Title

The Convergence of Linguistics and Computer Science

Journal

Japanese/Korean Linguistics, 32(1)

Authors

O'Grady, William
Zhang, Haopeng
Lee, Miseon

Publication Date

2026-07-31

DOI

10.5070/J7.67149

The Convergence of Linguistics and Computer Science*

WILLIAM O'GRADY
University of Hawaii at Manoa

HAOPENG ZHANG
University of Hawaii at Manoa

MISEON LEE
Hanyang University

1 Introduction

The relationship between linguistics and computer science arguably began in earnest at a symposium on information theory held at MIT in September of 1956, an event that the distinguished psychologist George Miller once referred to as the 'conception of cognitive science' (2003:142). In the opinion of Howard Gardner (1985:28), three key presentations were particularly impactful. One was Herbert Simon and Allen Newell's 'logic theory machine,' which demonstrated the first ever complete proof of a theory by a computing machine. The second was Miller's influential talk on 'the magical number 7,' which suggested that human short-term memory was limited to approximately seven items. The third was a lecture by Noam Chomsky, demonstrating that language could be analyzed with the formal precision of mathematics, a founding principle of generative grammar. Miller later made the following observation:

* The 32nd meeting of the Japanese/Korean Linguistics Conference was preceded by a satellite workshop on the growing importance of artificial intelligence for the study of language. The brief report presented here, co-authored by three contributors, is an extension of the remarks made at the conference by the first author. We are grateful for the comments of Miho Choo, Sungdai Cho, John Whitman and members of the audience.

Japanese/Korean Linguistics 32.
Edited by Youngdong Cho, Young-Hoon Kim, and John Whitman.

> 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 … (cited by Gardner 1985:29)

Miller depicted the field of cognitive science as follows, assigning it the same status as linguistics and the other fields of inquiry that have a stake in the study of cognition.

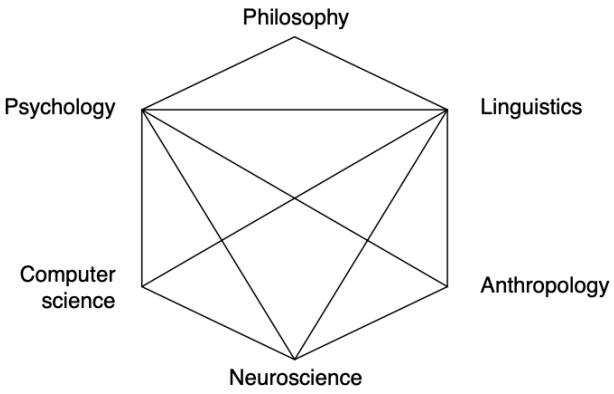

Figure 1: Miller's representation of cognitive science

Years later Chomsky (2022) recalled the emergence of cognitive science in much the same way.

> In the early days … AI was basically not distinct from cognitive science. It was just one of the ways of approaching the problems of cognitive science, using achievements … that had been made both intellectually and technologically ….

Over the course of several decades, computer science made great progress in a variety of domains, ranging from email to laptop computers. However, it was not until recently that a series of unexpected developments brought computer science and linguistics together in a way that may well change how we think about language, perhaps even allowing us to attain what Linzen (2020:5210) calls a classic goal – 'building a model that acquires language as efficiently as humans.'

Not surprisingly there are major differences of opinion with regard to the likelihood of such a development. On the one hand it has been suggested that this objective has already been achieved, thanks to the supposed rejection of the 'rules and principles that have dominated generative approaches to language' (Piantadosi 2023:9).

> … large language models have attained remarkable success at discovering grammar without using any of the methods that some in linguistics insisted were necessary ... (ibid.:1) …

> In fact, none of the principles and innate biases that Chomsky believes in (e.g. binding principles, binary branching, island constraints, Empty Category Principle, etc.) are necessary. (ibid.:15)

In contrast, Chomsky (2023) has painted a very different picture.

> [Language models] are designed in such a way that in principle they can tell us nothing about language, about learning, about intelligence, about thought. Nothing. There's a lot of sophisticated programming, but basically what it comes down to is sophisticated high-tech plagiarism.

These differences of opinion have initiated vigorous debate and, more importantly, launched various lines of empirical research.

> [It] is really exciting to use language models to study humans in controlled experiments… A case for using language models in linguistics and cognitive science is that they can help us answer [the question]: what do language learners look like in general? (Warstadt 2025)

> … compared to … human minds and human brains, language models are actually incredibly transparent… We can train a model and know with certainty every single word that went into the training. We can know exactly how the architecture is put together … (Mahowald 2025)

> By asking what can be acquired by data-driven learning algorithms without any obvious domain-specific biases [UG, for instance], we map out a lower-bound for learnability and clarify what (if anything) remains as a good candidate for top-down innate constraints. (Wilcox, Futrell and Levy 2024: 805-48)

In what follows, we will report on an experiment that we conducted on the syntax of negative concord (O'Grady, Zhang and Lee 2025), using it as an example of how a blend of AI and traditional linguistics can shed light on a curious phenomenon of considerable developmental and typological interest.

2 Negative Concord

Negative concord arises when simple negation is expressed by a construction consisting of two negators – a clausal/sentential negative and a negative pronoun. Korean and Japanese both manifest constructions of this type.

(1) Korean:
Jiseon-i amwukesto an mek-ess-ta.
Jiseon-NOM nothing not eat-PST-DECL
'Jiseon ate nothing.'

(2) Japanese:
Michiko-ga nanimo tabe-na-katta.
Michiko-NOM nothing eat-not-PST
'Michiko ate nothing.'

Interestingly, negative concord can also be observed in English-speaking children who do not encounter it in their environment, as first noted by Bellugi (1967) in her study of 'Adam'.[1]

[1] We set to the side here varieties of English that allow negative concord; see Robinson & Thoms (2021).

(3) a. I didn't do nothing. (file 63, age 3;5)

b. I didn't call him nothing. (file 72, age 3;8)

c. Because nobody didn't broke it. (file 107; age 4;5)

A more recent study by Hein et al. (2023) reveals that this phenomenon is in fact quite widespread. Drawing on a data set of 328,972 utterances produced by seven chlildren whose families did not employ negative concord, Hein and his colleagues identified 909 sentences that contained a negative pronoun or determiner (*nobody, nothing, no*, etc.). Of these, 178 (19.6%) manifested negative concord.

(4) a. I don't care about nothing. (Ross 5;04)

b. No one's not drying him, mum. (Fraser 3;00)

c. No tigers don't bit you? (Mark 2;08)

A study by Thornton et al. (2016) revealed that negative concord is also evident in how chidren interpret potentially ambiguous sentences such as the one below.

The girl (who skipped) did**n't** buy **nothing**.

Negative Concord Interpretation	Double Negative Interpretation
The two negators produce simple negation. = 'The girl bought nothing.'	The two negators cancel each other out. = 'The girl bought something.'

Figure 2: Two potential interpretations

Of the twenty children (aged 3 to 5) in Thornton et al.'s experiment, fifteen chose the negative concord interpretation. In contrast, only two of the fifteen adults in the study opted for such an interpretation.

This brings us to the question of how and why children come to accept negative concord in standard English. One view, common in work on Universal Grammar, proposes that negative concord occurs when a language has a negative marker that functions as a syntactic head (e.g. Zeijlstra 2004, 2008; Thornton and Tesan 2013:386).

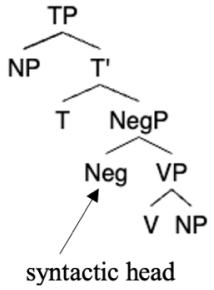

Figure 3: A negative marker functioning as syntactic head

On this view, the appearance of negative concord reflects the fact that English is unusual in having negative heads (e.g. *n't*) but not permittng negative concord – a state of affairs that creates temporary confusion among young language learners.

In the next section, we propose an alternative explanation for the sporadic emergence of negative concord in English.

3 A Cognitive Explanation

The particular idea that we will pursue focuses on biuniqueness, a cognitive force that places a high value on a one-to-one relationship between form and meaning (O'Grady in press). A simple example involves the formation of plurals in English, where the suffix *-s* systematically marks a plural noun and a plural noun is systematically marked by the suffix *-s*.

car	cars
dog	dogs
cup	cups
tree	trees
etc.	

As the following plural examples help illustrate, biuniqueness need not be absolute.

sheep, mice, feet, children

Crucially, however, such departures from the norm are known to invite over-regularization by young language learners (e.g. *sheeps*, *mouses*, *foots, childs*), consistent with a preference for biuniqueness. As we will see next, biuniqueness may also have a role to play in negation.

3.1 Strategies for Negation

There are two basic strategies for negating a proposition in English – reliance on the clausal negator *not* and recourse to negative pronouns such as *no one*, *nobody* and *nothing*. Thus all three of the sentences below can be used to express roughly the same meaning.

Not-type Negative Sentence	Pronoun-type Negative Sentence
The people did**n't** eat the food.	**Nobody** ate the food. The people ate **nothing**.

Table 1: Two strategies for expressing negation

There is, though, a very substantial difference in the relative frequency of the two strategies. Whereas *not* and its variants (*n't*, *didn't*, *don't* and *can't*) are among the most frequent words in speech to children (Quick et al. 2019), negative pronouns are quite rare. As a result, there is a very strong correlation between *not* and the expression of negation – a classic instantiation of biuniqueness.

not <==> negation

This fact opens the door to an interesting possibility. It is not unreasonable to think that just as *not*

encodes negation, so the presence of negation may trigger the production of *not* – a state of affairs that invites over-regularization in certain cases. In order to investigate this matter further we turned to a small language model.

3.2 The Language Model

A key feature of small language models is that they are trained on data comparable in size to the input that children receive in the first years of their lives. The 'strict-small' version of Warstadt et al.'s (2023) model (that we also used) was trained on a corpus consisting of approximately 10 million words (corresponding to 1,244,397 sentences). This is roughly the input that children receive over a period of two to five years, depending on the talkativeness of their caregivers (Gilkerson et al. 2017). The data set contained 126,611 negative sentences, the vast majority of which were negated by *not*. Only 1,129 sentences (about 1% of all negatives) had a negative pronoun as their negator. There were no instances of negative concord.[2]

Our design used a decoder-only architecture based on the GPT-2 model. It includes 12 Transformer decoder blocks with a hidden size of 768, resulting in approximately 124 million parameters. The model was pre-trained from scratch using the Adam optimizer, with Byte-Pair Encoding (BPE) for sub-word tokenization.

Our goal was to determine how a language model with the properties described above behaves when presented with negative concord patterns. Would the co-occurrrence of *not* and a negative pronoun be treated as acceptable or unacceptable?

Acceptability was assessed by measuring perplexity, a metric that quantifies how 'surprised' a model is by a sequence of words. Low perplexity indicates that a sequence of words has a high probability of occurring and is therefore likely to be acceptable. In contrast, high perplexity signals that a sequence of words is unexpected and hence more likely to be unnatural.

In calculating perplexity scores, it is common to work with minimal pairs that differ in just one variable, as in the examples below from Warstadt et al. (2020).

Low Perplexity	High Perplexity
Many people mistrust **themselves**.	Many people mistrust **himself**.
These casseroles **disgust** Kayla.	These casseroles **disgusts** Kayla.
Brett knew **what** many waiters find.	Brett knew **that** many waiters find.

Table 2: Perplexity contrasts

3.3 Testing the Biuniqueness Hypothesis

Our first set of test items involves negative pronouns in subject position, with and without negative concord.

[2] In order to encode the presence of negation, we added the symbol N* to the beginning of each sentence that contained a negator (*not*, *nobody*, *nothing*, etc.).

Negative Concord	Stand-alone Negative Pronouns
Nobody didn't remember her.	Nobody remembered her.
Nobody didn't laugh at the clown.	Nobody laughed at the clown.
Nobody didn't stay in the library.	Nobody stayed in the library.
Nobody didn't watch the elephant.	Nobody watched the elephant.
Nobody didn't jump over the fence.	Nobody jumped over the fence.
Nothing didn't please the teacher.	Nothing pleased the teacher.
Nobody didn't play in the park.	Nobody played in the park.
Nobody didn't cry during the movie.	Nobody cried during the movie.
Nobody didn't push the button.	Nobody pushed the button.
Nothing didn't look good.	Nothing looked good.

Table 3: Negative pronouns in subject position

As illustrated below, the test results are surprising in that the negative concord pattern has a much lower perplexity score than the pattern with a stand-alone negative pronoun. In other words, negative concord is more likely to be allowed by the language model – the exact opposite of what we find in standard English, but similar to what children sometimes do.[3] (The numbers in parentheses indicate 95% confidence intervals.)

Negative Concord	Stand-alone Negative Pronouns
128.90 (80.24, 193.57)	1850.57 (642.24, 3644.96)

Table 4: Mean perplexity scores for negative pronouns in subject position

A second set of test items involves negative pronouns in a complement position, with and without negative concord.

Negative Concord	Stand-alone Negative Pronouns
The nurse didn't spill nothing.	The nurse spilled nothing.
The teacher didn't change nothing.	The teacher changed nothing.
The elephant didn't hurt nobody.	The elephant hurt nobody.
The men didn't wash nothing.	The men washed nothing.
The doctor didn't talk to nobody.	The doctor talked to nobody.
The rich people didn't need nothing.	The rich people needed nothing.
The man didn't watch nobody.	The man watched nobody.
The tiger doesn't like nobody.	The tiger likes nobody.
The boy didn't drop nothing.	The boy dropped nothing.
The girl didn't look at nothing.	The girl looked at nothing.

Table 5: Negative pronouns in complement position

[3] The results reported below are taken from O'Grady, Zhang and Lee (2025).

Here too, the negative concord pattern has a much lower perplexity score than the pattern with a stand-alone negative pronoun – once again the exact opposite of what we find in standard English.[4]

Negative Concord	Stand-alone Negative Pronouns
216.64 (133.79, 316.56)	12608.50 (7657.53, 17340.39)

Table 6: Mean perplexity scores for negative pronouns in complement position.

3.4 What the Results Suggest

These results suggest that the clausal negative *not* is strongly associated with negation, as one would expect given the overwhelming percentage (99%) of negated sentences in which it occurs in our data set. To the extent that a comparable dominance arises in human language, the likelihood that children might on occasion overuse *not* in sentences containing a negative pronoun seems plausible. Indeed, this is exactly what one would expect if biuniqueness is a relevant factor.

Also in play here is an interesting similarity between language models and the human language faculty, which is that both are computational systems that rely heavily on frequency-based prediction. As Michaelove et al. (2024:109) note, language models are 'designed to predict the probability of a word … based on the statistics of language,' echoing the observation by Hawkins (2014:232) that 'frequency, predictability, expectedness and surprisal' are important factors in human sentence processing. In this regard at least, there is a substantial similarity between the two mechanisms.

3.5 Korean and English

A variety of issues call for further attention. One involves the fact that negative concord is an unusually common occurrence in the world's languages, attested in 170 of the 206 languages surveyed in the *World Atlas of Language Structures* (Haspelmath 2013). In Korean, for instance, every clause containing the negative pronoun *amwuto* 'no one' or *amwukesto* 'nothing' must be accompanied by a clausal negative (e.g., *an* V, V *ci-anh-*, or an inherently negative verb such as *eps-ta* 'not exist'). The widespread preference for a similar syntax of negation in so many other languages suggests that negative concord is the default option in human language.

Consistent with this thesis, an interesting contrast arises in the early stages of development in children's acquisition of Korean and English. Children learning Korean are consistent in their production of negative concord, even at a very young age. In a study of 100 children aged three to six, Song (2003:102) found just three omissions of a clausal negator in a total of 2000 elicited sentences containing a negative pronoun – an error rate of a mere .15%. This accurately reflects the input to which Korean children are exposed.

However, matters are quite different in children's acquisition of English, where the input from adults calls for stand-alone negative pronouns. Crucially, instead of faithfully complying with that input, children produce negative concord at an overall rate of approximately 20% according to

[4] We set to the side for now possible reasons for why the magnitude of the contrasts within the subject and complement patterns differ to such a high degree.

Hein et al. This suggests the presence of significant early pressure in favor of negative concord, presumably associated with the biuniqueness factor discussed in §3.3.

4 Concluding Remarks

A final matter requires consideration. At some point and in some way, children learning standard English must ultimately reject negative concord. A plausible scenario for how corrections of this sort take place involves the fact that listerners subconsciously predict upcoming words and patterns in the speech of others, as noted in §3.4. This practice has especially important consequences for language acquisition.

> 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 (Kaan and Grüter 2021: 4; see also Phillips and Ehrenhofer 2015, among many others).

In the case at hand, the erroneous expectation that a negative pronoun is to be accompanied by clausal negation should be noted and adjusted over time as children process sentences such as *Nobody cried* when they expect *Nobody didn't cry* or *I said nothing* when they anticipate *I didn't say nothing*.[5]

Could something similar happen in the case of language models, devices which are specifically trained to predict upcoming words? Experiments on this subject that we have conducted reveal that doubling the size of the corpus on which the language model is trained results in a marked lowering of the perplexity scores associated with stand-alone negative pronouns. Although this is an adjustment in the right direction, the model maintains a substantial preference for negative concord. This finding calls for further investigation now underway.

Regardless of the final outcome, it is evident that language models offer opportunities to simulate developmental scenarios that would not otherwise be accessible, creating the potential for new insights into the nature of language and learning. This is surely a line of research that is worth pursuing.

References

Bellugi, U. 1967. The Acquisition of the System of Negation in Children's Speech. Doctoral dissertation, Harvard University.

Chomsky, N. 2022. Debunking the Great AI Lie. (video interview: 29:46–30:03); available at https://www.youtube.com/watch?v=PBdZi_JtV4c.

Chomsky, N. 2023. Noam Chomsky on Artificial Intelligence (video interview:1:14–1:35); available at https://www.youtube.com/watch?v=_04Eus6sjV4.

Gardner, H. 1985. *The Mind's New Science: A History of the Cognitive Revolution*. New York: Basic Books.

[5] Variation is to be expected in individual learners, depending on their frequency of exposure to sentences containing a negative pronoun.

Gilkerson, J., J. Richards, S. Warren, J. Montgomery, C. Greenwood, D. K. Oller, J. Hansen, & T. Paul. 2017. Mapping the Early Language Environment Using All-Day Recordings and Automated Analysis. *American Journal of Speech-Language Pathology 26*, 248–265.

Haspelmath, M. 2013. Negative Indefinite Pronouns and Predicate Negation. *The World Atlas of Language Structures Online,* eds. M. Dryer & M. Haspelmath, Chapter 115. Available online: https://wals.info/chapter/115 (accessed on 18 March 2025).

Hein, J., C. Bill, I. Driemel, A. Gonzalez, I. Ilić, & P. Jeretič. 2023. Negative Concord in the Acquisition of English and German: Some Results from a Corpus Study. *Proceedings of the Chicago Linguistic Society* 58, eds. L. Fagen, S. Gray, Quain, S. Reyes, & I. Tang. Available at: https://www.johannes-hein.de/documents/CLS58Proceedingspaper.pdf.

Hawkins, J. 2004. *Efficiency and Complexity in Grammars*. Oxford, UK: Oxford University Press.

Hutchins, W. J. 2012. Victor H. Yngve. *Computational Linguistics* 38, 461-67.

Kaan, E. & T. Grüter. 2021. Advances and Directions. *Prediction in Second Language Processing and Learning*, 1-24. Amsterdam: John BenjaminsPublishing Company.

Linzen, T. 2020. How Can We Accelerate Progress Towards Human-like Linguistic Generalization? *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 5210–5217. https://doi.org/10.18653/v1/2020.acl-main.465

Mahowald, K. 2025. Psycholinguistic Insights from Small Language Models. HSP Online Seminar (9:10–9:38): Available at https://stanford.zoom.us/rec/share/T3nJqGOjPyJ7m6t2mmt-nIUOrnF9_pet3jHGtGU6qSyCBcOiifd-IE1OrSb4B7JhD.H6taAZJzQrzOKq_K

Michaelove, J., M. Bardolph, C. Van Petten, B. Bergen, & S. Coulson. 2024. Strong Prediction: Language Model Surprisal Explains Multiple N400 Effects. *Neurobiology of Language* 5, 107-35.

Miller, G. 2003. The Cognitive Revolution: A Historical Perspective. *Trends in Cognitive Science* 7, 141-44.

O'Grady, W. In press. Natural Syntax: A Case Study of Negation. *Cambridge Handbook of Natural Linguistics*, eds. K. Dziubalska-Kołaczyk & W. Dressler. Cambridge, UK: Cambridge University Press.

O'Grady, W., H. Zhang, & M. Lee. 2025. The Negative Concord Mystery: Insights from a Language Model. *Information* 16, 710-20.

Phillips, C. & L. Ehrenhofer. 2015. The role of language processing in language acquisition. *Linguistic Approaches to Bilingualism* 5, 409-53.

Piantadosi, S. 2023. Modern Language Models Refute Chomsky's Approach to Language, *From Fieldwork to Linguistic Theory: A Tribute to Dan Everett*, eds. E. Gibson & M. Poliak, 353-413. Berlin: Language Sciences Press.

Quick, N., K. Erickson, & J. Mccright. 2019. The Most Frequently Used Words: Comparing Child-directed Speech and Young Children's Speech to Inform Vocabulary Selection for Aided Input. *Augmentative and Alternative Communication* 35, 120-31.

Robinson, M. & G. Thoms. 2021. On the Syntax of English Variable Negative Concord. *University of Pennsylvania Working Papers in Linguistics* 27, Article 24.

Song, M. 2003. The First and Second Language Acquisition of Negative Polarity Items in English and Korean. Doctoral dissertation, University of Hawaii at Manoa.

Thornton, R., A. Notley, V. Moscati & S. Crain. 2016. Two Negations for the Price of One. *Glossa* 1, article 45.

Thorton, R. & G. Tesan. 2013. Sentential Negation in Early Child Language. *Journal of Linguistics* 49, 367-411.

Warstadt, A. 2025. Why It Matters That Babies and Language Models are the Only Known Language Learners. Talk presented at the Workshop on LLMs, Cognitive Science, Linguistics and Neuroscience. Simons Institute (3:05–4:20). Available at https://www.youtube.com/watch?v=Oyf-sS2oEvo

Warstadt, A., A. Parrish, H. Liu, A. Mohananey, W. Peng, S.-F. Wang, & S. Bowman. *BLiMP: The Benchmark of Linguistic Minimal Pairs for English. Transactions of the Association for Computational Linguistics* 8, 377–392.

Warstadt, A., I. Choshen, A. Mueller, E. Wilcox, A. Williams, & C. Zhuang. 2023. The BabyLM Challenge: Sample-efficient Pretraining on a Developmentally Plausible Corpus. *arXiv*:2301.11796.

Wilcox, E., R. Futrell, R. Levy. 2024. Using Computational Models to Test Syntactic Learnability. *Linguistic Inquiry* 55, 805-48.

Zeijlstra, H. 2004. *Sentential Negation and Negative Concord*. Utrecht: LOT.

Zeijlstra, H. 2008. Negative Concord Is Syntactic Agreement. Ms., University of Amsterdam. Lingbuzz/000645.