# Here’s an open-source tool you can use to search satellite photos for pretty much anything

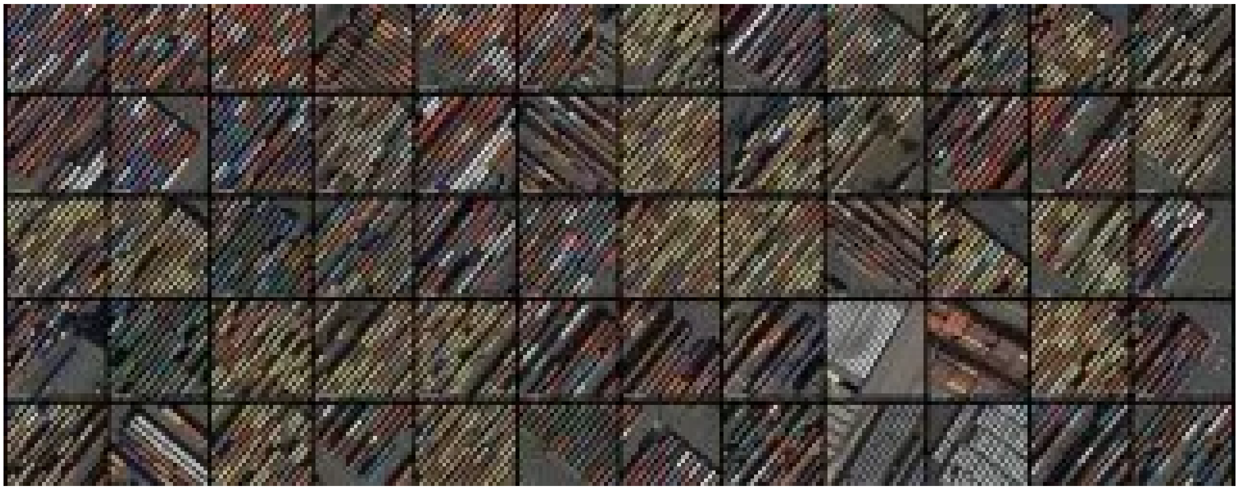

August 25, 2016

Golan Levin designed Terrapattern as a piece of contemporary art. Pick a satellite image, find others like it, tile them into a pattern. It was created to inspire others to make better use of satellite images, but was never itself meant for practical purposes.

But in just a few months of existence, it’s become a first iteration of what could be an essential digital weapon for humanitarian agencies, environmentalists, and civic activists.

Billed as the first open-source tool to perform “similar-image searches” for satellite photos, Terrapattern works intuitively. Browse a Google Earth-like map, click on something you’re interested in, and Terrapattern will return similar images.

My first test was a baseball diamond. I zoomed straight into Yankee Stadium in the Bronx, and clicked on home plate. Then I was looking at dozens of fields throughout the New York metro area.

But Levin, a Carnegie Mellon University professor, artist, and engineer, now thinks the best use of Terrapattern is to find more hidden features—things “that might be of interest to journalists, citizen scientists or NGOs.”

For example, he offers, we could use it to uncover heretofore hidden logging roads in the Amazon. These roads are just 10 or 12 feet (3-4 meters) across, but with Terrapattern’s high resolution—at one foot per pixel, it’s working with the most detailed satellite images available today—the tool could spot the thin lines that are harbingers of devastating deforestation to come.

It’s an ambitious vision, but not grandiose. Already, satellite images have become a key tool in managing displaced population camps. But though satellite imagery from private companies is now plentiful, finding patterns in it is impossible without heavy-duty computer algorithms to analyze it.

In recent years, a number of analytics companies have sprung up to provide these services. Orbital Insight, for example, uses satellite images and artificial intelligence to extract information that it can sell to Wall Street; for instance, it may look at how busy the parking lots are at big-box stores to predict the stores’ quarterly earnings.

Terrapattern is the first open-access tool for doing the same thing. “I wanted an Orbital Insight for the rest of us,” says Levin. “We’re a bunch of second-rate hackers who made this for $35,000. The purpose was to make it available for the public.”

Terrapattern is easy to use, but the underlying software is about as sophisticated as it gets.

Basic machine learning is fairly straightforward. You can manually train a computer, to, say, look through 50,000 pictures of 27-year-old men, and recognize all 15 that are of your cousin Tim. You simply give it a bunch of images, labelled “yes” (Tim) or “no” (not Tim). After seeing enough of those, the computer can look at any other image and determine the “yes/no” on its own. This is called “supervised training,” says Tom Mitchell, chair of CMU’s Machine Learning Department, who partnered on the project.

But while that method trains a computer to find all the Tims—or all the baseball fields, shipping container yards, solar panels, or whatever other category you choose—it doesn’t train it to find anything it is asked for.

That requires “unsupervised training.” For Terrapattern, says Mitchell, the computer was given about half a million satellite image tiles—each itself made up of hundreds of thousands of pixels—and was then forced to figure out a way to recreate each tile using a shorter string of code than that tile initially contained. To do that, the algorithm would break down each tile in the satellite image, and figure out all sorts of information (color, shape and contrast) and figure out patterns.

For example, a swimming pool might have hundreds or thousands of pixels of blue. Terrapattern learns that whenever it sees that many blue pixels, surrounded by a border of a different color, it doesn’t need to keep comparing images pixel by pixel. It can take the pattern as a whole, and then quickly hunt down every match in the database.

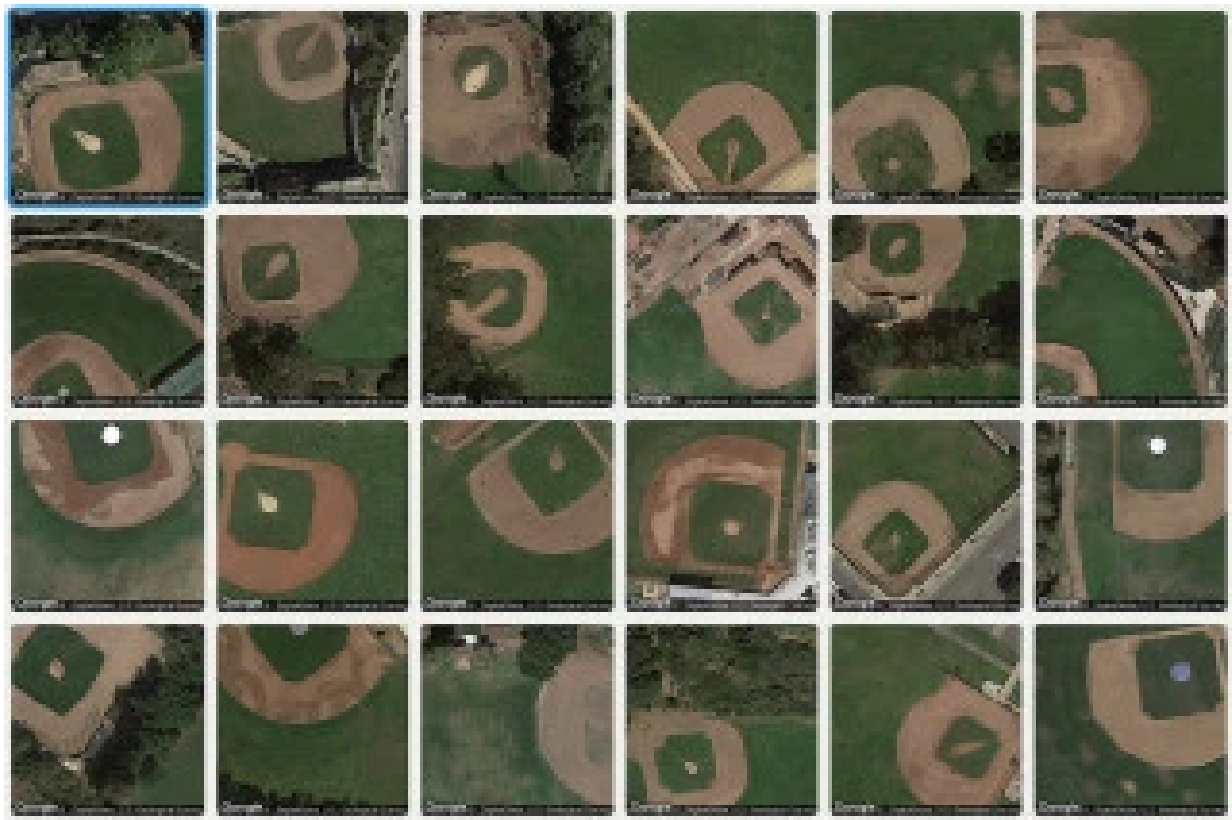

The software underlying Terrapattern can learn to recognize a baseball diamond as a whole unit—as opposed to searching every pixel to find matches. These images are from Pittsburgh. (Terrapattern)

Training a system like this can take weeks, says Mitchell, but once done, it’s highly efficient: Terrapattern can sift through all the images in its memory bank almost instantly.

But for Terrapattern to really be useful, it will have to expand significantly. Right now, it covers seven cities: Austin, Berlin, Detroit, Miami, New York, Pittsburgh and San Francisco. “I specifically did not immediately go to a conflict zone or a recent disaster area,” says Levin. “I’m really not qualified to do that. I would be stumbling around like a bozo.” To determine the right areas to map, he said, “I would want to carefully collaborate with a humanitarian organization with a reputation for acting neutral.”

However, Terrapattern is growing fast—it had just four cities when it first went public in May. “Now that we’ve gotten some people’s attention, we’re ready to talk about scaling up and working with new partners,” says Levin, to provide more satellite images and computing power.

For now, though, Levin and his team—developer David Newbury, media artist Kyle McDonald, and CMU students Irene Alvarado, Aman Tiwari and Manzil—mostly want their alpha launch to spark creativity among its users.

“I made this as a question for the public,” says Levin. “‘Here’s a thing, how does it make you think?’” Those answers could drive the technology forward. “Our ultimate goal was to make something that could portend the tools of the future.”

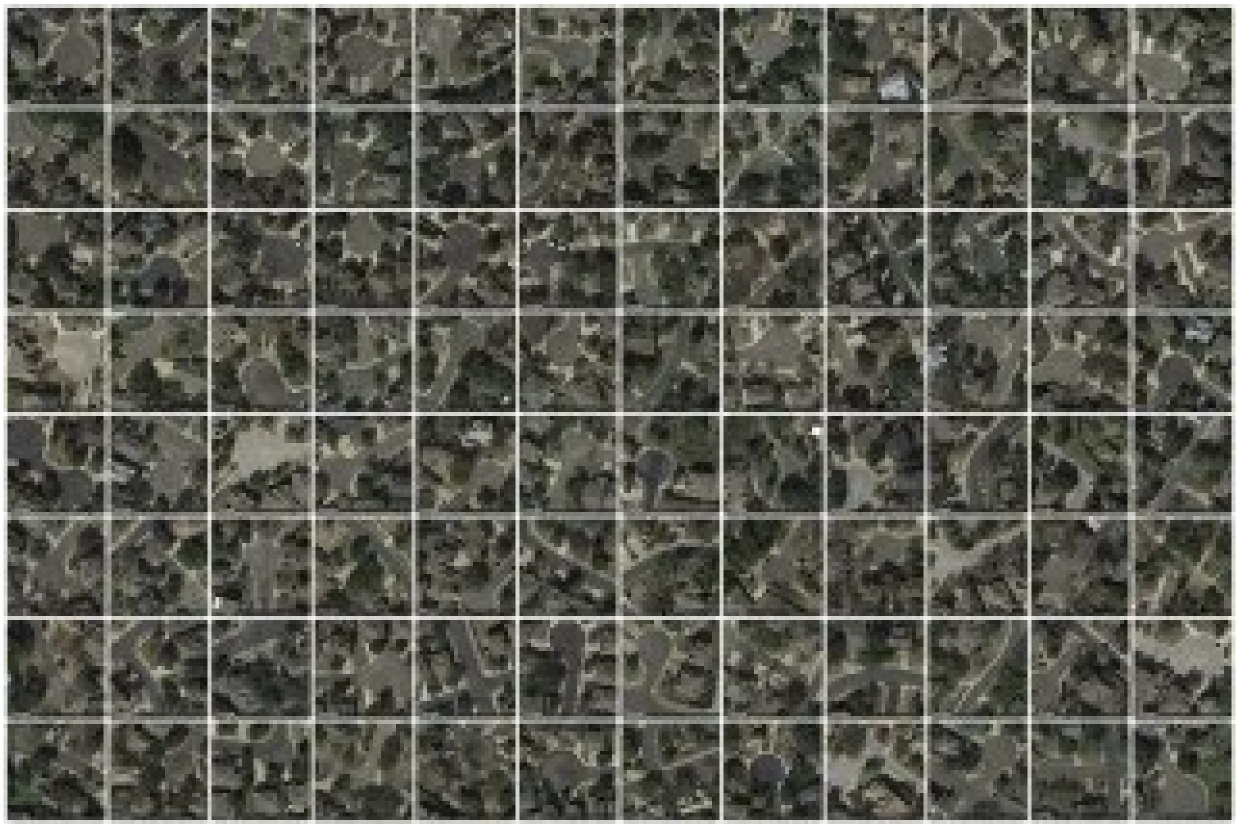

Cul-de-sacs in Pittsburgh (Terrapattern)

Airplanes in Miami (Terrapattern)

Playgrounds with sand surfaces in Berlin (Terrapattern)