Cognitive Neuroscience
Current Debates, Research & Reports

ISSN: 1758-8928 (Print) 1758-8936 (Online) Journal homepage: www.tandfonline.com/journals/pcns20

Beyond prediction: comments on the format of natural intelligence

Elliot Murphy

To cite this article: Elliot Murphy (18 Jun 2025): Beyond prediction: comments on the format of natural intelligence, Cognitive Neuroscience, DOI: 10.1080/17588928.2025.2521403

To link to this article: https://doi.org/10.1080/17588928.2025.2521403

Published online: 18 Jun 2025.

COGNITIVE NEUROSCIENCE
https://doi.org/10.1080/17588928.2025.2521403

COMMENTARY

Beyond prediction: comments on the format of natural intelligence

Elliot Murphy[a,b]

[a]Vivian L. Smith Department of Neurosurgery, McGovern Medical School, UTHealth, Houston, TX, USA; [b]Texas Institute for Restorative Neurotechnologies, UTHealth, Houston, TX, USA

Parr et al. (2025) treat syntax as rules over temporal sequences – a non-Markovian process that can be captured by autoregressive or hierarchical-temporal models. This move quietly collapses into two distinct representational layers:

Syntax: Hierarchical constituent geometry formed in an abstract workspace. This is *never linear*.

Externalized phonology: How syntax is transformed from a hierarchical tree-structure to a string. This is *always linear*.

By equating syntax with non-Markovian statistics over strings, the authors risk inheriting two problems:

(1) **Linguistic adequacy**: Many phenomena (binding, displacement, island effects, structural ambiguity, *wh*-movement, topicalization, recursion limits) are defined over trees, not strings (Adger, 2003; Everaert et al., 2015; Marcolli & Larson, 2025; Murphy, 2020, 2025; Murphy & Leivada, 2022; Reuland, 2011), and are not computed over sequential rules, but over *structural/hierarchical distance* (Figure 1).

(2) **Neurobiological adequacy**: In addition to tracking surface statistics, the brain *also* tracks constituent boundaries and tree-geometric depth (McCarty et al., 2023; Murphy, 2024; Weissbart & Martin, 2024; Woolnough et al., 2023).

This human bias to impose tree-structures may be innate (Crain et al., 2017; Perkins & Lidz, 2021; Yang et al., 2017). I will assume here (not without contention) that trees are a privileged representational format (see Dehaene et al., 2022; Senturia & Frank, 2023). In contrast, transformer-based language models seem to have a more immediate bias for imposing linear relations (Huang et al., 2024; McCoy et al., 2020; but see also Ahuja et al., 2025).

Recent surveys indicate that OpenAI's o3-mini-high fails to reliably distinguish between syntax and semantics (Murphy, Leivada, et al., 2025), with other studies offering conservative assessments of the prospects for LLMs to capture human-like linguistic intelligence (Dentella et al., 2024; Murphy, de Villiers, et al., 2025; for alternative perspectives, see Goldberg et al., 2025). Some of the neurobiological evidence cited by Parr and colleagues has been argued to be problematic through overlooked confounds (Hadidi et al., 2025).

Meanwhile, Slaats and Martin (2025) demonstrate why surprisal does not lead to natural language syntax, arguing that 'surprisal and entropy are not, by themselves, suitable as *explanans* for the core capacities of language;' probabilistic information is rather a 'cue to the next level of abstraction' (Slaats & Martin, 2025). We might consult here hierarchical generative syntax where probabilities sit on trees, not strings (Hunter, 2019).

Parr and colleagues could potentially enrich their research program by importing syntactic/parsing tree-generating operations into their latent state space, preserving the explanatory successes of generative grammar while retaining probabilistic machinery for prediction and memory. Deep temporal models can host typed variables which renders the transition to modeling certain features of syntactic recursion easier. It is true that a higher-level latent state can remain unchanged, while a lower-level sequence unfolds – this gives us temporal abstraction. What it does not yet buy us is *configurational simultaneity*, where multiple terminals are co-present in the same derivational moment even though they are pronounced at different times (necessary for agreement and binding). Perhaps, the higher-level sentence state could spawn an event frame within which lower-level factors are synchronized but unordered. If we also have an explicit linearization policy conditioned on the syntactic workspace state, a workspace tensor, an internal symbolic stack, or pointer variables, we might begin to approximate something like a push-down automata, with variational inference replacing deterministic stack operations. One trick here might be to make the latent state at each time-step

CONTACT Elliot Murphy elliot.murphy@uth.tmc.edu Vivian L. Smith Department of Neurosurgery, McGovern Medical School, UTHealth, 1133 John Freeman Blvd, Houston, TX 77030, USA

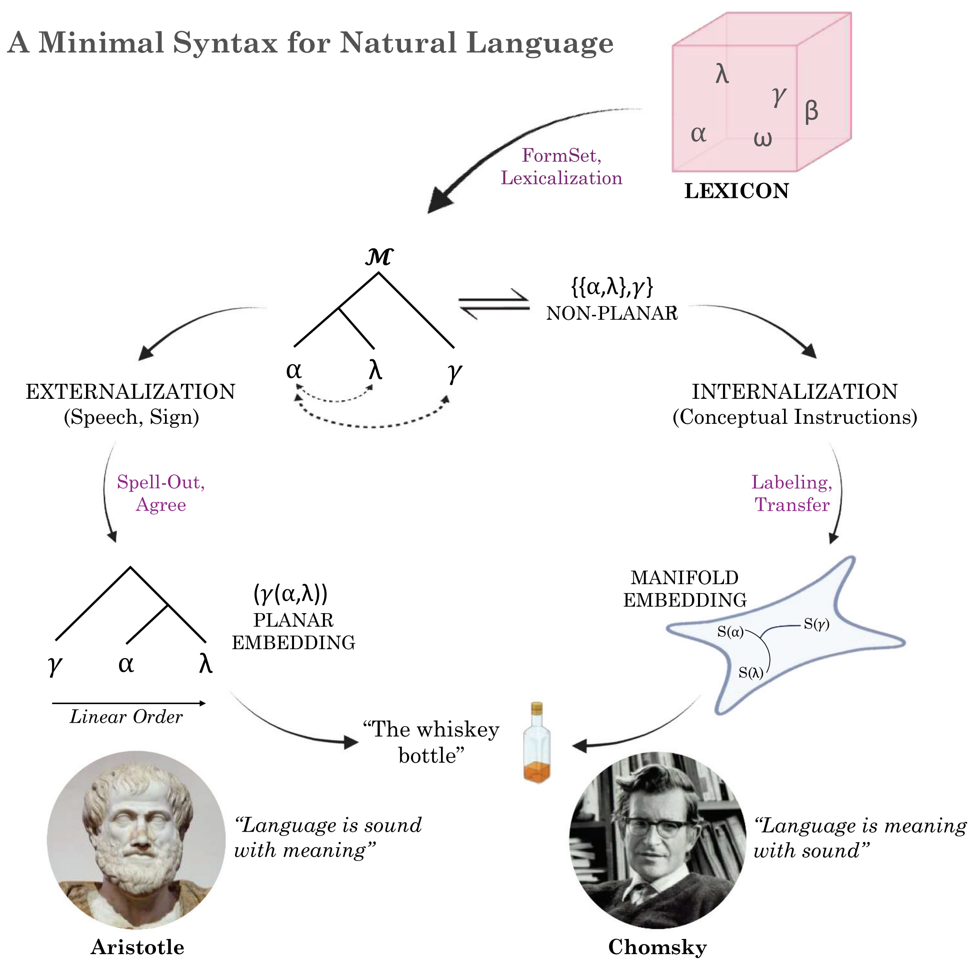

Figure 1. The architecture of the human language faculty (presented here at a higher altitude of generalization than specialist presentations in Chomsky et al., 2023; Marcolli et al., 2025; Murphy et al., 2024). This model is strictly atemporal; MERGE is a *free non-associative commutative magma*, and generates nonplanar trees. This algebraic formulation also provides potential (though not strict) constraints on what the neural code for syntax should look like. During internalization, structures are mapped to distinct conceptual systems (likely 'core knowledge systems;' Spelke & Kinzler, 2007) with their own distinct manifold embeddings. Lexical elements $(\alpha, \lambda, \gamma)$ are mapped to a workspace via MERGE ($\mathcal{M}$), which is then sent to interpretive systems and categorized ('labeled'), and/or are 'spelled-out' at externalization systems for speech, gesture, etc. For example, 'the whiskey bottle' would be formed via {*whiskey*, *bottle*} being merged with {*the*}, to form a hierarchical object that is externalized in different orders depending on the language in question (i.e., not all languages adhere to the English 'Determiner-Noun' order, but all human speakers converge on the same coordinates in semantic space for the phrase). The quotations at the bottom contrast opposing views on language design (see Chomsky, 2011). Portrait photographs reproduced from creative commons (Wikimedia Commons; commons.Wikimedia.org): public domain photo of marble bust of Aristotle (left). Public domain photo of Chomsky from Andrea Womack (c. 1973) (right).

a representation of 'current derivational state,' not just 'what word came before' – a left-corner minimalist grammar parser might then fall out naturally from the same inference machinery that already drives attention and memory.

While it is possible that human syntax might algorithmically be implemented by transformers (Marcolli et al., 2025), there may be a kind of competence-performance gap, or the format of representation may be entirely distinct. It is not immediately clear how we might begin to iron out these concerns – but the exciting approach from Parr and colleagues is nevertheless a fruitful step in this direction. For example, there may be certain ways to reconcile predictive processing with the present observations about syntax (e.g., non-Markovian memory reconciled with syntactic derivational history stack; attention gain reconciled with derivation phrase head selection; hidden state reconciled with syntactic workspace).

With all this said, I should also acknowledge the possibility that the algebraic recipe for syntax provided by modern mathematical linguistics (Figure 1) may be insufficient for mapping the psycholinguistic and neural properties of language. The brain may not in fact construct tree-structures in real time, and theoretical models of syntax may end up describing our knowledge of language at some higher level of generalized abstraction. As Captain Barbossa once described the pirates' code, it may have less in common with hard-coded rules and be more akin to a set of 'guidelines.'

Disclosure statement

No potential conflict of interest was reported by the author(s).

Funding

The author(s) reported that there is no funding associated with the work featured in this article.

References

Adger, D. (2003). *Core syntax: A minimalist approach*. Oxford University Press.

Ahuja, K., Balachandran, V., Panwar, M., He, T., Smith, N. A., Goyal, N., & Tsvetkov, Y. (2025). Learning syntax without planting trees: Understanding hierarchical generalization in transformers. *Transactions of the Association for Computational Linguistics*, *13*, 121–141. https://doi.org/10.1162/tacl_a_00733

Chomsky, N. (2011). Language and other cognitive systems. What is special about language? *Language Learning and Development*, *7*(4), 263–278. https://doi.org/10.1080/15475441.2011.584041

Chomsky, N., Seely, T. D., Berwick, R. C., Fong, S., Huybregts, M. A. C., Kitahara, H., McInnerney, A., & Sugimoto, Y. (2023). *Merge and the strong minimalist thesis*. Elements in Generative Syntax. Cambridge University Press.

Crain, S., Koring, L., & Thornton, R. (2017). Language acquisition from a biolinguistic perspective. *Neuroscience & Biobehavioral Reviews*, *81*(B), 120–149. https://doi.org/10.1016/j.neubiorev.2016.09.004

Dehaene, S., Al Roumi, F., Lakretz, Y., Planton, S., & Sablé-Meyer, M. (2022). Symbols and mental programs: A hypothesis about human singularity. *Trends in Cognitive Sciences*, *26*(9), 751–766. https://doi.org/10.1016/j.tics.2022.06.010

Dentella, V., Günther, F., Murphy, E., Marcus, G., & Leivada, E. (2024). Testing AI on language comprehension tasks reveals insensitivity to underlying meaning. *Scientific Reports*, *14*(1), 28083. https://doi.org/10.1038/s41598-024-79531-8

Everaert, M. B. H., Huybregts, M. A. C., Chomsky, N., Berwick, R. C., & Bolhuis, J. J. (2015). Structures, not strings: Linguistics as part of the cognitive sciences. *Trends in Cognitive Sciences*, *19*(12), 729–743. https://doi.org/10.1016/j.tics.2015.09.008

Goldberg, A. E., Rakshit, S., Hu, J., & Mahowald, K. (2025). A suite of LMs comprehend puzzle statements as well as humans. *arXiv*, 2505.08996. https://doi.org/10.48550/arXiv.2505.08996

Hadidi, N., Feghhi, E., Song, B. H., Blank, I. A., & Kao, J. C. (2025). Illusions of alignment between large language models and brains emerge from fragile methods and overlooked confounds. *bioRxiv*. https://doi.org/10.1101/2025.03.09.642245

Huang, K. J., Arehalli, S., Kugemoto, M., Muxica, C., Prasad, G., Dillon, B., & Linzen, T. (2024). Large-scale benchmark yields no evidence that language model surprisal explains syntactic disambiguation effects. *Journal of Memory and Language*, *137*, 104510. https://doi.org/10.1016/j.jml.2024.104510

Hunter, T. (2019). Left-corner parsing of minimalist grammars. In R. C. Berwick & E. P. Stabler (Eds.), *Minimalist parsing* (pp. 125–158). Oxford University Press.

Marcolli, M., Chomsky, N., & Berwick, R. C. (2025). *Mathematical structure of syntactic merge: An algebraic model for generative Linguistics*. MIT Press.

Marcolli, M., & Larson, R. K. (2025). Theta theory: Operads and coloring. *arXiv*, 2503.06091. https://doi.org/10.48550/arXiv.2503.06091

McCarty, M. J., Murphy, E., Scherschligt, X., Woolnough, O., Morse, C. W., Snyder, K., Mahon, B. Z., & Tandon, N. (2023). Intraoperative cortical localization of music and language reveals signatures of structural complexity in posterior temporal cortex. *iScience*, *26*(7), 107223. https://doi.org/10.1016/j.isci.2023.107223

McCoy, T., Frank, R., & Linzen, T. (2020). Does syntax need to grow on trees? Sources of hierarchical inductive bias in sequence-to-sequence networks. *Transactions of the Association for Computational Linguistics*, *8*, 125–140. https://doi.org/10.1162/tacl_a_00304

Murphy, E. (2020). Language design and communicative competence: The minimalist perspective. *Glossa: A Journal of General Linguistics*, *5*(1), 2. https://doi.org/10.5334/gjgl.1081

Murphy, E. (2024). ROSE: A neurocomputational architecture for syntax. *Journal of Neurolinguistics*, *70*, 101180. https://doi.org/10.1016/j.jneuroling.2023.101180

Murphy, E. (2025). The nature of language and the structure of reality. In A. Benítez-Burraco, I. F. López, M. Férnandez-Pérez, & O. Ivanova (Eds.), *Biolinguistics at the cutting edge: Promises, achievements, and challenges* (pp. 207-236). De Gruyter Mouton.

Murphy, E., de Villiers, J., & Morales, S. L. (2025). A comparative investigation of compositional syntax and semantics in DALL·E and young children. *Social Sciences & Humanities Open*, *11*, 101332. https://doi.org/10.1016/j.ssaho.2025.101332

Murphy, E., Holmes, E., & Friston, K. (2024). Natural language syntax complies with the free-energy principle. *Synthese*, *203*(5), 154. https://doi.org/10.1007/s11229-024-04566-3

Murphy, E., & Leivada, E. (2022). A model for learning strings is not a model of language. *Proceedings of the National Academy of Sciences of the United States of America*, *119*(23), e2201651119. https://doi.org/10.1073/pnas.2201651119

Murphy, E., Leivada, E., Dentella, V., Günther, F., & Marcus, G. (2025). Fundamental principles of linguistic structure are not represented by o3, *arXiv*, 2502.10934. https://doi.org/10.48550/arXiv.2502.10934

Parr, T., Pezzulo, G., & Friston, K. (2025). Beyond Markov: Transformers, memory, and attention. *Cognitive Neuroscience*, 1–19. https://doi.org/10.1080/17588928.2025.2484485

Perkins, L., & Lidz, J. (2021). Eighteen-month-old infants represent nonlocal syntactic dependencies. *Proceedings of the National Academy of Sciences of the United States of

America, *118*(41), e2026469118. https://doi.org/10.1073/pnas.2026469118
Reuland, E. (2011). *Anaphora and language design*. MIT Press.
Senturia, I., & Frank, R. (2023). On the spectra of syntactic structures. In *Proceedings of the Society for Computation in Linguistics (SCiL) 2023* June 15-17 Amherst, MA Tim Hunter, Brandon Prickett (Association for Computational Linguistics) (pp. 199–211). https://aclanthology.org/2023.scil-1.17/
Slaats, S., & Martin, A. E. (2025). What's surprising about surprisal. *Computational Brain & Behavior*, *8*(2), 233–248. https://doi.org/10.1007/s42113-025-00237-9
Spelke, E. S., & Kinzler, K. D. (2007). Core knowledge. *Developmental Science*, *10*(1), 89–96. https://doi.org/10.1111/j.1467-7687.2007.00569.x
Weissbart, H., & Martin, A. E. (2024). The structure and statistics of language jointly shape cross-frequency neural dynamics during spoken language comprehension. *Nature Communications*, *15*(1), 8850. https://doi.org/10.1038/s41467-024-53128-1
Woolnough, O., Donos, C., Murphy, E., Rollo, P. S., Roccaforte, Z. J., Dehaene, S., & Nitin, T. (2023). Spatiotemporally distributed frontotemporal networks for sentence reading. *Proceedings of the National Academy of Sciences of the United States of America*, *120*(17), e2300252120. https://doi.org/10.1073/pnas.2300252120
Yang, C., Crain, S., Berwick, R. C., Chomsky, N., & Bolhuis, J. J. (2017). The growth of language: Universal grammar, experience, and principles of computation. *Neuroscience & Biobehavioral Reviews*, *81*(B), 103–119. https://doi.org/10.1016/j.neubiorev.2016.12.023