



Commentary

Studying second language acquisition in the age of large language models: Unlocking the mysteries of language and learning, A commentary on “Age effects in second language acquisition: Expanding the emergentist account” by Catherine L. Caldwell-Harris and Brian MacWhinney

This is an exciting time for those who study language and language learning. Large Language Models, or LLMs, such as Open AI ChatGPT, Microsoft Bing, Google Bard, and a score of other recently-released artificial intelligence tools are upending much of what we thought we knew about language learning and language function. What makes LLMs especially interesting is that the artificial intelligence that is changing the world at an unprecedented rate is essentially a language learning model that follows rules that linguists, psychologists, computer scientists, and other cognitive scientists have been studying for decades. By relying on statistical probabilities to make predictions, LLMs are rapidly advancing from text autocomplete to providing logical answers to complex questions and generating new content.

Historically, the main theories of language advocated for an innate language learning mechanism, for universal grammar, for a critical period, and for a human factor that no other species possessed, one anchored in communicative intent, consciousness, and higher-order pragmatic and cognitive reasoning. These may still be defining characteristics, but they may also be challenged to the core as LLMs evolve. Already now, LLMs can do much more than anyone expected they would at this stage, and they are improving at a rapid pace. If statistical learning and prediction can accomplish so much so fast, how much more will LLMs be able to do as they evolve and build on the body of knowledge acquired over the course of humanity's existence? Is probabilistic learning enough to learn language and become a competent language user? If no, why not, what else is needed? At what point in the evolution of LLMs may they acquire pragmatic skills and develop communicative intent, if at all? And from there on, where do the boundaries between consciousness and non-consciousness lie? These are just some of the questions philosophers of language grapple with, and consensus is elusive. For that matter, we don't even agree on what language is and is not.

Within this broader zeitgeist, the study of second language acquisition (SLA) provides a special lens for investigating language and learning. SLA researchers can manipulate variables that cannot be manipulated during first language learning. This makes it possible to reveal factors that drive successful and unsuccessful language learning, the interactions among these factors, and their effects on the cognitive and neural systems. In other words, studying SLA makes it possible to uncover patterns and principles of language and learning more broadly, with implications for machine learning and artificial intelligence.

Language and learning are tightly connected. With each new language learned, it becomes easier to acquire new languages, and the more languages a cognitive architecture has at its disposal, the better it

becomes at further learning, in a virtuous cycle (Marian, 2023). Learning more than one language has notable consequences for cognition (Marian et al., 2009). Language is connected not only to learning, but to thought broadly defined. Some have even proposed that language and thought are two sides of the same coin (see the Sapir-Whorf hypothesis, Sapir, 1921; Whorf, 1956). Although most scientists who study language and thought do not believe that the two coincide, there is little arguing that they mutually influence each other, and that studying language gives us a unique window into the mind.

It is against this historic backdrop that Catherine Caldwell-Harris and Brian MacWhinney wrote their excellent keynote on age effects in SLA. I found it to be a very interesting and thorough analysis. The authors not only provide an up-to-date review of the literature on and adjacent to SLA, but also propose a framework that moves beyond maturational factors in thinking about what age effects in SLA mean and what they tell us about language and learning. As is characteristic of emergentist accounts of language, their discussion recognizes that language learning is a dynamic process that consists of the gradual emergence of linguistic knowledge through the interaction of multiple factors. Their expanded framework is grounded in empirical research and recognizes that there are multiple paths to language proficiency. I particularly appreciated the breadth of the approach, the attempt to structure existing knowledge into a cohesive three-legged theory, the many examples, and the parallels to phenomena outside language. The authors even draw on proteomics and discuss the four levels of protein folding to illustrate how emergentism operates on a different scale.

The proposed framework incorporates three groups of factors – environmental supports, cognitive abilities, and motivational forces. The authors propose that successful language acquisition is shaped by both protective and risk factors that vary across the developmental lifespan (infancy and early toddlerhood, early childhood, middle childhood, adolescence and young adulthood, adulthood, older age) and across linguistic processes (speech perception and audition, articulation, lexicon, morphology, syntax, and conversation). The broad scope of the discussion, which also incorporates language competition within and across languages (Fernandez-Duque et al., 2023), as well as cognitive and motivational variables, is appealing in its non-reductiveness. The result is a more accurate understanding of age effects in SLA, one that challenges the overly simplistic maturational and biological accounts that have been traditionally considered responsible for these effects.

This means that SLA outcomes are not due to an open or closed sensitive period, but rather a result of an interconnected set of variables that work in concert, mutually influencing and building upon each

other. Recognizing this underscores the value of the field of SLA in explicating the mysteries of language and learning, with implications for both human and artificial intelligence (with some modifications, for example to the motivational component, at least for now and depending on how it is defined).

The authors make convincing arguments against the critical period hypothesis, the sensitive period hypothesis, and maturational accounts more generally. The critical and sensitive periods hypotheses have had a good run, but it is time to retire them in favor of a broader framework that considers interactions among multiple factors over the course of an organism's lifespan (and possibly beyond, as Caldwell-Harris and MacWhinney point to the potential role of epigenetic shaping of aptitude). Moving beyond the "closing of a sensitive period" is necessary not only when considering SLA, but also native language learning, and language and learning more broadly. You can teach old dogs new tricks, as the authors themselves allude to when they discuss that a return to plasticity IS possible in other animal models (cats, rats, etc.). Evidence confirms the ability to learn a second language to fluency at any age, subject to the interactive effects of environmental, cognitive, and motivational factors. It is increasingly clear that language learning is a life-long process and not a finite task that takes place within a constrained time window.

Multiple factors are discussed in support of an emergentist account and against maturational mechanisms. Even aspects that have traditionally been seen as supporting a maturational account could be challenged. For example, I am not convinced that the attrition of L1 in young adoptees is similar to the phenomenon of "catastrophic interference" demonstrated in neural nets when the training set switches to a new set of input-output pairs, nor that it is an example of a "complete loss." Research employing the savings paradigm suggests that young adoptees who appear to have completely lost a language (and may not even know what language they once knew) learn the language they were once exposed to during infancy with greater ease than a language they were not exposed to, or than other people who were not exposed to that language in early development (Isurin, 2015; Piece et al., 2014). Regardless of whether "catastrophic interference" yields complete loss of previous input-output pairs in machine learning, the data suggest that in human learning, some knowledge is preserved even after decades and even without use, further challenging theories that L1 attrition and L2 success are tied to the same maturational mechanisms.

The downside to the flexible, inclusive, and multidisciplinary approach of emergentist theories of language is their under-specification and lack of clarity. There are many unknowns. We have yet to delineate how exactly the multitude of variables build upon each other in this integrated emergentist account to result in successful SLA (or language learning in general). To their credit, rather than hiding behind the complexity in a dismissive hand-wavy way, the authors make an honest attempt to enumerate and organize the many contributing variables into meaningful groupings. And although it is not yet possible to reach a level of precision that can be operationalized mathematically, collating the relevant variables and structuring them is a step in the right direction and one that is appropriate for the current state of affairs in the field. It is hard to pull together the vast SLA literature spanning decades and containing many inconsistent findings. Behavioral sciences, like the natural sciences, do not always yield clear and easily interpretable results. As much as researchers labor towards external validity and reproducibility, each individual study gives us only a snapshot of a specific population sample at a specific point in time and under a specific set of circumstances. Experiments on language learning have often tackled individual factors like motivation, inhibitory control, language experience, and cognitive abilities, as well as input factors like lexical form, language similarity, language co-activation, and visual-to-auditory or form-to-meaning mappings. The results of these studies inform us about second language learning in native speakers of English and of a handful of other languages, but they tell us less about SLA in populations that speak languages others than those spoken by WEIRD (Western, Educated, Industrialized, Rich, Democratic) nations. This is

unfortunately true about much of the SLA literature, which is precisely what made universal accounts like nativism and universal grammar so appealing. In our eagerness to uphold universal principles of language and learning, we may have rushed to adopt overly simplistic explanations. The aphorism "the simplest solution is almost always the best" is often but not always true and, when it comes to language, there is still much to uncover.

Which brings us full circle to LLMs and what they reveal about language and learning. LLMs provide a testing ground for some of the most important theories of language. LLMs learn fast and the outcomes can be observed nearly instantaneously. That is not necessarily the case for human learning, and even less so for *research* on human learning and development. Empirical research with humans can take years or even decades, especially when conducted longitudinally. Some of the cross-linguistic and cross-cultural developmental research on parent-child communication currently conducted in our laboratory (Rochana-vibhata & Marian, 2022, 2023) is rooted in conversations with Urie Bronfenbrenner about his ecological systems theory that took place in the nineteen-nineties. We can expect hypothesis testing to take place much faster in LLMs than in naturalistic studies of bilingual parent-child dyads in under-studied communities. LLMs can explicate the language learning process through continued learning and their own emergentism.

To be sure, although LLMs can test theories much faster and can help us learn about SLA, they also differ from humans on many of the environmental, cognitive, and socio-motivational factors discussed by Caldwell-Harris and MacWhinney. LLMs may very well evolve in ways that *stray* from human thought, not just due to speed, but also size, physical differences, and lack of embodiment (for now, though development of carbon versions is in progress). Machine learning operates under different constraints than human learning (which is limited by the speed at which neurons fire in the brain, among other things), and while we can learn from these differences, caution is advised when considering similarities and differences between machine and human language learning and making extrapolations. For example, LLMs are still subject to the WEIRD bias due to the differences between the databases available for training across high-vs low-resource languages, whereas the "slow" cross-cultural human research has the merit of being able to capture sociocultural and ecological influences on language acquisition that may be lost in machine learning.

Modern LLMs reinforce emergentist theories about the dynamic nature of language and challenge nativist assumptions of universal grammar and critical periods. Whether the probabilistic nature of learning from large-scale corpora that is at the basis of LLMs will evolve to exhibit emergent abilities that had not been directly trained or expected was, until not long ago, an open question. Most recently, evidence points to the affirmative, with several documented instances of LLMs exhibiting emergent abilities (Ornes, 2023) and confirmed examples of LLMs spontaneously and unexpectedly performing arithmetic operations, generating executable computer code, summarizing passages, and deducing answers to questions. While some call the emergent abilities of LLMs a mirage (Schaeffer et al., 2023), the breakneck speed at which LLMs advance gives us reasons to believe that emergentist accounts of language learning are on the right track and that analyses of age effects in SLA, like the ones presented by Caldwell-Harris and MacWhinney, are a prescient preview of things to come.

Acknowledgments

The preparation of this manuscript was supported in part by the Eunice Kennedy Shriver National Institute of Child Health and Human Development of the National Institutes of Health under Award Number R01HD059858. The content is solely the responsibility of the author and does not necessarily represent the official views of the National Institutes of Health. The author thanks Sirada Rochanavibhata, Ashley Chung-Fat-Yim, Matias Fernandez-Duque, Siqi Ning, and Aya Williams

for comments on an earlier draft. Correspondence about this article should be addressed to Professor Viorica Marian, Department of Communication Sciences and Disorders, Northwestern University, Evanston, Illinois 60208, USA. Email: v-marian@northwestern.edu.

References

- Fernandez-Duque, M., Hayakawa, S., & Marian, V. (2023). Speakers of different languages remember visual scenes differently. *Science Advances*, 9(33), Article eadh0064. <https://doi.org/10.1126/sciadv.adh0064>. Adv.9,eadh0064(2023).
- Isurin, L., & Seidel, C. (2015). Traces of memory for a lost childhood language: The savings paradigm expanded. *Language Learning*, 65(4), 761–790. <https://doi.org/10.1111/lang.12133>
- Marian, V. (2023). *The power of language: How the codes we use to think, speak, and live transform our minds*. Dutton, 9780593187074.
- Marian, V., Faroqi-Shah, Y., Kaushanskaya, M., Blumenfeld, H. K., & Sheng, L. (2009). Bilingualism: Consequences for language, cognition, development, and the brain. *American Speech Language Hearing Association Leader*, 14(13), 10–13. <https://doi.org/10.1044/leader.FTR2.14132009.10>
- Ornes, S. (2023). The unpredictable abilities emerging from large AI models. *Quanta Magazine*.
- Pierce, L. J., Klein, D., Chen, J. K., Delcenserie, A., & Genesee, F. (2014). Mapping the unconscious maintenance of a lost first language. *Proceedings of the National Academy of Sciences of the United States of America*, 111(48), 17314–17319. <https://doi.org/10.1073/pnas.1409411111>
- Rochanavibhata, S., & Marian, V. (2022). Culture at play: A cross-cultural comparison of mother-child communication during toy play. *Language Learning and Development*, 18(3), 294–309. <https://doi.org/10.1080/15475441.2021.1954929>
- Rochanavibhata, S., & Marian, V. (2023). Thai and American mothers socialize preschoolers' emotional development differently. *Scientific Reports*, 13, Article 12719. <https://doi.org/10.1038/s41598-023-39947-0>
- Sapir, E. (1921). *Language. Harcourt, Brace, and World*.
- Schaeffer, R., Miranda, B., & Koyejo, S. (2023). *Are emergent abilities of large language models a mirage?*. <https://doi.org/10.48550/arXiv.2304.15004>. arXiv preprint arXiv:2304.15004.
- Whorf, B. (1956). In J. B. Carroll (Ed.), *Language, thought, and reality: Selected writings of Benjamin lee Whorf*. MIT Press.

Viorica Marian
Northwestern University, United States
E-mail address: v-marian@northwestern.edu.