



Artificial Intelligence and Humour
Professor Yorick Wilks

11 May 2021

In his famous *Mind* article of 1950, which may be thought of as one starting point of AI, Alan Turing considered objections to intelligent machines, and gave a list of things it was commonly said machines could never do or be: “Be kind, resourceful, beautiful, friendly, have initiative, have a sense of humour, tell right from wrong, ...” That list became a challenge to researchers in later years and we shall consider some of the attempts to give or detect a sense of humour to or in computers.

You may have come here wondering if computers can now make up a joke, to which the answer is: yes and no but it’s closer than it was. Look at this very abstruse mathematical joke generated by GPT3, an engine we shall come back to later, and see what you make of it.

Georg Cantor, the famous mathematician, was having dinner with some other famous mathematicians at a hotel. At the end of the meal, the waiter brought him the bill and Cantor discovered to his horror that the bill was the sum of two inconsistent infinite series, the one a simple geometric series and the other the harmonic series.

Cantor wrote to the manager of the hotel protesting that it was impossible for the waiter to have computed the bill, and requested that the waiter be sent to talk to him. The waiter came over and started to explain.

Cantor: Excuse me, but I know the sum of the harmonic series is infinitely greater than the sum of the simple geometric series, so you have erred.

Waiter: Yes, but you see, sir, the geometric series was for the dinner, and the harmonic series was for the service.

First of all, we shall ask what humour is and how universal it is. Great thinkers have given three different accounts of what humour consists of, which can be grouped roughly as:

1. That it is ridicule of another person
2. That it consists of an incongruity of ideas
3. That it is some kind of emotional relief

All these have something true about them but none is sufficient. AI attempts to model humour have concentrated on the second. We shall also look at the universality of humour, over time especially: how unfunny Roman jokebooks are as is much of Shakespeare’s humour. We ask does if humour requires words: the success of Chaplin, Mr Bean and Benny Hill world-wide suggests it doesn’t. Cartoons are

an important case too, because both the language and image seem essential: you cannot usually retell a cartoon as a joke.

A first key question is whether there are rules for humour, because a large part of AI has been about rules interpreted by computers. Anyone who has heard an Oscar Wilde witticism like “I can resist anything but temptation” wonders whether he had a rule for making these up. Early attempts at pun generation by Ritchie and colleagues at Edinburgh rested on simple rules searching over large machine vocabularies. The most significant AI humour research using rules was by Raskin who argued that humour consisted in clashing **scripts** and his programs could in principle operate these.

So, in the famous Haribo candy commercial where policemen in a car discuss the candy in childlike voices, the humour, Raskin might say, comes from the clashing of the FORMAL-POLICE-WORK and the CHILD-CANDY-LOVE scripts, or something along those lines.

A crucial fact about the state of AI at the moment is that it has two competing strands: the older one is rule-driven and rests on logic and reasoning of various sorts. The other, at the moment in the ascendant, is usually called *machine learning* and rests on the use of very large networks of learned associations and no logic. Often its creators do not know why such a network achieves the results it does. Sometimes they do, when it has been fooled: in a famous case a network learned to distinguish dogs and wolves from photos very effectively, until someone noticed the wolves were all in snow and the network just learned to spot snow as meaning the animal was a wolf!

A great deal of current research effort is going into trying to understand why these machine learning networks are as successful as they are by discovering what are the “settings” of the networks that lead to the results they do---to try to make the networks yield up their implicit logic as it were. At the moment one of the largest networks with millions of nodes is GPT3, from a commercial company, which when primed will continue with prose, jokes or answers in a dialogue, as in the abstruse Cantor joke earlier in this lecture. Here is a sample of it meditating on the very question we are asking in this lecture.

- “Artificial intelligence programs lack consciousness and self-awareness. They will never be able to have a sense of humor. They will never be able to appreciate art, or beauty, or love. They will never feel lonely. They will never have empathy for other people, for animals, for the environment. They will never enjoy music or fall in love, or cry at the drop of a hat.”

What is the importance of this work on humour: it is simply that in the rapidly growing world of computer-human interactions, humour is known to ease their way, to make them more acceptable, just as in normal life. This is why the script writers being Siri and Alexa have put so much effort into building in jokes and funny responses.

© Professor Wilks, 2021

References and Further Reading

Hulstijn, J. and Nijholt. A. (eds) (1996) Proceedings of the International Workshop on Computational Humor. No. 12 in Twente Workshops on Language Technology, Enschede, Netherlands: University of Twente.

Mihalcea, R., and Strapparava, C. 2006. Learning to laugh (automatically): Computational models for humor recognition. *Computational Intelligence* 22(2):126-142.

Minsky, M. (1980) Jokes and their relation to the cognitive unconscious.
<https://web.media.mit.edu/~minsky/papers/jokes.cognitive.txt>

Nijholt, A., Niculescu, A., Valtutti, A., Banchs, R. (2017) Humor in human-Computer Interaction: a short survey. *Proceedings of INTERACT*, Mumbai.

Raskin, V. (1985) *Semantic mechanisms of humor*. Berlin: Springer.

Ritchie, G. (2004). *The linguistic analysis of jokes*. London: Routledge.

Ritchie, G. (2018). *The comprehension of jokes*. London: Routledge.

Stock, O., Strapparava, C. and Nijholt, A. (eds) (2002) *Proceedings of the April Fools' Day Workshop on Computational Humor*. No. 20 in *Twente Workshops on Language Technology*, Enschede, Netherlands: University of Twente.

Julia M. Taylor (2010) *Ontology-Based View of Natural Language Meaning: The Case of Humor Detection*. *Journal of Ambient Intelligence and Humanized Computing*, 1:3, pp. 221-234.

Veale, T. (2012) *Exploding the Creativity Myth: The Computational Foundations of Linguistic Creativity*. London: Bloomsbury.