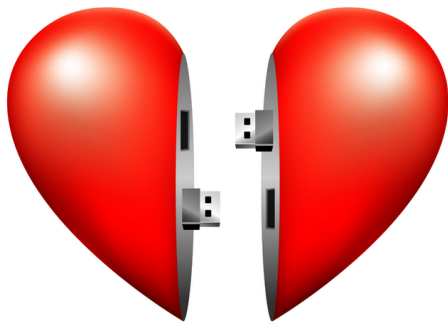


'Can I see your parts list?' What AI's attempted chat-up lines tell us about computer-generated language

28 April 2021, by Stefanie Ullmann



Credit: Sararoom Design/Shutterstock

Have you ever wondered what flirting with artificial intelligence would look like? Research scientist and engineer Janelle Shane has given us an idea by training a [neural network](#)—an algorithm loosely inspired by biological brain structures—to produce [chat-up lines](#).

Some of the results are hilarious and completely nonsensical, such as the inelegant: "2017 Rugboat 2-tone Neck Tie Shirt". But some of them turned out pretty well. At least, if you're a robot:

"I can tell by your red power light that you're into me. You look like a thing and I love you. Can I see your parts list?"

But how were these lines generated, and why do the results vary so much in terms of quality and cohesiveness? That's down to the types of neural networks Shane worked with: all based on [GPT-3](#), the world's largest language [model](#) to date.

Language modelling

GPT stands for generative pre-trained transformer. Its current version, developed by [OpenAI](#), is the third in a line of ever-improving [natural language](#) processing systems trained to produce human-like text or speech.

[Natural language processing](#), or NLP, refers to the application of computers to process and generate large amounts of coherent spoken or written text. Whether you ask Siri for a weather update, request for Alexa to turn on the lights, or you use Google to translate a message from French into English, you're able to do so because of developments in NLP.

It takes a variety of NLP tasks—from speech recognition to picking apart sentence structures—for applications such as Siri to successfully requests. The virtual assistant, much like any other language-based tool, is trained using many thousands of sentences, ideally as varied and diverse as possible.

Because human language is extremely complex, the best NLP applications rely increasingly on pre-trained models that allow "[contextual bidirectional learning](#)". This means considering a word's wider context in a sentence, scanning both left and right of any given word to identify the word's intended meaning. More recent models can even pay attention to more nuanced features of [human language, such as irony and sarcasm](#).

Computer compliments

GPT-3 is such a successful language-generating AI because it doesn't need retraining over and over again to complete a new task. Instead, it uses what the model has already learned about language and

applies it to a [something new](#)—such as [writing articles](#) and [computer code](#), generating [novel dialogue in video games](#), or formulating chat-up lines.

Compared to its predecessor GPT-2, the third-generation model is 116 times bigger and has been trained on billions of words of data. To generate its chat-up lines, GPT-3 was simply asked to automate the text for an article headlined: "These are the top pickup lines of 2021! Amaze your crush and get results!"

Because GPT-3's training updates have been added gradually over time, this same prompt could also be used on smaller, more basic variants—generating weirder and less coherent chat-up lines:

"Hey, my name is John Smith. Will you sit on my breadbox while I cook or is there some kind of speed limit on that thing? It is urgent that you become a professional athlete. CAPE FASHION."

But GPT-3's "DaVinci" variant—its largest and most competent iteration to date—delivered some more convincing attempts which might actually pass for effective flirting—with a little fine-tuning:

"You have the most beautiful fangs I've ever seen. I love you. I don't care if you're a doggo in a trenchcoat. I have exactly 4 stickers. I need you to be the 5th."

The latest variant of GPT-3 is currently the largest contextual language model in the world and is able to complete a number of highly impressive tasks. But is it smart enough to pass as a human?

Almost human

As one of the pioneers of modern computing and a firm believer in true artificial intelligence, Alan Turing developed the "Imitation Game" in 1950—today known as the "[Turing Test](#)". If a computer's performance is indistinguishable from that of a human, it passes the Turing Test. In language generation alone, GPT-3 could soon pass Alan Turing's test.

But it doesn't really matter if GPT-3 passes the Turing Test or not. Its performance is likely to depend on the specific task the model is used for—which, judging by the technology's flirting, should probably be something other than the delicate art of the chat-up line.

And, even if it were to pass the Turing Test, in no way would this make the model truly intelligent. At best, it would be extremely well trained on [specific semantic tasks](#). Maybe the more important question to ask is: do we even want to make GPT-3 more human?

Learning from humans

Shortly after its reveal in summer 2020, GPT-3 [made headlines](#) for spewing out shockingly sexist and racist content. But this was hardly surprising. The language generator was trained on vast amounts of text on the internet, and without remodelling and retraining it was doomed to replicate the [biases](#), harmful language and misinformation that we know to exist online.

Clearly, [language](#) models such as GPT-3 do not come without [potential risks](#). If we want these systems to be the basis of our digital assistants or conversational agents, we need to be more rigorous and selective when giving them reading material to learn from.

Still, [recent research](#) has shown that GPT-3's knowledge of the internet's dark side could actually be used to automatically detect online hate speech, with up to 78% accuracy. So even though its chat-up lines look unlikely to kindle more love in the world, GPT-3 could may be set, at least, to reduce the hate.

This article is republished from [The Conversation](#) under a Creative Commons license. Read the [original article](#).

Provided by The Conversation

APA citation: 'Can I see your parts list?' What AI's attempted chat-up lines tell us about computer-generated language (2021, April 28) retrieved 20 January 2022 from <https://techxplore.com/news/2021-04-ai-chat-up-lines-computer-generated-language.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.