Computational Modeling of Language Learning in the Era of Generative AI: A Response to Open Peer Commentaries

Ping Li and Qihui Xu

1 The Hong Kong Polytechnic University

2 BCBL: Basque Centre on Cognition, Brain and Language

Keywords: computational modeling, bilingualism, generative AI, LLMs, cross-disciplinary integration

In our article we reviewed computational models that have been used to gain insights into the learning and representation of multiple languages, and pointed out a number of new directions along which we could move the field forward. We appreciate the open commentaries from colleagues who suggested how we can build better models and integrate old and new models in this highly interdisciplinary enterprise. Here we provide some further thoughts and responses in addressing several issues raised in the commentaries. We position our response within the large context of what computational models can do in general, especially against the backdrop of today's new generative AI models.

The Scope of Bilingualism and Bilingual Research

In the last few years, researchers have realized that bilingualism is not a unitary concept but a phenomenon on a continuum (DeLuca et al., 2019). As also noted by Marian in her commentary, bilingualism is not an isolated island but rather a captivating component within a vast and interconnected landscape of other cognitive functions. How can we offer a theoretical account of the complex bilingual learning and representation across individuals who learn L1 and L2 in different contexts, for different purposes, and with different people (Grosjean et al., 2013; Li & Jeong, 2020)? This question entails that we need to understand a number of key questions such as how can bilinguals, especially late bilinguals, integrate new knowledge without disrupting or interfering with the old? What mechanisms allow for rapid learning for early but perhaps not for late L2 learners? And what role does statistical learning play in the dynamic language acquisition of two languages? On the socio-cultural front, how and why does active learning and immersion in an L2 environment facilitate easier acquisition and mitigate interference from the L1? How do the dynamic interactions between the individual, the language, and the environment shape the unique linguistic profiles of bilingual speakers?

At a first glance, it seems to be a daunting task to answer these questions given the current research progress on bilingual and bilingual modeling. However, if we look beyond and extend our scope to other research fields such as L1 acquisition research and Natural Language Processing (NLP), we may have already made some progress in addressing them. Tsoukala et al. (2017) demonstrated how recent computational cognitive models of sentence processing, originally designed for monolingual contexts, can be employed to study bilingualism. This approach allows for a continuum-based understanding of both monolingualism and bilingualism. Additionally, the application of Reinforcement Learning with Human Feedback (RLHF) has played a key role in allowing large language models (LLMs) to perform generative language tasks, aligning them more closely with human-like information values and sentence productions (OpenAI, 2023; Ouyang et al., 2022). This highlights the potential role of active interaction and motivation in shaping language learning. At the very least, understanding what learning mechanisms are important for L1 may inform what is missing in the learning of L2 on top of L1: if L2 learning can be made more like L1 learning (e.g., through social L2 learning with technology; see Jeong & Li, 2023), one might be surprised with what positive neurocognitive effects it may have on L2 representation and processing.

As highlighted in the commentaries (Dijkstra and van Heuven, 2022; Kachergis et al., 2022; Marian, 2022), the development of multiple diverse models can be challenging and resource-intensive. When we advocate for plurality, we do not suggest a fishing expedition exercise. Instead, we emphasize that the theoretical framework for understanding bilingualism should be situated within a broader theoretical framework of learning and cognitive science. As noted by Marian (2022), bilingualism is not an isolated island but rather a captivating component within a vast and interconnected landscape of other cognitive functions. Instead, the models that we develop should allow us to more flexibly and accurately examine how specific parameters or variables may affect bilingual language learning and representation. For example, Penaloza and colleagues highlighted the role of the language learning context and how it might impact the maintenance and decay of linguistic competency in younger and older populations and individuals with language disorders. We cannot agree more with this view. Indeed, in line with the arguments of Peñaloza and colleagues, Claussenius Kalman et al. (2021) suggested that researchers should examine bilingual expertise, ecosystem, and emergentism to further the study of dynamic developmental bilingualism. At the core of the three E's is the bilingual ecosystem, in which we should study the diversity and habits of language use, nature and amount of input, and methods of learning, so that we can better understand learning success as a result of the learner within the ecosystem. With regard to the expertise, the individual abilities and characteristics of the learner (e.g., age, aptitude, working memory, cognitive control, auditory perceptual ability) are all brought to bear in the learning task, and will interact with the context of learning to determine individual performance and trajectories. Computational models, due to their unique features (see discussion in Li and Xu's target article), are ideally suited for bringing these parameters and variables under investigation in a systematic and scalable manner.

Given the field's increased focus on second-person neuroscience (Redcay & Schilbach, 2019) to study social learning and human communication through naturalistic neuroimaging, we believe that the study of the bilingual learning context will hold significant promise for the future. The important question is how we might leverage computational models to study context-based learning and multimodal information processing in the ever-changing naturalistic learning environment. Rapid developments in AI and technology could help in this regard. For example, using immersive virtual reality-based methods, Legault et al. (2019) showed that simulated immersive learning environments can capture natural learning contexts for late adult learners and promote effective learning performance. In the era of generative AI, we might consider how to implement systems that present naturalistic learning platforms in which LLM-based virtual agents can interact with student learners, such that the virtual environments can simulate realistic and naturalistic environments of learning while at the same time allowing researchers to study complex interactions of learner-environment with a high degree of experimental control (see Li & Lan, 2022; Peeters, 2019).

Interpretability vs. Predictive power

While we all agree that "cross-disciplinary work is not a luxury but a necessity for success," researchers differ in their views as to how best we can make use of the power of computational modeling for understanding bilingual learning and representation. Dijkstra and van Heuven (2022) suggest that it is important to begin with specific core mechanisms of bilingual learning in order to attain more interpretable insights into the overall picture. On the other hand, Marian (2022) and Kachergis et al. (2022) advocate for an integrated perspective with greater predictive power, viewing the phenomenon of bilingualism as a continuous process closely intertwined with integrated cognitive functioning. In this regard, we discussed in our article predictive power and interpretability as two key features of good computational models. We believe that it is important to elucidate the roles of interpretability and predictive power in modeling research. When examining current models, a tradeoff becomes apparent: models with strong predictability often exhibit limited predictive power, particularly in relation to broader domains. Conversely, models with robust predictive power tend to be computationally complex and intricate, making it challenging to delve into the underlying states. One possible approach is to continue refining existing models until the theories of computational modeling become clearer, allowing for a meaningful progression towards more cross-disciplinary methodologies while still prioritizing interpretability. Alternatively, a more ambitious solution would involve extending our modeling efforts beyond language learning to encompass other interconnected cognitive components. By doing so, we can progressively cultivate a more comprehensive understanding characterized by a range of interpretable components and an overall heightened predictive power. Kachergis et al. explicitly argue that the time is ripe for achieving broader and more powerful predictions capable of capturing multiple language learning phenomena through a unified model.

While we concur with Dijkstra and van Heuven's viewpoint that a sound model should be interpretable to have a meaningful implication for bilingual processing, deep neural networks (DNNs) have demonstrated their value and potential contributions, even with their inherent black box nature. Just like conventional models, many deep learning models, especially those that are open-sourced, offer higher accessibility and manipulability compared to human subjects, making them valuable tools for investigating human cognition. Additionally, the "black box" problem in deep neural networks (DNNs) may be relatively less challenging than a similar problem in human subjects. Computational models provide direct access to internal representations, which is more cost-effective and granular than studying neural activities in the human brain. By precisely examining the activation of each layer of units after processing inputs, interpreting a DNN's internal states through ensembles of units and linking them to the model's behavior becomes more feasible than understanding neural activities in the human brain (see a few examples in Goldstein et al., 2022; see also Kriegeskorte & Douglas, 2018 for review).

Although there are cases where certain DNNs are not open-sourced or too large to be manipulated without expensive computational resources, such as ChatGPT, they remain valuable due to their exceptionally high predictive power (Binz & Schulz, 2023; OpenAI, 2023). This predictive power can stimulate a rethinking of various issues in human cognition, even by merely understanding their training methods (e.g., predicting the next word in a sequence) and overall architecture (e.g., transformer). Recent examples include a conference debate surrounding the use of ChatGPT's progress to reexamine the longstanding debate about embodiment: "Do large language models require sensory grounding for meaning and understanding?" (2023), along with empirical studies discussing the similar issue (Marjieh et al., 2023; Xu et al., 2023).

Bilingualism, being a crucial aspect of human cognition, can also benefit from this high predictive power. For instance, the multilingual pre-trained language model (MPLM) is renowned for its ability to comprehend multiple languages and its impressive zero-shot cross-lingual capabilities (Wu & Dredze, 2019; Artetxe et al., 2020). This zero-shot cross-lingual ability presents exciting opportunities to explore bilingual processing and understanding, such as in terms of language transfer and interaction between two languages. Unlike traditional approaches that required strictly aligned bilingual text for conducting bilingual tasks, the zero-shot learning capability of MPLM models enables knowledge transfer between languages without the need for explicit bilingual training. This advancement allows researchers to investigate how information and patterns acquired in one system can be effectively utilized and applied to another system.

While a trade-off between interpretability and predictive power may seem evident, it does not necessarily imply that choosing one necessitates rejecting the value of the other. Instead, we believe that a comprehensive understanding of bilingual learning and processing requires both the big unified pictures and the small nuanced pieces. We prioritize predictive power due to a perceived lack of emphasis on that aspect in bilingual modeling, possibly stemming from concerns surrounding the black box issue. However, it is important to note that the black box problem is not insurmountable. Researchers in the past have already used induced 'lesion' methods to probe into the black box with great success (e.g., Plaut et al., 1996) and it remains to be seen how this method can be scaled up to new heights with models such as LLMs (e.g., Lepori et al., 2023).

Plurality vs. Unification

Plurality and unification can be compatible if we interpret plurality as a methodological approach and unification as a final goal. This principle is exemplified in the field of NLP. Before the advent of LLMs, various models were distinct, each designed for a specific task such as sentiment analysis, classification, or machine translation. Now, a single LLM can handle multiple tasks, representing unification. However, without the pluralistic efforts and insights provided by earlier, task-specific models, the rapid development of such a unified model wouldn't have been possible.

However, one question that arises is whether the time is ripe for unification (Kachergis et al., 2022). Is unification the definitive answer, or could it serve as a catalyst for the emergence of a new generation of plural accounts, depending on how we define plurality? Considering the trajectory of NLP models, although a single model now appears to replace others for a vast array of language processing tasks, it remains uncertain if this is the optimal time to embrace such unification. There is indeed research indicating that Recurrent Neural Networks trained on large amounts of data can achieve comparable performance as ChatGPT (Peng et al., 2023). We may never truly know if it is the right moment for unification until a more unified model surpasses others in an unequivocal manner.

Nevertheless, we agree with Kachergis et al. (2022) that there is a need for a more comprehensive understanding of bilingual learning; in other words, enhancing the predictive power of models to account for a wider range of phenomena. Kachergis and colleagues (2022) demonstrate the use of Item Response Theory to capture both individual variation and common patterns in L1 acquisition, an essential aspect that we believe bilingual models should also incorporate, given the diversity among bilingual populations in terms of languages, language distance, and exposure to different cultures.

Scaling up the predictive power of models necessitates the establishment of large-scale data resources. As highlighted by Kachergis et al. (2022), the field of L1 acquisition research benefits from multiple open-resource large-scale databases, encompassing naturalistic conversations, vocabulary usage indexes, eye-tracking data, and more (such as the TalkBank and WordBank databases; Frank et al., 2017; MacWhinney, 2000). These resources facilitate the development and fitting of more integrated and powerful models. However, in the domain of bilingual research, such large-scale databases are comparatively limited, and model evaluations often rely on small datasets in experimental settings. Hence, we agree with Kachergis et al. (2022) and would like to highlight the necessity of large datasets for the development of robust bilingual models.

Large Language Models and Bilingual Language Learning

While we have touched upon the potential impacts of large language models (LLMs) on bilingual research in previous sections, it is important to further consider the significant contributions LLMs have made to language learning theories in general. When we wrote our article, we did not yet anticipate the popularity and impacts of LLMs, especially ChatGPT. The rise of LLMs has caught many by surprise and impacted many others in their work and life. For instance, an extensive discourse has emerged suggesting that the success of Large Language Models (LLMs) may have offered an answer to the ongoing debate on language learnability, refuting the Universal Grammar approach of Chomsky (1981) and the Language Instinct perspective (Pinker, 1994). This debate has been intense, up to today, between those who argue that language acquisition is impossible without a genetically programmed innate capacity, and those who believe that the statistical features inherent in the input of the learning environment are sufficiently rich to support learning (Chomsky et al., 2023; Contreras Kallens et al., 2023; Piantadosi, 2023; Yang, 2004). Such debates show how LLMs significantly influence our understanding of language and cognition.

But from our perspective, how can LLMs serve as testable computational models of the human brain in acquiring and representing multiple languages? Further, can the understanding of bilingual language learning also inform the design of more effective multilingual LLMs (although this is not our major concern for now)?

Computational modeling is particularly suited for modeling the complex scenario in bilingual language learning. One specific advantage is that it allows us to identify individual differences in a controlled manner (e.g., given the same input observing different output as a function of different L2 onset time). As Dijkstra and van Heuven point out, bilingual representations can differ significantly across languages and individual learners depending on, for example, task requirements and language properties/similarities. LLMs are built on aggregated huge amounts of text data, and therefore have no individualized behaviors or representations. Nor do LLMs currently consider learner-specific cognitive characteristics or abilities (see earlier discussion on the ecosystem). Another perspective that we advocate, as do Dijkstra and van Heuven (see both the BIA-d and DevLex-II models as discussed), is that computational models of bilingual learning should be developmental, reflecting the learning trajectories of each individual in representing and processing the new language across developmental stages. Again, LLMs currently lack such developmental profiles due to their approach in building sensible representations (e.g., word-and sentence-embeddings of large-scale text data). Both problems of LLMs also relate to what Dijkstra and van Heuven refer to as the black box problem, as discussed earlier.

We believe that although LLMs are currently lacking in both their abilities to simulate individual differences and in providing developmental profiles, these are not insurmountable problems for LLMs. For example, to solve the first problem, we need to train LLMs on specific datasets that vary by languages and by task demands. Individual cognitive abilities can be manipulated by embedding models using different window sizes, which may also solve the developmental issue in providing developmentally more complex input over time (Elman, 1993, “starting small”). To address the second issue, initiatives have been undertaken to develop open-source, smaller-scale Large Language Models (LLMs) such as Alpaca (Taori et al., 2023). The goal is to make these models more developmentally feasible and easily accessible for researchers. This would facilitate their replication or modification, thereby benefiting the wider research community. It also has the potential to allow researchers to deal with the black box problem of LLMs to some extent.

Unlike in human language learning, LLMs do not have a body or embodied actions (as yet), as they are relying for the most part on large-scale text data. This is quite different from human language learning in both L1 and L2. In L1, it has long been established that children do not learn from linguistic/visual materials only (e.g., by watching DVDs) but require an interactive context in which children and adults are socially engaged in communication (Kuhl et al., 2003; Tomasello, 2004; see also our earlier discussion learning context and ecosystem). In bilingual language learning, it seems that adults do not necessarily require embodiment and can acquire an L2 vocabulary quickly through associations or translations with items in their L1. However, research has provided evidence that bilingual learning through embodied experiences activates a broad brain network in the cortical and subcortical regions, in both the left and right hemispheres, thereby potentially enhancing the quality of representation and long-term memory retrieval (Jeong et al., 2021; Legault et al., 2019b; Li & Jeong, 2020). Furthermore, immersion in an L2 environment reduces interference from the first language (Linck et al., 2009). Alongside the discussion on conceptual representations derived from grounded experiences (Mollo & Millière, 2023; Xu et al., 2023), another potential avenue of exploration with contemporary deep learning models is to determine the extent to which multimodal learning aids in the acquisition of a new system built upon an existing one.

In conclusion, the commentators offered invaluable insights regarding the scope of research on bilingual modeling and the merits and drawbacks of employing pluralist approaches to fully leverage modeling in the context of bilingualism. The path to understanding bilingualism via modeling presents its unique set of challenges, yet these discussions indisputably set the stage for wider and more thorough investigations. Moreover, the conversations have provided us an opportunity to further elucidate our views on the necessity and promises of extending the theoretical and methodological horizons of bilingual modeling to wider fields of cognitive science. We eagerly anticipate new discoveries that will move our current perspectives to a new level of understanding, particularly in the cross-disciplinary integration of neurocognitive and computational models for the understanding of bilingualism.

Acknowledgements

Preparation of this article has been partially supported by the following grants: Hong Kong Research Grants Council (Project #PolyU15601520), Sin Wai Kin Foundation endowment grant, and the BERC 2022-2025 program Funded by the Spanish State Research Agency through BCBL Severo Ochoa excellence accreditation CEX2020-001010/AEI/10.13039/501100011033. We thank Jenny Geng for her editorial assistance to the manuscript.

References

Artetxe, M., Labaka, G., & Agirre, E. (2020). Translation Artifacts in Cross-lingual Transfer Learning. *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 7674–7684. https://doi.org/10.18653/v1/2020.emnlp-main.618

Binz, M., & Schulz, E. (2023). Using cognitive psychology to understand GPT-3. *Proceedings of the National Academy of Sciences*, *120*(6), e2218523120. https://doi.org/10.1073/pnas.2218523120

Chomsky, N. (1981). Lectures on Government and Binding. *Dordrecht: Foris*.

Chomsky, N., Roberts, L., & Watumull, J. (2023, March). Noam Chomsky: The False Promise of ChatGPT. *The New York Times*.

Claussenius-Kalman, H., Hernandez, A. E., & Li, P. (2021). Expertise, ecosystem, and emergentism: Dynamic developmental bilingualism. *Brain and Language, 222*, 105013. https://doi.org/10.1016/j.bandl.2021.105013

Contreras Kallens, P., Kristensen-McLachlan, R. D., & Christiansen, M. H. (2023). Large Language Models Demonstrate the Potential of Statistical Learning in Language. *Cognitive Science*, *47*(3), e13256. https://doi.org/10.1111/cogs.13256

DeLuca, V., Rothman, J., Bialystok, E., & Pliatsikas, C. (2019). Redefining bilingualism as a spectrum of experiences that differentially affects brain structure and function. *Proceedings of the National Academy of Sciences*, *116*(15), 7565–7574. https://doi.org/10.1073/pnas.1811513116

Dijkstra, T., & van Heuven, W. J. B. (2022). Inventing and Reinventing the Cog: A Commentary on "Computational Modeling of Bilingual Language Learning: Current Models and Future Directions." *Language Learning*, lang.12532. https://doi.org/10.1111/lang.12532

Elman, J. L. (1993). Learning and development in neural networks: The importance of starting small. *Cognition*, *48*(1), 71–99. https://doi.org/10.1016/0010-0277(93)90058-4

Frank, M. C., Braginsky, M., Yurovsky, D., & Marchman, V. A. (2017). Wordbank: An open repository for developmental vocabulary data. *Journal of Child Language*, *44*(3), 677–694. https://doi.org/10.1017/S0305000916000209

Goldstein, A., Zada, Z., Buchnik, E., Schain, M., Price, A., Aubrey, B., Nastase, S. A., Feder, A., Emanuel, D., Cohen, A., Jansen, A., Gazula, H., Choe, G., Rao, A., Kim, C., Casto, C., Fanda, L., Doyle, W., Friedman, D., … Hasson, U. (2022). Shared computational principles for language processing in humans and deep language models. *Nature Neuroscience*, *25*(3), 369–380. https://doi.org/10.1038/s41593-022-01026-4

Grosjean, F., Li, P., & Bialystok, E. (2013). *The psycholinguistics of bilingualism*. Wiley-Blackwell/John Wiley & Sons.

Jeong, Y., Cho, H.-Y., Kim, M., Oh, J.-P., Kang, M. S., Yoo, M., Lee, H.-S., & Han, J.-H. (2021). Synaptic plasticity-dependent competition rule influences memory formation. *Nature Communications*, *12*(1), 3915. https://doi.org/10.1038/s41467-021-24269-4

Kachergis, G., Marchman, V. A., & Frank, M. C. (2022). Toward a “Standard Model” of Early Language Learning. *Current Directions in Psychological Science*, *31*(1), 20–27. https://doi.org/10.1177/09637214211057836

Kachergis, G., Tan, A. W. M., & Frank, M. C. (2022). Plurality Is a Good Start, but It’s Time for Unification: A Commentary on “Computational Modeling of Bilingual Language Learning: Current Models and Future Directions.” *Language Learning*, lang.12531. https://doi.org/10.1111/lang.12531

Kriegeskorte, N., & Douglas, P. K. (2018). Cognitive computational neuroscience. *Nature*

Neuroscience, *21*(9), 1148–1160. https://doi.org/10.1038/s41593-018-0210-5

Kuhl, P. K., Tsao, F.-M., & Liu, H.-M. (2003). Foreign-language experience in infancy: Effects of short-term exposure and social interaction on phonetic learning. *Proceedings of the National Academy of Sciences*, *100*(15), 9096–9101. https://doi.org/10.1073/pnas.1532872100

Legault, J., Zhao, J., Chi, Y.-A., Chen, W., Klippel, A., & Li, P. (2019). Immersive Virtual Reality as an Effective Tool for Second Language Vocabulary Learning. *Languages*, *4*(1), 13. https://doi.org/10.3390/languages4010013

Lepori, M. A., Serre, T., & Pavlick, E. (2023). *Break It Down: Evidence for Structural Compositionality in Neural Networks* (arXiv:2301.10884). arXiv. http://arxiv.org/abs/2301.10884

Li, P., & Jeong, H. (2020). The social brain of language: Grounding second language learning in social interaction. *Npj Science of Learning*, *5*(1), 8. https://doi.org/10.1038/s41539-020-0068-7

Li, P., & Lan, Y.-J. (2022). Digital Language Learning (DLL): Insights from Behavior, Cognition, and the Brain. *Bilingualism: Language and Cognition*, *25*(3), 361–378. https://doi.org/10.1017/S1366728921000353

Linck, J. A., Kroll, J. F., & Sunderman, G. (2009). Losing Access to the Native Language While Immersed in a Second Language: Evidence for the Role of Inhibition in Second-Language Learning. *Psychological Science*, *20*(12), 1507–1515. https://doi.org/10.1111/j.1467-9280.2009.02480.x

MacWhinney, B. (2000). *The CHILDES project: Tools for analyzing talk* (3rd ed). Lawrence Erlbaum.

Marian, V. (2022). To Bilingualism and Beyond! Modeling Bilingualism Requires Looking Beyond Language: A Commentary on “Computational Modeling of Bilingual

Language Learning: Current Models and Future Directions.” *Language Learning*, lang.12530. https://doi.org/10.1111/lang.12530

Marjieh, R., Sucholutsky, I., van Rijn, P., Jacoby, N., & Griffiths, T. L. (2023). *What Language Reveals about Perception: Distilling Psychophysical Knowledge from Large Language Models* (arXiv:2302.01308). arXiv. http://arxiv.org/abs/2302.01308

Mollo, D. C., & Millière, R. (2023). *The Vector Grounding Problem* (arXiv:2304.01481). arXiv. http://arxiv.org/abs/2304.01481

OpenAI. (2023). *GPT-4 Technical Report* (arXiv:2303.08774). arXiv. http://arxiv.org/abs/2303.08774

Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C. L., Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A., Schulman, J., Hilton, J., Kelton, F., Miller, L., Simens, M., Askell, A., Welinder, P., Christiano, P., Leike, J., & Lowe, R. (2022). *Training language models to follow instructions with human feedback*. https://doi.org/10.48550/ARXIV.2203.02155

Peeters, D. (2019). Virtual reality: A game-changing method for the language sciences. *Psychonomic Bulletin & Review*, *26*(3), 894–900. https://doi.org/10.3758/s13423-019-01571-3

Peñaloza, C., Grasemann, U., Miikkulainen, R., & Kiran, S. (2023). Modeling Bilingualism as a Dynamic Phenomenon in Healthy and Neurologically Affected Speakers Across the Lifespan: A Commentary on “Computational Modeling of Bilingual Language Learning: Current Models and Future Directions.” *Language Learning*, lang.12566. https://doi.org/10.1111/lang.12566

Peng, B., Alcaide, E., Anthony, Q., Albalak, A., Arcadinho, S., Cao, H., Cheng, X., Chung, M., Grella, M., GV, K. K., He, X., Hou, H., Kazienko, P., Kocon, J., Kong, J., Koptyra, B., Lau, H., Mantri, K. S. I., Mom, F., … Zhu, R.-J. (2023). *RWKV:*

Reinventing RNNs for the Transformer Era (arXiv:2305.13048). arXiv. http://arxiv.org/abs/2305.13048

Piantadosi, S. T. (2023). *Modern language models refute Chomsky's approach to language*.

Pinker, S. (1994). *The language instinct*.

Redcay, E., & Schilbach, L. (2019). Using second-person neuroscience to elucidate the mechanisms of social interaction. *Nature Reviews Neuroscience*, *20*(8), 495–505. https://doi.org/10.1038/s41583-019-0179-4

Taori, R., Gulrajani, I., Zhang, T., Dubois, Y., Li, X., Guestrin, C., Liang, P., & Hashimoto, T. B. (2023). *Alpaca: A Strong, Replicable Instruction-Following Model* [Stanford Center for Research on Foundation Models]. https://crfm.stanford.edu/2023/03/13/alpaca.html

Tomasello, M. (2004). Learning through Others. *Daedalus*, *133*(No.1), 51–58. https://www.jstor.org/stable/20027896

Tsoukala, C., Frank, S. L., & Broersma, M. (2017). *"He's pregnant": Simulating the confusing case of gender pronoun errors in L2 English*. 3392–3397.

Wu, S., & Dredze, M. (2019). Beto, Bentz, Becas: The Surprising Cross-Lingual Effectiveness of BERT. *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 833–844. https://doi.org/10.18653/v1/D19-1077

Xu, Q., Peng, Y., Wu, M., Xiao, F., Chodorow, M., & Li, P. (2023). *Does Conceptual Representation Require Embodiment? Insights From Large Language Models* (arXiv:2305.19103). arXiv. http://arxiv.org/abs/2305.19103

Yang, C. D. (2004). Universal Grammar, statistics or both? *Trends in Cognitive Sciences*, *8*(10), 451–456. https://doi.org/10.1016/j.tics.2004.08.006