

Review



Cite this article: Calmus R, Wilson B, Kikuchi Y, Petkov CI. 2019 Structured sequence processing and combinatorial binding: neurobiologically and computationally informed hypotheses. *Phil. Trans. R. Soc. B* **375**: 20190304.
<http://dx.doi.org/10.1098/rstb.2019.0304>

Accepted: 4 November 2019

One contribution of 16 to a theme issue 'Towards mechanistic models of meaning composition'.

Subject Areas:

neuroscience, cognition, computational biology

Keywords:

computational modelling, sequence learning, serial order, chunking, language, binding

Authors for correspondence:

Ryan Calmus

e-mail: r.m.calmus@ncl.ac.uk

Christopher I. Petkov

e-mail: chris.petkov@ncl.ac.uk

Structured sequence processing and combinatorial binding: neurobiologically and computationally informed hypotheses

Ryan Calmus, Benjamin Wilson, Yukiko Kikuchi and Christopher I. Petkov

Newcastle University Medical School, Framlington Place, Newcastle upon Tyne, UK

id RC, 0000-0002-7826-9355; BW, 0000-0003-2043-7771; YK, 0000-0002-0365-0397; CIP, 0000-0002-4932-7907

Understanding how the brain forms representations of structured information distributed in time is a challenging endeavour for the neuroscientific community, requiring computationally and neurobiologically informed approaches. The neural mechanisms for segmenting continuous streams of sensory input and establishing representations of dependencies remain largely unknown, as do the transformations and computations occurring between the brain regions involved in these aspects of sequence processing. We propose a blueprint for a neurobiologically informed and informing computational model of sequence processing (entitled: Vector-symbolic Sequencing of Binding INstantiating Dependencies, or VS-BIND). This model is designed to support the transformation of serially ordered elements in sensory sequences into structured representations of bound dependencies, readily operates on multiple timescales, and encodes or decodes sequences with respect to chunked items wherever dependencies occur in time. The model integrates established vector symbolic additive and conjunctive binding operators with neurobiologically plausible oscillatory dynamics, and is compatible with modern spiking neural network simulation methods. We show that the model is capable of simulating previous findings from structured sequence processing tasks that engage fronto-temporal regions, specifying mechanistic roles for regions such as prefrontal areas 44/45 and the frontal operculum during interactions with sensory representations in temporal cortex. Finally, we are able to make predictions based on the configuration of the model alone that underscore the importance of serial position information, which requires input from time-sensitive cells, known to reside in the hippocampus and dorsolateral prefrontal cortex.

This article is part of the theme issue 'Towards mechanistic models of meaning composition'.

1. Introduction

Natural environments are richly structured in both space and time. Substantial progress has been made in understanding the neurobiological bases of learned relationships between spatially or temporally separated elements [1–3]. Moreover, prior research has established the importance of serial order for the brain [4], and *binding problems*, whereby distinct sensory events are combined for perception, decision and action [5], have attracted considerable interest and empirical enquiry [6,7].

Establishing relationships, or *dependencies*, between elements over time allows us to extract the structure of the sensory world and to make predictions about future events. However, understanding how the brain binds complex information distributed in time, building temporally organized structures that represent multiple linked dependencies, remains a considerable challenge facing the neuroscientific community. This sort of cognitive structure-building is challenging for the brain to achieve because complex input must be discretized in time, resultant discrete items chunked and stored in memory, dependencies identified and

related items bound for perception (or other purposes), and representations of multiple dependencies maintained concurrently in memory to be further manipulated [8,9].

Human language—written, spoken or signed—is a salient example of the complexity of the binding problem, because it features syntactically organized dependencies between semantic units [10]. Yet the problem of building complex representations is also relevant to complex action sequences [8,11,12], music [8,13,14], mathematics [9,15] and cognition in general [11,13]. Moreover, some of these systems are not unique to humans, since songbirds can construct complex vocalization sequences [16], an ability supported by a forebrain neural system [17], and correspondences have been established between humans and a number of species in processing adjacent and non-adjacent sequencing dependencies [18–20]. Thus, advancing our understanding of how complex structure can be built from sequential input in computational and neural systems is important for developing better machine and animal models to understand both general principles and species-specific aspects of combinatorial binding.

In this article, we propose a blueprint for a neurobiologically informed and informing computational model of sequence processing (entitled: Vector-symbolic Sequencing of Binding INSTANTIATING Dependencies, or VS-BIND). The VS-BIND approach integrates: (1) advances in modelling combinatorial binding within simulated neural systems using vector symbolic operations; (2) insights from neuroimaging and neurophysiological evidence in human and non-human primates on neural correlates of structured sequence processing and working memory; and (3) dynamic mechanisms for manipulating population codes that can be incorporated in modern spiking neural networks [21,22]. The approach allows us to plausibly transform internal representations, rendering these into both mathematically idealized and neurally simulated site-specific activity unfolding over time. Building on these foundational mechanisms, we focus on modelling chunk encoding and the binding of sensory items to represent adjacent, non-adjacent and more complex (hierarchically) structured sequencing dependencies. Our key objectives here are to motivate the approach, ground it in the relevant literature, and use it to generate distinct mechanistic predictions, the form of which, as will be seen below, depends on both the specific binding operations used and their configuration.

It is important to note that any model using combinatorial operators can only be described as *classically compositional* if combinatorial representations precisely reflect the meaning of all constituents and the relations between them. The operators used here are not classically compositional, but nevertheless serve an important purpose in allowing us to generate falsifiable predictions of neural mechanisms and correlates of structured sequence processing ripe for neurobiological testing across the species. To assist in this process, we also share MATLAB (MathWorks) code, including a demonstration (doi:10.5281/zenodo.3464607) [23]. Even on the basis of its structure alone and initial modelling, the VS-BIND approach posits a number of intriguing predictions.

2. Foundations of descriptive and computational models of structured sequence processing

Language relies on semantic and syntactic knowledge, supported by the detection of dependencies between phonemes,

morphemes, words and phrases in sentences. The language binding problem features the rapid detection of lexical symbols and the encoding of complex syntactic regularities at multiple scales and temporal granularities.

A large volume of neurobiological data implicate a fronto-temporal brain system in various aspects of language processing [11,24,25]. Neurobiological signals associated with the chunking and parsing of speech with respect to phrase boundaries have been identified in these regions [26–28]. Moreover, the temporal structure of speech or language content at different timescales (phonemic, syllabic, word or phrasal) produces stimulus- or context-driven neural entrainment at the relevant oscillatory frequency bands [29,30].

Behaviourally, a number of sequencing processes are now known to be evolutionarily conserved, including entrainment to rhythmic sensory input [29,31]. There is also information from artificial grammar learning paradigms, which are used to establish dependencies between otherwise arbitrary auditory or visual items in a sequence, either via statistical or rule learning [32,33]. Humans and a number of non-human animals can learn dependencies between sensory items next to each other in a sequence (*adjacent dependencies*), as well as dependencies further separated in time and by intervening items (*non-adjacent dependencies*; reviewed in [18]). The learning of hierarchically organized dependencies by non-human animals is, however, contentious and it remains to be seen whether this ability is uniquely human [34,35].

Comparative neuroimaging work has identified brain regions in human and monkey frontal and temporal cortex involved in processing sequencing dependencies [19,20]. This has led to descriptive models of the brain-bases of structured sequence processing and the relationship with, and distinctions from, neurobiological processes involved in language [34,36,37]. For the purposes of this paper, we will focus on the Wilson *et al.* [34] descriptive neurobiological model of human and non-human primate structured sequence processing, shown in figure 1. Relatively simple sequencing relationships, such as between items that occur next to each other in a sequence, can be learned by both humans and monkeys and are seen to engage corresponding brain regions, particularly the ventral frontal and opercular cortex (figure 1*a*, vFOC, bottom row) [34]. For non-adjacent dependencies (which increase working memory demands: store an item in memory long enough to link to its matched pair), it is less clear whether the frontal operculum, other inferior frontal areas (such as areas 44/45), and/or the dorso-lateral prefrontal cortex (DLPFC) are more involved. More complex (including hierarchical) dependencies engage inferior frontal areas 44/45 in humans (Broca's area; see figure 1*a*, middle and top rows) [38]. The hippocampus has also been implicated in structured sequence processing [39,40] and implicit learning [41], but its mechanistic role within these contexts remains incompletely understood.

Computational models of language or the processing of serial information provide compelling simulations of behavioural data [42,43]. However, modelling underlying neural mechanisms presents additional challenges. David Marr's tri-level framework [44] famously defines three levels of description that are still widely applied in characterizing any given model of the brain: the goals of the system (*computational level*), the cognitive processes required to reach this goal (*algorithmic level*), and the neural mechanisms required to instantiate them (*implementational level*). Poggio [45] extended

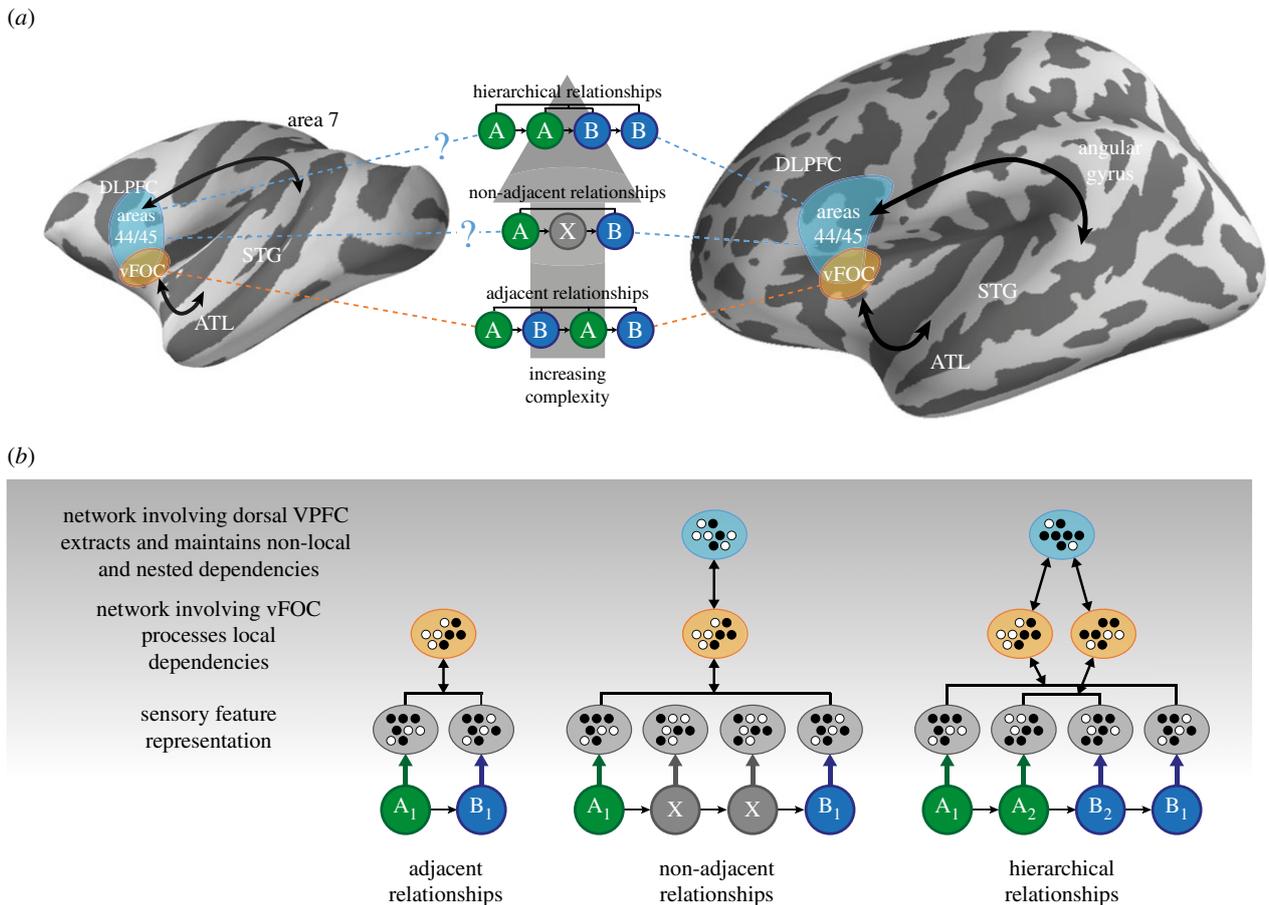


Figure 1. Neurobiologically informed heuristic model of structured sequence processing, by Wilson, Marslen-Wilson and Petkov. (a) Fronto-temporal regions involved in sequence processing, from [34]. DLPFC, dorsolateral prefrontal cortex; vFOC ventral frontal opercular cortex; ATL, anterior temporal lobe; STG, superior temporal gyrus. (Copyright © 2017 Benjamin Wilson, William D. Marslen-Wilson and Christopher I. Petkov, CC BY 4.0.) (b) Predicted combinatorial codes illustrated as neural patterns implemented by coordination between different regions. (Online version in colour.)

this framework by suggesting that models should also offer insights into learning processes and the evolutionary path that yields the system [46]. Although advances have been made in understanding many neuropsychological phenomena at individual levels of description, it remains desirable to advance understanding on multiple levels through holistic modelling approaches [47].

The fields of neuroscience and machine learning have been converging, in particular through the use of models incorporating functionally and anatomically distinct subpopulations of artificial neurons [48]. Artificial neural networks (ANNs), including deep and recurrent neural networks (DNNs and RNNs), are the dominant connectionist modelling paradigm in use today, using iterative training procedures to tune synaptic weights between artificial neurons and establish network-level computations. RNNs are relevant to sequence learning since recurrent feedback allows them to integrate information over time [49], while DNNs have revolutionized machine learning, increasingly inform neuroscientific analyses, and can generate neural correlates [50].

A likewise informative, but paradigmatically distinct approach is to model the brain at an algorithmic level, explaining behaviour and cognition in terms of computational processes that combine and transform cognitive symbols (see [51]). Reconciling neural and cognitive perspectives is a longstanding challenge, but there exists a computational modelling subfield that has made considerable strides in this direction: *symbolic connectionism*. This approach seeks to

produce ANN models with explicit support for combinatorial and symbolic operations (*neural-symbolic* networks). This is the subfield we look to in furthering our modelling aims, specifically using *vector symbolic architectures* (VSAs; see below), which can be used alone, in ANNs, or in dynamic spiking neural networks. Our use of VSAs has the benefit of generating predictions on neural mechanisms of combinatorial binding throughout the fronto-temporal system involved in structured sequence processing.

3. Computationally modelling structured representations in neural systems: an overview of approaches

There exist a number of symbolic connectionist solutions, and many non-symbolic or non-connectionist cognitive architectures, each addressing various aspects of neural binding. Here, we briefly overview symbolic connectionist approaches before justifying our use of VSAs. For brevity, we restrict discussion here only to approaches that specify ways to build structured representations while explicitly supporting neurobiologically plausible implementation.

Three distinct methods of modelling binding predominate, although they can in many ways be viewed as complementary [52]. The first uses coordinated temporal synchrony (see LISA [53,54] and SHRUTI [55]) or asynchrony (see DORA [56]) to

unify constituents. The second uses uniform grids of integrating circuits to create a stable memory for bindings (neural blackboard architectures, NBAs [57]). The third, encompassing VSAs, principally uses conjunctive spatial coding of abstract vector representations to associate items, via tensor products [58–60], circular convolution [61] or other defined transforms [62,63].

Although commonly contrasted, spatial and temporal binding mechanisms are not mutually exclusive. Conjunctive spatial coding and temporal synchrony/asynchrony are known to be complementary and can be thought of as different perspectives on the same dynamic process [64]. Furthermore, conjunctive coding is considered appropriate for long-term storage in temporally coordinating models [52,65], whilst vector-based methods can operate within dynamic frameworks likewise subject to temporal influences [22], which we will consider in more detail later. Thus, spatial and temporal binding approaches are not diametrically opposed but rather place different emphases and explanatory burdens on mutually informing aspects of neural coding. Although VSAs tend to be conceived of as static systems, we specifically advocate their use within a dynamic framework (for example, Nengo [21,22]) to incorporate advantages of both temporal and population coding.

The neural population codes employed by these models vary. In this regard, *localist* representations (which exhibit one-to-one or many-to-one mappings between features and neural activation) are typically contrasted with *distributed* representations (which exhibit many-to-many mappings). Distributed representations may be *dense* or *sparse*, which in the latter case means that a low proportion of neurons are active within a population at any one time. There is evidence that both localist [66,67] and sparse distributed representations are used by the brain [68–70], and there are advantages to both encoding strategies for neural systems. For instance, localist representations exhibit the lowest possible interference between encodings, whilst sparse distributed representations exhibit graceful degradation in performance in the presence of increasing noise [71]. Sparse distributed vectors exist on a representational continuum that allows them to demonstrate characteristics of either localist or distributed encodings depending on their sparsity [72]. This flexibility motivates the use of abstract vector systems that define combinatorial operators over sparse distributed representations. VSAs [73], which we use to implement VS-BIND, accomplish this.

These models also differ in terms of how semantic representations are instantiated [74]. Some define the semantic structure of relational encodings at the neural level, generating explicit *role-filler* [53,55] or *symbol-argument* bindings [59]. These models provide clear mechanisms to support compositional relational encodings of semantic knowledge, where the whole perfectly reflects its parts, a feature considered important for linguistic modelling [53–56]. By comparison, VSAs do not inherently specify neural implementations at all, but are rather supported in this regard by broad theoretical frameworks that specify mappings between abstract vectors and neural population codes (for example, the semantic pointer architecture [22], used in Nengo [21], or integrated connectionist/symbolic architecture, ICS [75]).

Structured sequence processing typically focuses on ordering relationships and can refer to operations on meaningless items (e.g. nonsense words or abstract visual shapes). We

consider cognitive architectures incorporating VSAs [58,61–63] to have particular strengths appropriate to their use in a neurobiologically plausible model of structured sequence processing. Firstly, using existing tools, VSAs can act as a bridge to multiple modelling paradigms. For example, Nengo, a library for large-scale dynamic neural simulation [21], supports VSA encodings within spiking neural networks through the use of the Semantic Pointer Architecture. Bayesian computations [76] and attractor dynamics [77] have also been instantiated within this spiking neural network system. Secondly, VSAs are highly scalable solutions possessing substantial storage capacity for high-dimensional information [78]. Thirdly, as we describe shortly, VSAs permit the definition of relationships at the algorithmic level using an easily interpretable algebra. Finally, VSAs remain neurocomputationally plausible [22] whilst being relatively straightforward to implement [61,62,79]. To describe our approach, we next outline VSA principles and operators, before explaining how temporal dynamics can control such operations. We conclude by describing the role of these mechanisms within a neurobiologically plausible model of structured sequence processing, VS-BIND.

4. Combinatorial population coding with vector symbolic architectures

In VSA models, the basic units of representation are high-dimensional vectors. These typically sparse representations (figure 2, vectors **A** and **B**) can be visualized directly (figure 2, top left panel) or encoded as neural activity using multidimensional tuning functions (figure 2, top right panel) [21,83].

Symbolic vectors can be recombined using specific, reversible combinatorial operators to create new representations containing information on their constituents and the relations between them [58,61,62]. Inputs to these operations can either be atomic vectors (those that are not compound representations) or the results of previous operations. Atomic vectors are often randomly generated in VSA models, but can also be generated by compressing even higher dimensional input. Using this latter approach, semantically similar concepts cluster together in vector space [22].

The power of VSA approaches lies in their combinatorial operators. Although these are broadly common to all VSAs, here we use only one VSA, that of Plate [61]. This VSA uses the following operators: *superposition* (or *bundling*; point-wise vector addition of inputs), *binding* (a conjunctive operator) and its inverse, *unbinding* (also termed *release*, which here relies on *inversion*, equivalent to logical ‘not’), and a vector *comparison* operator for readout of results. Note that within VSA terminology, only the conjunctive operation is known as *binding*, but both this specific operator and superposition fit the wider definition of binding in the broad context of neural binding problems [5]. The more specific VSA nomenclature, which we adhere to hereafter, is helpful because it imposes constraints on the potential neural mechanisms involved in each of these combinatorial processes.

Using the full set of VSA operators, algorithmic manipulations are undertaken easily. The operators are best demonstrated with just two inputs (here **A** and **B**; typical VSA operators are illustrated in the top left panel of figure 2). **A** + **B** (*superposition*) yields a vector *correlated* with both **A** and **B**. This is simple vector addition, essentially overlaying the

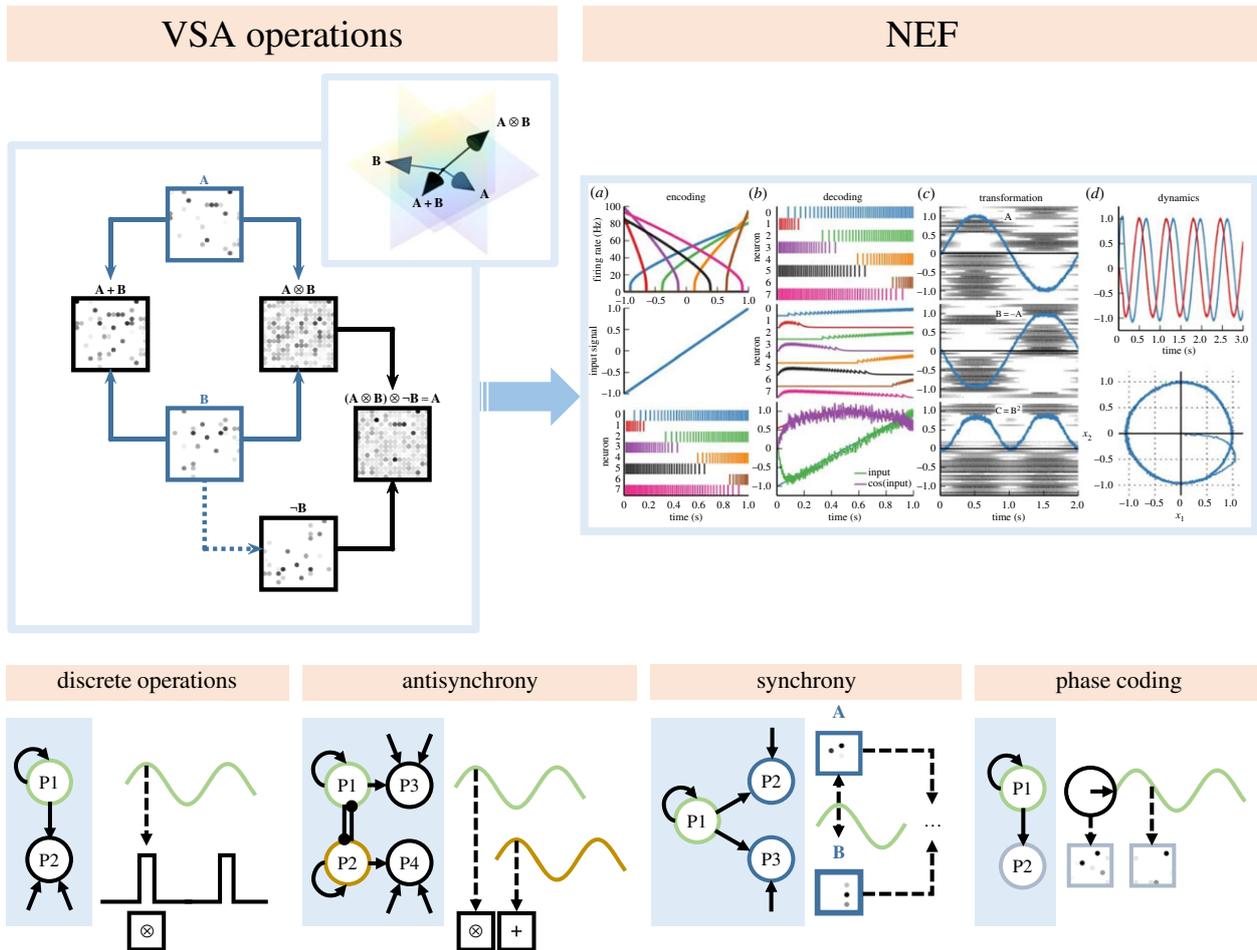


Figure 2. Spatial and temporal coding within a spiking model. *Top row, left panel:* vector symbolic architecture (VSA) operations using circular convolution to accomplish binding ($A \otimes B$) [61]. Sparse, high-dimensional random vectors represent distinct symbols A and B (blue text). To aid visual comparison, vectors are shown reshaped into squares, meaning the dimensionality of each vector equals the number of ‘pixels’ in each box (here, 256-dimensional, plotted as 16×16). These vectors can also be considered *directions* in high-dimensional space (inset, projected down to 3D using principal components analysis). Results of the VSA operators are shown (upper left main panel, black text, clockwise from left: *superposition*, *binding*, *unbinding* and *inversion*; see text, §4), with arrows indicating the flow of operands. A noisy recovered vector (rightmost square) can be cleaned up with an autoassociative memory to produce A (top square). *Top row, right panel* (from [21]; copyright © 2014 Bekolay, Bergstra, Hunsberger, DeWolf, Stewart, Rasmussen, Choo, Voelker and Eliasmith, CC BY 3.0): Core properties of the Neural Engineering Framework (NEF) [22] as implemented in Nengo [21]. To the left are tuning curves of individual neurons (a , top plot). In vector terms, each neuron fires maximally to its own *preferred direction*. Nonlinear *encoding* of an input signal (a , middle plot) yields spike trains for each neuron (a , bottom plot). *Decoding* (b) is possible using linear methods. Combining decoding with encoding, one can determine synaptic weights representing transformations between populations (c). Here, VSA representations are simply high-dimensional signals encoded like any other. Operations like convolution can be learned by simulated spiking networks incorporating spike-timing-dependent plasticity (STDP) in Nengo [22]. Finally, dynamic signals can be represented (d), of relevance for understanding oscillatory mechanisms [26,28,29]. Nengo is agnostic about neural models, with many spiking models available [80–82]. *Bottom row:* since the NEF provides mechanisms for spatial (VSA operations) and temporal (dynamic) manipulation of representations, possibilities exceed that of a static system. Simple interactions between segregated populations (*networks shown in blue boxes*) lead to controlled functional relationships. Thresholded dynamic activity, e.g. arising from an oscillator (P1, *leftmost panel*) can trigger discrete combinatorial operations. These operations can be segregated over time (*middle left panel*) through control by interacting, antisynchronous oscillations (P1 and P2). Likewise, common driving signals can synchronously strengthen representations in disparate regions (*middle right panel*, P1 multiplicatively modulating P2 and P3) for downstream processing such as feature binding. Finally, the *phase* of an oscillator, rather than its amplitude (*rightmost panel*) can drive downstream encodings such as those of relative position. (Online version in colour.)

sparse inputs. $A \otimes B$ (*binding*) is a conjunctive operator that yields a vector approximately *orthogonal* to both A and B , and as such is poorly correlated with either input. In this VSA, binding is calculated through circular convolution [61], which is just one possible way to create an output vector of the same length as one input alone [62,63]. This feat is possible because the operator encodes a *reduced representation* of both inputs, which can be unbound as described below.

Reduced representations are an important feature of a number of VSAs. Without them, conjunctive operations on input vectors of length N each result in output vectors of length N^2 (leading to exponential increases in vector size, a

combinatorial explosion). This is a characteristic of the ancestral VSA, Smolensky’s *tensor product binding* model [58,75]. Plate’s [61] VSA, by contrast, overcomes this scaling problem by making use of compressive representations known as *holographic reduced representations* (HRRs). Reduced representations allow a single model to support repeated (or recursive) operations on vectors without dimensionality increasing; the limiting factor in using reduced representations is instead that the process is lossy, so repeat operations cumulatively degrade the output vector. It is partly for this reason that this approach cannot be considered classically symbolic or perfectly compositional. However, lossy encoding

recapitulates natural limits on working memory and the depth of recursively nested structures that can be constructed or comprehended in natural language [84].

Here, *unbinding* is in essence binding, but with a change to one of the operands. Values may be unbound or *released* from a bound representation by computing a new binding between it and the *approximate inverse* of one of the original inputs with respect to the binding operator (from hereon in, just named the *inverse*, ‘ \neg ’; see $\neg\mathbf{B}$, figure 2). Inversion is accomplished by simply permuting all but the first dimension of \mathbf{B} , which at a neural level means rerouting input dimensions using a distinct pattern of synaptic weights. It is easier to understand the inversion operator by demonstrating its use in unbinding; the vector symbolic formula $(\mathbf{A} \otimes \mathbf{B}) \otimes \neg\mathbf{B}$ shown in figure 2 depicts this process. Here, we bind the bracketed representation with a ‘key’ containing the inverse of the element we know to be linked to the vector we wish to recover. Without such a key it is impossible to retrieve the contents of a binding via circular convolution. This means that VSA operations must be configured carefully to ensure that bound information can be retrieved in a plausible way.

To interrogate the result of unbinding, or any VSA operation, we can use the comparison operator to assess the similarity of the output to a defined vocabulary of representations. This can be useful as part of an autoassociative memory, or to control downstream actions or operations. For example, let $\mathbf{R} = (\mathbf{A} \otimes \mathbf{B}) \otimes \neg\mathbf{B}$. Here, undertaking a comparison between \mathbf{R} and a set of two vectors $\{\mathbf{A}, \mathbf{B}\}$ would reveal that \mathbf{R} is highly similar to \mathbf{A} and highly dissimilar to \mathbf{B} . Comparisons can be conducted on real-valued vectors using correlation or the vector dot product (where a smaller angle between input vectors indicates stronger similarity). The dot product in particular could be implemented by any neuron; it is equivalent to simply summing all presynaptic potentials weighted by their respective synaptic weights over a short period of time [70].

Use of VSA operations is by no means restricted to elementary computations. Mathematically, the superposition (+) and binding (\otimes) operators are designed to exhibit associativity, commutativity and distributivity akin to their scalar analogues, addition and multiplication [61]. For example, if \mathbf{A} , \mathbf{B} and \mathbf{C} each represent vectors, the vector symbolic formula $\mathbf{A} \otimes (\mathbf{B} + \mathbf{C})$ is equivalent to $(\mathbf{A} \otimes \mathbf{B}) + (\mathbf{A} \otimes \mathbf{C})$, just as one would expect algebraically. This makes the results of multiple VSA operations predictable and transparently interpretable. In a similar fashion, unbinding still functions if we superpose multiple bindings, as in $\mathbf{R} = (\mathbf{A} \otimes \mathbf{B}) + (\mathbf{C} \otimes \mathbf{D})$. In this case, $\mathbf{R} \otimes \neg\mathbf{D}$ will result in the value $(\mathbf{C} + \textit{noise})$, recovering a representation highly similar to \mathbf{C} .

Are such operations neurobiologically plausible? Plausibility of the neuronal arithmetic at a basic level is well supported; additive and multiplicative functions, which are sufficient to compute all of the VSA operations described here, including circular convolution, are established neural processes [85]. Multiplicative and divisive functions abound in the cognitive neurobiological literature, as feedback influences on neural responses [86,87]. The circular convolution operator could also be substituted for a number of alternative conjunctive distributed operators, for example, vector-derived transformation binding [63]. There is a good deal of evidence consistent with the presence of conjunctive distributed encodings in associative areas such as retrosplenial cortex [88], the ventral visual stream [89], and hippocampal CA1 and CA3 subregions [90].

How do we distinguish superposition from multiplicative [72] operations like convolutional binding in high dimensions? We expect that neurons instantiating additive operations should demonstrate linear responses to linear combinations of features of constituent representations, whilst multiplicative operations should result in nonlinearly selective responses. Empirical evidence already suggests that superposition cannot be the only combinatorial function, since linear and nonlinear mixed selectivity are *both* prevalent within the lateral prefrontal cortex and elsewhere [91,92]. Thus, there is evidence for the broad classes of functions to which the VSA operators belong in regions supporting cognitive function.

Simulations likewise demonstrate plausibility; Plate’s VSA operators, discussed here and shown in figure 2 [61], form the combinatorial backbone of Eliasmith’s Semantic Pointer Architecture, a neurobiologically plausible representational framework underlying Nengo, a Python library supporting the construction of spiking neural models [21,22]. The SPA proposes that high dimensional representational vectors arising in sensory cortex are compressed by similarity-preserving dimensionality reduction, and that the brain undertakes vector symbolic operations on these reduced representations. Nengo supports the transcoding of these cognitive representations into spiking neural representations by applying a set of fundamental neural encoding and decoding principles, the Neural Engineering Framework (NEF) [21], summarized with reference to the original paper in figure 2 (upper right panel and figure legend). Under this system, each spiking neuron contributes to a distributed encoding of an underlying latent vector representation; that is, there is a many-to-many mapping between the dimensions of a cognitive vector (itself already a distributed representation) and individual neurons. The resulting spiking activity is inherently dynamic, and thus suitable for manipulation over time by temporal mechanisms within the same dynamic framework.

5. Dynamically coordinating combinatorial operations with temporal mechanisms

As we have seen, population coding and temporal mechanisms are both functionally important to any account of domain-general structure-building [26,72]. Here, we introduce basic oscillatory principles that serve as temporal mechanisms within VS-BIND.

Neural oscillations in the brain reflect temporally coordinated responses of neural populations [93,94]. Oscillatory signals can also result in oscillatory *coupling* across frequency bands that reflects coordinated interactions within and between brain regions [94,95]. Theta-gamma coupling, for example, is associated with cognitive function and is also seen during structured sequence processing tasks [94,96].

We can instantiate dynamic relationships, including coupling, within a spiking neural model such that they undertake functionally useful coordinating or multiplexing roles in neurobiologically plausible ways (figure 2, bottom row). For example, a self-connected population (leftmost panel, P1) can generate oscillatory dynamics that, when thresholded, serve as the trigger for discrete vector symbolic operations on multiple inputs (P2). Likewise, interacting oscillators (P1 and P2, middle left) can exhibit antisynchrony, segregating discrete operations such that they do not interfere. Temporal synchrony can also be simulated (middle right) through top-down

influences on multiple neural populations, for example by controlling gain, such that downstream operators (for example, VSA operators) act on coordinated inputs. Here, temporal coordination serves to functionally associate two vector representations that would otherwise remain separate.

The dynamic context in which these spatial operators act is crucial, because it serves to mitigate the concern that, owing to their multiplicative interactions, such spatial bindings are too variable to support generalization over classes of their inputs, and thus insufficient as relational encodings, a problem characterized in the literature as *violation of role-filler independence* [72]. The explicit segregation of representations in space (figure 2, A versus B) and time (figure 2, top left panel, upstream/downstream) means that multiple neural ensembles concurrently instantiate different components of a combinatorial representation. Variability in the downstream binding of A and B would not prevent either of these two populations from generalizing over their inputs. Constituents may be dynamically bound or unbound as needed to segregate or aggregate information.

Finally, oscillatory signals can unidirectionally coordinate activity elsewhere, for example on the basis of phase (figure 2, bottom row, rightmost panel), thus potentially exhibiting phase–amplitude coupling effects. This produces behaviour consistent with conceptual models on the role of theta–gamma coupling [97] observed in the brain [94,98]. VS-BIND exhibits phase–amplitude coupling as an explicit functional property of relative position coding.

6. Network-level mechanistic hypotheses derived from VS-BIND

Our objective within the remainder of this paper is to outline the specific combinations of operations required to support dependency encoding during sequence processing, generating neurobiological hypotheses for adjacent, non-adjacent and hierarchical dependency encoding. The MATLAB demo provided at doi:10.5281/zenodo.3464607 further illustrates key principles outlined here.

(a) Adjacent dependencies

We first consider the binding of two items following each other closely in time. At the vector symbolic level, ordered sequences can be unambiguously represented using one of two principal methods. The first is by encoding elements with respect to each other (*chaining*). However, behavioural findings in multiple species do not convincingly support this approach [99]. An alternative method, encoding each element with respect to a serial *positional tag* ($Sequence = 1^{st} * Item1 + 2^{nd} * Item2$, and so on), has found greater support. Behavioural results in songbirds suggest a reliance on positional cues during sequence recognition [100]. Likewise, during recall, humans and non-human primates are more likely to confuse items in different memorized sequences if those items share the same ordinal position across sequences [9]. Positional influences on recall provide behavioural evidence that ordinality is incorporated into the encoding of sequences even when position is not explicitly featured in a task, offering support for positional tagging.

Detection of ordinal serial position, an essential component of sequence encodings within VS-BIND, we posit involves DLPFC, and motor and premotor cortex. Electrophysiological

findings in non-human primates align with this account, revealing populations of cells in each of these regions involved in consistently encoding serial position irrespective of stimulus identity [101,102]. Studies of the hippocampus likewise reveal temporal coding relative to stimulus presentation, hypothesized to form part of the context for later retrieval [103]. There is thus a neurobiological basis for the crucial role played by explicit positional tags within our model (figure 3, light blue–grey boxes). Within VS-BIND, positional tags are considered to be deterministically but flexibly generated by the brain.

Here, tags follow the nomenclature ‘primary’ and ‘secondary’ (1° and 2°) rather than the ordinal absolutes ‘first’ or ‘second’, in line with human and non-human primate behavioural [104–106] and electrophysiological [101,103] evidence suggesting that, within sequences of words or actions, ordinal position is encoded relative to sequence boundaries rather than in absolute terms. Within our model, positional tags are anchored to the boundaries of perceptual chunks (figure 3, box inset) by decoding stimulus-entrained oscillatory phase (figure 2) such that both the sequence items and their relative positional tags are derived from the sensory input. It is important that each tag remains orthogonal to the last, which is a requirement for later unambiguous recovery of specific elements in a sequence.

To encode a sequence, continuous input over time is first discretized. These discrete sensory items are bound to distinct positional tags (to form position–item representations) and superposed in a decaying, recurrently connected working memory buffer. This *sequence buffer*, likely supported by SMA and pre-SMA [107], therefore maintains a linear, ordered representation of the input sequence (i.e. $Sequence = 1^\circ \otimes Item1 + 2^\circ \otimes Item2 + 3^\circ \otimes Item3$). From this, individual sensory representations can subsequently be retrieved (via $Item \approx Sequence \otimes \neg Position$) and recoded to reflect dependencies between items or chunks (figure 3, bottom row, showing retrieved items). During this recoding process, maintained representations (A, B, or irrelevant intervening items, X) serially accumulate within a *dependency buffer* over time (i.e. moving rightwards) where they can be used in increasingly complex binding operations (moving upwards). The SMA encoding steps are comparable to the VSA approach in the ordinal serial encoding (OSE) working memory model of Choo & Eliasmith [99], which is capable of modelling behavioural characteristics of serial recall such as working memory *primacy* and *recency* effects. However, unlike the OSE model, VS-BIND incorporates centrally coordinating oscillatory activity and uses boundary-relative (rather than absolute) ordinal codes. Moreover, subsequent to serial encoding within working memory, VS-BIND describes the encoding of various dependencies, which need not be linear in organization (figure 3).

To encode dependencies, items are retrieved from sequence memory and bound with new positional tags (figure 3, light blue–grey boxes). As in sequence memory, bindings form over time between tags and corresponding items, which are superposed to form representations of specific dependencies (figure 3, topmost representation, all diagrams). The timing of each binding operation is associated with and coordinated by stimulus-related oscillatory activity (figure 2) [29,31].

(b) Non-adjacent dependencies

To support efficient encoding, and permit generalization from learned adjacent dependencies, simple non-adjacent relationships may be encoded in an equivalent manner to adjacent

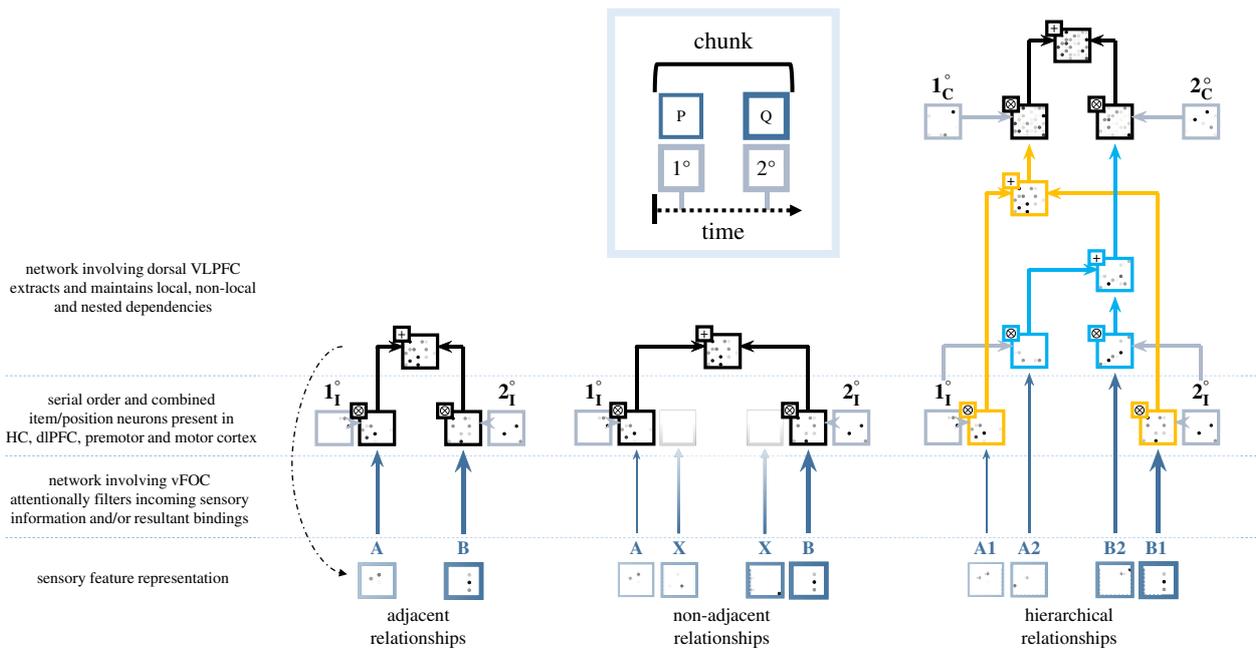


Figure 3. Neurobiologically informed vector symbolic encoding of sequence structure. Vector symbolic operators can account for the processing of a variety of sequencing dependencies. The solid arrows in these charts indicate the flow of information during *encoding* of a stimulus sequence only. These describe transformations of latent vector symbolic representations, as opposed to neural activation patterns. Representation strength is denoted by the shading and thickness of each box border. For clarity, representations are shown separated along the horizontal axis, though separate boxes do not necessarily imply separate neural populations are engaged, especially if describing identical computations, which could be undertaken by neurons of a single region. We suggest sensory representations (bottom row) are maintained within the supplementary motor area (not shown) and retrieved as needed. Operations unfold dynamically following principles outlined in figure 2. The final encoded sequence representation is found at the top of each diagram. Each is a reduced representation whose constituents can be inspected without serially unpacking all bindings; the superposed final result of the *adjacent relationship* encoding (leftmost diagram), for example, can be interrogated to recover its secondary element by simply binding it with $\neg 2^\circ$. Serial elements packaged into a single representation are considered to be *chunked* in the traditional sense (inset box), but identical operations can be applied to non-adjacent (middle diagram) or nested pairs of elements (rightmost diagram), using separate item (I) and chunk (C) position encodings. Selective fading along the vertical axis represents salience filtering in the non-adjacent example (middle). Finally, the dashed, curved arrow shows just one case in which *sub-symbolic* feedback from a cognitively abstract cortical region might ultimately influence the representation of individual elements in sensory cortex (there may be many such pathways, but one exemplar is shown). Thus, although the figure, for simplicity, suggests VS-BIND is largely a feedforward model, feedback influences feature and can, for example, allow certain areas to influence sensory cortical representations. (Online version in colour.)

dependencies. To accomplish this, all constituent items must be maintained in the sequence buffer long enough to be integrated, and maintained representations need to be selectively propagated to the dependency buffer. This process involves salience filtering (e.g. by repetition suppression, expectancy or attentional processes [108]; figure 3, centre diagram), possibly supported by regions such as ventral frontal cortex including the frontal operculum. This is based on findings that the frontal operculum appears to be more active during the presentation of infrequent or novel auditory cues [109] and responds preferentially to violations of adjacent dependencies [38]. We suggest that this region integrates information from working memory as soon as it is available, but only maintains it over a relatively short time period. In this case, preferential responses to adjacent dependency violations could be explained by frontal operculum actively inhibiting representations of both low probability items and short, low probability *n-grams* of contiguous elements (for example, those in which the constituent elements are not suitably ordered).

Salience filtering enables non-adjacent dependencies to be encoded into representations identical to their adjacent counterparts, for any length of dependency fitting into working memory. That is, for any salient **A** and **B** and *n* irrelevant **X** elements, the sequence **A-Xⁿ-B** may be rendered into the same

encoding as **AB** alone by selectively inhibiting downstream encodings of irrelevant **X** representations (figure 3, identical top representations, left and middle diagrams). Discarding irrelevant items results in a more efficient representational code, and in state-transition terms allows for grammaticality judgements to be made on *n*th-order non-adjacent dependencies using only a first-order Markov process.

(c) Hierarchical dependencies

As shown, within our model a dependency comprises multiple superposed *position* \otimes *item* bindings. Although the dependency can incorporate items retrieved from either contiguous or discontinuous positions in serial working memory, the recoded dependency representation can in both cases be considered a chunk (in figure 3, left diagram, the chunk is **A-B**; in the right diagram, the two chunks are **A1-B1** and **A2-B2**). To encode hierarchical dependencies (figure 3, right diagram, showing nested dependencies), every *dependency* needs to be bound with a unique positional tag, as for individual items. Like individual items, these can also be superposed, forming a superchunk containing a higher-order dependency representation. This process can be recursively repeated to form a single reduced representation of the hierarchical structure of the entire input

sequence, integrating progressively increasing amounts of information at higher hierarchical levels of encoding.

The above system is sufficient to compress hierarchical structure into a reduced representation. However, we must encode the reduced representation using more than item positional tags if we are to support unambiguous recovery of specific constituents and comfortably discard the original representations. To avoid generating identical codes for dissimilar structures, we can define sets of positional tags specific to each level of the hierarchy, for example 1_1^1 and 2_1^1 for the first and second items, 1_2^1 and 2_2^1 for the first and second chunks, and so on (figure 3, right diagram). For this reason, it can be computationally beneficial to define all positional tags as convolutional powers (i.e. $\text{base}^{\text{exponent}}$) of a given *base vector* through repeated self-binding (i.e. $2^\circ = 1^\circ \otimes 1^\circ = (1^\circ)^2$). This, or a similarly invariant function, can be learned by a network such that tags are encoded as a function of not just position (by varying the exponent), but also context (by varying the base vector). Crucially, by binding items only to a finite set of deterministically generated positional tags, the modelled system can always undertake unbinding by re-instantiating the same set of keys; by iterating through every possible tag, all constituent items can be retrieved in sequence.

A given dependency coding network can retrieve and recode items from arbitrary positions in the linear sequence buffer, provided it has sufficient integrative power. Therefore, this representational scheme is capable of encoding not just nested dependencies, but other types of hierarchical dependency such as crossed dependencies. For example, figure 3 (right diagram) illustrates retrieval of items 1 and 4 (A1 and B1) from sequence memory, which are chunked through binding and superposition; and retrieval of items 2 and 3 (A2 and B2), which are likewise chunked. These two chunks form a superchunk representing a nested dependency. Cross-serial dependencies between words are notably also present in some human languages. Like nested dependencies, these can also be readily represented by our model. We could, for example, have retrieved and bound items 1 and 3, and 2 and 4, respectively, to form a crossed dependency structure. This is illustrated within our coded demonstration (doi:10.5281/zenodo.3464607). The inclusion of cross-serial dependencies alters the minimum computational requirements for any parsing agent [110], and thus it is important that a domain-general representational model has the potential to account for them as well as other language-like hierarchical constructions.

Having recoded a linear sequence in terms of its dependencies, it is possible to recover one or more specific items from the representation of hierarchical structure by unbinding using keys that specify context-specific position(s). With reference to figure 3 and using this method: the key $\neg 1_2^1$ will recover all position-item bindings of the first chunk; the key $\neg 2_1^1$ will recover the second item of every chunk (with each still bound to information on the requisite chunk position); and the key $\neg(2_1^1 \otimes 1_2^1)$ will recover the second item of the first chunk. This flexible decoding scheme is important because it readily supports manipulation of entire chunks and generalization of dependencies to multiple timescales.

A natural consequence of the above encoding scheme is that the superposition of many dependency representations over time can gradually give rise to a memory trace. This trace will be influenced by the probabilistic distribution of dependencies over the set of input sequences, analogous to

implicit learning. It is now possible to computationally model functional characteristics of hippocampal subregions, and by this method, it has been proposed that a monosynaptic (entorhinal cortex to CA1) pathway possesses the relevant properties to support implicit learning [40]. This is relevant for sequence learning, where dependencies are established over many trials. Single-trial learning, by contrast, requires processing by further hippocampal subregions. This suggests a prominent role for parts of the hippocampal system in sequence processing, in line with recent findings [39,40]. However, this is not to say that we should expect activation of the hippocampus to be constant throughout. Indeed, the human neurobiology literature suggests that there is a decrease in hippocampal involvement over time when learning an artificial grammar [111], or acquiring a novel semantically meaningful lexicon [112]. These observations can be interpreted by way of a predictive coding account, in terms of the degree of mismatch between stored and incoming sensory encodings [113]. We can relate this associative mismatch account of hippocampal involvement to the idea of a superposed memory trace. Namely, any newly presented sequence is likely to be more similar to a superposition of encountered sequences (weighted to simulate rehearsal) as time progresses and the number of constituents in the superposition grows. Thus, the degree of mismatch, and any activation requisite for such an encoding, is likely to decrease over time.

We propose that hierarchical structure-building is one of the key roles of the dorsal aspect of ventrolateral prefrontal cortex (dorsal VLPFC, incorporating Brodmann areas 44 and 45). This position is supported by human neurobiological evidence on syntactic processes prominently featuring hierarchical dependencies [34,38]. It has been proposed that at least BA44 supports a recursive, multi-dependency management process, a fact that would explain increases in its activation observed with increasing depths of hierarchical dependency nesting [114]. The recursive reuse of a consistent architecture by VLPFC would be consistent with our recursive use of vector symbolic operations during the encoding of hierarchical dependencies. Repeated superposition of sparse dependency representations will manifest as an increase in local activity as increasing numbers of neurons support the representation; as an example of this effect, consider the increased activity represented at the top of figure 3, relative to the bottom. A 'reset' (or re-sparsification) of the buffer will result in a sudden drop in population activity. Such neural accumulation and reset activity has been identified within human intracranial recordings in subjects listening to sentences containing words that accumulate into phrases [28]. Furthermore, representations of constituents might not persist beyond the need to encode the reduced dependency representation, a fact that highlights an interesting property of the model: detectable neural delay activity does not need to persist throughout working memory maintenance of the input sequence, consistent with recently reported findings on the neurobiology of working memory [115].

It must be emphasized that the flow diagrams in figure 3 only show *symbolic* information flow. There are likely to be sub-symbolic influences acting over time to hone the associated neural activation patterns, requiring feedback and feedforward interactions between neural ensembles in the model. These are briefly alluded to in figure 3 as a single exemplar (curved, dashed arrow, left) denoting top-down influences acting on sensory cortical representations. These sub-symbolic influences

may be investigated further by instantiating a dynamic neural implementation of the model via the NEF [83], where learning algorithms employing, for instance, spike-timing-dependent plasticity can be applied to learn functions over time [116].

7. In conclusion: predictions emerging from the structure of VS-BIND

Vector symbolic operations implemented in artificial neural networks, as we have established, have the potential to further our understanding of combinatorial binding at neural, cognitive and behavioural levels, generating site-specific neural correlates ripe for testing. Testing falsifiable predictions made by this model is important to provide evidence about the plausibility of the combinatorial operators and processes modelled here.

The components of VS-BIND we have outlined here suggest the following predictions:

- *Superposition of sparse vectors* manifests as steadily increasing net activation [28].
- *Binding produces orthogonal vectors* by re-coding, which might be observed as reductions in neural responses, for example, reductions in high gamma activity during neural local field potential recordings [28].
- *Chunking operations* are intrinsic to linguistic bracketing and both produce and depend upon dynamic patterns of both increasing and decreasing oscillatory activity [26].
- *Relative ordinal position* is a crucial neural code when binding serial input in artificial and biological neural systems, indicating that time-sensitive or serial-position-sensitive cells, wherever they reside [101], are indispensable for complex combinatorial binding in sequence processing.
- *Hippocampal system involvement* is not only required for recurrent activity to establish associations during Hebbian learning, but components of this system are also involved in structured sequence processing, establishing rules and dependencies across temporal scales through interactions with at least inferior frontal cortex.

In summary, the VS-BIND model describes how sequences can be processed and represented using established VSA representations and operations, with the main goal of generating falsifiable predictions of neural correlates. While it was beyond the scope of this paper to model the learning processes, VSA-based models such as ours could

now be extended with spike-timing-dependent plasticity or other learning and memory algorithms.

It is also important to note that, although our model can represent novel sequences of arbitrary items, the model does not bind arbitrary items together, which risks producing uninterrogable or unrecoverable bound representations. Instead, VS-BIND only uses the binding operator to combine an item with a known positional or contextual tag. Critically, these positional tags are known to the model, being drawn from a small, stable set of codes that are deterministically re-instantiated by dIPFC neurons or hippocampal ‘time cells’. This process allows retrieval of arbitrary items without the need for model-specific knowledge or a look up table: all that is needed is the set of positional tags and the sequence representation. Such models could, in the future, be combined with systems that explicitly support the binding together of arbitrary items, for example, to flexibly manipulate linguistic relations in working memory [52].

We have here presented a blueprint for a neurobiologically plausible computational model, VS-BIND, outlining its principal mechanisms for encoding sequence dependencies. We have also highlighted aspects of VS-BIND that support ongoing efforts to simulate cognitive functions relevant to segmentation, chunking, recall and prediction using VSAs within spiking neural networks. Further developments, in conjunction with other models, such as those that are classically compositional approaches aimed at modelling language-specific or linguistic properties, carry tremendous potential to better understand fundamental aspects of cognition and to guide the pursuit of neurobiological correlates of complex mental structures.

Data accessibility. Supporting MATLAB code can be found at doi:10.5281/zenodo.3464607 [23].

Authors' contributions. This paper is part of R.C.'s PhD thesis work in the Laboratory of Comparative Neuropsychology. C.I.P. was lead supervisor and B.W. and Y.K. co-supervisors of the work. R.C. conceived the ideas and wrote the paper with C.I.P. and input from B.W. and Y.K. R.C. implemented and provided the MATLAB code.

Competing interests. We declare we have no competing interests.

Funding. This work was supported by a Wellcome Trust PhD Studentship (R.C.), a Wellcome Trust Investigator Award (C.I.P., WT092606AIA), the BBSRC (C.I.P., BB/J009849/1), a Sir Henry Wellcome Fellowship (B.W., WT110198), the European Research Council (C.I.P., ERC CoG, MECHIDENT) and the NIH (with M. Howard, R01-DC04290).

Acknowledgements. We thank Elizabeth Buffalo, Pascal Fries, Karl Friston, Tim Griffiths, David Poeppel, Mark Stokes and Chris Summerfield for inspiring discussions on complex combinatorial binding.

References

1. Kraus BJ, Robinson RJ, White JA, Eichenbaum H, Hasselmo ME. 2013 Hippocampal ‘time cells’: time versus path integration. *Neuron* **78**, 1090–1101. (doi:10.1016/j.neuron.2013.04.015)
2. Moser M-B, Rowland DC, Moser EI. 2015 Place cells, grid cells, and memory. *Cold Spring Harb. Perspect. Biol.* **7**, a021808. (doi:10.1101/cshperspect.a021808)
3. O'Keefe J. 1976 Place units in the hippocampus of the freely moving rat. *Exp. Neurol.* **51**, 78–109. (doi:10.1016/0014-4886(76)90055-8)
4. Lashley KS. 1951 The problem of serial order in behavior. In *Cerebral Mechanisms in Behavior; The Hixon Symposium* (ed. LA Jeffress), pp. 112–146. Oxford, UK: Wiley.
5. Feldman J. 2013 The neural binding problem(s). *Cogn. Neurodyn.* **7**, 1–11. (doi:10.1007/s11571-012-9219-8)
6. Gray CM. 1999 The temporal correlation hypothesis of visual feature integration: still alive and well. *Neuron* **24**, 31–47. (doi:10.1016/S0896-6273(00)80820-X)
7. Singer W. 1999 Neuronal synchrony: a versatile code for the definition of relations? *Neuron* **24**, 49–65. (doi:10.1016/S0896-6273(00)80821-1)
8. Fitch WT, Martins MD. 2014 Hierarchical processing in music, language, and action: Lashley revisited. *Ann. NY Acad. Sci.* **1316**, 87–104. (doi:10.1111/nyas.12406)
9. Dehaene S, Meyniel F, Wacongne C, Wang L, Pallier C. 2015 The neural representation of sequences: from transition probabilities to algebraic patterns and linguistic trees. *Neuron* **88**, 2–19. (doi:10.1016/j.neuron.2015.09.019)
10. Tettamanti M, Perani D. 2012 The neurobiology of structure-dependency in natural language grammar. In *The handbook of the*

- neuropsychology of language, 2 volume Set.* John Wiley & Sons.
11. Bickerton D, Szathmáry E. (eds). 2009 *Biological foundations and origin of syntax*, 1st edn, 492p. Cambridge, MA: The MIT Press.
 12. Rouault M, Koehlin E. 2018 Prefrontal function and cognitive control: from action to language. *Curr. Opin. Behav. Sci.* **21**, 106–111. (doi:10.1016/j.cobeha.2018.03.008)
 13. Jeon H-A. 2014 Hierarchical processing in the prefrontal cortex in a variety of cognitive domains. *Front. Syst. Neurosci.* **8**, 223. (doi:10.3389/fnsys.2014.00223)
 14. Koelsch S, Siebel WA. 2005 Towards a neural basis of music perception. *Trends Cogn. Sci.* **9**, 578–584. (doi:10.1016/j.tics.2005.10.001)
 15. Dehaene S, Molko N, Cohen L, Wilson AJ. 2004 Arithmetic and the brain. *Curr. Opin. Neurobiol.* **14**, 218–224. (doi:10.1016/j.conb.2004.03.008)
 16. Sainburg T, Theilman B, Thielk M, Gentner TQ. 2019 Parallels in the sequential organization of birdsong and human speech. *Nat. Commun.* **10**, 3636. (doi:10.1038/s41467-019-11605-y)
 17. Yu AC, Margoliash D. 1996 Temporal hierarchical control of singing in birds. *Science* **273**, 1871–1875. (doi:10.1126/science.273.5283.1871)
 18. Wilson B, Spierings M, Ravignani A, Mueller JL, Mintz TH, Wijnen F, van der Kant A, Smith K, Rey A. In press. Non-adjacent dependency learning in humans and other animals. *Top. Cogn. Sci.* (doi:10.1111/tops.12381)
 19. Wang L, Uhrig L, Jarraya B, Dehaene S. 2015 Representation of numerical and sequential patterns in macaque and human brains. *Curr. Biol.* **25**, 1966–1974. (doi:10.1016/j.cub.2015.06.035)
 20. Wilson B *et al.* 2015 Auditory sequence processing reveals evolutionarily conserved regions of frontal cortex in macaques and humans. *Nat. Commun.* **6**, 8901. (doi:10.1038/ncomms9901)
 21. Bekolay T, Bergstra J, Hunsberger E, DeWolf T, Stewart TC, Rasmussen D, Choo X, Voelker AR, Eliasmith C. 2014 Nengo: a python tool for building large-scale functional brain models. *Front. Neuroinform.* **7**, 48. (doi:10.3389/fninf.2013.00048)
 22. Eliasmith C. 2013 *How to build a brain: a neural architecture for biological cognition*. Oxford, UK: Oxford University Press.
 23. Calmus R. 2019 rcalmus/vs-bind: VS-BIND supplementary material (version v1.0.0). *Zenodo*. (doi:10.5281/zenodo.3464607)
 24. Friederici AD, Chomsky N. 2017 *Language in our brain: the origins of a uniquely human capacity*, 1st edn. Cambridge, MA: MIT Press.
 25. Hagoort P. 2013 MUC (memory, unification, control) and beyond. *Front. Psychol.* **4**, 416. (doi:10.3389/fpsyg.2013.00416)
 26. Ding N, Melloni L, Zhang H, Tian X, Poeppel D. 2016 Cortical tracking of hierarchical linguistic structures in connected speech. *Nat. Neurosci.* **19**, 158–164. (doi:10.1038/nn.4186)
 27. Ghitza O. 2011 Linking speech perception and neurophysiology: speech decoding guided by cascaded oscillators locked to the input rhythm. *Front. Psychol.* **2**, 130. (doi:10.3389/fpsyg.2011.00130)
 28. Nelson MJ *et al.* 2017 Neurophysiological dynamics of phrase-structure building during sentence processing. *Proc. Natl Acad. Sci. USA* **114**, E3669–E3678. (doi:10.1073/pnas.1701590114)
 29. Barczak A, O'Connell MN, McGinnis T, Ross D, Mowery T, Falchier A, Lakatos P. 2018 Top-down, contextual entrainment of neuronal oscillations in the auditory thalamocortical circuit. *Proc. Natl Acad. Sci. USA* **115**, E7605–E7614. (doi:10.1073/pnas.1714684115)
 30. Giraud A-L, Poeppel D. 2012 Cortical oscillations and speech processing: emerging computational principles and operations. *Nat. Neurosci.* **15**, 511–517. (doi:10.1038/nn.3063)
 31. Lakatos P, Karmos G, Mehta AD, Ulbert I, Schroeder CE. 2008 Entrainment of neuronal oscillations as a mechanism of attentional selection. *Science* **320**, 110–113. (doi:10.1126/science.1154735)
 32. Friederici AD, Bahlmann J, Heim S, Schubotz RI, Anwander A. 2006 The brain differentiates human and non-human grammars: functional localization and structural connectivity. *Proc. Natl Acad. Sci. USA* **103**, 2458–2463. (doi:10.1073/pnas.0509389103)
 33. Kikuchi Y, Sedley W, Griffiths TD, Petkov CI. 2018 Evolutionarily conserved neural signatures involved in sequencing predictions and their relevance for language. *Curr. Opin. Behav. Sci.* **21**, 145–153. (doi:10.1016/j.cobeha.2018.05.002)
 34. Wilson B, Marslen-Wilson WD, Petkov CI. 2017 Conserved sequence processing in primate frontal cortex. *Trends Neurosci.* **40**, 72–82. (doi:10.1016/j.tins.2016.11.004)
 35. Jiang X, Long T, Cao W, Li J, Dehaene S, Wang L. 2018 Production of supra-regular spatial sequences by macaque monkeys. *Curr. Biol.* **28**, 1851–1859. (doi:10.1016/j.cub.2018.04.047)
 36. Friederici AD. 2004 Processing local transitions versus long-distance syntactic hierarchies. *Trends Cogn. Sci.* **8**, 245–247. (doi:10.1016/j.tics.2004.04.013)
 37. Petkov CI, Wilson B. 2012 On the pursuit of the brain network for proto-syntactic learning in non-human primates: conceptual issues and neurobiological hypotheses. *Phil. Trans. R. Soc. B* **367**, 2077–2088. (doi:10.1098/rstb.2012.0073)
 38. Friederici AD, Fiebach CJ, Schlesewsky M, Bornkessel ID, von Cramon DY. 2006 Processing linguistic complexity and grammaticality in the left frontal cortex. *Cereb. Cortex* **16**, 1709–1717. (doi:10.1093/cercor/bhj106)
 39. Opitz B, Friederici AD. 2004 Brain correlates of language learning: the neuronal dissociation of rule-based versus similarity-based learning. *J. Neurosci.* **24**, 8436–8440. (doi:10.1523/JNEUROSCI.2220-04.2004)
 40. Schapiro AC, Turk-Browne NB, Botvinick MM, Norman KA. 2017 Complementary learning systems within the hippocampus: a neural network modelling approach to reconciling episodic memory with statistical learning. *Phil. Trans. R. Soc. B* **372**, 20160049. (doi:10.1098/rstb.2016.0049)
 41. Jablonowski J, Taesler P, Fu Q, Rose M. 2018 Implicit acoustic sequence learning recruits the hippocampus. *PLoS ONE* **13**, e0209590. (doi:10.1371/journal.pone.0209590)
 42. Cui Y, Ahmad S, Hawkins J. 2016 Continuous online sequence learning with an unsupervised neural network model. *Neural Comput.* **28**, 2474–2504. (doi:10.1162/NECO_a_00893)
 43. Gasser B, Arbib MA. 2017 A neuro-computational model of sequence learning in macaques: the simultaneous chaining paradigm. *Adapt. Behav.* **25**, 195–213. (doi:10.1177/1059712317715324)
 44. Marr D. 1982 *Vision: a computational approach*. San Francisco: Freeman.
 45. Poggio T. 2012 The levels of understanding framework, revised. *Perception* **41**, 1017–1023. (doi:10.1068/p7299)
 46. Fitch WT. 2014 Toward a computational framework for cognitive biology: unifying approaches from cognitive neuroscience and comparative cognition. *Phys. Life Rev.* **11**, 329–364. (doi:10.1016/j.plrev.2014.04.005)
 47. Eliasmith C, Kolbeck C. 2015 Marr's attacks: on reductionism and vagueness. *Top. Cogn. Sci.* **7**, 323–335. (doi:10.1111/tops.12133)
 48. Marblestone AH, Wayne G, Kording KP. 2016 Toward an integration of deep learning and neuroscience. *Front. Comput. Neurosci.* **10**, 94. (doi:10.3389/fncom.2016.00094)
 49. Hochreiter S, Schmidhuber J. 1997 Long short-term memory. *Neural Comput.* **9**, 1735–1780. (doi:10.1162/neco.1997.9.8.1735)
 50. Kar K, Kubilius J, Schmidt K, Issa EB, DiCarlo JJ. 2019 Evidence that recurrent circuits are critical to the ventral stream's execution of core object recognition behavior. *Nat. Neurosci.* **22**, 974–983. (doi:10.1038/s41593-019-0392-5)
 51. Samsonovich AV. 2010 Toward a unified catalog of implemented cognitive architectures. In *Biologically Inspired Cognitive Architectures 2010: Proc. 1st Annu. Meeting BICA Soc.* (eds AV Samsonovich, KR Jóhannsdóttir, A Chella, B Goertzel), pp. 195–244. Amsterdam, The Netherlands: IOS Press.
 52. Hummel JE, Holyoak KJ, Green C, Doumas LAA, Devnich D, Aniket K, Kalar DJ. 2004 A solution to the binding problem for compositional connectionism. *AAAI*.
 53. Hummel JE, Holyoak KJ. 2003 A symbolic-connectionist theory of relational inference and generalization. *Psychol. Rev.* **110**, 220–264. (doi:10.1037/0033-295X.110.2.220)
 54. Hummel JE, Holyoak KJ. 1997 Distributed representations of structure: a theory of analogical access and mapping. *Psychol. Rev.* **104**, 427–466. (doi:10.1037/0033-295X.104.3.427)
 55. Shastri L, Ajjanagadde V. 1993 From simple associations to systematic reasoning: a connectionist representation of rules, variables and dynamic bindings using temporal synchrony. *Behav. Brain Sci.* **16**, 417–451. (doi:10.1017/S0140525X00030910)
 56. Doumas LAA, Hummel JE, Sandhofer CM. 2008 A theory of the discovery and predication of relational concepts. *Psychol. Rev.* **115**, 1–43. (doi:10.1037/0033-295X.115.1.1)

57. van der Velde F, de Kamps M. 2006 Neural blackboard architectures of combinatorial structures in cognition. *Behav. Brain Sci.* **29**, 37–70; discussion 70–108. (doi:10.1017/S0140525X06009022)
58. Smolensky P. 1990 Tensor product variable binding and the representation of symbolic structures in connectionist systems. *Artif. Intell.* **46**, 159–216. (doi:10.1016/0004-3702(90)90007-M)
59. Halford GS, Wilson WH, Phillips S. 1998 Processing capacity defined by relational complexity: implications for comparative, developmental, and cognitive psychology. *Behav. Brain Sci.* **21**, 803–831; discussion 831–864. (doi:10.1017/S0140525X98001769)
60. Wilson WH, Halford GS, Gray B, Phillips S. 2001 The STAR-2 model for mapping hierarchically structured analogs. pp. 125–160. Cambridge, MA: MIT Press.
61. Plate TA. 1995 Holographic reduced representations. *IEEE Trans. Neural Netw.* **6**, 623–641. (doi:10.1109/72.377968)
62. Kanerva P. 1994 The spatter code for encoding concepts at many levels. In *ICANN '94* (eds M Marinaro, PG Morasso), pp. 226–229. London, UK: Springer.
63. Gosmann J, Eliasmith C. 2019 Vector-derived transformation binding: an improved binding operation for deep symbol-like processing in neural networks. *Neural Comput.* **31**, 849–869. (doi:10.1162/neco_a_01179)
64. Hadley RF. 2007 Synchronous versus conjunctive binding: a false dichotomy? *Connection Sci.* **19**, 111–130. (doi:10.1080/09540090701219031)
65. Martin AE, Doumas LAA. 2017 A mechanism for the cortical computation of hierarchical linguistic structure. *PLoS Biol.* **15**, e2000663. (doi:10.1371/journal.pbio.2000663)
66. Roy A. 2017 The theory of localist representation and of a purely abstract cognitive system: the evidence from cortical columns, category cells, and multisensory neurons. *Front. Psychol.* **8**, 186. (doi:10.3389/fpsy.2017.00186)
67. Bowers JS. 2009 On the biological plausibility of grandmother cells: implications for neural network theories in psychology and neuroscience. *Psychol. Rev.* **116**, 220–251. (doi:10.1037/a0014462)
68. Wixted JT, Squire LR, Jang Y, Papez MH, Goldinger SD, Kuhn JR, Smith KA, Treiman DM, Steinmetz PN. 2014 Sparse and distributed coding of episodic memory in neurons of the human hippocampus. *Proc. Natl Acad. Sci. USA* **111**, 9621–9626. (doi:10.1073/pnas.1408365111)
69. Willmore BDB, King AJ. 2009 Auditory cortex: representation through sparsification? *Curr. Biol.* **19**, R1123–R1125. (doi:10.1016/j.cub.2009.11.003)
70. Rolls ET, Treves A. 2011 The neuronal encoding of information in the brain. *Prog. Neurobiol.* **95**, 448–490. (doi:10.1016/j.pneurobio.2011.08.002)
71. Werning M, Hinzen W, Machery E. 2012 *The Oxford handbook of compositionality*. Oxford, UK: Oxford University Press.
72. Doumas LAA, Hummel JE. 2005 Approaches to modeling human mental representations: what works, what doesn't, and why. In *The Cambridge handbook of thinking and reasoning* (eds KJ Holyoak, RG Morrison), pp. 73–91. New York, NY: Cambridge University Press.
73. Gayler RW. 2003 Vector Symbolic Architectures answer Jackendoff's challenges for cognitive neuroscience. See <http://cogprints.org/3983/>.
74. Logie R. 2007 *The cognitive neuroscience of working memory*. Oxford, UK: Oxford University Press.
75. Smolensky P, Legendre G. 2011 Cognitive architecture. *The harmonic mind: from neural computation to optimality-theoretic grammar*, vol. 1. Cambridge, MA: MIT Press.
76. Sharma S, Voelker A, Eliasmith C. 2017 A spiking neural Bayesian model of life span inference. In *Proc. 39th Annu. Meeting Cogn. Sci. Soc.* (eds G Gunzelmann, A Howes, T Tenbrink, EJ Davelaar), pp. 3131–3136. Austin, TX: Cognitive Science Society.
77. Gilra A, Gerstner W. 2017 Predicting non-linear dynamics by stable local learning in a recurrent spiking neural network. *eLife* **6**, e28295. (doi:10.7554/eLife.28295)
78. Crawford E, Gingerich M, Eliasmith C. 2016 Biologically plausible, human-scale knowledge representation. *Cogn. Sci.* **40**, 782–821. (doi:10.1111/cogs.12261)
79. Levy SD, Gayler R. 2008 Vector Symbolic Architectures: a new building material for artificial general intelligence. In *Proc. 2008 Conf. Artificial General Intelligence* (eds P Wang, B Goertzel, S Franklin), pp. 414–418. Amsterdam, The Netherlands: IOS Press.
80. Hodgkin AL, Huxley AF. 1952 A quantitative description of membrane current and its application to conduction and excitation in nerve. *J. Physiol.* **117**, 500–544. (doi:10.1113/jphysiol.1952.sp004764)
81. Izhikevich EM. 2003 Simple model of spiking neurons. *IEEE Trans. Neural Netw.* **14**, 1569–1572. (doi:10.1109/TNN.2003.820440)
82. Stewart TC, Tripp B, Eliasmith C. 2009 Python scripting in the Nengo simulator. *Front. Neuroinform.* **3**, 7. (doi:10.3389/neuro.11.007.2009)
83. Eliasmith C, Anderson CH, Sejnowski TJ, Poggio TA. 2004 *Neural engineering: computation, representation, and dynamics in neurobiological systems*, 2nd edn. Cambridge, MA: MIT Press.
84. de Vries M, Christiansen M, Petersson KM. 2001 Learning recursion: multiple nested and crossed dependencies. *Biolinguistics* **5**, 10–35.
85. Silver RA. 2010 Neuronal arithmetic. *Nat. Rev. Neurosci.* **11**, 474–489. (doi:10.1038/nrn2864)
86. Kikuchi Y, Ip J, Lagier G, Mossom JC, Kumar S, Petkov CI, Barraclough NE, Vuong QC. 2019 Interactions between conscious and subconscious signals: selective attention under feature-based competition increases neural selectivity during brain adaptation. *J. Neurosci.* **39**, 5506–5516. (doi:10.1523/JNEUROSCI.3052-18.2019)
87. Reynolds JH, Heeger DJ. 2009 The normalization model of attention. *Neuron* **61**, 168–185. (doi:10.1016/j.neuron.2009.01.002)
88. Beyeler M, Rounds EL, Carlson KD, Dutt N, Krichmar JL. 2019 Neural correlates of sparse coding and dimensionality reduction. *PLoS Comput. Biol.* **15**, e1006908. (doi:10.1371/journal.pcbi.1006908)
89. Erez J, Cusack R, Kendall W, Barense MD. 2016 Conjunctive coding of complex object features. *Cereb. Cortex* **26**, 2271–2282. (doi:10.1093/cercor/bhw081)
90. Komorowski RW, Manns JR, Eichenbaum H. 2009 Robust conjunctive item–place coding by hippocampal neurons parallels learning what happens where. *J. Neurosci.* **29**, 9918–9929. (doi:10.1523/JNEUROSCI.1378-09.2009)
91. Rigotti M, Barak O, Warden MR, Wang X-J, Daw ND, Miller EK, Fusi S. 2013 The importance of mixed selectivity in complex cognitive tasks. *Nature* **497**, 585–590. (doi:10.1038/nature12160)
92. Parthasarathy A, Herikstad R, Bong JH, Medina FS, Libedinsky C, Yen S-C. 2017 Mixed selectivity morphs population codes in prefrontal cortex. *Nat. Neurosci.* **20**, 1770–1779. (doi:10.1038/s41593-017-0003-2)
93. Whittington MA, Cunningham MO, LeBeau FEN, Racca C, Traub RD. 2011 Multiple origins of the cortical gamma rhythm. *Dev. Neurobiol.* **71**, 92–106. (doi:10.1002/dneu.20814)
94. Buzsáki G, Wang X-J. 2012 Mechanisms of gamma oscillations. *Annu. Rev. Neurosci.* **35**, 203–225. (doi:10.1146/annurev-neuro-062111-150444)
95. Fries P. 2015 Rhythms for cognition: communication through coherence. *Neuron* **88**, 220–235. (doi:10.1016/j.neuron.2015.09.034)
96. Kikuchi Y *et al.* 2017 Sequence learning modulates neural responses and oscillatory coupling in human and monkey auditory cortex. *PLoS Biol.* **15**, e2000219. (doi:10.1371/journal.pbio.2000219)
97. Lisman JE, Jensen O. 2013 The theta-gamma neural code. *Neuron* **77**, 1002–1016. (doi:10.1016/j.neuron.2013.03.007)
98. Canolty RT, Edwards E, Dalal SS, Soltani M, Nagarajan SS, Kirsch HE, Berger MS, Barbaro NM, Knight RT. 2006 High gamma power is phase-locked to theta oscillations in human neocortex. *Science* **313**, 1626–1628. (doi:10.1126/science.1128115)
99. Choo F-X, Eliasmith C. 2010 A spiking neuron model of serial-order recall. In *Proc. 32nd Annu. Conf. Cogn. Sci. Soc.* (eds S Ohlsson, R Catrambone), pp. 2188–2193. Austin, TX: Cognitive Science Society.
100. Comins JA, Gentner TQ. 2010 Working memory for patterned sequences of auditory objects in a songbird. *Cognition* **117**, 38–53. (doi:10.1016/j.cognition.2010.06.009)
101. Carpenter AF, Baud-Bovy G, Georgopoulos AP, Pellizzer G. 2018 Encoding of serial order in working memory: neuronal activity in motor, premotor, and prefrontal cortex during a memory scanning task. *J. Neurosci.* **38**, 4912–4933. (doi:10.1523/JNEUROSCI.3294-17.2018)
102. Petrides M. 1991 Functional specialization within the dorsolateral frontal cortex for serial order memory. *Proc. R. Soc. Lond. B* **246**, 299–306. (doi:10.1098/rspb.1991.0158)

103. Long NM, Kahana MJ. 2019 Hippocampal contributions to serial-order memory. *Hippocampus* **29**, 252–259. (doi:10.1002/hipo.23025)
104. Henson RNA. 1999 Positional information in short-term memory: relative or absolute? *Mem. Cognit.* **27**, 915–927. (doi:10.3758/BF03198544)
105. Endress AD, Carden S, Versace E, Hauser MD. 2010 The apes' edge: positional learning in chimpanzees and humans. *Anim. Cogn.* **13**, 483–495. (doi:10.1007/s10071-009-0299-8)
106. Endress AD, Nespor M, Mehler J. 2009 Perceptual and memory constraints on language acquisition. *Trends Cogn. Sci.* **13**, 348–353. (doi:10.1016/j.tics.2009.05.005)
107. Cona G, Semenza C. 2017 Supplementary motor area as key structure for domain-general sequence processing: a unified account. *Neurosci. Biobehav. Rev.* **72**, 28–42. (doi:10.1016/j.neubiorev.2016.10.033)
108. Summerfield C, de Lange FP. 2014 Expectation in perceptual decision making: neural and computational mechanisms. *Nat. Rev. Neurosci.* **15**, 745–756. (doi:10.1038/nrn3838)
109. Rong F, Isenberg AL, Sun E, Hickok G. 2018 The neuroanatomy of speech sequencing at the syllable level. *PLoS ONE* **13**, e0196381. (doi:10.1371/journal.pone.0196381)
110. Stabler EP. 2004 Varieties of crossing dependencies: structure dependence and mild context sensitivity. *Cogn. Sci.* **28**, 699–720. (doi:10.1207/s15516709cog2805_4)
111. Opitz B, Friederici AD. 2003 Interactions of the hippocampal system and the prefrontal cortex in learning language-like rules. *Neuroimage* **19**, 1730–1737. (doi:10.1016/S1053-8119(03)00170-8)
112. Breitenstein C, Jansen A, Deppe M, Foerster A-F, Sommer J, Wolbers T, Knecht S. 2005 Hippocampus activity differentiates good from poor learners of a novel lexicon. *Neuroimage* **25**, 958–968. (doi:10.1016/j.neuroimage.2004.12.019)
113. Schiffer A-M, Ahlheim C, Wurm MF, Schubotz RI. 2012 Surprised at all the entropy: hippocampal, caudate and midbrain contributions to learning from prediction errors. *PLoS ONE* **7**, e36445. (doi:10.1371/journal.pone.0036445)
114. Friederici AD. 2011 The brain basis of language processing: from structure to function. *Physiol. Rev.* **91**, 1357–1392. (doi:10.1152/physrev.00006.2011)
115. Lundqvist M, Herman P, Miller EK. 2018 Working memory: delay activity, yes! Persistent activity? Maybe not. *J. Neurosci.* **38**, 7013–7019. (doi:10.1523/JNEUROSCI.2485-17.2018)
116. Aubin S, Voelker AR, Elias Smith C. 2017 Improving with practice: a neural model of mathematical development. *Top. Cogn. Sci.* **9**, 6–20. (doi:10.1111/tops.12242)