

# The hazards of presumptive computing

April 1 2015, by Michael Cowling

---



Until machines become truly intelligent, they're going to make a lot of mistakes when they try to help us. Credit: Steve Rainwater/Flickr, CC BY-SA

Have you ever texted somebody saying how "ducking annoyed" you are at something? Or asked Siri on your iPhone to call your wife, but somehow managed to be connected to your mother-in-law?

If you have, you may have been a victim of a new challenge in computing: that fine line where we trust a computer to make predictions for us despite the fact that it sometimes gets them wrong.

For one hapless administrator with the Australian Immigration department, this level of trust has almost certainly led to [major embarrassment](#) (or worse), with it being revealed that during November

last year they accidentally sent the personal details of the G20 leaders to the organisers of the Asian Cup Football tournament due to an autofilled e-mail address that went horribly wrong.

We trust the machines, but sometimes the machines let us down. So, what's happening? Are the machines too dumb to get what we mean? Or are they just getting too smart for their own good?

## **The uncanny valley of computing prediction**

It feels like we're entering an [uncanny valley](#) of computer prediction. This is where computers seem *almost* human, make us start to trust them, but then suddenly make a mistake so galling that we get uneasy that we've trusted a machine so completely.

The problem is that it's all just so convenient. My typing speed has increased immeasurably since I started to trust my iPhone to autocorrect the vague words I type into it and just went with the flow. And services like Google Now that predict the information you want before you even ask for it are even more useful.

But the trade-off is that sometimes it gets it wrong. And sometimes I find that I've inadvertently sent the wrong message to my wife, or had the phone make ridiculous suggestions like suggesting that my office is "home" (that went down well with the aforementioned wife!).

So, why is it so hard for a computer to be human?

## **Fool me once, computer...**

The challenge of making a computer seem human has been with us for quite a while. Ever since Alan Turing invented his computation machine

to break the Enigma code during the second world war, we've striven to make a computer that can think like a human and act like a human.

So much so, that we have even derived a test, called the Turing Test, to determine whether a computer can successfully fool somebody into thinking they are human.

In his paper that proposes the [Turing Test](#), Turing suggested that we don't need to make a computer that can genuinely think – whatever that means – but rather just build a computer simulation for which we can positively answer the question: "can [machines](#) do what we (as thinking entities) can do?", as cognitive scientist Stevan Harnad puts it.

Through a test he called the "imitation game", a human judge engages in natural language conversations with a human and a machine using a text-only channel. If the judge cannot tell the machine from the human, the machine is said to have passed the test.

Since Turing's original paper, [many variations](#) on the test have been proposed, adding perceptual capabilities like vision and audio, as well as extending the test with robotics.

But so far, no computer has definitively passed the original Turing Test. Every time we come close, they stumble into that uncanny valley, fall short in some way that makes us start to feel uneasy, and then the whole tower of cards falls.

This is not surprising. We are trying to make a machine deal with all the complexity of human processing and it's bound to make mistakes. A classic example of this is the tank parable by [Elieler Yudkowsky](#).

## **Tanks, but no tanks**

To demonstrate the problem of teaching a computer to be human, Yudkowsky [describes a situation](#) where US Army researchers train a computer to recognise whether or not a scene has a tank in it. To teach the computer this, the researchers show it many images, some with tanks in them, some without, and tell the computer whether or not each image contains a tank.

Through their testing, they determine that the computer has learnt to identify each scene correctly so they hand the system to the Pentagon, which then says it's people couldn't get it to work.

After some head scratching, the researchers discover that the photos of tanks had been taken on cloudy days and the photos without tanks had been taken on sunny days. So rather than learning to see tanks, the system had learnt to spot cloudy or sunny days!

Such are the hazards of teaching a computer a skill when it doesn't have sufficient context to understand what you want it to do.

## **Teaching a computer to know what we mean, not what we say**

So, after my mobile phone helpfully informed me that my workplace was "home" and I adjusted the address accordingly, I noticed my wife was quite quiet on the way home. I looked over at her and asked what was up and she said "nothing, I'm fine", at which point I knew I was in trouble!

But of course, that's not what she said. She said she was "fine", and a [computer](#), without context, would take her at her word. Context is everything, whether it's dealing with tanks or especially when dealing with a grumpy spouse.

Sometimes context is easy, such as the system Google implemented a couple of years ago that checks if you say the word "attached" in an email and then whether you've actually added an attachment, and [warns you if you haven't done both](#).

But sometimes context is harder, like when you type "Ian" and let it autocomplete, but end up with the wrong "Ian". After all, how is Gmail supposed to know which Ian you wanted without a host of other knowledge based on the content of your email and what you know about who you're emailing?

Nonetheless, computers are getting better at it. The iPhone autocomplete now adds "well" without an apostrophe until it detects a few words later that you meant "we'll" with an apostrophe, at which point it changes it. So it might not be long before it can tell you that you're e-mailing the wrong "Ian" too.

But for now we still need to be careful, because until computers can understand all the context of what we mean and what we do as humans – and there is no guarantee they ever will – we are still in that [uncanny valley](#) of presumptive computing.

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

Source: The Conversation

Citation: The hazards of presumptive computing (2015, April 1) retrieved 25 April 2024 from <https://phys.org/news/2015-04-hazards-presumptive.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.