Routledge
Taylor & Francis Group

Cognitive Neuroscience
Current Debates, Research & Reports

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

Necessity and plurality of neural dynamics under ROSE

Elliot Murphy

To cite this article: Elliot Murphy (17 Aug 2026): Necessity and plurality of neural dynamics under ROSE, Cognitive Neuroscience, DOI: 10.1080/17588928.2026.2718716

To link to this article: https://doi.org/10.1080/17588928.2026.2718716

Published online: 17 Aug 2026.

COGNITIVE NEUROSCIENCE
https://doi.org/10.1080/17588928.2026.2718716

Routledge
Taylor & Francis Group

Necessity and plurality of neural dynamics under ROSE

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

ABSTRACT

I am grateful to the authors of the commentary articles for identifying productive points of pressure for any neurocomputational account of syntax: whether semantic interpretation can proceed without a full syntactic derivation; how dynamical motifs are selected and coordinated; whether proposed mechanisms generalize across languages and modalities; how evolution may have repurposed older memory circuitry; and how neural dynamics distinguish types and tokens. Here, I clarify that ROSE is an architecture for implementing hierarchical syntactic computation, bringing with it no commitment that its full code is obligatory for every act of meaning construction. The commentaries suggest a number of compelling concrete extensions: task-dependent gating of S/E; factorized R/O subspaces; representation-relative rather than language-specific complexity measures; comparative tests of neural reuse; and occurrence-sensitive phase or state-space addresses. These thoughtful extensions sharpen the central aim of ROSE to formulate multiscale, falsifiable links between formal properties of language and neural dynamics.

ARTICLE HISTORY
Received 3 August 2026

KEYWORDS
ROSE; neural dynamics; syntax; dynamical motifs

The commentary articles responding to Murphy (2025) are thoughtful, generous, and appropriately demanding. Each of them offers creative approaches to assessing neurocomputational architectures for natural language. Taken together, they encourage a reading of ROSE not as a fixed catalog, but as a multiscale research program and conceptual framework whose linking hypotheses must be tested against alternative mechanisms, diverse languages and modalities, and the broader computational capacities of animal brains. One emerging theme throughout is that core properties of syntax like recursive composition should be seen neither as a neural primitive that can simply be read off oscillatory activity nor as a mechanism that must be necessarily invoked whenever language conveys meaning. In that unrestricted form, recursive composition risks becoming explanatorily moribund. What is needed are specifications as to when recursive structure building is computationally necessary, how its component operations might be implemented through more general neural dynamics, and which empirical results would distinguish those mechanisms from shallower semantic or predictive routes. Below, I address each commentary contribution in turn.

1. Reply to Baggio: Syntax on demand and the scope of ROSE

Baggio (2025) raises an important question concerning autonomous semantics, and I thank him for outlining – with his usual application of lucid prose and masterful historical insight – a range of interesting evidence for this important intervention. His questions about meaning construction are highly relevant, and the first topic that requires clarification here is that his initial thesis attributes to ROSE a stronger commitment than the model requires. ROSE is an account of the neurocomputational architecture of syntax: it asks how hierarchical structure can be constructed and maintained when the language system must preserve distinctions that depend on non-associative grouping and phrasal constituency. It is not a general theory claiming that every meaningful utterance, lexical association, pragmatic inference, or conceptual combination must recruit the full S/E machinery. Indeed, the target article explicitly treats the proposed oscillatory relations as principal rather than exclusive drivers, permits context-dependent movement between more predictive and more hierarchical parsing strategies, and allows shallow or 'good-enough' parsing when statistical information is sufficient (Murphy, 2025). It even leaves open the possibility that

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

future work may replace the low-frequency phase code or reduce the role of syntactic categories altogether. Still, 'good-enough' parsing supposes some level of syntax, and while Murphy (2025) rejected obligatory deployment of the full hierarchical code, Baggio is surely right that I left the matter insufficiently explicit regarding whether some acts of meaning construction can bypass syntactic computation altogether.

I will address each of these issues specifically. Baggio (2025) argues that 'models of language in the brain should make room for the possibility that meaning may also be derived in the absence of any syntactic computation, be it hierarchical or predictive.' Murphy (2025) briefly discusses this, noting that 'in everyday cases the use of statistical inferences may often suffice to extract a type of shallow parse (a "good-enough" parse) for linguistic structure'; further, 'many predictions can be successful simply by virtue of lexicosemantic associations [. . .], although semantic prediction mechanisms can also prime words that may *not* fit the correct syntactic category [. . .]. ROSE therefore needs to flexibly adapt to how individuals access and exploit the resources of their "I-language."' There is nothing inherent to ROSE that *requires* meaning to be constructed always via its syntax-driven apparatus or via predictive processing – this is why I also invoke processes like delta-gamma domain-general combinatorics, and the lexical workspace vs. structural workspace distinction. Baggio is right to point out that I kept the focus on syntactic vs. predictive processing, but many of the neural generators associated with ROSE can carry the weight of nonhierarchical semantics, and this topic was certainly deserving of some explicit treatment in the target article. Any linguistic meaning that is built 'below' the full S/E complex will be suited to these forms of less complex idiomatic or 'linear grammar'-type Jackendoffian objects, amongst others relevant to compositional semantics. As such, autonomous semantic composition is wholly compatible with ROSE.

According to Baggio's characterization of the 'Brain-as-Linguist' view, there is a 'commitment to the idea that syntax is obligatorily represented for any expression that is acquired, generated, or interpreted by the brain.' While Baggio is right to probe this important issue, I do not see anything *obligatory* for general meaning construction in the architecture of ROSE that necessitates this. Certainly, a core prediction is that when the full ROSE architecture is deployed, syntax will always be found – but, as the target article briefly noted, there are forms of linguistic processing that will trigger lexico-syntactic feature spaces (if we assume, as I do, that every word has a syntactic category) *without* necessitating the generation of hierarchical constituency structure via the S/E levels. Semantic processing, as Baggio notes, can proceed often without the need to infer hierarchical constituency structure, and indeed there are many mechanisms within the ROSE model to accommodate this. Going even further, as Murphy (2025) also notes, citing Elbourne's work, ROSE must additionally be open to the possibility that 'syntactic category' *is not even used at all* during parsing – this would shift the focus of what representational tokens are accessed via R/O without necessarily forcing a substantial revision of how S/E coordinate such features. Murphy (2025) notes that this possibility is not only a plausible direction for future work, but may 'be more directly relatable to neurocognitive architectures like ROSE in which conceptual representations are primary (i.e., the system "starts" with lexico-semantic features at R).' Baggio likewise asks us to consider cases in which we 'start' the system with semantics, not syntax, and indeed this is what ROSE does (Murphy, 2024).

We might distinguish here three possible options for meaning generation:

(1) **Full hierarchical syntax**: constituency and recursive S/E structure.
(2) **Sparse syntax**: lexico-syntactic category, local dependency, or linear grammatical constraints without a complete phrase marker.
(3) **No syntactic computation**: semantic constructions operating entirely within conceptual or associative representations.

ROSE delivers clear predictions for which of these trigger distinct mechanisms across the four levels, from R to E, and Baggio's comments provide interesting directions for a more psycholinguistically plausible and comprehensive assessment.

Perhaps, then, a useful clarification would be to distinguish the availability of hierarchical syntax from its degree of recruitment on a particular occasion. The R/O levels can support lexical-semantic activation, contextual updating, and associative prediction without requiring every such process to be recoded as a complete phrase marker. The S/E levels become critical when alternative groupings generate different interpretations (especially cases of direct structural ambiguity, like 'old men and women'), when distant elements must be related (especially non-local dependencies anchored in hierarchical relations), or when a completed constituent must be made available for further recursive computation. On this view, 'syntax on demand' can be implemented within ROSE as *task- and context-dependent gain* on

the pathways linking R/O to S/E, rather than as a rival architecture. Whether apparently syntax-free interpretation is literally devoid of syntactic information or instead exploits sparse category and dependency constraints remains an open question. Recent work likewise points to distinguishable but interacting lexical and compositional meaning spaces, rather than a single obligatory route to interpretation (Brouwer, 2026; Venhuizen & Brouwer, 2025).

In brief, ROSE certainly assumes – via the labeling hypothesis (Murphy, 2015) and my definition of the brain's language network (Murphy, 2025) – that labeling and 'narrow syntax' are obligatory for certain syntactic compositional processes (Murphy, Woolnough et al., 2022), but not for *all logically possible* forms of meaning-formation, of which there are certainly many. Baggio's proposal identifies an important extension that was left insufficiently addressed in Murphy (2025): R/O dynamics must sometimes be capable of combining conceptual content without invoking category, dependency, or hierarchical constituency representations. The empirical challenge is to distinguish such cases from processing that exploits sparse syntactic constraints.

2. Reply to Labancová and Kazanina: Dynamical motifs as reusable R/O routines

Building on (Driscoll et al., 2024; Labancová and Kazanina 2025) provides a valuable operational interpretation of dynamical motifs. I thank them for their thoughtful and innovative approach to this issue, and hope that my brief reflections here do justice to their thesis.

Labancová and Kazanina review the role of dynamical motifs in Murphy (2025) and rightly note that 'details of this process remain unclear.' This issue is both a theoretical and empirical one: ROSE must specify what constitutes a motif, how motifs are selected, and which observations would distinguish reusable factorization from generic mixed selectivity. In response to this, the appeal the authors make to factorized subspaces is especially timely. Tafazoli et al. (2026) show that macaque brains reuse task-relevant sensory and motor subspaces across compositionally related tasks, engaging and transforming them according to the animal's current task belief. For ROSE, the most useful analogue is unlikely to be a rigid one-word/one-motif dictionary. A motif should instead be understood as a reusable local transformation or attractor geometry that binds a content representation to a currently relevant grammatical category, thematic role, or other relational state. The same lexical item can therefore enter different trajectories, while the same role-binding computation can generalize across different lexical contents. This factorization gives the R/O levels both stability and flexibility.

The surrounding dynamics queried by Labancová and Kazanina are precisely where the higher ROSE levels can add explanatory value. Indeed, Labancová and Kazanina's speculations about how dynamical motifs may enforce thematic role assignments, shaped by a generative model of human 'beliefs' about parsing states, are particularly interesting. More broadly, parser state, contextual expectation, and task goals can bias which local subspace is excitable; low-frequency phase can open and close the relevant integration windows; spike-phase coupling can stabilize the selected feature bundle; and a phase reset or traveling wave transition can terminate one motif and route its output to the next workspace state. In this context, the temporary subspaces reported by Tian et al. (2024) provide a plausible intermediate mechanism through which factorized representations can be transferred and reorganized as processing demands change. Thus, dynamical motifs need not compete with ROSE's oscillatory scaffold; they provide the population-level computations that the scaffold selectively recruits and organizes.

3. Reply to Trettenbrein: Universality requires modality and typology

I fully agree with Trettenbrein (2025) that claims concerning how a Universal Neural Grammar ultimately stands or falls on cross-linguistic and cross-modal evidence. His commentary reviews a range of fascinating evidence, with great care and insight. As he notes, '[t]he assumption of modality-independent processes for structure building lies at the heart of ROSE, but the proposed correlates for hierarchical and sequential operations must be subjected to empirical test across languages and modalities in the future.'

Within this context, 'universal' should refer to an abstract computational format and to shared constraints on implementation. Trettenbrein highlights that sign languages are particularly decisive because they preserve hierarchical linguistic organization while radically altering the sensory and articulatory interface. Existing evidence that deaf signers recruit left posterior inferior frontal and posterior temporal regions for grammatical processing supports the possibility of a modality-independent core network, while also showing that this core interacts with modality-specific visual, spatial, and motor systems (Trettenbrein et al., 2021, 2024). Tactile sign languages provide an even stronger test because both the input statistics and the externalization channel differ from speech.

Trettenbrein is also right that node count cannot be treated as a language-neutral primitive. Under ROSE, it is best regarded as a parser- and representation-relative proxy for the number and timing of structure-building actions, not as a direct measure of Merge (Murphy, 2020; Murphy, Hoshi, et al., 2022). A postnominal adjective, a morphologically complex sign, and simultaneous manual and non-manual cues may instantiate the same abstract dependency while distributing evidence across different temporal windows and representational units. Cross-linguistic tests should therefore compare formal operations (composition operators that comply with the algebraic properties of syntax, labeling, closure, retrieval, and workspace access to substructures during dependency resolution) under grammars and parsers appropriate to each language. Comparable structural transitions should be recoverable from shared frontotemporal mechanisms after modality-specific timing and cue structure are modeled. Large-scale evidence for a broadly conserved language network across spoken languages (Malik-Moraleda et al., 2022), together with emerging sign-language results, makes this a plausible hypothesis. Trettenbrein's suggested research direction is consequently one of the most important falsification tests for any neurocomputational model of natural language syntax-semantics.

4. Reply to Shi: Neural reuse without reducing syntax to memory

Shi (2025) places ROSE within an evolutionary framework of neural reuse and offers a number of fascinating areas for further discussion. I agree that theta-gamma coordination is an ancient and domain-general resource, especially well characterized in hippocampal systems for sequencing, relational binding, and memory. The species-relevant claim concerns the organization of these generic mechanisms: the representational content they coordinate, their relation to human frontotemporal circuits, their interaction with cortical workspaces, and their capacity to preserve specifically linguistic invariants. Recent human recordings strengthen the case for a memory-language interface and other forms of language-control interfaces (Woolnough et al., 2026). Many other recent findings, too extensive for me to review here, establish that hippocampal circuitry is available for some elements of linguistic processing; but this line of research does not show that the hippocampus computes the core hierarchical operations attributed to the cortical language network. For similar criticisms of a hippocampal-centric neurocomputational model of language, I refer the reader to the excellent discussion in van Bree (2024).

With respect to formal details, Shi's feature-free proposal (entertained also in Murphy, 2020) is compatible with ROSE insofar as labeling continues to be modeled as a form of competition or symmetry-breaking process among candidate representations (Murphy, 2015), orchestrated in Murphy (2025) via phase-amplitude coupling. One remaining question is whether such dynamics can recover the category and dependency distinctions required online.

ROSE can therefore accommodate Shi's exaptation proposal while retaining a division of labor. Hippocampal and medial temporal mechanisms may index events, bind content across episodes, and support the transfer or reinstatement of completed linguistic objects (Murphy et al., 2023; Murphy, Hoshi, et al., 2022). Frontotemporal cortical dynamics may impose the more rapidly updated, category-sensitive, and recursively addressable structure required for online syntax. Increased cortico-hippocampal coordination is a plausible evolutionary and developmental hypothesis, but it should be tested rather than built in as the sole innovation. The strongest developmental tests would track various forms of oscillatory coupling longitudinally while contrasting linguistic hierarchy with non-linguistic event and sequence tasks. A simple coincidence between hippocampal maturation and vocabulary growth cannot distinguish shared maturation, causal reuse, and parallel development.

5. Reply to Uriagereka: From occurrences to dynamical systems

Uriagereka (2025) identifies the type-token-occurrence problem as a central challenge for the neurobiology of language. I thank him for raising this fascinating issue, which is too often overlooked. Repeated words such as 'fish' cannot be understood by measuring an isolated response to a physical token, because each occurrence inherits a different derivational history and occupies a different position in the developing structure. This important point can further clarify what 'atomic' means in ROSE. R-level features are not Lego-like neurons or spatially sealed pieces, but are rather computationally atomic relative to a particular operation: distributed, multiplexed representations that can participate in different population trajectories (Murphy, 2024). A stable subspace or attractor may encode aspects of a type, while a token becomes an occurrence through its context-sensitive trajectory, phase address, workspace position, and pattern of relations to other active objects. Internal Merge would then involve the selective reactivation of content at a new structural address, not the

creation of a second anatomical copy. Forthcoming work attempts to model these dynamics more carefully.

Uriagereka claims that 'Murphy inherits from the tradition [of generative grammar] [. . .] that "neurocomputational architecture for syntax . . . scales from single units to inter-areal dynamics" [quoting Murphy (2025)].' He is right to note that I inherited the 'atoms of language' approach, whereby morphemes are composed of structured bundles of features, but generative grammar did not pass down the notion that neurobiological scales of complexity relate to various levels of linguistic structure. Like some other researchers in the theoretical syntax space who aim to use tools from physics productively, Uriagereka then points to 'Hilbert spaces,' 'superpositions,' 'entangled dependencies,' and 'the known universe.'

Moving forward with this approach, he claims that 'biology seems to work from mundane and fundamental mechanisms, building molecular complexity on essentials so close to (largely subatomic) physics that they hardly have any choices.' The phrases 'largely subatomic' and 'hardly have any choices' overstate the case. I appreciate Uriagereka's intuitions about the importance of grounding the neurobiology in known physical constraints, and have expressed viewpoints wholly sympathetic to this (Murphy et al., 2024), but we also need to walk through this logic carefully:

(1) Biological systems must obey physics. True.
(2) Biological mechanisms are constrained by chemistry, energetics, geometry, and evolutionary history. True.
(3) Those constraints leave biology almost no alternative mechanisms. Generally false.

Physics determines what is possible, but it rarely *uniquely* determines what biology will build. Even at the protein level, the same function can often be implemented by different sequences, structures, or biochemical routes. Starr and colleagues, for example, found hundreds of alternative protein sequences using diverse biochemical mechanisms to produce the same derived function as an ancestral transcription factor (Starr et al., 2017). Certainly, the alternatives are not unlimited. Thermodynamics and molecular structure can strongly favor some solutions: one analysis found hundreds of chemically feasible alternatives to part of glycolysis, while the pathway used by organisms produced the greatest flux under typical physiological conditions (Court et al., 2015). The accurate picture is *constrained plurality*, not 'hardly any choices.'

'Largely subatomic physics' is also potentially misleading. Biological molecular mechanisms are ultimately compatible with fundamental quantum processes, but their useful *explanatory level* is usually atoms, molecules, electrochemical gradients, protein conformations, and cellular networks. Saying that biology is close to subatomic physics is like saying that syntax is close to quantum field theory because brains contain electrons. It is ontologically true at the bottom, but generally uninformative about the organization under investigation. A complex systems perspective is useful here (Krakauer, 2024).

Uriagereka also risks conflating (i) physical realization (linguistic computation must be implemented by matter), (ii) biophysical constraint (neurons cannot implement arbitrary operations), and (iii) mechanistic determination (fundamental physics directly dictates the neural code). ROSE accepts the first two, but the third does not follow. Forthcoming work explores how there may be many neurally admissible implementations of a representational or computational invariant, with evolution selecting one historically contingent solution from a restricted space.

A defensible replacement for Uriagereka's thesis would be something like: Biology constructs complex systems from evolutionarily ancient mechanisms whose operation is constrained by physics and chemistry, although those constraints typically permit multiple historically contingent implementations.

Lastly, Uriagereka writes: 'To make progress in that approach, we must be willing to test other animals, difficult as this may be, and go down to elementary formalisms (like Fourier analysis of relevant waves and applications of the Fourier Transform), to examine how such fundamentals could carry any representation (or whatever it is brains are binding at such basic levels).' Fourier analysis itself does not take us closer to fundamental biological mechanisms. All it does is decompose a signal into frequency components; it does not establish what those components represent, what generated them, or how they perform a computation. Uriagereka's appeal to Fourier methods therefore supports analyzing elementary dynamics, but not the stronger claim that low-level physics leaves biology essentially one representational option. Fourier and time-frequency analyses remain indispensable for identifying candidate carrier dynamics, but spectral decomposition alone cannot determine what a rhythm represents or whether it preserves non-associative grouping. It must be combined with models that distinguish alternative structural histories for tokens. The broader lesson is that dynamical implementation and symbolic invariance are not alternatives. The difficult task is to show how a continuous, distributed system realizes distinctions that can be stated discretely and tested causally.

6. Further clarifications about parsing

Some additional clarifications are needed concerning the predictions about ROSE for detecting signatures of parsing state transitions across neural scales. Due to a certain level of ambiguity about this issue in (Murphy, 2024; Slaats and Hervais-Adelman, 2026), we reasonably conclude that ROSE is undermined by the discovery of high-frequency gamma signatures for sentence-level parsing transitions. This does not follow, however. Murphy (2024) leaned away from attributing causal-explanatory power to mean gamma power for constituent/structural identity but also noted that possible gamma signatures of syntax had indeed been detected, albeit sparsely, in Woolnough et al. (2023). Slaats and Hervais-Adelman (2026) report high-frequency gamma signatures of node-count in multiple parsing models during naturalistic speech comprehension, and argue that this presents a problem for ROSE.

Yet, finding mean gamma signatures spread across the cortex for node structure when the effects of word frequency and other statistical features are regressed out certainly means that those effects are very likely not driven by word frequency – but they *can* also be driven by many other things beyond word frequency and statistics, not exclusively syntactic parsing state, in particular given the cortically widespread distribution of effects. Especially during natural speech comprehension, besides just word frequency and other statistical components there are many potential drivers of why regions across the cortex might show apparent interest in 'node structure' beyond narrowly syntactic inferences (see discussion in Murphy & Woolnough, 2024). This concern becomes especially acute when we begin to report effects of 'syntax' in occipital, motor, dorsal parietal and many other widespread regions.

ROSE assumes that any documented gamma effects of node structure reflect the local accessing of relevant lexico-semantic and syntactic features pertinent to driving inferences about that specific constituent, as outlined in Murphy (2024, 2025). This is not the same as positing that local gamma power is the sole causal driver of constituency structure – hence the introduction of the 'S' and 'E' levels.

Crucially, carrying information about a syntactic constituent is not the same as implementing the operation by which the constituent was constructed.

I argued in Murphy (2024), and maintain here, that mean gamma power by itself is not sufficient to implement or explain the relevant algebraic distinctions of syntax – gamma increasing or decreasing does not yield an explanatory account of constituency, but rather offers just another correlate. This does not lead to the prediction that no cortical region will show gamma-sensitivity for bottom-up or top-down parsing models.

More specifically, ROSE does not require gamma to be statistically insensitive to syntactic structure (Murphy, 2024 already reviews evidence for sparse encoding of syntactic processing in gamma); rather, it proposes that gamma primarily reflects the activation and transformation of linguistic feature bundles, while hierarchical structural identity is causally implemented through other means (S, E). Node count predictors could therefore explain variance in gamma power because changes in parsing state recruit different local operations or because low-frequency structural dynamics modulate gamma amplitude – it is not a case of either/or. Indeed, Slaats and Hervais-Adelman's observation of delta/theta-scale temporal patterns and their proposed role for phase-amplitude coupling are broadly consistent with this interpretation from ROSE.

To summarize falsification criteria discussed elsewhere (Murphy, 2024, 2025) and in forthcoming work: ROSE would be challenged if local gamma dynamics alone encoded and causally preserved non-associative grouping across changes in lexical content, parser strategy, and surface order, without any necessary contribution from slower phase organization or inter-areal state transitions.

All this to say, the target article under discussion (Murphy, 2025) details the role of gamma in connection to specific parsing models and places a very clear role of gamma in syntactic parsing. Under ROSE, gamma carries/transforms locally instantiated feature content used in syntactic computation, whereas the *relational organization* of those contents into recursive hierarchical structures is principally implemented through a phase-content code and network-level coordination.

7. Conclusion

Taken together, the commentaries help direct the search for the neural code for language in productive ways. ROSE must specify when hierarchical computation is recruited; treat motifs as reusable and controlled population routines; define universality over operations rather than surface timing; separate ancestral carrier mechanisms from their human linguistic organization; and model tokens as derivationally situated occurrences. Much work remains to done. As such, I am grateful to the commentary authors for generously engaging with my proposal, and their ideas can help make ROSE more explicitly pluralistic about mechanism while preserving its central demand: candidate neural implementations must be constrained by the computations that language actually requires.

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

Baggio, G. (2025). Autonomous semantics and syntax on-demand in neurocomputational models of language. *Cognitive Neuroscience*, *16*(1–4), 81–82. https://doi.org/10.1080/17588928.2025.2561588

Brouwer, H. (2026). Mapping meaning in the brain's language. *Cortex*, *194*, 12–21. https://doi.org/10.1016/j.cortex.2025.10.012

Court, S., Waclaw, B., & Allen, R. (2015). Lower glycolysis carries a higher flux than any biochemically possible alternative. *Nature Communications*, *6*(1), 8427. https://doi.org/10.1038/ncomms9427

Driscoll, L. N., Shenoy, K. V., & Sussillo, D. (2024). Flexible multitask computation in recurrent networks utilizes shared dynamical motifs. *Nature Neuroscience*, *27*(), 1349–1363. https://doi.org/10.1038/s41593-024-01668-6

Krakauer, D. (2024). *The complex world: An introduction to the foundations of complexity science*. SFI Press.

Labancová, K., & Kazanina, N. (2025). Dynamical motifs for computations in language. *Cognitive Neuroscience*, *16*(1–4), 83–84. https://doi.org/10.1080/17588928.2025.2581573

Malik-Moraleda, S., Ayyash, D., Gallée, J., Affourtit, J., Hoffmann, M., Mineroff, Z., Jouravlev, O., & Fedorenko, E. (2022). An investigation across 45 languages and 12 language families reveals a universal language network. *Nature Neuroscience*, *25*(8), 1014–1019. https://doi.org/10.1038/s41593-022-01114-5

Murphy, E. (2015). Labels, cognomes, and cyclic computation: An ethological perspective. *Frontiers in Psychology*, *6*, 715. https://doi.org/10.3389/fpsyg.2015.00715

Murphy, E. (2020). *The oscillatory nature of language*. Cambridge University Press. https://doi.org/10.1017/9781108864466.

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). ROSE: A universal neural grammar. *Cognitive Neuroscience*, *16*(1–4), 49–80. https://doi.org/10.1080/17588928.2025.2523875

Murphy, E., Forseth, K. J., Donos, C., Snyder, K. M., Rollo, P. S., & Tandon, N. (2023). The spatiotemporal dynamics of semantic integration in the human brain. *Nature Communications*, *14*(1), 6336. https://doi.org/10.1038/s41467-023-42087-8

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., Hoshi, K., & Benítez-Burraco, A. (2022). A. Subcortical syntax: Reconsidering the neural dynamics of language. *Journal of Neurolinguistics*, *62*, 101062. https://doi.org/10.1016/j.jneuroling.2022.101062

Murphy, E., & Woolnough, O. (2024). The language network is topographically diverse and driven by rapid syntactic inferences. *Nature Reviews Neuroscience*, *25*(10), 705. https://doi.org/10.1038/s41583-024-00852-8

Murphy, E., Woolnough, O., Rollo, P. S., Roccaforte, Z., Segaert, K., Hagoort, P., & Tandon, N. (2022). Minimal phrase composition revealed by intracranial recordings. *The Journal of Neuroscience*, *42*(15), 3216–3227. https://doi.org/10.1523/JNEUROSCI.1575-21.2022

Shi, E. R. (2025). How the brain recycled memory circuits for language: An evolutionary perspective on the ROSE model. *Cognitive Neuroscience*, *16*(1–4), 88–90. https://doi.org/10.1080/17588928.2025.2561587

Slaats, S., & Hervais-Adelman, A. (2026). Linguistic structure and probability are jointly encoded in high gamma power. bioRxiv. https://doi.org/10.64898/2026.07.21.739808

Starr, T., Picton, L., & Thornton, J. (2017). Alternative evolutionary histories in the sequence space of an ancient protein. *Nature*, *549*(7672), 409–413. https://doi.org/10.1038/nature23902

Tafazoli, S., Bouchacourt, F. M., Ardalan, A., Markov, N. T., Uchimura, M., Mattar, M. G., Daw, N. D., & Buschman, T. J. (2026). Building compositional tasks with shared neural subspaces. *Nature*, *650*(), 164–172. https://doi.org/10.1038/s41586-025-09805-2

Tian, Z., Chen, J., Zhang, C., Min, B., Xu, B., & Wang, L. (2024). Mental programming of spatial sequences in working memory in the macaque frontal cortex. *Science*, *385*(), eadp6091. https://doi.org/10.1126/science.adp6091

Trettenbrein, P. C. (2025). Grounding the computational principles of language in neurobiology requires cross-modal and cross-linguistic data. *Cognitive Neuroscience*, *16*(1–4), 85–87. https://doi.org/10.1080/17588928.2025.2561581

Trettenbrein, P. C., Meister, N.-K., Slivac, K., Finkbeiner, T. A., Steinbach, M., Friederici, A. D., & Zaccarella, E. (2024). Modality-independent core brain network for language as proved by sign language. PsyArXiv. https://doi.org/10.31234/osf.io/hnzgp

Trettenbrein, P. C., Papitto, G., Friederici, A. D., & Zaccarella, E. (2021). Functional neuroanatomy of language without speech: An ALE meta-analysis of sign language. *Human Brain Mapping*, *42*(3), 699–712. https://doi.org/10.1002/hbm.25254

Uriagereka, J. (2025). Craving the ROSE and grasping the thorn. *Cognitive Neuroscience*, *16*(1–4), 91–92. https://doi.org/10.1080/17588928.2025.2561597

van Bree, S. (2024). Why the neural ingredients for a language of thought are not like spatial cells (commentary on Kazanina & poeppel, 2023). *The European Journal of Neuroscience*, *59*(10), 2552–2555. https://doi.org/10.1111/ejn.16329

Venhuizen, N. J., & Brouwer, H. (2025). Two models of meaning: Revisiting the principle of compositionality from the neurocognition of language. *Psychology of Learning and Motivation*, *83*, 269–305. https://doi.org/10.1016/bs.plm.2025.07.007

Woolnough, O., Donos, C., Murphy, E., Rollo, P. S., Roccaforte, Z. J., Dehaene, S., & Tandon, N. (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

Woolnough, O., Murphy, E., Dehaene, S., & Tandon, N. (2026). Low-frequency modulations bridge the language and default mode networks. *Progress in Neurobiology*, *264*, 102944. https://doi.org/10.1016/j.pneurobio.2026.102944