Large Language Models: The best linguistic theory, a wrong linguistic theory, or no linguistic theory at all?

To appear in Zeitschrift für Sprachwissenschaft
Draft

Comments welcome!

Stefan Müller

September 19, 2024

Abstract

This paper discusses Ambridge & Blything's claim (2024) that Large Language Models are the best linguistic theory we currently have. It discusses claims that LLMs are wrong linguistic theories and concludes that they are not a linguistic theory at all. It is pointed out that Chomsky's claims about innateness, about transformations as underlying mechanisms of the language faculty and about plausible representations of linguistic knowledge are known to be flawed by quite some time by now and that we would not have needed LLMs for this. Chomsky's theories are not refuted by LLMs in their current form, since they are different in many aspects from human brains. However, the tremendous success of LLMs in terms of applications makes it more plausible to linguists and laymen that the innateness claims are wrong.

It is argued that the use of LLMs is probably limited when it comes to typological work and cross-linguistic generalizations. These require work in theoretical linguistics.

In a recent paper in a special issue of Theoretical Linguistics containing "Reflections on Theoretical Linguistics" on the occasion of the 50th anniversary of the journal Ambridge & Blything (2024) claim that "large language models are better than theoretical linguists at theoretical linguistics" (p. 33). The authors examine the output of an LLM with regard to the argument structures of verbs, and are impressed that the model predicts the same as Ambridge found out in experiments with students. The authors claim that LLMs are a theory of language. The best one we have right now:

> large language models (LLMs) are already the leading current theories of how speakers learn and represent these restrictions. Of course,

> they are not perfect theories [...] but they're better theories than any others that have been proposed. Ambridge & Blything (2024: 34)

LLMs are very interesting and you can do a lot of impressive things with them.[1] But are they theories? Do they help in any way to get a better understanding of language?

The authors claim that large language models are theories of language acquisition and representation and that they are instantiations of Construction Grammar (Goldberg 2006) approaches:

> Large language models [...] constitute theories of language acquisition and representation; theories that instantiate exemplar-, input- and construction-based approaches, though only very loosely. (Ambridge & Blything 2024: 33)

The authors claim that models are a theory (see also Piantadosi 2024: 360):

> OK, so the model makes the right predictions but – we hear you ask – where is the theory? That's the point: the model is the theory. (Ambridge & Blything 2024: 39)

This shows some confusion in terminology. Modeltheoretic approaches assume that there is a theory and that the theory can be used to build models (Richter 2021). A theory should use primitives that are appropriate for a certain domain and it should contain statements about these primitives. Large language models are neuronal nets that have been organized and trained in a certain way. Nodes of such nets can be examined and we can even find certain information in them (Manning et al. 2020, Zhang & Bowman 2018) but it is not a theory.

Later in the paper and contradicting their earlier claim, the authors claim that the programs that generate the LLMs are a theory:

> "But", critics object, "we have no idea what it's doing" (e.g., Kodner et al. 2023; Milway 2023). Quite the opposite: Unlike for traditional linguistic theories, every last detail of the model's assumptions and operation is written out in black and white, in thousands of lines of computer code. This code is a theory of the acquisition of (among other things) verb argument structure; it's even – like traditional linguistic theories – written in a language, albeit an artificial programming language, rather than a natural language like English. We know exactly what the model is doing. (Ambridge & Blything 2024: 39–40)

This very quote is an instance of mixing levels. We know what the *code* is doing. We do not know what the trained net, the model, is doing. This depends on the training data and even if we knew the training data, we could not predict what specific nodes in the net would do. The issue is just too complex for us humans

[1] For example, ChatGPT can explain prime factorizations in Trump-style (Piantadosi 2024: 356–357).

and the training data is too vast. This can be compared with Definite Clause Grammars: this is a notation that can be used to write down phrase structure grammars. Most Prolog interpreters come with a component that parses Definite Clause Grammars directly (see Clocksin & Mellish 1984: Chapter 9 for more information on DCGs and Müller 2023c: Task 10 on p. 81 for more information on working with DCGs online). Clocksin & Mellish (1984: 268–270) provide two pages of code for the translation of DCGs into Prolog code. The resulting Prolog program does Parsing as Deduction. In this case, we know what the code is doing. It reflects our theory about language. For a more elaborate example of Parsing as Deduction based on Government & Binding see Johnson (1989).

Fox & Katzir (2024) published a response to Ambridge & Blything (2024) in the same issue of Theoretical Linguistics. They write:

> The distinction between competence and performance and between correctness and likelihood are parts of all the best theories of human linguistic cognition, as are the aspects of linguistic representation that we briefly reviewed (modularity, constituency, and entailment). […] the LLM Theory does not even come close to approximating the relevant observations. Obviously it cannot derive these properties of human linguistic cognition and without doing so it cannot be considered a scientific theory at all. (Fox & Katzir 2024: 75)

The authors claim that LLMs cannot be a theory, since they do not make the competence-performance distinction, since they do not adhere to modularity and since they do not capture constituency. If these failures to capture certain properties of language would indeed entail that LLMs are not scientific theories, then neither Construction Grammar nor any flavor of Mainstream Generative Grammar (MGG)[2] would be. The distinction between competence and performance is rejected in Usage-Based Construction Grammar (Diessel 2015: 297). I personally think that this is a mistake (Müller 2023c: Chapter 15), but nevertheless approaches in Usage-based Construction Grammar are theories. The alternative approaches in MGG do not fare any better. All basic architectural assumptions in all of Chomsky's approaches are highly implausible from a psycholinguistic point of view. The Derivational Theory of Complexity, which assumed that sentences involving more transformations in their analysis are more difficult to process than sentences with fewer transformations has been proven wrong (Fodor et al. 1974: 320–328, Müller 2023c: Chapter 15.1). The T-model with its autonomous components of syntax, phonological and logical form has been proven wrong resulting in spectacular analyses in Cartography (Cinque & Rizzi 2010) to circumvent the autonomy of syntax restriction (Müller 2023b: Section 4.10.2). Derivation by Phase (Chomsky 2008) and other Minimalist architectures (Richards 2015: 812, 830) are entirely implausible as architectures

[2]The term MGG goes back to Culicover & Jackendoff (2005: 3). It refers to all proposals developed by Chomsky. Government & Binding (Chomsky 1981) and Minimalism Chomsky 1995 are the most recent incarnations.

for human language (Borsley & Müller 2021: Section 3.6), since they are incompatible with incremental processing (Marslen-Wilson 1975, Tanenhaus et al. 1996, Labelle 2007). If the argumentation by Fox & Katzir (2024) was valid, it would follow that all approaches in Usage-based Construction Grammar and MGG were not scientific theories. This would be a very strange conclusion, but it is not warranted. They are scientific theories, but they are bad ones.

Concerning the other points raised in the above quote: the claims about modularity are probably wrong (Pulvermüller 1999, Pulvermüller et al. 2013) and there are theories in which constituency does not play a role but dependency does (Tesnière 1959). And Manning et al. (2020) show that dependency information is encoded in LLMs.[3]

So Fox & Katzir (2024) argue that the theory, if it existed in LLMs, would be wrong. I argue that there is no theory about language in it. I believe that Ambridge & Blything (2024) are fundamentally wrong. To show this let us do a thought experiment. LLMs are neural networks. Their architecture is inspired by what we find in brains. They differ from brains in various ways, but let us assume that one could develop a perfect replica of a brain one day. To quote Norbert Wiener, the founder of cybernetics: "The best model of a cat is a cat, preferably the same cat." So, if we have a perfect copy of somebodies brain, what can we do with it? The artifical brain can then do exactly what the 48 Liverpool students mentioned in the Ambridge & Blything (2024) paper can do. Perhaps a bit more smoothed, because this replica can be fed much, much, much more data than all Liverpool students will ever see in their 48 lives combined. Now, the question is: What does this mean for linguistics? Is a replica of a brain a theory about language? No. It is a masterpiece of engineering. Nothing more. To build such masterpieces, you need theories about how brains work. You can then take parts of these theories and use them to build artificial brains. The code that people write to train the data structures they have created is code that is motivated by theories about the brain. It is not a theory, and certainly not a theory about language. The criticism that Ambridge & Blything (2024) reject is justified: LLMs are not theories about language; the information contained in LLMs is accessible only indirectly. Just as you cannot directly access the information in brains. You can only research the behavior of people. That's what we've been doing for hundreds of years. We look at what people say and write. We conduct experiments with people. We ask them about the acceptability of sentences. We test where they look when certain sentences are uttered. We measure brain activity (event-related potentials, cerebral bloodflow, etc.) We investigate what happens when certain areas of the brain are damaged. This gives us information about the processes and representations of linguistic knowledge in the brain. From this we can then draw conclusions for plausible theories.

[3]It is important to note that Manning et al. (2020) were able to discover the fact that dependencies are represented in LLMs because they knew the concept of dependencies, which was developed by Tesnière in 1937. So the linguistic theory and related concepts were a prerequisite to find linguistic structure in the neural networks. This point will be taken up again below in the discussion of typological work.

What is it like with LLMs? They are like brains: black boxes. We could start playing around with them now and try to find out what is stored where and how. But what good would that do? Actually such a research field exists already. Bender & Koller (2020: 5185) call it "BERTology": Engineers and linguists are playing with LLMs and check what they can do. This is interesting, but irrelevant for linguistics.[4]

Conclusion: We (as humanity) have created a technical masterpiece, but we know no more about our cognitive abilities than we did before.

Maybe the last sentence needs a bit of qualification. Piantadosi (2024) claims that Chomskyan linguistics has failed, that it was proven wrong by Large Language Models. As Piantadosi (2024: 366) writes himself LLMs "are trained on truly titanic datasets compared to children, by a factor of at least a few thousand". So, if linguistics is dealing with human capabilities, we are not quite there yet. To model language acquisition, we would need grounded input, we would need a realistic amount of training data, we would have to simulate the development of brains and the growth of cognitive capacities in early childhood. But what the success of LLMs suggests is that an elaborated component of Universal Grammar is not needed, that the argument of the Poverty of the Stimulus was flawed and so on. Above I wrote that "we know no more about our cognitive abilities than we did before". And this is true. We knew in 1974 (50 years ago) that transformations are psycholinguistically implausible (Fodor, Bever & Garrett 1974: 320–328). Psycholinguists sympathetic with the Chomskyan paradigm suggested that we have our linguistic knowledge represented as a Transformational Grammar, but that it then gets compiled out into a set of templates that are equivalent ot the constructions of Construction Grammar (Frazier & Clifton 1996: 27). But this of course begs the question why one should not work in Construction Grammar or a related framework like Constructional HPSG (Sag 1997, Müller et al. 2021) from the beginning. What is the evidence for some underlying transformation-based representation of linguistic knowledge? The various architectures that were proposed over the years were psycholinguistically implausible too. The T-Model (Chomsky 1981, 1986) was implausible (Müller 2023c: Section 15.2) and this got only worse with Phase-based variants of Minimalism (Chomsky 1995, 2008, Richards 2015: 812, 830, Borsley & Müller 2021: Section 5). But if the theories are incompatible with empirical facts like incremental processing, how can they tell us anything about human cognition and inateness? The Principle & Parameter model of language

[4]This was a bit of a hyperbole. LLMs may be used to play around with data and to check what these models need as input to get certain facts about language right. This can help linguists to discover relations and dependencies between linguistic phenomena that are plausible parts of a linguistic theory (generalizations, constructions, families of constructions and the relations between constructions). For example, Misra & Mahowald (2024) show that LLMs perform above chance on phrases like *a five beautiful days*, provided certain other constructions are in the training corpus. So, the place of LLMs in linguistics seems to be the one of subjects that one can feed arbitrary training material and that one can interrogate without them getting tired and without the need of an ethics vote. Since LLMs are different from real humans, the resulting theories should be checked with reference to actual data and actual human behavior, but they can serve as a first inspiration.

acquisition Chomsky (1981: 6) failed in various respects. It was assumed that one parameter was related to many properties of a language and worked like a switch (Rizzi 1986, Chomsky 2000: 8), but none of the suggested correlations held up (Haider 1994, 2001: Section 2.2, Müller 2023c: Section 16.1). The way parameterization was conceptualized was biologically implausible. For example, it was assumed that Subjacency was a universal principle and the parameterization concerned the part of speech of certain boundaries (Chomsky 1973: 271, 1986: 40, Baltin 1981, 2006). First, it could be shown that subjacency does not hold in Dutch, German and English (Koster 1978: 52, Müller 1999: 211, 2004, 2007: Section 3, Strunk & Snider 2013) and second, it is biologically absolutely implausible that part of speech information is encoded in our genes (Bishop 2002, Dąbrowska 2004: Section 6.4.2.2, Fisher & Marcus 2005). This was realized by Hauser, Chomsky & Fitch (2002). What remained as property that was assumed to be part of Universal Grammar was Merge, an operation for combining linguistic material. Somehow a triviality (Müller 2023c: 475). A triviality that caused another linguistic war (Pullum 2024).

There is one important aspect of research in the Principles & Parameters era: the systematic search for universals, for commonalities and differences lead to a much improved knowledge about variation. We know much more about language as such, that is, about structures that are similar in principle. For example, the German sentence in (1) is parallel to the English translation.

(1) dass die Straßenbahnen um die Ecke quietschen
that the trams around the corner squeak
'that the trams squeak around the corner'

As Müller (2013) pointed out, it is possible to develop analyses that capture the commonalities although the linearization of the constituents differs in German and English (English is an SVO language and German is SOV). Typological research is fascinating and requires the comparison of many very different languages on a theoretical level. I doubt that the results of cross-linguistic research can be derived from LLMs, without any interaction with theoretical linguistics. Training LLMs on multilingual material will be non-trivial and discovering cross-linguistic generalizations in network representations is probably impossible without a theoretically informed clue on what to look for (see also footnote 3). A suggestion for a methodological clean way of deriving cross-linguistic generalizations that differs from the MGG approach is assumed in the CoreGram project (Müller 2015).

Chomsky claimed that there would be language universals but there are no plausible candidates for syntactic universals left (Evans & Levinson 2009, Müller 2023c: Section 13.1). There are tendencies, for sure, but this is not sufficient for positing innate knowledge of language. Recently, I discovered a universal, but it is not syntactic but rather on the text level: the Festschrift Universal. The reader is referred to Müller (2024) for details.

The strongest argument for innate linguistic knowledge seemed to be the Argument from the Poverty of the Stimulus, but it was never actually correctly carried out (Pullum & Scholz 2002, Scholz & Pullum 2002). Chomsky

repeated his favourite argument with question formation as auxiliary inversion throughout several decades (Chomsky 1971: 29–33, 2013: 39). Bod (2009) showed how frequencies of subtrees can be used to learn structures of auxiliary inversion even though the examples that Chomsky (wrongly[5]) claimed to be non-existent in the data were not used in the learning procedure. Chomsky ignored these insights (Berwick, Pietroski, Yankama & Chomsky 2011, Chomsky 2013: 39) and so we find the auxiliary inversion claim again two years later in the same journal that also published Bod's paper. Similarly pattern-based modeling language acquisition research was much more successful in explaining cross-linguistic differences in acquisition than alternative accounts couched in Chomskyan frameworks (Freudenthal et al. 2006, 2007, 2009).

So, we knew that Chomskyan approaches to language and language acquisition failed in terms of their basic assumptions (transformations), they failed in terms of their architecture with respect to psycholinguistic evidence (separation of syntax and phonology and semantics in various forms) and they failed in respect to assumptions about genetics. Everybody working in non-Chomskyan paradigms knows this for more than a decade (see Müller 2010, 2016 for a summary of the respective discussions in German and English, respectively). We did not need LLMs for this, but maybe the actual usefulness of these networks is that they make the possibility that we do not need any innate domain-specific knowledge plausible to everybody: linguists and laymen. To show that LLMs can acquire languages like humans do, they have to be more human-like. To reach this goal, we probably need more knowledge about brains. As I pointed out above, if we manage to reach the goal of creating more human-like models, we know how brains work, but we do not necessarily know how languages work.

I believe that linguistic theories should contain rules and symbols. A linguistic theory can to some extent be derived from large corpora using automatic methods. Both the categories can be obtained via class formation and rules or valence patterns and the corresponding lexical entries can be derived automatically. The parts of speech and features like case, gender and number that are currently used in linguistic theories are basically the outcome of a distributional analysis that was done "by hand" during the last centuries. Grammar rules and also feature-value pairs may be assigned probabilities (Jurafsky 1996). These can also be derived from corpora. This is complicated and the mathematics is not fully understood yet. But one can train the system on large amounts of data. The training procedure contains assumptions about language: there are categories, there are constituents. There will be a residuum of infrequent phenomena that will not be captured this way (for example apparent multiple frontings, see Müller 2003). Some fine-tuning will be required and this is where the linguist comes in: rare data may decide between various alternative theories of a language (Müller 2023a: Chapter 6).

What would be missing in such grammars is the meaning component: a distributional analysis provides one with distribution classes, with syntactic regularities of the language under consideration. This is true for LLMs and

[5]See Pullum & Scholz (2002: 41–45).

any other outcome of a distributional analysis unless semantic information is explicitly encoded in the input and linked to real world experiences.[6] Therefore it is really surprising to see Construction Grammarians praising large language models as theories of human language. Wasn't it Construction Grammariens who told everybody in the field that human cognition is grounded (Barsalou 2008) and that language is not just abstract syntax and cannot be learned as such without a connection to semantics and the real world (Klein 1986: 44, Tomasello 2003: 113, Ambridge & Lieven 2011: Section 4.2.3, 4.2.8)? With grasping the communicative intention and attention sharing? Klein pointed out in 1986 already that no human being could learn Chinese by sitting an a room continually exposed to Chinese from loodspeakers. This just would not work. But this is how LLMs learn: they just see masses and masses of text. As Bender & Koller (2020) pointed out: BERT and ChatGPT and the like do not have a clue about what they are "saying". Their representations do not have any connection to semantics, they are not grounded. ChatGPT is a bullshit machine in the sense of Hicks et al. (2024), it is not and it does not contain a linguistic theory, not even a wrong one.

Acknowledgements

I thank Mark Felfe, Joachim Jacobs, Hubert Haider, Tibor Kiss, Oliver Schallert and Giuseppe Varaschin for comments and Konstantin Schulz, Anke Lüdeling, and Martin Klotz for discussion and Tine Mooshammer for raising questions about Ambridge & Blything (2024).

Oliver Schallert suggested adding more discussion of Piantadosi's paper. Thanks for this, I think it really improved the paper.

References

Ambridge, Ben & Liam Blything. 2024. Large language models are better than theoretical linguists at theoretical linguistics. *Theoretical Linguistics* 50(1–2). 33–48. DOI: 10.1515/tl-2024-2002.

Ambridge, Ben & Elena V. M. Lieven. 2011. *Child language acquisition: Contrasting theoretical approaches*. Cambridge, UK: Cambridge University Press. DOI: 10.1017/CBO9780511975073.

Baltin, Mark. 1981. Strict bounding. In Carl Lee Baker & John J. McCarthy (eds.), *The logical problem of language acquisition* (Cognitive Theory and Mental Representation 1), 257–295. Cambridge, MA: MIT Press.

[6]There is semantic knowledge implicit in LLMs. Piantadosi (2024: 358) points out that it is interwoven with syntactic information. The important point when it comes to human cognition is that the semantic knowledge in LLMs is not grounded. See Chang & Maia (2001) for computational experiments on language acquisition with grounding in the framework of Construction Grammar and Steels (2003) for experiments with grounded communication in robotics.

Baltin, Mark. 2006. Extraposition. In Martin Everaert & Henk van Riemsdijk (eds.), *The Blackwell companion to syntax*, 1st edn. (Blackwell Handbooks in Linguistics), 237–271. Oxford: Blackwell Publishers Ltd. DOI: 10.1002/9780470996591.ch25.

Barsalou, Lawrence W. 2008. Grounded cognition. *Annual Review of Psychology* 59. 617–645. DOI: 10.1146/annurev.psych.59.103006.093639.

Bender, Emily M. & Alexander Koller. 2020. Climbing towards NLU: On meaning, form, and understanding in the age of data. In Dan Jurafsky, Joyce Chai, Natalie Schluter & Joel Tetreault (eds.), *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 5185–5198. Online: Association for Computational Linguistics. DOI: 10.18653/v1/2020.acl-main.463.

Berwick, Robert C., Paul Pietroski, Beracah Yankama & Noam Chomsky. 2011. Poverty of the Stimulus revisited. *Cognitive Science* 35(7). 1207–1242. DOI: 10.1111/j.1551-6709.2011.01189.x.

Bishop, Dorothy V. M. 2002. Putting language genes in perspective. *TRENDS in Genetics* 18(2). 57–59. DOI: 10.1016/S0168-9525(02)02596-9.

Bod, Rens. 2009. From exemplar to grammar: Integrating analogy and probability in language learning. *Cognitive Science* 33(5). 752–793. DOI: 10.1111/j.1551-6709.2009.01031.x.

Borsley, Robert D. & Stefan Müller. 2021. HPSG and Minimalism. In Stefan Müller, Anne Abeillé, Robert D. Borsley & Jean-Pierre Koenig (eds.), *Head-Driven Phrase Structure Grammar: The handbook* (Empirically Oriented Theoretical Morphology and Syntax 9), 1253–1329. Berlin: Language Science Press. DOI: 10.5281/zenodo.5599874.

Chang, Nancy C. & Tiago V. Maia. 2001. Grounded learning of grammatical constructions. In *Papers from the 2001 AAAI Spring Symposium on Learning Grounded Representations* (AAAI Technical Report SS-01-05). AAAI.

Chomsky, Noam. 1971. *Problems of knowledge and freedom*. New York, NY: Pantheon Books.

Chomsky, Noam. 1973. Conditions on transformations. In Stephen R. Anderson & Paul Kiparsky (eds.), *A Festschrift for Morris Halle*, 232–286. New York, NY: Holt, Rinehart & Winston.

Chomsky, Noam. 1981. *Lectures on government and binding* (Studies in Generative Grammar 9). Dordrecht: Foris Publications. DOI: 10.1515/9783110884166.

Chomsky, Noam. 1986. *Barriers* (Linguistic Inquiry Monographs 13). Cambridge, MA: MIT Press.

Chomsky, Noam. 1995. *The Minimalist Program* (Current Studies in Linguistics 28). Cambridge, MA: MIT Press. DOI: 10.7551/mitpress/9780262527347.001.0001.

Chomsky, Noam. 2000. *New horizons in the study of language and mind*. Cambridge, UK: Cambridge University Press. DOI: 10.1017/CBO9780511811937.

Chomsky, Noam. 2008. On phases. In Robert Freidin, Carlos P. Otero & Maria Luisa Zubizarreta (eds.), *Foundational issues in linguistic theory:*

Essays in honor of Jean-Roger Vergnaud (Current studies in linguistics series 45), 133–166. Cambridge, MA: MIT Press. DOI: 10.7551/mitpress/9780262062787.003.0007.

Chomsky, Noam. 2013. Problems of projection. *Lingua* 130. 33–49. DOI: 10.1016/j.lingua.2012.12.003.

Cinque, Guglielmo & Luigi Rizzi. 2010. The cartography of syntactic structures. In Bernd Heine & Heiko Narrog (eds.), *The Oxford handbook of linguistic analysis* (Oxford Handbooks in Linguistics), 51–65. Oxford: Oxford University Press. DOI: 10.1093/oxfordhb/9780199544004.013.0003.

Clocksin, William F. & Christopher S. Mellish. 1984. *Programming in Prolog*. 5th edn. Berlin: Springer-Verlag. DOI: 10.1007/978-3-642-55481-0.

Culicover, Peter W. & Ray Jackendoff. 2005. *Simpler Syntax* (Oxford Linguistics). Oxford: Oxford University Press. DOI: 10.1093/acprof:oso/9780199271092.001.0001.

Dąbrowska, Ewa. 2004. *Language, mind and brain: Some psychological and neurological constraints on theories of grammar*. Washington, D.C.: Georgetown University Press.

Diessel, Holger. 2015. Usage-based Construction Grammar. In Ewa Dąbrowska & Dagmar Divjak (eds.), *Handbook of cognitive linguistics* (Handbücher zur Sprach- und Kommunikationswissenschaft 39), 296–322. Berlin: Mouton de Gruyter. DOI: 10.1515/9783110292022-015.

Evans, Nicholas & Stephen C. Levinson. 2009. The myth of language universals: Language diversity and its importance for cognitive science. *The Behavioral and Brain Sciences* 32(5). 429–448.

Fisher, Simon E. & Gary F. Marcus. 2005. The eloquent ape: Genes, brains and the evolution of language. *Nature Reviews Genetics* 7(1). 9–20. DOI: 10.1038/nrg1747.

Fodor, Jerry A., Thomas G. Bever & Merrill F. Garrett. 1974. *The psychology of language: An introduction to psycholinguistics and Generative Grammar* (McGraw-Hill series in psychology). New York, NY: McGraw-Hill Book Co.

Fox, Danny & Roni Katzir. 2024. Large language models and theoretical linguistics. *Theoretical Linguistics* 50(1–2). 71–76. DOI: 10.1515/tl-2024-2005.

Frazier, Lyn & Charles Clifton Jr. 1996. *Construal* (Language, Speech, and Communication 3). Cambridge, MA: MIT Press.

Freudenthal, Daniel, Julian M. Pine, Javier Aguado-Orea & Fernand Gobet. 2007. Modeling the developmental patterning of finiteness marking in English, Dutch, German, and Spanish using MOSAIC. *Cognitive Science* 31(2). 311–341. DOI: 10.1080/15326900701221454.

Freudenthal, Daniel, Julian M. Pine & Fernand Gobet. 2006. Modeling the development of children's use of optional infinitives in Dutch and English using MOSAIC. *Cognitive Science* 30(2). 277–310. DOI: 10.1207/s15516709cog0000_47.

Freudenthal, Daniel, Julian M. Pine & Fernand Gobet. 2009. Simulating the referential properties of Dutch, German, and English root infinitives in MOSAIC. *Language Learning and Development* 5(1). 1–29. DOI: 10.1080/15475440802502437.

Goldberg, Adele E. 2006. *Constructions at work: The nature of generalization in language* (Oxford Linguistics). Oxford: Oxford University Press.

Haider, Hubert. 1994. (Un-)heimliche Subjekte: Anmerkungen zur Pro-drop Causa, im Anschluß an die Lektüre von Osvaldo Jaeggli & Kenneth J. Safir, eds., The Null Subject Parameter. *Linguistische Berichte* 153. 372–385.

Haider, Hubert. 2001. Language typology and language universals: An international handbook. In Martin Haspelmath, Ekkehard König, Wulf Oesterreicher & Wolfgang Raible (eds.), *Language typology and language universals: An international handbook*, vol. 1 (Handbooks of Linguistics and Communication Science 20), 283–293. Berlin: Mouton de Gruyter. DOI: 10.1515/9783110194036-009.

Hauser, Marc D., Noam Chomsky & W. Tecumseh Fitch. 2002. The faculty of language: What is it, who has it, and how did it evolve? *Science* 298(5598). 1569–1579. DOI: 10.1126/science.298.5598.1569.

Hicks, Michael Townsen, James Humphries & Joe Slater. 2024. ChatGPT is bullshit. *Ethics and Information Technology* 26(2). 1–10. DOI: 10.1007/s10676-024-09775-5.

Johnson, Mark. 1989. Parsing as deduction: The use of knowledge of language. *Journal of Psycholinguistic Research* 18(1). 105–128.

Jurafsky, Daniel. 1996. A probabilistic model of lexical and syntactic access and disambiguation. *Cognitive Science* 20(2). 137–194. DOI: 10.1207/s15516709cog2002_1.

Klein, Wolfgang. 1986. *Second language acquisition* (Cambridge Textbooks in Linguistics). Cambridge, UK: Cambridge University Press. DOI: 10.1017/CBO9780511815058.

Koster, Jan. 1978. *Locality principles in syntax* (Studies in Generative Grammar 5). Dordrecht: Foris Publications. DOI: 10.1515/9783110882339.

Labelle, Marie. 2007. Biolinguistics, the Minimalist Program, and psycholinguistic reality. *Snippets* 14. 6–7. http://www.ledonline.it/snippets/allegati/snippets14002.pdf (21 February, 2022).

Manning, Christopher D., Kevin Clark, John Hewitt & Omer Levy. 2020. Emergent linguistic structure in artificial neural networks trained by self-supervision. *Proceedings of the National Academy of Sciences* 48(117). 30046–30054. DOI: 10.1073/pnas.1907367117.

Marslen-Wilson, William D. 1975. Sentence perception as an interactive parallel process. *Science* 189(4198). 226–228. DOI: 10.1126/science.189.4198.226.

Misra, Kanishka & Kyle Mahowald. 2024. *Language models learn rare phenomena from less rare phenomena: The case of the missing AANNs*. https://arxiv.org/abs/2403.19827.

Müller, Stefan. 1999. *Deutsche Syntax deklarativ: Head-Driven Phrase Structure Grammar für das Deutsche* (Linguistische Arbeiten 394). Tübingen: Max Niemeyer Verlag. DOI: 10.1515/9783110915990.

Müller, Stefan. 2003. Mehrfache Vorfeldbesetzung. *Deutsche Sprache* 31(1). 29–62.

Müller, Stefan. 2004. Complex NPs, subjacency, and extraposition. *Snippets* 8. 10–11.

Müller, Stefan. 2007. Qualitative Korpusanalyse für die Grammatiktheorie: Introspektion vs. Korpus. In Gisela Zifonun & Werner Kallmeyer (eds.), *Sprachkorpora – Datenmengen und Erkenntnisfortschritt* (Institut für Deutsche Sprache Jahrbuch 2006), 70–90. Berlin: Walter de Gruyter. DOI: 10.1515/9783110439083-006.

Müller, Stefan. 2010. *Grammatiktheorie* (Stauffenburg Einführungen 20). Tübingen: Stauffenburg Verlag.

Müller, Stefan. 2013. Unifying everything: Some remarks on Simpler Syntax, Construction Grammar, Minimalism and HPSG. *Language* 89(4). 920–950. DOI: 10.1353/lan.2013.0061.

Müller, Stefan. 2015. The CoreGram project: Theoretical linguistics, theory development and verification. *Journal of Language Modelling* 3(1). 21–86. DOI: 10.15398/jlm.v3i1.91.

Müller, Stefan. 2016. *Grammatical theory: From Transformational Grammar to constraint-based approaches*. 1st edn. (Textbooks in Language Sciences 1). Berlin: Language Science Press. DOI: 10.17169/langsci.b25.167.

Müller, Stefan. 2023a. *German clause structure: An analysis with special consideration of so-called multiple fronting* (Empirically Oriented Theoretical Morphology and Syntax). Berlin: Revise and resubmit Language Science Press.

Müller, Stefan. 2023b. *Germanic syntax: A constraint-based view* (Textbooks in Language Sciences 12). Berlin: Language Science Press. DOI: 10.5281/zenodo.7733033.

Müller, Stefan. 2023c. *Grammatical theory: From Transformational Grammar to constraint-based approaches*. 5th edn. (Textbooks in Language Sciences 1). Berlin: Language Science Press. DOI: 10.5281/zenodo.7376662.

Müller, Stefan. 2024. Anarchy, power, festschrifts, and universals. In Edward Gibson & Moshe Poliak (eds.), *From fieldwork to linguistic theory: A tribute to Dan Everett* (Empirically Oriented Theoretical Morphology and Syntax 15), ix–xxii. Berlin: Language Science Press. DOI: 10.5281/zenodo.12665903.

Müller, Stefan, Anne Abeillé, Robert D. Borsley & Jean-Pierre Koenig (eds.). 2021. *Head-Driven Phrase Structure Grammar: The handbook* (Empirically Oriented Theoretical Morphology and Syntax 9). Berlin: Language Science Press. DOI: 10.5281/zenodo.5543318.

Piantadosi, Steven T. 2024. Modern language models refute Chomsky's approach to language. In Edward Gibson & Moshe Poliak (eds.), *From fieldwork to linguistic theory: A tribute to Dan Everett* (Empirically Oriented Theoretical Morphology and Syntax 15), 353–414. Berlin: Language Science Press. DOI: 10.5281/zenodo.12665933.

Pullum, Geoffrey K. 2024. Daniel Everett on Pirahã syntax. In Edward Gibson & Moshe Poliak (eds.), *From fieldwork to linguistic theory: A tribute to Dan Everett* (Empirically Oriented Theoretical Morphology and Syntax 15), 23–74. Berlin: Language Science Press. DOI: 10.5281/zenodo.12665907.

Pullum, Geoffrey K. & Barbara C. Scholz. 2002. Empirical assessment of stimulus poverty arguments. *The Linguistic Review* 19(1–2). 9–50. DOI: 10.1515/tlir.19.1-2.9.

Pulvermüller, Friedemann. 1999. Words in the brain's language. *Behavioral and Brain Sciences* 22(2). 253–279. DOI: 10.1017/S0140525X9900182X.

Pulvermüller, Friedemann, Bert Cappelle & Yury Shtyrov. 2013. Brain Basis of Meaning, Words, Constructions, and Grammar. In Thomas Hoffmann & Graeme Trousdale (eds.), *The Oxford handbook of Construction Grammar* (Oxford Handbooks), 397–416. Oxford: Oxford University Press. DOI: 10.1093/oxfordhb/9780195396683.013.0022.

Richards, Marc. 2015. Minimalism. In Tibor Kiss & Artemis Alexiadou (eds.), *Syntax – theory and analysis: An international handbook*, vol. 2 (Handbooks of Linguistics and Communication Science 42), 803–839. Berlin: Mouton de Gruyter. DOI: 10.1515/9783110363708-001.

Richter, Frank. 2021. Formal background. In Stefan Müller, Anne Abeillé, Robert D. Borsley & Jean-Pierre Koenig (eds.), *Head-Driven Phrase Structure Grammar: The handbook* (Empirically Oriented Theoretical Morphology and Syntax 9), 89–124. Berlin: Language Science Press. DOI: 10.5281/zenodo.5599822.

Rizzi, Luigi. 1986. Null objects in Italian and the theory of *pro*. *Linguistic Inquiry* 17(3). 501–577.

Sag, Ivan A. 1997. English relative clause constructions. *Journal of Linguistics* 33(2). 431–483. DOI: 10.1017/S002222679700652X.

Scholz, Barbara C. & Geoffrey K. Pullum. 2002. Searching for arguments to support linguistic nativism. *The Linguistic Review* 19(1–2). 185–223. DOI: 10.1515/tlir.19.1-2.185.

Steels, Luc. 2003. Evolving grounded communication for robots. *Trends in Cognitive Science* 7(7). 308–312.

Strunk, Jan & Neal Snider. 2013. Subclausal locality constraints on relative clause extraposition. In Gert Webelhuth, Manfred Sailer & Heike Walker (eds.), *Rightward movement in a comparative perspective* (Linguistik Aktuell/Linguistics Today 200), 99–143. Amsterdam: John Benjamins Publishing Co. DOI: 10.1075/la.200.

Tanenhaus, Michael K., Michael J. Spivey-Knowlton, Kathleen M. Eberhard & Julie C. Sedivy. 1996. Using eye movements to study spoken language comprehension: Evidence for visually mediated incremental interpretation. In Toshio Inui & James L. McClelland (eds.), *Information integration in perception and communication* (Attention and Performance XVI), 457–478. Cambridge, MA: MIT Press.

Tesnière, Lucien. 1959. *Eléments de syntaxe structurale*. Paris: Librairie C. Klincksieck. Republished as *Elements of structural syntax*. Translated by Timothy Osborne and Sylvain Kahane. Amsterdam: John Benjamins Publishing Co., 2015. DOI: 10.1075/z.185.

Tomasello, Michael. 2003. *Constructing a language: A usage-based theory of language acquisition*. Cambridge, MA: Harvard University Press.

Zhang, Kelly & Samuel Bowman. 2018. Language modeling teaches you more than translation does: Lessons learned through auxiliary syntactic task analysis. In Tal Linzen, Grzegorz Chrupała & Afra Alishahi (eds.), *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, 359–361. Brussels, Belgium: Association for Computational Linguistics. DOI: 10.18653/v1/W18-5448.