

# 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](#).