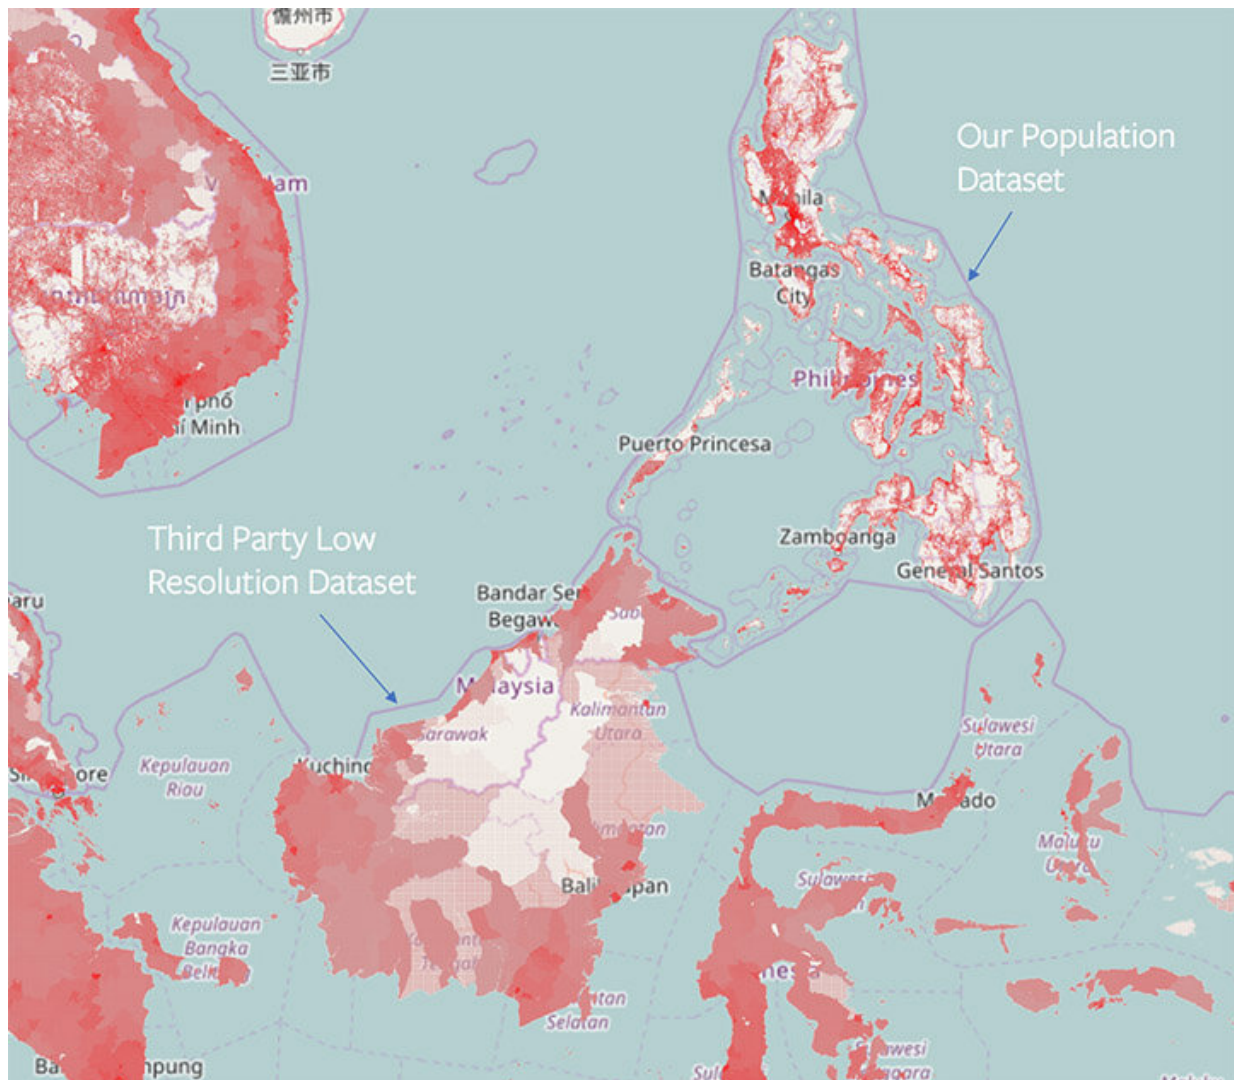# Facebook's mapping team aims to help aid workers know where help is needed

April 11 2019, by Nancy Cohen



Credit: Facebook

Artificial intelligence researchers and data scientists at Facebook have created population density maps. What's so special is that they are more accurate and with higher resolution than any of their predecessors. Derrick Bonak, Derrick Bonafilia, James Gill, Danil Kirsanov and Jason Sundram turned to the Facebook blog on Tuesday to write about their work.

These maps are designed for a significant purpose—mapping for humanitarian aid and development. Think aid workers in disease control and disaster preparedness.

"Building on our previous publication of similar high-resolution population maps for 22 countries, we're now releasing new maps of the majority of the African continent, and the project will eventually map nearly the whole world's population."

Ben Paynter in *Fast Company* laid out the terrific challenges in workers getting aid to the people who need it and how the maps can help.

"Vaccines, disease-battling insecticides, and new advancements in solar technology can all help people in developing countries stay healthier and have better-quality lives. That is, if you can locate them. In many places, smaller communities are spread out over vast and relatively uncharted terrain."

This was developed by the company's Boston-based World.AI team. The Facebook team had some evidence that their maps could achieve its purpose. They said that in Malawi, the Facebook maps were used to inform a measles and rubella campaign. The Red Cross was able to deploy trained local volunteers to specific areas in need.

The technology that serves their goal is a mixture of machine learning techniques, high-resolution satellite imagery, and population data. (They

wrote that the satellite maps in this project "were generated using commercially available satellite images from DigitalGlobe—the same type of imagery made available via publicly accessible mapping services.")

Their method involved mapping "hundreds of millions of structures" that are distributed across vast areas. They used that to extrapolate the local population density.

Another Facebook blog post further described the process. "For Africa alone, for example, the system examined 11.5 billion individual images to determine whether they contained a structure. Their approach found approximately 110 million structure locations in just a few days."

To explain just how AI was unleashed, a post by Karen Hao in "The Download" of *MIT Technology Review* walked readers through.

"First, a team at Facebook's World.AI group had to train a neural network to recognize whether a patch of land within a satellite image contained a home. To do this, the researchers created a training data set by overlaying more than 100 million crowdsourced coordinates of homes from OpenStreetMap onto satellite images. They also used old-school computer vision tricks to verify that the images labeled without homes didn't contain any telltale polygon-shaped objects."

Credit: Facebook

Satellite images of the African continent were divvied up into 100-foot-by-100-foot areas. They used the neural network to create an accurate, high-resolution population density map.

The blogs are worth reading, for sure, not only to find out what they achieved but to appreciate the great challenge that besets population mapping. It's a challenge, as they said, suited for deep learning.

"A country's census shows how many people live in a particular census tract, but it doesn't indicate where people live in these tracts—and sometimes the tracts encompass hundreds of square miles. Africa alone has 1.2 billion people across nearly 16 million square miles; its largest census tract is 150,000 square miles with 55,000 people. If researchers knew where the houses or other buildings were located in these tracts, they could create extremely accurate density maps by allocating the population proportionally to each one."

So, given a "massive imbalance," what did they do?

"Most of the world's land does not contain a building, so we have often dealt with negative-to-positive class imbalances of 100,000-to-1. We used a preprocessing step using classical computer vision techniques with near-perfect recall (at the cost of low precision) to discard most areas that did not contain a building. This left us with candidate ~30x30-meter (64x64-pixel) patches of satellite imagery."

(A caption in the report explains that "Our pipeline first sets aside locations that couldn't contain a building. Then the neural net ranks each remaining location according to the likelihood that it does contain a building.")

They moved on to the next challenge, they wrote, and that was classifying which patches contained a building. "While greatly reduced by the preprocessing, the ratio of empty squares to those with buildings was still 10-to-1 or even 1,000-to-1. This created an imbalanced binary classification problem, and we therefore evaluated our results using the F1 score, which is the harmonic mean of the precision and recall."

"The unprecedented resolution, scale, and accuracy of our newest offerings should continue to aid humanitarian relief and development efforts around the world."

The team has made the data set available for [download](#).

What's next: They plan on releasing high-resolution population maps of more places in the coming months. The project aims to keep adding new continents and countries.

"The Download" meanwhile referred to something called "Deep Geography" and its roundup suggests a general scientific intertest in extracting information from satellite imagery using deep learning. One of the examples in the post was of Microsoft, which last year "trained a deep-learning [model](#) to build a comprehensive data set of all the building footprints in the US."

  **More information:** [ai.facebook.com/blog/mapping-t … supervised-learning/](#)

© 2019 Science X Network