

Deep-learning algorithm recommends ingredients and recipes based on a photo of food

July 21 2017, by Adam Conner-Simons



Pic2Recipe, an artificial intelligence system developed at MIT, can take a photo of an entree and suggest a similar recipe to it. Credit: Jason Dorfman/MIT CSAIL



There are few things social media users love more than flooding their feeds with photos of food. Yet we seldom use these images for much more than a quick scroll on our cellphones.

Researchers from MIT's Computer Science and Artificial Intelligence Laboratory (CSAIL) believe that analyzing photos like these could help us learn recipes and better understand people's eating habits. In a new paper with the Qatar Computing Research Institute (QCRI), the team trained an <u>artificial intelligence</u> system called <u>Pic2Recipe</u> to look at a photo of food and be able to predict the <u>ingredients</u> and suggest similar recipes.

"In computer vision, food is mostly neglected because we don't have the large-scale datasets needed to make predictions," says Yusuf Aytar, an MIT postdoc who co-wrote a paper about the system with MIT Professor Antonio Torralba. "But seemingly useless photos on <u>social media</u> can actually provide valuable insight into health habits and dietary preferences."

The paper will be presented later this month at the Computer Vision and Pattern Recognition conference in Honolulu. CSAIL graduate student Nick Hynes was lead author alongside Amaia Salvador of the Polytechnic University of Catalonia in Spain. Co-authors include CSAIL postdoc Javier Marin, as well as scientist Ferda Ofli and research director Ingmar Weber of QCRI.

How it works

The web has spurred a huge growth of research in the area of classifying food data, but the majority of it has used much smaller datasets, which often leads to major gaps in labeling foods.

In 2014 Swiss researchers created the "Food-101" dataset and used it to



develop an algorithm that could recognize images of food with 50 percent accuracy. Future iterations only improved accuracy to about 80 percent, suggesting that the size of the dataset may be a limiting factor.

Even the larger datasets have often been somewhat limited in how well they generalize across populations. A database from the City University in Hong Kong has over 110,000 images and 65,000 recipes, each with ingredient lists and instructions, but only contains Chinese cuisine.

The CSAIL team's project aims to build off of this work but dramatically expand in scope. Researchers combed websites like All Recipes and Food.com to develop "Recipe1M," a database of over 1 million recipes that were annotated with information about the ingredients in a wide range of dishes. They then used that data to train a neural network to find patterns and make connections between the food images and the corresponding ingredients and recipes.





Pic2Recipe predicts recipes from photos using a neural network, a way to achieve machine learning in which a computer learns to perform some task by analyzing examples. Credit: Jason Dorfman/MIT CSAIL

Given a photo of a food item, Pic2Recipe could identify ingredients like flour, eggs, and butter, and then suggest several recipes that it determined to be similar to images from the database. (The team has an online demo where people can upload their own food photos to test it out.)

"You can imagine people using this to track their daily nutrition, or to photograph their meal at a restaurant and know what's needed to cook it at home later," says Christoph Trattner, an assistant professor at MODUL University Vienna in the New Media Technology Department who was not involved in the paper. "The team's approach works at a similar level to human judgement, which is remarkable."

The system did particularly well with desserts like cookies or muffins, since that was a main theme in the database. However, it had difficulty determining ingredients for more ambiguous foods, like sushi rolls and smoothies.

It was also often stumped when there were similar recipes for the same dishes. For example, there are dozens of ways to make lasagna, so the team needed to make sure that system wouldn't "penalize" recipes that are similar when trying to separate those that are different. (One way to solve this was by seeing if the ingredients in each are generally similar before comparing the recipes themselves).

In the future, the team hopes to be able to improve the system so that it can understand food in even more detail. This could mean being able to



infer how a food is prepared (i.e. stewed versus diced) or distinguish different variations of foods, like mushrooms or onions.

The researchers are also interested in potentially developing the system into a "dinner aide" that could figure out what to cook given a dietary preference and a list of items in the fridge.

"This could potentially help people figure out what's in their <u>food</u> when they don't have explicit nutritional information," says Hynes. "For example, if you know what ingredients went into a dish but not the amount, you can take a photo, enter the ingredients, and run the model to find a similar <u>recipe</u> with known quantities, and then use that information to approximate your own meal."

More information: Learning Cross-modal Embeddings for Cooking Recipes and Food Images: <u>im2recipe.csail.mit.edu/im2recipe.pdf</u>

Provided by Massachusetts Institute of Technology

Citation: Deep-learning algorithm recommends ingredients and recipes based on a photo of food (2017, July 21) retrieved 3 May 2024 from <u>https://techxplore.com/news/2017-07-deep-learning-algorithm-ingredients-recipes-based.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.