Current Directions in Psychological Science
2022, Vol. 31(2) 124–130

DOI: 10.1177/09637214211049233
www.psychologicalscience.org/CDPS

The Structure of Systematicity in the Brain

Randall C. O'Reilly[1], Charan Ranganath[2,3], and Jacob L. Russin[2,3]
[1]Department of Computer Science, [2]Department of Psychology, and [3]Center for Neuroscience, University of California, Davis

Abstract

A hallmark of human intelligence is the ability to adapt to new situations by applying learned rules to new content (systematicity) and thereby enabling an open-ended number of inferences and actions (generativity). Here, we propose that the human brain accomplishes these feats through pathways in the parietal cortex that encode the abstract structure of space, events, and tasks and pathways in the temporal cortex that encode information about specific people, places, and things (content). Recent neural network models show how the separation of structure and content might emerge through a combination of architectural biases and learning, and these networks show dramatic improvements over previous models in the ability to capture systematic, generative behavior. We close by considering how the hippocampal formation may form integrative memories that enable rapid learning of new structure and content representations.

Keywords

brain networks, computational models, deep neural networks, knowledge transfer, variable binding

The ability to directly transfer existing knowledge to new situations by applying learned rules to new content (*systematicity*) and thereby generate an open-ended number of different behaviors (*generativity*) is particularly advanced in humans relative to other animals. However, this ability is nevertheless shared in a limited but instructive way with even relatively simple computer programs. For example, it is easy to write a function that takes two arbitrary strings as arguments and systematically returns the concatenation of those two strings. With just a bit more logic, calling such a function repeatedly can easily construct an open-ended number of sentences. However, learning to do something like this in a neurally plausible manner is much more challenging, and neural network models continue to be criticized as lacking in these signature human abilities; much of their recent success is potentially attributable to something closer to rote memorization of increasingly large data sets (Fodor & Pylyshyn, 1988; Lake & Baroni, 2017; Marcus, 2018; O'Reilly et al., 2014; Plaut et al., 1996). Likewise, empirical cognitive neuroscience is only beginning to uncover the brain areas involved in these abilities (Constantinescu et al., 2016; Frankland & Greene, 2020; Park et al., 2020; Summerfield et al., 2020).

The essential trick employed by the simple concatenation function is that arbitrary content can be routed into it and operated upon generically, independently of any details of the content, through the ability to bind variables to any content. The cognitive equivalent of this, which is widely recognized as critical for human systematic behavior (Fodor & Pylyshyn, 1988), is a separation of *structure* (i.e., the rule-processing system, which is akin to the function) from *content*. In language, syntax has traditionally been thought of as an example of content-independent structure, and people's ability to accurately judge grammatical correctness of language with arbitrary content is demonstrated by Chomsky's (1957) famous example: "Colorless green ideas sleep furiously" (p. 15). Understanding how something like content-independent syntactic structure can be learned and represented via known biological mechanisms is thus a critical step toward advancing understanding of the neural basis of signature human abilities.

Corresponding Author:
Randall C. O'Reilly, Department of Computer Science, University of California, Davis
Email: oreilly@ucdavis.edu

In this article, we review cognitive and neural evidence consistent with the separation between structure and content across various cognitive domains and recent neural network models that demonstrate how such a separation might emerge through a combination of architectural biases and learning, thereby producing significantly greater systematicity. Critically, structure and content also need to interact and be integrated in various ways, and certain brain areas appear to be specialized for this integration, enabling the brain to apply multiple representational strategies in parallel.

Structure and Content in Cognitive Models

The FINST (fingers of instantiation) framework of Pylyshyn (1989) provides an early, simple model for how structure can be represented independently of specific content and also maps well onto cognitive-neuroscience data reviewed in the next section. The key idea is that abstract, content-independent, pointer-like "fingers" can index a small set (up to about 4) of different visual locations at a time. These content-independent pointers were originally proposed to serve as placeholders for encoding relationships among items in a scene (e.g., `INSIDE(a,b)`, as a propositional encoding of object a being inside b).

The core ideas in the FINST framework can be extended to represent any kind of abstract structure that specifies the relationships between different indexed elements. For example, the action of "giving" involves distinct functional roles (giver, recipient, item), and each of these could be represented using a separate FINST-like thematic role pointer, instead of using the semantic content specific to a particular situation (e.g., "Ashley gave Alan a gift"). If "giving" is encoded in terms of content-independent roles indexed by FINST-like pointers, these representations can automatically generalize to new contents, and additional inferences could be made about the properties and implications of the structural relationship (e.g., the giver no longer has the item and may have some expectation of reciprocity depending on the nature of the transaction and relationship). Encoding structural knowledge sufficiently rich for such inferences requires much more than FINST-like representations, but additional mechanisms to encode this knowledge in a relatively content-independent manner would clearly support systematic reasoning.

In short, FINST-like indexes provide a plausible attention-based neural mechanism for role-filler variable slots in the context of classical symbolic representational frameworks. At a cognitive level, the kinds of elaborated structural representations and processes that have been studied cover a range of different levels of complexity, from verb-based (action-based) elements, as in the "giving" example (Boylan et al., 2015), to more elaborated *schemas* or *scripts* describing longer sequences of events (e.g., the schema associated with a child's birthday party; see Binding of Structure and Content in the Medial Temporal Lobes). The ability to transfer structural knowledge across content domains is also central to established models of analogical reasoning (Gentner & Holyoak, 1997; e.g., the relational structure of "orbiting" can be transferred from the planetary to the atomic scale).

Representation of Structure in the Human Brain

There is a well-established distinction between spatial and object processing in the human brain, and this distinction can be reframed as one example of how the brain separates structure and content via distinct, but interacting, pathways. Visual (and auditory) networks in the brain route sensory input into distinct dorsal- and ventral-stream pathways (Ungerleider & Mishkin, 1982). The ventral visual pathway extends from early visual cortex to inferotemporal cortex and is characterized as the "what" pathway, specialized for object or scene recognition (i.e., visual content). The dorsal pathway through the parietal lobe is specialized for spatial "where" processing, as indicated by extensive evidence that this pathway represents spatial and relational information in a relatively content-independent manner.

The idea that the dorsal-stream pathway has the potential to support systematic structure-sensitive processing was already well developed by Pylyshyn (1989) and has been incorporated into psycholinguistic theories (Frankland & Greene, 2020; Landau & Jackendoff, 1993). Furthermore, there is substantial evidence that posterior parietal cortex encodes nonspatial structural information as well. First, it is clear that the parietal lobe plays a critical role in the sensory guidance of action performance (Orban et al., 2004), as captured in the "what"-"how"(ventral-dorsal) framework (Goodale & Milner, 1992). In humans, this action coding extends to the representation of verb-based argument structure (as in the "giving" example) in the inferior parietal area (angular gyrus) and supports the core of structure at the sentence level (Binder & Desai, 2011; Boylan et al., 2015; Palmer et al., 2016). More abstract structural and linguistic concepts, including event representations in higher-order parietal areas, are thought to build directly upon parietal action, space, and time foundations (Frankland & Greene, 2020; Landau & Jackendoff, 1993).

It is well established that parietal spatial and action representations anticipate the effects of eye, head, or body movements (Cavanagh et al., 2010). This suggests that acquisition of structure in the parietal lobe could

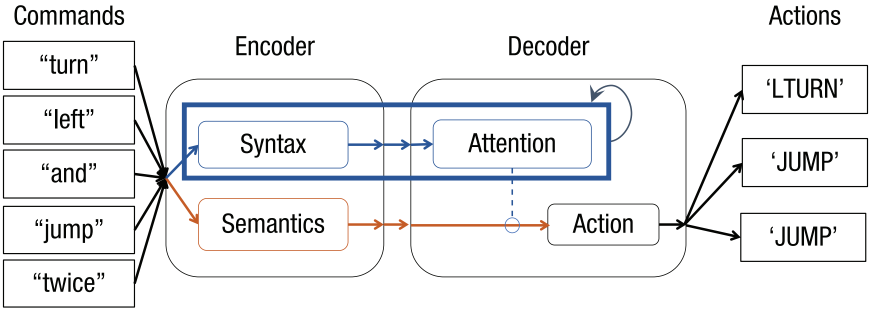

Fig. 1. Diagram depicting a neural network model (Russin et al., 2020) with separate pathways for learning structure (syntax) and content (semantics). The input command on the left ("turn left and jump twice") was mapped to a sequence of corresponding actions (LTURN, JUMP, JUMP), using a modeling framework developed originally for translating between languages (Lake & Baroni, 2017). The structure pathway (in blue) had full access to the temporal ordering of the words in the command. However, it could influence the output actions taken at each time step only through an attention mechanism (indicated by the blue dashed line and open circle). The character of the representations in these different pathways facilitated systematic generalization, but was not built in, and emerged as a result of the architectural constraints imposed by the separation of the pathways.

be based on *predictive learning* (O'Reilly et al., 2021). Specifically, learning driven by the difference between a predicted sensory outcome of an action and the subsequent sensory input can drive improved predictions and shape the formation of more abstract structural representations to more efficiently generate these predictions. Although it is fairly straightforward to learn to predict specific sensory outcomes from motor actions, an important area of current research is to determine the extent to which more abstract, structural representations can emerge, to capture the consistent, generalizable relationships that hold across a large number of such actions (Summerfield et al., 2020).

Even if the dorsal and ventral streams separately encode structure and content in the brain, fundamental questions remain about how these streams interact at the level of detailed neural mechanisms, to support systematic cognitive function. A recent neural network model provides a useful example of how FINST-like attentional pointers can operate on newly learned content information in the context of separate structure and content representations (Russin et al., 2020). Unlike many existing models that have relied on various hand-coded mechanisms to directly emulate programming-language-like variable-binding functionality in neural hardware, this model learned entirely via error-driven learning, with only very broad, biologically plausible architectural constraints between two processing pathways.

One pathway had full access to the temporal ordering of words within a sentence, whereas the other was able to process only the single word that was current at any given time (Fig. 1). Furthermore, the pathway that was sensitive to temporal order could influence the network output only via attentional modulation of the other pathway (thus functioning similarly to the FINST-like attentional pointers). Distinctions between pathways such as these could plausibly derive from evolved differences in the initial wiring of the neural architecture, providing affordances upon which subsequent learning operates.

With these constraints in place, the network learned syntax-like structural representations in the temporal-order pathway and semantics-like content information in the other pathway, and in so doing, exhibited significantly more systematic behavior on challenging out-of-domain generalization tests than unstructured models did (Lake & Baroni, 2017). Specifically, when the model was trained on proper action outputs for examples generated from a simple phrase structure grammar (e.g., "turn left twice"), the syntax-like pathway was able to direct attention to the proper sequence of content items in new commands (e.g., "jump twice"). This attentional indexing of the novel content items is consistent with the simple FINST framework described above and with the central importance of such mechanisms in the neural basis of structure-content dynamics. Furthermore, this model demonstrates how learned, distributed neural representations can take on a syntax-like role without traditional explicit symbolic grammar trees. However, this model and other related ones are simplified, and more work is needed to understand the nature of structure representations in the brain that support the more complex and sophisticated forms of human cognition.

A similar distinction between structure and content pathways was achieved in a fairly different way using a variable-binding system based on an explicit slot-filler (key-value) lookup table mechanism (Webb et al., 2021). In this model, the structure-processing pathway learned to control this lookup table independently of the specific content stored there, while the content pathway learned representations of individual items. This architectural separation of structure and content enabled the model to exhibit much more systematic behavior than comparison models without such a separation. (See Akyürek & Andreas, 2021, for discussion of other related architectures.)

Although these models provide important initial demonstrations of how a separation between structure and content can support systematicity, it is critical to appreciate that many real-world situations also require strong interactions between the two. This has perhaps been most extensively studied in the case of the pronunciation of English words: As everyone learning English as a second language knows too well, the number of exceptions and subregularities is mind-spinning, and monolithic neural network pathways (i.e., pathways without separation of structure and content) may provide the most appropriate mechanism for learning these mappings (Plaut et al., 1996). Likewise, the difference between parsing "Fruit flies like a banana" and "Time flies like an arrow" requires an interaction between syntax and semantics. In the visual domain, these interactions can be illustrated by the case of viewing a kitchen scene: Incoming content information may activate structural representations of typical kitchens, which then guide visual attention and semantic processing (Hayes & Henderson, 2019). In other words, human cognition likely represents a combination of separated and interacting streams (Franklin et al., 2020; O'Reilly et al., 2014).

Structure and Content in Episodic Memory

So far, we have used the classic model of the dorsal and ventral streams as one example of the separation of structure and content, but structure exists across multiple scales. For example, structural knowledge is required to parse and comprehend individual sentences, but comprehending a sentence within the context of a larger text passage often requires a mental representation of the structure of events depicted in that passage. Recent functional MRI (fMRI) research has uncovered a candidate set of brain regions that may represent the abstract structure of events—the *posterior medial* (PM) network—and regions that may represent the characteristics of the specific people, places, and things that serve as event content—the *anterior temporal* (AT) network (Fig. 2). The PM and AT networks can be considered as anatomically higher up from the classic dorsal and ventral streams, respectively (Kravitz et al., 2011; Ranganath & Ritchey, 2012). Although these networks are distinct, the hippocampus is in a key position to integrate information across them. Regions in the AT network project predominantly to the lateral entorhinal cortex, and regions in the PM network mostly project to the medial entorhinal cortex (MEC). The hippocampus proper then binds information from these streams together into unitary episodic memories (Ranganath & Ritchey, 2012). Considerable evidence suggests that the PM and AT networks differentially represent information about structure and content at higher, more complex levels and over longer spatiotemporal scales than the more posterior dorsal and ventral streams. For example, the PM network encodes the structure of events and spatial contexts and participates in the reconstruction of past events and the simulation of future events (Ranganath & Ritchey, 2012), and it is also central for discourse comprehension (Martín-Loeches et al., 2008).

One can think of the hippocampus, PM network, and AT network as complementary learning systems (Russin et al., 2021). The PM and AT networks are specialized for slow, integrative learning leading to highly generalizable representations and complement the rapid context-specific memory supported by the hippocampus. By capturing states of PM and AT network regions at critical moments in time, the hippocampus is in a unique position to rapidly form memories that combine information about structure and content that is encountered during a specific period of time (i.e., an *episode*). When a similar item or situation is encountered later on, hippocampal pattern completion can lead to reinstatement of previous states in the PM and AT networks, thereby facilitating the generation of a mental model of the current situation in a new or ambiguous context (Franklin et al., 2020).

Other computational models (Stachenfeld et al., 2017; Whittington et al., 2020) embody principles that are largely consistent with the ones proposed here but have placed more emphasis on the role of the hexagonally tiled *grid cells* in the MEC, which are thought to represent the topological structure of spatial environments, and hippocampal *place cells*, which are thought to encode specific locations within an environment. These models suggest that the MEC and hippocampus can represent any kind of task in a 2D state-space and are consistent with fMRI studies showing MEC activity patterns characteristic of grid cells during the performance of novel tasks (Constantinescu et al., 2016). Notably, however, these and other studies of humans suggest that grid cells are present and grid cell–like fMRI activity patterns occur throughout the PM network and in the medial prefrontal cortex (Jacobs et al., 2013).

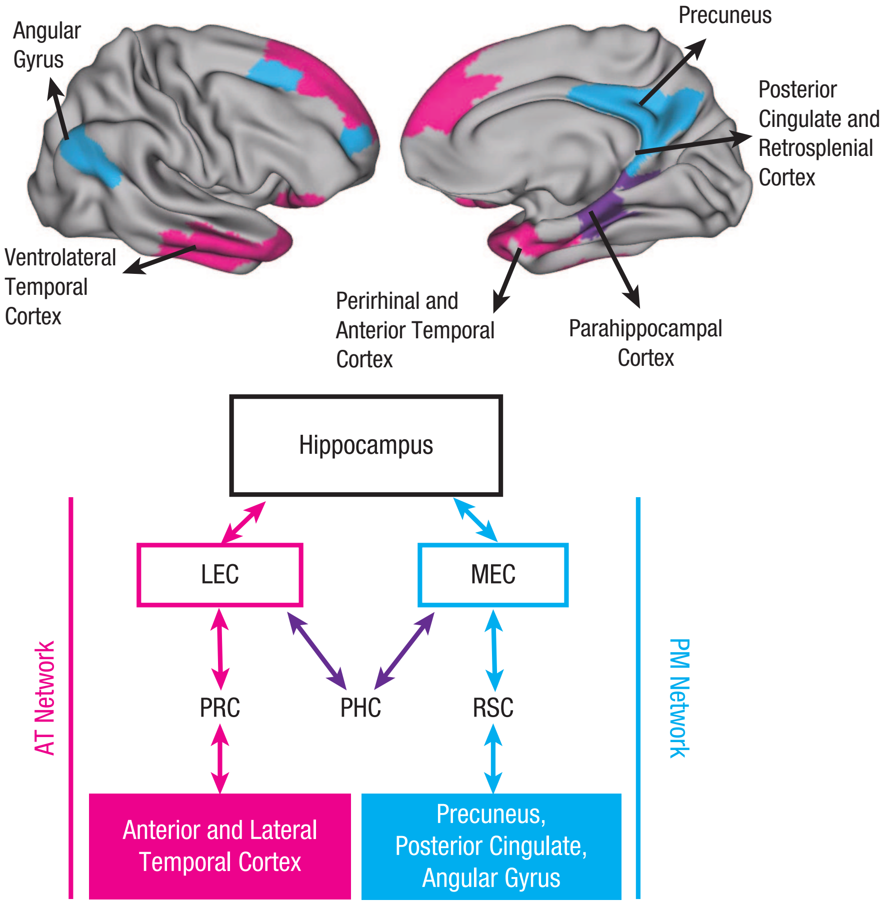

Fig. 2. Two cortico-hippocampal networks. The surface renderings at the top depict the neocortical regions in the posterior medial (PM) network (cyan), implicated in schema representation, and the anterior temporal (AT) network (magenta), implicated in representations of people and things. The diagram at the bottom schematically depicts the pathways by which information from the AT and PM networks is directed to parallel pathways in the medial temporal lobes before fully converging in the hippocampus. PRC = perirhinal cortex; PHC = parahippocampal cortex; RSC = retrosplenial cortex; LEC = lateral entorhinal cortex; MEC = medial entorhinal cortex.

Further work is needed to disentangle the unique contributions of the MEC, relative to earlier parietal-lobe pathways, in structure representation. One possibility is that the use of structural knowledge to guide flexible behavior is driven by the PM network, and that the MEC plays an important role in the use of episodic memory to support rapid structure learning by the PM network. For instance, it is possible that the MEC's network of grid cells performs an initial step of *pattern separation* on parietal inputs from the PM network, moving them apart in representational space and thus making them easier to learn without interference (Frankland & Cohen, 2022). This idea is consistent with results from lesion studies suggesting that hippocampal or MEC lesions can significantly impair spatial learning, while generally sparing navigation in familiar environments (Hales et al., 2014; Kolarik et al., 2016).

Conclusions

There is good reason to believe that the architecture of the neocortex is optimized to support systematic and generative behavior by bifurcating sensory information into separate structure and content pathways along the dorsal and ventral streams, respectively. The hierarchically organized pathways through the parietal lobe generate different levels of representations that encode structural relationships reliably associated with different actions and events. Such representations capture the general "logic" of actions and events—how things move, transform, and interact over space and time—in a way that can be readily applied to novel content. Significant work remains to be done in order to establish the nature of representations of structure. By understanding the evolution of learning and dynamics across interacting brain networks, researchers can make significant progress toward understanding the nature of the representations that give rise to uniquely human intelligence.

Recommended Reading

Kravitz, D. J., Saleem, K. S., Baker, C. I., & Mishkin, M. (2011). (See References). An integrated review of three parietal lobe pathways supporting looking, reaching, and navigating.

Pylyshyn, Z. (1989). (See References). An original proposal for a mechanism that separates structure from content on the basis of attention, which we think may still have currency in neural network models of the brain.

Ranganath, C., & Ritchey, M. (2012). (See References). A review of evidence from humans and animal models supporting the idea that the anterior temporal (AT) network represents information about specific entities and the posterior medial (PM) network represents information about event structure that is used to generate mental models of particular situations.

Summerfield, C., Luyckx, F., & Sheahan, H. (2020). (See References). A recent synthetic review of a wide range of data and theory supporting the idea that the parietal lobe represents content-independent structure.

Whittington, J. C. R., Muller, T. H., Mark, S., Chen, G., Barry, C., Burgess, N., & Behrens, T. E. J. (2020). (See References). A presentation of a computational model implementing separation of structure from content in the hippocampal system, with nice illustrations of different forms of structure and how they can support systematic behavior in different domains.

Transparency

Action Editor: Robert L. Goldstone
Editor: Robert L. Goldstone
Declaration of Conflicting Interests

R. C. O'Reilly is Chief Scientist at eCortex, Inc., and the Astera Obelisk lab, which may derive indirect benefit from the work presented here. The author(s) declared that there were no other potential conflicts of interest with respect to the authorship or the publication of this article.

Funding

This work was supported by Office of Naval Research Grants N00014-20-1-2578, N00014-19-1-2684/N00014-18-1-2116, N00014-18-C-2067, N00014-17-1-2961, and N00014-15-1-0033, awarded to R. C. O'Reilly and C. Ranganath.

ORCID iD

Randall C. O'Reilly https://orcid.org/0000-0003-0322-4600

Acknowledgments

We thank Jon Cohen, Alex Petrov, Erie Boorman, Ignacio Saez, Taylor Webb, Stephen Frankland, James Antony, John Rohrlich, Mary Zolfaghar, and members of the Computational Cognitive Neuroscience Lab, Dynamic Memory Lab, and Cohen lab for discussion and comments.

References

Akyürek, E., & Andreas, J. (2021). *Lexicon learning for few-shot neural sequence modeling*. arXiv. http://arxiv.org/abs/2106.03993

Binder, J. R., & Desai, R. H. (2011). The neurobiology of semantic memory. *Trends in Cognitive Sciences*, *15*(11), 527–536. https://doi.org/10.1016/j.tics.2011.10.001

Boylan, C., Trueswell, J. C., & Thompson-Schill, S. L. (2015). Compositionality and the angular gyrus: A multi-voxel similarity analysis of the semantic composition of nouns and verbs. *Neuropsychologia*, *78*, 130–141. https://doi.org/10.1016/j.neuropsychologia.2015.10.007

Cavanagh, P., Hunt, A. R., Afraz, A., & Rolfs, M. (2010). Visual stability based on remapping of attention pointers. *Trends in Cognitive Sciences*, *14*(4), 147–153. https://doi.org/10.1016/j.tics.2010.01.007

Chomsky, N. (1957). *Syntactic structures*. Mouton.

Constantinescu, A. O., O'Reilly, J. X., & Behrens, T. E. J. (2016). Organizing conceptual knowledge in humans with a gridlike code. *Science*, *352*(6292), 1464–1468. https://doi.org/10.1126/science.aaf0941

Fodor, J. A., & Pylyshyn, Z. W. (1988). Connectionism and cognitive architecture: A critical analysis. *Cognition*, *28*(1–2), 3–71. https://doi.org/10.1016/0010-0277(88)90031-5

Frankland, S. M., & Cohen, J. D. (2022). *Pattern separation in the entorhinal cortex* [Manuscript in preparation]. Department of Neuroscience, Princeton University.

Frankland, S. M., & Greene, J. D. (2020). Concepts and compositionality: In search of the brain's language of thought. *Annual Review of Psychology*, *71*, 273–303. https://doi.org/10.1146/annurev-psych-122216-011829

Franklin, N. T., Norman, K. A., Ranganath, C., Zacks, J. M., & Gershman, S. J. (2020). Structured Event Memory: A neuro-symbolic model of event cognition. *Psychological Review*, *127*(3), 327–361. https://doi.org/10.1037/rev0000177

Gentner, D., & Holyoak, K. J. (1997). Reasoning and learning by analogy: Introduction. *American Psychologist*, *52*(1), 32–34. https://doi.org/10.1037/0003-066X.52.1.32

Goodale, M. A., & Milner, A. D. (1992). Separate visual pathways for perception and action. *Trends in Neurosciences*, *15*(1), 20–25. https://doi.org/10.1016/0166-2236(92)90344-8

Hales, J. B., Schlesiger, M. I., Leutgeb, J. K., Squire, L. R., Leutgeb, S., & Clark, R. E. (2014). Medial entorhinal cortex lesions only partially disrupt hippocampal place cells and hippocampus-dependent place memory. *Cell Reports*, *9*(3), 893–901. https://doi.org/10.1016/j.celrep.2014.10.009

Hayes, T. R., & Henderson, J. M. (2019). Scene semantics involuntarily guide attention during visual search. *Psychonomic Bulletin & Review*, *26*(5), 1683–1689. https://doi.org/10.3758/s13423-019-01642-5

Jacobs, J., Weidemann, C. T., Miller, J. F., Solway, A., Burke, J. F., Wei, X.-X., Suthana, N., Sperling, M. R., Sharan, A. D., Fried, I., & Kahana, M. J. (2013). Direct recordings of grid-like neuronal activity in human spatial navigation. *Nature Neuroscience*, *16*(9), 1188–1190. https://doi.org/10.1038/nn.3466

Kolarik, B. S., Shahlaie, K., Hassan, A., Borders, A. A., Kaufman, K. C., Gurkoff, G., Yonelinas, A. P., & Ekstrom, A. D. (2016). Impairments in precision, rather than spatial strategy, characterize performance on the virtual Morris Water Maze: A case study. *Neuropsychologia*, *80*, 90–101. https://doi.org/10.1016/j.neuropsychologia.2015.11.013

Kravitz, D. J., Saleem, K. S., Baker, C. I., & Mishkin, M. (2011). A new neural framework for visuospatial processing. *Nature Reviews Neuroscience*, *12*(4), 217–230. https://doi.org/10.1038/nrn3008

Lake, B. M., & Baroni, M. (2017). *Generalization without systematicity: On the compositional skills of sequence-to-sequence recurrent networks*. arXiv. https://arxiv.org/abs/1711.00350v1

Landau, B., & Jackendoff, R. (1993). "What" and "where" in spatial language and spatial cognition. *Behavioral & Brain Sciences*, *16*(2), 217–238. https://doi.org/10.1017/S0140525X00029733

Marcus, G. (2018). *Deep learning: A critical appraisal*. arXiv. http://arxiv.org/abs/1801.00631

Martín-Loeches, M., Casado, P., Hernández-Tamames, J. A., & Álvarez-Linera, J. (2008). Brain activation in discourse comprehension: A 3t fMRI study. *NeuroImage*, *41*(2), 614–622. https://doi.org/10.1016/ j.neuroimage.2008.02.047

Orban, G. A., Van Essen, D., & Vanduffel, W. (2004). Comparative mapping of higher visual areas in monkeys and humans. *Trends in Cognitive Sciences*, *8*(7), 315–324. https://doi.org/10.1016/j.tics.2004.05.009

O'Reilly, R. C., Petrov, A. A., Cohen, J. D., Lebiere, C. J., Herd, S. A., & Kriete, T. (2014). How limited systematicity emerges: A computational cognitive neuroscience approach. In I. P. Calvo & J. Symons (Eds.), *The architecture of cognition: Rethinking Fodor and Pylyshyn's systematicity challenge* (pp. 191–226). MIT Press.

O'Reilly, R. C., Russin, J. L., Zolfaghar, M., & Rohrlich, J. (2021). Deep predictive learning in neocortex and pulvinar. *Journal of Cognitive Neuroscience*, *33*(6), 1158–1196. https://doi.org/10.1162/jocn_a_01708

Palmer, M., Bonial, C., & Hwang, J. D. (2016). VerbNet: Capturing English verb behavior, meaning and usage. In S. E. F. Chipman (Ed.), *The Oxford handbook of cognitive science* (pp. 315–336). Oxford University Press.

Park, S. A., Miller, D. S., Nili, H., Ranganath, C., & Boorman, E. D. (2020). Map making: Constructing, combining, and inferring on abstract cognitive maps. *Neuron*, *107*(6), 1226–1238. https://doi.org/10.1016/j.neuron.2020.06.030

Plaut, D. C., McClelland, J. L., Seidenberg, M. S., & Patterson, K. (1996). Understanding normal and impaired word reading: Computational principles in quasi-regular domains. *Psychological Review*, *103*(1), 56–115. https://doi.org/10.1037/0033-295X.103.1.56

Pylyshyn, Z. (1989). The role of location indexes in spatial perception: A sketch of the FINST spatial-index model. *Cognition*, *32*(1), 65–97. https://doi.org/10.1016/0010-0277(89)90014-0

Ranganath, C., & Ritchey, M. (2012). Two cortical systems for memory-guided behaviour. *Nature Reviews Neuroscience*, *13*(10), 713–726. https://doi.org/10.1038/nrn3338

Russin, J., Jo, J., O'Reilly, R. C., & Bengio, Y. (2020). Systematicity in a recurrent neural network by factorizing syntax and semantics. In S. Denison, M. Mack, Y. Xu, & B. C. Armstrong (Eds.), *Proceedings for the 42nd Annual Meeting of the Cognitive Science Society* (pp. 109–115). Cognitive Science Society.

Russin, J., Zolfaghar, M., Park, S. A., Boorman, E., & O'Reilly, R. C. (2021). Complementary structure-learning neural networks for relational reasoning. In T. Fitch, C. Lamm, H. Leder, & K. Teßmar-Raible (Eds.), *Proceedings for the 43nd Annual Meeting of the Cognitive Science Society* (pp. 1560–1566). Cognitive Science Society.

Stachenfeld, K. L., Botvinick, M. M., & Gershman, S. J. (2017). The hippocampus as a predictive map. *Nature Neuroscience*, *20*(11), 1643–1653. https://doi.org/10.1038/nn.4650

Summerfield, C., Luyckx, F., & Sheahan, H. (2020). Structure learning and the posterior parietal cortex. *Progress in Neurobiology*, *184*, Article 101717. https://doi.org/10.1016/j.pneurobio.2019.101717

Ungerleider, L. G., & Mishkin, M. (1982). Two cortical visual systems. In D. J. Ingle, M. A. Goodale, & R. J. W. Mansfield (Eds.), *The analysis of visual behavior* (pp. 549–586). MIT Press.

Webb, T. W., Sinha, I., & Cohen, J. D. (2021). *Emergent symbols through binding in external memory*. arXiv. http://arxiv.org/abs/2012.14601

Whittington, J. C. R., Muller, T. H., Mark, S., Chen, G., Barry, C., Burgess, N., & Behrens, T. E. J. (2020). The Tolman-Eichenbaum machine: Unifying space and relational memory through generalization in the hippocampal formation. *Cell*, *183*(5), 1249–1263. https://doi.org/10.1016/j.cell.2020.10.024