

Deep-learning algorithm can weigh up a neighborhood better than humans

September 25 2014, by Adam Conner-Simons



A new algorithm can outperform humans at predicting which of a series of photos is taken in a higher-crime area, or is closer to a McDonald's restaurant.

Human beings have a remarkable ability to make inferences based on their surroundings. Is this area safe? Where might I find a parking spot? Am I more likely to get to a gas station by taking a left or a right at this stoplight?

Such decisions require us to look beyond our "visual scene" and weigh an exceedingly complex set of understandings and real-time judgments. This begs the question: Can we teach computers to "see" in the same way? And once we teach them, can they do it better than we can?

The answers are "yes" and "sometimes," according to research out of MIT's Computer Science and Artificial Intelligence Laboratory (CSAIL). Researchers have developed an algorithm that can look at a pair of [photos](#) and outperform humans in determining things like which scene has a higher crime rate, or is closer to a McDonald's restaurant.

An [online demo](#) puts you in the middle of a Google Street View with four directional options and challenges you to navigate to the nearest McDonald's in the fewest possible steps.

While humans are generally better at this specific task than the algorithm, the researchers found that the [computer](#) consistently outperformed humans at a variation of the task in which users are shown two photos and asked which scene is closer to a McDonald's.

To create the algorithm, the team—which included PhD students Aditya Khosla, Byoungkwon An, and Joseph Lim, as well as CSAIL principal investigator Antonio Torralba—trained the computer on a set of 8 million Google images from eight major U.S. cities that were embedded with GPS data on crime rates and McDonald's locations. They then used deep-learning techniques to help the program teach itself how different qualities of the photos correlate. For example, the [algorithm](#) independently discovered that some things you often find near McDonald's franchises include taxis, police vans, and prisons. (Things you don't find: cliffs, suspension bridges, and sandbars.)

"These sorts of algorithms have been applied to all sorts of content, like inferring the memorability of faces from headshots," said Khosla. "But

before this, there hadn't really been research that's taken such a large set of photos and used it to predict qualities of the specific locations the photos represent."

The researchers presented a paper about the work at the IEEE International Conference on Computer Vision and Pattern Recognition (CVPR) this summer.

While the project was mostly intended as proof that computer algorithms are capable of advanced scene understanding, Khosla has brainstormed potential uses ranging from a navigation app that avoids high-crime areas, to a tool that could help McDonald's determine future franchise locations.

Khosla previously helped develop an [algorithm that can predict a photo's popularity](#).

More information: The study is available online: mcdonalds.csail.mit.edu/

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: Deep-learning algorithm can weigh up a neighborhood better than humans (2014, September 25) retrieved 19 April 2024 from <https://phys.org/news/2014-09-deep-learning-algorithm-neighborhood-humans.html>

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.