


PizzaGAN gets the picture on how to make a pizza

20 June 2019, by Nancy Cohen



Is nothing sacred? Who would dare to even attempt to talk about a machine-learning experiment that results in the perfect (gasp) pizza? It is difficult to contemplate, but a research quintet did not shy away from trying, and they worked to teach a machine how to make a great pie.

Say hello to PizzaGAN, a compositional layer-based generative model that was aimed to mirror the step-by-step procedure of pizza-making.

Their goal was to teach the machine by building a generative [model](#) that mirrors an ordered set of instructions. How they proceeded: "Each operator is designed as a Generative Adversarial Network (GAN). Given only weak image-level supervision, the operators are trained to generate a visual layer that needs to be added to or removed from the existing image. The proposed model is able to decompose an image into an ordered sequence of layers by applying sequentially in the right order the corresponding removing modules."

(Generative adversarial networks can [do](#) a lot of

things, Victoria Song remarked in *Gizmodo*. She said it was "basically the type of machine learning used to generate realistic AI faces and deepfakes.")

Results? Suffice to say they reported making a model to their satisfaction. "Experimental results on synthetic and real pizza images demonstrate that our proposed model is able to: (1) segment pizza toppings in a weakly-supervised fashion, (2) remove them by revealing what is occluded underneath them (i.e., inpainting), and (3) infer the ordering of the toppings without any depth ordering supervision."

The team talked about their synthetic and real pizza datasets.

"Pizza is the most photographed food on Instagram with over 38 million posts using the hashtag #pizza," they said. They downloaded half a million images from Instagram using several popular pizza-related hashtags. They filtered out undesired images using a CNN-based classifier trained on a set of manually labeled pizza/non-pizza images.

They crowd-sourced image-level labels for the pizza toppings on Amazon Mechanical Turk (AMT) for 9,213 pizza images.

For their synthetic pizza dataset, they used clip-art-style pizza images. "There are two main advantages of creating a dataset with synthetic pizzas. First, it allows us to generate an arbitrarily large set of pizza examples with zero human annotation cost. Second and more importantly, we have access to accurate ground-truth ordering information and multi-layer pixel segmentation of the toppings."

So, in the bigger picture, what contribution have they made, if any, to humankind? Victoria Song made a point, when she wrote, "In the long run, one could imagine a neural network being able to scan a photo and spit out a pretty accurate recipe based

on ingredients, how thoroughly it's cooked, and even barely visible spices."

After all is said (and done), "the research is mostly just demonstrating an AI's ability to differentiate between a confusing pile of ingredients." They knew this when they set out to focus on pizza. Think "archetypal example" of something that needs the sequential addition of ingredients in a specific order.

In the bigger picture, pizza is not the only item that could use their approach. "Though we have evaluated our model only in the context of pizza, we believe that a similar approach is promising for other types of foods that are naturally layered such as burgers, sandwiches, and salads."

For more information on their research, their paper is titled, "How to make a [pizza](#): Learning a compositional layer-based GAN [model](#)," by Dim Papadopoulos, Youssef Tamaazousti, Ferda Ofli, Ingmar Weber and Antonio Torralba. The paper is up on arXiv. The paper was submitted earlier this month.

More information: How to make a pizza:

Learning a compositional layer-based GAN model, arXiv:1906.02839 [cs.CV] arxiv.org/abs/1906.02839

pizzagan.csail.mit.edu/

© 2019 Science X Network

APA citation: PizzaGAN gets the picture on how to make a pizza (2019, June 20) retrieved 25 October 2021 from <https://techxplore.com/news/2019-06-pizzagan-picture-pizza.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.