

The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search http://ageconsearch.umn.edu aesearch@umn.edu

Papers downloaded from **AgEcon Search** may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.

No endorsement of AgEcon Search or its fundraising activities by the author(s) of the following work or their employer(s) is intended or implied.



Satellite Data in Agricultural and Environmental Economics: Theory and Practice

David Wüpper¹, Wyclife Agumba Oluoch¹, Hadi¹

1: Land Economics Group, University of Bonn Corresponding author email: wuepper@uni-bonn.de

Abstract

Agricultural and environmental economists are in the fortunate position that a lot of what is happening on the ground is observable from space. Most agricultural production happens in the open and one can see from space when and where innovations are adopted, crop yields change, or forests are converted to pastures, to name just a few examples. However, converting images into measurements of a particular variable is not trivial, as there are more pitfalls and nuances than "meet the eye". Overall, however, research benefits tremendously from advances in available satellite data as well as complementary tools, such as cloud-based platforms for data processing, and machine learning algorithms to detect phenomena and mapping variables. The focus of this keynote is to provide agricultural and environmental economists with an accessible introduction to working with satellite data, show-case applications, discuss advantages and weaknesses of satellite data, and emphasize best practices. This is supported by extensive Supplementary Materials, explaining the technical foundations, describing in detail how to create different variables, sketch out work flows, and a discussion of required resources and skills. Last but not least, example data and reproducible codes are available online.

JEL Codes: Q16, Q57, Q15, O13, Q55, L86



Copyright 2024 by David Wüpper, Wyclife Agumba Oluoch, and Hadi. All rights reserved. Readers may make verbatim copies of this document for non-commercial purposes by any means, provided that this copyright notice appears on all such copies.

Acknowledgements The authors acknowledge support by the European Research Council (ERC), Grant No. 101075824 (ERC Starting Grant LAND-POLICY), and the German Research Foundation (DFG), under Germany's Excellence Strategy EXC 2070, Grant No. 390732324 - PhenoRob. We thank Matin Qaim, Luke Sanford, Joel Ferguson, Jonathan Proctor, and Daniel Millimet for their suggestions and feedback.

1. Introduction

Satellite data has several features of interest to agricultural and environmental economists. It can provide measurements in high frequency and high resolution, over long periods of time and with wide spatial scope. It is also commonly more objective and methodologically unified than e.g., survey data. At times, it provides the only cost-effective and reliable measurement available (Burke et al., 2021; Donaldson & Storeygard, 2016; Lobell et al., 2020). Other times, it is a pivotal cross-check to verify other data, or it enables to add a critical variable to an existing dataset (Wuepper and Finger, 2023).



Fig.1. Croplands as seen from Space. (a) United States, (b) Uruguay, (c) Egypt, (d) India. Source: USGS Landsat via Google Earth Engine.



Fig.2. Quantifying Change over Time. Left: Croplands Near the Border Between South Africa and Lesotho in 2001. Right: Croplands Near the Border Between South Africa and Lesotho in 2018. Blue line indicates a section of the border between the two countries. Source: Wuepper et al. (2023).

Satellite data, however, cannot replace ground data and it is commonly only useful in combination (Nakalembe and Kerner, 2023; Kerner et al, 2024). Moreover, it comes with its own list of issues and pitfalls, such as new sources of measurement errors and inconsistencies, to name just two (Jain, 2020). As such, understanding the potential and limitations of satellite data has become an important skill for agricultural and environmental economists, to make the most out of this data source, and to avoid misconceptions.

The workflow of the here covered literature regularly involve at least two consecutive steps. First, satellite measurements come in the form of spatial layers e.g., showing croplands from different regions (**Fig. 1**) or different time periods (**Fig. 2**). The second step is to take these data and turn them into a table that can then be analyzed. This step is fundamentally straightforward, as there is a close correspondence between the two formats, and the same information can be stored alternatively in multiple layers or in a single table (**Fig. 3**).



Fig.3. Correspondence between Maps and Tables. (**A**) Each layer stores a variable. (**B**) each column stores unique identifier of a pixel or grid-cell and its coordinates (longitude and latitude) as well as all variables extracted. Source: (**A**) adapted from Vasiljević et al. (2020), (**B**) this study.

In contrast to survey data collection, satellite data can be "collected" relatively inexpensively and as frequently as commonly needed. The advancements in satellite technology have significantly reduced the cost of acquiring high resolution spatial data. As opposed to groundbased surveys that can be time consuming and expensive due to labor, travel, and equipment cost, satellite data offers a more affordable and efficient alternative. The economies of scale in satellite operations combined with the ability to cover large and remote areas efficiently make it highly cost-effective. Additionally, the continuous improvement in satellite sensors and data processing techniques has enhanced the accuracy and usability of satellite-derived information, further increasing its value as a survey tool. Finally, and especially useful, satellite data can always be collected close to spontaneously, in contrast to e.g., survey data, for which the data collection must be planned long in advance and once the data is collected, the data is what it is. With satellite data in contrast, a researcher might create a dataset and later realizes that an important variable is missing. This variable can be separately processed and added to the initial data. During peer-review, further data might be asked for (further outcome variables, e.g., or data on specific mechanisms). These too, can be added at any point in time.

An obvious and important limitation of satellite data is its inherent inability to measure immaterial variables such as preferences and beliefs, prices and contracts, property rights, and so on. The second important limitation of satellite data is that it critically requires ground measurements for validation, calibration, and training of algorithms. In the end, satellites mostly provide pictures and quite general measurements (things like colors, reflectance, etc.), and these need to be interpreted to become meaningful variables for economic research.

The work with satellite data often involves machine learning, e.g., to detect specific patterns in the data (classification), or to transform measures into the actual variables of interest (prediction). Examples include the identification of deforestation drivers, because deforestation looks different depending on its source (Curtis et al., 2018) and the mapping of cropland dynamics of over time (Potapov et al., 2022).

The aim of this keynote is to provide an accessible introduction and a coherent framework to using satellite data for research in agricultural and resource economics. We showcase different applications, cover satellite data advantages and disadvantages, potentials and risks, and distill best practices, together with more detailed explanations in the Supplementary Materials, and example code and data provided online.

Below, we showcase recent applications of using satellite data in agricultural and environmental economics (Section 2), discuss known pitfalls and potential remedies (Section 3) and conclude in Section 4. In the Supplementary Materials, we provide detailed descriptions of technical foundations, workflows, variable construction, and required skills and resources. On GitHub, we provide links to data and codes¹.

¹¹ https://github.com/land-economics-ilr-uni-bonn/sat-agri-econ

2. Applications

In this section, we discuss research using satellite data to measure outcomes, treatments, for causal identification, to test assumptions, and for practical applications. For reference, **Fig.4** exemplifies the differences in resolution between three of the most used satellite platforms, MODIS, Landsat, and Sentinel. Depending on the spatial scale of the analysis, required temporal frequency of measurement, and what the researcher is trying to measure, frequently at least one of these platforms is used. In addition, sometimes, very high-resolution data is needed, such as available from Planet Labs (2024).



Fig.4. Spatial resolutions among three satellite platforms; MODIS, Landsat and Sentinel-2, zoomed in to different levels of detail represented by panels A, B, and C. Both MODIS and Landsat images were acquired on the same day (30th June 2017) while Sentinel image was acquired a day earlier. They are displayed as red-green-blue (true color) composites. Source: (Luo et al., 2020)

a. Using Satellite Data to Measure Outcomes

A variable of major research and policy relevance is crop yields. Alternatively to the use of satellite measurements, one can measure crop yields on the ground, ask farmers to report them, or use official statistics, such as provided by the FAO (2024a). When relying on ground measures or farmer reports, the data is usually patchy and commonly only available for some selected farms. Official statistics on the other hand, are only available highly aggregated. Satellite data provides relatively rough but globally consistent, high resolution and high frequency, wall-to-wall crop yield measurements. This is usually based on either a vegetation index directly, such as the NDVI or the EVI (NASA LP DAAC, 2024a) or a derived product, such as NPP (NASA LP DAAC, 2024b), measured on global croplands, as e.g., identified by Potapov et al. (2022). An example of a large, global study with satellite measured crop yields as outcome variable is Wuepper et al. (2023). They are interested in how crop yields change in response to institutional changes and they measure crop yields globally as the annual maximum² of the EVI on 1 km² cropland grid-cells over 20 years. They estimate that a 10% improvement in countries' institutions improves crop yields by 2% on average globally.

Strobl and Strobl (2011) are interested in the impact of large dams in Africa, and they measure crop yield with the net primary production (NPP) on crop fields over time. NPP approximates biomass production, in terms of grams of carbon per m2. This can be converted into kilocalories, and hence to nutritional value. They estimate that upstream dams have on average provided up to 12% of the daily minimum per capita kilocalorie demand in downstream communities and increased agricultural production by 1%.

Asher et al. (2023) estimate the long-run development impact of India's irrigation canals. Their measure of crop yield is again based on the EVI and they find a larger crop yield effect in the dry season (+7.1%) and a smaller effect in the rainy season (+1.7%). The main long-run effect is a population increase around irrigated villages, advancing structural change.

To estimate the effect of Russia's Ukraine invasion on agricultural production, Deininger et al. (2023) use 4 years of panel data based on satellite imagery and extensive ground truthing data. Using machine learning techniques, they are then able to identify both location and extent of all conflict related damages to crops and fields, and both extent of area grown with crops and, for cereals fields, peak NDVI as crop yield measure. As a final step, they use statistical data to convert NDVI into actual cereal yield, measured in tons per hectare.

² To measure crop yield, it is usually more reliable to use maximum values and not the mean, as the mean can be biased by the measured color of a field when there is no full crop cover (fallow period, color being dominated by color of soil and not of the crop after planting, residue burning turning fields dark, etc.). A drawback is that this does not capture the yield obtained from multiple harvests.

Another example of a study that quantifies crop yields from satellite data is provided by Rampa and Lovo (2023). They evaluate the effect of a land tenure reform in Ethiopia. They rely on the NDVI as yield measure, at a resolution of 80 km² and over 13 years. They find that the land tenure reform increased crop yields by about 2% relative to the mean of the treated region.

Complimentary to the measurement of crop yields on current cropland is the identification of cropland changes, i.e., expansion and abandonment. In the study of Wuepper et al. (2023), the global land cover classifications from Friedl and Sulla-Menashe (2019) are used to estimate whether institutional changes lead to changes in the extent of countries' cropland area. The data provides 17 different land cover classes, which include two types of cropland. Wuepper et al. (2023) find that countries' cropland area (both classes together) tends to expand in response to institutional improvements, mostly at the expense of countries' forests.

Olsen et al. (2021) focus specifically on the impact of conflict-driven cropland abandonment on food insecurity in South Sudan. They combine Sentinel-1, Sentinel-2, Landsat-8 and MODIS imagery, as well as a Copernicus land-cover product, and use a random forest algorithm to identify all abandoned croplands. They find that the conflict caused 16% of croplands to be abandoned between 2016 and 2018.

Finally, Sukhtankar (2016) is an example for yet another related outcome: Identifying the area planted with a specific crop (sugar cane). Sukhtankar (2016) empirically investigates whether privately owned sugar mills affect farmers' crop choices and welfare differently than cooperative owned and public mills. It is found that private ownership of mills increases sugarcane production and farm incomes.

Huang et al. (2021) are interested in the impact of an anti-poverty program in Rural Kenya. To measure poverty, they combine satellite data with deep learning. What they catch are visible changes in housing quality that reflect improving household welfare. The alternative to their measurement approach would have been repeated in-person visits for interviews with a large number of people who can be difficult to reach. Huang et al. (2021) re-analyze a program for which such interviews exist too, and are able to show that they recover the same information. Whereas at the national level, night-lights data can already capture economic change, particularly in rural areas this does not work well, because there is often too little night light. Daytime imagery, however, shows what materials roofs are made of, and this in turn, is a reliable proxy for economic change, as Huang et al. (2021) show. Using their approach, they find the same, positive treatment effect of the program as does the comparison evaluation based on extensive survey work – at a much lower cost.

Ratledge et al. (2022) develop this further and estimate causal effect of electricity access on poverty in rural Uganda. They solve several issues, such as first, making sure that they

modelling of the outcome (poverty) is independent from their treatment (electrification). Would they have used night-lights, e.g., to predict poverty, they would have "build-in" their treatment effect into their outcome data. Secondly, they address the issue of attenuation bias that is often a feature of modelled data that contains less variance than actual outcome data.

Smith (2023) uses satellite data to study land values. Specifically, he is interested in the longrun effect of land concentration on the American frontier and he finds that historic land concentration leads to an increased reliance on tenant farming, lower investments, and lower land values that persist to today. To construct the land use values, he especially uses data from USDA (2024) (to measure land use) and FAO (2024b)(to measure productivity).

Ali et al. (2020) use Landsat 7 imagery over Ethiopia to construct pixel level annual and seasonal averages of normalized difference vegetation index (NDVI), soil-adjusted vegetation index (SAVI), and land surface water index (LSWI) to evaluate the performance of soil and water conservation measures. They find large treatment effects which they can also cross-verify with ground measures of river sediment loads.

Focusing on pastures in China, Hou et al. (2023) examine the effect of a land tenure reform on grassland quality, as measured by the NDVI. According to their estimates, the combination of privatization of land use rights and security protection improved grassland conditions by 3%.

Another outcome of major policy relevance are land-cover changes and particularly forest changes (deforestation, forest degradation, forest restoration). Satellite data combined with machine learning provides globally consistent forest measurements in high resolution and a long period of time, allowing policy evaluations based on difference-in-differences, difference in discontinuities, and similar. An exemplary study that uses such data is that of Groom et al. (2022). They evaluate Indonesia's moratorium on forest concessions, using difference-in-differences with the forest change dataset provided by Hansen et al. (2013). They find a small, positive effect on forest cover and that the moratorium was quite cost effective.

Wuepper et al. (2024) globally evaluate the effectiveness of all major public policies that countries implemented from 2001 to 2017. They use differences in discontinuities. They consider deforestation, forest degradation and improvement, and net primary production, at a spatial resolution of 1 km². Like Groom et al. (2022), they use the data of Hansen et al. (2013) to measure tree cover change, and changes in forest conditions are quantified with the Enhanced Vegetation Index (EVI), which measures the health of forests based on their color (NASA LP DAAC, 2024a), and Net Primary Productivity (NPP), which approximates how much carbon the vegetation accumulates (NASA LP DAAC, 2024b). Wuepper et al. (2024) estimate that countries' public policies reduced global tree cover loss by 4 percentage points - but with large between-country differences, mostly explained by policy stringency and enforcement.

Changes in tree and forest cover were among the first satellite measured outcomes economists analyzed and they continue to be highly utilized. For examples, Jayachandran et al. (2017) conducted a large randomized control trial in Uganda, analyzing how much deforestation can be reduced with payments for ecosystems (a lot), and West et al. (2023) how much voluntary carbon offset projects have actually contributed to forest conservation (a little). Studying 26 countries in Sub-Saharan Africa, Lundberg and Abman (2022) show that increased maize price volatility reduces deforestation rates. In Mexico, Chort and Öktem (2022) show that a coffee pest incentivized the diffusion of a resistant new plant variety that is adapted to sun exposure and does not need tree cover anymore, which fueled deforestation.

Another study, that of Abman and Lundberg (2024), also analyzes deforestation using the satellite derived data of Hansen et al. (2013) and additionally uses satellite data to actually come up with their study's control group. They use previously collected survey data on participating oil palm farmers in Ghana in a program that helped them improve their market access. The important limitation of their survey data was that it does not cover non-participating households or villages. To overcome this, they relied on remotely sensed settlement location data (GRID3, 2024), created Thiessen polygons around these coordinates (because exact settlement boundaries were not available), and used alternative approaches to select control observations from those settlements that were in the vicinity but not in their treatment group. They then use versions of difference-in-differences and the synthetic control method to estimate whether improved market access led to increased forest loss in the participating areas compared to non-participating but otherwise similar areas. This is indeed what they find.

Wilebore et al. (2019) use very high-resolution land cover data in their randomized control trial in which they provided unconditional conservation payments for communities near the Gola Rainforest Park on the border of Sierra Leone and Liberia. To measure agricultural land clearance, they use RapidEye multispectral satellite imagery with a resolution of 5x5m for the period before and after the intervention. The classification of land covers was based on a maximum likelihood classifier, using optical bands, texture metrics, and vegetation indices.

Also, agricultural air and water pollution can be derived from satellite measurements, enabling research e.g. on potential policy levers to mitigate these issues. Agricultural air pollution, e.g., is a major human health hazard. Nian (2023) uses high-resolution satellite data to measure agricultural fires at a resolution of 1 km² for all of China. The data is provided by China's Ministry of Ecology and Environment (MEE) which collects daily straw burning data from the moderate resolution imaging spectroradiometers (MODIS) of NASA's Satellites TERRA and AQUA. These satellites overpass China four times a day and a fire point is identified when a thermal anomaly is detected within a pixel using an algorithm that exploits the mid-infrared radiation from fires. Nian (2023) then estimates and compares the effectiveness of two different

mitigation levers: economic incentives and regulation. He estimates that demand from new biomass power plants reduced agricultural fires by 30% but a ban around airports was ineffective.



around harvest time

just after burning

one week later

Fig.5. Crop Residue Burning as seen on Sentinel-2 Imagery. Left image shows fields around harvest time, middle field shows the same fields directly after residues were burned, and the right image shows again the same fields one week later. Source: Jack et al. (2022)

In India, Jack et al. (2022) use high-resolution satellite imagery from Planet Labs (2024) and Sentinel-2 to identify residue burning on crop fields (See **Fig.5** above). Combining these two different measures is helpful, as PlanetScope data is measured more frequently (roughly every 2-3 days), so less likely to miss a burning event, whereas Sentinel-2 data is collected less frequently (every week to 10 days) but additionally provides information in the mid-infrared range, which helps to separate burned and unburned plots. Jack et al. (2022) then combine monitoring data and spot checks to train a random forest model to categorize fields into burned and not burned. This is used in a randomized control trial in which farmers are being paid not to burn their crop residues. To address liquidity constraints and farmer distrust, a portion of the money was paid unconditionally upfront. They find that this strategy reduces straw burning by 10 percentage points, in contrast to the standard contract, which has no effect.

Tang et al. (2023) are interested in the effect of intensive agriculture on water quality in China and exploit the fact that nutrient pollution in water creates algae blooms that are detectable on satellite imagery. The data on algae blooms is provided by Wang et al. (2023) who developed an algorithm to automatically detect algae blooms on MODIS Aqua images. Tang et al. (2023) find that a 1% increase in intensively farmed cropland leads to a 0.5 percent increase in the size of algae blooms and a 0.24% increase in their duration. A nation-wide policy to reduce fertilizer use is found to have mitigated the problem to some extent.

Advances in the availability of satellite data and complimentary machine learning tools has also reinvigorated and improved descriptive research. Often, this research is performed at the global level and the main product is a detailed map of a policy-relevant phenomenon, such as human development (Sherman et al., 2023) or the profitability of farming practices (Kamau et al., 2023).

For example, for a long time, the human development index has been measured at the country level using surveys and official statistics. More recently, finer measurements at the level of sub-national administrative units have been added (at high costs, based on extensive surveys). Sherman et al. (2023) use a machine learning based downscaling approach to produce a 0.1 \times 0.1- degree grid at the global level (see **Fig. 6** below). An interesting finding from this exercise is that apparently, aggregation bias has so far led to an erroneous HDI assignment of more than half of the global population within each country. For their computations, Sherman et al. (2023) build on the MOSAIKS approach (Rolf et al., 2021) and use both day and night time satellite imagery from Planet Labs (2024).



Fig.6. A High-Resolution Map of the Human Development Index. The Human Development Index (HDI) is a summary measure of average achievement in key dimensions of human development (income, education, health). Using statistical data, this is available at the national level. With satellite data and machine learning, it is available at the grid-cell level. Yellow indicates higher, purple lower development. Source: Sherman et al. (2023)

Other recent studies have mapped agricultural GDP (Ru et al., 2023), by first disaggregated different administrative statistics of agricultural GDP and then using satellite-derived indicators of crop, livestock, fishery, hunting and forestry production to produce a global gridded dataset showing local agricultural GDP all around the world, and the global mapping of the profitability of diversified versus simplified farming (Kamau et al. 2023).

b. Using Satellite Data to Measure Treatments

A treatment that is regularly measured using satellite data is agricultural pollution. He et al. (2020), e.g., use satellite data to detect agricultural straw burning in China (relying on the same

data as Nian (2023), which we discuss above). He et al. (2020) are interested in the health impact of the straw burning and find a considerable increase in deaths from cardiorespiratory diseases.

Another angle to the effect of air pollution is shown by Proctor (2021), who globally estimates the effect of changes in sunlight due to changes in clouds on crop productivity. He find that's average maize and soy yields have been reduced by 1% and 0.1% due to air pollution and could further decrease by 1.8% and 0.4% under further climate change. The main data source is the remotely sensed cloud data of Rossow et al. (2016).

In contrast, Taylor and Schlenker (2021) estimate how much CO_2 fertilization has contributed to agricultural productivity increases over time in the US and identify it as the main driver. The data comes from NASA (2024) and for identification, they exploit wind patterns, year-to-year anomalies from county-specific trends, and the spatial first-differences approach of Druckenmiller and Hsiang (2018).

Skidmore et al. (2023) estimate the effect of pesticide pollution on the incidence of childhood cancer in Brazil, focusing on the gradual spread of high-intensity soy production and exploiting that pesticide pollution always moves downstream and never upstream within a watershed. To do so, they use remotely sensed landcover maps from Mapbiomas (2023), which map soy, sugarcane, all other temporary crops, pasture, mining, and remaining natural vegetation, and they use remotely sensed watershed data from Instituto Brasileiro de Geografia e Estatística (2023). Skidmore et al. (2023) indeed find a robust connection between upstream agricultural intensification relying on toxic pesticides and downstream cancer incidences.

Another group of variables that is regularly measured with satellite data are weather-variables. For example, to correctly predict climate change impacts on global agricultural productivity, it is important to understand the impact of water supply on global crop yields and its relation to temperature stress. To answer this research question, Proctor et al. (2022) use satellite-measured soil moisture data measured by the European Space Agency (Dorigo et al., 2017; Gruber et al., 2019). This data provides surface soil moisture measures up to a depth of approximately 5 cm at daily $0.25^{\circ} \times 0.25^{\circ}$ resolution. Other used data includes crop yields from FAO (2024a) and temperature and rainfall data from NOAA (2024). Proctor et al. (2022) predict yield damages of -9% to -32% by the end of the century from changes in temperature and soil moisture, with each contributing roughly equally. In another study, Proctor et al. (2018) study the potential effect of geoengineering on global crop yields, using natural volcanic eruptions as a natural experiment. Their main satellite data are measures of stratospheric aerosol optical depth and cloud fraction data (Sato, 1993; World Radiation Data Centre, 2015). According to their estimates, the potential to increase crop yields with geoengineering is rather limited.

Anderson et al. (2021) combine a range of satellite measured variables such as soil moisture (from the Global Land Evaporation Amsterdam Model), precipitation (from the Climate Hazards group Infrared Precipitation with Stations), and the enhanced vegetation index (EVI, from the MODIS satellite MOD13A1 V6 product) with data on locusts, conflicts, and food insecurity. They then quantify the roles of conflict, drought, and locusts on food insecurity in Sub-Saharan Africa. One of their findings is that violent conflict exacerbated drought-related food insecurity.

Malacarne and Paul (2022) examine the impact of adopting improved management practices on food security and nutritional diversity in Tanzania and Mozambique – in general, and in the context of drought resilience. To measure household specific drought conditions, they alternatively use cumulative rainfall from Funk et al. (2015) and the Global Vegetation Health Index (VHI) from NOAA (2018). With both they find that improved management practices lead to better nutritional outcomes, but much less so under drought conditions.

Ali et al. (2018) estimate the costs and benefits of land fragmentation in Rwanda and measure their treatment variable, land fragmentation, as least-cost walking distance between homestead and parcels, using the global topography layer provided by the ASTER Science Team (2018). They find that land fragmentation negatively affects yield but also reduces yield variability, so it is unclear whether reducing land fragmentation would actually be a real improvement for Rwanda's farmers.

Deines et al. (2023) use satellite data to estimate the effect of cover crop adoption on maize and soybean yields for over 90,000 fields in the US between 2019 and 2020. Using causal forests analysis, they estimate an average yield loss of 5.5% for maize, and 3.5% for soy. Their identification of cover crop adoption is based on daily measures of the greenness of fields (using a fusion of Landsat and Modis data, with dynamic, phenology-based thresholds for the identification of cover cropping (Zhou et al., 2022).

Deines et al. (2019) use satellite data to estimate the impact of conservation tillage on crop yields. A map of tillage practices is provided by Azzari et al. (2019), who used a random forest classifier trained on ground data from 5866 fields in the region. For maps of crop yields, the Scalable Crop Yield Mapper (SCYM) of Lobell et al. (2015) was used. This is based on two steps. First, regionally parameterized crop models are used to statistically predict yields from crop phenology and climate covariates. Secondly, to generate yield estimates for each pixel, the models from step one are applied to satellite imagery and gridded climate datasets based on crop type maps. Deines et al. (2019) find a small, positive effect.

Chen et al. (2023) also estimate the impact of no-till adoption, but they focus on agricultural land values instead of yields. Their no-till adoption is measured by satellite too, but in a different way. Their measurement uses the percentage share of crop residues that high resolution

satellite imagery shows. This is used to classify cropland into categories such as no-till (more than 50% crop residues), or reduced tillage (less than 50% residues, but more than 15%), etc. (Hagen et al., 2020). Chen et al. (2023) estimate that no-till adoption significantly increases land values in the US.

c. Satellite Data for Causal Identification

Causal identification is regularly among the main challenges in applied economics research. In many data sets, it is not obvious whether there exists a feasible source of exogenous variation that allows to judge whether an observed relationship in the data is causal or not. Satellite data frequently provides such a source of exogenous variation. A classic example is the study by Duflo and Pande (2007), in which they estimate the agricultural impacts of irrigation dams in India. Clearly, the correlation between the regional number of dams and agricultural productivity has no causal interpretation because regions with more dams differ from regions with fewer dams in many dimensions (e.g., agricultural potential). What is available for causal identification here is an instrumental variable: The natural suitability of a river for the construction of a dam. The river gradient can be too high or too low for dam construction, and this is largely orthogonal to the plausible confounding factors one might worry about in the study's context. The earth's topography including the gradients of its rivers can be quantified with satellite data products, and Duflo and Pande (2007) use this as an instrumental variable to come to the conclusion that dam construction causally leads to a decline in poverty downstream of the dam, and an increase in poverty right where the dam is built. This strategy of exploiting arguably random variation in geography and the natural environment as instrumental variable is popular. Macchiavello and Morjaria (2020) estimate the effect of competition among Rwandan coffee mills on the performance of relational contracts and exploit geographic variation in the regional suitability to a build coffee mill. Rubio-Ramos (2023) estimates the effect of violence on social capital in the context of fluctuating coca production in Colombia. Here, the instrumental variable is an interaction between coca-market shocks and a suitability score for coca production, based on ecological conditions to grow coca, derived from satellite measurements.

A related strategy exploits satellite measured weather conditions to construct an instrumental variable. One example is the study of Crost and Felter (2019). They are interested in the effect of export crop revenue on inner-country conflict and they specifically focus on the case of banana exports on the Philippines. Here, insurgent groups use banana export revenues to finance their operations, but the issue is that the generated export revenues could potentially be endogenous, so they need exogenous shifts in banana prices. This is provided by random weather variation in Ecuador, the world's leading banana exporter. When growing conditions in Ecuador are bad, the global banana price goes up, increasing revenues on the Philippines.

Axbard (2016) studies the effect of fishermen's income on sea piracy. As an instrumental variable for fishermen's income, he uses satellite measured, exogenous changes in environmental fishing conditions. Satellite data provides chlorophyll-a concentrations and sea surface temperatures of the ocean. The constructed instrumental variable then measures how much of a given area had optimal values of each in a given month. The data comes from the NASA Modis Aqua satellite (Acker & Leptoukh, 2007).

Satellite measured geographical and environmental variation is equally helpful to identify causality in the context of regression discontinuity designs (see also Wuepper & Finger, 2023). An example is the study of Jones et al. (2022), in which they estimate the agricultural impacts of hillside irrigation systems in Rwanda. As in all example above, a naïve comparison of irrigated and non-irrigated farms and fields would be misleading, as there are other relevant differences. However, the studied irrigation systems use gravity to transport water to the fields, which can be exploited to set up a convincing regression discontinuity design, with the main irrigation channel marking the cut-off, and then there is a downhill side that has access to irrigation, and an uphill side, that does not. The discontinuity that arises between these two sides identifies the causal impact of irrigation. They find a large profit boost but limited adoption because of binding local labor market constraints.

The literature on agricultural pollution generally has the advantage that studies can exploit the directionality of the pollution, when the pollution is transported by wind or water currents. Sometimes, it can even be exploited that there are random deviations, e.g., wind directions can abruptly change.

A first example study is that of Ferguson and Govaerts (2024), who study the impact of adopting conservation agriculture in Mexico relying on various satellite data sources, measuring, among others, air quality and random wind direction changes. The latter is their source of exogenous variation. They combine this with complimentary data on the roll-out of relevant extension services, as well as agricultural profits, and infant deaths in down-win locations. They find that the extension services did successfully diffuse conservation agriculture and this was profitable for the farmers and advantageous for health and environmental outcomes down-wind. For their study, the important advantages of using satellite data is that they can rely on measurements from exactly the right locations (up-wind vs. down-wind) and timing (there is strong seasonality in agricultural burning). At the global level, Pullabhotla et al. (2023) estimate the burden of biomass fires on infant mortality, using georeferenced data on more than 2 million births matched to satellite-derived burned area exposure. Also in this study, the source of exogenous variation in exposure to agricultural air pollution are random wind direction changes. They find that each additional square kilometer of burning is associated with nearly 2% higher infant mortality in nearby downwind locations.

As demonstrated by Rangel and Vogl (2019) in Brazil, combing high-frequency changes in wind directions and fire locations is key in this literature, because agricultural activities, fires, and economics outcomes are so intertwined, that the pure panel variation in agricultural burnings is not sufficient to recover actual causal effects. Substantially, also in Brazil, the health effects of agricultural air pollution are large and important (Rangel & Vogl, 2019).

A different empirical strategy for causal identification is to explicitly model the counterfactual to an observed outcome and compare the two. Adamopoulos and Restuccia (2022) for example construct three global maps based on rich geospatial data mostly provided by the FAO (2024b), which show current agricultural yields, potential agricultural yields for the same crops, and potential yields for an optimized re-allocation of crops. They find that the currently observed global differences in agricultural productivity are not much explained by differences in natural growing conditions.

In a similar fashion, Wuepper et al. (2021) quantify global anthropogenic land degradation. Specifically, Wuepper et al. (2021) focus on four land degradation dimensions: soil erosion, deforestation, and carbon above and below ground. For each, they show two maps: One map of the current status, and one natural/historic benchmark. The difference between the two quantifies anthropogenic land degradation. To model the counterfactual scenario without anthropogenic land degradation, a key-input is the global land cover mix without humans. This is provided by Bastin et al. (2019), who use satellite imagery from protected areas all around the world to train a random forest algorithm which environmental characteristics are predictive of the natural vegetation seen in these protected areas. They then used this algorithm to "fill in the gaps" between the protected areas, to create a global land cover layer without humans.

d. Cross Validations and Assumptions Testing

Other useful applications of satellite data are to cross-validate an important measure or to create a variable that can be used to falsify an important modelling assumption. For example, it is highly common to rely on farmer self-reported data in agricultural economics, e.g. to measure field and farm sizes, crop yields, labor use, etc. - even though it is well known that these self-reports are affected by all sorts of measurement errors, including systematic ones (Abay et al., 2019; Gourlay et al., 2019).

From other outcome variable we know too that measures can deviate, e.g. we know that officially reported deforestation rates can be lower than actual ones to hide illegal deforestation, especially if the government is somehow involved (Balboni et al., 2021; Burgess et al., 2012), and that democracies report more reliable economic performance measures than dictatorships, with the latter having a positive bias (Martinez, 2022).

Having an alternative, satellite derived measure can be used to understand the reliability and accuracy of the reported measure of a variable. In addition, this even allows to disentangle misreporting and misperception, which is useful because sometimes misperception affects farmer behavior in a way that misreporting does not (Abay et al., 2021). Specifically focused on the measurement of crop yields, Lobell et al. (2020) demonstrate how well crop yields can be captured using satellite data so when this is an important variable and there could be measurement bias in the non-satellite measured crop yield data, it is advisable to compare measures.

As an example for a satellite data enabled test of an identifying assumption, Wuepper et al. (2024) use a spatial difference-in-discontinuities design to estimate the global impact of public policies on forest conservation. They show that their identified discontinuities in forest conditions at international border are not simply the result of discontinuities in forest potential using the data of Bastin et al. (2019), who in turn used photo-interpretations of high-resolution satellite imagery showing tree cover in protected areas, which were then used to train a machine learning algorithm to produce a global map of natural tree cover potential. In their study on the effect of countries on the global rate of soil erosion, Wuepper et al. (2020) use a similar approach but model their counterfactual even more, going from natural tree cover globally to natural land cover more generally and converting this to a natural rate of soil erosion globally at a spatial resolution of 1km². They show that globally, there would be no discontinuity in soil erosion at current international borders, so the actual discontinuity is caused by the countries.

Finally, satellite measured exogenous, environmental variables can also directly be used in placebo tests without any additional modeling of a specific counterfactual. This is common practice for example in regression discontinuity design applications (Wuepper & Finger, 2023). To give one specific example, Burgess et al. (2018) exploit deforestation discontinuities at Brazil's borders to estimate the effect of Brazil's forest policies. To test whether Brazil's border actually presents an exogenous treatment variable, they estimate whether there are discontinuities in the slope of the land or the distance to urban areas, water, or roads. Because they find no discontinuities in these exogenous variables, they conclude that the forest discontinuities they find are politically caused and not environmentally.

e. Satellite Data for Extension and Agricultural Risk Management

The satellite data products discussed thus far were created for research purposes. However, there are also active lines of research that analyze satellite data products with a practical purpose, e.g., informing insurance products, credit schemes, or extension services. A first example is the study of Jain et al. (2019), who demonstrate how to utilize high-resolution

satellite data (from SkySat and Planet Labs) to identify which fields would benefit the most from improving agricultural practices. With a randomized trial conducted in India, they show that the impact of an agricultural innovation (switching from manual fertilizer applications to a fertilizer spreader) could be doubled, simply by targeting the fields with the most potential instead of a uniform treatment. A second application example that shall be mentioned is index insurance. For a more in-depth review see the overview of Benami et al. (2021).

Vroege et al. (2021) demonstrate how freely available satellite data can be used to design an index insurance that effectively mitigates farmers' financial drought risk exposure. Using data from eastern Germany they also show that their insurance product based on satellite data outperforms a competing product based on meteorological measurements at ground stations.

With a randomized control trial, Boucher et al. (2024) explore the impact of providing drought tolerant seeds and satellite-based index insurance for farmers in Tanzania and Mozambique. They find that combining these two complementary technologies reduces a good share of farmers' drought risk exposure. Nevertheless, inducing sustained uptake proved to be a challenge, as a share of farmers who did not experience a drought early enough dis-adopted these technologies again.

The final example is the study of Möllmann et al. (2020), who investigate whether remotely sensed vegetation health indices can predict agricultural credit risk in Madagascar. They indeed find this to be the case, which indicates that it is possible to reduce the default risk to rural banks and thereby make them expand credit access.

3. Pitfalls and Remedies

a. Measurement Errors

Measurement error in satellite derived variables is ubiquitous (see Alix-Garcia & Millimet, 2023; Gordon et al., 2024; Proctor et al., 2023). Depending on whether the mismeasured variable is on the left or the right regression side and which kind of measurement error dominates (classical or systematic, and among those, mean reverting, or differential), estimated effects might be exaggerated or attenuated, and even sign flips are possible. Also computed standard errors can be too large or too small, depending on the kind of dominant measurement error and where the variable is located in the regression.

Researchers are especially worried about systematic measurement error, i.e., mismeasurement that correlates with the treatment variable. For example, the often-used global deforestation data of Hansen et al. (2013) contains measurement error that

systematically varies with topography, climate, and type of deforestation (Alix-Garcia & Millimet, 2023; Gordon et al., 2024; Pendrill et al., 2022). Depending on the application, this can seriously bias estimated effects. For example, it might be the case that the "treatment-forest" is in the wet tropics, and the "control-forest" is in the dry tropics, and the treatment coincides with the season in which the trees in the dry forest lose their leaves and are thus harder to detect. This could lead to the conclusion that the treatment conserved a lot of forest, even though only measurement error increased differentially.

Classical measurement error comes e.g., from random clouds blocking the view of the satellite, or because an available measure is just imprecise. For example, a satellite measurement might have a resolution of 100 hectares (1km²) whereas the actual variation of interest happens on agricultural fields with an average size of 3 or 4 hectares.

The more processed the used satellite data, the higher is the risk that additional measurement errors have been introduced. First of all, predicted variables tend to vary less than the actual phenomenon they are supposed to reflect. This is because extreme values get systematically underestimated, causing mean reverting error (Proctor et al., 2023; Ratledge et al., 2022). Secondly, the predictor variable used to create a variable of interest might directly include or exclude the effect of interest. For example, nighttime lights are a popular predictor to produce high resolution maps of economic development. However, in their study on the effect of rural electrification on economic development, Ratledge et al. (2022) emphasize that they cannot rely on nighttime lights to create their economic development variable, as nighttime-lights already directly measure electricity access. If one includes the treatment variable in the prediction model for the outcome variable, and then regresses this outcome on the treatment again, one does not estimate the causal effect of the treatment, but only mechanically recovers the previous prediction step. The opposite mistake is being made when a variable is modeled in such a way that the effect of interest is explicitly excluded (e.g. one wants to analyze whether more sustainable farming practices increase on-field biodiversity but the available biodiversity measure only models biodiversity habitat quality based on land cover types, so withinlandcover variability cannot possibly be detected).

There are a various ways to correct for the different kinds of measurement error in different contexts: Analyses of binary outcomes such as pixel-level deforestation coded as yes or no are especially likely to be biased by systematic measurement error. If one has some knowledge about the source of the measurement error (e.g., it might be connected to topography or cloud cover), one can use the misclassification logit or scobit models proposed by Alix-Garcia and Millimet (2023). These model the measurement error as a function of the problematic

covariates and off-the-shelf code is available for Stata³. Generally, the scobit does particularly well for the analysis of rare events (outcomes with many zeroes) but getting the model to converge can be challenging because an additional shape parameter needs to be estimated. To circumvent this issue, Alix-Garcia and Millimet (2023) recommend to treat the shape parameter of the model as unidentified and conduct a grid search across the range of potential values. The logit does not have this issue because the shape is fixed. A big advantage of using either of these models to correct for measurement error is that the correction is ex-post, in the analysis stage, and does not require any change to the data itself. Many other approaches correct the measurement ex-ante, during the dataset construction, as discussed below.

One example of an ex-ante correction measure is the augmented loss function of Ratledge et al. (2022), which is specifically set-up to correct for mean-reverting measurement error when modeling variables based on satellite data. Ratledge et al. (2022) add a penalty term to their loss function, which penalizes for bias in each quintile of the distribution. Thus, instead of having a global loss function that prioritizes the mean of the distribution to minimize prediction error, this augmented loss function explicitly treats each part of the distribution as equally important. The code is available for R^4

An especially general and versatile way to reduce bias from measurement error is multiple imputation (Proctor et al., 2023). This can be implemented with standard code in R, Stata, and Python⁵. The main requirement is that one has at least a small amount of ground truth data. As Proctor et al. (2023) demonstrate, not much of this calibration data is actually needed and it can come from locations somewhat far away.Fundamentally, the approach of Proctor et al. (2023) consists of two steps: First, the calibration data is used to estimate the structure of the measurement error present in the satellite data. Secondly, the estimated structure of the measurement error is used to correct parameter estimates and uncertainty measures in the subsequent regression analysis.

A final example for how one can mitigate measurement error induced bias is the approach proposed by Gordon et al. (2024), based on adversarial debiasing. The original inspiration for adversarial debasing was the recognized problem outside of agricultural economics that machine learning predictions can be biased (e.g., discriminating against personal characteristics such as race or gender in a job or credit application). To solve this problem, adversarial debiasing was invented, which, like the approach of Proctor et al. (2023), consists of two steps: First, a standard prediction model attempts to minimize prediction error for the

³ https://people.smu.edu/dmillimet/stata-code/

⁴ https://github.com/nwrat/RCQSB2022_public

⁵ In Stata, one can use the mi commands, in R one can use the mice package, in Python one can use the IterativeImputer within the scikit-learn library.

variable of interest (e.g. crop yield, soil erosion). However, non-standard, this first model has a loss function with a penalty that is informed by a second model. This second model – the adversary – is fed with the measurement errors from the first model and tries to predict treatment status for each observation. The better the second model can predict each observations treatment status, the larger a penalty is added to the loss function of the first model. Thus, the approach of Gordon et al. (2024) aims to minimize prediction error while also uncorrelating errors from treatment status. An Advantage of the adversarial debiasing approach – that it shares with the multiple imputation approach – is that there is no need to have any information about the source of the measurement bias. A drawback of the multiple imputation approach compared to the adversarial debiasing approach is that it is effectiveness depends on the specified functional relationship between measurement error and the dependent and independent variables (essentially, non-calibration data is treated as missing at random, conditional on covariates). However, the adversarial debiasing approach is computationally quite demanding, must be specifically adapted for each individual variable, and there is not yet an off-the-shelf code available.

b. Temporal Inconsistencies

While satellite data provides indispensable long-term, observational data going up to four decades back in time, users of the data ought to exercise caution when analyzing and inferring long-term trends from the data. This is because while data collection by satellite is by design a relatively consistent process (as compared to statistical surveys that might employ different sampling strategies, sample realization, and less objective self-reports), satellite-based time series data and the derived products are not without the issue of temporal inconsistencies. These temporal inconsistencies can be caused instrumentally by technical changes in the satellite platform and sensors, or analytically by changes in the processing and interpretation algorithms applied to transform the raw satellite data into specific thematic information, as well as if data from different sensors are fused together. In the first case, natural sensor degradation or satellite platform orbit drift has been known to affect long-term NDVI and leaf area index (LAI) trend from the AVHRR satellites (~9.2km resolution) widely used in studies examining trends from the 1980s, which has been dealt with in several data products (Cao et al., 2023; Wang et al., 2022). As an example, Landsat mission provides high-resolution (30m) data continuity with Landsat 5 (March 1984 to May 2012), Landsat 7 (July 1999 to April 2022), Landsat 8 (April 2013 to present), and Landsat 9 (October 2021 to present). While the subsequent Landsat sensors have been designed with continuity in mind (similar spectral and spatial resolution), fusing data from the different Landsat sensors requires careful consideration (also attention to the optimal processing level i.e., Landsat Collection 2 Level 2 Science Products (USGS, n.d.)). The scan line corrector failure in Landsat 7 has caused

systematic missing data (22% data loss per scene) since May 2003. Before Landsat 7 (i.e., before 1999), the earlier Landsat missions did not have a systematic global data acquisition plan, hence resulting in very sparse data record before 2000. In 2012, the number of captured images are particularly low due to decommissioning of Landsat 5 in 2011, and thus leaving only Landsat 7 acquiring images (typically two Landsat satellites are in operation at the same time) in the meantime before the launch of Landsat 8 in 2013. More generally, the different data density ultimately results in varying number and timing of clear-sky observations and hence presenting a source of variability in the monthly, seasonal, or yearly composites (e.g., mean), or phenology reconstruction. Landsat 8 also has a radiometrically improved sensor leading to better signal-to-noise ratio, which increases its sensitivity to the land surface features being targeted. In the second case, as developers of satellite data products improve their processing chain and interpretation algorithm (preprocessing e.g., cloud masking, or training data used in machine learning based product, or expert rules combining integrating other satellite based ancillary maps as they become available), temporal inconsistencies in the data products between subsequent updates (versions) can be expected. A notable example is the note for caution with regards to possible temporal inconsistencies in the Global Forest Change annual tree cover loss data products (Hansen et al., 2013), which may have a better detection of small-scale tree cover loss after 2013 in some regions due to the improved data from Landsat 8, as well as enhanced detection particularly from 2015 onwards due to improved algorithm (Weisse & Potapov, 2021). Thus, special care is warranted when interpreting longterm, year-to-year trend from the product (by triangulating/corroborating with multiple lines of evidence including official statistics), depending on the phenomenon of interest, inference objectives, and scales (Breidenbach et al., 2022; Ceccherini et al., 2020, 2021, 2022; Palahí et al., 2021). Similar cautions apply to other satellite-based time series data products that are increasingly produced such as annual land cover maps, recognizing probable misclassification errors on a per pixel level, in each year of the map (note however data products that use algorithms based on temporal segments that are by design more temporally consistent (Fried) et al., 2022; Gorelick et al., 2023; Zhang et al., 2024)). This is especially if the developers of the data products did not apply some temporal stabilization procedure e.g., by logical temporal transition (e.g., of land cover) filters in the algorithm design. Finally, beyond satellite-based data products, hybrid products (fusion of satellite, census, surveys, and models) also present temporal inconsistencies due to changes in input data used i.e., the underlying sub-national statistics, ancillary data (including satellite-based data such as cropland extent), or the downscaling/spatial allocation algorithm.

c. Irreversible Outcomes

A special case is the analysis of binary, nonrepeatable outcomes. One example is pixel-level

deforestation data that only contains the first year in which this happened (e.g., like in the popular Hansen et al. (2013) forest change dataset). Other examples could be adoption of irrigation (i.e., a pixel takes value zero if not yet equipped with irrigation infrastructure and one starting in the first year it gets equipped with irrigation infrastructure), or road construction. As Garcia and Heilmayr (2024) emphasize and demonstrate, this data structure can cause bias in common empirical frameworks such as difference-in-difference. The issue is the inclusion of individual unit fixed effects. This common practice interacts with the uncommon data structure such that one does not actually estimate the average treatment effect on the treated, but recovers an ex-post difference in outcomes between the treatment and control group. Fortunately, for this issue there are two alternative, simple solution: pixel-level specifications with spatially aggregated unit fixed effects and specifications with spatially aggregated units of analysis recover the true treatment effect (Garcia & Heilmayr, 2024).

d. Unobservables

A very clear weakness of satellite data is that many things cannot be seen on an image taken from space. This is yet another reason why satellite data is complementary to other data and is seldom sufficient as the sole data-source for an analysis. For example, for many questions in agricultural and resource economics, it is extremely useful to be able to objectively measure what happens on the fields in a region. One might record with high spatial and temporal resolution when and how a field was prepared, how it is managed throughout the season, and what agricultural and environmental outcomes can be observed at different points in time. What is missing, however, is to which farm the field belongs. Sometimes, fields are far away from the homestead and scattered widely, so one needs reliable information about the boundaries of each farm, if one wants to conduct a farm-level analysis. The farm-level is usually the most relevant unit to study farmers' decision making, so this is common. Also many variables that are relevant explanation for farmers' choices cannot be seen on any imagery. Preferences, behavioral biases, social interactions, culture, and personality are all important explanations for the behavior of farmers (Wuepper, Bukchin-Peles, Just, Zilberman 2023) and none of these can be measured with satellite data. Yet, combining cadastral data, farmer surveys, ground truth data, and then remote sensing measures can provide the kind of rich analyses that build on all the individual data strengths while compensating for all the individual data weaknesses.

4. Conclusion

Satellite data usually cannot be the only data-source but is immensely useful for many research endeavors when combined with other data (e.g., for ground truthing and to cover everything that cannot be seen on a satellite, such as the boundaries of individual farms, or farmer preferences, or property rights). Here, we have discussed many applications of satellite data and have attempted to build a basic understanding where this data comes from, and what to consider when working with it. We provide online code and data for practical exercises⁶ and there is an extensive Supplementary Materials file providing detailed discussion of work flows, construction of various variables, and a discussion of required resources and skills. Our hope is that this enables and inspired many more uses of satellite data in agricultural and environmental economics, as well as raise the average reliability of research in agricultural and environmental economics incorporating satellite data.

References

Abay, K. A., Abate, G. T., Barrett, C. B., & Bernard, T. (2019). Correlated non-classical measurement errors, 'Second best' policy inference, and the inverse size-productivity relationship in agriculture. Journal of Development Economics, 139, 171-184.

Abay, K. A., Bevis, L. E. M., & Barrett, C. B. (2021). Measurement Error Mechanisms Matter: Agricultural Intensification with Farmer Misperceptions and Misreporting. American Journal of Agricultural Economics, 103(2), 498-522. https://doi.org/https://doi.org/10.1111/ajae.12173

Abman, R., & Lundberg, C. (2024). Contracting, market access and deforestation. Journal of Development Economics, 103269.

Acker, J. G., & Leptoukh, G. (2007). Online analysis enhances use of NASA earth science data. Eos, Transactions American Geophysical Union, 88(2), 14-17.

Adamopoulos, T., & Restuccia, D. (2022). Geography and agricultural productivity: Crosscountry evidence from micro plot-level data. The Review of Economic Studies, 89(4), 1629-1653.

Ali, D. A., Deininger, K., & Monchuk, D. (2020). Using satellite imagery to assess impacts of soil and water conservation measures: Evidence from Ethiopia's Tana-Beles watershed. Ecological Economics, 169, 106512.

Ali, D. A., Deininger, K., & Ronchi, L. (2018). Costs and Benefits of Land Fragmentation: Evidence from Rwanda. The World Bank Economic Review, 33(3), 750-771. https://doi.org/10.1093/wber/lhx019

Alix-García, J., & Millimet, D. L. (2023). Remotely incorrect? Accounting for nonclassical measurement error in satellite data on deforestation. Journal of the Association of Environmental and Resource Economists, 10(5), 1335-1367.

Amatulli, G., Domisch, S., Tuanmu, MN. et al. (2018). A suite of global, cross-scale topographic variables for environmental and biodiversity modeling. Sci Data 5, 180040. https://doi.org/10.1038/sdata.2018.40

Anderson, W., Taylor, C., McDermid, S., Ilboudo-Nébié, E., Seager, R., Schlenker, W., Cottier, F., De Sherbinin, A., Mendeloff, D., & Markey, K. (2021). Violent conflict exacerbated drought-related food insecurity between 2009 and 2019 in sub-Saharan Africa. Nature food, 2(8), 603-615.

Asher, S., Campion, A., Gollin, D., & Novosad, P. (2023). The Long-Run Development Impacts of Agricultural Productivity Gains: Evidence from Irrigation Canals in India. Working Paper, Version September 26, 2023.

⁶ https://github.com/land-economics-ilr-uni-bonn/sat-agri-econ

ASTER Science Team. (2018). ASTER Global Digital Elevation Model V003. 2018, distributed by NASA EOSDIS Land Processes DAAC. https://doi.org/10.5067/ASTER/ASTGTM.003.

Axbard, S. (2016). Income opportunities and sea piracy in Indonesia: Evidence from satellite data. American Economic Journal: Applied Economics, 8(2), 154-194.

Azzari, G., Grassini, P., Edreira, J. I. R., Conley, S., Mourtzinis, S., & Lobell, D. B. (2019). Satellite mapping of tillage practices in the North Central US region from 2005 to 2016. Remote Sensing of Environment, 221, 417-429.

Balboni, C., Burgess, R., Heil, A., Old, J., & Olken, B. A. (2021). Cycles of fire? politics and forest burning in indonesia. AEA Papers and Proceedings,

Bastin, J.-F., Finegold, Y., Garcia, C., Mollicone, D., Rezende, M., Routh, D., Zohner, C. M., & Crowther, T. W. (2019). The global tree restoration potential. Science, 365(6448), 76-79.

Bilodeau, M.F., Esau, T.J., Zaman, Q.U. et al. (2024). Enhancing surface drainage mapping in eastern Canada with deep learning applied to LiDAR-derived elevation data. Sci Rep 14, 10016. https://doi.org/10.1038/s41598-024-60525-5

Borrelli, P., Robinson, D. A., Fleischer, L. R., Lugato, E., Ballabio, C., Alewell, C., Meusburger, K., Modugno, S., Schütt, B., & Ferro, V. (2017). An assessment of the global impact of 21st century land use change on soil erosion. Nature communications, 8(1), 2013.

Borrelli, P., Robinson, D. A., Panagos, P., Lugato, E., Yang, J. E., Alewell, C., Wuepper, D., Montanarella, L., & Ballabio, C. (2020). Land use and climate change impacts on global soil erosion by water (2015-2070). Proceedings of the National Academy of Sciences, 117(36), 21994-22001.

Boucher, S. R., Carter, M. R., Flatnes, J. E., Lybbert, T. J., Malacarne, J. G., Mareyna, P. P., & Paul, L. A. (2024). Bundling Genetic and Financial Technologies for More Resilient and Productive Small-Scale Farmers in Africa. The Economic Journal. https://doi.org/10.1093/ej/ueae012

Breidenbach, J., Ellison, D., Petersson, H., Korhonen, K. T., Henttonen, H. M., Wallerman, J., Fridman, J., Gobakken, T., Astrup, R., & Næsset, E. (2022). Harvested area did not increase abruptly—How advancements in satellite-based mapping led to erroneous conclusions. Annals of Forest Science, 79(1), 2. https://doi.org/10.1186/s13595-022-01120-4

Burgess, R., Costa, F. J., & Olken, B. A. (2018). Wilderness conservation and the reach of the state: evidence from national borders in the Amazon.

Burgess, R., Hansen, M., Olken, B. A., Potapov, P., & Sieber, S. (2012). The political economy of deforestation in the tropics. The Quarterly Journal of Economics, 127(4), 1707-1754.

Burke, M., Driscoll, A., Lobell, D. B., & Ermon, S. (2021). Using satellite imagery to understand and promote sustainable development. Science, 371(6535), eabe8628.

Cao, S., Li, M., Zhu, Z., Wang, Z., Zha, J., Zhao, W., Duanmu, Z., Chen, J., Zheng, Y., Chen, Y., Myneni, R. B., & Piao, S. (2023). Spatiotemporally consistent global dataset of the GIMMS leaf area index (GIMMS LAI4g) from 1982 to 2020. Earth System Science Data, 15(11), 4877–4899. https://doi.org/10.5194/essd-15-4877-2023

Ceccherini, G., Duveiller, G., Grassi, G., Lemoine, G., Avitabile, V., Pilli, R., & Cescatti, A. (2020). Abrupt increase in harvested forest area over Europe after 2015. Nature, 583(7814), 72–77. https://doi.org/10.1038/s41586-020-2438-y

Ceccherini, G., Duveiller, G., Grassi, G., Lemoine, G., Avitabile, V., Pilli, R., & Cescatti, A. (2021). Reply to Wernick, I. K. et al.; Palahí, M. et al. Nature, 592(7856), E18–E23. https://doi.org/10.1038/s41586-021-03294-9

Ceccherini, G., Duveiller, G., Grassi, G., Lemoine, G., Avitabile, V., Pilli, R., & Cescatti, A. (2022). Potentials and limitations of NFIs and remote sensing in the assessment of harvest rates: A reply to Breidenbach et al. Annals of Forest Science, 79(1), 31. https://doi.org/10.1186/s13595-022-01150-y

Chantarat, S., Mude, A. G., Barrett, C. B., & Carter, M. R. (2013). Designing index-based livestock insurance for managing asset risk in northern Kenya. Journal of Risk and Insurance, 80(1), 205-237.

Chen, L., Rejesus, R. M., Aglasan, S., Hagen, S., & Salas, W. (2023). The impact of no-till on agricultural land values in the United States Midwest. American Journal of Agricultural Economics, 105(3), 760-783.

Chort, I., & Öktem, B. (2022). Agricultural shocks, coping policies and deforestation: evidence from the coffee leaf rust epidemic in mexico. Coping Policies and Deforestation: Evidence from the Coffee Leaf Rust Epidemic in Mexico.

Crost, B., & Felter, J. H. (2019). Export Crops and Civil Conflict. Journal of the European Economic Association, 18(3), 1484-1520. https://doi.org/10.1093/jeea/jvz025

Curtis, P. G., Slay, C. M., Harris, N. L., Tyukavina, A., & Hansen, M. C. (2018). Classifying drivers of global forest loss. Science, 361(6407), 1108-1111.

Deines, J. M., Guan, K., Lopez, B., Zhou, Q., White, C. S., Wang, S., & Lobell, D. B. (2023). Recent cover crop adoption is associated with small maize and soybean yield losses in the United States. Global change biology, 29(3), 794-807.

Deines, J. M., Wang, S., & 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.

Deininger, K., Ali, D. A., Kussul, N., Shelestov, A., Lemoine, G., & Yailimova, H. (2023). Quantifying war-induced crop losses in Ukraine in near real time to strengthen local and global food security. Food Policy, 115, 102418.

Donaldson, D., & Storeygard, A. (2016). The view from above: Applications of satellite data in economics. Journal of Economic Perspectives, 30(4), 171-198.

Dorigo, W., Wagner, W., Albergel, C., Albrecht, F., Balsamo, G., Brocca, L., Chung, D., Ertl, M., Forkel, M., & Gruber, A. (2017). ESA CCI Soil Moisture for improved Earth system understanding: State-of-the art and future directions. Remote Sensing of Environment, 203, 185-215.

Druckenmiller, H., & Hsiang, S. (2018). Accounting for unobservable heterogeneity in cross section using spatial first differences.

Duflo, E., & Pande, R. (2007). Dams. The Quarterly Journal of Economics, 122(2), 601-646.

DWD. (2021). Climate Data Center. Deutscher Wetterdienst (DWD). https://www.dwd.de/DE/klimaumwelt/cdc/cdc_node.html.

FAO. (2024a). FAOSTAT: Crops and livestock products. https://www.fao.org/faostat/en/#data/QCL.

FAO. (2024b). Global Agro-Ecological Zones (AEZ) Database. https://gaez.fao.org/.

Ferguson, J., & Govaerts, B. (2024). Agricultural Intensification Without Environmental Degradation: Evidence from Mexico. Working Paper, UC Berkeley.

Friedl, M., & Sulla-Menashe, D. (2019). MCD12Q1 MODIS/Terra+ aqua land cover type yearly L3 global 500m SIN grid V006. NASA EOSDIS Land Processes DAAC, 10, 200.

Funk, C., Peterson, P., Landsfeld, M., Pedreros, D., Verdin, J., Shukla, S., Husak, G., Rowland, J., Harrison, L., & Hoell, A. (2015). The climate hazards infrared precipitation with stations—a new environmental record for monitoring extremes. Scientific Data, 2(1), 1-21.

Garcia, A., & Heilmayr, R. (2024). Impact evaluation with nonrepeatable outcomes: The case of forest conservation. Journal of Environmental Economics and Management, 125, 102971.

Gordon, M., Ayers, M., Stone, E., & Sanford, L. (2024). Remote Control: Debiasing Remote Sensing Predictions for Causal Inference. Working Paper Yale School of the Environment.

Gourlay, S., Kilic, T., & Lobell, D. B. (2019). A new spin on an old debate: Errors in farmerreported production and their implications for inverse scale-Productivity relationship in Uganda. Journal of Development Economics, 141, 102376.

GRID3. (2024). GRID3 DATA HUB. https://data.grid3.org/.

Groom, B., Palmer, C., & Sileci, L. (2022). Carbon emissions reductions from Indonesia's moratorium on forest concessions are cost-effective yet contribute little to Paris pledges. Proceedings of the National Academy of Sciences, 119(5), e2102613119.

Gruber, A., Scanlon, T., van der Schalie, R., Wagner, W., & Dorigo, W. (2019). Evolution of the ESA CCI Soil Moisture climate data records and their underlying merging methodology. Earth System Science Data, 11(2), 717-739.

Hadi, Yowargana, P., Zulkarnain, M. T., Mohamad, F., Goib, B. K., Hultera, P., Sturn, T., Karner, M., Dürauer, M., See, L., Fritz, S., Hendriatna, A., Nursafingi, A., Melati, D. N., Prasetya, F. V. A. S., Carolita, I., Kiswanto, Firdaus, M. I., Rosidi, M., & Kraxner, F. (2022). A national-scale land cover reference dataset from local crowdsourcing initiatives in Indonesia. Scientific Data, 9(1), 574. https://doi.org/10.1038/s41597-022-01689-5

Hagen, S. C., Delgado, G., Ingraham, P., Cooke, I., Emery, R., P. Fisk, J., Melendy, L., Olson, T., Patti, S., & Rubin, N. (2020). Mapping conservation management practices and outcomes in the corn belt using the operational tillage information system (OpTIS) and the denitrification–decomposition (DNDC) model. Land, 9(11), 408.

Hansen, M. C., Potapov, P. V., Moore, R., Hancher, M., Turubanova, S. A., Tyukavina, A., Thau, D., Stehman, S. V., Goetz, S. J., & Loveland, T. R. (2013). High-resolution global maps of 21st-century forest cover change. Science, 342(6160), 850-853.

He, G., Liu, T., & Zhou, M. (2020). Straw burning, PM2.5, and death: Evidence from China. Journal of Development Economics, 145, 102468. https://doi.org/https://doi.org/10.1016/j.jdeveco.2020.102468

Homolová, L., Malenovský, Z., Clevers, J. G. P. W., García-Santos, G., & Schaepman, M. E. (2013). Review of optical-based remote sensing for plant trait mapping. Ecological Complexity, 15, 1–16. https://doi.org/10.1016/j.ecocom.2013.06.003

Hou, L., Liu, P., & Tian, X. (2023). Grassland tenure reform and grassland quality in China. American Journal of Agricultural Economics, 105(5), 1388-1404.

Huang, L. Y., Hsiang, S. M., & Gonzalez-Navarro, M. (2021). Using satellite imagery and deep learning to evaluate the impact of anti-poverty programs.

Instituto Brasileiro de Geografia e Estatística. (2023). Bacias Hidrográficas do Brasil – BHB250. Bacias e Divisões Hidrográficas do Brasil.

https://www.ibge.gov.br/geociencias/cartas-emapas/informacoes-ambientais/31653-baciase-divisoes-hidrograficas-do-brasil.html. Jack, B. K., Jayachandran, S., Kala, N., & Pande, R. (2022). Money (Not) to Burn: Payments for Ecosystem Services to Reduce Crop Residue Burning.

Jain, M. (2020). The benefits and pitfalls of using satellite data for causal inference. Review of Environmental Economics and Policy.

Jayachandran, S., De Laat, J., Lambin, E. F., Stanton, C. Y., Audy, R., & Thomas, N. E. (2017). Cash for carbon: A randomized trial of payments for ecosystem services to reduce deforestation. Science, 357(6348), 267-273.

Jain, M., Rao, P., Srivastava, A. K., Poonia, S., Blesh, J., Azzari, G., McDonald, A. J., & Lobell, D. B. (2019). The impact of agricultural interventions can be doubled by using satellite data. Nature sustainability, 2(10), 931-934.

Jones, M., Kondylis, F., Loeser, J., & Magruder, J. (2022). Factor market failures and the adoption of irrigation in rwanda. American Economic Review, 112(7), 2316-2352.

Kamau, H., Roman, S., & Biber-Freudenberger, L. (2023). Nearly half of the world is suitable for diversified farming for sustainable intensification. Communications Earth & Environment, 4(1), 446.

Kazemi Garajeh, M., Salmani, B., Zare Naghadehi, S. et al. (2023). An integrated approach of remote sensing and geospatial analysis for modeling and predicting the impacts of climate change on food security. Sci Rep 13, 1057. https://doi.org/10.1038/s41598-023-28244-5

Keller, J. B., & Saitone, T. L. (2022). Basis risk in the pasture, rangeland, and forage insurance program: Evidence from California. American Journal of Agricultural Economics, 104(4), 1203-1223.

Kerner, H., Nakalembe, C., Yang, A. et al. (2024). How accurate are existing land cover maps for agriculture in Sub-Saharan Africa?. Sci Data 11, 486. https://doi.org/10.1038/s41597-024-03306-z

Lesiv, M., Laso Bayas, J. C., See, L., Duerauer, M., Dahlia, D., Durando, N., Hazarika, R., Kumar Sahariah, P., Vakolyuk, M. y., Blyshchyk, V., Bilous, A., Perez-Hoyos, A., Gengler, S., Prestele, R., Bilous, S., Akhtar, I. u. H., Singha, K., Choudhury, S. B., Chetri, T., . . . Fritz, S. (2019). Estimating the global distribution of field size using crowdsourcing. Global Change Biology, 25(1), 174-186. <u>https://doi.org/https://doi.org/10.1111/gcb.14492</u>

Lobell, D. B., Azzari, G., Burke, M., Gourlay, S., Jin, Z., Kilic, T., & Murray, S. (2020). Eyes in the sky, boots on the ground: Assessing satellite-and ground-based approaches to crop yield measurement and analysis. American Journal of Agricultural Economics, 102(1), 202-219.

Lobell, D. B., Thau, D., Seifert, C., Engle, E., & Little, B. (2015). A scalable satellite-based crop yield mapper. Remote Sensing of Environment, 164, 324-333.

Lundberg, C., & Abman, R. (2022). Maize price volatility and deforestation. American Journal of Agricultural Economics, 104(2), 693-716.

Luo, Y., Guan, K., Peng, J., Wang, S., & Huang, Y. (2020). STAIR 2.0: A generic and automatic algorithm to fuse modis, landsat, and Sentinel-2 to generate 10 m, daily, and cloud-/gap-free surface reflectance product. Remote Sensing, 12(19), 3209.

Macchiavello, R., & Morjaria, A. (2020). Competition and Relational Contracts in the Rwanda Coffee Chain. The Quarterly Journal of Economics, 136(2), 1089-1143. https://doi.org/10.1093/qje/qjaa048

Malacarne, J. G., & Paul, L. A. (2022). Do the benefits of improved management practices to nutritional outcomes "dry up" in the presence of drought? Evidence from East Africa. Food Policy, 113, 102332. https://doi.org/https://doi.org/10.1016/j.foodpol.2022.102332

Mapbiomas. (2023). Mapbiomas Collection 7.0. https://brasil.mapbiomas.org/en/colecoes-mapbiomas/.

Martinez, L. R. (2022). How much should we trust the dictator's GDP growth estimates? Journal of Political Economy, 130(10), 2731-2769.

Mello, F.A.O., Demattê, J.A.M., Bellinaso, H. et al. (2023). Remote sensing imagery detects hydromorphic soils hidden under agriculture system. Sci Rep 13, 10897. https://doi.org/10.1038/s41598-023-36219-9

Möllmann, J., Buchholz, M., Kölle, W., & Musshoff, O. (2020). Do remotely-sensed vegetation health indices explain credit risk in agricultural microfinance? World Development, 127, 104771. https://doi.org/https://doi.org/10.1016/j.worlddev.2019.104771

Nakalembe, C., & Kerner, H. (2023). Considerations for AI-EO for agriculture in Sub-Saharan Africa. Environmental Research Letters, 18(4), 041002.

NASA LP DAAC. (2024a). MOD13A2.061 Terra Vegetation Indices 16-Day Global 1km. https://developers.google.com/earth-engine/datasets/catalog/MODIS_061_MOD13A2.

NASA LP DAAC. (2024b). MOD17A3HGF.061: Terra Net Primary Production Gap-Filled Yearly Global 500m. https://developers.google.com/earthengine/datasets/catalog/MODIS 061 MOD17A3HGF.

National Aeronautics and Space Administration (NASA). (2024). About MODIS.

Nian, Y. (2023). Incentives, penalties, and rural air pollution: Evidence from satellite data. Journal of Development Economics, 161, 103049.

NOAA. (2018). Global Vegetation Health Products. STAR, Center for Satellite Applications and Research, NOAA, USA.

NOAA. (2024). CPC Global Daily Temperature and Precipitation. https://psl.noaa.gov/data/gridded/data.cpc.globaltemp.html and https://psl.noaa.gov/data/gridded/data.cpc.globalprecip.html.

Ojha, T., Misra, S., Raghuwanshi, N. S. (2015). Wireless sensor networks for agriculture: the state-of-the-art in practice and future challenges. Comput. Electron. Agric. 118 pp. 66 - 84 https://doi.org/10.1016/j.compag.2015.08.011

Olsen, V. M., Fensholt, R., Olofsson, P., Bonifacio, R., Butsic, V., Druce, D., Ray, D., & Prishchepov, A. V. (2021). The impact of conflict-driven cropland abandonment on food insecurity in South Sudan revealed using satellite remote sensing. Nature Food, 2(12), 990-996.

Palahí, M., Valbuena, R., Senf, C., Acil, N., Pugh, T. A. M., Sadler, J., Seidl, R., Potapov, P., Gardiner, B., Hetemäki, L., Chirici, G., Francini, S., Hlásny, T., Lerink, B. J. W., Olsson, H., González Olabarria, J. R., Ascoli, D., Asikainen, A., Bauhus, J., ... Nabuurs, G.-J. (2021). Concerns about reported harvests in European forests. Nature, 592(7856), E15–E17. https://doi.org/10.1038/s41586-021-03292-x

Pendrill, F., Gardner, T. A., Meyfroidt, P., Persson, U. M., Adams, J., Azevedo, T., Bastos Lima, M. G., Baumann, M., Curtis, P. G., & De Sy, V. (2022). Disentangling the numbers behind agriculture-driven tropical deforestation. Science, 377(6611), eabm9267.

Planet Labs. (2024). https://www.planet.com/.

Potapov, P., Turubanova, S., Hansen, M. C., Tyukavina, A., Zalles, V., Khan, A., Song, X.-P., Pickens, A., Shen, Q., & Cortez, J. (2022). Global maps of cropland extent and change show accelerated cropland expansion in the twenty-first century. Nature Food, 3(1), 19-28.

Proctor, J. (2021). Atmospheric opacity has a nonlinear effect on global crop yields. Nature Food, 2(3), 166-173. https://doi.org/10.1038/s43016-021-00240-w

Proctor, J., Carleton, T., & Sum, S. (2023). Parameter recovery using remotely sensed variables. NBER Working Paper 30861, http://www.nber.org/papers/w30861

Proctor, J., Hsiang, S., Burney, J., Burke, M., & Schlenker, W. (2018). Estimating global agricultural effects of geoengineering using volcanic eruptions. Nature, 560(7719), 480-483.

Proctor, J., Rigden, A., Chan, D., & Huybers, P. (2022). More accurate specification of water supply shows its importance for global crop production. Nature Food, 3(9), 753-763. https://doi.org/10.1038/s43016-022-00592-x

Pullabhotla, H. K., Zahid, M., Heft-Neal, S., Rathi, V., & Burke, M. (2023). Global biomass fires and infant mortality. Proceedings of the National Academy of Sciences, 120(23), e2218210120.

Rampa, A., & Lovo, S. (2023). Revisiting the effects of the Ethiopian land tenure reform using satellite data. A focus on agricultural productivity, climate change mitigation and adaptation. World Development, 171, 106364.

https://doi.org/https://doi.org/10.1016/j.worlddev.2023.106364

Rangel, M. A., & Vogl, T. S. (2019). Agricultural fires and health at birth. Review of Economics and Statistics, 101(4), 616-630.

Ratledge, N., Cadamuro, G., de la Cuesta, B., Stigler, M., & Burke, M. (2022). Using machine learning to assess the livelihood impact of electricity access. Nature, 611(7936), 491-495.

Rolf, E., Proctor, J., Carleton, T., Bolliger, I., Shankar, V., Ishihara, M., Recht, B., & Hsiang, S. (2021). A generalizable and accessible approach to machine learning with global satellite imagery. Nature communications, 12(1), 4392.

Rossow, W., Walker, A., Golea, V., Knapp, K., Young, A., Inamdar, A., & Hankins, B. (2016). International Satellite Cloud Climatology Project Climate Data Record, H-Series, NOAA National Centers for Environmental Information.

Ru, Y., Blankespoor, B., Wood-Sichra, U., Thomas, T. S., You, L., & Kalvelagen, E. (2023). Estimating local agricultural gross domestic product (AgGDP) across the world. Earth Syst. Sci. Data, 15(3), 1357-1387. https://doi.org/10.5194/essd-15-1357-2023

Rubio-Ramos, M. (2023). Trust, violence, and coca. Journal of Development Economics, 103216. https://doi.org/https://doi.org/10.1016/j.jdeveco.2023.103216

Sato, M., Hansen, J. E., McCormick, M. P. & Pollack, J. B. (1993). Stratospheric aerosol optical depths 1850–1990. J. Geophys. Res. 98, 22987–22994.

Sherman, L., Proctor, J., Druckenmiller, H., Tapia, H., & Hsiang, S. M. (2023). Global High-Resolution Estimates of the United Nations Human Development Index Using Satellite Imagery and Machine-learning.

Skidmore, M. E., Sims, K. M., & Gibbs, H. K. (2023). Agricultural intensification and childhood cancer in Brazil. Proceedings of the National Academy of Sciences, 120(45), e2306003120. https://doi.org/doi:10.1073/pnas.2306003120

Smith, C. (2023). Land concentration and long-run development in the frontier united states. Working Paper University of Maryland.

Strobl, E., & Strobl, R. O. (2011). The distributional impact of large dams: Evidence from cropland productivity in Africa. Journal of Development Economics, 96(2), 432-450. https://doi.org/https://doi.org/10.1016/j.jdeveco.2010.08.005 Sukhtankar, S. (2016). Does firm ownership structure matter? Evidence from sugar mills in India. Journal of Development Economics, 122, 46-62.

Tang, C., Guo, Y., Feng, L., & Keiser, D. (2023). Quantifying the Water Quality Impacts of Intensive Farming in China: A Satellite Data Approach. Available at SSRN 4498467.

Taylor, C. A., & Schlenker, W. (2021). Environmental Drivers of Agricultural Productivity Growth: CO₂ Fertilization of US Field Crops.

USDA. (2024). CropScape.

https://www.nass.usda.gov/Research_and_Science/Cropland/SARS1a.php.

Vasiljević, N., Vignaroli, A., Bechmann, A., & Wagner, R. (2020). Digitalization of scanning lidar measurement campaign planning. Wind Energy Science, 5(1), 73-87.

Vroege, W., Bucheli, J., Dalhaus, T., Hirschi, M., & Finger, R. (2021). Insuring crops from space: the potential of satellite-retrieved soil moisture to reduce farmers' drought risk exposure. European Review of Agricultural Economics, 48(2), 266-314. https://doi.org/10.1093/erae/jbab010

Wang, Y., Feng, L., & Hou, X. (2023). Algal Blooms in Lakes in China Over the Past Two Decades: Patterns, Trends, and Drivers. Water Resources Research, 59(10), e2022WR033340. https://doi.org/10.1029/2022WR033340

Wang, Z., Wang, H., Wang, T., Wang, L., Liu, X., Zheng, K., & Huang, X. (2022). Large discrepancies of global greening: Indication of multi-source remote sensing data. Global Ecology and Conservation, 34, e02016. https://doi.org/10.1016/j.gecco.2022.e02016

Weiss, M., Jacob, F., & Duveiller, G. (2019). Remote sensing for agricultural applications: A meta-review. Remote Sensing of Environment, 236, 111402. https://doi.org/10.1016/j.rse.2019.111402

Weisse, M., & Potapov, P. (2021, April 28). Assessing Trends in Tree Cover Loss Over 20 Years of Data. https://www.globalforestwatch.org/blog/data/tree-cover-loss-satellite-data-trend-analysis/

West, T. A. P., Wunder, S., Sills, E. O., Börner, J., Rifai, S. W., Neidermeier, A. N., Frey, G. P., & Kontoleon, A. (2023). Action needed to make carbon offsets from forest conservation work for climate change mitigation. Science, 381(6660), 873-877. https://doi.org/doi:10.1126/science.ade3535

Wilebore, B., Voors, M., Bulte, E. H., Coomes, D., & Kontoleon, A. (2019). Unconditional transfers and tropical forest conservation: Evidence from a randomized control trial in Sierra Leone. American Journal of Agricultural Economics, 101(3), 894-918.

Wolfert, S., Ge, L., Verdouw, C., Bogaardt, M.-J. (2017). Big data in smart agriculture - review Agric. Syst., 153 pp. 69 - 80 https://doi.org/10.1016/j.agsy.2017.01.023

World Radiation Data Centre. (2015). Global radiation. Daily sums, monthly sums and means. World Meteorological Organization http://wrdc.mgo.rssi.ru/.

Wuepper, D., & Finger, R. (2023). Regression discontinuity designs in agricultural and environmental economics. European Review of Agricultural Economics, 50(1), 1-28.

Wuepper, D., Borrelli, P., & Finger, R. (2020). Countries and the global rate of soil erosion. Nature sustainability, 3(1), 51-55.

Wuepper, D., Borrelli, P., Panagos, P., Lauber, T., Crowther, T., Thomas, A., & Robinson, D. A. (2021). A 'debt' based approach to land degradation as an indicator of global change. Global Change Biology, 27(21), 5407-5410.

Wuepper, D., Bukchin-Peles, S., Just, D., & Zilberman, D. (2023). Behavioral agricultural economics. Applied Economic Perspectives and Policy, 45(4), 2094-2105.

Wuepper, D., Crowther, T., Lauber, T., Routh, D., Le Clec'h, S., Garrett, R. D., & Börner, J. (2024). Public policies and global forest conservation: Empirical evidence from national borders. Global Environmental Change, 84, 102770. https://doi.org/https://doi.org/10.1016/j.gloenvcha.2023.102770

Wuepper, D., Wang, H., Schlenker, W., Jain, M., & Finger, R. (2023). Institutions and Global Crop Yields. National Bureau of Economic Research.

Zhou, Q., Guan, K., Wang, S., Jiang, C., Huang, Y., Peng, B., Chen, Z., Wang, S., Hipple, J., & Schaefer, D. (2022). Recent rapid increase of cover crop adoption across the US Midwest detected by fusing multi-source satellite data. Geophysical Research Letters, 49(22), e2022GL100249.

SUPPLEMENTARY MATERIALS

Satellite Data in Agricultural and Environmental Economics: Theory and Practice

Keynote ICAE 2024, New Delhi

David Wuepper^{1,*}, Wyclife Agumba Oluoch¹, Hadi¹

¹ Land Economics Group, University of Bonn

* Corresponding author: wuepper@uni-bonn.de

28.6.2024

CONTENTS

S1 Foundations (pages 1 -6) S2 Typical Work Flows (pages 6 - 11) S3 Variable Construction (pages 11 - 25)

S4 Required Skills and Resources (pages 25 - 32)

S1 Foundations

The notable growth in remote sensing relevant for agriculture dates back to 1972 with the development of Earth Resources Technology Satellite (ERTS-1) - later in 1975 called Land Satellite (Landsat) - to better reflect its mission to observe and monitor the earth surface. Landsat brought the greatest change for civilian use due to open mapping at 30 m spatial resolution with about 8 days revisit time (same location imaged every 8 days).

Since then, a number of improvements have been achieved, especially in terms of spatial and spectral resolutions. This has made it possible to investigate a number of aspects of land surface features such as crop types, stage of growth, species, health condition, drought stress, conversion of land from one use to another (for instance deforestation for agricultural expansion) among others at finer details that were not possible before. Here we look at a selected satellite platforms with sensors at spectral and spatial resolutions that are relevant for agriculture. We pay specific attention to free and open access data from both Sentinel and MODIS which everyone can access and use without proprietary constraints. We pay more attention to these because of their difference in spatial and temporal resolutions with Sentinel having higher spatial resolution but relatively lower temporal resolution while MODIS has higher temporal resolution but lower spatial resolution. There are also other distinct advantages and drawbacks of these and other platforms. Below, we discuss briefly the Sentinel platform (a), the MODIS platform (b), and different direct measures, called signatures (c)

a. SENTINEL

Sentinel is a series of missions for earth observation that were developed by the European Space Agency (ESA) in a collaborative project with the European Commission as part of the European Union's Copernicus program (ESA 2024). The aim of the project was to provide timely and reliable information for monitoring and managing the environment, climate change, and civil security. There are already six Sentinel missions designed for specific use cases depending on the specifications of their sensors and orbital alignment.

Sentinel-1: This is equipped with a c-band Synthetic Aperture Radar (SAR) instrument that enables it to monitor the earth surface irrespective of weather conditions since 2014. It penetrates the cloud, the major problem for optical sensors, and is thus primarily preferred in regions of high cloud cover, detecting land cover changes, ground motion, and emergency response strategies during natural disasters like floods (Kumar et al. 2022). During flood events, there is normally corresponding high cloud cover that renders optical sensors inappropriate to monitor flood extent on the ground and assess damage, for instance within cropland. Sentinel-1 has three spatial resolutions (10, 25, and 40 m). The scenes come pre-
corrected for thermal noise, radiometric calibration and terrain, normally achieved using the 30 m Shuttle Radar Topography Mission (SRTM 30) product.

Sentinel-2: This has a multispectral imaging instrument with the ability to capture very high resolution (10 m) optical imagery since 2017. It is specifically useful for monitoring agricultural lands, land cover, and other environmental changes. Every scene of Sentinel 2 image has a swath of 290 km hence quite appropriate for regional crop mapping expeditions. This is far wider and at a finer spatial resolution (10 m) and spectral resolution (13 bands) than for Landsat 5 or 7 that was 185 km at coarser spatial resolution of 30 m and limited to eight spectral bands. The revisit time for this satellite is 10 days while the whole constellation has a shorter revisit period of 5 days.

Sentinel-3: This sentinel mission has Ocean and Land Color Instrument (OLCI) and is thus focused on ocean and land surface monitoring since 2016. It has a collection of instruments including radiometers and altimeters as well as optical and microwave sensors to achieve its missions. It gives important data for monitoring of sea surface temperature, ocean color, sea level variations and land cover changes. However, it captures data at spatial resolution of 300 m and has global temporal resolution of 2 days. This is making it suitable for regional scale analyses of sea and land temperatures, environmental monitoring, including wildfire. while it may not be the primary Sentinel mission for agriculture, it can be very useful for farms located close by coastal areas where the influence of sea or ocean on crop performance is instrumental.

Sentinel-4 and 5: These two missions are focused on monitoring composition of the earth's atmosphere. while Sentinel-4 focuses on monitoring of aerosols and trace gas concentrations within the troposphere of Europe. Sentinel-5 focusses on pollutants and greenhouse gasses monitoring throughout the globe. The spatial resolution is about 1 km but captures global data on a daily basis. This is therefore a useful platform for monitoring air quality parameters in regions with intense agricultural practices. It can also help in assessing the greenhouse gas emissions within agro-systems at finer temporal resolutions.

Sentinel-6: This is concerned with monitoring of sea level changes. It is designed to have high precision in the monitoring process which is vital for climate change investigations. A number of agricultural activities have been blamed on the evolving changes in climate culminating into melting of glaciers and consequent rise in sea levels.

b. MODIS

Moderate Resolution Imaging Spectroradiometer (MODIS) is a sensor aboard Terra and Aqua satellites for monitoring land and water, respectively (NASA 2024). The temporal resolution of the sensor is 1 to 2 days and gathers images at 36 different spectral bands. The data gathered by the sensors are useful for global monitoring for changes across land water spaces. The images have a swath of about 2330 km cross track at 250 m (bands 1 and 2), 500 m (bands 3 to 7) and 1000 m (bands 8 to 36) spatial resolutions. Although the spatial resolutions are quite coarse, compared to Landsat or Sentinel, they are very useful for monitoring environmental changes at regional or global scales. These are especially useful where more frequent coverages are crucial such as changes in burnt area in case of wild fires. Generally, all the aspects of land that can be monitored by Sentinel can also be monitored by MODIS more frequently but at a coarser spatial resolution.

c. Spectral, Temporal, and Spatial Signatures

The primary satellite observable is referred to as the *spectral* signature, which is the reflectance of the Earth surface in discrete wavelength bands (blue, green, red, near-infrared (NIR), shortwave-infrared (SWIR) domains) which conveys diagnostic information about the properties of earth's surface cover (Atzberger, 2013; Weiss et al., 2019).

Mapping land cover, such as the extent of crop land, is based on differences in the remotely sensed *spectral* signature between land cover types. The *temporal* signature captures different phenological characteristics, such as distinct intra-annual temporal developments associated e.g., with plant growth cycles.

Going beyond the qualitative spectral-temporal differences between land cover types (such as cropland versus non-cropland, or different crop types, e.g.), and to a lesser extent, between crop types, some primary variables that can be physically retrieved from remotely sensed radiometric measurements include: (1) plant density, leaf area, green cover, leaf biochemical content, leaf orientation, canopy height, the fraction of absorbed photosynthetically active radiation, and albedo; (2) land surface temperature (vegetation/soil) and thus plant water in the emitted thermal domain; and (3) soil moisture and surface roughness in the microwave domain (Homolová et al., 2013; Weiss et al., 2019).

The satellite observable canopy optical reflectance (or radar backscattering) contains the compound (aggregate) effect of these leaf/canopy biophysical and biochemical properties, which influence is wavelength-dependent (their optical-reactivity, or absorption features, occur in certain wavelengths). In practical terms, then, whether e.g., certain farming practices or outcomes can be detected by satellites depends on whether such practices present soil/plant diagnostic features that belong to the above state variables. Even if that is the case, the

challenge is however to extract (isolate) from the remotely sensed signal, parts of the signal that uniquely express the variations in the state variable of interest. Spectrally, with broad waveband multispectral sensors in the current operational Earth observing satellites, measurement in a spectral channel (band) may contain absorption features of different leaf constituents (e.g., plant water, lignin, cellulose, and nitrogen-containing proteins in the shortwave infrared band) (Kokaly, 2001). Spatially, the signal recorded for a given satellite image pixel contains contributions from the different surface materials (e.g., growing crops, soil, crop residues) that are present within that measurement footprint. Indeed, in the science of remote sensing, a variety of spectral vegetation indices (Main et al., 2011; Montero et al., 2023), or/and the combination of different sensors (optical, synthetic aperture radar SAR, thermal) (Pinter, Jr. et al., 2003), serve to maximize sensitivity of the remote sensing indicator/predictor to the cropping system variable of interest, whilst minimizing sensitivity to other sources of variations that contribute to the system's remotely sensed signal.

Vegetation indices are derived from combinations of two or more bands of the electromagnetic spectrum, for instance red, blue, green, and infra-red bands. Using different combinations and calibrations, these signals can be converted to a range of different indices that can be used to measure a range of variables, such as crop yields, irrigation, forest degradation, and many more.

One of the most widely utilized vegetation indices in agriculture is the Normalized Difference Vegetation Index (NDVI). NDVI quantifies the health of green vegetation in an area by measuring the difference between the near-infrared (NIR) and red bands of the electromagnetic spectrum. It therefore serves as a reliable indicator of vegetation vigor, biomass production, and photosynthetic activity that could be difficult to detect from a single band. By analyzing NDVI values obtained from satellite or drone imagery, farmers can assess the spatial variability of crop health within their fields, identify areas of stress or nutrient deficiency, and implement targeted interventions such as irrigation, fertilization, or pest control, among others.

In addition to NDVI, the Enhanced Vegetation Index (EVI) and Soil Adjusted Vegetation Index (SAVI) are indispensable alternatives in agricultural monitoring. EVI, for instance, minimizes atmospheric influences and background noise, thereby providing more accurate assessments of vegetation dynamics in areas prone to atmospheric interference or soil reflectance variations. Similarly, SAVI adjusts for soil brightness, making it particularly useful in arid or semi-arid regions where soil background can significantly affect vegetation indices' accuracy.

Different vegetation indices have been formulated for different applications. For crop yields, e.g., Burke & Lobell, (2017) use the Normalized Difference Vegetation Index (NDVI), the green

chlorophyll vegetation index (GCVI), and the enhanced vegetation index (EVI). The formula for the NDVI is (NIR-red) / (NIR + red), for the GCVI it is (NIR / green) – 1, and for the EVI it is 2.5 * (NIR – red) / (NIR + 6 * red + 7 * blue + 1).

The Land Surface Water Index (LSWI = (NIR - SWIR) / (NIR + SWIR)), along with NDVI and land surface temperature has been used to identify irrigated fields (Massari et al., 2021). The Normalized Difference Tillage Index (NDTI = (SWIR1 - SWIR2) / (SWIR1 + SWIR2), where SWIR1 and SWIR2 are the two shortwave infrared bands common in multispectral satellite system) has been used to differentiate crop residues from bare soil, and thus detect soil tillage practices (Zheng et al., 2014). Here, thanks to the ease for satellite data access and computing these days, researchers have the flexibility to compute different satellite indicators like vegetation indices if not readily available and test the different indices in their application.

Furthermore, vegetation indices contribute to sustainable land management practices by monitoring land cover changes, deforestation, and habitat fragmentation. By tracking vegetation dynamics over time, stakeholders can assess the impact of land use policies, climate change, and natural disasters on ecosystems and biodiversity, to name just a few examples.

d. Evolving advances in remote sensing for agriculture

Improvements in artificial intelligence especially in computer vision is further advancing the limits of what we can achieve with remotely sensed data. Supervised machine learning techniques, for instance, rely on large amount of training samples to calibrate models that can be used with some level of uncertainty to predict on unseen data. While this has found important usage in still images like detecting objects in fotos, the same techniques have been used successfully with spatially explicit remotely sensed data. For instance, mapping the extent of cocoa plantations in Ghana and Ivory Coast (Kalischek et al., 2023), crop classification (Teixeira et al., 2023), oil palm in (Zortea et al., 2018; Dalagnol et al., 2022), and detection of forest degradation (Dalagnol et al., 2023). These techniques are cost effective and useful especially in areas where observation data are inadequate. For instance, the cocoa plantations mapping study by Kalischek et al (2023), showed not only the extent of cocoa plantations, but also how much the plantations have contributed to forest loss.

While these were national level investigations, the advances can be upscaled to global level investigations. For instance, they can help in detection of how much coffee or tea is coming from deforested land across the globe. Such findings fit well into enactment of such policies as the European Union Deforestation Regulation when such crops as coffee fields are mapped globally. In addition, the techniques are going beyond the need for a lot of annotated data in

building the predictive models towards reliance on foundational models to self-learn and map the features of interest. Such self-learning models have shown good results in agriculture too (Li et al., 2024; Sornapudi and Singh, 2024). The future of leveraging AI and remote sensing data holds great potential for many agricultural challenges across the globe.

S2 Typical Work Flows

Here we describe a typical end-to-end workflow in preparing geospatial data in general, and satellite data in particular, into an analysis-ready data for agriculture and resource economics research. We illustrate the workflow as technically implemented primarily in Google Earth Engine (Gorelick et al., 2017), however conceptually the workflow is generic and transferable across different computational environments (software and libraries) and framework, which we further detail in section **S3** (Required Skills and Resources) below.

The starting point of the workflow is the research question, econometric design, and in turn the data required to carry out the analysis. The data requirements, which are based on the spatial and temporal scale and coverage of the phenomenon of interest, determine the adequate spatial and temporal resolution and coverage of the satellite data or/and derived thematic data products. Additionally, the econometric design, i.e., whether the identification is based on cross-sectional regression model or panel regression model, determines whether the analysis can be done with satellite-based data products that have been created as a one-time-off demonstration, or it must rely on those that are created as a time series (e.g., annual, from 2000 to the present). For the same thematic variable, there can exist several data products, one product employing the latest improved input satellite data and algorithm but cannot be generated as a long-term time series, while another product generated as a long-term time series prioritizing long-term temporal consistency. Various satellite-based data products on agricultural variables are reviewed in section **S2**. (Variables).

When working in Earth Engine, the researchers then search for the satellite-based data or data products in the Earth Engine public data catalog (<u>https://developers.google.com/earth-engine/datasets</u>), or the Earth Engine community datasets (Roy & Swetnam, 2024). If the required data is not available in these catalogs i.e., the researchers download the data from other repositories (e.g., where authors of scientific publications have uploaded their data), or if the researchers have their own data, they can "upload" or ingest the data into their Earth Engine private asset (standard quota 250GB, number of assets 10,000). Their own data can be *raster* data i.e. gridded continuous surface (in Earth Engine the object class ee.Image and ee.ImageCollection), or *vector* data i.e. points, lines, or polygons (in Earth Engine the object class et owork with them are described in section **S3** (Required Skills and Resources). If no data exist

that satisfy the researcher's analysis requirement, or if the available existing data products are not reliable enough for the analysis, or if the researchers wish to improve upon the existing data, the researchers may also consider creating their own gridded dataset by combining several existing gridded datasets or/and by using standard machine learning methods (available in Earth Engine). In the latter, the researchers would need to collect reference ("ground truth") data for model training and validation. The reference data can be a sample of field survey locations (point coordinates, as vector data), or a reference map from which we sample the training and validation points (pixels).

The gridded data downloaded from some external repositories may need to be pre-processed prior to ingesting them into Earth Engine (for analysis combining with the plethora of gridded datasets readily available in Earth Engine, and for easily scalable computation afforded by Earth Engine supercomputer cluster). For example, the researchers have downloaded global raster data of static-climatology monthly precipitation at 1km resolution (not available in Earth Engine data catalog), and would need only the annual sum of precipitation, for the country of Ethiopia. Instead of ingesting into Earth Engine asset all the original twelve raster files with global extent, which can take up unnecessarily substantial storage space in the Earth Engine asset, the researchers can crop the rasters into the spatial extent of Ethiopia, and temporally aggregate (reduce in Earth Engine lingo) the rasters by calculating per-pixel sum of the values, in their local machine i.e. their own laptops e.g. using R software. In another example, the researchers have downloaded a table of country-level socioeconomic variables, annually, from 1961 to 2023. They then would like to add other variables from satellite based gridded data products available in Earth Engine data catalog, to each country in the country-level table. Thus, they georeference the data table by adding the country boundary geometry to the data table, converting it into geographic vector data (i.e., basically a data table with a special column for the geometry). To fill missing data, they calculate the average of the socioeconomic variable during some number of years. They also need to add continent information, and *intersect* the country boundary map with the continent boundary map. Then, they ingest the processed vector data into Earth Engine asset. The concepts and building blocks (algorithm primitives) of geocomputation (e.g., spatial join) for raster data, vector data, and raster-vector interaction are described in section Required Skills and Resources.

Back in Earth Engine, after they identify the satellite data or data products that satisfy their analysis requirements, they import them, and apply spatial-temporal query i.e., filter by location (rectangular bounding box of the spatial region of interest), by time period, and by metadata (e.g., estimated cloud cover percent in the whole image/scene). (As will be described below, common Earth Engine workflows consist of three fundamental operations namely *filter, map, reduce*). Satellite data and data products often come with per-pixel quality assurance/quality

control (QA/QC) indicators, and data users are strongly advised to read carefully the provided documentations about, at minimum, how to interpret the QA/QC indicator (and the unit of the measurement, whether some scaling factors need to be applied), and, often necessary for deeper understanding, how the data product is made, how the product has been assessed for quality and what is the expected accuracy (algorithm theoretical basis document, user guide, validation report, known issues, scientific journal publications). The per-pixel QA/QC information is usually stored in the form of bitmask (bitwise flags), that is the binary digits (bits) in the binary representation of QA/QC value (decimal) stores information about some quality attributes of the data e.g. cloud, cloud's shadow, radiometric saturation, data fill (whether backup algorithm is used) and so on. For example, Landsat collection 2 level 2 (surface reflectance) product comes with QA_PIXEL bitmask layer (band), where bit 0 indicates whether the pixel status is data fill, bit 1 dilated cloud, bit 2 cirrus cloud, bit 3 cloud, and bit 4 cloud shadow (USGS, n.d.). If however the standard QA/QC band is not sufficiently accurate in the researcher's geographical area of interest, the researchers might employ a custom QA/QC band. For example, it is known that the cloud and shadow QA/QC bands provided with the official Sentinel-2 level 2A (surface reflectance) product is not as reliable as the Landsat counterpart (due to, in parts, the absence of thermal bands in Sentinel-2), and the remote sensing community is actively developing a better cloud and cloud's shadow masks (Pasquarella, 2023), which the researchers should take advantage of. Satellite based thematic data products that are produced using machine learning also increasingly come with QA/QC layer that informs the prediction uncertainty (see section Required Skills and Resources).

Once the selected satellite-based data and data products are filtered for high-quality observations, the researchers may need to apply some pre-processing to the data, such as computing a derived variable such as NDVI, calculating summary statistics for a given time window, or applying a further mask to the gridded data to keep only the observations (pixels) valid for the analysis. In the first preprocessing use case, the researchers create a custom function to perform some calculation combining different bands (raster algebra, or band math operation, e.g. calculating NDVI from the red and near-infrared bands) in an image. The custom function is then applied to each of the images (e.g. Landsat scenes, where scenes correspond to different spatial tiles or image acquisition dates), which in Earth Engine is known as a map operation i.e. we map the function over the image collection. In the second preprocessing use case, the researchers may want to compute, for each pixel, the annual (or other time windows) average air temperature or sum of precipitation. This is a data aggregation (a reducer operation in Earth Engine lingo) across time dimension (also known as "compositing"). Finally, in the third preprocessing use case, the researchers may want to focus their analysis only on cropland, and therefore use a land cover map to mask the image to keep only cropland pixels (non-cropland pixels are set as no-data).

Ultimately, the analysis-ready data for the economics/econometrics research design is typically in the form of a table, where the row corresponds to the observational unit that can be a grid cell, or a country, or a field survey area, and so on. To obtain this analysis-ready data table, the researchers extract the data from the preprocessed satellite-based data and data products as source, for each of the observation units (each grid cell, or each country, or each field plot) as the destination/target. In GIS parlance, this raster-to-vector data extraction is usually termed "zonal statistics", and in Earth Engine accomplished by the function/method reduceRegion or reduceRegions, which takes as input the image (containing the stack of layers representing the data variables of interest) and the feature collection of the spatial observation units, and return as output a table where the row corresponds to the individual spatial observation unit and the columns correspond to the data variables (layers/bands in the image). This output table can be exported to Google Drive, and analyzed in R or Stata for the econometrics analysis. A special case is when the observation unit is the grid cell, for which it is more efficient to export the data out of Earth Engine as raster data (as an image) and perform the necessary conversion of the raster to a table (a row is a grid cell, optionally storing the point coordinate of the cell's centroid) in local machine. In general, working with raster (i.e., in the image space) is more efficient, particularly when processing data in Earth Engine.

When the workflow involves machine learning applications, the researchers need reference data ("ground truth" data) to train the machine learning model, and to assess the accuracy of the model's prediction. In some cases, for example for mapping land cover types, deforestation, or forest restoration, the reference data can be collected by visual interpretation of high-resolution (10-30m) public Earth observation satellite imagery, spectral time series profile, and very high resolution (<1m) basemaps such as Google Maps/Earth satellite basemap or other basemaps from other providers (Pasquarella et al., 2016). One feature of Earth Engine namely the Earth Engine Apps and user interface API can aid in this process of reference data collection, as well as in sharing the prediction maps as web-based interactive maps to non-GIS specialists, and solicit feedback. While Earth Engine offers functionalities to train and apply some standard classical machine learning algorithms, for complex machine learning workflow that requires rigorous feature selection, model hyperparameter tuning, and advanced spatial cross validation, the researchers can export the training data with the covariates extracted in Earth Engine, to Google Drive, to then run the machine learning experimentations using machine learning libraries in R or Python. The optimal model hyperparameter and features set can then be applied in Earth Engine at inference time to predict for the whole study area wall-to-wall.

Given the series of data preprocessing steps typically required to ensure high-quality observations (pixels), it is recommended to visually check the resulting data layers at each

step to avoid errors. For example, in the workflow that necessitates masking out non-cropland areas from the image, the researchers display the masked image to verify indeed only cropland areas remain as valid-data pixels. The Earth Engine web-based interactive development environment IDE (*code editor*), which uses the Earth Engine JavaScript client library (API), is particularly designed to facilitate this highly interactive and iterative workflow. In addition to displaying computed objects (e.g, image) on the map panel in the code editor, researchers can interactively inspect the pixel's values, and print the objects in the code editor console to verify the object contains the elements as expected as a result of the computation (e.g., images that satisfy the date filter, or that the output image/table contains the bands/columns requested).

Here, it is worth paying attention to the limits of this *interactive (synchronous, online) computation environment*, versus the *batch (asynchronous) computation environment* in Earth Engine. The *interactive* environment allows for computation that takes maximum five minutes of processing time, and tens of MB of memory use. This means, the researchers need to create a small testing (or debugging) use case from their end-to-end processing chain that can fit that interactive environment jobs limit. Once the researchers are confident that the workflow gives the correct outputs, they can scale the workflow and execute it over the entire dataset, by exporting the outputs. This will use the *batch* environment, in which a job (task) has maximum lifetime effectively four days on average, allowing up to 3,000 tasks submitted per user (but each individual user is limited to a small number of concurrently running tasks). When the batch jobs need to be programmed as a large number of individual tasks, the Earth Engine Python client library (API) is more convenient to use (export jobs automatically executed) and the researchers may need to translate their Earth Engine JavaScript API code into the Earth Engine Python API code (which is relatively easy as they have very similar coding syntax).

Even in the batch computing environment, depending on the size of the computation jobs, the researchers may need to try out different parallelism designs in their Earth Engine code until the batch export can succeed, to fit the scaling limits of the Earth Engine infrastructure. The scaling limits are in place considering that while as a whole Earth Engine can manage extremely large computations (e.g., processing global 30m Landsat data to produce the Global Forest Watch tree cover loss maps (Hansen et al., 2013)), fundamentally it is composed of clusters of low-end servers on one hand, and is a shared resource among the Earth Engine users. In effect, when scaling their workflow and sending the batch tasks to the Earth Engine server, researchers may encounter various scaling-related errors (not an exhaustive list; here we note the typical remedy) such as "computation timed-out", "user memory limit exceeded" (increase tileScale parameter in the function, to accommodate the limit of 2GB stack depth per tile), "too many input pixels per output pixel" (typically encountered when performing reduceResolution, which can be solved by running the reduceResolution calls in two steps i.e.,

with an intermediate resolution), "too many aggregations" (break down computations into separate export tasks/jobs), "too many concurrent aggregations" (change tileScale), "computed object (or value) is too large" (caused by the 100MB cache limit, typically encountered when exporting a table output e.g., training points with the predictor variables; may be solved by exporting data for a subset of the points at a time), and the generic "internal server error" (Google Earth Engine, n.d.-a, n.d.-b). In practice, the researchers need to think carefully about the size of their computation (which depends on the spatial extent, spatial resolution, time span, and number of data layers i.e. bands), and design how to divide the data extraction jobs into separate batch tasks, whether by data variable (e.g., separate tasks for precipitation, for air temperature, and so on), by region (e.g, separate tasks per country), or by year. Finally, two generic solutions that can be attempted are, first, to tile the study area, and create the export task per tile, and second, to store intermediate computation results. For additional data storage needs, and for specialized workflow such as deep learning, researchers can consider Earth Engine integration with services provided by Google Cloud Platform.

S3 Variable Construction

Here we provide an overview of both pure satellite-based data products and hybrid data products and how to construct different variables. A remark on the general advantages and disadvantages of the two types of data products is useful in advance. Satellite data offers the only means for continuous monitoring of agricultural areas in a spatially distributed manner. Satellite-based data products on agricultural variables have the advantages of having high spatial and temporal resolutions (in theory, up to the nominal resolution of the raw satellite imagery, e.g., at 10-30m, every few days with Sentinel-2/1 and Landsat (Claverie et al., 2018); however, the usable data frequency depends on clear-sky observations availability). However, in practice, for cropland variables, temporally, the satellite measurements need to be integrated across time to mine sufficient diagnostic information about cropland (e.g., in an area cropland and grassland looks similar at one time, but looks different at different times; total crop biomass estimated by accumulating the signal throughout the growing season). This, in addition to the issue of un-usable cloudy observations, and inherently noisy individual-date observation (temporal aggregations serve to reduce noise level), means high-confidence satellite-based cropland data products are typically made, operationally, at annual update cadence, or, at best seasonally. This data update cadence is of course better than traditional agricultural statistics (typically every 5 to 10 years). Thematically, however, not all cropland variables are physically detectable from the satellite observable. As will be elaborated later in this section, for instance differentiating a large number of crop types is often not possible using only satellite data, if the crop types considered are not sufficiently separable from their spectral,

temporal, and spatial (textual) signature, along with other environmental characteristics, as observed by the satellites. In this situation, hybrid data products provide higher thematic accuracy, by integrating data from satellites and thematically-rich data from official statistics (e.g., per-crop production quantity). Almost all the hybrid data products are essentially based on FAOSTAT national-level crop production data (FAO, 2024) carried out under the FAO World Program for the Census of Agriculture (WCA), with data records available from 1961, supplemented by sub-national official statistics sources. These crop production data typically include harvested area (in hectares), production quantity (in tonnes), and yield (in tonnes per hectare), for each of the FAO 175 standardized crops. The hybrid data products then use the so-called spatial allocation, or spatial downscaling, procedure to disaggregate/distribute the administrative unit aggregated statistics to the cropland grid cells, based on some criteria and various gridded covariates, resulting in gridded representation of the crop production data. There exist now several hybrid gridded crop data products. Notable examples are the IFPRI Spatial Production Allocation Model, SPAM (Yu et al., 2020), the FAO IIASA Global Agro-Ecological Zones (Fischer et al., 2021), M3Crops (Monfreda et al., 2008), and MIRCA2000 (Portmann et al., 2010). These hybrid data products are similar in that they apply spatial downscaling to the crop production statistics, into 5' (five minutes) resolution or about 9.2km at the equator. They differ, fundamentally, in terms of the underlying aggregated crop production statistics they use, i.e., the sub-national political units (or, Statistical Reporting Units, SRU, which data compilation represents the major effort), the spatial downscaling (allocation) procedure (SPAM applies a cross-entropy model to distribute the political unit level agricultural statistics into cropland grid cells based on a satellite-based cropland mask, and other ancillary covariates namely irrigated areas, population density, potential crop suitability and yields (based on GAEZ), and market access (You et al., 2014; Yu et al., 2020); the other data products distribute the political unit level statistics uniformly in all grid cells belonging to the respective political unit or weighted by the grid-cell cropland area shares), whether distinctions between production or cropping systems (irrigated vs rainfed, low-input vs high-input vs substinence farming systems in SPAM, single-cropping vs multiple cropping in GAEZ) are made, whether distinctions between seasons are made (e.g., MIRCA2000 applies temporal downscaling to estimate monthly growing areas), and the number of crop classes (42 crop classes in SPAM, 23 crop classes in GAEZ, 175 crop classes in M3Crops, 26 crop classes in MIRCA2000).

While hybrid data products provide more crop types and more accurate crop yield data than what can reliably be predicted using satellite data alone, given the way they are made and the input data sources, they possess inherent limitations. First, spatially, the accuracy of the spatial distribution is uncertain, while consistency between the political units is also uncertain because census methodologies and dissemination vary. Secondly, temporally, as agricultural statistics

data collection is typically conducted every 5 to 10 years due to the prohibitive cost, the hybrid crop data products are not updated frequently and tend to be outdated (i.e., ca. 2000 or/and 2010); but recent efforts updating to ca. 2015 e.g., an update to GAEZ by (Grogan et al., 2022), and to ca. 2020 e.g., the SPAM 2020 beta release (International Food Policy Research Institute (IFPRI), 2024), and CropGRIDS (Tang et al., 2024) as an update to M3Crops (Monfreda et al., 2008). Furthermore, despite the occasional updates, for researchers interested in changes through time, it is not advisable to infer changes between the hybrid data products generated for different years, due to the lack of temporal consistency and continuity of the inputs to the spatial allocation model.

The accuracy of the hybrid cropland data products is dependent on the accuracy of the underlying agricultural censuses and surveys, which are known to frequently contain errors from uneven data collection, misreporting, or incorrect tabulation. Therefore, all data sources have errors or imperfection, combining different data sources may reduce the errors, and what is important is to account for the errors when interpreting analysis results and deriving robust conclusions (e.g., robustness checks with different data sources). Yet another option that is now accessible to researchers is to generate their own custom geospatial datasets using easily accessible satellite data and machine learning computation, especially to take advantage of the study-specific training data and their domain-specific expert knowledge in the design of the machine learning implementation (include predictive features that are known to be correlated with the prediction outcome variable according to established theory i.e., in favor of explainable model decision making, and avoiding the use of features that may correlate with treatment or/and control variables that may invalidate causal inference when applicable).

a. Agricultural Structures and Landscape Compositions

Agricultural structures refer to the spatial configurations of the agricultural landscape which are shaped by various drivers such as climate change and farming practices. As such, the fundamental information to characterize agricultural structures is the spatially explicitly cropland physical/planting areas (acreage), or the cropland *mask* (a grid cell is labelled as either cropland, or non-cropland). Identifying cropland areas, versus other land cover types, using satellite data is now a routine task with generally high expected accuracy (albeit some challenges remain in smallholder agriculture landscape (Kerner et al., 2024)). That said, the above mentioned challenges in mapping cropland areas in smallholder agricultural landscape is also due to these physical limitations of the remotely sensed radiometric signals: in these areas the cropped field is small relative to satellite image nominal spatial resolution, as well as is interspersed with other vegetation in an intercropping practice, and thus presenting a problem of mixed-land-cover spectral signature making it more difficult to distinguish cropland

and non-cropland. This can be solved by higher spatial resolution imagery (i.e., sub-meter imagery from commercial providers), greater amount of training data to teach the machine learning algorithm in this region, and advances in the machine deep learning algorithms themselves that can better encode, and generalize (across geographies, times, downstream tasks, resolutions, sensors, and image acquisition conditions; "foundational models"), the unique *spatial/contextual textural* features and temporal signatures of cropland (Burke et al., 2021). The textural features are especially useful in mapping perennial crops like fruit trees, vineyards, and plantations, that do not exhibit strong annual growth cycles (phenology) like annual crops.

Owing to the decades of research by the remote sensing community, both in academia and industry, and their partnership, as well as advances in satellite data spatial, spectral, temporal resolutions, and advances in machine learning (classification or semantic segmentation task; note this is to be differentiated from crop field delineation which is an *instance segmentation* task, that requires sub-meter imagery and hence has yet to be demonstrated at global scale) and high-performance computing, we have now several global high-resolution cropland land cover maps at our disposal. Notable global cropland mask data products (also readily available in Earth Engine data catalog) that have become recently available are (note here we noted the reported accuracy from product validation by the data providers but the accuracy numbers, from the validation sample used, are representative of the global areas; reported accuracies among data products are not directly comparable due to different map validation methodological details; we noted the data updates, as of this writing): (1) the cropland-specific University of Maryland (UMD) Global Land Analysis and Discovery (GLAD) cropland extent (30m, available for 2003, 2007, 2011, 2015, and 2019 three-year epochs) (Potapov et al., 2021); the new-generation 10m generic land cover maps (with cropland as one of the land cover types) based on Sentinel-2/1 namely (2) the European Space Agency (ESA) WorldCover, 10m, 2020 and 2021, annual, with reported overall accuracy (see Olofsson et al. (2014) for details on accuracy assessment terminology) 74.4%, and for cropland class users' accuracy UA (1 – commission error) of 81.1% and producers' accuracy PA (1 – omission error) of 76.7% (Buchhorn et al., 2020), (3) the Esri LULC (Land Use and Land Cover), 10m, 2017-2022, annual, with reported overall accuracy 85% (Karra et al., 2021), (4) the Google Dynamic World, 10m, 2015-2024, classified every Sentinel-2 satellite image scene (acquisition date), with accuracy 88.9% reported for the cropland class (Brown et al., 2022), (5) the ESA WorldCereal, 10m, 2020 and 2021, annual temporary cropland extent and seasonal maps for cereals, with reported UA and PA of 88.5% and 92.1% for temporary cropland extent (Van Tricht et al., 2023), (7) MODIS (sensor, on board of the Aqua/Terra satellite platforms) MCD12Q1 land cover, 500m, 2001-2022, annual, with reported overall accuracy 73.6% (Sulla-Menashe et al., 2019), (8) the GLC_FCS30D land cover map, 30m, annual during 2000-2022,

roughly five-yearly during 1985-2000, with reported PA and UA of >85% for cropland (X. Zhang et al., 2024), and **(9)** the Global Food-and-Water Security-support Analysis Data (GFSAD) Global Cropland-Extent Product at 30-m Resolution (GCEP30), ca. 2015, in which the cropland definition includes plantations, and with reported producer's accuracy 83.4 percent, and user's accuracy 78.3 percent, for the cropland class (Thenkabail et al., 2021).

Note that while we here specifically refer to cropland, the generic land cover datasets mentioned above also map land cover class(es) pertaining to grassland or pasture. However, general-purpose land cover maps do not differentiate grassland management systems i.e. between pastures/rangelands and unmanaged lands (Parente et al., 2024). Ramankutty et al. (2008) generated the pasture extent (and cropland extent) map with 5 arcmin resolution for ca. 2000 by harmonizing satellite-derived grassland land cover and ground-based national and sub-national agricultural inventory on pasture (and cropland). HYDE annual (1961-2019, annual after 2000, decadal before 1990), 5 arcmin, dataset differentiates between more intensively managed pasture, and extensively managed rangeland, based on population density and aridity index (Klein Goldewijk et al., 2017). An attempt to create 30m annual (2000-2022) cultivated grassland versus natural/semi-natural grassland is demonstrated in the upcoming Global Pasture Watch dataset, which illustrates the particular challenge in harmonizing the definitions of grassland as a very complex system in previous mapping efforts.

In choosing which cropland mask (extent) datasets best suits a specific study context, researchers would then need to consider the appropriate spatial resolution, temporal resolution and span, and cropland definition. Concerning definition, as land cover is a semantic representation of the land, data providers use not precisely the same cropland definition. For example, the GLAD map includes herbaceous cropland but not woody plantations, and includes fallow land, whereas the WorldCereal map includes temporary cropland (active cropland in a year) and thus excludes fallow land. Some satellite-derived cropland extent data products include perennial crops, while others do not (Potapov et al., 2023). Care must also be taken when reconciling satellite-derived cropland extent against agricultural censuses and surveys (e.g., FAO categories of "arable land", "permanent crops", "temporary crops", "temporary fallow", "temporary meadow and pastures") (Tubiello et al., 2023), while noting that national and sub-national reporting agencies also may use different operational definitions despite the FAO standardization framework. With regards the spatial and temporal resolution (and span), naturally the optimal choice depends on the expected detectability (spatial footprint on the ground, and length of time since intervention, also comparing pre- and postintervention) of the interventions under study. Temporally, the MODIS 500m annual map, and the very recently introduced GLC_FCS30D 30m annual map, provides an annual panel data spanning over two decades, and is thus suitable for panel studies looking at year-to-year

variations (necessary for tracing treatment effects after some time has passed, or for when the treatments occurred at different times), or long-difference research designs. Particularly for these latter types of applications, researchers need to be aware of the expected year-to-year accuracies of the cropland masks, i.e., the temporal stability of the classification. The MODIS and GLC_FCS30D accounted for the year-to-year stability in the design of the algorithm, and provided the multiyear accuracy assessment, and thus higher quality assurance that the temporal trend reflects the actual changes in the status of the lands, and not primarily year-toyear misclassification (which can be expected from satellite data e.g., due to varying amount and intra-annual distribution of clear-sky optical satellite observations, as well as imperfect satellite image pre-processing e.g., atmospheric correction). Researchers needing data from before the year 2000 (high-resolution satellites data availability is sparse before 2000), can consider hybrid data products such as the History Database of the Global Environment (HYDE) which provides annual cropland from 10000 before Common Era (BCE) to 2015 Common Era (CE) at 5' (five minutes, or ~9.2km at the equator) resolution (Klein Goldewijk et al., 2017), based on satellite-based land cover maps (ESA-CCI, 300m resolution, 1992-2020, annual (ESA, 2017)), official statistics, and other geospatial datasets.

It is important to note that, while satellites can provide the cropland *physical or planting* area, researchers working with crop production accounting often need cropland *harvested* area instead, and here hybrid data products that harmonize satellite-based cropland *physical* area and crop production official statistics remain the most reliable data for this. In practice, for example in M3Crops dataset, the national and sub-national statistics of harvested area is spatially distributed/allocated to cropland grid cells based on criteria such as the potential cropping intensities (Monfreda et al., 2008). The potential cropping intensities are in turn modelled based on criteria such as climate conditions and irrigation such as the FAO and IIASA's Global Agroecological Zone Assessment suite of gridded crop data (Fischer et al., 2021). Here, the recently available satellite-based 30m cropping intensities data (CI, equals the ratio between harvested area and physical area) during 2016-2018 (M. Zhang et al., 2021) holds promise to improve the estimation of cropland harvested area.

While cropland mask is an operational capability using satellite data up to global scale, distinguishing crop *types* on the other hand is yet to achieve similar application readiness. Fundamentally, this is because the large number of different crop types may not have clearly separable spectral or/and temporal signature that the satellite multispectral sensors (the Operational Land Imager sensor on board Landsat-8 satellite has eleven spectral bands, the MultiSpectral Instrument (MSI) sensor on board Sentinel-2 satellite has thirteen bands with the novel red edge bands) can resolve. In the machine learning classification task, this issue is known as low interclass variability and high intraclass variability, and means the success of

the crop type classification algorithm depends on the availability of a large amount of local training data as the relationship between spectral-temporal signature and crop types are highly site-specific and not easily transferable geographically. The implication is while some considerably successful local-to-regional demonstration of crop type mapping exist (some large-scale examples: continental scale in Europe (d'Andrimont et al., 2021), national scale in Germany (Griffiths et al., 2019), and the United States Department of Agriculture (USDA) Cropland Data Layer (CDL) (USDA National Agricultural Statistics Service, n.d.)), no global complex-set crop type maps currently exist and will likely remain so in the foreseeable future. Becker-Reshef et al. (2023) compiled and harmonized 24 national and regional crop type datasets derived from satellite imagery from 21 sources covering 66 countries to develop a Global Best Available Crop Specific Masks (GEOGLAM-BACS) at 5km resolution for four main crops. That said, if such a sizable global training data of complex-set crop types, and satellite hyperspectral data become available, coupled with advances in deep learning algorithm, the capability of generating a global crop type map remains to be seen. This, however, likely would require substantial ground knowledge particularly pertaining to the expected crop types that are known to grow in certain areas, and the crop calendar for the large number of crops which is to date not available (but to note the global crop calendar for wheat and maize by Franch et al. (2022) within the ESA WorldCereal framework). Previous global crop calendar datasets based on national-level observations and sub-national observations in some countries were outdated, namely MIRCA2000 (ca. 2000, 26 crops, distinction between rainfed and irrigated crops), and SAGE (1990s or early 2000s, 19 crops) (Sacks et al., 2010). Here, the hybrid data products are advantageous as agricultural statistics disaggregate among a large number of crops, albeit with limitations as noted above (uncertain accuracy of the spatial distribution, relatively coarse spatial resolution of ~9.2km, and being outdated and not frequently updated).

In terms of total crop production, compared to harvested areas, crop *yield* (ratio of the total mass of harvested product to the total harvested area used to grow the crop) is a much more dynamic agricultural variable as it can substantially fluctuate year by year due to weather condition. Over the last half century, it is the growth in crop yields, rather growth in harvested areas, that has contributed predominantly to increased production (Blomqvist et al., 2020). Improving crop yields, or closing the yield gaps, on existing cropland (i.e., sustainable intensification, as opposed to extensification by cropland expansion) is a main lever to achieve global food security. Spatially explicit, at high spatial resolution, crop yield data is therefore a critical information need. Absolute crop yield estimation using satellite data depend on availability of site-specific and crop-specific training data, with successful local use cases but no operational large-scale up to global scale crop yield data product (Atzberger, 2013; Rembold et al., 2013; Weiss et al., 2019). To note however, depending on the studies, absolute crop yield may be required, or relative changes in crop yield as proxied by the satellite

indicators may suffice. Indeed, currently, for operational large-scale (regional to global) crop yield monitoring, such systems evaluate trends in crop yield in a qualitative manner by for example computing the anomalies on remotely sensed vegetation indices, which are then investigated by expert analysts to infer the likely impacts on crop yield (Fritz et al., 2019; Nakalembe et al., 2021).

To circumvent the need of local yield calibration data, some studies make use of point-based crop growth models into which the spatially continuous satellite derived variables is assimilated, with some successful large-scale applications e.g., in the U.S., and some case study regions in India and Zambia (Azzari et al., 2017; Lobell et al., 2015). In Azzari et al. (2017), the obtained accuracies as assessed against county-level statistics in the U.S., distinct level in India, and province level in Zambia, were R2 consistently above 0.5 and RMSE 0.85 ton per ha (9% of average yields) in the U.S., R2 above 0.45 and RMSE 0.56 ton per ha (19% of average yields) in India, and R2 0.55 and RMSE 0.4 ton per ha (26% of average yields) in Zambia. In Burke & Lobell (2017) study in Kenya, the agreement (R2) with field plot level yield (full plot crop cutting, which had the highest agreement with the satellite yield as compared to sub-plot cutting and self-report) is up to 0.4. In interpreting the agreement (r = 0.04) with the objective crop cutting-based measures (Lobell et al., 2020).

The successful use of crop growth models however requires accurate crop growth models, which needs to be calibrated (with many agronomy model parameters) for each crop and potentially for each region. Therefore, such use of crop growth models to estimate crop yield using satellite data is similarly mostly in research rather than in operational applications up to global scale (Weiss et al., 2019). Global-scale analysis of crop yield, especially for a large number of crop classes, thus rely on hybrid crop data products such the above mentioned IFPRI SPAM (which is now updated to 2020), with the caveat that the spatial allocation of the input political-level aggregated yield into the grid cells is uncertain (compared to the per-pixel direct yield estimation using satellite data), and the rather coarse resolution (~9.2km). For an analysis requiring a time series of crop yield, two data products offer a unique long-term, annual, historical perspective. Ray et al. (2019) compiled annual crop yield and harvested area from 1974 to 2013 for the top ten global crops at ~20,000 political units globally (pure census based i.e., not based on satellite data). The global dataset of historical yields for major crops (GDHY) by lizumi & Sakai (2020) provides annual (and seasonal) yield from 1981 to 2016 for four major crops at 0.5 degrees (~55 km) resolution globally. GDHY is a hybrid product based on census data and satellite based FPAR, LAI and solar radiation, as well as other hybrid crop data products M3Crops. Finally, it is worth noting that the ground data collection of crop yields

used in censuses is often based on farmer self-report which can be unreliable; imperfect satellite-based crop yield can be just as reliable (Burke & Lobell, 2017; Lobell et al., 2020; Paliwal & Jain, 2020). In this sense, satellite-derived yield and yield from agricultural statistics can be used synergistically, for example to cross-check each other.

Finally, the delineation of farm boundaries is regularly important. While this used to be done manually by navigating around farms, advances in remote sensing and machine learning techniques have made this possible digitally (Bilodeau et al., 2024). Edge detection techniques are making it possible to rapidly determine boundaries of many cropland farms for optimized precision agriculture. This is normally supported by ground metadata such as information about farm owners and others. This can also be helpful in detecting changes in the field boundaries that could imply encroachment into forested or protected lands and could help in enforcing regulations such as European Union Deforestation Regulation that hinders trade on products coming from deforested lands.

b. Cropping System, Farming Practices and Agricultural Technology

Cropping system refers to the planting arrangement in time and space on a piece of land, and associated farming practices and agricultural technology. Relevant agricultural variables under this category include variables describing the crop succession (temporal component > 1 year), the cropping pattern (crop sequence and spatial arrangement within 1 year), and the crop management techniques such as applications of fertilizer, irrigation, pesticide, machinery, soil tillage, and harvest practices (Bégué et al., 2018). Overall, compared to the variables describing agricultural structure, relatively few satellite remote sensing studies focused on cropping system. Most studies were conducted at local scale, highly dependent on local reference (training) data, and requires significant knowledge on local agricultural practices. As a result, there is presently a lack of operational large-scale (up to global-scale), systematically updated, geospatial data products are available under these themes. Therefore, in practice, researchers may need to create their own gridded map for the cropping system variables they are interested in, using machine learning. As mentioned before, creating such maps is now relatively low-effort thanks to cloud-based data access and computation platform such as the Google Earth Engine (Gorelick et al., 2017), although availability of high-quality local training data is the pre-requisite, and the accuracy of the generated map needs to be evaluated carefully to ensure that the map is fit-for-purpose. Below we provide an overview of cropping system variables that have been successfully estimated in the remote sensing community and the methods employed.

Crop *successions* describe the multiyear farming practices in terms of crop rotation, fallowing, and shifting cultivation (including abandonment). Linking the observable farming practices to

the farmers' socioeconomic condition and decision making, crop rotation, as opposed to monocropping, may indicate the presence of low-input versus high-input farms (crop rotation is necessary in low-input farms to maintain soil fertility). The temporal granularity and continuity (e.g., Landsat missions data record going back as far as 1980s, with next Landsat missions already planned) of satellite data allows to track crop successions, such as demonstrated for Europe (Estel et al., 2015) and several study areas in the world (Yin et al., 2020). While crop fallows can be inferred from annual maps of cropland mask (extent) that are readily available, crop rotation on the other hand requires annual crop type maps, which, as previously noted, remains challenging, and thus is yet to be available as an operational product. This is due to the large spatial and temporal heterogeneity in the spectral-temporal signature of the same crop (which is in turn due to heterogeneity in management practices) or/and inadequate spectral-temporal separability between different crop types with the current generation satellites (which may be improved by recently launched hyperspectral satellites such as EnMAP (Chabrillat et al., 2022)). Studies that do attempt to produce annual crop type maps, to infer crop successions, therefore should carefully account for the relatively high chance of crop type misclassification errors in each map, in turn leading to erroneously identified changes year to year. In practice, this can be done by applying logical crop type yearly transition rule based on expert knowledge.

Cropping pattern denote practices such as single cropping, multiple cropping, where multiple cropping can take the form of sequential cropping (multiple harvests during one growing season), or intercropping (e.g., agroforestry). Sequential cropping can be identified by examining the cropland phenology using high temporal resolution intra-annual time series data from. MODIS sensor data, despite the moderate resolution of 250m at most, has powered this type of application owing to its daily revisit (compared to 8-day revisit of two simultaneously operational 30m Landsat satellites, 5-day revisit of two 10-20m Sentinel-2 i.e. 2A and 2B satellites; but promising new capabilities are expected from the 2-3 days temporal resolution with the 30m harmonized Landsat and Sentinel-2 data), which allows for more clear-sky observations and thus to better reconstruct the phenological cycle (Arvor et al., 2011; J. C. Brown et al., 2007; Gumma et al., 2016). Sentinel-1 radar is not affected by clouds, and thus provides a denser time series, however the individual Sentinel-1 observation is inherently noisy due to speckle noise as well as its sensitivity to soil moisture. The individual optical satellite observation also has noise due to imperfect preprocessing (atmospheric, geometric, terrain correction). In practice, applications relying on dense time series need to apply smoothing to the time series data, or to fit a statistical time series model (Atzberger & Eilers, 2011; Kong et al., 2019). Once the temporal profiles are obtained, interpreting e.g., classifying the temporal profiles into the types of sequential cropping system however require field knowledge of the local sequential cropping system and crop calendar, hence making large-scale (up to globalscale) mapping of sequential cropping system challenging. While sequential cropping can be resolved by the high temporal resolution of the currently operational satellites, *intercropping* (crop diversification) typically can only be resolved with imagery with at least 1m spatial resolution (commercial satellites such as Ikonos, Geoeye, Quickbird, Worldview, and Pleiades), and is thus only cost-effective for small areas (however, there are some successful demonstrations with PlanetScope 3m and RapidEye 5m imagery for large farmland trees, see Brandt et al., 2024). Interpretation of intercropping in the very high spatial resolution imagery typically is best performed by manual visual interpretation by a local expert on satellite image visual interpretation and on local ground conditions.

Crop management techniques include the processes from soil preparation until crop harvesting and post-harvest practices, including application of agrochemical inputs such as fertilizer, pesticide, and herbicide, as well as application of irrigation. Concerning the specific macronutrients input from fertilizer, namely nitrogen (N), phosphorus (P), and potassium (K), physically speaking, only N has a detectable absorption feature, via chlorophyll and proteins that contain N. P and K on the other hand do not exhibit detectable absorption features in their spectra, and thus fundamentally cannot be estimated from satellite data, at least directly. Further complicating the detection, macronutrients in leaves do not necessarily respond to treatments of the corresponding macronutrients in the soil; crops in areas treated with the highest amount of a certain nutrient does not necessarily have highest foliar content of that nutrient, and have instead the highest foliar content of another nutrient (Siedliska et al., 2021). This, however, also means that, application of agrochemical inputs can be indirectly deduced from its covariation with other leaf constituents that are spectrally detectable, especially chlorophyll that some studies have shown to respond to changes in also P and K inputs (Lu et al., 2020; Siedliska et al., 2021). Indeed, a successful N estimation has been demonstrated in many local, crop-specific studies, via N covariation with chlorophyll (Berger et al., 2020; Homolová et al., 2013). The successful applications however largely make use of field hyperspectral measurements, and rarely use satellite data (Berger et al., 2020) as hyperspectral data is necessary to uniquely isolate the narrow absorption features of the leaf constituent, and historically there were no operational (for global mapping) satellite hyperspectral missions. The operational capabilities of the newly launched hyperspectral such as EnMAP (Chabrillat et al., 2022; Storch et al., 2023), PRISMA (Cogliati et al., 2021) and HISUI (Matsunaga et al., 2020), along with future planned missions, remain to be seen. Due to the above challenges, an operational large-scale monitoring of farming agrochemical inputs is yet to be realized, as it requires knowing local applications of various farming inputs that together affect the crop conditions (feasible for small study areas within which some farming inputs are the same or are known). We note however a first global-scale attempt of using current operational (multispectral) satellites (Moreno-Martínez et al., 2018) to create a global

map of specific leaf area (SLA), leaf dry matter content (LDMC), leaf nitrogen content per dry mass (LNC), and leaf phosphorus content per dry mass (LPC) at 500m resolution, for a typical year during 2000-2010. They employed a machine learning prediction model to leverage the TRY global leaf traits database, and obtained accuracy of RMSE \leq 20% and Pearson's correlation R = 0.78. We recommend however to interpret the estimated leaf traits, especially in the context of local variability in agricultural areas, by accounting for other information such as crop types from hybrid crop data products overviewed before.

Two global-scale hybrid data products (thus relatively coarse resolution i.e., ~9.2km) exist for agrochemical inputs application rate. The first is the crop-specific (138 crops, also pasture) nitrogen (N), phosphorus (P), and potassium (K) application rates for ca. 2000 (Mueller et al., 2012), created based on national and sub-national fertilizer statistics, and the hybrid (satellite and census statistical fusion) cropland extent (Monfreda et al., 2008). The product applied income-based extrapolation to fill data gaps, and as with most hybrid datasets, apply assumptions/constraints to scale/harmonize between reported agricultural statistics. More recently, gridded data (~9.2km) on different nitrogen components at annual time step (1860-2019) became available (Tian et al., 2022). This dataset is based on the HYDE gridded historical cropland extent, which in turn used as input the satellite-based ESA CCI generalpurpose land cover maps (ESA, 2017). A hybrid gridded crop-specific pesticide application rates at 5 arcmin resolution (~9.2km) for 2015 is available (Maggi et al., 2019), which performed a global spatial inference i.e., extrapolating adequately-detailed data in the US to global-scale, using a set of gridded soil, hydroclimatic, agricultural, and socio-economic variables (some of which are derived from satellite data), conditioned on pesticide regulations and FAOSTAT country-level pesticide database.

On the farming input of irrigation, satellite data have proven capable to distinguish between irrigated cropland and rainfed cropland. This capability is again owing to the high temporal resolution of satellite data, as the key to identifying irrigation is the time window of observation: the greenness (NDVI), or the surface moisture (Land Surface Water Index LSWI = (NIR - SWIR) / (NIR + SWIR)), and the land surface temperature, differ between irrigated and rainfed fields at the time of the peak of growing season (Deines et al., 2019; Massari et al., 2021). In notoriously cloudy tropical regions, SAR data which are not affected by cloud cover is commonly used. A global map of rainfed versus irrigated cropland at 30m resolution for ca. 2015 is available (Teluguntla et al., 2023), with reported global producer's and users' accuracy above 80% for both irrigated and rainfed class. Beyond qualitative identification of irrigated vs rainfed cropland areas, estimating the amount/volumes of applied irrigation is however yet to be demonstrated (but, different irrigation techniques may be detectable based on in-season timing of field's maximum greenness). For this continuous measurement of irrigation, a hybrid

data product taking advantage of satellite data and agricultural statistics on the area equipped for irrigation (AEI) has been produced by Siebert et al. (2015). AEI represents the area of land that is equipped with infrastructure to provide water to crops, excluding rainwater harvesting (note may not be the same as area actually irrigated (AAI), but AEI and AAI typically share the same temporal trends). They created gridded historical AEI at ~9.2km resolution, for the period 1900 to 2005 (10-year steps until 1980 and 5-year steps from 1980 onward), based on countrylevel FAOSTAT, and sub-national level Aquastat and Eurostat database of AEI as well as other sub-national data sources. The hybrid (satellite and census) cropland and pasture extent maps from the HYDE and Earthstat were used for data gaps infilling or consistency constraints.

Soil *tillage* (related to the use of agricultural machinery) can be detected by satellites as tillage operations modify two land surface features, namely the amount of crop residue (plant litter) cover, and the surface roughness (Zheng et al., 2014). The physical basis behind crop residues remote mapping is that crop residues can be spectrally distinguished from bare soil in the SWIR electromagnetic spectrum due to known absorption features from cellulose and lignin in crop residues, such as using the Normalized Difference Tillage Index (NDTI=(SWIR1-SWIR2)/(SWIR1+SWIR2), where SWIR1 and SWIR2 are the two shortwave infrared bands common in multispectral satellite system). Surface roughness can be detected by Sentinel-1 C-band radar with wavelength 5.6cm and hence sensitive to cm-level variations in soil height. The empirical relationship between the satellite indicators and the tillage features, like other cropping system variables described in this section, is however not universal (under the influence of spatial heterogeneity of other management practices, soil characteristic, and weather) and thus depend on local calibration. Pure satellite-based global mapping of tillage has not been demonstrated, but large-scale mapping has been demonstrated in the U.S. (Azzari et al., 2019). In the absence of reference data for model calibration, expert knowledge rule-based determination of tillage types was conducted by Porwollik et al. (2019) based on gridded crop type, soil depth, field size (Fritz et al., 2015), erosion, aridity, and country-level conservation agriculture area and income.

c. Environmental and Geographical Variables

Majority of agricultural activities are dependent on environmental variables such as temperature, precipitation, humidity, and solar radiation. Agricultural yield is pegged on the genome of the item under production, environment and management practices employed by the farmers. The environmental variables are normally obtained via weather instruments sited at specific ground stations. However, efforts have been made to capture high resolution data based on remote sensing capabilities via UAVs, airborne or satellite as well as proximal sensing techniques (Ojha et al., 2015). At the same time, developments around data storage, computation and efficient algorithms have also grown tremendously (Weiss et al. 2019).

Indeed, the use of these remotely sensed techniques has boosted the profitability of agricultural ventures in the past few decades (Wolfert et al., 2017).

Ability to rapidly observe the environmental variables over time and across space is making it possible for farmers to explain the performance of most farm produce. It also has the advantages including gathering data from rather inaccessible locations or under extreme weather conditions. Temperature, for example, is gathered using satellite images to help farmers decide when to conduct specific agricultural practices. For example, the preceding number of high temperature days can help a farmer decide when to sow. They also include soil variables such as soil moisture, soil temperature, soil organic carbon content, and soil pH. On the other hand, geographical variables as captured by satellites refer to the physical features and spatial characteristics of land or earth surface. They include topography such as elevation, slope, and aspect, field boundaries, drainage patterns, and road networks among others. There are also land cover variables such as cropland, urban, forest, bare soil, water bodies among others. Other geographical variables include proximity variables such as distance to water bodies, urban centers, distance to roads, distance to borders and other infrastructure that affect agricultural and environmental outcomes. Here, we give an overview of some of these variables.

In many analyses, elevation play a role. This refers to the height of the land surface or a point above the sea level and e.g., can affect the climate of a place as well as drainage and erosion (Amatulli et al., 2018) that are important in agricultural activities. It is normally associated with temperature with high altitudes having lower temperature values than higher elevation points (Kazemi Garajeh et al, 2023). This is captured by satellite sensors that have active sensors to receive backscatter from the ground to record height of a place with respect to time it took to send and receive back the signal. Other derivatives of elevation such as slope and aspect are also of interest in agriculture. Slope (the steepness or inclination of land) is an important factor influencing surface runoff and soil erosion. It is thus an important consideration in agriculture including consideration for mechanization and types of crops to grow at specific slope angles. Aspect (the direction that a slope faces) also plays an important role in regulating microclimate of a place and how much exposure a place is to sunlight. Both of these play a role in determining the type of crop to grow at a given location as some crops need more solar exposure than others. All these can be obtained from satellite remote sensing information to help farmers make the right decisions on the kind of agricultural activities to venture into.

The hydrology and drainage patterns of are often important in agriculture as water plays a central role in farming. Precise mapping of the natural as well as man-made drainage systems including rivers and irrigation channels is crucial for managing water resources in a farm (Bilodeau et al., 2024). This is also helpful in preventing or controlling flooding as well as planning irrigation systems. Both optical and radar imagery from satellite sensors help identify

and monitor the drainage patterns and have proved valuable in detecting hydromorphic soils for conservation in agricultural fields of Brazil (Mello et al., 2023). For instance, long term monitoring of the behavior of a river in case of flooding could help a farmer design dykes that can withstand the worst possible flooding scenario based on the historical observed flooding in the region.

Finally, almost always need researchers work with land cover. This is often right at the start, e.g., one might need to exclude all land cover types that are irrelevant to the research, e.g., all non-crop land, or all non-forest land. Alternatively, sometimes one simply needs to account for different land-covers, e.g., a vegetation index reflects a different phenomenon on cropland and on grassland, or maybe land cover changes are the actual research focus.

Satellite remote sensing and machine learning techniques allow the identification of different land cover types such as forests, grasslands, and croplands (Song et al. 2016).

S4 Required Skills and Resources

Geospatial data expertise involves conceptual and practical understanding of geospatial data models/structures, and algorithms (i.e. what established operations can be performed on geospatial data). Geospatial data are georeferenced i.e., are placed in the real world to relate the spatial elements (entities) in the data with the surface in the real world. That is, they have a coordinate reference system (CRS). The CRS can be geographic CRS ("longitude-latitude", typically in decimal degrees) or projected CRS ("easting-northing" in meters from a local coordinate system origin/reference point). Keeping track of CRS is essential as workflows very often involve data with different CRS, and thus often require reprojecting data into different CRS. Some vector data (data models described below) geometric operations are better performed to the data in projected CRS. On the data models, geospatial data represent the world as raster data and vector data. Raster data represents the world as a continuous grid of cells, i.e., as an image, fields, surfaces, dividing the surface up into cells of constant size. Raster data is therefore suitable for spatially smoothly continuous phenomena such as elevation and rainfall. Raster is basically an array of numbers, which can be a two-dimensional grid (i.e., single-layer, or a matrix) where the rows and columns indices represent real world coordinates such as longitude and latitude and the value of each array element represents the data (e.g., elevation value at that coordinate), or three-dimensional (multi-layer) with additional time dimension (time series of raster e.g. monthly average rainfall for every 1km grid cell within a study area), or four-dimensional (a generic "data (hyper)cube") with an additional dimension representing different data variables or "bands" (e.g., different spectral bands in satellite image), and so on. Raster data can store only one attribute in one layer (the pixel value). In terms of data types, raster data can be discrete (categorical) values for example to represent soil types, or can be continuous for example to represent elevation. Vector data on the other

hand represent the world as discrete objects with well-defined borders, namely as points (e.g., cities), lines (e.g., rivers) and polygons (e.g., country boundaries or field measurement plots). Geographic vector data are composed by the geographic representation, and the attributes or properties associated with each geographic entity. In other words, vector data can be understood as basically a table, with a special column to store the "geometry". In contrast to raster (single-layer) data, vector data can store multiple attributes as columns in the data (e.g., countries vector data with attributes country-level population and GDP). Raster data can be converted to vector data (*vectorization*), and vice versa vector data can be converted to raster data (*rasterization*). For example, a researcher may decide each grid cell is their observational unit of analysis, and therefore needs to convert the raster into a table, where each cell becomes a row in the table (with the cell id), optionally storing the coordinate of each cell's centroid (raster to points conversion).

There are numerous file formats for raster data and vector data, with the most commonly used ones are GeoTIFF (.tif, basically a TIFF image format but with special header storing the geospatial metadata information) and netCDF (.nc) for raster data, whereas ESRI-shapefile (.shp) and GeoPackage (.gpkg, which is preferred over shapefile) for vector data.

The above conceptual data models are implemented using different designs in different software and libraries that support geospatial data, including Geospatial Data Abstraction Library (GDAL) which is the backbone of many GIS software especially in terms of read/write support for the myriad of raster data formats (Gandhi, n.d.; GDAL/OGR contributors, 2024), R (R Core Team, 2023), Earth Engine (Gorelick et al., 2017), and Python (Van Rossum & Drake, 2009). Understanding the specific data structure, and the algorithms (methods), in the specific implementation is key to efficient handling of geospatial data in the chosen software and library. Use of built-in data models, object classes, and methods is recommended over custom self implementations to avoid potential mistakes due to mis-aligned dimensions or geolocations of the data. In R, raster data are supported by the SpatRaster (and SpatRasterDataset, SpatRasterCollections) object class in terra library (Hijmans, 2023b), which is the successor the raster library (with RasterLayer, RasterStack, and RasterBrick object class; Cohen, n.d.; Hijmans, 2023a), and as spatiotemporal raster data cube in the stars library which is designed for easier handling of complex multidimensional rasters i.e., raster data (hyper)cubes with many layers (e.g., bands), for many moments in time (e.g., months), and many attributes (e.g., sensor type A and sensor type B). For data wrangling, subsetting, and manipulation, the stars library allows to apply dplyr verbs (Wickham, François, et al., 2023). Terra is in most cases more efficient (computational and memory wise). NetCDF file typically has rich metadata which in R can be read with ncdf4 library (Pierce, 2023). In the Python world (which, notably, has better support for deep learning, and parallel processing use cases), raster data are supported by the DataArray and Dataset object class in the xarray package (Hoyer et al., 2024; Hoyer & Hamman, 2017). Raster file formats can be read into xarray objects using rioxarray library (Snow et al., 2024) which is based on rasterio library (Gillies & others, 2013). In Earth Engine, the corresponding object class for raster data are ee.Image and ee.ImageCollection. Vector data models, in R, are implemented as SpatVector object class in terra, and simple features in sf library (Pebesma, 2018; Pebesma & Bivand, 2023; which is successor to sp library). sf has better support for vector data and allows to apply dplyr verbs for data wrangling. In Python, vector data can be imported as GeoDataFrame object class in GeoPandas package (Bossche et al., 2024; Jordahl et al., 2020) which is based on shapely library (Gillies et al., 2024) to encode the geometry. Finally, in Earth Engine, ee.Feature and ee.FeatureCollection are used to represent vector data. An analysis often needs to employ methods available in different libraries, which is facilitated by existing functions to convert between object classes used in different libraries, especially among the geospatial libraries in the R ecosystem.

Geospatial algorithms or geocomputation that are typically applied to raster data include: (i) data aggregation ("reducer" operation in Earth Engine terminology) in space dimension (local i.e., per-cell, zonal i.e., per neighborhood window or other zones, or global operations i.e., perraster e.g., creating raster of distance to rivers) or/and time dimension (e.g. per-pixel annual mean); (ii) modifying spatial extent or resolution of raster(s) with operations such as mosaic, merge, crop, mask, and resample; and (iii) reprojecting raster(s) to another CRS. Reprojecting a raster involves reprojecting the grid and resampling of the per-pixel values to populate the new grid, and thus changing the values. Care must therefore be given in specifying the appropriate resampling algorithm, such as nearest neighbor for categorical variable, and bilinear interpolation (i.e., based on distance-weighted average of four closest pixels) for continuous values, depending on whether preserving original values or having a smooth spatial variation is more important (as well as whether resolution changes are from coarser to finer resolution, or from finer to coarser resolution). To take advantage of computationally efficient matrix algebra operations (array broadcasting), raster layers that share the same spatial resolution and extent, can be stacked into a single multi-layer raster object in most software and libraries. Identifying functions which apply an operation across pixels in a single raster, versus those that apply an operation across multiple rasters on a per-pixel basis, is key. Further, many of the raster processing methods/functionalities then automatically apply a given per-raster operation to all the raster layers in the raster object (i.e., for loop is not necessary). On to vector data, typical algorithms/geocomputation for them include: (i) geometric queries, or spatial join/relationship, such as adjacency, distance computation, finding nearest unit, intersection, containment, and other spatial predicates; (ii) buffering; and (iii) simplifying geometry and fixing invalid geometry. Spatial join between point (destination/target) and polygon (source) vector data for example entails the process of finding for each of the point in

the target data, the polygon in the source data that the point intersects, and then assign the value of the attributes associated with the polygon, to the point. An example use case for this is to assign country information to points observations representing the centroids of ground measurement plots (which country each point falls in).

More broadly, spatial join can be considered a generic operation that can be applied between a vector data and another vector data (e.g., point-on-polygon overlay described above), between a raster data and another raster data (e.g., re-gridding a raster data into another raster data template that may involve reprojecting, resampling, changing extent, and aggregating pixels if resolutions differ, in the source raster), but also between a raster data and a vector data (Lovelace et al., 2019). The last type of operation involving raster-vector interaction is particularly important in preparing analysis-ready data for econometric analysis. For example, researchers may need to reduce the data they work with by cropping some raster data with polygon vector data delineating their study area. This raster-polygon cropping applies spatial subsetting to the raster data to keep grid cells that fall within the geographic extent (rectangular bounding box) of the geometry of the vector polygon data. Grid cells that are not inside the exact geometry of the polygon can also masked (i.e., set to no-data). Another important raster-vector operation is raster data extraction to vector spatial elements ("extract" function in terra, "reduceRegion" and "reduceRegions" methods in Earth Engine). A practical example use case is when the researchers need to add information about environmental variables obtained/derived from satellite data, to their analysis unit which can be survey areal sample unit. Raster-points data extraction entails the process in which for each of the points, it finds which raster cell the point is located within, and assigns the value of the cell to the point. The assigned value can be the value of the single cell the point falls in, or spatial interpolation of several nearby cells (weighted by distance). Raster-polygons data extraction identifies for each of the polygons, all the raster cells that intersect with the polygon, and assigns a vector of the cell values to the polygon. The vector of the cell values can also be summarized during the process i.e., we compute the mean of all the cells within each polygon. exact extractr (Daniel Baston, 2023) is a specialized, highly performant library in R for raster data extraction with polygons. It is worth noting that spatial join using rasters is usually faster than using only vectors. For example, if one wishes to add country-level population (polygon vector data) to each cell in a 1km raster, rather than converting the raster into points (centroids) and then perform points-on-polygons spatial join, it can be more efficient to rasterize the country polygons data into the 1km raster grid as template, and then stack them at this stage of processing chain. Once the econometric analysis-ready data is created, for a large data table, the data.table (Dowle & Srinivasan, 2023) library in R can be a faster alternative to regular dplyr data frame (dtplyr library by Wickham, Girlich, et al., (2023) allows to write dplyr code for data.table object).

To take the benefits from high resolution satellite data and data products, and to investigate research questions at a large scale, the research team requires expertise in big data processing. The technical challenges in analyzing big data are in processing data which are beyond the available computer's memory ("out-of-memory" computation), and in parallelizing the computation across computer cores ("out-of-core" computation) to generate output in feasible processing time. Fortunately, tools that facilitate relatively easy set up for out-ofmemory and out-of-core geocomputation are increasingly available. To deal with memory limitation, the tools adopt a lazy evaluation framework, in which data is only loaded into memory as needed. That is, when a large file is read, only the necessary metadata such as the dimensions (the "container") of the raster data, or the columns names and data type of the vector data, is read while the actual content of the data is not loaded into memory as yet. Applying a series of geoprocessing functions to these "proxy" objects merely builds a computational graph description, which is not automatically executed, but rather executed on demand via calling certain functions. This represents the change in paradigm from "file-based" processing, to "metadata-based" processing, and different providers of Earth Observation data are converging to use a standardized data catalog and metadata specification (SpatioTemporal Asset Catalogs, STAC) to allow a unified interface for data access and query. Lazy evaluation is performed by the terra library in R, by the xarray package in Python (when chunks argument is specified it uses lazy Dask (Rocklin, 2015) array instead of eager numpy array) for raster data. For vector (tabular) data, this is for example supported by the arrow (Richardson et al., 2023) library in R, and by Dask DataFrame (Dunnington, 2022) and daskgeopandas packages in Python. Likewise, Earth Engine also builds computational graph description before explicit execution calls (e.g., printing objects in the console, displaying objects in the map panel, or exporting computed objects). The modern geospatial tool stack also offers an easy interface to request that the geoprocessing, once being asked to be executed, is to be parallelized (i.e., the work distributed) across compute cores. This is achieved for example by setting up a Dask cluster in Python, and in R by reading a subset of raster data at a time (Krzysztof, 2023) and utilizing standard framework for parallel processing namely foreach (Weston, 2022) and doParallel (Weston, 2022). At times, however, instead of setting up multi-core processing, faster processing time can be achieved by using algorithms natively written in C++, such as using GDAL tools (which supports also "virtual" raster object), accessible within R e.g. via gdal utils function in the sf library. Also, enhanced file formats more suitable for big data e.g., cloud-optimized GeoTIFF (COG) for raster data, and GeoParquet for vector data, can improve performance as they are specifically designed to support streaming parts of the data content (partial read, parallel read). For raster data, careful data compression (e.g., lossless LZW or DEFLATE algorithm) and data type specification (e.g., integers vs floats) that reduces data size on disk can improve computational performance. For

large vector data, use of database e.g., PostGIS (Committee & others, 2018) and duckdb (Mühleisen & Raasveldt, 2024) with the spatial extension (Gabrielsson, 2023) can be considered.

When computing resources in the researcher's machine are not adequate, they need to scale their workflow to a cluster of computers. This can be an in-house high-performance computing facility, or cloud-hosted virtual machine instances that can be rented from the Amazon Web Services, Google Cloud Platform, Microsoft Azure, and others. Cloud-based scaling is advantageous as that means there is no need to move (download) the data (different cloud providers readily host satellite data in their cloud storage), thus facilitating a cloud-native workflow that brings the algorithm to the data, instead of bringing data to the algorithm. Naturally, if researchers need to use such cloud services, some knowledge of cloud computing and devops would be necessary. Google Earth Engine alleviates the burden of substantial IT management and cloud expertise and effort of which researchers often do not have the luxury of having. When using Earth Engine, users are liberated from the details of working in a parallel processing environment. This however means that users cannot influence the parallelism, and they need to have a good understanding of the parallelism paradigm in Earth Engine to harness the best performance out of the available, shared, resources. Here, it is helpful to have an awareness of the fundamental difference in Earth Engine geocomputation models, and scaling limits, as compared to traditional (desktop-based) GIS and remote sensing workflow. In the traditional geospatial processing paradigm, typically one needs to explicitly, spatially align/harmonize raster data into the same, common, grid (extent, resolution, projection system), throughout the processing chain. These resampling and reprojection are expensive operations. In contrast, when working in Earth Engine, users only need to specify the common output grid at the end of the processing chain, and Earth Engine will internally perform the necessary resampling and reprojection on the fly. In other words, the projection, resolution and spatial extent in which computations take place, are determined on a "pull" basis (i.e., determined by the arguments scale/crsTransform, crs, region in export calls, or interactively by the map display viewport visible tiles and zoom level), and input data are requested only for the respective tile and resolution (Earth Engine internally pre-computes pyramids of reducedresolution images). Note however users can also explicitly control the resampling and reprojection operation (with reduceResolution, reproject, or/and resample) when desired, for example when the users wish to use bilinear resampling instead of the default nearest neighbor, or when they wish to upsample a finer resolution image to a coarser resolution image by computing specifically the sum (Earth Engine's default pyramiding policy is average for continuous-valued image) of finer resolution pixels in each of the coarser resolution pixel. Some notes on scaling workflows to fit Earth Engine infrastructure and parallelism are provided in section **S1** (Workflows).

Another required skill when working with geospatial data is geodata visualization in both static and interactive modes. Frequent visual inspection is recommended to make sense of the data, catch potential errors/artifacts, throughout processing chain i.e., to examine whether the output from an important processing step (e.g., masking an image so undesired pixels are correctly assigned no-data and ignored from subsequent computations, or reprojecting data or specifying its CRS so they spatially align, etc.) is not out of expectation. The Earth Engine webbased IDE/code editor (JavaScript client library/API) is especially designed for this purpose with its powerful map panel component. Other geodata visualization tools include QuantumGIS QGIS (QGIS Development Team, 2024); ggplot (with geom_sf and geom_raster) (Wickham, 2016), mapview (Appelhans et al., 2023), and tmap (Tennekes, 2018) in R; and geemap (Wu, 2020) and HoloViz (Yang et al., 2022) in Python.

Finally, equally important as the technical skills enumerated above is domain knowledge expertise concerning the satellite-derived data (measurements), and the intended application (inference), under the specific theme of the research. When the researchers are thinking to use original satellite data to derive some measures (e.g., create a map of crop type using machine learning), it is beneficial to have appropriate expectation on the potentials and limitations of the available satellite data (in terms of spatial, temporal, spectral resolution i.e. whether it can sufficiently detect the outcome of the treatment in relation to the situations on the ground) how the satellite observable physically responds to the target variable of interest (i.e., the data generating process). To best leverage the many decades of remote sensing science, researchers need to be able to identify the most reliable satellite data processing levels (e.g., whether atmospheric, view angle corrections have been performed for optical satellite data; whether speckle, terrain corrections have been applied to SAR data), and apply preprocessing using the QA/QC layers provided in the product appropriately (e.g., filtering out clouds and shadows). When the researchers intend to use satellite-based thematic data products, such as a map of discrete-valued land cover types, or continuous-valued vegetation biomass, they need to have some understanding of the algorithm used to create the product (by consulting the Algorithm Theoretical Basis Document or User Guide documents), and be informed about the data product accuracy (in validation report or relevant scientific publications, though assessing accuracy in local study areas is additionally recommended). By understanding the nature of the errors in the satellite data/data products (which can stem from the physical sensitivity of the satellite indicators used to predict the variable of interest, or from the thematic processing e.g., quality, quantity, and spatial distribution of training data used in machine learning based products), the researchers can assess the impact of the errors downstream on the econometrics analysis. For example, remote sensing predictions may be differentially accurate for positive and negative outcome in a binary classification task (e.g., deforestation is over-predicted, i.e. has higher commission error but lower omission error,

whereas non-deforestation is under-detected, i.e. has higher omission error but lower commission error), or for different class labels (categories) in a multiclass classification task (e.g., grass and shrub are more often incorrectly classified compared to forest, water and builtup), or for different parts of a distribution for continuous outcomes (e.g., the tendency of MODIS Vegetation Continuous Field and Landsat Percent Tree Cover product to underestimate high percent tree cover values and overestimate low percent tree cover values (Gerard et al., 2017; Staver & Hansen, 2015), or of the radar-based aboveground biomass map to underestimate high biomass values (Sinha et al., 2015)). Fortunately, increasingly data products developers provide uncertainty maps along with their prediction maps. Researchers however need to carefully interpret the uncertainty maps as they are based on various techniques such as quantile regression, number of valid observations (input pixels), distance to second class, ensemble-based prediction variance, or area of model applicability (extrapolation of feature space) (Meyer & Pebesma, 2021; Mo et al., 2023; Singh et al., 2024). For hybrid data that combines satellite based data and agricultural censuses or surveys, in addition to understanding the errors in the satellite based maps used in the product (how they are used and thus how and the extent of which the errors propagate to the end product), the researchers need to have the skills account for uncertainties due to the underlying statistical data, interpolation/extrapolation model fit, and data gaps infilling criteria. Awareness of the various gridded data layers used in the processing chain of hybrid products also informs whether circularity (outcome and treatment data share the same predictor) can be a major issue. Ultimately, the researchers need to evaluate the fit-for-purposeness of the satellite data and data products within the specific context of the study and the inference objectives. For example, the researchers need to consider whether the absolute accuracy of the satellite based predicted outcomes permits the research questions to be answered with sufficient confidence, or the relative accuracy is acceptable i.e., can the satellite-based predicted outcomes sufficiently detect the responses to the treatments of interest (Burke & Lobell, 2017; Lobell et al., 2020).

References

Appelhans, T., Detsch, F., Reudenbach, C., & Woellauer, S. (2023). mapview: Interactive Viewing of Spatial Data in R [Computer software]. https://CRAN.R-project.org/package=mapview

Arvor, D., Jonathan, M., Meirelles, M. S. P., Dubreuil, V., & Durieux, L. (2011). Classification of MODIS EVI time series for crop mapping in the state of Mato Grosso, Brazil. International Journal of Remote Sensing, 32(22), 7847–7871. https://doi.org/10.1080/01431161.2010.531783

Atzberger, C. (2013). Advances in Remote Sensing of Agriculture: Context Description, Existing Operational Monitoring Systems and Major Information Needs. Remote Sensing, 5(2), 949–981. <u>https://doi.org/10.3390/rs5020949</u>

Atzberger, C., & Eilers, P. H. C. (2011). A time series for monitoring vegetation activity and phenology at 10-daily time steps covering large parts of South America. International Journal of Digital Earth, 4(5), 365–386. https://doi.org/10.1080/17538947.2010.505664

Azzari, G., Grassini, P., Edreira, J. I. R., Conley, S., Mourtzinis, S., & Lobell, D. B. (2019). Satellite mapping of tillage practices in the North Central US region from 2005 to 2016. Remote Sensing of Environment, 221, 417–429. https://doi.org/10.1016/j.rse.2018.11.010

Azzari, G., Jain, M., & Lobell, D. B. (2017). Towards fine resolution global maps of crop yields: Testing multiple methods and satellites in three countries. Remote Sensing of Environment, 202, 129–141. https://doi.org/10.1016/j.rse.2017.04.014

Becker-Reshef, I., Barker, B., Whitcraft, A., Oliva, P., Mobley, K., Justice, C., & Sahajpal, R. (2023). Crop Type Maps for Operational Global Agricultural Monitoring. Scientific Data, 10(1), 172. https://doi.org/10.1038/s41597-023-02047-9

Bégué, A., Arvor, D., Bellon, B., Betbeder, J., De Abelleyra, D., P. D. Ferraz, R., Lebourgeois,V., Lelong, C., Simões, M., & R. Verón, S. (2018). Remote Sensing and Cropping Practices:A Review. Remote Sensing, 10(2), 99. https://doi.org/10.3390/rs10010099

Berger, K., Verrelst, J., Féret, J.-B., Wang, Z., Wocher, M., Strathmann, M., Danner, M., Mauser, W., & Hank, T. (2020). Crop nitrogen monitoring: Recent progress and principal developments in the context of imaging spectroscopy missions. Remote Sensing of Environment, 242, 111758. https://doi.org/10.1016/j.rse.2020.111758

Blomqvist, L., Yates, L., & Brook, B. W. (2020). Drivers of increasing global crop production: A decomposition analysis. Environmental Research Letters, 15(9), 0940b6. https://doi.org/10.1088/1748-9326/ab9e9c

Bossche, J. V. den, Jordahl, K., Fleischmann, M., Richards, M., McBride, J., Wasserman, J.,
Badaracco, A. G., Snow, A. D., Ward, B., Tratner, J., Gerard, J., Perry, M., Farmer, C., Hjelle,
G. A., Taves, M., Hoeven, E. ter, Cochran, M., Bell, R., rraymondgh, ... Kaushik. (2024).
geopandas/geopandas: V0.14.4 (v0.14.4) [Computer software]. Zenodo.
https://doi.org/10.5281/zenodo.11080352

Brandt, M., Gominski, D., Reiner, F., Kariryaa, A., Guthula, V. B., Ciais, P., Tong, X., Zhang, W., Govindarajulu, D., Ortiz-Gonzalo, D., & Fensholt, R. (2024). Severe decline in large

farmland trees in India over the past decade. Nature Sustainability, 1–9. https://doi.org/10.1038/s41893-024-01356-0

Brown, C. F., Brumby, S. P., Guzder-Williams, B., Birch, T., Hyde, S. B., Mazzariello, J., Czerwinski, W., Pasquarella, V. J., Haertel, R., Ilyushchenko, S., Schwehr, K., Weisse, M., Stolle, F., Hanson, C., Guinan, O., Moore, R., & Tait, A. M. (2022). Dynamic World, Near real-time global 10 m land use land cover mapping. Scientific Data, 9(1), 251. https://doi.org/10.1038/s41597-022-01307-4

Brown, J. C., Jepson, W. E., Kastens, J. H., Wardlow, B. D., Lomas, J. M., & Price, K. P. (2007). Multitemporal, Moderate-Spatial-Resolution Remote Sensing of Modern Agricultural Production and Land Modification in the Brazilian Amazon. GIScience & Remote Sensing, 44(2), 117–148. https://doi.org/10.2747/1548-1603.44.2.117

Buchhorn, M., Lesiv, M., Tsendbazar, N.-E., Herold, M., Bertels, L., & Smets, B. (2020). Copernicus Global Land Cover Layers—Collection 2. Remote Sensing, 12(6), Article 6. https://doi.org/10.3390/rs12061044

Burke, M., & Lobell, D. B. (2017). Satellite-based assessment of yield variation and its determinants in smallholder African systems. Proceedings of the National Academy of Sciences, 114(9), 2189–2194. https://doi.org/10.1073/pnas.1616919114

Burke, M., Driscoll, A., Lobell, D. B., & Ermon, S. (2021). Using satellite imagery to understand and promote sustainable development. Science, 371(6535), eabe8628. https://doi.org/10.1126/science.abe8628

Chabrillat, S., Guanter, L., Kaufmann, H., Foerster, S., Beamish, A., Brosinsky, A., Wulf, H., Asadzadeh, S., Bochow, M., Bohn, N., Boesche, N., Bracher, A., Brell, M., Buddenbaum, H., Cerra, D., Fischer, S., Hank, T., Heiden, U., Heim, B., ... Segl, K. (2022). EnMAP Science Plan. In EnMAP Technical Report; (p. 87 pages, 2 MB) [Pdf]. GFZ Data Services. https://doi.org/10.48440/ENMAP.2022.001

Claverie, M., Ju, J., Masek, J. G., Dungan, J. L., Vermote, E. F., Roger, J.-C., Skakun, S. V., & Justice, C. (2018). The Harmonized Landsat and Sentinel-2 surface reflectance data set. Remote Sensing of Environment, 219, 145–161. https://doi.org/10.1016/j.rse.2018.09.002

Cogliati, S., Sarti, F., Chiarantini, L., Cosi, M., Lorusso, R., Lopinto, E., Miglietta, F., Genesio, L., Guanter, L., Damm, A., Pérez-López, S., Scheffler, D., Tagliabue, G., Panigada, C., Rascher, U., Dowling, T. P. F., Giardino, C., & Colombo, R. (2021). The PRISMA imaging spectroscopy mission: Overview and first performance analysis. Remote Sensing of Environment, 262, 112499. https://doi.org/10.1016/j.rse.2021.112499

Cohen, J. (2022). Shifting from Raster to Terra | OHI. https://oceanhealthindex.org/news/raster_to_terra/

Committee, P. P. S. & others. (2018). PostGIS, spatial and geographic objects for postgreSQL [Computer software]. https://postgis.net

d'Andrimont, R., Verhegghen, A., Lemoine, G., Kempeneers, P., Meroni, M., & Van Der Velde, M. (2021). From parcel to continental scale – A first European crop type map based on Sentinel-1 and LUCAS Copernicus in-situ observations. Remote Sensing of Environment, 266, 112708. <u>https://doi.org/10.1016/j.rse.2021.112708</u>

Dalagnol, R., Wagner, F. H., Galvão, L. S., Braga, D., Osborn, F., da Conceição Bispo, P., ... & Saatchi, S. (2023). Mapping tropical forest degradation with deep learning and Planet NICFI data. Remote Sensing of Environment, 298, 113798. https://doi.org/10.1016/j.rse.2023.113798

Dalagnol, R., Wagner, F. H., Emilio, T., Streher, A. S., Galvão, L. S., Ometto, J. P., & Aragao, L. E. (2022). Canopy palm cover across the Brazilian Amazon forests mapped with airborne LiDAR data and deep learning. Remote Sensing in Ecology and Conservation, 8(5), 601-614. https://doi.org/10.1002/rse2.264

Daniel Baston. (2023). exactextractr: Fast Extraction from Raster Datasets using Polygons [Computer software]. https://CRAN.R-project.org/package=exactextractr

Deines, J. M., Kendall, A. D., Crowley, M. A., Rapp, J., Cardille, J. A., & Hyndman, D. W. (2019). Mapping three decades of annual irrigation across the US High Plains Aquifer using Landsat and Google Earth Engine. Remote Sensing of Environment, 233, 111400. https://doi.org/10.1016/j.rse.2019.111400

Dijk, M. V., Wood-Sichra, U., Ru, Y., Guo, Z., & You, L. (2023). mapspamc: An R package to create crop distribution maps for country studies using a downscaling approach. https://doi.org/10.21203/rs.3.rs-2497136/v1

Dowle, M., & Srinivasan, A. (2023). data.table: Extension of `data.frame` [Computer software]. https://CRAN.R-project.org/package=data.table

Dunnington, D. (2022). Profiling point-in-polygon joins in R | Dewey Dunnington. https://dewey.dunnington.ca/post/2022/profiling-point-in-polygon-joins-in-r/

ESA (2017). Land Cover CCI Product User Guide Version 2. Tech. Rep. maps.elie.ucl.ac.be/CCI/viewer/download/ESACCI-LC-Ph2-PUGv2_2.0.pdf

European Space Agency (ESA) (2024). "Sentinel-2 Mission." Accessed online on 12 June 2024. https://www.esa.int/Applications/Observing_the_Earth/Copernicus/Sentinel-2.

Estel, S., Kuemmerle, T., Alcántara, C., Levers, C., Prishchepov, A., & Hostert, P. (2015). Mapping farmland abandonment and recultivation across Europe using MODIS NDVI time series. Remote Sensing of Environment, 163, 312–325. https://doi.org/10.1016/j.rse.2015.03.028

FAO (2024). FAOSTAT Statistical Database [dataset]. https://www.fao.org/faostat/en/#data

Fischer, G., Nachtergaele, F., van Velthuizen, H., Chiozza, F., Franceschini, G., Henry, M., Muchoney, D., & Tramberend, S. (2021). Global Agro-Ecological Zones v4—Model Documentation. IIASA/FAO. https://www.fao.org/geospatial/resources/detail/en/c/1410694/

Franch, B., Cintas, J., Becker-Reshef, I., Sanchez-Torres, M. J., Roger, J., Skakun, S., Sobrino, J. A., Van Tricht, K., Degerickx, J., Gilliams, S., Koetz, B., Szantoi, Z., & Whitcraft, A. (2022). Global crop calendars of maize and wheat in the framework of the WorldCereal project. GIScience & Remote Sensing, 59(1), 885–913. https://doi.org/10.1080/15481603.2022.2079273

Friedl, M. A., Woodcock, C. E., Olofsson, P., Zhu, Z., Loveland, T., Stanimirova, R., Arevalo,
P., Bullock, E., Hu, K.-T., Zhang, Y., Turlej, K., Tarrio, K., McAvoy, K., Gorelick, N., Wang, J.
A., Barber, C. P., & Souza, C. J. (2022). Medium Spatial Resolution Mapping of Global Land
Cover and Land Cover Change Across Multiple Decades From Landsat. Frontiers in Remote
Sensing, 3. https://doi.org/10.3389/frsen.2022.894571

Fritz, S., See, L., Bayas, J. C. L., Waldner, F., Jacques, D., Becker-Reshef, I., Whitcraft, A., Baruth, B., Bonifacio, R., Crutchfield, J., Rembold, F., Rojas, O., Schucknecht, A., Van der Velde, M., Verdin, J., Wu, B., Yan, N., You, L., Gilliams, S., ... McCallum, I. (2019). A comparison of global agricultural monitoring systems and current gaps. Agricultural Systems, 168, 258–272. https://doi.org/10.1016/j.agsy.2018.05.010

Fritz, S., See, L., McCallum, I., You, L., Bun, A., Moltchanova, E., Duerauer, M., Albrecht, F., Schill, C., Perger, C., Havlik, P., Mosnier, A., Thornton, P., Wood-Sichra, U., Herrero, M., Becker-Reshef, I., Justice, C., Hansen, M., Gong, P., ... Obersteiner, M. (2015). Mapping global cropland and field size. Global Change Biology, 21(5), 1980–1992. https://doi.org/10.1111/gcb.12838

Gabrielsson, M. (2023, April 28). PostGEESE? Introducing The DuckDB Spatial Extension. DuckDB. https://duckdb.org/2023/04/28/spatial.html

Gandhi, U. (n.d.). Mastering GDAL Tools (Full Course Material). Retrieved June 12, 2024, from https://courses.spatialthoughts.com/gdal-tools.html

GDAL/OGR contributors. (2024). GDAL/OGR Geospatial Data Abstraction software Library [Computer software]. Open Source Geospatial Foundation. https://doi.org/10.5281/zenodo.5884351

Gerard, F., Hooftman, D., van Langevelde, F., Veenendaal, E., White, S. M., & Lloyd, J. (2017). MODIS VCF should not be used to detect discontinuities in tree cover due to binning bias. A comment on Hanan et al. (2014) and Staver and Hansen (2015). Global Ecology and Biogeography, 26(7), 854–859. https://doi.org/10.1111/geb.12592

Gillies, S. & others. (2013). Rasterio: Geospatial raster I/O for Python programmers [Computer software]. Mapbox. https://github.com/rasterio/rasterio

Gillies, S., van der Wel, C., Van den Bossche, J., Taves, M. W., Arnott, J., Ward, B. C., & others. (2024). Shapely (2.0.4) [Computer software]. Zenodo. https://doi.org/10.5281/zenodo.10982792

Google Earth Engine (n.d.-a). Coding Best Practices | Google Earth Engine | Google for Developers. Retrieved June 15, 2024, from https://developers.google.com/earth-engine/guides/best_practices

Google Earth Engine (n.d.-b). Debugging guide | Google Earth Engine | Google for Developers. Retrieved June 15, 2024, from https://developers.google.com/earth-engine/guides/debugging

Gorelick, N., Hancher, M., Dixon, M., Ilyushchenko, S., Thau, D., & Moore, R. (2017). Google Earth Engine: Planetary-scale geospatial analysis for everyone. Remote Sensing of Environment, 202, 18–27. https://doi.org/10.1016/j.rse.2017.06.031

Gorelick, N., Yang, Z., Arévalo, P., Bullock, E. L., Insfrán, K. P., & Healey, S. P. (2023). A global time series dataset to facilitate forest greenhouse gas reporting. Environmental Research Letters, 18(8), 084001. https://doi.org/10.1088/1748-9326/ace2da

Griffiths, P., Nendel, C., & Hostert, P. (2019). Intra-annual reflectance composites from Sentinel-2 and Landsat for national-scale crop and land cover mapping. Remote Sensing of Environment, 220, 135–151. https://doi.org/10.1016/j.rse.2018.10.031

Grogan, D., Frolking, S., Wisser, D., Prusevich, A., & Glidden, S. (2022). Global gridded crop harvested area, production, yield, and monthly physical area data circa 2015. Scientific Data, 9(1), 15. https://doi.org/10.1038/s41597-021-01115-2

Gumma, M. K., Thenkabail, P. S., Teluguntla, P., Rao, M. N., Mohammed, I. A., & Whitbread, A. M. (2016). Mapping rice-fallow cropland areas for short-season grain legumes intensification in South Asia using MODIS 250 m time-series data. International Journal of Digital Earth, 9(10), 981–1003. https://doi.org/10.1080/17538947.2016.1168489
Hansen, M. C., Potapov, P. V., Moore, R., Hancher, M., Turubanova, S. A., Tyukavina, A.,
Thau, D., Stehman, S. V., Goetz, S. J., Loveland, T. R., Kommareddy, A., Egorov, A., Chini,
L., Justice, C. O., & Townshend, J. R. G. (2013). High-Resolution Global Maps of 21st-Century
Forest Cover Change. Science, 342(6160), 850–853. https://doi.org/10.1126/science.1244693

Hijmans, R. J. (2023a). raster: Geographic Data Analysis and Modeling. https://CRAN.R-project.org/package=raster

Hijmans, R. J. (2023b). terra: Spatial Data Analysis [Computer software]. https://CRAN.R-project.org/package=terra

Homolová, L., Malenovský, Z., Clevers, J. G. P. W., García-Santos, G., & Schaepman, M. E. (2013). Review of optical-based remote sensing for plant trait mapping. Ecological Complexity, 15, 1–16. https://doi.org/10.1016/j.ecocom.2013.06.003

Hoyer, S., & Hamman, J. (2017). xarray: N-D labeled Arrays and Datasets in Python. Journal of Open Research Software, 5(1). https://doi.org/10.5334/jors.148

Hoyer, S., Roos, M., Joseph, H., Magin, J., Cherian, D., Fitzgerald, C., Hauser, M., Fujii, K., Maussion, F., Imperiale, G., Clark, S., Kleeman, A., Nicholas, T., Kluyver, T., Westling, J., Munroe, J., Amici, A., Barghini, A., Banihirwe, A., ... Wolfram, P. J. (2024). Xarray (v2024.05.0) [Computer software]. Zenodo. https://doi.org/10.5281/zenodo.11183201

lizumi, T., & Sakai, T. (2020). The global dataset of historical yields for major crops 1981– 2016. Scientific Data, 7(1), 97. https://doi.org/10.1038/s41597-020-0433-7

International Food Policy Research Institute (IFPRI), I. F. P. R. I. (IFPRI). (2024). Global Spatially-Disaggregated Crop Production Statistics Data for 2020 Version 1.0 (Version V3) [dataset]. Harvard Dataverse. https://doi.org/10.7910/DVN/SWPENT

Joglekar, A. K. B., Wood-Sichra, U., & Pardey, P. G. (2019). Pixelating crop production: Consequences of methodological choices. PLOS ONE, 14(2), e0212281. https://doi.org/10.1371/journal.pone.0212281

Jordahl, K., Bossche, J. V. den, Fleischmann, M., Wasserman, J., McBride, J., Gerard, J., Tratner, J., Perry, M., Badaracco, A. G., Farmer, C., Hjelle, G. A., Snow, A. D., Cochran, M., Gillies, S., Culbertson, L., Bartos, M., Eubank, N., maxalbert, Bilogur, A., ... Leblanc, F. (2020). geopandas/geopandas: V0.8.1 (v0.8.1) [Computer software]. Zenodo. https://doi.org/10.5281/zenodo.3946761

Kalischek, N., Lang, N., Renier, C., Daudt, R. C., Addoah, T., Thompson, W., ... & Wegner, J. D. (2023). Cocoa plantations are associated with deforestation in Côte d'Ivoire and Ghana. Nature Food, 4(5), 384-393. https://doi.org/10.1038/s43016-023-00751-8

Karra, K., Kontgis, C., Statman-Weil, Z., Mazzariello, J. C., Mathis, M., & Brumby, S. P. (2021). Global land use / land cover with Sentinel 2 and deep learning. 2021 IEEE International Geoscience and Remote Sensing Symposium IGARSS, 4704–4707. https://doi.org/10.1109/IGARSS47720.2021.9553499

Kerner, H., Nakalembe, C., Yang, A., Zvonkov, I., McWeeny, R., Tseng, G., & Becker-Reshef, I. (2024). How accurate are existing land cover maps for agriculture in Sub-Saharan Africa? Scientific Data, 11(1), 486. https://doi.org/10.1038/s41597-024-03306-z

Klein Goldewijk, K., Beusen, A., Doelman, J., & Stehfest, E. (2017). Anthropogenic land use estimates for the Holocene – HYDE 3.2. Earth System Science Data, 9(2), 927–953. https://doi.org/10.5194/essd-9-927-2017

Kong, D., Zhang, Y., Gu, X., & Wang, D. (2019). A robust method for reconstructing global MODIS EVI time series on the Google Earth Engine. ISPRS Journal of Photogrammetry and Remote Sensing, 155, 13–24. <u>https://doi.org/10.1016/j.isprsjprs.2019.06.014</u>

Kokaly, R. F. (2001). Investigating a Physical Basis for Spectroscopic Estimates of Leaf Nitrogen Concentration. Remote Sensing of Environment, 75(2), 153–161. https://doi.org/10.1016/S0034-4257(00)00163-2

Kumar, V., Huber, M., Rommen, B. et al. (2022). Agricultural SandboxNL: A national-scale database of parcel-level processed Sentinel-1 SAR data. Sci Data 9, 402. https://doi.org/10.1038/s41597-022-01474-4

Krzysztof, D. (2023). Parallel raster processing in stars. <u>https://kadyb.github.io/stars-parallel/Tutorial.html</u>

Li, J., Xu, M., Xiang, L., Chen, D., Zhuang, W., Yin, X., & Li, Z. (2024). Foundation models in smart agriculture: Basics, opportunities, and challenges. Computers and Electronics in Agriculture, 222, 109032. https://doi.org/10.1016/j.compag.2024.109032

Lobell, D. B., Azzari, G., Burke, M., Gourlay, S., Jin, Z., Kilic, T., & Murray, S. (2020). Eyes in the Sky, Boots on the Ground: Assessing Satellite- and Ground-Based Approaches to Crop Yield Measurement and Analysis. American Journal of Agricultural Economics, 102(1), 202–219. https://doi.org/10.1093/ajae/aaz051

Lobell, D. B., Thau, D., Seifert, C., Engle, E., & Little, B. (2015). A scalable satellite-based crop yield mapper. Remote Sensing of Environment, 164, 324–333. https://doi.org/10.1016/j.rse.2015.04.021

Lovelace, R., Nowosad, J., & Muenchow, J. (2019). Geocomputation with R. Chapman and Hall/CRC. https://r.geocompx.org/

Lu, J., Yang, T., Su, X., Qi, H., Yao, X., Cheng, T., Zhu, Y., Cao, W., & Tian, Y. (2020). Monitoring leaf potassium content using hyperspectral vegetation indices in rice leaves. Precision Agriculture, 21(2), 324–348. https://doi.org/10.1007/s11119-019-09670-w

Maggi, F., Tang, F. H. M., La Cecilia, D., & McBratney, A. (2019). PEST-CHEMGRIDS, global gridded maps of the top 20 crop-specific pesticide application rates from 2015 to 2025. Scientific Data, 6(1), 170. <u>https://doi.org/10.1038/s41597-019-0169-4</u>

Main, R., Cho, M. A., Mathieu, R., O'Kennedy, M. M., Ramoelo, A., & Koch, S. (2011). An investigation into robust spectral indices for leaf chlorophyll estimation. ISPRS Journal of Photogrammetry and Remote Sensing, 66(6), 751–761. https://doi.org/10.1016/j.isprsjprs.2011.08.001

Montero, D., Aybar, C., Mahecha, M. D., Martinuzzi, F., Söchting, M., & Wieneke, S. (2023). A standardized catalogue of spectral indices to advance the use of remote sensing in Earth system research. Scientific Data, 10(1), 197. https://doi.org/10.1038/s41597-023-02096-0

Massari, C., Modanesi, S., Dari, J., Gruber, A., De Lannoy, G. J. M., Girotto, M., Quintana-Seguí, P., Le Page, M., Jarlan, L., Zribi, M., Ouaadi, N., Vreugdenhil, M., Zappa, L., Dorigo, W., Wagner, W., Brombacher, J., Pelgrum, H., Jaquot, P., Freeman, V., ... Brocca, L. (2021). A Review of Irrigation Information Retrievals from Space and Their Utility for Users. Remote Sensing, 13(20), Article 20. https://doi.org/10.3390/rs13204112

Matsunaga, T., Iwasaki, A., Tachikawa, T., Tanii, J., Kashimura, O., Mouri, K., Inada, H., Tsuchida, S., Nakamura, R., Yamamoto, H., & Iwao, K. (2020). Hyperspectral Imager Suite (HISUI): Its Launch and Current Status. IGARSS 2020 - 2020 IEEE International Geoscience and Remote Sensing Symposium, 3272–3273. https://doi.org/10.1109/IGARSS39084.2020.9323376

Meyer, H., & Pebesma, E. (2021). Predicting into unknown space? Estimating the area of applicability of spatial prediction models. Methods in Ecology and Evolution, 12(9), 1620–1633. https://doi.org/10.1111/2041-210X.13650

Mo, L., Zohner, C. M., Reich, P. B., Liang, J., De Miguel, S., Nabuurs, G.-J., Renner, S. S., Van Den Hoogen, J., Araza, A., Herold, M., Mirzagholi, L., Ma, H., Averill, C., Phillips, O. L., Gamarra, J. G. P., Hordijk, I., Routh, D., Abegg, M., Adou Yao, Y. C., ... Crowther, T. W. (2023). Integrated global assessment of the natural forest carbon potential. Nature, 624(7990), 92–101. https://doi.org/10.1038/s41586-023-06723-z

Monfreda, C., Ramankutty, N., & Foley, J. A. (2008). Farming the planet: 2. Geographic distribution of crop areas, yields, physiological types, and net primary production in the year

 2000.
 Global
 Biogeochemical
 Cycles,
 22(1),
 2007GB002947.

 https://doi.org/10.1029/2007GB002947

 2007GB002947.

Moreno-Martínez, Á., Camps-Valls, G., Kattge, J., Robinson, N., Reichstein, M., van Bodegom, P., Kramer, K., Cornelissen, J. H. C., Reich, P., Bahn, M., Niinemets, Ü., Peñuelas, J., Craine, J. M., Cerabolini, B. E. L., Minden, V., Laughlin, D. C., Sack, L., Allred, B., Baraloto, C., ... Running, S. W. (2018). A methodology to derive global maps of leaf traits using remote sensing and climate data. Remote Sensing of Environment, 218, 69–88. https://doi.org/10.1016/j.rse.2018.09.006

Mueller, N. D., Gerber, J. S., Johnston, M., Ray, D. K., Ramankutty, N., & Foley, J. A. (2012). Closing yield gaps through nutrient and water management. Nature, 490(7419), 254–257. https://doi.org/10.1038/nature11420

Mühleisen, H., & Raasveldt, M. (2024). duckdb: DBI Package for the DuckDB Database Management System [Computer software]. https://r.duckdb.org/

Nakalembe, C., & Kerner, H. (2023). Considerations for AI-EO for agriculture in Sub-Saharan Africa. Environmental Research Letters, 18(4), 041002. https://doi.org/10.1088/1748-9326/acc476

Nakalembe, C., Becker-Reshef, I., Bonifacio, R., Hu, G., Humber, M. L., Justice, C. J., Keniston, J., Mwangi, K., Rembold, F., Shukla, S., Urbano, F., Whitcraft, A. K., Li, Y., Zappacosta, M., Jarvis, I., & Sanchez, A. (2021). A review of satellite-based global agricultural monitoring systems available for Africa. Global Food Security, 29, 100543. https://doi.org/10.1016/j.gfs.2021.100543

NASA. (2024). atmospheric carbon dioxide (CO2) measurements by the Orbiting Carbon Observatory-2. https://disc.gsfc.nasa.gov/datasets/.

Olofsson, P., Foody, G. M., Herold, M., Stehman, S. V., Woodcock, C. E., & Wulder, M. A. (2014). Good practices for estimating area and assessing accuracy of land change. Remote Sensing of Environment, 148, 42–57. https://doi.org/10.1016/j.rse.2014.02.015

Paliwal, A., & Jain, M. (2020). The Accuracy of Self-Reported Crop Yield Estimates and Their Ability to Train Remote Sensing Algorithms. Frontiers in Sustainable Food Systems, 4. https://doi.org/10.3389/fsufs.2020.00025

Parente, L., Sloat, L., Mesquita, V., Consoli, D., Stanimirova, R., Hengl, T., Bonannella, C.,
Teles, N., Wheeler, I., Ehrmann, S., Hunter, M., Ferreira, L., Mattos, A. P., Oliveira, B., Meyer,
C., Şahin, M., Witjes, M., Fritz, S., Malek, Ž., & Stolle, F. (2024). Mapping global grassland

dynamics 2000—2022 at 30m spatial resolution using spatiotemporal Machine Learning. https://doi.org/10.21203/rs.3.rs-4514820/v1

Pasquarella, V. (2023). All Clear with Cloud Score+. By Valerie Pasquarella, Research... | by Google Earth | Google Earth and Earth Engine | Medium. https://medium.com/google-earth/all-clear-with-cloud-score-bd6ee2e2235e

Pasquarella, V. J., Holden, C. E., Kaufman, L., & Woodcock, C. E. (2016). From imagery to ecology: Leveraging time series of all available Landsat observations to map and monitor ecosystem state and dynamics. Remote Sensing in Ecology and Conservation, 2(3), 152–170. https://doi.org/10.1002/rse2.24

Pebesma, E. (2018). Simple Features for R: Standardized Support for Spatial Vector Data. The R Journal, 10(1), 439. https://doi.org/10.32614/RJ-2018-009

Pebesma, E., & Bivand, R. (2023). Spatial Data Science: With Applications in R (1st ed.). Chapman and Hall/CRC. <u>https://doi.org/10.1201/9780429459016</u>

Pierce, D. (2023). ncdf4: Interface to Unidata netCDF (Version 4 or Earlier) Format Data Files [Computer software]. <u>https://CRAN.R-project.org/package=ncdf4</u>

Pinter, Jr., P. J., Hatfield, J. L., Schepers, J. S., Barnes, E. M., Moran, M. S., Daughtry, C. S.
T., & Upchurch, D. R. (2003). Remote Sensing for Crop Management. Photogrammetric
Engineering & Remote Sensing, 69(6), 647–664. https://doi.org/10.14358/PERS.69.6.647

Portmann, F. T., Siebert, S., & Döll, P. (2010). MIRCA2000—Global monthly irrigated and rainfed crop areas around the year 2000: A new high-resolution data set for agricultural and hydrological modeling. Global Biogeochemical Cycles, 24(1), 2008GB003435. https://doi.org/10.1029/2008GB003435

Porwollik, V., Rolinski, S., Heinke, J., & Müller, C. (2019). Generating a rule-based global gridded tillage dataset. Earth System Science Data, 11(2), 823–843. https://doi.org/10.5194/essd-11-823-2019

Potapov, P., Hansen, M. C., Turubanova, S., Tyukavina, A., Zalles, V., Song, X.-P., & Khan, A. (2023). Reply to: Measuring the world's cropland area. Nature Food, 4(1), 33–34. https://doi.org/10.1038/s43016-022-00668-8

Potapov, P., Turubanova, S., Hansen, M. C., Tyukavina, A., Zalles, V., Khan, A., Song, X.-P., Pickens, A., Shen, Q., & Cortez, J. (2021). Global maps of cropland extent and change show accelerated cropland expansion in the twenty-first century. Nature Food, 3(1), 19–28. https://doi.org/10.1038/s43016-021-00429-z QGIS Development Team (2024). QGIS Geographic Information System [Computer software]. Open Source Geospatial Foundation Project. http://qgis.osgeo.org

R Core Team. (2023). R: A Language and Environment for Statistical Computing [Computer software]. R Foundation for Statistical Computing. https://www.R-project.org/

Ramankutty, N., Evan, A. T., Monfreda, C., & Foley, J. A. (2008). Farming the planet: 1. Geographic distribution of global agricultural lands in the year 2000. Global Biogeochemical Cycles, 22(1), 2007GB002952. https://doi.org/10.1029/2007GB002952

Ray, D. K., West, P. C., Clark, M., Gerber, J. S., Prishchepov, A. V., & Chatterjee, S. (2019). Climate change has likely already affected global food production. PLOS ONE, 14(5), e0217148. https://doi.org/10.1371/journal.pone.0217148

Rembold, F., Atzberger, C., Savin, I., & Rojas, O. (2013). Using Low Resolution Satellite Imagery for Yield Prediction and Yield Anomaly Detection. Remote Sensing, 5(4), 1704–1733. https://doi.org/10.3390/rs5041704

Richardson, N., Cook, I., Crane, N., Dunnington, D., François, R., Keane, J., Moldovan-Grünfeld, D., Ooms, J., & Apache Arrow. (2023). arrow: Integration to "Apache" "Arrow" [Computer software]. https://CRAN.R-project.org/package=arrow

Rocklin, M. (2015). Dask: Parallel computation with blocked algorithms and task scheduling. Proceedings of the 14th Python in Science Conference, 130–136.

Ross, N. (2023). fasterize: Fast Polygon to Raster Conversion [Computer software]. https://CRAN.R-project.org/package=fasterize

Roy, S., & Swetnam, T. (2024). samapriya/awesome-gee-community-datasets: Community Catalog (2.6.0) [Computer software]. Zenodo. https://doi.org/10.5281/zenodo.11118613

Sacks, W. J., Deryng, D., Foley, J. A., & Ramankutty, N. (2010). Crop planting dates: An analysis of global patterns. Global Ecology and Biogeography, 19(5), 607–620. https://doi.org/10.1111/j.1466-8238.2010.00551.x

Siebert, S., Kummu, M., Porkka, M., Döll, P., Ramankutty, N., & Scanlon, B. (2015). A global data set of the extent of irrigated land from 1900 to 2005. https://doi.org/10.13019/M20599

Siedliska, A., Baranowski, P., Pastuszka-Woźniak, J., Zubik, M., & Krzyszczak, J. (2021). Identification of plant leaf phosphorus content at different growth stages based on hyperspectral reflectance. BMC Plant Biology, 21(1), 28. https://doi.org/10.1186/s12870-020-02807-4

Singh, G., Moncrieff, G., Venter, Z., Cawse-Nicholson, K., Slingsby, J., & Robinson, T. B. (2024). Uncertainty quantification for probabilistic machine learning in earth observation using conformal prediction (arXiv:2401.06421). arXiv. https://doi.org/10.48550/arXiv.2401.06421

Sinha, S., Jeganathan, C., Sharma, L. K., & Nathawat, M. S. (2015). A review of radar remote sensing for biomass estimation. International Journal of Environmental Science and Technology, 12(5), 1779–1792. https://doi.org/10.1007/s13762-015-0750-0

Snow, A. D., RichardScottOZ, Raspaud, M., Brochart, D., Kouzoubov, K., Braun, R., Henderson, S., Amici, A., stefank0, Chegini, T., Bell, R., Taves, M., pmallas, apiwat-chantawibul, Fröhlich, Y., Nogueira, X., Mathuranayagam, V. B., Augspurger, T., Badger, T. G., ... Carroll, I. (2024). corteva/rioxarray: 0.15.5 Release (0.15.5) [Computer software]. Zenodo. https://doi.org/10.5281/zenodo.11035866

Staver, A. C., & Hansen, M. C. (2015). Analysis of stable states in global savannas: Is the CART pulling the horse? – a comment. Global Ecology and Biogeography, 24(8), 985–987. https://doi.org/10.1111/geb.12285

Sornapudi, S., & Singh, R. (2024). Self-Supervised Backbone Framework for Diverse Agricultural Vision Tasks. arXiv preprint arXiv:2403.15248. https://doi.org/10.48550/arXiv.2403.15248

Song, X. P., Hansen, M. C., Stehman, S. V., Potapov, P. V., Tyukavina, A., Vermote, E. F., & Townshend, J. R. (2018). Global land change from 1982 to 2016. Nature, 560(7720), 639-643.

Storch, T., Honold, H.-P., Chabrillat, S., Habermeyer, M., Tucker, P., Brell, M., Ohndorf, A., Wirth, K., Betz, M., Kuchler, M., Mühle, H., Carmona, E., Baur, S., Mücke, M., Löw, S., Schulze, D., Zimmermann, S., Lenzen, C., Wiesner, S., ... Fischer, S. (2023). The EnMAP imaging spectroscopy mission towards operations. Remote Sensing of Environment, 294, 113632. https://doi.org/10.1016/j.rse.2023.113632

Sulla-Menashe, D., Gray, J. M., Abercrombie, S. P., & Friedl, M. A. (2019). Hierarchical mapping of annual global land cover 2001 to present: The MODIS Collection 6 Land Cover product. Remote Sensing of Environment, 222, 183–194. https://doi.org/10.1016/j.rse.2018.12.013

Tang, F. H. M., Nguyen, T. H., Conchedda, G., Casse, L., Tubiello, F. N., & Maggi, F. (2024). CROPGRIDS: A global geo-referenced dataset of 173 crops. Scientific Data, 11(1), 413. https://doi.org/10.1038/s41597-024-03247-7 Teixeira, I., Morais, R., Sousa, J. J., & Cunha, A. (2023). Deep learning models for the classification of crops in aerial imagery: A review. Agriculture, 13(5), 965. https://doi.org/10.3390/agriculture13050965

Teluguntla, P., Thenkabail, P., Oliphant, A., Gumma, M., Aneece, I., Foley, D., & Mccormick, R. (2023). The GFSAD Landsat-derived Global Rainfed and Irrigated-Cropland Product at nominal 30m of the World (GFSADLGRIP30WORLD) [dataset]. NASA EOSDIS Land Processes DAAC. https://doi.org/10.5067/Community/LGRIP/LGRIP30.001

Tennekes, M. (2018). tmap: Thematic Maps in R. Journal of Statistical Software, 84(6), 1–39. https://doi.org/10.18637/jss.v084.i06

Thenkabail, P. S., Teluguntla, P. G., Xiong, J., Oliphant, A., Congalton, R. G., Ozdogan, M., Gumma, M. K., Tilton, J. C., Giri, C., Milesi, C., Phalke, A., Massey, R., Yadav, K., Sankey, T., Zhong, Y., Aneece, I., & Foley, D. (2021). Global cropland-extent product at 30-m resolution (GCEP30) derived from Landsat satellite time-series data for the year 2015 using multiple machine-learning algorithms on Google Earth Engine cloud. In Professional Paper (1868). U.S. Geological Survey. https://doi.org/10.3133/pp1868

Tian, H., Bian, Z., Shi, H., Qin, X., Pan, N., Lu, C., Pan, S., Tubiello, F. N., Chang, J., Conchedda, G., Liu, J., Mueller, N., Nishina, K., Xu, R., Yang, J., You, L., & Zhang, B. (2022). History of anthropogenic Nitrogen inputs (HaNi) to the terrestrial biosphere: A 5 arcmin resolution annual dataset from 1860 to 2019. Earth System Science Data, 14(10), 4551–4568. https://doi.org/10.5194/essd-14-4551-2022

Tubiello, F. N., Conchedda, G., Casse, L., Pengyu, H., Zhongxin, C., De Santis, G., Fritz, S., & Muchoney, D. (2023). Measuring the world's cropland area. Nature Food, 4(1), 30–32. https://doi.org/10.1038/s43016-022-00667-9

USDA National Agricultural Statistics Service (n.d.). Cropland Data Layer [dataset]. https://nassgeodata.gmu.edu/CropScape

USGS (n.d.). Landsat Collection 2 Level-2 Science Products | U.S. Geological Survey. Retrieved June 12, 2024, from https://www.usgs.gov/landsat-missions/landsat-collection-2level-2-science-products

Van Rossum, G., & Drake, F. L. (2009). Python 3 Reference Manual. CreateSpace.

Van Tricht, K., Degerickx, J., Gilliams, S., Zanaga, D., Battude, M., Grosu, A., Brombacher, J.,
Lesiv, M., Bayas, J. C. L., Karanam, S., Fritz, S., Becker-Reshef, I., Franch, B., Mollà-Bononad,
B., Boogaard, H., Pratihast, A. K., Koetz, B., & Szantoi, Z. (2023). WorldCereal: A dynamic

open-source system for global-scale, seasonal, and reproducible crop and irrigation mapping. Earth System Science Data, 15(12), 5491–5515. https://doi.org/10.5194/essd-15-5491-2023

Weiss, M., Jacob, F., & Duveiller, G. (2019). Remote sensing for agricultural applications: A meta-review. Remote Sensing of Environment, 236, 111402. https://doi.org/10.1016/j.rse.2019.111402

Weston, S. (2022a). doParallel: Foreach Parallel Adaptor for the "parallel" Package [Computer software]. https://CRAN.R-project.org/package=doParallel

Weston, S. (2022b). foreach: Provides Foreach Looping Construct [Computer software]. https://CRAN.R-project.org/package=foreach

Wickham, H. (2016). ggplot2: Elegant Graphics for Data Analysis. Springer-Verlag New York. https://ggplot2.tidyverse.org

Wickham, H., François, R., Henry, L., Müller, K., & Vaughan, D. (2023). dplyr: A Grammar of Data Manipulation [Computer software]. https://CRAN.R-project.org/package=dplyr

Wickham, H., Girlich, M., Fairbanks, M., & Dickerson, R. (2023). dtplyr: Data Table Back-End for "dplyr" [Computer software]. https://CRAN.R-project.org/package=dtplyr

Wu, Q. (2020). geemap: A Python package for interactive mapping with Google Earth Engine. Journal of Open Source Software, 5(51), 2305. https://doi.org/10.21105/joss.02305

Yang, S., Madsen, M. S., & Bednar, J. A. (2022). HoloViz: Visualization and Interactive Dashboards in Python. Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, 4846–4847. https://doi.org/10.1145/3534678.3542621

Yin, H., Brandão, A., Buchner, J., Helmers, D., Iuliano, B. G., Kimambo, N. E., Lewińska, K.
E., Razenkova, E., Rizayeva, A., Rogova, N., Spawn, S. A., Xie, Y., & Radeloff, V. C. (2020).
Monitoring cropland abandonment with Landsat time series. Remote Sensing of Environment, 246, 111873. https://doi.org/10.1016/j.rse.2020.111873

You, L., Wood, S., Wood-Sichra, U., & Wu, W. (2014). Generating global crop distribution maps: From census to grid. Agricultural Systems, 127, 53–60. https://doi.org/10.1016/j.agsy.2014.01.002

Yu, Q., You, L., Wood-Sichra, U., Ru, Y., Joglekar, A. K. B., Fritz, S., Xiong, W., Lu, M., Wu, W., & Yang, P. (2020). A cultivated planet in 2010 – Part 2: The global gridded agricultural-production maps. Earth System Science Data, 12(4), 3545–3572. https://doi.org/10.5194/essd-12-3545-2020

46

Zhang, M., Wu, B., Zeng, H., He, G., Liu, C., Tao, S., Zhang, Q., Nabil, M., Tian, F., Bofana, J., Beyene, A. N., Elnashar, A., Yan, N., Wang, Z., & Liu, Y. (2021). GCI30: A global dataset of 30 m cropping intensity using multisource remote sensing imagery. Earth System Science Data, 13(10), 4799–4817. https://doi.org/10.5194/essd-13-4799-2021

Zhang, X., Zhao, T., Xu, H., Liu, W., Wang, J., Chen, X., & Liu, L. (2024). GLC_FCS30D: The first global 30 m land-cover dynamics monitoring product with a fine classification system for the period from 1985 to 2022 generated using dense-time-series Landsat imagery and the continuous change-detection method. Earth System Science Data, 16(3), 1353–1381. https://doi.org/10.5194/essd-16-1353-2024

Zheng, B., Campbell, J. B., Serbin, G., & Galbraith, J. M. (2014). Remote sensing of crop residue and tillage practices: Present capabilities and future prospects. Soil and Tillage Research, 138, 26–34. https://doi.org/10.1016/j.still.2013.12.009

Zortea, M., Nery, M., Ruga, B., Carvalho, L. B., & Bastos, A. C. (2018). Oil-palm tree detection in aerial images combining deep learning classifiers. In IGARSS 2018-2018 IEEE International Geoscience and Remote Sensing Symposium (pp. 657-660). IEEE. https://doi.org/10.1109/IGARSS.2018.8519239