2016 • Press for “Terrapattern" (English)

Golan Levin

Press articles and critical reviews pertaining to the project, “Terrapattern":

002 Meyer, Robinson. “The Thrill of Terrapattern, a New Way to Search Satellite Imagery". *The Atlantic,* 5/27/2016.
007 Twilley, Nicola. “Meet Terrapattern, Google Earth's Missing Search Engine". *The New Yorker,* 5/25/2016.
011 Mandelbaum, Ryan F. “This website lets you find the hidden similarities in big cities". *PopSci.com / Popular Science,* 5/25/2016.
013 Stinson, Liz. “Terrapattern is Like a Search Engine for Satellite Imagery". *WIRED.com,* 5/27/2016.
016 Brownlee, John. “This Neural Network Reveals Your City's Secret Patterns". *Fast Company,* 6/1/2016.
023 Singleton, Micah. “Terrapattern is the first open-access visual search engine for satellite maps". *The Verge,* 5/27/2016.
025 Voon, Claire. “A Visual Search Engine for the Aerial Patterns of Cities". *Hyperallergic*, 7/6/2016.
033 Visnjic, Filip. “Terrapattern – Neural network visual search tool for satellite imagery". *CreativeApplications.net,* 5/25/2016.
035 Blakemore, Erin. “This New Satellite Project Helps People Find Patterns in City Spaces". *Smithsonian,* 6/9/2016.
037 Manaugh, Geoff. “Subverting Our New Space Overlords". *The Atlantic,* 6/9/2016.
044 Kottke, Jason. “Satellite imagery search engine". *Kottke.org,* 5/25/2016.
045 Anzilotti, Eillie. “Mapping the Hidden Patterns of Urban Life". Citylab, 6/7/2016.
048 Yau, Nathan. “Visual search tool for satellite imagery". *FlowingData*, 6/1/2016.
049 Dempsey, Caitlin. “Terrapattern: Search Engine for Satellite Imagery". *GIS Lounge,* 5/28/2016.
050 “Cool Tech! TerraPattern, open-source tool for discovering ‘patterns of interest’". *GISUser,* 5/31/2016.
051 Whitehead, Timothy. “Terrapattern, the search engine for imagery". *Google Earth Blog,* 5/27/2016.
053 Whitehead, Timothy. “Image recognition and Google Earth". *Google Earth Blog,* 7/28/2016.
055 Clarke, Keir. “Maps Mania: Searching for Map Patterns". *GoogleMapsMania*, 5/26/2016.
056 Toobin, Adam. “Kill Some Time With Terrapattern, Which Turns Satellite Images Into Art". *Inverse.com,* 5/27/2016.
061 Wolfson, Elijah. “There's now a tool you can use to scour satellite images for criminal activity and environmental devastation". *Quartz,* 8/25/2016.
066 Coldewey, Devin. “Terrapattern is reverse image search for maps, powered by a neural network". *TechCrunch,* 5/25/2016.
068 Reese, Hope. “Google Maps meets AI: Carnegie Mellon's Terrapattern can find and map every pool in New York City". *TechRepublic*, 6/8/2016.
070 Kohlstedt, Kurt. “Terrapattern: Satellite Image Search Engine Matches Similar Places". *Weburbanist*, 5/30/2016.
072 Brown, Eileen. “Terrapattern search engine finds patterns in the Google Earth landscape". *ZDNet*, 7/14/2016.

The Atlantic

The Thrill of Terrapattern, a New Way to Search Satellite Imagery

Finding empty pools for skateboarding has never been so easy.

Astronaut Scott Kelly snapped this photo of the San Francisco peninsula from the International Space Station in 2015.

NASA

ROBINSON MEYER

11:32 AM ET | TECHNOLOGY

TEXT SIZE

Sometimes, a new tool comes into the world that is so expansively, obviously useful that you can't do anything but sit back and think: Wow.

For me, at least, that's Terrapattern, a visual search engine for satellite

imagery, released this week by a team of artists and geographers at Carnegie Mellon University. It is Google's "reverse image-search" tool for maps, basically: Click on a spot you find interesting, and Terrapattern will show you other spots on the map like it.

This means you can find all the cul-de-sacs in Pittsburgh:

Terrapattern

Or all the football fields:

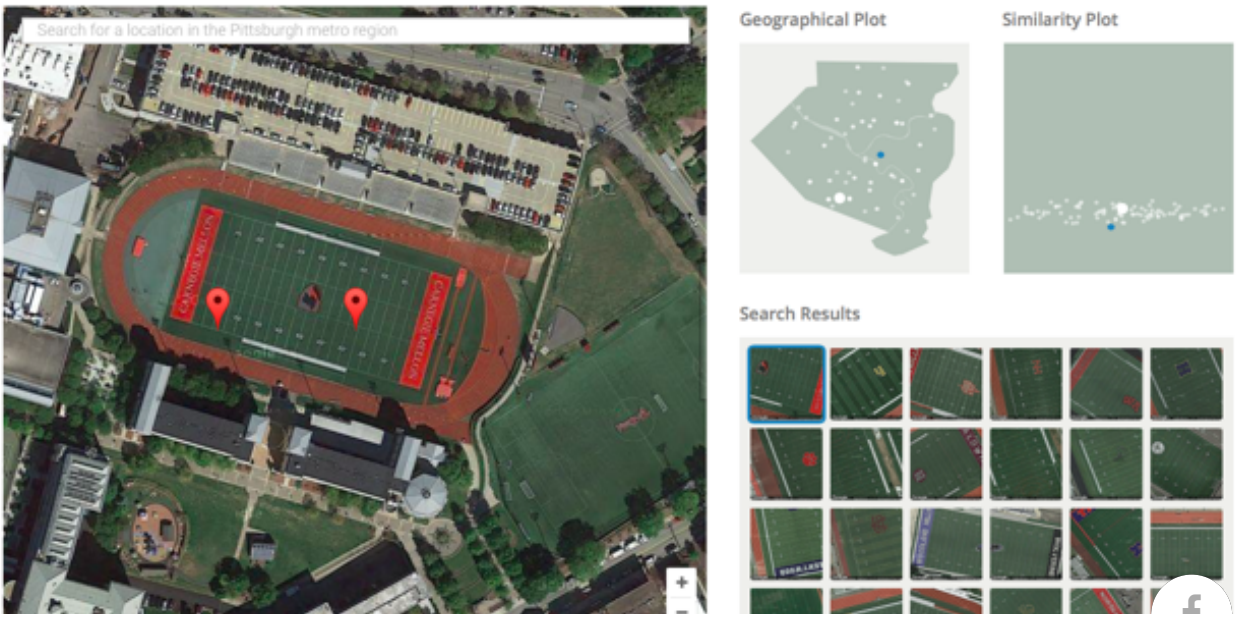

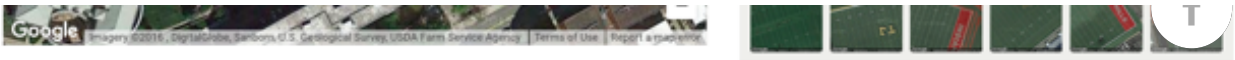

Credit

Or all the solar panels in New York City. Or all the bridges in the Bay Area.

"One of our friends is using it to find disused swimming pools for guerilla skateboarding," writes the Terrapattern team on their site. They built the tool for "discovering 'patterns of interest' in unlabeled satellite imagery," they say, a way to explore "the unmapped and the unmappable."

Right now, Terrapattern only covers four American cities: Pittsburgh, Detroit, San Francisco, and New York City. Terrapattern is so computing-hungry that it is effectively a proof of concept right now, at least for a team of artists working with less than $35,000. Each metro region takes about 10 gigabytes of RAM—not storage, but active memory.

That said, Terrapattern is relatively technically straightforward. It's constructed from a convolutional neural network and CoverTree, an algorithm that remembers some descriptions and allows the searches to happen quickly.

"It took the neural net about five days to train," said Golan Levin, an artist and engineer, in an email. Levin led the team that developed Terrapattern. "It was as easy as pointing the neural net at the map tiles," he said, though the map tiles were associated with place descriptions through OpenStreetMap.

"To be perfectly honest, most of our time was spent moving files around from place to place. Things get slow when you're moving hundreds of gigabytes," he said.

Levin and the rest of the Terrapattern team thinks that the tool is especially adept at finding "nonbuilding structures and other forms of otherwise unremarkable soft infrastructure that aren't usually indicated on maps," the

team writes. Think of those empty swimming pools—or a wind turbine, or an inflated sports dome.

Or, for that matter, a bridge damaged by an earthquake. Dale Kunce, who manages international digital mapping at the American Red Cross, told me that Terrapattern likely had applications in humanitarian situations. He imagined a situation where Terrapattern (or software like it) could process a satellite image of a disaster area and produce a "first pass" list of damaged structures. Then a human editor could come in and cull that list by hand.

"I am not usually impressed by stuff these days, but I was impressed by this," Kunce said. He told me it fits into the advance of applying digital maps to disaster relief over the past decade: moving from Google Maps, which included satellite imagery of most places; to OpenStreetMap, which let anyone make and use digital street data for free; to using OpenStreetMap in professional and humanitarian situations.

Now, software like Terrapattern and Facebook's population-estimating algorithm let volunteers apply their skills of discernment faster and at a greater scale. Estimating storm damage or population centers might be the next step in crowdsourced disaster-relief mapping. "The most powerful supercomputer in the world is not as good at recognizing things as the human brain. No one's build Watson for satellite imagery," he told me.

Levin said it was hard to know when Terrapattern might be ready for humanitarian deployment. "Currently our prototype only works in four cities. San Francisco is not currently suffering from a humanitarian crisis, in any reasonable sense of the word," he said.

But Terrapattern is also not likely the only technology of its kind. Right now, a number of startups—including Planet Labs and Terra Bella, which is owned by Google—are tossing dozens of small imaging satellites into orbit. They're

doing this because imagery-deciphering technology is expected to mature in the next few years, meaning they could algorithmically read the amount of oil in oil wells for financial firms. Descartes Labs also intends to apply machine learning to imagery in order to estimate agricultural yields.

But there hasn't been a working product quite like Terrapattern yet—or at least one available to the general public to play around with. It is, as they put it, geospatial-intelligence analysis "for the rest of humanity." Terrapattern is an experiment to see if "visual 'query-by-example' for satellite imagery might become a part of our everyday future," they write. "Remember, you saw it here first."

ABOUT THE AUTHOR

ROBINSON MEYER is an associate editor at *The Atlantic*, where he covers technology.

Twitter Facebook Email

ELEMENTS

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.

POPULAR SCIENCE

This Website Lets You Find The Hidden Similarities In Big Cities

Terrapattern is your source for shipwrecks, bus depots and abandoned warehouses

By RYAN F. MANDELBAUM Posted MAY 25, 2016

Let's say you're looking for a type of place, rather than a specific place, to visit. You could be a skateboarder looking for an empty pool to ride in, an urban explorer looking for abandoned warehouses to explore... or you could be a bored internet user looking for satellite photos of shipwrecks off the coast of New York City just because it's cool.

A new website, Terrapattern, can do that for you.

Start by clicking anywhere on the map in one of four cities: New York, Detroit, Pittsburgh and San Francisco. Terrapattern will analyze a square measuring 256x256 pixels on the map, which is equivalent to 50x50 meters in real life, and find every other place in the city that looks like that square (from an aerial view, that is). Click a bus depot in New York, and boom, all of the other bus depots will appear with a photo and their location on the map.

Terrapattern uses a Deep Convolutional Neural Network (DCNN) to perform the search -- which should be familiar if you remember Google's Deepdream artificial intelligence program, which turned ordinary photos into colorful blobs of dogs and eyes.

In this case, the researchers first needed to "train" the network by showing it hundreds of thousands of maps, which took about five days (thanks to the labeled satellite images from the crowdsourced

OpenStreetMap project). The trained network could then look at a new picture, notice all of the curves, textures and other features, and use them to locate anything that matches. It's kind of like the way your brain knows a new mass of bread, meat and cheese is a cheeseburger because you've already seen thousands of other cheeseburgers.

"I wanted a way we could open this technology to everyone: citizen scientists, journalists and artists, or just everyday people who want to understand the world in a better way," said Golan Levin, Terrapattern project director, in an interview with *Popular Science*. He collaborated with developer David Newbury and artist Kyle McDonald, alongside Carnegie Mellon students Irene Alvarado, Aman Tiwari and Manzil Zaheer. Levin hopes people will see something like Terrapattern and offer feedback, and that we'll eventually see Terrapattern-like integration on platforms like Google Maps.

Searching four cities alone already requires 10GB of RAM; searching the whole US would take almost 2000 times that amount of memory, which would take much more sophisticated software. Until then, you can explore Detroit, New York, Pittsburgh or San Francisco with Terrapattern here.

LIZ STINSON DESIGN 05.27.16 7:00 AM

TERRAPATTERN IS LIKE A SEARCH ENGINE FOR SATELLITE IMAGERY

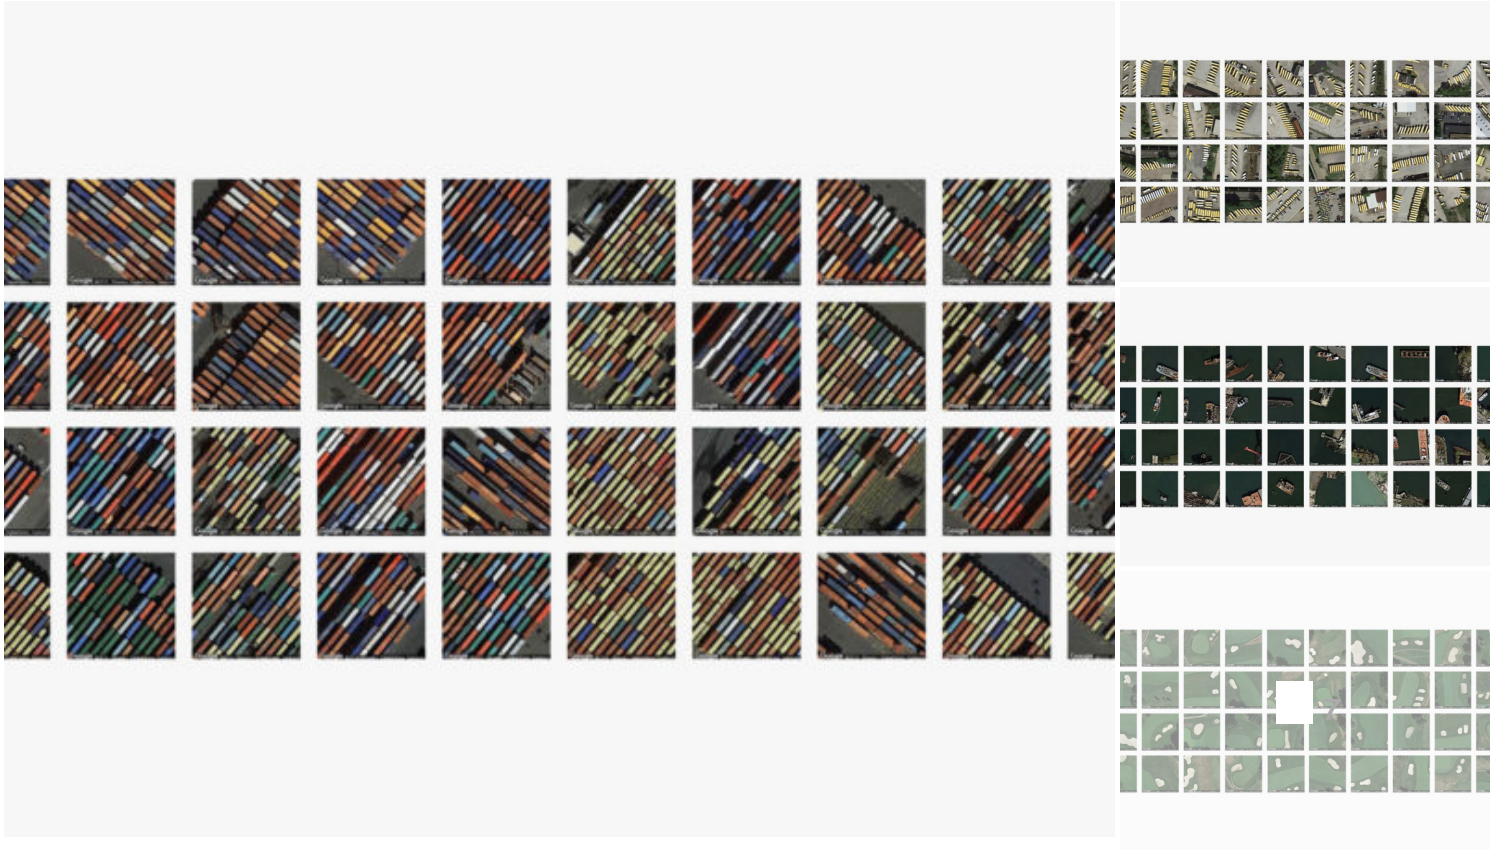

IN 2008, THROUGH something of a happy accident, a team of zoologists from the University of Duisburg-Essen in Germany discovered that grazing cows and deer tend to align their bodies with magnetic north. It was an odd thing to notice, particularly because the researchers had been perusing satellite imagery for something else entirely. But that's what happens when you look at something from 400 miles above the Earth's surface—change your perspective, and you'll change what you see.

When Golan Levin, an artist and professor at Carnegie Mellon University, heard about the cow discovery, he found it "to be simultaneously wonderful and very inspiring and totally useless." He was

also overcome, he says, by the desire to make similar discoveries. So he, along with artists and researchers Kyle McDonald, David Newbury, Irene Alvarado, Aman Tiwari, and Manzil Zaheer, created Terrapattern, a tool that lets people search for patterns otherwise hidden in troves of publicly available satellite imagery.

Terrapattern works a lot like Google's reverse image search: click on a map tile, and the tool searches for satellite imagery with similar visual attributes. For example, if you happen to click on a purple tennis court in San Francisco, Terrapattern will surface dozens of examples of similar tennis courts in the area and show their locations on a map. Same goes for crosswalks, pools, baseball diamonds—pretty much anything with distinctive visual characteristics. "Our tool is extremely good at finding various things that are un-mappable or unmapped or unconventionally unworthy of mapping," Levin says.

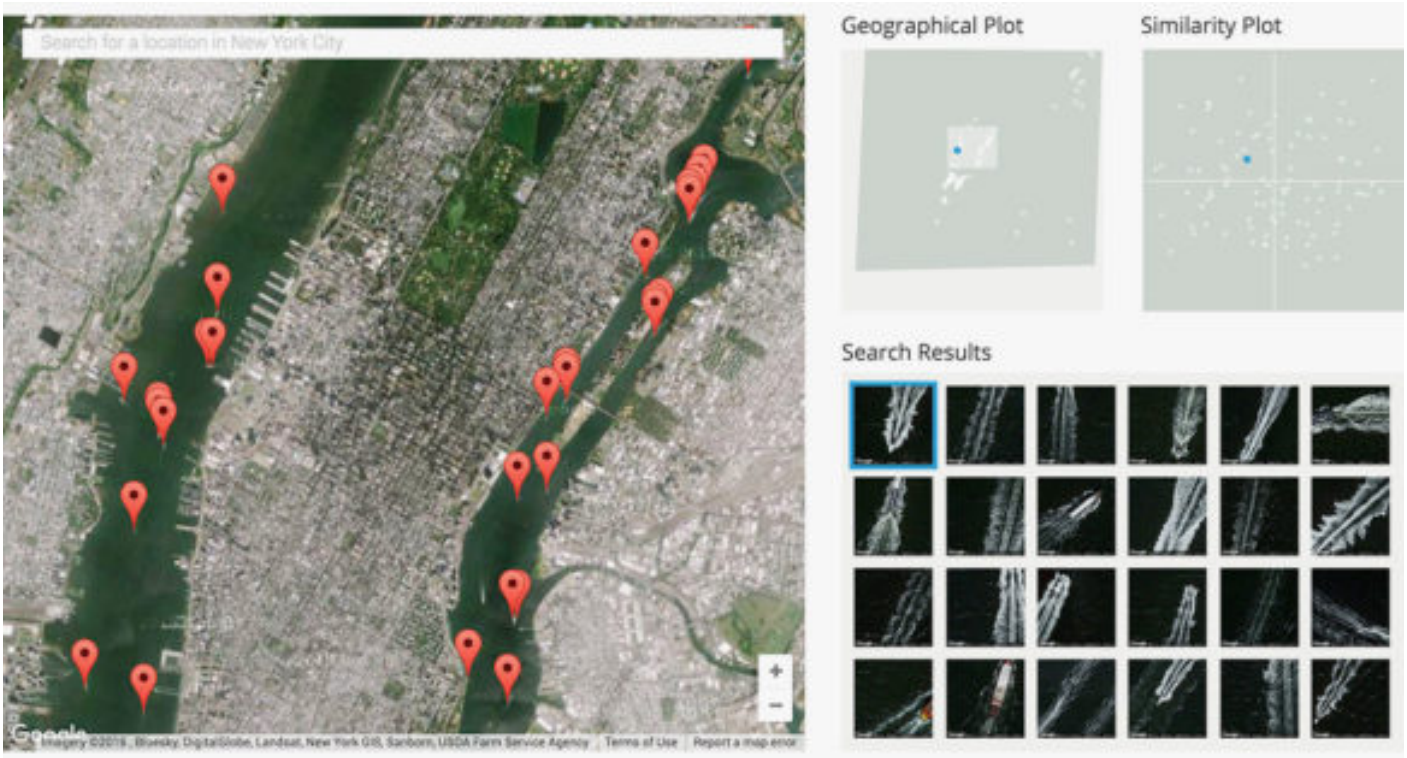

TERRAPATTERN

Like many recent computer vision projects, Terrapattern uses something called a convolutional neural network to make sense of what it's looking at. The system analyzes images in layers: Lower layers identify basic visual characteristics like edges and curves; while higher layers make sense of the shapes identified by layers below them. But unlike other convolutional neural networks, which might eventually tell you the thing it's looking at is a school bus, or a picture of a cat, Terrapattern's network doesn't assign a final classification to any of the images it analyzes. Instead, using half a million satellite images from OpenStreetMap, Levin and his team trained the network to recognize the descriptions of features in an image—roundness, smoothness, color, and so on. "When we want to discover places that look similar to your query, we just have to find places whose descriptions are similar to those of the tile you selected," Levin explains on the tool's website.

Using aerial observation to uncover unseen patterns isn't a new idea. Researchers have been using the technique for decades to track animal migration, monitor deforestation, and unearth stunning archaeological discoveries. Recently, companies like Orbital Insights have even combined satellite imagery and computer vision to collect (and sell) market-impacting intelligence. But Levin's vision is more egalitarian. He wants to create an Orbital Insights for the average person.

Of course, it remains to be seen just how the average person will use Terrapattern. Levin used it to discover graveyards of rotting boats along the coasts of New York City. He imagines journalists could use it to aid their investigations, and humanitarians to better coordinate relief efforts—though he concedes "there is a visual pleasure in using the project which I think goes beyond utility." The way he sees it, it's an artist's job to push the boundaries of new technologies, to investigate how they might impact society sometime in the future. It's not so much about providing obvious applications. "I give myself permission to explore things that may not have obvious use," he says.

3 MINUTE READ

This Neural Network Reveals Your City's Secret Patterns

Terrapattern is a new search engine for satellite imagery that could open the door to a real-time Google Earth.

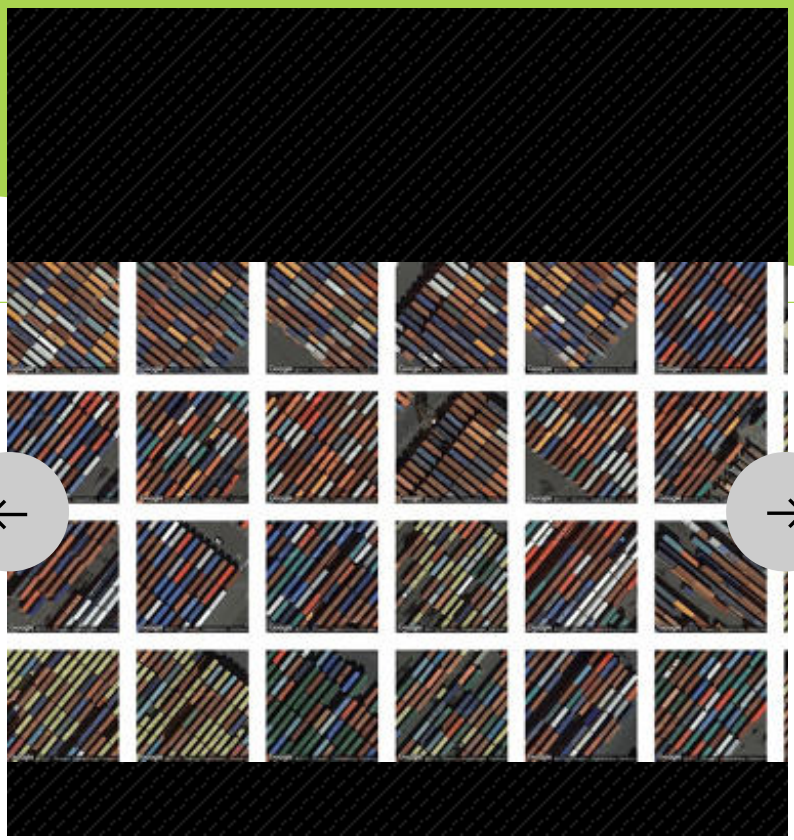

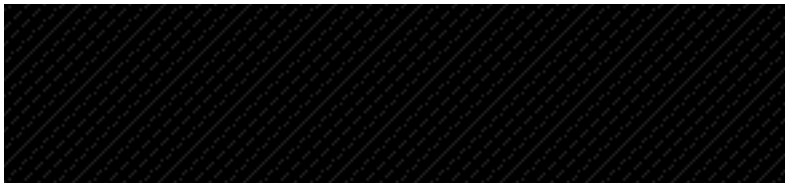

01 /07

JOHN BROWNLEE | 06.01.16 | 7:00 AM

What if you could search the topography of the real world as easily as you search the Internet, all in real time? Terrapattern is a new search engine that does just that. Created by a team led by data researcher Golan Levin, Terrapattern combines satellite data and machine learning to let you search four major American cities for interesting patterns and points of interest, whether that's all the Christmas tree farms in San Francisco or the fracking wells of Pittsburgh.

While it's a great time waster for anyone, Terrapattern could have a big impact on the workflows of journalists, government agencies, urban designers, and more, within the next few years.

Like the neural network that powers Google's Deep Dream AI, Terrapattern uses a deep convolutional neural network to group patterns in satellite map data by type. For example, it can tell just by looking at a satellite map that two tennis courts, or sports stadiums, are of the same type. All neural networks like this need to be "trained" to identify content using a library of data. In the case of Terrapattern, this training library was provided by OpenStreetMap, a collaborative map of the world (complete with a satellite component) that volunteers have been building out since 2004. Using OpenStreetMap's data, Levin was able to train his neural network to "look" at a tile of satellite imagery, and "see" what it shows.

You can use it to search for all the churches in Detroit, or all the topiary gardens in New York City.

Unlike Google Photos, you can't just search Terrapattern for all examples of, say, swimming pools or skate parks. You need to know where one example of what you're looking for is located first—so you have to browse by city, dragging your way through a satellite map and clicking on features that interest you. Once you do so, Terrapattern suddenly shows off its power: You can use it to search for all the churches in Detroit, or all the topiary gardens in New York City, and so on.

Terrapattern has a couple limitations right now. One, it only works in four areas: Pittsburgh, New York, San Francisco, and Detroit, cities chosen by the team because they either lived there, or had friends there. In theory, there's no reason Terrapattern couldn't be expanded to encompass the whole world, but the bottleneck is computer power. Powering a single city's search in Terrapattern requires 10 GB of active RAM; Levin estimates they would require

2,000 times more to store and serve a searchable model of the whole United States.

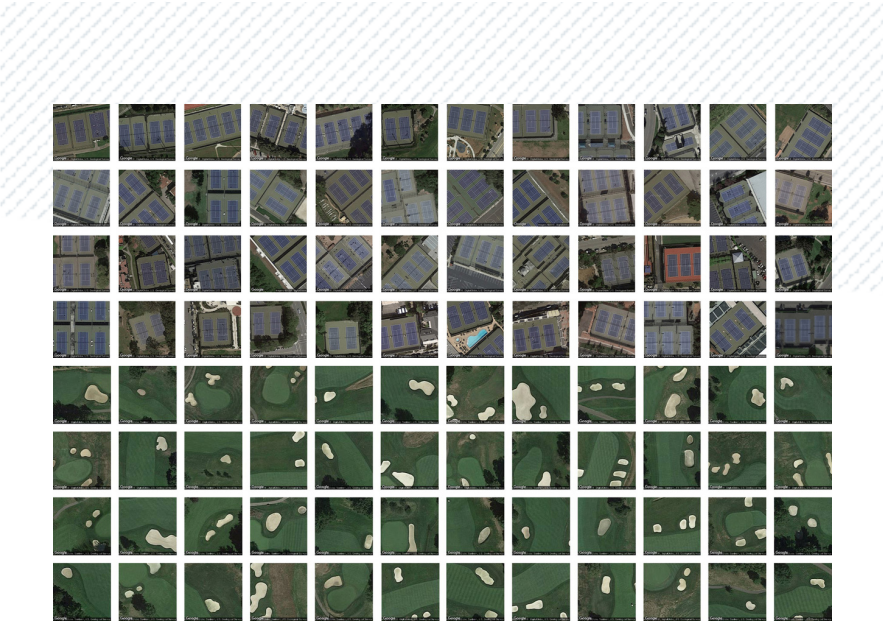

The second limitation has to do with the size of a search area. Terrapattern isn't great about searching for features smaller than 10 meters (around 33 feet), or larger than 100 meters (328 feet), but that isn't a hard limit. Future versions of Terrapattern past the alpha prototype stage could well enable larger or smaller scale searches.

According to Terrapattern's creators, there is every reason to believe that within the next few years, access to whole-earth satellite imagery that is updated on a daily basis will be widely available online. Compare that to today,

where the satellite imagery of your area in Google Maps might be weeks or months out of date. Once daily satellite imagery is widely available, Levin and his team believe that there will be a lot of interest in tools that can help make sense of it.

Levin hopes it will serve as a useful baseline for the planet-sized search engines to come.

Humanitarian organizations could use satellite data to identify mass migrations of refugees as they happen, and predict where they're going before they get there. Journalists could also use this data to identify patterns for stories, proving (for instance) there's been a 30% increase in solar panel installations in the last six months, a 20% decrease in forest fires from a year ago, or where armies were building up their forces. Skaters could find new zones to conquer by watching what swimming pools in their area have been allowed to go empty over the summer.

Terrapattern is an attempt to prototype the sort of tool that will be needed when

satellite imagery is updated regularly. It's a thought experiment that anticipates what we'll want from a real-time Google Earth when the time comes, and while it's not there yet—Terrapattern can't compare satellite imagery over time, nor can you use natural language to search, like you can in Google Photos—Levin hopes it will at least prove useful as a baseline for the planet-sized search engines to come.

You can try it for yourself here.

All Images: via Terrapattern

THE VERGE

Terrapattern is the first open-access visual search engine for satellite maps

Machine learning is everywhere

By Micah Singleton on May 27, 2016 12:19 pm

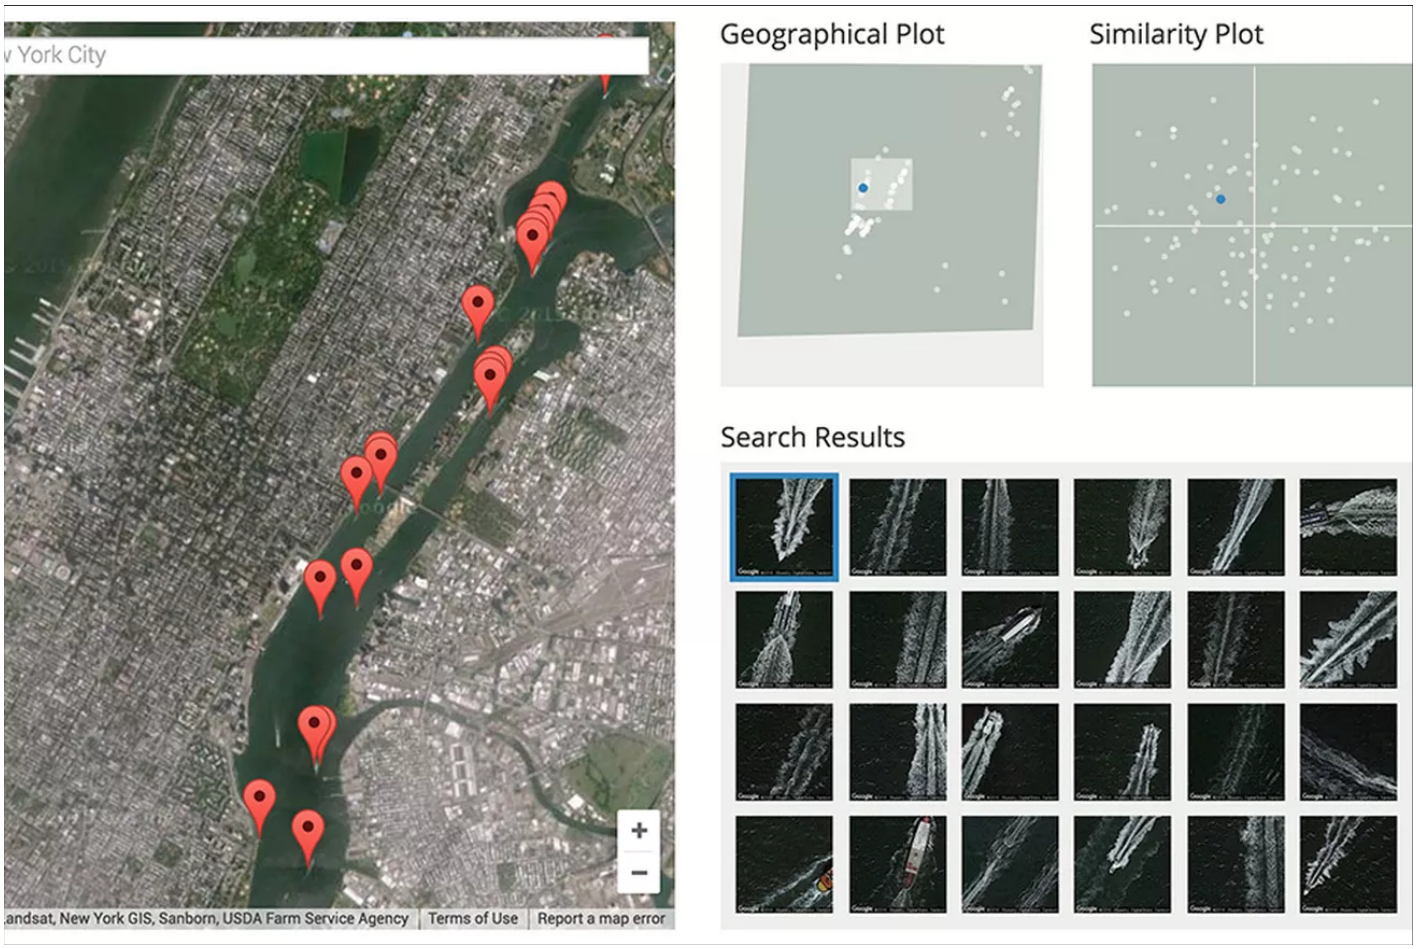

Terrapattern is a amazing visual search engine for satellite imagery, and somehow the first of its kind. It's very easy to use, click on a section of the map, and Terrapattern will show you all similar geographical features or landmarks in the area. A football field, bus station, outdoor pool — it doesn't matter, Terrapattern can pinpoint the related image and location with surprising accuracy.

The program was created by Golan Levin, David Newbury, and Kyle McDonald, with funding from the John S. and James L. Knight Foundation. Terrapattern is built on a Deep Convolutional Neural Network (DCNN), and has been trained to recognize geographical features within small squares in four cities — New York, San Fransisco, Pittsburgh, and Detroit.

In an interview with *Popular Science*, Levin said he and the team built Terrapattern to open up the visual mapping search to a larger audience (the US Military has had similar tech for years), and hopes that it would eventually make it into other mapping platforms like Google Maps. "I wanted a way we could open this technology to everyone: citizen scientists, journalists and artists, or just everyday people who want to understand the world in a better way," Levin said.

You can test out Terrapattern here and check out the code on GitHub.

HYPERALLERGIC
Sensitive to Art & its Discontents

A Visual Search Engine for the Aerial Patterns of Cities

by **Claire Voon** on July 6, 2016

Pools in Pittsburgh (all images screenshots by the author)

Thanks to a small team of artists and coders, you may now explore cities through patterns of infrastructure as captured in aerial photography. **Terrapattern**, developed at the Carnegie Mellon **Frank-Ratchye STUDIO for Creative Inquiry**, is the first open-access visual search tool for satellite imagery. It is currently available for Pittsburgh, San Francisco, New York City, Detroit, Austin, Miami, and Berlin. This means you may scan these cities' landscapes for common forms of your particular interest that are not conventionally labelled on a map: circular backyard pools or cul-de-sacs, perhaps, or even dilapidated nautical wrecks. All you have to do is find the tile of topography that intrigues you, and dozens of search results of similar views will arrive courtesy of machine learning algorithms trained to sift through images from **OpenStreetMap**. You can then export these images as a geographic text file.

There is an alluring and satisfying poetry in the composite images formed by the results of scattered sites brought together, but you're probably wondering what useful purpose Terrapattern might serve. STUDIO for Creative Inquiry's director and new-media artist, **Golan Levin**, who headed the project, emphasizes that the team did not create Terrapattern with a specific objective in mind. Rather, working with developer **David Newbury**, artist **Kyle McDonald**, and students **Irene Alvarado**, **Aman Tiwari**, and **Manzil**

Zaheer, he hopes their tool will allow users to do whatever they would like with it, whether that means using it to understand the environment or for humanitarian projects or even for pure recreation. One of Levin's friends is hunting for empty swimming pools to jump into for guerrilla skateboarding. Online, the team **references** initiatives by **Monitoring of the Andean Amazon Project** and by **DataKind.org** as precedents. In this sense, Terrapattern is intended, as its developers put it, to "democratize geospatial intelligence," providing the everyday person with a power that lies largely in the hands of state actors or big corporations. But Levin also considers it an artwork that provides people with new insight into their cities.

"My hope for this project [is] that it is an influential prototype that allows people to suddenly think in a new way about satellite imagery," Levin told Hyperallergic. "It's what I would call a revelatory artistic practice, in which I'm trying to allow people to see the world in a new way. To give people this kind of view — this kind of panoptic perceptron that allows them to see connections in the landscape that they couldn't see before — is a power that I'm really pleased to be able to present in the form of an interactive networked artwork."

Purple tennis courts in San Franciso

It's easy to while away time on Terrapattern, even if you don't really have a set intention. Clicking around invariably leads to some interesting, visual understandings of urbanism, even if they don't necessarily carry great social meaning. For instance, I could easily use GoogleMaps to search for tennis courts in San Francisco, but Terrapattern allows me to see how many purple ones the city has. Perhaps more helpful to some is how easy the project makes it to find buildings with solar panels on their roofs. In Pittsburgh, you'll find neighborhoods rife with round yard pools and cul-de-sacs; contrasting with these indicators of suburbia are the shipping container yards of New York City, with cars neatly lined up like colorful bits of unused chalk; or the areas in Detroit where expressways intersect, which, when compiled, form a giddy snapshot of urban transportation. I was also able to locate the sections of New York City's overflowing cemeteries that are divided by wide roads, a collection of images that alludes to the city's history of

negotiating the relationship between its dead and its living.

Of course, users are not restricted to only tracking infrastructure. The Carnegie Mellon team has collected images of boat wakes in rivers; one of my first searches was for clusters of yellow taxi cabs. Such tiled images of ephemeral forms exemplify Terrapattern's potential for all sorts of discovery. Terrapattern is currently in alpha mode, and its developers are working to roll out more cities soon. On deck next are London and Johannesburg.

"Mostly I want to give people this fun experience, where they spend time clicking around and think this is fascinating even if they don't really know what it's good for," Levin said. "If someone spent two hours with it, that indicates this is something profound."

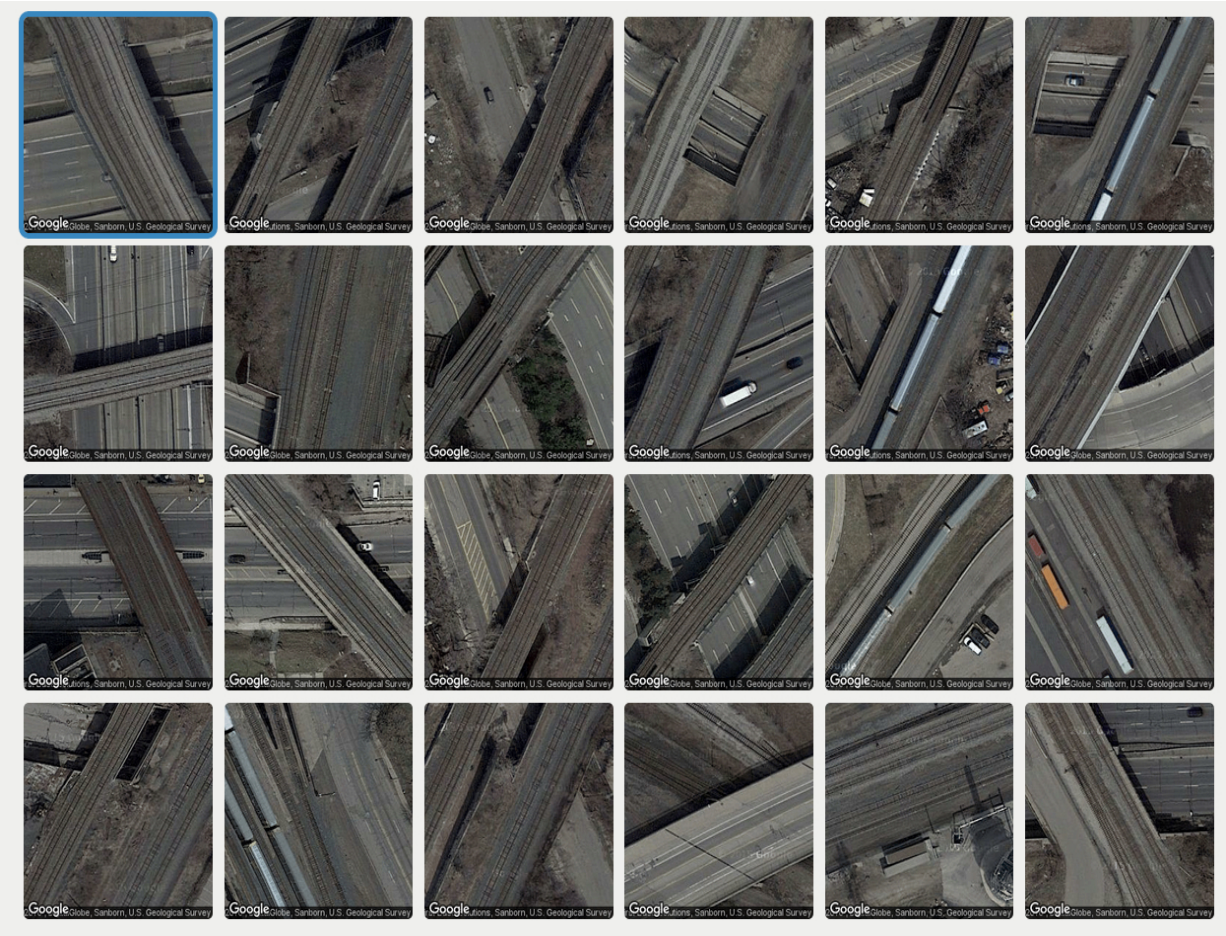

Intersecting expressways in Detroit

Areas of cemeteries divided by roads in New York City

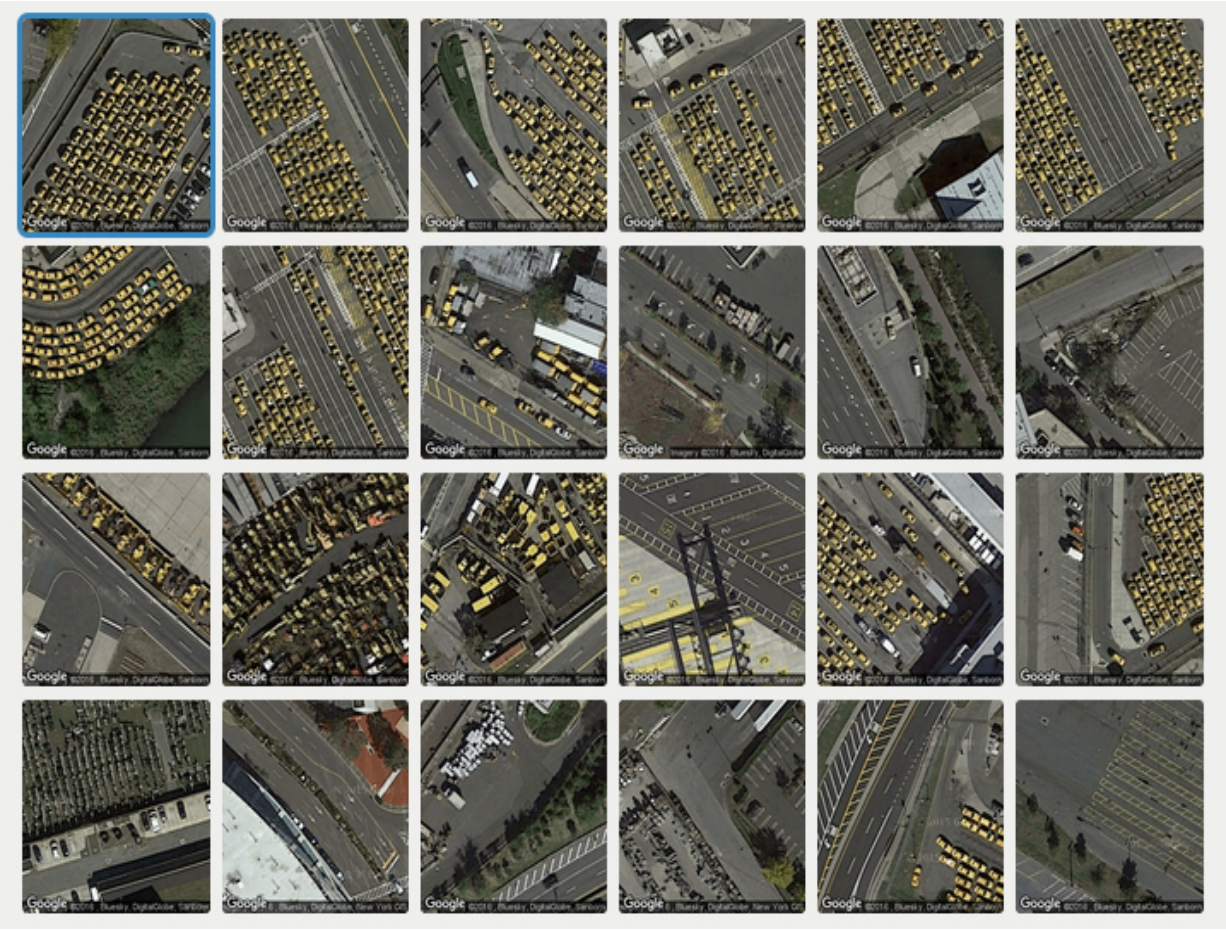

Clusters of cabs in New York City

Dilapidated plots in Detroit

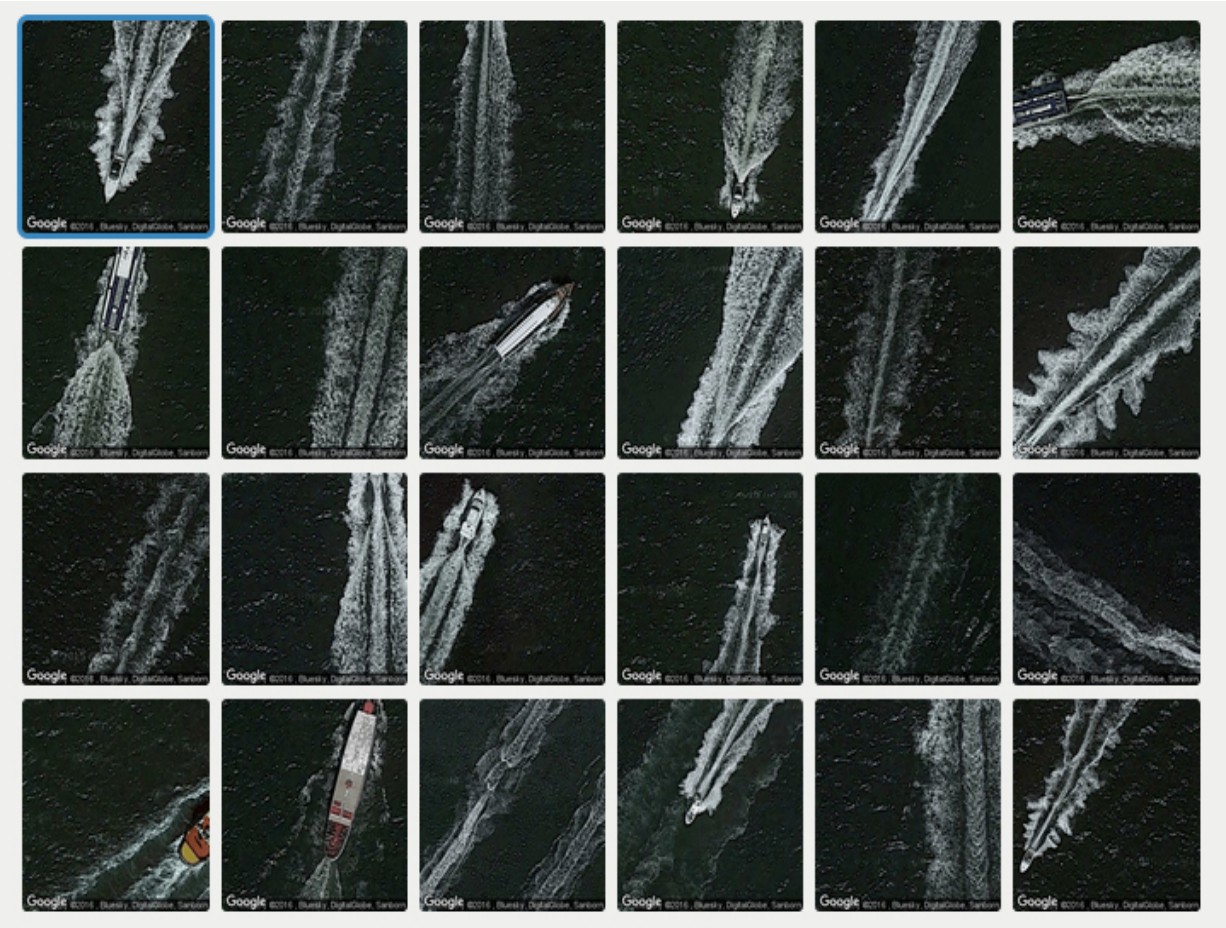

Boat wakes in New York City rivers

Cul-de-sacs in Pittsburgh

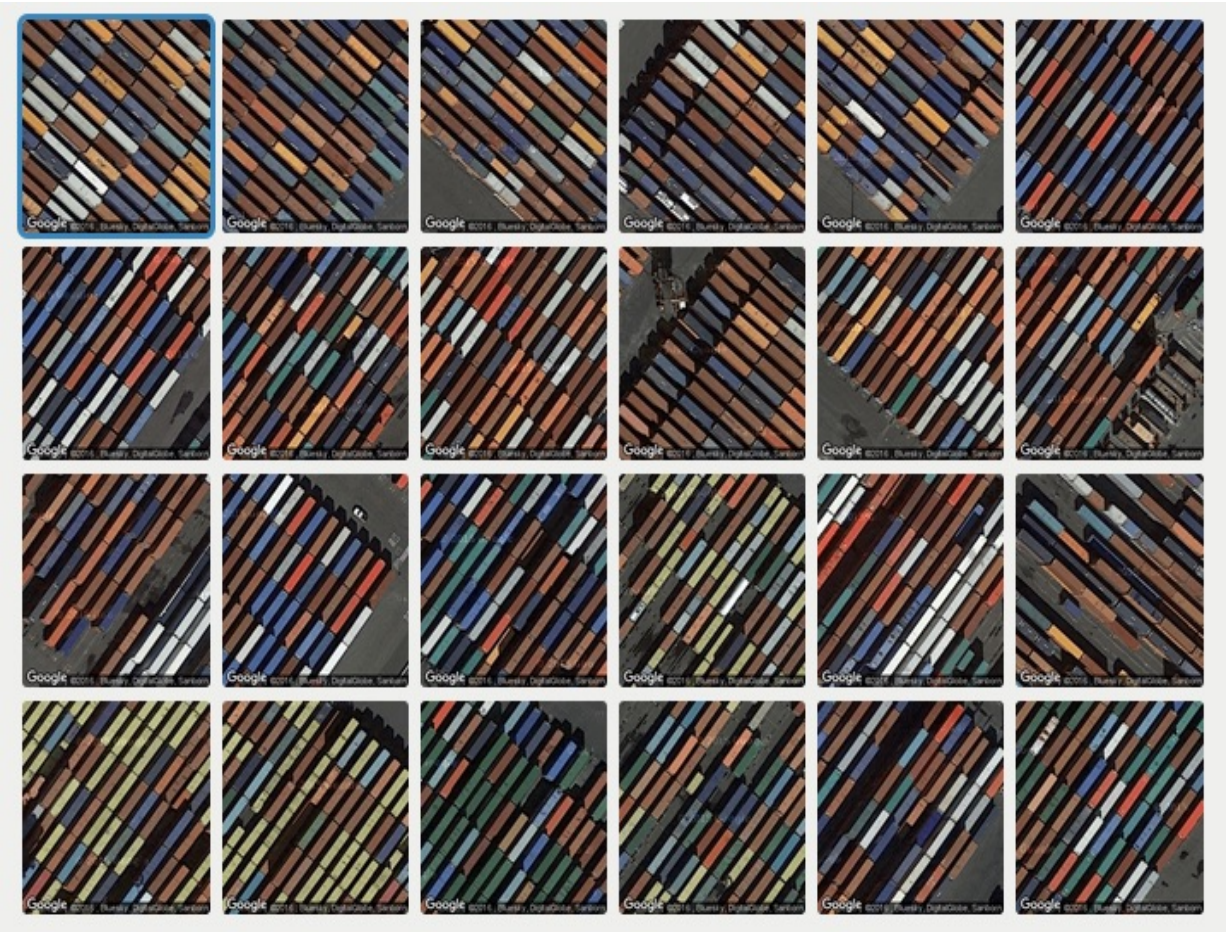

Shipping container yards in New York City

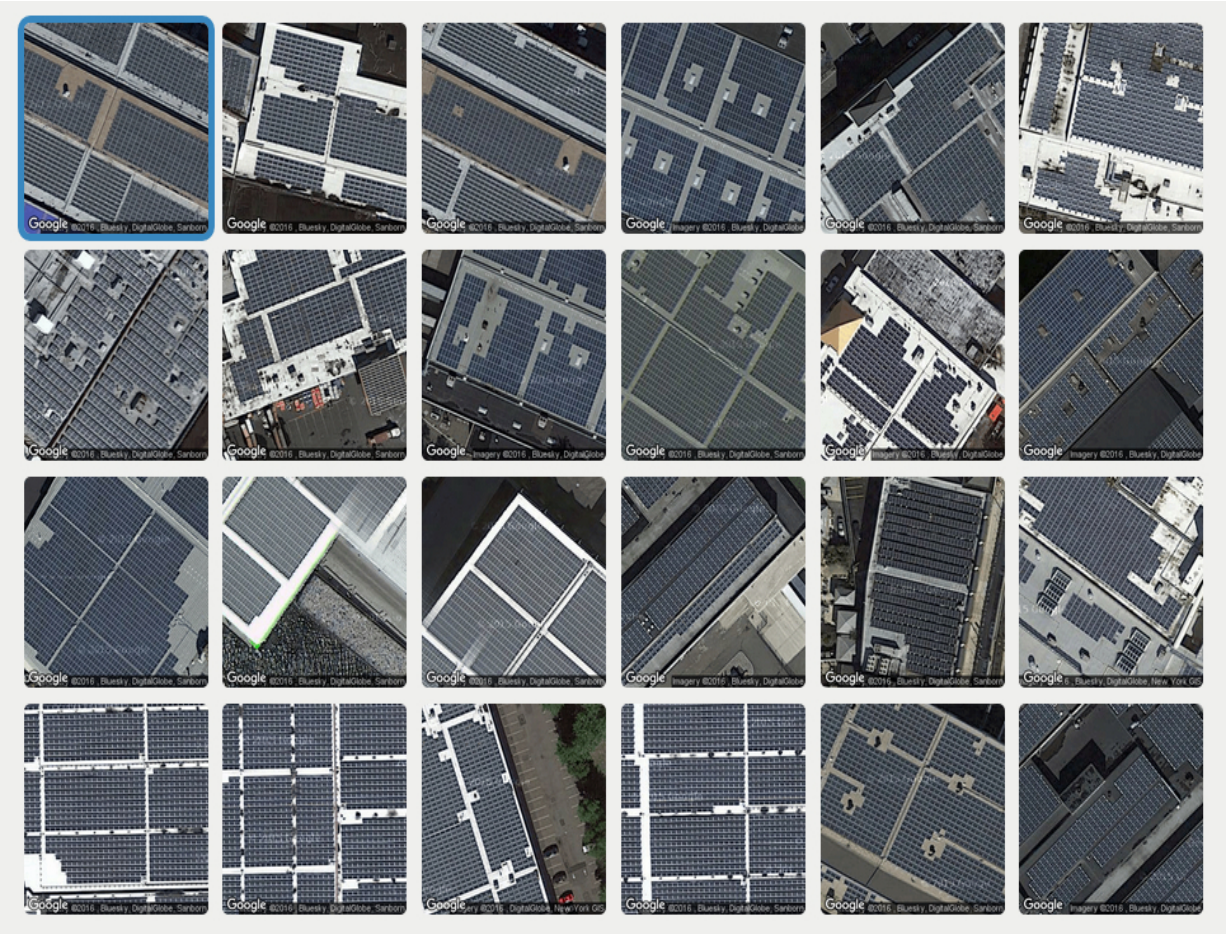

Solar panels in New York City

aerial photographyAman TiwariDavid NewburyGolan LevinIrene AlvaradoKyle McDonaldManzil ZaheerTerrapattern

Terrapattern – Neural network visual search tool for satellite imagery

Created by Golan Levin, David Newbury, and Kyle McDonald, with the assistance of Golan's students at CMU, Irene Alvarado, Aman Tiwari, and Manzil Zaheer, *Terrapattern is a* visual search tool for satellite imagery. The project provides journalists, citizen scientists, and other researchers with the ability to quickly scan large geographical regions for specific visual features.

Terrapattern uses a deep convolutional neural network (DCNN), based on the ResNet ("Residual Network") architecture developed by Kaiming He et al. The team trained a 34-layer DCNN using hundreds of thousands of satellite images labeled in OpenStreetMap, teaching the neural network to predict the category of a place from a satellite photo. In the process, their network learned which high-level visual features (and combinations of those features) are important for the classification of satellite imagery. They used 466 of the Nominatim categories (such as "airport", "marsh", "gas station", "prison", "monument", "church", etc.), with approximately 1000 satellite images per category. Their resulting model, which took 5 days to compute on an nVidia 980 GPU, has a top-5 error rate of 25.4%.

Public housing projects in New York City

Rusting fuel tanks in New Jersey

Electrical transformer stations in Pittsburgh

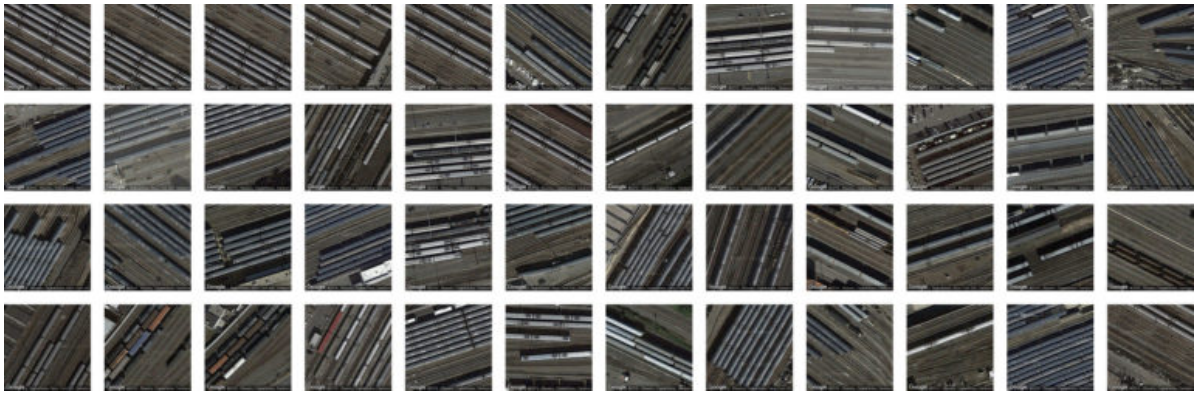

Rail yards in northern New Jersey

The *Terrapattern* search tool features three visualizations: a slippy map, for specifying visual queries; an "Geographical Plot" (or minmap), which shows the locations of search responses in the surrounding metro region; and a "Similarity Plot", which organizes the returned results within an abstract 2D space using Principal Component Analysis, or PCA. The Terrapattern website is built using Ruby and JavaScript, with satellite imagery from Google Maps, while the Geographical PLot and Similarity Plot were created in JavaScript using p5.js.

There has been many machine learning projects produced over the last year – especially drawing on computer vision. We see openCV and similar pixel based CV tools being replaced with neural networks. Even though the technology has been around for a few years now, thanks to the quick rise of opensource tools, artists are beginning to utilise techniques that were until now only available to the researchers. It is important to note that *Terrapattern* is not intended to allow users to locate their long lost teddy or other 'personal' searches since *Terrapattern* relies on map tiles to seek out similar items. This means that objects larger than 20m meters are much easier to find than (say) an iPhone. Nonetheless, it does shed considerable light on where this technology is heading, or even more worrisome what already exists, and just how easy it is with sufficient resources and computing power to locate pretty much anything left out in the open.

For more technical information on *Terrapattern*, including their open-source code, models and datasets, please see this list of technical references and their Github repository.

Smithsonian.com

SmartNews Keeping you current

This New Satellite Project Helps People Find Patterns in City Spaces

Terrapattern turns a mad world into a satisfying, matchy-matchy nirvana

Ahhh...that was satisfying. (Terrapattern (Flickr/Creative Commons))

By Erin Blakemore
smithsonian.com
June 9, 2016 2:34PM

What *can't* satellites do? They keep an eye on animals, track humanitarian crises, even help predict famines—and their above-ground perspective lets ground-bound observers find unexpected beauty in their surroundings. Now, writes Eillie Anzilotti for *CityLab*, a new project is helping people find lovely patterns and strange similarities in cities around the globe.

Terrapattern launched last month, Anzilotti reports, and its concept is deceptively simple: Use satellite images to track specific visual features around large geographical regions. The open-source project uses machine learning to help people find places that look the same.

"We are particularly keen to help people identify, characterize and track indicators which have not been detected or measured previously," the founders write, "and which have sociological, humanitarian, scientific, or cultural significance."

That's cool in theory, but addictive in practice. Users can pick a visual feature from one of thousands of high-res satellite images from five metro areas around the globe. A neural network then scans other cities for the same kinds of images.

Like baseball diamonds or Christmas tree farms? You can use Terrapattern to find a dizzying number of similar examples. But you don't just have to look at recognizable objects like train tracks or runways—you can simply focus in on an area that has your favorite color or an interesting design.

The site was created in part by Golan Levin, a Carnegie Mellon art professor who's obsessed with how humans interact with technology. His art does everything from help fonts evolve to add fingers to hands using creepy interactive software. He tells Anzilotti that he hopes the technology could be used to quickly locate disparities or environmentally damaging activities.

But for people obsessed with symmetry, visual style and the oddly satisfying, the project serves up so many soothing similarities that its potential benefits to the world are just a bonus. Can't get enough? Head to Terrapattern's stunning Flickr page for even more examples of Earth's most gratifying patterns.

About Erin Blakemore

Erin Blakemore is a Boulder, Colorado-based journalist. Her work has appeared in publications like *The Washington Post*, *TIME*, *mental_floss*, *Popular Science* and *JSTOR Daily*. Learn more at .

|

The Atlantic

Subverting Our New Space Overlords

Governments and hedge funds are pulling economic data from daily satellite images of ports, farms, and even mall parking lots—here's how they might be fooled.

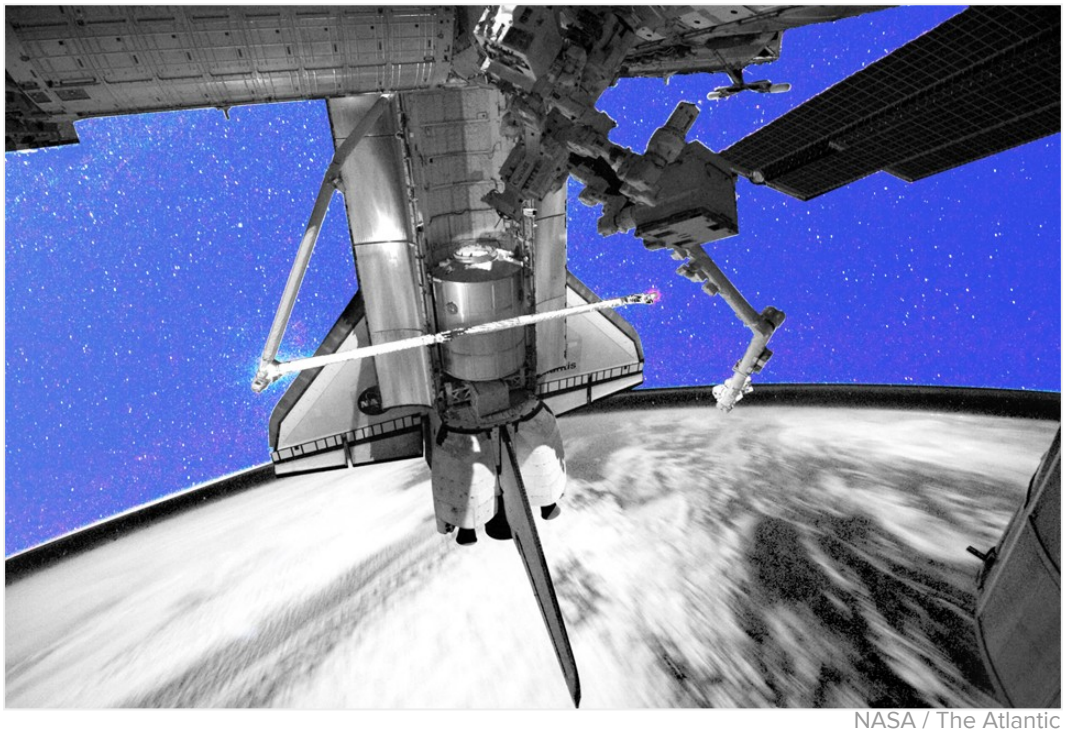

NASA / The Atlantic

GEOFF MANAUGH

10:03 AM ET | TECHNOLOGY

TEXT SIZE

Complex financial information is hidden in plain sight all over the planet, according to James Crawford, CEO of Orbital Insight. The number of ships docked at a Malaysian port, even the color of a wheat field in western Nebraska, are actually signs, Crawford explained to me, visible indicators of economic activity, not just for

a local region but for an entire global industry.

Satellites, balloons, lasers — the coming information age in space.
Read more

Seen this way, mundane landscapes previously deemed unworthy of analysis can, in fact, be meticulously—and profitably—scrutinized. This newfound appreciation is not aesthetic, of course, but fiscal, as even the growing shadows of a Chinese construction site can be interpreted as valuable clues about the strength of the underlying economy.

Crawford's company, Orbital Insight, is one of a new breed of market-research firms pioneering the use of high-resolution satellite imagery. This is called geo-analytics, or geography crossed with the algorithmic firepower of Big Data. With access to satellite images—refreshed on a daily basis and available at a scale of one meter per pixel—companies such as Orbital Insight use artificially intelligent deep-learning algorithms to sort through the data and look for patterns.

Often this is just about change-detection: that is, looking for a particular pixel that has flipped from one color to another, thus indicating a new agricultural condition or the beginning of major construction work. Other times, it is all about quantity: methodically counting the number of cars parked outside a shopping mall in Minneapolis, or the trucks lined up outside a Chinese steel yard. In either case, it's about combining machine vision with data science, or giving computational power a large enough visual dataset to work with.

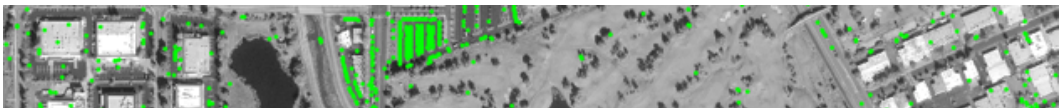

The Super Bowl, as seen by Orbital Insight.

Visual evidence captured by satellites is thus now subject to narrative interpretation for the purpose of extracting potential financial insight—and this potential financial insight can then be sold to paying customers. This, in fact, is Orbital Insight's operating business model, marketing its geo-analytic expertise to hedge funds, U.S. government agencies, and nonprofits alike for what those groups might be able to learn from satellite data. Nonprofits, for example, might look for the level of water in a remote desert lake or threatened reservoir, or the true extent of a developing city as it sprawls into a nearby national park, for indications of what policies they might next pursue; hedge funds might see this same data and decide to short certain commodities.

Not one to mince words, Crawford refers to this as developing "new understandings of the Earth."

Crawford's own background here is instructive for understanding what is really happening: he was once engineering director of Google's book-scanning project, another daunting task in which huge amounts of visual data had to be sorted by machines so that tiny marks on a page could be recognized as letters, then those letters as words, then those words as searchable sentences. "It was a similar problem," he explained to me. "It's just a tremendous number of pictures. You have to figure out where the text is and what it all means; you have to put it into Google search and so on."

The challenge is now in applying this same form of graphic analysis to landscape: learning to read the surface of the Earth such that topographic features, building shapes, and the presence of shadows can all begin to reveal their deeper meanings. It is, in other words, about transforming the world into a hieroglyph—and then cracking this new terrestrial language, making it legible not only by human beings but by semi-intelligent machines. Not one to mince words, Crawford refers to this as developing "new understandings of the Earth."

Speaking to Crawford as he reels off future applications of his firm's technology is like listening to a one-man litany of geographic features and their potential implications. He is a reader in search of a text. Crawford will spin through long and fluid lists of everything from how many houses are going up in the suburbs of Phoenix to the number of train cars clacking across a Chinese railroad network—to the perceived attendance of suburban churches in the U.S. and the

corn yields of large-scale agribusiness in the American heartland to rice production in southeast Asia.

"We view the Earth as really big and really complicated—but fundamentally understandable," he explained. "Using this combination of AI technology and humans, working in partnership, we can handle that volume of imagery to extract the insights that might be implicit in it. The question we ask ourselves is: where are the information gaps about the Earth?"

Rather than character recognition, then, it is what we might call landscape recognition.

* * *

At the end of May 2016, a small team led by Golan Levin at Carnegie Mellon University released a new online tool called Terrapattern. Terrapattern, as *The New Yorker* described it, can be thought of as "Google Earth's missing search engine." It is capable of scanning—and automatically finding visual similarities amongst—thousands of satellite images of the Earth's surface. You can click on any visible feature—Brooklyn rooftops, Detroit parks—and Terrapattern will instantly locate any other nearby landscapes that resemble it. It will, in other words, match the pattern.

In essence, Terrapattern simply gives us—the general public—a glimpse of some of the proprietary analytic tools that firms such as Orbital Insight already have at their disposal. For now, Terrapattern's geographic resourcefulness is limited to Pittsburgh, San Francisco, Detroit, Berlin, and New York City, but its goal is global: an entire planet subject to autonomous visual analysis. A parking lot in Knoxville, a football field in Fresno, a factory complex near Salt Lake City: these organized landscapes will eventually find

their visual twins in places as far afield as South Africa, Venezuela, and Japan.

Companies such as Orbital Insight can count the cars in the parking lot around a building.

It is intriguing to consider how this interpretive process might be flipped on its head, so to speak, or turned around from back to front, to look at how these machine-learning algorithms might be made to fail in intriguing ways. Doing so would seem to suggest a range of peculiar new landscape design opportunities, both for governments hoping to maintain secrecy and for private firms wishing to disguise their everyday activities.

A Walmart franchise hoping to look better-attended than it really is,

for example, could spoof the satellites with a parking lot full of fake cars, or a new residential building could be playfully designed so that its roofscape looks like a neighborhood park, throwing off the watchful eyes of Terrapattern.

Consider the work of artist Adam Harvey. Harvey has made a name for himself exploring, among other things, how cosmetics might be used "to break apart the continuity of a face," in his words, so that that person can no longer be identified by facial-recognition algorithms. He calls this "camouflage from face-detection technology."

There is no reason to believe this sort of effect could not be scaled up to the level of architectural or landscape design, like dazzle paintings on the hulls of World War I ships, deliberately employed as a kind of corporate smokescreen to evade financial interpretation from above. Duping the algorithm, so to speak—or disrupting the hieroglyph—suggests a way to subvert the era of total visibility that Orbital Insight and Terrapattern both imply.

Those eyes soaring high above the clouds might always be watching us, in other words, but they needn't always understand what they see.

ABOUT THE AUTHOR

GEOFF MANAUGH writes regularly at *BLDGBLOG*.

kottke.org posts about Terrapattern

Satellite imagery search engine MAY 25 2016

Terrapattern is a search engine for satellite images. You click on a specific feature of interest on a map and the site returns results that match it. For instance, here are the locations of solar panels in NYC.

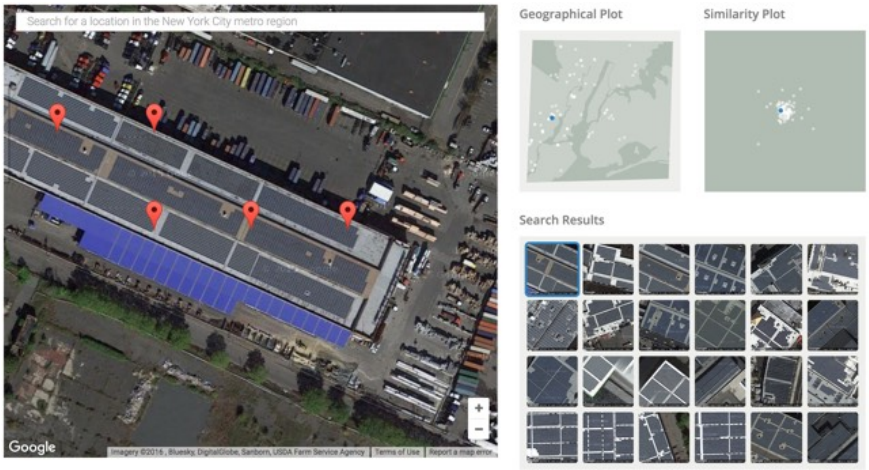

You can also use Terrapattern to find school bus depots, fracking wells, Air Force bombers, baseball diamonds, train tracks, and much more.

There are only four cities currently represented (Pittsburgh, New York, San Francisco, and Detroit) but this is already super cool to play around with. (via @genmon)

THIS IS KOTTKE.ORG
Front page
About + contact
Site archives

You can follow kottke.org on Twitter, Facebook, Tumblr, Feedly, or RSS.

From The Atlantic
CITYLAB

MAPS

Mapping the Hidden Patterns of Urban Life

Terrapattern collects visually similar features from satellite images into one searchable platform.

EILLIE ANZILOTTI | @eillieanzi | Jun 7, 2016 | Comments

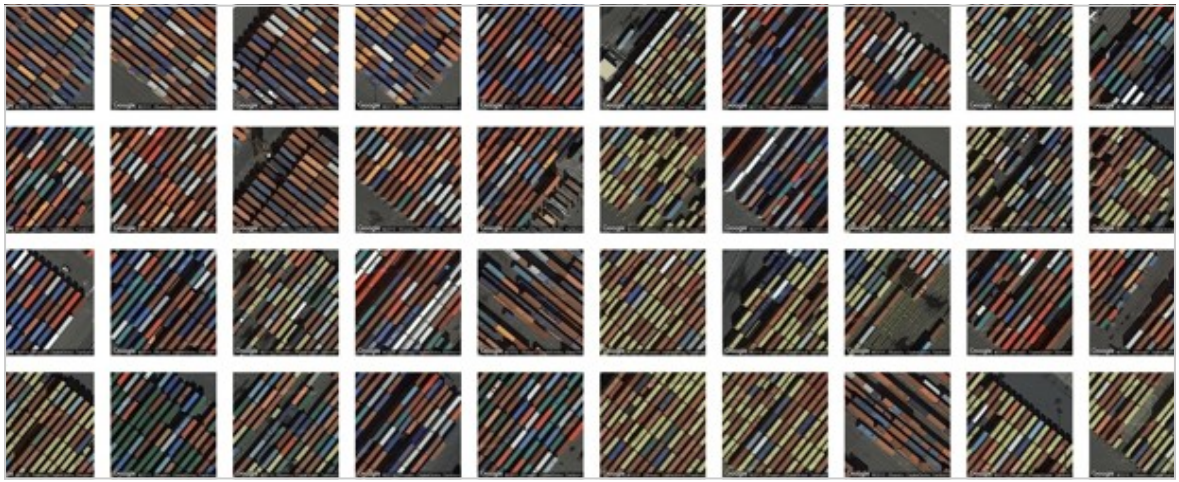

Shipping containers in San Francisco. (Terrapattern/Flickr)

Think about the view from an airplane: the landscape organized into neat tiles of farmland, freeway roundabouts drawn into the earth like children's scribbles, rooftop gardens dotted across the tops of urban skyscrapers. There's a satisfying pattern to it all.

But that bird's-eye view also offers an information-loaded perspective that's now, accessible to the earth-bound. Terrapattern, an open-source prototype project that launched last month, employs satellite imagery to track visually similar occurrences in five cities: Pittsburgh, San Francisco, Detroit, New York, and Berlin. It will soon expand to several more, including Miami and Philadelphia.

The platform is "an interface for finding '*more like this, please*,' in satellite photos," according to its website. It's more exploration-driven than a

simple image-search engine: you can't just load a map of Detroit, type in "urban gardens," and receive a slew of similar plots. Instead, you can look through the city street views, select a feature of interest—say, a yellow cross-walk in San Francisco—and Terrapattern will offer up the visually similar sites throughout the region.

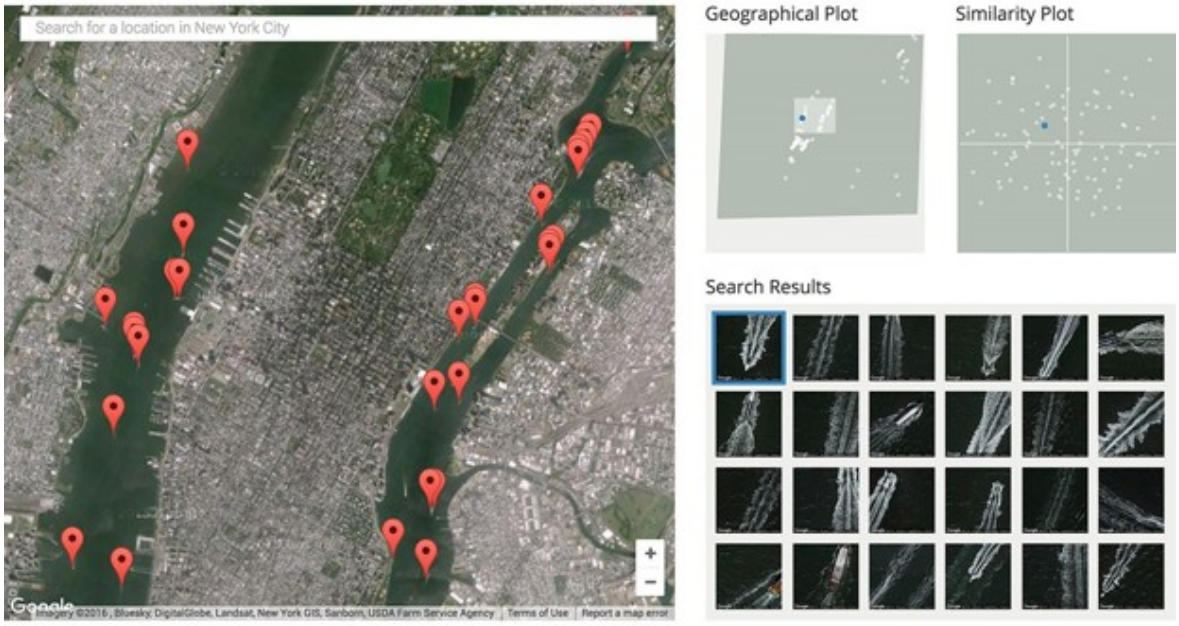

A search for boat waves in the New York City waters. (Terrapattern/Flickr)

The developers feed the Terrapattern hundreds of thousands of satellite images sourced from OpenStreetMap, teaching the system to learn the visual cues important to identifying a particular place or feature. The platform contains a huge wealth of data—each metro area requires at least 10 gigabytes of RAM for active memory. It synthesizes the patterns into concrete, searchable data points, says the Carnegie Mellon University professor Golin Levin, who along with David Newbury, Kyle McDonald, and students Irene Alvarado, Aman Tiwari, and Manzil Zaheer, developed Terrapattern.

School buses in Pittsburgh. (Terrapattern/Flickr)

Searching the platform is addicting, and there's a whimsical appeal to it: from above, the school buses lined up in Pittsburgh's depots look like rows of chalk.

But this way of mapping, Levin says, may also shed a new light on issues of economic disparity, as well as humanitarian and ecological causes. From above, Levin adds, researchers could track illegal logging roads in rainforests as precursors to deforestation. The tool could also be used in disaster relief efforts to assess the extent of flood or fire damage, for example.

Terrapattern launched on a budget of less than $35,000; there is no way, Levin says, that this platform alone could present a searchable view of the entire United States—that would require 2,000 times more computer storage, *Fast Company* reported. But the point, Levin says, is "to inspire people's ideas about what's possible." Several other startups—including the Google-owned Terra Bella—have sent small imaging satellites into orbit; such tools provide up-to-date pictures of the earth that, when combined with Terrapattern's visual similarity tracking technology, could present a categorical view of intentional and natural land-use changes over time.

Even just the five metro areas currently searchable on Terrapattern show a strangely voyeuristic look into trends in human behavior. The number of solar panels on New York City roofs is as encouraging as the presence of fracking wells in Pittsburgh is bleak. "Where there is a regularized way that people treat the world, that will be reflected in the landscape," Levin says. Never has that been more obvious.

Visual search tool for satellite imagery

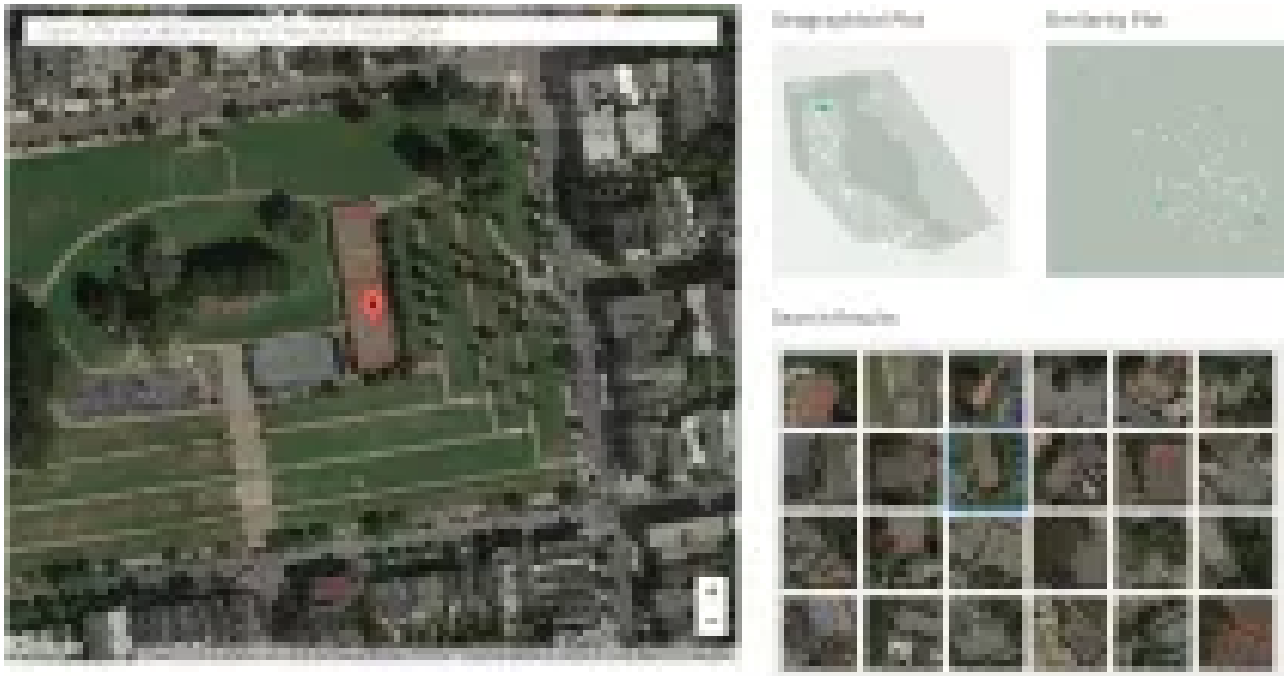

Terrapattern is a fun prototype that lets you search satellite imagery simply by clicking on a map. For example, you can click on a tennis court, and through machine learning, the application looks for similar areas.

> Terrapattern uses a deep convolutional neural network (DCNN), based on the ResNet ("Residual Network") architecture developed by Kaiming He et al. We trained a 34-layer DCNN using hundreds of thousands of satellite images labeled in OpenStreetMap, teaching the neural network to predict the category of a place from a satellite photo. In the process, our network learned which high-level visual features (and combinations of those features) are important for the classification of satellite imagery.

Terrapattern: Search Engine for Satellite Imagery ~ GIS Lounge

Terrapattern is a way to search visually within satellite imagery. Developed using an open source search tool, Terrapattern is the brainchild of Golan Levin, David Newbury, and Kyle McDonald who received funding from the John S. and James L. Knight Foundation Prototype Fund.

What Terrapattern does is perform "similar-image" searches in unlabeled satellite imagery using deep learning machine vision techniques. To train the system, a deep convolutional neural network (DCNN) was fed hundreds of thousands of categorized OpenStreepMap imagery. After the DCNN has learned which visual features are useful for classification, descriptions for millions of satellite imagery tiles were then computed. Since searching can take a long time and a significant amount of RAM, the developers use the CoverTree algorithm to precompute relationships between the descriptions, reducing user search time to 1-2 seconds.

This approach allows users to search for particular features across a set of satellite imagery with nonbuilding structures (structures such a bridges, bus shelters, and swimming pools) and soft infrastructure yielding the best search results. For example, users can search for images that show yellow crosswalks, purple tennis courts, grassy fields, and more. Terrapattern lets users search visually for geographic features that may be ephemeral or not typically mapped out. Results from searches are then exportable as GeoJSON files.

Terrapattern is currently only available for these four major U.S. cities: New York, Detroit, Pittsburgh, and San Francisco. An alpha prototype, the developers are limited by the enormous about of active RAM (about 10GB per major city) that terrapattern takes to store the model data. The developers are interested in hearing back from "citizen scientists, data journalists, humanitarian researchers, and other domain experts to tell us about how our app is, or could be, of use." Fill out the survey to provide your input.

Visit: **Terrapattern** and access the code on **GitHub.**

Share:

Cool Tech! TerraPattern, Open-source Tool For Discovering "patterns Of Interest"

Here's some pretty nifty imaging technology that rolls in some slick image recognition and a ton of data! Enter TerraPattern.

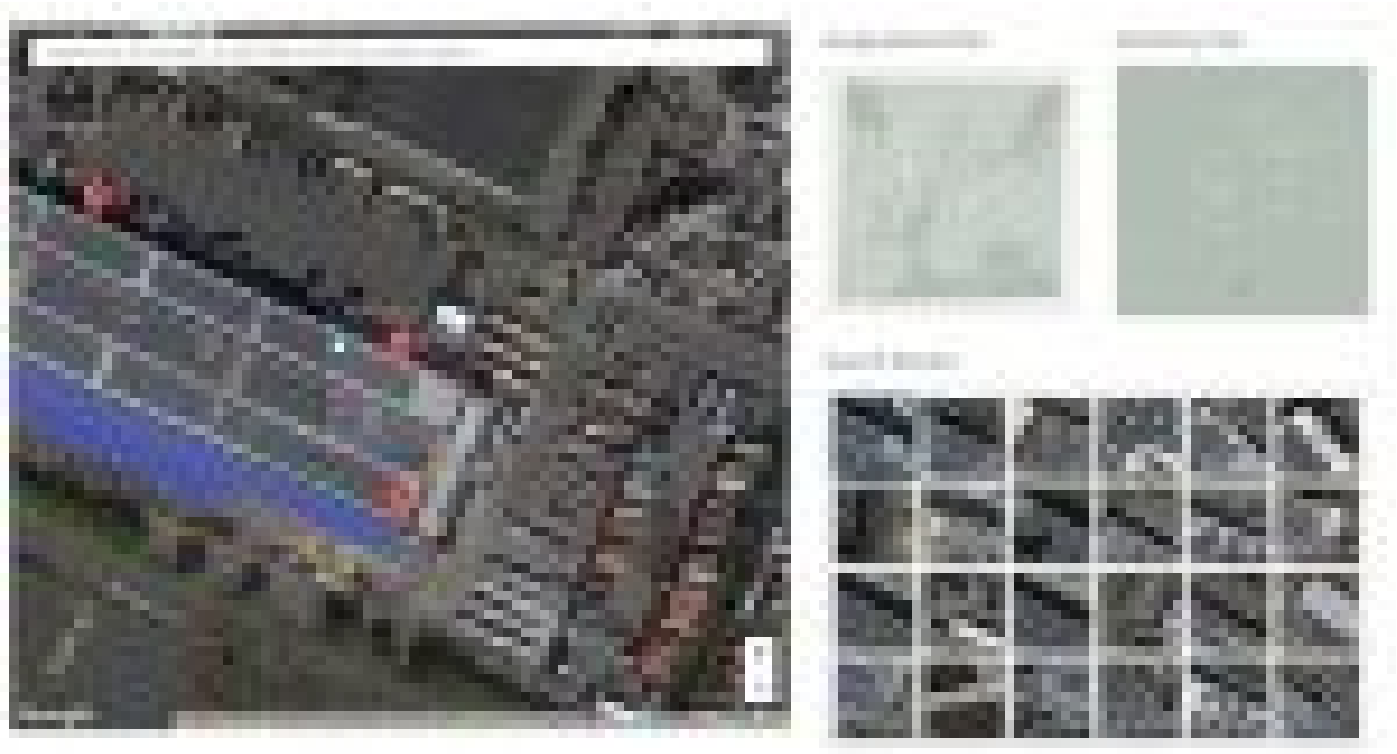

Our tool is ideal for locating specialized 'nonbuilding structures' and other forms of soft infrastructure that aren't usually indicated on maps.
From TerraPattern they answer the question, "What is TerraPattern"?

> *It's an open-source tool for discovering "patterns of interest" in unlabeled satellite imagery—a prototype for exploring the unmapped, and the unmappable.*

Currently in an early "Alpha" stage, TerraPattern provides the technology for users to locate an area of interest on an image, for example over a building with solar panel rooftop. The service then seeks out areas with a similar image footprint. The search results are displayed in a pleasant result box along with a map that shows where the similar places are found geographically compared to the selected point/place on the map.

The service provides image recognition for several cities including: New York, San Francisco, Pittsburgh, and Detroit. Give it a go at http://www.terrapattern.com

Thanks to Paul B for the cool Tip!

Terrapattern, the search engine for imagery - Google Earth Blog

May 27, 2016

Terrapattern is a very interesting new online search engine for aerial and satellite imagery. It is still an alpha version and only covers a few cities in the US, but the potential is very significant. To try it out, just go to www.terrapattern.com, select a city, then select a map tile of interest and it will find other similar map tiles.

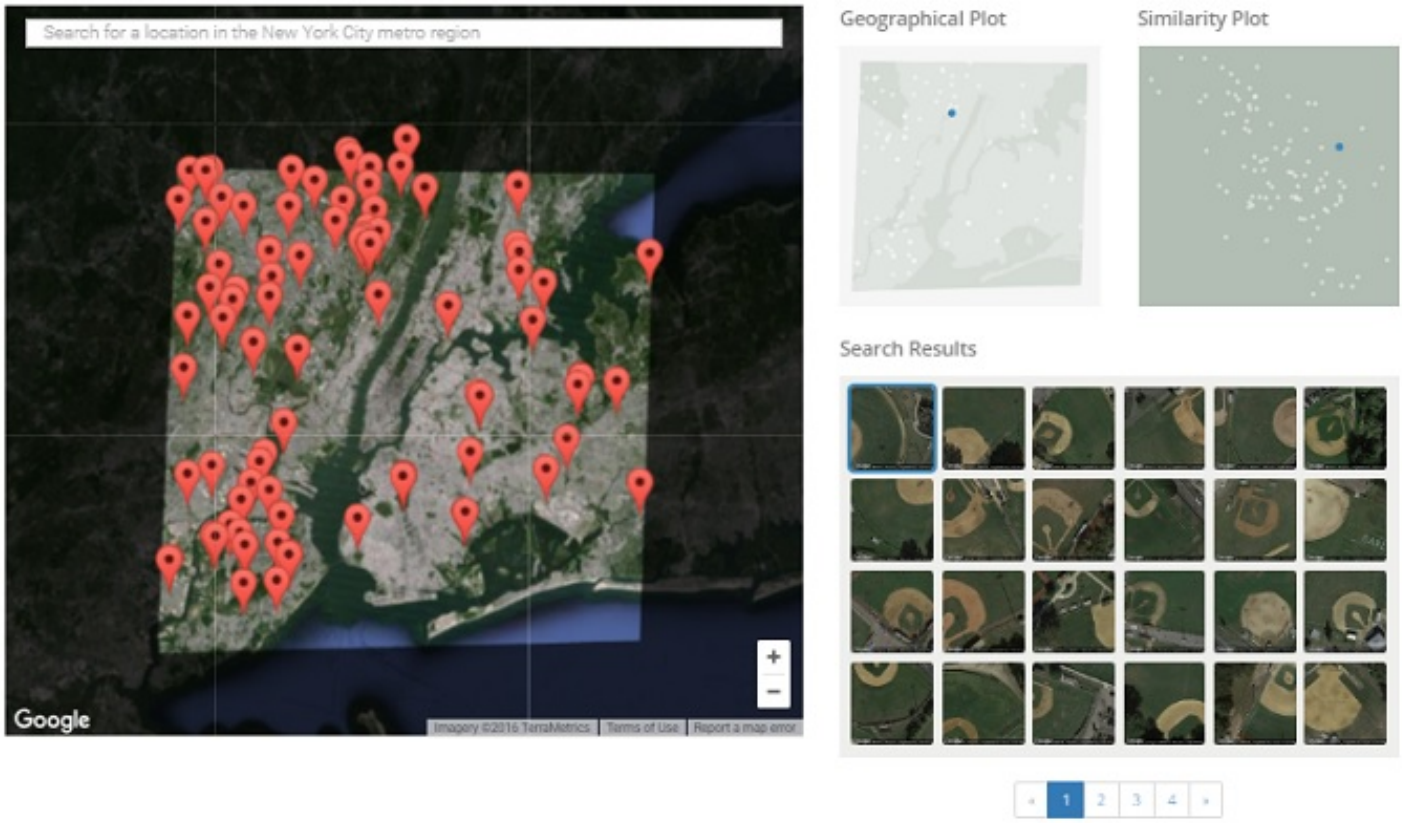

The baseball fields of New York.

We found it remarkably good for finding various types of sports fields (football fields, tennis courts, baseball fields, golf courses). We suggest also trying things like container storage areas, parking lots, junk yards, different types of roofs, etc. We found that if we selected a section of a bridge it could find other bridges as well as piers, but not with 100% accuracy. This is because, as explained in the FAQ, it works on individual tiles and large features like bridges, which cover many tiles, are less accurately recognised. We foresee, however, a future version working on different scales of detail to classify structures of different sizes.

It is open source, so if anyone has computing power and other necessary resources they can use the code for other parts of the world. We are not sure if explicit permission is required before using Google Maps imagery for something like this, so be sure to read the Geoguidelines before implementing such a project.

They are hardly the first to think of using image recognition on aerial and satellite imagery. For example, we have previously had a look at German design studio Onformative who applied face recognition to imagery with remarkable results. It is probable that all major map makers have considered the possibility of using imagery to identify features relevant to maps and highly likely that several companies have built systems to attempt to do this. If any of our readers knows of any such projects and how successful they have been, please let us know in the comments. Google has a sophisticated image recognition artificial intelligence (AI) that they use with photos to provide both Google image search and tagging capability in Google Photos. It is almost certain that they have tried it out at some point on satellite / aerial imagery too.

Read more about Terrapattern here.

About Timothy Whitehead

Timothy has been using Google Earth since 2004 when it was still called Keyhole before it was renamed Google Earth in 2005 and has been a huge fan ever since. He is a programmer working for Red Wing Aerobatx and lives in Cape Town, South Africa.

Image recognition and Google Earth - Google Earth Blog

July 28, 2016

The last few years have seen major advances in computer artificial intelligence (AI). One area where AI is starting to show practical use is in imagery recognition. Google Earth and Street View imagery combined with image recognition has a wide range of possible applications. We have in the past had a look at Terrapattern, an experimental search engine for aerial and satellite imagery. They are adding new areas with time, so be sure to keep an eye on them.

We recently came across this story about a Caltech researcher that is helping the city of Los Angeles to count its trees with the help of a combination of Google Earth imagery and Street View. In this case they are trying to not only count individual trees but also identify the species.

The idea of using imagery for surveys of vegetation is of course far from new. Google Earth Engine, for example, is designed around such large scale analysis. When you wish to simply determine whether there is vegetation cover or possibly the overall health of the vegetation, a much better option than Google Earth imagery is to use false colour imagery – and satellites are typically designed with this in mind.

Another example of people using image recognition on Street View imagery is this one about identifying fire-hydrants and mentioned in that article is a project using Street View to study gentrification, which uses historical Street View to measure changes in buildings over time.

There is also this project, which uses Street View to geolocate an image. You could potentially take a photo with your mobile phone camera and the system could tell you where you were with accuracy similar to GPS. At present, this sort of thing is often done by crowd-sourcing rather than an automated system. The potential for automated systems has both potential benefits and serious privacy concerns.

Google itself applies some image recognition to Street View. The best known is identifying licence plates and faces, which are blurred for privacy reasons. However, it also reads house numbers and various street signs, and this information is used to improve Google Maps.

If Google were to add infrared to their Street View cameras, maybe it would make it easier to distinguish between faces of people who need privacy and faces of statues who need publicity.

Having infrared Street View has other uses and has been thought of already.

About Timothy Whitehead

Timothy has been using Google Earth since 2004 when it was still called Keyhole before it was renamed Google Earth in 2005 and has been a huge fan ever since. He is a programmer working for Red Wing Aerobatx and lives in Cape Town, South Africa.

Maps Mania: Searching for Map Patterns

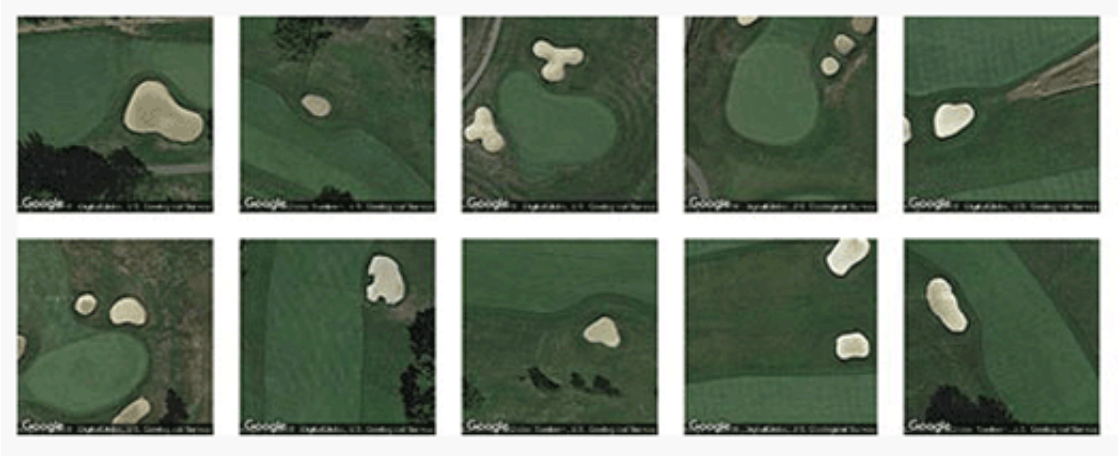

A few years ago Onformative developed an algorithm, called **GoogleFaces**, that scans Google Maps satellite imagery looking for patterns in the landscape that we might recognize as resembling human faces.

Terrapattern has taken this idea of detecting patterns in aerial imagery and developed it into something that is actually very useful. Using deep learning machine vision techniques Terrapattern is able to accept a user input (a selected area on a satellite map) and search for other locations which look the same.

For example, if you click on a golf course sand trap on the aerial map Terrapattern will instantly find other locations with golf course sand traps. Click on a stretch of a nice river bank with tree cover and you will be shown other locations where Terrapattern recognizes the same patterns in the imagery.

At the moment Terrapattern only works for Pittsburgh, San Francisco, New York and Detroit. However other cities are coming soon. The Terrapattern **about** page has other examples that you might want to try searching for, such as baseball diamonds, airplanes or solar panels. Part of the fun of Terrapattern however is just clicking on the map to see how quickly it finds other similar looking locations.

That 'about' page also includes a lengthy 'How it Works' explanation of the neural network behind Terrapattern.
Posted by Keir Clarke at 12:51 PM

Kill Some Time With Terrapattern, Which Turns Satellite Images Into Art

Prepare to get addicted.

Adam Toobin (/user/69-adam-toobin) | May 27, 2016

A new program called Terrapattern (http://www.terrapattern.com/) searches satellite images of certain regions to find visually similar features. All you have to do is pick out a small slice of an area, like a baseball field, piers on the water, or a group of parked school buses, and Terrapattern returns any areas that look the same. Currently only available for Pittsburgh, San Francisco, New York, and Detroit, the program's goal is to develop technology that makes scouring the massive quantity of satellite photos of earth that are coming in every day more reasonable.

Call it big data for satellite images or a harbinger of a potentially intrusive new way of life, Terrapattern is certainly an excellent way to waste time on the internet. Cities like New York often appear unknowably large, but suddenly we have the ability to scan it for similar features,

as if we were “Control-F” searching a document for common words. Whether you’re a data scientist, citizen journalist, or just a curious browser, Terrapattern doesn’t get boring quickly, with a database the size of some of the greatest American cities. And the best part is that it’s arriving just in time.

“It has been predicted that, within the next three years, access to daily-updated, whole-earth satellite imagery with sub-meter resolution will become widely available online,” Terrapattern’s website reads (http://www.terrapattern.com/about). “We are particularly keen to help people identify, characterize and track indicators which have not been detected or measured previously, and which have sociological, humanitarian, scientific, or cultural significance.”

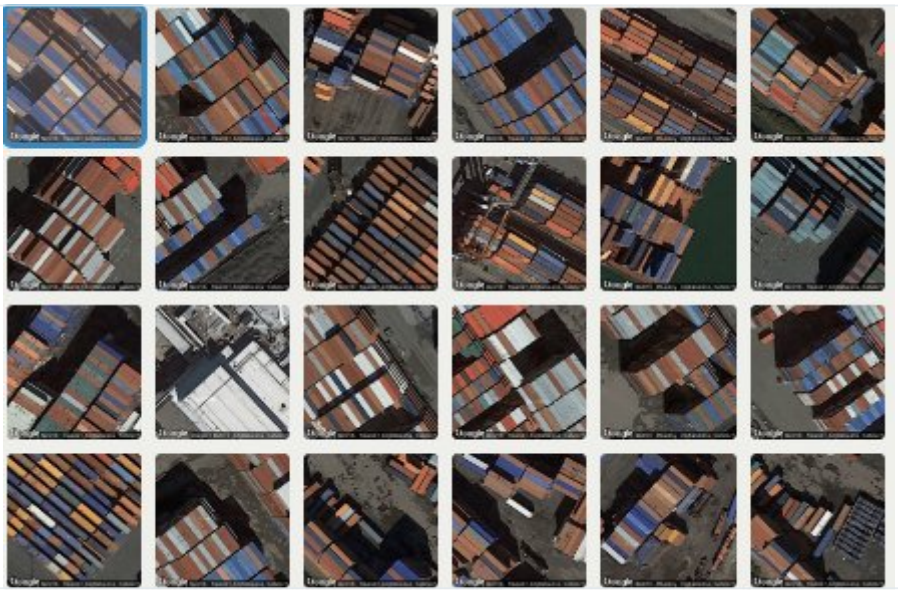

jonkeegan
@jonkeegan
Follow

Found some colorful shipping container patterns in NYC using #terrapattern nyc.terrapattern.com/?lat=40.703659…

11:59 AM - 25 May 2016

4 15

A journalist could gather a sense of a region's economic vitality by scanning shipping containers over a period of time and checking whether they are growing or shrinking in number. A watchdog can better keep an eye on a city's huge number of bridges, keeping an eye out for obvious signs of decays without ever leaving their couch. Or a politician, who feels their neighborhood has too few parks, can study the issue on Terrapattern and get real data to back up their assertions.

Bill Morris
@vtcraghead

Follow

New York Metro cemeteries, extracted using #Terrapattern:

nyc.terrapattern.com/?lat=40.692686…
11:25 AM - 25 May 2016
5

The profusion of satellite imagery that will be arriving in the coming years will only enhance these capabilities, and undoubtedly new and less reputable ideas for how to manage such a powerful tool will emerge as well. "Sub-meter" satellite resolution is a pretty powerful force, and while a single satellite snapshot of someone's home (/article/13754-find-your-house-with-nasa-s-help-using-aster-s-2-95-million-space-photos) may not feel too invasive, if that photo is updated on a daily basis, people may start to balk at giving permission.

Google Street View remains controversial around the world, with several European countries having banned the service on the basis of laws that prohibit filming someone without their permission. As everyone on the internet knows, the Street View car has amassed troves of private moments (http://thumbpress.com/maps-shall-expose-google-maps-catches-people-in-hilarious-situations-20-pics/thief-exposed-on-google-maps/), including sun bathers, amorous couples, drug deals, lots of people peeing, as well as thousands of the strange, bizarre, and inexplicable (http://www.hackread.com/funny-creepy-strange-google-street-view-images/) events that occur on world streets every day. Terrapattern wants to make it easier to search for similarities across these images, but it can be hard to

separate acceptable and inappropriate uses of such a platform.

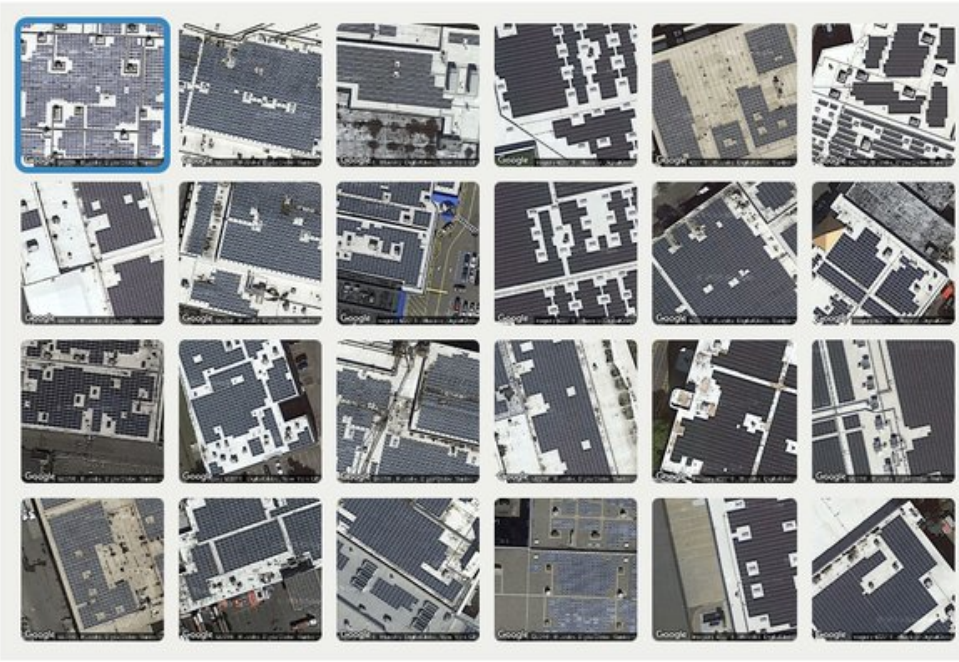

Bill Morris
@vtcraghead

Follow

GOT IT. NYC metro large rooftop solar installations, extracted with #Terrapattern: nyc.terrapattern.com/?lat=40.744414…

11:40 AM - 25 May 2016

9 23

It's a sign of Terrapattern's potential (and addictiveness) that it's only been active a few days, and #Terrapattern (https://twitter.com/hashtag/terrapattern) already returns dozens of entrancing photos of city scenes in harmony on Twitter. Be sure to check it out (http://www.terrapattern.com/), but beware you may be sacrificing your productivity for a week or two.

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

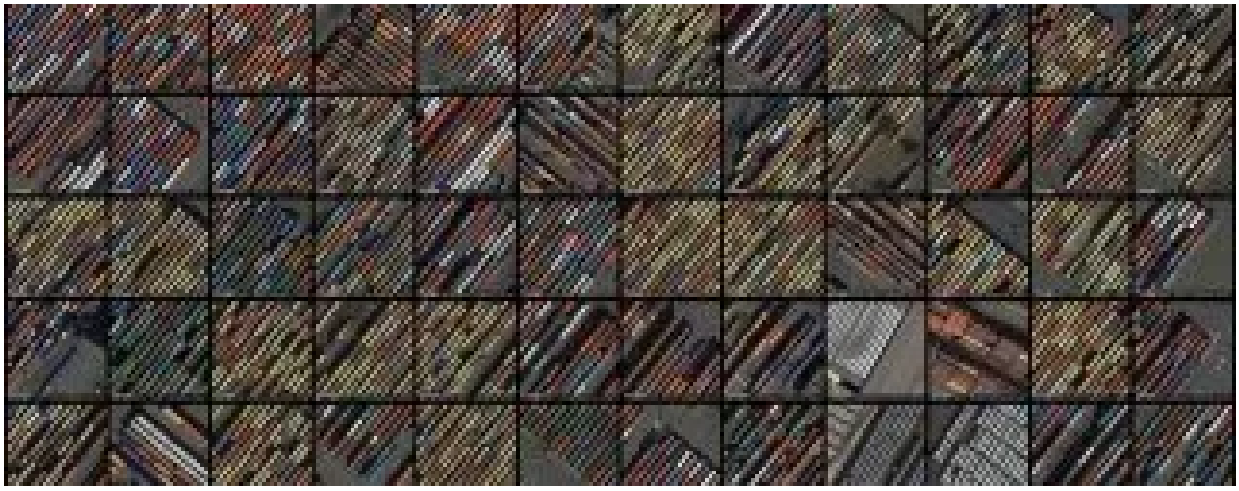

August 25, 2016

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.

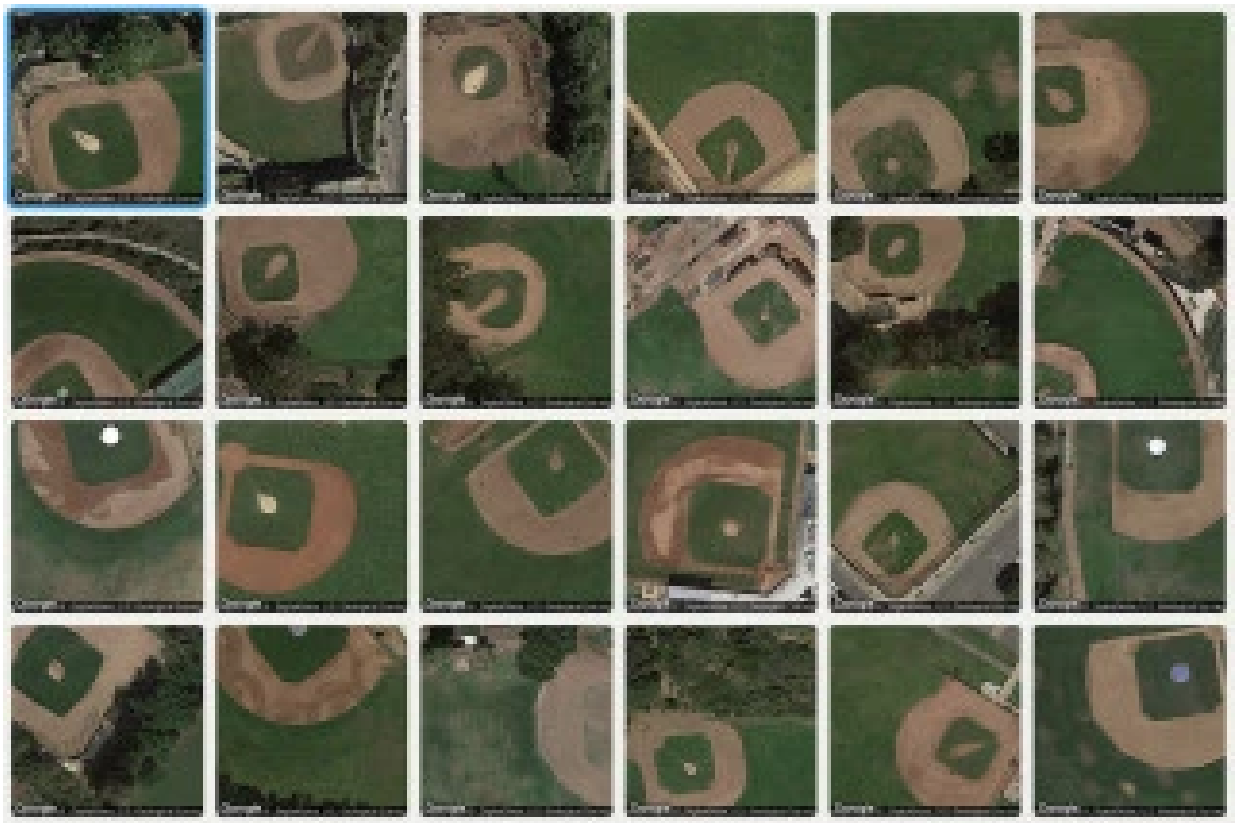

The software underlying Terrapattern can learn to recognize a baseball diamond as a whole unit—as opposed to searching every pixel to find matches. These images are from Pittsburgh. (Terrapattern)

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.”

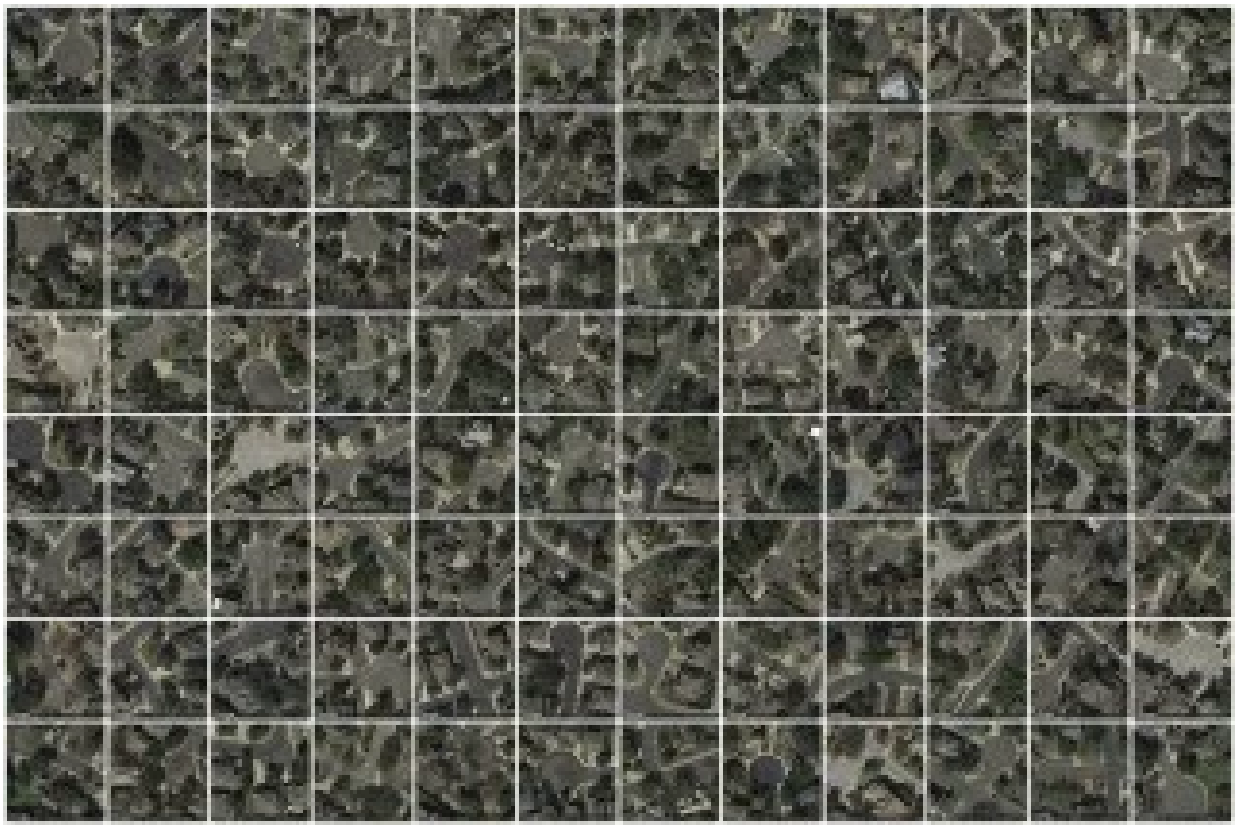

Cul-de-sacs in Pittsburgh (Terrapattern)

Airplanes in Miami (Terrapattern)

Playgrounds with sand surfaces in Berlin (Terrapattern)

Terrapattern is reverse image search for maps, powered by a neural network

Posted May 25, 2016 by Devin Coldewey

- **240**

SHARES

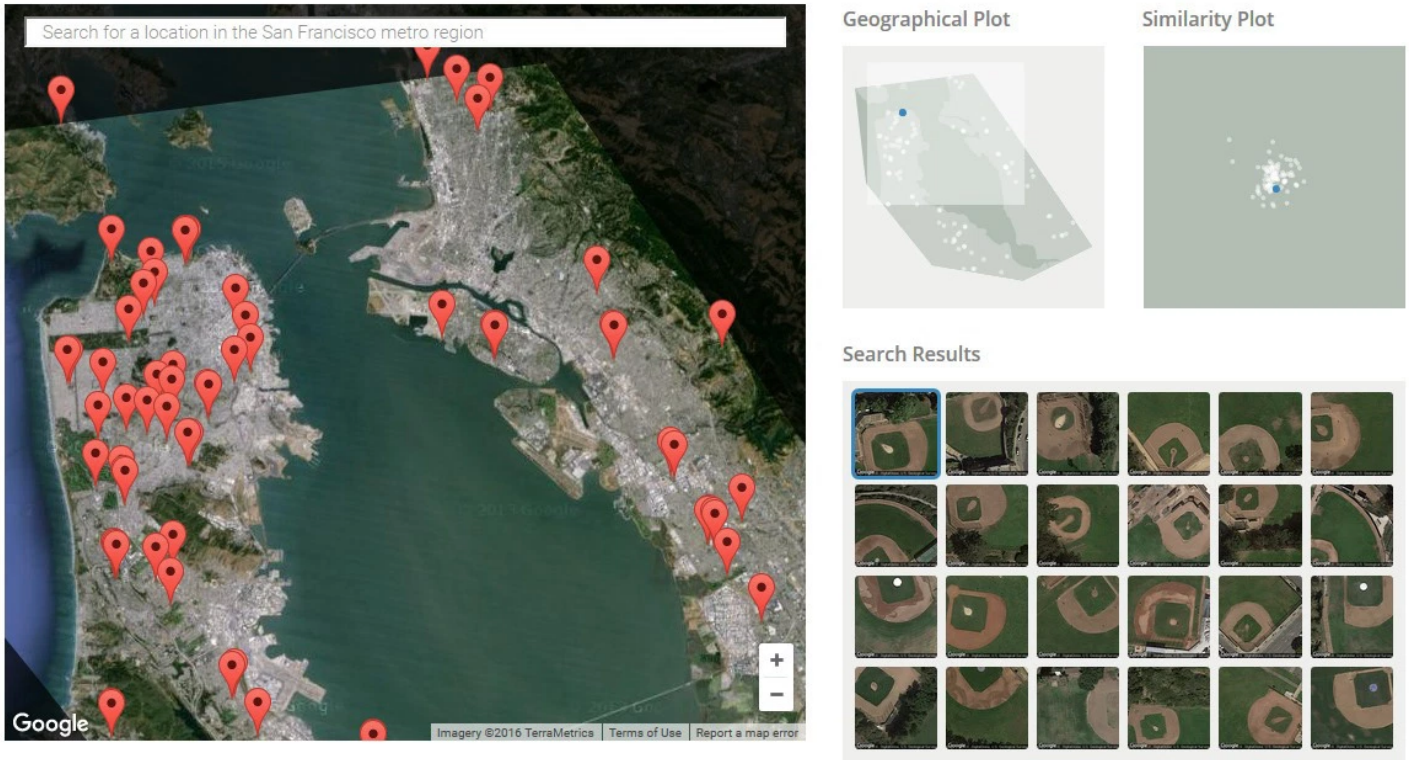

Terrapattern is a visual search engine that, from the first moment you use it, you wonder: Why didn't Google come up with this 10 years ago? Click on a feature on the map — a baseball diamond, a marina, a roundabout — and it immediately highlights everything its algorithm thinks looks like it. It's remarkably fast, simple to use and potentially very powerful.

Go ahead and give it a try first to see how natural it is to search for something. How does that work? And how did a handful of digital artists and developers create it — and for under $35,000?

The secret, as with so many other interesting visual computing projects these days, is a convolutional neural network. It's essentially an AI-like program that extracts every little detail from an image and looks for patterns at various levels of organization — similar to how our own visual system works, though the brain is infinitely more subtle and flexible.
In Terrapattern's case, the neural network was trained to look at small squares of the landscape and,

comparing those patterns to a huge database of tagged map features from OpenStreetMap, it learned to associate them with certain concepts.

Think of how a camera recognizes a face and knows when it is blinking or smiling. It doesn't actually "know" what faces, smiles and eyes are, but it associates them with certain patterns of pixels, and can reliably pick them out.

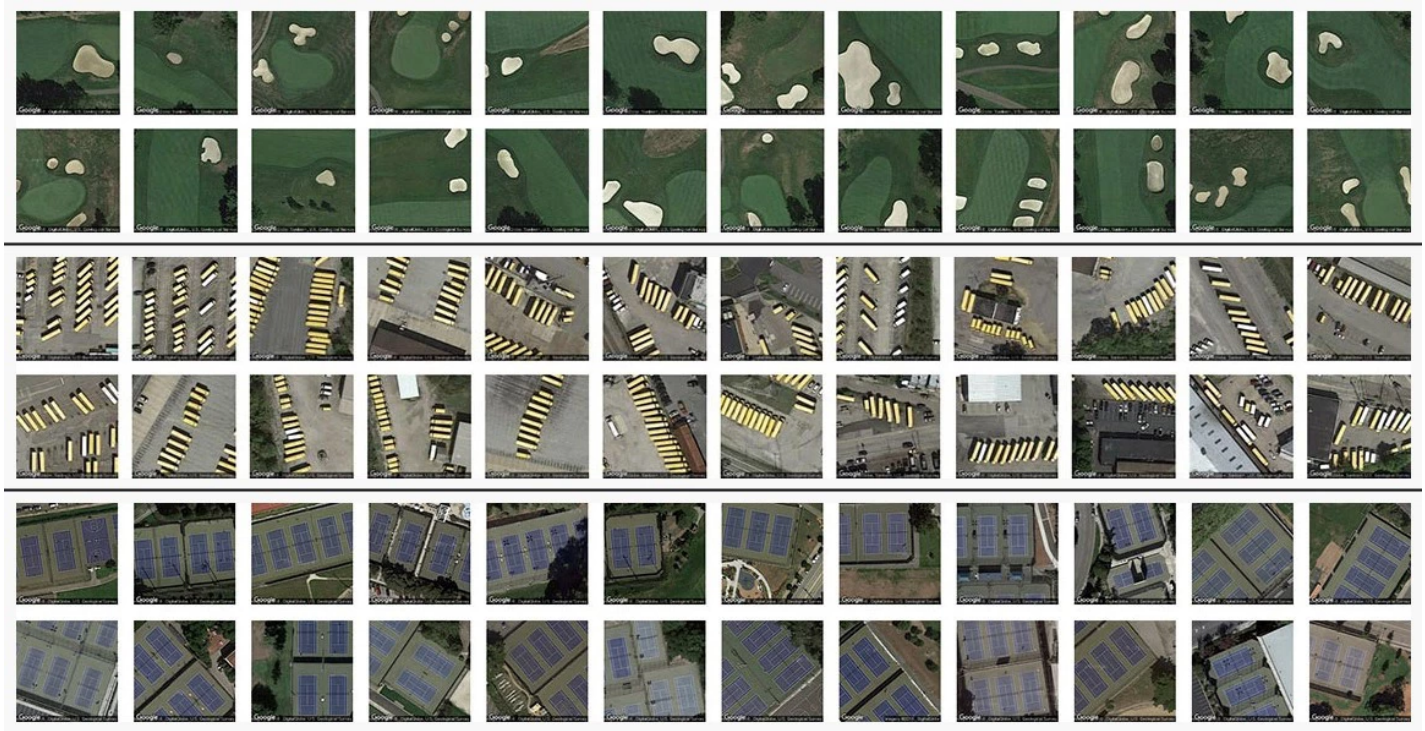

Once Terrapattern had been trained to recognize and categorize all manner of geographical features, from boats to water towers, its creators set it free on detailed maps of the greater New York, Pittsburgh, Detroit and San Francisco areas. It scoured the landscape and built a huge database of features and similarities — which can be quickly queried and the results returned immediately (the neural network isn't doing any "thinking" when you click on a feature — its work is done for this dataset).

Of course, you could just search for "tennis fields in Oakland" or the like and get perfectly good results, but this allows one to search for things that may not be listed so formally. What if you were looking for houses in the middle of fields, or cul de sacs, or dead lawns, or circular parking lots? Terrapattern knows where those are just as much as it knows where the airports and ferry terminals are. They're all just assemblages of features to the neural network.

Terrapattern was made by Golan Levin, David Newbury and Kyle McDonald, with money from the Knight Foundation's Prototype Fund. With the resources they have, they were able to map the four cities mentioned, but more are coming soon. And with luck, feature detection at higher and lower levels. It's easy to find a ballpark, but hard to find, say, four-way stops (at the small level) or prison complexes (at a larger one).

The work is free under a Creative Commons 4.0 license, and you can check out their code over at GitHub.

Google Maps meets AI: Carnegie Mellon's Terrapattern can find and map every pool in New York City - TechRepublic

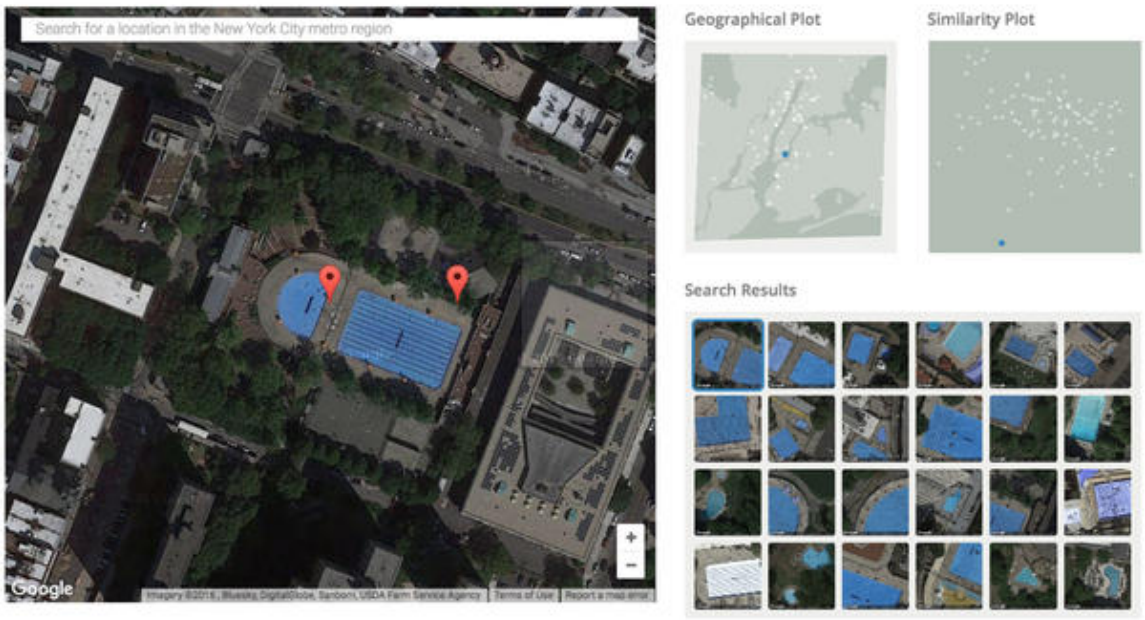

Selection of pools in New York City area

Image: screenshot, Terrapattern.com

A group of students and professors at Carnegie Mellon recently unveiled an AI tool that scans satellite photos and matches the images to similar-looking locations in the nearby geographic area. While image recognition using AI is nothing new—machine learning techniques are frequently used to identify, classify, and categorize objects, and even faces, from pools of millions of photos—the new tool called Terrapattern offers something unique: The ability to pinpoint the GPS coordinates of similar landmarks.

So, how does it work? You simply click part of a Google map, (currently, you can choose between New York City, San Francisco, Pittsburgh, and Detroit—presumably, it isn't too difficult to add new cities to the list), and the program spits out every other similar-looking geographic area in the region, as well as places the pins on a map. It does this by using a relatively low level of detail, looking at areas that are similar in terms of color, or shape.

The tool, which was created through a grant in Media Innovation from the Knight Foundation Prototype Fund, offers many uses at a personal level. Want to find an apartment near a lake? Or a school near a baseball field? You can use the tool to scan for these objects all over the city.

What else could the tool be used for? It may, one day, enable a route planner for autonomous vehicles. If, for example, you preferred driving near the coast, or wanted to avoid bridges, the features could potentially help map a route based on images, rather than going strictly by fastest route.

According to Manuela Veloso, head of Carnegie Mellon's Machine Learning Department, the program offers a valuable addition to current image recognition technology. On top of being novel, she said, "it's also extremely compelling, at the performance level. Imagine how many millions and trillions of images exist, and how much computing processing it takes to search this tremendous amount of data effectively."

And, beyond personal interest in parks or schools, Terrapattern has a plethora of uses in other areas as well. Imagine the world of statistics and data, for instance. The tool could be used not to find a pool for you to live near, but how many pools there are in a certain city. And what part of town they're in. And, perhaps, how run-down they look.

This kind of information could be relevant to all kinds of people and businesses. Sociologists. People doing inventory. Governments researching poverty or economics of a location. Developers who are planning new apartment complexes, or urban infrastructure.

Through this tool, Veloso said, "we can understand the dynamics of a region. How are things changing? Are there more tennis courts? Is the population getting richer? The images capture the economics of the country."

Terrapattern: Satellite Image Search Engine Matches Similar Places

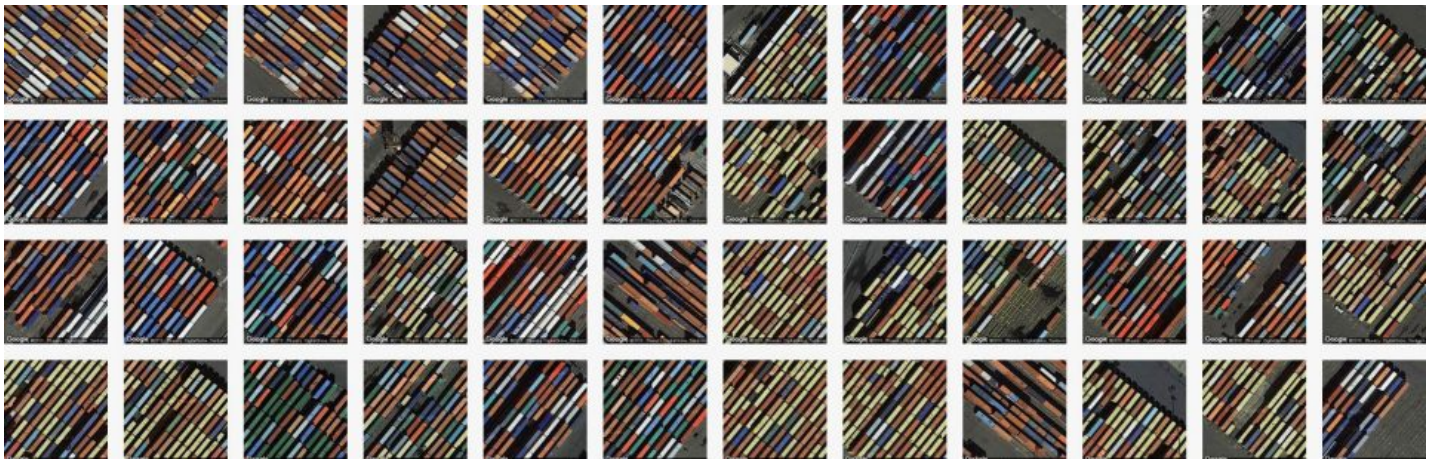

A powerful tool for artists, designers and researchers, Terrapattern lets users seek out similar-looking locations from an aerial perspective, finding connections and patterns between disparate landscapes and built environments.

The premise is simple: start with a single place, be it a park or street, stadium or shipyard, then let the tool work its magic. The results are uncanny: colors, textures and shapes tied together by computer vision and clever algorithms. The broader use cases are infinite, but specific ones are possible too, like: a user could look for abandoned ships floating around the island of Manhattan.

The system works by looking at its subjects in layers, looking for identifying features like curves, edges and shadows that indicate height. In a way, its task is similar than some pattern recognition software since it is not called upon to identify the subject, just match it.

“For our purposes,” explain the creators, “‘interesting’ features are anthropogenic or natural phenomena that are not only socially or scientifically meaningful, but also visually distinctive?thus lending themselves ideally to machine recognition. Examples could include things like animal herds, methane blowholes, factories, destroyed homes, or logging roads. Many other patterns await discovery.”

The system draws on data from OpenStreetMap, combing through hundreds of thousands of images looking for something like whatever you submitted. Researchers can use tools like this to monitor natural habitats or make archaeological finds, but ordinary people can employ this tool to create art or make inquiries about the cities they live in. Even a quick tour around the engine reveals emergent macro-patterns from individual tiles, some worthy of wall art treatment.

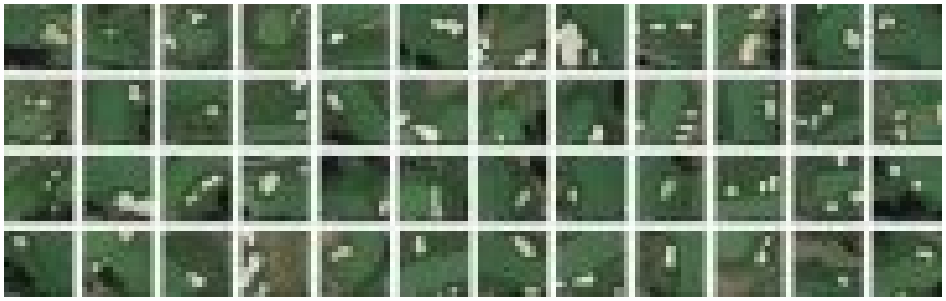

Terrapattern’s creators are indeed excited for more non-standard and unexpected uses: “Terrapattern is ideal for discovering, locating and labeling typologies that aren’t customarily indicated on maps. These might include ephemeral or temporally-contingent features (such as vehicles or construction sites), or the sorts of banal infrastructure (like fracking wells or smokestacks) that only appear on specialist blueprints, if they appear at all.”

Terrapattern search engine finds patterns in the Google Earth landscape

Imagine Google Earth injected with AI capabilities, scanning geographical regions for specific visual features and patterns.

By Eileen Brown for Social Business | July 14, 2016 -- 16:39 GMT (09:39 PDT) | Topic: Big Data Analytics

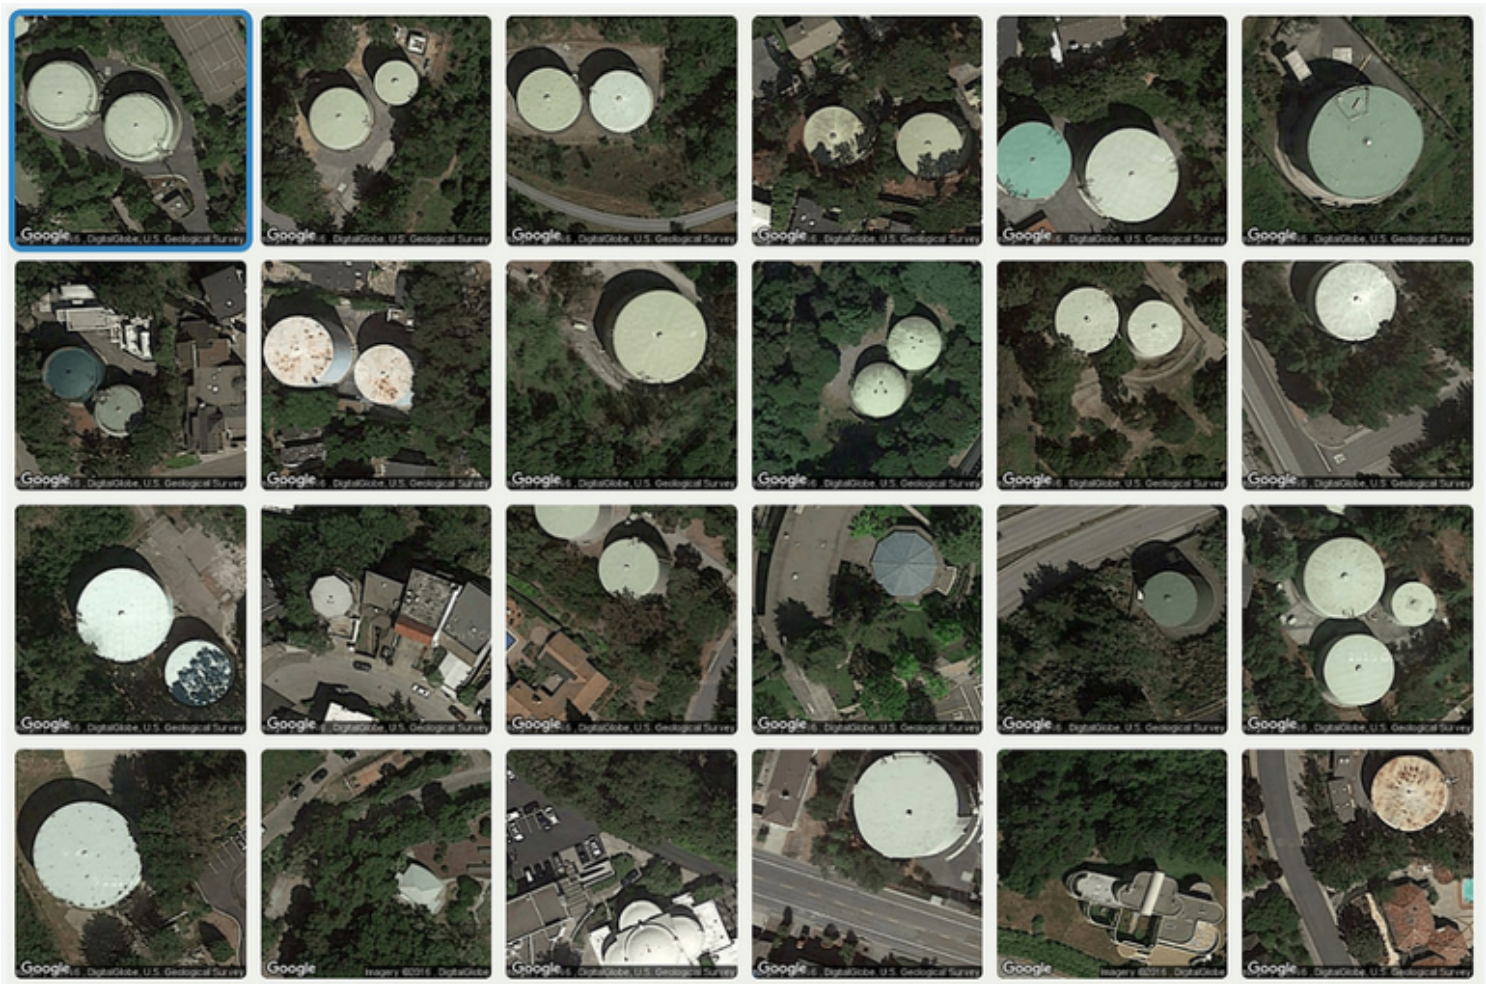

Terrapattern

You can spend a long time searching satellite images for interesting locations. Now imagine a tool that can not only show you the location, but scans large geographical areas to find specific features that are similar, and then it presents these results in a pattern-like format.

In 2002, Carnegie Mellon's (CMU) School of Computer Science launched what it calls the world's

"first PhD program in Machine Learning". It attempted to learn how to program systems to automatically learn and use its experience to improve its results.

A group of CMU professors and students have now created a visual search tool for satellite imagery called Terrapattern (http://www.terrapattern.com/).

It is an interface for finding what the team calls "more like this, please" in satellite photos. Not a company or start-up, Terrapattern is an "experimental research prototype, developed in a university setting".

It was developed at the Frank-Ratchye Studio for Creative Inquiry at Carnegie Mellon University, with support from the John S. and James L. Knight Foundation Prototype Fund.

Its creators are Golan Levin, David Newbury, and Kyle McDonald, along with Carnegie Mellon students Irene Alvarado, Aman Tiwari, and Manzil Zaheer.

Described as a "visual search engine for satellite imagery", the project uses a Deep Convolutional Neural Net (DCNN) to assist with image recognition.

The open-source, open-access project was created by a collaborative team of artists, creative technologists, and students. It is particularly useful for locating things that aren't usually indicated on maps.

(http://www.zdnet.com/article/google-earth-ventures-into-thin-air/)

Google Earth ventures 'into thin air'
(http://www.zdnet.com/article/google-earth-ventures-into-thin-air/)

In a landmark project, Google Earth has teamed up with locals to provide a fully immersive experience of the Sherpa community and their mountain home in the Sagarmatha region in Nepal, home to the tallest mountain in the world, Everest.

Read More
(http://www.zdnet.com/article/google-earth-ventures-into-thin-air/)

The project includes a list of cities, including Pittsburgh, San Francisco, New York City, Detroit, Berlin, Miami, and Austin.

Clicking an interesting spot on Terrapattern's map will find other geographical locations that look similar to your selection. You can download a list of these locations in GeoJSON format.

The tool is useful for locating specialized "non-building structures" and other forms of otherwise "unremarkable soft infrastructure" that are not usually called out on maps.

It learns which visual features are important for classifying satellite imagery and makes these features searchable.

It then computes descriptions for millions more satellite photos that cover various regions of interest in the area. Its algorithm pre-computes relationships between the descriptions, allowing searches to take just a second or two.

You can look for all the parks in an area, search for unusual swimming pool shapes, or have a look at golf courses in the region. The project is intended to present a new way of exploring and discovering "patterns of interest" to understand and organize the world.

More locations will be added to the project "soon", but for now, have a look at all the similarities in the city you choose. Unique features are not as unique as they first seem.

Related content:

- Indeed Job Search will overtake LinkedIn search soon according to report (http://www.zdnet.com/article/indeed-job-search-will-overtake-linkedin-search-soon-according-to-report/)
- Does Google's RankBrain machine learning improve search results for users? (http://www.zdnet.com/article/does-googles-rankbrain-machine-learning-improve-search-results-for-users/)
- Valossa secures funding for real time, AI powered video search technology (http://www.zdnet.com/article/valossa-secures-funding-for-real-time-ai-powered-video-search-technology/)