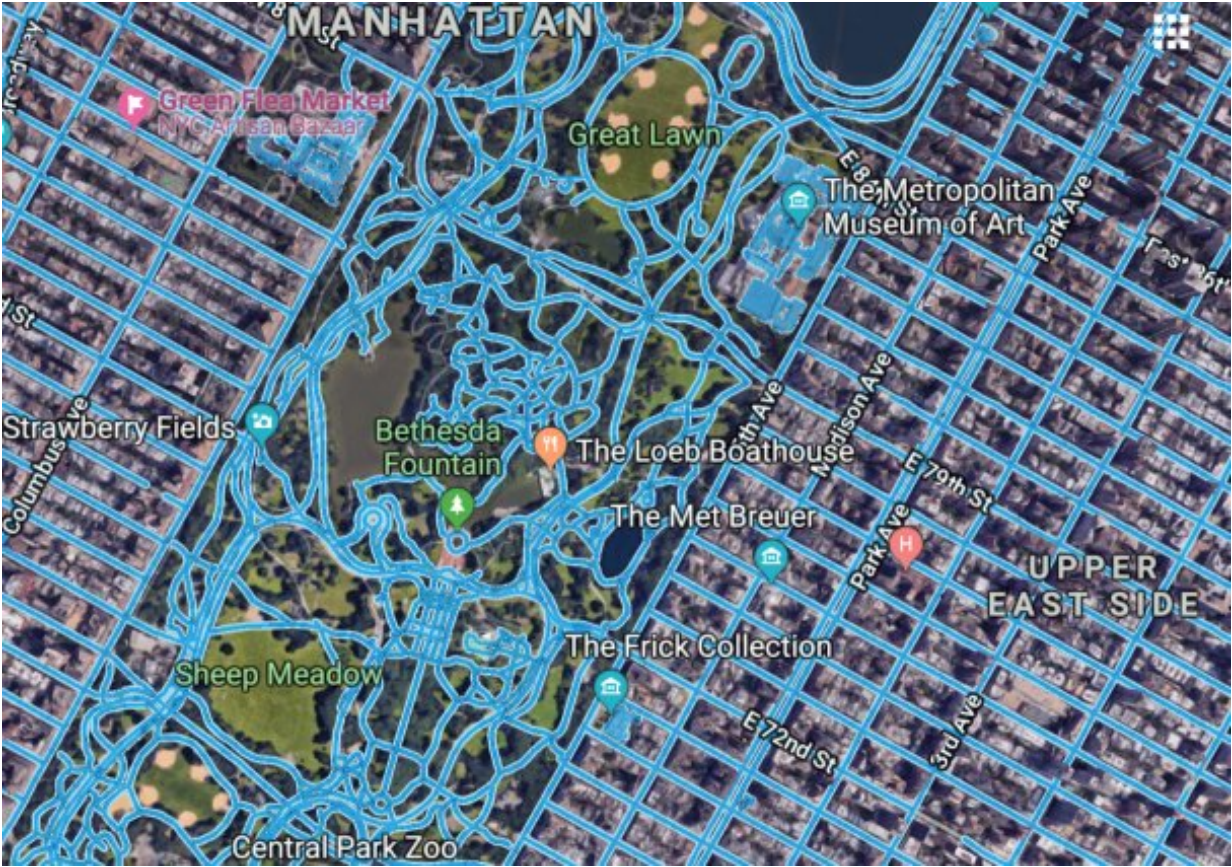# A new way to build road maps from aerial images automatically

April 17 2018, by Adam Conner-Simons



To create complete road maps companies like Google often have to have employees take existing aerial images with incomplete information and manually fill in the gaps. Credit: Massachusetts Institute of Technology

Map apps may have changed our world, but they still haven't mapped all

of it yet. Specifically, mapping roads can be difficult and tedious: Even after taking aerial images, companies still have to spend many hours manually tracing out roads. As a result, even companies like Google haven't yet gotten around to mapping the vast majority of the [more than 20 million miles of roads](#) across the globe.

Gaps in maps are a problem, particularly for systems being developed for self-driving cars. To address the issue, researchers from MIT's Computer Science and Artificial Intelligence Laboratory (CSAIL) have created RoadTracer, an automated method to build road maps that's 45 percent more accurate than existing approaches.

Using data from aerial images, the team says that RoadTracer is not just more accurate, but more cost-effective than current approaches. MIT professor Mohammad Alizadeh says this work will be useful both for tech giants like Google and for smaller organizations without the resources to curate and correct large amounts of errors in maps.

"RoadTracer is well-suited to map areas of the world where maps are frequently out of date, which includes both places with lower population and areas where there's frequent construction," says Alizadeh, one of the co-authors of a new paper about the system. "For example, existing maps for remote areas like rural Thailand are missing many roads. RoadTracer could help make them more accurate."

For example, looking at aerial images of New York City, RoadTracer could correctly map 44 percent of its road junctions, which is more than twice as effective as traditional approaches based on image segmentation that could map only 19 percent.

The paper, which will be presented in June at the Conference on Computer Vision and Pattern Recognition (CVPR) in Salt Lake City, Utah, is a collaboration between CSAIL and the Qatar Computing

Research Institute (QCRI).

Alizadeh's MIT co-authors include graduate students Fayven Bastani and Songtao He, and professors Hari Balakrishnan, Sam Madden, and David DeWitt. QCRI co-authors include senior software engineer Sofiane Abbar and Sanjay Chawla, who is the research director of QCRI's Data Analytics Group.

Current efforts to automate maps involve training neural networks to look at aerial images and identify individual pixels as either "road" or "not road." Because aerial images can often be ambiguous and incomplete, such systems also require a post-processing step that's aimed at trying to fill in some of the gaps.

Unfortunately, these so-called "segmentation" approaches are often imprecise: If the model mislabels a pixel, that error will get amplified in the final road map. Errors are particularly likely if the aerial images have trees, buildings, or shadows that obscure where roads begin and end. (The post-processing step also requires making decisions based on assumptions that may not always hold up, like connecting two road segments simply because they are next to each other.)

Meanwhile, RoadTracer creates maps step-by-step. It starts at a known location on the road network, and uses a neural network to examine the surrounding area to determine which point is most likely to be the next part on the road. It then adds that point and repeats the process to gradually trace out the road network one step at a time.

"Rather than making thousands of different decisions at once about whether various pixels represent parts of a road, RoadTracer focuses on the simpler problem of figuring out which direction to follow when starting from a particular spot that we know is a road," says Bastani. "This is in many ways actually a lot closer to how we as humans

construct mental models of the world around us."

The team trained RoadTracer on aerial images of 25 cities across six countries in North America and Europe, and then evaluated its mapping abilities on 15 other cities.
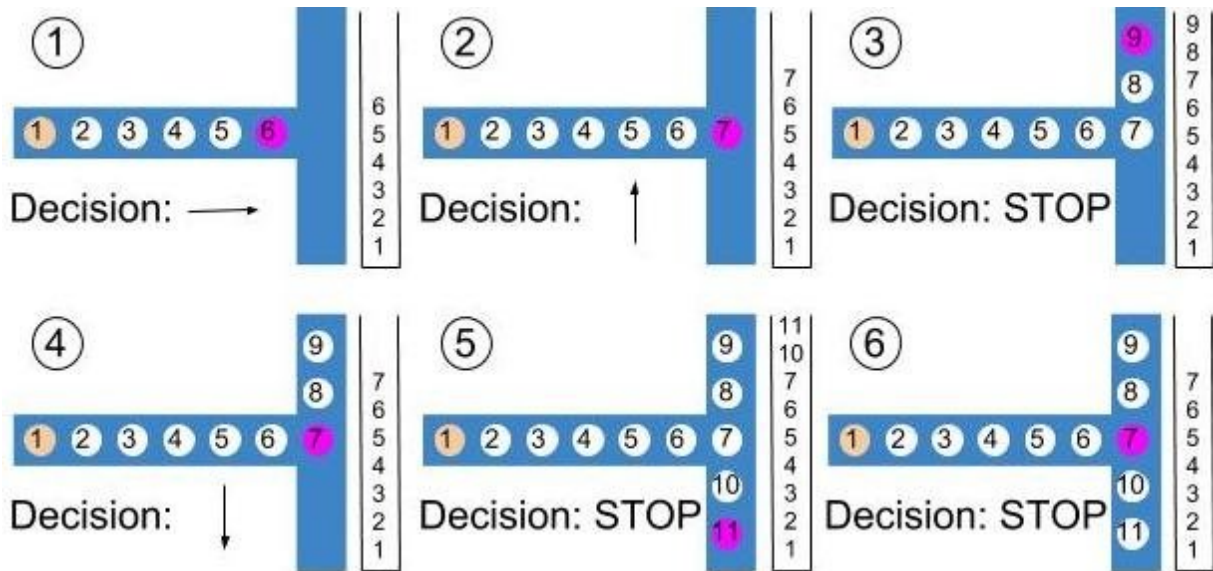


To create a road map, the system starts at a known location and uses a neural network to examine the surrounding area to determine which point is most likely to be the next part on the road. It then adds that point and repeats the process to gradually trace out the road one step at a time. Credit: Massachusetts Institute of Technology

"It's important for a mapping system to be able to perform well on cities it hasn't trained on, because regions where automatic mapping holds the most promise are ones where existing maps are non-existent or inaccurate," says Balakrishnan.

Bastani says that the fact that RoadTracer had an error rate that is 45

percent lower is essential to making automatic mapping systems more practical for companies like Google.

"If the error rate is too high, then it is more efficient to map the roads manually from scratch versus removing incorrect segments from the inferred map," says Bastani.

Still, implementing something like RoadTracer wouldn't take people completely out of the loop: The team says that they could imagine the system proposing [road](#) maps for a large region and then having a human expert come in to double-check the design.

"That said, what's clear is that with a system like ours you could dramatically decrease the amount of tedious work that humans would have to do," Alizadeh says.

Indeed, one advantage to RoadTracer's incremental approach is that it makes it much easier to correct errors; human supervisors can simply correct them and re-run the algorithm from where they left off, rather than continue to use imprecise information that trickles down to other parts of the map.

Of course, aerial images are just one piece of the puzzle. They don't give you information about roads that have overpasses and underpasses, since those are impossible to ascertain from above. As a result, the team is also separately developing algorithms that can create maps from GPS data, and working to merge these approaches into a single system for mapping.

*This story is republished courtesy of MIT News ([web.mit.edu/newsoffice/](http://web.mit.edu/newsoffice/)), a popular site that covers news about MIT research, innovation and teaching.*