

Once upon a bot: can we teach computers to write fiction?

They submit manuscripts on time. They never suffer writer's block. And they don't spend hours Googling their Amazon sales. There's just one thing wrong with robot authors - their stories stink. **Tom Meltzer** talks to the scientists teaching creative writing to the next generation of androids, while **Nicholas Lezard** reviews the latest robot fiction

Tom Meltzer

Tuesday 11 November 2014
17.29 GMT



Tell us a story ... computers find metaphor and sarcasm challenging. Photograph: Popperfoto/Getty

This month, several thousand aspiring authors are attempting to write a novel in 30 days. They are taking part in an annual event known as NaNoWriMo, National Novel Writing Month, in the hope that the time pressure will spur them on. For a small community of computer programmers, though, NaNoWriMo has a lighthearted sister competition: National Novel Generating Month, the goal of which is to teach a computer to write a novel for you.

However, finished NaNoGenMo projects are unlikely to trouble Booker judges. They include a version of *Moby-Dick* in which the words have been swapped for meows of the same length (immortal opening line: Meow me Meeeeeow); another version in which a few key words have been swapped out for emoji; and a novel made up of unconnected excerpts from an online database of teenage girls' accounts of their dreams.

Mark Riedl, an associate professor at the Georgia Institute of Technology, is not taking part in NaNoGenMo. But he and his colleagues are among the many computer scientists working on far more sophisticated digital storytellers. For the past two years, they have been tinkering with a program called Scheherazade, which learns how to describe tasks by analysing crowd-sourced human accounts, and then attempts to produce plausible short stories about, say, going to the movies or a restaurant.

At its best, Scheherazade writes fairly convincing vignettes: “You entered the movie theater ... You find the seats as indicated on your movie ticket ... You sat comfortably in your seats.” But it’s prone to telltale errors. “For Scheherazade, a successful story is one in which people will read the story and recognise the activity and not find too many obvious errors,” says Riedl.

Novels require more than that, of course. Part of the challenge is teaching a computer not merely to describe, but to imagine. This is the goal of the What-If Machine (Whim) project, a venture involving teams at five universities across Europe. Like Scheherazade, the Whim, as the program is affectionately known, seeks to understand what is possible by analysing vast databases of human prose. It then inverts or twists what it has learned to produce a new idea that could serve as the premise of a story.

The results would not look out of place as taglines for children’s books: what if there was a little whale who forgot how to swim? What if there was a little cat who learned how to write? The next step is for the computer to figure out which of its ideas will resonate with human readers, by feeding it reams of data about how we respond.

“The What-If Machine was trying to do that with one of our experiments,” says Teresa Llano, of the team at Goldsmiths, University of London. “We were doing a survey about ideas for characters for Disney films. We had ‘What if there was a little goat that was afraid of eating?’ and people didn’t like it. They didn’t want to see a goat dying.”

Then there is metaphor, sarcasm and all the many ways we avoid saying precisely what we mean. Computers struggle enough with the literal. How can we teach them to talk figuratively? At the University of Dublin, Whim member Tony Veale is working on just that. His program Metaphor Magnet is designed to produce metaphorical insights and ironies by inverting and contrasting stereotypes harvested from the web. The results vary from the oddly witty (“When the sandals that are worn by humble monks are worn by the smuggest hipsters”) to distinctly robotic definitions (“Referee. noun. A fussy fusion of 5 parts evaluation-performing valuer to 4 parts jersey-wearing runner.”) “I suppose the outputs of @MetaphorMagnet might be described as cynical fortune cookies,” says Veale.

There are many other challenges, from the rhythmic qualities of the prose to character arcs and plotting. The hardest to crack will be the elements of great writing we ourselves struggle to explain: the poetic force of the sentences, the unique insights of the author, the sense of a connection. “It’s not hard to generate a story,” says Riedl. “It’s not hard to tell a story. It’s hard to tell good stories. How do you get a computer to understand what good means?”

Figuring this out involves stripping human creativity down into its numerous constituent parts, devising algorithms for every device an author might deploy. In building a computer that can write, we are exposing the computer within the writer.



Science fiction ... the What-If Machine.

'I want to unplug the machine' : robot fiction reviewed by Nicholas Lezard Metaphor Magnet

How might a soothing mother become a threatening bank robber? What if soothing mothers were to give up their babies, plot bank robberies and become threatening bank robbers?

The most daring bank robber is not more demanding than the most scared child.

Bank robbers embrace criminality. Profiteers profit from the wars that encourage criminality. Who is worse?

When the bank robbers that break into calm vaults hide behind livid masks.

What if a principled scientist isn't crooked enough to perform expensive research? Might this methodical scientist embrace criminality, become a methodical bank robber and steal the necessary money?

NL: The Metaphor Magnet looks the most formally intriguing of these, and the most plausible, but only because (a) it is in note form, and (b) its repetitions and stylistic stiffness suggest an avant-garde literary project that exploits repetition and stiffness, or deliberately seeks to emulate the kind of prose that might come out of a machine. Had these been presented to me as an early experimental work by, say, Ben Marcus, I may have been fooled, had it not been for the general sense of ineptitude. Then again, I rather liked "When the bank robbers that break into calm vaults hide behind livid masks" - "calm vaults" reminds me of an Emily Dickinson poem.

What-If Machine

What if there was a banker who lost her bank and couldn't get money? But then she discovered that a robbery would solve her problem, so she forgot all about her old bank.

What if there was a doctor who lost her medicine and couldn't heal? But then she discovered

that a spirit would solve her problem, so she forgot all about her old medicine.

What if there was a poor, sluggish bird that couldn't sing because he didn't have enough practice? But he decided to perform anyway.

What if there was a banker whose regulator ran away and she suddenly became as excessive as traffic?

What if there was a clumsy banker who lost his hat and suddenly became as conservative as a nun?

What if there was a clumsy banker who lost her crony and suddenly became as tough as a brute?

What if there was a banker who woke up in an alley as a cat but could still use the phone?

What if there was a nurse who woke up in an aquarium as a fish but could still talk?

NL: The What-If Machine clearly signals itself as something whose first language is hexadecimal. "Couldn't get money" means what? That is: what does "get" mean? Nothing, or nothing specific enough. There is perhaps a surreal charm in some of these scenarios, but even someone who did not know that they were computer-generated would experience something like the effect of the "uncanny valley", whereby the attempt to get as close to human expression as possible only results in a sensation of great creepiness. Once you notice the formal similarity between the banker and the doctor in paragraphs one and two, of course, the jig is up (the bird who couldn't sing etc). Three sounds like the proposal for a children's story that will - indeed must - never be published, while "excessive as traffic" in four is word salad, not invention. As for the "clumsy banker", at this point I want to unplug the machine from its power source to stop it from irritating us ever again, which is, rightly or wrongly, not an option available to writers made of flesh and blood.

Scheherazade

John got into his car with his disguise, gun and note in his knapsack and headed towards the Old Second in the next town over, repeating his rehearsed demands silently over and over in his head.

John watched while a little old lady left the bank and walked to her car and then slipped on his gloves, slipped his gun into his coat pocket, grabbed his mask and strode determinedly to the lobby door and pulled it open.

John looked at his reflection in the glass of the door, gave himself a little smirk and covered his face. John took another deep breath as he wondered if this was really a good idea, and entered the bank.

John looked around the bank, making sure his timing was right.

John spotted a young blond teller, Sally, behind the counter.

John stood behind the lady and toddler and politely waited his turn, noticing the nameplate on the counter ... “Sally”.

When it was his turn, John, wearing his Obama mask, approached the counter. Sally saw Obama standing in front of her and she felt her whole body tense up as her worst nightmare seemed to be coming true.

Once Sally began to run, John pulled out the gun and directed it at the bank guard.

John wore a stern stare as he pointed the gun at Sally.

Sally screamed hysterically which alerted other people in the bank.

NL: My first thought was: “Oh look, it’s an extract from Dan Brown’s new novel.” Then I realised it was even clumsier than the master of turning rubbish into money. But not that much clumsier. I suspect that Scheherazade may even have been programmed using algorithms determined by genre fiction in general and Brown in particular, so relentless is the parade of cliches, redundant modifiers, and dimwit expositions. “Sally screamed hysterically which alerted other people in the bank” is a killer of a closing sentence, isn’t it?

The disturbing thing is that a little tweaking of the program, such as getting the machine to learn that you don’t begin six consecutive sentences with the same word, especially if it’s “John”, could have turned this into something that might have been written by a very stupid human being with a tin ear; and there is plenty enough of that around. But even if one day the computer will pass muster at the level of the sentence, there is, on this evidence, no foreseeable way as yet that it will be able to construct a narrative that is both plausible and gripping. You may breathe easy. Unless you are Dan Brown.

. This article was amended on 1 December 2014 to remove a quote from professor Mark Riedl, who, after further investigation into NaNoGenMo, had changed his view that no one was taking it seriously.