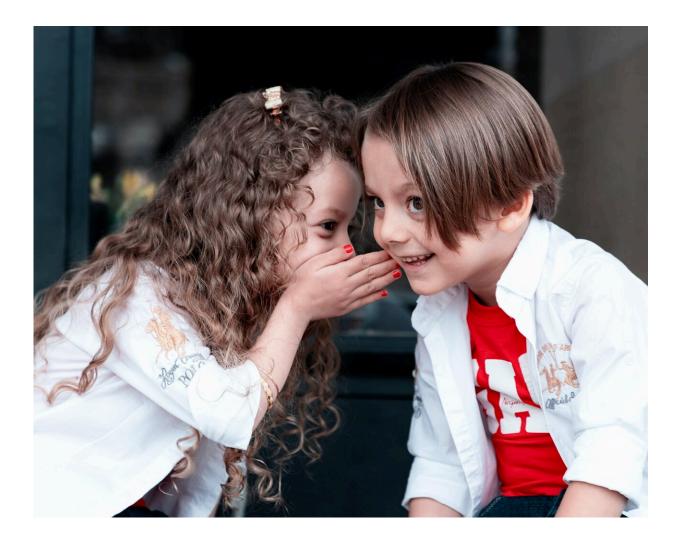# AI is changing scientists' understanding of language learning—and raising questions about an innate grammar

October 19 2022, by Morten H. Christiansen and Pablo Contreras Kallens



Credit: Unsplash/CC0 Public Domain

Unlike the carefully scripted dialogue found in most books and movies, the language of everyday interaction tends to be messy and incomplete, full of false starts, interruptions and people talking over each other. From casual conversations between friends, to bickering between siblings, to formal discussions in a boardroom, authentic conversation is chaotic. It seems miraculous that anyone can learn language at all given the haphazard nature of the linguistic experience.

For this reason, many language scientists—including Noam Chomsky, a founder of modern linguistics—believe that language learners require a kind of glue to rein in the unruly nature of everyday language. And that glue is grammar: a system of rules for generating grammatical sentences.

Children must have a grammar template wired into their brains to help them overcome the limitations of their language experience—or so the thinking goes.

This template, for example, might contain a "super-rule" that dictates how new pieces are added to existing phrases. Children then only need to learn whether their native language is one, like English, where the verb goes before the object (as in "I eat sushi"), or one like Japanese, where the verb goes after the object (in Japanese, the same sentence is structured as "I sushi eat").

But new insights into language learning are coming from an unlikely source: artificial intelligence. A new breed of large AI language models can write newspaper articles, poetry and computer code and answer questions truthfully after being exposed to vast amounts of language input. And even more astonishingly, they all do it without the help of grammar.

## Grammatical language without a grammar

Even if their choice of words is sometimes strange, nonsensical or contains racist, sexist and other harmful biases, one thing is very clear: the overwhelming majority of the output of these AI language models is grammatically correct. And yet, there are no grammar templates or rules hardwired into them—they rely on linguistic experience alone, messy as it may be.

GPT-3, arguably the most well-known of these models, is a gigantic deep-learning neural network with 175 billion parameters. It was trained to predict the next word in a sentence given what came before across hundreds of billions of words from the internet, books and Wikipedia. When it made a wrong prediction, its parameters were adjusted using an automatic learning algorithm.

Remarkably, GPT-3 can generate believable text reacting to prompts such as "A summary of the last "Fast and Furious' movie is…" or "Write a poem in the style of Emily Dickinson." Moreover, GPT-3 can respond to SAT level analogies, reading comprehension questions and even solve simple arithmetic problems—all from learning how to predict the next word.

## Comparing AI models and human brains

The similarity with human language doesn't stop here, however. Research published in *Nature Neuroscience* demonstrated that these artificial deep-learning networks seem to use the same computational principles as the human brain. The research group, led by neuroscientist Uri Hasson, first compared how well GPT-2—a "little brother" of GPT-3—and humans could predict the next word in a story taken from the podcast "This American Life": people and the AI predicted the exact same word nearly 50% of the time.

The researchers recorded volunteers' brain activity while listening to the

story. The best explanation for the patterns of activation they observed was that people's brains—like GPT-2—were not just using the preceding one or two words when making predictions but relied on the accumulated context of up to 100 previous words. Altogether, the authors conclude: "Our finding of spontaneous predictive neural signals as participants listen to natural speech suggests that active prediction may underlie humans' lifelong language learning."

A possible concern is that these new AI language models are fed a lot of input: GPT-3 was trained on linguistic experience equivalent to 20,000 human years. But a preliminary study that has not yet been peer-reviewed found that GPT-2 can still model human next-word predictions and brain activations even when trained on just 100 million words. That's well within the amount of linguistic input that an average child might hear during the first 10 years of life.

We are not suggesting that GPT-3 or GPT-2 learn language exactly like children do. Indeed, these AI models do not appear to comprehend much, if anything, of what they are saying, whereas understanding is fundamental to human language use. Still, what these models prove is that a learner—albeit a silicon one—can learn language well enough from mere exposure to produce perfectly good grammatical sentences and do so in a way that resembles human brain processing.

## Rethinking language learning

For years, many linguists have believed that learning language is impossible without a built-in grammar template. The new AI models prove otherwise. They demonstrate that the ability to produce grammatical language can be learned from linguistic experience alone. Likewise, we suggest that children do not need an innate grammar to learn language.

"Children should be seen, not heard" goes the old saying, but the latest AI language models suggest that nothing could be further from the truth. Instead, children need to be [engaged in the back-and-forth of conversation](#) as much as possible to help them develop their language skills. Linguistic experience—not grammar—is key to becoming a competent [language](#) user.

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

Provided by The Conversation