

Stanford SOCIAL INNOVATION Review

Informing and inspiring leaders of social change

Technology

Can Machine Learning Double Your Social Impact?

A look at what it takes to successfully deploy machine learning tools for social good and the most exciting opportunities ahead.

By [Ben Brockman](#), [Andrew Fraker](#), [Jeff McManus](#) & [Neil Buddy Shah](#) | Feb. 20, 2019



Students in Jahazpur, Rajasthan, settle in for lunch during the Educate Girls development impact bond evaluation. (Photo by Ryan Fauber)

The next big thing in the social sector has officially arrived. Machine learning is now at the center of [international conferences](#), [\\$25 million dollar funding competitions](#), [fellowships at prestigious universities](#), and [Davos-launched initiatives](#). Yet amidst all of the hype, it can be difficult to understand which social sector problems machine learning is best positioned to solve, how organizations can practically use it to enhance their impact, and what kind of sector-wide investments can enable the ambitious use of it for social good in the future.

questions.

Our work at IDinsight, a nonprofit that uses data and evidence to help leaders in the social sector combat poverty, and the work of other organizations offer some insights into these

What Is Machine Learning and What Questions Can It Help Answer?

Machine learning uses data (usually a lot) and statistical algorithms to predict something unknown. In the private sector, for example, ride sharing apps use traffic data to predict customer wait times. Online streaming companies use customer history to predict which videos customers will want to watch next.

In the social sector, machine learning is particularly ripe for use in addressing two kinds of problems. The first is *prevention problems*. If an organization focused on conflict resolution can [predict where violent conflict is likely to breakout](#), for example, it can double-down on peacebuilding interventions. If a health NGO can predict where [disease is most likely to spread](#), it can prioritize distribution of public health aid. The second is *data-void problems*. The data governments and nonprofits use to target social programs is rarely granular, recent, or accurate enough to pinpoint the specific regions or communities that would benefit most, and collecting more-comprehensive data is often prohibitively expensive. As a result, many of the people who need a program the most don't receive it, and vice versa. If, however, an NGO fighting hunger in a rural state in Ethiopia knows which villages have the highest malnutrition rates, it can focus its outreach efforts in those communities, instead of oversaturating a different region that has fewer needs.



An empty classroom in Bhilwara district, Rajasthan, in a school supported by Educate Girls volunteers. (Photo by Kate Sturla)

Consider [Educate Girls](#), a nonprofit in India tackling gender and learning gaps in primary education. IDinsight previously ran a [randomized evaluation of Educate Girl's program](#) as a part of a [development impact bond](#), and found that the program has large, positive effects on both school enrollment and learning outcomes. Yet despite very ambitious plans to reach millions more children over the next five years, Educate Girls could not immediately expand its program to every one of India's 650,000 villages. The question, then, became how to prioritize villages to reach as many out-of-school girls as possible. Educate Girls' records showed that more than half of all out-of-school girls in its current program areas were concentrated in just 10 percent of

villages. But there were no up-to-date, comprehensive, reliable data sources that indicated where the most out-of-school girls were concentrated in other areas, and knocking on tens of millions of doors would be prohibitively costly. Hence, Educate Girls had a data-void problem that machine learning could help it overcome.

What Are the Practical Requirements for Using Machine Learning?

There are four requirements for those looking to use machine learning: good predictors, high-quality outcomes data, the capacity to act on predictions, and the ability to maintain the machine-learning algorithms.

1. Good predictors.

Suppose a nonprofit has a prediction problem. First, it will need a list of every person (or village or farm or school) it wants to make a prediction about, along with some information on them. This is called "predictor data." As the name suggests, predictor data (such as school attendance or student demographics) helps predict the "outcome" (such as student dropout risk) the organization cares about.

Generally, the more predictor data there is and the more strongly it correlates with the outcome variable, the more accurate the predictions will be. Social sector organizations will likely find the best predictor data in comprehensive, granular datasets, such as national censuses, satellite imagery, cellphone call records, or management information systems of a large nonprofit or government program. In our work with Educate Girls, for instance, we compiled 300 predictors from publicly available data sources in India, including the 2011 census and the annual census of school facilities.

Before charging forward, it's important that governments and organizations have predictor data for every person or place they want to make a prediction about. Some public data sources systematically undercount vulnerable groups, and excluding them from the prediction can lead to the denial of much-needed programs and services.

2. High-quality outcomes data.

To build and test a prediction model, organizations also need *some* outcome data for some individuals or places. Without it, it's impossible to know which predictors are relevant or to rigorously assess the accuracy of predictions.

When we started our machine-learning work with Educate Girls, its team had already collected the enrollment status for girls from more than 1 million households while rolling out its program to the first 8,000 program villages. We linked this data to the publicly available predictor data for the same group of villages, and used machine learning methods to uncover and pressure-test patterns that predicted enrollment. While a census of girls in 8,000 villages is a lot of data, it's a small fraction of the hundreds of thousands of villages where Educate Girls could work in the future. Even if it had not already collected this data as part of programming, it would have been a worthwhile investment to collect some outcome data so that it could build machine learning models to inform expansion strategy.

To be useful, this outcome data needs to:

- **Include enough data to uncover true patterns.** This can mean thousands or tens of thousands of outcome data points. Without enough data, a few cases may be falsely interpreted as a general pattern, or other important patterns may be overlooked, resulting in worse predictions. In the case of Educate Girls, the data from 8,000 villages was enough to both build and test models to figure out which socioeconomic and educational indicators mattered most.
- **Match the granularity of decision-making.** The outcome data must be at the same granularity as the relevant program decision. For Educate Girls, knowing the number of out-of-school girls in a couple of states would have provided little guidance for making village-level expansion decisions. Making village-level decisions required that we use of village-level predictor data and village-level outcome data.
- **Link predictor and outcome data.** To build a prediction model, organizations must be able to link each record in the outcome dataset to the same record in the predictor dataset. For instance, we needed to link each of the 8,000 villages in Educate Girls' dataset of out-of-school girls to the correct village in the public datasets we used. This mundane process can be particularly challenging when working across multiple data sources that use different codes or names to denote the same villages or people, and can sometimes be the most time-consuming step of a machine learning project.
- **Represent the target population.** Most importantly, the outcome dataset must reflect the diversity of the population the project is focused on. For instance, if we had built our prediction model for Educate Girls using only enrollment data from urban areas, we would likely have under-predicted the number of out-of-school girls in remote rural areas, where variables like distance to school matter more. Because machine learning predictions can have tremendous impacts on the lives of vulnerable people, we must put systems in place to measure and mitigate bias. Open source tools such as the [Aequitas project, at University of Chicago](#) can help measure bias risks and should be a standard part of every toolkit.

3. The capacity to act on predictions.

Making accurate predictions is only half the battle. To truly drive social impact with machine learning, social sector leaders must be willing and able to change how their organization operates based on predictions. Governments and nonprofits are often accustomed to one-size-fits-all programs. Shifting to targeted approaches means it must be legally and politically feasible to prioritize action for some people or communities over others, based on need or estimated risk.


This is often easier for nonprofits. For example, Educate Girls can decide to concentrate its resources in clusters of villages where it expects there are a large number of out-of-school girls, rather than work in all villages within a given administrative district. By contrast, governments interested in using machine learning may find that moving from a universal approach to a targeted one is more **politically difficult** and in some cases impossible.

4. The ability to maintain the machine.

Like any machine, prediction machines need maintenance. For example, since the factors that predict school enrollment in one state in India may be different from the factors in another state, Educate Girls must update its prediction algorithm each time it expands to a new geography.

Yet, contrary to popular belief, most machine-learning algorithms do not get smarter over time without human help. Instead, they require a human data scientist to load fresh training data, re-run the algorithm, adjust algorithm parameters based on the results, and deploy new models. Organizations that use machine learning need to set up long-term partnerships or build internal capacity to maintain high performance over time.

With these four requirements in hand, the next question is: How do you actually do machine learning? It's a complicated process involving merging datasets, testing different algorithms, and using statistical methods to fine-tune and select the best model. This work is best done by an expert. The good news is that more and more organizations—including [DataKind](#), [ML for Social Good at Carnegie Mellon](#), and [Data Science for Social Good Initiative at University of Chicago](#), and us at IDinsight—are set up to help others with their machine learning problems.



In this map, larger circles denote villages with more predicted out-of-school girls, small dots denote villages with few out-of-school girls (not recommended for targeting), and polygons show optimized clusters of villages that a single field team could viably cover.

Using the data and process described above, we generated village-level predictions that will allow Educate Girls to reach between 50 and 100 percent more out-of-school girls for the same budget, depending on the scale and geographic diversity of their expansion. We validated the accuracy of this model by comparing our predictions to Educate Girls data collected after we made our predictions for a given area. This past fall, Educate Girls used these predictions to decide which 1,800 villages to expand to in 2019, based on district, block, and village-level predictions.

Importantly, we delivered the predictions to Educate Girls in a way that fits its existing operational model, in geographically compact clusters of villages. The figure below shows the district-level predictions for Educate Girls' next expansion region, as well as the recommended clusters of villages that maximize the number of out-of-school girls reached without disrupting Educate Girls' operational model. Over the next 5 years, we estimate that Educate Girls would be able to reach around 1,000,000 out-of-school girls with its current expected budget and previous approach. However, we expect that by using the predictions generated by machine learning, it will be able to reach around 600,000 additional girls for roughly the same cost.

What Kinds of Investment Does Machine Learning Need?

One big challenge that stands in the way of replicating this approach is that relatively few government and nonprofit programs have the data they need to do it. Agriculture in the developing world, for example, is rife with prediction problems. What if we could estimate the crop production for every smallholder farmer in Africa? Or forecast crop failure risk for every farmer in South Asia? Governments, agribusinesses, and nonprofits could offer a huge variety of services for poor farmers, including better insurance products; tailored advice; and targeted, in-person, farm-extension-worker support when too much or too little rain falls.

Much of this is [technologically possible](#) using modern satellite data and machine learning; the trick is having enough geo-tagged data on crop production to build prediction models in different places and for different crops. But collecting this data is expensive and difficult, leading most organizations and companies to jealousy guard the data they collect, and putting a damper on sector-wide innovation.

One model that could unlock machine learning innovation across the entire sector is a philanthropically funded, open-data set with labeled, globally representative agriculture data. The model could be patterned after the successful [ImageNet dataset](#), which has [sparked a revolution](#) in researchers and companies' ability to build computer vision models. A consortium of researchers could jointly survey a representative set of tens of thousands of farmers who grow the 5-10 most important crops in each region of Asia, Africa, Latin America, and the Middle East. Using a standardized survey, the group could collect data on the crop type, planting patterns, and crop yields, and geo-tag each farm boundary. The data could be de-identified and available to anyone who wants to use satellite imagery for global agriculture work—including governments, NGOs, the World Food Program, and agribusinesses—and could dramatically lower the cost of innovation to better serve poor farmers.

A few related efforts are already underway. [The Radiant Earth Foundation](#), for example, recently launched a platform to organize and host agricultural data from around the world called MLHub.Earth. But while it uses crowdsourced data from existing research projects, there is so far no collective pooling of resources to fund a broader data-gathering effort. IDinsight is meanwhile developing a platform called Data on Demand, which seeks to drastically reduce data collection cost and timelines to make collecting this type of data more feasible. But we cannot do it alone. For these efforts to be truly impactful, philanthropy and governments have an important role to play in funding the collection of accurate and geographically representative data.

Of course, this type of thinking need not be limited to agriculture; the same type of challenges and sector-wide data collaboration opportunities exist in global health, infrastructure, and poverty targeting. We are excited about the potential for machine learning in the

broader development sector, because there are many other prediction applications that will allow NGOs and governments to provide better healthcare, education, and other services. Where the data and operating environments allow it, machine learning applications can help organizations that face prevention and data-void problems drastically improve their efficiency today. And in the longer term—if the social sector answers the call to invest in better data collection in more places—we believe machine learning can live up to the hype and become an even more powerful tool to improve billions of lives.



Ben Brockman is a senior manager at **IDinsight** (@IDinsight) and leads the organization's machine learning work, including its partnership with **Educate Girls**.



Andrew Fraker is a co-founder of **IDinsight** and head of the organization's innovation team. The innovation team is currently focused on building better data monitoring systems, developing a machine learning practice, and radically increasing the efficiency of at scale data collection.



Jeff McManus is a senior economist at **IDinsight**, where he designs and implements evaluations of social sector programs in India and Africa.



Neil Buddy Shah is the CEO and co-founder of **IDinsight**.

If you like this article enough to print it, be sure to subscribe to SSIR!

Copyright © 2020 Stanford University. Designed by Arsenal, developed by Hop Studios