

Using artificial intelligence to enrich digital maps

January 23 2020, by Rob Matheson



An AI model developed at MIT and Qatar Computing Research Institute that uses only satellite imagery to automatically tag road features in digital maps could improve GPS navigation, especially in countries with limited map data. Credit: Google Maps/MIT News

A model invented by researchers at MIT and Qatar Computing Research



Institute (QCRI) that uses satellite imagery to tag road features in digital maps could help improve GPS navigation.

Showing drivers more details about their routes can often help them navigate in unfamiliar locations. Lane counts, for instance, can enable a GPS system to warn drivers of diverging or merging lanes. Incorporating information about parking spots can help drivers plan ahead, while mapping bicycle lanes can help cyclists negotiate busy city streets. Providing updated information on <u>road conditions</u> can also improve planning for disaster relief.

But creating detailed maps is an expensive, time-consuming process done mostly by big companies, such as Google, which sends vehicles around with cameras strapped to their hoods to capture video and images of an area's roads. Combining that with other data can create accurate, up-to-date maps. Because this process is expensive, however, some parts of the world are ignored.

A solution is to unleash <u>machine-learning models</u> on <u>satellite images</u> —which are easier to obtain and updated fairly regularly—to automatically tag <u>road</u> features. But roads can be occluded by, say, trees and buildings, making it a challenging task. In a paper being presented at the Association for the Advancement of Artificial Intelligence conference, the MIT and QCRI researchers describe "RoadTagger," which uses a combination of neural network architectures to automatically predict the number of lanes and road types (residential or highway) behind obstructions.

In testing RoadTagger on occluded roads from digital maps of 20 U.S. cities, the model counted lane numbers with 77 percent accuracy and inferred road types with 93 percent accuracy. The researchers are also planning to enable RoadTagger to predict other features, such as parking spots and bike lanes.



"Most updated digital maps are from places that big companies care the most about. If you're in places they don't care about much, you're at a disadvantage with respect to the quality of map," says co-author Sam Madden, a professor in the Department of Electrical Engineering and Computer Science (EECS) and a researcher in the Computer Science and Artificial Intelligence Laboratory (CSAIL). "Our goal is to automate the process of generating high-quality digital maps, so they can be available in any country."

The paper's co-authors are CSAIL graduate students Songtao He, Favyen Bastani, and Edward Park; EECS undergraduate student Satvat Jagwani; CSAIL professors Mohammad Alizadeh and Hari Balakrishnan; and QCRI researchers Sanjay Chawla, Sofiane Abbar, and Mohammad Amin Sadeghi.

Combining CNN and GNN

Quatar, where QCRI is based, is "not a priority for the large companies building digital maps," Madden says. Yet, it's constantly building new roads and improving old ones, especially in preparation for hosting the 2022 FIFA World Cup.

"While visiting Qatar, we've had experiences where our Uber driver can't figure out how to get where he's going, because the map is so off," Madden says. "If navigation apps don't have the right information, for things such as lane merging, this could be frustrating or worse."

RoadTagger relies on a novel combination of a convolutional neural network (CNN)—commonly used for images-processing tasks—and a graph neural network (GNN). GNNs model relationships between connected nodes in a graph and have become popular for analyzing things like social networks and molecular dynamics. The model is "endto-end," meaning it's fed only raw data and automatically produces



output, without human intervention.

The CNN takes as input raw satellite images of target roads. The GNN breaks the road into roughly 20-meter segments, or "tiles." Each tile is a separate graph node, connected by lines along the road. For each node, the CNN extracts road features and shares that information with its immediate neighbors. Road information propagates along the whole graph, with each node receiving some information about road attributes in every other node. If a certain tile is occluded in an image, RoadTagger uses information from all tiles along the road to predict what's behind the occlusion.

This combined architecture represents a more human-like intuition, the researchers say. Say part of a four-lane road is occluded by trees, so certain tiles show only two lanes. Humans can easily surmise that a couple lanes are hidden behind the trees. Traditional machine-learning models—say, just a CNN—extract features only of individual tiles and most likely predict the occluded tile is a two-lane road.

"Humans can use information from adjacent tiles to guess the number of lanes in the occluded tiles, but networks can't do that," He says. "Our approach tries to mimic the natural behavior of humans, where we capture local information from the CNN and global information from the GNN to make better predictions."

Learning weights

To train and test RoadTagger, the researchers used a real-world map dataset, called OpenStreetMap, which lets users edit and curate digital maps around the globe. From that dataset, they collected confirmed road attributes from 688 square kilometers of maps of 20 U.S. cities—including Boston, Chicago, Washington, and Seattle. Then, they gathered the corresponding satellite images from a Google Maps dataset.



In training, RoadTagger learns weights—which assign varying degrees of importance to features and node connections—of the CNN and GNN. The CNN extracts features from pixel patterns of tiles and the GNN propagates the learned features along the graph. From randomly selected subgraphs of the road, the system learns to predict the road features at each tile. In doing so, it automatically learns which image features are useful and how to propagate those features along the graph. For instance, if a target tile has unclear lane markings, but its neighbor tile has four lanes with clear lane markings and shares the same road width, then the target tile is likely to also have four lanes. In this case, the model automatically learns that the road width is a useful image feature, so if two adjacent tiles share the same road width, they're likely to have the same lane count.

Given a road not seen in training from OpenStreetMap, the model breaks the road into tiles and uses its learned weights to make predictions. Tasked with predicting a number of lanes in an occluded tile, the model notes that neighboring tiles have matching pixel patterns and, therefore, a high likelihood to share information. So, if those tiles have four lanes, the occluded tile must also have four.

In another result, RoadTagger accurately predicted lane numbers in a dataset of synthesized, highly challenging road disruptions. As one example, an overpass with two lanes covered a few tiles of a target road with four lanes. The model detected mismatched pixel patterns of the overpass, so it ignored the two lanes over the covered tiles, accurately predicting four lanes were underneath.

The researchers hope to use RoadTagger to help humans rapidly validate and approve continuous modifications to infrastructure in datasets such as OpenStreetMap, where many maps don't contain lane counts or other details. A specific area of interest is Thailand, Bastani says, where roads are constantly changing, but there are few if any updates in the dataset.



"Roads that were once labeled as dirt roads have been paved over so are better to drive on, and some intersections have been completely built over. There are changes every year, but <u>digital maps</u> are out of date," he says. "We want to constantly update such road attributes based on the most recent imagery."

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

Provided by Massachusetts Institute of Technology

Citation: Using artificial intelligence to enrich digital maps (2020, January 23) retrieved 27 April 2024 from <u>https://techxplore.com/news/2020-01-artificial-intelligence-enrich-digital.html</u>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.