

# Oh great, Facebook wants to know you're being sarcastic

September 26 2013, by Matthew Higgs

---



“Cat! No, wait, dog! Oh I give up.” Credit: Travis S.

You might think social networks couldn't possibly gather more information on you than they already do. That in a world where your every move is tagged, flagged and logged, there is nothing more that could possibly be gleaned from your digital footprints. You'd be wrong.

It was [revealed](#) this week that Facebook has followed Google into the world of "deep learning", a new frontier in [artificial intelligence](#) (AI) that uses networks of computers to mimic the [human brain](#) to better understand the subtle nuances that make us tick.

Internet companies specialise in tailoring services according to what they believe we like. If you spend a lot of time looking at holidays then you'll get shown a lot of [adverts](#) about travel services. If you communicate with particular people more than others on Facebook, they will tailor your news feed to see more of one friend and less of another.

But while this tailoring is remarkable in many respects, it is still clunky in others. Facebook can't detect [sarcasm](#), for example. If you post on your wall that you "just can't wait for more news" about your long-estranged school friend's [daily calorie intake](#) as they pursue yet another fad diet, it will take you at your word and send you "helpful" links to celebrity weight loss programmes. Deep learning is an attempt to tap into these more subtle aspects of human thought, building on what we have already achieved in [machine learning](#).

## **Machine can haz lolcatz**

The best known [example](#) of deep learning came out of Google X last year. A network of computer processors was let loose on YouTube and quickly learned to identify videos of cats without being told what a cat actually looks like.

Deep learning builds on a long history of machine learning. In this field of AI, the concepts you are trying to find often have labels, such as "picture of a cat" or "picture of a dog". These are objective labels that are based on socially accepted facts.

A cat is a cat and a dog is a dog, so when you explain to a machine what

a cat looks like, it can search out pictures of them with ease. The idea behind many machine learning algorithms is that if you show the machine enough labelled data, it can learn how to identify the labels of future unlabelled data.

But human nature is more complex. Traits like sarcasm and fickle shifts in emotion and desires all make understanding us a lot harder than recognising a wet nose and whiskers in an online gallery.

We have to question how well an algorithm can learn in this setting. In machine learning, the performance of an algorithm is often judged using its prediction accuracy. The machine is tested by looking at how many cats and dogs it correctly classifies from a test set. But how do we measure accuracy in a more subjective context?

Traditional machine learning is often divided into two distinct types of method: supervised and unsupervised. Supervised methods focus on problems in which there are labels like "cat" or "dog" and work on the assumption that there is a "teacher" supervising the learning process, showing the computer what a cat is in the first place. In unsupervised learning, there are no such labels and the aim is to simply find interesting features and representations of the data.

Deep learning sits somewhere in between these two approaches and utilises "semi-supervised" methods. The main idea is that the information in unlabelled data can be used to leverage the information in the labelled data.

Assume that an unsupervised algorithm separates a set of pictures into two distinct groups, then it is told that one of the data points from one group is a cat and one data point from the other group is a dog. If the data set consists of only pictures of cats and dogs, then it is likely that the label in each of those labelled examples applies to the whole group.

This is the fundamental idea in many deep learning applications.

In the same way that different parts of the brain focus on different functions, a deep learning network can be partitioned into different areas of activity. Generally, the lower layers in a deep learning network focus on discovering the interesting features in unlabelled data, while the top layers connect these features to observable labelled data.

This is analogous to the way in which the brain receives and processes visual information through the eye. By learning the interesting features of input information, deep learning networks are able to make better predictions with fewer labels.

## **You and your feed**

Facebook holds a huge amount of information and a certain amount of it indicates the preferences of the user. Deep learning can leverage all the additional information to better predict what the user really wants.

Whether or not someone is being sarcastic in their posts depends inherently upon the context of use. Who they are speaking to, the preceding comment, the topic of the thread, even the time of day, are all examples of features that may be important in determining whether someone is being sarcastic or not. To process all this data and to capture high level concepts requires the sophisticated machinery of deep learning. Then, after these features have been learned, maybe a simple variable such as the number of likes to your comment can be used to indicate how well your comment was received.

These methods are not just limited to Facebook. Anywhere we leave a digital footprint, machine learning methods can be applied to find the traits of our behaviour within the digital environment.

This does however open up ethical issues about how our data is collected

and used. Companies are becoming better and better at accurately representing us and telling us what they think we want. They can only go so far though, so if you don't like the way AI is heading, you can just beat the system by leading a more spontaneous life.

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

Source: The Conversation

Citation: Oh great, Facebook wants to know you're being sarcastic (2013, September 26)  
retrieved 27 April 2024 from <https://phys.org/news/2013-09-great-facebook-youre-sarcastic.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>
--