MEET TERRAPATTERN, GOOGLE EARTH'S MISSING SEARCH ENGINE

By Nicola Twilley, MAY 25, 2016

VIEW FULL SCREEN

The wakes of boats in the rivers around New York City.

"Why don't you click on the tennis court?" Golan Levin, an associate professor of art at Carnegie Mellon University, suggested. I was looking at a satellite image of the school's campus in Pittsburgh, embedded in the home page of Levin's latest online project, Terrapattern. "What you should immediately see are all the most tennis-court-ish patches of Allegheny County," he said. I clicked. With gratifying speed, the right-hand side of my screen filled with dozens and dozens of tennis courts—solo or in pairs or in clusters of six, white on green, purple on green, green on red. A confusingly painted parking lot ended up in the mix, too. "Now try the football field," Levin said.

Terrapattern is the first open-access search tool for satellite imagery. Choose something that catches your eye—a dish farm, a gravel pit, a traffic circle—and

Terrapattern will find other things that look like dishes, pits, or circles, pinpoint them on a map, and serve them up as a downloadable set of G.P.S. coördinates. (One of Levin's friends, a skateboarder, is already using the site to locate empty swimming pools.) Terrapattern was designed by a small team, including Levin, several of his students, the artist Kyle McDonald, and the developer David Newbury. It is intended as a prototype. "Our budget, in terms of the computing power we can afford, makes about twenty-five hundred square miles of the American landscape searchable," Levin said. For the moment, Terrapattern is limited to Pittsburgh, San Francisco, and New York. According to Newbury, it seems to have a particular gift for finding marinas and bridges.

National militaries and intelligence agencies have long had the capacity to read the global landscape from above. Since the beginning of the Cold War, the United States government has kept a close eye on naval deployments, missile movements, and meteorological phenomena from space. In the nineteen-seventies, NASA set up gigantic mainframe computers to help make sense of the data, using basic signal-matching algorithms. More recently, corporations have begun to extract their own information from satellite imagery, developing machine-learning models that can monitor, say, the number of cars in retail parking lots across the United States, or the level of oil in floating-lid storage tanks (based on the depth of shadow on their interiors). With Terrapattern, Levin and his colleagues hope to make this sort of information more accessible to the general public. He cited as an influence the Harvard Humanitarian Initiative's Signal Program on Human Security and Technology, which documented war crimes in Sudan by using artificial intelligence to identify and tally the visual signature of burned huts. Similar tools have been developed by non-governmental organizations and academic researchers to locate illegal gold mines, track Antarctic penguin populations, and reveal lost Egyptian monuments.

Terrapattern works according to the same concept—A.I. applied to satellite imagery—but with a significant difference. Rather than designing a detector for specific features, like access roads or penguin guano, Levin and his team focussed

on making their service open-ended and customizable. To build it, they started with the type of A.I. that is most commonly used for image-recognition tasks: a deep convolutional neural network. For such an algorithm to be effective, it needs to be trained on many examples. The neural net that reads your checks at the A.T.M., for example, learned to do so by being shown thousands of handwritten numerals alongside their typed-out iterations until it could reliably match certain configurations of curves and edges with a corresponding set of labels. At first, the Terrapattern team downloaded an open-source neural net that had been pre-trained on ImageNet, a database of more than a million images sorted into thousands of categories. But when they set the A.I. loose on satellite photos, McDonald told me, it identified a Quonset hut as toilet paper and a transformer station as a pirate ship.

To familiarize Terrapattern's net with landscape features rather than everyday objects, McDonald turned to OpenStreetMap, a kind of cartographer's Wikipedia, in which hundreds of streets, rivers, and parks have been painstakingly labelled by human volunteers. After just four days' practice, Terrapattern's A.I. had developed a powerful method for reading satellite imagery, separating it into small tiles and decomposing each one into information about shape, color, contrast, and texture, then reassembling it into meaning through layers of probability and comparison, all in a matter of seconds. The Terrapattern team has uploaded their model to an open-source "model zoo," making it the first A.I. trained on satellite imagery that is freely available to use and adapt.

For something put together by a small group of artists and coders on a modest grant of thirty-five thousand dollars, Terrapattern works remarkably well. Levin seemed surprised himself. "It actually does the thing we wanted it to do," he said. At the same time, he added, "we don't exactly know what it's good for." Levin hopes that Terrapattern will make it easier for activists and citizen scientists to research the land-use issues that they care about, but he's aware that a cynic might see the site as a solution in search of a problem—a toy rather than a tool. And someone more security-minded might well wonder about its potential for misuse.

"Big things like nuclear-power plants, they're already on maps," Levin said. But Terrapattern "does make it a lot easier to spot soft infrastructure—things like oil tanks."

Levin is an artist as well as an engineer, and ultimately it is as artwork that Terrapattern is most successful. Perhaps the same satellites that gave us precision-guided munitions and long-range weather forecasting will now inspire more personal questions and insights, reconfiguring our understanding of the landscape. In the meantime, the opportunity to see how Terrapattern responds to human queries reveals something else—the biases and shortcomings of artificially intelligent analysis. "We built Terrapattern to allow people to search the world according to how they see it," Levin said. "But it also puts them in dialogue with the landscape as seen by machines. And that's just a very uncanny experience."

Nicola Twilley is the author of the blog Edible Geography, a co–host of the Gastropod podcast, and a contributing writer at newyorker.com.

WATCH: Manu Prakash, a biophysicist at Stanford and the creator of the Foldscope, muses on the beauty of the microscopic world.