

Article

# Natural Syntax, Artificial Intelligence and Language Acquisition

William O'Grady<sup>1,\*</sup> and Miseon Lee<sup>2</sup><sup>1</sup> Department of Linguistics, University of Hawai'i at Manoa, Honolulu, HI 96822, USA<sup>2</sup> Department of English Language & Literature, Hanyang University, Seoul 04763, Republic of Korea; mlee@hanyang.ac.kr

\* Correspondence: ogrady@hawaii.edu; Tel.: +1-808-956-8602

**Abstract:** In recent work, various scholars have suggested that large language models can be construed as input-driven theories of language acquisition. In this paper, we propose a way to test this idea. As we will document, there is good reason to think that processing pressures override input at an early point in linguistic development, creating a temporary but sophisticated system of negation with no counterpart in caregiver speech. We go on to outline a (for now) thought experiment involving this phenomenon that could contribute to a deeper understanding both of human language and of the language models that seek to simulate it.

**Keywords:** artificial intelligence; language acquisition; language models; Natural Syntax

## 1. Introduction

A striking feature of contemporary large language models (LLMs)—ChatGPT, GPT4, BERT and others—is their remarkable grammatical accuracy, which seems to surpass the abilities of even the most articulate of humans.

The key point that we want to make is that the output of [large language models] is, almost without exception, grammatical. Even when examining experiments that are designed to exhibit some of the model's shortcomings . . . , one cannot help but notice that the content is grammatically correct . . . [1] (p. 3)

How could such a level of success have been achieved? The answer is surprising: LLMs are trained on gigantic bodies of text to predict the next word that they will encounter.

[LLMs] are trained only on text prediction. This means that the models form probabilistic expectations about the next word in a text and they use the true next word as an error signal to update their latent parameters. [2] (p. 5)

Training of this sort is not out of line with various proposals that have been made about language learning in humans, which also seems to be sensitive to predictions and feedback. (As a reviewer notes, the availability of input and feedback becomes even more vital in the case of multilingual acquisition, where exposure to speech is distributed over two or more languages, causing well-documented delays in the acquisition of each; see, e.g., [3].)

A learner makes predictions. If these predictions are not borne out, the difference between the predicted and actual input serves as implicit corrective feedback, pushing the learner more towards the target language. [4] (p. 4)

. . . if a learner can use his current knowledge of the language and the context to predict how the sentence will unfold, then the comparison of what he expects with what actually occurs could provide valuable additional information. [5]

[The model] learns from its predictions at each word position. Incremental prediction increases the amount of information that can be used for learning



**Citation:** O'Grady, W.; Lee, M. Natural Syntax, Artificial Intelligence and Language Acquisition. *Information* **2023**, *14*, 418. <https://doi.org/10.3390/info14070418>

Academic Editor: Katsuhide Fujita

Received: 15 May 2023

Revised: 8 July 2023

Accepted: 10 July 2023

Published: 20 July 2023



**Copyright:** © 2023 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

and focuses learning on the particular representations that made the incorrect prediction. The prediction-based approach is thus a potential solution to the problem of learning from sparse input. [6] (p. 263)

However, a deeper and more general question now arises. Proponents of generative grammar have long insisted that not all features of language can be acquired from the available input, especially when phenomena such as anaphora, quantifier scope, and filler-gap dependencies are in play. In such cases, they claim, some properties are not instantiated with sufficient frequency and clarity to be acquired from experience—they must be derived from inborn principles of universal grammar.

*The Argument from Poverty of Stimulus*

... crucial components of our tacit linguistic knowledge are not learned through experience but are given by our biological/genetic specifications [Universal Grammar]. [7] (p. 151)

In sharp contrast to this view, some cognitive scientists are now taking the position that LLMs are more than simple wonders of engineering; they embody a theory of data-based language acquisition that is at odds with the UG model.

The success of large language models is a failure for generative theories because it goes against virtually all of the principles these theories have espoused. In fact, none of the principles and innate biases that Chomsky and those who work in his tradition have long claimed necessary needed to be built into these models. [2] (pp. 14–15)

We suggest that the most recent generation of Large Language Models (LLMs) might finally provide the computational tools to determine empirically how much of the human language ability can be acquired from linguistic experience [1] (p. 1).

Rapid progress in machine learning for natural language processing has the potential to transform debates about how humans learn language [8] (p. 18).

In our view, both perspectives represent extreme positions, neither of which is likely to maximize our understanding of the human language faculty. In their stead, we will outline an approach to language and learning which posits biases that go beyond the capacity to learn from experience but do not involve reference to universal grammar.

We will proceed as follows. In the next section, we will draw on the syntax of negation to outline an approach to language and learning that is sensitive to the economy of effort, a core feature of human cognition in general. Section 3 then examines a puzzling stage in the acquisition of negation that seems to reflect the presence of inborn processing biases, opening the door for an empirical test whose form and design we go on to outline as a thought experiment. We end, in Sections 4 and 5, with some brief concluding remarks and a post-script on the relationship between syntax and AI.

## 2. Natural Syntax

Our starting point is the widely held view that complex systems, including language, arise from the interaction of simpler and more basic forces and factors. In the case of the syntax of human language, these appear to consist largely of processing pressures that reduce the cost of the operations required for speech and comprehension—an idea that has been developed in considerable detail in work by [9–12], among others.

*Processing Determinism*

The algorithms that map form onto meaning (and vice versa) are shaped by forces that seek to minimize processing costs.

The advantages of cost minimization are both obvious and widely acknowledged.

Speed in communicating the intended message from speaker to hearer and minimal processing effort in doing so are the two driving forces of efficiency ... [9] (p. 9), [10] (p. 48)

... the [processor] should operate in the most efficient manner possible, promptly resolving dependencies so that they do not have to be held any longer than necessary. This is a standard assumption in work on processing, where it is universally recognized that sentences are built in real time under conditions that favor quickness. [13] (p. 7)

Efficiency in communication means that successful communication can be achieved with minimal effort on average by the sender and receiver [14].

It is important to distinguish this view from the claims underlying UG-based theories, whereas UG, as standardly understood, includes inborn principles of *grammar*, a processing cost bias is a broadly manifested feature of cognition in general.

... the avoidance [and minimization] of cognitive demand is a fundamental principle of cognition ... [15] (p. 2), citing [16]

A spate of recent work demonstrates that humans seek to avoid the expenditure of cognitive effort, much like physical effort ... [17] (p. 92)

We will henceforth use the term ‘Natural Syntax’ to refer to theories of sentence formation that adopt this perspective; see [11] for details or [12] for an abridged discussion (Ref. [18] (p. 135) defines ‘natural’ as ‘synonymous to cognitively simple, easily accessible (especially to children), elementary and therefore universally preferred, i.e., derivable from human nature ... ’). The syntax of negation provides an opportunity to explore the advantages of such an approach.

### *The Syntax of Negation*

The most common strategy for expressing negation in English makes use of the sentential negative *not*, which has the effect of denying the occurrence of an event or state.

Jane didn’t buy a Tesla. (=‘There was no buying event involving Jane and a Tesla.’)

A second strategy expresses negation with the help of a negative pronoun that denotes a null set. As illustrated in the example below, these items negate an event by indicating that (for example) its agent or patient denotes a null set and therefore has no referent.

Jane bought **nothing**.

Under certain circumstances, a sentential negative can co-occur with a negative pronoun. At least two types of patterns can be identified—double negation and negative concord.

### **Double Negation**

Double negation occurs when two negatives in the same clause cancel each other out, resulting in a positive interpretation. As [19] (pp. 58–59) and many others have noted, double negation is a quite rare occurrence, largely restricted to situations in which the speaker wishes to deny a negative claim made by someone else, as in the following example.

- A. Once again, you did nothing.
- B. I didn’t do *nothing*—I washed the car!  
(not nothing = something)

As illustrated here, the end result is an interpretation that can be paraphrased as ‘I did something’. For this to happen, the default null-set (=∅) interpretation of *nothing* must be transformed under the influence of the sentential negative—at an additional computational cost. (This is an instance of the phenomenon known as ‘scope’, in which one logical operator affects the interpretation of another).

Double negative interpretation (=‘I did something.’):

I didn’t do **nothing**.



Due to its scopal interaction with *not*, the negative pronoun takes on a ‘not nothing’ interpretation (= ‘something’).

### Negative Concord

Negative concord produces a very different result: the two negatives combine to yield a single negative interpretation. Korean is a case in point (TOP = topic, CON = connective, DECL = declarative).

*Korean*

Kunye-nun **amwukesto** ha-ci **anh**-ass-ta.  
 she-TOP nothing do-COMP not-PAST-DECL

‘She did nothing.’

Patterns of this type also occur in earlier forms of English and are still common in various non-standard varieties of Modern English (e.g., [20]).

*Middle English*

Ther **nys nat oon** kan war by other be.  
 There not-is not one can aware by other be  
 =‘There is no one who can be warned by another.’  
 (Chaucer: Troilus 1, 203; cited by [21] (p. 86))

*Non-standard Modern English*


Don’t blame me. I didn’t do **nothing** wrong.  
 =‘I did nothing wrong.’/‘I didn’t do anything wrong.’

A key feature of negative concord patterns is that there is no interaction between the two negatives; each simply carries out the function that it would have in sentences where it is the sole negative.

- i. The sentential negative signals the non-occurrence of an event, just as it does in sentences such as *I didn’t go* and their equivalent in other languages;
- ii. The negative pronoun has the same null-set interpretation that it has in stand-alone contexts (e.g., *What did you see? Nothing!*).

Negative concord interpretation (=‘I did nothing.’):

I didn’t do **nothing**.



As illustrated here, *nothing* manifests a stable null-set interpretation; there is no modification of its denotation, despite the presence of *not* in the same sentence. As we will see next, this fact sheds light on a curious phenomenon that arises in the acquisition of English.

### 3. The Acquisition of Negation

At least two quite distinct developmental paths appear to define the course of language acquisition. The first, which LLMs are able to navigate with considerable ease, involves what might be called ‘mundane acquisition’—the mastery of patterns that are widely manifested in the input (the ‘training set’). The second path, although less evident, is arguably more relevant for understanding the true character of language acquisition in humans.





- i. The Type A language model accurately mirrors children's two-stage acquisition process, with the data set for 4-year-olds, leading to the generation of negative concord patterns and the data set for older children correcting that tendency;
- ii. The Type A model fails, but the Type B language model, with its processing cost bias, accurately mirrors the two-stage process that we see in child language acquisition. (For now, we leave to the side the technical issue of precisely how a bias for processing cost should be incorporated into the language model.)

The first result would point toward the ability of a language model to simulate human-like development without incorporating a processing-cost bias, raising the question of whether the same might somehow be true of children as well. In contrast, the second outcome would have consequences in the opposite direction, suggesting that pure induction-based language models fall short of simulating the developmental process that produces language in children.

#### 4. Conclusions

According to the recollection of pioneering psychologist and neuroscientist George Miller [33], the field of cognitive science was conceived on 11 September 1956, at a symposium held at MIT under the auspices of the 'Special Interest Group in Information Theory'. As Miller noted, linguistics and computer science were central to the prospects of the newly created field.

I went away from the Symposium with a strong conviction, more intuitive than rational, that human experimental psychology, theoretical linguistics, and computer simulation of cognitive processes were all pieces of a larger whole, and that the future would see progressive elaboration and coordination of their shared concerns . . . [34], cited by [35].

In the spirit of Miller's remarks, we have taken seriously the possibility that induction-based language models have reached the point where their capacity to simulate child language acquisition deserves to be assessed empirically. With that in mind, we have proposed a test involving a basic and fundamental feature of language—the expression of negation. As we have noted, the interest in this particular phenomenon (and many others like it) lies in its well-documented developmental trajectory, which appears to be shaped by more than just exposure to language use.

Our test turns to a simple question. Can a language model with no built-in bias related to processing cost (or other comparable factors) accurately simulate children's acquisition of negation given exposure only to input comparable to what human learners receive in the first years of life? We hope and believe that the answer to this question will shed light on how and whether the joint study of human language and artificial intelligence might contribute to a deeper understanding of the issues confronting each field.

#### 5. Post-Script—A Syntax for AI?

Contemporary Natural Syntax has roots in three relatively understudied traditions within linguistics. The first, dating back to the 1970s, had as its starting point the observation by David Stampe [36] that humans, the only species capable of language, are in fact poorly equipped for speech due to limitations on their articulatory and perceptual mechanisms and on the computational capacity of their minds. Those limitations, Stampe argued, end up shaping the properties of language and the manner in which they are acquired by children. Cambridge University Press is scheduled to publish a comprehensive handbook of natural linguistics in 2024.

A second foundation for Natural Syntax lies in the field of sentence processing, whose traditional line of thinking now has a competitor. Instead of focusing on how syntactic principles define the processing strategies involved in sentence comprehension, there is serious consideration of the possibility that the reverse might be true: processing efficiency could shape syntactic principles. This idea has been pursued with particular success by John Hawkins (e.g., [9,10]).

*Performance-Grammar Correspondence Hypothesis* (paraphrased):

Grammars generate sentences whose properties reflect preferences relating to the ease and efficiency of processing. [10] (p. 3)

Finally, the roots of Natural Syntax can be traced to linguistic emergentism, especially the pioneering research of scholars such as Elizabeth Bates, Jeff Elman and Brian MacWhinney.

Just as the conceptual components of language may derive from cognitive content, so might the computational facts about language stem from . . . the multitude of competing and converging constraints imposed by perception, production, and memory for linear forms in real time [37] (pp. 189–190).

. . . in very significant ways, language is a radically new behavior. At a phenomenological level, it is quite unlike anything else that we (or any other species) do. It has features that are remarkable and unique. The crucial difference between this view and the view of language as a separable domain-specific module . . . is that the uniqueness emerges out of an interaction involving small differences in domain-nonspecific behaviors. [38] (p. 25)

In order to qualify as emergentist, an account of language functioning must tell us where a language behavior “comes from”. In most cases, this involves accounting for a behavior in a target domain as emerging from some related external domain. [39] (p. xii)

Contemporary Natural Syntax is built on three principles—direct mapping, algorithmic orientation and processing determinism [11,12].

*Direct Mapping*

The mapping between a sentence’s form and its meaning is direct; there is no role for syntactic structure. Instead, forms (strings of words) are mapped directly onto ‘proto-semantic’ representations consisting largely of predicate–argument complexes.

FORM		MEANING
<i>Harry left</i>	↔	LEAVE <h>

*Algorithmic Orientation*

The mapping between form and meaning is regulated by algorithms that operate in real time in the course of speaking and understanding.

*Processing Determinism*

The algorithms involved in mapping form onto meaning (and vice versa) are largely shaped by forces that contribute to processing efficiency.

There is a partial fit here with work on language-related AI that might well be worth pursuing, especially if the outcome of the thought experiment outlined earlier confirms the relevance of processing-cost biases to models of language acquisition. If that turns out to be so, Natural Syntax would align with the design of language models in at least three ways:

- i. Natural Syntax does not presuppose or make use of syntactic structure; sentences are just what they appear to be—strings of words, as also assumed in language models;
- ii. The algorithms of Natural Syntax are designed to operate in a way compatible with real-time processing—from the first word to the last, a common premise in language models as well;
- iii. The algorithms of Natural Syntax are shaped by processing considerations that (depending on the outcome of experiments such as the one suggested here) could open the door to a deeper understanding of the syntax of human language in general, including language acquisition and language change—also a goal of contemporary AI research.

As things now stand, of course, there are also important differences between Natural Syntax and language models—most obviously the need for at least a primitive semantic



representation and the commitment to the primacy of processing cost. Fortunately, though, these are empirical issues for which the possibility of discovering genuine answers is coming ever closer.

**Author Contributions:** Conceptualization, W.O. and M.L.; Writing—original draft, W.O.; Writing—review & editing, M.L. The authors contributed equally to this work. All authors have read and agreed to the published version of the manuscript.

**Funding:** This research received no external funding.

**Data Availability Statement:** No new data were created in this study.

**Acknowledgments:** We acknowledge with gratitude the comments provided by Peter MacKinnon and the two reviewers.

**Conflicts of Interest:** The authors declare no conflict of interest.

## References

1. Contreras Kallens, P.; Kristensen-McLachland, R.D.; Christiansen, M.H. Large language models demonstrate the potential of statistical learning in language. *Cogn. Sci.* **2023**, *47*, e13256. [CrossRef] [PubMed]
2. Piantadosi, S. Modern Language Models Refute Chomsky's Approach to Language. Unpublished ms. Department of Psychology, University of California, Berkeley, CA, USA. 2023. Available online: <http://colala.berkeley.edu/> (accessed on 1 May 2023).
3. Hoff, E.; Core, C.; Place, S.; Rumiche, R.; Señor, M.; Parra, M. Dual language exposure and bilingual acquisition. *J. Child Lang.* **2012**, *39*, 1–27. [CrossRef] [PubMed]
4. Kaan, E.; Grüter, T. Prediction in second language processing and learning: Advances and directions. In *Prediction in Second Language Processing and Learning*; Kaan, E., Grüter, T., Eds.; John Benjamins: Amsterdam, The Netherlands, 2021; pp. 1–24.
5. Phillips, C.; Ehrenhofer, L. The role of language processing in language acquisition. *Linguist. Approaches Biling.* **2015**, *5*, 409–453. [CrossRef]
6. Chang, F.; Dell, G.S.; Bock, K. Becoming syntactic. *Psychol. Rev.* **2006**, *113*, 234–272. [CrossRef] [PubMed]
7. Legate, J.; Yang, C. Empirical re-assessment of stimulus poverty arguments. *Linguist. Rev.* **2002**, *19*, 151–162. [CrossRef]
8. Warstadt, A.; Bowman, S. What artificial neural networks can tell us about human language acquisition. In *Algebraic Structures in Natural Language*; Lappin, S., Bernardy, J.-P., Eds.; Taylor & Francis Group: New York, NY, USA, 2023; pp. 17–59.
9. Hawkins, J. *Efficiency and Complexity in Grammars*; Oxford University Press: Oxford, UK, 2004.
10. Hawkins, J. *Cross-Linguistic Variation and Efficiency*; Oxford University Press: Oxford, UK, 2014.
11. O'Grady, W. *Natural Syntax: An Emergentist Primer*, 3rd ed.; 2022; Available online: <http://ling.hawaii.edu/william-ogrady> (accessed on 12 August 2022).
12. O'Grady, W. Working memory and natural syntax. In *The Cambridge Handbook of Working Memory and Language*; Schwieter, J., Wen, Z., Eds.; Cambridge University Press: Cambridge, UK, 2022; pp. 322–342.
13. O'Grady, W. *Syntactic Carpentry: An Emergentist Approach to Syntax*; Erlbaum: Mahwah, NJ, USA, 2005.
14. Gibson, E.; Futrell, R.; Piantadosi, S.; Dautriche, I.; Mahowald, K.; Bergen, L.; Levey, R. How efficiency shapes human language. *Trends Cogn. Sci.* **2019**, *23*, 389–407. [CrossRef] [PubMed]
15. Christie, S.T.; Schrater, P. Cognitive cost as dynamic allocation of energetic resources. *Front. Neurosci.* **2015**, *9*, 289. [CrossRef] [PubMed]
16. Kool, W.; McGuire, J.; Rosen, Z.; Botvinick, M. Decision making and the avoidance of cognitive demand. *J. Exp. Psychol. Gen.* **2010**, *139*, 665–682. [CrossRef]
17. Otto, A.R.; Daw, N. The opportunity cost of time modulates cognitive effort. *Neuropsychologia* **2019**, *123*, 92–105. [CrossRef]
18. Dressler, W. What is natural in natural morphology? *Prague Linguist. Circ. Pap.* **1999**, *3*, 135–144.
19. Zeijlstra, H. Sentential Negation and Negative Concord. Ph.D. Thesis, University of Amsterdam, Amsterdam, The Netherlands, 2004.
20. Robinson, M.; Thoms, G. On the syntax of English variable negative concord. *Univ. Pa. Work. Pap. Linguist.* **2021**, *27*, 24.
21. Fischer, O.; van Kemenade, A.; Koopman, W.; van der Wurff, W. *The syntax of Early English*; Cambridge University Press: Cambridge, UK, 2001.
22. Bellugi, U. The Acquisition of the System of Negation in Children's Speech. Ph.D. Thesis, Harvard University, Cambridge, MA, USA, 1967.
23. Thornton, R.; Notley, A.; Moscati, V.; Crain, S. Two negations for the price of one. *Glossa* **2016**, *1*, 45. [CrossRef]
24. Moscati, V. Children (and some adults) overgeneralize negative concord: The case of fragment answers to negative questions in Italian. *Univ. Pa. Work. Pap. Linguist.* **2020**, *26*, 169–178.
25. Zhou, P.; Crain, S.; Thornton, R. Children's knowledge of double negative structures in Mandarin Chinese. *J. East Asian Ling.* **2014**, *23*, 333–359. [CrossRef]

26. van Schijndel, M.; Mueller, A.; Linzen, T. Quantity doesn't buy quality syntax with neural language models. In Proceedings of the Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), Hong Kong, China, 3–7 November 2019; pp. 5831–5837.
27. Linzen, T. How can we accelerate progress toward human-like linguistic generalization? In Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, Online, 5–10 July 2020; pp. 5210–5217.
28. Zhang, Y.; Warstadt, A.; Li, H.-S.; Bowman, S. When do you need billions of words of pretraining data? In Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, Bangkok, Thailand, 1–6 August 2021; Volume 1, pp. 1112–1125.
29. Hosseini, E.; Schrimpf, M.; Zhang, Y.; Bowman, S.; Zaslavsky, N.; Fedorenko, E. Artificial Neural Network Language Models Align Neurally and Behaviorally with Humans Even after a Developmentally Realistic Amount of Training. Unpublished ms. 2022. Available online: <https://www.biorxiv.org/content/10.1101/2022.10.04.510681v1> (accessed on 1 May 2023).
30. Warstadt, A.; Choshen, L.; Mueller, A.; Wilcox, E.; Williams, A.; Zhuang, C. Call for Papers—the BabyLM Challenge: Sample-Efficient Pretraining on a Developmentally Plausible Corpus. 2023. Available online: <https://babylm.github.io/> (accessed on 1 May 2023).
31. Ambridge, B. A computational simulation of children's language acquisition. In Proceedings of the Third Conference on Language, Data and Knowledge (LDK 2021), Zaragoza, Spain, 1–3 September 2021; pp. 4:1–4:3.
32. Gilkerson, J.; Richards, J.; Warren, S.; Montgomery, J.; Greenwood, C.; Oller, D.K.; Hansen, J.; Terrence, P. Mapping the early language environment using all-day recordings and automated analysis. *Am. J. Speech Lang. Pathol.* **2017**, *26*, 248–265. [[CrossRef](#)] [[PubMed](#)]
33. Miller, G. The cognitive revolution: A historical perspective. *Trends Cogn. Sci.* **2003**, *17*, 141–144. [[CrossRef](#)] [[PubMed](#)]
34. Miller, G. The Society Theory. In *Artificial Intelligence: An MIT Perspective*; Winston, P.H., Brown, R.H., Eds.; MIT Press: Cambridge, MA, USA, 1979; Volume 1, pp. 423–450.
35. Gardner, H. *The Mind's New Science: The History of the Cognitive Revolution*; Basic Books: New York, NY, USA, 1985.
36. Stampe, D. A dissertation on Natural Phonology. Ph.D. Thesis, Department of Linguistics, University of Chicago, Chicago, IL, USA, 1973.
37. Bates, E. Bioprograms and the innateness hypothesis. *Behav. Brain Sci.* **1984**, *7*, 188–190. [[CrossRef](#)]
38. Elman, J. The emergence of language: A conspiracy theory. In *The Emergence of Language*; MacWhinney, B., Ed.; Erlbaum: Mahwah, NJ, USA, 1999; pp. 1–27.
39. MacWhinney, B. Preface. In *The Emergence of Language*; MacWhinney, B., Ed.; Erlbaum: Mahwah, NJ, USA, 1999; pp. ix–xvii.

**Disclaimer/Publisher's Note:** The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.