

# A Response to Piantadosi (2023)

Dan Milway (dan@milway.ca)

It is perhaps an axiom of criticism that one should treat the object of criticism on its own terms. Thus, for instance, a photograph should not be criticized for its lack of melody. This axiom makes it difficult to critique a recent paper by Piantadosi (2023)—hereafter SP—as it is difficult to determine what its terms are. It is ostensibly the latest installment of the seeming perennial class of papers that argue on the basis of either a new purported breakthrough in so-called AI or an exotic natural language dataset that rationalist theories of grammar are dead wrong, but it actually is a curious mix of criticism of Generative Grammar, promissory notes, and promotion for OpenAI’s proprietary ChatGPT chatbot.

The confusion begins with the title of the paper in (1) which doubles as its thesis statement and contains a category error.

(1) Modern language models refute Chomsky’s approach to language.

To refute something is show that it is false, but approaches do not have truth values. One can refute a claim, a theory, or a hypothesis, and one can show an approach to be ineffective, inefficient, or counterproductive, but one cannot refute an approach. The thesis of the paper under discussion, then, is neither true nor false, and we could be excused for ignoring the paper altogether.

Another axiom of criticism, though, is the principle of charity, which dictates that we present the best possible version of the object of our criticism. To that end we can split up (1) into two theses (2) and (3).

(2) Modern language models refute Chomsky’s theories language.

(3) Modern language models show Chomsky’s approach to language to be obsolete.

It is these theses that I address below.

The general shape of SP’s argument is as follows: (A) Chomsky claims that adult linguistic competence cannot be attained or simulated on the basis of data and statistical analysis alone. (B) The model powering ChatGPT simulates adult linguistic competence on the basis of data and statistical analysis alone. Therefore, (C) The model powering ChatGPT shows Chomsky’s claims to be false. To support his argument, SP presents queries and outputs from ChatGPT and argues that each refutes or approaches a refutation of a specific claim of Chomsky’s—each argument is of the form “Chomsky claims a purely statistical model could never do X, but ChatGPT can do (or can nearly do) X.”

As the hedging in this summary indicates, SP admits there are some phenomena for which ChatGPT does not exhibit human-like behaviour. For instance, when SP prompts the chatbot to generate ten sentences like (4), the program returns ten sentences all of which share the syntactic structure of (4), none of which are wholly meaningless like (4).

(4) Colorless green ideas sleep furiously.

SP explains this as away, writing “[w]e can note a weakness in that it does not as readily generate wholly meaningless sentences ..., likely because meaningless language is rare in the training data.” Humans can generate meaningless language, despite the fact that is “rare in the training data” for us too. The autonomy of syntax, then, is an instance where OpenAI’s language model does not exhibit human-like behaviour.

Furthermore, SP notes that current models require massive amounts of data to achieve their results—amounts far outstripping the amount of data available to a child. He also notes that the data is qualitatively different

from that available to a child<sup>1</sup> In doing so, he admits that modern language models (MLMs) are not good models of the human language faculty, contradicting one of the premises of his argument.

Though these empirical shortcomings of models like the one powering ChatGPT quite plainly refute (2), we do not even need such evidence to do so, as (2) is self-refuting. It is self-refuting because it does not address theoretical claims that Chomsky or, to my knowledge, any Generative theoretician has made. Far from claiming that MLMs could never do the things that ChatGPT can do, Chomsky has repeatedly claimed the opposite—that with enough data and computing power, a statistical model would almost certainly outperform any scientific theory in terms of empirical predictions. Indeed, this is the point of one the quotes that SP includes:

You can't go to a physics conference and say: I've got a great theory. It accounts for everything and is so simple it can be captured in two words: "Anything goes."

All known and unknown laws of nature are accommodated, no failures. Of course, everything impossible is accommodated also.

Furthermore, Generative theories are about a component of human cognition<sup>2</sup>, and nowhere does SP claim that "modern language models" are good models of human cognition. Indeed, this is an extension of the above discussion of the data requirements of MLMs, and logically amounts to a claim that the supposed empirical successes of MLMs are illusory without biological realism.

So, SP does not show that MLMs refute Chomsky's theory, but what of his approach to language? Here we can look at the purported successes of MLMs. For instance, SP presents ChatGPT data showing grammatical aux-inversion in English, but provides no explanation as to how it achieves this. Such an explanation though, is at the core of Chomsky's approach to language. If MLMs do not provide an explanation, then how can they supplant Chomsky's approach?

The failure of MLMs to supplant Chomsky's approach can be demonstrating by extending one of SP's metaphors. According to SP, the approach to science used by MLMs is the same that is used to model and predict hurricanes and pandemics. Let's assume this is true, it is also true that meteorological and epidemiological models have at their cores, equations arrived at by theoretical/explanatory work done by physicists and biologists respectively. If MLMs supplant theoretical/explanatory linguistics, then hurricane and pandemic models should supplant physics and biology. No serious person would make this argument about physics or biology, yet it is fairly standard in linguistics.

Thus far we have been taking SP's data at face-value, and while there is absolutely no reason to believe that SP has falsified it in any way, there is still a serious problem with it—it is, practically speaking, unreplicable, since we have no access to the model that generated it. The data in the paper was generated by ChatGPT in early 2023. When it was initially released, ChatGPT worked with the GPT 3.5 model, and has since been migrated to GPT 4—both of which are closed-source. So, while SP adduces ChatGPT data as evidence in favour of the sort of models that he has developed as his research program, there is no way to know whether ChatGPT uses the same sort of model. Indeed, ChatGPT could be built atop a model based on Generative theories of language for all we know.

Returning to the axiom I started with—that one should criticize something on its own terms—The ultimate weakness of SP's paper, is its failure to follow it. Chomsky's main critique of MLMs—alluded to in the quote above—is not that they are unable to produce grammatical expressions. It's that if they were to be trained on data from an impossible language—a language that no human could acquire—they would "learn" that language just as easily as, say, English. One does not need to look very far to find Chomsky saying exactly this. Take, for instance, the following quote in which Chomsky responds to a request for his critique of current so-called AI systems.<sup>3</sup>

There's two ways in which a system can be deficient. One way is it's not strong enough—[it] fails to do certain things. The other way is it's too strong—it does what it shouldn't do. Well,

---

<sup>1</sup>SP also wrongly implies that the data that informs actual language acquisition consists of child-directed speech.

<sup>2</sup>This is the crux of the I-/E-language distinction that Chomsky often discusses.

<sup>3</sup>Taken from extemporaneous speech. Edited to remove false starts and other disfluencies. Source: [https://www.youtube.com/watch?v=PBdZi\\_JtV4c](https://www.youtube.com/watch?v=PBdZi_JtV4c)

my own interests happen to be language and cognition—language specifically. So take GPT. Gary Marcus others have found lots of ways in which the system’s deficient—this system and others—[it] doesn’t do certain things. That can in principle at least be fixed—you add another trillion parameters double the number of terabytes and maybe do better. When a system is too strong it’s unfixable typically and that’s the problem with GPT and the other systems.

So if you give a database to the GPT system which happens to be from an impossible language—one that violates the rules of language—they’ll do just as well—often better because the rules can be simpler. For example one of the fundamental properties of the way language works—there’s good reasons for it—is that the rules the core rules ignore linear order of words—they ignore everything that you hear. They attend only to abstract structures that the mind creates. So it’s very easy to construct impossible languages which use very simple procedures involving linear order of words [The] trouble is that’s not language but GPT will do just fine with them. so it’s kind of as if somebody were to propose uh say a revised version of the of the periodic table which included all the elements all the possible elements and all the impossible elements and didn’t make any distinction between them that wouldn’t tell us anything about elements. And if a system works just as well for impossible languages as for possible ones by definition not telling us anything about language. And that’s the way these systems—work it generalizes the other systems too. So the deep problem that concerns me is too much strength. I don’t see any conceivable way to remedy that.

The key notion here is that of an “impossible language” which, though it seems to have an *a priori* flavour to it, is actually an empirical notion. Generative theory, like every scientific theory, predicts not only what is possible, but what is impossible? For instance, generative theory predicts that linear order is not available to syntax, and therefore no language has grammatical rules based on linear order. SP indirectly addresses this concern:

It’s worth thinking about the standard lines of questioning generative syntax has pursued—things like, why don’t kids ever say “The dog is believed’s owners to be hungry” or “The dog is believed is hungry” [...]. The answer provided by large language models is that these are not permitted under the best theory the model finds to explain what it does see. Innate constraints are not needed.

Following this standard empiricist reasoning, there are no impossible languages, only languages which have yet to be seen.<sup>4</sup> If all we had to go on was description of actually existing languages, then the empiricist and rationalist accounts would be equally plausible. Luckily for us, we are not limited in this way, we have experimental results that directly support the rationalist accounts—Smith and Tsimpli (1995), for instance, provides evidence that, while we can learn “impossible languages”, we do so in a fundamentally different way than we learn possible languages, with the former treated like puzzles rather than languages.

To summarize, SP purports to show that MLMs refute Chomsky’s approach to language—a logical impossibility. What he does show is that there are multiple aspects adult English competence that ChatGPT is unable to simulate, and the in the cases where ChatGPT was able to mimic an adult English speaker, there is no explanation as to how. Neither of these results are germane to either Chomsky’s approach to language or his theories of language, as Chomsky studies the human capacity for language, which MLMs tell us nothing about. More importantly, SP does not even address Chomsky’s actual critique of MLM *qua* models of language competence.

## References

- Chomsky, Noam, Gary Marcus, and Jeremy Kahn. 2022. “Debunking the Great Ai Lie.” Youtube. 2022. [https://www.youtube.com/watch?v=PBdZi\\_JtV4c](https://www.youtube.com/watch?v=PBdZi_JtV4c).
- Piandadosi, Steven T. 2023. “Modern Language Models Refute Chomsky’s Approach to Language.” <https://ling.auf.net/lingbuzz/007180>.

---

<sup>4</sup>Setting aside languages which are logical impossibilities, like a language which has and lacks determiners.

Smith, Neilson Voyne, and Ianthi-Maria Tsimpli. 1995. *The Mind of a Savant: Language Learning and Modularity*. Blackwell Publishing.