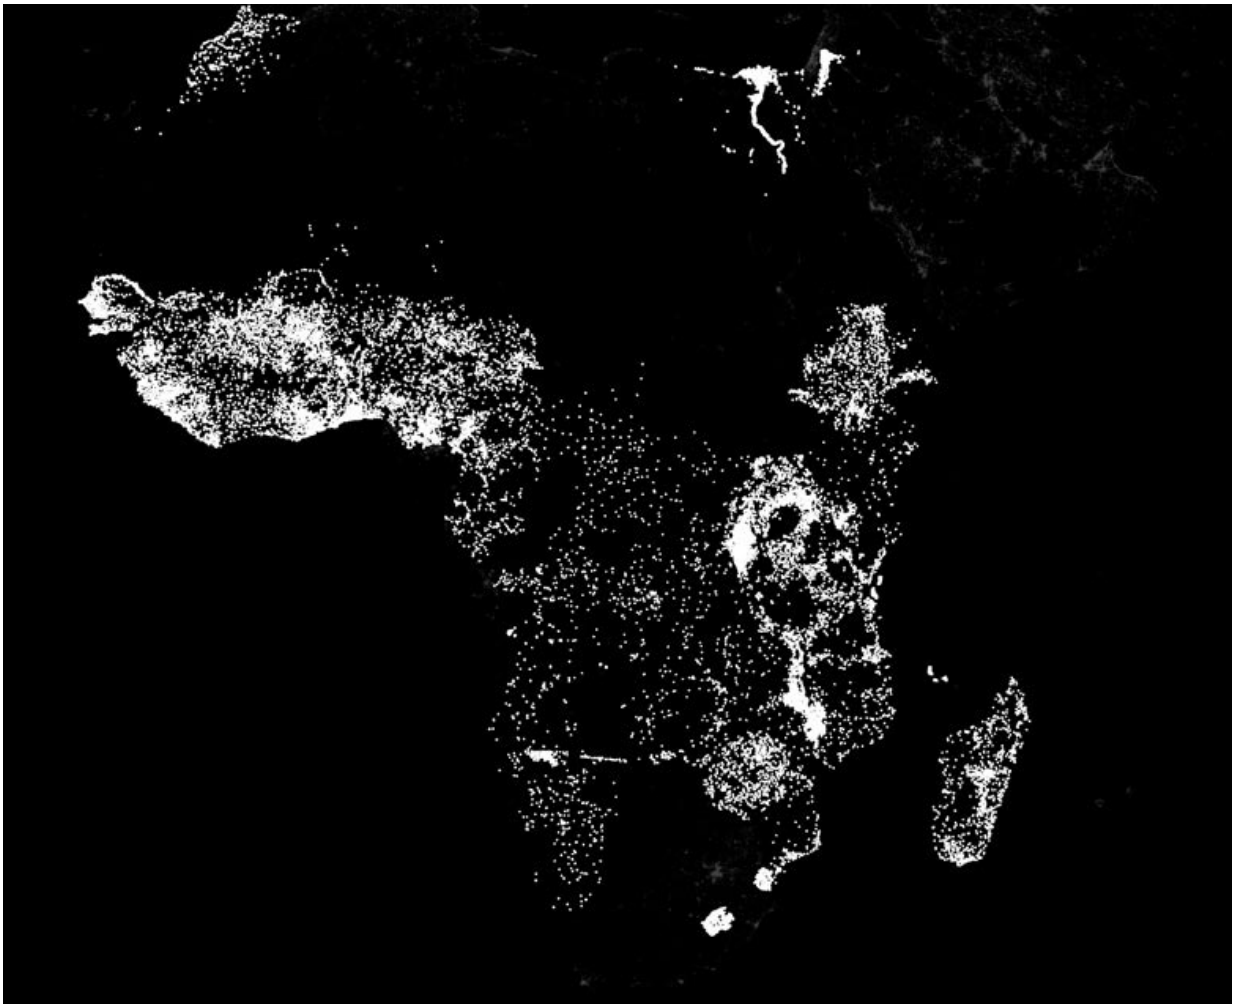# Researchers use dark of night and machine learning to shed light on global poverty

February 25 2016, by Glen Martin

Stanford researchers use machine learning to compare the nighttime lights in Africa – indicative of electricity and economic activity – with daytime satellite images of roads, urban areas, bodies of water and farmland.

One of the biggest challenges in fighting poverty is the lack of reliable information. In order to aid the poor, agencies need to map the dimensions of distressed areas and identify the absence or presence of infrastructure and services. But in many of the poorest areas of the world such information is rare.

"There are very few data sets telling us what we need to know," said Marshall Burke, an assistant professor in Stanford's Department of Earth System Science. "We have surveys of a limited number of households in some countries, but that's about it. And conducting new surveys in hard-to-reach corners of the world, such as parts of sub-Saharan Africa, can be extremely time-consuming and expensive."

A new poverty mapping-technique developed by interdisciplinary researchers at Stanford offers cause for hope. The technique is based on millions of high-resolution satellite images of likely poverty zones. To analyze these images, the researchers used machine learning, a discipline within the broader field of artificial intelligence. In machine learning, scientists provide a computational model with raw data and an objective – but do not directly program the system to solve the problem. Instead, the idea is to design an algorithm that learns how to solve the puzzle by combing through the data without direct human intervention.

The researchers began their poverty-mapping project knowing that nighttime lights provide an excellent proxy for economic activity by revealing the presence of electricity and the creature comforts it represents. That was half of the raw data that their system needed.

"Basically, we provided the machine-learning system with daytime and nighttime satellite imagery and asked it to make predictions on poverty," said Stefano Ermon, assistant professor of computer science. "The system essentially learned how to solve the problem by comparing those two sets of images."

Learning to spot poverty

Burke, Ermon and fellow team members David Lobell, an associate professor of Earth system science, undergraduate computer science researcher Michael Xie and electrical engineering PhD student Neal Jean detailed their approach [in a paper](link) for the proceedings of the 30th AAAI Conference on Artificial Intelligence.

Their basic technique – directing a model to compare images to predict a specific value – is a variant of machine learning known as transfer learning. Ermon likens this to how the skills for driving a car are transferable to riding a motorcycle. In the case of poverty mapping, the model used daytime imagery to predict the distribution and intensity of nighttime lights – and hence relative prosperity.

It then "transferred" what it learned to the task of predicting poverty. It did this by constructing "filters" associated with different types of infrastructure that are useful in estimating poverty. The system did this time and again, making day-to-night comparisons and predictions and constantly reconciling its machine-devised analytical constructs with details it gleaned from the data.

"As the model learns, it picks up whatever it associates with increasing light in the nighttime images, compares that to daytime images of the same area, correlates its observations with data obtained from known field-surveyed areas and makes a judgment," Lobell said.

Those judgments were exceptionally accurate. "When we compared our model with predictions made using expensive field-collected data, we found the performance levels were very close," Ermon said.

Highly effective machine-learning models can be very complex: The model the team developed has more than 50 million tunable, data-

learned parameters. So although the researchers know what their mapping model is doing, they don't know exactly how it is doing it. "To a very real degree we only have an intuitive sense of what it is doing," said Lobell. "We can't say with certainty what associations it is making, or precisely why or how it is making them."

Next generation of surveying

Ultimately, the researchers believe, this model could supplant the expensive and time-consuming ground surveys currently used for poverty mapping.

"This offers an unbelievable opportunity for cheap, scalable and surprisingly accurate measurement of [poverty](#)," Burke said. "And the beauty with developing and working with these huge [data sets](#) is that the models should do a better and better job as they accumulate more and more information."

The availability of information is something of a limiting factor. Right now satellite coverage of impoverished areas is spotty. More imagery, acquired on a more consistent basis, would be needed to give their system the raw material to take the next step and predict whether locales are inching toward prosperity or getting further bogged down in misery.

But such data restraints could soon be lifted – or at least mitigated.

"There's a huge number of new high-resolution satellite images that are being taken right now that should be available in the next 18 months," Burke said. "That should help us predict in time as well as space. Also, there are several micro-sat companies that plan to provide images of the planet almost daily, and we're rapidly getting enough satellites up to do that. I don't think it will be too long before we're able to do cheap, scalable, highly accurate mapping in time as well as space."

Even as they consider what they might be able to do with more abundant satellite imagery, the Stanford researchers are contemplating what they could do with different raw data – say, mobile phone activity. Mobile phone networks have exploded across the developing world, said Burke, and he can envision ways to apply machine-learning systems to identify a wide variety of prosperity indicators.

"We won't know until we try," Lobell said. "The beauty of machine learning in general is that it's very useful at finding that one thing in a million that works. Machines are quite good at that."

  More information: Transfer Learning from Deep Features for Remote Sensing and Poverty Mapping. arxiv.org/pdf/1510.00098.pdf

Provided by Stanford University