

What the dog-fish and camel-bird can tell us about how our brains work

July 7 2015, by Jessica Birkett



One of the psychedelic nightmares generated by Google's Inceptionism system.
Credit: Google Research

You may have seen some of the "[nightmarish](#)" images generated by Google's aptly named [Inceptionism project](#). Here we have freakish fusions of dogs and knights (as in the image above), dumbbells with arms attached (see below) and a menagerie of Hieronymus Bosch-ian creatures:

But these are more than just computerised curiosities. The process that generated these images can actually tell us a great deal about how our own minds process and categorise images – and what it is we have that computers still lack in this regard.

Digging deep

Artificial neural networks, or "[deep learning](#)", have enabled terrific progress in the field of machine learning, particularly in [image classification](#).


Conventional approaches of machine learning typically relied on top-down rule-based programming, with explicit stipulation of what features particular objects had. They have also typically been inaccurate and error-prone.

An alternative approach is using [artificial neural networks](#), which evolve bottom-up through experience. They typically have several interconnected information processing units, or neurons. A programmer weights each neuron with certain functions, and each function interprets information according to an assigned mathematical model telling it what to look for, whether that be edges, boundaries, frequency, shapes, etc.


The neurons send information throughout the network, creating layers of interpretation, eventually arriving at a conclusion about what is in the image.

Google's [Inceptionism](#) project tested the limits of its neural network's image recognition capacity. The Google research team trained the network by exposing it to millions of images and adjusting network parameters until the program delivered accurate classifications of the objects they depicted.

Then they turned the system on its head. Instead of feeding in a image – say, a banana – and having the neural network say what it is, they fed in random noise or an unrelated image, and had the network look for bananas. The resulting images are the network's "answers" to what it's learned.




Coming soon to a nightmare near you. Credit: Google Research



Google's artificial neural network discovered all sorts of bizarre creatures lurking in the clouds. Credit: Google Research

What it tells us about machine-learning

The results of the Inceptionism project aren't just curiosities. The psychedelic interpretations made by the program indicate that *something* is missing that is unique to information processing in biological systems. For example, the results show that the system is vulnerable to *over-generalising* features of objects, as in the case of the dumbbell requiring an arm:



Dumbbells often have arms attached, but not like this.

This is similar to believing that cherries only occur atop ice cream sundaes. Because the neural network operates on correlation and probability (most dumbbells are going to be associated with arms), it lacks a capacity to distinguish contingency from necessity in forming stable concepts.

The project also shows that the over-reliance on feature detection leads to problems with the network's ability to identify probable co-occurrence. This results in a tendency towards *over-interpretation*, similar to how Rorschach tests reveal images, or inmates in Orange is the New


Black see [faces in toast](#).

Similarly, Google's neural network sees creatures in the sky, as with the strange creatures like the "Camel-Bird" and "Dog-Fish" above. It even picks up oddities within the Google homepage:




More than meets the eye. Credit: Google

A stable classification mechanism so far eludes deep learning networks. As described by the researchers at Google:



Not all the images generated by Inceptions are sinister. Credit: Google Research

We actually understand surprisingly little of why certain models work and others don't. [...] The techniques presented here help us understand and visualize how neural networks are able to carry out difficult classification tasks, improve network architecture, and check what the network has learned during training.



Starting with random noise, Google's artificial neural network found some

bananas. Credit: Google Research

What it tells us about ourselves

The Inceptionism project also tells us a little about how our own neural networks function. For humans like us, perceptual information about objects is integrated from various inputs, such as shape, colour, size and so on, to then be transformed into a concept about that thing.

For example, a "cherry" is red, round, sweet and edible. And as you discover more things like a cherry, your [neural network](#) creates a category of things like cherries, or to which a cherry belongs, such as "fruit". Soon, you can picture a cherry without actually being in the presence of one, owing to your authority over what a cherry is like at the conceptual level.

Conceptual organisation enables us to perceive drawings, photos and symbols of a cloud as referring to the same "cloud" concept, regardless of how much the cloud's *features* may suggest the appearance of Dog-Fish.

It also enables you to communicate about abstract objects, despite never having *experienced* them directly, such as [unicorns](#).

One implication that arises from this research by Google is that simulating intelligence requires an *additional* organisational component beyond just consolidated feature detection. Yet it's still unclear how to successfully replicate this function within deep learning models.

While our experimental artificial neural networks are getting better at image recognition, we don't yet know how they work – just like we don't

understand how our own brains work. But by continuing to test how artificial neural networks fail, we will learn more about them, and us. And perhaps generate some pretty pictures in the process.

This story is published courtesy of [The Conversation](#) (under Creative Commons-Attribution/No derivatives).

Source: The Conversation

Citation: What the dog-fish and camel-bird can tell us about how our brains work (2015, July 7) retrieved 20 April 2024 from <https://phys.org/news/2015-07-dog-fish-camel-bird-brains.html>

<p>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.</p>
--