One Acre Fund

Measuring Yields From Space

Goldilocks Toolkit
Innovations for Poverty Action
poverty-action.org/goldilocks

Right-fit monitoring and evaluation (M&E) systems embody the principles of Credible, Actionable, Responsible, and Transportable, or CART. In the Goldilocks case study series, we examine the M&E systems of several innovative organizations and explore how the CART Principles can work in practice.

Acknowledgements: The Goldilocks toolkit was authored by Mary Kay Gugerty, Dean Karlan, and Delia Welsh. Lead authors for this toolkit resource also include Giovanni Castaldo (UC Berkeley) and Temina Madon (CEGA), with editing and design by the IPA Communications team (David Batcheck, Laura Burke, Heidi McAnnally-Linz, Cara Vu).

One Acre Fund: Using Satellite Imagery to Evaluate Smallholder Farmers' Productivity

A substantial proportion of the world's poor are smallholder farmers, and measuring their productivity is a basic first step in understanding where livelihoods are improving (and where additional help is needed). Unfortunately, comprehensive data on farmers' productivity do not exist. They are collected piecemeal by researchers and governments in limited areas of the globe, and over limited periods, typically using household surveys which are both expensive and time consuming.

One Acre Fund (1AF) is a social enterprise working with such farmers in East Africa. 1AF's mission is to address long-standing barriers to farmer productivity. 1AF provides seeds and fertilizers on credit, delivers agricultural inputs within walking distance of farmers' homes, trains farmers on how to use inputs most effectively, helps them store their crops safely, and links them to local traders to provide access to markets.

This comprehensive approach is geared towards maximizing farm profitability.

Like many organizations, 1AF is eager to demonstrate and improve the impact of its programs. To track its progress, the organization relies essentially on field surveys. They assess productivity gains by comparing the costs of farmers' agricultural inputs to the market value of outputs (and for the latter, 1AF manually weighs a small portion of farmers' harvests). In 2015, 1AF weighed over 16,000 harvests across four countries.

Despite recent M&E revamping efforts – staffing up each country office with a dedicated person, equipping field surveyors with tablets, and reducing selection bias in the identification of control groups– 1AF is continuously looking to professionalize its evaluation practices. More efficient M&E practices would enable an increase in sample sizes (providing the capacity to look at a broader set of plots within and outside the organization's programs) and would help strengthen the link of causality between 1AF's intervention and its impact. Since field surveying is particularly time-consuming, expensive, and prone to measurement errors, 1AF is always looking for new methods for generating actionable data faster, cheaper, and more accurately.

Technology Solution

A large body of research has shown that satellite imagery can provide accurate estimates of crop yields for large farms in both the developed and developing world[1]. Yet attempts to expand this work to the much smaller plot sizes in East Africa have been hampered by a lack of high-resolution imagery, and by a lack of geo-referenced ground-level data needed to calibrate and test remote sensing methods.

Instead, traditional estimates of farmer- and field-level yields have been obtained from ground surveys that are a) expensive to conduct; b) difficult to follow up on; and c) prone to measurement errors, such as recall bias among respondents (see Figure 1 for an example of errors in self-reported area under cultivation). Currently, 1AF relies primarily on crop cuts, which constitute the least biased and most reliable estimate of harvests. Crop cuts over small areas are not ideal because farmers' fields are not homogenous: one part may be extremely rich and productive, while another area may be overcrowded and stressed. Nevertheless, with a large enough sample and random placement, crop cuts are a very robust way to estimate program impact by comparing 1AF and comparison farmers.

This case study explores the use of satellite imagery to predict yields on

FIGURE 1. MEASUREMENT ERROR, A PERSISTENT CHALLENGE WITH SURVEY-BASED RESEARCH

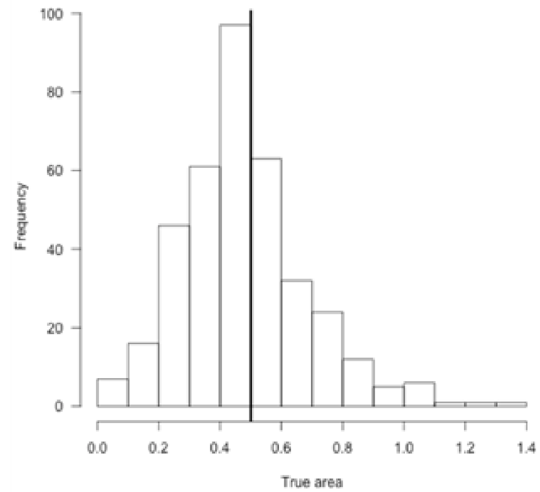

The histogram shows field sizes (independently measured) for a sample of plots that Ugandan farmers reported as 0.5 acres. This demonstrates the discrepancy between farmer-reported field size, and independent measurement of the same plots (Source: LSMS).

1 For a recent review, see David Lobell, "The use of satellite data for crop yield gap analysis", *Field Crops Research* 143 (2013).

small-holder plots. With an increasing number of start-ups and larger companies commercializing satellite imagery over the past couple of years, the quality and volume of imagery has substantially increased. This has generated more accurate and scalable images at an increasingly accessible cost. At the same time, unprecedented large-scale surveying efforts currently underway in certain parts of the developing world offer a unique opportunity for groundtruthing remote sensing estimates.

Technology Application

CEGA researchers Marshall Burke and David Lobell, along with Karthik Rajkumar (all at Stanford University) partnered with Skybox Imagery and other companies to estimate smallholder yields from high quality satellite imagery (Figure 2). The ground data ('ground truth') comprised information on the geolocation of farmers' field boundaries provided by 1AF (Figure 3), along with surveys asking farmers about the crops they grew, how much food they produced and what they planned to do with it.

To extract yield information from the satellite imagery, the research team derived a vegetation index (VI) value for each pixel in an image. The index is used to distinguish plant material from other features like terrain, land, water or urban areas. The index is based on the capacity

FIGURE 2. EXAMPLES OF IMAGERY USED IN THE PROJECT

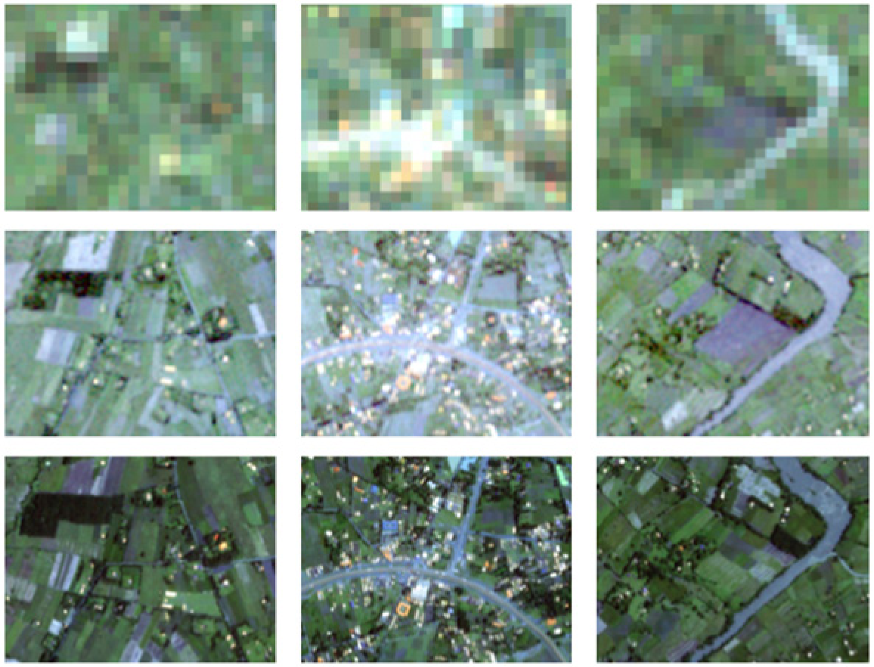

These are all sites in Western Kenya and the matrix is divided in rows with images from the same satellite source and resolution. From top to bottom, the sources are Landsat (30 m resolution), RapidEye (5m), and Skybox (1m). The dates these images were taken are all different, which explains the heterogeneity in vegetation cover across them.

of materials to absorb light across different wavelengths. Since plants and bare land absorb and reflect in different bands of the light spectrum, and at different intensities, it is possible to use the reflectance data from multiple bands of satellite sensors to distinguish the two.

FIGURE 3. SOME OF THE FIELDS BEING MONITORED IN AN ONGOING EXPERIMENT IN KENYA WITH 1AF

The boundaries for each smallholder plot are obtained by walking along the rim of the field and recording its shape and location using a handheld GPS device.

Results and Recommendations

The first step in using imagery to predict crop yields is to process and calibrate the raw satellite data, to correct for geolocation errors and atmospheric scattering. The CEGA team next developed algorithms to estimate plot-specific yields, based on VI. The researchers finally compared their plot-level estimates (generated from the Skybox raster from June 2014) with ground-truth data from 1AF, finding encouraging correlations.

In the figures that follow, plot-level prediction of maize yields from satellite data obtained in June 2014 (x-axis) are directly compared with yields measured during 1AF's crop cutting exercise (y-axis).

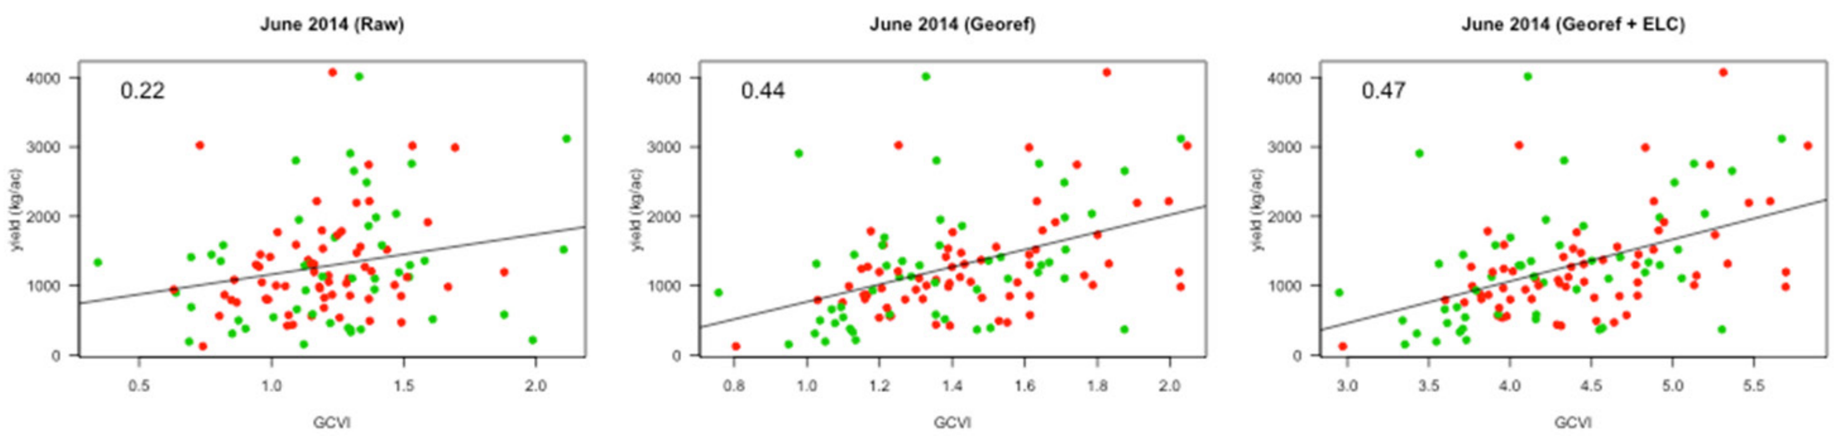

FIGURE 4. COMPARISON OF THE VARIOUS STAGES IN THE PREPROCESSING OF THE RASTER DATA

Each figure shows the plot-level yield estimates based on satellite imagery (x-axis) plotted against the measured yield using crop-cut data from 1AF. Each dot represents one farm plot.

The yield predictions are cleared of (removable) error terms by using 1AF's GPS-enabled area measures wherever possible. As shown in the leftmost panel of Figure 4, important relationships are between the two estimates are obscured when raw data with no processing are used. However, the middle panel shows the same comparison after the satellite data have been georeferenced (i.e., a geometric correction has been applied), lifting the correlation coefficient between vegetative index (VI) and yield from 0.22 to 0.44. Further corrections for atmospheric scattering improve the correlation to 0.47.

Goldilocks Principles

The use of satellite imagery to evaluate small farmers' productivity is at early stages of development. The satellite-based yield predictions are not perfectly correlated with the estimates provided by current 1AF's unshelled weights. However, as this case study suggests, the approach does adhere to the Goldilocks CART principles.

Credible: Collect high quality data and accurately analyze the data.

By comparing high quality ground-truth data from One Acre Fund with satellite estimates of yield, the CEGA team determined that visual indices calculated with satellite imagery are valid and reliable in detecting small plots' yields. Relying on reproducible analyses of satellite images—rather than lengthy surveys that depend on respondents' behavior and self-reports—makes yield measurement less biased and more credible. There is also less risk of enumerator error. Still, it is not clear yet if this method can replace crop cuts in the near future. 1AF's latest analysis has found that it is possible to predict shelled maize using only unshelled weight. Certainly, the efficiency gains of going from 2 weights to 1 weight are not as great as moving from doing no crop cuts to satellites. Nevertheless, for now switching to satellite imagery would mean losing a lot of data accuracy.

Actionable: Commit to act on the data you collect.

One of the greatest advantages of implementing satellite imagery as a way to evaluate smallholder farmers' productivity is the reduction in time needed to provide an actionable information stream. Even at this early stage, pictures taken from space can be quickly assessed with light groundtruthing contributions, or even by using data from other sites or crop simulations to calibrate the predictions. This translates in the potential for faster and more effective action. Understanding where and how to target 1AF services becomes a much easier task when frequent yield estimates are available.

Responsible: Ensure the benefits of data collection outweigh the costs.

Leveraging low-cost, readily available satellite imagery for M&E exemplifies the principle of responsibility: it provides thorough and objective data collection, and potentially minimizes costs (especially with the advent of low-cost microsatellites). Field surveying being particularly expensive, satellite images allow decision makers to work on actionable data faster and with fewer costs. With an increasing number of start-ups and larger companies commercializing satellite imagery, competition in the space has substantially increased. There are fixed costs of calibrating imagery and designing algorithms to detect fields and estimate yields; however, these are transportable across contexts and are declining over time as image providers improve their products. Overall the benefits and the scope of this data collection strategy largely outweigh the costs incurred.

Transportable: Collect data that will generate knowledge for other programs.

The scope of remote sensing is very broad. The data collected and analyzed by satellites for example can be used in a number of different applications, by a variety of actors. Government, aid organizations and researchers are all deeply interested in readily available productivity and wealth data. This makes the information stream coming from satellite images extremely transportable.

Going Forward

It is important to note that crop yields might not be the only relevant variable in determining the overall productivity of rural households. In the future, remote sensing might be used more holistically to distinguish between crop types, or to determine other dimensions of wellbeing for rural households. The development community can only benefit by exploring these monitoring and evaluation challenges through the lens of remote sensing.

In the immediate future, satellite imagery can be viewed an independent, cost-effective complement to current data collection methods for organizations addressing agricultural or environmental issues. While the technology is not ripe for full implementation right now, modest changes in current surveying practices – in 1AF's case, georeferencing plot size when they collect crop cuttings– could be the first step in generating more of the ground-truth data needed to improve satellite analytics. Ultimately, this investment in ground-truth will streamline data collection and improve impact measurement for the broader development community.