

Students develop algorithm that categorizes user-generated restaurant photos

8 June 2016, by Adam Zewe



TripAdvisor features millions of user-generated restaurant meal photos like this one, but there is no rhyme or reason to the photos that appear when a user clicks on a restaurant listing.

Whether they are surfing the web in search of surf and turf, or browsing for the perfect grass-fed beef burger, more than 90 percent of U.S. consumers search online for restaurants, according to a study by market research firm Chadwick Martin Bailey.

Helping tech-savvy, would-be diners get a better picture of what might actually end up on their plates was the aim of a group of students in the computational science and engineering master's program offered by the Institute for Applied Computational Science (IACS) at the Harvard John A. Paulson School of Engineering and Applied Sciences. They worked with TripAdvisor to make the travel website's restaurant photos more relevant for visitors.

Virgile Audi, S.M. '16, Crystal Lim, S.M. '16, Reinier Maat, S.M. '16, and Leonhard Spiegelberg, S.M. '16, collaborated with representatives from Needham, Massachusetts-based TripAdvisor for their final project in the masters program, supervised by Pavlos Protopapas, IACS scientific

director.

TripAdvisor [images](#), the bulk of which are uploaded directly by the worldwide community of reviewers, are not currently sorted by the website, so there is no rhyme or reason to the photos that appear when a user clicks on a restaurant listing.

"Our ultimate goal is to present the user with a very curated list of authentic, real images that give a good impression of what the restaurant and the food actually look like," Spiegelberg explained.

To accomplish that, the students developed a classifier [algorithm](#) that automatically sorts uploaded photos into one of five categories: food, drinks, restaurant interior, restaurant exterior, and menus.

The team developed the algorithm using a technique called convolutional neural networks, which is inspired by the neural networks inside the visual cortex of the human brain. The algorithm examines an image, maps all the pixels, and determines which of the five categories it falls into using clues such as light intensity, edges, and colors.



The team's algorithm classified user-generated restaurant photos into one of five categories. This photo was especially difficult to classify, since it could be placed in two categories (drinks and exterior). Credit: Virgile Audi

But the [machine-learning algorithm](#) needed some help from the students to get started, in the form of a training set of pre-classified images that would serve as a baseline. After debating how best to label 20,000 unclassified TripAdvisor images for the training set, the students chose to rely on the competitive spirit of their classmates.

"We crowd-sourced the labeling process," Maat explained. "We created a competitive platform where users were presented with an image that they were asked to label into one of the categories. It was a race to label the most images, and we created a real-time leader board so our classmates could see how they were doing during the competition."

This enterprising delegation paid off; the students quickly compiled a collection of 20,000 labeled images that they used to initially train the algorithm.

"This is a technique we call 'supervised machine learning,'" Audi said. "The pictures that confuse the algorithm, we looked at them and manually labeled

them, and then the algorithm starts to re-learn based on the pictures that were manually added."

The algorithm learns the most when an image is on the borderline between two categories, so it is important to manually classify those borderline images correctly, Maat explained. One of the biggest challenges the students faced was figuring out how to classify images that don't fit any of the categories well, or whether to simply exclude those pictures.

As they worked, they found that the algorithm was so complex that it took more than seven hours to run on a laptop. Their contacts at TripAdvisor came to the rescue, lending them one of the best graphical video cards available, which reduced the run-time to about 20 minutes.

After training and refining their classification algorithm, the students were able to automatically sort images with 87 percent accuracy. They are proud of what they accomplished, and excited that TripAdvisor has expressed interest in continuing to work on their algorithm, Audi said.

"This project felt a lot like a startup," said Lim. "There are so many different aspects to this problem. It was very interesting and rewarding to see how applicable all the things we have learned in class are to a real-world situation."

Provided by Harvard University

APA citation: Students develop algorithm that categorizes user-generated restaurant photos (2016, June 8) retrieved 27 January 2022 from <https://techxplore.com/news/2016-06-students-algorithm-categorizes-user-generated-restaurant.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.