

How AI systems use Mad Libs to teach themselves grammar

July 24 2020, by Edmund L. Andrews



Credit: Pixabay/CC0 Public Domain

Imagine you're training a computer with a solid vocabulary and a basic knowledge about parts of speech. How would it understand this sentence: "The chef who ran to the store was out of food."

Did the chef run out of food? Did the store? Did the chef run the store that ran out of food?

Most human English speakers will instantly come up with the right answer, but even advanced artificial intelligence systems can get confused. After all, part of the sentence literally says that "the store was out of food."

Advanced new machine learning models have made enormous progress on these problems, mainly by training on huge datasets or "treebanks" of sentences that humans have hand-labeled to teach grammar, syntax and other linguistic principles.

The problem is that treebanks are expensive and labor intensive, and computers still struggle with many ambiguities. The same collection of words can have widely different meanings, depending on the sentence structure and context.

But a pair of new studies by artificial intelligence researchers at Stanford find that advanced AI systems can figure out linguistic principles on their own, without first practicing on sentences that humans have labeled for them. It's much closer to how human children learn languages long before adults teach them grammar or syntax.

Even more surprising, however, the researchers found that the AI model appears to infer "universal" grammatical relationships that apply to many different languages.

That has big implications for [natural language](#) processing, which is increasingly central to AI systems that answer questions, translate languages, help customers and even review resumes. It could also facilitate systems that learn languages spoken by very small numbers of people.

The key to success? It appears that machines learn a lot about [language](#) just by playing billions of fill-in-the-blank games that are reminiscent of "Mad Libs." In order to get better at predicting the missing words, the systems gradually create their own models about how words relate to each other.

"As these models get bigger and more flexible, it turns out that they actually self-organize to discover and learn the structure of human language," says Christopher Manning, the Thomas M. Siebel Professor in Machine Learning and professor of linguistics and of computer science at Stanford, and an associate director of Stanford's Institute for Human-Centered Artificial Intelligence (HAI). "It's similar to what a human child does."

Learning Sentence Structure

The first study reports on experiments by three Stanford Ph.D. students in computer science—Kevin Clark, John Hewitt and Urvashi Khandelwal—who worked with Manning and with Omer Levy, a researcher at Facebook Artificial Intelligence Research.

The researchers began by using a state-of-the-art language processing model developed by Google that's nicknamed BERT (short for "Bidirectional Encoder Representations from Transformers"). BERT uses a Mad Libs approach to train itself, but researchers had assumed that the model was simply making associations between nearby words. A sentence that mentions "hoops" and "jump shot," for example, would prompt the model to search for words tied to basketball.

However, the Stanford team found that the system was doing something more profound: It was learning [sentence structure](#) in order to identify nouns and verbs as well as subjects, objects and predicates. That in turn improved its ability to untangle the true meaning of sentences that might

otherwise be confusing.

"If it can work out the subject or object of a blanked-out verb, that will help it to predict the verb better than simply knowing the words that appear nearby," Manning says. "If it knows that 'she' refers to Lady Gaga, for example, it will have more of an idea of what 'she' is likely doing."

That's very useful. Take this sentence about promotional literature for mutual funds: "It goes on to plug a few diversified Fidelity funds by name."

The system recognized that "plug" was a verb, even though that word is usually a noun, and that "funds" was a noun and the object of the verb—even though "funds" might look like a verb. Not only that, the system didn't get distracted by the string of descriptive words—"a few diversified Fidelity"—between "plug" and "funds."

The system also became good at identifying words that referred to each other. In a passage about meetings between Israelis and Palestinians, the system recognized that the "talks" mentioned in one sentence were the same as "negotiations" in the next sentence. Here, too, the system didn't mistakenly decide that "talks" was a verb.

"In a sense, it's nothing short of miraculous," Manning says. "All we're doing is having these very large neural networks run these Mad Libs tasks, but that's sufficient to cause them to start learning grammatical structures."

Discovering Universal Language Principles

In a [separate paper](#) based largely on work by Stanford student Ethan Chi, Manning and his colleagues found evidence that BERT teaches itself

universal principles that apply in languages as different as English, French and Chinese. At the same time, the system learned differences: In English, an adjective usually goes in front of the noun it's modifying, but in French and many other languages it goes after the noun.

The bottom line is that identifying cross-language patterns should make it easier for a system that learns one language to learn more of them—even if they seem to have little in common.

"This common grammatical representation across languages suggests that multilingual models trained on 10 languages should be able to learn an eleventh or a twelfth language much more easily," Manning says.

"Indeed, this is exactly what we are starting to find."

More information: Christopher D. Manning et al. Emergent linguistic structure in artificial neural networks trained by self-supervision, *Proceedings of the National Academy of Sciences* (2020). [DOI: 10.1073/pnas.1907367117](https://doi.org/10.1073/pnas.1907367117)

Provided by Stanford University

Citation: How AI systems use Mad Libs to teach themselves grammar (2020, July 24) retrieved 24 April 2024 from <https://techxplore.com/news/2020-07-ai-mad-lib-grammar.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.