

So a Computer Walks Into a Bar...

The future of artificial intelligence depends on teaching machines to talk like human beings—complete with irony, sarcasm and puns




Illustration: Ruth Gwily

By James Geary

The Wall Street Journal

Sept. 21, 2018 9:55 a.m. ET

If you were a stand-up comic, you definitely wouldn't want a review of your act to include the line "succeeds in generating pieces of text that are recognizably jokes, but some of them are not very good jokes." That's how researchers in the U.K. described their attempt in the early 1990s to create one of the first computer programs to tell jokes. Using as source material the "Crack-A-Joke Book," a collection of one-liners beloved by British schoolchildren, the team trained their system to generate simple puns. One of its most successful zingers: "What kind of tree can you wear? A fir coat."

As voice becomes the consumer interface of choice—some 39 million Americans own a voice-activated artificial-intelligence device like Amazon’s Alexa or Google Assistant—conversing with computers needs to be as natural, intuitive and frictionless as conversing with people. That will require AI systems to master irony, sarcasm, ambiguity, humor and puns— aspects of conversation that humans handle with ease but that still flummox even the most sophisticated AI. In other words, to really connect with people, AI needs to be witty.

Research is under way to accomplish exactly that. To help AI systems recognize sarcasm, Princeton computer scientists have compiled some 1.3 million examples of it, from a source where it is in plentiful supply—the website Reddit. At Carnegie Mellon University, researchers created an AI system that generates amusing memes (involving cats, of course), including this witticism overlaid across a picture of a bespectacled, bow-tied feline in front of a blackboard: “Come stay at Chemistry Cat Inn: Great day rates and even better nitrates.”

One of the most formidable challenges for AI is the “world knowledge problem.” AI systems are good at responding to specific goal-directed tasks, like looking up a recipe for chicken soup. But to understand even the simplest joke—a skeleton walks into a bar and says, “I’ll have a beer and a mop, please”—“you need to know an awful lot about the world,” said Kim Binsted, part of the team that trained that early computer to make puns and now a professor in the Information and Computer Sciences Department at the University of Hawaii at Manoa.

‘Endowing AI systems with the element of surprise is perhaps the ultimate technical challenge.’

About a decade later, Dr. Binsted helped to create a different computer model called Witty Idiomatic Sentence Creation Revealing Ambiguity In Context (WISCRAIC—get it?) that generated witticisms based on a database of words that sound alike and a general knowledge of idiomatic phrases. WISCRAIC searched its resources to find punning opportunities in combinations of names (e.g., Johnny),

occupations (deer-keeper) and adjectives (poor), cranking out knee-slappers like this: “Your mate Johnny is a hard up deer-keeper. He really needs doe!”

Tony Veale and colleagues at University College Dublin used a similar method to create Idiom Savant, a computer program that detects and interprets puns. Idiom Savant can identify a phrase like “acupuncture is a job well done” as a pun by recognizing the familiar idiom “a job well done,” noting the phonetic overlap between “job” and “jab,” and recognizing that the latter has been substituted for the former.

This, in fact, is pretty much the way that Sigmund Freud theorized that wit works—through what he called a “peculiar process of condensation and fusion.” By merging two distinct meanings—“dough” (slang for money) and “doe” (a deer, a female deer)—the pun blends domains of understanding that in conventional thinking remain separate. Puns point to the essence of all true wit: the ability to hold two different ideas about the same thing at the same time.

The human brain is really good at this, thanks largely to what neuroscientists call the “default network,” which kicks in during mental states such as daydreaming and free association. At these moments, our attention is diffuse, and the brain gleans information from a broad array of sources. In contrast, the brain’s “executive network” specializes in specific goal-directed tasks, like reasoning, planning and problem-solving. Current AI systems perform best when modeling the executive network: “Alexa, what’s the recipe for chicken soup?”

True wit, however, will most likely be found by modeling the default network. Consider the strange condition of Witzelsucht, which appears in certain cases of neuropsychiatric disease and brain damage. It is an abnormal compulsion to make jokes and puns (the term Witzelsucht is derived from the German words for “wit” and “sickness” or “addiction”). A possible explanation for Witzelsucht is that the executive network has been damaged, freeing the default network to indulge in the unbridled associations that lead to joking and punning.

More Essays

[We Can’t Afford the Drugs That Could Cure Cancer](#)

September 20, 2018

[Lessons From Tesla \(the Man, Not the Car\)](#)

September 14, 2018

[The Supreme Court Was Never Meant to Be Political](#)

September 14, 2018

[Advice to the Next Generation: Nap When You Can](#)

September 14, 2018

Endowing AI systems with the element of surprise is perhaps the ultimate technical challenge. “Personality is not strictly functional,” says Emma Coats, leader of the team tasked with endowing Google Assistant with a sense of personality. “So the more Google Assistant sounds like a person, the higher the expectations users will have that communication will be intuitive.”

Kory Mathewson, a Ph.D. candidate in computer science at the University of Alberta, is working toward intuitive human-AI communication. In collaboration with Piotr Mirowski, an AI researcher from London, he has developed a program called A.L.Ex., short for “artificial language experiment.” Trained on subtitles

drawn from roughly 100,000 films, A.L.Ex. composes its own responses to human performers and communicates these through speech recognition and text-to-speech software.

So far, A.L.Ex. is no Dorothy Parker. Mr. Mathewson considers it a victory if the system just stays on the same subject as the human, as this excerpt from one live improv comedy sketch shows:

Human: Captain, the ship is under attack. The frigate is arriving...

A.L.Ex.: You see, I don't feel the same way. I just don't want to stay in the position to get burnt.

Human: We're all going to get burned by cannon fire...

A.L.Ex.: While you're in there, tell me. When I get dressed, I'll go up there and see what happens.

Mr. Mathewson compares performing with A.L.Ex. to playing alongside a novice improviser: "The AI has no notion of narrative arc. It just generates randomness. So the system pushes the human to be a better improviser because you have to justify the randomness." Still, he's confident that with enough exposure to advanced banter the system will get wittier. Faster would also be nice, Mathewson notes, since in comedy, timing is everything.

How will we know when AI systems with true wit have arrived? Maybe when you ask one "What do you get when you cut a comedian in two?" and the AI shoots back, without missing a beat, "A half-wit."

— *Mr. Geary is the deputy curator of the Nieman Foundation for Journalism at Harvard and the author of "Wit's End: What Wit Is, How It Works and Why We Need It," which W.W. Norton will publish on Nov. 13.*