Essay

Does artificial intelligence genuinely capture the essence of language?

Jan Apolline D. Estrella

If you have been around Facebook, Twitter, and Tiktok, then you most probably have heard about the viral artificial intelligence (AI) bot, ChatGPT. Publicly released just this November 2022, this state-of-the-art technology by OpenAI has recently taken the internet by storm by making science fiction a reality. It can talk to you in a very natural and human-like way, and it can answer almost any question that you might have.

In less than a week, it caught the attention of over a million users not just for how human-like it is or for how similar it is to Tony Stark's AI butler J.A.R.V.I.S. in the Iron Man movies, but also for how it has proven itself to be useful to working professionals who use ChatGPT as a tool for better and faster writing of articles, academic papers, job applications, and even computer code (Aydın & Karaarslan, 2022). It has also been used by students to automatically write their essays for them. However, this is a controversial use case of plagiarism that requires an entirely different conversation for another time.

ChatGPT is not the first of its kind. There is actually an area in artificial intelligence that is dedicated solely to making computers understand and use human language. This field in computer science is aptly called natural language processing (NLP). Perhaps it is the first time you have heard of it, but it is definitely not the first time you have experienced it. NLP is behind the helpful—but sometimes annoying—autocomplete feature and spell checker on your phones. It is also the mechanism that carefully assists you in writing emails on Gmail and papers on Google Docs by predicting the next words that you want to say or correcting your grammar. You might have even relied on NLP to finish your song translation class project as you made use of Google Translate. Beyond academic and professional usage, you might still find yourself depending on NLP as you count on Netflix's recommendation system to lead you to your next movie or series. Even before ChatGPT became famous, NLP has been around you all along.

Now that you are aware of what NLP is, you might be asking yourself, "But how can a computer understand human language?" A short history of NLP will lead you to conclude that there have been multiple answers to this question over the years. NLP emerged in the late 1940s when people began to hope for an automatic machine translator. During this time, most researchers came from a background in the study of linguistics and language. NLP research in this period was primarily focused on syntax.

Linguists felt the need to form an explicit, complete, and formal characterization of language that can be applied to computer algorithms (Jones, 1994). One notable researcher from this phase is Noam Chomsky, the father of modern linguistics. In 1957, he published the book *Syntactic Structures*, where he recognized that for a language to be understandable to a computer, the sentence structure would have to be changed (Foote, 2019). With this, he introduced an elegant style of representing grammar called Phrase Structure Grammar, which models how a sentence can be broken into parts due to the recursive nature of language. For example, a sentence can be broken down into a noun phrase and a verb phrase; a verb phrase can then be broken down into a verb and a noun phrase; a noun phrase can then be broken down into a determiner, noun, and prepositional phrase; and so on.

Figure 1
An Example of a Phrase Structure Tree

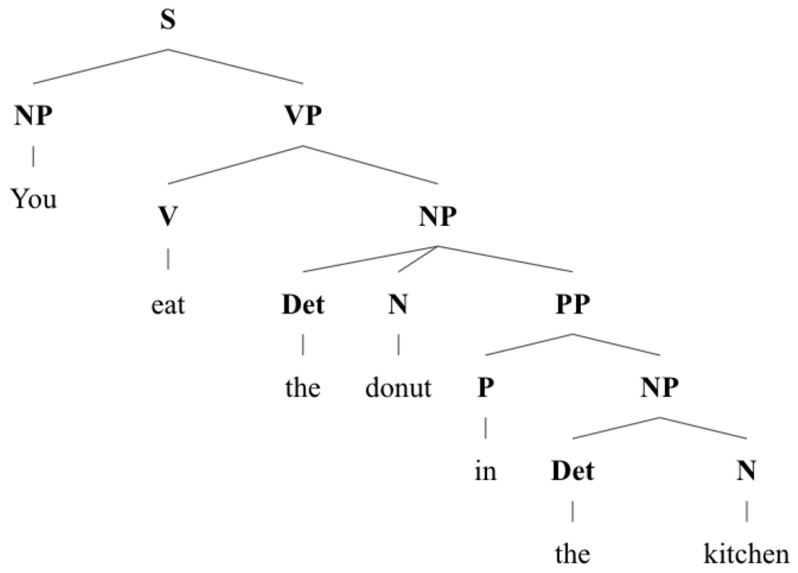

In other words, NLP technologies back then focused on the implementation of these graceful yet brute-force handwritten textbook rules. Unfortunately, these rules can be too rigid, and making them can be exhaustive for researchers. On top of that, the computers they used had very limited storage and were extremely slow to take in all these syntactic rules. NLP was struggling and it is easy to see why. After 12 years of research and 20 million dollars, the task of machine translation was still more costly than manual translation. So in 1966, the U.S. National Research Council officially halted the funding of NLP.

It took 14 years for the past failure of NLP to be redeemed with a new revolution: the focus shifted from syntactic to stochastic. To put it simply, the aforementioned tedious rules were replaced with statistical and probabilistic methods (Foote, 2019). This has

been made possible by the existence of more powerful computers. Before things get too abstract for you, let me share with you a tangible example of an NLP technology that utilizes statistical reasoning: Google Translate. This is how it works: by finding phrase or sentence patterns in a humongous dataset of translated texts, it tries to find pieces of text most likely to be associated with the text you want to be translated (Grajales, 2015). Forming sentences depend on statistically predicting what the next word after a sequence of words will be by copying from examples.

These statistical methods served as a stepping stone to more modern artificial intelligence that is better at learning the patterns found in different languages. Such well-known new technologies include something called transformers, which is behind the genius of ChatGPT. With mathematical equations, a transformer tries to observe the relationships between all the words in a sentence, and figures out the specific contexts that certain words or phrases are often used in. For example, let us look at two English sentences:

> The chicken didn't eat the strawberry because **it** was full.
>
> The chicken didn't eat the strawberry because **it** was rotten.

In the first sentence, *it* refers to *chicken*. But in the second sentence, *it* refers to *strawberry*. In the translation of these two sentences into French, the translation of *it* varies depending on the gender of the subject it points to. If the masculine noun *chicken* is the focus, then *it* is translated to *il*. But if the feminine noun *strawberry* is the antecedent, then the translation is *elle*.

> Le **poulet** n'a pas mangé la fraise car **il** était plein.
>
> Le poulet n'a pas mangé la **fraise** car **elle** était pourrie.

By paying attention to the surrounding context of each word, a transformer is able to navigate through such ambiguities in a way that previous approaches could not. Hence, with the success of such methods, stochastic models have since then dominated the task of "understanding" human language in the field of artificial intelligence.

Not to sound too philosophical, but the question now is, is this "understanding" really understanding? With the unquestionable excellence of artificial intelligence in performing language tasks, this is a question that people already forget to ask. Although their performance is indeed groundbreaking, stochastic technologies like Google Translate and ChatGPT do not make use of any grammar rules, dictionaries, nor even guidance from linguists or language experts. One can also wonder why these models have to take in datasets as large as 36 million sentences to "learn" a language (Vaswani et al., 2017) while a child can learn a language by hearing just a few sentences from his parents. This makes one realize that these technologies seem to have no understanding of language after all! As the father of modern linguistics who advocates for an elegant and simple theory of language through his works, Chomsky ridicules machine learning that uses purely statistical methods to mimic a certain behavior without really understanding the meaning behind that behavior (Gold, 2011). AI bots

sometimes act like clueless students who choose to find the answer to a math problem on Google without trying to understand the solution. This drove me to ask a question that allowed me to explore my hopes and establish my convictions as a wandering computer science undergraduate in the field of natural language processing: does artificial intelligence genuinely capture the essence of language?

For a long time, I have admired the elegance of Chomsky's philosophy and have sought to uncover what it means for a computer to "understand" language even if it means deviating from how present-day statistical and probabilistic technology works. I have only been able to create and interpret ideas regarding this artificial intelligence debate from an algorithmic and computational perspective: language understanding by simplicity versus by stochasticity. But after a brief experience of being a linguistics student under an elective course, I found that this tension—when translated to a language used by linguists—equates to a battle between the prescriptive and descriptive approaches to language.

The prescriptive approach views grammar as a set of rules that define the "proper" use of language (Yule, 2010). It basically focuses on how language should be used. For instance, one must not end a sentence with a preposition. Sounds familiar? (Hint: Chomsky appears to have a prescriptive approach to NLP.) On the other hand, the descriptive approach revolves around how language is used. Researchers gather and analyze samples of language usage and attempt to characterize language from there. They focus more on the patterns and habits of people when using language. Does it ring any bells? (Hint: In some way, AI tools seem to act like these descriptive linguists.)

In the beginning, I mistakenly thought I immediately had to take a side. But over time, I learned that it is not the answers that bring us to the breakthrough, but the questions.

> What does language understanding mean?
>
> Beyond what it looks like, what is intelligence?
>
> Will we be able to create more successful language technologies by understanding how the human brain learns language?
>
> How can machines extract meaning from language?
>
> What does it take for artificial intelligence to genuinely capture the essence of language?

Just as Douglas Hofstadter or "the man who would teach machines to think" would remind, "For now, what is important is not finding the answer, but looking for it."

References

Aydın, Ö., & Karaarslan, E. (2022). OpenAI ChatGPT generated literature review: Digital twin in healthcare. In Ö. Aydın (Ed.), *Emerging computer technologies 2* (pp. 22–31). İzmir Akademi Derneği. https://doi.org/10.2139/ssrn.4308687

Foote, K. D. (2019, May 22). *A brief history of natural language processing (NLP)*. Dataversity. https://www.dataversity.net/a-brief-history-of-natural-language-processing-nlp/

Gold, K. (2011, June 21). *Norvig vs. Chomsky and the fight for the future of AI*. Tor. https://www.tor.com/2011/06/21/norvig-vs-chomsky-and-the-fight-for-the-future-of-ai/

Grajales, C. A. G. (2015, June 23). *The statistics behind Google Translate*. Stats & Data Science Views. https://www.statisticsviews.com/article/the-statistics-behind-google-translate/

Jones, K. S. (1994). Natural language processing: A historical review. In A. Zampolli, N. Calzolari, & M. Palmer (Eds.), *Current issues in computational linguistics: In honour of Don Walker* (pp. 3–16). Springer. https://doi.org/10.1007/978-0-585-35958-8_1

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. In I. Guyon, U. von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. V. N. Vishwanathan, & R. Garnett (Eds.), *Advances in neural information processing systems 30 (NIPS 2017)* (pp. 5999–6009). Curran Associates, Inc.

Yule, G. (2010). *The study of language* (4th ed.). Cambridge University Press. https://doi.org/10.1017/CBO9780511757754