


New system estimates geographic location of photos

June 18 2008



A method developed by Carnegie Mellon University researchers can estimate where a photo was taken by matching it to similar, GPS-tagged photos in the Flickr online photo collection. Credit: James Hays/Carnegie Mellon University

Researchers at Carnegie Mellon University have devised the first computerized method that can analyze a single photograph and determine where in the world the image likely was taken. It's a feat made possible by searching through millions of GPS-tagged images in the Flickr online photo collection.

The IM2GPS algorithm developed by computer science graduate student James Hays and Alexei A. Efros, assistant professor of computer science and robotics, doesn't attempt to scan a photo for location clues, such as types of clothing, the language on street signs, or specific types of vegetation, as a person might do. Rather, it analyzes the composition of

the photo, notes how textures and colors are distributed and records the number and orientation of lines in the photo. It then searches Flickr for photos that are similar in appearance.

"We're not asking the computer to tell us what is depicted in the photo but to find other photos that look like it," Efros said. "It was surprising to us how effective this approach proved to be. Who would have guessed that similarity in overall image appearance would correlate to geographic proximity so well?"

Hays and Efros found they could accurately geolocate the images within 200 kilometers for 16 percent of more than 200 photos in their test set — up to 30 times better than chance. And even if their algorithm failed to identify the specific location, they often found that it could narrow the possibilities, such as by identifying the locale as a beach or a desert.

"It seems there's not as much ambiguity in the visual world as you might guess," said Hays, who will present the research at the IEEE Computer Society Conference on Computer Vision and Pattern Recognition June 24-26 in Anchorage, Alaska. "Estimating geographic information from images is a difficult, but very much a doable, computer vision problem."

Identifying the locale of a photo could enhance image search techniques, making them less dependent on captions or associated text. A computer system for geolocating photos could be useful in finding family photos from a specific trip and in some forensic applications. Determining the location of photos also makes it possible to combine them with geographic data bases related to climate, population density, vegetation, topography and land use.

Knowing the locale also can aid in such computer vision tasks as object identification, Hays said. If a computer recognizes that a photo likely was taken in Japan, for instance, the computer will have a better idea of

what a taxicab should look like.

Hays said many online photos have some sort of geographic label, but these human descriptions can often be incorrect, or overly broad, such as a photo of the Grand Canyon labeled "U.S." The growing number of online photos that have GPS tags, by contrast, are unambiguous regarding their location, even though many are photos of rooms, people or events such as birthday parties that are useless for geolocation tasks. By using photos with both geographic keywords and GPS coordinates, Hays and Efros were able to find more than six million photos that were useful and accurately geolocated.

The IM2GPS algorithm readily located photographs of such landmarks as the Cathedral of Notre Dame in Paris. More surprisingly, it was able to recognize that a narrow street in Barcelona was typical of Mediterranean villages, rather than an American alleyway.

But some odd matches also occurred. The architecturally unique Sydney Opera House seemed to the computer to be similar to a hotel in Mississippi as well as a bridge in London. A shot of the Eiffel Tower at dusk was matched to other Eiffel Tower shots, but also to San Francisco's Coit Tower and New York's Statue of Liberty, both shot at dusk.

One reason for this confusion, Hays explained, is that the algorithm is not designed to recognize specific objects so much as it is to recognize geographic areas. For instance, an image of Utah's Monument Valley caused the IM2GPS algorithm to successfully retrieve a number of other images from Monument Valley and the American Southwest, rather than images of a specific rock formation.

For more information, see the IM2GPS project Web site:
graphics.cs.cmu.edu/projects/im2gps/

Source: Carnegie Mellon University

Citation: New system estimates geographic location of photos (2008, June 18) retrieved 19 April 2024 from <https://phys.org/news/2008-06-geographic-photos.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.