Stanford University
Human-Centered
Artificial Intelligence

APRIL 2022

Using Satellite Imagery to Understand and Promote Sustainable Development

Marshall Burke, Anne Driscoll,
David Lobell, and Stefano Ermon

THE NUMBER OF NON-MILITARY SATELLITES IN ORBIT IS RAPIDLY GROWING. Each of these satellites offers unprecedented access to imagery to help measure sustainable development outcomes, such as eliminating hunger, promoting health and well-being, and building sustainable communities. Images previously held only by governments and a few corporations are now widely available to academics, civil society organizations, and individuals. Researchers can analyze these images for a wide range of purposes, including to measure agricultural productivity, urban population density, and rural economic activity. A single satellite image, for example, might be able to tell the story of a village's economic health—its crop yields, its agricultural diversity, and its infrastructure development.

Artificial intelligence (AI)-powered machine learning (ML) tools can extract and assess such information from satellite imagery, making them an intriguing and valuable addition to the sustainable development toolkit. Yet many questions remain. In particular, researchers and policymakers need to better understand how well these models can map satellite image inputs to sustainable development outcomes and what limits the models' performance. Our paper in *Science*, "Using Satellite Imagery to Understand and Promote Sustainable Development," outlines how researchers have used ML models to estimate sustainable development outcomes, assess methods for model training, examine the challenges hindering models' improvement, and consider models' future applications. We conclude by identifying current limitations to these approaches and ways to respond.

KEY TAKEAWAYS

- The data needed to inform policymaking for sustainable development is often lacking or inaccurate.
- Machine learning analysis of satellite imagery could help estimate sustainable development outcomes—broadening the availability of existing high-quality development data.
- Currently, there is relatively limited adoption of satellite imagery analysis in many sustainable development domains, with the primary exceptions being population and agricultural measurement.
- Policymakers and researchers should explore using synthetic data and pursue more work on model explainability and scalability to ensure ML models can be appropriately trained for satellite images.

Introduction

Assessing sustainable development outcomes is critical to policymakers. These measures, which concern everything from water equality to food availability, quite literally affect human life, health, and well-being—as well as the broader health and sustainability of communities and the planet. However, traditional approaches for measuring these outcomes are time-consuming, expensive, and inaccurate; they often rely on extrapolating incomplete household or field-level data to broader populations. In Africa alone, 34 percent of countries have not had an updated agricultural survey in more than 15 years; in half of all African nations, nationally representative livelihood surveys occur once every six years. Development data that is unreliable or simply unavailable hamstrings the implementation of policies, programs, and research on sustainable development.

Satellite imagery, which is rapidly becoming more available at higher resolutions, offers a potential solution to data scarcity and unreliability. In our paper, we compared the availability and frequency of imagery from 200 random sample sites—100 in the United States and the European Union (EU) and 100 in Africa—between 2010 and 2019. We found a substantial increase in both high and moderate-to-low resolution captures over the past decade. Some sites that were once imaged a couple of times a year are now captured multiple times a week. What's more, the quality of this imagery is striking: Even at moderate resolution, sensors are capturing highly localized human activity, including infrastructure growth and crop productivity.

Stockpiles of images, however, are not helpful on their own. Machine learning models, functioning as decoders, amplify images' value. ML models can map raw inputs onto sustainable development outputs, transforming satellite images into predictive indicators. Researchers can harness flexible ML models, which learn and improve through training, to draw inferences about, for example, a village's economic productivity from the usage and distribution of the village's lights at night, or from the counts of objects present in daytime imagery. It was our goal to map and better understand these trade-offs and their policy implications.

Each of these satellites offers unprecedented access to imagery to help measure sustainable development outcomes, such as eliminating hunger, promoting health and well-being, and building sustainable communities.

Research Outcomes

Our survey of existing ML models for analyzing satellite imagery had four main findings: 1) Satellite images' utility in predicting development outcomes appears to be improving—and is already strong as is. 2) One of the largest constraints to using ML for this purpose is not satellite imagery itself but the training data used to construct the model. 3) Satellite imagery analysis is more likely to augment existing efforts, like collecting data on the ground, than to replace those methods altogether. 4) There is relatively limited adoption of satellite imagery analysis in many sustainable development domains, with the primary exceptions of population and agricultural measurement.

We also identified three policy benefits and three limitations for consideration by policymakers working on sustainable development, as well as researchers, civil society organizations, and other individuals and entities trying to leverage ML and satellite imagery to promote development.

Policy Benefit 1: Accurate estimates mean better interventions

Estimates of many sustainable development factors—such as population size, crop growth, and economic activity—dictate aid packages and targeted interventions. Given the complexities of measuring these variables, however, miscalculation and inaccuracy abound compromise the design and execution of good policy. Because ML models that harness satellite imagery may provide more spatial and temporal coverage of areas, they may equip policymakers with additional data points and more accurate information. For example, traditional censuses, particularly in lower-resourced countries, can be infrequent—but researchers can use satellite-based estimates of buildings, nighttime lights, and other markers to provide policymakers with more accurate and timely approximations of local population size.

Because ML models that harness satellite imagery may provide more spatial and temporal coverage of areas, they may equip policymakers with additional data points and more accurate information.

Policy Benefit 2: Machine learning models are fairly accurate in capturing complex variables

Many sustainable development variables are inherently difficult to measure. Consider the size and structure of informal settlements. In lower-resourced countries, an estimated 30 percent of the urban population lives in informal settlements, making them a key focus area for sustainable development efforts. Yet their diverse structures complicate attempts to map their density and true size. We examined 20 studies that attempted to measure the location of informal settlements from satellite imagery, and we found that ML model accuracy (as measured against ground data) surpassed 80 percent for all but three. Similarly, we analyzed 11 studies using ML on satellite imagery to predict economic livelihood and found high overall performance across many geographies, with satellite-based models explaining more than 70 percent of the variation in the ground-measured wealth at the local level. In short, the models were fairly accurate—and can work with cheap, widely available, and continually updated satellite images.

Policy Benefit 3: Scarce data and unreliable data sources need not be obstacles

Adding to the problem of data plagued by miscalculation and inaccuracy, some sustainable development contexts have data that is either scarce or entirely inaccessible. In some cases, researchers are forced to rely on self-reported data rather than empirically collected information. For instance, over 2.5 billion individuals living in poverty—over half of the global poor—farm small plots of land to survive. But our analysis suggests that assessments based on farmers' self-reported data are less accurate than imagery models trained on unbiased ground data. As such, ML models may enable policy assessments of sustainable development when ground data is noisy—where data is corrupted or distorted by faulty data collection instruments, human or computer errors, or other limitations. They could also help when there is a deficit of high-quality data: We found that high volumes of training data were insufficient for ML analysis of satellite images; instead, it took 30 to 50 *quality* training samples to sufficiently stabilize performance.

Policy Limitations

Foremost, satellite-based methods should complement, not replace, existing measurement methods. A number of factors drive this reasoning. First, certain sustainable development variables are simply ill-suited to measurement through satellite imagery. Gender discrimination and education attainment, for example, are difficult to capture and assess with satellite images. Relying exclusively on satellite-based estimates would leave policymakers empty-handed on these kinds of variables. Second, satellite-derived assessments raise questions that can affect their usability. For instance, if AI models cannot explain their decision processes and researchers don't fully understand them, this will undermine user trust and complicate policy issues around algorithmic bias and inaccuracy. Third, public- and private-sector partnerships may be critical to scaling and operationalizing satellite-based ML models. Absent these partnerships, efforts for sustainable development may become reliant on a technology that is not functioning optimally. Satellite-derived estimates should be incorporated into policymakers' toolboxes; after all, our analysis suggests they can approach or exceed the accuracy of traditional survey-based measures in certain domains, at a dramatically lower cost.

Policy Discussion

In looking to the future, satellite-based models face several core challenges—most notably, securing ample access to high-quality training data and overcoming concerns about noisy data. ML models are often only as good as the data on which they are trained, and satellite-based ML models are no different—a challenge in the sustainable development context where such data is often hard to obtain.

We identify a number of strategies to mitigate this challenge. First, researchers can train satellite models on synthetic data—data that is artificially created instead of generated by real-world events. This approach requires no ground data for training. In agricultural settings, a primary use case for satellite images in development, researchers have also found that synthetically trained models perform as well as, or better than, models calibrated from limited field data.

Another option is “transfer learning,” where models use large quantities of readily available data to learn a task similar to the target task of interest; once the model has mastered this similar task, it is “fine-tuned” to accomplish its intended task. The third approach is to explore training models on large amounts of unlabeled satellite imagery through unsupervised or semi-supervised learning, where models extract patterns from unlabeled data (sometimes combined with small amounts of labeled data) without substantial human input. As satellite imagery grows more abundant, these approaches may help mitigate training data challenges, improving ML models’ capabilities and accuracy. Policymakers can develop policies, funnel investments, and raise awareness of these issues in the sustainable development context.

Lastly, researchers and policymakers should focus on the problem of noisy training data that can diminish a model’s ability to learn, and if used to test models, it can make otherwise accurate models appear error-prone. To address this, policymakers and researchers should ensure that a small amount of very high-quality data is used for model testing. Researchers can also identify variables that are highly correlated to outcomes of interest (like rainfall and crop growth) and test their relationships to understand what happens when algorithms are trained on noisy data.

ML models are often only as good as the data on which they are trained, and satellite-based ML models are no different—a challenge in the sustainable development context where such data is often hard to obtain.

The original article, **“Using Satellite Imagery to Understand and Promote Sustainable Development,”** can be accessed at: https://www.science.org/doi/10.1126/science.abe8628.

Stanford University’s Institute on Human-Centered Artificial Intelligence (HAI), applies rigorous analysis and research to pressing policy questions on artificial intelligence. A pillar of HAI is to inform policymakers, industry leaders, and civil society by disseminating scholarship to a wide audience. HAI is a nonpartisan research institute, representing a range of voices. The views expressed in this policy brief reflect the views of the authors. For further information, please contact **HAI-Policy@stanford.edu.**

Marshall Burke is an associate professor in the department of earth system science and deputy director of the Center on Food Security and the Environment at Stanford University.

Anne Driscoll is a research data analyst at the Center on Food Security and the Environment at Stanford University.

David Lobell is a professor in the department of earth system science and the Gloria and Richard Kushel Director of the Center on Food Security and the Environment at Stanford University.

Stefano Ermon is an assistant professor in the department of computer science at Stanford University.

Stanford University
Human-Centered
Artificial Intelligence

Stanford HAI: Cordura Hall, 210 Panama Street, Stanford, CA 94305-4101
T 650.725.4537 **F** 650.123.4567 **E** HAI-Policy@stanford.edu **hai.stanford.edu**