


Algorithm for identifying object boundaries in digital images 50,000 times more efficient than predecessor

May 31 2011, by Larry Hardesty



Courtesy of Jason Chang

Determining the boundaries of objects is one of the central problems in computer vision. It's something humans do with ease: We glance out the window and immediately see cars as distinct from the sidewalk and street and the people walking by, or lampposts as distinct from the facades of the buildings behind them. But duplicating that facility in silicon has proven remarkably difficult.

One of the best ways for a computer to determine [boundaries](#) is to make lots of guesses and compare them; the boundaries that most of the guesses agree on are likeliest to be accurately drawn. Until now, that process has been monstrously time consuming. But Jason Chang, a

graduate student in the Department of [Electrical Engineering](#) and Computer Science, and John Fisher, a principal research scientist at MIT's [Computer Science](#) and Artificial Intelligence Lab (CSAIL), have figured out how to make it at least 50,000 times more efficient. Their findings could help improve systems for [medical imaging](#), for tracking [moving objects](#) and for 3-D [object-recognition](#), among others.

One reason that boundary determination — or as it's more commonly known, image segmentation — is such a hard problem is that there's no one right answer. Ask 10 people to trace the boundaries of objects in a digital image, and you'll likely get 10 different responses. "We want an algorithm that's able to segment images like humans do," Chang says. "But because humans segment images differently, we shouldn't come up with one segmentation. We should come up with a lot of different segmentations that kind of represent what humans would also segment."

Populating the field

To generate its set of candidate segmentations, Chang and Fisher's algorithm strikes different balances between two measures of segmentation quality. One measure is the difference between the parts of the image on opposite sides of each boundary. The most obvious way to gauge difference is by color value: A segmentation with blue pixels on one side of the boundary and red pixels on the other would be better than one that featured slightly different proportions of the same 30 shades of blue on both sides. But researchers have devised other, more subtle measures of difference, and Chang and Fisher's algorithm can use any of them.



Courtesy of Jason Chang

The other measure of a segmentation's quality is its simplicity. If a computer is trying to segment a street scene that features a car, for instance, you probably don't want it to draw boundaries around every separate gleam of different-colored light on the car's hood. Simplicity and difference in appearance tend to be competing measures: It's easy to maximize color difference, for instance, if you draw a boundary around every pixel that's different from its neighbors, but such a segmentation would be ridiculously complex.

Chang and Fisher's algorithm assigns each segmentation a total score based on both simplicity and difference in appearance. But different segmentations could have roughly equivalent total scores: A segmentation that's a little too complex but has superb color difference, for instance, could have the same score as a segmentation that's pretty good on both measures. Chang and Fisher's algorithm is designed to find candidates with very high total scores. That ensures that none of the candidates will be outrageously bad, but it also makes the computation

that much more complicated.

Other researchers have adopted the same general approach, but to generate their candidates, they adapted algorithms originally designed to find the one segmentation with the highest total score. Chang and Fisher realized that, because they were considering so many suboptimal possibilities anyway, they could use a less precise algorithm that runs much more efficiently. Although it's not essential to their approach that they find the highest-scoring segmentation, it's still likely that the many candidates they produce will include a few that are very close to it.

"There are a lot of competing methodologies out there, so it's hard for me to say that this is going to revolutionize segmentation," says Anthony Yezzi, a professor of electrical and computer engineering at the Georgia Institute of Technology. But, Yezzi says, the way in which Chang and Fisher's algorithm represents images is "an interesting new vehicle that I think could get a lot of mileage, even beyond [segmentation](#)." The same technique, Yezzi says, could be applied to problems of object tracking — whether it's the motion of an object in successive frames of video or changes in a tumor's size over time — and pattern matching, where the idea is to recognize the similarity of objects depicted from slightly different angles or under different lighting conditions.

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: Algorithm for identifying object boundaries in digital images 50,000 times more efficient than predecessor (2011, May 31) retrieved 26 April 2024 from

<https://phys.org/news/2011-05-algorithm-boundaries-digital-images-efficient.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.