve optimum value.

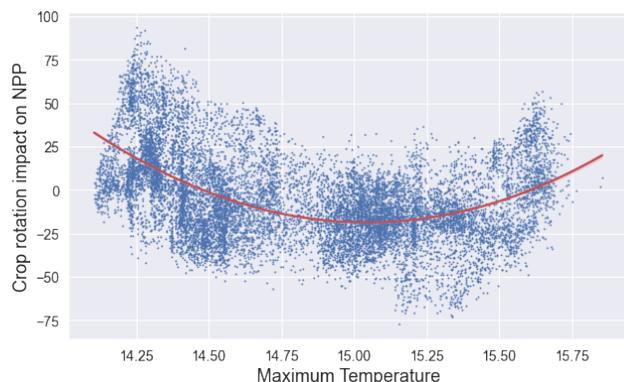


Figure 5. Crop rotation impact against max. temperature (°C).

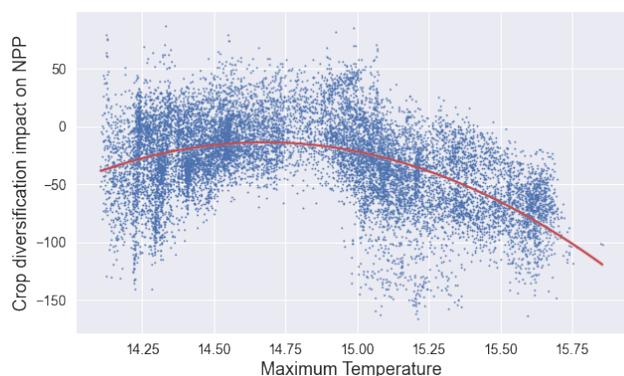


Figure 6. Landscape crop diversity impact against max. temperature (°C).

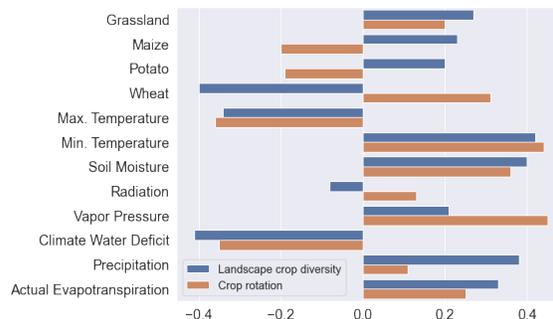


Figure 7. Barplot showing Spearman correlation coefficients of estimated CATEs with all features for both practices.

## 5. Conclusions

We presented an approach for assessing agricultural land suitability using causal machine learning. We estimated the heterogeneous treatment effects of CR and LCD on climate regulation (NPP), accounting for historical crop and environmental data. We used earth observations (MODIS), climate data (TerraClimate) and openly available yearly crop type maps covering North Belgium (Flanders) from 2010 to 2020. CATE results reveal significant impact heterogeneity in space, highlighting the usefulness of extended spatio-temporal data coverage, the importance of spatially targeted measures and the relevance of CATE estimates as land suitability scores. The ATE for CR is insignificant, while for LCD it is significant but relatively small compared to the influence environmental drivers have on NPP [73]. It would be difficult to extract such subtle effects from localized experiments of confined samples.

Significant challenges remain towards the application of the proposed approach in real-life planning. Observational studies enable the manifestation of heterogeneity, which is vital for spatially explicit recommendations, but suffer from unobserved confounding. For this reason, our results should be used along with insights from local field experiments. Furthermore, we offered insights on how our two treatments affect climate regulation but we did not provide detailed information on the contribution of specific crop types. For the CR treatment, we do not consider the crop types involved in the transition but simply the transition itself. Similarly, for the LCD all crop type mixtures are treated equally. Additional analyses with crop-aware treatments can offer enhanced understanding for evidence-based decision making.

## Acknowledgements

This work has been supported by the EIFFEL project (EU Horizon 2020 - GA No. 101003518) and the MICROSERVICES project (2019-2020 BiodivERSA joint call, under the BiodivClim ERA-Net COFUND programme, and with the GSRI, Greece - No. T12ERA5-00075).

## References

- [1] John T. Abatzoglou, Solomon Z. Dobrowski, Sean A. Parks, and Katherine C. Hegewisch. TerraClimate, a high-resolution global dataset of monthly climate and climatic water balance from 1958–2015. *Scientific Data*, 5(1):170191, Dec. 2018. **4**
- [2] Jason Abrevaya, Yu-Chin Hsu, and Robert P Lieli. Estimating conditional average treatment effects. *Journal of Business & Economic Statistics*, 33(4):485–505, 2015. **3**
- [3] Komlavi Akpoti, Amos T Kabo-bah, and Sander J Zwart. Agricultural land suitability analysis: State-of-the-art and outlooks for integration of climate change analysis. *Agricultural systems*, 173:172–208, 2019. **2**
- [4] Susan Athey and Guido Imbens. Recursive partitioning for heterogeneous causal effects. *Proceedings of the National Academy of Sciences*, 113(27):7353–7360, July 2016. **6**
- [5] Susan Athey, Julie Tibshirani, and Stefan Wager. Generalized random forests. *The Annals of Statistics*, 47(2), Apr. 2019. **2**
- [6] Keith Battocchi, Eleanor Dillon, Maggie Hei, Greg Lewis, Paul Oka, Miruna Oprescu, and Vasilis Syrgkanis. EconML: A Python Package for ML-Based Heterogeneous Treatment Effects Estimation. <https://github.com/microsoft/EconML>, 2019. Version v0.12.0. **5, 6**
- [7] J. Bengtsson, J. M. Bullock, B. Egoth, C. Everson, T. Everson, T. O’Connor, P. J. O’Farrell, H. G. Smith, and R. Lindborg. Grasslands—more important for ecosystem services than you might think. *Ecosphere*, 10(2):e02582, 2019. **6**
- [8] David A Bohan, Reto Schmucki, Abrha T Abay, Mette Termansen, Miranda Bane, Alice Charalabidis, Rong-Gang Cong, Stephane AP Derocles, Zita Dorner, Matthieu Forster, et al. Designing farmer-acceptable rotations that assure ecosystem service provision in the face of climate change. In *Advances in ecological research*, volume 65, pages 169–244. Elsevier, 2022. **4**
- [9] Timothy M Bowles, Maria Mooshammer, Yvonne Socolar, Francisco Calderón, Michel A Cavigelli, Steve W Culman, William Deen, Craig F Drury, Axel Garcia y Garcia, Amélie CM Gaudin, et al. Long-term evidence shows that crop-rotation diversification increases agricultural resilience to adverse growing conditions in north america. *One Earth*, 2(3):284–293, 2020. **4**
- [10] Huigang Chen, Totte Harinen, Jeong-Yoon Lee, Mike Yung, and Zhenyu Zhao. Causalml: Python package for causal machine learning, Mar. 2020. **5**
- [11] Victor Chernozhukov, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins. Double/debiased machine learning for treatment and structural parameters. 21(1):C1–C68, 2018. **2, 3**
- [12] Chengjin Chu, Megan Bartlett, Youshi Wang, Fangliang He, Jacob Weiner, Jérôme Chave, and Lawren Sack. Does climate directly influence npp globally? *Global Change Biology*, 22(1):12–24, 2016. **4**
- [13] Carlos Cinelli, Andrew Forney, and Judea Pearl. A Crash Course in Good and Bad Controls. *SSRN Electronic Journal*, 2020. **3, 5**
- [14] Liesbeth Colen, Sergio Gomez y Paloma, Uwe Latacz-Lohmann, Marianne Lefebvre, Raphaële Préget, and Sophie Thoyer. Economic experiments as a tool for agricultural policy evaluation: insights from the european cap. *Canadian Journal of Agricultural Economics/Revue canadienne d’agroeconomie*, 64(4):667–694, 2016. **1**
- [15] Richard K Crump, V Joseph Hotz, Guido W Imbens, and Oscar A Mitnik. Nonparametric tests for treatment effect heterogeneity. *The Review of Economics and Statistics*, 90(3):389–405, 2008. **2**
- [16] Paula Cullen, Pierre Dupraz, James Moran, Pat Murphy, Roman O’Flaherty, Cathal O’Donoghue, Robert O’Shea, and Mary Ryan. Agri-environment scheme design: Past lessons and future suggestions. *EuroChoices*, 17(3):26–30, Dec. 2018. **1**
- [17] Alexandre Danvi, Thomas Jütten, Simone Giertz, Sander J Zwart, and Bernd Diekkrüger. A spatially explicit approach to assess the suitability for rice cultivation in an inland valley in central benin. *Agricultural water management*, 177:95–106, 2016. **2**
- [18] Jillian M Deines, Sherrie Wang, and David B Lobell. Satellites reveal a small positive yield effect from conservation tillage across the US Corn Belt. *Environmental Research Letters*, 14(12):124038, Dec. 2019. **1, 2**
- [19] Davide Del Prete, Laura Forastiere, and Valerio Leone Sciabolazza. Causal inference on networks under continuous treatment interference: an application to trade distortions in agricultural markets. *Available at SSRN: <https://ssrn.com/abstract=3363173> or <http://dx.doi.org/10.2139/ssrn.3363173>*, Mar. 2019. **1**
- [20] Miroslav Dudík, Dumitru Erhan, John Langford, and Lihong Li. Doubly robust policy evaluation and optimization. *Statistical Science*, 29(4):485–511, 2014. **2**
- [21] Lukas Egli, Zia Mehrabi, and Ralf Seppelt. More farms, less specialized landscapes, and higher crop diversity stabilize food supplies. *Environmental Research Letters*, 16(5):055015, 2021. **2, 4**
- [22] AA El Baroudy. Mapping and evaluating land suitability using a gis-based model. *Catena*, 140:96–104, 2016. **2**
- [23] Christian Folberth, Nikolay Khabarov, Juraj Balkovič, Rastislav Skalský, Piero Visconti, Philippe Ciais, Ivan A Janssens, Josep Peñuelas, and Michael Obersteiner. The global cropland-sparing potential of high-yield farming. *Nature Sustainability*, 3(4):281–289, 2020. **2**
- [24] Bojie Fu, Shuai Wang, Junze Zhang, Zengqian Hou, and Jinghai Li. Unravelling the complexity in achieving the 17 sustainable-development goals. *National Science Review*, 6(3):386–388, 2019. **1**
- [25] Nicole B Gabler, Naihua Duan, Diana Liao, Joann G Elmore, Theodore G Ganiats, and Richard L Kravitz. Dealing with heterogeneity of treatment effects: is the literature up to the challenge? *Trials*, 10(1):1–12, 2009. **2**
- [26] Chengcheng Gang, Yanzhen Zhang, Zhaoqi Wang, Yizhao Chen, Yue Yang, Jianlong Li, Jimin Cheng, Jianguo Qi, and Inakwu Odeh. Modeling the dynamics of distribution, extent, and npp of global terrestrial ecosystems in response to future climate change. *Global and Planetary Change*, 148:153–165, 2017. **4**

- [27] Madhav D Gholkar, Sheshakumar Goroshi, Raghavendra P Singh, and Jai Singh Parihar. Influence of agricultural developments on net primary productivity (npp) in the semi-arid region of india: a study using glopem model. *The International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences*, 40(8):725, 2014. 4
- [28] Binyam T Hailu, Eduardo Eiji Maeda, Janne Heiskanen, and Petri Pellikka. Reconstructing pre-agricultural expansion vegetation cover of ethiopia. *Applied Geography*, 62:357–365, 2015. 4
- [29] Justine S. Hastings, Thomas J. Kane, and Douglas O. Staiger. Preferences and Heterogeneous Treatment Effects in a Public School Choice Lottery. NBER Working Papers 12145, National Bureau of Economic Research, Inc, Apr. 2006. 2
- [30] JL Havlin, DE Kissel, LD Maddux, MM Claassen, and JH Long. Crop rotation and tillage effects on soil organic carbon and nitrogen. *Soil Science Society of America Journal*, 54(2):448–452, 1990. 6
- [31] Paul Hünermund and Elias Bareinboim. Causal inference and data fusion in econometrics. *arXiv preprint arXiv:1912.09104*, 2019. 1
- [32] Paul Hünermund and Dirk Czarnitzki. Innovation policy and causality. *ifo DICE Report*, 17(04):3–6, 2019. 1
- [33] Andrew Jesson, Peter Manshausen, Alyson Douglas, Duncan Watson-Parris, Yarin Gal, and Philip Stier. Using Non-Linear Causal Models to Study Aerosol-Cloud Interactions in the Southeast Pacific. *arXiv:2110.15084 [physics]*, Nov. 2021. 2, 3
- [34] Andrew Jesson, Sören Minderhmann, Yarin Gal, and Uri Shalit. Quantifying Ignorance in Individual-Level Causal-Effect Estimates under Hidden Confounding. *arXiv:2103.04850 [cs, stat]*, Aug. 2021. 2
- [35] Douglas L Karlen, Eric G Hurley, Susan S Andrews, Cynthia A Cambardella, David W Meek, Michael D Duffy, and Antonio P Mallarino. Crop rotation effects on soil quality at three northern corn/soybean belt locations. *Agronomy journal*, 98(3):484–495, 2006. 6
- [36] David M Kent, Peter M Rothwell, John PA Ioannidis, Doug G Altman, and Rodney A Hayward. Assessing and reporting heterogeneity in treatment effects in clinical trials: a proposal. *Trials*, 11(1):1–11, 2010. 2
- [37] Dan M Kluger, Art B Owen, and David B Lobell. Combining randomized field experiments with observational satellite data to assess the benefits of crop rotations on yields. *arXiv preprint arXiv:2112.13700*, 2021. 1
- [38] Richard L Kravitz, Naihua Duan, and Joel Braslow. Evidence-based medicine, heterogeneity of treatment effects, and the trouble with averages. *The Milbank Quarterly*, 82(4):661–687, 2004. 2
- [39] Sören R. Künzel, Jasjeet S. Sekhon, Peter J. Bickel, and Bin Yu. Metalearners for estimating heterogeneous treatment effects using machine learning. *Proceedings of the National Academy of Sciences*, 116(10):4156–4165, Mar. 2019. 2
- [40] Rattan Lal. Managing soils and ecosystems for mitigating anthropogenic carbon emissions and advancing global food security. *BioScience*, 60(9):708–721, 2010. 6
- [41] Nicolas Lampkin, Gerald Schwarz, and Stéphane Bellon. Policies for agroecology in europe, building on experiences in france, germany and the united kingdom. *Journal of sustainable and organic agricultural systems*, 70(2):103–112, 2021. 1
- [42] Sadeeka Layomi Jayasinghe, Lalit Kumar, and Janaki Sandamali. Assessment of potential land suitability for tea (*Camellia sinensis* (L.) O. Kuntze) in Sri Lanka using a GIS-based multi-criteria approach. *Agriculture*, 9(7):148, 2019. 2
- [43] Myoung-jae Lee. Non-parametric tests for distributional treatment effect for randomly censored responses. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 71(1):243–264, 2009. 2
- [44] Łukasz Mikołajczyk, Ryszard Laskowski, Elżbieta Ziólkowska, and Agnieszka J Bednarska. Species-specific landscape characterisation method in agro-ecosystems. *Ecological Indicators*, 129:107894, 2021. 4
- [45] E Kathryn Morris, Tancredi Caruso, François Buscot, Markus Fischer, Christine Hancock, Tanja S Maier, Torsten Meiners, Caroline Müller, Elisabeth Obermaier, Daniel Prati, et al. Choosing and using diversity indices: insights for ecological applications from the German biodiversity exploratories. *Ecology and Evolution*, 4(18):3514–3524, 2014. 4
- [46] Hillary Mugiyu, Vimbayi GP Chimonyo, Mbulisi Sibanda, Richard Kunz, Cecilia R Masemola, Albert T Modi, and Tafadzwanashe Mabhaudhi. Evaluation of land suitability methods with reference to neglected and underutilised crop species: A scoping review. *Land*, 10(2):125, 2021. 2
- [47] Maria Niedertscheider, Thomas Kastner, Tamara Fetzl, Helmut Haberl, Christine Kroisleitner, Christoph Plutzer, and Karl-Heinz Erb. Mapping and analysing cropland use intensity from a npp perspective. *Environmental Research Letters*, 11(1):014008, 2016. 1
- [48] Yude Pan, Richard Birdsey, John Hom, Kevin McCullough, and Kenneth Clark. Improved estimates of net primary productivity from MODIS satellite data at regional and local scales. *Ecological Applications*, 16(1):125–132, 2006. 4
- [49] Judea Pearl. *Causality*. Cambridge University Press, 2009. 2
- [50] Peter M Robinson. Root-n-consistent semiparametric regression. *Econometrica: Journal of the Econometric Society*, pages 931–954, 1988. 3
- [51] Maria Rousi, Vasileios Sitokostantinou, Georgios Meditskos, Ioannis Papoutsis, Ilias Gialampoukidis, Alkiviadis Koukos, Vassilia Karathanassi, Thanassis Drivas, Stefanos Vrochidis, Charalampos Kontoes, et al. Semantically enriched crop type classification and linked earth observation data to support the common agricultural policy monitoring. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 14:529–552, 2020. 4
- [52] Donald B Rubin. Causal inference using potential outcomes: Design, modeling, decisions. *Journal of the American Statistical Association*, 100(469):322–331, 2005. 2, 3
- [53] Steven W Running, Ramakrishna R Nemani, Faith Ann Heinsch, Maosheng Zhao, Matt Reeves, and Hirofumi Hashimoto. A continuous satellite-derived measure of global terrestrial primary production. *Bioscience*, 54(6):547–560, 2004. 4
- [54] Fabio Rubio Scarano. Ecosystem-based adaptation to climate change: concept, scalability and a role for conserva-

- tion science. *Perspectives in Ecology and Conservation*, 15(2):65–73, 2017. 1
- [55] Marius Schmidt, Tim G Reichenau, Peter Fiener, and Karl Schneider. The carbon budget of a winter wheat field: An eddy covariance analysis of seasonal and inter-annual variability. *Agricultural and Forest Meteorology*, 165:114–126, 2012. 6
- [56] Feliu Serra-Burriel, Pedro Delicado, Andrew T. Prata, and Fernando M. Cucchietti. Estimating heterogeneous wildfire effects using synthetic controls and satellite remote sensing. *Remote Sensing of Environment*, 265:112649, Nov. 2021. 2
- [57] Zongyao Sha, Ruren Li, Jonathan Li, and Yichun Xie. Estimating carbon sequestration potential in vegetation by distance-constrained zonal analysis. *IEEE Geoscience and Remote Sensing Letters*, 18(8):1352–1356, 2020. 2
- [58] Claude Elwood Shannon. A mathematical theory of communication. *The Bell system technical journal*, 27(3):379–423, 1948. 4
- [59] Thomas Stocker. *Climate change 2013: the physical science basis: Working Group I contribution to the Fifth assessment report of the Intergovernmental Panel on Climate Change*. Cambridge university press, 2014. 7
- [60] Ge Sun, Dennis Hallema, and Heidi Asbjornsen. Ecohydrological processes and ecosystem services in the anthropocene: a review. *Ecological Processes*, 6(1):1–9, 2017. 4
- [61] Giovanni Tamburini, Riccardo Bommarco, Thomas Cherico Wanger, Claire Kremen, Marcel GA van der Heijden, Matt Liebman, and Sara Hallin. Agricultural diversification promotes multiple ecosystem services without compromising yield. *Science advances*, 6(45):eaba1715, 2020. 4
- [62] Eric J Tchetgen Tchetgen and Tyler J VanderWeele. On causal inference in the presence of interference. *Statistical methods in medical research*, 21(1):55–75, 2012. 5
- [63] European Union. The post-2020 common agricultural policy: environmental benefits and simplification. 2018. Available at: <https://ec.europa.eu/info/food-farming-fisheries/key-policies/common-agricultural-policy/future-cap>. 1
- [64] Stefan Wager and Susan Athey. Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, 113(523):1228–1242, 2018. 2
- [65] Stefan Wager and Susan Athey. Estimation and Inference of Heterogeneous Treatment Effects using Random Forests. *Journal of the American Statistical Association*, 113(523):1228–1242, July 2018. 2, 5
- [66] Richard J Willke, Zhiyuan Zheng, Prasun Subedi, Rikard Althin, and C Daniel Mullins. From concepts, theory, and evidence of heterogeneity of treatment effects to methodological approaches: a primer. *BMC medical research methodology*, 12(1):1–12, 2012. 2
- [67] Qing Yang, Gengyuan Liu, Biagio F Giannetti, Feni Agostinho, Cecília MVB Almeida, and Marco Casazza. Energy-based ecosystem services valuation and classification management applied to china’s grasslands. *Ecosystem Services*, 42:101073, 2020. 4
- [68] Zhe Yuan, Yongqiang Wang, Jijun Xu, and Zhiguang Wu. Effects of climatic factors on the net primary productivity in the source region of yangtze river, china. *Scientific Reports*, 11(1):1–11, 2021. 4
- [69] Azin Zarei, Abel Chemura, Stephanie Gleixner, and Holger Hoff. Evaluating the grassland npp dynamics in response to climate change in tanzania. *Ecological Indicators*, 125:107600, 2021. 5
- [70] Maosheng Zhao, Faith Ann Heinsch, Ramakrishna R Nemani, and Steven W Running. Improvements of the modis terrestrial gross and net primary production global data set. *Remote sensing of Environment*, 95(2):164–176, 2005. 4
- [71] Min Zheng, Jesseca K Marsh, Jeffrey V Nickerson, and Samantha Kleinberg. How causal information affects decisions. *Cognitive Research: Principles and Implications*, 5(1):1–24, 2020. 1
- [72] Yanyan Zhou, Dongxia Yue, Jianjun Guo, Guanguang Chen, and Dong Wang. Spatial correlations between landscape patterns and net primary productivity: A case study of the shule river basin, china. *Ecological Indicators*, 130:108067, 2021. 8
- [73] Yanyan Zhou, Dongxia Yue, Chen Li, Xinliang Mu, and Jianjun Guo. Identifying the spatial drivers of net primary productivity: A case study in the bailong river basin, china. *Global Ecology and Conservation*, 28:e01685, 2021. 4, 8